The uncountable spectra of countable theories

Bradd Hart
Ehud Hrushovski
Michael C. Laskowski*

April 11, 2000

Abstract

Let T be a complete, first-order theory in a finite or countable language having infinite models. Let $I(T,\kappa)$ be the number of isomorphism types of models of T of cardinality κ. We denote by μ (respectively $\hat{\mu}$) the number of cardinals (respectively infinite cardinals) less than or equal to κ.

Theorem $I(T,\kappa)$, as a function of $\kappa > \aleph_0$, is the minimum of 2^κ and one of the following functions:

1. 2^κ;
2. the constant function 1;
3. $\left\{ \begin{array}{ccl} |\hat{\mu}^n/{\sim_G}| - |(\hat{\mu}-1)^n/{\sim_G}| & \hat{\mu} < \omega; & \text{for some } 1 < n < \omega \text{ and} \\ \hat{\mu} & \hat{\mu} \geq \omega; & \text{some group } G \leq Sym(n) \end{array} \right.$
4. the constant function $\beth_2$;
5. $\beth_{d+1}(\mu)$ for some infinite, countable ordinal d;
6. $\sum_{i=1}^{d} \Gamma(i)$ where d is an integer greater than 0 (the depth of T);

$$\Gamma(i) \text{ is either } \beth_{d-i-1}(\mu^{\hat{\mu}}) \text{ or } \beth_{d-i}(\mu^{\sigma(i)} + \alpha(i))$$

where $\sigma(i)$ is either $1, \aleph_0$ or $\beth_1$, and $\alpha(i)$ is 0 or $\beth_2$; the first possibility for $\Gamma(i)$ can occur only when $d - i > 0$.

*The first author was partially supported by the NSERC. The third author was partially supported by the NSF, Grant DMS-9704364. All the authors would like to thank the MSRI for their hospitality.
AMS Subject Classification: 03C45

The cases (2), (3) of functions taking a finite value were dealt with by Morley and Lachlan. Shelah showed (1) holds unless a certain structure theory (superstability and extensions) is valid. He also characterized (4) and (5), and showed that in all other cases, for large values of κ, the spectrum is given by $\beth_{d-1}(\mu^{<\sigma})$ for a certain σ, the "special number of dimensions".

The present paper shows, using descriptive set theoretic methods, that the continuum hypothesis holds for the special number of dimensions. Shelah's superstability technology is then used to complete the classification of the all possible uncountable spectra.

1 Introduction

A *theory* is a set of sentences - finite statements built from the function and relation symbols of a fixed language by the use of the Boolean connectives ("and", "not", etc.) and quantifiers ("there exists" and "for all"). The usual axioms for rings, groups and real closed fields are examples of theories. Associated to any theory is its class of *models*. A model of a theory is an algebraic structure that satisfies each of the sentences of the theory. For a theory T and a cardinal κ, $I(T,\kappa)$ denotes the number of isomorphism classes of models of T of size κ. The *uncountable spectrum* of a theory T is the map $\kappa \mapsto I(T,\kappa)$, where κ ranges over all uncountable cardinals. As examples, any theory of algebraically closed fields of fixed characteristic has $I(T,\kappa) = 1$ for all uncountable κ, while any completion of Peano Arithmetic has $I(T,\kappa) = 2^\kappa$.

With Theorem 6.1 we enumerate the possible uncountable spectra of complete theories in a countable language. Examples of theories possessing each of these spectra are given in [7].

Starting in 1970, Shelah placed the uncountable spectrum problem at the center stage of model theory. His goal was to develop a taxonomy of complete theories in a fixed countable language. Shelah's thesis was that the equivalence relation of 'having the same uncountable spectrum' induces a partition of the space of complete theories that is natural and useful for other applications. Over a span of twenty years he realized much of this research program. In addition to the results mentioned in the abstract, he showed that the uncountable spectrum $I(T,\kappa)$ is non-decreasing for all complete theories T and that the divisions between spectra reflect structural properties. Shelah found a number of dividing lines among complete theories. The definitions of these dividing lines do not mention uncountable objects, but collectively they form an important distinction between the associated classes of uncountable models. On one hand, he showed that if a theory is on the 'non-structure' side of at least one of these lines then models of the theory embed a certain amount of set theory; as a consequence their spectrum is maximal (i.e., $I(T,\kappa) = 2^\kappa$ for all uncountable κ). This is viewed as a negative feature, ruling out the possibility of a reasonable structure theorem for the class of models of the theory.

On the other hand, for models of theories that are on the 'structure' side of each of these lines, one can associate a system of combinatorial geometries. The isomorphism type of a model of such a theory is determined by local information (i.e., behavior of countable substructures) together with a system of numerical invariants (i.e., dimensions for the corresponding geometries). It follows that the uncountable spectrum of such a theory cannot be maximal. Thus, the uncountable spectrum of a complete theory in a countable language is not maximal if and only if every model of the theory is determined up to isomorphism by a well-founded, independent tree of countable substructures.

Our work is entirely contained in the stability theoretic universe created by Shelah. We offer three new dividing lines on the space of complete theories in a (fixed) countable language (see Definitions 3.2 and 5.23) and show that these divisions, when combined with those offered by Shelah, are sufficient to characterize the uncountable spectra of all such theories. These new divisions partition the space of complete theories into Borel subsets (with respect to the natural topology on the space). The first two of these divisions measure how far the theory is from being totally transcendental, while the third division makes a much finer distinction between two spectra.

A still finer analysis in terms of geometric stability theory is possible. We mention for instance that any model of a complete theory whose uncountable spectrum is $\min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha+\omega|+\beth_2)\}$ for some finite $d>1$ interprets an infinite group. This connection turns out not to be needed for the present results, and will be presented elsewhere.

The main technical result of the paper is the proof of Theorem 3.3, which asserts that the embeddability of certain countable configurations of elements into some model of the theory gives strong lower bounds on the uncountable spectrum of the theory. The proof of this theorem uses techniques from descriptive set theory along with much of the technology developed to analyze superstable theories.

We remark that the computation of $I(T, \aleph_0)$ is still open. To wit, it remains unknown whether any countable, first-order theory T has $I(T, \aleph_0) = \aleph_1$, even when the continuum hypothesis fails. Following our work this instance is the only remaining open question regarding the possible values of $I(T, \kappa)$.

2 Background

Work on the spectrum problem is quite old. Morley's categoricity theorem [14], which asserts that if $I(T, \kappa) = 1$ for some uncountable cardinal κ, then $I(T, \kappa) = 1$ for all uncountable κ is perhaps the most familiar computation of a spectrum. However, some work on the spectrum problem predates this. If T has an infinite model, then for every infinite cardinal κ the inequality $1 \leq I(T, \kappa) \leq 2^\kappa$ follows easily from the Löwenheim-Skolem theorem. Improving on this, Ehrenfeucht [4] discovered the notion of what is now called an unstable theory and showed that $I(T, \kappa) \geq 2$ for certain uncountable κ whenever T is unstable.

One cannot overemphasize the impact that Shelah has had on the the uncountable spectrum problem. Much of the creation of the subfield of stability theory is singlehandedly due to him and was motivated by this problem. The survey of definitions and theorems of his that follow establish the framework for this paper and indicate why it is sufficient for us to work in the very restrictive setting of classifiable theories of finite depth.

Call a complete theory T with an infinite model *classifiable* if it is superstable, has prime models over pairs, and does not have the dimensional order property. The following two theorems of Shelah indicate that the this notion is a very robust dividing line among the space of complete theories.

Theorem 2.1 *If a countable theory T is not classifiable then $I(T, \kappa) = 2^\kappa$ for all $\kappa > \aleph_0$.*

Proof. If T is not superstable then the spectrum of T is computed in VIII 3.4 of [18]; this is the only place where this spectrum is computed for all uncountable cardinals κ. Hodges [9] contains a very readable proof of this for regular cardinals. Shelah's computation of the spectrum of a superstable theory with the dimensional order property is given in X 2.5 of [18]. More detailed proofs are given in Section 3 of Chapter XVI of Baldwin [1] and Theorem 2.3 of

Harrington-Makkai [5]. Under the assumptions that T is countable, superstable, and does not have the dimensional order property, the property of prime models over pairs is equivalent to T not having the omitting types order property (OTOP). That the omitting types order property implies that T has maximal spectrum was proved by Shelah in Chapter 12, Section 4 of [18]. Another exposition of this fact is given in [6]. ∎

In order to state the structural consequences of classifiability we state three definitions.

Definition 2.2 1. M is an *na-substructure* of N, $M \subseteq_{na} N$, if $M \subseteq N$ and for every formula $\varphi(x, y)$, tuple a from M and finite subset F of M, if N contains a solution to $\varphi(x, a)$ not in M then M contains a solution to $\varphi(x, a)$ that is not algebraic over F.

2. An *ω-tree* $(I, \lessdot)$ is a partial order that is order-isomorphic to a nonempty, downward closed subtree of ${}^{<\omega}J$ for some index set J, ordered by initial segment. The root of I is denoted by $\langle\rangle$ and for $\eta \neq \langle\rangle$, η^- denotes the (unique) predecessor of η in the tree.

3. An *independent tree of models* of T is a collection $\{M_\eta : \eta \in I\}$ of models of T indexed by an ω-tree I which is independent with respect to the order on I.

4. A *normal tree of models* of T is a collection $\{M_\eta : \eta \in I\}$ of models of T indexed by an ω-tree I satisfying:

   - $\eta \lessdot \nu \lessdot \tau$ implies $M_\tau/M_\nu \perp M_\eta$;
   - for all $\eta \in I$, $\{M_\nu : \eta = \nu^-\}$ is independent over M_η.

5. A *tree decomposition* of N is a normal tree of models $\{M_\eta : \eta \in I\}$ with the properties that, for every $\eta \in I$, M_η is countable, $M_\eta \subseteq_{\mathrm{na}} N$ and if $\eta \lessdot \nu$ implies $M_\eta \subseteq_{\mathrm{na}} M_\nu$ and $wt(M_\nu/M_\eta) = 1$.

Theorem 2.3 *1. Any normal tree of models is an independent tree of models.*

2. If T is classifiable then there is a prime model over any independent tree of models of T.

3. *Every model N of a classifiable theory is prime and minimal over any maximal tree decomposition contained in N.*

Proofs. The proof of (1) is an exercise in tree manipulations and orthogonality. Details can be found in Chapter 12 of [18] or Section 3 of Harrington-Makkai [5].

The proof of (2) only relies on T having prime models over pairs. Its proof can be found in Chapter 12 of [18] or in [6].

The proof of (3) is more substantial. In [18] Shelah proves that any model of a classifiable theory has a number of tree decompositions of various sorts. However, the fact that a model of a classifiable theory admits a decomposition using countable, *na*-submodels is the content of Theorem C of Shelah-Buechler [19]. ∎

The two parts of Theorem 2.3 provide us with a method of producing upper bounds on spectra. Namely, $I(T, \kappa)$ is bounded above by the number of labelled ω-trees of size κ. Since the components of the tree decompositions are countable, we may assume that the set of labels has size at most $2^{\aleph_0}$. In Section 5 we obtain better upper bounds in a number of cases by adding more structural information which decrease the set of labels. However, at this point, there are still too many ω-trees of size κ to obtain a reasonable upper bound on $I(T, \kappa)$. A further reduction is available by employing the following additional definitions and theorems of Shelah.

Definition 2.4 An ω-tree $(I, \lessdot)$ is *well-founded* if it does not have an infinite branch. The *depth of a node* η of a well-founded tree I is defined inductively by

$$dp_I(\eta) = \sup\{dp_I(\nu) + 1 : \eta \lessdot \nu\}$$

and the *depth of I, $dp(I)$* is equal to $dp_I(\langle\rangle)$. A theory T is *deep* if some model of T has tree decomposition indexed by a non-well-founded tree. A (classifiable) theory T is *shallow* if it is not deep. The *depth $dp(T)$* of a shallow theory T is the supremum of the depths of decomposition trees of models of T.

We remark that this definition of the depth of a theory differs slightly from the one given in [18]. The following theorem of Shelah is a consequence of Theorems X 5.1, X 4.7, and X 6.1 of [18]. Other proofs appear in Harrington-Makkai [5] and Baldwin [1].

Theorem 2.5 *1. If T is classifiable and deep then $I(T,\kappa) = 2^\kappa$ for all $\kappa > \aleph_0$.*

2. If T is shallow then $dp(T) < \omega_1$ and, if $\omega \le dp(T) < \omega_1$, then

$$I(T,\aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{dp(T)+1}(|\alpha+\omega|)\}.$$

As a consequence of these results, we are justified in making the following assumption:

> *All theories in the rest of this paper are countable, classifiable and of finite depth d.*

For such theories, one obtains the naive upper bound of

$$I(T,\aleph_\alpha) \le \beth_{d-1}(|\alpha+\omega|^{2^{\aleph_0}})$$

by induction on d, simply by counting the number of labelled trees in the manner described above.

In general, obtaining lower bounds on spectra is a challenging enterprise. The difficulty is due to the fact that the tree decomposition of a model given in Theorem 2.3 is typically not canonical. The method of *quasi-isomorphisms* introduced by Shelah and streamlined by Harrington-Makkai and Baldwin-Harrington (see e.g., Section 3 of [5], Section 3 of [3], or Theorem XVII.4.7 of [1]) can be employed to show that if two models have 'sufficiently disparate' decomposition trees, then one can conclude that the models are non-isomorphic. From this, one obtains a (rather weak) general lower bound on the spectrum of a classifiable theory of finite depth $d > 1$, namely

$$I(T,\aleph_\alpha) \ge \min\{2^{\aleph_\alpha}, \beth_{d-2}(|\alpha+\omega|^{|\alpha+1|})\}.$$

A proof of this lower bound is given in Theorem 5.10(a) of [16].

Accompanying Shelah's 'top-down' analysis of the spectrum problem is work of Lachlan, Saffe, and Baldwin who computed the spectra of theories satisfying much more stringent constraints.

In [11] and [12], Lachlan classifies the spectra of all ω-categorical, ω-stable theories. We use this classification verbatim at the end of Section 5. With [16], Saffe computes the uncountable spectra of all ω-stable theories. A more detailed account of the analysis of this case is given by Baldwin in [1]. Aside from a few specific facts, we do not make use of this analysis here.

The history of this paper is modestly complicated. Shelah knew the value of $I(T, \aleph_\alpha)$ for large values of α (reported in Chapter XIII of [18]) modulo a certain continuum hypothesis-like question known as the SND (special number of dimensions) problem. In 1990, Hrushovski solved the SND problem; he also announced a calculation of the uncountable spectra. This calculation included a gap related to the behavior of non-trivial types but nonetheless a framework for the complete computation was introduced. The project lay fallow for several years before being taken up by the current authors; their initial work was reported in [7]. Hart and Laskowski recast the superstructure of the argument in a way to avoid the earlier gap and the work was completed while the three authors were in residence at MSRI.

We assume that the reader is familiar with stability theory. All of the necessary background can be found in the union of the texts by Baldwin [1] and Pillay [15]. Our notation is consistent with these texts. In Sections 3–6 we assume that we have a fixed, classifiable theory in a countable language of some finite depth. (The sufficiency of this assumption is explained in Section 2.) We work in T^{eq}. In particular, every type over an algebraically closed set is stationary. As well, throughout the paper we denote finite tuples of elements by singletons. As usual in the study of stable theories, we assume that we are working within the context of a large, saturated model $\mathcal{C}$ of T. All sets are assumed to be subsets of $\mathcal{C}$, and all models are assumed to be elementary submodels of $\mathcal{C}$. In particular, the notation $M \subseteq N$ implies $M \preceq N$.

3 More dividing lines

In this section we provide a local analysis of a classifiable theory of finite depth d. In particular, we ask how far away the theory is from being totally transcendental. Towards this end, we make the following definitions.

Definition 3.1 1. For any $n \leq dp(T)$, a *chain of length n, $\overline{\mathcal{M}}$,* is a sequence $M_0 \subseteq \ldots \subseteq M_{n-1}$ of n countable models of T, where M_{i+1}/M_i has weight 1 and $M_{i+1}/M_i \perp M_{i-1}$ for $i > 0$.

2. A chain $\overline{\mathcal{M}}$ is an *na-chain* if, in addition, each $M_i \subseteq_{na} \mathcal{C}$.

3. For $\overline{\mathcal{M}}$ a chain of length n, the set of *relevant regular types* is the set

$$R(\overline{\mathcal{M}}) = \{p \in S(M_{n-1}) : p \text{ is regular and } p \perp M_{n-2}\}.$$

   When $n = 1$, $R(\overline{\mathcal{M}})$ is simply the set of regular types over M_0.

4. A type $p \in R(\overline{\mathcal{M}})$ is *totally transcendental (t.t.) over* $\overline{\mathcal{M}}$ if there is a strongly regular $q \in R(\overline{\mathcal{M}})$, $q \not\perp p$ with a prime model $\overline{\mathcal{M}}(q)$ over M_{n-1} and any realization of q.

It is clear that the notion of a relevant type being t.t. depends only on its non-orthogonality class. Our new dividing lines concern the presence or absence of a relevant type that fails to be t.t. and whether there is a trivial 'bad' type.

Definition 3.2 A theory T is *locally t.t. over* $\overline{\mathcal{M}}$ if every type in $R(\overline{\mathcal{M}})$ is t.t. over $\overline{\mathcal{M}}$. We say T admits a *trivial failure* (of being t.t.) *over* $\overline{\mathcal{M}}$ if some trivial type $p \in R(\overline{\mathcal{M}})$ is not t.t. over $\overline{\mathcal{M}}$.

Our notation is consistent with standard usage, as any totally transcendental theory is locally t.t. over any chain. The heart of the paper will be devoted to showing that these conditions provide dividing lines for the spectra. In particular, the proof of the lower bounds offered below follows immediately from 3.23, 3.27, and 5.10.

Theorem 3.3 *1. If T is not locally t.t. over some chain of length n then*

$$I(T, \aleph_\alpha) \geq \begin{cases} \min\{2^{\aleph_\alpha}, \beth_2\} & \text{if } n = 1 \\ \min\{2^{\aleph_\alpha}, \beth_{n-2}(|\alpha + \omega|^{\beth_2})\} & \text{if } n > 1 \end{cases}$$

for all ordinals $\alpha > 0$.

2. If T admits a trivial failure over some chain of length n then

$$I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{n-1}(|\alpha + \omega|^{\beth_2})\}$$

for all ordinals $\alpha > 0$.

Complementing this theorem, in Subsection 3.4 we show that if T is locally t.t. over $\overline{\mathcal{M}}$, then there is a strong structure theorem for the class of weight-one models over M_{n-1} that are orthogonal to M_{n-2}, especially when the chain $\overline{\mathcal{M}}$ has length equal to the depth of the theory.

The proof of Theorem 3.3 is broken into a number of steps. For the most part, the two parts of the theorem are proved in parallel. In Subsection 3.1 we define the crucial notions of diverse and diffuse families of leaves. Then, in the next two subsections we analyze the two ways in which a theory could fail to be locally t.t. over a chain. For each of these, we will show that the failure of being locally t.t. over some chain of length n implies the existence of a diverse family of leaves of size continuum over some na-chain of length n. In addition, if there is a trivial failure of being t.t., then the family of leaves mentioned above will actually be diffuse. Then, in Section 4, we establish some machinery to build many non-isomorphic models from the existence of a diverse or diffuse family of leaves. Much of this is bookkeeping, but there are two important ideas developed there. Foremost is the Unique Decomposition Theorem (Theorem 4.1), which enables us to preserve non-isomorphism as we step down a decomposition tree. The other idea which is used is the fact that decomposition trees typically have many automorphisms. This fact implies that models that are built using such trees as skeletons have desirable homogeneity properties (see Definition 4.3). Finally, we complete the proof of Theorem 3.3 in Section 5.

3.1 Diverse and diffuse families of leaves

We begin by introducing some convenient notation for prime models over independent trees of models. First of all, suppose that N_1 and N_2 are two submodels of our fixed saturated model which are independent over a common submodel N_0. By $N_1 \oplus_{N_0} N_2$ we will mean a prime model over $N_1 \cup N_2$; this exists because we are assuming our theory has prime models over pairs. For our purposes, the exact model that we fix will not matter because we are only interested in its isomorphism type. In abstract algebraic terms, we want to think of this as an "internal" direct sum.

Now suppose that M_0 is any model and $M_0 \subseteq M_i$ for $i = 1, 2$, not necessarily independent over M_0. By $M_1 \oplus_{M_0} M_2$ we will mean the "external" direct sum i.e., we choose M_i' isomorphic to M_i over M_0 and such that M_1' is independent over M_2' over M_0 and form $M_1' \oplus_{M_0} M_2'$ in the internal sense defined above. We similarly define $\bigoplus_M \mathcal{F}$ for a family of models, each of which contains a fixed model M.

Suppose that $\mathcal{M} = \langle M_\zeta : \zeta \in I \rangle$ is an independent family of models with respect to a tree ordering $\langle I, \lessdot \rangle$ such that if $\eta \lessdot \zeta \in I$ then $M_\eta \subseteq M_\zeta$. We say that a family of elementary maps $\langle f_\zeta : \zeta \in I \rangle$ is *compatible* with $\mathcal{M}$ if whenever $\zeta \in I$ then $\mathrm{dom}(f_\zeta) = M_\zeta$ and if $\eta \lessdot \zeta \in I$ then $f_\zeta \restriction_{M_\eta} = f_\eta$.

Definition 3.4 If $\overline{\mathcal{M}}$ is an *na*-chain of length n, then the *set of leaves of* $\overline{\mathcal{M}}$*, Leaves*$(\overline{\mathcal{M}})$, is a set containing one representative up to isomorphism over $\overline{\mathcal{M}}$ of all chains $\overline{\mathcal{N}}$ of length $(n+1)$ extending $\overline{\mathcal{M}}$. If $\overline{\mathcal{M}}$ is an *na*-chain of length n and $Y \subseteq Leaves(\overline{\mathcal{M}})$, then an $(\overline{\mathcal{M}}, Y)$*-tree* is an independent tree of models $\mathcal{M} = \langle M_\zeta : \zeta \in I \rangle$ where I has height at most $n+1$ together with a distinguished copy of $\overline{\mathcal{M}}$ and a family of elementary maps $\langle f_\zeta : \zeta \in I \rangle$ compatible with $\mathcal{M}$ such that f_ζ maps M_ζ onto $N_{lg(\zeta)}$ for some $\overline{\mathcal{N}} \in Y$. (In particular, note that if $lg(\zeta) \leq n$ then f_ζ maps onto $M_{lg(\zeta)}$.) An $(\overline{\mathcal{M}}, Y)$-*model* is a model which is prime over an $(\overline{\mathcal{M}}, Y)$-tree; the copy of $\overline{\mathcal{M}}$ in this tree will be distinguished as a chain of submodels of this $(\overline{\mathcal{M}}, Y)$-model.

We make an important convention: Suppose N_1 and N_2 are two $(\overline{\mathcal{M}}, Y)$-models and we wish to form $N_1 \oplus_{M_k} N_2$ for some $k < n$. We declare that this sum is an "external" sum as discussed above. If we wish to view this new model as an $(\overline{\mathcal{M}}, Y)$-model, we need to determine which distinguished copy of $\overline{\mathcal{M}}$ will now be distinguished. Our convention will be that it will be the one from the left most summand.

As notation, if Z is a set of $\overline{\mathcal{M}}$-leaves then

$$N^*(Z) = \bigoplus_{M_{n-1}} \{N : \overline{\mathcal{N}} \in Z \text{ where } \overline{\mathcal{N}}(n) = N\}.$$

We next isolate the two crucial properties of a set Y of leaves. In Section 5 we will show the effects on lower bound estimates for spectra given that there are large families of leaves with these properties. Lemma 4.2 of Section 4 will show that a diffuse family is diverse.

Definition 3.5 Let $\overline{\mathcal{M}}$ be an *na*-chain of length n.

1. A set $Y \subseteq Leaves(\overline{\mathcal{M}})$ is *diffuse* if

$$N \oplus_{M_{n-1}} V \ncong_V N' \oplus_{M_{n-1}} V$$

for all distinct $N, N' \in Y$ and any $(\overline{\mathcal{M}}, Y)$-model V.

2. A set $Y \subseteq Leaves(\overline{\mathcal{M}})$ is *diverse* if

$$N^*(Z_1) \oplus_{M_{n-2}} V \ncong N^*(Z_2) \oplus_{M_{n-2}} V$$

over V and M_{n-1} for all distinct $Z_1, Z_2 \subseteq Y$ and any $(\overline{\mathcal{M}}, Y)$-model V. If $n = 1$ we omit the model V and the condition becomes: for distinct $Z_1, Z_2 \subseteq Y$,

$$N^*(Z_1) \ncong_{M_0} N^*(Z_2)$$

The following lemma provides us with an easy way of producing a diffuse family of leaves.

Lemma 3.6 *Suppose that $\{p_i : i \in I\}$ is a set of pairwise orthogonal types in $R(\overline{\mathcal{M}})$, and that $\{N_i : i \in I\} \subseteq Leaves(\overline{\mathcal{M}})$ is a set of models such that N_i is dominated by a realization of p_i over M_{n-1}. Then $\{N_i : i \in I\}$ is diffuse.*

Proof. In fact, for any model V containing M_{n-1}, if $N_i \oplus_{M_{n-1}} V \cong_V N_j \oplus_{M_{n-1}} V$ then p_i is not orthogonal to p_j so $i = j$. ∎

We conclude this subsection by introducing the notion of a *special* type, showing that they are present in every non-orthogonality class of $R(\overline{\mathcal{M}})$, and proving two technical lemmas that will be used in Subsections 3.2 and 3.3 to obtain diverse or diffuse families of leaves.

As notation for a regular strong type p, let $[p]$ denote the collection of strong types (over any base set) non-orthogonal to p. We let $R^\infty([p]) = \min\{R^\infty(q) : q \in [p]\}$. It is easy to see that $R^\infty([p])$ is the smallest ordinal α such that p is non-orthogonal to some formula θ of R^∞-rank α, which is also the smallest ordinal β such that p is foreign to some formula ψ of R^∞-rank β. The following lemma is general and holds for any superstable theory.

Lemma 3.7 *Let $M^- \subseteq M$ be models of a superstable theory and suppose that a regular type $p \in S(M)$ is orthogonal to M^-, but is non-orthogonal to some $\theta(x, b)$, where $R^\infty(\theta(x, b)) = R^\infty([p])$. Then $\sigma(p)$ is foreign to $\theta(x, b)$ for any automorphism $\sigma \in Aut_{M^-}(\mathcal{C})$ satisfying $\sigma(M) \underset{M^-}{\perp} b$.*

Proof. Suppose that $\theta(c, b)$ holds and let X be any set. We claim that $\mathrm{tp}(c/Xb) \perp \sigma(p)$. To see this, first note that $\sigma(p) \perp M^- b$, since we assumed

$\sigma(p) \perp M^-$ and $\sigma(M) \underset{M^-}{\perp} b$. There are now two cases. If $R^\infty(c/Xb) = R^\infty([p])$, then $\mathrm{tp}(c/Xb)$ does not fork over b, so $\sigma(p) \perp \mathrm{tp}(c/Xb)$ by the note above. On the other hand, if

$$R^\infty(c/Xb) < R^\infty([p]) = R^\infty([\sigma(p)]),$$

then $\mathrm{tp}(c/Xb) \perp \sigma(p)$ as well. ∎

We now turn our attention back to a particular chain $\overline{\mathcal{M}}$ of length n.

Definition 3.8

1. If $p \in R(\overline{\mathcal{M}})$ then q is a *tree conjugate* of p if for some $k < n-1$ there is an automorphism σ fixing M_k pointwise such that $\sigma(p) = q$ and $\sigma(M) \underset{M_k}{\perp} M_{n-1}$. (If $n = 1$ then p does not have any tree conjugates.)
2. A type $p \in R(\overline{\mathcal{M}})$ is *special via* $\varphi(x,e)$ if $\varphi(x,e)$ is p-simple, $\varphi(x,e) \in p$, and the tree conjugates of p are foreign to $\varphi(x,e)$. A type $p \in R(\overline{\mathcal{M}})$ is *special* if it is special via some formula.

As promised, we show that special types exist in every non-orthogonality class of $R(\overline{\mathcal{M}})$. The proof of the following lemma is an adaptation of the proof of Lemma 8.2.19 of [15], which in turn is adapted from arguments in [18].

Lemma 3.9 *Let a be any realization of a type $p \in R(\overline{\mathcal{M}})$, where $\overline{\mathcal{M}}$ is a chain of length n. There is an $a' \in \mathrm{acl}(M_{n-1}a) \setminus M_{n-1}$ such that $p' = \mathrm{tp}(a'/M_{n-1}) \in R(\overline{\mathcal{M}})$ is special.*

Proof. Choose a formula $\theta(x,b)$ non-orthogonal to p with $R^\infty(\theta(x,b)) = R^\infty([p])$ and choose a (regular) type q non-orthogonal to p containing $\theta(x,b)$. Choose a set $A \supseteq M_{n-1}$ and a non-forking extension $r \in S(A)$ of q such that $a \underset{M_{n-1}}{\perp} A$, r is stationary and there is a realization c of r with $a \underset{A}{\not\perp} c$. Now choose $a' \in Cb(\mathrm{stp}(bc/Aa)) \setminus \mathrm{acl}(A)$ Since $\mathrm{tp}(a/A)$ does not fork over M_{n-1}, $a' \in \mathrm{acl}(M_{n-1}a) \setminus M_{n-1}$, hence $p' = \mathrm{tp}(a'/M_{n-1})$ is regular and non-orthogonal to p. It remains to find an $L(M_{n-1})$-formula witnessing that $\mathrm{tp}(a'/M_{n-1})$ is special. Choose a Morley sequence $I = \langle b_n c_n : n \in \omega\rangle$ in $\mathrm{stp}(bc/Aa)$ with $b_0c_0 = bc$. Since $a' \in \mathrm{dcl}(\overline{b}\overline{c})$ for some initial segment of I,

$a' = f(\overline{b}, \overline{c})$ for some $\emptyset$-definable function f. Note that $a \underset{M_{n-1}}{\perp} A\overline{b}$. Choose a finite $A_0 \subseteq A$ on which everything is based, and let $w = \mathrm{tp}(A_0\overline{b}/M_{n-1})$ and let $\varphi(x, e)$ be the $L(M_{n-1})$-formula

$$\varphi(x, e) := d_w \overline{y}\, \exists \overline{z} \left(\bigwedge_i \theta(z_i, y_i) \wedge x = f(\overline{y}, \overline{z}) \right).$$

For notation, assume that w is based on e. Clearly $\varphi(a', e)$ holds and it follows easily that $\varphi(x, e)$ is p'-simple. Hence if $n = 1$ we are done. So suppose that $n > 1$ and $\varphi(a^*, e)$ holds. Fix $k < n - 1$ and an automorphism σ over M_k such that $\sigma(M_{n-1}) \underset{M_k}{\perp} M_{n-1}$. Choose $A^*\overline{b}^*$ realizing $w|M_{n-1}\sigma(M_{n-1})$ and choose $\overline{c}^*$ such that $a^* = f(\overline{b}^*, \overline{c}^*)$ and $\theta(c_i^*, b_i^*)$ holds for each i. Since p' is not orthogonal to p, $\theta(x, b_i^*)$ is non-orthogonal to p' and has least R^∞-rank among all formulas non-orthogonal to p'. Thus, as $b_i^* \underset{M_k}{\perp} \sigma(M_{n-1})$, it follows from Lemma 3.7 that $\sigma(p')$ is foreign to $\theta(x, b_i^*)$ for all i. Hence $\mathrm{tp}(a^*/e)$ is hereditarily orthogonal to $\sigma(p')$. That is, the tree conjugates of p' are foreign to $\varphi(x, e)$. ∎

The following lemma is simply a restatement of Lemma 3.9 together with an application of the Open Mapping Theorem.

Lemma 3.10 *Suppose that $\overline{\mathcal{M}}$ is a chain of length $n > 1$ and $p \in R(\overline{\mathcal{M}})$ is special via φ. Further, suppose that a model $N/M_{n-1} \perp M_{n-2}$, and U is dominated over W by an independent set of conjugates of p. Then any realization of φ in $N \oplus_{M_{n-2}} U$ is contained in $N \oplus_{M_{n-2}} W$.*

Proof. Suppose that $a \in N \oplus_{M_{n-2}} U$ realizes φ. Then $\mathrm{tp}(a/NU)$ is isolated. As the tree conjugates of p are foreign to φ,

$$a \underset{NW}{\perp} d$$

for any $d \in U$ realizing a conjugate of p. Thus $a \underset{NW}{\perp} U$, since U is dominated over W by an independent set of realizations of conjugates of p. Hence $\mathrm{tp}(a/NW)$ is isolated, which implies that $a \in N \oplus_{M_{n-2}} W$. ∎

If, in addition, our special type is trivial then we can say more. The lemma that follows is one of the main reasons why we are able to build a diffuse family instead of a diverse family when the 'offending' type is trivial.

Lemma 3.11 *Suppose that $\overline{\mathcal{M}}$ is a chain of length n, $N \in Leaves(\overline{\mathcal{M}})$, $p \in R(\overline{\mathcal{M}})$ is any trivial, special type via φ such that some realization of p dominates N over M_{n-1}. Suppose further that $M_{n-1} \subseteq W \subseteq U$, where U is dominated by W-independent realizations of non-forking extensions of tree conjugates of p over W and regular types not orthogonal to p. Then if an element a satisfies φ, $a \in N \oplus_{M_{n-1}} U$, $\mathrm{tp}(a/M_{n-1})$ is regular and $a \underset{M_{n-1}}{\perp} U$ then $tp(a/NW)$ is isolated.*

Proof. Note first that the assumptions imply that $tp(a/M_{n-1})$ is not orthogonal to p. Since p is trivial and $a \underset{M_{n-1}}{\perp} U$, it follows that a forks with N over M_{n-1}. Since $\varphi(a)$ holds, it follows that $tp(a/N)$ (and any extension of this type) is orthogonal to p and all tree conjugates of p. It follows that $a \underset{NW}{\perp} U$. Since $tp(a/NU)$ is isolated, it follows that $tp(a/NW)$ is isolated. ∎

3.2 The existence of prime models

Throughout this section we assume that there is some chain $\overline{\mathcal{M}}$ of length n together with a type $r \in R(\overline{\mathcal{M}})$ for which there is no prime model over $M_{n-1}c$, where c is a realization of r. By choosing an extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ which is minimal in a certain sense, we will construct a highly disparate family $Y = \{N_\eta : \eta \in {}^{\omega}2\}$ of $Leaves(\overline{\mathcal{M}}')$ and a family of types $\{s_\eta(x, z)\}$ over M'_{n-1} that will witness this disparity. Then, following the construction of the family in Proposition 3.13, we argue in Corollary 3.23 that if the original type was special, then this set of leaves is diverse. Further, if in addition the type r is trivial, we show that this family is actually diffuse. We begin by specifying what we mean by a free extension of a chain.

Definition 3.12 An na-chain $\overline{\mathcal{M}}'$ *freely extends* the chain $\overline{\mathcal{M}}$ if both chains have the same length (say n), $M_0 \subseteq M'_0$, $M'_{i-1} \cup M_i \subseteq M'_i$, and $M'_{i-1} \underset{M_{i-1}}{\perp} M_i$ for all $0 < i < n$.

It is readily checked that if $r \in R(\overline{\mathcal{M}})$ is special, then for any free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$, the non-forking extension of r to $R(\overline{\mathcal{M}}')$ will be special as well. The bulk of this subsection is devoted to the proof of Proposition 3.13. In order to state the proposition precisely, we require some notation.

Suppose that Y is a family of $Leaves(\overline{\mathcal{M}})$ that is indexed by ${}^{\omega}2$. Fix an $(\overline{\mathcal{M}}, Y)$-model V and a sequence $\eta \in {}^{\omega}2$. We can decompose V into two pieces,

$$V = V_\eta \bigoplus_{W_V} V_{\text{no } \eta}$$

where W_V is the model prime over the tree truncated below level n,

$$V_\eta = \bigoplus_{W_V} \{N_i : N_i \text{ conjugate to } N_\eta\}$$

and

$$V_{\text{no } \eta} = \bigoplus_{W_V} \{N_i : N_i \text{ not conjugate to } N_\eta\}.$$

Proposition 3.13 *Assume that $\overline{\mathcal{M}}$ is a chain of length n and $r \in R(\overline{\mathcal{M}})$ is a type such that there is no prime model over $M_{n-1}c$ for c a realization of r. Then there is a free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ and a family $Y = \{N_\eta : \eta \in {}^{\omega}2\}$ of $Leaves(\overline{\mathcal{M}}')$, along with a family $\{s_\eta(x, z) : \eta \in {}^{\omega}2\}$ of types over M'_{n-1} such that each N_η realizes s_η, yet for any $(\overline{\mathcal{M}}', Y)$-model V, V omits $s_\eta(x, c^*)$ for all $c^* \in V_{\text{no } \eta}$ where c^* realizes $r|M'_{n-1}$.*

Proof. The lack of a prime model over $M_{n-1}c$ implies the lack of a prime model over $acl(M_{n-1}c)$. Look at all possible quadruples $(\overline{\mathcal{M}}', r', \theta, \psi)$ where $\overline{\mathcal{M}}'$ is a free extension of $\overline{\mathcal{M}}$, r' is the non-forking extension of r to M'_{n-1}, and (fixing a realization c of r' and letting $C = acl(M'_{n-1}c)$) $\theta(x)$ is a formula in $L(C)$ with no isolated extensions over C and ψ is a formula in $L(M'_{n-1})$ such that every realization of $\theta(x)$ is ψ-internal. Among all such quadruples, fix one for which $R^\infty(\psi)$ is minimal. To ease notation in what follows, we denote $\overline{\mathcal{M}}'$ by $\overline{\mathcal{M}}$ and r' by r. As well, fix a realization c of r and let $C = acl(M_{n-1}c)$.

We will construct the family of leaves and types simultaneously by building successively better finite approximations. We adopt the notation of forcing (i.e., partial orders and filters meeting collections of dense sets). However, as we will only insist that our filters meet countably many dense sets, a 'generic object' will already be present in the ground model and each of these constructions could just as easily be considered as a Henkin construction.

Our set of forcing conditions $\mathcal{P}$ is the set of finite functions $p : {}^{<\omega}2 \to L(C)$ such that $p(\eta) \vdash p(\mu)$ whenever $\mu \lessdot \eta$. (In what follows, we write

p_η for $p(\eta)$). For $p, q \in \mathcal{P}$ we say $p \leq q$ if and only if $\mathrm{dom}(p) \supseteq \mathrm{dom}(q)$ and $p_\eta \vdash q_\eta$ for all $\eta \in \mathrm{dom}(q)$. If $G \subseteq \mathcal{P}$ is any filter and $\eta \in {}^\omega 2$, let $p_\eta(G) = \{\theta \in L(C) : p(\mu) \vdash \theta$ for some $p \in G$ and some $\mu \lessdot \eta\}$.

We first list a set of basic density conditions that we want our filter to meet. Let $p^* \in \mathcal{P}$ have domain $\{\langle\rangle\}$ and range $\{\theta(x)\}$.

1. For each $\varphi \in L(C)$ and each $\eta \in {}^{<\omega}2$,

$$D_{\varphi,\eta} = \{p \in \mathcal{P} : p(\eta) \vdash \varphi \text{ or } p(\eta) \vdash \neg\varphi\};$$

2. For each $\psi(y, z) \in L(C)$ and each $\eta \in {}^{<\omega}2$,

$$D_{\exists z\psi,\eta} = \{p \in \mathcal{P} : p_\eta \vdash \neg\exists z\psi(y, z) \text{ or } p_\eta \vdash \psi(y, u) \text{ for some variable } u\}.$$

It is easy to see that every basic condition is dense in $\mathcal{P}$. As well, if G is a filter meeting all of the conditions mentioned above (i.e., $G \cap D \neq \emptyset$ for each D) then for each $\eta \in {}^\omega 2$, $p_\eta(G)$ is a complete type over C and if $\bar{b}$ is a realization of $p_\eta(G)$ then $N_\eta = C \cup \bar{b}$ is a leaf of $\overline{\mathcal{M}}$ Additionally, if $p^* \in G$ then $\mathrm{tp}(b_x/C)$ will extend $\theta(x)$ where b_x corresponds to x in $\bar{b}$. In what follows, we will take $s_\eta(x, z)$ to be $\mathrm{tp}(b_x c/M_{n-1})$.

Before stating the crucial density conditions that will ensure that the family of types satisfies the conclusion of the proposition, we pause to set notation.

Definition 3.14 A *finite approximation* $\mathcal{F}$ consists of a finite independent tree $\mathcal{N} = \langle N_\zeta : \zeta \in I\rangle$ and a family of maps $\mathbf{f} = \langle f_\zeta : \zeta \in I\rangle$ compatible with $\mathcal{N}$ such that f_ζ is an elementary isomorphism from B_ζ to $M_{l(\zeta)}$ if $l(\zeta) < n$ and to C if $l(\zeta) = n$. Let $g_\zeta = f_\zeta^{-1}$ and let $I^+ = \{i \in I : l(i) = n\}$ and $I^- = I \setminus I^+$.

$\mathcal{F}$ also comes with a number $m \in \omega$ and an assignment ρ from nodes in I^+ to elements of 2^m. For $\eta \in 2^m$, let $I_\eta = \{i \in I^+ : \rho(i) = \eta\}$. We also have a language associated with $\mathcal{F}$; this language is $L(B)$ where $B = \bigcup_\zeta N_\zeta$ and special variables x^i where $i \in I^+$ and x is a variable in the language $L(C)$.

Suppose that $p \in \mathcal{P}$. We adopt the convention that for any $\eta \in 2^{<\omega}$ $p(\eta) = p(\bar{\eta})$ where $\bar{\eta}$ is the greatest element in the $dom(p)$ less than or equal to η. In this way we may assume that $2^m \subseteq dom(p)$. Now suppose that σ

is any assignment of partial types over C to the elements of 2^m. We define $F(\sigma)$ using $\mathbf{f}$ as follows:

$$F(\sigma) = \{g_{\rho(i)}(\sigma(\rho(i))(x_1^i, \ldots, x_k^i)) : i \in I^+\}$$

where the free variables in $\sigma(\rho(i))$ are $x_1, \ldots, x_k$ respectively. A particular instance of this which will be used below is: fix $E = \langle \bar{\zeta} : \zeta \in 2^m \rangle$ where $\bar{\zeta} \geq \zeta$ and is maximal in $dom(p)$. We define $f_E(p) = F(\sigma)$ where $\sigma(\zeta) = p(\bar{\zeta})$.

We now introduce the critical set of density conditions. Each of these conditions depends on a finite approximation $\mathcal{F}$, $\eta \in 2^m$ where m is the number from $\mathcal{F}$, and formulas $\chi(y, \bar{u}_1)$ and $\varphi(x, y, \bar{u})$ where $\bar{u}$ is some finite sequence of variables x^i where $i \in I^+$ and $\bar{u}_1$ includes all the variables in $\bar{u}$ except those for which $i \in I_\eta$. Moreover, there will be at least one element of I^+ which is not labelled η; let's call one such element ϵ and identify $\overline{\mathcal{M}}$ with $\langle M_{\epsilon\restriction_i} : i < n \rangle$.

Density Condition 3.15 Let $D = D(\mathcal{F}, \eta, \chi, \varphi)$ be the set of all $p \in \mathcal{P}$ such that, for every $E = \langle \bar{\zeta} : \zeta \in 2^m \rangle$ where $\bar{\zeta} \geq \zeta$ and is maximal in $dom(p)$, either

1. $f_E(p) \vdash \exists y(\chi(y, \bar{u}_1) \wedge \neg\delta(y))$ for some $\delta \in r$ or
2. $f_E(p) \vdash \neg\exists x \exists y(\chi(y, \bar{u}_1) \wedge \varphi(x, y, \bar{u}))$ or
3. $f_E(p) \vdash \exists x \exists y(\chi(y, \bar{u}_1) \wedge \varphi(x, y, \bar{u}) \wedge \neg p_{\bar{\eta}}(x, y))$ where in $p_{\bar{\eta}}$, the variable y corresponds to c and all other variables and parameters are suppressed.

We check that if G is a filter that contains p^*, meets the basic conditions, and intersects each of the sets $D(\mathcal{F}, \eta, \chi, \varphi)$, then the sets of leaves $Y = \{N_\eta : \eta \in {}^\omega 2\}$ and types $\{s_\eta(x, c) : \eta \in {}^\omega 2\}$ satisfy the conclusion of the proposition.

Toward this end, fix an $(\overline{\mathcal{M}}, Y)$-tree, an $\eta \in {}^\omega 2$ and suppose V ($V_{\text{no } \eta}$) is prime over this tree (respectively over the leaves not conjugate to N_η). Further, suppose that $c^* \in V_{\text{no } \eta}$ realizes r and $b \in V$ realizes $s_\eta(x, c^*)$. Now in fact c^* and b are isolated over a finite part of the given $(\overline{\mathcal{M}}, Y)$-tree; c^* is isolated over this tree by a formula $\chi(y)$ and b is isolated over c^* and the tree by a formula $\varphi(x, c^*)$. We suppress the parameters from the tree to ease notation.

There is a number m such that if conjugates of N_ζ and N_μ appear in the finite tree needed to isolate c^* and b then if $\zeta \restriction_m = \mu \restriction_m$ then $\zeta = \mu$ (there are only finitely many N_ζ's involved altogether). Now this finite tree together with m, $\eta \restriction_m$, χ and φ lead naturally to a finite approximation $\mathcal{F}$ and a density condition $D = D(\mathcal{F}, \eta \restriction_m, \chi, \varphi)$ which meets the filter G. Now from the conditions for D, clearly the first condition must have failed since $\chi(y)$ implies r. The second condition must also have failed because $\varphi(x,y) \wedge \chi(y)$ is consistent. But then the third condition would have told us that $\varphi(x, c^*)$ could not even isolate $tp(b/M_{n-1}c^*)$ which contradicts the choice of φ. This final contradiction tells us that there is in fact no $b \in V$ which realizes $s_\eta(x, c^*)$.

Thus, to complete the proof, it suffices to show that each of these sets are dense below p^*. Before doing this, we pause to give a definition and prove three lemmas that are of independent interest.

Definition 3.16 We say a type $p \in S(A)$ is *d-isolated* (definitionally isolated) if for every φ there is a formula $\psi \in p$ such that if for any $q \in S(\mathcal{C})$ which does not fork over A , if $\psi \in q$ then $p|\mathcal{C}$ and q have the same φ-definition.

Lemma 3.17 *If T is countable and A is algebraically closed then the d-isolated types in $S(A)$ are dense.*

Proof. Fix a consistent formula $\theta_0(x) \in L(A)$; we will construct a d-isolated type containing θ_0 by induction. Enumerate all L-formulas as φ_i. Suppose that by induction we have defined a consistent formula $\theta_n \in L(A)$. Choose θ_{n+1} to imply θ_n with the least $R(-, \varphi_n, \aleph_0)$-rank. Since A is algebraically closed, we can assume that $Mult(\theta_{n+1}, \varphi_n, \aleph_0) = 1$. We claim that any complete type p in $S(A)$ which contains θ_n for all n is d-isolated.

To see this, pick a formula φ; $\varphi = \varphi_n$ for some n. Let $\psi = \theta_{n+1}$. Now suppose that $q \in S(\mathcal{C})$ contains ψ and does not fork over A. Now

$$R(q, \varphi, \aleph_0) = R(q \restriction_A, \varphi, \aleph_0) = R(\psi, \varphi, \aleph_0)$$

where the first equality follows from the non-forking and the second by the choice of ψ. So the φ-types of q and $p|\mathcal{C}$ must be the same for otherwise we would have a contradiction to the multiplicity of ψ being 1. Hence $p|\mathcal{C}$ and q have the same φ-definition. ∎

Lemma 3.18 *If $a_1, \ldots, a_n$ are independent over A, $tp(a_i/A)$ is d-isolated for all i with $\theta_i \in tp(a_i/A)$, $\varphi(a_1, \ldots, a_n)$ holds, and A is algebraically closed then there are $\theta_i^* \in tp(a_i/A)$ which imply θ_i such that if $b_1, \ldots, b_n$ are independent over A and $\theta_i^*(b_i)$ holds for every i then $\varphi(b_1, \ldots, b_n)$ holds.*

Proof. We prove this by induction on n. The case $n = 1$ is clear so suppose that n is greater than 1. Suppose that $\psi(y_1, \ldots, y_{n-1})$ is the φ-definition of $tp(a_n/A)$. Since $tp(a_n/A)$ is stationary, $\psi(a_1, \ldots, a_{n-1})$ holds. Now suppose that $\theta_n^*(x)$ is a formula in $tp(a_n/A)$ which is stronger than θ_n and such that any other type which contains θ_n^* has the same φ-definition as $tp(a_n/A)$. Now by induction there are $\theta_1^*, \ldots, \theta_{n-1}^*$ which hold for $a_1, \ldots, a_{n-1}$ respectively and such that if $b_1, \ldots, b_{n-1}$ are independent over A and $\theta_i^*(b_i)$ holds for all i then $\psi(b_1, \ldots, b_{n-1})$ holds. Now suppose that $b_1, \ldots, b_n$ are independent over A and $\theta_i^*(b_i)$ holds for all i. By the choice of θ_n^*, ψ is the φ-definition for $tp(b_n/Ab_1, \ldots, b_{n-1})$. By assumption, $\psi(b_1, \ldots, b_{n-1})$ holds and b_n is independent from $b_1, \ldots, b_{n-1}$ over A so $\varphi(b_1, \ldots, b_n)$ holds. ∎

Lemma 3.19 *Suppose that $\psi \in L(M)$ for some model M.*

1. *Suppose that every realization of $\varphi(x)$ is ψ-internal. Then there is a number k so that every realization of $\varphi(x)$ is in the definable closure of k realizations of ψ over M.*

2. *If every realization of $\varphi(x)$ is ψ-internal then c, the canonical parameter of φ is ψ-internal.*

Proof. Note that since we are working over a model, an element a is ψ-internal if and only if a is in the definable closure of $M \cup \psi(\mathcal{C})$. Thus, the proof of the first fact is just compactness. To see the second, note that c is fixed by any automorphism of $\mathcal{C}$ which fixes M and $\psi(\mathcal{C})$.

∎

We now complete the proof of the Proposition 3.13 by showing that each of the sets $D = D(\mathcal{F}, \eta, \chi, \varphi)$ is dense below p^*. The way that we will proceed is to fix $p \leq p^*$ and E as in the definition of D and produce $p_E \leq p$ in $\mathcal{P}$ so that Density Condition 3.15 is satisfied for the given E. By repeating this process now successively for all the finitely many possible E's, we will produce a $q \in D$ such that $q \leq p$. So it suffices to concentrate on one particular E.

By an $\mathcal{F}$-potential extension of p we will mean a $\bar{q} = \langle q_\zeta : \bar{\zeta} \in 2^m \rangle$ such that q_ζ is a d-isolated extension of $p_{\bar{\zeta}}$ together with a realization $\bar{a}$ of $F(\bar{q})$. We concentrate on three cases which correspond to the three conditions in Density Condition 3.15.

Case one: Does there exist an $\mathcal{F}$-potential extension of p, $(\bar{q}, \bar{a})$ such that $\exists y(\chi(y, \bar{a}) \wedge \neg\delta(y))$ holds for some $\delta \in r$? If so then by Lemma 3.18, for every $\zeta \in 2^m$ we can find $\alpha_{\bar{\zeta}} \in q_\zeta$ stronger than $p(\bar{\zeta})$ so that if we form p_E from p by strengthening $p_E(\bar{\zeta})$ to $\alpha_{\bar{\zeta}}$ then $f_E(p_E) \vdash \exists y(\chi(y, \bar{u}_1) \wedge \neg\delta(y))$ for some $\delta \in r$.

Case two: Does there exist an $\mathcal{F}$-potential extension of p, $(\bar{q}, \bar{a})$ such that $\neg\exists x \exists y(\chi(y, \bar{a}) \wedge \varphi(x, y, \bar{a}))$ holds? If so, as in the first case, we can use Lemma 3.18 to define $p_E \leq p$ such that $f_E(p_E) \vdash \neg\exists x \exists y(\chi(y, \bar{u}_1) \wedge \varphi(x, y, \bar{u}))$.

Case three: Does there exist an $\mathcal{F}$-potential extension of p, $(\bar{q}, \bar{a})$ such that $\exists x \exists y(\chi(y, \bar{a}) \wedge \varphi(x, y, \bar{a}) \wedge \neg\delta(x, y))$ holds for some $\delta \in q_\eta$)? If so then using Lemma 3.18, we can define $p_E \leq p$ so that

$$f_E(p_E) \vdash \exists x \exists y(\chi(y, \bar{u}_1) \wedge \varphi(x, y, \bar{u}) \wedge \neg(p_E)_{\bar{\eta}}(x, y))$$

So we are left with the possibility that all three cases fail. That is, for any choice of d-isolated q_ζ and $\bar{q}$ as above and any realization $\bar{a}$ of $F(\bar{q})$, $\chi(y, \bar{a})$ implies $r(y)$, $\chi(y, \bar{a}) \wedge \varphi(x, y, \bar{a})$ is consistent and implies $q_\eta(x, y)$. We will show that this is an impossibility.

Let M' be any extension of M_{n-1} which contains C and is d-isolated over C. Increase B_i for $i \in I^+$ to models and extend the maps f_i so that f_i maps onto M'. Let the prime model over $\bigcup\{B_i : i \in I^- \text{ or } i \notin I_\eta\}$ be called $N_{\text{no } \eta}$. For $i \in I_\eta$, let N_i be prime over $N_{\text{no } \eta}$ together with B_i and finally let N be prime over $N_{\text{no } \eta}$ together with all the N_i's.

Now under the the conditions we are working, we can find $c^* \in N_{\text{no } \eta}$ which realizes r (it is a witness for χ with suitably chosen parameters). Moreover, there are $a_i \in N_i$ for $i \in I_\eta$ such that a_i realizes $\gamma_i = g_i(p_{\bar{\eta}})$ and $b \in N$ such that $\varphi(b, c^*, \bar{a})$ holds where $\bar{a}$ is a sequence which contains all the a_i's and we have suppressed all the parameters from $N_{\text{no } \eta}$. By our assumptions on φ we can find $d_1, \ldots, d_k \in N \setminus N_{\text{no } \eta}$ all of which fork with $N_{\text{no } \eta}$ over B_ϵ and all of which satisfy ψ. Under these circumstances then there is a formula $\beta(x) \in L(N_{\text{no } \eta})$ which is satisfied by all the d_j's and for which $R^\infty(\beta) < R^\infty(\psi)$. We may assume that any realization of $\varphi(x, c^*, \bar{a})$ is β-internal. We now wish to use β to get a contradiction to the minimality of the rank of ψ. In order to get this contradiction we must produce the other three parts of the quadruple mentioned at the beginning of the proof. To produce the n-chain, we do the following: fix any $i \in I_\eta$. In order to form an n-chain $\bar{N}^i$, for $j < n$, choose N^i_j so that

1. N^i_j is prime over $\{B_\mu : \mu \not\geq i \restriction_{j+1}\}$, and
2. $N^i_j \subseteq N^i_{j+1} \subseteq N_{\text{no } \eta}$ for all $j < n$.

It is easy to check that $\bar{N}^i$ is an n-chain and β is defined over N^i_{n-1}; note that $N^i_{n-1} = N_{\text{no } \eta}$. Let $r_i = g_i(r)|N_{\text{no } \eta}$. Of course, r_i is realized in N_i by say c_i and c_i dominates N_i over $N_{\text{no } \eta}$.

We now have three components of a quadruple which will contradict the minimality of $R^\infty(\psi)$. We need to find a formula $\tau_i \in L(C_i)$ where $C_i = acl(N^i_{n-1} c_i)$ such that every realization of τ_i is β-internal and τ_i has no isolated extension over C_i. In fact, τ_i will be realized by a quotient of a_i.

From the finite approximation F, we have been given we have an independent family of models $\bar{N} = \langle N_\mu : \mu \in I\rangle$ together with a family of maps $\{f_\mu : \mu \in I\}$ compatible with $\bar{N}$. If $\zeta \in I^-$ then $f_\mu : N_\mu \to M_{l(\mu)}$ and moreover, there are models M_ν for $\nu \in 2^m$ which are d-atomic over C so that if $i \in I^+$ then $f_i : N_i \to M_{\sigma(i)}$. We let $g_\mu = f_\mu^{-1}$ and let $C_i = g_i(C)$ for $i \in I^+$.

Definition 3.20 We say that $\bar{b} = \langle b_i : i \in I_\eta\rangle$ is an η-sequence satisfying $\tau \in L(C)$ if there is a $p \in S(C)$ which is d-isolated, contains τ and $f_i(b_i/C_i) = p$ for all $i \in I_\eta$. We denote this p as $p_{\bar{b}}$.

At this point in the argument we know that we have a formula $\varphi(x, c^*, \bar{y})$ and a formula $p_\eta \in L(C)$, which implies a fixed formula $\theta(x_0)$ which has no isolated extension over C, so that if we choose an η-sequence $\bar{b}$ which satisfies p_η with $b_i \in N_i$ for all $i \in I_\eta$ then $\varphi(x, c^*, \bar{b})$ is β-internal. Moreover, whenever $\varphi(b, c^*, \bar{b})$ holds then $f_\epsilon(b/C_\epsilon)$ as a type in the variable x_0 is contained in $p_{\bar{b}}$.

Lemma 3.21 *Under the above assumptions, there are formulas τ and E in $L(C)$ together with a formula φ' so that τ implies p_η and for any η-sequence $\bar{b}$, $\varphi(x, c^*, \bar{b})$ and $\varphi(x, c^*, \bar{b}')$ are equivalent where $b'_i = b_i/E$ and b'_i is β-internal for all $i \in I_\eta$.*

Proof. For this proof we define the notion of a quotient of an η-sequence.

Definition 3.22 A formula $\tau \in L(C)$ together with a sequence of equivalence relations $\langle E_i : i \in I_\eta \rangle$ is called a *quotient sequence* if

1. τ implies p_η and
2. there is a formula $\varphi'(x, c^*, \bar{y})$ so that for any η-sequence $\bar{b}$ which satisfies τ, $\varphi(x, c^*, \bar{b})$ is equivalent to $\varphi'(x, c^*, \bar{b}')$ where $b'_i = b_i/E_i$.

The β-number of a quotient sequence τ together with $\langle E_i : i \in I_\eta \rangle$ is the cardinality of the set of $i \in I_\eta$ such that for every η-sequence $\bar{b}$, b_i/E_i is β-internal; such an i is said to be β-internal.

If we succeed in showing that there is a quotient sequence τ and $\langle E_i : i \in I_\eta \rangle$ with a β-number of $|I_\eta|$ then the required formulas in the lemma will be τ and the conjunction of $g_i(E_i)$ over all $i \in I_\eta$.

Suppose $|I_\eta| = k$ and fix a quotient sequence τ and $\langle E_i : i \in I_\eta \rangle$ with maximal β-number. If this number is k we are done so suppose it is less than k. Let φ' be the formula in the definition of quotient sequence. Fix an η-sequence $\bar{b}$ which satisfies τ with $b_i \in N_i$. Let $i_1, \ldots, i_k$ be some enumeration of I_η such that i_k corresponds to an index which is not β-internal. Let $r_j = tp(C_{i_j})$ and $q_j = tp(b'_{i_j}/C_{i_j})$ where $b'_{i_j} = b_{i_j}/E_{i_j}$; note that $b'_{i_j} \in N_{i_j}$ and so q_j is d-isolated.

Define a sequence of formulas ψ_j for $j < k$ by induction: Take ψ_0 to be

$$\forall x(\varphi'(x, v, y_1, w_k, z_{k-1}, w_{k-1}, \ldots, z_1, w_1) \leftrightarrow \varphi'(x, v, y_2, w_k, z_{k-1}, w_{k-1}, \ldots, z_1, w_1))$$

and

$$\psi_{j+1} = d_{r_j} w_j d_{q_j} z_j(\psi_j).$$

Finally, let

$$E(y_1, y_2, w_k) = d_r v \psi_{k-1}.$$

It is straightforward to show that $E(y_1, y_2, w_k)$ defines an equivalence relation for any choice of w_k. In fact, the canonical parameter of $\varphi'(x, c^*, \bar{b}, \bar{c})$ is β-internal and since from it, together with $b_{i_1} C_{i_1} \ldots b_{i_{k-1}} C_{i_{k-1}}$ we can define b_{i_k}/E over C_{i_k}, b_{i_k} is β-internal. It is now clear that we can define the necessary φ'' for this instance of $\bar{b}$ and increase the β-number. We need to uniformize this argument and to do that we now work on strengthening τ.

Let $\chi_{i_j} \in q_j$ be such that any type which contains χ_{i_j} has the same ψ_j-definition as q_j. This exists because q_j is d-isolated. Now let $\tau_i(w)$ for $i \neq i_k$ be the formula expressing the fact that w/E_i satisfies χ_i. We see that $f_i(\tau_i) \in p_{\bar{b}}$ and so we let τ' be the conjunction of τ with all the $f_i(\tau_i)$'s. It is now fairly clear that if we choose any $\bar{d}$, an η-sequence which satisfies τ' then the argument in the paragraph after the definition of E goes through and so we have found a quotient sequence with higher β-number. ∎

We now conclude the proof that D is dense. Consider the formula $\alpha(y) = \exists x \tau(x) \wedge y = x/E$. α has no isolated extension over C for suppose $\bar{\alpha}$ is such an isolated extension of α. Consider the formula $\Gamma(x,y) = \tau(x) \wedge \bar{\alpha}(y) \wedge y = x/E$. This formula has no isolated extension over C and in particular there are certainly two η-sequences $\bar{d}$ and $\bar{e}$ which satisfy $\exists y \Gamma(x,y)$ and disagree in the x_0-variable i.e. there is some $\delta(x_0) \in p_{\bar{d}} \setminus p_{\bar{e}}$. But according to the previous Lemma, $\varphi(x, c^*, \bar{d})$ is equivalent to $\varphi(x, c^*, \bar{d}')$. By independence and the fact that $\bar{\alpha}$ isolates a complete type, $\varphi'(x, c^*, \bar{d}')$ is equivalent to $\varphi'(x, c^*, \bar{e}')$. However the latter is equivalent to $\varphi(x, c^*, \bar{e})$ which contradicts the fact that $\bar{d}$ and $\bar{e}$ disagree in the x_0-variable. From this contradiction, we conclude that α in fact has no isolated extensions. Let $\alpha' = g_i(\alpha)$ for some $i \in I_\eta$. By the Open Mapping Theorem, it is clear that α' has no isolated extension over $acl(N_{\text{no }\eta} C_i)$. Now by the previous Lemma, all realizations of α' are β-internal and so this finally contradicts the minimality of $R^\infty(\psi)$. ∎

We now show that if a special type fails to have a prime model over a chain, then the set of leaves constructed in Proposition 3.13 is diverse. Further, if the type is trivial as well, then the set is diffuse.

Corollary 3.23 *1. If there is a chain $\overline{\mathcal{M}}$ of length n and a special type $r \in R(\overline{\mathcal{M}})$, with no prime model over $M_{n-1}c$ for a realization c of r, then there is a free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ with a diverse family of Leaves$(\overline{\mathcal{M}}')$ of size continuum.*

2. If there is a chain $\overline{\mathcal{M}}$ of length n and a trivial, special type $r \in R(\overline{\mathcal{M}})$, with no prime model over $M_{n-1}c$ for a realization c of r, then there is a free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ with a diffuse family of Leaves$(\overline{\mathcal{M}}')$ of size continuum.

Proof. (1) Fix any special type $r \in R(\overline{\mathcal{M}})$ such that there is no prime model over $M_{n-1}c$ for some realization c of r. Choose the free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$

and the family $Y = \{N_\eta : \eta \in {}^\omega 2\}$ constructed in Proposition 3.13. We claim that Y is diverse. Indeed, let Z_1, Z_2 be distinct subsets of ${}^\omega 2$. Without loss, there is $\eta \in Z_1 \setminus Z_2$. As $s_\eta(x,z)$ is realized in N_η, it is surely realized in $N^*(Z_1) \oplus_{M_{n-2}} V$ for any $(\overline{\mathcal{M}}, Y)$-model V. However, it follows immediately from Proposition 3.13 that s_η is omitted in $N^*(Z_2)$. So, if $n=1$, there is nothing more to prove. On the other hand, if $n > 1$, then since the tree conjugates of r are foreign to some $\varphi \in r$, Lemma 3.10 tells us that any realization c^* of r in $N^*(Z_2) \oplus_{M_{n-2}} V$ is contained in $N^*(Z_2) \oplus_{M_{n-2}} W$ where W is the truncation of V. But Proposition 3.13 tells us that $s_\eta(x, c^*)$ is omitted in $N^*(Z_2) \oplus_{M_{n-2}} V$ for any such c^*. Hence Y is diverse.

(2) Here, fix a trivial, special type $r \in R(\overline{\mathcal{M}})$ such that there is no prime model over $M_{n-1}c$ for some realization c of r. Then as above, the family $Y = \{N_\eta : \eta \in {}^\omega 2\}$ constructed in Proposition 3.13 is diffuse. To see this, choose $\eta \neq \nu$ and some $(\overline{\mathcal{M}}, Y)$-model V. In order to show that $N_\eta \oplus_{M_{n-1}} V \not\cong_V N_\mu \oplus_{M_{n-1}} V$ it suffices to show that $s^*_\eta(x,z)$ is omitted in $N_\mu \oplus_{M_{n-1}} V$ where $s^*_\eta(x,z)$ is the non-forking extension of s_η to V. By way of contradiction, assume that d^*c^* realizes s^*_η in $N_\mu \oplus_{M_{n-1}} V$. Then, by Lemma 3.11 c^* would be in $N_\mu \oplus_{M_{n-1}} W$ where W is the truncation of V, but as before Proposition 3.13 says that $s_\eta(x, c^*)$ is omitted in $N_\mu \oplus_{M_{n-1}} V$, which yields our contradiction. ∎

3.3 The existence of strongly regular types

We begin by describing a general forcing construction that allows us to find a large subset of any perfect set of types over any countable, algebraically closed set in a stable theory in which no type in the set is isolated over any independent set of realizations of the others.

Lemma 3.24 *Assume that $A \subseteq B$, where A is countable and algebraically closed, B realizes countably many complete types over A, and $Q \subseteq S(A)$ is a perfect subset. Then there is a subset $R \subseteq Q$ of size $2^{\aleph_0}$ such that for any distinct $r_1, \ldots, r_n, s \in R$, if $\{c_1, \ldots, c_n\}$ is independent over A, does not fork with B over A, and each c_i realizes r_i, then s is not isolated over $Bc_1 \ldots c_n$.*

Proof. Call a formula $\theta(x)$ *Q-perfect* if $\{q \in Q : \theta \in q\}$ is perfect and let $\mathcal{Q}$ be the set of Q-perfect formulas. Let $\mathcal{P}$ denote the set of all finite functions

$p : {}^{<\omega}2 \to \mathcal{Q}$ such that $p(\mu) \vdash p(\eta)$ for all $\eta \lessdot \mu$ in $\mathrm{dom}(p)$. As before, we write p_η for $p(\eta)$ whenever $\eta \in \mathrm{dom}(p)$. For $p, q \in P$, we say $p \leq q$ if and only if $\mathrm{dom}(p) \supseteq \mathrm{dom}(q)$ and $p_\eta \vdash q_\eta$ for all $\eta \in \mathrm{dom}(q)$.

We describe two countable lists of dense subsets of $\mathcal{P}$ that we wish our filter G to meet. The first list of dense sets will ensure that we build complete types (hence elements of Q) in the limit. For each $\theta \in \mathcal{Q}$ and $\eta \in {}^{<\omega}2$, let

$$D_{\theta,\eta} = \{p \in \mathcal{P} : p_\eta \vdash \theta \text{ or } p_\eta \vdash \neg\theta\}.$$

Since p_η is Q-perfect, each $D_{\theta,\eta}$ is dense in $\mathcal{P}$.

Definition 3.25 Suppose that $p \in \mathcal{P}$ with $\mathrm{dom}(p) \subseteq {}^{\leq m}2$ and $\eta_1, \dots, \eta_k$ are (not necessarily distinct) elements from ${}^m2$. An $L(A)$-formula $\varphi(x_1, \dots, x_k)$ is *decided positively by p at $\eta_1, \dots, \eta_k$* if $\varphi(c_1, \dots, c_k)$ holds for all A-independent sets $\{c_1, \dots, c_k\}$ satisfying $p_{\eta_i} \in \mathrm{tp}(c_i/A) \in Q$ for each i. The formula φ is *decided by p at $\eta_1, \dots, \eta_k$* if either φ or $\neg\varphi$ is decided positively by p at $\eta_1, \dots, \eta_k$.

Fact. For every $L(A)$-formula $\varphi(x_1, \dots, x_k)$, every $p \in \mathcal{P}$ with $\mathrm{dom}(p) \subseteq {}^{\leq m}2$, and every collection $\eta_1, \dots, \eta_k$ from ${}^m2$, there is $q \leq p$ that decides φ at $\eta_1, \dots, \eta_k$ and satisfies $\mathrm{dom}(q) \subseteq {}^{\leq m}2$. In addition, q may be chosen so that $q_\mu = p_\mu$ for all $\mu \in {}^m2$ distinct from $\eta_1, \dots, \eta_k$.

Proof. By induction on k. Assume that the $L(A)$-formula $\varphi(x_0, \dots, x_k)$, a condition p satisfying $\mathrm{dom}(p) \subseteq {}^{\leq m}2$, and $\eta_0, \dots, \eta_k$ from ${}^m2$ are given. Choose a Q-perfect formula θ of least $R(-, \varphi, \omega)$-rank subject to

$$\theta \vdash p_{\eta_0} \qquad \text{and} \qquad Mult_\Delta(\theta) = 1.$$

(Since A is algebraically closed and p_{η_0} is Q-perfect, such a θ exists.) It follows that

$$d_r x_0 \varphi \equiv d_s x_0 \varphi$$

for all $r, s \in Q$ with $\theta \in r \cap s$. Let $\psi(x_1, \dots, x_k)$ be the $L(A)$-formula $d_r x_0 \varphi$ for any such r and let $p' \leq p$ be such that $\mathrm{dom}(p') \subseteq {}^{\leq m}2$ and $p'_{\eta_0} \vdash \theta$. Then, by applying the inductive hypothesis to ψ, we get $q \leq p'$ such that q decides ψ at $\eta_1, \dots, \eta_k$. It follows that q decides φ at $\eta_0, \dots, \eta_k$. ∎

As notation, for every pair of $L(A)$-formulas $\delta(x, y_1, \ldots, y_k, z)$ and $\theta(x)$, and every type $r \in S(A)$ realized in B, let

$$\Gamma_{\delta,r,\theta}(y_1, \ldots, y_k) := d_r z \forall x(\delta(x, y_1, \ldots, y_k, z) \to \theta(x)),$$

For every $L(A)$-formula δ and every type $r \in S(A)$ realized in B, let

> $D_{\delta,r} = \{p \in \mathcal{P}$: there is m such that $\mathrm{dom}(p) \subseteq {}^{\leq m}2$ and for all $\eta_1, \ldots, \eta_k, \mu$ from ${}^m 2$ satisfying $\mu \neq \eta_i$ for each i, there is an $L(A)$-formula θ such that
>
> $$\neg [\Gamma_{\delta,r,\theta}(c_1, \ldots, c_k) \leftrightarrow \theta(d)]$$
>
> for all d realizing p_μ and all A-independent $\{c_1, \ldots, c_k\}$ satisfying $p_{\eta_i} \in \mathrm{tp}(c_1/A) \in Q$ for each $i\}$.

Claim. Each $D_{\delta,r}$ is dense in $\mathcal{P}$.

Proof. Choose $p \in \mathcal{P}$ arbitrarily. Choose m such that $\mathrm{dom}(p) \subseteq {}^{\leq m}2$. It suffices to handle each choice of $\eta_1, \ldots, \eta_k, \mu \in {}^m 2$ separately. So fix some such $\eta_1, \ldots, \eta_k, \mu$. Since Q is perfect, there is a formula θ such that both $p_\mu \wedge \theta$ and $p_\mu \wedge \neg\theta$ are Q-perfect. From the Fact above, there is $p' \leq p$ that decides $\Gamma_{\delta,r,\theta}$ at $\eta_1, \ldots, \eta_k$ and satisfies $\mathrm{dom}(p') \subseteq {}^{\leq m}2$ and $p'_\mu = p_\mu$.

There are two cases. In either case, put $q_{\eta_i} = p'_{\eta_i}$ for each i and put $q_\gamma = p'_\gamma$ for all $\gamma \notin \{\eta_1, \ldots, \eta_k, \mu\}$. Put

$$q_\mu = \begin{cases} p_\mu \wedge \neg\theta & \text{if } \Gamma_{\delta,r,\theta} \text{ is decided positively by } p \text{ at } \eta_1, \ldots, \eta_k \\ p_\mu \wedge \theta & \text{otherwise} \end{cases}$$

Now $q \leq p$ and q meets our requirement for $\eta_1, \ldots, \eta_k, \mu$. After repeating this process for all sequences $\eta_1, \ldots, \eta_k, \mu \in {}^m 2$ with $\mu \neq \eta_i$ for each i we obtain some $q^* \leq p$ with $q^* \in D_{\delta,r}$. ∎

Now fix any filter G such that $G \cap D \neq \emptyset$ for each of the dense sets D mentioned above. For each $\eta \in {}^\omega 2$, let

$$p_\eta(x) = \{\theta : p_{\eta^*} \vdash \theta \text{ for some } p \in G \text{ and some } \eta^* \lessdot \eta\}.$$

It follows from the first collection of dense sets that each p_η is a complete type. Also, as each p_η consists of Q-perfect formulas, each $p_\eta \in Q$. As any type trivially isolates itself, the second collection of dense sets ensures us that $p_\eta \neq p_\mu$ for all $\eta \neq \mu$. Now suppose that $\eta_1, \ldots, \eta_k, \mu \in {}^\omega 2$ are distinct. Pick $\{c_i : 1 \leq i \leq k\}$ to be independent over A, non-forking with B over A, where each c_i realizes p_{η_i}. Suppose that $\delta(x, c_1, \ldots, c_k, b)$ isolates a type in Q. We claim that this type is not p_μ. To see this, let $r = \mathrm{tp}(b/A)$ and choose m such that $\eta_1 \restriction m, \ldots, \eta_k \restriction m, \mu \restriction m$ are distinct. Pick $p \in G \cap D_{\delta,r}$ and let θ be the formula witnessing that $p \in D_{\delta,r}$ for $\eta_1 \restriction m, \ldots, \eta_k \restriction m, \mu \restriction m$. If θ is in the type isolated by $\delta(x, c_1, \ldots, c_k, b)$ then $\Gamma_{\delta,r,\theta}(c_1, \ldots, c_k)$ holds, hence $\Gamma_{\delta,r,\theta}$ is decided positively by p at $\eta_1 \restriction m, \ldots, \eta_k \restriction m$. So $p_{\mu\restriction m} \vdash \neg\theta$. On the other hand, if θ is not in the type isolated by $\delta(x, c_1, \ldots, c_k, b)$ then $\neg\Gamma_{\delta,r,\theta}$ is decided positively by p at $\eta_1 \restriction m, \ldots, \eta_k \restriction m$, so $p_{\mu\restriction m} \vdash \theta$. In either case the type isolated by $\delta(x, c_1, \ldots, c_n, b)$ is not p_μ. ∎

Lemma 3.26 *Suppose that we are given an na-chain $\overline{\mathcal{M}}$ of length n, an $L(M_{n-1})$-formula φ, and a perfect set $Q \subseteq R(\overline{\mathcal{M}})$ such that every $q \in Q$ is special via φ.*

1. *There is a diverse family $Y \subseteq Leaves(\overline{\mathcal{M}})$ of size continuum.*
2. *If, in addition, every $q \in Q$ is trivial, then there is a diffuse family $Y \subseteq Leaves(\overline{\mathcal{M}})$ of size continuum.*

Proof. (1) Suppose that $\overline{\mathcal{M}}$, φ, and Q are given. Choose $R \subseteq Q$ as in Lemma 3.24, taking M_{n-1} for A and W for B. (Note that the isomorphism type of W does not depend on the choice of Y.) We may assume that for any $r \in R$, there is a prime model N_r over $M_{n-1}c$ for any realization c of r, lest there would already be a diverse family of size continuum by Corollary 3.23(1). We claim that $Y = \{N_r : r \in R\}$ is diverse. To see this, suppose $Z_1, Z_2 \subseteq R$ with $s \in Z_1 \setminus Z_2$. It suffices to show that s is omitted in $N^*(Z_2) \oplus_{M_{n-2}} V$ for any $(\overline{\mathcal{M}}, Y)$-model V. When $n = 1$ this follows immediately from Lemma 3.24, while if $n > 1$, then Lemma 3.10 tells us that any potential realization of s must lie in $N^*(Z_2) \oplus_{M_{n-2}} W_V$, which it does not.

(2) Now suppose that the types are trivial as well. Arguing as above, apply Lemma 3.24 to obtain a subset $R \subseteq Q$ for M_{n-1} and W. Further, by Corollary 3.23(2) we may assume that there is a prime model N_r over any realization of r. In this case, the family $Y = \{N_r : r \in R\}$ will be diffuse. To

see this, choose $s \neq r$ and let s^* be the non-forking extension of s to V, some $(\overline{\mathcal{M}}, Y)$-model. We claim that s^* is omitted in $N_r \oplus_{M_{n-1}} V$. If it were realized by an element a, then Lemma 3.11 would imply that $a \in N_r \oplus_{M_{n-1}} W_V$, which would contradict Lemma 3.24. ∎

Proposition 3.27 *Fix a chain $\overline{\mathcal{M}}$ of length n.*

1. *If some free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ has no diverse family of Leaves$(\overline{\mathcal{M}}')$ of size continuum, then for every $p \in R(\overline{\mathcal{M}})$ there is a strongly regular $q \in R(\overline{\mathcal{M}})$ non-orthogonal to p with a prime model over $M_{n-1}c$ for any realization c of q.*

2. *If some free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ has no diffuse family of Leaves$(\overline{\mathcal{M}}')$ of size continuum, then for every trivial $p \in R(\overline{\mathcal{M}})$ there is a strongly regular $q \in R(\overline{\mathcal{M}})$ non-orthogonal to p with a prime model over $M_{n-1}c$ for any realization c of q.*

Proof. We first prove (2). Suppose that the na-chain $\overline{\mathcal{M}}'$ is a free extension of $\overline{\mathcal{M}}$ that has no diffuse family of *Leaves*$(\overline{\mathcal{M}}')$ of size continuum. Fix a trivial type p in $R(\overline{\mathcal{M}})$. Choose a formula $\varphi \in L(M_{n-1})$ of least R^∞-rank among all special formulas non-orthogonal to p and let

$$X = \{q \in S(M_{n-1}) : q \text{ trivial, weight-1, } \varphi \in q, \text{ and } q \perp M_{n-2}\}.$$

If $n = 1$ we delete this last condition. By Lemmas B.2 and B.5, X is a non-empty G_δ subset of $S(M_{n-1})$.

Claim. There are only countably many non-orthogonality classes represented in X.

Proof. If not, then since non-orthogonality is a Borel equivalence relation, Lemma B.1(1) implies that X would contain a pairwise orthogonal family $\{q_i : i \in 2^{\aleph_0}\}$. Let Z be the set of non-forking extensions of each q_i to M'_{n-1}. For each i, choose a regular type r_i domination-equivalent to q_i. Since M'_{n-1} is an na-substructure of the universe, there is a regular type s_i over M'_{n-1} non-orthogonal to q_i. Clearly, each $s_i \in R(\overline{\mathcal{M}}')$, so we have a pairwise orthogonal family of regular types in $R(\overline{\mathcal{M}}')$ of size continuum, hence there is a diffuse family of *Leaves*$(\overline{\mathcal{M}}')$ by Lemma 3.6, which is a contradiction. ∎

Let $\{r_i : i \in j \leq \omega\}$ be a maximal pairwise orthogonal subset of X in which every $r_i \perp p$ and let

$$X_p = \{q \in X : q \not\perp p\}.$$

Since $q \in X_p$ if and only if $q \in X$ and $q \perp r_i$ for all $i \leq j$, it follows from Lemma B.4 that X_p is a G_δ subset of $S(M_{n-1})$ as well. Since φ was chosen to be special of least R^∞-rank, it follows from the proof of Proposition 8.3.5 of [15] that every element of X_p is regular as well. So, if X_p were uncountable, then again by Lemma B.1(1) it would contain a perfect subset. By taking non-forking extension of each of these to M'_{n-1}, we would obtain a perfect subset of $R(\overline{\mathcal{M}}')$, where each element is special via φ. But this contradicts Lemma 3.26(2). Thus, we may assume that X_p is a non-empty, countable G_δ subset of $S(M_{n-1})$. Hence, by Lemma B.1(2) there is an $L(M_{n-1})$-formula $\psi \vdash \varphi$ isolating some $q \in X_p$. By Proposition D.15 of [13], this q is a strongly regular type via ψ. As well, since q is special, if there were no prime model over $M_{n-1}c$ for a realization c of q, then Corollary 3.23(2) would give us a diffuse family of $Leaves(\overline{\mathcal{M}})$ of size continuum.

We now prove (1). Again, assume that $\overline{\mathcal{M}}'$ is a *na*-chain freely extending $\overline{\mathcal{M}}$ with no diverse subset of $Leaves(\overline{\mathcal{M}}')$ of size continuum. Fix a type p in $R(\overline{\mathcal{M}})$. If p is trivial, then as we will see in Lemma 4.2 that every diffuse family is diverse, we are done by (2). Thus, we may assume that p is non-trivial. Hence, by Lemma 8.2.20 of [15], there is a type $p' \in R(\overline{\mathcal{M}})$ non-orthogonal to p and a formula $\theta \in p'$ witnessing that p' is special and, in addition, the p-weight of θ is 1, and p-weight is definable inside θ. Without loss, assume that $p = p'$. Let

$$X = \{q \in S(M_{n-1}) : \theta \in q, w_p(q) = 1\}.$$

Since $w_p(\theta) = 1$ and p-weight is definable inside θ, X is a closed subset of $S(M_{n-1})$. Now fix a $q \in X$. Since $w_p(q) = 1$, $q \not\perp p$. As well, since q is p-simple, q is regular. Hence, if $n > 1$ then $q \perp M_{n-2}$. That is,

$$X = \{q \in R(\overline{\mathcal{M}}) : q \text{ regular}, \theta \in q, \text{ and } q \not\perp p\}$$

is closed, hence a G_δ. The proof is now analogous to (1). If X were uncountable, then by Lemma B.1(1) it would contain a perfect subset, so by taking non-forking extensions to M'_{n-1} there would be a diverse family of

$Leaves(\overline{\mathcal{M}}')$ of size continuum by Lemma 3.26. So, we may assume that X is a non-empty, countable, G_δ subset of $S(M_{n-1})$ and hence there is an $L(M_{n-1})$-formula $\psi \vdash \theta$ isolating some $q \in X$. Again by Proposition D.15 of [13], this q is a strongly regular type via ψ. As well, since q is special, if there were no prime model over $M_{n-1}c$ for a realization c of q, then Corollary 3.23(1) would yield a diverse family of $Leaves(\overline{\mathcal{M}}')$ of size continuum. ∎

3.4 Structure theorems for locally t.t. theories

In this subsection we analyze some of the positive consequences of a theory being locally t.t. over an n-chain $\overline{\mathcal{M}}$.

The following three lemmas are true for any superstable theory. There is nothing novel about their statements or proofs as they are central in the analysis of models of a totally transcendental theory. They are given here simply to indicate that the assumption of being totally transcendental can be weakened to include our context.

Lemma 3.28 *Suppose that $p, q \in S(M)$ where p is regular and q is strongly regular. Then p and q are orthogonal if and only if they are almost orthogonal over M.*

Proof. Suppose that q is strongly regular via ψ and that $p \not\perp q$. Choose a finite e and realizations a of $p|Me$ and b of $q|Me$ respectively such that $a \underset{M}{\not\perp} eb$. Choose a formula $\varphi(x; y, z)$ over M such that $\varphi(a; e, b)$ holds and $\varphi(a; y, z)$ forks over M. Thus

$$\exists z[\varphi(a; e, z) \wedge \psi(z)].$$

Since M is a model and $a \underset{M}{\perp} e$, there is $e' \in M$ such that $\exists z[\varphi(a; e, z) \wedge \psi(z)]$. Let b' be any witness to this formula. As b' forks with a over M, $\mathrm{tp}(b'/M) \not\perp p$, hence $\mathrm{tp}(b'/M) = q$ since q is strongly regular via ψ. Thus p and q are not almost orthogonal over M. ∎

Lemma 3.29 *Suppose that $p, q \in S(M)$ are non-orthogonal, where p is regular and q is strongly regular. Then any model containing a realization of p contains a realization of q.*

Proof. Let a realize p and let $N \supseteq Ma$. Again suppose that q is strongly regular via ψ. Since p and q are not almost orthogonal over M, there is a formula $\varphi(x, y)$ over M such that $\varphi(a, y)$ forks over M and $\varphi(a, b)$ holds for some b realizing q. So $\exists y[\varphi(a, y) \wedge \psi(y)]$ holds in the monster model, hence by elementarity, N contains a witness b' to $\varphi(a, y) \wedge \psi(y)$. As before, it follows from the strong regularity of q and the fact that $\varphi(a, y)$ forks over M that $\mathrm{tp}(b'/M) = q$. ∎

Notation: Suppose that $p \in S(M)$. We adopt the notation $M(p)$ for the prime model over M and any realization of p, if this prime model exists.

Lemma 3.30 *Suppose $q \in S(M)$ is strongly regular and $M(q)$ exists. Then $M(q)$ is atomic over M and any $a \in M(q)$ such that* $\mathrm{tp}(a/M)$ *is regular.*

Proof. Suppose that $M(q)$ is atomic over Mb, where b realizes q and that q is strongly regular via $\psi(y)$. Choose $a \in M(q) \setminus M$ arbitrarily such that $tp(a/M)$ is regular. It suffices to show that $\mathrm{tp}(b/Ma)$ is isolated. Let $\theta(x, b)$ isolate $\mathrm{tp}(a/Mb)$. Since $a \notin M$, it follows from the Open Mapping Theorem that a forks with b over M, so let $\varphi(a, y) \in \mathrm{tp}(b/Ma)$ fork over M. Then the formula

$$\alpha(a, y) := \theta(a, y) \wedge \varphi(a, y) \wedge \psi(y)$$

isolates $\mathrm{tp}(b/Ma)$. To see this, it is evident that $\alpha(a, b)$ holds. So suppose that $\alpha(a, c)$ holds for some element c. We will show that $\mathrm{tp}(b/Ma) = \mathrm{tp}(c/Ma)$. Since $\varphi(a, y)$ forks over M, $\mathrm{tp}(a/M)$ is regular and q is strongly regular via ψ, it must be that $\mathrm{tp}(c/M) = q$. Hence, $\theta(x, c)$ isolates a complete type over Mc. That is, $\mathrm{tp}(ab/M) = \mathrm{tp}(ac/M)$, so $\mathrm{tp}(b/Ma) = \mathrm{tp}(c/Ma)$ as desired. ∎

When we combine the three lemmas above with the existence of prime models over independent trees of models, we obtain a reasonable structure theory for the class of models of models extending a fixed chain. We illustrate this by continuing our analogy with the analysis of totally transcendental theories. The proofs of the following two corollaries follow immediately from the lemmas above.

Corollary 3.31 *If M is countable and $p, q \in S(M)$ are both strongly regular and are non-orthogonal, then if $M(q)$ exists then $M(p)$ exists as well. Further, $M(p)$ and $M(q)$ are isomorphic over M.*

Definition 3.32 I is a strongly regular sequence over $\overline{\mathcal{M}}$ if I is independent over M and every $a \in I$ realizes a strongly regular type in $R(\overline{\mathcal{M}})$.

Corollary 3.33 *If T is locally t.t. over a chain $\overline{\mathcal{M}}$ of length n and I is a strongly regular sequence over $\overline{\mathcal{M}}$, then there is a prime model N over $M_{n-1} \cup I$. Further, if p and q are non-orthogonal, strongly regular types from $R(\overline{\mathcal{M}})$, then* $\dim(p, N) = \dim(q, N)$.

The following lemma will be used in conjunction with the previous corollary to obtain good upper bounds.

Lemma 3.34 *Suppose that $\overline{\mathcal{M}}$ is a d-chain, where d is the depth of T. If T is locally t.t. over $\overline{\mathcal{M}}$, then any model $N \supseteq M_{d-1}$ with $N/M_{d-1} \perp M_{d-2}$ (when $d > 1$) is prime over M_{d-1} and any maximal strongly regular sequence over M_{d-1}.*

Proof. Let I be any maximal strongly regular sequence over $\overline{\mathcal{M}}$ inside N and, from the previous corollary, let $N' \subseteq N$ be prime over $M_{d-1} \cup I$. Suppose, by way of contradiction, that $N' \neq N$. Choose $a \in N \setminus N'$ such that $p = \text{tp}(a/N')$ is regular. Since T has depth d, p is not orthogonal to M_{d-1}, so choose a regular type q over M_{d-1} non-orthogonal to p. Since T is locally t.t. over $\overline{\mathcal{M}}$, we may assume that q is strongly regular. Let q' be the non-forking extension of q to N'. It follows from Lemma 3.29 that q' is realized in N, but this contradicts the maximality of I. ∎

We conclude this subsection by giving a general lemma that will be used to aid our counting in some cases when the depth of the theory is small (typically $d = 1$ or 2). As notation, note that any countable model M can be viewed as a chain of length 1. In this case, we let $R(M)$ be the set of regular types over M and we say that T is locally t.t. over M if it is locally t.t. over the associated chain of length 1.

Lemma 3.35 *Suppose that T is countable, superstable, and over any countable model M, and T is locally t.t. over M. If in addition, $R(M)$ contains only countably many non-orthogonality classes, then T is ω-stable.*

Proof. To show that T is ω-stable, for every countable model M_0, we will find a countable model M_ω realizing every type over M_0. So fix a countable model M_0. We construct M_ω to be a union of a chain of models M_n, where M_0 is given and for each n, M_{n+1} is chosen to be prime over $M_n \cup I_n$, where I_n is a strongly regular sequence over M_n where each non-orthogonality class of $R(M_n)$ has dimension $\aleph_0$. (Such a sequence I_n exists since T is locally t.t. over M_n and is countable since $R(M_n)$ is.)

We claim that M_ω realizes every type over M_0; in fact, it realizes every type over M_n for all n. If not, then choose $p \in S(M_n)$ of least rank such that p is omitted in M_ω. Let a be a realization of p and let $N \supseteq M_n a$ be any countable model that is dominated by a over M_n. (For instance, N could be taken to be l-isolated over $M_n a$.) Let $q \in S(M_n)$ be any regular type realized in $N \setminus M_n$. Since T is locally t.t. over M_n there is a strongly regular type $r \in R(M_n)$ non-orthogonal to q, which by Lemma 3.29 is also realized in N, say by b. Since N is dominated by a over M_n, a forks with b over M_n, so $R^\infty(a/M_n b) < R^\infty(p)$. However, we can easily assume that $b \in M_{n+1}$ and so $R^\infty(a/M_{n+1}) < R^\infty(p)$ and certainly $tp(a/M_{n+1})$ is not realized in M_ω which is a contradiction. ∎

4 Unique Decompositions and Iteration

Before we can state one of the key theorems of this section, we introduce a very useful model U. Fix an na-chain $\overline{\mathcal{M}}$ of length n and a set Y of leaves of $\overline{\mathcal{M}}$. For every uncountable cardinal λ, we define an $(\overline{\mathcal{M}}, Y)$-model U_λ which can be thought of as a 'λ-saturated $(\overline{\mathcal{M}}, Y)$-model' of size $\lambda + 2^{\aleph_0}$. That is, we take U_λ to be prime over an $(\overline{\mathcal{M}}, Y)$-tree $\mathcal{M} = \langle M_\zeta : \zeta \in I \rangle$ with ordering $\lessdot$ on I. We demand that I be λ-branching above every $\eta \in I$ with $lg(\eta) < n$ i.e. $\{\nu : \nu^- = \eta\}$ is of size λ. Moreover, if $\langle f_\zeta : \zeta \in I \rangle$ is the family of elementary maps compatible with $\mathcal{M}$ demonstrating that it is an $(\overline{\mathcal{M}}, Y)$-tree then for every $\zeta \in I$ with $lg(\zeta) = n$ and every $\bar{N} \in Y$, there are λ-many $\nu \in I$ with $\nu^- = \zeta$ such that f_ν maps onto N_n. For brevity, we denote $U_{\aleph_1}$ by U. The importance of this $(\overline{\mathcal{M}}, Y)$-model is given by the following theorem, whose proof follows from the main theorem of [8].

Theorem 4.1 [Unique Decomposition Theorem] *If $Q = \bigoplus_{M_k}\{P_i : i \in I\}$ and $Q' = \bigoplus_{M_k}\{P'_i : i \in I'\}$ where each P_i, P'_i is an $(\overline{\mathcal{M}}, Y)$-model with weight one over M_k and $Q \oplus_{M_k} V \cong_V Q' \oplus_{M_k} V$ for some $(\overline{\mathcal{M}}, Y)$-model V then there is a bijection $f : I \to I'$ such that $P_i \oplus_{M_k} U \cong_U P'_{f(i)} \oplus_{M_k} U$ for all $i \in I$.*

Proof. There is no loss in assuming that in fact $N = Q \oplus_{M_k} V = Q' \oplus_{M_k} V$ so in the terminology of [8] $\{P_i : i \in I\}$ and $\{P'_i : i \in I'\}$ are sets of independent M_k-components of N. There is a bijection $f : I \to I'$ so that for any $i \in I$, $\{P_i\} \cup \{P'_i : i \in I'\} \setminus \{P'_{f(i)}\}$ is a set of independent M_k-components of N. So if we fix $i \in I$ and let $\hat{V} = (\bigoplus_{M_k} \{P'_i : i \in I'\} \setminus \{P'_{f(i)}\}) \oplus_{M_k} V$ then $N = P_i \oplus_{M_k} \hat{V} = P'_{f(i)} \oplus_{M_k} \hat{V}$. Now choose λ large enough so that $\hat{V}$ embeds into U_λ. By freely joining U_λ with N over $\hat{V}$ we obtain $P_i \oplus_{M_k} U_\lambda = P'_{f(i)} \oplus_{M_k} U_\lambda$. As well, note that $U \subseteq^{\mathcal{P}}_{\aleph_1} U_\lambda$ in the notation of [8]. Thus, by Theorem 2.9 of [8], $P_i \oplus_{M_k} U \cong_U P'_{f(i)} \oplus_{M_k} U$ which is the conclusion of the Theorem. ∎

Our first application of this theorem gives the promised implication between the two central notions of this section.

Lemma 4.2 *If Y is diffuse then Y is diverse.*

Proof. Assume that Y is not diverse, i.e., there are distinct $Z_1, Z_2 \subseteq Y$ with $N^*(Z_1) \oplus_{M_{n-2}} V \cong N^*(Z_2) \oplus_{M_{n-2}} V$ over M_{n-1} and V for some $(\overline{\mathcal{M}}, Y)$-model V. Let V' be the $(\overline{\mathcal{M}}, Y)$-model formed by taking the prime model over M_{n-1} and V. We can assume that $N^*(Z_1) \oplus_{M_{n-1}} V' \cong_{V'} N^*(Z_2) \oplus_{M_{n-1}} V'$. Since $Z_1 \neq Z_2$ it follows from the Unique Decomposition Theorem that there are distinct $N, N' \in Y$ such that $N \oplus_{M_{n-1}} U \cong_U N' \oplus_{M_{n-1}} U$, which implies that Y is not diffuse. ∎

The following notion captures a notion of homogeneity that our models will possess.

Definition 4.3 If $W \subseteq Q$, a set H *reflects* Q *over* W if $W \subseteq H \subseteq Q$ and for every countable $A \subseteq Q$ there is an automorphism σ of Q over W with $\sigma(A) \subseteq H$.

The utility of this notion is given by the lemma below, which will be used in many different contexts in what follows.

Lemma 4.4 *Suppose that $\mathcal{F} = \{P_i : i \in \theta\}$ is a family of models, pairwise non-isomorphic over $V \cup W$, where V is countable, and there is a cardinal $\lambda \geq |W|$ such that $\lambda^{\aleph_0} < \theta$ and every P_i has a reflecting subset over W of size at most λ. Then there is a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ of size θ such that $P \not\cong_W P'$ for distinct $P, P' \in \mathcal{F}'$.*

Proof. Say $i \sim j$ if and only if $P_i \cong_W P_j$. It suffices to show that every $\sim$-class has size at most $\lambda^{\aleph_0}$. So, by way of contradiction, assume that some $\sim$-class $C \subseteq \theta$ has size greater than $\lambda^{\aleph_0}$. Without loss, assume $0 \in C$. Let $H \subseteq P_0$ be reflecting over W. For each $i \in C$, let $f_i : P_i \to P_j$ be an isomorphism over W. By composing each f_i by an automorphism over W if needed, the fact that H is reflecting allows us to assume that $f_i(V) \subseteq H$ for all $i \in C$. Since $|C| > \lambda^{\aleph_0}$, there are distinct $i, j \in C$ such that $f_i(v) = f_j(v)$ for each $v \in V$. Thus, $f_j^{-1} \circ f_i : P_i \to P_j$ is an isomorphism over $V \cup W$, which is a contradiction. ∎

The following definition is arranged to allow us us to 'step down' trees and to achieve a lower bound on $I(T, \aleph_\alpha)$. The reader should observe that the hypotheses on κ, λ imply that the existence of an iterable family of models implies that each model has size greater than continuum.

Definition 4.5 A family $\mathcal{F} = \{P_i : i \in \theta\}$ of $(\overline{\mathcal{M}}, Y)$-models is *$k$-iterable* if there are cardinals $\kappa, \lambda \geq \aleph_0$ such that:

1. $\lambda^{\aleph_0} < \kappa$ and $\lambda^{\aleph_0} < \theta$;
2. each P_i has size κ and has a reflecting subset H_i over M_k of size λ; and
3. $P_i \oplus_{M_k} U \not\cong_U P_j \oplus_{M_k} U$ for all distinct $i, j \in \theta$.

Lemma 4.6 *If there is a k-iterable family of θ models, each of size $\aleph_\alpha$, then $I(T, \aleph_\alpha) \geq \theta$.*

Proof. Let $\mathcal{F}$ be such a family. It follows immediately from the definition that the models are pairwise non-isomorphic over M_k and there is a cardinal λ satisfying $\lambda^{\aleph_0} < \aleph_\alpha$ such that every $P \in \mathcal{F}$ has a reflecting subset of size λ over M_k. Thus, $I(T, \aleph_\alpha) \geq \theta$ follows immediately from Lemma 4.4 (taking $W = \emptyset$). ∎

The intuition is that if a family is k-iterable, then we can use the lemma below to 'step down' the tree k times, roughly exponentiating the number of models at each step.

Lemma 4.7 *If $\mathcal{F} = \{P_i : i \in \theta\}$ is a k-iterable family of models each of size $\aleph_\alpha$ such that $k > 0$, $\theta \leq \aleph_\alpha$, and $\theta^{\aleph_0} < |\alpha + \omega|^\theta$, then there is a $(k-1)$-iterable family of $|\alpha + \omega|^\theta$ models, each of size $\aleph_\alpha$.*

Proof. Fix a cardinal $\lambda \geq \aleph_0$ such that $\lambda^{\aleph_0} < \aleph_\alpha$ and, for each $i \in \theta$ choose a subset H_i reflecting P_i over M_k. By enlarging each H_i slightly we may take it to be the prime model over a subtree of a decomposition tree of P_i.

We first claim that there is a subfamily $\mathcal{F}'$ of $\mathcal{F}$ of size θ such that $P \oplus_{M_{k-1}} U \not\cong_U P' \oplus_{M_{k-1}} U$ for all distinct $P, P' \in \mathcal{F}'$. For, if this were not the case, there would be a set $X \subseteq \theta$ of size θ such that $P_i \oplus_{M_{k-1}} U \cong_U P_j \oplus_{M_{k-1}} U$ for all $i, j \in X$. It is easily checked that $H_i' = H_i \oplus_{M_{k-1}} U$ is of size at most $\lambda \cdot 2^{\aleph_0}$ and reflects $P_i \oplus_{M_{k-1}} U$ over U. So, by applying Lemma 4.4 to the subfamily there would be $i \neq j$ such that $P_i \oplus_{M_{k-1}} U \cong_{UM_k} P_j \oplus_{M_{k-1}} U$. If we let V be the prime model over M_k and U, we can arrange that

$$P_i \oplus_{M_k} V \cong_V P_j \oplus_{M_k} V$$

But this would imply $P_i \oplus_{M_k} U \cong_U P_j \oplus_{M_k} U$ by the Unique Decomposition Theorem which would contradict the k-iterability of $\mathcal{F}$. Thus, by reindexing we may assume that our original family $\mathcal{F}$ satisfies $P_i \oplus_{M_{k-1}} U \not\cong_U P_j \oplus_{M_{k-1}} U$ for all distinct $i, j \in \theta$.

Let S be an independent (over M_{k-1}) family of $\aleph_\alpha$ copies of every P_i for every $i \in \theta$ and for each $P \in S$ choose a submodel H_P prime over a subtree of a decomposition tree of P of size at most λ. Let J denote the set of all cardinal-valued functions $f : \theta \to \aleph_\alpha^+$ such that $f(0) = \aleph_\alpha$. For each $f \in J$, let $S_f \subseteq S$ consist of $f(i)$ copies of each P_i, and let $S_f^* \subseteq S_f$ consist of $\min\{f(i), \aleph_0\}$ copies of each P_i. Let $Q_f = \bigoplus_{M_{k-1}} S_f$. Each Q_f has size $\aleph_\alpha$. Let $K_f = \bigoplus_{M_{k-1}} \{H_P : P \in S_f^*\}$. As any collection of automorphisms of distinct components of Q_f extends to an automorphism of Q_f over M_{k-1}, it follows that each K_f reflects Q_f over M_{k-1}. As well, $|K_f| \leq \theta \cdot \lambda$ and it follows from our cardinality assumptions that both $\lambda^{\aleph_0}$ and $\theta^{\aleph_0}$ are strictly less than $|\alpha + \omega|^\theta$.

Finally, suppose that $Q_f \oplus_{M_{k-1}} U \cong_U Q_g \oplus_{M_{k-1}} U$ for some distinct $f, g \in J$. Then, by the Unique Decomposition Theorem there are distinct $i, j \in \theta$ such that $P_i \oplus_{M_{k-1}} U \cong_U P_j \oplus_{M_{k-1}} U$, which is contrary to our additional hypothesis on $\mathcal{F}$ mentioned above. ∎

By iterating Lemma 4.7 we obtain the following:

Lemma 4.8 *If for some $m \geq 1$ there is an m-iterable family of θ models, each of size $\aleph_\alpha$, where $\theta^{\aleph_0} = \theta$, then $I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{m-1}(|\alpha + \omega|^\theta)\}$.*

Proof. Define a sequence of cardinals $\theta_0 < \theta_1 < \ldots < \theta_m$ by letting $\theta_0 = \theta$ and $\theta_{k+1} = |\alpha + \omega|^{\theta_k}$. Note that $\theta_k^{\aleph_0} = \theta_k$ for each k. There are now two cases. First, if $\theta_{m-1} \leq \aleph_\alpha$, then by applying Lemma 4.7 m times, one gets a family of $\theta_m = \beth_{m-1}(|\alpha + \omega|^\theta)$ 0-iterable models over M_0. Hence, $I(T, \aleph_\alpha) \geq \beth_{m-1}(|\alpha + \omega|^\theta)$.

For the second case, assume that $\aleph_\alpha < \theta_{m-1}$. Choose k least such that $\aleph_\alpha < \theta_k$ (k may be zero). By applying Lemma 4.7 k times, we obtain a family $\mathcal{G}$ of θ_k $(m-k)$-iterable models, each of size $\aleph_\alpha$. Now if $2^{\aleph_\alpha} \leq \theta_k$, then $I(T, \aleph_\alpha) = 2^{\aleph_\alpha}$ by applying Lemma 4.6 to this family. However, if $\theta_k < 2^{\aleph_\alpha}$, then as $\aleph_\alpha^{\aleph_0} \leq \theta_k^{\aleph_0} = \theta_k < 2^{\aleph_\alpha}$, we can apply Lemma 4.7 one more time to a subfamily of $\mathcal{G}$ of size $\aleph_\alpha$. This produces an $(m-k-1)$-iterable family of size $2^{\aleph_\alpha}$, hence $I(T, \aleph_\alpha) = 2^{\aleph_\alpha}$ by Lemma 4.6. ∎

5 The counting

In this section, we combine the dichotomies from Section 2 with the machinery in Section 3 to obtain lower and upper bounds in many situations. Our results will be strong enough to compute the uncountable spectra in all cases. As noted in the introduction, we need only concern ourselves with (countable) classifiable theories of finite depth. For notation, *assume that such a theory T has depth d.* We recall the nomenclature of [7].

Definition 5.1 Suppose $1 \leq n \leq d$.

- We say that $TT(n)$ *holds* if, T is locally t.t. over every chain $\overline{\mathcal{M}}$ of length n.
- $TF(n)$ holds if T admits a trivial failure over some chain of length n (i.e., some trivial $p \in R(\overline{\mathcal{M}})$ is not t.t. over $\overline{\mathcal{M}}$ for some chain $\overline{\mathcal{M}}$ of length n.)
- $NTF(n)$ holds (read: non-trivial failure) if there is a chain $\overline{\mathcal{M}}$ of length n and a non-trivial type $p \in R(\overline{\mathcal{M}})$ that is not totally transcendental over $\overline{\mathcal{M}}$.
- We write $\#RD(n) = 2^{\aleph_0}$ if there is some chain $\overline{\mathcal{M}}$ of length n where $R(\overline{\mathcal{M}})$ contains a family of continuum pairwise non-orthogonal types.

- We write $\#RD(n) = \aleph_0$ if there is some chain $\overline{\mathcal{M}}$ of length n for which $R(\overline{\mathcal{M}})$ contains infinitely many non-orthogonal types, yet there is no chain $\overline{\mathcal{N}}$ of length n with continuum many non-orthogonality classes represented in $R(\overline{\mathcal{N}})$.

- We say $\#RD(n)$ is finite if only finitely many non-orthogonality classes are represented in $R(\overline{\mathcal{M}})$ for every chain $\overline{\mathcal{M}}$ of length n.

- We write $\#RD(n) = 1$ if all types in $R(\overline{\mathcal{M}})$ are non-orthogonal for all chains $\overline{\mathcal{M}}$ of length n.

The notation $\#RD$ stands for the number of relevant dimensions. Since for any n-chain $\overline{\mathcal{M}}$, $S(M_{n-1})$ is a Polish space and the relation of non-orthogonality is a Borel equivalence relation, the number of non-orthogonality classes represented in $R(\overline{\mathcal{M}})$ is either countable or of size $2^{\aleph_0}$; hence at least one of the last four conditions hold for each n.

In the first subsection we use the machinery established earlier to obtain a number of lower bounds. Then, in the subsections that follow, we put on more and more conditions on our theory to obtain better and better upper bounds.

5.1 Lower Bounds

The goal of this subsection is to obtain good lower bounds from instances of 'non-structure' of the theory. The first of these is due to Shelah and has been known for some time. It will be used in two places, both where the depth of the theory is quite low. For a proof, see either Theorem 1.20 of Chapter IX of [18] or Theorem C of [2].

Lemma 5.2 *If T is superstable but not ω-stable then $I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_2\}$ for all $\alpha > 0$.*

The next result is the 'General lower bound' mentioned in the Introduction. It is proved by the method of quasi-isomorphisms. (See either Theorem 5.10(a) of [16] or the discussion on page 396 of [1] for a proof.)

Lemma 5.3 *If T is classifiable of depth $d > 1$ then*

$$I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{d-2}(|\alpha + \omega|^{|\alpha+1|})\}$$

for all $\alpha > 0$.

For each of the next five lemmas and propositions, *assume that $\overline{\mathcal{M}}$ is an na-chain of length n.* We begin by considering the effect of a diffuse family of leaves.

Lemma 5.4 *If there is a diffuse family of Leaves($\overline{\mathcal{M}}$) of size $\mu \geq \aleph_0$, then for any ordinal α satisfying $\aleph_\alpha > 2^{\aleph_0}$ and $\mu^{\aleph_0} < |\alpha+\omega|^\mu$, there is an $(n-1)$-iterable family of $|\alpha+\omega|^\mu$ models, each of size $\aleph_\alpha$.*

Proof. Let $\mathcal{F} = \{P_i : i \in 2^{\aleph_0}\}$ be diffuse. Let J be the set of all cardinal-valued functions $f : 2^{\aleph_0} \to \aleph_\alpha^+$ such that $f(0) = \aleph_\alpha$. It follows immediately from the Unique Decomposition Theorem that the family $\mathcal{G} = \{Q_f : f \in J\}$, where

$$Q_f = \bigoplus_{M_{n-1}} \{P_i^{(f(i))} : i \in 2^{\aleph_0}\}$$

satisfies $Q_f \oplus_{M_{n-1}} U \not\cong_U Q_g \oplus_{M_{n-1}} U$ for distinct $f, g \in J$. As well, the substructure $H_f = \bigoplus_{M_{n-1}} \{P_i^{(\hat{f}(i))} : i \in 2^{\aleph_0}\}$, where $\hat{f}(i) = \min\{f(i), \aleph_0\}$, reflects Q_f over M_{n-1} and has size continuum. ∎

The following lemma is routine.

Lemma 5.5 *If there is a diverse family of Leaves($\overline{\mathcal{M}}$) of size $\mu \geq \aleph_0$, then for any cardinal $\kappa \geq \mu$ there is a family of 2^μ pairwise non-isomorphic models over M_{n-1}, each of size κ.*

Proposition 5.6 *If there is a diffuse family of Leaves($\overline{\mathcal{M}}$) of size continuum then $I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{n-1}(|\alpha+\omega|^{2^{\aleph_0}})\}$ for all ordinals $\alpha > 0$.*

Proof. Fix $\alpha > 0$. There are three cases. First, assume $2^{\aleph_\alpha} = 2^{\aleph_0}$. The existence of a diffuse family of size continuum clearly implies that the theory T is not $\aleph_0$-stable, hence $I(T, \aleph_\alpha) \geq 2^{\aleph_0} = 2^{\aleph_\alpha}$ by Lemma 5.2. Second, assume $\aleph_\alpha \leq 2^{\aleph_0}$ and $2^{\aleph_\alpha} > 2^{\aleph_0}$. Then by Lemma 5.5 there is a family of $2^{\aleph_\alpha}$ models over M_{n-1}, each of size $\aleph_\alpha$, that are pairwise non-isomorphic over M_{n-1}. Our cardinal assumptions imply that $\aleph_\alpha^{\aleph_0} < 2^{\aleph_\alpha}$, so $I(T, \aleph_\alpha) = 2^{\aleph_\alpha}$.

Finally, if $\aleph_\alpha > 2^{\aleph_0}$, then by Lemma 5.4 there is an $(n-1)$-iterable family of $|\alpha+\omega|^{2^{\aleph_0}}$ models, each of size $\aleph_\alpha$. If $n = 1$, then $I(T, \aleph_\alpha) \geq |\alpha+\omega|^{2^{\aleph_0}}$ by Lemma 4.6. However, if $n > 1$, then by Lemma 4.8

$$I(T, \aleph_\alpha) \geq \beth_{n-2}\left(|\alpha+\omega|^{|\alpha+\omega|^{2^{\aleph_0}}}\right) = \beth_{n-1}(|\alpha+\omega|^{2^{\aleph_0}}).$$

We now turn our attention to the existence of a diverse family of leaves.

Lemma 5.7 *If there is a diverse family of size* $\mu \geq \aleph_0$ *and* $n > 1$*, then for any* $\aleph_\alpha > 2^{\aleph_0}$ *there is an* $(n-2)$*-iterable family of size* $\min\{2^{\aleph_\alpha}, |\alpha+\omega|^{2^\mu}\}$*, where each model has size* $\aleph_\alpha$*.*

Proof. Choose a cardinal $\theta \leq \min\{\aleph_\alpha, 2^\mu\}$ such that

$$|\alpha+\omega|^\theta = \min\{2^{\aleph_\alpha}, |\alpha+\omega|^{2^\mu}\}, \quad \theta^{\aleph_0} < \aleph_\alpha, \quad \text{and} \quad \theta^{\aleph_0} < |\alpha+\omega|^\theta.$$

(Such a θ can be chosen from $\{2^{\aleph_0}, \beth_2, \aleph_\alpha\}$.) It follows directly from the diversity of the family of leaves that there is a collection $\mathcal{F} = \{P_i : i \in \theta\}$ of models over M_{n-1}, each of size at most $\aleph_\alpha$, that satisfies $P_i \oplus_{M_{n-2}} U \not\cong_U P_j \oplus_{M_{n-2}} U$ for all distinct $i, j \in \theta$.

Let S be an independent (over M_{n-2}) family of $\aleph_\alpha$ copies of every P_i for every $i \in \theta$. Let J denote the set of all cardinal-valued functions $f : \theta \to \aleph_\alpha^+$ such that $f(0) = \aleph_\alpha$. For each $f \in J$, let $S_f \subseteq S$ consist of $f(i)$ copies of each P_i, and let $S_f^* \subseteq S_f$ consist of $\min\{f(i), \aleph_0\}$ copies of each P_i. Let $Q_f = \bigoplus_{M_{n-2}} S_f$. Each Q_f has size $\aleph_\alpha$. Let $K_f = \bigoplus_{M_{n-2}}\{P : P \in S_f^*\}$. Again, as any collection of automorphisms of distinct components of Q_f extends to an automorphism of Q_f over M_{n-2}, it follows that each K_f reflects Q_f over M_{n-2}. Since $|K_f| = \theta$ our hypotheses on θ insure that this family satisfies the cardinal hypotheses of iterability.

So, suppose that $Q_f \oplus_{M_{n-2}} U \cong_U Q_g \oplus_{M_{n-2}} U$ for some distinct $f, g \in J$. Then, by the Unique Decomposition Theorem there are distinct $i, j \in \theta$ such that $P_i \oplus_{M_{n-2}} U \cong_U P_j \oplus_{M_{n-2}} U$, which is impossible. ∎

Proposition 5.8 *If there is a diverse family of Leaves*$(\overline{\mathcal{M}})$ *of size continuum then*

$$I(T, \aleph_\alpha) \geq \begin{cases} \min\{2^{\aleph_\alpha}, \beth_2\} & \text{if } n = 1 \\ \min\{2^{\aleph_\alpha}, \beth_{n-2}(|\alpha+\omega|^{\beth_2})\} & \text{if } n > 1 \end{cases}$$

for all ordinals $\alpha > 0$*.*

Proof. First, if $\aleph_\alpha \leq 2^{\aleph_0}$ then $I(T, \aleph_\alpha) = 2^{\aleph_\alpha}$ by splitting into the same two cases as in the proof of Proposition 5.6. As well, if $\aleph_\alpha \geq 2^{\aleph_0}$ and $n = 1$, then Lemma 5.5 implies $I(T, \aleph_\alpha) \geq \beth_2$.

So assume $\aleph_\alpha > 2^{\aleph_0}$ and $n > 1$. Then by Lemma 5.7 there is an $(n-2)$-iterable family of $\min\{2^{\aleph_\alpha}, |\alpha+\omega|^{\beth_2}\}$ models, each of size $\aleph_\alpha$. Thus, the proposition follows from Lemma 4.6 if $n = 2$, and from Lemma 4.8 if $n > 2$. ∎

We finish this section by obtaining lower bounds that arise from having suitably large collections of pairwise orthogonal types.

Proposition 5.9 *If $R(\overline{\mathcal{M}})$ contains an infinite family of pairwise orthogonal types for some na-chain $\overline{\mathcal{M}}$ of length n, then*

$$I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{n-1}(|\alpha + \omega|^{\aleph_0})\}$$

for all ordinals $\alpha > 0$.

Proof. First, we claim that $I(T, \aleph_\alpha) \geq 2^{\aleph_0}$ for all $\alpha > 0$. To see this we split into cases: If T is not ω-stable, this follows from Lemma 5.2. If T is ω-stable this can be verified by examining the spectra given by Saffe [16].

Next, fix an na-chain $\overline{\mathcal{M}}$ of length n for which $R(\overline{\mathcal{M}})$ contains infinitely many pairwise orthogonal types. By Lemma 3.6 there is a diffuse family of $Leaves(\overline{\mathcal{M}})$ of size $\aleph_0$. We split into two cases.

If $|\alpha + \omega|^{\aleph_0} > 2^{\aleph_0}$, then $\alpha > 2^{\aleph_0}$, hence $\aleph_\alpha > 2^{\aleph_0}$. Thus, Lemma 5.4 provides us with a family of $|\alpha+\omega|^{\aleph_0}$ $(n-1)$-iterable models over M_{n-1}, each of size $\aleph_\alpha$. So, we obtain our lower bound via Lemma 4.6 or Lemma 4.8.

On the other hand, assume that $|\alpha+\omega|^{\aleph_0} = 2^{\aleph_0}$. If $n = 1$, then $I(T, \aleph_\alpha) \geq 2^{\aleph_0}$ as noted above, so we assume that $n > 2$. There are now three subcases.

- If $\aleph_\alpha > 2^{\aleph_0}$, then there is an $(n-2)$-iterable family of size $\beth_3$, where each model has size $\aleph_\alpha$. As before, the bound follows from Lemmas 4.6 and Lemma 4.8.

- If $\aleph_\alpha \leq 2^{\aleph_0}$, but $2^{\aleph_\alpha} > 2^{\aleph_0}$, then the existence of a diverse family of size $\aleph_0$ implies that there is a family $\mathcal{F} = \{P_i : i \in \aleph_\alpha\}$ of countable models over M_{n-1} such that $P_i \oplus_{M_{n-2}} U \not\cong_U P_j \oplus_{M_{n-2}} U$ for distinct $i, j \in \aleph_\alpha$. Varying the dimensions of each of these yields a family of $2^{\aleph_\alpha}$ models, pairwise non-isomorphic over M_{n-2}, each of of size $\aleph_\alpha$. However, as $\aleph_\alpha^{\aleph_0} = 2^{\aleph_0}$ in this case, we obtain $I(T, \aleph_\alpha) = 2^{\aleph_\alpha}$.

- Finally, if $2^{\aleph_\alpha} = 2^{\aleph_0}$ then $I(T, \aleph_\alpha) = 2^{\aleph_0}$ as in the first paragraph. ∎

Using the notation given at the beginning of this section, we can easily summarize our lower bound results.

Theorem 5.10 *Let T be a countable, classifiable theory of depth $d \geq n$.*

1. *If $\#RD(n) = 2^{\aleph_0}$ then $I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{n-1}(|\alpha + \omega|^{2^{\aleph_0}})\}$ for all ordinals $\alpha > 0$.*

2. *If $\#RD(n) = \aleph_0$ then $I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{n-1}(|\alpha + \omega|^{\aleph_0})\}$ for all ordinals $\alpha > 0$.*

3. *If $TF(n)$ holds then $I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{n-1}(|\alpha + \omega|^{2^{\aleph_0}})\}$ for all ordinals $\alpha > 0$.*

4. *If $NTF(n)$ holds then*

$$I(T, \aleph_\alpha) \geq \begin{cases} \min\{2^{\aleph_\alpha}, \beth_2\} & \textit{if } n = 1 \\ \min\{2^{\aleph_\alpha}, \beth_{n-2}(|\alpha + \omega|^{\beth_2})\} & \textit{if } n > 1 \end{cases}$$

   for all ordinals $\alpha > 0$.

Proof. (1) Choose a chain $\overline{\mathcal{M}}$ of length n such that $R(\overline{\mathcal{M}})$ contains a pairwise orthogonal family of types. By passing to a free extension of $\overline{\mathcal{M}}$, we may assume that $\overline{\mathcal{M}}$ is an *na*-chain. Thus, it follows from Lemma 3.6 that there is a diffuse family of $Leaves(\overline{\mathcal{M}})$ of size continuum, so the lower bound follows immediately from Proposition 5.6.

(2) If some chain $\overline{\mathcal{M}}$ has an infinite family of pairwise orthogonal types in $R(\overline{\mathcal{M}})$, then this property will be inherited by any *na*-chain that freely extends $\overline{\mathcal{M}}$, so the result follows immediately from Proposition 5.9.

(3) Suppose that some chain $\overline{\mathcal{M}}$ of length n supports a trivial type $p \in R(\overline{\mathcal{M}})$ that is not totally transcendental above $\overline{\mathcal{M}}$. By Lemma 3.9 we may assume that p is special. Since p is not totally transcendental, it follows from Corollary 3.23(2) and Proposition 3.27(2) that there is a free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ with a diffuse family of $Leaves(\overline{\mathcal{M}}')$ of size continuum. So the bound follows from Proposition 5.6.

(4) This is analogous to (3). If $p \in R(\overline{\mathcal{M}})$ is not totally transcendental over $\overline{\mathcal{M}}$, then, again using Lemma 3.9, it follows from Corollary 3.23(1) and Proposition 3.27(1) that there is a free extension $\overline{\mathcal{M}}'$ of $\overline{\mathcal{M}}$ with a diverse family of $Leaves(\overline{\mathcal{M}}')$ of size continuum. Thus, the lower bound is given by Proposition 5.8. ∎

5.2 Upper bounds

In this short subsection we state some definitions and recall a very useful theorem for obtaining upper bounds.

Definition 5.11 Suppose that $\overline{\mathcal{M}}$ is an n-chain.

- An $\overline{\mathcal{M}}$-*component* is any model N such that $M_{n-1} \subseteq_{na} N$ and $wt(N/M_{n-1}) = 1$.
- We write $\#C^{\alpha}_{\overline{\mathcal{M}}}$ for the number of $\overline{\mathcal{M}}$-components of size at most $\aleph_\alpha$, up to isomorphism over M_{n-1} and
- write $\#C^{\alpha}_{n} = \sup\{\#C^{\alpha}_{\overline{\mathcal{M}}} : \overline{\mathcal{M}}$ is a $(d-n)$-chain$\}$ *unless* $n = 0$ and $\#C^{\alpha}_{\overline{\mathcal{M}}}$ is finite for all d-chains. In this case, put $\#C^{\alpha}_{0} = 1$.

Proposition 5.12 *1.* $\#C^{\alpha}_{0} \leq 2^{\aleph_0}$ *for all* α.

2. If $TT(d)$ *holds then* $\#C^{\alpha}_{0} = \#RD(d)$ *unless the latter is finite in which case* $\#C^{\alpha}_{0} = 1$.

The second part of the above Proposition follows directly from the definition of $TT(d)$ and Corollary 3.31.

The following theorem is proved by inductively counting the number of components as we step down a decomposition tree.

Theorem 5.13 *If* T *is a countable, classifiable theory with finite depth* d *then*

$$\#C^{\alpha}_{i+1} \leq |\alpha + \omega|^{\#C^{\alpha}_{i}} + 2^{\aleph_0}$$

and

$$I(T, \aleph_\alpha) \leq \beth_{d-i-1}(|\alpha + \omega|^{\#C^{\alpha}_{i}} + 2^{\aleph_0})$$

Proof. By downward induction on i, using the fact that every model is prime over a normal tree of countable, na-substructures. ∎

As a Corollary, we obtain the naive upper bound given in the Introduction.

Corollary 5.14 *If T is countable, classifiable and has finite depth d then*

$$I(T, \aleph_\alpha) \leq \beth_{d-1}(|\alpha + \omega|^{2^{\aleph_0}})$$

Combining this upper bound with Theorem 5.10 yields the following spectrum:

Corollary 5.15 *If $TF(d)$ holds or $\#RD(d) = 2^{\aleph_0}$ then*

$$I(T, \aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha + \omega|^{2^{\aleph_0}})\}.$$

In what follows, the assumptions of the given subsection are indicated in the subsection heading. Under these assumptions, a general upper bound will be derived and at the end of each subsection, we will indicate under what conditions this upper bound is met.

5.3 $TF(d)$ fails, $\#RD(d) \leq \aleph_0$ and $d > 1$

The hypotheses imply that we have control of the trivial components at level d. In order to get a better upper bound, we record a theorem of Shelah that allows us to control the number of non-trivial components as well. The following lemma is Lemma 4.5 of Chapter XIII of [18].

Lemma 5.16 *Suppose that T is countable and classifiable. If $M_0 \subseteq_{\aleph_1} M_i$ and $tp(M_i/M_0)$ has weight one and is non-trivial for $i = 1, 2$ and $tp(M_1/M_0)$ and $tp(M_2/M_0)$ are not orthogonal then $M_1 \cong M_2$ over M_0.*

Fix a $(d-1)$-chain $\overline{\mathcal{M}}$ and N, an $\overline{\mathcal{M}}$-component of size at most $\aleph_\alpha$. We wish to find a set of invariants which determines the isomorphism type of N over M_{d-2}.

Choose $M_{d-1} \subseteq_{na} N$ so that M_{d-1} is countable and properly contains M_{d-2}. Now let I be a maximal M_{d-1}-independent collection of countable models contained in N, where each is a weight one, na-extension of M_{d-1}. Since T has depth d, N is prime over I. Since $TF(d)$ fails, each $N' \in I$ for which $tp(N'/M_{d-1})$ is not orthogonal to a trivial regular type, is actually determined, up to isomorphism over M_{d-1} by this information alone (see Corollary 3.33). However, we do not have such control over the non-trivial types. In order to remedy this, we do the following.

Choose $I_0 \subseteq I$, $|I_0| \leq 2^{\aleph_0}$ such that M', the prime model over I_0 is an $\aleph_1$-substructure of N. By Fact 5.16, the isomorphism types of the non-trivial components are determined, up to isomorphism over M', by their non-orthogonality class. Since the non-trivial (and trivial) components present in I are all based on M_{d-1}, there are $|\alpha+\omega|^{\#RD(d)}$ possibilities for the isomorphism type of N over M'. We did have to fix a countable model M_{d-1} over M_{d-2} and a model M' of size at most $2^{\aleph_0}$. Clearly, there are at most $\beth_2$ many possibilities for these choices, up to isomorphism over M_{d-2}.

In summary then, if $TF(d)$ fails and $d>1$, then $\#C_1^\alpha \leq |\alpha+\omega|^{\#RD(d)}+\beth_2$. Thus, by Theorem 5.13, we have the following:

$$I(T,\aleph_\alpha) \leq \beth_{d-2}(|\alpha+\omega|^{|\alpha+\omega|^{\#RD(d)}+\beth_2}+2^{\aleph_0}) = \beth_{d-1}(|\alpha+\omega|^{\#RD(d)}+\beth_2).$$

Combining these upper bounds with the lower bounds of Theorem 5.10 yield the following corollaries.

Corollary 5.17 *1. If $TF(d)$ fails, $NTF(d)$ holds, $\#RD(d)=\aleph_0$, and $d>1$ then*

$$I(T,\aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha+\omega|^{\aleph_0}+\beth_2).\}$$

2. If $TF(d)$ fails, $NTF(d)$ holds, $\#RD(d)$ is finite, and $d>1$ then

$$I(T,\aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha+\omega|+\beth_2).\}$$

Proof. For both cases, note that $I(T,\aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{d+1}\}$ since $NTF(d)$ holds. If $\#RD(d)=\aleph_0$ then the bound $I(T,\aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha+\omega|^{\aleph_0})\}$ follows from Theorem 5.10(2), Combining these lower bounds with the upper bound mentioned above yields the spectrum in (1).

On the other hand, if $\#RD(d)$ is finite, then by combining the lower bound of the preceding paragraph with the general lower bound (Lemma 5.3), we match the upper bound mentioned above. ∎

Thus, except for the case when $d=1$ which we handle in Subsection 5.8, we have computed the spectra of all (classifiable of finite depth d) theories for which $TT(d)$ fails or $\#RD(d)=2^{\aleph_0}$.

5.4 $TT(d)$ holds and $\#RD(d) \leq \aleph_0$

Since $TT(d)$ holds, the positive results from Subsection 3.4 come into play and yield substantially better upper bounds. In particular, Proposition 5.12 implies $\#C_0^\alpha = \aleph_0$ if $\#RD(d) = \aleph_0$ and 1 if $\#RD(d)$ is finite.

Corollary 5.18 *If $TT(d)$ holds and $\#RD(d) = \aleph_0$ then*

$$I(T, \aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha + \omega|^{\aleph_0})\}.$$

Proof. Since $\#C_0^\alpha = \aleph_0$, it follows from Theorem 5.13 that an upper bound is

$$\beth_{d-1}(|\alpha + \omega|^{\aleph_0} + 2^{\aleph_0}) = \beth_{d-1}(|\alpha + \omega|^{\aleph_0}).$$

However, the matching lower bound is immediate from Theorem 5.10(2). ∎

Suppose that $\#RD(d)$ is finite. Then Theorem 5.13 yields an upper bound of

$$\beth_{d-1}(|\alpha + \omega| + 2^{\aleph_0}).$$

Note that if $d > 1$ then we can obtain a matching lower bound if either $TT(d-1)$ fails or $\#RD(d-1) = 2^{\aleph_0}$.

Corollary 5.19 *If $TT(d)$ holds, $\#RD(d)$ is finite, $d > 1$, and either $TT(d-1)$ fails or $\#RD(d-1) = 2^{\aleph_0}$, then*

$$I(T, \aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha + \omega| + 2^{\aleph_0})\}.$$

Proof. If either $TF(d-1)$ holds or $\#RD(d-1) = 2^{\aleph_0}$, then the lower bound follows immediately from Theorem 5.10. If $NTF(d-1)$ holds, then the lower bound is obtained by combining the lower bound of Theorem 5.10(4) with the general lower bound of Lemma 5.3. ∎

5.5 Obtaining na-inclusion

Before continuing, we give a technical lemma and a construction that are needed to establish the more fussy upper bounds.

Lemma 5.20 *Suppose that $\overline{\mathcal{M}}$ is a chain of length $n > 1$ and $TT(n)$ holds. If $M_{n-1} \subseteq N$ such that $tp(N/M_{n-1})$ is orthogonal to M_{n-2} and every strongly regular type over M_{n-1} which is realized in $N \setminus M_{n-1}$ has infinite dimension in M_{n-1} then $M_{n-1} \subseteq_{na} N$.*

Proof. Suppose that $\varphi(x, \bar{a})$ has a solution in $N \setminus M_{n-1}$ where $\bar{a} \in M_{n-1}$. By relativizing the proof that some regular type is realized between any pair of models of a superstable theory, there is an element $d \in dcl(\varphi(N)) \setminus M_{n-1}$ such that $q = \text{tp}(d/M_{n-1})$ is regular. Since $TT(n)$ holds, there is a strongly regular type $p \in R(\overline{\mathcal{M}})$ non-orthogonal to q. Without loss, we may assume that p is based on $\overline{a}$. So, by Lemma 3.29 there a realization c of p in $N \setminus M_{n-1}$ that depends on d over M_{n-1}. Choose $b \in \varphi(N) \setminus M_{n-1}$ that depends on c over M_{n-1} and choose $\overline{e} \supseteq \overline{a}$ from M_{n-1} such that b and c are dependent over $\bar{e}$. Let the formula $\chi(b, c, \bar{e})$ witness this dependence. Let I be an infinite Morley sequence in the type of $p|\bar{a}$ inside M_{n-1}. Choose $c' \in I$ so that c' and $\bar{e}$ are independent over $\overline{a}$. We have

$$\exists \chi(x, c, \bar{e}) \wedge \varphi(x, \bar{a})$$

which is also true when c' replaces c. So pick b' so that

$$\chi(b', c', \bar{e}) \wedge \varphi(b', \bar{a})$$

holds. Since b' and c' are dependent over $\bar{e}$, b' cannot be in the algebraic closure of $\bar{a}$ and so we finish. ∎

We next describe a construction that will be used in the next subsections.

Construction 5.21 Fix an chain $\overline{\mathcal{M}}$ of length $(n-1)$ and an $\overline{\mathcal{M}}$-component N. Suppose that $n > 1$, $TT(n-1)$ and $TT(n)$ hold and $\#RD(n)$ is finite. Since $TT(n-1)$ holds, we can find $a \in N \setminus M_{n-2}$ such that $\text{tp}(a/M_{n-2})$ is strongly regular. Let N_0 be contained in N and prime over $M_{n-2}a$; again, N_0 exists because we are assuming $TT(n-1)$. Let $\overline{\mathcal{N}}_0$ be the chain of length n obtained by concatenating N_0 to $\overline{\mathcal{M}}$. We will define, by induction on i, countable models $N_i \subseteq N$, numbers n_i and strongly regular types $p^i_l \in R(N_i)$ for $l < n_i$ which are pairwise orthogonal.

So suppose that N_i has been defined. Let $\overline{\mathcal{N}}_i$ be the chain of length n obtained by concatenating N_i to $\overline{\mathcal{M}}$. Let n_i be the number of non-orthogonality classes in $R(\overline{\mathcal{N}}_i)$ which are realized in N and which are orthogonal to p^j_l for all $j < i$ and $l < n_j$. Since $TT(n)$ holds, we can find strongly regular representatives p^i_l for $l < n_i$ of these classes.

Now, for $l < n_i$, let I^i_l be a maximal Morley sequence in N for p^i_l if the dimension of this type in N is countable and otherwise, let I^i_l be any countable, infinite Morley sequence in N for p^i_l. Let I^i be the concatenation of the sequences I^i_l and finally, let N_{i+1} be prime over $N_i I^i$.

We next argue that this process must stop after finitely many steps. To see this, notice that if not, then if $N_\omega = \bigcup_i N_i$, $\overline{\mathcal{N}}_\omega$ together with all the p^i_l's exemplifies that $\#RD(n)$ is not finite. What have we achieved? If N_i is the last element of this chain then by Lemma 5.20, $N_i \subseteq_{na} N$. To see this, suppose that $b \in N \setminus N_i$ such that $q = tp(b/N_i)$ is strongly regular. Then for some j and l, q is not orthogonal to p^j_l and the dimension of this latter type in N_i is infinite.

5.6 $TT(d)$ and $TT(d-1)$ hold, $\#RD(d)$ is finite, $d > 1$, and $\#RD(d-1) \leq \aleph_0$

Our first goal is to obtain an upper bound for theories satisfying these hypotheses. This is accomplished by analyzing Construction 5.21 in detail.

Fix a chain $\overline{\mathcal{M}}$ of length $d-1$ and an $\overline{\mathcal{M}}$-component N of size at most $\aleph_\alpha$. Pick $a \in N$ such that $tp(a/M_{d-2})$ is strongly regular.

Let N' be the model described in Construction 5.21. $N' \subseteq_{na} N$ and so by Lemma 3.34, N is prime over N' and I, a strongly regular sequence over N'. By the construction of N' there are $|\alpha+1|^{\#RD(d)}$ possibilities for I up to isomorphism over N'. Now the construction of N' was accomplished in finitely many steps. We see that at each step i, there were at most countably many choices for I^i and at the first stage, since $\#RD(d-1) \leq \aleph_0$, countably many choices for N_0 so there are at most countably many choices for N'.

In summary, there are at most $|\alpha+\omega|$ many isomorphism types of N over M_{d-2}, i.e., $\#C^\alpha_1 \leq |\alpha+\omega|$. Hence, under the assumptions of this subsection, we get an upper bound of

$$\beth_{d-2}(|\alpha+\omega|^{|\alpha+\omega|} + 2^{\aleph_0}) = \beth_{d-1}(|\alpha+\omega|).$$

Note that this upper bound matches the general lower bound (Lemma 5.3) whenever α is infinite. So, in the computation of the spectra that follow, we are only interested in computing $I(T, \aleph_\alpha)$ when α is finite. We can dispense with a number of cases at this point.

Corollary 5.22 *Suppose that* $TT(d)$ *and* $TT(d-1)$ *hold,* $d > 1$ *and* $\#RD(d)$ *is finite. If any of the following three conditions hold*

- $d > 1$ *and* $\#RD(d-1) = \aleph_0$*;*
- $d > 2$ *and* $TT(d-2)$ *fails; or*
- $d > 2$ *and* $\#RD(d-2) = 2^{\aleph_0}$

then $I(T, \aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha + \omega|)\}$.

Proof. All three cases follow immediately by combining lower bounds from Theorem 5.10 with the general lower bound (Lemma 5.3) and matching the upper bound mentioned above. ∎

To distinguish between the spectra $\beth_{d-1}(|\alpha+\omega|)$ and $\beth_{d-2}(|\alpha+\omega|^{|\alpha+1|})\}$, we need one further dichotomy.

Definition 5.23 Suppose that T is a countable, classifiable theory with finite depth $d > 1$ and moreover, if both $TT(d)$ and $TT(d-1)$ hold and $\#RD(d)$ is finite. We say that T has the *final property* if for every chain $\overline{\mathcal{M}}$ of length d, there are only finitely many isomorphism types of models N over M_{d-2} of the form $Pr(M_{d-1} \cup J)$, where J is a countable, strongly regular sequence from $RD(\overline{\mathcal{M}})$.

We remark that in the case where T is ω-stable (and satisfies the other properties) T has the final property if and only if all types of depth $d-2$ are abnormal (of Type V) in the sense of Baldwin [1].

Corollary 5.24 *Suppose that* $d > 1$*,* $TT(d)$ *and* $TT(d-1)$ *hold, and* $\#RD(d)$ *is finite. If* T *does not have the final property, then*

$$I(T, \aleph_\alpha) = \min\{2^{\aleph_\alpha}, \beth_{d-1}(|\alpha + \omega|)\}.$$

Proof. First, if T is totally transcendental, then this is proved in Saffe [16] or Baldwin [1], so we assume that T is not totally transcendental. Hence, by Lemma 3.35, we may assume that $d > 2$. So, by the previous Corollary, we may further assume that $\#RD(d-1)$ is finite, $TT(d-2)$ holds, and that $\#RD(d-2) \leq \aleph_0$. As well, if $\aleph_\alpha \leq 2^{\aleph_0}$, then we are done by Lemma 5.2, so choose α such that $\aleph_\alpha > 2^{\aleph_0}$. As noted above, we may assume α is finite (else the upper bound already matches the general lower bound). We need to show that there are at least

$$I(T, \aleph_\alpha) \geq \min\{2^{\aleph_\alpha}, \beth_{d-1}\}.$$

It is easily checked that if the final property fails for some chain, then it fails for some na-chain, so choose $\overline{\mathcal{M}}$ an na-chain of length d for which the final property fails. Then, by using Lemma 5.6 of Chapter XVIII of [1] one obtains a family $\{P_i : i \in 2^{\aleph_0}\}$ of countable models over M_{d-2} such that $P_i \oplus_{M_{n-3}} U \ncong_U P_j \oplus_{M_{d-3}} U$ for distinct i, j. Then, arguing as in the last two paragraphs of Lemma 5.7, we obtain a $(d-3)$-iterable family of models of size $\beth_2$, each of which is of size $\aleph_\alpha$. Hence, we obtain our lower bound from Lemmas 4.6 and 4.8. ∎

5.7 The final case

The assumptions of this case are too many to put in the subsection heading. We will assume that $d > 2$, $TT(d)$, $TT(d-1)$ and $TT(d-2)$ hold, $\#RD(d)$ and $\#RD(d-1)$ are finite, $\#RD(d-2) \leq \aleph_0$, and assume that the final property holds. The case when $d = 2$ is handled in the next subsection.

In order to compute the best general upper bound in this case, we will compute $\#C_2^\alpha$. Toward this end, fix a chain $\overline{\mathcal{M}}$ of length $d-2$ and an $\overline{\mathcal{M}}$-component N of size at most $\aleph_\alpha$. Now fix $a \in N$ such that $tp(a/M_{d-3})$ is strongly regular and let N_0 be the prime model over $M_{d-3}a$. Since we are assuming that $TT(d-2)$ holds and $\#RD(d-2) \leq \aleph_0$, there are at most countably many choices for N_0 up to isomorphism over M_{d-3}. Now by using Construction 5.21, we can find $M_{d-2} \subseteq_{na} N$ and by using the argument from the previous subsection, there are at most countably many choices for M_{d-2} over N_0. Let $\overline{\mathcal{M}}'$ be the chain of length $d-1$ formed by concatenating M_{d-2} to $\overline{\mathcal{M}}$. Now in order to understand N over M_{d-2}, it suffices to understand the isomorphism types of $\overline{\mathcal{M}}'$-components inside N. Let N' be the model described in Construction 5.21, now working over $\overline{\mathcal{M}}'$. Again, $N' \subseteq_{na} N$ and N is prime over N' and a strongly regular sequence I over N'. As before, there are only $|\alpha+1|^{\#RD(d)}$ possibilities for I over N'. However, since $\#RD(d-1)$ is finite, there are only finitely many choices for the first model in the construction of N'. In addition, since the final property holds, then at each stage there are only finitely many choices for the i^{th} model. And, as we noted, the construction of N' is accomplished in finitely many steps and so by König's Lemma, there are only finitely many choices for N' in all.

In summary, there are at most $|\alpha+\omega|$ choices for M_{d-2} and $|\alpha+l|$ choices for N over M_{d-2}. Hence, over M_{d-2}, there are at most $|\alpha+\omega|^{|\alpha+1|}$ many such N up to isomorphism. We conclude then that

$$\#C_2^\alpha \leq |\alpha+\omega|^{|\alpha+1|} + \aleph_0 = |\alpha+\omega|^{|\alpha+1|}$$

and the upper bound in this case is

$$\beth_{d-3}(|\alpha+\omega|^{|\alpha+\omega|^{|\alpha+1|}} + 2^{\aleph_0}) = \beth_{d-2}(|\alpha+\omega|^{|\alpha+1|}).$$

As this agrees with the general lower bound, this is the spectrum in this case.

5.8 The case of d = 1 or 2

In this section we would like to take care of the case when $d = 1$ and finish the previous subsection in the case when $d = 2$.

To improve the upper bounds in the case $d = 1$ we make several remarks. If $TF(1)$ fails and $\#RD(1) = 1$ (this case was handled in subsection 5.3), we actually obtain the upper bound $\beth_2$. This can be seen even in the proof presented in that subsection. After fixing a model of size $2^{\aleph_0}$, there is only one dimension possible. Of course, the spectrum $\min\{2^{\aleph_\alpha}, \beth_2\}$ is achieved when T is unidimensional and not ω-stable.

Next, we improve the upper bounds when $TT(1)$ holds and $\#RD(1)$ is finite. It follows directly from Lemma 3.35 that such theories are totally transcendental, so we could quote Saffe [16]. However, we include a brief discussion for completeness. The upper bound then, from subsection 5.3 is

$$|\alpha+\omega|.$$

If $\#RD(1)$ is not 1 then in [11] and [12], Lachlan shows that this spectrum is correct under these assumptions unless T is ω-categorical. In those papers, he shows that if T is ω-categorical then for some number m and some $G \leq Sym(m)$,

$$I(T, \aleph_\alpha) = |(\alpha+1)^m/G| - |(\alpha)^m/G|$$

Of course, the case when $TT(1)$ holds and $\#RD(1) = 1$ is when T is $\aleph_1$-categorical and the upper bound is 1.

The case $d = 2$ only has to be distinguished in the previous subsection. There we were assuming $TT(d-1)$ and $\#RD(d-1)$ is finite so again T is totally transcendental. To obtain an upper bound in this case, we look to subsection 5.6 and so compute only $\#C_1^\alpha$ which in this case (assuming the final property) would be $|\alpha + l|$ for some finite number l. So the upper bound becomes

$$|\alpha + \omega|^{|\alpha+1|}$$

which agrees with the general lower bound in this case.

6 The spectra

Collecting together all the spectra from Section 4 with the spectra mentioned in the introduction, we obtain the following Theorem. This theorem was announced in [7], where examples of theories with each of these spectra were given.

Theorem 6.1 *For any countable, complete theory T with an infinite model, the uncountable spectrum $\aleph_\alpha \mapsto I(T, \aleph_\alpha)$ $(\alpha > 0)$ is the minimum of the map $\aleph_\alpha \mapsto 2^{\aleph_\alpha}$ and one of the following maps:*

1. $2^{\aleph_\alpha}$;
2. $\beth_{d+1}(|\alpha+\omega|)$ for some d, $\omega \le d < \omega_1$;
3. $\beth_{d-1}(|\alpha+\omega|^{2^{\aleph_0}})$ for some d, $0 < d < \omega$;
4. $\beth_{d-1}(|\alpha+\omega|^{\aleph_0}+\beth_2)$ for some d, $0 < d < \omega$;
5. $\beth_{d-1}(|\alpha+\omega|+\beth_2)$, for some d, $0 < d < \omega$;
6. $\beth_{d-1}(|\alpha+\omega|^{\aleph_0})$, for some d, $0 < d < \omega$;
7. $\beth_{d-1}(|\alpha+\omega|+2^{\aleph_0})$, for some d, $1 < d < \omega$;
8. $\beth_{d-1}(|\alpha+\omega|)$, for some d, $0 < d < \omega$;
9. $\beth_{d-2}(|\alpha+\omega|^{|\alpha+1|})$, for some d, $1 < d < \omega$;
10. identically $\beth_2$;
11. $\begin{cases} |(\alpha+1)^n/\sim_G| - |\alpha^n/\sim_G| & \alpha < \omega; \text{ for some } 1 < n < \omega \text{ and} \\ |\alpha| & \alpha \ge \omega; \text{ some group } G \le Sym(n) \end{cases}$
12. identically 1.

A Appendix: On na-inclusions

In this first appendix we establish two facts about na-inclusions that are used in the text. To simplify the proofs, we extend the definition of an na-extension to arbitrary sets. (This more general definition is only used in this Appendix.)

Definition A.1 If $A \subseteq B$ then $A \subseteq_{na} B$ if whenever $\varphi(x) \in L(A)$ such that $\varphi(B) \setminus A$ is non-empty and $F \subseteq A$ is any finite set then $\varphi(A) \setminus acl(F)$ is non-empty.

The following lemma follows immediately from a union of chains argument.

Lemma A.2 *For any set B and any $A \subseteq B$, there is $A' \subseteq_{na} B$ such that $A \subseteq A'$ and $|A'| \le |A| + |L|$.*

Lemma A.3 *If B is independent from C over A and $A \subseteq_{na} B$ then $C \subseteq_{na} BC$.*

Proof: Suppose that we fix $c \in C$ and $b \in B$ so that $\theta(b, c)$ holds and $b \notin acl(C)$. It suffices to find $b' \in C$ (in fact we will find it in A) so that $\theta(b', c)$ holds and $b' \notin acl(c)$. Let $p = stp(c/B)$. Then p is based on some $a \in A$ and b satisfies $d_p y\theta(x, y)$ which is a formula almost over a. Choose $b' \in A \setminus acl(a)$ which satisfies the same formula. Then immediately $\theta(b', c)$ holds and $b' \notin acl(c)$. ∎

Lemma A.4 *If $M \subseteq_{na} A$ and A dominates B over M then $M \subseteq_{na} B$.*

Proof. Suppose we fix $m \in M$ and $b \in B \setminus M$ so that $\theta(b)$ holds for some $\theta \in L(M)$. Let $c = Cb(stp(b/A)$. $c \in acl(A) \setminus M$ since $b \notin M$. Moreover, $c \in dcl(\bar{d})$ for a finite sequence of realizations of θ, $\bar{d}$. Choose $c' \in M \setminus acl(m)$ and $\bar{d}' \in \theta(M)$ so that $c' \in dcl(\bar{d}')$. Since $c' \notin acl(m)$, one of these realizations must also not be in $acl(m)$. ∎

Corollary A.5 *If M_1 and M_2 are independent over M_0, $M_0 \subseteq_{na} M_1$ and $N = M_1 \oplus_{M_0} M_2$ then $M_2 \subseteq_{na} N$.*

Definition A.6 If $M \subseteq A$ then we say that A is na-extendible over M if there is a model N, $A \subseteq N$ and $M \subseteq_{na} N$.

The following lemma is implicit in [19].

Lemma A.7 *Suppose that $M \subseteq A$ and A is na-extendible over M. For any consistent formula $\varphi(x) \in L(A)$, if $\psi(x) \in L(A)$ is a consistent formula of least R^∞-rank which implies $\varphi(x)$ then if $\psi(b)$ holds and $M \subseteq_{na} Ab$ then Ab is dominated by A over M.*

B Appendix: G_δ subsets of Stone spaces

In this appendix we note two facts from descriptive set theory and then establish that several subsets of the Stone space $S(A)$ are G_δ with respect to the usual topology when A is algebraically closed. These results are used in the proof of Proposition 3.27.

Lemma B.1 *Let X be any Polish (i.e., separable, complete metric) space.*

1. *Let E be a Borel equivalence relation on X. Either E has countably many classes or there is a perfect set of pairwise E-inequivalent elements of X.*

2. *Any non-empty subset of a countable, G_δ subset of X has an isolated point.*

Proof of 2. Suppose $A = \bigcap_{n\in\omega} U_n$ is non-empty and countable, where each U_n is open in X. Since every subset of A is also a G_δ, it suffices to show that A has an isolated point. Let $\bar{A}$ denote the topological closure of A in X. As $\bar{A}$ is Polish, it is a Baire space. However, every $U_n \cap \bar{A}$ is dense in $\bar{A}$ and

$$\bigcap_{n\in\omega}(U_n \cap \bar{A}) \cap \bigcap_{a\in A} \bar{A} \smallsetminus \{a\} = \emptyset,$$

so it follows that $\bar{A} \smallsetminus \{a\}$ is not dense in $\bar{A}$ for some $a \in A$. That is, this a is isolated in A. ∎

Fix A countable and algebraically closed. Then the space of types $S(A)$ (with the usual topology) is a Polish space.

Lemma B.2 *$\{p \in S(A) : p$ is trivial and $wt(p) = 1\}$ is a G_δ subset of $S(A)$.*

Proof. We use the following characterization, which is implicit in Theorem 1.8 of [3].

A type p is trivial of weight 1 if and only if

1. For all countable $B \supseteq A$, all weight 1 types q and r over B, and all realizations $\langle b, c\rangle$ of $q \otimes r$, if a realizes $p|B$, $a \underset{B}{\perp} b$, and $a \underset{B}{\perp} c$, then $a \underset{B}{\perp} bc$.
2. p has weight at least 1, i.e., for all $B \supseteq A$, for all types q, r over B, and all realizations $\langle b, c\rangle$ of $q \otimes r$, if a realizes $p|B$ and $a \underset{B}{\not\perp} b$, then $a \underset{B}{\perp} c$.
3. p is not algebraic.

Now, for all three of these conditions, a specific instance of its negation is witnessed by a formula. Hence, a type p is not trivial, weight 1 if and only if it is not contained in one of the "bad" formulas. At first glance, it appears like there are continuum constraints, but since the language is countable there really are only countably many. Hence the negation of our set is an F_σ, so we finish. ∎

Lemma B.3 *$X_B = \{p : p$ determines a complete type over $B\}$ is a G_δ subset of $S(A)$ for every countable $B \supseteq A$. Hence, $\{p : p \underset{A}{\perp}^a q\}$ is G_δ in $S(A)$ for every strong type q that is based on A.*

Proof. The second sentence follows immediately from the first. Fix a countable $B \supseteq A$. For every $L(B)$-formula $\delta(x)$ we say that the $L(A)$-formula *decides* δ if

$$\varphi(x) \wedge \varphi(y) \quad \text{implies} \quad \delta(x) \leftrightarrow \delta(y).$$

Let $D(\delta) = \{p : \text{there is some } \varphi \in p \text{ that decides } \delta\}$. Clearly, each set $D(\delta)$ is open. Thus, it suffices to show that $p \in X_B$ if and only if $p \in D(\delta)$ for all $\delta(x) \in L(B)$. However, if $p \in X_B$, then

$$p(x) \wedge p(y) \quad \text{implies} \quad \operatorname{tp}(x/B) = \operatorname{tp}(y/B),$$

so it follows by compactness that some $\varphi \in p$ decides δ for every $\delta \in L(B)$.

Conversely, suppose $p \in D(\delta)$ for all $\delta \in L(B)$. Then, given a, b each realizing p, $\delta(a) \leftrightarrow \delta(b)$ for all $\delta \in L(B)$, hence $\operatorname{tp}(a/B) = \operatorname{tp}(b/B)$. ∎

Corollary B.4 *$\{p : p$ is trivial, $wt(p) = 1$ and $p \perp q\}$ is a G_δ for every such q.*

Proof. Recall that if p is trivial and $wt(p) = 1$, then $p \perp q$ if and only if $p \underset{A}{\perp}^a q^{(\omega)}$. ∎

Lemma B.5 *$\{p \in S(A) : p \underset{A}{\perp}^a A'\}$ is a G_δ subset of $S(A)$ for any algebraically closed $A' \subseteq A$. In particular,*

$$\{p : p \textit{ is trivial, } wt(p) = 1 \textit{ and } p \perp A'\}$$

is a G_δ for any such A'.

Proof. Since $p \mathop{\perp}\limits_{A}^{a} A'$ implies $p \perp A'$ among trivial, weight 1 types, the second sentence follows from the first. To conclude the first, say that the $L(A)$-formula $\psi(x)$ *decides* $\delta(x,y) \in L(A)$ if

$$\forall x_1 \forall x_2 \forall y \left(\psi(x_1) \wedge \psi(x_2) \wedge y \mathop{\perp}\limits_{A'} A \to [\delta(x_1,y) \leftrightarrow \delta(x_2,y)] \right)$$

and let $D(\delta) = \{p \in S(A) : \text{some } \psi \in p \text{ decides } \delta(x,y)\}$. Clearly, each $D(\delta)$ is an open subset of $S(A)$. Thus, it suffices to prove that $p \mathop{\perp}\limits_{A}^{a} A'$ if and only if $p \in D(\delta)$ for all $\delta(x,y) \in L(A)$.

To see this, first assume that $p \mathop{\perp}\limits_{A}^{a} A'$. Then $c \mathop{\perp}\limits_{A} d$ for any c realizing p and any $d \mathop{\perp}\limits_{A'} A$. That is,

$$\forall x_1 \forall x_2 \forall y \left(p(x_1) \wedge p(x_2) \wedge y \mathop{\perp}\limits_{A'} A \to [\mathrm{tp}(x_1 y/A) = \mathrm{tp}(x_2 y/A)] \right),$$

so by compactness for every $\delta(x,y) \in L(A)$ there is a $\psi(x) \in p$ deciding δ.

Conversely, suppose $p \in D(\delta)$ for every $\delta(x,y) \in L(A)$. Take c realizing p and d freely joined from A over A'. Choose c^* realizing $p|Ad$. Since $p \in D(\delta)$ for each $\delta(x,y)$, it follows that $\mathrm{tp}(cd/A) = \mathrm{tp}(c^*d/A)$, so $c \mathop{\perp}\limits_{A} d$ as required. ∎

References

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

[2] J. Baldwin, Diverse classes, *Journal of Symbolic Logic* **54** (1989), 875-893.

[3] J. Baldwin and L. Harrington, Trivial pursuit: remarks on the main gap, *Annals of Pure and Applied Logic* **34** (1987), 209–230.

[4] A. Ehrenfeucht, On theories categorical in power, *Fundamenta Mathematicae*, **44** (1957), 241–248.

[5] L. Harrington and M. Makkai, An exposition of Shelah's 'Main Gap' – counting uncountable models of ω-stable and superstable theories, *Notre Dame Journal of Formal Logic*, **26** (1985), 139–177.

[6] B. Hart, An exposition of OTOP, In: *Classification Theory: Chicago, 1985*, ed. by J. Baldwin, Springer-Verlag (1987), 107–126.

[7] B. Hart and M.C. Laskowski, A survey of the uncountable spectra of countable theories, In: *Algebraic Model Theory*, B. Hart et al. (eds.), Kluwer Academic Publishers (1997), 107-118.

[8] B. Hart, E. Hrushovski, and M.C. Laskowski, Unique decomposition in classifiable theories, submitted to the *Journal of Symbolic Logic*.

[9] W. Hodges, *Model Theory*, Cambridge University Press, 1993.

[10] E. Hrushovski, Unidimensional theories are superstable, *Annals of Pure and Applied Logic* **50** (1990), 117-138.

[11] A.H. Lachlan, Theories with a finite number of models in an uncountable power are categorical, *Pacific Journal of Mathematics*, 61 (1975), 465–481.

[12] A.H. Lachlan, Spectra of ω-stable theories, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 24 (1978), 129–139.

[13] M. Makkai, A survey of basic stability theory with particular emphasis on orthogonality and regular types, *Israel Journal of Mathematics* **49** (1984), 181-238.

[14] M. Morley, Categoricity in power, *Transactions of the AMS* **114** (1965), 514–538.

[15] A. Pillay, *Geometric Stability Theory*, Oxford University Press, 1996.

[16] J. Saffe, The number of uncountable models of ω-stable theories, *Annals of Pure and Applied Logic*, **24** (1983), 231–261.

[17] S. Shelah, The number of non-isomorphic models of an unstable first-order theory, *Israel Journal of Mathematics*, **9** (1971), 473–487.

[18] S. Shelah, *Classification Theory*, (revised edition) North Holland, Amsterdam, 1990.

[19] S. Shelah and S. Buechler, On the existence of regular types, *Annals of Pure and Applied Logic* **45** (1989), 277–308.