EXAMPLES IN DEPENDENT THEORIES

ITAY KAPLAN AND SAHARON SHELAH

Abstract. We show a counterexample to a conjecture by Shelah regarding the existence of indiscernible sequences in dependent theories.

1. Introduction

In the summer of 2008, Saharon Shelah announced in a talk in Rutgers that he had proved some very important results in dependent (NIP) theories. One of these was the existence of indiscernible, an old conjuncture of his. Here is the definition:

Definition 1.1. Let T be a theory. For a cardinal κ, $n \leqslant \omega$ and an ordinal δ, $\kappa \to (\delta)_{T,n}$ means: for every sequence $\langle a_\alpha \,|\alpha \in \kappa \rangle \in {}^{\kappa}(\mathfrak{C}^n)$, there is a non-constant sub-sequence of length δ which is an indiscernible sequence.

In stable theories, it is known that for any λ satisfying $\lambda = \lambda^{|T|}$, $\lambda^+ \to (\lambda^+)_{T,n}$ (proved by Shelah in [10], and follows from local character of non-forking). The first such claim of this kind was proved by Morley in [3] for ω-stable theories. Later, in [9], Shelah proved

Theorem 1.2. *If T is strongly dependent, then for all $\lambda \geqslant |T|$, $\beth_{|T|^+}(\lambda) \to (\lambda^+)_{T,n}$ for all $n < \omega$.*

Shelah conjectured that this is true as well under just the assumption of NIP. This conjuncture is connected to a result by Shelah and Cohen: in [1], they proved that a theory is stable iff it can be presented in some sense in a free algebra in a fix vocabulary but allowing function symbols with infinite arity. If this result could be extended to: a theory is dependent iff it can be represented as an algebra with ordering, then this could be used to prove existence of indiscernibles.
Despite announcing it, there was a mistake in the proof for dependent theories, and here we shall see a counter-example. In this paper, we shall show that

Theorem 1.3. *There is a countable dependent theory T such that if κ is smaller than the first inaccessible cardinal, then for all $n \in \omega$, $\kappa \not\to (\omega)_{T,n}$.*

It appears in a more precise way as theorem 3.8 below.
An even stronger result can be obtained, namely

Theorem 1.4. *For every θ there is a dependent theory T of size θ such that for all κ and δ, $\kappa \to (\delta)_{T,1}$ iff $\kappa \to (\delta)^{<\omega}_{\theta}$.*

where

The authors would like to thank the Israel Science Foundation for partial support of this research (Grant No. 710/07). Publication no. 946 on Shelah's list.

Definition 1.5. $\kappa \to (\delta)^{<\omega}_{\theta}$ means: for every coloring $c : [\kappa]^{<\omega} \to \theta$ there is an homogeneous sub-sequence of length δ (i.e. there exists $\langle \alpha_i \,|i < \delta \rangle \in {}^{\delta}\lambda$ and $\langle c_n \,|n < \omega \rangle \in {}^{\omega}\theta$ such that $c\,(\alpha_{i_0}, \ldots, \alpha_{i_{n-1}}) = c_n$ for every $i_0 < \ldots < i_{n-1} < \delta$).

One can see that $\kappa \to (\delta)^{<\omega}_{\theta}$ implies $\kappa \to (\delta)_{T,1}$, so this is the best result possible. However, the proof is considerably harder, so we will give the proof in a subsequent paper.

The second part of the paper is devoted to giving a related example in the field of real numbers, $\mathbb{R}$. By 1.2, and as $Th\,(\mathbb{R})$ is strongly dependent, we know we cannot prove the same theorem, so we allow dealing with sequences of infinite tuples of elements of the model (ω-tuples in fact). The exact statement is:

Theorem 1.6. *If λ is smaller than the first inaccessible cardinal, then for all $n \in \omega$, $\kappa \not\to (\omega)_{RCF,\omega}$.*

This is theorem 4.4.

The third part of the paper is devoted to another kind of counterexample. In a series of papers ([7, 6, 8, 5]), Shelah has proved (among other things) that dependent theories give rise to a "generic pair" of models (and in fact this characterize dependent theories). This is explained in more details in the section. The natural question is whether the theory of the pair is again dependent. The answer is no. We present an example of an ω-stable theory all of whose generic pair (and even weakly generic – see there) have the independence property.

Notes. Many strong results that have been announced in that talk in 2008 remain true. For instance, the main theorem there is that a theory is dependent iff the number of types up to isomorphism is small and there is a so called polarized theorem (see more in [5]).
We also should note that a related result can be found in an unpublished paper in Russian by Kudajbergenov that states that for every ordinal α there exists a dependent theory (but it may be even strongly dependent) T_α such that $|T_\alpha| = |\alpha| + \aleph_0$ and $\beth_\alpha\,(|T_\alpha|) \not\to (\aleph_0)_{T_\alpha,1}$ and thus seem to show that the bound in 1.2 is tight.

Idea of the construction. The counterexample is a "tree of trees" with functions connecting the different trees. For every η in the tree ${}^{\omega\geqslant}2$ we shall have a predicate P_η and an ordering $<_\eta$ such that $(P_\eta, <_\eta)$ is a dense tree. In addition we shall have functions $G_{\eta,\eta^\frown\{i\}} : P_\eta \to P_{\eta^\frown\langle i\rangle}$ for $i = 0, 1$. The idea is to prove that $\kappa \nrightarrow (\mu)_{T,1}$ by induction on κ, i.e. to prove that in $P_{\langle\rangle}$ there are no indiscernibles. To use the induction hypothesis, we push the counter examples we already have for smaller κ-s to deeper levels in the tree ${}^{\omega\geqslant}2$.

2. Preliminaries

2.1. **Notation.** We use standard notation. a, b, c are elements, and $\bar{a}, \bar{b}, \bar{c}$ are finite or infinite tuples of elements.
$\mathfrak{C}$ will be the monster model of the theory.
$S_n\,(A)$ is the set of complete types over A, and $S^{qe}_n\,(A)$ is the set of all quantifier free complete types over A.
For a finite set of formulas with a partition of variables, $\Delta\,(\bar{x};\bar{y})$, $S_{\Delta(\bar{x};\bar{y})}\,(A)$ is the set of all Δ-types over A, i.e. maximal consistent subsets of

$\{\varphi(\bar{x},\bar{a}), \neg\varphi(\bar{x},\bar{a}) \,|\, \varphi(\bar{x},\bar{y}) \in \Delta \,\&\, \bar{a} \in A\}$. Similarly we define $\mathrm{tp}_{\Delta(\bar{x};\bar{y})}\left(\bar{b}/A\right)$ as the set of formulas $\varphi(\bar{x},\bar{a})$ such that $\varphi(\bar{x},\bar{y}) \in \Delta$ and $\mathfrak{C} \models \varphi\left(\bar{b},\bar{a}\right)$.

2.2. **Dependent theories.** For completeness, we give here the definitions and basic facts we need on dependent theories.

Definition 2.1. A first order theory T is dependent if it does not have the independence property which means: there is no formula $\varphi(\bar{x},\bar{y})$ and elements in $\left\langle \bar{a}_i, \bar{b}_s \,|\, i<\omega, s \subseteq \omega \right\rangle$ in $\mathfrak{C}$ such that $\varphi\left(\bar{a}_i,\bar{b}_s\right)$ iff $i \in s$.

We shall need the following fact (see [10, II, 4])

Fact 2.2. *Let* T *be any theory. Then for all* $n<\omega$, T *is dependent iff* $\Box_n$ *iff* $\Box_1$ *where for any* $k<\omega$,

$\Box_k$ For every finite set of formulas Δ, and $\varphi(\bar{x};\bar{y})$ such that $\lg(\bar{x}) = k$, there is a polynomial f such that for every finite set $A \subseteq M \models T$, $\left|S_{\Delta(\bar{x};\bar{y})}(A)\right| \leqslant f(|A|)$.

Also, later, we shall need the notion of $p^{(\omega)}$ for an invariant type p and a basic fact about them:

Definition 2.3. Suppose $p(x)$ and $q(y)$ are global invariant types, i.e. invariant over a small set A. Then $p \otimes q(x,y)$ is a global invariant type defined as follows: for any $B \supseteq A$, let $a \models p|_B$ and $b \models q|_{Ba}$, then $p \otimes q = \bigcup_{B \supseteq A} \mathrm{tp}(a,b/B)$. One can easily check that it is well defined and invariant over A. Let $p^{(n)} = p \otimes p \ldots \otimes p$ where the product is done n times. So $p^{(n)}$ is a type in $(x_0,\ldots,x_{n-1})$, and $p^{(\omega)} = \bigcup_{n<\omega} p^{(n)}$ and it is a type in $(x_0,\ldots,x_n,\ldots)$.

Fact 2.4. (*see* [2, Lemma 2.5]) *If* T *is NIP then for a set* A *the function with domain global invariant types over* A *that takes* p *to* $p^{(\omega)} \restriction A$ *is injective.*

3. An example of lack of indiscernibles

Let S_n be the finite binary tree ${}^{n\geqslant}2$. On a well ordered tree such as S_n, we define $<_{\mathrm{suc}}$ as follows: $\eta <_{\mathrm{suc}} \nu$ if ν is a successor of η in the tree.
Let L_n be the following language:

$$L_n = \{P_\eta, <_\eta, \wedge_\eta, G_{\eta,\nu} \,|\, \eta,\nu \in S_n, \eta <_{\mathrm{suc}} \nu\}.$$

Where:

- P_η is a unary predicate ; $<_\eta$ is a binary relation symbol ; $\wedge_\eta$ is a binary function symbol ; $G_{\eta,\nu}$ is a unary function symbol.

Let $T_n^\forall$ be the following theory:

- $P_\eta \cap P_\nu = \emptyset$ for $\eta \neq \nu$.
- $(P_\eta, <_\eta, \wedge_\eta)$ is a tree, where $\wedge_\eta$ is the intersection function on P_η, i.e. $x \wedge_\eta y = \max\{z \in P_\eta \,|\, z \leqslant_\eta x \,\&\, z \leqslant_\eta y\}$ (so if for example, $x \notin P_\eta$, then $x \wedge_\eta y = x$).
- $G_{\eta,\nu} : P_\eta \to P_\nu$, meaning that outside P_η, $G_{\eta,\nu}$ is the identity, and no further restrictions.

Thus we have:

Claim 3.1. $T_n^\forall$ is a universal theorem.

Claim 3.2. $T_n^\forall$ has JEP and AP.

Proof. Easy to see. □

From this we deduce,

Corollary 3.3. *$T_n^\forall$ has a model completion, T_n which eliminates quantifiers, and moreover: if $M \models T_{n+1}^\forall$, $M' = M \restriction L_n$ and $M' \subseteq N' \models T_n^\forall$ then N' can be enriched to a model N of $T_{n+1}^\forall$ so that $M \subseteq N$. Hence if M is an existentially closed model of $T_{n+1}^\forall$, then M' is an e.c. model of $T_n^\forall$. Hence $T_n \subseteq T_{n+1}$.*

Proof. The moreover part: for each $\eta \in S_{n+1} \backslash S_n$, we define $P_\eta^N = P_\eta^M$ and in the same way $\wedge_\eta$. The functions $G_{\eta,\nu}$ for $\eta \in S_n$ and $\nu \in S_{n+1}$ will be extensions of $G_{\eta,\nu}$. □

Now we show that T_n is dependent, but before that, a few easy remarks:

Observation 3.4.

(1) If $A \subseteq M \models T_0^\forall$ is a finite substructure (so just a tree, with no extra structure), then for all $b \in M$, the structure generated by A and b is $A \cup \{b\} \cup \{\max\{b \wedge a \,|a \in A\}\}$.

(2) If $M \models T_n^\forall$ and $\eta \in {}^{n\geqslant}2$, we can define a new structure $M_\eta \models T_{n-\lg(\eta)}^\forall$ whose universe is $\bigcup \left\{ P_{\eta\hat{}\nu}^M \,\middle|\nu \in {}^{n-\lg(n)\geqslant}2 \right\}$ by: $P_\nu^{M_\eta} = P_{\eta\hat{}\nu}^M$, and in the same way we interpret every other symbol (for instance, $G_{\nu_1,\nu_2}^{M_\eta} = G_{\eta\hat{}\nu_1,\eta\hat{}\nu_1}^M$). For every formula $\varphi(\bar{x}) \in L_{n-\lg(\eta)}$ there is a formula $\varphi'(\bar{x}) \in L_n$ such that for all $\bar{a} \in M_\eta$, $M \models \varphi'(\bar{a})$ iff $M_\eta \models \varphi(\bar{a})$ (we get φ' by concatenating η before any symbol).

(3) For M as before and $\eta \in {}^{n\geqslant}2$, for any $k < \omega$ there is a bijection between $\left\{p(x_0, \ldots, x_{k-1}) \in S_k^{qe}(M) \,|\forall i < k \,(P_\eta(x_i) \in p)\right\}$ and $\left\{p(x_0, \ldots, x_{k-1}) \in S_k^{qe}(M_\eta) \,\middle|\forall i < k \left(P_{\langle\rangle}(x_i) \in p\right)\right\}$.

Proof. (3): The bijection is given by (2). This is well defined, meaning that if $p(x_0, \ldots, x_{k-1})$ is a type over M_η such that $\forall i < k \left(P_{\langle\rangle}(x_i) \in p\right)$, then $\{\varphi'(\bar{x}) \,|\varphi(\bar{x}) \in p\}$ determines a complete type over M, such that $\forall i < k \,(P_\eta(x_i) \in p)$. This is because any atomic formula of the form $F(\bar{x}, \bar{a}) <_\nu G\left(\bar{x}, \bar{b}\right)$ where $\bar{a}, \bar{b} \in M$ and F, G are terms is either trivially false, as in the case $\nu \ngeqslant \eta$, or trivially equivalent to a formula with no involvement of the $\bar{x}$s at all (as in the case $a \wedge_\eta x <_\nu b$ where $a \notin P_\eta$) or it is equivalent to a formula of the form $F'(\bar{x}, \bar{a}') <_\nu G'\left(\bar{x}, \bar{b}'\right)$ where $\bar{a}', \bar{b}'$ are from M_η (for example, if $F = x_0 \wedge G_{\nu,\eta}(a)$, then $a' = G_{\nu,\eta}(a)$). Obviously it is injective and onto. □

Proposition 3.5. *T_n is dependent.*

Proof. We use fact 2.2. It is sufficient to find a polynomial $f(x)$ such that for every finite set A, $\left|S^1(A)\right| \leqslant f(|A|)$.
First we note that for a set A, the size of the structure generated by A is bounded by a polynomial in $|A|$: it is generated by applying $\wedge_{\langle\rangle}$ on $P_{\langle\rangle} \cap A$, applying $G_{\langle\rangle,\langle 1\rangle}$ and $G_{\langle\rangle,\langle 0\rangle}$, and then applying $\wedge_{\langle 0\rangle}, \wedge_{\langle 1\rangle}$ and so on. Every step in the process is polynomial, and it ends after n steps.
Hence we can assume that A is a substructure, i.e. $A \models T_n^\forall$.
The proof is by induction on n. To easy notation, we shall omit the subscript η

from $<_\eta$ and $\wedge_\eta$.
First we deal with the case $n = 0$. In T_0, $P_{\langle\rangle}$ is a a tree with no extra structure, while outside $P_{\langle\rangle}$ there is no structure at all. The number of types outside $P_{\langle\rangle}$ is bounded by $|A| + 1$ (because there is only one non-algebraic type). In the case that $P_{\langle\rangle}(x) \in p$ for some type p over A, we can characterize p by characterizing the type of $x' := \max\{a \wedge x \,|\, a \in A\}$, i.e. the cut that x' induces on the tree, and by knowing whether $x' = x$ or $x > x'$ (we note that in general, every theory of a tree is dependent by [4]).
Now assume that the claim is true for n. Suppose $\eta \in {}^{n+1\geqslant}2$ and $1 \leqslant \lg(\eta)$. By 3.4(3), there is a bijection between the types $p(x)$ over A where $P_\eta(x) \in p$ and the number of types $p(x)$ in $T_{n+1-\lg(\eta)}$ over A_η where $P_{\langle\rangle} \in p$. $A_\eta \models T^{\forall}_{n+1}$, and so by the induction hypothesis, the number of types over A_η is bounded by a polynomial in $|A_\eta| \leqslant |A|$. As the number of types $p(x)$ such that $P_\eta(x) \notin p$ for all η is bounded by $|A| + 1$ as in the previous case, we are left with checking the number of types $p(x)$ such that $P_{\langle\rangle}(x) \in p$.
In order to describe p, we first have to describe p restricted to the language $\left\{<_{\langle\rangle}, \wedge_{\langle\rangle}\right\}$, and this is polynomially bounded. Let $x' = \max\{a \wedge x \,|\, a \in A\}$. By the remark above, if $A \cup \{x\}$ is not closed under $\wedge_{\langle\rangle}$, x' is the only new element in the structure generated by $A \cup \{x\}$ in $P_{\langle\rangle}$. Hence , we are left to determine the type of the pairs $\left(G_{\langle\rangle,\langle i\rangle}(x), G_{\langle\rangle,\langle i\rangle}(x')\right)$ over A for $i = 0, 1$ (if x' is not new, then it's enough to determine the type of $G_{\langle\rangle,\langle i\rangle}(x)$). The number of these types is equal to the number of types of pairs in T_n over $A_{\langle i\rangle}$. As T_n is dependent we are done by fact 2.2. □

Definition 3.6. Let $L = \bigcup_{n<\omega} L_n$, $T = \bigcup_{n<\omega} T_n$ and $T^\forall = \bigcup_{n<\omega} T^\forall_n$ (it follows that T is the model completion of $T^\forall$).

Note that we need the moreover part of 3.3 to know that $T_n \subseteq T_{n+1}$.
We easily have

Corollary 3.7. T *is complete, it eliminates quantifiers and is dependent.*

We shall prove the following theorem:

Theorem 3.8. *For any two cardinals* $\mu \leqslant \kappa$ *such that in* $[\mu, \kappa]$ *there are no inaccessible cardinals,* $\kappa \nrightarrow (\mu)_{T,1}$.

We shall prove a slightly stronger statement, by induction on κ:

Claim 3.9. Given μ and κ, such that either $\kappa < \mu$ or there are no (uncountable) strongly inaccessible cardinals in $[\mu, \kappa]$, there is a model $M \models T^\forall$ such that $\left|P^M_{\langle\rangle}\right| \geqslant \kappa$ and $P^M_{\langle\rangle}$ does not contain a non-constant indiscernible sequence (for quantifier free formulas) of length μ.

Remark 3.10. From now on, indiscernible will only mean "indiscernible for quantifier free formulas"; this clearly suffices.

Proof. Suppose we have μ. The proof is by induction on κ. We divide into cases:

Case 1. $\kappa < \mu$. Clear.

Case 2. $\kappa = \mu = \aleph_0$. Denote $\eta_j = \langle 1 \dots 1\rangle$, i.e. the constant sequence of length j and value 1. Find $M \models T^\forall$ such that its universe contains a set $\{a_{i,j} \,|\, i, j < \omega\}$ where $a_{i,j} \neq a_{i',j'}$ for all $(i, j) \neq (i', j')$, $a_{i,j} \in P_{\eta_j}$ and in addition $G_{\eta_j,\eta_{j+1}}(a_{i,j}) =$

$a_{i,j+1}$ if $j < i$ and $G_{\eta_j,\eta_{j+1}}(a_{i,j}) = a_{0,j+1}$ otherwise. We also need that $P^M_{\langle\rangle} = \{a_{i,0} | i < \omega\}$. Any model satisfying these properties will do (so no need to specify what the tree structures are). Now. if in $P^M_{\langle\rangle} = \{a_{i,0} | i < \omega\}$ there is a non-constant indiscernible sequence, $\langle a_{i_k,0} | k < \omega \rangle$, then there is a large enough j such that

$$G_{\eta_j,\eta_{j+1}} \circ \ldots \circ G_{\eta_0,\eta_1}(a_{i_0,0}) = G_{\eta_j,\eta_{j+1}} \circ \ldots \circ G_{\eta_0,\eta_1}(a_{i_1,0}).$$

But for every large enough k, $G_{\eta_j,\eta_{j+1}} \circ \ldots \circ G_{\eta_0,\eta_1}(a_{i_1,0}) \neq G_{\eta_j,\eta_{j+1}} \circ \ldots \circ G_{\eta_0,\eta_1}(a_{i_k,0})$ – contradiction.

Case 3. κ is singular. Suppose $\kappa = \bigcup_{i<\sigma} \lambda_i$ where $\sigma, \lambda_i < \kappa$ for all $i < \sigma$. By the induction hypothesis, for $i < \sigma$ there is a model $M_i \models T^\forall$ (for λ_i) such that in $P^{M_i}_{\langle\rangle}$ there is no non-constant indiscernible sequence of length μ. Also, there is a model N (for σ) such that in $P^N_{\langle\rangle}$ there is no non-constant indiscernible sequence of length μ. We may assume that the universes of all these models are pairwise disjoint and disjoint from κ.
So there are enumerations without repetition $\{a_i | i < \sigma\} \subseteq P^N_{\langle\rangle}$, and
$\{b_j | \bigcup\{\lambda_l | l < i\} \leqslant j < \lambda_i\} \subseteq P^{M_i}_{\langle\rangle}$ that witness $\left|P^N_{\langle\rangle}\right| \geqslant \sigma$,
$\left|P^{M_i}_{\langle\rangle}\right| \geqslant \lambda_i \geqslant |\lambda_i \backslash \bigcup\{\lambda_l | l < i\}|$ resp. Let $\bar{M}$ be a model extending each M_i and hence containing the disjoint union of the sets $\bigcup_{i<\sigma} M_i$ (exists by JEP).
Define a new model $M \models T^\forall$: $\left(P^M_{\langle\rangle}, <_{\langle\rangle}\right) = (\kappa, \in)$ (so $\wedge_{\langle\rangle} = \min$); $\left(P^M_{\langle 1\rangle \frown \eta}, <_\eta\right) = \left(P^N_\eta, <_\eta\right)$ and $\left(P^M_{\langle 0\rangle \frown \eta}, <_{\langle 0\rangle,\eta}\right) = \left(P^{\bar{M}}_\eta, <_\eta\right)$. In the same way define $\wedge_\eta$ for all η of length $\geqslant 1$. The functions are also defined in the same way: $G^M_{\langle 1\rangle \frown \eta, \langle 1\rangle \frown \nu} = G^N_{\eta,\nu}$ and $G^M_{\langle 0\rangle \frown \eta, \langle 0\rangle \frown \nu} = G^{\bar{M}}_{\eta,\nu}$. We are left to define $G_{\langle\rangle,\langle 0\rangle}$ and $G_{\langle\rangle,\langle 1\rangle}$. So let: $G_{\langle\rangle,\langle 1\rangle}(\alpha) = a_{\min\{i | \alpha \in \lambda_i\}}$ and $G_{\langle\rangle,\langle 0\rangle}(\alpha) = b_\alpha$ for all $\alpha < \kappa$.
Note that if I is an indiscernible sequence contained in $P^M_{\langle 1\rangle}$ then I is an indiscernible sequence in N contained in $P^N_{\langle\rangle}$, and the same is true for $\langle 0\rangle$ and $\bar{M}$ resp.
Assume $\langle \alpha_j | j < \mu \rangle$ is an indiscernible sequence in $P^M_{\langle\rangle}$. Then $\left\langle G_{\langle\rangle,\langle 1\rangle}(\alpha_j) | j < \mu \right\rangle$ is a constant sequence (by the choice of N). So there is $i < \sigma$ such that $\bigcup\{\lambda_l | l < i\} \leqslant \alpha_j < \lambda_i$ for all $j < \mu$. So $\left\langle G_{\langle\rangle,\langle 0\rangle}(\alpha_j) = b_{\alpha_j} | j < \mu \right\rangle$ is a constant sequence (it's indiscernible in $P^{\bar{M}}_{\langle\rangle}$ and in fact contained in $P^{M_i}_{\langle\rangle}$), hence $\langle \alpha_j | j < \mu \rangle$ is constant, as we wanted.

Case 4. κ is regular uncountable. By the hypothesis of the claim, κ is not strongly inaccessible, so there is some $\lambda < \kappa$ such that $2^\lambda \geqslant \kappa$. By the induction hypothesis on λ, there is a model $N \models T^\forall$ such that in $P^N_{\langle\rangle}$ there is no non-constant indiscernible sequence of length μ. Let $\{a_i | i \leqslant \lambda\} \subseteq P^N_{\langle\rangle}$ witness that $\left|P^N_{\langle\rangle}\right|$ is at least λ.
Define $M \models T^\forall$ as follows: $P^M_{\langle\rangle} = {}^{\lambda \geqslant}2$ and the ordering is inclusion (equivalently, the ordering is by initial segment). $\wedge_{\langle\rangle}$ is defined naturally: $f \wedge_{\langle\rangle} g = f \restriction \min\{\alpha | f(\alpha) \neq g(\alpha)\}$.
For all η, let $P^M_{\langle 1\rangle \frown \eta} = P^N_\eta$, and the ordering and the functions are naturally induced from N. The main point is that we set $G_{\langle\rangle,\langle 1\rangle}(f) = a_{\lg(f)}$. Let $P^M_{\langle 0\rangle \frown \eta}$, $G_{\langle 0\rangle \frown \eta, \langle 0\rangle \frown \nu}$, etc. have any legal values according to $T_\forall$. Let $G_{\langle\rangle,\langle 0\rangle}$ be any function.
Suppose that $\langle f_i | i < \mu \rangle$ is an non-constant indiscernible sequence:

If $f_1 < f_0$ (i.e. $f_1 <_{\langle\rangle} f_0$), we shall have an infinite decreasing sequence in a well-ordered tree – a contradiction.
If $f_0 < f_1$, $\langle f_i \,|i < \mu\rangle$ is increasing, so $\left\langle G^M_{\langle\rangle,\langle 1\rangle}(f_i) = a_{\lg(f_i)} \,|i < \mu \right\rangle$ is non-constant – contradiction (as it is an indiscernible sequence in M and hence in $P^N_{\langle 0\rangle}$).
Let $h_i = f_0 \wedge f_{i+1}$ for $i < \mu$ (where $\wedge = \wedge_{\langle\rangle}$). This is an indiscernible sequence, and by the same arguments, it cannot increase or decrease, but as $h_i \leqslant f_0$, and $\left(P_{\langle\rangle}, <_{\langle\rangle}\right)$ is a tree, it follows that h_i is constant.
Assume $f_0 \wedge f_1 < f_1 \wedge f_2$, then $f_{2i} \wedge f_{2i+1} < f_{2(i+1)} \wedge f_{2(i+1)+1}$ for all $i < \mu$, and again $\langle f_{2i} \wedge f_{2i+1} \,|i < \mu\rangle$ an increasing indiscernible sequence and we have a contradiction.
By the same reasoning, it cannot be that $f_0 \wedge f_1 > f_1 \wedge f_2$. As both sides are less or equal than f_1, it must be that $f_0 \wedge f_2 = f_0 \wedge f_1 = f_1 \wedge f_2$. But that is a contradiction (because if $\alpha = \lg(f_0 \wedge f_1)$, then $|\{f_0(\alpha), f_1(\alpha), f_2(\alpha)\}| = 3$).

□

4. RCF

Here we give a related theorem about the theory of Real closed fields, i.e. $\mathsf{Th}(\mathbb{R}, +, \cdot, 0, 1, <)$. Fix RCF as our theory, so $\mathfrak{C} \models \mathsf{RCF}$.
For our proof we need a more lenient version of the arrow relation (see 4.2 below), this helps with the induction hypothesis.

Notation 4.1. The set of all open intervals (a, b) (where $a < b$ and $a, b \in \mathfrak{C}$) is denoted by $\mathfrak{I}$.

Definition 4.2. For a cardinal κ, $n \leqslant \omega$ (i.e. $n < \omega$ or $n = \omega$) and an ordinal δ, $\kappa \to (\delta)^{\text{interval}}_n$ means: for every sequence of (non-empty, open) n-tuples of intervals $\left\langle \bar{I}_\alpha \,|\alpha < \kappa \right\rangle \in {}^{\kappa}(\mathfrak{I}^n)$ (so for each α, $\bar{I}_\alpha = \langle I^\varepsilon_\alpha \,|\varepsilon < n\rangle$), there is a subset $u \subseteq \kappa$ of order type δ, and $\left\langle \bar{b}_\varepsilon \,|\varepsilon \in u \right\rangle$ such that $\bar{b}_\varepsilon \in \bar{I}_\varepsilon$ (i.e. $\bar{b}_\varepsilon = \left\langle b^0_\varepsilon, \ldots, b^{n-1}_\varepsilon \right\rangle$ and $b^i_\varepsilon \in I^i_\varepsilon$) such that $\left\langle \bar{b}_\varepsilon \,|\varepsilon \in \delta \right\rangle$ is a non-constant indiscernible sequence.

Remark 4.3. Note that

(1) If $\kappa \to (\delta)^{\text{interval}}_n$ then $\kappa \to (\delta)^{\text{interval}}_m$ for all $m \leqslant n$.
(2) If $\kappa \not\to (\delta)^{\text{interval}}_n$ then $\kappa \not\to (\delta)_{\mathsf{RCF},n}$ (why? if $\left\langle \bar{I}_\alpha \,|\alpha < \kappa \right\rangle$ witness that $\kappa \not\to (\delta)^{\text{interval}}_n$, then choose $\bar{b}_\alpha \in \bar{I}_\alpha$ (as above) such that $\left\langle \bar{b}_\alpha \,|\alpha < \kappa \right\rangle$ is without repetitions, and by definition it will not have an indiscernible sub-sequence of length δ).
(3) If $\lambda < \kappa$ and $\kappa \not\to (\delta)^{\text{interval}}_n$ then $\lambda \not\to (\delta)^{\text{interval}}_n$.

We shall prove the following theorem:

Theorem 4.4. *For any two cardinals $\mu \leqslant \kappa$ such that in $[\mu, \kappa]$ there are no inaccessible cardinals,* $\kappa \nrightarrow (\mu)^{\text{interval}}_\omega$.

The proof follows from a sequence of claims:

Claim 4.5. If $\kappa < \mu$ then $\kappa \not\to (\mu)^{\text{interval}}_n$ for all $n \leqslant \omega$.

Proof. Obvious. □

Claim 4.6. If $\kappa = \mu = \aleph_0$ then $\kappa \not\to (\mu)_1^{\text{interval}}$.

Proof. For $n < \omega$, let $I_n = (n, n+1)$. □

Claim 4.7. Suppose $\kappa = \bigcup\{\lambda_i \,|i < \sigma\}$ and $n \leqslant \omega$. Then, if $\sigma \not\to (\mu)_n^{\text{interval}}$ and $\lambda_i \not\to (\mu)_n^{\text{interval}}$ then $\kappa \not\to (\mu)_{2+2n}^{\text{interval}}$.

Proof. By assumption, we have a sequence of intervals $\left\langle \bar{R}_i \,|i < \sigma \right\rangle$ that witness $\sigma \not\to (\mu)_n^{\text{interval}}$ and for each $i < \sigma$ we have $\left\langle \bar{S}_\beta \,\middle| \bigcup_{j<i} \lambda_j < \beta < \lambda_i \right\rangle$ that witness $\lambda_i \not\to (\mu)_n^{\text{interval}}$.
First, fix an increasing sequence of elements $\langle b_i \,|i < \sigma \rangle$.
For $\alpha < \kappa$, let $\beta = \beta(\alpha) = \min\{i < \sigma \,|\alpha < \lambda_i\}$ and for $i < 2+2n$, define:

- If $i = 0$, let $I_\alpha^i = \left(b_{2\beta(\alpha)}, b_{2\beta(\alpha)+1}\right)$.
- If $i = 1$, let $I_\alpha^i = \left(b_{2\beta(\alpha)+1}, b_{2\beta(\alpha)+2}\right)$.
- If $i = 2k+2$, let $I_\alpha^i = R_{\beta(\alpha)}^k$.
- If $i = 2k+3$, let $I_\alpha^i = S_\alpha^k$.

Suppose $u \subseteq \kappa$ of order type μ and $\left\langle \bar{b}_\varepsilon \,|\varepsilon \in u \right\rangle$ is a non-constant indiscernible sequence such that $\bar{b}_\varepsilon \in \bar{I}_\varepsilon$ (as above). Denote $\bar{b}_\varepsilon = \left\langle b_0^\varepsilon, \ldots, b_{2+2n-1}^\varepsilon \right\rangle$. Note that $b_1^\varepsilon < b_0^{\varepsilon'}$ iff $\beta(\varepsilon) < \beta(\varepsilon')$ (we need two intervals for the only if direction).
Hence $\langle \beta(\varepsilon) \,|\varepsilon \in u \rangle$ is increasing or constant. But if it is increasing then we have a contradiction to the choice of $\left\langle \bar{R}_i \,|i < \sigma \right\rangle$. So it is constant, and suppose $\beta(\varepsilon) = i_0$ for all $\varepsilon \in u$. But then $u \subseteq \lambda_{i_0} \setminus \bigcup_{j<i_0} \lambda_j$ and we get a contradiction to the choice of $\left\langle \bar{S}_\beta \,\middle| \bigcup_{j<i_0} \lambda_j < \beta < \lambda_{i_0} \right\rangle$. □

Claim 4.8. Suppose $\lambda \not\to (\mu)_n^{\text{interval}}$. Then $2^\lambda \not\to (\mu)_{4+2n}^{\text{interval}}$.

Proof. Suppose $\left\langle \bar{I}_\alpha \,|\alpha < \lambda \right\rangle$ witnesses that $\lambda \not\to (\mu)_n^{\text{interval}}$.
By adding two intervals to each $\bar{I}_\alpha$, we can ensure that it has the extra property that if $\bar{c}_1 \in I_{\alpha_1}$ and $\bar{c}_2 \in \bar{I}_{\alpha_2}$ then $c_1^1 < c_2^0$ iff $\alpha_1 < \alpha_2$ (as in the previous claim).
Notation: write $\bar{c}_1 <^* \bar{c}_2$ for $c_1^1 < c_2^0$, it is not really an ordering (it is not transitive) but on tuples $\bar{c}$ that belong to some $\bar{I}_\alpha$, it is transitive.
By this we have increased the length of $\bar{I}_\alpha$ to $2+n$.
We shall find below a definable 4-place function f such that:

♡ For every 2 ordinals, δ, ζ, If $\left\langle \bar{R}_\alpha \,|\alpha < \delta \right\rangle$ is a sequence of ζ-tuples of intervals, then there exists a sequence of 2ζ-tuples of intervals, $\left\langle \bar{S}_\eta \,\middle|\eta \in {}^\delta 2 \right\rangle$ such that for all $k < \zeta$ and $\eta_1 \neq \eta_2$, if $b_1 \in S_{\eta_1}^{2k}, b_2 \in S_{\eta_1}^{2k+1}$ and $b_3 \in S_{\eta_2}^{2k}, b_4 \in S_{\eta_2}^{2k+1}$ then $f(b_1, b_2, b_3, b_4)$ is in $R_{\lg(\eta_1 \wedge \eta_2)}^k$.

Apply ♡ to our situation to get $\left\langle \bar{J}_\eta \,\middle|\eta \in {}^\lambda 2 \right\rangle$ such that $\bar{J}_\eta = \langle J_\eta^i \,|i < 4+2n \rangle$ and for all $k < 2+n$ and $\eta_1 \neq \eta_2$, if $b_1 \in J_{\eta_1}^{2k}, b_2 \in J_{\eta_1}^{2k+1}$ and $b_3 \in J_{\eta_2}^{2k}, b_4 \in J_{\eta_2}^{2k+1}$ then $f(b_1, b_2, b_3, b_4)$ is in $I_{\lg(\eta_1 \wedge \eta_2)}^k$.
This is enough (the reasons are exactly as in the regular case of the proof of Theorem 3.8, but we shall repeat it for clarity):
To simplify notation, we regard f as a function on tuples: in each pair of consecutive

coordinates it acts as f, so if $\bar{b}_1 \in \bar{J}_{\eta_1}, \bar{b}_2 \in \bar{J}_{\eta_2}$ then $f\left(\bar{b}_1, \bar{b}_2\right)$ is in $\bar{I}_{\lg(\eta_1 \wedge \eta_2)}$ (more precisely, $f\left(\bar{b}_1, \bar{b}_2\right) = \langle a_k \,|k < 2+n\rangle$ where $a_k = f\left(b^1_{2k}, b^1_{2k+1}, b^2_{2k}, b^2_{2k+1}\right) \in I^i_{\lg(\eta_1 \wedge \eta_2)}$ for $k < 2+n$).
Suppose $u = \langle \eta_i \,|i < \mu\rangle \subseteq {}^{\lambda}2$ and $\left\langle \bar{b}_{\eta_i} \,|i < \mu\right\rangle$ is a non-constant indiscernible sequence such that $\bar{b}_{\eta_i} \in \bar{J}_{\eta_i}$ (as above).
Let $h_i = \eta_0 \wedge \eta_{i+1}$ for $i < \mu$. If $\lg(h_0) < \lg(h_1)$ then $f\left(\bar{b}_{\eta_0}, \bar{b}_{\eta_{i+1}}\right) <^* f\left(\bar{b}_{\eta_0}, \bar{b}_{\eta_{i+2}}\right)$ for all $i < \mu$, so $\langle \lg(h_i) \,|i < \mu\rangle$ is increasing, and so $f\left(\bar{b}_{\eta_0}, \bar{b}_{\eta_{i+1}}\right)$ contradicts our choice of $\left\langle \bar{I}_\alpha \,|\alpha < \lambda\right\rangle$. Hence (because $h_i \leqslant \eta_0$) h_i is constant.
Assume $\eta_0 \wedge \eta_1 < \eta_1 \wedge \eta_2$, then $f\left(\bar{b}_{\eta_0}, \bar{b}_{\eta_1}\right) <^* f\left(\bar{b}_{\eta_1}, \bar{b}_{\eta_2}\right)$ so $f\left(\bar{b}_{\eta_1}, \bar{b}_{\eta_2}\right) <^* f\left(\bar{b}_{\eta_2}, \bar{b}_{\eta_3}\right)$, and so $\lg(\eta_0 \wedge \eta_1) < \lg(\eta_1 \wedge \eta_2) < \lg(\eta_2 \wedge \eta_3)$ hence, $f\left(\bar{b}_{\eta_0}, \bar{b}_{\eta_1}\right) <^* f\left(\bar{b}_{\eta_2}, \bar{b}_{\eta_3}\right)$ and it follows that $\langle \lg(\eta_{2i} \wedge \eta_{2i+1}) \,|i < \mu\rangle$ is increasing. And this is again a contradiction.
By an analog analysis, it cannot be that $\eta_0 \wedge \eta_1 > \eta_1 \wedge \eta_2$. As both sides are less or equal than η_1, it must be that $\eta_0 \wedge \eta_2 = \eta_0 \wedge \eta_1 = \eta_1 \wedge \eta_2$. But that is impossible (because if $\alpha = \lg(\eta_0 \wedge \eta_1)$, then $|\{\eta_0(\alpha), \eta_1(\alpha), \eta_2(\alpha)\}| = 3$).

Claim. ♡ is true.

Proof. Let $f(x, y, z, w) = (x - z) / (y - w)$ (do not worry about division by 0, we shall explain below).
It is enough, by the nature of ♡, to assume $\zeta = 1$ (we treat each family of intervals separately). By compactness, we may assume that δ is finite, and to avoid confusion, denote it by m. So we have a finite tree, ${}^{m\geqslant}2$, and a sequence of intervals $\langle R_i \,|i < m\rangle$. Each R_i is of the form (a_i, b_i). Let $c_i = (b_i + a_i)/2$. Let $d \in \mathfrak{C}$ be any element greater than any member of $A := \operatorname{acl}(a_i, b_i \,|i < m)$, and $0 < e \in \mathfrak{C}$ such that $e < d^{-k}$ for all $k \in \mathbb{N}$. For each $\eta \in {}^m2$, let $a_\eta = \sum_{i<m} \eta(i)\, c_i d^{m-i}$, and $b_\eta = \sum_{i<m} \eta(i)\, d^{m-i}$.
Let $S^0_\eta = (a_\eta - e, a_\eta + e)$ and $S^1_\eta = (b_\eta - e, b_\eta + e)$.
This works:
Assume $b_1 \in S^0_{\eta_1}, b_2 \in S^1_{\eta_1}$ and $b_3 \in S^0_{\eta_2}, b_4 \in S^1_{\eta_2}$.
We have to show $(b_1 - b_3)/(b_2 - b_4) \in R_{\lg(\eta_1 \wedge \eta_2)}$. Denote $k = \lg(\eta_1 \wedge \eta_2)$.
$a_{\eta_1} - a_{\eta_2}$ is of the form $\varepsilon c_k d^{m-k} + F(d)$ where $\varepsilon \in \{-1, 1\}$, and $F(d)$ is a polynomial over A of degree $\leqslant m-k-1$. $b_{\eta_1} - b_{\eta_2}$ is of the form $\varepsilon d^{m-k} + G(d)$, where ε is the same for both (and G is a polynomial over $\mathbb{Z}$ of degree $\leqslant m-k-1$). Now, $b_1 - b_3 \in (a_{\eta_1} - a_{\eta_2} - 2e, a_{\eta_1} - a_{\eta_2} + 2e)$, and $b_2 - b_4 \in (b_{\eta_1} - b_{\eta_2} - 2e, b_{\eta_1} - b_{\eta_2} + 2e)$, and hence we know that $b_2 - b_4 \neq 0$. It follows that $(b_1 - b_3)/(b_2 - b_4)$ is inside an interval whose endpoints are $\left\{\left(\varepsilon c_k d^{m-k} + F(d) + 2e\right) / \left(\varepsilon d^{m-k} + G(d) \pm 2e\right)\right\}$.
But

$$\left(\varepsilon c_k d^{m-k} + F(d) + 2e\right) / \left(\varepsilon d^{m-k} + G(d) \pm 2e\right) \in R_k$$

by our choice of d and e, and we are done. □

□

The proof of Theorem 4.4 now follows by induction on κ: Fix μ, and let κ be the first cardinal for which the theorem fails. Then by 4.5, $\kappa \geqslant \mu$. By claim 4.6,

$\aleph_0 < \kappa$. By claim 4.7, κ cannot be singular. By claim 4.8, κ cannot be regular, because if it were, there would be a $\lambda < \kappa$ such that $2^\lambda \geqslant \kappa$ (because κ is not strongly inaccessible). Note that we did use claim 4.5 to deal with cases where we couldn't use the induction hypothesis (for example, in the regular case, it might be that $\lambda < \mu$).

More remarks. Theorem 4.4 can be generalized to allow parameters:
Suppose $\mathfrak{C} \models \mathsf{RCF}$, and $A \subseteq \mathfrak{C}$.

Definition 4.9. $\kappa \to_A (\mu)^{\text{interval}}_{\omega}$ means the same as in definition 4.2, but we demand indiscernible sequences to be indiscernible over A.

Then we have:

Theorem 4.10. *For any cardinal $|A| \leqslant \kappa$ such that in $[|A|, \kappa]$ there are no inaccessible cardinals, $\kappa \nrightarrow_A (\aleph_0)^{\text{interval}}_{\omega}$.*

Proof. The proof goes exactly as in the proof of Theorem 4.4, but the base case for the induction is different. There it was $\mu = \kappa = \aleph_0$. Here it is $\kappa = |A|$. Indeed, enumerate $A = \{a_i \,|\, i < \mu\}$. Let $\varepsilon \in \mathfrak{C}$ be greater than 0 but smaller than any element in $\operatorname{acl}(A)$. For $\alpha < \mu$, let $I_\alpha = (a_i, a_i + \varepsilon)$.
Suppose there is a subset $u = \{\varepsilon_i \,|\, i < \omega\} \subseteq \mu$ and $\langle b_i \,|\, i < \omega\rangle$ such that $b_i \in I_{\varepsilon_i}$ and $\langle b_i \,|\, i < \omega\rangle$ is non-constant and indiscernible. Let $i_0 = i_{\varepsilon_0}, i_1 = i_{\varepsilon_1}$. Suppose without loss of generality that $a_{i_0} < a_{i_1}$, then $b_0 \in (a_{i_0}, a_{i_0} + \varepsilon)$, so $b_0 < a_{i_1}$. But $a_{i_1} < b_1$ – contradiction. □

5. Generic pair

Here we give an example of an ω-stable theory, such that for all weakly generic pairs of structures $M \prec M_1$ the theory of the pair (M, M_1) has the independence property. Here is the definition:

Definition 5.1. A pair (M, M_1) as above is *weakly generic* if for all formula $\varphi(x)$ with parameters from M, if φ has infinitely many solutions in M, then it has a solution in $M_1 \backslash M$.

This definition is driven by the well known "generic pair conjecture" (see [7, 5]), and it is worth while to give the precise definitions.

Definition 5.2. Assume $\lambda = \lambda^{<\lambda} > |T|$ (in particular, λ is regular), $2^\lambda = \lambda^+$, and for all $\alpha < \lambda^+$, $M_\alpha \models T$ is of cardinality λ. Suppose $\langle M_\alpha \,|\, \alpha < \lambda^+\rangle$ is an increasing continuous sequence. Furthermore, $M = \bigcup_{\alpha<\lambda^+} M_\alpha$ is a saturated model of cardinality λ^+. The generic pair property says that there exists a club $E \subseteq \lambda^+$ such that for all $\alpha < \beta \in E$, the pair (M_α, M_β) has the same isomorphism type.

Proposition 5.3. *This is a property of M, and does not depend on the particular choice of $\langle M_\alpha\rangle$.*

Proof. Suppose M satisfies the definition 5.2. The means that $|M| = \lambda^+$, M is saturated, there are $\langle M_\alpha \,|\, \alpha < \lambda^+\rangle$ such that $M = \bigcup_{\alpha<\lambda^+} M_\alpha$, and there is a club $E \subseteq \lambda^+$ such that for all $\alpha, \beta \in E$ such that $\alpha < \beta$, the pair (M_α, M_β) has the same isomorphism type. Suppose $M = \bigcup_{\alpha<\lambda^+} N_\alpha$ for another increasing continuous sequence $\langle N_\alpha \,|\, \alpha < \lambda^+\rangle$. Let $E_0 = \{\delta < \lambda^+ \,|\, N_\delta = M_\delta\}$. This is a club of λ^+, and so $E \cap E_0$ is also a club of λ^+ showing that $\langle N_\alpha\rangle$ has the same property. □

Justifying definition 5.1 we have

Claim 5.4. Assume that M has the generic pair property, and $N \prec N_1 \models T$.
Then (1) implies (2) where:

(1) (N, N_1) have the isomorphism type of a pair (M_α, M_β) as in definition 5.2 above.
(2) (N, N_1) is weakly generic as in definition 5.1.

Proof. So suppose that E, $\langle M_\alpha \,|\alpha < \lambda^+ \rangle$ is as in definition 5.2. Suppose $\alpha, \beta \in E$ and $\alpha < \beta$. We are given a formula $\varphi(x)$ with parameter from M_α, such that $\aleph_0 \leqslant |\varphi(M_\alpha)|$. But by saturation of M, $\lambda^+ = |\varphi(M)|$, and $M = \bigcup_{\alpha \in E} M_\alpha$, hence there is some $\alpha < \beta' \in E$ such that $\varphi(M_\beta) \neq \emptyset$, but as $(M_\alpha, M_\beta) \cong (M_\alpha, M_{\beta'})$, we are done. ☐

The generic pair property is very important in dependent theories. In fact, they characterize them: In [6, 5, 7, 8] it is proved that T is dependent iff for large enough regular λ every model of T has the generic pair property (assuming GCH).
Hence it makes sense to ask whether the theory of the pair is dependent.
The answer is no:

Theorem 5.5. *There exists an ω-stable theory such that for every weakly generic pair of models $M \prec M_1$, the theory of the pair (M, M_1) has the independence property.*

We shall describe the example:
Let $L = \{P, R, Q_1, Q_2\}$ where R, P are unary predicates and Q_1, Q_2 are binary relations.
Let $\tilde{M}$ be the following structure for L:

(1) The universe is

$$\begin{aligned} \tilde{M} \;=\; & \{u \subseteq \omega \,||u| < \omega\} \cup \\ & \{(u, v, i) \,|u, v \subseteq \omega, |u| < \omega, |v| < \omega, i \in \omega \,\&\, u \subseteq v \Rightarrow i < |v| + 1\}. \end{aligned}$$

(2) The predicates are interpreted as follows:
- $P^{\tilde{M}} = \{u \subseteq \omega \,||u| < \aleph_0\}$,
- $R^{\tilde{M}}$ is $M \backslash \left(P^M\right)$,
- $Q_1^{\tilde{M}} = \left\{(u, (u, v, i)) \,\middle|\, u \in P^{\tilde{M}}\right\}$ and
- $Q_2^{\tilde{M}} = \left\{(v, (u, v, i)) \,\middle|\, v \in P^{\tilde{M}}\right\}$.

Let $T = Th\left(\tilde{M}\right)$.
As we shall see in the next claim, T gives rise to the following definition:

Definition 5.6. We call a structure $(B, \cup, \cap, -, \subseteq, 0)$ a pseudo Boolean algebra (PBA) when it satisfies all the axioms of a Boolean algebra except: There is no greatest element 1 (i.e. remove all the axioms concerning it).

Pseudo Boolean algebra can have atoms like in Boolean algebras (nonzero elements that do not contain any smaller nonzero elements).

Definition 5.7. Call a PBA of finite type if every element is a union of finitely many atoms.

Notation 5.8. For a PBA A, and $C \subseteq A$ a sub-PBA, let
$A_C := \left\{ a \in A \,\middle|\, \exists c \in C \left(a \subseteq^A c \right) \right\}$, and for a subset $D \subseteq A$, let $at(D)$ be the set of atoms contained in D.

Proposition 5.9. *Every PBA of finite type is isomorphic to* $(\mathcal{P}_{<\infty}(\kappa), \cup, \cap, -, \subseteq, \emptyset)$ *for some* κ *where* $\mathcal{P}_{<\infty}(\kappa)$ *is the set of all finite subsets of* κ*. Moreover: Assume* A, B *are infinite PBAs of finite type and* $C \subseteq A, B$ *is a common sub-PBA. Then, if*

1. $|\mathrm{at}(A \backslash A_C)| = |\mathrm{at}(B \backslash B_C)|$.
2. *For every* $c \in C$*,* A *and* B *agree on the size of* c *(the number of atoms it contains).*

Then there is an isomorphism of PBAs $f : A \to B$ *such that* $f \restriction C = \mathrm{id}$*.*

Proof. The first part follows from the easy observation that in a PBA of finite type, every element has a unique presentation as a union of finitely many atoms. So if A is a PBA, and its set of atoms is $\{a_i \,|\, i < \kappa\}$, then take a_i to $\{i\}$.
For the moreover part, first we do a back-and-forth, adding elements $a \in A$ and $b \in B$ to the domain and range of a partial isomorphism (starting with $\mathrm{id} \restriction C$) to extend it to an isomorphism from A_C to B_C.
Let $a \in A$. First assume a is an atom and that $a \subseteq c$ for some $c \in C$. Let $d \subseteq c$ be of maximal size such that $d \in C$. Then $a \subseteq^A c - d$, and there must be some $b \in B$ such that $b \subseteq^B c - d$ (because the size of c and d is the same). Take a to b.
If a is not an atom, but $a \subseteq c$ for some $c \in C$, then by the previous case we may assume that the domain of the isomorphism already contains all of c's atoms. So $c \in C$ (as their union) and we are done.
Now, $|at(A \backslash A_C)| = |at(B \backslash B_C)|$, so any bijection between the set of atoms induces an isomorphism. □

Claim 5.10. T is ω-stable.

Proof. First add new relations to the language, which are all definable –

$$\{S_n, \subseteq_n, \pi_1, \pi_2, \cap_n, \cup_n, -_n, e \,|\, n \geqslant 1\}$$

where S_n is a unary relation defined on P, $\subseteq_n$ is a binary relation defined on P, π_1, π_2 are two unary functions from R to P, $\cap_n, -_n$ are binary functions from S_n to S_n and e is a constant in P. Their interpretation in $\tilde{M}$ are as follows:

- $\pi_1((u, v, i)) = u$, $\pi_2((u, v, i)) = v$.
- For each $1 \leqslant n < \omega$, $S_n(v) \Leftrightarrow |v| \leqslant n$.
- For each $1 \leqslant n$, $u \subseteq_n v$ iff $|u| \leqslant n$, $|v| \leqslant n$ and $u \subseteq v$.
- $u \cap_n v = u \cap v$ for all $u, v \in S_n$.
- $u -_n v = u \backslash v$ for $v \subseteq_n u$.
- $u \cup_n v = u \cup v$ for $u, v \in S_n$.
- $e = \emptyset$.

Note that they are indeed definable:

1. $\pi_1(x)$ is the only y such that $Q_1(y, x)$, and similarly for π_2.
2. Let $E(x, y)$ by an auxiliary equivalence relation defined by $\pi_1(x) = \pi_1(y) \wedge \pi_2(x) = \pi_2(y)$.
3. e is the unique element in P such that there exists exactly one element $z \in R$ such that $\pi_1(z) = x = \pi_2(z)$.

(4) $x \subseteq_n y$ is defined by "$S_n(x), S_n(y)$ and the number of elements in the E class of some (equivalently any) element z such that $\pi_1(z) = x$, $\pi_2(z) = y$ is at most $n+1$".
(5) $S_n(x)$ is defined by "$P(x)$ and $e \subseteq_n x$" (In particular, $e \in S_n$ for all n).
(6) $\cap_n$ and $-_n$ are then naturally definable using $\subseteq_n$. For instance $x -_n y = z$ iff x, y, z are in S_n, $z \subseteq_n x$ and for each $e \neq w \subseteq_n y$, $w \nsubseteq_n z$.
(7) $x \cup_n y = z$ iff $x, y \in S_n$, $z \in S_{2n}$, $x, y \subseteq_{2n} z$ and $z -_{2n} x \subseteq_{2n} y$.

Furthermore, $\subseteq_k \restriction S_n = \subseteq_n$ for $n \leqslant k$. Hence every model M of T gives rise naturally to an induced PBA: $B^M := \left(\bigcup_n S_n^M, \cup^M, \cap^M, -^M, \subseteq^M, e^M\right)$ where $\cup^M = \bigcup \left\{\cup_n^M | n < \omega\right\}$, and similarly for $\subseteq^M, -^M$ and $\cap^M$ (see definition 5.6).

Claim. In the extended language, T eliminates quantifiers.

Proof. Suppose $M, N \models T$ are saturated models, $|M| = |N|$ and $A \subseteq M, N$ is a common substructure (where $|A| < |M|$). It's enough to show that we have an isomorphism from M to N fixing A.
By Proposition 5.9, we have an isomorphism f from B^M to B^N preserving A (by saturation and the choice of language, the condition of the proposition are satisfied). On $P^M \setminus \left(B^M \cup A\right)$ there is no structure and it has the same size as $P^N \setminus \left(B^N \cup A\right)$ (namely $|N|$), so we can extend the isomorphism f to cover P^M.
We are left with covering R^M: let $a \in R^M$, and $a_i = \pi_i(a)$ for $i = 1, 2$. Then, we already know $f(a_1), f(a_2)$. Suppose $a_1 \subseteq_n a_2$ for minimal n. Then there are exactly n elements in $z \in M$ with $\pi_1(z) = a_1, \pi_2(z) = a_2$, also in N, and the number of such z-s not in A is the same for both M, N. Hence we can take this E-equivalence class from M to the appropriate class in N.
If not, i.e. $a_1 \nsubseteq_n a_2$ for all n, then there are infinitely many elements z in N and in M with $\pi_1(z) = a_1$, $\pi_2(z) = a_2$, and again we take this E-class in M outside of A to the appropriate E-class in N. □

Now we can conclude by a counting types argument. Let M be a countable model of T. Let $p(x)$ be a non-algebraic type over M. There are some cases:

Case 1. $S_n(x) \in p$ for some n. Then the type is determined by the maximal element c in M such that $c \subseteq_n x$ (this is easy, but also follows from Proposition 5.9).

Case 2. $S_n(x) \notin p$ for all n but $P(x) \in p$. Then x is already determined – there is nothing more we can say on x.

Case 3. $R(x) \in p$. Then the type of x is determined by the type of $(\pi_1(x), \pi_2(x))$ over M.

So the number of types over M is countable. □

Proposition 5.11. *A generic pair of models of* T *has the independence property.*

Proof. Suppose (M, M_1) is a generic pair (see definition 5.1). We think of it as a structure of the language L_Q, where Q is interpreted as M. Consider the formula

$$\varphi(x, y) = P(x) \wedge P(y) \wedge \exists z \notin Q \left(Q_1(x, z) \wedge Q_2(y, z)\right).$$

This formula has IP: Let $\{a_i | i < \omega\} \subseteq M$ be elements from P^M such that $a \in S_1^M$ (as in the language of the proof of 5.10), i.e. they are atoms in the induced PBA, and $a_i \neq a_j$ for $i \neq j$. For any finite $s \subseteq \omega$ of size n, there is an element $b_s \in P^M$ be such that $a_i \subseteq_n^M b_s$ for all $i \in s$. Then for all $i \in \omega$, $\varphi(a_i, b_s)$ iff $i \notin s$:

If $\varphi(a_i, b_s)$ there are infinitely many z-s in M such that $Q_1(a_i, z) \wedge Q_2(b_s, z)$ (otherwise they would all be in M). This means that $a_i \not\subseteq_n^M b_s$ – a contradiction. On the other direction, the same exact argument works, but this time use the fact that the pair is generic. □

References

[1] Moran Cohen and Saharon Shelah. Stable theories and representation over sets, 2009. http://www.citebase.org/abstract?id=oai:arXiv.org:0906.3050.

[2] Ehud Hrushovski and Anand Pillay. On nip and invariant measures. *preprint*.

[3] Michael Morley. Categoricity in power. *Transaction of the American Mathematical Society*, 114:514–538, 1965.

[4] Michel Parigot. Théories d'arbres. *J. Symbolic Logic*, 47(4):841–853 (1983), 1982.

[5] Saharon Shelah. A dependent dream and recounting types. *preprint*. 950 in Shelah archive.

[6] Saharon Shelah. Dependent t and existence of limit models. *preprint*. 877 in Shelah archive.

[7] Saharon Shelah. Dependent theories and the generic pair conjecture. *Communications in Contemporary Mathematics*, submitted. math.LO/0702292.

[8] Saharon Shelah. No limit model in inaccessibles. *preprint*. 906 in Shelah archive.

[9] Saharon Shelah. Strongly dependent theories. *submitted*.

[10] Saharon Shelah. *Classification theory and the number of nonisomorphic models*, volume 92 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, second edition, 1990.