

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Characterization Theorems by Generalized Indiscernibles

Permalink

<https://escholarship.org/uc/item/8n85b4n7>

Author

Scow, Lynn Cho

Publication Date

2010

Peer reviewed|Thesis/dissertation

Characterization Theorems by Generalized Indiscernibles

by

Lynn Cho Scow

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Thomas Scanlon, Chair
Professor Leo A. Harrington
Professor George Mark Bergman
Professor John MacFarlane

Fall 2010

Characterization Theorems by Generalized Indiscernibles

Copyright 2010
by
Lynn Cho Scow

Abstract

Characterization Theorems by Generalized Indiscernibles

by

Lynn Cho Scow

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Thomas Scanlon, Chair

We generalize the Unstable Formula Theorem characterization of stable theories in [She78]: that a theory T is stable just in case any infinite indiscernible sequence in a model of T is an indiscernible set. We use a generalized form of indiscernibles from [She78, She09]: in our notation, a sequence of parameters from an L -structure M , $(\bar{a}_i : i \in I)$, indexed by an L' -structure I is *L' -generalized indiscernible in M* if $\text{qftp}^{L'}(\bar{i}; I) = \text{qftp}^{L'}(\bar{j}; I)$ implies $\text{tp}^L(\bar{a}_{\bar{i}}; M) = \text{tp}^L(\bar{a}_{\bar{j}}; M)$ for all same-length, finite \bar{i}, \bar{j} from I . Let T_g be the theory of linearly ordered symmetric graphs with no loops in the language with signature $\{<, R\}$, L_g . Say that a *quantifier-free weakly-saturated* model of an L -theory T is some model $M \models T$ that embeds realizations of all quantifier-free L -types consistent with T . We show that a theory T is NIP if and only if every quantifier-free weakly-saturated L_g -generalized indiscernible in a model of T is an indiscernible sequence.

In the process of proving this characterization, a result is introduced relating the utility of generalized indiscernibles indexed by a structure I to the Ramsey-type properties of the structures in $\text{age}(I)$. Results about the previous development of generalized indiscernibles and known results about certain Ramsey-type properties are explicated in the text.

To my mother, in memory.

Contents

1	Introduction	1
1.1	Model theory definitions and conventions	3
1.1.1	Numerical conventions	3
1.1.2	Languages and Structures	3
1.1.3	Types	5
1.2	Amalgamation Classes	7
1.2.1	Ordered graphs	9
1.3	The random ordered graph	9
1.4	Concepts from classification theory	13
1.4.1	Stability and the Independence Property (IP)	14
1.5	Characterizing classes of theories	15
2	Generalized indiscernibles	16
2.1	Motivation	16
2.2	Basic definitions for generalized indiscernibles	18
2.3	Mechanics	20
3	The modeling property	22
3.1	Definitions	22
3.2	Ramsey classes	23
3.2.1	Preliminaries	23
3.2.2	Definition of a Ramsey Class	24
3.2.3	Results around Ramsey classes	26
3.3	Ramsey Classes and the Modeling Property	30
3.3.1	Characterization: sufficiency	32
3.3.2	Characterization: necessity	38
4	Trees	46
4.1	Ramsey theorems for trees: introduction	46
4.2	Full binary tree indiscernibles: overview	46
4.3	Definitions	47

4.4	Halpern-Läuchli results:	51
4.5	Argument	51
4.5.1	Some further explanations	54
4.6	The Ramsey theorem and fbtis	60
4.7	Further directions	61
4.7.1	in the direction of 1-fbti's	61
4.7.2	In the direction of 2-fbti's	62
5	Characterizing NIP theories	63
5.1	Definitions	63
5.2	Characterization of NIP T	63
	Bibliography	68

Acknowledgments

I want to thank Tom Scanlon, my advisor, for his constant encouragement, for his interest in esoteric claims, and for his practical approach. I want to thank Leo Harrington, for treating the strange as familiar in mathematics and inviting me to do the same. I want to thank Cafe 504 in Oakland, without which this dissertation would never have been written. I want to thank my model-theoretic siblings, my family, and Scott.

Chapter 1

Introduction

This paper aims to generalize a well-known result from [She78], the result that a theory T is stable if and only if for any model M of T , every infinite indiscernible sequence in M is an indiscernible set. We use slightly different notation to express a notion of generalized indiscernibility put forth in [She78, She09]. Consider a sequence $a_i = f(i)$ given by an injection, $f : I \rightarrow M$ (or M^n for some fixed n), where the indexing model, I , is an L' -structure and M is an L -structure, for some languages L, L' . Fix a sublanguage $L'' \subseteq L'$. We say this sequence $(a_i : i \in I)$ is L'' -generalized indiscernible in M if any same-length tuples of indices having the same complete quantifier-free L'' -type in I map to tuples of elements having the same complete L -type in M . (See [BSon, DS04, LS03] for more examples of L'' -indiscernibles.)

Say that a property of theories is a “combinatorial weakness property” if it is given by forbidding a certain collection of finite graphs whose edges are given by definable relations but whose set of vertices is not necessarily definable. If we examine the classical result more closely, we see that it characterizes a “combinatorial weakness-property” of T (e.g., stability) by way of the following: that for $\rho' = \{<\}$, any ρ' -indiscernible indexed by an infinite linear order in any model M of T is an $\{\}$ -indiscernible, i.e. an indiscernible set. The ambient theory of M forces a reduction on the theory indexing the indiscernible.

A theory is stable if and only if there is no formula that defines a linear order on some set of parameters in a model of that theory. The property of being stable is a strong dividing line for theories (the parameters are allowed to be themselves sequences of elements from the model.) Not having the independence property (equivalently, having NIP) is known to be a more general property than being stable, and provides a robust dividing line even among unstable theories. A theory has NIP just in case there is no formula whose instances define every subset of an infinite set of parameters in some model of the theory.

This paper aims to generalize on the reduction theorem stated above in the case that our “combinatorial weakness property” is NIP, instead of the property of stability. The general form of our characterization is as follows. We work with models of an L' -theory T' as our index structures, and we distinguish a sublanguage $L'' \subseteq L'$. We then characterize NIP

L -theories T in terms of whether, for “sufficiently” saturated models I , any L' -generalized indiscernible sequence in a model $M \models T$ indexed by I is in fact L'' -generalized indiscernible.

Let L_g be the first order language with signature ρ_g , where $\rho_g = \{<, R\}$, interpreted as an order and an edge relation, respectively. Let T_g be the L_g -theory of linearly ordered symmetric graphs with no loops. Say I is a *qf-weakly-saturated* (“quantifier-free-weakly-saturated”) model of a theory T' if it realizes all the quantifier-free types consistent with T' . In the following, refer to a ρ_g -indiscernible as an *ordered graph-indiscernible*; we call it a *qf-weakly-saturated ordered graph-indiscernible* if it is an ordered graph-indiscernible indexed by a qf-weakly-saturated model $I \models T_g$.

We prove the following characterization of NIP theories:

Theorem 5.2.10. (*NIP Characterization Theorem*) *A theory T has NIP if and only if any quantifier-free weakly-saturated ordered graph-indiscernible in a model of T is an indiscernible sequence.*

In Chapter 3 we introduce elements from Nešetřil’s theory of Ramsey Classes. [Neš05] Consider a class \mathcal{U} of finite relational structures. For the purposes of our discussion, by an *A-subobject of B* we will mean a substructure of B isomorphic to A . A class of finite relational structures \mathcal{U} is a Ramsey class if for any finite integer k and objects A, B in \mathcal{U} , there is a C in \mathcal{U} satisfying the following: for any k -coloring of the A -subobjects of C we can find $B' \subseteq C$, B' isomorphic to B , so that all the A -subobjects of B' are colored the same color. This is a natural extension of the classical Ramsey theorem for finite sets, which states that for any integers k, n, m there is a number N so that for any k -coloring of the n -element subsets of $[N]$, we can find a subset of $[N]$ of size m so that all the n -element sets from this subset are colored the same under this coloring.

In the course of proving Theorem 5.2.10, we will develop the notion of L' -generalized indiscernibles having the “modeling property”, with respect to some set of finite L' -structures, \mathbf{K} , and prove a characterization of this property in terms of Ramsey classes (see Ch. 3 for all relevant definitions):

Theorem 3.3.27. *Let L' be a finite relational language containing a binary relation symbol for order, $<$, and let \mathcal{U} be some nonempty collection of finite L' -structures that are linearly ordered by $<$. Let \mathcal{U}^* be the set of isomorphism types from \mathcal{U} . Suppose that \mathcal{U} has JEP and is closed under isomorphism and substructures.*

\mathcal{U} is a Ramsey Class if and only if L' -generalized indiscernibles have the modeling property for \mathcal{U}^ .*

We give model-theoretic definitions and conventions in this chapter, present the theory of the random ordered graph, and a lemma for stretching I -indiscernibles; in Chapter 3 we define versions of the “modeling property”, give combinatorial definitions and prove our main combinatorial result, Theorem 3.3.27; in Chapter 4 we address the special case of “full binary tree indiscernibles” from [DS04] as an illustration of the modeling property; in Chapter 5 we prove the main result, Theorem 5.2.10.

We wish to say that much of the discussion of generalized indiscernibles is prefigured by material in Ch.VII § 2 of [She78], in particular the lemma on stretching generalized indiscernibles amounts to a clarification of Lemma 2.2, [She78].

1.1 Model theory definitions and conventions

1.1.1 Numerical conventions

When we talk about all i such that $i \leq k$, we mean the integers $1 \leq i \leq k$; when we talk about all i such that $i \in n$, we mean the set $\{0, 1, \dots, n-1\}$. Generally, by i, j, k, m, n we denote natural numbers and by A, B, C we denote sets. We mean s, t, x, y, v to denote variables when they occur in the context of formulas, e.g. $\varphi(s_1, \dots, s_m)$. By a *countable* set we mean a set with cardinality $\leq \aleph_0$. Given a finite tuple \bar{a} by $(\bar{a})_i$ we mean the i -th element of the tuple. By $l(\bar{a})$ we mean the length of the tuple.

By B^A we mean all the functions $f : A \rightarrow B$. By $[1, n]$ we mean the set $\{1, \dots, n\}$ of integers between 1 and n . By $\text{Sym}(n)$, we mean the symmetric group on $\{1, \dots, n\}$.

For $\vec{i} := \langle i_1, \dots, i_n \rangle$ and parameters $(a_i : i \in I)$, by $\bar{a}_{\vec{i}}$ we mean $\langle a_{i_1}, \dots, a_{i_n} \rangle$. For a set A , the set $[A]^n$ is defined to be all n -element subsets of A .

1.1.2 Languages and Structures

Here we merely list some of the conventions we will use, but for a rigorous development of languages and structures see [CK73]. All languages L are assumed to be first-order, and all structures M are assumed to be nonempty. All languages are assumed to have a relation symbol for equality. Given a first-order formula φ , we use the convention that $\varphi^0 := \varphi$ and $\varphi^1 := \neg(\varphi)$. For some condition $P(i)$ such as “ $i \leq n$ ” and φ containing parameters indexed by i , by $\varphi^{P(i)}$ we will mean φ^0 if $P(i)$ holds, and φ^1 if $\neg P(i)$ holds. For an L -structure M , by $|M|$ we mean the underlying set of this model, and by $||M||$ we mean the cardinality of this set. Given a sublanguage $L' \subseteq L$ and an L -structure M , by $M \upharpoonright L'$ we mean the reduct of M to the L' sublanguage (a structure with the same underlying set that only interprets relations from L' and only in the same way as M).

We mean the notation L to denote *all* the first-order formulas in some signature, ρ (sometimes denoted $\text{Fm}(L)$, but we will not do so.) We assume that every L -formula is built out of the functions, relations and constants from the signature of L as well as the variables $x_1, x_2, \dots, x_n, \dots$. By abuse of notation, we list signatures ρ not as sequences of integers, but explicitly as some (usually finite) set of symbols of the language, e.g. $\rho := \{R_1, R_2, f\}$ where we specify which of these symbols are functions or relations and also their arities. By an *atomic formula* we mean an instantiation of a relation symbol by variables or terms or else a statement of equality between terms.

An L -theory T is a consistent set of L -sentences, i.e., formulas with no open variables. We consider theories to be closed under logical consequence. An *axiomatization* of T is a set of sentences whose closure under logical consequence is the set of all sentences in T . A *universal* theory is a theory axiomatized by universal formulas, i.e. of the form $(\forall x_1) \dots (\forall x_n)[\phi(x_1, \dots, x_n)]$ where ϕ is a quantifier-free formula. For an L -theory T , by T_\forall we mean all the consequences of T that are universal L -sentences.

A complete theory is an L -theory, for some language L , that is maximal consistent for L -sentences. For an L -structure M , by $\text{Th}(M)$ we denote the full first-order theory of M , namely every L -formula ψ such that $M \models \psi$. The full first order theory of any structure is always a complete theory. By $\text{Th}_{|M|}(M)$ we mean the full first order theory of M with constants for all the elements in $|M|$.

As for notation, we may refer to the language in signature ρ as L_ρ . For a structure I , we may refer to its language as $L(I)$. An L -formula φ is n -ary if its variables are among x_1, \dots, x_n (sometimes written as $\varphi = \varphi(x_1, \dots, x_n)$.) By $L_{\text{at}}(x_1, \dots, x_n)$ we mean all atomic, n -ary L -formulas. And we define $L_{\text{at}} := \bigcup_n L_{\text{at}}(x_1, \dots, x_n)$.

We have a relation of satisfaction between L -formulas, L -structures M and elements of M : say $M \models \varphi(a_1, \dots, a_n)$ if $\langle a_1, \dots, a_n \rangle \in \varphi^M$, i.e. the interpretation of φ as a subset of M^n contains the tuple \bar{a} . A set $A \subseteq M^k$ is called *definable* if it is equal to the set ψ^M for some k -ary L' -formulae ψ . For a collection of k -ary L -formulas, $\Psi := \{\psi_j : j \in J\}$ and elements $\bar{i} := \langle i_1, \dots, i_k \rangle$ in an L -structure M , by $\langle i_1, \dots, i_k \rangle \models \Psi(x_1, \dots, x_k)$ (or $\bar{i} \models \Psi$) we mean that for the relevant structure M , $M \models \psi_j(i_1, \dots, i_k)$, for all $j \in J$.

We will be using the model-theoretic *algebraic closure*, a generalization of algebraic closure in fields:

Definition 1.1.1. Given a subset $A \subseteq |M|$ of some L -structure M , we say that $b \in \text{acl}(A)$ (b is in the *algebraic closure* of A) if there is some L -formula $\varphi(x; \bar{y})$ and some tuple of elements \bar{a} from A such that

1. $M \models \varphi(b; \bar{a})$, and
2. $\{c : M \models \varphi(c; \bar{a})\}$ is finite

Definition 1.1.2. ([Hod93]) We say that an L -structure M has *no algebraicity* if for any subset $A \subseteq |M|$, $\text{acl}(A) = A$.

We will be making some use of ultraproducts, so we restate one of the main theorems here, though details may be found in [CK73]. First we need one definition:

Definition 1.1.3. Given an ultrafilter $D \subseteq \mathcal{P}(I)$ and $\bar{a} := \langle a_i : i \in I \rangle$ and $\bar{b} := \langle b_i : i \in I \rangle$, say that $\bar{a} \sim \bar{b}$ modulo the ultrafilter D if $\{i \in I : a_i = b_i\} \in D$.

Given an ultrafilter $\mathcal{D} \subseteq \mathcal{P}(I)$ for some set I , and L -structures $(M_i : i \in I)$ we may form a new L -structure, the ultraproduct of the M_i with respect to \mathcal{D} , $\mathcal{M} = \prod_{i \in I} M_i / \mathcal{D}$. Elements in \mathcal{M} are tuples $\bar{a} = \langle a_i : i \in I \rangle$ modulo the ultrafilter.

Theorem 1.1.4. (*Loś*) Fix an n -ary formula φ and elements \bar{a}^i for $i \leq n$ of $\prod_{i \in I} M_i/D$.

$$\prod_{i \in I} M_i/D \models \varphi(\bar{a}^1, \dots, \bar{a}^n) \Leftrightarrow \{i \in I : M_i \models \varphi(a_i^1, \dots, a_i^n)\} \in \mathcal{D}$$

As a consequence, the ultraproduct satisfies a sentence just in case almost all the M_i do so.

Special maps

We work almost exclusively in relational languages, and the following holds for relational languages, L :

Definition 1.1.5. For two L structures A, B a map $f : A \rightarrow B$ is

1. an L -homomorphism if for all relations R in L , $R^A(\bar{a}) \Rightarrow R^B(f(\bar{a}))$, where R^A, R^B are the interpretations of R in A, B respectively.
2. an L' -isomorphism if it is a 1-1 and onto map and for all relations R in L , $R^A(\bar{a}) \Leftrightarrow R^B(f(\bar{a}))$ (an L -automorphism is an L -isomorphism from some structure A onto itself.)
3. an L -embedding if it is an L -homomorphism $f : A \rightarrow B$ and there is an L -homomorphism $g : \text{ran}(f) \rightarrow A$ such that $f \circ g = g \circ f = \text{id}$.

In other words, f is an isomorphism of A with $\text{ran}(f)$.

4. an L -elementary embedding if it is an L -embedding $f : A \rightarrow B$ and for all L -formulas $\varphi(\bar{x})$ (i.e., formulas with quantifiers):

$$\varphi^A(\bar{a}) \Leftrightarrow \varphi^B(f(\bar{a}))$$

1.1.3 Types

A type is a collection of formulas in some first-order language with some specified list of one or more free variables. If we mean the type to be consistent, or consistent with some theory, we will say so explicitly.

Notation 1.1.6. For a finite type p , by $\bigwedge p$ we will mean the formula obtained by the conjunction of all formulas in p .

The following are useful as definitions for sets of quantifier-free formulas:

Definition 1.1.7. Let L be any language, $m < \omega$, $\Delta \subseteq L$ some subset of L -formulas that is closed under negation and permutation of variables:

1. p is a (Δ, m) -type if p is a consistent set of m -ary formulas and for every φ in p , either φ or $\neg\varphi$ belongs to Δ .

2. p is a Δ -type if it is a (Δ, m) -type and Δ is understood to consist exclusively of m -ary formulas.
3. p is a *complete* (Δ, m) -type if it is a (Δ, m) -type that is a maximal consistent set of formulas from Δ . (Since p is assumed consistent, p is maximal consistent just in case for any ψ from Δ , exactly one of ψ and $\neg\psi$ is in p .)
4. p is a *quantifier-free* (L, m) -type if p is a (Σ, m) -type where Σ is the set of all quantifier-free, m -ary L -formulas.
5. p is a *quantifier-free* L -type if it is a quantifier-free (L, m) -type for some m .
6. p is a *complete quantifier-free* L -type if it is a complete quantifier-free (L, m) -type for some m .

The following are definitions for so-called “complete” types of elements in structures:

Definition 1.1.8. Fix an L -structure M and a subset $\Sigma \subseteq L$, and some finite sequence from M , $\bar{a} := \langle a_1, \dots, a_n \rangle$.

1. by $\text{tp}^\Sigma(\bar{a}; M)$ we mean the set of all formulas $\varphi \in \Sigma$ satisfied by \bar{a} in M .
2. by $\text{tp}^L(\bar{a}; M)$ we mean the set of all L -formulas satisfied by \bar{a} in M .
3. by $\text{qftp}^L(\bar{a}; M)$ we mean the set of all quantifier-free L -formulas satisfied by \bar{a} in M .

Definition 1.1.9. For $\Delta(y_1, \dots, y_m) \subseteq L$ some finite subset of L -formulas, define

$$S_m^\Delta(\emptyset; M)$$

to be the set of complete (Δ, m) -types realized by tuples from M .

Definition 1.1.10. We say that I is a *quantifier-free weakly-saturated* model of a theory T' if it realizes all quantifier-free types consistent with T' .

Definition 1.1.11. We say a theory T' is *quantifier-free* \aleph_0 -categorical if there are only finitely many quantifier-free n -types consistent with T' for each n , up to equivalence in T' .

Remark 1.1.12. Note that in the case that L' is a finite relational language, any L' -theory T' is automatically quantifier-free \aleph_0 -categorical.

1.2 Amalgamation Classes

In some applications of indiscernible sequences, we wish to take an indiscernible sequence indexed by an infinite linear order and extend it to an indiscernible sequence indexed by a dense linear order. The theory of dense linear orders with no endpoints (DLO) is related to the theory of linear orders in at least two ways. DLO is the model companion of the theory of orders, there is a unique countable model of DLO, and this model can be obtained by taking the Fraïssé limit of all finite linearly ordered sets.

In general, we may wish to find an index model J with a relationship to a given index model I analogous to the relationship of a dense linear order without endpoints to an infinite linear order. To do this in the context of graphs, we must first develop the theory of the “random ordered graph” and revisit some well-known model-theoretic definitions. For more references on these topics see [Hod93] and [Lac87].

Definition 1.2.1. Given any language L and class of L -structures, U , closed under isomorphism

1. U has the *joint embedding property (JEP)* if for any $B_1, B_2 \in U$ there is a $D \in U$ and L -embeddings h_1, h_2 such that $h_i : B_i \rightarrow D$.
2. U has the *amalgamation property (AP)* if for any $A, B_1, B_2 \in U$ and L -embeddings $f_i : A \rightarrow B_i$ (we refer to such a family A, B_1, B_2, f_1, f_2 as an amalgamation base), there is an *amalgam* $C \in U$ and L -embeddings $g_i : B_i \rightarrow C$ such that $g_1 \circ f_1 = g_2 \circ f_2$.
3. U has *strong amalgamation* if it has AP in addition to which, for any amalgamation base A, B_1, B_2, f_1, f_2 we may always choose a strong amalgamation: i.e. C, g_1, g_2 such that $g_1(B_1) \cap g_2(B_2) = (g_1 \circ f_1)(A)$.
4. U has *free amalgamation* if every amalgamation base A, B_1, B_2, f_1, f_2 has a strong amalgamation C, g_1, g_2 such that whenever $R^C(c_1, \dots, c_m)$ we have that either $c_i \in \text{ran}(g_1 \circ f_1)$ for all i , or else $c_i \in \text{ran}(g_2 \circ f_2)$ for all i .

Remark 1.2.2. Strong amalgamation says that there is no collapsing of elements from B_1, B_2 other than is absolutely necessary from the amalgamation condition (i.e., that corresponding points from A collapse.)

Definition 1.2.3. Let L be some language and U a class of finitely generated L -structures closed under isomorphism and substructure. We say that U is an *amalgamation class* if it has both JEP and AP.

Definition 1.2.4. For a class \mathcal{U} of structures, we say that a set \mathbf{K} is a *set of isomorphism types from \mathcal{U}* if for every structure $M \in \mathcal{U}$ there is exactly one structure $M' \in \mathbf{K}$ isomorphic to M .

We quote some helpful definitions from [Hod93]:

Definition 1.2.5. (age) A class J closed under isomorphism is the *age* of a structure M if it contains exactly the isomorphism types of finitely generated substructures of M .

Definition 1.2.6. A structure M is *ultrahomogeneous* if every isomorphism between finitely generated substructures of M extends to an automorphism of the model.

Definition 1.2.7. A structure M is *weakly homogeneous* if for any A, B finitely generated substructures of M , $A \subseteq B$ and any embedding $f : A \rightarrow M$, there is an embedding $g : B \rightarrow M$ extending f .

Note that ultrahomogeneous structures are easily weakly homogeneous.

We will be using the following result from Fraïssé:

Definition 1.2.8. ([Fra54], presentation from [Hod93]) Fix some countable language L , and let \mathbf{K} be a nonempty countable set of finitely generated L -structures, which is closed under substructures and has JEP and AP. Then there is an L -structure \mathfrak{A} , unique up to isomorphism, such that

1. $|\mathfrak{A}|$ is countable,
2. $\mathbf{K} = \text{age}(\mathfrak{A})$, and
3. \mathfrak{A} is ultrahomogeneous.

\mathfrak{A} is referred to as the *Fraïssé limit* of \mathbf{K} .

Remark 1.2.9. Note that for a finite relational language L' , all ages of L' -structures are countable.

Some useful results that can be read in [Hod93]:

Theorem 1.2.10. *For a countable ultrahomogeneous M , $\mathbf{K} = \text{age}(M)$ has strong amalgamation if and only if M has no algebraicity.*

Lemma 1.2.11. *Let L' be a finite relational language and \mathbf{K} some nonempty set of isomorphism types of finite L' -structures. Suppose that \mathbf{K} is an amalgamation class with Fraïssé limit \mathcal{I} . Suppose that J is an L' -structure with $\text{age}(J) \subseteq \mathbf{K}$. Then \mathcal{I} embeds a copy of J .*

1.2.1 Ordered graphs

First let us fix the theory T_g in $L_g := \{<, R\}$ that says that $A \models T_g$ is a symmetric graph (a graph with symmetric edge relation) with no loops that is linearly ordered by $<$. In other words, A is an undirected graph with a linear ordering. This is a universal theory:

$$T_g = \{(\forall x)(\forall y)(xRy \rightarrow yRx), (\forall x)(\neg xRx), (\forall x, y)(x < y \rightarrow \neg y < x), \\ (\forall x)(\neg x < x), (\forall x)(\forall y)(\forall z)(x \leq y \wedge y \leq z \rightarrow x \leq z)\}$$

Let \mathcal{K} be the class of finite structures embeddable in models of T_g . Note that since L_g is a finite relational language, the finitely-generated L_g -structures are exactly the finite L_g -structures. Moreover, finite ordered symmetric graphs are clearly closed under substructure, have JEP and AP. Thus, \mathcal{K} is an amalgamation class.

Since \mathcal{K} is an amalgamation class, it has a Fraïssé limit, \mathcal{R} . In the case of finite relational languages, this is the unique countable, ultrahomogeneous model whose age is equal to \mathcal{K} . Refer to this \mathcal{R} as the “random ordered graph”. We know from standard arguments that since L_g is finite and relational, $\text{Th}(\mathcal{R}) =: T_{\mathcal{R}}$ in L_g admits elimination of quantifiers and is \aleph_0 -categorical.

1.3 The random ordered graph

In this section we use Fraïssé’s theorem in the case of ordered symmetric graphs with no loops to develop the theory of what we will call the “random ordered graph”. We do this in analogy to the way in which the theory of dense linear orders without endpoints is obtained from the theory of linear orders. For illustration, we give the complete proof for ordered graphs, as we will use the random ordered graph as an index set in Chapter 5.

The case of the random graph is worked out in detail in Ch. VII of [Hod93], as is the case of the *random structure* in a language with finite signature. We include the case of the random ordered graph here for illustration. In the next few lemmas we will be finding an explicit description for the theory of the structure \mathcal{R} , from the previous section, and illustrating its relationship with T_g . First, we show that $T_g \subseteq T_{\mathcal{R}}$:

Claim 1.3.1. \mathcal{R} is a model of T_g

Proof. T_g is a set of universal axioms $\forall \bar{x} \Phi(\bar{x})$, where $\Phi(\bar{x})$ is quantifier free. If \mathcal{R} failed one of these axioms, there would be $A \subseteq \mathcal{R}$, a substructure, such that $\bar{a} \models \neg \Phi(\bar{x})$, for some enumeration $\bar{a} = A$. But this A must be in the age of \mathcal{R} , \mathcal{K} , by definition of \mathcal{R} , so it must satisfy all the axioms in T_g , in particular, elements from A must realize $\Phi(\bar{x})$. \square

Here is another way to think of the theory of \mathcal{R} , in analogy to the role the theory of dense linear orders plays in the use of indiscernible sequences. First let us give an axiomatization of $T_{\mathcal{R}}$.

Definition 1.3.2. By a complete ordered graph type we mean a complete quantifier-free type consistent with T_g .

Remark 1.3.3. A complete ordered graph type $p(x_1, \dots, x_k)$ is a combinatorial object. It can be associated with an ordered pair (σ, τ) where $\sigma \in \text{Sym}(k)$ and $\tau \in 2^{[1,k]^2}$ such that we have the following equivalence:

$$\left(\bigwedge p(x_1, \dots, x_n) \right) \leftrightarrow \left[\bigwedge (x_{\sigma(i)} < x_{\sigma(j)}) \wedge \left(\bigwedge_{\langle i,j \rangle \in [1,k]^2} R(x_i, x_j)^\tau \right) \right]$$

Remark 1.3.4. Note that complete ordered graph types are exactly the atomic diagrams of members of \mathcal{K} .

Definition 1.3.5. Call the following theory T_t (in the following, note that complete ordered graph types are formulas):

1. T_g
2. $p(\bar{z})$ a complete ordered graph type: $(\exists \bar{z})(\bigwedge p)(\bar{z})$
3. consistent pairs $p(\bar{x}), q(\bar{x}; \bar{y})$ of complete ordered graph types:
 $(\forall \bar{x})((\bigwedge p)(\bar{x}) \rightarrow (\exists \bar{y})(\bigwedge q)(\bar{x}; \bar{y}))$

Claim 1.3.6. T_t is $T_{\mathcal{R}}$.

Proof. It suffices to show that T_t and $T_{\mathcal{R}}$ have exactly the same countable models, since the theories of two uncountable models can be brought down to their elementary substructures. For one direction, since $T_{\mathcal{R}}$ is \aleph_0 -categorical, it suffices to show that $\mathcal{R} \models T_t$. We have argued in Claim 1.3.1 that $\mathcal{R} \models T_g$. To see that \mathcal{R} satisfies (2) and (3), note that the age of \mathcal{R} contains all finite ordered graphs and that \mathcal{R} is weakly homogeneous.

For the other direction, take any countable model $M \models T_t$. M must be isomorphic to \mathcal{R} , since (1) and (2) force the age of M to be exactly \mathcal{K} , and \mathcal{R} is the unique weakly homogeneous model with age \mathcal{K} . Thus $M \models T_{\mathcal{R}}$. \square

For the next two claims we will want the following notation. Note that using standard arguments, any partial order can be extended to a linear order.

Definition 1.3.7. For two ordered graphs B_1, B_2 containing the set A , let $B_1 \oplus_A B_2$ refer to the *amalgam of B_1 and B_2 over A* by which we mean the partially ordered graph whose universe is the disjoint union of $B_1, B_2 \setminus A$ and whose graph and order relation are simply the union of the corresponding relations on B_1, B_2 . By $B_1 \oplus B_2$ we mean the amalgam of B_1 and B_2 over $A = \emptyset$.

Refer to $\overline{B_1 \oplus_A B_2}$ as a *closure of B_1 and B_2 over A* , by which we mean the ordered graph obtained from the amalgam of B_1 and B_2 over A by extending the partial order on this structure in the following way to a linear order: in the interval (a_i, a_j) , $a_i, a_j \in A \cup \{\infty, -\infty\}$, all members of B_1 will be below all members of B_2 .

Claim 1.3.8. $T_{\mathcal{R}}$ is the model companion of T_g

Proof. (The following are standard arguments, repeated here since we could not find an appropriate reference.) We have already said that $T_{\mathcal{R}}$ admits elimination of quantifiers, and so it is model complete. By the above axiomatization, it is clear that $T_{\mathcal{R}} \supseteq T_g$. It remains to show any model of T_g can be extended to a model of $T_{\mathcal{R}}$. Let $M \models T_g$ we build a model $N \models T_{\mathcal{R}}$ extending M , in stages. In the following, we ask the reader to please be aware that two mentions of the formula “ p ” do not necessarily refer to the same formula, as, in an effort to cut down on subscripts, we instead refer to “formulas $(\bigwedge p)/(\bigwedge q)$ from the i th axiom in (2)/(3),” for (2) (3) from Definition 1.3.5, where we will make clear the nature of the enumeration.

Here are (2) and (3) restated:

$$(2) (\exists \bar{z})(\bigwedge p)(\bar{z})$$

(for $p(\bar{z})$ a complete ordered graph type)

$$(3) (\forall \bar{x})((\bigwedge p)(\bar{x}) \rightarrow (\exists \bar{y})(\bigwedge q)(\bar{x}; \bar{y}))$$

(for consistent pairs $p(\bar{x}), q(\bar{x}; \bar{y})$ of complete ordered graph types)

First take a countable enumeration of the formulas from (2). Suppose we have an extension M_i of M that realizes the first i formulas from (2). For $i \geq 1$, let B_{i+1} be a realization of $(\bigwedge p)(\bar{z})$, in some enumeration, for p in the $(i+1)$ st type from (2). Let $M_{i+1} := \overline{M_i \oplus B_{i+1}}$. The union $\bigcup_{i=1}^{\omega} M_i =: M_{\omega}$ models all the axioms in (2). Since M_{ω} is the increasing union of ordered graphs M_i , and the theory of ordered graphs is an AE-theory, M_{ω} is also an ordered graph.

Now take a countable enumeration of the axioms from (3) so that each axiom comes up infinitely often. M_{ω} from the previous paragraph has size κ for some cardinal κ . Fix $1 \leq i < \omega$. Suppose we have ordered graphs N_k of cardinality κ , for $1 \leq k \leq i$ such that $M_{\omega} \subseteq N_k \subseteq N_{k+1}$ for $k < i$ and for the $k+1$ st axiom in (3), $(\forall \bar{x})((\bigwedge p)(\bar{x}) \rightarrow (\exists \bar{y})(\bigwedge q)(\bar{x}; \bar{y}))$, for every B from N_k realizing p , in some enumeration, there is a $C \supseteq B$ in N_{k+1} realizing q , in some enumeration extending that of B .

Define N_{i+1} as follows. Let $(\forall \bar{x})((\bigwedge p)(\bar{x}) \rightarrow (\exists \bar{y})(\bigwedge q)(\bar{x}; \bar{y}))$ be the $i+1$ st axiom in (3). Enumerate the substructures of N_i that in some enumeration realize p as $(B_{\alpha}^{i+1} : \alpha \in \kappa)$ (N_i has cardinality κ , by induction.) Now we define N_{i+1} to be the union of the S_{α} for the following ordered graphs S_{α} , for $\alpha \in \kappa$:

- $S_0 := N_i$
- $S_{\alpha+1} := \overline{S_{\alpha} \oplus_{B_{\alpha}^{i+1}} C}$, where C realizes q from the $i+1$ st axiom in (3), in some enumeration extending that of B_{α}^{i+1} .

As the increasing union of ordered graphs, N_{i+1} is itself an ordered graph. By construction for every B from N_{i+1} realizing p from the $i+1$ st axiom in (3), in some enumeration, there

is a $C \supseteq B$ in N_{i+1} realizing q , in some enumeration extending that of B . We also see that N_{i+1} has cardinality κ and extends N_i , which extends M_ω .

Now let $N := \bigcup_{i=1}^\omega N_i$. We argue that N models $T_{\mathcal{R}}$. We already argued that $N \supseteq M_\omega$, and M_ω realizes all the axioms in (2). But then by persistence, so does N . As for the axioms in (3), take any such axiom $\eta := (\forall \bar{x})(\bigwedge p)(\bar{x}) \rightarrow (\exists \bar{y})(\bigwedge q)(\bar{x}; \bar{y})$ and a $B \subseteq N$ realizing p , in some enumeration. Since B is finite, there is an N_k containing all of B such that the $k+1$ st axiom in the enumeration is also η , and since p is quantifier-free, B also realizes p in N_k . But then there is C extending B contained in N_{k+1} that realizes q , in some enumeration extending that of B . But this C is contained in N_{k+1} which is contained in N . Since q is quantifier-free, C realizes q in N . We know that N models T_g since it is the increasing union of models that do so, as argued above. \square

Remark 1.3.9. As an index set, \mathcal{R} will play a similar role to that of dense linear orders in the traditional applications of indiscernible sequences.

Claim 1.3.10. *1. \mathcal{R} is densely ordered. 2. Graph types are dense in the order over finite sets of parameters: for any \bar{i} from \mathcal{R} , R -type $p(\bar{x}; \bar{i})$ and consistent order-type $q(\bar{x}; \bar{i})$ such that both p and q imply all their variables are distinct, we have that $p(\bar{x}; \bar{i}) \cup q(\bar{x}; \bar{i})$ has infinitely many realizations.*

Proof. We will verify these claims directly using quantifier-free weak saturation and weak homogeneity in \mathcal{R} .

1. Consider any $i_1 < i_3$. Since there is a model in \mathcal{K} satisfying $j_1 < j_2 < j_3$ in \mathcal{K} , an extension of the map $j_k \mapsto i_k$ gives a witness in \mathcal{R} to $i_1 < x < i_3$.

2. Given any order type in \mathcal{R} , $q(\bar{x}; \bar{i})$ and graph type $p(\bar{x}; \bar{i})$, we may assume that $q(\bar{x}; \bar{y})$ and $p(\bar{x}; \bar{y})$ are complete order- and R -types, respectively. There is a member C in \mathcal{K} realizing $p \cup q$ in the enumeration $(c_1, \dots, c_k; c_{k+1}, \dots, c_{k+n})$. \mathcal{R} embeds all members of \mathcal{K} , so the type of $(c_{k+1}, \dots, c_{k+n})$ is realized by some $(a_{k+1}, \dots, a_{k+n})$ in \mathcal{R} . The map $c_s \mapsto a_s$ for $k < s \leq (k+n)$ can be extended to all of C to give a realization $(b_1, \dots, b_k; a_{k+1}, \dots, a_{k+n}) \models p \cup q$ in \mathcal{R} .

Moreover, we can get infinitely many such realizations. Let A, B be the finite structures on \bar{a}, \bar{b} , respectively. The structure $B_1 := (B \cup A) \oplus_A (B \cup A)$ is a finite ordered graph, thus in \mathcal{K} , and $A \cup B \subseteq B_1$ by identification with the first coordinate in the amalgam. We have already said \mathcal{R} embeds a copy of $B \cup A$, thus by an argument similar to that in the previous paragraph, embeds a copy of \bar{B}_1 over $B \cup A$. Thus \mathcal{R} embeds two copies of B whose intersection is exactly their embedded copies of A . Iterate this procedure for arbitrarily large n , to get n copies of B in \mathcal{R} disjoint over A . \square

Remark 1.3.11. It is easy to see that T has the independence property. In fact, this is witnessed by $xRy \wedge x < y$ and by $xRy \wedge y < x$.

1.4 Concepts from classification theory

Classification theory is a subdiscipline of model theory in which an aim is to find dividing lines for theories. A property of a theory will be isolated such that having the property puts some constraint on the models of the theory whereas having the negation of the property offers a strong negation of this constraint. One interest in these properties is that one can study the existence of universal models for a theory or else the spectrum of models in a theory.

For example, in a model of a stable theory, for arbitrarily large cardinals λ , there are not more than λ types over any set of parameters of size λ ; in an unstable theory, for any cardinal λ there is a set A of size λ in some model of the theory such that there are $> \lambda$ distinct types over A . One ramifications of this particular divide is that for stable theories T , one can prove in ZFC that there are saturated models of T in arbitrarily large cardinalities.

In [She78] new combinatorial methods were introduced to quantify the complexity of first-order theories. One of these is the notion of *stability*

Definition 1.4.1. An L -theory T is *stable* if for any L -model $M \models T$ there does not exist L -formula $\varphi(\bar{x}; \bar{y})$ such that φ has the *order property* in M , i.e. such that there are parameters \bar{a}_i, \bar{b}_j from M for $i, j < \omega$ such that

$$\varphi(\bar{a}_i; \bar{b}_j) \Leftrightarrow i < j < \omega$$

What seems like a local syntactic property is in fact a robust dividing line for theories.

One important tool of classification theory is the indiscernible sequence, originally named so in [Mos56]:

Definition 1.4.2. (indiscernible sequences) Given a linear order $(I, <)$ a sequence $(a_i : i \in I)$ of parameters (or same-length tuples of parameters) from an L -structure M $(a_i : i \in I)$ is

1. an *indiscernible sequence* if for all $n > 0$ and for all $i_1 < \dots < i_n, j_1 < \dots < j_n$ from I ,

$$\text{tp}^L(a_{i_1}, \dots, a_{i_n}; M) = \text{tp}^L(a_{j_1}, \dots, a_{j_n}; M)$$

2. an *indiscernible set* if for all $n > 0$ and for all sets of n distinct parameters $\{i_1, \dots, i_n\}, \{j_1, \dots, j_n\}$ from I ,

$$\text{tp}^L(a_{i_1}, \dots, a_{i_n}; M) = \text{tp}^L(a_{j_1}, \dots, a_{j_n}; M)$$

It is easy to find a set of parameters that is an indiscernible sequence but not an indiscernible set. Just consider any nonempty subset of a linear order (J, \prec) , $I \subseteq J$, where we take the ordering $<$ on the indices for the parameters to come from the ordering \prec between those parameters on J . The fact that there is this gap between the two notions serves to characterize the entire class of stable theories.

Stability is a useful property in that it has several diverse characterizations, many showing up in the Unstable Formula Theorem – Chapter II, Theorem 2.2 of [She78]. Here is a sampling:

Theorem 1.4.3. ([She78]) *The following are equivalent*

1. T is stable
2. Every infinite indiscernible sequence in a model $M \models T$ is an indiscernible set.

1.4.1 Stability and the Independence Property (IP)

Another notion of complexity introduced in [She78] is the independence property:

Definition 1.4.4. (NIP) A theory T has the *independence property (IP)* if there is a model $M \models T$ and a formula $\psi(\bar{x}; \bar{y})$ such that for all n there is a set of parameters from M , $\{\bar{a}_i : i \in n\}$ with $l(\bar{a}_i) = l(\bar{y})$ such that:

$$M \models \exists \bar{x} \bigwedge_{s \subseteq n} [\psi(\bar{x}; \bar{a}_i)]^{i \in s}$$

We say that a theory T has *not the independence property* if T does not have IP.

Remark 1.4.5. It is known that a theory has IP just in case it has IP witnessed by some formula $\psi(x; \bar{y})$ where $l(x) = 1$. [She78] Thus, we can restate as follows: T has IP if and only there is a model $M \models T$ and subsets $A \subseteq |M|$ of arbitrarily large size so that every subset B is defined by some instance $\psi(x; \bar{c}_B)$ of $\psi(x; \bar{y})$.

Definition 1.4.6. (strict order property) A theory T has the *strict order property* if there is a model $M \models T$ and a formula $\varphi(\bar{x}; \bar{y})$ and tuples $\langle a_i : i \in \omega \rangle$ from M with $l(\bar{a}_i) = l(\bar{y})$ such that for $i < j$

$$M \models (\forall x)[\varphi(\bar{x}, \bar{a}_i) \rightarrow \varphi(\bar{x}, \bar{a}_j)] \wedge (\exists x)[\varphi(\bar{x}, \bar{a}_j) \wedge \neg \varphi(\bar{x}, \bar{a}_i)]$$

It is shown in [She78] that every unstable theory has either the independence property or the strict order property. Thus, the class of NIP theories includes the class of stable theories. In fact, some results for stable theories have been extended to NIP theories. Our project is similarly to take the indiscernible-sequence characterization of stable theories and extend the method to obtain a characterization of NIP theories by way of a generalization of indiscernible sequences.

1.5 Characterizing classes of theories

There are several reasons why the generalization from stable to NIP goes through. For one thing, the properties of being stable and NIP share several important aspects in common. Both are given by forbidding a certain infinite sequence of finite graphs in models M of the theory, such that the edge relation is definable in the model, but the set of points that constitute the vertices of the graph is not necessarily definable. These aspects are well-characterized by taking an infinite structure I , whose submodels consist of the forbidden subgraphs, and allowing arbitrary injections into the underlying set of the model. Suppose further that these injections allow that the positive edge relation in I is captured by some definable relation in M (this is the “indiscernibility”.) Then, should M forbid that such configurations as are present in I be definable in its universe, it will enforce a kind of reduct onto I , at least as far as it is able to capture the definable relations of I in its own language. This is the “reduction theorem.”

Chapter 2


Generalized indiscernibles

2.1 Motivation

The project in this paper is largely to investigate a generalized sort of indiscernible sequence, point out its origins and uses in the literature, provide an exposition of some of the pre-existing material and present some new results. In this section we explain a little of the history of this concept.

In our discussion of L' -generalized indiscernibles we will have in mind an “index” model I , some L' -structure, and a “target model” M , some L -structure. We assume that L' contains a binary relation $\{<\}$ that is always interpreted as a linear order in I .

It is good to keep in mind a dichotomy: on one side, the index model of the index theory in the index language (often denoted with primes); on the other side, the target model of the target theory in the target language.



The elements of a generalized indiscernible, $\bar{a}_i := f(i)$, are given by a function (notation typically suppressed) $f : I \rightarrow M^n$, for some fixed n . So as not to overburden the notation, in this chapter we make the convention that elements a_i in the target model M refer to tuples, and variables x in the target language L refer to finite sequences of variables, unless we expressly say otherwise.

The following is an informal definition for generalized indiscernibles:

Definition 2.1.1. (informal) We say that $(a_i : i \in I)$ in M is *generalized-indiscernible* if for all finite tuples \bar{i}, \bar{j} from I , whenever $\text{qftp}^{L(I)}(\bar{i}; I) = \text{qftp}^{L(I)}(\bar{j}; I)$ we have that $\text{tp}^{L(M)}(\bar{a}_{\bar{i}}; M) = \text{tp}^{L(M)}(\bar{a}_{\bar{j}}; M)$.

The above definition is in analogy to the case of indiscernible sequences, which are a special case of generalized indiscernibles if we consider the indexing linear order to be a structure in its own right.

The notion of a generalized sort of indiscernible sequence was first presented in Ch. VII of [She78] in the discussion of “indiscernible indexed sets.” In definition 2.4, Shelah defines a set indexed by a model I , $\{b_s : s \in I\}$ to be “ Δ - n -indiscernible” if

1. for all s, t of the same quantifier-free type, $l(b_s) = l(b_t)$
2. whenever \bar{s}, \bar{t} are two same-length tuples from I having the same quantifier-free type, then $\bar{b}_{\bar{s}}, \bar{b}_{\bar{t}}$ agree on all the formulas from Δ .

For the purposes outlined in Ch. VII, Shelah assumes that every complete quantifier-free type realized in the “index-model”, I , is equivalent to a quantifier-free formula. In our applications, we will typically require that the language of I be finite and relational, and so the latter condition is always true (in fact, all of our models are quantifier-free \aleph_0 -categorical.) It is an interesting problem to investigate what effect these assumptions have on the mechanics of the generalized indiscernible. It seems probable that many results go through for infinite relational languages, but the addition of function symbols creates problems unless there is a finite bound in I on the model theoretic closure of a finite set of parameters. We hope to return to these questions in a later work.

Two important concepts are introduced early in Ch. VII, §2. Fix some language L' . Fixing some representative structure I , the class $K(I)$ is defined as all $L(I)$ -structures J such that any quantifier-free type realized in J is realized somewhere in I , itself. In particular, Shelah identifies ω with the usual ordering as the appropriate representative for the class of linear orders. The idea is that indiscernibles indexed by J can be obtained from indiscernibles indexed by I in such a way that the type structure in the target model is preserved locally.

Choosing T' to be a universal theory in a finite relational language has the advantage that for $I \models T'$, restrictions of the language of I to finite subsets $A \subseteq I$ are substructures that model T' . This is particularly useful when we need to pass from an I -indexed set to a J -indexed indiscernible, and we would like for J to inherit the theory T' . Our methods require us to pass through the model theoretic closure of finite subsets of I , in an application of a Ramsey-type theorem: much the same way indiscernible sequences are obtained from sets of parameters. That complete quantifier-free types are equivalent to formulas has an added advantage in bringing much of the argument for Theorem 3.3.8 into the definable realm.

In this paper, we choose to isolate this property of the “representative” I of $K(I)$ as follows: we fix the theory T' in the finite relational language L' , and let $K(I)$ be all models

of T' . Then for I a quantifier-free weakly-saturated model of T' , the identical desirable relationship of I to $K(I)$ is obtained. Our lemma about “stretching” the generalized indiscernible is essentially an explication of a case of Lemma 2.2 in [She78].

Lemma 2.1.2 (first half 2.2, [She78]). *If $I_1 \in K(I)$ and $\{b_s : s \in I\}$ is an indiscernible indexed set in a model $M \models T_1$, then there is an indiscernible indexed set $\{a_i : i \in I_1\}$ in a model $M_1 \models T_1$ such that if $s \in I$, $t \in I_1$ $qftp^{L(I)}(s; I) = qftp^{L(I)}(t; I_1)$, then $tp^{L(M)}(b_s; M) = tp^{L(M)}(a_i; M_1)$.*

Another important concept from Ch. VII is the question of whether nontrivial generalized indiscernibles exist, indexed by a particular class $K(I)$ of L' -structures. For example, in Lemma 2.3, Shelah is essentially stating that for any target theory T , nontrivial indiscernible sequences (indexed by ω) exist in some model of T . However, this is not so for $K(I)$ where R is the (infinite) random graph in the pure graph language:

Example 2.1.3. Let R be the random graph (without a linear ordering) in the pure graph language $L' := \{R\}$. (L', I) -generalized indiscernibles do not exist within a target model M that is linearly ordered. Any potential I -indiscernible $(a_i : i \in I)$ in M would have the problem that for i and j contained in an edge of I , we would need to have $tp(a_i, a_j) = tp(a_j, a_i)$. However, for any distinct length- n tuples \bar{a}, \bar{b} from M , one of $(\bar{a})_1, (\bar{b})_1$ must be less than the other. Thus there will be some relation in the type of $\langle \bar{a}, \bar{b} \rangle$ that is not in the type of $\langle \bar{b}, \bar{a} \rangle$. Thus $tp(a_i, a_j) = tp(a_j, a_i)$ yields a contradiction.

We start by presenting our notation and pointing out some key observations about generalized indiscernibles. In this chapter we also introduce definitions around the modeling property for generalized indiscernibles, with a view towards introducing the concepts we will need for our proofs of the main results Theorem 3.3.8 in Ch. 3 and Theorem 5.2.10 in Ch. 5.

2.2 Basic definitions for generalized indiscernibles

We state our notation again for clarity, along with some conventions:

Definition 2.2.1 (generalized indiscernibles). Fix an index language L' , a target language L and a sublanguage $L'' \subseteq L'$. Suppose we are given parameters $a_i := f(i)$ where $f : I \rightarrow M^n$ for some fixed n .

We say that $(a_i : i \in I)$ is (L'', I) -generalized-indiscernible in M if for all finite n ,

$$qftp^{L''}(i_1, \dots, i_n; I) = qftp^{L''}(j_1, \dots, j_n; I) \Rightarrow tp^L(a_{i_1}, \dots, a_{i_n}; M) = tp^L(a_{j_1}, \dots, a_{j_n}; M)$$


In the case that L' -structure I and L -structure M are clear from context, we may just say that $(a_i : i \in I)$ is L'' -generalized indiscernible.

By an I -indexed indiscernible we mean an $(L(I), I)$ -generalized indiscernible sequence $(a_i : i \in I)$ in some target structure M .

If we say that a sequence $(a_i : i \in I)$ is *generalized indiscernible*, we mean that it is an I -indexed indiscernible.

We will always assume that our generalized indiscernibles are *nontrivial*, i.e. that for $i \neq j$, that $a_i \neq a_j$.

Here is a visualization of an L' -generalized indiscernible $(a_i : i \in I)$ in M :



On the left are tuples of indices \bar{i} from the index I . The arrows indicate the complete type $q_i(\bar{x})$ of the tuples $a_{\bar{i}}$ in the target, M . For the purposes of visualization, M is imagined as a graph that eliminates quantifiers. In this example, if the index tuple satisfies $v_1 < v_2 < v_3$, then it maps to a target type $q_1(\bar{x})$; if the index tuple satisfies $v_2 < v_1 < v_3$, then it maps to a target type $q_2(\bar{x})$.

We will also want some notation for how two L' -generalized indiscernibles relate. Given I -indexed indiscernible $(a_i : i \in I)$, we want to know the type of $\langle a_{i_1}, \dots, a_{i_n} \rangle$ in the target, not just for every n , but for every complete quantifier-free (L', n) -type $\eta(x_1, \dots, x_n)$ such that $\langle i_1, \dots, i_n \rangle \models \eta$. The following is notation inspired by the presentation in [Mar02].

Definition 2.2.2. Fix a language L' in finite relational signature $\rho' \supseteq \{<\}$ and fix a sub-language L'' given by a ρ'' such that $\{<\} \subseteq \rho'' \subseteq \rho'$.

Give an L' -structure I linearly ordered by $<$, (L', I) -generalized indiscernible $(a_i : i \in I)$, and complete quantifier-free L'' -type $\eta(v_1, \dots, v_n)$ realized in I and consistent with $v_1 < \dots, < v_n$, define:

$$\text{tp}_{L''}(I) := \{p^\eta(\langle a_i : i \in I \rangle) : \eta \text{ is an } n\text{-ary complete quantifier-free}$$

$$L''\text{-type realized in } I \text{ and consistent with } v_1 < \dots, < v_n, n < \omega\}$$

$$\text{where } p^\eta(\langle a_i : i \in I \rangle) = \{\psi(x_1, \dots, x_n) : \text{there exists } i_1 < \dots < i_n \text{ from } I$$

such that $(i_1, \dots, i_n) \models \eta(v_1, \dots, v_n)$ and $\langle a_{i_1}, \dots, a_{i_n} \rangle \models \psi(x_1, \dots, x_n)$.

By generalized-indiscernibility, the identical set is defined by replacing “there exists” by “for all”.

By $\text{tp}(\langle a_i : i \in I \rangle)$ we denote $\text{tp}_{L(I)}(\langle a_i : i \in I \rangle)$.

Definition 2.2.3. (based on) Fix I an L' -structure and some sublanguage $L'' \subseteq L'$. For a given set of parameters $(a_i : i \in I)$ in an L -structure $M \models T$ and an L' -structure J , we say that the L' -indiscernible sequence $(b_i : i \in J)$ in $M_1 \models T$ is L'' -based on the a_i if for any finite set of L -formulas, Σ and for any finite tuple from I , s_1, \dots, s_n there exists an L'' -isomorphic tuple t_1, \dots, t_n in J such that $\text{tp}^\Sigma(\bar{b}_s; M_1) = \text{tp}^\Sigma(\bar{a}_t; M)$.

If we say that an I -indexed indiscernible $(b_i : i \in I)$ is based on some parameters, we mean that it is $L(I)$ -based on them.

Remark 2.2.4. Note that for two L' -generalized indiscernibles $(a_i : i \in I)$ and $(b_j : j \in J)$, that the b_j are based on the a_i just in case

1. every complete quantifier-free L'' -type realized in J is realized in I , and
2. for every complete quantifier-free L'' -type η realized in J , $p^\eta(\langle b_j : j \in J \rangle) = p^\eta(\langle a_i : i \in I \rangle)$.

For applications in the case that $L'' := L'$, it is important to note that the condition “every complete quantifier-free L' -type realized in J is realized in I ” is equivalent to the condition $\text{age}(J) \subseteq \text{age}(I)$.

2.3 Mechanics

In the case of indiscernible sequences, or I -indexed indiscernibles where I is a pure linear order, we often wish to “stretch” the indiscernible. To “stretch” a given indiscernible sequence $(a_i : i \in I)$ is, informally, to extend it to an indiscernible sequence $(b_j : j \in J)$ for some linear order J (perhaps with $I \subseteq J$ or $I \preceq J$.) In “stretching” the indiscernible, we pass to the J -indexed indiscernible in such a way that $p^\eta(\langle b_j : j \in J \rangle) = p^\eta(\langle a_i : i \in I \rangle)$. In other words, we seek to find a J -indexed indiscernible that is based on the a_i . As in the case of linear orders as index structures, if we require the parameters $(b_j : j \in J)$ to satisfy a type that say the b_j are based on the L' -generalized indiscernible a_i , this is sufficient to make the b_j L' -generalized indiscernible. We will use this fact in the next lemma.

There is an extra condition in the next lemma that is not explicitly seen in the indiscernible sequence case. An indiscernible sequence $(a_i : i \in I)$ can be “stretched” to a new index $(J, <)$ provided that I is infinite (e.g., given $I = (\omega, \in)$ we can “stretch” to $J = (\kappa, \in)$ for infinite cardinal κ .) We require that I be not only infinite, but that $\text{age}(I) \supseteq \text{age}(J)$. We should keep in mind that any infinite linear orders I, J automatically have the same age. Thus we must state this condition more explicitly in the general case.

Lemma 2.3.1 (Stretching the (L', I) -generalized indiscernible). *Suppose I is an L' -structure for finite relational L' . Suppose we have an (L', I) -generalized indiscernible $(a_i : i \in I)$ in some L -structure M .*

For any L' -structure J with $\text{age}(J) \subseteq \text{age}(I)$, we can find (L', J) -indiscernible $(b_j : j \in J)$ in some structure $M_1 \models \text{Th}(M)$ that is L' -based on the a_i .

Proof. By Remark 2.2.4, we need only show that we can find an L' -indiscernible $(b_j : j \in J)$ such that for every complete quantifier-free L' -type η realized in J , $p^\eta(\langle b_j : j \in J \rangle) = p^\eta(\langle a_i : i \in I \rangle)$.

It suffices to realize the following type Σ in new constants $\{c_j : j \in J\}$. This will yield an (L', J) -indiscernible with $p^\eta(\langle b_j : j \in J \rangle) = p^\eta(\langle a_i : i \in I \rangle)$ for all complete quantifier-free L' -types η realized in J . It is convenient to enumerate the c_j as $\{c_{j(k)} : k < \alpha\}$ for some ordinal α .

$$\Sigma := \text{Th}(M) \cup \{\theta(c_{j(k_1)}, \dots, c_{j(k_m)}) : m < \omega; \text{distinct } k_1, \dots, k_m \in \alpha;$$

$$\langle j_{k(1)}, \dots, j_{k(m)} \rangle \models \eta(v_1, \dots, v_m); \text{ and } \theta(x_1, \dots, x_m) \in p^\eta(I)\}$$

It remains to show the above type is finitely satisfiable. Choose any finite subtype of Σ ; we may assume this is of the form $\text{Th}(M) \cup F$. Then there is some finite list of constants and L -formulas occurring in F . We may assume that these are $c_{j(0)}, \dots, c_{j(k-1)}$ and $\theta_0, \dots, \theta_{l-1}$.

Let $\text{qftp}^{L'}(j(0), \dots, j(k-1); J) =: \eta$. This type is realized by some sequence $\langle i(0), \dots, i(k-1) \rangle$ from I , as $\text{age}(J) \subseteq \text{age}(I)$. Assign the $a_{i(s)}$ to be the interpretations of the $c_{j(s)}$ in M . Then the expansion of M to the $\{c_{j(0)}, \dots, c_{j(k-1)}\}$ satisfies $\text{Th}(M) \cup F$. \square

Corollary 2.3.2. *If I is a quantifier-free weakly-saturated model of T' , then given some I -indexed indiscernible in M and $J \models T'$, we can find a J -indexed indiscernible in some elementary extension of M based on the I -indexed indiscernible.*

In particular, for I a quantifier-free weakly-saturated ordered graph, for any ordered graph J we can find (L_g, J) -indiscernibles based on (“stretching”) the (L_g, I) -generalized indiscernible.

Remark 2.3.3. We will use this in the case that $J := \mathcal{R}$, the random ordered graph.

In the next chapter, we define the “modeling property” and Ramsey classes of structures and pointing out some their key properties.

Chapter 3

The modeling property

3.1 Definitions

By the next definition, we wish to separate the question of when there exists an I -indiscernible in some model of a theory T (in [She78] this is referred to as the existence of an elementary type that is “proper” for (I, T)) from the question of when an I -indiscernible both exists in some model of T *and* we have quite a bit of control over what its finite subtypes end up being in this model.

Definition 3.1.1. I -indexed indiscernibles have the *modeling property* if given any parameters $(a_i : i \in I)$ in any model M , there exist L' -generalized indiscernible $(b_i : i \in I)$ in some model $M_1 \models \text{Th}(M)$ based on the a_i .

The definition of the *modeling property* helps us isolate a notion of strength for a generalized indiscernible. This is important because we have the following phenomenon.

Proposition 3.1.2. *Given two signatures $\rho_1 \subseteq \rho_2$ any L_{ρ_1} -generalized indiscernible is automatically L_{ρ_2} -generalized indiscernible.*

Proof. This is straightforward from the definitions. □

In particular, for a signature ρ_2 containing a relation $<$ that linearly orders I , we have the above result for the case $\rho_1 := \{<\}, \rho_2$. Thus, it is always possible to use Ramsey’s theorem to obtain (L', I) -generalized indiscernibles for $L' = L_{\rho_2}$. However, this does not pin down the strength of the indiscernible. What we want to know is, given an L' -structure I , can we find an (L', I) -generalized indiscernible that distinguishes between elements of the *finest* partition on I by quantifier-free L' -types.

We will also want a slight generalization of the modeling property:

Definition 3.1.3. Fix a finite relational language L' . Given an age \mathbf{K} of L' -structures we say that L' -generalized indiscernibles have the *modeling property for \mathbf{K}* if given any set of

parameters $(a_i : i \in I)$ in a model M indexed by an L' -structure I such that $\text{age}(I) = \mathbf{K}$, we may find L' -generalized indiscernible $(b_j : j \in J)$ in a model $M_1 \models \text{Th}(M)$ based on the a_i , indexed by L' -structure J such that $\text{age}(J) = \mathbf{K}$.

Here we give an application of the modeling property in previous literature, and one which illustrates the utility of the distinctions we are making:

Example 3.1.4. In [DS04] 2-full binary tree indiscernibles (2-fbtis) are shown to have the modeling property. 2-full binary tree indiscernibles are I -indiscernibles where $|I| := \omega^{>2}$. The language of I , L_T has a constant for the root, and a relation that holds on η, ν in $\omega^{>2}$ whenever η is an initial segment of ν . This yields the result that given a tree of parameters $(a_i : i \in I)$ witnessing SOP_2 in some target model, M , there is a 2-full binary tree indiscernible witnessing the same. A set of parameters indexed by $|I|$ witnesses SOP_2 if there is a formula $\varphi(x, y)$ in the target language, such that $\{\varphi(x; a_\eta), \varphi(x; a_\nu)\}$ is consistent in M just in case η, ν are comparable in the partial tree order and the set of restrictions of φ along any path is consistent.

Note that in the case of finding a 2-fbti witness to SOP_2 , it would not be enough to use the modeling property for *order* indiscernibles (for example using the lexicographic order on I .) If we apply the modeling property for order indiscernibles to a set $(a_\eta : \eta \in \omega^{>2})$, this merely guarantees that for every formula φ and increasing pairs η', ν' in the lexicographic order, the φ -type of $(b_{\eta'}, b_{\nu'})$ matches the φ -type of (a_η, a_ν) for some (η, ν) with the same *order*-type as (η', ν') in I . Supposing that $(a_\eta : \eta \in \omega^{>2})$ witness SOP_2 by way of $\varphi(x, y)$, our resulting indiscernible tree $(b_\eta : \eta \in \omega^{>2})$ will either have all pairs $\{\varphi(x; a_\eta), \varphi(x; a_\nu)\}$ consistent for $\eta < \nu$, or it will have all such pairs inconsistent, neither of which cases witness SOP_2 by φ .

The previous example will be discussed at greater length in Ch. 4.

3.2 Ramsey classes

3.2.1 Preliminaries

At this juncture, it is helpful to introduce the Nešetřil-Rödl notion of “Ramsey class”¹ and related definitions. We will be using a partition property of certain classes of finite structures proved simultaneously in [NR77, AH78]. In this chapter, L' will be a finite, relational language containing a binary relation $\{<\}$ for order. When we wish to talk about an arbitrary first-order language, we use the notation \mathcal{L} . Let V be the class of all finite L' -structures. We will be considering the question of which subclasses U of V are Ramsey classes. We let L', V, U play the above roles throughout this chapter. In the case that U

¹This notation is from [Neš05], though the same notion is originally referred to as “partition category” in [NR77]

is the age of our index structure, I , we can conclude that I -indiscernibles have certain nice properties, in particular, the modeling property.

The property of being a Ramsey class can be stated more generally for certain categories, and we will give this presentation in subsection 3.2.3. In our applications elsewhere than in subsection 3.2.3, for a subclass U of V to be a Ramsey class means the following:

Definition 3.2.1. (informal) U is a Ramsey class if for any structures A, B in U and integer $k > 0$, there is a structure C in U so that for any coloring of the substructures of C isomorphic to A with k colors, there is a $B' \subseteq C$ isomorphic to B , so that all substructures of B' isomorphic to A are colored the same color.

More exact definitions are presented in subsection 3.2.2.

The modeling property is a generalization of a property that indiscernible sequences naturally inherit by way of the Ramsey theorem for finite sequences. This is the property that for a given infinite sequence of parameters in some model (indexed by a linear order that is not necessarily definable) we can find an indiscernible sequence that is “finitely modeled” on this given sequence: i.e., for any finite increasing sequence from the indiscernible and any finite fragment of the language Δ , this sequence has the same Δ -type as some increasing sequence from the original set of parameters. We will give a generalization of this notion for generalized indiscernibles in section 3.3.

For the application we are most interested in – characterizing NIP theories in Chapter 5 – we happen to know that our generalized indiscernibles have the modeling property, due to special properties of the index model, a quantifier-free weakly-saturated ordered graph.

3.2.2 Definition of a Ramsey Class

Define an A -substructure of C to be a substructure $A' \subseteq C$ isomorphic to A where we do not reference a particular enumeration of A' . More formally,

Definition 3.2.2. An A -substructure of C is an equivalence class $[f]_E$ of L' -embeddings $f : A \rightarrow C$, under the equivalence relation fEg if $f = gh$ for some $h : A \rightarrow A$ a L' -automorphism.

We refer to the set of A -substructures of C as $\binom{C}{A}$.

Definition 3.2.3. For an integer $k > 0$, by a k -coloring of $\binom{C}{A}$ we mean a function $f : \binom{C}{A} \rightarrow \eta$, where η is some set of size k , typically $\eta := \{1, \dots, k\}$.

Remark 3.2.4. In the case that all structures in U are linearly ordered by a relation $<$ in L' , the equivalence classes in Definition 3.2.2 have size 1, since the only L' -automorphism of A is the identity map. In this case, a copy of A , $A' \subseteq B$ can be identified with the unique embedding that maps A into B . Thus, coloring A -substructures of C is equivalent to coloring embeddings of A in C .

Here we state the definition of the extension of the Ramsey arrow notation to L' -structures.

Definition 3.2.5. Let A, B, C be objects in U and k some positive integer.

1. By

$$C \rightarrow (B)_k^A$$

we mean that for all k -colorings of $\binom{C}{A}$, there is a $B' \subseteq C$, where B' is U -isomorphic to B and all A -substructures of B' are colored the same color, under the restriction of the k -coloring to B' : i.e. $f \upharpoonright \binom{B'}{A}$ is a constant function.

2. If, for a particular coloring $f : \binom{C}{A} \rightarrow k$ we have a $B' \subseteq C$ such that $f \upharpoonright \binom{B'}{A}$ is a constant function, we say that B' is *homogeneous for this coloring (homogeneous for f)*.

We may say that C is *Ramsey for (B, A, k)* , reading $(B)_k^A$ clockwise, from the left.

Now, we are ready to state the definition of *Ramsey class* as we will use it:

Definition 3.2.6. Let U be a class of L' -structures. U is a *Ramsey class* if for any $A, B \in U$ and positive integer k , there is a C in U such that $C \rightarrow (B)_k^A$.

We need a few more definitions before we present some known results about Ramsey classes:

Definition 3.2.7. Given an n -ary relation $R(x_1, \dots, x_n)$ in \mathcal{L} and an \mathcal{L} -structure M we say that

1. R is *antireflexive on M* if $\neg R(a_1, \dots, a_n)$ holds in M whenever the a_k are from M and $a_i = a_j$ for some $1 \leq i \neq j \leq n$.
2. R is *symmetric on M* if $R(a_1, \dots, a_n)$ implies $R(a_{\sigma(1)}, \dots, a_{\sigma(n)})$ whenever the a_k are from M and σ is a permutation on $\{1, \dots, n\}$.
3. An \mathcal{L} -structure A is an *(\mathcal{L} -)hypergraph structure* if all relations from \mathcal{L} are symmetric and antireflexive on A .

The following is a particular case of the Nešetřil-Rödl theorem:

Theorem 3.2.8. ([NR77]) *If L^* is a finite relational language in a signature ρ^* containing a binary relation $\{<\}$ for order, and V^* is the class of all finite L^* -structure whose reduct to $L_{(\rho^* \setminus \{<\})}$ is a hypergraph structure, then V^* is a Ramsey Class.*

Corollary 3.2.9. *For \mathcal{K} the set of all isomorphism types of finite ordered graphs (symmetric, with no loops), \mathcal{K} is a Ramsey class.*

The conditions of Theorem 3.2.8 very conveniently fit the particular case of graphs in Corollary 3.2.9 that we will use to characterize NIP theories in Chapter 5. In the next section, we will go over some facts about Ramsey classes that apply to more general situations.

3.2.3 Results around Ramsey classes

It is possible to state the notion of Ramsey class more generally for certain categories. In this subsection I will describe how this is done in the Nešetřil and Rödl papers, making primary use of the more updated notation and some of the newer results from [Neš05].

In this section we will be using a finite relational signature Δ that will not include the order relation $<$. Thus we make the convention that

Definition 3.2.10. For Δ a finite relational signature, let $L[\Delta]$ be the language in signature $\Delta \cup \{<\}$.

Notation and definitions

We work primarily in full subcategories of either $\text{Rel}(\Delta)$ or $\text{Set}(\Delta)$.² These latter categories are defined as follows.

Definition 3.2.11. Fix Δ some finite relational signature.

Define $\text{Set}(\Delta)$ to be the category whose objects are all tuples

$$A = (A(X), \langle R_i^A \rangle_{i \in I})$$

where $A(X)$ is some finite set and $R_i^A \subseteq [A(X)]^{\delta_i}$

Morphisms are L_Δ -homomorphisms between structures, i.e. functions $f : A \rightarrow B$, such that

$$1. (a_j : j \leq \delta_i) \in R_i \Rightarrow (f(a_j) : j \leq \delta_i) \in S_i$$

Remark 3.2.12. Note that the objects of $\text{Set}(\Delta)$ are in bijection with the finite L_Δ -hypergraph structures.

It is important not to confuse $\text{Rel}(\Delta)$ for the category of finite $L[\Delta]$ -structures whose reduct to L_Δ is a hypergraph structure, with $L[\Delta]$ -homomorphisms as maps.

We define an associated category for any full subcategory $\mathbf{K} \subseteq \text{Set}(\Delta)$.

Definition 3.2.13. Let \mathbf{K} be any full subcategory of $\text{Set}(\Delta)$. By $\overrightarrow{\mathbf{K}}$ we mean the category whose objects are

$$\{\langle A, \leq_A \rangle : \text{for any order } \leq_A \text{ on } A \in \mathbf{K}\}$$

and whose morphisms $f : \langle A, \leq_A \rangle \rightarrow \langle B, \leq_B \rangle$ are \mathbf{K} -morphisms $f : A \rightarrow B$ such that whenever $a_1, a_2 \in A$:

$$a_1 \leq_A a_2 \Rightarrow f(a_1) \leq_B f(a_2).$$

We point out an easy but useful equivalence:

²These are $\text{OSoc}(\Delta)$ and $\text{Soc}(\Delta)$, respectively in [NR77]

Claim 3.2.14. *Fix a finite relational language L^- . Let $L := L^- \cup \{<\}$ for a new binary relation symbol, $<$. Let U be the category of finite L^- -structures with L^- -homomorphisms and let $\mathcal{M}(U)$ be all the expansions of structures in U to L -structures linearly ordered by $<$, with L -homomorphisms. The following are equivalent:*

1. \vec{U} is a Ramsey category
2. $\mathcal{M}(U)$ is a Ramsey category
3. the class of structures $\mathbf{Ob}(\mathcal{M}(U))$ forms a Ramsey class.

Remark 3.2.15. The above equivalence can be useful when investigating which (full) subcategories \mathbf{K} of $\text{Set}(\Delta)$ give Ramsey categories, $\vec{\mathbf{K}}$, and then translating these results to the model-theoretic context.

One obvious application of the above remark is to the instances in [Neš05] where a class \mathbf{K} is stated to be Ramsey “with embeddings as subobjects” – as, by the latter statement, the author means that $\vec{\mathbf{K}}$ is a Ramsey category, and so we may conclude that $\mathcal{M}(\mathbf{K})$ is a Ramsey class.

Known results

We address two different directions in which Theorem 3.2.8 might be generalized, and the problems with each. In this section, Δ is finite and relational, as usual.

(I) DELETING OBJECTS:

First, some definitions. In the following “Forb” stands for “forbidden”.

Definition 3.2.16. For \mathcal{F} a subset of objects from $\text{Set}(\Delta)$, by $\text{Forb}_\Delta(\mathcal{F})$ we mean the subclass:

$$\text{Forb}_\Delta(\mathcal{F}) = \{B \in \mathbf{K} : \nexists \text{ a monomorphism } f : A \rightarrow B, \forall A \in \mathcal{F} \}$$

In $\text{Set}(\Delta)$ a monomorphism is equivalent to an injective L_Δ -homomorphism.

Definition 3.2.17. A structure in $\text{Rel}(\Delta)$ is *irreducible* if any two points are contained in a hyperedge.

Equivalently, a structure is irreducible if it cannot be expressed as the *free* amalgam of distinct structures in $\text{Rel}(\Delta)$.

We have the following results

Theorem 3.2.18. *(4.2 in [Neš05]) If \mathbf{K} is closed under substructures and isomorphism and has JEP and $\vec{\mathbf{K}}$ is a Ramsey class³, then \mathbf{K} has the amalgamation property (in fact, \mathbf{K} has strong amalgamation)*

³equivalently, “ \mathbf{K} is Ramsey with embeddings as subobjects”

Theorem 3.2.19. (by Prop 5.2, Theorem 5.3 of [Neš05]) *If \mathbf{K} is a full subcategory of $\text{Set}(\Delta)$ closed under monomorphisms, having JEP and strong amalgamation, then $\vec{\mathbf{K}}$ is a Ramsey Class.*

(Moreover, \mathbf{K} can be expressed as $\text{Forb}_\Delta(\mathcal{F})$ for some class of irreducible objects \mathcal{F} of $\text{Set}(\Delta)$.)

So in the case of subclasses of $\text{Set}(\Delta)$ closed under monomorphism, isomorphism and with strong amalgamation, if such a subclass \mathbf{K} additionally has the joint embedding property, then it *must* arise as $\text{Forb}_\Delta(\mathcal{F})$ for some class of irreducible objects \mathcal{F} of $\text{Set}(\Delta)$. (Theorem 5.3 of [Neš05])

Let us explicate the conditions of Theorem 3.2.19 in the case of ordered graphs (symmetric, with no loops – $\Delta_g = \{R\}$.) For a full subcategory $\mathbf{K} \subseteq \text{Set}(\Delta_g)$, to be closed under monomorphisms is to contain all non-induced subgraphs. (However, this is not to say that we in any way relax the condition that each $A \in \vec{\mathbf{K}}$ be linearly ordered. It is important to remember that in [Neš05], the linear order is not part of the language.)

In our notation, letting \mathcal{K} the class of all finite ordered graphs (symmetric, no loops):

Corollary 3.2.20. *Let U be any subclass of \mathcal{K} closed under not-necessarily induced subgraphs. If U is closed under isomorphism, has the joint embedding property and strong amalgams, then \vec{U} is a Ramsey class.*

Remark 3.2.21. In other words, given ultrahomogeneous ordered graph I with no algebraicity and closed under not-necessarily induced subgraphs, the expansion I^* to a linearly ordered graph has the modeling property.

(II) NO LINEAR ORDER:

Let \mathbf{K} be a full subcategory of $\text{Set}(\Delta)$. We know under certain hypotheses (those in Corollary 3.2.20) that $\vec{\mathbf{K}}$ is a Ramsey class. What if we do not wish to impose a linear order on the structures in \mathbf{K} ? For example, we may wish for colorings of size- n structures $A \in \mathbf{K}$ to come from n -types in a target model M which have some amount of symmetry, i.e. types $p(x_1, \dots, x_n)$ such that $\langle a_1, \dots, a_n \rangle \models p$ and for all $\sigma \in G \subseteq S_n$:

$$\langle a_{\sigma(1)}, \dots, a_{\sigma(n)} \rangle \models p.$$


In a future work, we hope to pursue Nešetřil’s notion of *admissible orderings* from [Neš05] to find enrichments \mathbf{K}^* of \mathbf{K} by definable relations such that \mathbf{K}^* is a Ramsey class, though \mathbf{K} may not be.

Suffice it to say that if we impose no linear ordering or extra structure at all, then we have the A -partition property for only a few A in $\text{Set}(\Delta)$:

Theorem 3.2.22 (Theorem 1, §6 [NR77]). *For \mathcal{F} a class of irreducible L_Δ -structures, $\mathbf{K} = \text{Forb}_\Delta(\mathcal{F})$, \mathbf{K} has the A -partition property if and only if A is totally symmetric.*


In the case of failure of total symmetry, we have the following counterexample:

Example 3.2.23. Consider the case where A is the wedge:




Clearly A is not a totally symmetric graph since there is a permutation of the vertices of A that fails to be an automorphism. We wish to show that the A -partition property fails for $\mathbf{K} =$ the class of all finite symmetric graphs with no loops (*without* order.) That is to say, we wish to find a $B \in \mathbf{K}$ and a finite integer k (we will choose $k = 2$) such that for any choice of $C \in \mathbf{K}$, C fails to be Ramsey for (B, A, k) , i.e., there is a k -coloring $\sigma : \binom{C}{A} \rightarrow k$ such that for no $B' \subseteq C$ isomorphic to B do we have that all A -subobjects of B' are monocolored under σ .

We will choose B to be the triangle:






Now fix $C \in \mathbf{K}$. We build a coloring of the A -subobjects of C by way of imposing an arbitrary linear ordering on the vertices of C . Name our two colors be "red" and "blue". We color any copy of A in C red if, counting the vertices in increasing order, we are presented with an edge, an edge, and then a nonedge. We color copies of A in C blue otherwise. Note that A does have a non-identity automorphism, τ that switches the two poles away from the center:



Thus, according to the linear ordering, there are not $3! = 6$ ways to color copies of A in C but only $6/2 = 3$ (the case where the middle vertex is least, greatest, or in between relative to the outer vertices. In our case, two colors will be enough, and this is a proper coloring as it is invariant under automorphisms of A).

To see that this is indeed a counterexample, take any copy B' of the triangle B in C . There will be both a red and a blue copy of A in B' : enumerate the vertices of B' by





Discussion 3.2.24. This is the problem when we lie in the intermediary zone where some permutation of an object A other than the identity yields an automorphism, and yet not all of them do.

3.3 Ramsey Classes and the Modeling Property

For the moment, let's forget the notation we introduced to talk about Ramsey categories and focus on Ramsey classes. In the next two lemmas, we will be developing a characterization of the relationship between Ramsey classes and generalized indiscernibles with the modeling property.

Fix L' throughout this subsection, a first-order language in finite, relational signature ρ' .

Notation 3.3.1. Though formally our variables in any first language are among x_1, x_2, \dots , we will often use s_j, t_j to represent a variable x_{i_j} in the case that x_{i_j} is in the interpretation of Q , and y_j in the case that x_{i_j} is in the interpretation of P , for greater ease in reading.

There are many ways to set up notation to represent generalized indiscernibles in an L -structure M indexed by an L' -structure I . We prefer notation that illustrates to the greatest extent which conditions are first-order definable in which models. Rather than using a many-sorted language, we introduce new predicates for the “index-sort” and “target-sort”, Q, P , respectively.

Definition 3.3.2. Given a finite relational language L' with signature ρ' , and a language L with signature ρ

1. let $\mathcal{L}(L', L)$ be the language with signature $\rho' \cup \rho \cup \{P, Q, f\}$ where P, Q are new unary predicate symbols and f is a new unary function symbol.
2. We will use as our base theory T_0 which will be defined as
 - (a) everything is in P or Q but nothing is in both P and Q
 - (b) f restricts to an injective function from P into Q

For the sake of completeness, we show this theory is universal:

$$\forall x[(P(x) \vee Q(x)) \wedge \neg(P(x) \wedge Q(x))] \wedge \forall x \forall y(x \neq y \rightarrow f(x) \neq f(y))$$

$$\forall x(Q(x) \rightarrow P(f(x)))$$

(we don't care what f does on the complement of Q .)

Remark 3.3.3. In the case that we wish to model a $Q(I)$ -indiscernible $(\bar{a}_j : j \in I)$ whose tuples have length k ($l(\bar{a}_j) = k$) it is easy to require that $P \subseteq M^k$ for some power of the universe.

As usual, we will assume the tuples are of length 1 for notational simplicity.

Given an L' -structure I , an L -structure M and a sequence of parameters in M indexed by I , $(a_i : i \in I)$, there is a canonical way to obtain an $\mathcal{L}(L', L)$ -structure \mathfrak{C} such that

1. $|Q^{\mathfrak{C}}| = |I|$ and $Q^{\mathfrak{C}} \cong_{id} I$ with the induced language
2. $|P^{\mathfrak{C}}| = |M|$ and $P^{\mathfrak{C}} \cong_{id} M$ with the induced language
3. $a_i = f^{\mathfrak{C}}(i)$

Definition 3.3.4. By $\mathfrak{C}(I, M, \langle a_i : i \in I \rangle)$ we mean the $\mathcal{L}(L', L)$ -structure \mathfrak{C} such that

1. $\mathfrak{C} \models T_0$
2. $|Q^{\mathfrak{C}}| = |I|$ and $Q^{\mathfrak{C}} \cong_{id}^{L'} I$
3. $|P^{\mathfrak{C}}| = |M|$ and $P^{\mathfrak{C}} \cong_{id}^L M$
4. $a_i = f^{\mathfrak{C}}(i)$

By $I(\mathfrak{C})$ we will refer to the L' -reduct of \mathfrak{C} to universe $Q^{\mathfrak{C}}$; by $M(\mathfrak{C})$ we will refer to the L -reduct of \mathfrak{C} to universe $P^{\mathfrak{C}}$. Note the slight difference between referring to the structure-properties of $I(\mathfrak{C})$ versus merely the set, $Q^{\mathfrak{C}}$.

It will be useful to have the following notion.

Definition 3.3.5. Fix an $\mathcal{L}(L', L)$ -structure $\mathcal{M} \models T_0$. For an L' -formula θ , we define θ^Q , the *relativization of θ to Q* recursively, as follows:

- For relations $R(x_1, \dots, x_m) = \theta$, $\theta^Q(x_1, \dots, x_m) = R(x_1, \dots, x_m)$.
- For formulas $\varphi(\bar{x})$ and $\psi(\bar{x})$
 1. for $\theta(\bar{x}) = \varphi(\bar{x}) \wedge \psi(\bar{x})$, $\theta^Q(\bar{x}) = \varphi^Q(\bar{x}) \wedge \psi^Q(\bar{x})$
 2. for $\theta(\bar{x}) = \neg\varphi(\bar{x})$, $\theta^Q(\bar{x}) = \neg(\varphi^Q(\bar{x}))$
- For formulas $\psi(x)$ with free variable x and $\theta = \exists x\psi(x)$,
 $\theta^Q = \exists x(Q(x) \wedge \psi(x))$

These are some definitions that will help us capture the notion that “the b_j are L' -modeled on the a_i ” in a set of sentences:

For a theory T , by $(T)^Q$ we mean $\{\theta^Q : \theta \in T\}$.

Definition 3.3.6. Given an $\mathcal{L}(L', L)$ -structure $\mathcal{N} \models T_0$, $m < \omega$, $q \in Q_m(I(\mathcal{N}))$ and a finite $\Delta(y_1, \dots, y_m) \subseteq L$, define the

(Δ, q) -profile in \mathcal{N}

to be:

$$P_{(\Delta, q)}(\mathcal{N}) = \{p \in S_m^\Delta(\emptyset; M(\mathcal{N})) : \mathcal{N} \models (\exists x_1, \dots, x_m) (\bigwedge_{\theta \in q} \theta(x_1, \dots, x_m) \wedge \bigwedge_{\psi \in p} [\psi(f(x_1), \dots, f(x_m))])\}$$

Definition 3.3.7. Given an $\mathcal{L}(L', L)$ -structure $\mathcal{N} \models T_0$, $m < \omega$, $q \in Q_m(I(\mathcal{N}))$ and a finite $\Delta(y_1, \dots, y_m) \subseteq L$, define the

(Δ, q) -profile in \mathcal{N}

to be:

$$P_{(\Delta, q)}(\mathcal{N}) = \{p \in S_m^\Delta(\emptyset; M(\mathcal{N})) : \mathcal{N} \models (\exists x_1, \dots, x_m) (\bigwedge_{\theta \in q} \theta(x_1, \dots, x_m) \wedge \bigwedge_{\psi \in p} [\psi(f(x_1), \dots, f(x_m))])\}$$

In the next two subsections we prove the following characterization of L' -generalized indiscernibles having the modeling property for \mathcal{U} :

Theorem 3.3.27. *Let L' be a finite relational language containing a binary relation symbol for order, $<$, and let \mathcal{U} be some nonempty collection of finite L' -structures that are linearly ordered by $<$. Let \mathcal{U}^* be the set of isomorphism types from \mathcal{U}^* . Suppose that \mathcal{U}^* has JEP and is closed under isomorphism and substructures.*

\mathcal{U} is a Ramsey Class if and only if L' -generalized indiscernibles have the modeling property for \mathcal{U}^ .*

3.3.1 Characterization: sufficiency

Lemma 3.3.8. *Let L' be a finite relational language containing a binary relation symbol for order, $<$, and let \mathcal{U} be some nonempty collection of finite L' -structures that are linearly ordered by $<$. Let \mathcal{U}^* be the set of isomorphism types from \mathcal{U} . Suppose that \mathcal{U} has JEP and is closed under isomorphism and substructures.*

*If \mathcal{U} is a Ramsey Class, then L' -generalized indiscernibles have the modeling property for \mathcal{U}^**

Proof. Fix L' , \mathcal{U} , \mathcal{U}^* as in the assumptions above. Suppose that \mathcal{U} is a Ramsey Class. Fix L' -structure I with $\text{age}(I) = \mathcal{U}^*$. Recall Definition 3.1.3:

Definition 3.1.3. Fix a finite relational language L' . Given an age \mathbf{K} of L' -structures we say that L' -generalized indiscernibles have the modeling property for \mathbf{K} if given any set of parameters $(a_i : i \in I)$ in a model M indexed by an L' -structure I such that $\text{age}(I) = \mathbf{K}$, we may find L' -generalized indiscernible $(b_j : j \in J)$ in a model $M_1 \models \text{Th}(M)$ based on the a_i , indexed by L' -structure J such that $\text{age}(J) = \mathbf{K}$.

Thus, we fix a set of parameters in an L -structure M that are indexed by I , $(a_i : i \in I)$ and we wish to show that we can find (L', J) -indiscernibles $(b_j : j \in J)$ in some $M_1 \models \text{Th}(M)$ that are based on the a_i , for $\text{age}(J) = \mathcal{U}$.

We will define a new sequence $(b_j : j \in J) := (f^{\mathfrak{C}}(j) : j \in I(\mathfrak{C}))$ for some $\mathcal{L}(L', L)$ -structure $\mathfrak{C} \models T_0$ realizing the following $\mathcal{L}(L', L)$ -theory, S . We will have that $(b_j : j \in J)$ is an L' -generalized indiscernible sequence in L -structure $M(\mathfrak{C}) \models \text{Th}(M)$ L' -based on the a_i , for $\text{age}(J) = \mathcal{U}$. As usual, $\mathfrak{C} = (I(\mathfrak{C}), M(\mathfrak{C}), f^{\mathfrak{C}}, \dots)$ where \dots lists the interpretations of symbols from the signatures of L', L , respectively. In the following presentation, it helps to keep in mind that in the following the s - and t -variables live in Q and the y -variables live in P .

Here is the $\mathcal{L}(L', L)$ -theory S , listed by subtheories $(A)_S - (D)_S$:

$$S = (\text{Th}(M))^P \cup (A)_S \cup (B)_S \cup (C)_S \cup (D)_S :=$$

$$(A)_S := \{$$

$$(\forall s_1, \dots, s_n)(\forall t_1, \dots, t_n)([\bigwedge_{1 \leq i \leq n} Q(s_i) \wedge Q(t_i)] \rightarrow$$

$$\left[\underbrace{\bigwedge_{\theta \in L'_{\text{at}}(x_1, \dots, x_n)} \theta}_{a} [\theta(s_1, \dots, s_n) \leftrightarrow \theta(t_1, \dots, t_n)] \rightarrow [\varphi(f(s_1), \dots, f(s_n)) \leftrightarrow \varphi(f(t_1), \dots, f(t_n))] \right])$$

$$: n < \omega, \varphi(y_1, \dots, y_n) \in L(y_1, \dots, y_n)\}$$

$$(B)_S := \{$$

$$(\forall t_1, \dots, t_m)([\bigwedge_{1 \leq i \leq m} Q(t_i)] \rightarrow \left[\underbrace{\bigvee_{q \in Q_m(I)} (\bigwedge q(t_1, \dots, t_m))}_{b} \right])$$

$$: m < \omega\}$$

$$\begin{aligned}
 (C)_S := \{ & \\
 (\forall t_1, \dots, t_m) ([\bigwedge q(t_1, \dots, t_m)] \rightarrow & \left[\underbrace{\bigvee_{p \in P_{(\Delta, q)}(I)} (\bigwedge p(f(t_1), \dots, f(t_m)))}_c \right]) \\
 : m < \omega, q \in Q_m(I), \text{ finite } \Delta(y_1, \dots, y_m) \subseteq L(y_1, \dots, y_m) \} & \\
 \\
 (D)_S := \{ & \\
 (\exists t_1, \dots, t_m) ([\bigwedge_{1 \leq i \leq m} Q(t_i)] \wedge (\bigwedge q(t_1, \dots, t_m))) & \\
 : m < \omega, q \in Q_m(I) \} &
 \end{aligned}$$

Remark 3.3.9. Note that conjunctions/disjunctions a, b, c are all finite. As L' is quantifier-free \aleph_0 -categorical, Q_m is finite, and since Δ is finite, $P_{(\Delta, q)}$ is finite. Moreover, the types $q \in Q_m(I)$, $p \in P_{(\Delta, q)}(I)$ are in fact formulas, again since L' is quantifier-free \aleph_0 -categorical, and Δ is finite.

Claim 3.3.10. *For any $\mathcal{L}(L', L)$ -structure $\mathfrak{C} \models T_0$ modeling S , $(f^{\mathfrak{C}}(j) : j \in I(\mathfrak{C}))$ is an L' -generalized indiscernible in $M(\mathfrak{C})$, $I(\mathfrak{C})$ is an L' -structure with $\text{age} = \mathcal{U}^*$, and $M(\mathfrak{C}) \models T$.*

Proof. Let \mathfrak{C} be an $\mathcal{L}(L', L)$ -structure realizing S . Put $J := I(\mathfrak{C})$, and $b_j := f^{\mathfrak{C}}(j)$ for $j \in J$. The sequence $(b_j : j \in J)$ is clearly generalized indiscernible because of the conditions in $(A)_S$. Moreover, $M(\mathfrak{C}) \models \text{Th}(M)$, as this is indicated by the first set of sentences in S . By $(D)_S$, every type from $\text{age}(I) = \mathcal{U}^*$ is realized in J . By $(B)_S$, every m -tuple from J realizes some type from $Q_m(I)$, i.e. $\text{age}(J) \subseteq \text{age}(I)$. Thus $\text{age}(J) = \text{age}(I) = \mathcal{U}^*$.

Finally, $(b_j : j \in J)$ must be based on $(a_i : i \in I)$ since, not only does any m -tuple $\langle j_1, \dots, j_m \rangle$ realize some type q^* from $Q_m(I)$, but by a collection of conditions from $(C)_S$, for any finite $\Delta \subseteq L$, \bar{j} satisfies some type p in the (Δ, q^*) -profile of I , namely it satisfies the same complete Δ -type as some m -tuple from I satisfying q^* . \square

Claim 3.3.11. *S is finitely satisfiable.*

Proof. Take a finite subset $F \subseteq S$. Let $(A)_F$ denote the intersection of $(A)_S$ and F , $(B)_F := (B)_S \cap F$, and $(C)_F := (C)_S \cap F$. Let $\mathcal{M} := \mathfrak{C}(I, M, \langle a_i : i \in I \rangle)$ for I and M and $(a_i : i \in I)$ as fixed at the beginning of the proof. We will show that $(A)_F$ is satisfiable in some large-enough substructure $\mathcal{M}_0 \subseteq \mathcal{M}$. In fact, we will choose \mathcal{M}_0 so that $I(\mathcal{M}_0) \subseteq I(\mathcal{M})$ is a finite substructure and $M(\mathcal{M}_0) = I(\mathcal{M})$.

First note some consequences.

Claim 3.3.12. *For any $\mathcal{L}(L', L)$ -structure \mathcal{M}_0 such that $I(\mathcal{M}_0) := I_0$ for some finite substructure $I \subseteq I(\mathcal{M})$, $M(\mathcal{M}_0) := M$ and $f^{\mathcal{M}_0}(i) := a_i$, we have that $\mathcal{M}_0 \models T_0 \cup (\text{Th}(M))^P \cup (B)_F \cup (C)_F$.*

Proof. Let $\mathcal{M}_0 \subseteq \mathcal{M}$ be an $\mathcal{L}(L', L)$ -substructure such that $I(\mathcal{M}_0) := I_0$ for some finite $I_0 \subseteq_{\text{finite}} I(\mathcal{M})$ and $M(\mathcal{M}_0) = M(\mathcal{M})$. Then

1. By the set-up, that $I_0 \subseteq I$ is a substructure, $M(\mathcal{M}_0) = M$ and $f^{\mathcal{M}_0}(i) := a_i$, we have that $\mathcal{M}_0 \subseteq \mathcal{M}$ is an $\mathcal{L}(L', L)$ -substructure.
2. $\mathcal{M}_0 \models T_0$, since \mathcal{M}_0 is a substructure of \mathcal{M} , and T_0 is universal.
3. Clearly $\text{age}(I_0) \subseteq \text{age}(I)$ as $I_0 \subseteq I$. Thus $\mathcal{M}_0 \models (B)_F$
4. Fix a finite $\Delta \subseteq L$ and tuple \bar{j} from J . Let p be the complete Δ -type of $f(\bar{j})$ in $M(\mathcal{M}_0)$. It remains to show that p is in the (Δ, q) -profile of I .

By the previous line, the complete quantifier-free L' -type of \bar{j} , $q^*(\bar{x})$, is in $Q(I)$. Thus, by definition, for $p := \text{tp}^\Delta(f(\bar{j}); M)$, $p \in P_{(\Delta, q^*)}(I)$. Thus $\mathcal{M}_0 \models (C)_F$.

□

If we assume that our \mathcal{M}_0 must have the properties that

1. $I(\mathcal{M}_0) = I_0$ for some finite $I_0 \subseteq I(\mathcal{M})$, and
2. $M(\mathcal{M}_0) = M$
3. $f^{\mathcal{M}_0}(i) = a_i$

then we need only search for our substructure $I_0 \subseteq I$ and the rest of the structure $\mathcal{M}_0 := (I_0, M(\mathcal{M}), \langle a_i : i \in I_0 \rangle)$ is completely determined.

It remains to show that we can choose a finite substructure $I_0 \subseteq I$ so that \mathcal{M}_0 also satisfies $(A)_F \cup (D)_F$.

Finding $I_0 \subseteq I$

There is some finite sublanguage $L_0 \subseteq L$ such that all formulas in either $(A)_F$ or $(D)_F$ are from the sublanguage $\mathcal{L}(L', L_0) \subseteq \mathcal{L}(L', L)$. Recall that L' is assumed finite relational, and so $\mathcal{L}(L', L_0)$ is itself also finite, and in some finite list of variables x_1, \dots, x_r .

We may enumerate:

1. the complete (L_0, r) -types η_1, \dots, η_k , for some k
2. the complete quantifier-free (L', r) -types, q_1, \dots, q_m , for some m

Thus we may assume (by adding dummy variables if necessary) that the only types q that occur in $(D)_F$ are among the q_i ; also, that the only φ that occur in conditions of $(A)_F$ are $\varphi \in L_0$.

Remark 3.3.13. For simplicity, we assume that all types q_i have variables listed in $<_{L'}$ -increasing order, and that if a structure $D \in \mathcal{U}$ “satisfies q_i ”, that is to say that $\langle d_1, \dots, d_r \rangle \models q_i$ where the d_i are listed in $<_{L'}$ -increasing order.

Discussion 3.3.14. In order to satisfy $(D)_F$, it suffices to make sure that our choice of model I_0 has the property that

(\dagger) I_0 realizes all the q_i , for all $i \leq m$.

In order to satisfy $(A)_F$, it suffices to make sure that for any complete (L', r) -type q realized in I_0 (thus, one of the q_i) that:

- ($*$) $_q$ 1. for any $\langle i_1, \dots, i_r \rangle$ and $\langle j_1, \dots, j_r \rangle$ in I_0 realizing $q(\bar{x})$, that
 2. $M \models (\varphi(\bar{a}_{\bar{i}}) \leftrightarrow \varphi(\bar{a}_{\bar{j}}))$, for all $\varphi(x_1, \dots, x_r)$ in L_0 .

Satisfying $(A)_F \cup (D)_F$

First let us look at a special case. In the following example we show how to realize (\dagger) and just the one ($*$) $_{q_1}$ with a finite $I_0 \subseteq I$:

Example 3.3.15. $[[\dagger] \ \& \ (*_{q_1}]$ We will show how to define our substructure $I_0 =: E'$.

Every $q_i(\bar{x})$ can be realized by a finite ordered graph, D_i in $\text{age}(I)$, thus can be realized by a member of $\mathcal{U}^* = \text{age}(I)$. Since \mathcal{U} has JEP, there is E in \mathcal{U} embedding all D_i , thus realizing all the q_i . Any E' in \mathcal{U} containing E will work to satisfy (\dagger).

In order to satisfy ($*$) $_{q_1}$, it suffices to find some $E' \in \mathcal{U}$ containing E and an (L_0, r) -type p_1 (this will be a type p_1 in $P_{(L_0, q_1)}(I)$) such that

$$\text{for all } \bar{j} \text{ in } E' \text{ satisfying } q_1(\bar{x}) \Rightarrow \bar{a}_{\bar{j}} \text{ satisfies } p(\bar{y})$$

Let D_1 in \mathcal{U} realize q_1 . We have said there are k (L_0, r) -types, η_j . Since \mathcal{U} is assumed to be a Ramsey class, there is C in \mathcal{U} , $C \rightarrow (E)_k^{D_1}$. Since $\mathcal{U} = \text{age}(I)$, there is an embedded copy of C , $C' \subseteq I$. Color all copies of D_1 in C' , $D' \subseteq C'$ according to the quantifier-free (L_0, r) -types, η_j , of the tuple $\langle a_{d_1}, \dots, a_{d_r} \rangle$, for $d_1 < \dots < d_r$ in D' . Since C is Ramsey for (E, D_1, k) , there is a copy of E in C' , $E' \subseteq C'$, such that all copies of D_1 in E' are colored the same color. We have found finite $E' \subseteq C' \subseteq I$. This $E' =: I_0$ works to satisfy ($*$) $_{q_1}$.

For the general case, we have the sequence of (L', r) -types, q_1, \dots, q_m . To satisfy (\dagger) and ($*$) $_{q_i} : i \leq m$, we must ensure that we pick $I_0 := E' \subseteq I$ such that

1. E' embeds a realization of each q_i
2. for all i , there is a fixed (L_0, r) -type p_i such that for any realization \bar{j} of q_i from E' , $\bar{a}_{\bar{j}}$ realizes p_i .

We first define a sequence of structures from \mathcal{U} that will play a role with respect to q_i similar to the role that C' had with respect to q_1 , above. Let the D_i be realizations of the q_i , thus, members of \mathcal{U} .

Claim 3.3.16. *There is a Y in \mathcal{U} such that Y embeds realizations of each q_i , $i \leq m$, and such that for any tuples \bar{i}, \bar{j} of length r from Y with the same complete quantifier-free L' -type, $tp^{L_0}(\bar{a}_{\bar{i}}; M(\mathcal{M})) = tp^{L_0}(\bar{a}_{\bar{j}}; M(\mathcal{M}))$. (Thus Y satisfies properties (\dagger) and $\langle (*)_{q_i} : i \leq m \rangle$.)*

Proof. Let E be as in Example 3.3.15, a member of \mathcal{U} that embeds all the q_i , $i \leq m$. Since we assume that \mathcal{U} is a Ramsey class, for each $1 \leq i \leq m$, for any $B \in \mathcal{U}$ and $i \leq m$, there is a $C \in \mathcal{U}$ such that:

$$C \rightarrow (B)_k^{D_i}$$

Moreover, any member $D \in \mathcal{U}$ has a copy in I , from the fact that $\text{age}(I) = \mathcal{U}$. We will utilize these facts to define a sequence from \mathcal{U} : $\langle Z_i : 0 \leq i \leq m \rangle$.

(stage $n = 0$): Let $Z_0 := E$

(stage $0 < n \leq m$): Let Z_n be such that $Z_n \rightarrow (Z_{n-1})_k^{D_n}$.

Now obtain Y by induction on n . At stage n , we define Y_{m-n} and Z_{m-n} for $0 \leq n \leq (m-1)$.

(stage $n = 0$): Let Y_m be a copy of Z_m in I .

(stage $0 < n < m$): By the previous stage, we have a substructure $Y_{m-n+1} \subseteq I$ isomorphic to Z_{m-n+1} . Color copies S of D_m , $S \subseteq Y_{m-n+1}$ according to the L_0 -type in M of $\langle f^{\mathcal{M}}(s_i) : i \leq r \rangle$ where $\langle s_i : i \leq r \rangle$ is an $<_{L'}$ -increasing enumeration of S . By definition of Z_{m-n+1} , there is a copy Y_{m-n} of Z_{m-n} , $Y_{m-n} \subseteq Y_{m-n+1}$, such that Y_{m-n} is homogeneous for D_{m-n+1} .

For all i , $1 \leq i \leq m$, we have that $Y_i \subseteq Y_{i+1}$, and Y_i is homogeneous for copies of D_{i+1} . Thus by an easy induction, for every $i > 0$, Y_i is homogeneous for copies of D_j for all j such that $(i+1) \leq j \leq m$. At stage $n = m-1$, we obtain Y_1 , a copy of Z_1 homogeneous for all copies of D_1, \dots, D_m . By the definition of Z_1 , there is a copy E' of E in Y_1 . This E' embeds all the q_i . Moreover, since $E' \subseteq I$, for any tuples \bar{i}, \bar{j} from E' of length r with the same L' -type, those tuples must satisfy one of the types q_{i_0} . Thus, in increasing enumeration, \bar{i}, \bar{j} are isomorphic to D_{i_0} , and since $E' \subseteq Y_1$, all copies of D_{i_0} in E' are colored the same color. \square

Define $I_0 := E'$ from the conclusion of Claim 3.3.16. By Claim 3.3.16, $I(\mathcal{M}_0)$ satisfies (\dagger) and $\langle (*_{q_i} : i \leq m) \rangle$.

This implies that $\mathcal{M}_0 \models (A)_F \cup (D)_F$, for $\mathcal{M}_0 := (I_0, M, \langle a_i : i \in I_0 \rangle)$. By Claim 3.3.12, such \mathcal{M}_0 satisfies $T_0 \cup (\text{Th}(M))^P \cup (B)_F \cup (C)_F$. Thus \mathcal{M}_0 satisfies $T_0 \cup F$ as desired.

We have argued that $F \cup T_0$ is satisfiable for arbitrary finite subsets $F \subseteq S$, and so $S \cup T_0$ is satisfiable, as desired. \square

By the satisfiability of S and Claim 3.3.10, the lemma is proved. \square

3.3.2 Characterization: necessity

Lemma 3.3.17. *Let L' be a finite relational language containing a binary relation symbol for order, $<$, and let \mathcal{U} be some nonempty collection of finite L' -structures that are linearly ordered by $<$. Let \mathcal{U}^* be the set of isomorphism types from \mathcal{U} . Suppose that \mathcal{U} has JEP and is closed under isomorphism and substructures.*

If L' -generalized indiscernibles have the modeling property for \mathcal{U}^ , then \mathcal{U} is a Ramsey class*

Proof. First, we wish to show that two properties are equivalent. Recall that \mathcal{U} is a Ramsey Class if

(1) for any $A, B \in \mathcal{U}$ and positive integer k , there is $C \in \mathcal{U}$ such that for any k -coloring $g : \binom{C}{A} \rightarrow k$, there is a $B' \subseteq C$ isomorphic to B , homogeneous for this coloring.

We claim that this is equivalent to,

(2) for any $A, B \in \mathcal{U}$ and positive integer k , for any L' -structure I such that $\text{age}(I) = \mathcal{U}^*$ and k -coloring $g : \binom{I}{A} \rightarrow k$, there is a $B' \subseteq I$ isomorphic to B , homogeneous for this coloring.

Claim 3.3.18. (1) \Rightarrow (2)

Proof. Given $A, B \in \mathcal{U}$ and $k > 0$, we start with an L' -structure I such that $\text{age}(I) = \mathcal{U}$ and a k -coloring g of the A -substructures of I . By (1), there exists a $C \in \mathcal{U}$ with certain properties. As $\text{age}(I) = \mathcal{U}^*$, there is a copy of this C , $C' \subseteq I$. By properties of C , any k -coloring of the A -substructures of C' yields a substructure isomorphic to B and homogeneous for this coloring. Note that there is a bijection between the A -substructures of C' and the A -substructures of C , as \mathcal{U} is closed under substructures. In particular, for the restriction of g to the A -substructures of C' , we get $B' \subseteq C' (\subseteq I)$, $B' \cong B$, such that B' is homogeneous for this coloring. Thus (2) is proved. \square

The converse is also true:

Claim 3.3.19. (2) \Rightarrow (1)

Proof. Let $(E_i : i < \omega)$ be an enumeration of \mathcal{U}^* , starting with a structure of size at least k . For every $\alpha < \omega$, there is a structure in \mathcal{U} embedding all the E_i for $i < (\alpha + 1)$, by JEP. Let C_α be one such structure (not necessarily unique):

$$C_\alpha := \bigoplus_{i < \alpha} E_i$$

All we need to know is that

1. each C_α is in \mathcal{U} , for all $\alpha < \omega$
2. for every structure $E \in \mathcal{U}$, there is a $\beta_E < \omega$ such that E embeds into the C_α for all $\alpha \geq \beta_E$.

Suppose (1) fails, witnessed by A, B, k where A has cardinality m and B has cardinality l . Thus each C from \mathcal{U} has a “bad” coloring $g_C : \binom{C}{A} \rightarrow \{1, \dots, k\}$ such that no $B' \subseteq C$ isomorphic to B has all A -substructures the same color. Expand L' to L^+ containing a new m -ary function symbol h and new constant symbols d_0, d_1, \dots, d_k .

Extend the meaning of g_C so that for a finite tuple $\bar{a} = \langle a_1, \dots, a_m \rangle$,

$$g_C(\bar{a}) = g_C(\{a_1, \dots, a_m\}).$$

Say that an interpretation of f in an L^+ -structure C has property $(*)$ if:

$$h^C(a_1, \dots, a_m) = \begin{cases} d_{g_C(\bar{a})}^C, & \text{if } \{a_i : i \leq m\} \text{ is isomorphic to } A \\ d_0^C, & \text{otherwise} \end{cases} \quad (3.1)$$

For each α , expand structure C_α so that

1. the $d_i^{C_\alpha}$ are (any) distinct elements from the model, and
2. h^{C_α} has property $(*)$.

Property $(*)$ is a definable condition. First of all, for arbitrary finite L' -structure C , $\|C\| = n_C$, define

$$p_C := \{\theta(x_1, \dots, x_{n_C}) : C \models (\exists s_1, \dots, s_{n_C})(\theta(s_1, \dots, s_{n_C}) \wedge \bigwedge_{i < j} s_i < s_j)\}$$

Remark 3.3.20. Note that p_C is a finite type, as it describes the (increasing) L' -type of a finite structure, where L' is finite relational

For all n_C -element L' -structures, Q , Q is isomorphic to C , for $\|C\| = n_C$ just in case

$$\langle q_1, \dots, q_{n_C} \rangle \models p_C(x_1, \dots, x_{n_C})$$

where $q_1 < \dots, < q_{n_C}$ are the elements of Q in $<_{L'}^Q$ -increasing enumeration.

Claim 3.3.21. For a finite L^+ -structure C , h^C has property $(*)$ just in case

$$C \models (\forall s_1, \dots, s_{n_C}) ((\bigwedge_{i < j} s_i < s_j) \rightarrow [\neg \bigwedge p_A(s_1, \dots, s_{n_C}) \wedge h(s_1, \dots, s_{n_C}) = d_0])$$

$$\vee [\bigwedge p_A(s_1, \dots, s_{n_C}) \wedge \bigvee_{1 \leq i \leq n_C} h(s_1, \dots, s_{n_C}) = d_j]$$

Proof. clear. □

Let \mathcal{D} be a nontrivial ultrafilter extending the cofinite filter on ω , and let $\mathcal{I} := \prod_{\alpha < \omega} C_\alpha / \mathcal{D}$, the ultraproduct of the C_α with respect to \mathcal{D} .

Claim 3.3.22. $\text{age}(\mathcal{I}) \subseteq \mathcal{U}$

Proof. For any finite L' -structure C not included in \mathcal{U} , there is a sentence φ_C expressing that no structure isomorphic to C exists as a substructure of any structure modeling φ_C . Thus every $A \in \mathcal{U}$ models φ_C for any $C \notin \mathcal{U}$. By Łoś, $\mathcal{I} \models \varphi$ for any φ in the common theory of the structures in \mathcal{U} . Thus $\text{age}(\mathcal{I}) \subseteq \mathcal{U}$ as desired. □

Claim 3.3.23. $\text{age}(\mathcal{I}) \supseteq \mathcal{U}$

Proof. Let $A \in \mathcal{U}$. By construction, there is a $\beta_A < \omega$ such that

$$C_\alpha \models (\exists \bar{x}) (\bigwedge q_A(\bar{x}))$$

Thus, $\{\alpha < \omega : C_\alpha \models (\exists \bar{x}) (\bigwedge q_A(\bar{x}))\}$ is in the filter, and so $\mathcal{I} \models (\exists \bar{x}) (\bigwedge q_A(\bar{x}))$ as well, implying that \mathcal{U} embeds a copy of A . □

By the previous claims, $\text{age}(\mathcal{I}) = \mathcal{U}$.

By Łoś, we have that:

$$\begin{aligned} \mathcal{I} \models (\forall s_1, \dots, s_{n_C}) ((\bigwedge_{i < j} s_i < s_j) \rightarrow [\neg \bigwedge p_A(s_1, \dots, s_{n_C}) \wedge \\ h(s_1, \dots, s_{n_C}) = d_0] \vee [\bigwedge p_A(s_1, \dots, s_{n_C}) \wedge \bigvee_{1 \leq i \leq n_C} h(s_1, \dots, s_{n_C}) = d_j]) \end{aligned}$$

Thus, we have a definable coloring on the n -element substructures of \mathcal{I} given by the interpretation of h . In fact, this is a k -coloring of A -substructures of \mathcal{U} , as every n -element substructure $Q \subseteq \mathcal{U}$ isomorphic to A gets sent to one of the $d_i^{\mathcal{I}}$, for $1 \leq i \leq k$.

Thus, by (2), there is $B' \subseteq \mathcal{U}$ isomorphic to B homogeneous with respect to this coloring. Thus, for some fixed $1 \leq k_0 \leq k$

$$\mathcal{I} \models (\exists x_1, \dots, x_l) (\bigwedge p_B(x_1, \dots, x_l) \wedge$$

$$\left[\bigwedge_{(1 \leq i_1 < \dots < i_m \leq l)} (\bigwedge p(x_{i_1}, \dots, x_{i_m}) \rightarrow h(x_{i_1}, \dots, x_{i_m}) = d_{k_0}) \right])$$

But then cofinitely many of the C_α model this sentence as well. Choose any one of these C_{α_0} . This model has a substructure B' isomorphic to B all of whose A -substructures are colored the same color, k_0 , contradicting our initial hypothesis. \square

So it suffices to show the following, in order to prove Lemma 3.3.17:

Claim 3.3.24. *If L' -generalized indiscernibles have the modeling property for \mathcal{U} , then (2) holds.*

Proof. To establish (2), fix L' -structure I with $\text{age} = \mathcal{U}$ and A, B, k and $g : \binom{I}{A} \rightarrow \{1, \dots, k\}$. We need to find $B' \subseteq I$ isomorphic to B , homogeneous for this coloring.

We want to reflect this as a coloring given by definable subsets of a target model, M . A has some finite size n . Let L be the language with k n -ary relations, R_1, \dots, R_k . and construct an $\mathcal{L}(L', L)$ -structure \mathcal{M} as follows:

1. $I(\mathcal{M}) = I$
2. $M(\mathcal{M})$ is a structure with underlying set, $|\mathcal{M}(M)| = |I|$
3. $f^{\mathcal{M}}(i) = i$
4. The relation R_s , $1 \leq s \leq k$, is interpreted as follows:

For i_1, \dots, i_n from $|\mathcal{M}(M)|$,

$$R_s^{\mathcal{M}}(i_1, \dots, i_n) \Leftrightarrow$$

- (a) $I \models \bigwedge_{(1 \leq l < m \leq n)} (i_l < i_m)$
- (b) $I \models (\bigwedge p_A)(i_1, \dots, i_n)$
- (c) $g(i_1, \dots, i_n) = s$

By assumption, we can find an L' -indexed indiscernible $(b_j : j \in J)$ in $M_1 \models \text{Th}(M)$ based on the $(a_i : i \in I)$ for L' -structure J with $\text{age} = \mathcal{U}$. Within the indiscernibility is the homogeneity of the coloring that we are looking for.

Remark 3.3.25. In the following discussion, it helps to remember that there are, in effect, two $\mathcal{L}(L', L)$ -structures at work here: first there is $\mathcal{M} = (I, M, \langle a_i : i \in I \rangle)$, then there is the new indiscernible in a separate structure, $\mathfrak{C}(J, M_1, \langle b_j : j \in J \rangle)$. What these two $\mathcal{L}(L', L)$ -structures have in common is that they both represent generalized indiscernibles mapping into models of the theory of M , which generalized indiscernibles happen to be indexed by models having the same age.

We are trying to transfer information between the two generalized indiscernibles, using the b_j for their “homogeneity” and the a_i for their information about modeling.

Since $\text{age}(J) = \text{age}(I) = \mathcal{U}$, there is a substructure $B' \subseteq J$ isomorphic to B , as fixed in the beginning of this proof. We enumerate B' by $B' =: \{j_k : k \leq N\} \subseteq J$ where $j_k < j_m$, if $1 \leq k < m \leq N$. Now use the modeling property: for $\Delta := L$, there is some i_1, \dots, i_N such that

$$\text{qftp}(i_1, \dots, i_N; I(\mathcal{M})) = \text{qftp}(j_1, \dots, j_N; J), \text{ and}$$

$$\text{tp}^\Delta(b_{j_1}, \dots, b_{j_n}; M(\mathcal{M})) = \text{tp}^\Delta(a_{i_1}, \dots, a_{i_n}; M(\mathcal{M}))$$

In other words, for all l , and increasing $1 \leq t(1) < \dots < t(n) \leq N$:

$$(1) \quad M(\mathcal{M}) \models R_l(a_{i_{t(1)}}, \dots, a_{i_{t(n)}}) \leftrightarrow R_l(b_{j_{t(1)}}, \dots, b_{j_{t(n)}})$$

Claim 3.3.26. *In fact, $B'' := \{a_{i_p} : 1 \leq p \leq N\} \cong B' \cong B$ is a copy of B in I that is homogeneous for the coloring, g .*

Proof. Since the b_j are generalized indiscernible, any two increasing n -tuples \bar{i}, \bar{j} index tuples $b_{\bar{i}}, b_{\bar{j}}$ in the target model with the same complete type. Let A''_1, A''_2 be two copies of A in B''

$$A''_1 =: \langle i_{s(1)}, \dots, i_{s(n)} \rangle$$

$$A''_2 =: \langle i_{t(1)}, \dots, i_{t(n)} \rangle$$

Say that A''_1 is colored $g(A''_1) =: c_1$ and A''_2 is colored $g(A''_2) = c_2$. We wish to show that $c_1 = c_2$.

Consider the similarly-indexed sets in J :

$$A'_1 =: \langle j_{s(1)}, \dots, j_{s(n)} \rangle$$

$$A'_2 =: \langle j_{t(1)}, \dots, j_{t(n)} \rangle$$

By generalized indiscernibility of the b_j :

$$(1) \quad R_{s_1}(j_{k(1)}, \dots, j_{k(n)}) \text{ and } R_{s_2}(j_{m(1)}, \dots, j_{m(n)})$$

We wish to show that $c_1 = c_2$.

Because $g(A''_1) = c_1$, and $A''_1 \cong A$, we have that $R_{c_1}^{M(\mathcal{M})}(i_{s(1)}, \dots, i_{s(n)})$. Likewise, since $g(A''_2) = c_2$, and $A''_2 \cong A$ we have that $R_{c_2}^{M(\mathcal{M})}(i_{t(1)}, \dots, i_{t(n)})$.

By modeling, we must have that

$$(2) \quad \text{tp}^L(b_{j_{s(1)}}, \dots, b_{j_{s(n)}}; M(\mathcal{M})) = \text{tp}^L(a_{i_{s(1)}}, \dots, a_{i_{s(n)}}; M(\mathcal{M}))$$

$$(3) \quad \text{tp}^L(b_{j_{t(1)}}, \dots, b_{j_{t(n)}}; M(\mathcal{M})) = \text{tp}^L(a_{i_{t(1)}}, \dots, a_{i_{t(n)}}; M(\mathcal{M}))$$

by taking subsequences, as a consequence of B'' and B' having the same L -type in $M(\mathcal{M})$.

In particular, we have that for $l = c_1, c_2$:

$$(4) R_{c_1}^{M(\mathcal{M})}(b_{j_{s(1)}}, \dots, b_{j_{s(n)}}) \Leftrightarrow R_{c_1}^{M(\mathcal{M})}(a_{i_{s(1)}}, \dots, a_{i_{s(n)}})$$

$$(5) R_{c_2}^{M(\mathcal{M})}(b_{j_{t(1)}}, \dots, b_{j_{t(n)}}) \Leftrightarrow R_{c_2}^{M(\mathcal{M})}(a_{i_{t(1)}}, \dots, a_{i_{t(n)}})$$

So the following both hold

$$(6) R_{c_1}^{M(\mathcal{M})}(b_{j_{s(1)}}, \dots, b_{j_{s(n)}})$$

$$(7) R_{c_2}^{M(\mathcal{M})}(b_{j_{t(1)}}, \dots, b_{j_{t(n)}})$$

We know by indiscernibility:

$$\text{tp}^{M(\mathcal{M})}(b_{j_{s(1)}}, \dots, b_{j_{s(n)}}) = \text{tp}^{M(\mathcal{M})}(b_{j_{t(1)}}, \dots, b_{j_{t(n)}})$$

Thus, these tuples agree on the R_l :

$$(8) M(\mathcal{M}) \models R_l(b_{j_{k(1)}}, \dots, b_{j_{k(n)}}) \leftrightarrow R_l(b_{j_{m(1)}}, \dots, b_{j_{m(n)}})$$

By (6), (7), (8):

$$(9) R_{c_1}^{M(\mathcal{M})}(b_{j_{s(1)}}, \dots, b_{j_{s(n)}}) \text{ and } R_{c_1}^{M(\mathcal{M})}(b_{j_{t(1)}}, \dots, b_{j_{t(n)}})$$

The R_l were defined in $M(\mathcal{M})$ in such a way so that for $c_1 \neq c_2$, no tuple satisfies both R_{c_1} and R_{c_2} . Thus, by (9), it must be that $c_1 = c_2$. \square

Thus we have established that $B'' \subseteq I$ is homogeneous for the coloring g , and so we have established (2). \square

As we have shown that (2) \Rightarrow (1) in Claim 3.3.19, this ends the proof of Lemma 3.3.17 \square

Theorem 3.3.27. *Let L' be a finite relational language containing a binary relation symbol for order, $<$, and let \mathcal{U} be some nonempty collection of finite L' -structures that are linearly ordered by $<$. Let \mathcal{U}^* be the set of isomorphism types from \mathcal{U} . Suppose that \mathcal{U} has JEP and is closed under isomorphism and substructures.*

\mathcal{U} is a Ramsey Class if and only if L' -generalized indiscernibles have the modeling property for \mathcal{U}^ .*

Proof. By Lemma 3.3.8 and Lemma 3.3.17 \square

We can draw a few consequences for I -indexed indiscernibles and the modeling property:

Proposition 3.3.28. *Let L' be a finite relational language and \mathcal{U}^* the age of some L' -structure. Let \mathcal{U} be the closure of \mathcal{U}^* under isomorphism. Suppose L' -generalized indiscernibles have the modeling property for \mathcal{U} . Then for any L' -structure I with $\text{age}(I) = \mathcal{U}^*$, I -indexed indiscernibles have the modeling property.*

Proof. Fix L' -structure I with $\text{age}(I) = \mathcal{U}^*$ and parameters $(a_i : i \in I)$ in some model M . Since L' -generalized indiscernibles are assumed to have the modeling property for \mathcal{U} , there is an L' -structure J with $\text{age}(J) = \mathcal{U}^*$ and an L' -generalized indiscernible $(b_j : j \in J)$ based on the a_i in some structure $M_1 \models \text{Th}(M)$. We know by Lemma 3.3.17 that \mathcal{U} must be a Ramsey class, thus by Theorem 3.2.18, \mathcal{U}^* is an amalgamation class with a Fraïssé limit, \mathcal{I} . Since $\text{age}(\mathcal{I}) \subseteq \text{age}(J)$ (in fact, they are both equal to \mathcal{U}^*), we have by Lemma 2.3.1 an L' -generalized indiscernible $(c_i : i \in \mathcal{I})$ based on the b_j . But by Prop 1.2.11, since $\text{age}(I) \subseteq \mathcal{U}^*$, \mathcal{I} must embed a copy of I . If we restrict the parameters to this copy of I , getting $(c_i : i \in I)$, we obtain I -indexed indiscernibles based (by transitivity) on the initial a_i . \square

Corollary 3.3.29. *Let L' be a finite relational language and T' a universal L' -theory. Suppose the class of all finite L' -models of T' , \mathcal{U} , is a Ramsey class. Suppose I is a quantifier-free weakly saturated model of T' . Then I -indexed indiscernibles have the modeling property.*

Proof. Let \mathcal{U}^* be the set of all isomorphism types from \mathcal{U} . \mathcal{U} is easily closed under substructure (since the theory T' is universal) and has JEP. Since \mathcal{U} is also assumed to be a Ramsey class, by Lemma 3.3.8 we know that L' -generalized indiscernibles have the modeling property for \mathcal{U} . By Proposition 3.3.28, it suffices to show that $\text{age}(I) = \mathcal{U}^*$, in order to establish that I -indexed indiscernibles have the modeling property.

$(\text{age}(I) \subseteq \mathcal{U}^*)$: Given finite $B \subseteq I$, $B \models T'$ since T' is a universal theory and B is a substructure of $I \models T'$. Thus $B \in \mathcal{U}^*$.

$(\mathcal{U}^* \subseteq \text{age}(I))$: Suppose $A \in \mathcal{U}^*$. Thus $A \models T'$. Since $p_A(\bar{x})$ is realized in $A \models T'$, it is a complete quantifier-free L' -type consistent with T' . Since I is a quantifier-free weakly saturated model of T' , I contains a realization of $p_A(\bar{x})$, i.e. a copy of A . Thus A is in $\text{age}(I)$. \square

Corollary 3.3.30. *For any quantifier-free weakly saturated ordered graphs I , I -indexed indiscernibles have the modeling property.*

Proof. We know by 3.2.9 that the class of all finite ordered graphs is a Ramsey class. Let \mathcal{U} in Corollary 3.3.29 be the class of all finite ordered graphs. \square

Remark 3.3.31 (existence of generalized indiscernibles not equivalent to the modeling property). We have in [Neš05] the example of (let's say ordered) graphs of girth > 4 as a class \mathcal{U} that is not an amalgamation class (to see this, take two vertices in your amalgamation base,

each of which is connected to one point in extending graphs B_1, B_2 .) Because \mathcal{U} is not an amalgamation class, it cannot be a Ramsey class, by Theorem 3.2.18. We have the *existence* of I -indexed indiscernibles for all I with $\text{age} = \mathcal{U}$. By applying the classical Ramsey theorem to the reduct structure $(I, <)$, we obtain an order indiscernible indexed by I , which is, a fortiori, an I -indexed indiscernible. However, we do not have that I -indiscernibles have the modeling property for \mathcal{U} .

We see this by applying Theorem 3.3.27. First, it is clear that \mathcal{U} is closed under isomorphism and substructures; it has JEP because the disjoint union of graphs with girth greater > 4 has girth > 4 and any partial order may be extended to a linear order. Let $\mathcal{U}^* = (C_i : i < \omega)$ be an enumeration of isomorphism types from \mathcal{U} so that every member of \mathcal{U} embeds into cofinitely many of the C_i (as in the argument for Claim 3.3.19.) Let I be an ultraproduct of the members of \mathcal{U}^* with respect to some nontrivial ultrafilter on ω extending the cofinite filter. As in the argument for Claim 3.3.19, $\text{age}(I) = \mathcal{U}$.

Thus, by Theorem 3.3.27, since \mathcal{U} is not a Ramsey class, I -indiscernibles fail to have the modeling property for \mathbf{K} , even though I -indexed indiscernibles exist for all I with $\text{age} = \mathcal{U}$ in models of all theories, T .

Chapter 4

Trees

4.1 Ramsey theorems for trees: introduction

From Ch. 3 we know that having the modeling property for certain generalized indiscernibles is closely related to a Ramsey-type theorem on finite substructures. In this chapter we look at the case of trees as index structures. One well-known partition theorem for finite subtrees of an infinite, finitely-branching tree is the Halpern-Läuchli theorem (see [Tod10] for an excellent presentation of variants of this theorem.) The latter theorem is used in [DS04] to prove an early modeling property result for what the authors call *full binary tree indiscernibles*, a sort of generalized indiscernible indexed by $\omega^{>2}$. Other Ramsey theorems for trees may be found in [Bla81, Mil81, Mil79] and the Appendix to [She78]. Predating the full binary tree indiscernibles is a modeling property result for infinitary trees (such as $\omega^{>\omega}$) in Ch. VII of [She78].

In the following chapter we give a detailed exposition of the proof of the modeling property for full binary tree indiscernibles as it is written in [DS04]. We show how Halpern-Läuchli is applied to obtain the result, and we gesture towards some further directions that might be taken.

4.2 Full binary tree indiscernibles: overview

A type of generalized indiscernible named *1-full binary tree indiscernibles (1-fbti's)* is introduced in [DS04]. In this chapter we give the language and theory of the index model for this type of generalized indiscernible. Also discussed in the paper are *0-fbti's* and *2-fbti's*. In our discussion of this paper, by *full binary tree indiscernibles* we will always mean 1-fbti's; if we mean to talk about 0- or 2-fbti's, we will name them as such, specifically. In this chapter, we continue our convention that single variables or constants in the target model stand for tuples of variables or elements in the model.

Here we give some preliminary definitions:

Definition 4.2.1. (informal) A *full binary tree indiscernible* is an (L_1, I_1) -generalized indiscernible where I_1 is the structure with underlying set ${}^{\omega}>2$ (the full binary tree on ω) and L_1 is the language given by:

1. a constant c_0 denoting the root of ${}^{\omega}>2$, $\langle \rangle$.
2. a 4-ary relation $x \wedge y \triangleleft z \wedge w$ interpreted as “the meet of x and y is below the meet of z and w in the partial tree-order”;
3. a 4-ary relation $(x \wedge y) \frown \langle 0 \rangle \triangleleft z \wedge w$ interpreted as “the node immediately above and to the left of the meet of x and y is below the meet of z and w in the partial tree order.”

Remark 4.2.2. Note that 2 n -tuples $\bar{\eta}, \bar{\nu}$ from ${}^{\omega}>2$ have the same quantifier-free L_1 -type just in case their closures under meets have the same type in the language whose signature consists of a constant for the root; the partial tree-order, \triangleleft ; and the relation $x \frown \langle 0 \rangle \triangleleft y$. This is accomplished without using function symbols.

A modeling-property result is obtained for these 1-fbtis, in pp. 30-35 of [DS04]. The authors prove that for any sequence of parameters $(a_i : i \in {}^{\omega}>2)$ in any model M , there exist (L_1, I_1) -generalized indiscernibles based on the a_i in some $M_1 \models \text{Th}(M)$.¹

In this section we give an exposition of the argument as it appears in the paper.

4.3 Definitions

We fix notation for the full binary tree on ω , ${}^{\omega}>2$, the partial tree-order, \triangleleft , and the lexicographic tree-order, $<$. Elements $\eta \in {}^{\omega}>2$ are finite functions from some $n < \omega$ into 2. By $\eta \triangleleft \nu$ we mean that $\eta \subseteq \nu$ as functions. By $\eta \wedge \nu$ we mean the meet of η and ν in the partial tree order, equivalently described as the sequence given by the greatest common initial segment of η and ν . By $\eta < \nu$ we mean that $(\eta \wedge \nu) \frown \langle 0 \rangle \triangleleft \nu$.

We reiterate a definition of the structure I_1 , first defining the language L_1 in detail:

Definition 4.3.1. By L_1 we mean the language in signature ρ_1 where ρ_1 consists of

1. a constant c_0 ;
2. a 4-ary relation $x \wedge y \triangleleft z \wedge w$;
3. a 4-ary relation $(x \wedge y) \frown \langle 0 \rangle \triangleleft z \wedge w$.

By L_2 we mean the reduct of L_1 to the signature $\rho_2 :=$

1. a constant c_0 ;

¹As these authors work in a monster model of $T := \text{Th}(M)$, there is no explicit mention of models M , M_1 , though the result is equivalent.

2. a 4-ary relation $x \wedge y \triangleleft z \wedge w$

Remark 4.3.2. Note that L_2 is not equipped to distinguish directionality in the partial tree-order in the way that L_1 is.

Definition 4.3.3. By I_1 we mean the structure with underlying set $|I_1| := {}^{\omega>}2$ which interprets the constant c_0 as the root, $\langle \rangle$, and interprets all the relations in ρ_1 in the usual way.

By I_2 we mean the reduct of I_1 to the language L_2 .

In [DS04] the following is proved (a similar result is stated for 2-fbti's, but we address this result in a later section):

Theorem 4.3.4 ([DS04]). *Given a sequence of parameters $\langle \bar{a}_\eta : \eta \in {}^{\omega>}2 \rangle$ in some $M \models T$ there exist 1-fbti $\langle \bar{b}_\eta : \eta \in {}^{\omega>}2 \rangle$ in some $M_1 \models T$ such that*

() for any tuple $\bar{\eta}$ from T and finite set Δ of L -formulas there exists a tuple $\bar{\nu}$ such that $qftp^{L_1}(\bar{\nu}; I_1) = qftp^{L_1}(\bar{\eta}; I_1)$ and $tp^{L(M)}(\bar{b}_{\bar{\eta}}; M_1) = tp^{L(M)}(\bar{a}_{\bar{\nu}}; M)$.*

Here is a restatement of this theorem in our own notation:

Theorem 4.3.5. *(L_1, I_1) -generalized indiscernibles have the modeling property.*

We aim to give an exposition of the argument for Theorem 4.3.5 as it appears in the paper. One of the main tools of the argument is to pass to a stronger language L_0 , which we define in the next few paragraphs. We define this language by describing the signature in stages:

Definition 4.3.6. Given tuples $\bar{\eta}, \bar{\nu}$ of elements from ${}^{\omega>}2$ and $n, \gamma \in \omega$, we say that $\bar{\eta} \sim_{\gamma, n} \bar{\nu}$ if there is some finite d such that

1. $\bar{\eta} := \langle \eta_0, \dots, \eta_{d-1} \rangle$ and $\bar{\nu} := \langle \nu_0, \dots, \nu_{d-1} \rangle$ where the η_i and ν_i are from ${}^{\omega>}2$,

and all of the following conditions hold.

For the following conditions fix

1. notation $(m := d + d + \binom{d}{2})$:

$$\bar{\eta}' = \langle \eta'_l : l \in m \rangle :=$$

$$\langle \langle \rangle, \eta_0 \upharpoonright \gamma, \dots, \eta_{n-1} \upharpoonright \gamma; \eta_0, \dots, \eta_{n-1}; \eta_0 \wedge \eta_1, \dots, \eta_i \wedge \eta_j (i < j), \dots \rangle$$

$$\bar{\nu}' = \langle \nu'_l : l \in m \rangle :=$$

$$\langle \langle \rangle, \nu_0 \upharpoonright \gamma, \dots, \nu_{n-1} \upharpoonright \gamma; \nu_0, \dots, \nu_{n-1}; \nu_0 \wedge \nu_1, \dots, \nu_i \wedge \nu_j (i < j), \dots \rangle$$

2. a function for our own use, $u(\gamma, n)(\bar{\eta}) := \{l(\eta_i) : i \in d \wedge l(\eta_i) > \gamma\}$.² Applied to ν ,
 $u(\gamma, n)(\bar{\nu}) := \{l(\nu_i) : i \in d \wedge l(\nu_i) > \gamma\}$

These are the conditions:

- (i) for $k \leq \gamma$: $\eta'_l \in {}^k 2 \leftrightarrow \nu'_l \in {}^k 2$, and if both sides of biconditional are true, then $\eta'_l = \nu'_l$
- (ii) $n \geq |u(\gamma, n)(\bar{\eta}')| = |u(\gamma, n)(\bar{\nu}')|$
- (iii) if $lg(\eta'_l), lg(\eta'_k) \in u(\gamma, n)(\bar{\eta}')$ then $lg(\eta'_l) < lg(\eta'_k) \leftrightarrow lg(\nu'_l) < lg(\nu'_k)$
- (iv) $\eta'_l \triangleleft \eta'_k \leftrightarrow \nu'_l \triangleleft \nu'_k$
- (v) $\eta'_l \frown \langle 0 \rangle \triangleleft \eta'_k \leftrightarrow \nu'_l \frown \langle 0 \rangle \triangleleft \nu'_k$

Here we informally explain the condition that $\bar{\eta} \sim_{\gamma, n} \bar{\nu}$. We will use the terminology that η is *exceptional* at l if η'_l has length greater than γ (where γ is understood, we might say that $\bar{\eta}$ has *exceptional length* at l).

It is not hard to show that $\bar{\eta} \sim_{\gamma, n} \bar{\nu}$ if and only if

1. $\bar{\eta}$ and $\bar{\nu}$ agree as functions on all $k \in \gamma$. In other words, they have the same restriction to ${}^\gamma 2$,
2. or every l , η and ν agree on whether they are exceptional at l , and if they are not, then they are identical at l ,
3. η and ν have an identical number of exceptional lengths and n bounds this number,
4. length-of-indexed-element linearly orders the indices of $\bar{\eta}'$ and $\bar{\nu}'$ (not necessarily compatible with the order on the integers in m), and the identity map on m is an order-isomorphism with respect to this new order,
5. the partial tree order is preserved, and
6. directionality is preserved.

It is clear that for a fixed γ and n , there is a first-order relational language (countably infinite in this case) such that $\bar{\eta} \sim_{\gamma, n} \bar{\nu}$ if and only if $\bar{\eta}, \bar{\nu}$ have the same quantifier-free complete type in this language. So we define a language thus (not uniquely):

Definition 4.3.7. We say that $\bar{\eta}, \bar{\nu}$ have the same quantifier-free $L_{\gamma, n}$ -type if $\bar{\eta} \sim_{\gamma, n} \bar{\nu}$.

Definition 4.3.8. By L_0 we mean the language whose signature is the union of the signatures of the $L_{\gamma, n}$ for $n, \gamma < \omega$.

By I_0 we mean the (intended) expansion of I_1 to the language L_0 .

²referred to as u_1 in [DS04]

Remark 4.3.9. From the latter exposition, it is clear that for all parameters $(a_\eta : \eta \in {}^{\omega>}2)$ in any model M ,

$$\text{qftp}^{L_0}(\bar{\eta}; I) = \text{qftp}^{L_0}(\bar{\nu}; I) \Rightarrow \text{qftp}^{L_1}(\bar{\eta}; I) = \text{qftp}^{L_1}(\bar{\nu}; I)$$

Definition 4.3.10. A sequence of parameters $(a_\eta : \eta \in {}^{\omega>}2)$ is *0-full binary tree indiscernible (0-fbti)* if it is L_0 -generalized indiscernible.

Given a subset $\Delta \subseteq L$, and $i \in \{0, 1, 2\}$, say that the sequence is *i-fbti for Δ* if

$$\text{qftp}^{L_i}(\bar{\eta}; I) = \text{qftp}^{L_i}(\bar{\nu}; I) \Rightarrow \text{tp}^\Delta(\bar{a}_{\bar{\eta}}; M) = \text{tp}^\Delta(\bar{a}_{\bar{\nu}}; M)$$

Indiscernibility is obtained in stages for a tree of parameters. We have the following definitions from [DS04]:

Definition 4.3.11. A sequence of parameters $(b_\eta : \eta \in {}^{\omega>}2)$ is

1. (γ, n) -indiscernible if

$$\bar{\eta} \sim_{\gamma, n} \bar{\nu} \Rightarrow \text{tp}^\Delta(\bar{a}_{\bar{\eta}}; M) = \text{tp}^\Delta(\bar{a}_{\bar{\nu}}; M)$$

2. (ω, n) -indiscernible if it is (γ, n) -indiscernible for all $\gamma < \omega$

Remark 4.3.12. Note that a sequence of parameters indexed by ${}^{\omega>}2$ is 0-fbti if it is (ω, n) -indiscernible for all $n < \omega$.

We end this section with a dictionary for terminology from [DS04]:

Remark 4.3.13. For tuples of elements from ${}^{\omega>}2$, $\bar{\eta}$ and $\bar{\nu}$

1. $\bar{\eta} \sim_{\gamma, n} \bar{\nu}$ iff $\text{qftp}^{L_{\gamma, n}}(\bar{\eta}; I_0) = \text{qftp}^{L_{\gamma, n}}(\bar{\nu}; I_0)$
2. $\bar{\eta} \sim_0 \bar{\nu}$ iff $\text{qftp}^{L_0}(\bar{\eta}; I_1) = \text{qftp}^{L_0}(\bar{\nu}; I_1)$
3. $\bar{\eta} \sim_1 \bar{\nu}$ iff $\text{qftp}^{L_1}(\bar{\eta}; I_1) = \text{qftp}^{L_1}(\bar{\nu}; I_1)$
4. $\bar{\eta} \sim_2 \bar{\nu}$ iff $\text{qftp}^{L_2}(\bar{\eta}; I_2) = \text{qftp}^{L_2}(\bar{\nu}; I_2)$

Remark 4.3.14. It would be correct to put “ I_0 ” for “ I_1 ” or “ I_2 ” in any of the above conditions.

Remark 4.3.15. The sequence $(b_\eta : \eta \in {}^{\omega>}2)$ is

1. (γ, n) -fbti if it is $L_{\gamma, n}$ -generalized indiscernible.
2. (ω, n) -fbti if it is $L_{\omega, n}$ -generalized indiscernible.
3. 0-fbti if it is L_0 -generalized indiscernible
4. 1-fbti if it is L_1 -generalized indiscernible
5. 2-fbti if it is L_2 -generalized indiscernible

Remark 4.3.16. Note that $(a_\eta : \eta \in {}^{\omega>}2)$ is 0-fbti if it is (ω, n) -fbti for all n

4.4 Halpern-Läuchli results:

Here we give some of the technical results that we need to understand the argument for fbt's. All definitions in this section are quoted from [Tod10].

Definition 4.4.1. Given a tree T and $x \in T$, $T[x]$ is the set of elements in T related to x under the partial tree-order. $T(k)$ is the set of elements in T at level k .

Definition 4.4.2. A product of trees $\prod_{i < d} T_i$ contains a *somewhere-dense matrix* if there exist $x_i \in X_i \subset T_i$ and $k > \max \text{length}(x_i)$ such that each X_i is $k - x_i$ -dense in T_i :

i.e. X_i dominates (under the tree order, \leq_{T_i}) every member of $T_i[x_i] \cap T_i(k)$ – those are the level- k members of the subtree of elements of T_i related to x_i .

Theorem 4.4.1 (Halpern-Läuchli; as in [Tod10]). *For every finite sequence T_i ($i < d$) of finitely-branching, rooted trees and every finite partition of their product $\prod_{i < d} T_i$, one of the pieces of the partition must contain a somewhere-dense matrix.*

Corollary 4.4.3 (Finite Halpern-Läuchli; as in [Tod10]). *For every $l \geq 1$ there is an n such that any partition $c : \prod T_i \upharpoonright n \rightarrow l$ is constant on a somewhere dense matrix $\prod X_i$ with $X_i \subseteq T_i \upharpoonright n$.*

Definition 4.4.4. For a finitely-branching, rooted tree T , $S \subset T$ is a *strong subtree witnessed by A* if $A \subseteq \omega$ is some infinite set such that

- (1) $S \subset \cup \{T(n) : n \in A\}$ and $S \cap T(n) \neq \emptyset$ for all $n \in A$
- (2) If $m < n$ are successive in A then for every $s \in S \cap T(m)$ every immediate successor of s in T has exactly one extension to $S \cap T(n)$

Corollary 4.4.5 (Strong Subtree Halpern-Läuchli; as stated in [Tod10]). *For every finite partition of $\prod_{i < d} T_i$ there is a piece of the partition P and a sequence S_i of strong subtrees of T_i witness by the same set A such that $\cup_{n \in \omega} \prod_{i < d} S_i(n) \subset P$*

4.5 Argument

Here we give a summary of the argument for fbt's.

(Summary:) Fix some Δ . Suppose we can obtain some b's L_1 -based on the a's indexed by $\omega^{>2}$ that are 0-fbti (for Δ). Then we obtain 1-fbti b's as follows: we take a particular skew subtree of $\omega^{>2}$ of great enough height that is still perfect. In this subtree, all \sim_1 -equivalent tuples will have the property that they are automatically $\sim_{n,\gamma}$ -equivalent for some n and γ (γ will be the maximal height on which these tuples agree and n will be the number of exceptional values afterwards – finite, since it is bounded by the length of the tuple.) However, our tree was 0-fbti, so $\sim_{n,\gamma}$ -equivalence implies that the associated tuples are Δ -equivalent. Thus we will get the desired property that two \sim_1 -equivalent tuples of indices index Δ -equivalent tuples in the target model, at least we get this property locally, which is enough for a compactness argument.

Now we give an outline of the detailed argument, and in the subsection afterwards, we describe some of the more difficult points.

1. We are given a sequence of a's and we are trying to obtain b's that are indiscernible and based on the a's. This is the main outline of the argument.

- (a) We get indiscernibility in two parts. All throughout we are taking subtrees of the original tree of parameters that are L_1 -substructures

Example 4.5.1. We just wish to point out that it is possible to not do this. One would only need to take a subset of ${}^\omega 2$ and then decide to interpret “left” and “right” in the opposite way.

We want to show that:

- i. $\bar{\eta} \sim_1 \bar{\nu} \Rightarrow \bar{b}_{\bar{\eta}} \equiv \bar{b}_{\bar{\nu}}$ in a big enough L_1 -substructure of ${}^\omega 2$.
- ii. Note that by compactness it is enough to show that for every finite $\Delta(x_0, \dots, x_{d-1}) \subset L$:

$$\bar{\eta} \sim_1 \bar{\nu} \Rightarrow \bar{a}_{\bar{\eta}} \equiv_{\Delta} \bar{a}_{\bar{\nu}}$$

in a big enough subtree

- iii. In fact we can assume that the length of $\bar{\eta}$ matches that of the length of variables in Δ , and by compactness, we may also assume that we are taking tuples of this length from a full binary tree of height some N .
- (b) To get L_1 -generalized indiscernibility by way of L_0 -generalized indiscernibility, we need a large enough 0-fbti tree, L_1 -based on the original parameters, such that we can take a special skew subtree wherein 0-fbt-indiscernibility implies \sim_1 -indiscernibility.
- (c) To preserve the modeling condition, we will make sure that our transition from one tree to the next is by taking successive strong subtrees (note that this does not preserve the quantifier-free L_0 -definable structure, only for L_1 .) If we start out with ${}^\omega 2$, this will cause every subsequent tree to be full binary, but also L_1 -based on the previous tree.

At every step we are likely to identify our new substructure with ${}^\omega 2$.

- (d) That our subtree is full binary allows us to find inside of it the “big enough subtree of ${}^\omega 2$ ” that we mentioned in step 1.(b). We obtain our subtree in stages, level by level.

All the n th branchings of our subtrees occur at the same level of the original tree, so that we can be sure that, in their intersection, every node in our subtree eventually branches.

(e) Outline of the L_0 -generalized indiscernibility argument:

- i. Start with $(a_\eta : \eta \in {}^{\omega>}2)$, tree of parameters
- ii. Fix finite $\Delta(x_0, \dots, x_{d-1}) \subseteq L$.
- iii. We need only show we can find 0-fbtis for Δ that are L_1 -based on the a 's.
- iv. Δ limits the length of the tuples from E , which limits the size of n in the $\sim_{\gamma, n}$ -indiscernibility. So we need only show

$$\langle \eta_0, \dots, \eta_{d-1} \rangle \sim_0 \langle \eta_0 \dots, \eta_{d-1} \rangle \Rightarrow \text{tp}^\Delta(\bar{a}_{\bar{\eta}}; M) = \text{tp}^\Delta(\bar{a}_{\bar{v}}; M)$$

which actually amounts to

$$\langle \eta_0, \dots, \eta_{d-1} \rangle \sim_{\omega, d} \langle \eta_0 \dots, \eta_{d-1} \rangle \Rightarrow \text{tp}^\Delta(\bar{a}_{\bar{\eta}}; M) = \text{tp}^\Delta(\bar{a}_{\bar{v}}; M)$$

since d puts a bound on the exceptional lengths in each $\bar{\eta}$ of length d .

- A. (induction): Given $T(n)$, an (ω, n) -indiscernible tree L_1 -based on $T(n-1)$, we can find a subtree $T(n+1)$, an $(\omega, n+1)$ -indiscernible tree, L_1 -based on $T(n)$.
- B. The intersection of the $T(n)$ for $n \in M$ gives a 0-fbti for Δ tree, L_1 -based on the original set of parameters.

Fix n . This is how to obtain $T(n)$ from $T(n-1)$:

- $T_0^n := T(n-1)$
- Given T_γ^n , (γ, n) -indiscernible, we find a subtree, $T_{\gamma+1}^n$ that is $(\gamma+1, n)$ -indiscernible, but agrees with the original tree up to γ .
- Given the requirement of agreement up to γ at each stage γ , these trees T_γ^n intersect to form T_∞^n , which is (ω, n) -indiscernible. Now set this tree to be $T(n)$.

2. Here is the structure of the argument for 0-fbtis that are L_1 -modeled on the a_i 's in greater detail. We show how to proceed by induction from an $L_{(\omega, n)}$ -indiscernible for Δ tree L_1 -based on the original parameters, to an $L_{(\omega, n+1)}$ -indiscernible for Δ tree, with the same property.

- (a) Suppose we have $(\bar{a}_\eta : \eta \in {}^{\omega>}2)$ that is $L_{(\omega, n)}$ -indiscernible (i.e. it is $L_{(\gamma, n)}$ -indiscernible for all $\gamma < \omega$.) We want to find a strong subtree of ${}^{\omega>}2$ that is $L_{(\omega, n+1)}$ -indiscernible.
- (b) In our argument we define:

$$T_0^{n+1} := T(n)$$

where $T(n)$ is the $L_{(\omega, n)}$ -indiscernible tree.

- (c) In this case, T_i^{n+1} will be $L_{(i+1, n+1)}$ -indiscernible.

- (d) We show that we can go from T_γ^{n+1} to $T_{\gamma+1}^{n+1}$ such that the $(T_\gamma^{n+1})_{\gamma < \omega}$ form a fusion sequence.
- (e) In other words, each is a subtree of the previous and the levels $\{0, \dots, \gamma\}$ of T_γ^{n+1} and $T_{\gamma+1}^{n+1}$ are equal. Thus, the T_γ^{n+1} converge:

$$T_\infty^{n+1} = \bigcap_{\gamma \in \omega} T_\gamma^{n+1}$$

- (f) We set

$$T(n+1) := T_\infty^{n+1}$$

where $T(n+1)$ is the $L_{(\omega, n+1)}$ -indiscernible tree.

- (g) Since $n \leq d$ is bounded by Δ , after finitely many $T(i)$ we are done. Since each is a strong subtree of the previous, the intersection gives us the 0-fbti, countably-infinite full binary tree, L_1 -based on the original set of parameters.

3. We try to describe what the skew subtree looks like.

- (a) Recall in 1.(a)(iii) that we said we only need to realize \sim_1 -indiscernibility for Δ in a height N fully binary tree.
- (b) When we witness the finite subtype of the type that guarantees L_1 -generalized indiscernibility, as in 1.(a)(iii), we do so in a subtree of our 0-fbti tree on ${}^\omega 2$ that preserves L_1 -basedness and is “skew of height n ”
- (c) The “skew of height n ” tree will have a natural isomorphism with ${}^{\leq n} 2$ that preserves L_1 -basedness.
- (d) However, in the skew tree, any \sim_1 -similar pair will be automatically \sim_{γ, n^*} -similar for some γ and n^* .
- (e) Thus the 0-fbti-ness of the tree housing the skew subtree yields the 1-fbti-ness in the skew tree which is carried by L_1 -isomorphism to ${}^{\leq n} 2 \subseteq {}^\omega 2$

4.5.1 Some further explanations

(2.(a)): Suppose we have $(\bar{a}_\eta : \eta \in {}^\omega > 2)$ that is $L_{(\omega, n)}$ -indiscernible (i.e. it is $L_{(\gamma, n)}$ -indiscernible for all $\gamma < \omega$.) We want to find a strong subtree of ${}^\omega > 2$ that is $L_{(\omega, n+1)}$ -indiscernible. At each stage we thin out the parameters, a_η , to obtain a strong subtree.

By induction on $n < \omega$:

case $n=0$: given $(\bar{a}_\eta : \eta \in {}^\omega > 2)$ this sequence is automatically $(\omega, 0)$ -indiscernible, since this is to say that two identical tuples in ${}^\omega > 2$ have the same Δ -type. So we can set $T(0) := {}^\omega 2$.

case $n+1$: Suppose we are given $(\bar{a}_\eta : \eta \in T(n))$ which is (ω, n) -indiscernible. We want to get a strong subtree of $T(n)$, $T(n+1)$, which is $(\omega, n+1)$ -indiscernible.

By induction on $\gamma < \omega$:

Set $T_0^{n+1} := T(n)$ our starting tree. We will obtain trees $T_{\gamma+1}^{n+1}$ for $\gamma < \omega$ which will be

1. $(\gamma+1, n+1)$ -indiscernible,
2. a strong subtree of T_γ^{n+1} , and
3. agreeing with this previous tree on the levels $\{0, 1, \dots, \gamma\}$

Thus we preserves $(\gamma, n+1)$ -indiscernibility from T_γ^{n+1} to higher trees, $T_{\gamma+k}^{n+1}$.


(2.(d)): Here we argue from T_γ^{n+1} to $T_{\gamma+1}^{n+1}$:

We are given $(\bar{a}_\eta : \eta \in T_\gamma^{n+1})$ and this sequence is (ω, n) -indiscernible (as well as $(\gamma, n+1)$ -indiscernible, but this is not needed for the induction step.)

In the following, identify T_γ^{n+1} with ${}^\omega 2$. We are going to find a strong subtree $T_{\gamma+1}^{n+1}$ indexing a $(\gamma+1, n+1)$ -indiscernible sequence that agrees with T_γ^{n+1} on levels $\{0, 1, \dots, \gamma\}$

In other words, the γ -th level remains fixed from T_γ^{n+1} to $T_{\gamma+1}^{n+1}$, but the $\gamma+1$ -level will possibly be higher in the subtree than in the original tree. It's as if we fix the tree on levels $\{0, \dots, \gamma\}$, and then in our application of Halpern-Läuchli, we take all the subsequent levels and space them out in ${}^{\omega>2}$.

Here we take a quick break in the argument to provide an illustration:



(Explanation of illustration): The tree pictured is T_3^{n+1} , imagined as $\omega > 2$, and this tree is $(3, n + 1)$ -indiscernible. We are trying to obtain T_4^{n+1} . We apply Halpern-Läuchli to the product $T_1 \times \dots \times T_8$ of cones above top nodes in ${}^4 2$, in a manner that is explained subsequently.

Halpern-Läuchli returns subtrees of each T_i , where each level k of T_i has been moved up to some possibly higher level k^* . We may put these new trees atop the original nodes, except the root may have been moved up too. So the image of $\langle 0, 0, 1 \rangle$ is the node connected to $\langle 0, 0 \rangle$ by the “new edge” – the node that is second from the left in the

image of what was level 3 in the original tree. In fact, this is really the only level of nodes that will be preserved (and everything below it) when we transition from T_4^{n+1} to T_5^{n+1} .

Back to the argument:

In the new tree $T_{\gamma+1}^{n+1}$ we want that:

$$\eta_0, \dots, \eta_{d-1} \sim_{\gamma+1, n+1} \nu_0, \dots, \nu_{d-1} \Rightarrow \bar{b}_{\bar{\eta}} \equiv_{\Delta} \bar{b}_{\bar{\nu}}$$

Note that there are finitely many classes $Y = [\bar{\eta}] / \sim_{\gamma+1, n+1}$, Y_1, \dots, Y_p (this is by inspection of the language $L_{\gamma+1, n+1}$ which is bounded by functions in γ and n .)

We have in T_{γ}^{n+1} one copy of $\omega^{>2}$ rooted on each element of $\gamma 2$, Call these trees T_1, \dots, T_{2^γ} .

First we claim that for any $\sim_{\gamma+1, n+1}$ -class Y and $\bar{\eta} = \langle \eta_0, \dots, \eta_{d-1} \rangle$ in Y we have an equivalence on indices i, j from $\{0, \dots, d-1\}$. In the rest of this argument, “exceptional” always means with respect to the tree $\gamma 2$.

We define the equivalence relation:

$$i \sim_{Y, \bar{\eta}} j \Leftrightarrow \eta'_i \upharpoonright \min u(\gamma, n+1)(\bar{\eta}) = \eta'_j \upharpoonright \min u(\gamma, n+1)(\bar{\eta})$$

Claim 4.5.2. *This relation $\sim_{Y, \bar{\eta}}$ is independent of the choice of $\bar{\eta}$.*

Proof. Let $\bar{\nu}$ be also in Y . So $\bar{\nu} \sim_{\gamma+1, n+1} \bar{\eta}$. Let $\rho = \min u(\gamma+1, n+1)(\bar{\nu})$, $\mu = \min u(\gamma, n+1)(\bar{\eta})$.

We need to show:

$$\eta'_i \upharpoonright \mu = \eta'_j \upharpoonright \mu \Leftrightarrow \nu'_i \upharpoonright \rho = \nu'_j \upharpoonright \rho$$

This claim is symmetric, so we do right-to-left. Suppose $\eta'_i \upharpoonright \sigma \neq \eta'_j \upharpoonright \sigma$.

(case 0): $\eta'_i = \eta'_j$. Then the biconditional holds, since every $\sim_{Y, \bar{\eta}}$ is reflexive.

(case 1): $\eta'_i \neq \eta'_j$ and they have different restrictions to $\gamma+1$. The $\sim_{\gamma+1, n+1}$ -similarity of $\bar{\eta}, \bar{\nu}$ guarantees that ν_i, ν_j have the same property

(case 2): $\eta'_i \neq \eta'_j$ and they have the same restrictions to $\gamma + 1$. Then the meet $\eta'_i \wedge \eta'_j$ must be in ${}^s 2$ for $\gamma + 1 \leq s$.

However, if it is not $s = \gamma + 1$, then it must be $s = \mu$. Otherwise, since the meet is itself an η'_k and of length less than η'_i it would have exceptional length less than μ , a contradiction, since no η' has that property. But it cannot be $s = \mu$ because we assumed that $\eta'_i \neq \eta'_j$.

Thus $\eta'_i \wedge \eta'_j \in {}^{\gamma+1} 2$. So, $\nu'_i \wedge \nu'_j$ is guaranteed by the $\sim_{\gamma+1, n+1}$ -similarity to also be in ${}^{\gamma+1} 2$, which means that the restrictions of the $\nu'_i \upharpoonright \rho, \nu'_j \upharpoonright \rho$ must be different, since $\nu'_i \wedge \nu'_j$ has already branched earlier.

□

Let ν^Y be the set $\{[i] / \sim_Y : i \in d\}$. If we let $n_Y = |\nu^Y|$ we can list representatives of the classes in ν^Y as $\langle s_1^Y, \dots, s_{n_Y}^Y \rangle$. These are the indices i for the “stubs” of the branches η'_i in $\bar{\eta} \in Y$, having been cut down to the least exceptional length, without listing a stub twice.

We want to color every level tuple from the product $T_1 \times \dots \times T_{2^\gamma}$ and then apply Halpern-Läuchli. Suppose we have a tuple from level μ (this is level $\gamma + \mu$ according to T_γ^{n+1}),

$$\bar{e} = (e_1, \dots, e_{2^\gamma}) \in \underbrace{({}^\mu 2) \times \dots \times ({}^\mu 2)}_{2^\gamma}$$

(*) send $(e_1, \dots, e_{2^\gamma})$ to the color (c_1, \dots, c_p) where each $c_i = \text{tp}_\Delta(\bar{a}_{\bar{\eta}})$, where

- $\bar{\eta}$ is some tuple $\bar{\eta} = \langle \eta_0, \dots, \eta_{d-1} \rangle$ in Y_i , and
- $e_{s_k} = \eta'_{s_k} \upharpoonright \min u(\gamma, n + 1)(\bar{\eta})$.

Remark 4.5.3. We color level tuples, but the coloring only depends on n_{Y_i} elements of the tuple – an enumeration of the places where our $\bar{\eta}$ has a stub of an η' , of minimal exceptional length.

Claim 4.5.4. *The map at * is well-defined.*

Proof. Take any other $\bar{\nu} \sim_{\gamma+1, n+1} \bar{\eta}$, i.e. in Y_i . Let $e_{s_k} = \nu_{s_k} \upharpoonright \min u(\gamma, n + 1)(\bar{\nu})$. Fix $u(\gamma, n + 1)(\bar{\nu}) =: \rho$. We want to conclude that $\bar{\eta} \sim_{\mu+\rho, n} \bar{\nu}$. This will cause both $\bar{\eta}$ and $\bar{\nu}$ to be mapped to the same Δ -type, $\text{tp}_\Delta(\bar{a}_{\bar{\eta}}) = \text{tp}_\Delta(\bar{a}_{\bar{\nu}})$, by the assumption of (ω, n) -indiscernibility.

We only need to show that the two tuples look the same restricted to $\mu + \rho$. This is because (iv) and (v) are independent of γ and n ; for (iii), the relative heights is naturally

preserved; and for (ii), the bound on the exceptional heights follows immediately given that we are making the least exceptional height from before, $\gamma + \rho$, now non-exceptional.

However, for (i), it is also clear because we are implicitly assuming that $u(\gamma, n + 1)(\bar{\nu}) = u(\gamma, n + 1)(\bar{\eta})$, and all the restrictions to $\gamma + \rho$, index-by-index, are the common values e_{s_k} . (We know the indices of $\bar{\eta}, \bar{\nu}$ at which they occur are common, since they depend entirely on the $\sim_{\gamma, n+1}$ -class Y_i .) \square

(3.): We may assume that $(\bar{a}_\eta : \eta \in {}^{\omega>}2)$ are 0-fbti for Δ and that we are looking for $(\bar{b}_\eta : \eta \in {}^{\omega>}2)$ that are 1-fbti and L_1 -based on the a 's.

We use the construction in [DS04] to skew the tree to the left. What we actually do is find finite skew subtrees of arbitrary height n . This construction preserves the linear order as well as the tree order, however it is such that you can read all of (i)-(iii) off of these two orders. So in other words we have that $\bar{\eta}$ and $\bar{\nu}$ having the same \sim_1 -type in ${}^{\leq n}2$ have images in the skew tree of height n with the same \sim_{0, n^*} -type where n^* is big enough to bound the u_i of condition (iii).

Here is the construction:

for each n we have a function

$$h_n : {}^{n \geq 2} \rightarrow {}^{\omega > 2}$$

(this is going to be our isomorphism between ${}^{n \geq 2}$ and our skew subtree of height n .)

and we have a number (it is “+1” in [DS04], but let’s make it “+2” here, just for some extra room)

$$k_n = \max\{\lg(h_n(\eta)) + 2 : \eta \in {}^{n \geq 2}\}$$

The idea is that every time the original tree turns right, the skew tree copies the old tree, but when it turns left, the skew tree adds on enough zeroes so that no two forks are on the same level – in fact any fork (meaning the node preceding a fork) is above any fork to its right.

Here is the definition:

$$h_0(\langle \rangle) = \langle \rangle$$

$$h_{n+1}(\langle \rangle) = \langle \rangle$$

$$\begin{aligned} h_{n+1}(\langle 1 \rangle \hat{\ } \nu) &= \langle 1 \rangle \hat{\ } h_n(\nu) \\ h_{n+1}(\langle 0 \rangle \hat{\ } \nu) &= \langle 0, 0, \dots, 0 \rangle \hat{\ } h_n(\nu) \end{aligned}$$

where there are k_n zeroes in the last expression.

Thus given $\bar{\eta}, \bar{\nu}$ both in $n \geq 2$ and $n^* = \text{length of } cl(\bar{\eta})$,

$$(*^2) \quad \bar{\eta} \sim_1 \bar{\nu} \rightarrow h_n(\bar{\eta}) \sim_{0, n^*} h_n(\bar{\nu})$$

so that two tuples \sim_1 -similar in the skew tree have \sim_{0, n^*} -similar images as considered in the super- 0-fbti tree.

Verification of $(*^2)$:

(i) We only need to check for $\gamma = 0$. But if the images of η'_0, η'_1 have empty meet in the skew tree than the originals have empty meet in the old tree. This is because the skewing preserves the \sim_1 -language.

(ii) By definition of n^* on the basis of h_n .

(iii) Lengths decrease as one goes from left to right, and top to bottom, which latter properties \sim_1 can detect.

4.6 The Ramsey theorem and fbti

The finite Ramsey theorem guarantees the modeling property for indiscernible sequences, but it does not guarantee the modeling property for I -indiscernibles for general I . A good illustration of this is the 2-full binary tree indiscernibles from [DS04].

In [DS04], the authors prove that 1-fbtis have the modeling property. Recall the definition of SOP_2 :

Definition 4.6.1. Parameters $(a_\eta : \eta \in {}^\omega > 2)$ in M have SOP_2 if for some formula $\varphi(x; y)$ in $L(M)$,

1. for any $\eta, \nu \in {}^\omega > 2$ $\{\varphi(x; a_\eta), \varphi(x; a_\nu)\}$ is consistent just in case η and ν are comparable (either $\eta \triangleleft \nu$ or $\nu \triangleleft \eta$)
2. for every $\sigma \in {}^\omega 2$, $\{\varphi(x; a_{\sigma \upharpoonright n}) : n < \omega\}$ is consistent.

One application of tree indiscernibles in [DS04] is that given a set of parameters witnessing SOP_2 , we can find a tree indiscernible set witnessing the same. This is a case where the use of the Halpern-Läuchli Theorem, a Ramsey-type theorem for full binary trees, is an essential use.

Consider a given parameter set $C = \{a_\eta : \eta \in {}^{\omega>}2\}$ in some L -structure M , witnessing SOP_2 (for notational simplicity, we assume $\text{length}(a)=1$). To find tree indiscernibles witnessing the modeling property for C , we want to realize the following type in new parameters $(c_\eta : \eta \in {}^{\omega>}2)$:

$$\begin{aligned} \Gamma = & \{ \neg\theta(c_{i_1}, \dots, c_{i_n}) : (\forall \bar{b} \in C) (\text{qftp}_{\rho_t}^I(\bar{b}) = \text{qftp}_{\rho_t}^I(\bar{c}) \rightarrow \neg\theta(b_{i_1}, \dots, b_{i_n})) \} \\ & \cup \{ \varphi(c_{i_1}, \dots, c_{i_m}) \rightarrow \varphi(c_{j_1}, \dots, c_{j_m}) : \text{qftp}_{\rho_t}^I(\bar{i}) = \text{qftp}_{\rho_t}^I(\bar{j}) \} \end{aligned}$$

Suppose we impose the order on the b 's compatible with the linear order on ${}^{\omega>}2$, which has the same order type $(\mathbb{Q}, <)$. Now let us apply the finite Ramsey theorem for sets to $(b_i : i \in \mathbb{Q})$. Indeed, we will get a tree-indiscernible by Proposition.

However, the finite Ramsey theorem only tells us that increasing sequences of the c 's have the Δ -type (for any finite Δ a subset of the L -formulas) as increasing sequences of the b 's. Thus, even though $\{\varphi(x; b_\eta), \varphi(x; b_\nu)\}$ is inconsistent whenever $\eta \perp \nu$, $\{\varphi(x; c_\eta), \varphi(x; c_\nu)\}$ could very well be consistent based on the fact that $\eta < \nu$ (or vice versa) and $\{\varphi(x; b_{\eta_1}), \varphi(x; b_{\eta_2})\}$ is consistent whenever $\eta_1 \triangleleft \eta_2$ (since $x \triangleleft y \Rightarrow x < y$).

4.7 Further directions

4.7.1 in the direction of 1-fbti's

In one application of the modeling-property result for full binary tree indiscernibles, the authors obtain a very indiscernible witness to SOP_1 given an initial set of parameters $(a_\eta : \eta \in {}^{\omega>}2)$ witnessing SOP_1 . Recall the definition:

Definition 4.7.1. (SOP_1) A complete theory T has SOP_1 if for some formula $\varphi(\bar{x}; \bar{y})$ and parameters \bar{a}_η from ${}^{\omega>}2$ we have that

1. for any path $\rho \in {}^{\omega>}2$, $\{\varphi(\bar{x}; \bar{a}_{\rho \upharpoonright n}) : n < \omega\}$ is consistent.
2. if $\nu \frown \langle 0 \rangle \triangleleft \eta$, then $\{\varphi(\bar{x}; \bar{a}_\eta), \varphi(\bar{x}; \bar{a}_{\nu \frown \langle 1 \rangle})\}$ is inconsistent.

The argument given is by way of finding 1-fbtis based on the initial set of parameters. This is an interesting situation in which to look at the modeling property. We want to talk about the pair $\varphi(\bar{x}; \bar{a}_\eta), \varphi(\bar{x}; \bar{a}_\nu)$ for the condition $[(\eta \wedge \nu) \frown \langle 0 \rangle \triangleleft \eta] \wedge [(\eta \wedge \nu) \frown \langle 1 \rangle = \nu]$. But suppose

(*)³ that this is not a quantifier-free definable condition in L_1

(it might be, but since it seems tricky to show, let's say for the sake of argument that it is not.) In response to this issue, we may wish to add a new 2-ary relation, $P(\eta, \nu)$ exactly for this condition. Thus we have a new language L^+ in signature $\rho_1 \cup \{P(x, y)\}$ and an

expansion I^+ of I_1 to the intended structure in this language. We are still able to find L^+ -generalized indiscernibles, by virtue of being able to find L_1 -generalized indiscernibles, but we have not shown yet that (L^+, I^+) -generalized indiscernibles have the modeling property. Namely, though we are able to find L^+ -generalized indiscernibles that are L_1 -based on the initial parameters witnessing SOP_1 , we have not shown that we can find them L^+ -based on the initial set of parameters. Thus they are free to express any number of conditions in the L_1 -language. In fact, b_η, b_ν for $\neg(\nu \triangleleft \eta)$, $(\nu \wedge \eta) \frown \langle 0 \rangle \triangleleft \eta$ will always yield consistent $\{\varphi(x; b_\eta), \varphi(x; b_\nu)\}$, or they will always yield inconsistent $\{\varphi(x; b_\eta), \varphi(x; b_\nu)\}$. If we start with an SOP_1 tree of parameters that is as free as possible (i.e., such that for $(\nu \wedge \eta) \frown \langle 0 \rangle \triangleleft \eta$, $(\nu \wedge \eta)$ that $\{\varphi(x; b_\eta), \varphi(x; b_\nu)\}$ is consistent) it is possible to get a non- SOP_1 indiscernible set that is L_1 -based on the original set, assuming $(*)^3$.

The partition theorem that we use to model on the original set of parameters takes finite subtrees that do not agree with $\omega^{>2}$ on the P -relation. This might lead us to seek a stronger partition theorem on the subtrees of $\omega^{>2}$.

4.7.2 In the direction of 2-fbti's

There is also a worry that 2-fbti's may not have the L_2 -modeling property. For example, if we start with parameters $(a_\eta : \eta \in \omega^{>2})$ in the model $(\omega, <)$ under the usual linear ordering, 2-fbti $(b_\nu : \nu \in \omega^{>2})$ based on the a_η in a model of the theory of linear order would have to have the property that there is one quantifier-free $(L_{\{<\}}, 2)$ -type Δ such that whenever $\nu \neq \eta$, $\nu \triangleleft \eta$ that $\text{tp}(b_\eta, b_\nu) = \Delta$. But then this type must be symmetric in $(\omega, <)$, since by generalized indiscernibility, $\text{tp}(b_\eta, b_\nu) = \text{tp}(b_\nu, b_\eta)$, and we have argued that this structure has no symmetric quantifier-free n -types for any n .

This leads us to develop a more flexible theory of generalized indiscernibles. There is some hope for this in the ‘‘admissible orderings’’ of [Neš05]. When we associate to a finite structure A some set \preceq_A of admissible orderings, such that any order-preserving map among these orderings is an automorphism of A , we are keeping track of the permutations of A which yield automorphisms. This is a way to get away from the rigidity given by $<$ in the original Nešetřil-Rödl theorem. It is possible that some of the results in [Neš05] on partition properties for $(\text{age}(I), \preceq)$ may serve to extend the theory of generalized indiscernibles and the modeling property. We would aim for some sort of result where, given parameters in a model M whose types have a certain amount of symmetry, at least as much as that of the types in the index model I , we can find I -indiscernibles based on that set of parameters.

For example, given $(a_\eta : \eta \in E)$ witnessing SOP_2 , find 1-fbti $(b_\eta : \eta \in \omega^{>2})$ based on these parameters. First of all, the b_η retain the property of SOP_2 , since it is definable in the original quantifier-free language of L_1 . However, it is likely that they also happen to be 2-fbti, at least according to the reduct of M to φ , where φ witnesses SOP_2 for the a_η .

Chapter 5

Characterizing NIP theories

5.1 Definitions

Note that when we talk about subgraphs, we mean submodels of graphs, which is equivalent to the subgraphs being induced subgraphs. Often we will simply say “ordered graph” when we mean “linearly ordered symmetric graph with no loops.”

A *complete ordered graph type* will mean a complete quantifier-free L_g -type consistent with T_g . A *complete R -type* will mean a complete quantifier-free $L_g \upharpoonright \{R\}$ -type consistent with T_g ; similarly, a *complete order type* will mean a complete quantifier-free $L_g \upharpoonright \{<\}$ -type consistent with T_g . An AE-theory is a theory whose axioms are AE-formulas, i.e. formulas of the form $\forall \bar{x} \exists \bar{y} \theta(\bar{x}; \bar{y})$, where θ is quantifier-free. It is not hard to show that for any AE-theory, an increasing union of models of this theory is itself a model.

5.2 Characterization of NIP T

In this section we prove the two lemmas which will give our characterization of NIP theories by ordered graph indiscernibles.

Lemma 5.2.1 (\Leftrightarrow). *Suppose T has the independence property. Then there exists an infinite quantifier-free weakly-saturated ordered graph-indiscernible in a model of T which fails to be an indiscernible sequence.*

Proof. In fact we will find an ordered graph-indiscernible indexed by \mathcal{R} satisfying these conditions.

Let T be as above and choose $M \models T$ to be some \aleph_0 -saturated model of T . By Lemma 2.2 of [LS03], since T has IP there exists a formula $\varphi(x; y)$ that “codes graphs” in T . In other words, φ is a symmetric formula that embeds every finite graph relation. Let $G \subseteq |M|^2$ be the realizations of φ in M . Since M is countably saturated, by compactness there is countable $A \subseteq M$ so that $\mathcal{R} \upharpoonright \{R\} \cong_f \mathfrak{A} := (A, G \upharpoonright A \times A)$. Now consider as our initial

set of parameters the elements of A indexed by elements of \mathcal{R} by way of this mapping: $(a_g : f(g) = a)$.

From the previous section we know that I -indiscernibles exist and have the modeling property, where I is a weakly saturated model of T_g . In particular, for $I = \mathcal{R}$. Thus we can find \mathcal{R} -indiscernible $(b_g : g \in \mathcal{R})$ based on the a_g . It remains to show that this sequence is not order-indiscernible. For this we establish:

Claim 5.2.2. $M \models \varphi(b_g, b_h) \Leftrightarrow \mathcal{R} \models R(g, h)$.

Proof. Suppose $M \models \varphi(b_g, b_h)$, then there exists (a_e, a_f) , $\text{qftp}^{\mathcal{R}}(g, h) = \text{qftp}^{\mathcal{R}}(e, f)$, such that $M \models \varphi(a_e, a_f)$. However, $M \models \varphi(a_e, a_f)$ just in case $R^{\mathfrak{A}}(a_e, a_f)$, just in case $\mathcal{R} \models R(e, f)$. Thus we conclude that $\mathcal{R} \models R(g, h)$.

For the other direction, suppose $M \models \neg\varphi(b_g, b_h)$, then there exists (a_e, a_f) , $\text{qftp}^{\mathcal{R}}(g, h) = \text{qftp}^{\mathcal{R}}(e, f)$, such that $M \models \neg\varphi(a_e, a_f)$. However, $M \models \neg\varphi(a_e, a_f)$ just in case $R^{\mathfrak{A}}(a_e, a_f)$ fails, just in case $\mathcal{R} \models \neg R(e, f)$. Thus we conclude that $\mathcal{R} \models \neg R(g, h)$. \square

By the axioms in $T_{\mathcal{R}}$ we can easily find i_1, i_2, j_1, j_2 in \mathcal{R} such that

- $i_1 < i_2 \wedge i_1 R i_2$, and
- $j_1 < j_2 \wedge \neg j_1 R j_2$

Thus,

- $\varphi(\bar{b}_{i_1}, \bar{b}_{i_2})$, but
- $\neg\varphi(\bar{b}_{j_1}, \bar{b}_{j_2})$

showing that the b_g fail to be order-indiscernible. \square

Lemma 5.2.3 (\Rightarrow). *Suppose some quantifier-free weakly saturated ordered graph-indiscernible in a model $M \models T$ fails to be order-indiscernible. Then T has the independence property.*

Proof. Let T, M be as above and $(\bar{a}_i : i \in I)$ an I -indiscernible in M as above. We may assume that M is \aleph_0 -saturated by taking an elementary extension. By Claim 2.3.1, we may assume our indiscernible is an \mathcal{R} -indiscernible, by stretching our given sequence to an \mathcal{R} -indiscernible such that $p_n^{\eta}(\mathcal{R}) = p_n^{\eta}(I)$ for all n . This new indiscernible will retain the property of not being order-indiscernible. We will show that T is forced to have the independence property.

We've assumed that for some n and complete order-type $p = \{x_1 < \dots < x_n\}$, two n -tuples \bar{i}, \bar{j} satisfying p map to $\bar{b}_{\bar{i}}, \bar{b}_{\bar{j}}$ in M with different $\theta(\bar{x}_1, \dots, \bar{x}_n)$ -types for some θ in L . Without loss of generality, \bar{i} maps to the tuple realizing θ .

Definition 5.2.4. With respect to a complete R -type $q(x_1, \dots, x_n)$, by the R -truth value of the pair (x_s, x_r) we mean 0 if $R(x_s, x_r)$ occurs in q , 1, otherwise.

Claim 5.2.5. We may assume that the types of \bar{i}, \bar{j} in x_1, \dots, x_n are identical but for disagreement in the R -truth value of (x_k, x_l) , for some $k < l$: i.e. that
 $tp(i_k/\{i_s : s \neq k, l\}) = tp(j_k/\{j_s : s \neq k, l\})$, $tp(i_l/\{i_s : s \neq k, l\}) = tp(j_l/\{j_s : s \neq k, l\})$,
 $tp(i_1 \dots i_{k-1}, i_{k+1} \dots i_{l-1}, i_{l+1} \dots i_n) = tp(j_1 \dots j_{k-1}, j_{k+1} \dots j_{l-1}, j_{l+1} \dots j_n)$
 but $i_k R i_l$ and $\neg j_k R j_l$.

Proof. Consider the set X of all complete increasing R -types in x_1, \dots, x_n . Let $k := \binom{n}{2}$. If we enumerate the unordered pairs from n we can identify each R -type with a tuple from $\prod_k Z/2Z$. Identify an R -type in variables (x_1, \dots, x_n) with $\eta \in {}^k 2$ if for the i th pair, (x_r, x_t) , $\eta(i) = 0$ iff R holds on this pair. The group $G = \prod_k Z/2Z$ acts on X transitively: for any two R -types in X , it is clear we can proceed from one to the other by a finite list of alternations of R -truth values between pairs.

Our I -indiscernible gives a function that send quantifier-free types in I to $\{\theta, \neg\theta\}$, by $p \mapsto \theta$ if some/all \bar{i} with quantifier-free type p give $\theta(\bar{b}_{\bar{i}})$, $\neg\theta$, otherwise. This naturally gives a mapping from $X \rightarrow \{\theta, \neg\theta\}$.

Let A be the set of complete R -types in X that map to θ , B the set of types in X that map to $\neg\theta$. By failure of order-indiscernibility, each of A and B contain some element, (the types of \bar{i}, \bar{j}) call these elements a, b respectively. The group $\prod_k Z/2Z$ is generated by the set of $e_t = (g_1, \dots, g_k)$, for $t \leq k$, where $g_v = 1 \leftrightarrow v = t$. By transitivity, some group element g takes a to b . This element g is a product $h_s \cdot \dots \cdot h_1$ from the generating set, $\{e_v : v \leq k\}$. Since, $g \cdot a$ fails to be in A , there is some least m such that $h_m \cdot \dots \cdot h_1 \cdot a$ fails to be in A .

Call $h_{m-1} \cdot \dots \cdot h_1 \cdot a =: a'$, and $h_m \cdot \dots \cdot h_1 \cdot a =: b'$. We may replace \bar{i} with $\bar{i}' \models a'$, and \bar{j} with $\bar{j}' \models b'$. Thus we have two realizations of p whose representations in X differ at one coordinate, i.e. which satisfy two complete quantifier-free types in L_g that differ only in the alternation of an R -truth value at one pair of variables, call them (x_k, x_l) . \square

From now on assume \bar{i}, \bar{j} satisfy the properties in Claim 5.2.5.

Claim 5.2.6. We may assume that the j_s equal the i_s for $s \neq k, l$.

Proof. $tp(i_1 \dots i_{k-1}, i_{k+1} \dots i_{l-1}, i_{l+1} \dots i_n) = tp(j_1 \dots j_{k-1}, j_{k+1} \dots j_{l-1}, j_{l+1} \dots j_n)$, and \mathcal{R} is homogeneous, so we may extend the partial elementary embedding σ that sends $i_s \mapsto j_s$ for $s \neq k, l$ to an elementary embedding f that is also defined on i_k, i_l . Replace \bar{i} with its image under f , \bar{i}' and keep \bar{j} the same. \square

Per Claim 5.2.6, rename

$$(i_1 \dots i_{k-1}, i_{k+1} \dots i_{l-1}, i_{l+1} \dots i_n) = (j_1 \dots j_{k-1}, j_{k+1} \dots j_{l-1}, j_{l+1} \dots j_n)$$

as $=(a_3, \dots, a_n)$,

$$(i_k, i_l) \text{ as } =: (i_1^*, i_2^*), \text{ and}$$

(j_k, j_l) as $=: (j_1^*, j_2^*)$.

We have the following (θ' is obtained from θ by some permutation of variables):

1. $\text{qftp}_{\{<\}}^{\mathcal{R}}(i_1^*, i_2^*/a_3, \dots, a_n) = \text{qftp}_{\{<\}}^{\mathcal{R}}(j_1^*, j_2^*/a_3, \dots, a_n) =: p'(z_1, z_2)$
2. $\text{qftp}^{\mathcal{R}}(i_1^*/a_3, \dots, a_n) = \text{qftp}^{\mathcal{R}}(j_1^*/a_3, \dots, a_n) =: u(z_1)$
3. $\text{qftp}^{\mathcal{R}}(i_2^*/a_3, \dots, a_n) = \text{qftp}^{\mathcal{R}}(j_2^*/a_3, \dots, a_n) =: v(z_2)$
4. $R(i_1^*, i_2^*)$
5. $\neg R(j_1^*, j_2^*)$
6. $\theta'(\bar{b}_{i_1^*}, \bar{b}_{i_2^*}, \bar{b}_{a_3}, \dots, \bar{b}_{a_n})$
7. $\neg\theta'(\bar{b}_{j_1^*}, \bar{b}_{j_2^*}, \bar{b}_{a_3}, \dots, \bar{b}_{a_n})$

Claim 5.2.7. $p'(z_1, z_2, a_3, \dots, a_n) \cup u(z_1) \cup v(z_2) \cup R(z_1, z_2) =: F$ is a complete quantifier-free type in $S_2^{\mathcal{R}}(\{a_3, \dots, a_n\})$, as is $p'(z_1, z_2, a_3, \dots, a_n) \cup u(z_1) \cup v(z_2) \cup \neg R(z_1, z_2) =: G$

Proof. Since ρ_g is a signature consisting of two binary relations, every quantifier free L_g -type in (x_1, \dots, x_m) is determined by the quantifier-free types of the pairs (x_i, x_j) . Moreover, the union of a complete $<$ -type on (x_1, \dots, x_m) and a complete R -type on the same variables, yields a complete quantifier-free type in L_g . This is because no quantifier-free L_g formula is obtained by a composition, but only as a conjunction of positive or negative instances of $v_1 R v_2$ and $v_1 < v_2$.

It suffices to replace parameters a_i by variables z_i in F for $3 \leq i \leq n$, adjoin the complete quantifier-free type of (a_3, \dots, a_n) , and show that the resulting type F' is a complete quantifier-free type in (z_1, \dots, z_n) . We already know that the order-type is complete, from (1), above. As for the R -type, pairs between a_i, a_j are accounted for, and (2), (3) and (4) determine all pairs between z_i and a_j , or z_i and z_j .

The argument for G proceeds similarly. □

Definition 5.2.8. Split the variables of θ so that $\theta(x; \bar{y}) = \theta(x_1; x_2, x_3, \dots, x_n)$

Claim 5.2.9. $\theta(x; \bar{y})$ has the independence property.

Proof. Fix a finite integer m and enumerate the subsets of $\{1, \dots, m\}$ as $(w_s : s \leq 2^m)$. We need only verify that there are parameters \bar{c}_t in M for $t \leq n$ and instances b_s for $s \leq 2^m$ so that

$$\theta(b_s; \bar{c}_t) \Leftrightarrow t \in w_s$$

Note that $F(z_1, z_2)$ is a complete quantifier-free type and we have both $F(i_1^*, i_2^*, a_3, \dots, a_n)$ and $\theta'(\bar{b}_{i_1^*}, \bar{b}_{i_2^*}, \bar{b}_{a_3}, \dots, \bar{b}_{a_n})$. Similarly, $G(z_1, z_2)$ is a complete quantifier-free type and we have both $G(j_1^*, j_2^*, a_3, \dots, a_n)$ and $\neg\theta'(\bar{b}_{j_1^*}, \bar{b}_{j_2^*}, \bar{b}_{a_3}, \dots, \bar{b}_{a_n})$, and $G(z_1, z_2)$ is a complete type. Thus by \mathcal{R} -indiscernibility:

1. $F(z_1, z_2) \Rightarrow \theta(\bar{b}_{z_1}, \bar{b}_{z_2}, \bar{b}_{a_3}, \dots, \bar{b}_{a_n})$
2. $G(z_1, z_2) \Rightarrow \neg\theta(\bar{b}_{z_1}, \bar{b}_{z_2}, \bar{b}_{a_3}, \dots, \bar{b}_{a_n})$

We are done if we can show that there exist values for $y_1, \dots, y_{2^m}, z_1, \dots, z_m$ satisfying

- (4) $F(y_s, z_t)$ if $t \in w_s$
- (5) $G(y_s, z_t)$ if $t \notin w_s$

as then b_s can be taken to be y_s and \bar{c}_t can be taken to be $z_t \widehat{a_3 \dots a_n}$.

The conditions from (1), (2) specify a type $\Sigma(y_1, \dots, y_{2^m}; z_1, \dots, z_m; a_3, \dots, a_n)$ which it remains to show is satisfiable in \mathcal{R} . However, this type amounts to two components:

- (5) $R(y_s, z_t) \Leftrightarrow t \in w_s$, and
- (6) (i) $p'(y_s, z_t/a_3, \dots, a_n)$
 (ii) $u(y_s)$
 (iii) $v(z_t)$

First, we need to know that substituting variables u_i for the a_i that we can find some finite ordered graph B on

$$\{y_s\}_s \cup \{z_t\}_t \cup \{u_3, \dots, u_n\} \text{ realizing } \Sigma(\bar{y}; \bar{z}; \bar{u}) \cup \{ \text{the complete quantifier-free type of } a_3, \dots, a_n \text{ in } \bar{u} \}$$

The members of (2)(i) as well as the order type of the a_i specify a partial order on B that can be extended to a linear order. The graph type of the a_i specifies graph relations among the u_i , (1) specifies graph relations among the y_s and z_t , (2)(ii) specifies graph relations between the y_s and the u_i , and (2)(iii) specifies graph relations between the z_t and the u_i . No pair is mentioned twice, so there is no inconsistency in the graph type. Any completion of the partial order to a linear order serves as the complete graph type of B .

Name the ordered graph on $\{a_3, \dots, a_n\}$ by A . It is clear that A embeds into B as B realizes the complete ordered graph type of A on u_3, \dots, u_n . By weak homogeneity, since \mathcal{R} embeds A and $B \in \mathcal{K}$, it must embed B over A . Thus Σ is realized in \mathcal{R} . \square

We have shown that for $\theta(x; \bar{y})$ some L -formula, θ has IP in $M \models T$, thus T has IP. \square

Theorem 5.2.10 (NIP Characterization Theorem). *A theory T has NIP if and only if any quantifier-free weakly-saturated ordered graph-indiscernible in a model of T is an indiscernible sequence.*

The proof is by Lemmas 5.2.1 and 5.2.3.

Bibliography

- [AH78] F. Abramson and L. Harrington. Models without indiscernibles. *Journal of Symbolic Logic*, 43(3):572–600, 1978.
- [Bla81] A. Blass. A partition theorem for perfect sets. *Proceedings of the American Mathematical Society*, 82(2):271–277, 1981.
- [BSon] J. Baldwin and S. Shelah. The stability spectrum for classes of atomic models, (in preparation).
- [CK73] C. C. Chang and H. J. Keisler. *Model Theory*. North-Holland, Amsterdam, 1973.
- [DS04] M. Džamonja and S. Shelah. On \triangleleft^* -maximality. *Annals of Pure and Applied Logic*, 125(1-3):119–158, 2004.
- [Fra54] R. Fraïssé. Sur l’extension aux relations de quelques propriétés des ordres. *Annales Scientifiques de l’École Normale Supérieure*, 71(3):363–388, 1954.
- [Hod93] W. Hodges. *Model Theory*. Cambridge University Press, Cambridge, U.K., 1993.
- [Lac87] A. H. Lachlan. Homogeneous structures. In *Proceedings of the International Congress of Mathematicians 1986*, pages 314–321, Providence, RI, 1987. American Mathematical Society.
- [LS03] M. C. Laskowski and S. Shelah. Karp complexity and classes with the independence property. *Annals of Pure and Applied Logic*, 120:263–283, 2003.
- [Mar02] D. Marker. *Model Theory: An Introduction*. Springer-Verlag, New York, NY, 2002.
- [Mil79] Keith R. Milliken. A Ramsey theorem for trees. *Journal of Combinatorial Theory. Series A*, 26(3):215–237, 1979.
- [Mil81] K. R. Milliken. A partition theorem for the infinite subtrees of a tree. *Transactions of the American Mathematical Society*, 263(1):137–148, 1981.
- [Mos56] Ehrenfeucht Mostowski. Models of axiomatic theories admitting automorphisms. *Fundamenta Mathematica*, 43:50–68, 1956.

- [Neš05] J. Nešetřil. Homogeneous structures and ramsey classes. *Combinatorics, Probability and Computing*, 14:171–189, 2005.
- [NR77] J. Nešetřil and V. Rödl. Partitions of finite relational and set systems. *Journal of Combinatorial Theory (A)*, 22:289–312, 1977.
- [She78] S. Shelah. *Classification Theory and the number of non-isomorphic models*. North-Holland, Amsterdam-New York, 1978.
- [She09] S. Shelah. Nonstructure theory (in preparation), 2009.
- [Tod10] S. Todorcevic. *Introduction to Ramsey spaces*, volume 174 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2010.