The Stability spectrum for classes of atomic models

John T. Baldwin*
University of Illinois at Chicago

Saharon Shelah
Hebrew University of Jerusalem
Rutgers University

February 20, 2012

Abstract

We prove two results on the stability spectrum for $L_{\omega_1,\omega}$. Here $S_i^m(M)$ denotes an appropriate notion (at or mod) of Stone space of m-types over M. Theorem A. (unstable case) Suppose that for some positive integer m and for every $\alpha < \delta(T)$, there is an $M \in \boldsymbol{K}$ with $|S_i^m(M)| > |M|^{\beth_\alpha(|T|)}$. Then for every $\lambda \geq |T|$, there is an M with $|S_i^m(M)| > |M| = \lambda$. Theorem B. (strictly stable case) Suppose that for every $\alpha < \delta(T)$, there is $M_\alpha \in \boldsymbol{K}$ such that $\lambda_\alpha = |M_\alpha| \geq \beth_\alpha$ and $|S_i^m(M_\alpha)| > \lambda_\alpha$. Then for any μ with $\mu^{\aleph_0} > \mu$, $\boldsymbol{K}$ is not i-stable in μ. These results provide a new kind of sufficient condition for the unstable case and shed some light on the spectrum of strictly stable theories in this context. The methods avoid the use of compactness in the theory under study. In the Section 4, we expound the construction of tree indiscernibles for sentences of $L_{\omega_1,\omega}$. Further we provide some context for a number of variants on the Ehrenfeucht-Mostowski construction.

1 Context

For many purposes, e.g., the study of categoricity in power, the class of models of a sentence ϕ of $L_{\omega_1,\omega}$ can be profitably translated to the study of the class of models of a first order theory T that omit a collection Γ of first order types over the empty set. In particular, if ϕ is complete (i.e. a Scott sentence) Γ can be taken as the collection of all non-principal types and the study is of the atomic models of T. This translation dates from the 60's; it is described in detail in Chapter 6 of [Bal09]. The study of finite

*We give special thanks to the Mittag-Leffler Institute where this research was conducted. This is paper 959 in Shelah's bibliography. Baldwin was partially supported by NSF-0500841. Shelah thanks the Binational Science Foundation for partial support of this research.

diagrams (see below) is equivalent to studying sentences of $L_{\omega_1,\omega}$; the study of atomic models of a first order theory is equivalent to studying *complete* sentences of $L_{\omega_1,\omega}$.

The stability hierarchy provides a crucial tool for first order model theory. Shelah [She78a] and Keisler [Kei76] show the function $f_T(\lambda) = \sup\{|S(M)| : |M| = \lambda, M \models T\}$ has essentially only six possible behaviors (four under GCH). In [She70], Shelah establishes a similar result for *homogeneous* finite diagrams. The homogeneity assumption is tantamount to assuming amalgamation over all sets. This is a strong hypothesis that is avoided in Shelah's further investigation of categoricity in $L_{\omega_1,\omega}$ ([She83a, She83b]), which is expounded as Part IV of [Bal09]. Important examples, due to Marcus and Zilber, which do not satisfy the homogeneity hypothesis are also described in [Bal09]. As we explain below, this investigation begins by identifying the appropriate notion of type over a set (and thus of ω-stability). Shelah [She83a, Bal09] showed that ω-stability implies stability in all powers. And assuming $2^{\aleph_0} < 2^{\aleph_1}$, ω-stability was deduced from $\aleph_1$-categoricity. But further questions concerning the stability hierarchy for this notion of type for arbitrary sentences of $L_{\omega_1,\omega}$ had not been investigated. We do so now.

Remark 1.1 In fact our results hold for *arbitrary finite diagrams*, the class of models of first order theory that omit a given set of types over the empty set. Thus they generalize some of the results of [She70]. But our results are by no means as complete as in homogeneous case. We have chosen to formulate the results for atomic classes because of the connection with the strong results of [She83a, She83b]. But we are using only the property that atomic classes are defined by omitting all non-principal types over the empty set. We do not in this paper use the stronger properties of atomic classes exploited in the arguments of [She83a, She83b].

There are (at least) two *a priori* reasonable notions of Stone space for studying atomic models of a first order theory. (As noted, we could more generally replace 'atomic' by 'finite diagram'.) Recall that for a first order theory T (with a monster model $\mathbb{M}$) $A \subset \mathbb{M}$ is an *atomic set* if each finite sequence from A realizes a principal type over the empty set. An atomic set is an *atomic model* if it is also a model of the theory T.

Definition 1.2 *Let $\boldsymbol{K}$ be the class of atomic models of a complete first order theory.*

1. *Let A be an atomic set; $S_{\mathrm{at}}(A)$ is the collection of $p \in S(A)$ such that if $\boldsymbol{a} \in \mathbb{M}$ realizes p, $A\boldsymbol{a}$ is atomic.*
2. *Let A be an atomic set; $S_{\mathrm{mod}}(A)$ is the collection of $p \in S(A)$ such that p is realized in some $M \in \boldsymbol{K}$ with $A \subseteq M$.*

In [Bal09] we wrote S^* for the notion called S_{mod} here. The latter notation is more evocative. We will simultaneously develop the results for both notions of Stone space and indicate the changes required to deal with the two cases. We will write $S_i(M)$ where i can be either at or mod.

We sometimes write $|T|$ for $|\tau|$ where τ is the vocabulary of T. $\boldsymbol{K} = \boldsymbol{K}_T$ is the class of atomic models of T. We write $H = H(\mu)$ for the Hanf number for atomic models of all theories with $|T| = \mu$. By [She78a] H equals $\beth_{\delta(T)}$, where $\delta(T)$, the well-ordering number of the class of models of a theory T omitting a family of types, is defined in VII.5 of [She78a]. It is also shown there that if T is countable, H evaluates as $\beth_{\omega_1}$ while for uncountable T $H = \beth_{(2^{|T|})^+}$. Fix $\mu_\alpha = \beth_\alpha(|T|)$.

Remark 1.3 In [She70], Shelah's definition of stability makes a stronger requirement; it implies *by definition* the existence of homogeneous models in certain cardinals. We do not make that assumption here so we are considering a larger class of theories.

Definition 1.4

1. *$\boldsymbol{K}$ is i-stable in λ (for $i = \mathrm{at}$ or mod) if for every $m < \omega$, and $M \in \boldsymbol{K}$ with $|M| = \lambda$, $|S_i^m(M)| = \lambda$.*
2. *Stability classes. For either $i = \mathrm{at}$ or mod,*
   - (a) *$\boldsymbol{K}$ is i-stable if it is i-stable in some λ.*
   - (b) *$\boldsymbol{K}$ is i-superstable if it is i-stable in all $\lambda \geq H$.*
   - (c) *$\boldsymbol{K}$ is strictly i-stable if it is i-stable but not i-superstable.*

For any M, $S_{\mathrm{at}}(M)$ contains $S_{\mathrm{mod}}(M)$ so at-stability in λ implies mod-stability in λ. Thus for both notions ω-stability implies stability in all powers by results of [She83a, She83b], expounded in [Bal09].

We prove Theorem A in Section 2 and Theorem B in Section 3. The proof of Theorem B uses an application of omitting types in Ehrenfeucht-Mostowski models generated by trees of the form ${}^{<\omega}\lambda$. This is by no means new technology but we weren't able to locate an explicit statement of the result so we include a proof in Section 4.

We acknowledge helpful discussions with Tapani Hyttinen, Ali Enayat, Alexei Kolesnikov, and Lynn Scow.

2 Unstable $\boldsymbol{K}$

We first show that if there are cardinals λ_α in which $\boldsymbol{K}$ is 'sufficiently unstable', then $\boldsymbol{K}$ is not stable in any cardinal. This is a partial response to the following question.

Question 2.1 *Must an atomic class that is unstable in all λ have the order property?*

Definition 2.2 *A class of atomic models has the* order property *if in large enough models M in the class there is some first order ϕ:*

$$M \models (\phi(\mathbf{c}_t, \boldsymbol{a}_s) \equiv \phi(\mathbf{c}_t, \mathbf{b}_s)) \text{ iff } s <_I t.$$

and the set $\{\mathbf{c}_t : t \in I\}$ is an atomic set.

Theorem 2.3 falls short of answering this question in two ways: the hypothesis is stronger than mere instability; the sequence constructed in conditions 3) and 4) of Lemma 2.7 does not have *all* the sequences contained in a single atomic set but only requires each triple to be atomic.

The argument with splitting in Theorem 2.3 derives from Theorems I.2.6/I.2.7 of [She78b] showing that if in a stable first order theory, $\kappa(T)$ is finite in a stability cardinal it is finite in all cardinals.

Theorem 2.3 *Suppose that for some positive integer m and for every $\alpha < \delta(T)$, there is an $M_\alpha \in \boldsymbol{K}$ with $|S_i^m(M_\alpha)| > |M_\alpha|^{\beth_\alpha(|T|)}$. Then for every $\lambda \geq |T|$, there is an M with $|S_i^m(M)| > |M| = \lambda$.*

Remark 2.4 (Proof Sketch) Before the formal proof we outline the argument. We start with a sequence of models M_α and many distinct types over each of them. By an argument which is completely uniform in α, we construct triples $\langle \boldsymbol{a}_{\alpha,i}, \mathbf{b}_{\alpha,i}, \mathbf{d}_{\alpha,i}\rangle$ for $i < \mu_\alpha^+$ with the $\boldsymbol{a}_{\alpha,i}, \mathbf{b}_{\alpha,i} \in M_\alpha$ and $\mathbf{d}_{\alpha,i}$ in an elementary extension M'_α of M_α and so that $M_\alpha \mathbf{d}_{\alpha,i}$ is atomic and the distinctness of the types of the $\mathbf{d}_{\alpha,i}$ is explicitly realized by formulas. Then we apply Morley's omitting types theorem to the M'_α and extract from this sequence a countable sequence of order indiscernibles with desirable properties. Finally, this set of indiscernibles easily yields models of all cardinalities with the required properties.

Remark 2.5 The idea of the proof can be seen by ignoring the α and proving a slightly weaker result from one model of size $\beth_{\delta(T)}$.

Notation 2.6 *In this section, $\lambda_\alpha = |M_\alpha|^{\beth_{\alpha+2}(|T|)}$; $\mu_\alpha = \beth_\alpha(|T|)$; $\kappa_\alpha = \beth_{\alpha+2}(|T|)$.*

Lemma 2.7 *There is Φ, proper for linear orders, in a vocabulary τ_Φ extending τ with $|\tau_\Phi| = |\tau|$, with fixed additional unary predicates P, P_1 and binary R such that:*

1. *For every linear ordering I, $N_I = EM_\tau(I,\Phi) \models T$ and $M_I = EM_\tau(I,\Phi) \restriction P \in \boldsymbol{K}$. Naturally, $J \subset I$ implies $N_J \prec N_I$ and $M_J \prec M_I$.*
2. *The skeleton of N_I is $\langle \boldsymbol{a}_i \frown \mathbf{b}_i \frown \mathbf{c}_i : i \in I\rangle$ and $\lg(\mathbf{c}_i) = m$.*
3. *For some first order ϕ:*

$$N_I \models (\phi(\mathbf{c}_t, \boldsymbol{a}_s) \equiv \phi(\mathbf{c}_t, \mathbf{b}_s)) \text{ iff } s <_I t.$$

4. *$M_I \cup \mathbf{c}_i \subset N_I$ and is atomic.*
5. *For $S_{\mathrm{mod}}(M)$, we add the requirement that for each $s \in I$,*

$$M_{I,s} = N_I \restriction \{d : N_I \models R(d, \mathbf{c}_s)\}$$

is an atomic elementary submodel of N_I containing $M_I \mathbf{c}_s$.

Proof. The proof of Lemma 2.7 requires a number of steps. Fix for each $\alpha < \delta(T)$, $M_\alpha \in \boldsymbol{K}$ with $|M_\alpha| = \lambda_\alpha$ such that $|S_i^m(M_\alpha)| > \lambda_\alpha = |M_\alpha|^{\beth_{\alpha+2}(|T|)}$. Fix $p_{\alpha,i}$ for $i < \lambda_\alpha^+$, a list of distinct types in $S_i^m(M_\alpha)$. We work throughout in a monster model $\mathbb{M}$ of T.

Notation 2.8 *In the following construction, we choose by induction triples $\langle \boldsymbol{a}_{\alpha,i}, \mathbf{b}_{\alpha,i}, \mathbf{d}_{\alpha,i}\rangle$ for $i < \mu_\alpha^+$. We use the following notation for initial segments of the sequences.*

1. *$D_{\alpha,i} = \{\mathbf{d}_{\alpha,j} : j < i\}$.*

2. $X_{\alpha,i} = \{\boldsymbol{a}_{\alpha,j}, \mathbf{b}_{\alpha,j} : j < i\}$.

3. $q_{\alpha,i}$ *is the type of* $\mathbf{d}_{\alpha,i}$ *over* $X_{\alpha,i}$.

The following variant on splitting is crucial to carry out the construction. We call it ex-splitting (for external) because the elements which exemplify splitting are required to satisfy the same type over a set D which is not in (so external to) the model M and, in particular, is not required to be realized in an atomic set.

Definition 2.9 *Let M be a model, $X \subset M$ and $D \subset \mathbb{M}$. We say that $p \in S_i^m(M)$* ex-splits *over (D, X) if there exist $\boldsymbol{a}, \mathbf{b} \in M, \mathbf{f} \in \mathbb{M}$ so that $\mathbf{f}$ realizes $p \restriction X$, $\boldsymbol{a} \equiv_D \mathbf{b}$ but $(\boldsymbol{a}, \mathbf{f})$ and $(\mathbf{b}, \mathbf{f})$ realize different types over $\emptyset$.*

We will apply the next claim to M_α, $X_{\alpha,i}$, and $D_{\alpha,i}$ when carrying out the construction in paragraph 2.12. Note that this computation does not depend on $|M|$.

Claim 2.10 *For any model M, the number of types in $S_i^m(M)$ that do not ex-split over a pair (D, X) with $|X| = |D| \leq \mu_\alpha$ is at most $\mu_{\alpha+2}$.*

Proof. Let P denote the collection of $\mathrm{tp}(\mathbf{e}/M)$ with $\lg(\mathbf{e}) = m$ that do not ex-split over a pair (D, X). Each type r in P is determined by knowing $r \restriction X$ and for each formula $\phi_i(x_1, \ldots x_{k_i})$ for $i < |T|$ the restriction of r to one k_i-tuple from each equivalence class of the equivalence relation E_k on M defined by $\boldsymbol{a} E_{k_i} \mathbf{b}$ if $\boldsymbol{a}$ and $\mathbf{b}$ realize the same k_i-type over D. So, since $|D| = \mu_\alpha$, there are at most

$$2^{\mu_\alpha} \times (2^{2^{|D|}})^{|T|} = (2^{2^{\mu_\alpha}})^{|T|} = \mu_{\alpha+2}$$

possible such r. $\square_{2.10}$

As noted, for each M_α we will be constructing by induction on $i < \mu_\alpha^+$, sets $X_{\alpha,i}, D_{\alpha,i}$ of cardinality μ_α. We need to choose in advance a type p_α which does not ex-split over any $(X_{\alpha,i}, D_{\alpha,i})$ that arises. In order to do that we restrict the source of $D_{\alpha,i}$; clearly $X_{\alpha,i} \subset M_\alpha$. That is, we will fix M'_α with $M_\alpha \prec M'_\alpha$, $|M'_\alpha| = \lambda_\alpha$ and M'_α is μ_α^+-saturated and choose $D_{\alpha,i} \subset M'_\alpha$. (Note then that M'_α is *not* in general atomic.)

The number of types in $S_i^m(M_\alpha)$ that do not ex-split over *any* pair (D, X) with $|X| = |D| = \beth_\alpha$ is bounded by the number of such sets, $|M'_\alpha|^{\mu_\alpha}$, times the number of types in $S_i^m(M_\alpha)$ that do not ex-split over a particular choice of (D, X), which is $\mu_{\alpha+2}$ by Claim 2.10. That is, the bound is $|M'_\alpha|^{\mu_\alpha} \times \mu_{\alpha+2}$. Since this number is less than λ_α^+, we can fix a type $p_\alpha \in S_i^M(M_\alpha)$ which does not ex-split over any of the relevant (D, X).

Definition 2.11 *For each $\alpha < \delta(T)$, fix M'_α with $M_\alpha \prec M'_\alpha$, $|M'_\alpha| = \lambda_\alpha$, and M'_α is μ_α^+ saturated. Choose, by induction on $i < \mu_\alpha^+$, triples $\mathbf{e}_{\alpha,i} = \langle \boldsymbol{a}_{\alpha,i}, \mathbf{b}_{\alpha,i}, \mathbf{d}_{\alpha,i} \rangle$ where*

a) $\mathbf{d}_{\alpha,i} \in M'_\alpha$.

b) $\boldsymbol{a}_{\alpha,i}, \mathbf{b}_{\alpha,i}$ *are sequences of the same length from M_α that realize the same type over $D_{\alpha,i} = \{\mathbf{d}_{\alpha,j} : j < i\}$.*

c) The types over the empty set of $(\boldsymbol{a}_{\alpha,i}, \mathbf{d}_{\alpha,i})$ *and* $(\mathbf{b}_{\alpha,i}, \mathbf{d}_{\alpha,i})$ *differ.*

d) $q_{\alpha,i} = p_\alpha \restriction X_{\alpha,i} = \mathrm{tp}(\mathbf{d}_{\alpha,i}/X_{\alpha,i})$ *so if* $j < i$, $q_{\alpha,j} \subseteq q_{\alpha,i}$.

e) $M_\alpha \mathbf{d}_{\alpha_i}$ *is an atomic set for each* i. *(In the* mod*-version* $N_{\alpha,i}$ *is an atomic model containing* $M_\alpha \mathbf{d}_{\alpha_i}$*.)*

Construction 2.12 Choose $\mathbf{d}_{\alpha,i}$ to realize $p_\alpha \restriction X_{\alpha,i}$. By Claim 2.10 and since $|S_i(M_\alpha)| > \lambda_\alpha$ we can choose $\boldsymbol{a}_{\alpha,i}$ and $\mathbf{b}_{\alpha,i}$ to satisfy conditions b) and c). So we have

$$\mathrm{tp}(\mathbf{d}_{\alpha,i}, \boldsymbol{a}_{\alpha,j}) = \mathrm{tp}(\mathbf{d}_{\alpha,i}, \mathbf{b}_{\alpha,j}) \text{ if and only if } i < j. \tag{1}$$

We want this order condition for a single formula. For each $i < \mu_\alpha^+$, the types of $(\boldsymbol{a}_{\alpha,i}, \mathbf{d}_{\alpha,i})$ and $(\mathbf{b}_{\alpha,i}, \mathbf{d}_{\alpha,i})$ differ. That is, $\phi_{\alpha,i}(\boldsymbol{a}_{\alpha,i}, \mathbf{d}_{\alpha,i})$ and $\neg\phi_{\alpha,i}(\mathbf{b}_{\alpha,i}, \mathbf{d}_{\alpha,i})$ for some $\phi_{\alpha,i}$. By the pigeon-hole principal we may assume the $\phi_{\alpha,i}$ is always the same ϕ_α. (Further, since $|T|$ is not cofinal in $\delta(T)$, we can assume the ϕ_α is the same ϕ for all α.)

Now the construction is completed. We expand τ to a language $\tau_\Phi \supset \tau$ by adding predicates $P, <, R$ and Skolem functions. We add Skolem axioms to T to get a theory T_1 that admits quantifier elimination, requiring that these Skolem functions applied to elements of P (P_n) give an element of P (P_n) so that P (P_n) will pick out an elementary submodel. (We make a similar requirement for $R(x, \mathbf{y})$ in the mod-case.) Let M_α^+ be a model of T_1 (submodel of M'_α) with cardinality μ_α^+ containing M_α and all the $\mathbf{d}_{\alpha,i}$ ($N_{\alpha,i}$ in the mod-case). Interpret P as the model M_α, P_n as M_n^α, and the relation $<$ as the ordering on the triples $\langle \mathbf{e}_{\alpha,i} : i < \mu_\alpha \rangle$ imposed by ϕ_α.

Assign the Skolem functions so that the $\mathbf{e}_{\alpha,i}$ generate M'_α and interpret R by

$$R = \{e\hat{\ }\mathbf{d}_{\alpha,i}; e \in M_\alpha, i < \mu_\alpha^+\}.$$

(In the $S_{\mathrm{mod}}(M)$ case, interpret R as $\{e\hat{\ }\mathbf{d}_{\alpha,i} : i < \mu_\alpha^+, e \in N_{\alpha,i}\}$.)

Notation 2.13 *Let* Γ *be the collection of types* $\mathcal{P}_n \cup \mathcal{Q}_n$. *Each non-principal* n*-type* q *over the empty set determines one element of* $\mathcal{P}_n$ *and each non-principal* $n+m$*-type* q *determines one element of* $\mathcal{Q}_n$*:*

1. $\mathcal{P}_n = \{\bigwedge_{i<n} P(x_i)\} \cup \{q(\mathbf{x}) : q \text{ is a non-principal } n\text{-type }\}$
2. $\mathcal{Q}_n = \{\bigwedge_{i<n} R(x_i, y)\} \cup \{q(\mathbf{x}, \mathbf{y}) : q \text{ is a non-principal } n+m\text{-type}, m < \omega \}$

Now apply Morley's omitting types theorem[1] to the τ_Φ-theory T_1 and the collection of M_α^+ to get a countable sequence I of order indiscernibles and an extension Φ of T_1, (the EM-template) such that Φ is realized in each M_α^+ and such that for every linear order J, $EM_\tau(J, \Phi) \models T_1$ and omits Γ.

[1] See Appendix A.3.1 of [Bal09] for a precisely tailored version. See [She78a] or [Hod93], page 587 for a version with the role of the ordering more explicit. The latter two sources make the connection with the well-ordering number clear.

Remark 2.14 (Morley's Method) The next observation requires a little care in proving Morley's theorem rather than just quoting it. The M'_α are generated by the $\mathbf{e}_{\alpha,i}$ and we have interpreted $<$ so that these are exactly the domain of $<$. So in proving the omitting types theorem, all witnesses for the consistency of the template $\Phi(\mathbf{c})$ can be chosen from the domain of $<$. We use this fact below. It is this extra care that in the mind of the first author distinguishes "Morley's Method" from Morley's theorem. But this may be an idiosyncratic interpretation. The earliest mention of the phrase I have found is in [She74] and that refers to a standard application of the two cardinal theorem for cardinals far apart.

Note that any τ_Φ formula $\phi(\mathbf{x})$ is in Φ if it is true of every tuple $\langle \mathbf{e}_{\alpha,i_1}, \ldots \mathbf{e}_{\alpha,i_1} \rangle$ with $i_1 < i_2 < \ldots i_n$. We describe a crucial such sentence.

Let $\mathbf{x}^1\mathbf{x}^2\mathbf{x}^3$ be a triple of sequences with the first two having the same length as $\lg(\boldsymbol{a}) = \lg(\mathbf{b})$ and the third has length m. Let $\psi(\mathbf{x}, \mathbf{y})$ denote:

$$\phi(\mathbf{y}^3, \mathbf{x}^1) \equiv \phi(\mathbf{y}^3, \mathbf{x}^2).$$

Let ψ_1 be the assertion that ϕ defines a linear order on its domain; this directly translates precisely Lemma 2.7.3 and is true by the displayed statement 1. These structures clearly satisfy all the conditions of the requirements in Lemma 2.7 and we complete the proof.

$\square_{2.7}$

Proof of Theorem 2.3: To show instability in λ, let I be a dense linear ordering with cardinality λ which has more than λ cuts and choose $J \supset I$, that realizes more than $|I|$ cuts over I. Then $EM_\tau(J, \Phi)$ realizes more than λ types in $S_i^m(P(EM_\tau(I, \Phi))$. To see this, consider for any cut in I realized by an element $j \in J$ the type:

$$\{\psi(\langle \boldsymbol{a}_i, \mathbf{b}_i, \mathbf{c}_i \rangle, \mathbf{x},) : i < j\} \cup \{\psi(\langle \boldsymbol{a}_i, \mathbf{b}_i, \mathbf{c}_i \rangle, \mathbf{x},) : i \geq j\}.$$

Then $\langle \boldsymbol{a}_j, \mathbf{b}_j, \mathbf{c}_j \rangle$ realizes the type in $EM_\tau(J, \Phi)$ and $P(EM_\tau(J, \Phi))\mathbf{c}_j$ is an atomic set since $\mathcal{Q}$ was omitted. For the mod-case, use the interpretation of R to define $N_{\alpha,i}$.
$\square_{2.3}$

3 Strictly stable case

In Theorem III.5.15 of [She78b], Shelah determines the form of the least stability cardinal for a stable first order theory. Patterned on these arguments we take some steps towards such a characterization for atomic classes.

As the following examples show, it is easy to have superstable (i.e. eventually stable) (incomplete) sentences of $L_{\omega_1,\omega}$ that are not stable in arbitrarily large cardinals below the Hanf number H. Theorem 3.3 has two easily stated corollaries: If $\boldsymbol{K}$ is not eventually stable then it is not stable in every λ with $\lambda^\omega > \lambda$. If $\boldsymbol{K}$ is eventually stable then it is stable in some $\lambda < H$.

The results here are related to those in [GS86] but the combinatorics here is considerably simpler than in [GS86] for two related reasons. First, we construct tree indiscernibles indexed by ${}^{<\omega}\lambda$ while they are concerned with ${}^{\leq\omega}\lambda$; the limit node is much more difficult to handle. Second, they are constructing many non-isomorphic models, we only construct many different types. To obtain these stronger results, they assume the existence of large cardinals while this paper is in ZFC.

Example 3.1 For $\alpha < \omega_1$, let ϕ_α be Morley's sentence [Mor65b] that has a model in $\beth_\alpha$ but no larger model. It is easy to see that the sentences are not stable in the cardinalities where they have models. Let ψ be the Scott sentence of an infinite set with only equality. Now let ψ_α assert that either a structure has a nontrivial relation and obeys ϕ_α or just ψ. Then ϕ_α is $\beth_\alpha$-unstable but stable (indeed categorical) in all cardinals beyond $\beth_{\omega_1}$.

If one adds even joint embedding such trivial examples are no longer apparent.

Question 3.2 *Is there a complete sentence of $L_{\omega_1,\omega}$ which is stable beyond H (for either $\mod$ or at) but fails stability for some cardinals less than H?*

We retain the value of $\mu_\alpha = \beth_\alpha(|T|)$ from the first section but λ_α is redefined in the hypothesis of the next theorem to be a sufficiently large cardinal in which stability fails.

Theorem 3.3 *Suppose that for every $\alpha < \delta(T)$, there is $M^\alpha \in \boldsymbol{K}$ such that $\lambda_\alpha = |M^\alpha| \geq \mu_\alpha$ and $S^m_i(M^\alpha) > \lambda_\alpha$. Then for any μ with $\mu^{\aleph_0} > \mu$, $\boldsymbol{K}$ is not stable in μ.*

Proof. Fix for each $\alpha < \delta(T)$, $M^\alpha \in \boldsymbol{K}$ such that $|S^m_i(M^\alpha)| > \lambda_\alpha$. Fix $p_{\alpha,i}$ for $i < \lambda^+_\alpha$, a list of distinct types in $S^m_i(M^\alpha)$. We work throughout in a monster model $\mathbb{M}$ of T.

To prepare for the application of an appropriate version of Morley's omitting types theorem we construct a sequence of models and certain types. For this, we construct trees of types that arise from failure of stability. The combinatorics slightly extends the classical arguments and avoids compactness. Note that this stage of the construction takes place in the original language. We will apply the following general result uniformly to each M^α.

Fact 3.4 *Suppose $|M| \geq \mu_{\alpha+1}$ and $\mathcal{P}$ is a collection of $> \lambda_\alpha = |M|$ members of $S^m_i(M)$. Then there exists a sequence $\langle \mathbf{b}_j : j < \mu_\alpha\rangle$ with each $\mathbf{b}_j \in M$ and a formula $\phi(\mathbf{x}, \mathbf{y}) = \phi_{\mathcal{P}}$ such that for each $j < \mu_\alpha$,*

$$|\{p \in \mathcal{P} : i < j \rightarrow \phi(\mathbf{x}, \mathbf{b}_i) \in p \text{ and } \neg\phi(\mathbf{x}, \mathbf{b}_j) \in p\}| > \lambda_\alpha. \qquad (2)$$

Proof. We consider many possibilities for ϕ and prove one works. We choose $\{\phi_\eta : \eta \in T_i\}$ by induction on $i < \mu_\alpha$ where each T_i is a subset of ${}^i2$ and each $\mathbf{b}_\eta \in M^\alpha$ so that

1. $j < i$ and $\eta \in T_i$ implies $\eta \restriction j \in T_j$.
2. if $\eta \in T_i$ then $p_\eta = \{\phi_{\eta\restriction j}(\mathbf{x}, \mathbf{b}_{\eta\restriction j})^{\eta(j)} : j < i\}$ is included in $> \lambda_\alpha$ members of $\mathcal{P}$.
3. For limit i,

   $T_i = \{\eta \in {}^i2 : (\forall j < i)\eta \restriction j \in T_j \text{ and } p_\eta \text{ is included in } > \lambda_\alpha \text{ members of } \mathcal{P}\}$

4. if $i = j + 1$ then $T_i = \{\eta\hat{\ }0, \eta\hat{\ }1 : \eta \in T_j\}$.

For the successor step in the induction recall the following crucial observation of Morley. Suppose there are more than $|M|$ types over M extending a partial type p. Then there exists a formula $\phi(\mathbf{x}, \boldsymbol{a})$ with $\boldsymbol{a} \in M$ such that both $p \cup \{\phi(\mathbf{x}, \boldsymbol{a})\}$ and $p \cup \{\neg\phi(\mathbf{x}, \boldsymbol{a})\}$ have more than $|M|$ extensions to complete types over M. (We are extending Morley's analysis to types in $S_i^m(M)$ but the argument is just counting; there is a unique type which has more than λ_α extensions.)

The interesting point in the induction is the limit stage. We cannot guarantee that individual paths survive. But at each stage in the induction, we have defined types over a set of cardinality μ_α. So there are at most $\mu_{\alpha+1}$ types over $\{\mathbf{b}_\eta : \lg(\eta) < \delta\}$. So one of the paths must have more than λ_α extensions to $S_i^m(M)$.

So $T_{\mu_\alpha} \neq \emptyset$. Choose $\eta \in T_{\mu_\alpha}$. Let $\phi_j(\mathbf{x}, \mathbf{b}_j) = \phi_{\eta\restriction j}(\mathbf{x}, \mathbf{b}_{\eta\restriction j})^{\eta(j)}$ for $j < \mu_\alpha$. Since the path has length $\mu_\alpha = \beth_\alpha(T)$, by the pigeonhole principle we may assume there is a single formula ϕ. (The pigeonhole argument fails for $\alpha = 0$ but we need the result only for large α so we may assume $\alpha \neq 0$.) This completes the construction of the ϕ and the $\mathbf{b}_j$. We have the result by condition 4. $\square_{3.4}$

Now we apply this fact to construct from the original M^α given in the hypothesis of Theorem 3.3 a sequence of models $\hat{M}^\alpha$ and associated sequences $\mathbf{b}_{\alpha,\rho}$ and $\mathbf{c}_{\alpha,\rho}$ for $\rho \in {}^{<\omega}\mu_\alpha$.

Definition 3.5 *Let $\hat{M}^\alpha$ be a μ_α^+ saturated elementary extension of M^α. We construct for each α by induction on $n < \omega$, submodels M_n^α of M^α and types $\{q_\nu^\alpha : \nu \in {}^{<\omega}\mu_\alpha\}$ with $q_\nu^\alpha \in S_i^m(M_{\lg(\nu)}^\alpha)$ and realizations $\mathbf{c}_{\alpha,\nu} \in \hat{M}^\alpha$ of q_ν^α satisfying the following conditions.*

1. *$\langle M_n^\alpha : n < \omega\rangle$ is an increasing chain of submodels of M^α, each with cardinality μ_α.*
2. *If $k \leq n$ and $\nu \in {}^k\mu_\alpha$, then $q_\nu^\alpha \in S_i^m(M_k^\alpha)$.*
3. *Each $q_\nu^\alpha \in S_i^m(M_n^\alpha)$ has $> \lambda_\alpha$ extensions to $S_i^m(M^\alpha)$*
4. *Suppose $k < r \leq n$, $\nu \in {}^k\mu_\alpha$, $\rho \in {}^r\mu_\alpha$ and ρ extends ν:*
$$q_\nu^\alpha \subseteq q_\rho^\alpha.$$
5. *If $\nu \in {}^k\lambda_\alpha$, $k < n$, $i \neq j$, then*
$$q_{\nu\hat{\ }i}^\alpha \neq q_{\nu\hat{\ }j}^\alpha.$$
They are distinguished by the $\mathbf{b}_{\alpha,\rho}$, as specified in statement 3 below.
6. *$\mathbf{c}_{\alpha,\nu} \in \hat{M}_\alpha$ realizes q_ν^α. (In the* mod*-case, $N_{\alpha,\rho}$ is the universe of an atomic model containing $M\mathbf{c}_{\alpha,\rho}$.)*

Construction 3.6 We use Fact 3.4 to construct objects meeting this definition. Let the subscript x denote at or mod. By induction, for each $\rho \in {}^n\mu_\alpha$ the type $q^\alpha_\rho \in S^m_x(M^\alpha_n)$ has $> \lambda_\alpha$ extensions to $S^m_x(M^\alpha)$. Let $\mathcal{P}_\rho = \{r \in S_x(M_\alpha) : q^\alpha_\rho \subseteq r\}$ so $|\mathcal{P}_\rho| > \lambda_\alpha$. By Fact 3.4, we find $\langle \mathbf{b}_{\alpha,\rho\hat{}j} : j < \mu_\alpha\rangle$ and ϕ_ρ satisfying displayed statement 2.

Let M^α_{n+1} be a submodel of M^α with $M^\alpha_n \cup \{\mathbf{b}_{\alpha,\rho} : \rho \in {}^{n+1}(\mu_\alpha)\} \subseteq M^\alpha_{n+1}$ and with cardinality μ_α. $M^\alpha_{n+1} \subset M_\alpha$ so is an atomic model and each q^α_ρ extends to an atomic type over M^α.

For $\rho \in {}^n(\mu_\alpha)$ and $i < \mu_\alpha$ first define

$$p'_{\rho\hat{}i} = q^\alpha_\rho \cup \{\phi_\rho(x, \mathbf{b}_{\alpha,\rho\hat{}j}) : j < i\} \cup \{\neg\phi_\rho(x, \mathbf{b}_{\alpha,\rho\hat{}i})\}.$$

Since $\lambda_\alpha < |\{r \in S_x(M^\alpha) : p'_\rho \subseteq r\}|$, we can find $p^\alpha_{\rho\hat{}i} \in S^m_x(M^\alpha_n)$ extending $p'_{\rho\hat{}i}$ such that $\mathcal{P}_{\rho\hat{}i} = \{r \in S_x(M^\alpha) : p_{\rho\hat{}i} \subseteq r\}$ has cardinality $> \lambda_\alpha$. Note that

$$p^\alpha_{\rho\hat{}i} \supseteq q^\alpha_\rho \cup \{\phi_\rho(x, \mathbf{b}_{\alpha,\rho\hat{}j}) : j < i\} \cup \{\neg\phi_\rho(x, \mathbf{b}_{\alpha,\rho\hat{}i})\}.$$

This completes the $n+1$st stage of the construction. So we can construct the M^α_n and $\{q_{\nu,i} : \nu \in {}^{<\omega}\mu_\alpha\}\rangle$, $\hat{M}^\alpha$ and by μ^+_α-saturation choose $\mathbf{c}_{\alpha,\rho} \in \hat{M}^\alpha$. In the mod-case choose an atomic model $N_{\alpha,\rho}$ with $M^\alpha\mathbf{c}_{\alpha,\rho} \subset N_{\alpha,\rho} \prec \hat{M}^\alpha$. Note

$$\{\phi_\rho(\mathbf{c}_{\alpha,\rho\hat{}i}, \mathbf{b}_{\alpha,\rho\hat{}j}) : j < i\} \cup \{\neg\phi_\rho(\mathbf{c}_{\alpha,\rho\hat{}i}, \mathbf{b}_{\alpha,\rho\hat{}i})\}. \tag{3}$$

With the construction complete, we expand τ to a language $\tau_\Phi \supset \tau$ in two stages. Form τ' by adding predicates $P, P_n, <, <^*, R$ and Skolem functions. We add Skolem axioms to T to get a theory T' that admits quantifier elimination, requiring that these Skolem functions applied to elements of P give an element of P so that P will pick out an elementary submodel.

Let M^+_α be a model of T (submodel of $\hat{M}^\alpha$) with cardinality μ_α containing M^α_n for $n < \omega$ and all the $\mathbf{c}_{\alpha,\rho}$ Assign the τ'-Skolem functions so that $P(\hat{M}^\alpha) = M^\alpha = \bigcup_{n<\omega} M^\alpha_n$ is generated by the $\mathbf{b}_{\alpha,\rho}$ for $\rho \in {}^{<\omega}\mu_\alpha$. Let X_α be the tree with domain $\langle \mathbf{b}_{\alpha,\rho} : \rho \in {}^{<\omega}(\mu_\alpha)\rangle$ and the following relations. Interpret $<$ as the partial order on the $\langle \mathbf{b}_{\alpha,\rho} : \rho \in {}^{<\omega}(\mu_\alpha)\rangle$ given by inclusion on the ρ-indices. Let $<^*$ be a linear order of the $\langle \mathbf{b}_{\alpha,\rho} : \rho \in {}^{<\omega}(\mu_\alpha)\rangle$ given by lexiocographic order on the ρ-indices. Interpret R as

$$\{e\hat{}\mathbf{c}_{\alpha,\rho} : \rho \in {}^{<n}(\mu_\alpha), e \in \bigcup_{n<\omega} M^\alpha_n\}.$$

Form τ_Φ by adding function symbols F_n. Define $F_n(\mathbf{b}_{\alpha,\rho}) = \mathbf{c}_{\alpha,\rho}$. Now let T_1 be the collection of all $L(\tau_\Phi)$-sentences that are true in each $\hat{M}_\alpha$.

In the $S_{\mathrm{mod}}(M)$ case, we must do a bit more. Interpret R as

$$\{\mathbf{e}\hat{}\mathbf{c}_{\alpha,\rho} : \rho \in {}^{<n}(\mu_\alpha), \text{ and } \mathbf{e} \in N_{\alpha,\rho}\}.$$

Define the τ'-Skolem functions so that the Skolem closure of $M\mathbf{c}_{\alpha,\rho}$ is $N_{\alpha,\rho}$. This implies that if $R(\mathbf{e}, \mathbf{c}_{\alpha,\rho})$ holds then $\mathbf{e}$ is a sequence given by τ'-Skolem functions with arguments a finite number of members of $P(M^+_\alpha)$ and $\mathbf{c}_{\alpha,\rho}$.

By conditions 4-6 of Definition 3.5,

Claim 3.7 *For any finite linearly ordered initial $<$-segment of the tree with length $n+1$, enumerated by $\mathbf{x}_0, \ldots \mathbf{x}_n$, (so $P_i(\mathbf{x}_i)$):*

1. $\bigwedge_{i \leq n} [P_i(\mathbf{z}) \wedge \mathbf{z} <^* \mathbf{x}_i \rightarrow \phi_i(F_n(x_n), \mathbf{z})]$
2. $\bigwedge_{i \leq n} \neg\phi_i(F_n(x_n), \mathbf{x}_i)$.

The universal quantification of each such sentence is true in each $\hat{M}_\alpha$ and so is in T_1.

As in Notation 2.13 let Γ be the collection of types:

1. $\mathcal{P}_n = \{\bigwedge_{i<n} P(x_i)\} \cup \{q(\mathbf{x}) : q$ is a non-principal n-type $\}$
2. $\mathcal{Q}_n = \{\bigwedge_{i<n} R(x_i, \mathbf{y})\} \cup \{q(\mathbf{x}, \mathbf{y}) : q$ is a non-principal $n+m$-type, $m < \omega$ $\}$

Now apply the omitting types theorem (as stated in Section 4) to the τ_Φ-theory T_1 and the collection of M^+_α to get a countable set of tree-indiscernibles in order type ${}^{<\omega}\omega$ and an extension Φ of T_1, (the EM-template) such that for every tree of J of order ${}^{<\omega}\lambda$, $EM_\tau(J, \Phi) \models T_1$ and omits Γ.

Finally we must show there are many types; we separate the mod and at cases.

Claim 3.8 *If $\lambda^\omega > \lambda$ then there is an I with $|I| = \lambda$ such that $S^m_{\mathrm{at}}(M_I) > \lambda$, where $M_I = EM(I, \Phi) \restriction P$.*

Proof. Note that by displayed statement 3 and Claim 3.7 we have:

1. If $\rho, \rho\widehat{\ }i \in I$, $\mathrm{tp}(F_n(\rho)/P_n(M)) \subseteq \mathrm{tp}(F_{n+1}(\rho\widehat{\ }i)/P_{n+1}(M))$.
2. If $\rho \in I$ and $i \neq j$,

$$\mathrm{tp}(F_{n+1}(\rho\widehat{\ }j)/P_{n+1}(M)) \neq \mathrm{tp}(F_{n+1}(\rho\widehat{\ }i)/P_{n+1}(M)).$$

Now in any $M_I = EM(I, \Phi)$ for any $\rho \in J$ define $p_\rho \in S^m_{\mathrm{at}}(P_N(M_I) = \mathrm{tp}(F_n(\rho), P_n(M))$. Now letting $p_\eta \in S^m_{\mathrm{at}}(P(M_I))$ be $\bigcup_{i<\omega} p_{\eta\restriction n}$, we find λ^ω members of $S^m_{\mathrm{at}}(P(M_I))$. The definition of S^m_{at} guarantees the union is in S^m_{at}. $\square_{3.8}$

Now we extend this result to mod.

Claim 3.9 *If $\lambda^\omega > \lambda$ then there is an I with $|I| = \lambda$ such that $S^m_{\mathrm{mod}}(M_I) > \lambda$, where $M_I = EM(I, \Phi) \restriction P$.*

Proof. We need to construct an atomic model N_η containing $M_I\mathbf{c}_\eta$ (from the proof of Claim 3.8). The natural choice is the τ'-Skolem closure of $M_I\mathbf{c}_\eta$. The reason the reduct of this structure to τ is atomic is that any finite sequence is of the form $\boldsymbol{a}, \mathbf{b}$ where the $\boldsymbol{a}$ come from $P_n(M_I)$ (for a fixed n) and each of the $\mathbf{b}$ has the form $G(\boldsymbol{a}, \mathbf{c}_\eta)$ where G is a τ'-Skolem function. But then the τ type of $\boldsymbol{a}\mathbf{b}$ is the same as the τ-type of a sequence $\boldsymbol{a}'\mathbf{b}'$ where $\boldsymbol{a}' \in P_n(\hat{M}^\alpha)$ and each $b' \in N_{\alpha,\rho}$ is of the form $G(\boldsymbol{a}', \mathbf{c}_{\eta\restriction n})$. $\square_{3.9}$

Remark 3.10 We investigate the difference in hypotheses between Theorem 2.3 and Theorem 3.3[2]. We first study Theorem 2.3.

Let $\kappa = |M_\alpha|$.

Case 1. $\kappa \leq \beth_\alpha$: then $\kappa^{\beth_\alpha}$ is equal to $\beth_\alpha^{\beth_\alpha} = 2^{\beth_\alpha} = \beth_{\alpha+1}$. The assumption of the theorem is that $|S_i^m(M_\alpha)| > \kappa^{\beth_\alpha} = 2^{\beth_\alpha}$. This case is not possible since $|S_i^m(M_\alpha)| \leq 2^\kappa \leq 2^{\beth_\alpha}$.

Case 2. $\kappa > \beth_\alpha$. On one hand we have $\kappa \leq \kappa^{\beth_\alpha}$; on the other $\beth_{\alpha+1} = \beth_\alpha^{\beth_\alpha} \leq \kappa^{\beth_\alpha}$. Thus, $\kappa^{\beth_\alpha} \geq \max(\beth_{\alpha+1}, \kappa)$. The hypothesis in the theorem says that $|S_i^m(M_\alpha)| > \kappa^{\beth_\alpha}$, so $|S_i^m(M_\alpha)| > \max(\beth_{\alpha+1}, \kappa)$.

This leads to two cases:

Case 2a. $\kappa \geq \beth_{\alpha+1}$: then $|S_i^m(M_\alpha)| > \max(\beth_{\alpha+1}, \kappa^{\beth_\alpha}) \geq \kappa$. So the requirement is at least instability in κ.

Case 2b. $\beth_\alpha < \kappa < \beth_{\alpha+1}$: then $|S_i^m(M_\alpha)| > \max(\beth_{\alpha+1}, \kappa^{\beth_\alpha}) = \beth_{\alpha+1} > \kappa$. This yields instability in κ. (Under GCH, of course, this case is empty.)

In general, the hypothesis in case 2b) requires more than instability in κ: if κ has cofinality less than or equal to the cofinality of $\beth_\alpha$, then $\kappa^{\beth_\alpha} > \kappa$, and the number of types needs to be (possibly) much greater than κ.

Theorem 3.3 asserts that $\boldsymbol{K}$ is unstable in some cardinal then it is unstable in any λ with $\lambda^\omega > \lambda$ so it is analogous to the first order case. Further it asserts that the first stability cardinal for a superstable class is less than H.

Thus, in Theorem 2.3 we assume 'serious' instability and get instability everywhere and in Theorem 3.3 we assume "just" instability, and get instability for cardinals of countable cofinality only.

We further analyze case 2a under GCH. The possible values of $\kappa^{\beth_\alpha}$, given that $\kappa > \beth_{\alpha+1}$, become κ and $2^\kappa = \kappa^+$ (the first is the case when the cofinality of κ is greater than the cofinality of $\beth_\alpha$; otherwise, the second alternative holds).

Under the GCH the difference between 'serious' and 'just' instability disappears. Moreover, we can expect to find M_α satisfying the hypothesis only for $|M_\alpha|$ of cofinality greater than the cofinality of $\beth_\alpha$. So under the GCH, the difference between the hypotheses in 2.1 and 3.2 disappears, but the conclusion of 3.2 is weaker.

4 Tree Indiscernibility

The main result of this section is the existence of tree indiscernibles as needed in the previous section. But we take the occasion to discuss the role of various types of index sets for indiscernible collections and to make explicit the role of expanding the vocabulary when finding indiscernibles in various contexts.

The theorem reported here is implicit in the literature (e.g. [She78a, GS86]) but we could not find an explicit statement. Theorem VII.3.6 of [She78a] finds an indiscernible tree in the first order case on ${}^{\leq\omega}\omega$ but we want to omit types as well. The basic plan of the proof dates to Morley [Mor65b]. We indicate the modifications needed for the more complicated combinatorics to build models to omit types that are over indiscernible trees instead of over linear orders.

[2]This analysis was worked out by the first author and Alexei Kolesnikov.

Many variants of tree indiscernibles are used in various parts of model theory; we describe several of these applications to provide a context for the current version and to emphasize some important distinctions. Indiscernibles may be ordered by linear orders, or trees of the form ${}^{<\omega}2$, ${}^{<\omega}\lambda$ or even ${}^{\leq\omega}2$, ${}^{\leq\omega}\lambda$. We may want to find the ordering in the basic vocabulary τ (to witness unstability at some level) or not (to avoid introducing instability) but only in an expanded vocabulary τ^*.

We first consider situations where the (partial) ordering is explicitly added to the language. Ehrenfeucht and Mostowski (finding automorphisms) index the indiscernible by a linear order which is *not* in the base language. Morley's proof that $\aleph_1$-categoricity implies ω-stability pursues the same strategy. He is counting the number of τ types and there is certainly no ordering in the vocabulary τ. There are further applications using the extended vocabulary to two-cardinal models [She75, She76] and to Peano arithmetic [MP84]. Tree indiscernibles on ${}^{<\omega}2$ rely on Halpern-Lauchli; tree indiscernibles on ${}^{<\omega}\omega$ rely on Erdos-Rado.

To show an unstable first order theory has the maximal number of models, it is essential that Shelah is building order indiscernibles with respect to an ordering that is definable in the base vocabulary τ. To investigate the difference in stability spectrum for stable but not superstable theories, the (tree)-order must be expressible in τ and the tree is $\lambda^{<\omega}$. But to count the number of models of superstable theory involves trees of height $\omega + 1$. The proof in this paper differs from [She78a], where the number of models of an unsuperstable theory is computed, because in working with $L_{\omega_1,\omega}$, we must omit types. In VII.3.6 of [She78a], Erdos-Rado is applied to show the existence of a 'uniform' β-tree implies the existence of a tree of indiscernibles indexed by ${}^{<\omega}\omega$. Thus indiscernibles indexed by linear orders as well as the trees $\lambda^{<}$ and $\lambda^{\leq\omega}$ have a long history. More recently, the tree orders occur in [DǑ4, KKS, LS03, Sco]. The use of trees indexed by $2^{\leq\omega}$ to construct many models in $\aleph_1$ if a countable theory is not ω-stable appears in [She78a]. (The tree is found in VI.3.7; it is used to construct many models in VIII.1.2.) An exposition of this result and some extensions to uncountable languages occur in [Bal89]. The construction of many models from infinitary order properties in [GS86], with trees of the form $\lambda^{\leq\omega}$ requires large cardinal axioms for the combinatorics.

We see three steps in this kind of construction. The references in parentheses are to the application of this method to the proof of the strictly stable case in this paper.

1. Model theoretic construction of specific syntactic-combinatoric configurations on models. (Construction 3.6.)
2. Application of Erdos-Rado or Halpern-Lauchli and compactness to extract a countable family of indiscernibles. (Theorem 4.7.)
3. Application of Ehrenfeucht-Mostowski models to obtain models of arbitrary cardinality. (Claim 3.8.) This is sometimes called 'stretching'.

We first establish some background notation. The exact vocabulary for describing the partial order is significant; we follow [She78b] rather.

Notation 4.1 *1. A* tree $\boldsymbol{T}$ *is a subset of* ${}^{\leq\omega}\lambda$ *that is closed under initial segment.*

2. atp *means atomic (quantifier-free) type.*

3. *The vocabulary τ^* will denote the vocabulary for trees we use. It contains the partial order on the tree, $<$, the lexiocographic order on the tree $<^*$, $\wedge$ (meet) and the levels P_n. τ^*_n omits the P_i with $i > n$.*

4. *When elements $\boldsymbol{a}_\eta$ and $\boldsymbol{a}_\tau$ in a structure M are indexed by $\eta, \tau \in \boldsymbol{T}$ that realize the same quantifier free τ^*-type in the tree then $\boldsymbol{a}_\tau$ and $\boldsymbol{a}_\eta$ have the same length.*

5. *If ν is an n-element sequence from $\boldsymbol{T}$, $\boldsymbol{a}_\nu$ denotes $\langle \boldsymbol{a}_{\nu(0)}, \ldots \boldsymbol{a}_{\nu(n-1)} \rangle$.*

Definition 4.2 *For any vocabulary τ, let M be a τ-structure and Σ a set of τ-formulas.*
If $\mathrm{atp}_{\tau^*}(\eta/\emptyset) = \mathrm{atp}_{\tau^*}(\nu/\emptyset)$ implies $\mathrm{tp}_\Sigma(\boldsymbol{a}_\eta/\emptyset) = \mathrm{tp}_\Sigma(\boldsymbol{a}_\nu/\emptyset)$ in M then we call $\langle \boldsymbol{a}_\eta : \eta \in \boldsymbol{T} \rangle \subset M$ a set of Σ-tree indiscernibles*:*
We just say tree indiscernibles *if Σ contains all formulas in $L(\tau)$.*

We rely on a combinatorial lemma that follows from Erdos-Rado. The result is proved as Theorem 2.6 in the appendix to [She78a]. A stronger result (the bound on $k(m,n)$ is smaller) with a shorter proof is sketched in the appendix of [GS86]. Kim, Kim, and Scow have recently included a full argument in [KKS]; their paper includes a careful distinction of several notions of 'tree-indiscernibility'.

Lemma 4.3 ([She78a]) *For every $n, m < \omega$, there is a $k = k(n,m) < \omega$ such that if $\lambda = \beth_k(\chi)^+$ the following is true. For any function $f : [^{\leq n}\lambda]^m \rightarrow \chi$, there exists a $\boldsymbol{T} \subseteq {}^{\leq n}\lambda$ such that*

1. *Each $\eta \in \boldsymbol{T}$ has χ^+ immediate successors in $\boldsymbol{T}$.*

2. *If ν and τ are m-tuples from $\boldsymbol{T}$ with $\mathrm{atp}_{\tau^*}(\eta/\emptyset) = \mathrm{atp}_{\tau^*}(\nu/\emptyset)$, then*

$$f(\tau) = f(\eta).$$

We now prove the theorem on the existence of tree-indiscernibles. In order to be clear about the definability of the tree in the original vocabulary we extend Notation 4.1 and are quite pedantic about the vocabularies involved.

Notation 4.4 1. *τ_Φ includes both τ and τ^* and includes Skolem functions for τ_Φ, where the Skolem axioms and relations with crucial τ-formulas are axiomatized in a τ_Φ-theory T_1.*

2. *The set of constants C which guarantee the consistency of the order are added to τ_Φ.*

3. *Σ_i denotes the set of formulas $\phi \in \tau_\Phi - \{P_j : j > i\}$ with at most i free variables.*

Tree-indiscernibles are a special case of generalized indiscernibility as defined in VII.2 of [She78a]. Indiscernibles indexed by other types of structure appear for example in [LS03, Dǒ4, Sco]. The following notion of *modeling property*, based on one introduced by Scow[Sco] in a slightly different context is helpful for stating the results here. The point is that although the type of an infinite collection of indiscernibles may not be realized in any of the input models, each (even complete) type of a finite subsequence is. Thus properties of finite character (such as realizing a finite type) follow immediately if the indiscernibles have the modeling property. We use $\approx$ for isomorphic.

Definition 4.5 *Let Σ be a collection of τ_ϕ-formulas. A collection of Σ-tree-indiscernibles $B = \{b_\eta : \eta \in \boldsymbol{T}\}$ has the* modeling property *if it is derived from a sequence (M_α, X_α) (where $M_\alpha \supset X_\alpha = \{a_\eta : \eta \in \boldsymbol{T}_\alpha\}$, and $\boldsymbol{T}_\alpha \approx \boldsymbol{T}$ for $\alpha < H$) such that for every finite sequence ν from $\boldsymbol{T}$ and every sequence $\mathbf{b}_\nu$ from B and some α there is a sequence $\boldsymbol{a}_{\nu'} \in X_\alpha$ with ν' having the same τ^*-type as ν and such that $\boldsymbol{a}_{\nu'}$ and $\mathbf{b}_\nu$ have the same Σ-type.*

Note that in the argument below when the X_α are refined using Lemma 4.3 a tuple $\boldsymbol{a}_\nu \in X^i_{\alpha,n}$ was originally named $\boldsymbol{a}_{\nu'} \in X^0_{\alpha+m^*,n}$ (where $m^* < \omega$ can be easily computed). But, ν and ν' realize the same τ^*-type.

Remark 4.6 There are at least four approaches to the proof of Morley's omitting types theorem that differ subtly. 1) In [CK73, Mar02][3] the language is countable and there are separate steps to guarantee indiscernibility and omission of the types (meeting indiscernibility type omission requirements in turn for each formula and for each type). 2) The extension to uncountable languages is announced without proof in the original paper[Mor65a]. In the argument here, we use the Skolemization of the models M_α to deduce the omission of types from the indiscernibility. This argument strategy is forced because in dealing with uncountable languages, working with one formula at each step (as in [Mar02]) makes the induction too long to be actually carried out. We replace this long induction by working with all formulas with n-free variables at step n. 3) The arguments in [She78a, GL02, GS86] employ nonstandard-models of set theory. 4) Finally, the arguments in [Hod87, Kei71], work directly in infinitary logic using Hintikka sets or consistency properties. The arguments of [She78a, GS86, Hod87] make the connection with well-ordering numbers explicit.

Tsuboi [Tsu08] shows that a family of $< 2^{\aleph_0}$ *complete* types that is omitted up to $\aleph_{\omega_1}$ can be omitted in arbitrarily large models; this argument introduces some new combinatorial ideas.

Recall that $\mu_\alpha = \beth_\alpha(|T|)$. Writing μ_α rather than $\beth_\alpha$ and considering M_α for $\alpha < \delta(T) = (2^{|T|})^+$ is part of the price for dealing with uncountable T.

Note that when applying this theorem in Section 3, the M_α here are the M^+_α (as Skolemized) there.

Theorem 4.7 *Let T be a theory with Skolem functions in a vocabulary τ_Φ. Suppose for $\alpha < \delta(T)$, there exists a model M_α of T with $|M_\alpha| \geq \mu_\alpha$ such that M_α omits a family Γ of τ-types. τ_Φ contains the vocabulary τ^* and X_α is a set of elements in M_α that form a tree of type ${}^{<\omega}\mu_\alpha$ in M_α defined by the interpretations of $<, <^*, \wedge, P_n$, for*

[3] Compare comments on the proof in [CK73]. The stated result is the existence of large models omitting types without mentioning indiscernibility.

$n < \omega$. More restrictively, $X_{\alpha,n}$ is the restriction of X_α to the P_m for $m \leq n$; it has order type ${}^{\leq n}\mu_\alpha$.

Then, there is a countable set of tree-indiscernibles $C = \langle \mathbf{c}_\tau : \tau \in I \rangle$ with I of order type ${}^{<\omega}\omega$ such that C has the modeling property with respect to (M_α, X_α) and an extension Φ of T such that for every tree J of the form ${}^{<\omega}\lambda$, $EM_\tau(J, \Phi) \models T$, witnesses the universal τ_Φ-sentences that are true on all X_α, and omits Γ.

Proof. After expanding the language τ_Φ with new constants $\langle c_\rho : \rho \in {}^{<\omega}\omega \rangle$, we need to demonstrate the consistency of the following families of sentences.

1. $c_\rho \neq c_\eta$ if $\rho \neq \eta$.

2. For each τ_Φ-formula $\phi(\mathbf{v})$, for each quantifier-free τ^*-type r. If η, ν both realize r,
$$\phi(\mathbf{c}_\nu) \equiv \phi(\mathbf{c}_\eta).$$

3. For each ℓ-type $p \in \Gamma$, for each sequence of ℓ τ_Φ-terms $t_i(\mathbf{u})$ with $\lg(\mathbf{u}) = m$ ($\mathbf{t}(\mathbf{u}) = \langle t_0(\mathbf{u}), \ldots, t_{\ell-1}(\mathbf{u}) \rangle$) and each quantifier-free τ^*-m-type r, there is a $\phi_p(v_0, \ldots v_{\ell-1}) \in p$, such that if ν realizes r
$$\neg\phi_p(\mathbf{t}(\mathbf{c}_\nu)).$$

4. If $\chi(\mathbf{x}_1, \ldots \mathbf{x}_n)$ is a τ_Φ-formula that is true in all X_α (i.e on the substructure of the τ_Φ expansion of M_α with universe X_α for any n-tuple from X_α) then $\chi(\mathbf{c}_{\tau_1}, \ldots \mathbf{c}_{\tau_n}) \in \Phi$ for any $\tau_1, \ldots \tau_n \in {}^{<\omega}\lambda$[4].

Let $\boldsymbol{T} \subseteq {}^{<\omega}\lambda$ and $\boldsymbol{T}_n = \boldsymbol{T} \cap {}^{\leq n}\lambda$. We begin with pairs $(M_\alpha, X^0_{\alpha,n})$ for $n < \omega$, a model M_α, and a subset $X^0_{\alpha,n} = \{\boldsymbol{a}_\tau : \tau \in \boldsymbol{T}_{\alpha,n}\}$ which contains a sufficiently large tree as in the hypothesis of the theorem. Here, $\boldsymbol{T}_\alpha \subseteq {}^{<\omega}\mu_\alpha$ and $\boldsymbol{T}_{\alpha,n} = \boldsymbol{T} \cap {}^{\leq n}\mu_\alpha$.

We construct by induction for $i < \omega$ and for each n a pair $(M^i_\alpha, X^i_{\alpha,n})$ with $X^i_{\alpha,n} = \{\boldsymbol{a}_\tau : \tau \in \boldsymbol{T}^i_{\alpha,n}\} \subset \bigcup_{j \leq n} P_j(M^i_\alpha)$ with $(\boldsymbol{T}^i_{\alpha,n}, <, <^*) \approx {}^{\leq n}\mu_\alpha$. And we construct the diagram Φ, checking its finite consistency. Let Φ_0 include all τ_Φ sentences true in all $X^0_{\alpha,n}$ and the assertion that the c_ρ are distinct.

At stage i, we apply the next result, Claim 4.9, which needs some further notation.

Notation 4.8 *Let S_n be the collection of τ_Φ-n-types over the empty set which are realized in $\bigcup_{i \leq n} P_n(M_\alpha)$ (i.e. the Σ_n-types in the sense of Notation 4.4 that are realized).*

Claim 4.9 *The sequence $(M^i_\alpha, X^i_{\alpha,n})$ has the property that for each α:*

If $\eta, \nu \in \boldsymbol{T}^i_{\alpha,n}$ both realize the same quantifier-free τ^*-type r, and $n \leq i$ then for each $\phi \in \Sigma_n$

$$\phi(\mathbf{c}_\eta) \equiv \phi(\mathbf{c}_\nu). \tag{4}$$

Moreover, $(X^i_{\alpha,n}, <, <^*, \wedge) \approx ({}^{\leq n}\mu_\alpha, <, <^*, \wedge)$.

[4] An application of this observation is in the paragraph after Remark 2.14.

Proof. Consider $(M^i_{\alpha+k}, X^i_{\alpha+k,i})$ where $k = k(m,i)$. Let $f : [X^i_{\alpha+k,i}]^m \to S_n$, where $f(\nu) = s$ if $\mathrm{tp}_{\tau_\Phi}(\boldsymbol{a}_\nu) = s$. Now by Lemma 4.3, there is a $Y_{\alpha,i}$ (contained in $X^i_{\alpha+k,i} \subset \bigcup_{j\leq i} P_j(M_\alpha)$) and with $(Y_{\alpha,i}, <, <^*) \approx {}^{\leq n}\mu_\alpha$ and (4) is true on $Y_{\alpha,i}$. Denote $Y_{\alpha,i}$ as $X^{i+1}_{\alpha,i}$ and $M^i_{\alpha+k}$ as M^{i+1}_α. For $j \geq i$, let $X^{i+1}_{\alpha,j}$ be the elements of $X^i_{\alpha+k,j}$ that extend members of $Y_{\alpha,i} = X^{i+1}_{\alpha,i}$. $\Box_{4.9}$

To complete the proof of Theorem 4.7, we refine (and rename for convenience) the index set of ordinals to guarantee that for all α, each τ^*-type in S_n is given the same truth value for all tuples from $X^i_{\alpha,i}$ realizing r. This assignment gives us Φ_{n+1}. We can do this because at any stage, the number of Σ_n-theories is at most $2^{|T|}$ which is not cofinal in $(2^{|T|})^+$. Note that as i increases in this induction, the indiscernibility is being insured for larger Σ_i. Since the Σ_i are increasing this results in a consistent theory Φ giving tree-indiscernibility in $L(\tau_\Phi)$.

At stage i, we have assigned to each τ^*_i type r, a complete Σ_i-diagram in τ_Φ; each formula $\phi(\mathbf{v}) \in \Sigma_i$ has a fixed truth value for all $\mathbf{c}_\eta$ where η realizes r. In particular, since all M_α omit each ℓ-type $p \in \Gamma$ for any finite ℓ , for each sequence of ℓ-Skolem functions $\mathbf{t}$ in a most m-variable, and each η realizing a τ^*-type in m-variables there is a $\phi_p \in \Sigma_{\ell\cdot m}$ with $\phi_p \in p$ and $\neg\phi_p(\mathbf{t}(\mathbf{c}_\eta)$.

$\Box_{4.7}$

This completes the general proof for obtaining tree indiscernibles and so the proof of Theorem 3.3 is complete as well.

References

[Bal89] J. Baldwin. Diverse classes. *The Journal of Symbolic Logic*, 54:875–893, 1989.

[Bal09] John T. Baldwin. *Categoricity*. Number 51 in University Lecture Notes. American Mathematical Society, 2009. www.math.uic.edu/˜ jbaldwin.

[CK73] C.C. Chang and H.J Keisler. *Model theory*. North-Holland, 1973. 3rd edition 1990.

[DŎ4] Džamonja+Shelah. On $\lhd^*$-maximality. *Annals Pure and Applied Logic*, 125:119–158, 2004. item 692 in Shelah archive.

[GL02] R. Grossberg and Olivier Lessmann. Shelah's stability spectrum and homogeneity spectrum in finite diagrams. *Archive for Mathematical Logic*, 41:1–31, 2002.

[GS86] R. Grossberg and Saharon Shelah. A non structure theorem for an infinitary theory which has the unsuperstablity property. *Illinois J Math*, 30:364–390, 1986.

[Hod87] W. Hodges. What is a structure theory? *Bulletin of the London Mathematics Society*, 19:209–237, 1987.

[Hod93] W. Hodges. *Model Theory*. Cambridge University Press, 1993.

[Kei71] H.J Keisler. *Model theory for Infinitary Logic*. North-Holland, 1971.

[Kei76] H.J. Keisler. Six classes of theories. *J. Australian Math. Soc. Ser. A*, 25:257–266, 1976.

[KKS] Bjunghan Kim, Hyeung-Joon Kim, and Lynn Scow. Tree indiscernibility. preprint.

[LS03] Michael C. Laskowski and Saharon Shelah. Karp complexity and classes with the independence property. *Annals Pure and Applied Logic*, 120:263–283, 2003.

[Mar02] D. Marker. *Model Theory: An introduction*. Springer-Verlag, 2002.

[Mor65a] M. Morley. Categoricity in power. *Transactions of the American Mathematical Society*, 114:514–538, 1965.

[Mor65b] M. Morley. Omitting classes of elements. In Addison, Henkin, and Tarski, editors, *The Theory of Models*, pages 265–273. North-Holland, Amsterdam, 1965.

[MP84] George Mills and Jeff Paris. Regularity in models of arithmetic. *J. Symbolic Logic*, 49:272–280, 1984.

[Sco] L. Scow. Characterization of nip theories by random graph ordered indiscernibles. to appear: APAL.

[She70] S. Shelah. Finite diagrams stable in power. *Annals of Mathematical Logic*, 2:69–118, 1970.

[She74] S. Shelah. Categoricity of uncountable theories. In *Proc Tarski Symposium (Univ. of California, Berkeley, Calif., 1971)*, Proceedings of Symposia in Pure Mathematics XXV, pages 187–203. American Mathematical Society, 1974.

[She75] S. Shelah. A two-cardinal theorem. *Proc American Math Soc*, 48:207–213, 1975.

[She76] S. Shelah. A two-cardinal theorem and a combinatorial theorem. *Proc American Math Soc*, 62:134–136, 1976.

[She78a] S. Shelah. *Classification Theory and the Number of Nonisomorphic Models*. North-Holland, 1978.

[She78b] Saharon Shelah. *Classification Theory and the Number of Non-isomorphic Models*. North Holland, 1978.

[She83a] S. Shelah. Classification theory for nonelementary classes. I. the number of uncountable models of $\psi \in L_{\omega_1\omega}$ part A. *Israel Journal of Mathematics*, 46:3:212–240, 1983. paper 87a.

[She83b] S. Shelah. Classification theory for nonelementary classes. I. the number of uncountable models of $\psi \in L_{\omega_1\omega}$ part B. *Israel Journal of Mathematics*, 46;3:241–271, 1983. paper 87b.

[Tsu08] Akito Tsuboi. Models omitting given complete types. *Notre Dame Journal of Formal Logic*, 49:393–399, 2008.