Saturated models

Models realizing many types

Throughout these slides, T will be a complete theory in a countable language which has infinite models.

By the Compactness Theorem, any model of T has an elementary extension that realizes all types.

One expects such an extension to behave like a "completion" or "compactification" of the original model.

Defn. Call a model $\mathbb{S}$ of T *weakly saturated* if it realizes all types in $S_n(T)$ for all n.

Models realizing many types

The definition of "weakly saturated model" seems dual to the definition of atomic model, so in an ideal world, the following would be true:

1. Countable weakly saturated models of T would exist.
2. Any two would be isomorphic.
3. Any countable model of T would embed elementarily into the weakly saturated model.
4. Two tuples in a weakly saturated model would have the same type iff they differed by an automorphism.

But all of these statements are false.
The first statement becomes true provided $|S_n(T)| < 2^{\aleph_0}$ for all n. And then all statements become true with "ω-saturated" in place of "weakly saturated".

The Ehrenfeucht theory

Example. Let T be the theory of dense linear order without endpoints expanded by a strictly increasing ω-chain of constants.

1. Signature involves $<, c_0, c_1, \ldots$ only.
2. Axioms for T =
   (i) Axioms of dense linear orders without endpoints.
   (ii) $c_i < c_{i+1}$ for each i.
3. Theory has q.e. and is complete by modhw3p3 (Eblen, Murali, Ornstein).
4. ($I(T, \omega) = 3$.) Theory has three isomorphism types of countable models. Any countable model is isomorphic to one of the form $\langle \mathbb{Q}; <, c_0, c_1, \ldots \rangle$ where
   (i) (Model $\mathbb{M}_1$) The sequence $(c_i)_{i \in \omega}$ is unbounded.
   (ii) (Model $\mathbb{M}_2$) The sequence $(c_i)_{i \in \omega}$ has a least upper bound in the model.
   (iii) (Model $\mathbb{M}_3$) The sequence $(c_i)_{i \in \omega}$ has an upper bound in the model but has no least upper bound in the model.

The countable models $\mathbb{M}_1, \mathbb{M}_2, \mathbb{M}_3$

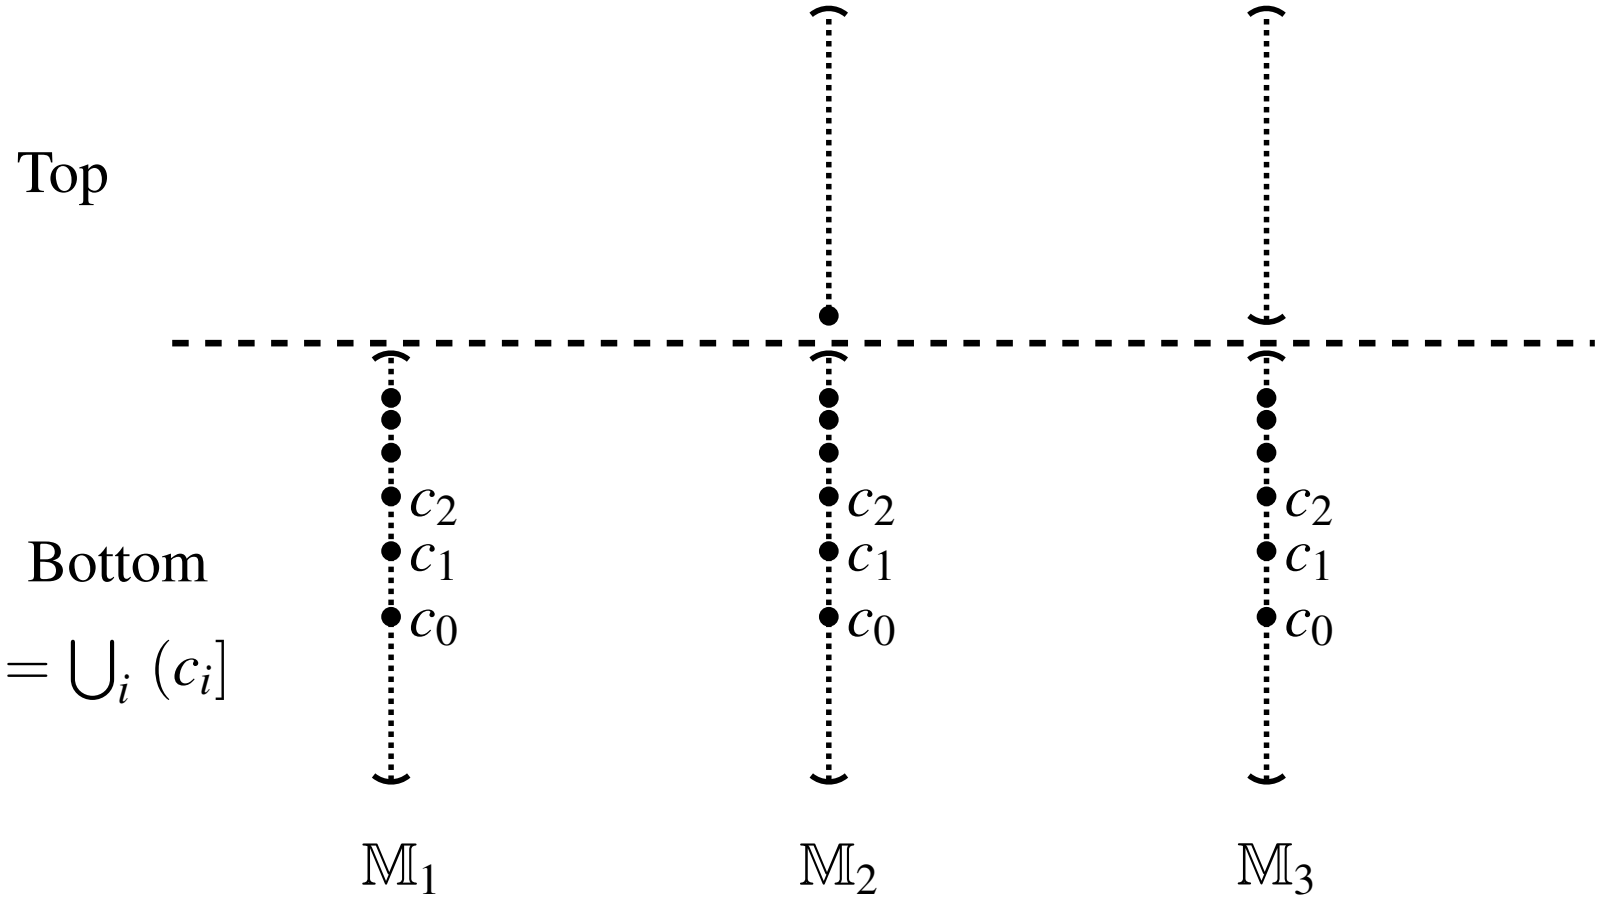

Figure: $\mathrm{Top}(x) = \{c_0 < x, \quad c_1 < x, \quad c_2 < x, \ \ldots\}$, nonisolated $p \in S_1(T)$

Observations

1. The fact that $I(T,\omega) = 3$ can be checked by noting that a countable model is determined up to isomorphism by the part above all the constants, and that part is a (possibly empty) dense linear order without top element.
2. All embeddings between models are elementary by q.e. $\mathbb{M}_1 \prec \mathbb{M}_2 \prec \mathbb{M}_3 \prec \mathbb{M}_2$.
3. $I(T,\omega) < 2^{\aleph_0}$ implies $S_n(T)$ is scattered for all n, so one of the models must be atomic. The only plausible candidate is $\mathbb{M}_1$.
4. All countable models embed elementarily into both $\mathbb{M}_2$ and $\mathbb{M}_3$. This is enough to prove that $\mathbb{M}_2$ and $\mathbb{M}_3$ are both weakly saturated.
5. The model $\mathbb{M}_2$ does not have the type-extension property. Let $p \in S_1(T)$ be the type $p(x_1) = \mathrm{Top}(x_1)$. Let $q \in S_2(T)$ be the type $q(x_1,x_2) = \mathrm{Top}(x_1) \cup \mathrm{Top}(x_2) \cup \{x_2 < x_1\}$. $q|_1 = p$. Let $a = \mathrm{lub}(c_i)$. The 1-tuple (a) realizes p, Some 1-tuples that realize p can be extended to 2-tuples that realize q. But the 1-tuple (a) cannot be extended to a 2-tuple that realizes q.

Tweaking the example by coloring the points

We introduce two new unary relations, $\mathrm{red}(x)$ and $\mathrm{blue}(x)$. Our goal is to construct a theory like Ehrenfeucht's, but with every point colored either red or blue, but not both colors.

1. Signature involves $<, \mathrm{red}(x), \mathrm{blue}(x), c_0, c_1, \ldots$ only.
2. Axioms for T =
   - (i) Axioms of linear orders without endpoints.
   - (ii) An axiom saying that each point has a unique color:

$$(\forall x)((\mathrm{red}(x) \wedge \neg\mathrm{blue}(x)) \vee (\neg\mathrm{red}(x) \wedge \mathrm{blue}(x))).$$

   - (iii) Both red points and blue points are dense:

$$(\forall w)(\forall x)((w < x) \rightarrow (\exists y)(\exists z)(\mathrm{red}(y) \wedge \mathrm{blue}(z) \wedge (w < y < x) \wedge (w < z < x))).$$

   - (iv) $c_i < c_{i+1}$ for each i.
   - (v) $\mathrm{red}(c_i)$ for each i.

This theory has four countable models

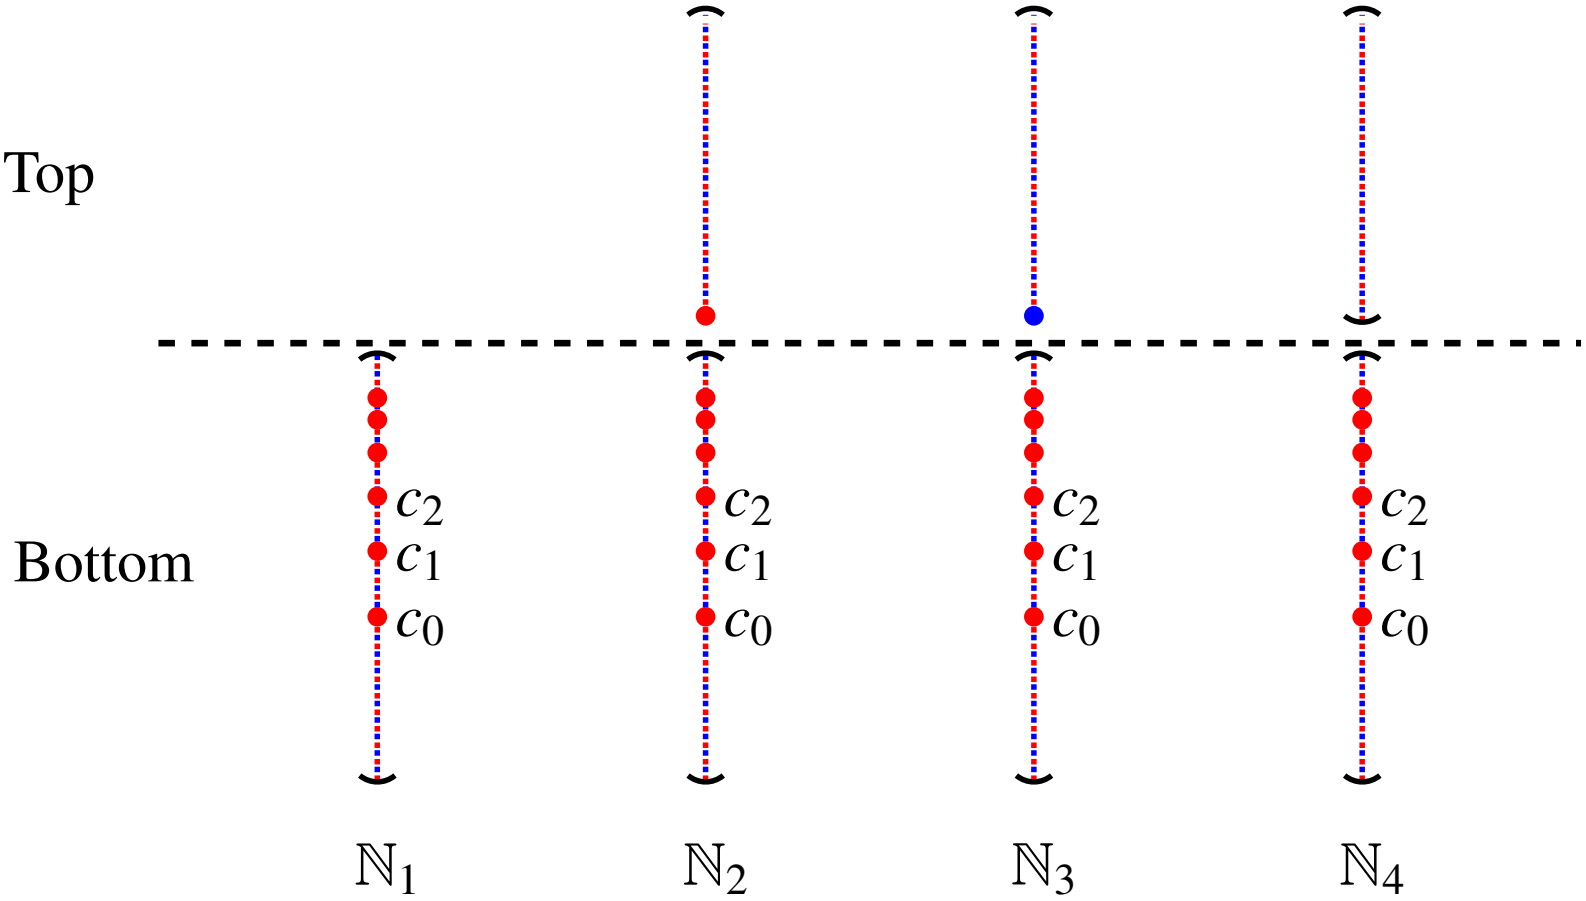

Figure: $\mathbb{N}_1 \prec \mathbb{N}_2 \prec \mathbb{N}_3 \prec \mathbb{N}_4 \prec \mathbb{N}_2$

More observations about the uncolored version

1. $\mathbb{M}_2$ (the model where $\mathrm{lub}(c_i)$ exists in the model) does not have the type-extension property. The problem involves the bound $a = \mathrm{lub}(c_i)$.
2. If $a = \mathrm{lub}(c_i)$, then $p = \{c_i < x < a \mid i \in \omega\}$ is a 1-type in $L(\{a\})$, which is not realized in $(\mathbb{M}_2)_a$. Thus, $\mathbb{M}_2$ is weakly saturated, while an expansion by a single constant is no longer weakly saturated.
3. All upper bounds of the sequence $(c_i)_{i\in\omega}$ have the same 1-type over the empty set (namely $\mathrm{Top}(x)$). But $a = \mathrm{lub}(c_i)$ does not differ from other realizations of $\mathrm{Top}(x)$ by an automorphism.
4. On the other hand, $\mathbb{M}_3$ does have the type-extension property, any expansion of $\mathbb{M}_3$ by finitely many constants is again weakly saturated, and any two tuples of the same type in $\mathbb{M}_3$ differ by an automorphism. $\mathbb{M}_3$ is ω-saturated.

Isomorphism

Let $\mathbb{A}$ and $\mathbb{B}$ be countable structures both enumerated by ω:
$\mathbb{A} = \{a_0, a_1, a_2, \ldots\}$
$\mathbb{B} = \{b_0, b_1, b_2, \ldots\}$
The assignment $a_i \mapsto b_i$ is an isomorphism iff it is type-preserving:

$$\mathrm{tp}(a_0 \cdots a_{n-1}) = \mathrm{tp}(b_0 \cdots b_{n-1}) \tag{1}$$

for all n.

Suppose we want to build an isomorphism one element at a time, by ensuring that, given equality of types of length-n initial segments $\mathbf{a}, \mathbf{b}$, as in (1), and given the choice for a_n, we can find a corresponding choice for b_n. If we work only with types over the empty set, then we need some form of the type-extension lemma. It is enough to assume $\mathbb{A}$ and $\mathbb{B}$ are weakly saturated PLUS any two tuples of the same type differ by an automorphism. OR, we can work with $\mathbb{A}_{\mathbf{a}}$ and $\mathbb{B}_{\mathbf{b}}$ and then deal only with types in the expanded language $L(\mathbf{a})$.

ω-saturation

Defn. Let T be a complete theory.

1. A model $\mathbb{M}$ of T is *ω-saturated* if, whenever $\mathbf{a} \in M^n$, $\mathbb{M}_{\mathbf{a}}$ realizes every type in $S_1(\mathbf{a})$. Often written "whenever $A \subseteq \mathbb{M}$, $|A| < \omega$, $\mathbb{M}_A$ realizes every type in $S_1(A)$". [1]
2. (Type extension) A model $\mathbb{M}$ of T is *ω-homogeneous* if, whenever $\mathbf{a}, \mathbf{b} \in M^n$, $\mathrm{tp}(\mathbf{a}) = \mathrm{tp}(\mathbf{b})$, and $c \in M$, then there exists $d \in M$ such that $\mathrm{tp}(\mathbf{a}c) = \mathrm{tp}(\mathbf{b}d)$.
3. A model $\mathbb{M}$ of T is *strongly ω-homogeneous* if, whenever $\mathbf{a}, \mathbf{b} \in M^n$, $\mathrm{tp}(\mathbf{a}) = \mathrm{tp}(\mathbf{b})$, then there is an automorphism α of $\mathbb{M}$ such that $\alpha(\mathbf{a}) = \mathbf{b}$.
4. A model $\mathbb{M}$ of T is *ω^+-universal* every countable model of T is elementarily embeddable in $\mathbb{M}$. (In particular, an ω^+-universal model will be weakly saturated.)

[1] Equivalently, $\mathbb{M}_A$ realizes every type in $S_n(A)$ for each finite n, Proposition 4.3.2, Marker.

Relationships

Theorem. Let T be a complete theory in a countable language. TFAE about a countable model $\mathbb{M}$ of T.

1. $\mathbb{M}$ is ω-saturated.
2. $\mathbb{M}$ is weakly saturated and ω-homogeneous. ($\mathbb{M}$ realizes all types over the empty set and has the type extension property.)
3. $\mathbb{M}$ is weakly saturated and strongly ω-homogeneous.
4. $\mathbb{M}$ is ω^+-universal and ω-homogeneous.
5. $\mathbb{M}$ is ω^+-universal and strongly ω-homogeneous.

Trivial implications.
ω^+-universality implies weak saturation.
Strong ω-homogeneity implies ω-homogeneity.

Not-too-hard implications.
ω-saturation implies strong ω-homogeneity. (Back and forth.)
ω-saturation implies ω^+-universality. (Forth.)

Weak saturation and ω-homogeneity in E's Theory

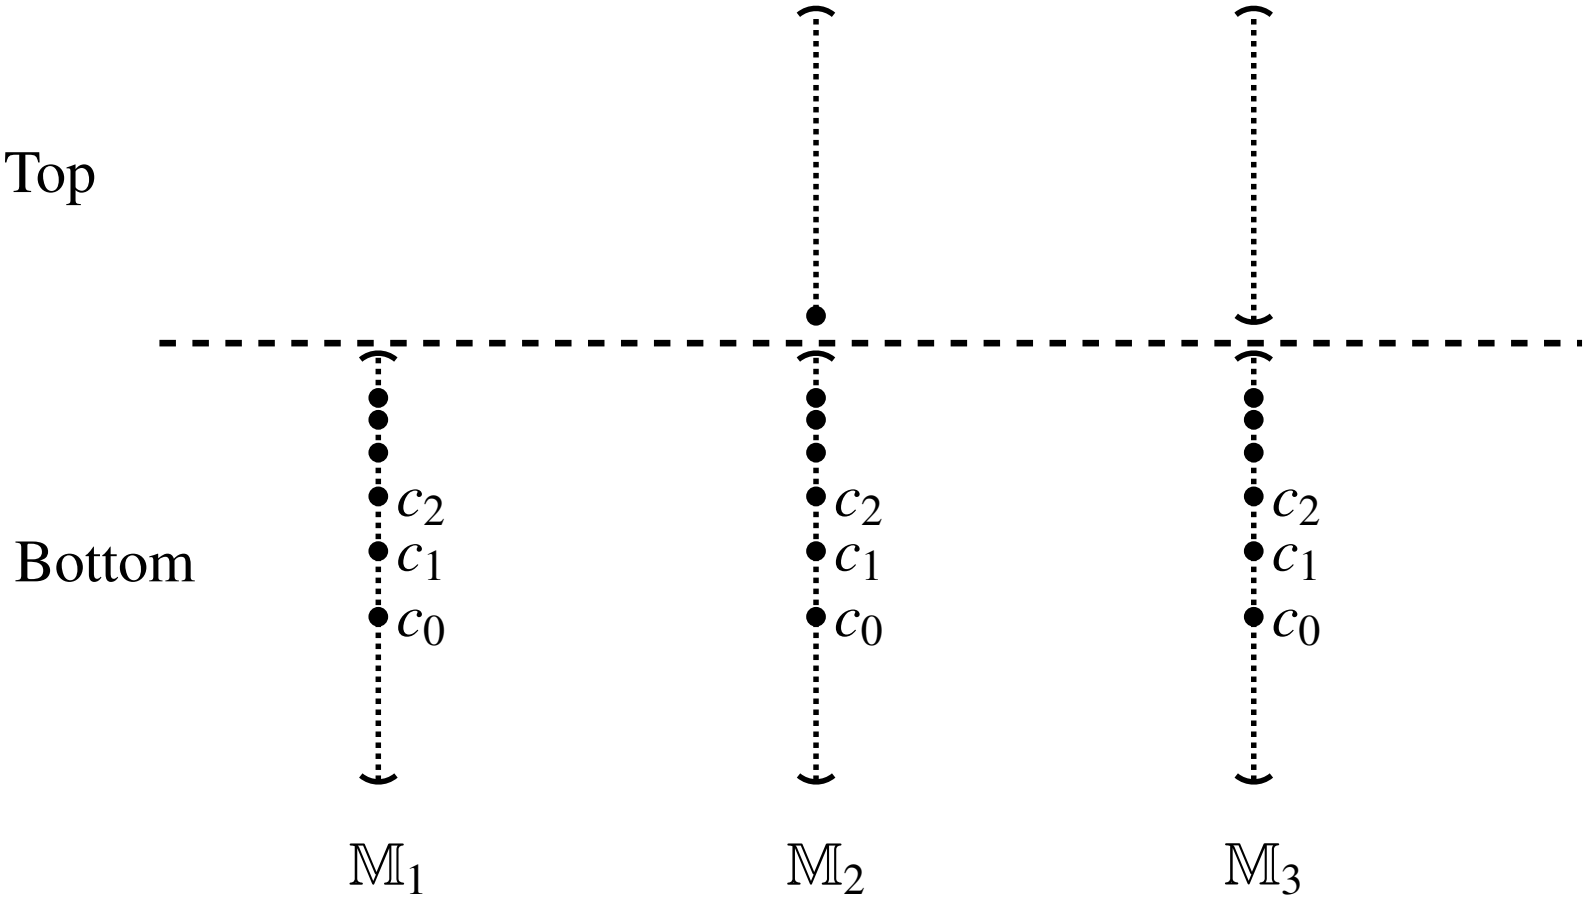

Figure: $\mathbb{M}_2, \mathbb{M}_3$ weakly saturated; $\mathbb{M}_1, \mathbb{M}_3$ ω-homogeneous

A proof sketch

Theorem. Two countable ω-saturated models of T are isomorphic. (Back and forth.)

Assume $\mathbb{A}$ and $\mathbb{B}$ are ω-saturated models of T.

Enumerate them.

Start back and forth: Assume that $f : \mathbf{a} \to \mathbf{b}$ is a partial isomorphism that we want to extend. At this point, $\mathrm{tp}^{\mathbb{A}}(\mathbf{a}) = \mathrm{tp}^{\mathbb{B}}(\mathbf{b})$. Equivalently, $\mathbb{A} \models \varphi(\mathbf{a})$ iff $\mathbb{B} \models \varphi(\mathbf{b})$. Equivalently, $\mathbb{A}_{\mathbf{a}} \equiv \mathbb{B}_{\mathbf{b}}$.

Assume it is our turn to extend the domain. Let $c \in \mathbb{A}$ be the least unconsidered element. Let $p = \mathrm{tp}^{\mathbb{A}_{\mathbf{a}}}(c)$. Let d be a realization of p in $\mathbb{B}_{\mathbf{b}}$. Thus, $\mathbb{A} \models \theta(\mathbf{a}c)$ iff $\mathbb{B} \models \theta(\mathbf{b}d)$. I.e., $\mathrm{tp}^{\mathbb{A}}(\mathbf{a}c) = \mathrm{tp}^{\mathbb{B}}(\mathbf{b}d)$. Extend f so that $f(c) = d$. $\square$

When $\mathbb{A} = \mathbb{B}$, this argument proves strong ω-homogeneity of ω-saturated models. Half of the argument proves ω^+-universality.

Existence and uniqueness

Theorem. Let T be a complete theory in a countable language. If $\mathbb{A}$ and $\mathbb{B}$ are countable ω-saturated models of T, then $\mathbb{A} \cong \mathbb{B}$.

Theorem. Let T be a complete theory in a countable language. TFAE.

1. T has a countable ω-saturated model.
2. T has a countable weakly saturated model.
3. T is "small". ($|S_n(T)| < 2^{\aleph_0}$ for all n.)

Part of proof.
(1) $\Rightarrow$ (2) $\Rightarrow$ (3) uses only ideas we have seen.
(3) implies that, for any model $\mathbb{M}$ of T, and any finite subset $A \subseteq \mathbb{M}$, ($|A| = m$, say), then $|S_n^{\mathbb{M}}(A)| \leq |S_{m+n}(T)| \leq \omega$.
Idea for the rest. Let $\mathbb{M}_0 = \mathbb{M}$. Find a countable elementary extension $\mathbb{M}_{i+1} \succ \mathbb{M}_i$ that realizes the countable set of 1-types over finite subsets of $\mathbb{M}_i$. Let $\widehat{\mathbb{M}}$ be the union of the $\mathbb{M}_i$. $\square$

Extensions to higher cardinalities

Defn. A model $\mathbb{M}$ of T is *κ-saturated* if whenever $A \subseteq M$ satisfies $|A| < \kappa$, then $\mathbb{M}_A$ realizes all $p \in S_1^{\mathbb{M}}(A)$. We say that $\mathbb{M}$ is *saturated* if it is $|\mathbb{M}|$-*saturated*.

To discuss this when $\kappa \neq \omega$, we need a concept of type for infinitely long tuples.

Some basic results.

1. κ-saturated = κ^+-universal and κ-homogeneous.
2. Formation of ultrapowers increases saturation.
3. An infinite model $\mathbb{M}$ satisfying $|\mathbb{M}| \leq 2^\kappa$ has a κ^+-saturated elementary extension of cardinality 2^κ.
4. There is a model $\mathbb{X}$ of ZFC, a theory T, and a model $\mathbb{M}$ of T such that $\mathbb{M}$ has no saturated elementary extension in $\mathbb{X}$.

Saturated models of ACF_p

Theorem. A model $\mathbb{K}$ of ACF_p is saturated if(f) it contains an algebraically independent subset of size $|\mathbb{K}|$.

Proof sketch. Let's explain why $\mathbb{C}$ is saturated.
Let $A \subseteq \mathbb{C}$ satisfy $|A| < |\mathbb{C}|$.
Let $\mathbb{F}$ be an algebraically closed subfield of $\mathbb{C}$ containing A and satisfying $|\mathbb{F}| < |\mathbb{C}|$.
By q.e., types over A are determined by $\pm$atomic part.
Any type over A with a "$+$atomic" part ($p(x, \mathbf{a}) = 0$) is realized in $\mathbb{F}$, hence in $\mathbb{C}$.
Any complete type over A with only "$-$atomic" describes an element transcendental over $\mathbb{F}$. $\mathbb{C}$ has such an element. $\square$.

Appendix: Ultrapowers are somewhat saturated

Let $\mathbb{A}$ be a structure and $B \subseteq \mathbb{A}$ be a subset. Let $\kappa = \|L(B)\|$. Let $I = \mathcal{P}_{\text{fin}}(\kappa)$ be the set of finite subsets of κ.

$I = \mathcal{P}_{\text{fin}}(\kappa)$ is directed by inclusion. The tail ends of this directed set form a proper filter on I, which can be extended to an ultrafilter $\mathcal{U}$ on I. Let's outline why $\prod_{\mathcal{U}} \mathbb{A}$ realizes every type in $S_1^{\mathbb{A}}(B)$.

Accept for now that every $p \in S_1^{\mathbb{A}}(B)$ has cardinality $\kappa = \|L(B)\|$, and choose a bijection $\beta_p : \kappa \to p$. There is an induced bijection from $I = \mathcal{P}_{\text{fin}}(\kappa)$ to $\mathcal{P}_{\text{fin}}(p)$, which we also call β_p. Thus, for each $i \in I$, there is assigned a set $\beta_p(i)$, which is a finite subset of p.

Since p is consistent with $\text{Th}(\mathbb{A}_B)$, for each i there is an element $a_i \in \mathbb{A}$ that satisfies all formulas in the finite set $\beta_p(i)$. Let $\mathbf{a} \in \mathbb{A}^I$ be the tuple satisfying $(\mathbf{a})_i = a_i$ for all i. For each $\varphi(x) \in p$ we have that $[\![\varphi[\mathbf{a}]]\!]$ contains the tail end in $I = \mathcal{P}_{\text{fin}}(\kappa)$ generated by $\beta_p^{-1}(\varphi(x))$. This tail end belongs to $\mathcal{U}$, hence $\prod_{\mathcal{U}} \mathbb{A} \models \varphi[\mathbf{a}]$. This is true for any $\varphi(x) \in p$, so $\mathbf{a}$ realizes p in $\prod_{\mathcal{U}} \mathbb{A}$.

Similarly, every $q \in S_1^{\mathbb{A}}(B)$ is realized in the same ultrapower. It is worth recording that $|\prod_{\mathcal{U}} \mathbb{A}| = |\mathbb{A}|^{\|L(B)\|}$.

Appendix to the appendix

On the previous slide it was claimed that all types in $S_1^{\mathbb{A}}(B)$ have the same cardinality, namely $\kappa = \|L(B)\|$. This fact was used so that we could correspond finite subsets of any $p \in S_1^{\mathbb{A}}(B)$ to finite subsets of the fixed set κ. This was needed so that a fixed ultrapower was able to realize all types in $S_1^{\mathbb{A}}(B)$ simultaneously.

One can prove that all types in $S_1^{\mathbb{A}}(B)$ have size $\kappa = \|L(B)\|$ as follows.

1. If $p \in S_1^{\mathbb{A}}(B)$, then $p \subseteq L(B)$, so $|p| \leq \|L(B)\| = \kappa$.
2. The existence of the map $\varphi(\mathbf{x}) \mapsto (\forall x_1) \cdots (\forall x_k)\varphi(\mathbf{x})$, which assigns to a formula its universal closure, is a finite-to-one map from $L(B)$ to a subset of the $L(B)$-sentences. This establishes that the set of $L(B)$-sentences has size at least κ. (I am basing this claim on the fact that if X and Y are infinite and there is a finite-to-one map from X into Y, then $|X| \leq |Y|$.)
3. Any $p \in S_1^{\mathbb{A}}(B)$ contains half of the $L(B)$-sentences, hence p has cardinality at least $\frac{1}{2}\kappa = \kappa$.