[*Held in 120../A000-wholething.. Last changed July 26, 2004*]

An introduction to
Good old fashioned model theory

[*Held in 120../Preamble.. Last changed July 26, 2004*]

The course will be taught

Each Tuesday from 2 – 5 pm in room J

with some lectures and some tutorial activities.

The main body of course will cover the following topics.

- Basic ideas of language, satisfation, and compactness
- Some examples of elimination of quantifiers
- The diagram technique, the characterization of $\forall_1$- and $\forall_2$-axiomatizable theories, and similar results
- Model complete theories, companion theories, existentially closed structures, and various refinements
- Atomic structures and sufficiently saturated structures
- The splitting technique and $\aleph_0$-categoricity

Depending on the time available, some of the following topics will be looked at.

- Henkin constructions
- Omiting types
- The back-and-forth technique
- Saturation methods.

Full course notes will be provided of which this is the first part.

A look at the contents page will indiacte how these stand at the moment. These notes will be modified as the course progresses.The main body of the course is contained in Part I. The extra topics are contained in Part II.

Remarks and comments *in this kind of type* indicate that something needs to be done before the final version is produced.

Contents

For those of you do not have the advantages of a well rounded education, I attach a list of the upper case gothic letters used togther with the roman equivalent.

$$\mathfrak{A} \quad \mathfrak{B} \quad \mathfrak{C} \quad \mathfrak{D} \quad \mathfrak{E} \quad \mathfrak{M} \quad \mathfrak{N} \quad \mathfrak{P} \quad \mathfrak{T} \quad \mathfrak{U} \quad \mathfrak{Z}$$

$$A \quad B \quad C \quad D \quad E \quad M \quad N \quad P \quad T \quad U \quad Z$$

Tick off each one as you meet it. There is a **very** small prize for the first person with a full set (to be exchanged in the the Welsh village of Llarygrub).

Introduction–*to be done*

0.1 Historical account–*to be done*

0.2 A survey of the literature–*to be re-done*

[*Held in 120../Survey.. Last changed July 26, 2004*]

It seems there are few modern text that fit this introductory course. Most text are either too old (and out of print), or cover far too much for a first course, or are on more specialized topics.

The three articles [2], [8], and [11] of the handbook [1] form a nice introduction to the subject. Most of the other chapters in the model theory part are also worth reading at some stage. (If you want to read any of [1], get a friend to help you pick it up. If you want to buy [1], you will have to form a commune.)

The best introduction to model theory used to be [4]. Unfortunately I have lost my copy, and can not locate another. From memory, I think it doesn't cover enough for the whole course, but it is still worth reading if you can get hold of it.

The book [3] is a bit dated, but still a good introduction. If you rip out chapters 1, 5, 6, 8, 12, 13, and 14 you will have a neat little text which covers much of the course. The book [5] is from the same period, but is much more comprehensive. It suffers from the same defect, too much stuff on topics that are no longer central to the subject. Nevertheless, it is still worth reading.

The book [13] was written by someone who wanted to learn some model theory, and because of this it is a nice introductory text. It also contains a discussion of stability theory (as it then stood), which is something not covered in this course but should looked at eventually.

The best modern book on model theory is [6], but this is not for the beginner. It contains almost everything you need to know about model theory, much more than a first course can contain. Unfortunately it is written in an encyclopedic rather than a progressive form, and this makes it a bit hard for a beginner to find what he needs. Nevertheless, you should aim to become familiar with the contents of this book.

need some chat about [7]

Another book worth reading is [12]. This is an excellent introduction to many part of mathematical logic suitable for a postgraduate. It is a bit eccentric in part, not least that it is written in French. However, it is still worth reading. It is being translated into English. [*This has now been translated and published -- by Springer -- so get your copies while they last*]

References

[1] J. Barwise (ed): *Handbook of Mathmatical Logic*, (North-Holland, 1977).

[2] J. Barwise: An introduction to first-order logic, Chapter A.1 of [1].

[3] J. L. Bell and A. B. Slomson: Models and Ultraproducts, (North-Holland, 1969).

[4] J. Bridge: ***Find a copy.***

[5] C. C. Chang and H. J. Keisler: *Model theory*, (North-Holland, originally published in 1973 with a third edition in 1990).

[6] W. Hodges: *Model Theory*, (C. U. P. 1993).

[7] W. Hodges: *A shorter model theory*, (C. U. P. 199*).

[8] H. J. Keisler: Fundamentals of model theory: Chapter A.2 of [1].

[9] H. J. Keisler: *Model theory for infinitary languages*, (North-Holland, 1971).

[10] G. Kreisel and J. L. Krivine: *Elements of mathematical logic*, (North-Holland, 1967).

[11] A. Macintyre: Model completeness, Chapter A.4 of [1].

[12] B. Poizat: *Course de Théorie des Modèles* (B. Poizat, 1985).

[13] G. E. Sacks: *Saturated model theory* (W. A. Benjamin, 1972).

[14] C. Smoryński: *Logical number theory 1*, (Springer, 1991).

Part I

The development

This part contains the main body of the course as outlined in the preamble. Of course, the lectures will not cover everything in these notes, but you should try to get some familiarity with the whole of this material.

[*Held in 120-../B10-bit.. Last changed July 26, 2004*]

1 Syntax and semantics

Model theory (and, in fact, much of mathematical logic) is concerned, in part, with the use of syntactic gadgetry (languages) to extract information about the objects under investigation. To do this efficiently we must have (at least at the back of our mind) a precise definition of the underlying language and the associated facilities. Setting up such a definition can look complicated, and is certainly tedious. However, the ideas involved are essentially trivial. The role of the definition is merely to delimit what can, and what can't, be done in a first order or elementary language.

In the next subsection we look at the details of this definition, and in the rest of this section we look at some of the consequential ramifications. Before we do that let's motivate the ideas behind the constructions.

A structure $\mathfrak{A}$ (of the kind we deal with in model theory) is a non-empty set A, called the carrier of the structure, furnished with some distinguished attributes each of which is an element of A, a relation on A, or an operation on A. Many objects used in mathematics are structures in this sense, but many are not. The crucial restriction here is that the structure is single sorted with just one carrier, this carrier is non-empty, and all the furnishings are first order over the carrier. In the wider scheme of things these are serious restrictions.

Often we are concerned with a whole family of structures each of the same similarity type or signature. For such a family there is a common language which can be used to talk about any and all of these structures. This language is generated in a uniform way from the signature.

This language L allows the use of the connectives

not and or implies

(and perhaps others of the same ilk). It also allows some quantification. However, in any instance of a quantification the bound variable can only, and must, range over the carrier (of the structure being talked about). The language allows quantification over elements of the carrier, but *not* over subsets of, lists from, or any other kind of gadget extracted from the carrier. This is the first order restriction.

In more advanced work other kinds of languages are used, but not here. At this level model theory is about the use of first order languages, and the exploitation of a central and distinguishing result, the compactness theorem.

[*Held in 120../B11-bit.. Last changed July 26, 2004*]

1.1 Signature and language

In this subsection we make precise the ideas outlined above. This will take several steps but, as explained, there is nothing very complicated going on.

1.1 DEFINITION. A signature is an indexed family of symbols where each is either

- a constant symbol K

- a relation symbol R with a nominated arity
- an operation symbol O with a nominated arity

respectively. Each arity is a non-zero natural number. We speak of an n-placed relation symbol or an n-placed operation symbol to indicate the the arity is n. ■

The word 'symbol' here indicates that eventually we will generate a formal language L of certain strings. The letters 'K', 'R', and 'O' have been chosen in a rather awkward manner to remind us that they are syntactic symbols.

Before we start to generate the full language L let's take a quick look at the kind of gadget that will provide the semantics for L.

1.2 DEFINITION. A structure $\mathfrak{A}$ for a given signature consists of the following.

- A non-empty set A, the carrier of $\mathfrak{A}$.
- For each constant symbol K, a nominated element $\mathfrak{A}[\![K]\!]$ of A.
- For each n-placed relation symbol R, a nominated n-placed relation $\mathfrak{A}[\![R]\!]$ on A.
- For each n-placed operation symbol O, a nominated n-placed operation $\mathfrak{A}[\![O]\!]$ on A.

These nominated gadgets are the distinguished attributes of $\mathfrak{A}$. ■

You should differentiate between the symbols K, R, and O of the signature and the interpretation $\mathfrak{A}[\![K]\!], \mathfrak{A}[\![R]\!]$, and $\mathfrak{A}[\![O]\!]$ of these in the particular structure $\mathfrak{A}$. There is only one language L of the given signature, but there are many different structures of that signature, each of which provides an interpretation of each symbol of the signature.

There are times in model theory when we need to take note of the size of a structure (or a language).

1.3 DEFINITION. The cardinality $|\mathfrak{A}|$ of a structure is the cardinality $|A|$ of it carrier. ■

Notice that a signature can be empty, in which case a structure (for that signature) is just a non-empty set. At the other extreme, a signature can be uncountably large. On the whole, in these notes we are concerned with signatures that have no more than countably many relation symbols and operation symbols. However, for technical reasons, it is convenient to allow the number of constant symbols to be arbitrarily large.

1.4 DEFINITION. Given a signature, the associated primitive symbols are as follows.

- The symbols of the signature.
- The equality symbol $\doteq$.
- An unlimited stock of variables v.
- The connectives $\neg, \wedge, \vee, \rightarrow$
- The quantifiers $\forall$ and $\exists$.

- The constant sentences which are true and false.
- The punctuation symbols (and).

There are no other primitive symbols.

A string is a finite list of primitive symbols. ■

In other words, the primitive symbols of the language L are the symbols of the signature together with a fixed collection of other symbols. These extra symbols are the same for each language.

Again, you should not confuse the formal symbol '$\doteq$' (which is a primitive of each language) with the informal symbol '$=$' used to indicate the equality of two gadgets.

A string is any finite list of primitive symbols, and can be complete gibberish

$$))\neg v \doteq (\wedge w) \rightarrow \exists$$

or can be potentially meaningful

$$(\forall v)(\exists w)((fvw \doteq v) \wedge \neg(fwv \doteq w))$$

(where v, w are variables and f is a 2-placed operation symbol). Our aim is to extract the potentially meaningful strings.

We do that in three steps. We define the terms t, the atomic formulas θ, and then the formulas ϕ. The formulas are the potentially meaningful strings. Each of these strings

$$t \qquad \theta \qquad \phi$$

has an associated support

$$\partial t \qquad \partial\theta \qquad \partial\phi$$

giving the set of variables occurring freely in the string. The support is generated at the same time as the parent string.

1.5 DEFINITION. Each signature has an associated set of terms and each such term t has an associated set ∂t of free variables. These are generated as follows.

- Each variable v is a term, and $\partial v = \{v\}$.
- Each constant symbol K is a term, and $\partial K = \emptyset$.
- For each n-placed operation symbol O and each list $t_1, \ldots, t_n$ of terms, the compound

  $$(Ot_1 \cdots t_n)$$

  is a term and

  $$\partial(Ot_1, \cdots t_n) = \partial t_1 \cup \cdots \cup \partial t_n$$

  is its set of free variables.

There are no other terms. ■

Notice that each term can be uniquely parsed (so that its construction can be displayed as a finite tree). That is the job of the brackets in the construction.

1.6 DEFINITION. For a signature the **atomic formulas** are those strings

$$\mathsf{true} \qquad \mathsf{false} \qquad (t_1 \doteq t_2) \qquad (Rt_1 \cdots t_n)$$

where $t_1, t_2, \ldots t_n$ are terms and R is an n-placed relation symbol. Each such atomic formula θ has a set $\partial\theta$ of free variables given by

$$\emptyset \qquad \emptyset \qquad \partial t_1 \cup \partial t_2 \qquad \partial t_1 \cup \cdots \cup \partial t_n$$

respectively. ■

Notice the difference between

$$(t_1 = t_2) \qquad (t_1 \doteq t_2)$$

where t_1, t_2 are terms. The first asserts that the two terms are the same (that is, the same string of primitive symbols), whereas the second is an atomic formula of the language which, in isolation, has no truth value.

Finally we can extract the potentially meaningful strings.

1.7 DEFINITION. Each signature has an associated set of **formulas** and each such formula ϕ has an associated set $\partial\phi$ of **free variables**. These are generated as follows.

- Each atomic formula is a formula with free variables, as given.
- For each formula ψ the string $\neg\psi$ is a formula with $\partial\neg\psi = \partial\psi$.
- For each pair θ, ψ of formulas, each of the strings

  $$(\theta \wedge \psi) \qquad (\theta \vee \psi) \qquad (\theta \rightarrow \psi)$$

  is a formula with free variables

  $$\partial\theta \cup \partial\psi$$

  in each case.
- For each formula ψ and variable v, each of the strings

  $$(\forall v)\psi \qquad (\exists v)\psi$$

  is a formula with free variables

  $$\partial\psi - \{v\}$$

  in both cases.

There are no other formulas. ■

Notice that once again each formula can be uniquely parsed.

The **language L given by a signature** is the set of all formulas associated with that signature. In practice we usually don't even mention the signature, but use phrases such as

structure suitable for the language L
term of the language L
formula of the language L

and so on.

In a way the formulas are not the most important strings associated with a signature.

1.8 DEFINITION. A sentence of a language is a formula σ of that language with no free variables, $\partial\sigma = \emptyset$. ■

Sentences are those strings which are either true or false in any particular structure. Formulas are really just a stepping stone in the construction of sentences.

It may not be clear what role the two contant sentences true and false play.

At times it is convenient to have a sentence which is trivially valid in every structure we meet and, as a string, is very simple. If the language has a constant K, then $(K \doteq K)$ is such a sentence. However, sometime there isn't a contant around, and then we have to look elsewhere. We could take the sentence $(\forall v)(v \doteq v)$, but that quantifier can be a nuisance. The primitive true is there so that we always have such a trivially true and simple sentence. The primitive false palys a similar role, except this one is a trivially false and simple sentence.

Finally, for this subsection, as indicated above we want to measure the size of a language.

1.9 DEFINITION. The cardinality $|L|$ of a language L is either $\aleph_0$ or the size of the signature, whichever is the larger. ■

The cardinal $|L|$ is infinite. If the signature is finite or countable, then $|L| = \aleph_0$, otherwise it is the cardinality of the signature. For the most part we will be interested in countable languages. However, even for such a language, one of the most common techniques of model theory involves the use of associated languages of larger cardinality. Thus we have to deal with the general case.

Exercises

1.1 Consider the signature with just one symbol $<$ which is a 2-placed relation symbol. We write this as an infix. Let u, v be a pair of distinct variables and set

$$\phi_0 := (v \doteq v) \qquad \theta_r := (\exists v)[(u \doteq v) \wedge \phi_r] \qquad \phi_{r+1} := (\exists u)[(u < v) \wedge \theta_r]$$

for each $r < \omega$ to obtain two ω-chains of formulas.

(a) Write down ϕ_0, ϕ_1, ϕ_2, and perhaps ϕ_3 until you can see what is going on.

(b) Calculate $\partial\phi_r$ and $\partial\theta_r$ for each $r < \omega$.

(c) Describe a different way of setting up equivalent formulas which makes the uses of the variables easier to see.

[*Held in 120../B12-bit.. Last changed July 26, 2004*]

1.2 Basic notions

In subsection we gather together a few more basic notions and some conventions which make the day to day handling of formulas a little less tedious.

We begin with the simplest comparison between two structures.

1.10 DEFINITION. Given a pair $\mathfrak{A}, \mathfrak{B}$ of structures (for the same language), we write

$$\mathfrak{A} \subseteq \mathfrak{B}$$

and say $\mathfrak{A}$ is a substructure of $\mathfrak{B}$ or $\mathfrak{B}$ is a superstructure of $\mathfrak{A}$, if the following hold.

- The carrier A of $\mathfrak{A}$ is a subset of the carrier B of $\mathfrak{B}$.
- For each constant symbol K of the signature, $\mathfrak{A}[\![K]\!] = \mathfrak{B}[\![K]\!]$.
- For each n-placed relation symbol R of the signature, the relation $\mathfrak{A}[\![R]\!]$ is the restriction to A of the relation $\mathfrak{B}[\![R]\!]$. In other words,

$$\mathfrak{A}[\![R]\!]a_1 \cdots a_n \iff \mathfrak{B}[\![R]\!]a_1 \cdots a_n$$

  holds for each $a_1, \dots, a_n \in A$.
- For each n-placed relation symbol O of the signature, the set A is closed under $\mathfrak{B}[\![O]\!]$, and the operation $\mathfrak{A}[\![O]\!]$ is the restriction to A of the operation $\mathfrak{B}[\![O]\!]$. In other words

$$\mathfrak{A}[\![O]\!]a_1 \cdots a_n = \mathfrak{B}[\![O]\!]a_1 \cdots a_n$$

  holds for each $a_1, \dots, a_n \in A$.

In short, A is closed under the attributes of $\mathfrak{B}$, and these give the attributes of $\mathfrak{A}$. ■

Thus, given a structure $\mathfrak{B}$, each substructure $\mathfrak{A}$ is completely determined by its carrier. However, not every subset of the carrier of $\mathfrak{B}$ is the carrier of a substructure.

In subsection 3.1 we will look at a generalization of this idea. We define the notion of an embedding

$$\mathfrak{A} \xrightarrow{\quad f \quad} \mathfrak{B}$$

of a structure into another using a function f between the carriers. When this function is an insertion we obtain $\mathfrak{A} \subseteq \mathfrak{B}$.

Two structures $\mathfrak{A}, \mathfrak{B}$ (of the same language) are isomorphic

$$\mathfrak{A} \cong \mathfrak{B}$$

if there is an isomorphism from one to the other. This is a bijection between the carriers which matches the distinguished attributes. We needn't write down the precise details of this, for it is a particular case of an embedding (which we look at later). However, we will use the notion.

Each language L consists of a set of formulas, some of which are sentences. Often we are interested in particular kinds of formulas, ones of a certain 'complexity'. The most common measure of formulas is by quantifier complexity.

There are various useful classifications of formulas. Let's look at two of these, one of which builds on top of the other.

Let L be an arbitrary language.

- An atom is just an atomic formula, as in Definition 1.6.

- A literal is an atom α or the negation $\neg\alpha$ of an atom.
- A formula δ is quantifier-free if it contains no quantifiers, no uses of $\forall$ or $\exists$.

Each quantifier-free formula δ can be rephrased in one of two useful normal forms.

- The conjunctive normal form in which δ is rephrased as a conjunction

$$D_1 \wedge \cdots \wedge D_m$$

  of disjuncts each of which is a disjunction of literals.

- The disjunctive normal form in which δ is rephrased as a disjunction

$$C_1 \vee \cdots \vee C_m$$

  of conjuncts each of which is a conjunction of literals.

Here 'rephrased' means 'is logically equivalent to'. The standard boolean manipulation of formulas enables us to move from δ to either of the normal forms.

In the same way, using the rules for manipulating quantifiers, we know that each formula can be rephrased in prenex normal form as

$$(Q_1 v_1) \cdots (Q_n v_n)\delta$$

where each Q in the prenex is a quantifier and the matrix δ is quantifier-free. By taking note of the alternations in the prenex we obtain the quantifier hierarchy of formulas.

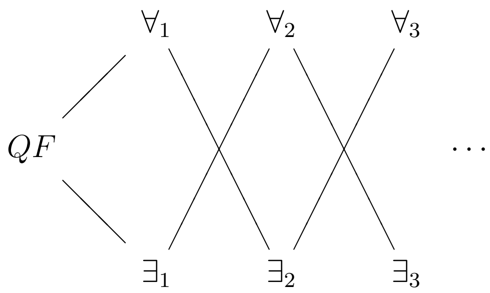

Thus $\forall_0 = \exists_0 = QF$ is the set of formulas each of which is logically equivalent to a quantifier-free formula. For each $n \in \mathbb{N}$ the sets

$$\forall_{n+1} \qquad \exists_{n+1}$$

consists of the formulas logically equivalent to

$$(\forall v_1, \ldots, v_n)\phi \qquad (\exists v_1, \ldots, v_n)\phi$$

where

$$\phi \in \exists_n \qquad \phi \in \forall_n$$

respectively. Inclusions between these sets are indicated in the diagram.

As you can imagine, handling formulas can be a bit tedious especially if we stick strictly to the letter of the law.

When we display a particular formula we often leave out some of the brackets or use different shapes of brackets to make the formula more readable. There are one or two other tricks we sometimes use.

Each formula ϕ has an associated set $\partial\phi$ of free variables. We often write

$$\phi(v_1, \ldots, v_n)$$

to indicate that $\partial\phi = \{v_1, \ldots, v_n\}$. Notice that at this level we are not concerned with the order or number of occurrences of each free variable in ϕ. In fact, a variable cam occur freely and bound in the same formula. Exercise 1.1 gives an extreme example of what can happen.

For much of what we do any finite list of variables can be treated as a single variable. We use some informal conventions to handle this. Thus we often write

$$\phi(v)$$

to indicate that v is a *list*

$$v_1, \ldots, v_n$$

of variables each of which occurs freely in ϕ. Thus, in this usage

$$\phi(v) \qquad \phi(v_1, \ldots, v_n)$$

mean the same thing. There is, of course, plenty of scope for confusion here. However, we will always make this situation clear.

In the same way we sometimes write

$$(\forall v)\phi(v) \quad \text{for} \quad (\forall v_1, \ldots, v_n)\phi(v_1, \ldots, v_n)$$

when the number of variables in the list v is not important.

As well as single formulas we also use sets of formulas. Let Γ be such a set of formulas, and consider

$$\partial\Gamma = \bigcup\{\partial\phi \mid \phi \in \Gamma\}$$

the set of all variables that occur free somewhere in Γ. This could be an infinite set. Often we restrict this support.

1.11 DEFINITION. A **type** is a set Γ of formulas such that $\partial\Gamma$ is finite. ■

Do not confuse this use of the word 'type' with other uses. Some of these have no relationship at all with this usage.

Exercises

1.2 Consider the following three posets.

	$\mathfrak{A}$	$\mathfrak{B}$	$\mathfrak{C}$
carrier	$A = \{a, b\}$	$B = \{a, b, c\}$	$D = \{a, b, c, d\}$
comparisons	$a \le b$	$a \le b \le c$	$d \le a \le c,\ d \le b \le c$

Draw a picture of each and determine those pairs of structures where one is a substructure of the other.

[*Held in 120../B13-bit.. Last changed July 26, 2004*]

1.3 Satisfaction

Having made the effort to set up the notion of the language L suitable for structures $\mathfrak{A}$ of some given signature, and in which the idea of a sentence σ is made precise, it is now patently obvious what

$$\text{the structure } \mathfrak{A} \text{ satisfies } \sigma$$

means. The whole of the rather tedious construction of L was driven with this in mind. Nevertheless, it is worth looking at the formal definition of this notion (not least because some people think that it has some content). In more advanced work various non-elementary languages are used, and then the internal workings of the language and its satisfaction relation are more important.

We are talking about the pivotal notion of model theory, so we can't keep writing it out in words. Accordingly we let

$$\mathfrak{A} \models \sigma$$

abbreviate the phrase above ($\mathfrak{A}$ satisfies σ). This, of course, is a relation between structures $\mathfrak{A}$ and sentences σ, so each instance is either true or false. This truth value is generated by recursion on the construction of σ. That is, we define outright the value for simple sentences, and then show how to obtain the value for compound sentences from the values for its components. This brings out a minor snag.

As we unravel the construction of σ, we will meet certain formulas, and these may contain free variables. But such variables have no interpretation in $\mathfrak{A}$. They are merely a linguistic device to indicate certain bindings in larger formulas. However, to push through the recursion, we need to show what to do with any free variables that arise as we unravel the sentence σ. We use a little trick.

1.12 DEFINITION. For a structure $\mathfrak{A}$, an $\mathfrak{A}$-assignment is a function x which attaches to each variable v an element $v\mathsf{x}$ of $\mathfrak{A}$. ■

(Notice that although we call x a function, we write its argument v on the left.)

The idea is that if we meet a free variable v which, apparently, has no interpretation, then we give it the value $v\mathsf{x}$. Before we see how this helps with the satisfaction relation, let's look at a similar use in a simpler situation.

Each term t (of the underlying language L) is built from certain constants K, certain operation symbols O, and certain free variables v. How can we give t a value in some structure $\mathfrak{A}$? We have interpretations $\mathfrak{A}[\![K]\!]$ and $\mathfrak{A}[\![O]\!]$ of K and O, but we have no interpretation of v. To get round this we use an assignment x, and define the value of t in $\mathfrak{A}$ at x.

1.13 DEFINITION. For each structure $\mathfrak{A}$ (suitable for a language L), each $\mathfrak{A}$-assignment x, and each term t (of L) the element

$$\mathfrak{A}[\![t]\!]\mathsf{x}$$

of $\mathfrak{A}$ is generated by recursion on the construction of t using the following clauses.

- If t is a variable v then

  $$\mathfrak{A}[\![t]\!]\mathsf{x} = v\mathsf{x}$$

  using the element assigned to v by x.

- If t is a constant symbol K then

$$\mathfrak{A}[\![t]\!]\mathsf{x} = \mathfrak{A}[\![K]\!]$$

(the interpretation of K in $\mathfrak{A}$).

- If t is a compound

$$(Ot_1 \dots t_n)$$

where O is an n-placed operation symbol and $t_1, \dots, t_n$ are smaller terms, then

$$\mathfrak{A}[\![t]\!]\mathsf{x} = \mathfrak{A}[\![O]\!]a_1 \cdots a_n$$

where

$$a_i = \mathfrak{A}[\![t_i]\!]\mathsf{x}$$

for each $1 \leq i \leq n$.

No other clause are required. ■

There is nothing in this definition. Each term t names, in an obvious way, a certain (compound) operation on (the carrier of) $\mathfrak{A}$. This construction merely evaluates this operation, in the obvious way, where the inputs are supplied by the assignment x. In particular, for each term t almost all of the assignment x is not needed.

1.14 LEMMA. *Let $\mathfrak{A}$ be a structure and let t be a term (of the underlying language). If x and y are $\mathfrak{A}$-assignments which agree on ∂t, that is if*

$$v\mathsf{x} = v\mathsf{y}$$

holds for each $v \in \partial t$, then

$$\mathfrak{A}[\![t]\!]x = \mathfrak{A}[\![t]\!]y$$

holds.

This is proved by the obvious induction over the construction of t.

To generate the satisfaction relation we use the same trick. We define a more general relation

$$\mathfrak{A} \models \phi\mathsf{x}$$

which says

$\mathfrak{A}$ satisfies the formula ϕ where each free variable v takes the value $v\mathsf{x}$

and then we show the irrelevancy of most of x in Lemma 1.16.

1.15 DEFINITION. For each structure $\mathfrak{A}$ (suitable for a language L), each $\mathfrak{A}$-assignment x, and each formula ϕ (of L) the truth value

$$\mathfrak{A} \models \phi\mathsf{x}$$

is generated by recursion on the structure of t using the following clauses.

$$
\begin{array}{lcll}
\mathfrak{A} \models (\mathsf{true})\mathsf{x} & \Longleftrightarrow & \text{true} & \\
\mathfrak{A} \models (\mathsf{false})\mathsf{x} & \Longleftrightarrow & \text{false} & \\
\mathfrak{A} \models (t_1 \doteq t_2)\mathsf{x} & \Longleftrightarrow & \mathfrak{A}[\![t_1]\!]\mathsf{x} = \mathfrak{A}[\![t_2]\!]\mathsf{x} & \\
\mathfrak{A} \models (Rt_1 \cdots t_n)\mathsf{x} & \Longleftrightarrow & \mathfrak{A}[\![R]\!]a_1 \cdots a_n & \text{where } a_i = \mathfrak{A}[\![t_i]\!]\mathsf{x} \\
\mathfrak{A} \models (\neg\psi)\mathsf{x} & \Longleftrightarrow & \text{not}(\mathfrak{A} \models \psi\mathsf{x}) & \\
\mathfrak{A} \models (\theta \wedge \psi)\mathsf{x} & \Longleftrightarrow & \mathfrak{A} \models \theta\mathsf{x} \text{ and } \mathfrak{A} \models \psi\mathsf{x} & \\
\mathfrak{A} \models (\theta \vee \psi)\mathsf{x} & \Longleftrightarrow & \mathfrak{A} \models \theta\mathsf{x} \text{ or } \mathfrak{A} \models \psi\mathsf{x} & \\
\mathfrak{A} \models (\theta \to \psi)\mathsf{x} & \Longleftrightarrow & \text{not}(\mathfrak{A} \models \theta\mathsf{x}) \text{ or } \mathfrak{A} \models \psi\mathsf{x} & \\
\mathfrak{A} \models ((\forall v)\psi)\mathsf{x} & \Longleftrightarrow & \mathfrak{A} \models \psi\mathsf{y} & \begin{array}{l}\text{for each } \mathfrak{A}\text{-assignment } \mathsf{y} \\ \text{which agrees with } \mathsf{x} \text{ except} \\ \text{possibly in the } v\text{-position}\end{array} \\
\mathfrak{A} \models ((\exists v)\psi)\mathsf{x} & \Longleftrightarrow & \mathfrak{A} \models \psi\mathsf{y} & \begin{array}{l}\text{for some } \mathfrak{A}\text{-assignment } \mathsf{y} \\ \text{which agrees with } \mathsf{x} \text{ except} \\ \text{possibly in the } v\text{-position}\end{array}
\end{array}
$$

No other clauses are required. ■

Notice that

$$\mathfrak{A} \models (\mathsf{true})\mathsf{x}$$

always holds, whereas

$$\mathfrak{A} \models (\mathsf{false})\mathsf{x}$$

never holds. This is the principal job of these two constant sentences.

To get the original satisfaction relation (for sentences) we make the following observation.

1.16 LEMMA. *Let $\mathfrak{A}$ be a structure and let ϕ be a formula (of the underlying language). If x and y are $\mathfrak{A}$-assignments which agree on $\partial\phi$, that is if*

$$x\mathsf{x} = v\mathsf{y}$$

holds for each $v \in \partial\phi$, then

$$\mathfrak{A} \models \phi\mathsf{x} \Longleftrightarrow \mathfrak{A} \models \phi\mathsf{y}$$

holds.

By definition, a sentence is a formula σ with $\partial\sigma = \emptyset$. Vacuously, for such a sentence, each two $\mathfrak{A}$-assignments x and y agree on $\partial\sigma$, and hence

$$\mathfrak{A} \models \sigma\mathsf{x} \Longleftrightarrow \mathfrak{A} \models \sigma\mathsf{y}$$

holds. In other words, either

$$\mathfrak{A} \models \sigma\mathsf{x}$$

for every $\mathfrak{A}$-assignment or for no $\mathfrak{A}$-assignment. Thus, we may write

$$\mathfrak{A} \models \sigma$$

to indicate that $\mathfrak{A} \models \sigma\mathsf{x}$ holds for every x.

In a similar way we may simplify the notation $\mathfrak{A} \models \phi\mathsf{x}$.

Consider a formula $\phi(v_1, \ldots, v_n)$, that is a formula ϕ with $\partial\phi = \{v_1, \ldots, v_n\}$. Given a structure $\mathfrak{A}$ and an assignment x, the truth value of

$$\mathfrak{A} \models \phi\mathsf{x}$$

depends only on the elements

$$a_1 = v_1\mathsf{x}, \ldots, a_n = v_n\mathsf{x}$$

selected from x by the free variables. Thus we often write

$$\mathfrak{A} \models \phi(a_1, \ldots, a_n)$$

in place of the official notation.

Sometime we go even further. By a point a of the structure $\mathfrak{A}$ we mean a list $a_1, \ldots, a_n$ of elements of $\mathfrak{A}$. We may then write

$$\mathfrak{A} \models \phi(a)$$

for the satisfaction relation. Of course, this assumes there is a match between the point a and the list v of free variables of ϕ.

In subsection 1.2 we introduced the idea of isomorphic structures

$$\mathfrak{A} \cong \mathfrak{B}$$

(of the same signature). There is a semantic analogue of this.

We write

$$\mathfrak{A} \equiv \mathfrak{B}$$

and say $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent if

$$\mathfrak{A} \models \sigma \Longleftrightarrow \mathfrak{B} \models \sigma$$

holds for each sentence σ (of the underlying language). Almost trivially

$$\mathfrak{A} \cong \mathfrak{B} \Longrightarrow \mathfrak{A} \equiv \mathfrak{B}$$

holds (but the proof of this is rather tedious). However, the converse is false (in general). In fact, as we will see in Theorem 1.27, it can happen that $\mathfrak{A} \equiv \mathfrak{B}$ but $|\mathfrak{A}| \neq |\mathfrak{B}|$ and so these structures can't be isomorphic.

Exercises

1.3 Sketch the proofs of Lemmas 1.14 and 1.16.

1.4 Consider the formulas ϕ_r generated in Exercise 1.1, and let $\mathfrak{N} = (\mathbb{N}, <)$.

(a) Characterize the elements a of $\mathfrak{N}$ such that $\mathfrak{N} \models \phi_r(a)$.

(b) Show that $\mathfrak{N} \models (\forall v)[\phi_{r+1} \to \phi_r]$ holds for each $r < \omega$.

(c) Describe the formulas $\psi(v)$ such that $\mathfrak{N} \models (\forall v)[\psi \to \phi_r]$ holds for each $r < \omega$.

1.5 Consider the language on the empty signature. In other words, consider the language of pure equality.

(a) Show that for each $n < \omega$ there are sentences σ_n and τ_n such that

$$\mathfrak{A} \models \sigma_n \iff |\mathfrak{A}| \geq n \qquad \mathfrak{A} \models \tau_n \iff |\mathfrak{A}| = n$$

holds for each structure $\mathfrak{A}$. What is the quantifier complexity of each of these sentences?

(b) Show there is a set Inf of sentences such that

$$\mathfrak{A} \models Inf \iff \mathfrak{A} \text{ is infinite}$$

holds for each structure $\mathfrak{A}$. Is there a finite set of such sentences?

(c) Is there a set Fin of sentences such that

$$\mathfrak{A} \models Fin \iff \mathfrak{A} \text{ is finite}$$

for each each structure $\mathfrak{A}$?

1.6 Suppose $\mathfrak{A} \subseteq \mathfrak{B}$.

Show that

$$\mathfrak{A} \models \delta(a) \iff \mathfrak{B} \models \delta(a)$$

for each quantifier-free formula $\delta(v)$ and point a of $\mathfrak{A}$ which matches the free variables v of δ.

Show that

$$\mathfrak{A} \models \theta(a) \Longrightarrow \mathfrak{B} \models \theta(a)$$

for each $\exists_1$-formula $\theta(v)$ and point a of $\mathfrak{A}$ which matches the free variables v of θ.

Find an example to show that this implication is not an equivalence.

[*Held in 120../B14-bit.. Last changed July 26, 2004*]

1.4 Consequence

Each language L has an associated satisfaction relation

$$\mathfrak{A} \models \sigma$$

between structures $\mathfrak{A}$ (for L) and sentences σ (of L). We can refine this.

1.17 DEFINITION. Let L be a language, let $\mathcal{K}$ be a class of structures for L, and let Σ be a set of sentences of L.

(a) The relation

$$\mathcal{K} \models \Sigma$$

holds if

$$\mathfrak{A} \models \sigma$$

holds for each $\mathfrak{A} \in \mathcal{K}$ and $\sigma \in \Sigma$.

(b) The theory $Th(\mathcal{K})$ of $\mathcal{K}$ is the set of all sentences σ such that $\mathcal{K} \models \sigma$.

(c) The models $\mathcal{M}(\Sigma)$ of Σ is the class of all structures $\mathfrak{A}$ such that $\mathfrak{A} \models \Sigma$. ■

The two assignments $\mathcal{M}(\cdot)$ and $Th(\cdot)$ form a galois connection. Thus

$$\mathcal{K} \subseteq \mathcal{M}(\Sigma) \Longleftrightarrow \Sigma \subseteq Th(\mathcal{K})$$

holds for each class $\mathcal{K}$ (of structures) and each set Σ (of sentences). Furthermore, either side holds precisely when

$$\mathcal{K} \models \Sigma$$

holds. In particular, both the composites $Th \circ \mathcal{M}$ and $\mathcal{M} \circ Th$ are closure operations and, as expected, we look at the closed gadgets.

1.18 DEFINITION. (a) A set T of sentences is a theory if $T = Th(\mathcal{K})$ for some class $\mathcal{K}$ of structures. Equivalently, T is a theory if (and only if) $T = Th(\mathcal{M}(T))$.

(b) A class $\mathcal{K}$ of structures is elementary or an elementary class if $\mathcal{K} = \mathcal{M}(\Sigma)$ for some set Σ of sentence. Equivalently, $\mathcal{K}$ is elementary if (and only if) $\mathcal{K} = \mathcal{M}(Th(\mathcal{K}))$. ■

These notions prompt some obvious questions.

- Are there any necessary and sufficient conditions for a class to be elementary?
- By definition, a class is elementary if it has the form $\mathcal{M}(\Sigma)$ for a set Σ of sentences. When is a class strictly elementary, that is when does it have the form $\mathcal{M}(\sigma)$ for a single sentence?
- What are the necessary and sufficient conditions for a set of sentences to be a theory?
- By definition, a set T is a theory if it has the form $Th(\mathcal{M}(\Sigma))$ for some set Σ. We then say Σ axiomatizes T. When does a theory have a finite set of axioms? When is a theory $\forall_n$-axiomatizable for some n?

To answer these and similar questions we need a tool, the compactness theorem. This is the pivotal method of model theory, and is discussed in detail in the next subsection. Here we see how it relates to another part of mathematical logic.

1.19 DEFINITION. A set Σ of sentences (of some language) is consistent or satisfiable if it has a model, that is if $\mathfrak{A} \models \Sigma$ for some structure $\mathfrak{A}$. ■

The theory T is inconsistent (not consistent) if $T = Th(\emptyset)$, in which case T is the set of all sentences of the language. Rarely do we need to consider this theory, so we often say 'a theory T' when we mean 'a consistent theory T'.

1.20 DEFINITION. A theory T is complete if it is consistent and $\mathfrak{A} \equiv \mathfrak{B}$ for all models $\mathfrak{A}, \mathfrak{B}$ of T.

A theory T is κ-categorical (for a cardinal κ) if $\mathfrak{A} \cong \mathfrak{B}$ for all models $\mathfrak{A}, \mathfrak{B}$ of T with $|\mathfrak{A}| = \kappa = |\mathfrak{B}|$. ■

It is an easy exercise to see that a theory is complete if and only if it has the form $Th(\mathfrak{A})$ for a structure $\mathfrak{A}$. Another easy exercise (but using a result we haven't yet seen) shows that if a consistent theory T in in a language L is κ-categorical for some $\kappa \geq |L|$, then it is complete.

Roughly speaking a complete theory is a large theory in the sense that it can't take in any more sentences without becoming inconsistent. At the other extreme, the pure logic of a language it the theory of the class of all structures for that language. Thus a sentence belongs to this theory if and only if it is universally valid.

1.21 DEFINITION. For a set Σ of sentences and a sentence σ we write

$$\Sigma \vdash \sigma$$

and say Σ entails σ or σ is a consequence of Σ if $\sigma \in Th(\Sigma)$, that is if $\mathfrak{A} \models \sigma$ for each model $\mathfrak{A} \models \Sigma$. ■

In this notation Σ is a set of axioms for a theory T exactly when

$$\sigma \in T \Longleftrightarrow \Sigma \vdash \sigma$$

holds for each sentence σ. Often we describe a theory by writing down a particular set of axioms. Then one of the problems is to characterize all (or a large amount of) the consequences of the axioms. At other times the problem can be to axiomatize the theory of some given class of structures (which is described in a non-elementary way).

Although it is not strictly part of model theory, at this point it is worth comparing this semantic consequence relation with the kind of consequence relation met in a course on predicate calculus.

The Definition 1.21 of the relation $\vdash$ involves an external quantification over a potentially large class of structures. The relation $\Sigma \vdash \sigma$ holds if …*for all structures* $\mathfrak{A}$. However, $\Sigma \vdash \sigma$ is a relation between syntactic objects, and the question arises of whether it can be characterized in purely combinatorial terms.

Gödel's completeness theorem shows that it can.

(You should not confuse the two different uses of 'completeness' here. They are related, but not the same.)

To analyse $\vdash$ we first set up a proof-theoretic relation

$$\Sigma \dot{\vdash} \sigma$$

between sets Σ of sentences and sentences σ. The essential feature of this is that it is entirely combinatorial. This relation holds if and only if there is a certain (finite) configuration of strings of symbols. The intended semantics is not referred to at all. Thus, $\Sigma \dot{\vdash} \sigma$ can be shown to hold by exhibiting a certain collection of symbols formatted in an appropriate way.

There are several different ways of setting up $\dot{\vdash}$, most of which are needed for one job or another (and some of which are entirely untainted by content and interest). Here we needn't worry about the precise details.

The analysis now investigates the relationship between

$$\dot{\vdash} \qquad \vdash$$

to produce two particular results, one minor and one major.

- The relation $\dot{\vdash}$ is sound, that is

$$\Sigma \dot{\vdash} \sigma \Longrightarrow \Sigma \vdash \sigma$$

holds for all Σ and σ. This is a relatively trivial observation.

- The relation $\dot{\vdash}$ is adequate, that is

$$\Sigma \vdash \sigma \Longrightarrow \Sigma \dot{\vdash} \sigma$$

holds for all Σ and σ. This requires quite a bit of work.

- The combination of these two results is the completeness theorem, that is

$$\Sigma \vdash \sigma \Longleftrightarrow \Sigma \dot{\vdash} \sigma$$

holds for all Σ and σ.

Because of the way $\dot{\vdash}$ is set up we observe that if

$$\Sigma \dot{\vdash} \sigma$$

then

$$\Gamma \dot{\vdash} \sigma$$

for some finite part Γ of Σ. This leads to the following result.

1.22 THEOREM. *Let Σ be a set of sentences (in some language). If each finite part of Σ has a model, then Σ has a model.*

Proof. We prove the contrapositive. Thus suppose Σ does not have a model. Then, vacuously we have

$$\Sigma \vDash \sigma$$

for each sentence σ. Consider the sentence false which does not have a model. We have

$$\Sigma \vDash \mathsf{false}$$

and hence

$$\Sigma \dot{\vdash} \mathsf{false}$$

by the adequacy of $\dot{\vdash}$. But now

$$\Gamma \dot{\vdash} \mathsf{false}$$

for some finite part Γ of Σ, and then

$$\Gamma \vDash \mathsf{false}$$

by the soundness of $\dot{\vdash}$. This shows that Γ does not have a model. ■

This result is a version of the compactness theorem.

Exercises

The first two exercises are concerned with the language with just one attribute, and that is a binary relation symbol. Thus a structure for this language has the form (A, R) where A is a non-empty set and R is binary relation on A.

1.7 An equivalence structure has the form (A, R) where R is an equivalence relation on the carrier A.

(a) Axiomatize the class of equivalence structures.

(b) For $m, n < \omega$, axiomatize the class of equivalence structures each having no more than m equivalence classes, and each of these has no more than n members.

(c) Axiomatize the class of equivalence structures having infinitely many equivalence classes and each of these is infinite.

(d) Axiomatize the class of equivalence structures which are such that if there is a finite equivalence class, then there is an equivalence class of each larger finite size.

1.8 (a) Axiomatize the classes of posets, linear orderings, dense linear orderings, and discrete linear orderings.

(b) Write down formulas $\theta(u, v, w), \psi(u, v, w), \phi(u, v)$ such that

$$\mathfrak{A} \models \theta(a, b, c) \iff c \text{ is the l.u.b. of } b \text{ and } c$$

$$\mathfrak{A} \models \psi(a, b, c) \iff a, b, c \text{ are linearly ordered in } \mathfrak{A}$$

$$\mathfrak{A} \models \phi(a, b) \iff \begin{array}{l} \text{if } a, b \text{ are comparable and not equal, then exactly} \\ \text{three elements lie strictly between } a, b, \text{ and these} \\ \text{three elements are pairwise incomparable} \end{array}$$

hold for each poset $\mathfrak{A}$ and elements a, b, c of $\mathfrak{A}$.

The next exercise uses the language suitable for structures

$$\mathfrak{A} = (A, *, e)$$

where $*$ is a binary operation on A and e is a distinguished element.

1.9 (a) Axiomatize the classes of groups, abelian groups, torsion-free abelian groups, and divisible abelian groups. Which of these classes are finitely axiomatizable.

Write down formulas $\theta(u), \psi(u, v), \phi(u, v)$ such that

$$\mathfrak{A} \models \theta(a) \iff a \text{ is a commutator}$$

$$\mathfrak{A} \models \psi(a, b) \iff a \text{ is in the centralizer of } b$$

$$\mathfrak{A} \models \phi(a, b) \iff \text{there is an inner automorphism taking } a \text{ to } b$$

for each group $\mathfrak{A}$ and elements a, b of $\mathfrak{A}$.

1.10 By using a suitable signature axiomatize the classes of rings (with 1), commutative rings (with 1), integral domains, integral domains of characteristic p (where p is a prime), integral domains of characteristic 0, fields, algebraically closed fields.

Which of the classes are $\forall_2$-axiomatizable, $\forall_1$-axiomatizable, finitely axiomatizable?

1.11 Consider the reals as a structure $(\mathbb{R}, +, \times, \leq, 0, 1)$ (with the obvious attributes. Look up the axioms which characterize this structure up to isomorphism. Observe that most of these are first order, but the crucial one isn't. What is this non-elementary axiom?

1.12 Let R be a ring with 1, and consider the right R-modules. Think of each of these as a structure

$$\mathfrak{A} = (A, +, 0, (f_r \mid r \in R))$$

where $(A, +, 0)$ is an abelian group and, for each $r \in R$, the 1-placed operation f_r is $a \mapsto ar$. Thus R is used to index part of the signature.

Write down axioms for this class of modules.

What changes need to be made to axiomatize the class of left R-modules?

1.13 Show that for each consistent theory T the following are equivalent.

(i) T is complete

(ii) For each sentence σ, if $T \cup \{\sigma\}$ is consistent, then $\sigma \in T$.

(iii) For each theory T', if $T \subseteq T'$ then either $T = T'$ or T' is inconsistent.

(iv) For each pair σ, τ of sentences, if $\sigma \vee \tau \in T$, then $\sigma \in T$ or $\tau \in T$.

The following exercise is quite tricky, but it solution uses an important technique which you should learn as soon as possible.

1.14 Show that if no finite extension of a (consistent) theory is complete, then the theory has at least $2^{\aleph_0}$ complete extensions.

Finally, here is almost all you need to know about galois connections.

1.15 Let A, S be a pair of posets with elements $a, b, c, \ldots$ and $r, s, t, \ldots$, respectively. Let

$$A \xrightarrow{\ \sharp\ } S \qquad\qquad A \xleftarrow{\ \flat\ } S$$

be a pair of assignments such that

$$a \leq \flat s \Longleftrightarrow s \leq \sharp a$$

holds for each $a \in A$ and $s \in S$.

(a) Show that both the composites $\flat \circ \sharp$ and $\sharp \circ \flat$ are inflationary.

(b) Show that $\sharp \circ \flat \circ \sharp = \sharp$ and $\flat \circ \sharp \circ \flat$ hold.

(c) Show that both $\flat$ and $\sharp$ are antitone.

(d) Show that both $\flat \circ \sharp$ and $\sharp \circ \flat$ are closure operations.

[*Held in 120.../B15-bit... Last changed July 26, 2004*]

1.5 Compactness

At the end of the previous subsection we obtained a statement (and an indication of a proof) of the compactness theorem. In this subsection we take a closer look at this result.

1.23 DEFINITION. A set Σ of sentences (of some language) is consistent or satisfiable if it has a model, that is if $\mathfrak{A} \models \Sigma$ for some structure $\mathfrak{A}$.

A set Σ of sentences is finitely satisfiable if each finite part has a model. ■

In this terminology, we saw that the completeness theorem implies the following.

1.24 THEOREM. (The crude compactness theorem) *If a set of sentences (of some language) is finitely satisfiable, then it is satisfiable.*

How should we prove this?

We have seen already one method of proof. We set up a proof-theoretic consequence relation $\dot{\vdash}$ and then prove a completeness result. The compactness result is an immediate consequence. However, this is not entirely satisfactory, for two reasons.

Firstly, we must set up all the machinery for $\dot{\vdash}$, and this takes some time. Furthermore, this machinery is never used again in model theory. (It may be used elsewhere in mathematical logic, but then $\dot{\vdash}$ will be the principal object of study, and it will be designed with some specific class of tasks in mind.)

Secondly, at the heart of the proof of completness a certain structure is constructed. The method of construction can be modified to give a direct proof of compactness (without a detour through $\dot{\vdash}$ and its properties).

The witnessing construction is described in section 8. This is a method of producing a structure out of a certain kind of family of sets of sentences (called a consistency property). In the first instance this construction gives us both compactness and completeness, virtually by the same proof. This method of construction is quite flexible, and gives us quite a lot of control over the end product. This is used to advantage in more advanced work. Furthermore, the same method can be lifted to higher order languages (but, of course, this requires a bit more work).

Another idea on how to prove compactness should have occurred to you.

Let Σ be a finitely satisfiable set of sentences, and let $\mathbf{\Delta}$ be the set of finite subset Δ of Σ. We are given a model $\mathfrak{A}(\Delta)$ of each such $\Delta \in \mathbf{\Delta}$. Is there a way of patching together, in a coherent fashion, all of these $\mathfrak{A}(\Delta)$ to produce a model of Σ? There is, and it is called the ultraproduct construction. This is described in section 14.

For the time being we do not need the details of the proof of Theorem 1.24, so let's look at some applications of compactness.

1.25 THEOREM. *Let L be any language.*

(a) The class of all infinite structures (for L) is elementary but not strictly elementary.

(b) The class of all finite structures is not elementary.

(c) A sentence (of L) holds in all infinite structures if and only if it holds in all sufficiently large structures.

(d) The theory of the class of finite structures has an infinite model.

(e) The theory of the class of infinite structures has no finite model.

Proof. (a) By Exercise 1.5 we know that for each $n < \omega$ there is a sentence σ_n such that

$$\mathfrak{A} \models \sigma_n \iff |\mathfrak{A}| \geq n$$

holds for each structure $\mathfrak{A}$. Let

$$Inf = \{\sigma_n \mid n < \omega\}$$

so that $\mathcal{M}(Inf)$ is exactly the class of infinite structures. In particular, this class is elementary.

By way of contradiction, suppose that this class is strictly elementary. Thus there is a single sentence τ such that

$$\mathfrak{A} \models \tau \iff \mathfrak{A} \text{ is infinite}$$

holds for each structure $\mathfrak{A}$. In particular,

$$\mathfrak{A} \models \neg\tau \iff \mathfrak{A} \text{ is finite}$$

holds for each structure $\mathfrak{A}$. We show that the set

$$Inf \cup \{\neg\tau\}$$

is consistent, which is the required contradiction.

Any finite subset of this set is a subset of

$$\{\sigma_0, \ldots, \sigma_n, \neg\tau\}$$

for some $n < \omega$. By considering a sufficiently large finite structure, we see that this subset has a model. Thus $Inf \cup \{\neg\tau\}$ is finitely satisfiable and hence, by the compactness property, is satisfiable, as required.

(b) By way of contradiction, suppose the class of finite structures is elementary. Thus there is a set Fin of sentences such that

$$\mathfrak{A} \models Fin \iff \mathfrak{A} \text{ is finite}$$

holds for each structure $\mathfrak{A}$. A slight modification of the argument used in (a) show that the set

$$Inf \cup Fin$$

is finitely satisfiable, and hence is satisfiable. This is not so, since no structure is both infinite and finite.

(c) Suppose the sentence τ holds in all infinite structures. Then

$$Inf \vdash \tau$$

and hence, by compactness, we have

$$\sigma_0, \ldots, \sigma_n \vdash \tau$$

for some $n < \omega$. Thus τ holds in any structure $\mathfrak{A}$ with $|\mathfrak{A}| \geq n$.

Conversely, suppose there is some $n < \omega$ such that the sentence τ holds in each structure $\mathfrak{A}$ with $|\mathfrak{A}| \geq n$. Then

$$\sigma_n \vdash \tau$$

and hence

$$Inf \vdash \tau$$

to show that τ holds in all infinite structures.

(d) Let T be the theory of the class of finite structures. For each $n < \omega$, any sufficiently large structure is a model of $T \cup \{\sigma_n\}$, and hence

$$T \cup Inf$$

is finitely satisfiable. By compactness, this set has a model, and hence T has an infinite model.

(e) Now let T be the theory of the class of infinite structure. Then $Inf \subseteq T$, and hence no finite structure can be a model of T. ■

Theorem 1.24 is the *crude* compactness result because it can be refined to extract more information. We use the cardinality $|L|$ of the underlying language.

1.26 THEOREM. (The refined compactness theorem) *Let Σ be a set of L-sentences (for some language). If Σ is finitely satisfiable, then it has a model $\mathfrak{A}$ with $|\mathfrak{A}| = |L|$.*

Notice that this refined version does not follow from completeness, as outline in the previous subsection. However, it is an immediate consequence of the witnessing construction, which allows us to control the size of the structure produced. This might not seem much, but it has some surprising consequences.

1.27 THEOREM. *The theory $Th(\mathfrak{N})$ of the natural numbers is not $\aleph_0$-categorical. That is there is a countable structure $\mathfrak{A}$ with $\mathfrak{A} \equiv \mathfrak{N}$ and $\mathfrak{A} \not\cong \mathfrak{N}$.*

Proof. Notice that we didn't specify which language $Th(\mathfrak{N})$ is formalized in. That is because is doesn't matter. The result holds no matter which language we used. However, it is useful to have numerals (constant terms) $\ulcorner n \urcorner$ in the language.

To prove the result we enrich the language by adding one new constant symbol a, say. Look at the set

$$Th(\mathfrak{N}) \cup \{(\ulcorner n \urcorner \neq a) \,|\, n \in \mathbb{N}\}$$

in this enriched language. This is finitely satisfiable. To see this notice that in any finite part there is a largest $n \in \mathbb{N}$ such that $\ulcorner n \urcorner$ occurs, and then

$$(\mathfrak{N}, n+1)$$

is a model.

By the refined compactness result, the set has a countable model $(\mathfrak{A}, a)$ where $\mathfrak{A} \equiv \mathfrak{N}$ and a is some distinguished element. There is a unique embedding $\mathfrak{N} \longrightarrow \mathfrak{A}$ (given by $n \longmapsto \mathfrak{A}[\![\ulcorner n \urcorner]\!]$) and the extra sentences ensure that a is not in the range of this. Thus $\mathfrak{A} \not\cong \mathfrak{N}$. ■

This is sometimes known as Skolem's paradox (even though it is not a paradox, just a surprise). Skolem's original proof used a kind of ultraproduct construction. Later the trick used in this proof will be turned into a powerful tool.

1.16 (a) Show that if $\Sigma \vdash \tau$ (where Σ is a set of sentences and τ is a sentence of the same language), then $\Gamma \vdash \tau$ for some finite $\Gamma \subseteq \Sigma$.

(b) Show that a consistent set Σ of sentences is a theory if and only if it contains all universally valid sentences and $\tau \in \Sigma$ whenever $\sigma, \sigma \to \tau \in \Sigma$.

1.17 Suppose

$$\mathcal{K} = \bigcap \{\mathcal{K}_r \mid r < \omega\}$$

where $\{\mathcal{K}_r \mid r < \omega\}$ is a strictly descending chain of strictly elementary classes. Show that $\mathcal{K}$ is elementary but not strictly elementary.

1.18 Let $\mathcal{K}$ be a strictly elementary class (for some language) and suppose

$$\mathcal{K} = \mathcal{L} \cup \mathcal{R} \qquad \mathcal{L} \cap \mathcal{R} = \emptyset$$

where both $\mathcal{L}$ and $\mathcal{R}$ are elementary. Show that both $\mathcal{L}$ and $\mathcal{R}$ are strictly elementary.

1.19 Let $\mathcal{F}, \mathcal{F}_0, \mathcal{F}_p, \mathcal{F}_f$ be the classes of fields, fields of characteristic zero, fields of characteristic p (for a given prime p), a fields of finite (no-zero) characteristic, respectively. Let T, T_0, T_p, T_f be the respective theories of these classes.

(a) Which of these classes are elementary and which are strictly elementary.

(b) Which if these theories are finitely axiomatizable.

(c) Show that each sentence $\tau \in T_0$ holds in each field of sufficiently large (prime) characteristic.

(d) Show that T_f has a model of characteristic zero.

1.20 Let $\mathfrak{R}$ be the real numbers viewed as a first order structure.

Show there is a countable structure $\mathfrak{A}$ with $\mathfrak{A} \equiv \mathfrak{R}$.

Can you say what this structure might be?

[*Held in 120../B20-bit.. Last changed July 26, 2004*]

2 The effective elimination of quantifiers

Strictly speaking, the topic of this section, quantifier elimination, is not a part of model theory proper. It is included here for two reasons, one minor and one major. Anyone who claims to have some familiarity with mathematical logic should know something about quantifier elimination. That is the minor reason. The major reason is that the topic had a considerable influence on the early development of model theory, and we will follow and idealized version of that path. It could be said the quantifier elimination is a recurring theme throughout these notes.

[*Held in 120../B21-bit.. Last changed July 26, 2004*]

2.1 The generalities of quantifier elimination

Suppose T is a theory in some language. It doesn't matter how T is describe. It could be given in the form $Th(\mathcal{K})$ for some class $\mathcal{K}$ of structures. It could be given as the consequences of some set of axioms. It could be given in some other way.

To understand T we need to know something about the way quantification behaves in (the models of) T.

2.1 DEFINITION. Let T be a theory in some language.

(a) Two formulas ϕ and ψ are T-equivalent if

$$T \vdash (\forall v)[\phi \leftrightarrow \psi]$$

where v is a list of variables which includes $\partial\phi \cup \partial\psi$.

(b) The theory T has EQ (elimination of quantifiers) if each formula is T-equivalent to some quantifier-free formula. ■

Of course, if a theory has EQ then it must be rather special. One of our long terms aims (which we achieve in section 4) is to characterize this speciality. In this section we begin with a few examples of this property.

How can we show that a theory T has EQ? The obvious way is to describe an algorithm which, when supplied with a formula ϕ, will return a quantifier-free formula ψ which is T-equivalent to ϕ. In this section we will describe, in reasonable but not full detail, two examples of such an algorithm. We will then survey some of the other algorithms of this kind.

The theories considered in this section have, what we term, 'effective elimination of quantifiers'. However, the qualifier 'effective' has very little content. In section 4 we will give a more general characterization of EQ. The word 'effective' is used here merely to distinguish these examples from this later characterization.

At first sight it looks rather complicated to organize an algorithm which eliminates quantifiers from a theory. This is because we have to handle all possible combinations of quantifiers. However, some of the basic results of logic help with this organization, and takes us to the heart of the problem.

2.2 THEOREM. *To eliminate quantifiers for a theory T it is sufficient (and necessary) to find a quantifier-free equivalent (relative to T) of each formula*

$$(\exists w)\delta(w, v_1, \ldots, v_k)$$

where δ is a conjunction of literals (in the indicated variables) and where the quantified variable w occurs in each such literal.

Proof. Suppose we can eliminate the quantifier $(\exists w)$ from each formula of the indicated kind. We show how to eliminate quantifiers from progressively larger classes of formulas until we have dealt with all formulas.

Consider first a formula

$$(\exists w)[\gamma \wedge \delta]$$

where each of γ and δ is a conjunction of literals, where w does not occur in γ, but w does occur in each conjunct of δ. This formula is logically equivalent to

$$\gamma \wedge (\exists w)\delta$$

so, by the given algorithm, we can eliminate the quantifier $(\exists w)$. In other words, we can eliminate the quantifier from any formula

$$(\exists w)\delta$$

where δ is any conjunction of literals (without any restrictions on the occurrences of w).

Consider any formula

$$(\exists w)[\delta_1 \vee \ldots \vee \delta_m]$$

where each δ_i is a conjunction of literals. (Thus every quantifier-free formula can be put in the disjunctive normal form of this matrix.) The whole formula is logically equivalent to

$$(\exists w)\delta_1 \vee \ldots \vee (\exists w)\delta_m$$

so, by the above algorithm, we can eliminate the quantified variable from each of these separate disjuncts. In other words, we can eliminate the quantifier from any formula

$$(\exists w)\delta$$

where δ is any quantifier-free formula.

Consider any formula

$$(\exists w_l, \ldots, w_1)\delta$$

where δ is quantifier-free. By considering

$$(\exists w_1)\delta \quad (\exists w_2, w_1)\delta \quad \cdots \quad (\exists w_l, \ldots, w_1)\delta$$

we can eliminate each quantifier in turn (from the inside) using the quantifier-free equivalents at the successive stages. In other words, we can eliminate the quantifiers from any formula

$$(\exists w)\delta$$

where δ is quantifier-free and $(\exists w)$ is any *block* of existentially quantified variables.

We now eliminate the quantifiers from each $\exists_{n+1}$-formula. We proceed by recursion on n. The base case, $n = 0$, is dealt with above. For the recursion step, $n \mapsto n+1$, consider any $\exists_{n+2}$-formula

$$\psi = (\exists w)\phi$$

where ϕ is a $\forall_{n+1}$-formula and w is a list of variables. The negation $\neg\phi$ is a $\exists_{n+1}$-formula so, by recursion, we obtain

$$T \vdash \neg\phi \leftrightarrow \delta$$

for some quantifier-free formula δ. In particular.

$$T \vdash \phi \leftrightarrow \neg\delta$$

so that

$$T \vdash \psi \leftrightarrow (\exists w)\neg\delta$$

and it suffices to apply the base algorithm to eliminate this last block of quantifiers. ■

In the next two subsections we look at two particular theories, and show that each has EQ.

Exercises

2.1 Let T be a theory with EQ formalized in language with a finite signature where there are no operation symbols.

What can you say about the size of the boolean algebra of sentences modulo T.

How many complete extensions does T have?

[*Held in 120../B22-bit.. Last changed July 26, 2004*]

2.2 The natural numbers

How can we characterize the natural numbers? Dedekind observed that the structure

$$\mathfrak{N} = (\mathbb{N}, S, 0)$$

(where S is the successor operation) is characterized by the induction property.

Each subset X of $\mathbb{N}$ which contains 0 and
is closed under S must be the whole of $\mathbb{N}$

Peano pointed out that some care must be taken with this idea, for we need to know which sets X are 'acceptable'. In the present context this means that the nature of the language in which the characterization is formalized has a significant impact on the result. For instance, by Theorem 1.27, if we use a first order language, then a characterization up to isomorphism is impossible. The best we can hope for is a characterization up to elementary equivalence. In other words we can not hope for much more than a characterization of $Th(\mathfrak{N})$.

Let's attempt to axiomatize this theory.

There are two trivial axioms.

$$(0)\quad (\forall v)[(Sv \neq 0)] \qquad\qquad (1)\quad (\forall u, v)[(Su \doteq Sv) \to (u \doteq v)]$$

which are the first two of Dedekind's axioms.

Next we want to add to these some analogue of the induction axiom (as stated above). We can not formalize this directly in our first order language since it involves a quantification over *subsets* of the carrier. However, many such subsets can be named in the language, and we can certainly state the induction property for each one of these.

Let $\phi(u_1, \ldots, u_n, v)$ be any formula in the indicated variables. We can think of this as a name for the set of all v for which the formula holds. Of course, this set depends on the parameters $u_1, \ldots, u_n$. In other words, the formula gives us a parameterized family of subsets. We can thus regard

$$(\forall u_1, \ldots u_n)[\phi(u, 0) \wedge (\forall v)[\phi(u, v) \to \phi(u, Sv)] \,.\!\!\to\!\!.\, (\forall v)\phi(u, v)]$$

as a statement of the induction property for this particular family of subsets.

Thus, we can look at the theory T^+ axiomatized by the two trivial axioms (0,1) together with all the induction axioms for all possible formulas ϕ. Certainly $\mathfrak{N} \models T^+$.

To analyse T^+ the first thing to do is to extract some useful consequence of the axioms. For this we need a bit of notation.

The terms of this language have a simple form. Each has one of the shapes

$$S^k 0 \qquad S^k w$$

where w is a variable and $k \in \mathbb{N}$. Here 'S^k' indicates a k-fold application of S to either 0 or w. We abbreviate

$$S^k 0 \quad \text{by} \quad \ulcorner k \urcorner$$

to obtain the **numerals**. Thus in the structure $\mathfrak{N}$ the numeral $\ulcorner k \urcorner$ is the canonical name of $k \in \mathbb{N}$. Notice that $\ulcorner 0 \urcorner$ and 0 are the same term.

2.3 LEMMA. *Both*

$$T^+ \vdash (\forall v)[(v \doteq \ulcorner 0 \urcorner) \vee (\exists w)[Sw \doteq v]] \qquad\qquad T^+ \vdash (\forall v)[S^{k+1} v \neq v]$$

hold (for each $k \in \mathbb{N}$).

These are proved by a mixture of internal and external induction. With this we can define a more amenable theory.

2.4 DEFINITION. Let T be the theory axiomatized by

(0) $(\forall v)[(Sv \neq 0)]$

(1) $(\forall u, v)[(Su \doteq Sv) \to (u \doteq v)]$

(2) $(\forall v)[(v \doteq \ulcorner 0 \urcorner) \vee (\exists w)[Sw \doteq v]]$

(3) $(\forall v)[S^{k+1} v \neq v]$

for each $k \in \mathbb{N}$. ■

Observe that T is $\forall_2$-axiomatizable. In fact, only the third axiom is a $\forall_2$-sentence, each of the others is a $\forall_1$-sentence. Lemma 2.3 shows that $T \subseteq T^+$. In particular, $\mathfrak{N} \models T$. We will show that T has EQ and hence, as a result, T is complete, so that $T = T^+ = Th(\mathfrak{N})$.

We need some more consequences of these axioms.

2.5 LEMMA. *For each $k \in \mathbb{N}$*

$$T \vdash (\forall v)[(\exists w)[S^{k+1}w \doteq v] \leftrightarrow (v \not\doteq \ulcorner 0 \urcorner) \wedge \cdots \wedge (v \not\doteq \ulcorner k \urcorner)]$$

holds.

One final observation before we get to the elimination algorithm. Each quantifier-free sentence of this language is equivalent to a combination of atomic sentence

$$(\ulcorner m \urcorner \doteq \ulcorner n \urcorner)$$

for various $m, n \in \mathbb{N}$. Each such compound sentence is either true or false in $\mathfrak{N}$. In fact, for each such sentence σ, either $\sigma \in T$ or $\neg\sigma \in T$.

With this we can show how to eliminate the quantifiers relative to the theory T.

2.6 THEOREM. *The theory T has EQ.*

Proof. How can we eliminate the bound variable w from the formula

$$\theta := (\exists w)[L_1 \wedge \cdots \wedge L_l]$$

where each conjunct L is a literal? On general ground we may assume that w occurs in each L. Thus each such literal is either an atomic formula or the negation of an atomic formula of the shape

$$(S^m w \doteq S^n t)$$

where $m, n \in \mathbb{M}$ and the term t is $w, 0$, or another variable. We consider all the various possibilities, and act accordingly.

Suppose there is an atomic formula α, perhaps negated, of the shape

$$(S^m w \doteq S^n w)$$

for $m, n \in \mathbb{N}$. By considering the cases

$$m = n \qquad m \neq n$$

we see that

$$T \vdash \alpha \leftrightarrow \mathsf{true} \qquad T \vdash \alpha \leftrightarrow \mathsf{false}$$

holds, respectively. In other words, either that conjunct can be disregarded, or

$$T \vdash \neg\theta$$

holds

The upshot of this is that we may assume that each occurring atomic formula has the shape

$$(S^m w \doteq s)$$

where $m \in \mathbb{N}$ and w does not appear in the term s.

Suppose one of the conjuncts L is positive, that is

$$L := (S^m w \doteq s)$$

for some term s. To eliminate the quantified variable $(\exists w)$ from θ we combine this particular conjunct L with each other conjunct M in turn. Each such conjunct M has one of the shapes

$$(S^p w \doteq t) \qquad (S^p w \not\doteq t)$$

depending on its parity. Remember that w does not appear in the term t. Consider the first shape. Then, working in T we have

$$\begin{aligned} L \wedge M &\leftrightarrow (S^{m+p} w \doteq S^p s) \wedge (S^{p+m} w \doteq S^m t) \\ &\leftrightarrow (S^{m+p} w \doteq S^p s) \wedge (S^p s \doteq S^m t) \qquad \leftrightarrow \quad L \wedge (S^p s \doteq S^m t) \end{aligned}$$

where now w does not appear in the second component. The case where M is negative can be handled in the same way, so we obtain

$$T \vdash L \wedge M \leftrightarrow L \wedge (S^p s \doteq S^m t) \qquad T \vdash L \wedge M \leftrightarrow L \wedge (S^p s \not\doteq S^m t)$$

for the positive case and negative case, respectively. From this we see that the matrix of θ is equivalent to

$$L \wedge \delta$$

for some quantifier-free formula δ in which w does not occur. Thus

$$T \vdash \theta \leftrightarrow ((\exists w)L) \wedge \delta$$

and it is now easier to eliminate this quantified variable.

Look at the shape of L. We have

$$s = S^n r$$

where r is 0 or another variable, and $n \in \mathbb{N}$. We need to consider whether $n \geq m$ or $n < m$. Setting

$$n = m + k \qquad m = n + k + 1$$

as appropriate, we have

$$T \vdash L \leftrightarrow (w \doteq S^k r) \qquad T \vdash L \leftrightarrow (S^{k+1} w \doteq r)$$

respectively. But

$$T \vdash (\exists w)[w \doteq S^k r] \leftrightarrow \mathsf{true} \qquad T \vdash (\exists w)[S^{k+1} w \doteq r] \leftrightarrow (r \not\doteq \ulcorner 0 \urcorner) \wedge \cdots \wedge (r \not\doteq \ulcorner k \urcorner)$$

where the second equivalence come from Lemma 2.5. In either case we see that $(\exists w)L$ is equivalent to a quantifier-free formula, and hence θ is equivalent to a quantifier-free formula.

This procedure works if there is at least one positive conjunct L. It remains to deal with the case where each conjunct is negative. In this case $\neg\theta$ is equivalent to

$$(\forall w)[M_1 \vee \cdots \vee M_l]$$

where each disjunct M has the shape

$$(S^m w \doteq S^n t)$$

where $m, n \in \mathbb{N}$ and the term t is 0 or a different variable.

Since $\neg\theta$ is universally quantified we may instantiate w by any numeral we please to obtain

$$T \vdash \neg\theta \rightarrow \beta$$

where β is a quantifier-free formula in which w does not occur. We may do this for a selection of numerals to obtain

$$T \vdash \neg\theta \rightarrow (\beta_0 \wedge \beta_1 \wedge \cdots \wedge \beta_l)$$

for appropriate $\beta_0, \beta_1, \ldots, \beta_l$. (The number of selections here, $1+l$, is deliberately chosen so that in a moment we may use a pigeon hole argument.) We show how to select the instantiating numerals so that

$$T \vdash (\beta_0 \wedge \beta_1 \wedge \cdots \wedge \beta_l) \rightarrow \mathsf{false}$$

and hence

$$T \vdash \theta$$

holds.

Consider any disjunct M. This has the shape

$$(S^m w \doteq S^n t)$$

for some $m, n \in \mathbb{N}$. Consider any $k \geq n$. We may set $w = \ulcorner k \urcorner$ so that the instantiated disjunct is equivalent to

$$(t \doteq \ulcorner a \urcorner)$$

for some $a \in \mathbb{N}$. In fact, $m + k = n + a$. In the same way, by setting $w = \ulcorner k+1 \urcorner$ this instantiation of the same disjunct is equivalent to

$$(t \doteq \ulcorner a+1 \urcorner)$$

for the same a as before.

By setting $w = \ulcorner k \urcorner$ for some sufficiently large k we obtain

$$T \vdash \neg\theta \rightarrow (t_1 \doteq \ulcorner a_1 \urcorner) \vee \cdots \vee (t_l \doteq \ulcorner a_l \urcorner)$$

for some $a_1, \ldots, a_l \in \mathbb{N}$.

By setting $w = \ulcorner k+1 \urcorner$ we obtain

$$T \vdash \neg\theta \rightarrow (t_1 \doteq \ulcorner a_1 + 1 \urcorner) \vee \cdots \vee (t_l \doteq \ulcorner a_l + 1 \urcorner)$$

for the same $a_1, \ldots, a_l \in \mathbb{N}$.

Repeating this for each of

$$w := \ulcorner k \urcorner, w := \ulcorner k+1 \urcorner, \dots, w := \ulcorner k+l \urcorner,$$

we obtain

$$T \vdash \neg\theta \to \gamma$$

where γ is

$$\begin{array}{ll} & (t_1 \doteq \ulcorner a_1 + 0 \urcorner) \vee \cdots \vee (t_l \doteq \ulcorner a_l + 0 \urcorner) \\ \wedge & \\ & (t_1 \doteq \ulcorner a_1 + 1 \urcorner) \vee \cdots \vee (t_l \doteq \ulcorner a_l + 1 \urcorner) \\ \wedge & \\ & \vdots \\ \wedge & \\ & (t_1 \doteq \ulcorner a_1 + l \urcorner) \vee \cdots \vee (t_l \doteq \ulcorner a_l + l \urcorner) \end{array}$$

for some $a_1, \dots, a_l \in \mathbb{N}$. This formula is a conjunction of disjunctions. We may rephrase it as a disjunction of conjunctions. Each such conjunction has the shape

$$\delta := (t_{j(0)} \doteq \ulcorner a_{j(0)} + 0 \urcorner) \wedge (t_{j(1)} \doteq \ulcorner a_{j(1)} + 1 \urcorner) \wedge \cdots \wedge (t_{j(l)} \doteq \ulcorner a_{j(l)} + l \urcorner)$$

where the indexes $j(0), j(1), \dots, j(l)$ are selected from $\{1, \dots, l\}$. Each conjunction arises from a different selection of indexes.

For each such δ there are $1+l$ indexes $j(\cdot)$ selected from a set of size l. Thus $j(r) = j(s) = j$ (say) for some $r \neq s$, and hence

$$T \vdash \delta \to (\ulcorner a_j + r \urcorner \doteq \ulcorner a_j + s \urcorner) \to (\ulcorner r \urcorner \doteq \ulcorner s \urcorner) \to \mathsf{false}$$

which leads to

$$T \vdash \gamma \to \mathsf{false}$$

as required. ■

An important by-product of this result is that we now have a complete axiomatization of the theory $Th(\mathfrak{N})$, and, what is more, we have got rid of the induction axioms. This axiomatization enables us to give a full description of all the structures $\mathfrak{A} \equiv \mathfrak{N}$. Which is nice.

Exercises

2.2 Prove Lemmas 2.3 and 2.5.

2.3 (a) Show that each model of T consists of a single copy of $\mathfrak{N}$ together with a family of disjoint copies of $(\mathbb{Z}, S, 0)$.

(b) Show that T is κ-categorical for each uncountable κ.

(c) Describe the spectrum of countable models of T.

2.4 (a) Show that

$$\mathfrak{N} \models \delta \Longleftrightarrow T \vdash \delta$$

holds for each quantifier-free sentence δ.

(b) Show that $T = Th(\mathfrak{N})$.

2.5 Exercise 2.4 shows that Definition 2.4 provides a simple axiomatization of $Th(\mathfrak{N})$. But Gödel's incompleteness theorem says there is no such axiomatization. Explain this.

[*Held in 120-.../B23-bit.. Last changed July 26, 2004*]

2.3 Lines

In subsection 2.2 we used a signature with two symbols, a constant symbol and a 1-placed operation symbol. In this subsection we use a signature with just one symbol, a 2-placed relation symbol which we write as an infix.

We look at structures

$$\mathfrak{A} = (A, \leq)$$

each of which is a dense linear order without end points. For short we call such a structure a line.

2.7 EXAMPLE. Both $\mathbb{Q}$ and $\mathbb{R}$ (with their natural orderings) are lines. In section 10 we see that $\mathbb{Q}$ is the only countable line (up to isomorphism). ■

A line (in this sense) is a linearly ordered set with no first point, no last point, and with no gaps. It is easy to see that these form an elementary class, by writing down the appropriate axioms. Furthermore, the theory T of this class is $\forall_2$-axiomatizable with a finite set of axioms. We need not write down all of these axioms, but we should look at some of them.

Each line $\mathfrak{A}$ is a poset which is linear. The distinguished attribute $\leq$ is reflexive, and (the universal closure of)

$$(u \leq v) \vee (v \leq u)$$

is the axiom which ensures linearity. It can be checked that

$$(v \not\leq u) \leftrightarrow (u \leq v) \wedge (u \neq v) \qquad (u \leq v) \leftrightarrow (v \not\leq v) \wedge (u \doteq v)$$

are consequence of this and the other axioms. It is convenient to let

$$u < v \quad \text{abbreviate} \quad v \not\leq u$$

and pretend that this is an atomic formula. In fact, we could axiomatize the class using a signature with two 2-placed relations $\leq$ and $<$, and add

$$(u < v) \leftrightarrow (v \not\leq u)$$

as an axiom. In some ways that is neater. Notice that

$$(u \neq v) \leftrightarrow (u < v) \vee (v < u)$$

is a consequence of these axioms.

So far we have used only the axioms of linearly ordered sets, and this leads to a useful observation.

2.8 LEMMA. *Relative to the theory of linearly ordered sets, each quantifier-free formula is equivalent to a* $\{\wedge, \vee\}$*-combination of formulas*

$$(u \leq v) \qquad (u < v) \qquad (u \doteq v)$$

for appropriate variables u, v.

In other words, provided we work in terms of both $\leq$ and $<$, then we can get rid of all uses of negation. In fact, we can go further. Since

$$(u \doteq v) \leq (u \leq v) \wedge (v \leq u)$$

we can get rid of all uses of the equality symbol.

All these axioms are $\forall_1$-sentences. However, a line has no first point and no last point, so both of

$$(\forall w)(\exists u)[u < w] \qquad (\forall w)(\exists v)[w < v]$$

are required as further axioms. It has no gaps, so

$$(\forall u, v)[(u < v) \rightarrow (\exists w)[(u < w) \wedge (w < v)]]$$

is another axiom. These extra axioms are $\forall_2$-sentences.

2.9 THEOREM. *The theory* T *of lines has EQ.*

Proof. How can we eliminate the bound quantifier from a formula

$$\theta := (\exists w)[L_1 \wedge \cdots \wedge L_l]$$

where the matrix is quantifier-free? From the remarks above we may assume that each conjunct has one of the shapes

$$(u \leq w) \quad (u < w) \qquad (w < v) \quad (w \leq v)$$

where w is the distinguished variable (we want to eliminate) and u, v are other variables. Of course, several other variables may occur throughout the conjuncts.

We need to consider three cases.

Suppose only conjuncts of the shapes

$$(u \leq w) \qquad (u < w)$$

occur. Then, since a line has no last point, we see that

$$T \vdash \theta$$

holds, and we are done.

Suppose only conjuncts of the shapes

$$(w < v) \qquad (w \leq v)$$

occur. Then, since a line has no first point, we see that

$$T \vdash \theta$$

holds, and we are done.

It remains to deal with the case where there is a mixture of 'left' and 'right' conjuncts. Consider all pairs of such conjuncts, one from the left and one from the right. There are four kinds of such pairs as listed below.

Pair	Replacement
$(u \leq w) \wedge (w \leq v)$	$(u \leq v)$
$(u \leq w) \wedge (w < v)$	$(u < v)$
$(u < w) \wedge (w \leq v)$	$(u < v)$
$(u < w) \wedge (w < v)$	$(u < v)$

We replace each such pair by the indicate formula, and let γ be the conjunction of these replacements. The variable w does not occur in γ. Remembering that a line has no gaps, we see that

$$T \vdash \theta \leftrightarrow \gamma$$

holds, and we are done. The implication $\theta \rightarrow \gamma$ is trivial, but the converse $\gamma \rightarrow \theta$ requires several moment's thought. ■

Consider any sentence σ. Then

$$T \vdash \sigma \leftrightarrow \delta$$

for some quantifier-free sentence δ. But, in this language, the only two quantifier-free sentences are true and false. Thus one of

$$T \vdash \sigma \qquad T \vdash \neg\sigma$$

must hold. This shows that T is complete.

Exercises

2.6 Show that, as linearly ordered sets, $\mathbb{Q} \equiv \mathbb{R}$.
What does this tell you about the (Dedekind) completeness of R?

2.7 Show that each linearly ordered set can be embedded in a line.

2.8 Show that the theory of lines is $\aleph_0$-categorical, but is not κ-categorical for any cardinal $\kappa > \aleph_0$.

2.9 Consider the class of all structures $(A, \leq, a, b)$ where $(A, \leq)$ is a dense linearly ordered set with first point a and last point b. Write down a set of axioms for this class, and show that the corresponding theory has EQ.

[*Held in 120../B24-bit.. Last changed July 26, 2004*]

2.4 Some other examples – *to be re-done*

There are several other theories of structures based on $\mathbb{N}$ or $\mathbb{Z}$ which have EQ. A discussion of these is given in pages 307 – 334 of [14]

Locate, or work out, the algorithm for the theory of two successors

Chapter 4 of [10] is devoted to EQ. It discuses several other theories of linear orderings with EQ, as well as some more sophisticated theories.

Pages 49 – 60 of [5] considers theories with EQ. The exercises give a a fairly long list of examples (without solutions).

It would be nice to have a comprehensive list of examples. I don't know of such a list, but presumably there is one somewhere.

Exercises–*needed*

[*Held in 120-../B30-bit.. Last changed July 26, 2004*]

3 Basic methods

We are now in a position to begin the development of model theory proper. This is the use of the compactness property to obtain information about structures and classes of structures.

[*Held in 120-../B31-bit.. Last changed July 26, 2004*]

3.1 Some semantic relations

We start with a description of the fundamental notions. We have met some of these earlier, but it is as well to have them defined all in one place.

3.1 DEFINITION. Let $\mathfrak{A}, \mathfrak{B}$ be structures (for the same language).
We write

$$\mathfrak{A} \equiv \mathfrak{B}$$

and say $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent if

$$\mathfrak{A} \models \sigma \Longleftrightarrow \mathfrak{B} \models \sigma$$

holds for each sentence σ (of the underlying language).
For each $n < \omega$ we write

$$\mathfrak{A} \equiv_n \mathfrak{B}$$

and say $\mathfrak{A}$ and $\mathfrak{B}$ are n-equivalent if

$$\mathfrak{A} \models \sigma \Longleftrightarrow \mathfrak{B} \models \sigma$$

holds for each $\forall_n$-sentence σ (of the underlying language). ■

It is easy to check that each of these relations is an equivalence relation. We often write

$$\equiv_\omega \quad \text{for} \quad \equiv \qquad\qquad \equiv_\infty \quad \text{for} \quad \cong$$

and this allows is to treat certain properties of a whole family of equivalence relations in one go.
Notice that if

$$\mathfrak{A} \models \sigma \Longrightarrow \mathfrak{B} \models \sigma$$

holds for each sentence then $\mathfrak{A} \equiv \mathfrak{B}$. Similarly if this implication holds for each quantifier-free sentence, then $\mathfrak{A} \equiv_0 \mathfrak{B}$ holds. However, a similar observation fails for $\equiv_n$ for $n \neq 0$.

3.2 DEFINITION. Let $\mathfrak{A}, \mathfrak{B}$ be structures (for the same language). For each $n < \omega$ we write

$$\mathfrak{A} \equiv\rangle(\forall_n)\, \mathfrak{B} \qquad \mathfrak{A} \equiv\rangle(\exists_n)\, \mathfrak{B}$$

if

$$\mathfrak{A} \models \sigma \Longrightarrow \mathfrak{B} \models \sigma$$

holds for each

$$\forall_n\text{-sentence} \qquad \exists_n\text{-sentence}$$

σ, respectively. ■

Thus $\equiv\rangle(\forall_0)$, $\equiv\rangle(\exists_0)$ and $\equiv_0$ are the same relation, but $\equiv\rangle(\forall_{n+1})$ and $\equiv\rangle(\exists_{n+1})$ are converse relations.

Recall that, from Definition 1.10, we have the notion

$$\mathfrak{A} \subseteq \mathfrak{B}$$

of one structure being a substructure of another (or one structure being a superstructure of another). We can refine this notion.

3.3 DEFINITION. Let $\mathfrak{A} \subseteq \mathfrak{B}$ be structures (for the same language).

We write

$$\mathfrak{A} \prec \mathfrak{B}$$

and say $\mathfrak{A}$ is an elementary substructure of $\mathfrak{B}$ if

$$\mathfrak{A} \models \phi(a) \Longleftrightarrow \mathfrak{B} \models \phi(a)$$

holds for each formula $\phi(v)$ (of the underlying language) and point a of $\mathfrak{A}$ which matches the free variables v of ϕ.

For each $n < \omega$ we write

$$\mathfrak{A} \prec_n \mathfrak{B}$$

and say $\mathfrak{A}$ is a n-substructure of $\mathfrak{B}$ if

$$\mathfrak{A} \models \phi(a) \Longleftrightarrow \mathfrak{B} \models \phi(a)$$

holds for each $\forall_n$-formula $\phi(v)$ (of the underlying language) and point a of $\mathfrak{A}$ which matches the free variables v of ϕ. ■

Thus $\prec_0$ and $\subseteq$ and are the same relation. As with $\equiv$ we sometimes write

$$\prec_\omega \quad \text{for} \quad \prec$$

so that we can treat certain properties of the relations together. (We don't have a relation $\prec_\infty$ since $\mathfrak{A} \prec_\infty \mathfrak{B}$ could only mean $\mathfrak{A} = \mathfrak{B}$.)

In subsection 1.2 we mentioned the informal notion of an isomorphism between two structure. We can now make that precise, and set up the more general notion of an embedding of one structure into another. As often happens this is rather tedious in the intial stages, but it becomes simpler later on.

3.4 DEFINITION. Let $\mathfrak{A}, \mathfrak{B}$ be structures (for the same language).

An embedding

$$\mathfrak{A} \xrightarrow{\quad f \quad} \mathfrak{B}$$

(from $\mathfrak{A}$ to $\mathfrak{B}$) is a function

$$f : A \longrightarrow B$$

between the carriers, as indicated, which is injective and is such that the following hold.

- For each constant symbol K of the signature

$$fa = b$$

where

$$a = \mathfrak{A}[\![K]\!] \qquad b = \mathfrak{B}[\![K]\!]$$

are the input and output of f.

- For each n-placed relation symbol R of the signature

$$\mathfrak{A}[\![R]\!]a_1 \cdots a_n \Longleftrightarrow \mathfrak{B}[\![R]\!]b_1 \cdots b_n$$

for each $a_1, \ldots, a_n$ from $\mathfrak{A}$ and $b_i = fa_i$ for each $1 \leq i \leq n$.

- For each n-placed operation symbol O of the signature

$$f(\mathfrak{A}[\![O]\!]a_1 \cdots a_n) = \mathfrak{B}[\![O]\!]b_1 \cdots b_n$$

for each $a_1, \ldots, a_n$ from $\mathfrak{A}$ and $b_i = fa_i$ for each $1 \leq i \leq n$.

There is one clause for each symbol of the signature. ■

This looks a bit complicated, but it is saying nothing more than the range $f[A]$ of f is (the carrier of a) substructure of $\mathfrak{B}$ and f is an ismorphism from $\mathfrak{A}$ to this substructure. In particular, an isomorphism is a surjective embedding.

Exercises

3.1 Show that for each $r \in \mathbb{N} \cup \{\omega, \infty\}$, the relation $\equiv_r$ is an equivalence.
Show that for each $r, s \in \mathbb{N} \cup \{\omega, \infty\}$,

$$\mathfrak{A} \equiv_s \mathfrak{B} \Longrightarrow \mathfrak{A} \equiv_r \mathfrak{B}$$

for $r \leq s$ (using the obvious comparison).

3.2 Show that

$$\mathfrak{A} \prec_r \mathfrak{B} \Longrightarrow \mathfrak{A} \equiv\rangle(\exists_{r+1})\, \mathfrak{B}$$

(for each $r < \omega$).

[`Held in 120../B32-bit..  Last changed July 26, 2004`]

3.2 The diagram technique

We concluded subsection 3.1 with the rather tedious definition of an embedding between structures. Using quantfier-free formulas we can characterize this notion in a more convenient way. (Before you read this proof you might like to have a look at Exercise 1.6.)

3.5 LEMMA. *Let $\mathfrak{A}, \mathfrak{B}$ be a pair of structures, and let*

$$f : A \longrightarrow B$$

be a function between the carriers. Then f is an emebedding if and only if

$$\mathfrak{A} \models \delta(a_1, \ldots, a_n) \iff \mathfrak{B} \models \delta(b_1, \ldots, b_n)$$

holds for each quantifier-free formula $\delta(v_1, \ldots, v_n)$ and elements $a_1, \ldots, a_n$ from $\mathfrak{A}$ and where $b_i = fa_i$ for each $1 \leq i \leq n$.

Proof. Suppose first that this semantic equivalence does hold. Then, using the atomic formulas

$$\begin{array}{ll} \delta(v, w) & := (v \doteq w) \\ \delta(v) & := (v \doteq K) \\ \delta(v_1, \ldots, v_n) & := Rv_1 \cdots v_n \\ \delta(v_0, v_1, \ldots, v_n) & := (v_0 \doteq Ov_1, \ldots, v_n) \end{array}$$

we see that f is injective and the various required signature clauses hold.

Conversely, suppose f is an embedding. By definition, we have the required equivalence for each atomic formula δ, and a simple induction gives it for each quantifier-free formula. ■

This result indicates how we can refine the notion of an embedding.

3.6 DEFINITION. Let $\mathfrak{A}, \mathfrak{B}$ be a pair of structures, and let

$$f : A \longrightarrow B$$

be a function between the carriers. Then, for each $r < \omega$ or $r = \omega$, f is a $\prec_r$-embedding if and only

$$\mathfrak{A} \models \phi(a_1, \ldots, a_n) \iff \mathfrak{B} \models \phi(b_1, \ldots, b_n)$$

holds for each $\forall_r$-formula $\phi(v_1, \ldots, v_n)$ and elements $a_1, \ldots, a_n$ from $\mathfrak{A}$ and where $b_i = fa_i$ for each $1 \leq i \leq n$. ■

Lemma 3.5 shows that a $\prec_0$-embedding is nothing more than an embedding. As r increases the notion of a $\prec_r$-embedding becomes more and more restrictive, until the case $r = \omega$ gives the notion of an elementary embedding.

Given two structures $\mathfrak{A}, \mathfrak{B}$, when can we say $\mathfrak{A}$ is embeddable in $\mathfrak{B}$? In other words, when does an embedding

$$\mathfrak{A} \longrightarrow \mathfrak{B}$$

exits? (We might want a restricted embedding, but for now let's stick to this simple case.)

So far, at any one time we have been concerned with just one parent language, and for the most part this has been left in the background. We are now going to use two laguages, one of which is an enrichment of the other. This must be done with some care, and since this is the first time we have used this trick, we will take it slowly.

Let L be the parent language (so that $\mathfrak{A}$ and $\mathfrak{B}$, above are L-structures). We will enrich L by adding a family of new constant symbols to the signature. This will generate a larger language L'

The new constant symbols are called the parameters of the enriched language L' (to distinguish them from the old constant symbols of L which also occur in L'). Furthermore, each paramater has an intended interpretation.

Let a be an enumeration of a part of $\mathfrak{A}$. We enrich L by adding a name for each one of the elements in a. These names are the parameters. In practice, we do not distinguish between an element a in $\mathfrak{A}$ and its name which is added to L. In particular, we write $L(\mathsf{a})$ for the enriched language. Furthermore, when a enumerates the whole of $\mathfrak{A}$, as it often does, we write $L(\mathfrak{A})$ for the enriched language.

Of course, this rather sloppy convention has some obvious pit-falls. However these are easily avoided once we become familiar with the convention and know how it is used. (The alternative is worse, and attractive only to the anally retentive.)

The idea is that $L(\mathsf{a})$ is designed to talk about $\mathfrak{A}$, and specifically about the part enumerated by a. Thus the structure $(\mathfrak{A}, \mathsf{a})$ (which is formed from $\mathfrak{A}$ by distinguishing each element in a) is a kind of canonical structure for $L(\mathsf{a})$. Of course, this language can talk about other structures of the form $(\mathfrak{B}, \mathsf{b})$ where $\mathfrak{B}$ is an L-strucure and b is an enumeration of certain elements of $\mathfrak{B}$.

3.7 DEFINITION. Let $\mathfrak{A}$ be an L-structure, let a be an enumeration of a part of $\mathfrak{A}$, and consider the enriched language $L(\mathsf{a})$.

The a-diagram

$$Diag(\mathfrak{A}, \mathsf{a})$$

of $\mathfrak{A}$ is the set of quantfier-free $L(\mathsf{a})$-sentences which hold in $(\mathfrak{A}, \mathsf{a})$.

When a enumerates the whole of $\mathfrak{A}$ we call $Diag(\mathfrak{A}, \mathsf{a})$ the diagram of $\mathfrak{A}$. ■

What does a sentence in $Diag(\mathfrak{A}, \mathsf{a})$ look like? Each such sentence has the shape

$$\delta(a_1, \ldots, a_n)$$

where

$$\delta(v_1, \ldots, v_n)$$

is a quantifier-free L-formula in the indicated variable, and $a_1, \ldots, a_n$ are taken from a. If we collapse these two finite list to a batch v and point a, then each member of $Diag(\mathfrak{A}, \mathsf{a})$ has the shape $\delta(a)$.

In this way we can flit between the two languages L and $L(\mathsf{a})$, and creatively confuse syntax with semantics. The following result illustrates how this is used.

3.8 LEMMA. *Let $\mathfrak{A}, \mathfrak{B}$ be two L-structures (for some language L). Let a be an enumeration of the whole of $\mathfrak{A}$. Then $\mathfrak{A}$ is embeddable in $\mathfrak{B}$ if and only if*

$$(\mathfrak{B}, \mathsf{b}) \models Diag(\mathfrak{A}, \mathsf{a})$$

for some enumeration b of a part of $\mathfrak{B}$.

Once the above ideas has been absorbed, this result is almost a triviality. The enumeration b gives the range of the required embedding, and the matching of a with b gives the required function.

Let's now move towards a similar result that is slightly more interesting.

3.9 DEFINITION. Let T be a L-theory, and let $\mathfrak{A}$ be a L-structure (for some language L. We set

$$T[\mathfrak{A}] = T \cup Diag(\mathfrak{A}, \mathsf{a})$$

where a is an enumeration of the whole of $\mathfrak{A}$ to produce a set of $L(\mathfrak{A})$-sentences. ■

Strictly speaking, $T[\mathfrak{A}]$ depends on which particular enumeration a is used. However, this will never be a cause for concern. Notice that although T is a L-theory, it is *not* a $L(\mathfrak{A})$-theory. Notice also that $T[\mathfrak{A}]$ need not be consistent.

3.10 DEFINITION. Let T be a L-theory, and let $\mathfrak{A}$ be a L-structure (for some language L).

(m) We say $\mathfrak{A}$ is a model of T if $\mathfrak{A} \models T$. We write $\mathcal{M}(T)$ for the class of models of T.

(s) We say $\mathfrak{A}$ is a submodel of T if $\mathfrak{A} \subseteq \mathfrak{B}$ for some $\mathfrak{B} \models T$. We write $\mathcal{S}(T)$ for the class of submodels of T. ■

Thus $\mathcal{M}(T) \subseteq \mathcal{S}(T)$. We will look at various other subclasses of $\mathcal{S}(T)$.

3.11 LEMMA. *Let T be a L-theory, and let $\mathfrak{A}$ be a L-structure (for some language L. The set $T[\mathfrak{A}]$ of $L(\mathfrak{A})$-sentences is consistent if and only if $\mathfrak{A} \in \mathcal{S}(T)$.*

Proof. If $\mathfrak{A} \in \mathcal{S}(T)$ then $\mathfrak{A} \subseteq \mathfrak{B}$ for some $\mathfrak{B} \models T$, and this provides the required model of $T[\mathfrak{A}]$.

Conversely, suppose $T[\mathfrak{A}]$ is consistent, and let $(\mathfrak{B}, \mathsf{b})$ be any model. Then $(\mathfrak{B}, \mathsf{b}) \models T$, and hence $\mathfrak{B} \models T$ (since the parameters don't appear in T). Since $(\mathfrak{B}, \mathsf{b}) \models Diag(\mathfrak{A}, \mathsf{a})$, there is an embedding $\mathfrak{A} \longrightarrow \mathfrak{B}$. By replacing $\mathfrak{B}$ by a suitable isomorphic copy, we may suppose $\mathfrak{A} \subseteq \mathfrak{B}$, and hence $\mathfrak{A} \in \mathcal{S}(T)$. ■

This kind of argument can be rather tedious. However, for the most part it only needs to be done once. From now on we can use without mention any of the tricks explained so far in this subsection.

To conclude, we have the promised slightly more interesting result.

3.12 DEFINITION. A theory T has JEP (the joint embedding property) if for each pair $\mathfrak{A}, \mathfrak{B}$ of models of T, there is a wedge of embeddings

$$\begin{tikzcd} \mathfrak{A} \arrow[dr] & \\ & \mathfrak{C} \\ \mathfrak{B} \arrow[ur] & \end{tikzcd}$$

to some model $\mathfrak{C}$ of T. ■

This structural property of the class of models of a theory corresponds to a certain primeness property of the corresponding consquence relation.

3.13 THEOREM. *A (consistent) theory T has JEP if and only if*

$$T \vdash \alpha \vee \beta \Longrightarrow T \vdash \alpha \text{ or } T \vdash \beta$$

holds for each pair α, β of $\forall_1$-sentences (of the underlying language).

Proof. Suppose that T has the primeness property, consider any pair $\mathfrak{A}, \mathfrak{B}$ of models of T. Let a be a enumeration of the whole of $\mathfrak{A}$, and let b be a enumeration of the whole $\mathfrak{B}$. Add these enumeration to the underlying laguage L to form a doubly enriched language $L(\mathsf{a}, \mathsf{b})$. Thus $L(\mathsf{a}, \mathsf{b})$ is the amalgam of the two languages $L(\mathfrak{A})$ and $L(\mathfrak{B})$. We may assume that the two sets of new parameters are disjoint. To produce the joint embedding it suffices to show that

$$T \cup Diag(\mathfrak{A}, \mathsf{a}) \cup Diag(\mathfrak{B}, \mathsf{b})$$

is consistent.

By way of contradiction, suppose this set is not consistent. Then, by compactness, it has a finite subset which is inconsistent. This gives a sentence $\gamma(a)$ from $Diag(\mathfrak{A}, \mathsf{a})$ and a sentence $\delta(b)$ from $Diag(\mathfrak{B}, \mathsf{b})$ such that

$$T \cup \{\gamma(a), \delta(b)\}$$

is inconsistent. Thus we have

$$\mathfrak{A} \models \gamma(a) \qquad \mathfrak{B} \models \delta(b) \qquad T \vdash \neg\gamma(a) \vee \neg\delta(b)$$

for some quantifier-free formulas $\gamma(v), \delta(w)$ of the underlying language, and points a from $\mathfrak{A}$ and b from $\mathfrak{B}$ matching the free variables v and w, respectively.

We use the third condition to show that

$$T \vdash (\forall v)\neg\gamma(v) \vee (\forall w)\neg\delta(w)$$

holds.

Consider any model $\mathfrak{D}$ of T. Consider any pair d, e of points of $\mathfrak{D}$ where d matches v and e matches w. We have $(\mathfrak{D}, d, e) \models T$, and hence

$$\mathfrak{D} \models \neg\gamma(d) \vee \neg\delta(e)$$

holds. (This involves interpreting the parameters that name a by d, and the parameters that name b by e.) Since d and e are arbitrary and do not occur in T, this gives

$$\mathfrak{D} \vdash (\forall v)\neg\gamma(v) \vee (\forall w)\neg\delta(w)$$

and hence

$$T \vdash (\forall v)\neg\gamma(v) \vee (\forall w)\neg\delta(w)$$

since $\mathfrak{D}$ is an arbitrary model of T.

Both γ and δ are quantifier-free formulas, so the assumed primeness property gives that one of

$$T \vdash (\forall v)\neg\gamma(v) \qquad T \vdash (\forall w)\neg\delta(w)$$

holds. But $\mathfrak{A}, \mathfrak{B} \models T$ and both

$$\mathfrak{A} \models \gamma(a) \qquad \mathfrak{B} \models \delta(b)$$

which is the required contradiction.

The proof of the converse implication is an easy exercise. ■

We made rather heavy weather of this proof by describing almost every detail. We did this to show exactly how it works. In future, similar proofs will not be so detailed. Also, this is not the simplest example of a use of the diagram technique, but it's probably the best one to start with. If you are not happy with it, try to write out the full details of the proof of Lemma 3.16 in the next subsection.

3.3 Let $\mathfrak{A}, \mathfrak{B}$ be arbitrary structures (for the same language).
Show that for each $n < \omega$,

$$\mathfrak{A} \equiv\rangle(\forall_{n+1})\, \mathfrak{B}$$

holds if and only if

$$\mathfrak{A} \xrightarrow{\ f\ } \mathfrak{C} \qquad\qquad \mathfrak{B} \xrightarrow{\ g\ } \mathfrak{C}$$

for some $\prec$-embedding f and some $\prec_n$-embedding g to a common structure $\mathfrak{C}$.

[*Held in 120-../B33-bit.. Last changed July 26, 2004*]

3.3 Restricted axiomatization

A set of axioms for a theory T is a subset $\Sigma \subseteq T$ such that

$$\sigma \in T \Longleftrightarrow \Sigma \vdash \sigma$$

holds for each sentence σ. Every theory T has a set of axioms, for we can always take $\Sigma = T$. However, we usually want a more interesting set of axioms.

3.14 DEFINITION. Let T be a theory (in some language L). For each $n < \omega$, we say T is $\forall_n$-axiomatizable is T is the set of all consequences of some set of $\forall_n$-sentences (of L). ■

We are most often concerned with $\forall_1$-axiomatizability and $\forall_2$-axiomatizability, but the general notion is useful. In practice we can often tell that a theory is $\forall_n$-axiomatizable by simple observing the axioms we have written down for it. On other occasions this isn't helpful, either because we have a set of axioms that is too complicated, or we do not even have a set of axioms.

3.15 EXAMPLE. Most theories that arise in algebra are $\forall_2$-axiomatizable or better. For instance, by choosing the right signature the theory of rings is $\forall_1$-axiomatizable and the theory of fields is $\forall_2$-axiomatizable. However, is it possible that within this signature, the theory of fields has a set of $\forall_1$-axioms?

Most of the axioms of Peano arithmetic are $\forall_2$ (and even $\forall_1$). However, the usual induction axioms uses sentences of arbitrary quantifier complexity. Is there some way by which we can replace these axioms by ones of bounded quantifier complexity? ■

How can we test a theory T for $\forall_n$-axiomatizability? Observe that if T is $\forall_n$-axiomatizable, then $T \cap \forall_n$ is such a set of axioms. Thus we need to investigate how $\mathcal{M}(T)$ and $\mathcal{M}(T \cap \forall_n)$ are related. Here is the answer for $n = 1$.

3.16 LEMMA. *Let T be a theory (in some language L). For each structure $\mathfrak{A}$ (suitable for L), the following*

$$(i)\quad \mathfrak{A} \in \mathcal{S}(T) \qquad\qquad (ii)\quad \mathfrak{A} \models T \cap \forall_1$$

are equivalent.

Proof. The implication $(i) \Rightarrow (ii)$ is immediate, so we may concentrate on the implication $(ii) \Rightarrow (i)$.

Suppose $\mathfrak{A} \models T \cap \forall_1$, and let a be some enumeration of $\mathfrak{A}$. We enrich L to $L(\mathfrak{A})$ by adding a as new parameters. It suffices to show that the set

$$T \cup Diag(\mathfrak{A}, \mathsf{a})$$

of $L(\mathfrak{A})$-sentences is consistent.

By way of contradiction, suppose the set $T \cup Diag(\mathfrak{A}, \mathsf{a})$ is not consistent. Then, by compactness, some finite subset does dot have a model, and hence

$$T \cup \{\delta(a)\}$$

is inconsistent for some quantifier-free L-formula $\delta(v)$ and some point a of $\mathfrak{A}$ where $\mathfrak{A} \models \delta(a)$. In other words

$$T \vdash \neg\delta(a)$$

holds within the language $L(\mathfrak{A})$. But a does not occur in T, so that

$$T \vdash \neg\delta(v)$$

holds within the language L, and hence

$$T \vdash (\forall v)\neg(\delta(v)$$

holds. By (ii) this gives $\mathfrak{A} \models (\forall v)\neg\delta(v)$, which is the contradiction. ■

This result gives us a characterization of $\forall_1$-axiomatizability.

3.17 THEOREM. *A theory T (in some language L) is $\forall_1$-axiomatizable if any only if each submodel of T is a model.*

Proof. Suppose first that T is $\forall_1$-axiomatizable. Then, in fact, $T \cap \forall_1$ is a set of axioms for T, and hence

$$\mathfrak{A} \models T \cap \forall_1 \Longrightarrow \mathfrak{A} \models T$$

holds (for all structures $\mathfrak{A}$). In particular,

$$\mathfrak{A} \in \mathcal{S}(T) \Longrightarrow \mathfrak{A} \models T \cap \forall_1 \Longrightarrow \mathfrak{A} \models T$$

holds, to show that $\mathcal{S}(T) = \mathcal{M}(T)$.

Conversely, suppose $\mathcal{S}(T) = \mathcal{M}(T)$. Then, by Lemma 3.16, we have

$$\mathfrak{A} \models T \cap \forall_1 \Longrightarrow \mathfrak{A} \in \mathcal{S}(T) = \mathcal{M}(T) \Longrightarrow \mathfrak{A} \models T$$

(for each structure $\mathfrak{A}$) to show that $T \cap \forall_1$ is set of axioms of T. ■

The same kind of argument gives us an interpolation result.

3.18 THEOREM. *For each theory T and sentences λ, ρ (in the same language), the following are equivalent.*

(i) There is an $\forall_1$-sentence σ show that both

$$T \vdash \lambda \to \sigma \qquad T \vdash \sigma \to \rho$$

hold.

(ii) The implication

$$\mathfrak{B} \models \lambda \Longrightarrow \mathfrak{A} \models \rho$$

for all models $\mathfrak{A}, \mathfrak{B}$ of T with $\mathfrak{A} \subseteq \mathfrak{B}$.

Proof. Only the implication $(ii) \Rightarrow (i)$ offers much resistance. To prove this let T' be the deductive closure of $T \cup \{\lambda\}$ and let $\Sigma = T' \cap \forall_1$. Thus, Σ is the set of all $\forall_1$-sentence σ such that

$$T \vdash \lambda \to \sigma$$

holds, and it suffices to show that

$$T \cup \Sigma \vdash \rho$$

holds.

To this end, consider any model $\mathfrak{A}$ of $T \cup \Sigma$. By Lemma 3.16 there is some model $\mathfrak{B} \models T'$ with $\mathfrak{A} \subseteq \mathfrak{B}$. Both $\mathfrak{A}$ and $\mathfrak{B}$ are models of T and $\mathfrak{B} \models \lambda$ (since $\lambda \in T'$). Thus, by (ii), we have $\mathfrak{A} \models \rho$, as required. ■

The last few results are concerned with $\forall_1$-sentences. All of them can be generalized to give analogues using $\forall_{n+1}$-sentences for $n < \omega$. These results are dealt with in the exercises.

The diagram technique is quite versitile and is used many times in model theory. Here is an application that produces a different kind of result. This is a variant of the result of Exercise 3.3.

3.19 LEMMA. *Let $\mathfrak{A} \subseteq \mathfrak{B}$ be a pair of structures (of the same signature). The following are equivalent.*

(i) $\mathfrak{A} \prec_1 \mathfrak{B}$

(ii) There is some structure $\mathfrak{C}$ with $\mathfrak{B} \subseteq \mathfrak{C}$ and $\mathfrak{A} \prec \mathfrak{C}$.

(iii) There is some structure $\mathfrak{C}$ with $\mathfrak{B} \subseteq \mathfrak{C}$ and $\mathfrak{A} \prec_1 \mathfrak{C}$.

Proof. The two implications $(ii) \Rightarrow (iii)$ and $(iii) \Rightarrow (i)$ are immediate. The implication $(i) \Rightarrow (ii)$ is the content of this result. To prove this consider any pair $\mathfrak{A} \prec_1 \mathfrak{B}$. Let a be an enumeration of $\mathfrak{A}$ and let b be an enumeration of $\mathfrak{B}$. (The elements of $\mathfrak{A}$ will be enumerated twice, once in a and once in b, but this doesn't matter.) It suffices to show that

$$Th(\mathfrak{A}, \mathsf{a}) \cup Diag(\mathfrak{B}, \mathsf{a}, \mathsf{b})$$

is consistent.

(Notice that this involves two enrichments of the underlying language; a first one by the addition of a, and then a further one by the addition of b.)

If this set is not consistent then

$$Th(\mathfrak{A}, \mathsf{a}) \vdash \neg\delta(b, a)$$

for some quantifier-free formula $\delta(w, v)$ (of the underlying language) some point a of $\mathfrak{A}$ and some point b of $\mathfrak{B}$. Since b does not occur in (the language of) $Th(\mathfrak{A}, \mathsf{a})$, we have

$$Th(\mathfrak{A}, \mathsf{a}) \vdash \neg\delta(w, a)$$

and hence

$$Th(\mathfrak{A}, \mathsf{a}) \vdash (\forall w)\neg\delta(w, a)$$

so that

$$\mathfrak{A} \models (\forall w)\neg\delta(w, a)$$

holds. But now (i) gives $\mathfrak{B} \models (\forall w)\neg\delta(w, a)$, which leads to a contradiction. ■

This result can be generalized to produce a characterization of the relation $\prec_{n+1}$ between structures.

Exercises

3.4 Let T be a theory (in some language L). Let $n < \omega$.

(a) Show that a structure $\mathfrak{A}$ is a model of $T \cap \forall_{n+1}$ if any only if there is a model $\mathfrak{B} \models T$ with $\mathfrak{A} \prec_n \mathfrak{B}$.

(b) Show that T is $\forall_{n+1}$-axiomatizable if and only if

$$\mathfrak{A} \prec_n \mathfrak{B} \models T \Longrightarrow \mathfrak{A} \models T$$

holds (for all structures $\mathfrak{A}, \mathfrak{B}$).

(c) Show that for each sentence λ, μ the following are equivalent.

(i) There is an $\forall_{n+1}$-sentence σ such that both $T \vdash \lambda \to \sigma$ and $T \vdash \sigma \to \rho$ hold.

(ii) The implication

$$\mathfrak{B} \models \lambda \Longrightarrow \models \rho$$

holds for all models $\mathfrak{A}, \mathfrak{B}$ of T with $\mathfrak{A} \prec_n \mathfrak{B}$.

3.5 Let $\mathfrak{A} \subseteq \mathfrak{B}$. Show that for each $n < \omega$, there is a structure $\mathfrak{C}$ such that

$$\mathfrak{B} \prec_n \mathfrak{C} \qquad \mathfrak{A} \prec \mathfrak{C}$$

if and only if $\mathfrak{A} \prec_{n+1} \mathfrak{B}$.

[*Held in 120-../B34-bit.. Last changed July 26, 2004*]

3.4 Directed families of structures

In this subsection we look at a method of combining many different structures into one structure. We then use this construction to obtain further results on the axiomatizability of theories.

3.20 DEFINITION. Let $\mathcal{A}$ be a non-empty family of structures (for a common language).

(a) The family is directed if for each $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$, there is some $\mathfrak{C} \in \mathcal{A}$ such that both $\mathfrak{A} \subseteq \mathfrak{C}$ and $\mathfrak{B} \subseteq \mathfrak{C}$ hold.

(b) The family is a chain if for each $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$, either $\mathfrak{A} \subseteq \mathfrak{B}$ or $\mathfrak{B} \subseteq \mathfrak{A}$ holds.

(c) The family is a λ-chain (for some ordinal λ) if

$$\mathcal{A} = \{\mathfrak{A}_i \,|\, i < \lambda\}$$

where $\mathfrak{A}_i \subseteq \mathfrak{A}_j$ for all $i \leq j < \lambda$.

(d) The family is an ω-chain if

$$\mathcal{A} = \{\mathfrak{A}_i \,|\, i < \omega\}$$

where $\mathfrak{A}_i \subseteq \mathfrak{A}_j$ for all $i \leq j < \omega$. ■

Clearly, every ω-chain is an example of a λ-chain (for the case $\lambda = \omega$), and every λ-chain is a chain. A few moment's thought shows that every chain is directed.

Our first job is to show that for each directed family $\mathcal{A}$ we can construct a structure $\bigcup \mathcal{A}$ which extends each member of $\mathcal{A}$, and is as small as possible. Thus we produce a certain set U and then furnish this to obtain a structure $\mathfrak{U}$ such that $\mathfrak{A} \subseteq \mathfrak{U}$ for each $\mathfrak{A} \in \mathcal{A}$.

3.21 CONSTRUCTION. Let $\mathcal{A}$ be a directed family of structures, as above. We produce $\mathcal{U} = \bigcup \mathcal{A}$ as follows.

- Consider first the carriers of the members of $\mathcal{A}$. Let U be the union of all these sets. Thus each $a \in U$ is also a member of at least one $\mathfrak{A} \in \mathcal{A}$. It may be in at least two members, say $\mathfrak{A}$ and $\mathfrak{B}$, where neither of these extends the other. However, since $\mathcal{A}$ is directed, there is some $\mathfrak{C} \in \mathcal{A}$ which extends both $\mathfrak{A}$ and $\mathfrak{B}$, and then a is a member of $\mathfrak{C}$. We will use this trick several times.

- Consider any constant K of the language. Each $\mathfrak{A} \in \mathcal{A}$ has a canonical interpretation $a = \mathfrak{A}[\![K]\!]$ of this constant, to give a member of U. How do these various members compare? Consider two members $\mathfrak{A}, \mathfrak{B}$ of $\mathcal{A}$ with interpretations $a = \mathfrak{A}[\![K]\!]$ and $b = \mathfrak{B}[\![K]\!]$ of the constant. We know there is some $\mathfrak{C} \in \mathcal{A}$ which extends both $\mathfrak{A}$ and $\mathfrak{B}$. Thus, working in $\mathfrak{C}$ we have

  $$a = \mathfrak{A}[\![K]\!] = \mathfrak{C}[\![K]\!] = \mathfrak{B}[\![K]\!] = b$$

  to show that a, b are the same element of U. In other words, this constant K determines a unique element of U. We let this be the interpretaion of K.

- Consider any n-placed relation symbol R if the language. We must decide whether or not $[\![R]\!]a_1 \cdots a_n$ is true for each $a_1, \ldots, a_n \in U$.

  Each a_i is an element of at least one member $\mathfrak{A}_i$ of $\mathcal{A}$. By repeated use of the directedness we see there is some $\mathfrak{A} \in \mathcal{A}$ which extends each of these $\mathfrak{A}_i$. Thus all of $a_1, \ldots, a_n$ are members of some single $\mathfrak{A} \in \mathcal{A}$. In that structure $\mathfrak{A}[\![R]\!]a_1 \cdots a_n$ is either true or false. However, as yet this truth value may depend on the holding structure $\mathfrak{A}$ used.

Consider any $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$, both of which contain $a_1, \ldots, a_n$. There is some $\mathfrak{C} \in \mathcal{A}$ which extends $\mathfrak{A}$ and $\mathfrak{B}$. Then

$$\mathfrak{A}[\![R]\!]a_1 \cdots a_n = \mathfrak{C}[\![R]\!]a_1 \cdots a_n = \mathfrak{B}[\![R]\!]a_1 \cdots a_n$$

to show that $[\![R]\!]a_1 \cdots a_n$ is independent of the choice of holding structure. This gives a single truth value for $[\![R]\!]a_1 \cdots a_n$.

- Consider any n-placed operation symbol O if the language. We must assign a value to $[\![O]\!]a_1 \cdots a_n$ for each $a_1, \ldots, a_n \in U$. As in the previous case we can find a common holding structure $\mathfrak{A} \in \mathcal{A}$ which contains each of $a_1, \ldots, a_n$. This structure gives a value $\mathfrak{A}[\![O]\!]a_1 \cdots a_n$ in $\mathfrak{A}$, but we need to check that this is independent of the choice of $\mathfrak{A}$. The argument for this is similar to the relation case.

This structure $\mathfrak{U} = \bigcup \mathcal{A}$ is the union of the directed system $\mathcal{A}$. Notice that $\mathfrak{A} \subseteq \mathfrak{U}$ holds for each $\mathfrak{A} \in \mathcal{A}$. ■

The construction $\mathcal{A} \longmapsto \bigcup \mathcal{A}$ has some preservation properties.

3.22 LEMMA. *Let $\mathcal{A}$ be a directed family of structures with union $\mathfrak{U}$. Then*

$$\mathfrak{U} \models Th(\mathcal{A}) \cap \forall_2$$

holds.

Proof. Consider any $\forall_2$-sentence σ with $\mathcal{A} \models \sigma$. We must show that $\mathfrak{U} \models \sigma$ holds. We have

$$\sigma = (\forall u)\phi(u)$$

for some $\exists_1$-formula $\phi(u)$ and list u of variables. Consider any point a of $\mathfrak{U}$ which matches u. We must show that $\mathfrak{U} \models \phi(a)$. Since $\mathcal{A}$ is directed there is at least one $\mathfrak{A} \in \mathcal{A}$ which contains all of a. But $\mathfrak{A} \models \sigma$ (by the choice of σ), and hence $\mathfrak{A} \models \phi(a)$. Also, $\mathfrak{A} \subseteq \mathfrak{U}$, and hence, since ϕ is $\exists_1$, we have $\mathfrak{U} \models \phi(a)$, as required. ■

Within a directed system $\mathcal{A}$ there are many inclusions between its members. Some of these may have stronger preservation properties. In an extreme case each of these inclusions may be elementary.

3.23 LEMMA. *Let $\mathcal{A}$ be a directed family of structures with union $\mathfrak{U}$. Suppose*

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec \mathfrak{B}$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$. Then

$$\mathfrak{A} \prec \mathfrak{U}$$

holds for all $\mathfrak{A} \in \mathcal{A}$.

Proof. We show each of the following

[n] For each $\forall_{2n}$-formula $\phi(u)$ and each $\mathfrak{A} \in \mathcal{A}$, the implication

$$\mathfrak{A} \models \phi(a) \Longrightarrow \mathfrak{U} \models \phi(a)$$

holds for all point a of $\mathfrak{A}$ (matching the free variables u of ϕ).

by induction on n.

The base case, $n = 0$, is trivial.

For the induction step, $n \mapsto n+1$, consider any $\forall_{2n+2}$-formula $\phi(u)$. We know that ϕ is

$$(\forall v)(\exists w)\psi(w, v, u)$$

for some $\forall_{2n}$-formula $\phi(w, v, u)$. Consider any $\mathfrak{A} \in \mathcal{A}$ and any point a of $\mathfrak{A}$ for which

$$\mathfrak{A} \models \phi(a)$$

holds. We must show that

$$\mathfrak{U} \models \phi(a)$$

holds. To this end, consider any point b of $\mathfrak{U}$ which matches v, we must produce some point c of $\mathfrak{U}$ which matches w for which

$$\mathfrak{U} \models \psi(c, b, a)$$

holds.

The point b comes from $\mathfrak{B} \in \mathcal{A}$, and then there is some $\mathfrak{C} \in \mathcal{A}$ which extends both $\mathfrak{A}$ and $\mathfrak{B}$, so that both a and b are points of $\mathfrak{C}$. We have

$$\mathfrak{A} \models \phi(a) \qquad \mathfrak{A} \prec \mathfrak{C}$$

so that $\mathfrak{C} \models \phi(a)$, and hence $\mathfrak{C} \models (\exists w)\psi(w, b, a)$ holds. This gives some point c of $\mathfrak{C}$ with $\mathfrak{C} \models \psi(c, b, a)$. The induction hypthesis $[n]$ now gives $\mathfrak{U} \models \psi(c, b, a)$, as required. ■

With this result we can obtain a characterization of $\forall_2$-axiomatizability that is rather different from the earlier characterization.

3.24 THEOREM. *For each theory T (in some language L) the following are equivalent.*

(i) T is $\forall_2$-axiomatizable.

(ii) $\mathcal{M}(T)$ is closed under unions of directed systems.

(iii) $\mathcal{M}(T)$ is closed under unions of ω-chains.

Proof. $(i) \Rightarrow (ii)$. This follows by Lemma 3.22.

$(ii) \Rightarrow (iii)$. This is trivial.

$(iii) \Rightarrow (i)$. Assuming (iii), consider any model $\mathfrak{A} \models T \cap \forall_2$. It is sufficient to show that $\mathfrak{A} \models T$. By Exercise 3.4 (which generalizes Lemma 3.16) we have $\mathfrak{A} \prec_1 \mathfrak{B}$ for some $\mathfrak{B} \models T$. Thus, by Lemma 3.19 we have

$$\mathfrak{A} \subseteq \mathfrak{B} \subseteq \mathfrak{A}' \quad \mathfrak{B} \models T \quad \mathfrak{A} \prec \mathfrak{A}'$$

for some structures $\mathfrak{B}, \mathfrak{A}'$. Notice that $\mathfrak{A}' \models T \cap \forall_2$.

By iterating this construction we obtain two ω-chains

$$\mathcal{A} = \{\mathfrak{A}_i \mid i < \omega\} \qquad \mathcal{B} = \{\mathfrak{B}_i \mid i < \omega\}$$

where $\mathfrak{A}_0 = \mathfrak{A}$ and

$$\mathfrak{A}_i \subseteq \mathfrak{B}_i \subseteq \mathfrak{A}_{i+1} \quad \mathfrak{B}_i \models T \quad \mathfrak{A}_i \prec \mathfrak{A}_{i+1}$$

for each $i < \omega$.

Each of these chains has a union. But the two chains interlace, so there is a single structure $\mathfrak{U}$ which is the union of both $\mathcal{A}$ and $\mathcal{B}$. By construction, $\mathcal{B}$ is a chain of models of T, and hence (*iii*) gives $\mathfrak{U} \models T$. Again by construction, $\mathcal{A}$ is a chain of elementary embedding, and hence Lemma 3.23 gives

$$\mathfrak{A} \prec \mathfrak{U} \models T$$

so that $\mathfrak{A} \models T$, as required. ■

To conclude this section we look at a situation which is slightly unusual in model theory. We look at intersections of structures.

Given any family $\mathcal{A}$ is structures (for some language) we can make sense of the intersection $\bigcap \mathcal{A}$. This is either empty, or a substructure of each of the members of $\mathcal{A}$. Note also, that if the underlying language contains at least one constant symbol, then $\bigcap \mathcal{A}$ can not be empty.

We show how to exhibit the union of a directed family of models as a binary intersection. This uses a different kind of enrichment of the underlying language.

3.25 LEMMA. *Let T be a theory (in some language L). Let $\mathcal{A}$ be a directed family of models of T, and let $\mathfrak{U}$ be $\bigcup \mathcal{A}$. Then there are models $\mathfrak{B}, \mathfrak{C}, \mathfrak{D}$ of T such that*

$$\mathfrak{U} \subseteq \mathfrak{B} \prec \mathfrak{D} \quad \mathfrak{U} \subseteq \mathfrak{C} \subseteq \mathfrak{D} \quad \mathfrak{U} = \mathfrak{B} \cap \mathfrak{C}$$

hold.

Proof. By Lemma 3.22 we have $\mathfrak{U} \models T \cap \forall_2$, and then Exercise 3.4 gives some model $\mathfrak{B}$ with $\mathfrak{U} \prec_1 \mathfrak{B} \models T$. Our problem is to construct $\mathfrak{C}$ and $\mathfrak{D}$.

Let b be an enumeration of $\mathfrak{B}$, and consider $Th(\mathfrak{B}, \mathsf{b})$ in the enriched language. We now form a further enrichment by adding a new 1-placed relation symbol R. Within the language $L + \{R\}$ let Γ be the set of sentences

> R is the carrier of a substructure (of the parent structure),
> and this substructure is a model of T

(which isn't too hard to formalize). Using this consider the set

$$Th(\mathfrak{B}, \mathsf{b}) \cup \Gamma \cup \{Ra \mid a \text{ from } \mathfrak{U}\} \cup \{\neg Rb \mid b \text{ from } \mathfrak{B} \text{ but not from } \mathfrak{U}\}$$

of sentences of the largest language. Any finite subset of this is a subset of

$$Th(\mathfrak{B}, \mathsf{b}) \cup \Gamma \cup \{Ra_1, \dots, Ra_m\} \cup \{\neg Rb \mid b \text{ from } \mathfrak{B} \text{ but not from } \mathfrak{U}\}$$

for some elements $a_1, \dots, a_m$ of $\mathfrak{U}$. These elements all belong to some $\mathfrak{A} \in \mathcal{A}$ And the setting $R = A$ (the carrier of $\mathfrak{A}$) produces a model of this smaller set. Thus the larger set is consistent.

Let $(\mathfrak{D}, C, \mathsf{b})$ be any model of this larger set. From $Th(\mathfrak{B}, \mathsf{b})$ we have $\mathfrak{B} \prec \mathfrak{D}$. From Γ the subset C is the carrier of a substructure $\mathfrak{C} \subseteq \mathfrak{D}$ which is a model of T. From the third component we have $\mathfrak{U} \subseteq \mathfrak{C}$, and the fourth component gives $\mathfrak{B} \cap \mathfrak{C} \subseteq \mathfrak{U}$, which is what we want. ■

Occasionally we meet a theory of the following kind.

3.26 DEFINITION. A theory T is convex if

$$\mathfrak{A} \cap \mathfrak{B} \models T$$

holds for all models $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$ of T with $\mathfrak{A} \subseteq \mathfrak{C}$, and $\mathfrak{B} \subseteq \mathfrak{C}$, and $\mathfrak{A} \cap \mathfrak{B} \neq \emptyset$. ■

The final result of this subsection is rather surprising.

3.27 THEOREM. *Each convex theory T is $\forall_2$-axiomatizable.*

Proof. By Lemma 3.25 the class $\mathcal{M}(T)$ is closed under unions of directed families. ■

We will use directed families of structures many times.

Exercises

3.6 Let $\mathcal{A}$ be a directed family of structures with union $\mathfrak{U}$. Show that if, for some $n < \omega$,

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_n \mathfrak{B}$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$, then

$$\mathfrak{A} \prec_n \mathfrak{U}$$

holds for all $\mathfrak{A} \in \mathcal{A}$.

3.7 Let $\mathcal{A}$ be a directed family of structures with union $\mathfrak{U}$. Suppose that for each pair $\mathfrak{A} \subseteq \mathfrak{B}$ of members of $\mathcal{A}$, there is some $\mathfrak{B} \subseteq \mathfrak{C} \in \mathcal{A}$ with $\mathfrak{A} \prec \mathfrak{C}$. Show that, under these circumstances, $\mathfrak{A} \prec \mathfrak{U}$ for each $\mathfrak{A} \in \mathcal{A}$.

3.8 Two directed families $\mathcal{A}$ and $\mathcal{B}$ of structures interlace if both

- For each $\mathfrak{A} \in \mathcal{A}$, there is some $\mathfrak{B} \in \mathcal{B}$ with $\mathfrak{A} \subseteq \mathfrak{B}$
- For each $\mathfrak{B} \in \mathcal{B}$, there is some $\mathfrak{A} \in \mathcal{A}$ with $\mathfrak{B} \subseteq \mathfrak{A}$

hold. Show that

$$\bigcup \mathcal{A} = \bigcup \mathcal{B}$$

for such families.

[*Held in 120../B35-bit.. Last changed July 26, 2004*]

3.5 The up and down techniques

In this final subsection of this section we obtain two consequences of refined compactness, Theorem 1.26, which are used all the time in model theory, often without even mentioning it. The results allow us to move between cardinalities, provided we stay above the cardinality of the underlying language.

3.28 THEOREM. *Let L be a language, let $\mathfrak{A}$ be a L-structure, and let κ_d, κ_u be cardinals such that*

$$|L| \le \kappa_d \le |\mathfrak{A}| \le \kappa_u$$

hold. Then there are structures $\mathfrak{A}_d, \mathfrak{A}_u$ (for L) such that

$$|\mathfrak{A}_d| = \kappa_d \qquad \mathfrak{A}_d \prec \mathfrak{A} \prec \mathfrak{A}_u \qquad |\mathfrak{A}_u| = \kappa_u$$

hold.

These are often referred to as the

downward Löwenheim-Skolem   upward Löwenheim-Skolem

results, respectively, but that can be a bit of a mouthful.

How can we produce these two structures $\mathfrak{A}_d$ and $\mathfrak{A}_u$? Surprisingly, the larger one is easier to construct.

The up technique

Given a structure $\mathfrak{A}$ for a language L and a cardinal κ, how can we produce a structure $\mathfrak{A} \prec \mathfrak{B}$ with $|\mathfrak{B}| = \kappa$? Clearly, to do this we need $|\mathfrak{A}| \le |\mathfrak{B}| = \kappa$, but there are other restrictions as well.

We use the elementary (full) version of the diagram technique.

Let a be an enumeration of the whole of $\mathfrak{A}$. Let b be an enumeration of a further κ new parameters. Unlike the members of a, these further parameters are not associated with any elements of $\mathfrak{A}$. We work in the enrichment $L(\mathsf{a}, \mathsf{b})$ of L. Note that

$$|L(\mathsf{a}, \mathsf{b})| = |L| + |\mathsf{a}| + |\mathsf{b}| = |L| + \kappa$$

since $|\mathsf{a}| = |\mathfrak{A}| \le \kappa = |\mathsf{b}|$. We assume $|L| \le \kappa$ to get $|L(\mathsf{a}, \mathsf{b})| = \kappa$.

Look at the set of sentences

$$Th(\mathfrak{A}, \mathsf{a}) \cup \text{ `the members of } \mathsf{b} \text{ are pairwise distinct'}$$

of $L(\mathsf{a}, \mathsf{b})$. Observe that this set is finitely satisfiable in $\mathfrak{A}$. You should check the details of this and note that only

$$\aleph_0 \le |\mathfrak{A}|$$

is needed.

Refined compactness gives a model $(\mathfrak{B}, \mathsf{a}, \mathsf{b})$ of this set where $|\mathfrak{B}| = |L(\mathsf{a}, \mathsf{b})| = \kappa$. The first component ensures that $\mathfrak{A} \prec \mathfrak{B}$, which is what we want.

You should observe the potential of this technique. In fact, we have already proved more than was asked for.

3.29 THEOREM. *Let L be a language, let $\mathfrak{A}$ be in infinite L-structure, and let κ be a cardinal with $|L| \le \kappa, |\mathfrak{A}| \le \kappa$. Then there is a structure $\mathfrak{B}$ (for L) such that $\mathfrak{A} \prec \mathfrak{B}$ and $|\mathfrak{B}| = \kappa$.*

As a further refinment we might want to arrange that the structure $\mathfrak{B}$ has several other properties. Maybe this can be done by controlling these properties by sentences of the enriched language.

The down technique

Given a structure $\mathfrak{A}$ for a language L and a cardinal κ, how can we produce a structure $\mathfrak{B} \prec \mathfrak{A}$ with $|\mathfrak{B}| = \kappa$? Clearly, to do this we need $\kappa = |\mathfrak{B}| \leq |\mathfrak{A}|$, but there are other restrictions as well.

To produce $\mathfrak{B}$ we work in an enrichment of L, but this time we adjoin certain new operation symbols.

Recall that a subset of B of the carrier A of $\mathfrak{A}$ is the carrier of a substructure $\mathfrak{B} \subseteq \mathfrak{A}$ if and only if B is closed under the distinguished attributes of $\mathfrak{A}$. In general, this closure does not ensure $\mathfrak{B} \prec \mathfrak{A}$. However it does if $Th(\mathfrak{A})$ has EQ, so we attempt to arrange this.

We 'skolemize' the language by adjoining certain skolem operation symbols (which are then interpreted as skolem functions).

Consider an arbitrary formula

$$\phi(v, u)$$

of the language L, where u is a list of n variables and v is a single nominated target variable. For each such formula and nomination we select a new n-placed operation symbol $f = f_\phi$, and look at the sentence

$$\text{Skol}(\phi) \qquad (\forall u)[(\exists v)\phi(v, u) \to \phi(fu, u)]$$

in the enriched language. We do this for each formula ϕ and nominated target variable v, using a different skolem operation symbol f for each such pair. Let L' be the enrichment of L obtained by adjoining each of the skolem operation symbols. Notice that $|L'| = |L|$, so that although L' has many more formulas, its cardinality is the same as L.

Let S' be the theory in L' axiomatized by the set of all $\text{Skol}(\phi)$.

Consider any L-structure $\mathfrak{A}$. We show how to enrich this to a L'-structure by adjoining an interpretation $\mathfrak{A}'[\![f]\!]$ for each skolem operation symbol f. This, of course, will be a skolem function.

Let $\phi(v, u)$ be the formula associated with f. Consider any point a of $\mathfrak{A}$ which matches u. We require a value $\mathfrak{A}'[\![f]\!]a$. There are two cases to consider.

Suppose $\mathfrak{A} \models (\exists v)\phi(v, a)$. Then $\mathfrak{A} \models \phi(b, a)$ for some element b of $\mathfrak{A}$, and we let $\mathfrak{A}'[\![f]\!]a$ be any such element.

Suppose $\mathfrak{A} \models \neg(\exists v)\phi(v, a)$. Then we let $\mathfrak{A}'[\![f]\!]a$ be any element whatsoever.

This ensures that $\mathfrak{A}' \models \text{Skol}(\phi)$. In fact

$$\mathfrak{A}' \models (\forall u)[(\exists v)\phi(v, u) \leftrightarrow \phi(fu, u)]$$

holds.

At first sight it seems that $Th(\mathfrak{A}')$ has EQ (for, as we have just seen, the operation symbol f can be used to eliminate certain quantifications $(\exists v)$). Another look reveals an error. In $Th(\mathfrak{A}')$ we can eliminated certain quantifiers, but only from formulas of the original language L. As yet, we can not eliminate quantifiers from all formulas of L', for such a formula may not have an associated skolem operation symbol. This is easy to correct.

We iterate the process

$$L \subseteq L' \subseteq L' \subseteq \cdots \subseteq L^{(r)} \subseteq \cdots$$

to produce an ascending ω-chain of languages each of cardinality $|L|$. Let L^* be the union of these languages. Thus $|L^*| = |L|$.

Each L-structure $\mathfrak{A}$ can be enriched, step by step, to a L^*-structure $\mathfrak{A}^*$. Furthermore, by accumulation, we see that the following holds.

> For each formula $\phi(v, u)$ of L^* where v is a single variable, there is an operation symbol f of L^* such that
>
> $$\mathfrak{A}^* \models (\forall u)[(\exists v)\phi(v, u) \leftrightarrow \phi(fu, u)]$$
>
> holds.

(The formula ϕ lives in some $L^{(r)}$, and then f is in $L^{(r+1)}$.) This ensures that $Th(\mathfrak{A}^*)$ has EQ.

With this we can produce the required substructure of $\mathfrak{A}$.

Suppose $|L| \leq \kappa \leq |\mathfrak{A}|$. Let $\mathfrak{A}^*$ be a skolem enrichment of $\mathfrak{A}$, as described above. Take any subset $X \subseteq A$ (the carrier of $\mathfrak{A}$) with $|X| \leq \kappa$. Let B be the closure of X in A under all the distinguished attributes of $\mathfrak{A}^*$. Note that $|B| = \kappa$. Furthermore, B is the carrier of a substructure $\mathfrak{B}^* \subseteq \mathfrak{A}^*$. Even more, since $Th(\mathfrak{A}^*)$ has EQ, we have $\mathfrak{B}^* \prec \mathfrak{A}^*$. Let $\mathfrak{B}$ be the L-structure obtained from $\mathfrak{B}^*$ by removing the extra skolem operations. We have $\mathfrak{B} \prec \mathfrak{A}$ and $|\mathfrak{B}| = \kappa$, as required.

Again this technique has quite a bit of potential. We have already proved the following.

3.30 THEOREM. *Let L be a language, let $\mathfrak{A}$ be a L-structure, and let κ be a cardinal with $|L| \leq \kappa \leq |\mathfrak{A}|$. for each subset X of $\mathfrak{A}$ with $|X| \leq \kappa$, there is a structure $\mathfrak{B} \prec \mathfrak{A}$ with $|\mathfrak{B}| \leq \kappa$ such that X is a subset of $\mathfrak{B}$.*

Perhaps, with some effort, we can make these skolem functions do other things.

Exercises

3.9 Let $\mathfrak{A} \subseteq \mathfrak{B}$ be a pair of structure and suppose the following holds. For each formula $\phi(v, u)$ with a nominated variable v, if $\mathfrak{B} \models (\exists v)\phi(v, a)$ for some point a from $\mathfrak{A}$, then $\mathfrak{B} \models \phi(b, a)$ for some element b from $\mathfrak{A}$ (not just $\mathfrak{B}$). Show that $\mathfrak{A} \prec \mathfrak{B}$.

[*Held in 120-../B40-bit.. Last changed July 26, 2004*]

4 Model complete and submodel complete theories

In Definition 2.1 we introduced the notion of a theory T having EQ (elimination of quantifiers). This looks like a rather syntactic idea. In subsections 2.2 and 2.3 we gave two examples of such theories. The methods of proof used there look rather intricate, and special to those particular theories. In this section we will show there is a rather general condition which ensures that a theory has EQ. Later [*say where*] we will show how these conditions can be verified by an analysis of the spectrum of models of a theory.

[*Held in 120-../B41-bit.. Last changed July 26, 2004*]

4.1 Model complete theories

Given a theory T and a submodel $\mathfrak{A} \in \mathcal{S}(T)$ we can form a consistent set $T[\mathfrak{A}]$ of sentences in the $\mathfrak{A}$-enriched langauge $L(\mathfrak{A})$ of the parent language L. This set $T[\mathfrak{A}]$ generates a (consistent) $L(\mathfrak{A})$-theory, and we can ask that this enrichment of T has a certain property. We can ask for such a property for a whole family of submodels.

4.1 DEFINITION. A theory T is model complete if the enriched theory (axiomatized by) $T[\mathfrak{A}]$ is complete for each $\mathfrak{A} \in \mathcal{M}(T)$. ■

This is the original definition of model completeness but, as the next result shows, the notion can be characterized in a rather more amenable fashion.

4.2 LEMMA. *A theory T is model complete if and only if*

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec \mathfrak{B}$$

holds for each pair $\mathfrak{A}, \mathfrak{B}$ of models of T.

Almost invariably it is this characterization that we use. Only occasionally do we return to the official Definition 4.1.

(You might want to think about what happens if we add to T the diagram of a *sub*model of T, or a special kind of submodel. We will return to this idea several times.)

Model completemess is one of the most important notions of model theory. Here is why.

4.3 THEOREM. *For each theory T the following are equivalent.*

(i) Each formula is T-equivalent to an $\forall_1$-formula.

(ii) Each formula is T-equivalent to an $\exists_1$-formula.

(iii) Each $\forall_1$-formula is T-equivalent to an $\exists_1$-formula.

(iv) The implication

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_1 \mathfrak{B}$$

holds for each pair $\mathfrak{A}, \mathfrak{B}$ of models.

(v) T is model complete.

Proof. The implication $(i) \Rightarrow (ii)$ holds by a simple use of negated formulas.

$(ii) \Rightarrow (iii)$. This is trivial.

$(iii) \Rightarrow (iv)$. This is immediate.

$(iv) \Rightarrow (v)$. Suppose (iv) holds. For each models $\mathfrak{A} \subseteq \mathfrak{B}$ we have $\mathfrak{A} \prec_1 \mathfrak{B}$ and hence, by Lemma 3.19, we have

$$\mathfrak{A} \prec \mathfrak{A}' \qquad \mathfrak{B} \subseteq \mathfrak{A}'$$

for some model $\mathfrak{A}'$ of T. But now we have $\mathfrak{B} \prec_1 \mathfrak{A}'$ so that a repeat of the argument gives

$$\mathfrak{B} \prec \mathfrak{B}' \qquad \mathfrak{A}' \subseteq \mathfrak{B}'$$

for some model $\mathfrak{B}'$ of T. By iterating this construction we produce two interlacing elementary chains

$$\mathcal{A} = \{\mathfrak{A}_i \mid i < \omega\} \qquad \mathcal{B} = \{\mathfrak{B}_i \mid i < \omega\}$$

with

$$\mathfrak{A}_i \prec \mathfrak{A}_{i+1} \quad \mathfrak{A}_i \subseteq \mathfrak{B}_i \subseteq \mathfrak{A}_{i+1} \subseteq \mathfrak{B}_{i+1} \quad \mathfrak{B}_i \prec \mathfrak{B}_{i+1}$$

for each $i < \omega$.

Since the chains interlace, they have a common union $\mathfrak{U}$.

Since the chains are elementary we have

$$\mathfrak{A} = \mathfrak{A}_0 \prec \mathfrak{U} \qquad \mathfrak{B} = \mathfrak{B}_0 \prec \mathfrak{U}$$

and hence $\mathfrak{A} \prec \mathfrak{B}$, as required.

$(v) \Rightarrow (i)$. Consider any formula $\phi(v)$ where v is the list of free variables of ϕ. We enrich the underlying language with parameters a matching the list v. Within this enriched language let $\Psi(a)$ be the set of $\forall_1$-sentences $\psi(a)$ such that

$$T \vdash \phi(a) \rightarrow \psi(a)$$

holds. If $T(a)$ is the theory in the enriched language axiomatized by $T \cup \{\phi(a)\}$, then $\Psi(a) = T(a) \cap \forall_1$. It suffices to show that

$$T \cup \Psi(a) \vdash \phi(a)$$

holds.

Consider any model $(\mathfrak{A}, a)$ of $T \cup \Psi(a)$. By Lemma 3.16, there is some model $(\mathfrak{B}, a)$ of $T(a)$ with $(\mathfrak{A}, a) \subseteq (\mathfrak{B}, a)$. But now

$$\mathfrak{A} \models T \quad \mathfrak{A} \subseteq \mathfrak{B} \quad \mathfrak{B} \models T \quad \mathfrak{B} \models \phi(a)$$

so that, by the assumption (iv), we have $\mathfrak{A} \prec \mathfrak{B}$ and hence $\mathfrak{A} \models \phi(a)$, which leads to the required result. ■

This result indicates that model completeness is a weakened form of quantifiers elimination. Thus the property EQ must be model completness together with some other property. What can this extra property be?

Exercises

4.1 Prove Lemma 4.2.

4.2 (a) Show that if a theory has EQ then it is model complete.
(b) Show that a model complete theory is $\forall_2$-axiomatizable.
(c) Show that if a model complete theory is $\forall_1$-axiomatizable, then it has EQ.

[`Held in 120-../B42-bit..  Last changed July 26, 2004`]

4.2 The amalgamation property

Recall that in subsection 3.1 we introduced the notion of a theory with JEP. This is a condition on the whole family of models of the theory, but we saw that it is equivalent to a kind of primeness of the $\forall_1$-part of the theory.

We now need a similar, but more intricate, property of the family of models.

4.4 DEFINITION. A theory T has AP (the amalgamation property) if for each wedge of embedding between models of T, as to the left of Table 1, there is a model $\mathfrak{D}$ of T and a commuting square of embeddings, as to the right of the table, which closes the wedge. ■

Notice that in Table 1 we have not given the embeddings names. Sometimes we need

$$\begin{tikzcd}[column sep=small] & \mathfrak{B} & \\ \mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & & \\ & \mathfrak{C} & \end{tikzcd} \qquad\qquad \begin{tikzcd}[column sep=small] & \mathfrak{B} \arrow[dr, "h"] & \\ \mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & & \mathfrak{D} \\ & \mathfrak{C} \arrow[ur, "k"'] & \end{tikzcd}$$

these as above and in in the first result. Both JEP and AP are concerned with embeddings between structures. However, sometimes we have embeddings with extra preservation properties, or require embeddings with extra prop1gerties.

4.5 LEMMA. *Given a wedge, as to the left above, where f is a $\prec_1$-embedding and g is an embedding, there is a commuting square, as to the right above where h is an embedding and k is an elementary embedding.*

Proof. By replacing $\mathfrak{B}$ and $\mathfrak{C}$ by suitable isomorphic copies, we can suppose that

$$\mathfrak{A} \prec_1 \mathfrak{B} \qquad \mathfrak{A} \subseteq \mathfrak{C} \qquad \mathfrak{B} \cap \mathfrak{C} = \mathfrak{A}$$

hold. Let $\mathsf{a}, \mathsf{b}, \mathsf{c}$ be enumerations of $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$, respectively. It suffices to show that

$$Th(\mathfrak{C}, \mathsf{c}, \mathsf{a}) \cup Diag(\mathfrak{B}, \mathsf{b}, \mathsf{a})$$

is consistent.

If this is not consistent then there is some quantifier-free formula $\delta(v, u)$ (of the underlying language) and appropriate points a, b such that

$$Th(\mathfrak{C}, \mathsf{c}, \mathsf{a}) \vdash \neg\delta(b, a) \qquad \mathfrak{B} \models \delta(b, a)$$

$$
\begin{tikzcd}[column sep=small, row sep=small]
 & \mathfrak{B} \\
\mathfrak{A} \arrow[ur] \arrow[dr] & \\
 & \mathfrak{C}
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}[column sep=small, row sep=small]
 & \mathfrak{B} \arrow[dr] & \\
\mathfrak{A} \arrow[ur] \arrow[dr] & & \mathfrak{D} \\
 & \mathfrak{C} \arrow[ur] &
\end{tikzcd}
$$

Table 1: A wedge and completed wedge

hold. The first of these gives

$$Th(\mathfrak{C}, \mathsf{c}, \mathsf{a}) \vdash (\forall v)\neg\delta(v, a)$$

and hence $\mathfrak{C} \models (\forall v)\neg\delta(v, a)$. But then $\mathfrak{A} \models (\forall v)\neg\delta(v, a)$, so that $\mathfrak{B} \models (\forall v)\neg\delta(v, a)$, and hence $\mathfrak{B} \models \neg\delta(b, a)$, which is a contradiction. ■

For a model complete theory, all embeddings between models are elementary, so we have the following consequence.

4.6 COROLLARY. *If a theory is model complete then it has AP.*

We now turn to the syntactic characterization of AP. This is similar to the characterization of JEP (as given in Theorem 3.13) except we must deal with formulas, not just sentences.

Remember that a set Γ is a type if

$$\partial\Gamma = \bigcup\{\partial\phi \mid \phi \in \Gamma\}$$

is finite. Given such a type we set

$$\neg\Gamma = \{\neg\phi \mid \phi \in \Gamma\}$$

which, of course, is also a type with $\partial(\neg\Gamma) = \partial\Gamma$.

4.7 THEOREM. *For each theory T the following are equivalent.*

(i) T has AP.

(ii) For each pair ψ, ϕ of $\forall_1$-formulas with

$$T \vdash \psi \vee \phi$$

there are $\exists_1$-formulas λ, ρ such that

$$T \vdash \lambda \vee \rho \quad T \vdash \lambda \to \psi \quad T \vdash \rho \to \phi$$

hold.

Proof. (i)$\Rightarrow$(ii). Consider any $\forall_1$-formulas ψ, ϕ such that

$$T \vdash \psi \vee \phi$$

holds. Let v be the list of free variables occurring in these formulas. We use the sets

$$\Sigma(T, \psi) \qquad \Sigma(T, \phi)$$

of $\exists_1$-formulas λ, ρ such that

$$T \vdash \lambda \rightarrow \psi \qquad T \vdash \rho \rightarrow \phi$$

hold (and $\partial\lambda \cup \partial\rho \subseteq v$). Both these sets are $\exists_1$-types, and so the two sets of negations $\neg\Sigma(T, \psi), \neg\Sigma(T, \phi)$ are $\forall_1$-types.

We will show that

$$T \cup \neg\Sigma(T, \psi) \cup \neg\Sigma(T, \phi)$$

is inconsistent.

How does this help? Observe that both $\Sigma(T, \psi)$ and $\Sigma(T, \phi)$ are closed under disjunctions. Thus, assuming the set above is inconsistent, we obtain $\lambda \in \Sigma(T, \psi)$ and $\rho \in \Sigma(T, \phi)$ such that $T \cup \{\neg\lambda, \neg\rho\}$ is inconsistent. In other words, $T \vdash \lambda \vee \rho$, and we are done.

It remains to show that the displayed set is inconsistent. This is where the assumption (i) is needed.

We make a preliminary observation.

Suppose we have a model $\mathfrak{A} \models T$ and a point a of $\mathfrak{A}$ which realizes $\neg\Sigma(T, \psi)$. A simple argument show that

$$T \cup \mathit{Diag}(\mathfrak{A}, \mathsf{a}) \cup \{\neg\psi(a)\}$$

is consistent (where a is an enumeration of $\mathfrak{A}$). You should make sure you can produce this argument. A use of compactness gives us an embedding of $\mathfrak{A}$ into some model $\mathfrak{B} \models T$ where $\mathfrak{B} \models \neg\psi(a)$.

With this we can complete the whole argument.

By way of contradiction, suppose the displayed set is consistent. Thus it is realized by some point a of some model $\mathfrak{A}$ of T. Using the two components we obtain a wedge of embeddings between models of T, as to the left of Table 1, where

$$\mathfrak{B} \models \neg\psi(a) \qquad \mathfrak{C} \models \neg\phi(a)$$

holds. Since T has AP (by (i)) this gives us a commuting square of embeddings between models of T, as to the right of that table. Both $\neg\psi$ and $\neg\phi$ are $\exists_1$-formulas, and hence $\mathfrak{D} \models \neg(\psi \vee \phi)(a)$. This is the contradiction (since $T \vdash \psi \vee \phi$).

(ii)$\Rightarrow$(i). Suppose (ii) holds and consider any wedge of embeddings between models of T, as to the left of Table 1. By replacing $\mathfrak{B}$ and $\mathfrak{C}$ by suitably isomorphic copies, we may suppose

$$\mathfrak{A} \subseteq \mathfrak{B} \quad \mathfrak{A} \subseteq \mathfrak{C} \quad \mathfrak{B} \cap \mathfrak{C} = \mathfrak{A}$$

hold. Let $\mathsf{a}, \mathsf{b}, \mathsf{c}$ be enumerations of $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$ respectively. It suffices to show that

$$T \cup \mathit{Diag}(\mathfrak{B}, \mathsf{b}, \mathsf{a}) \cup \mathit{Diag}(\mathfrak{C}, \mathsf{c}, \mathsf{a})$$

is consistent.

By way of contradiction, suppose this set is not consistent. Then we have quantifier-free formulas $\beta(v,u), \gamma(w,u)$ (of the underlying language) and points a, b, c such that

$$T \vdash \neg\beta(v,u) \vee \neg\gamma(w,u) \qquad \mathfrak{B} \models \beta(b,a) \quad \mathfrak{C} \models \gamma(c,a)$$

hold. The lists of variables u, v, w are disjoint so we have

$$T \vdash \psi \vee \phi$$

where

$$\psi(u) = (\forall v)\neg\beta(v,u) \qquad \phi(u) = (\forall w)\neg\gamma(w,u)$$

are the two components. Both of these are $\forall_1$-formulas, so that (ii) gives

$$T \vdash \lambda \vee \rho \qquad T \vdash \lambda \rightarrow \psi \quad T \vdash \rho \rightarrow \phi$$

for some $\exists_1$-formulas λ, ρ. The first of these gives $\mathfrak{A} \models (\lambda \vee \rho)(a)$, and hence $\mathfrak{A} \models \lambda(a)$, say. Since $\mathfrak{A} \subseteq \mathfrak{B}$, this gives $\mathfrak{B} \models \lambda(a)$, and hence $\mathfrak{B} \models \psi(a)$ follows from the second. But now $\mathfrak{B} \models \neg\beta(b,a)$, which is the required contradiction. ■

You will need time to digest this proof.

Exercises

4.3 Show that a theory T has AP if and only if $T \cap \forall_2$ has AP.

4.4 Let T be a theory such that $T \cap \forall_1$ has AP. Show that for each $\exists_1$-formula θ and $\forall_1$-formula ϕ with $T \vdash \theta \rightarrow \phi$, there is a quantifier-free formula δ (with the same free variables) such that both $T \vdash \theta \rightarrow \delta$ and $T \vdash \delta \rightarrow \phi$ hold.

4.5 The proof of Theorem 4.7 can be modified to produce a characterization of certain kinds of submodels of a theory.

A submodel $\mathfrak{A}$ of a theory T is an amalgamation base for T if for each wedge of embedding from $\mathfrak{A}$ to models of T, as to the left of Table 1, there is a model $\mathfrak{D}$ of T and a commuting square of embeddings, as to the right of that table, closing the given embeddings. Thus T has AP exactly when each model is an amalgamation base for T.

Show that a submodel $\mathfrak{A}$ of T is an amalgamation base for T if and only if for each pair ψ, ϕ of $\forall_1$-formulas with

$$T \vdash \psi \vee \phi$$

the structure $\mathfrak{A}$ omits the $\forall_1$-type

$$\neg\Sigma(T,\psi) \cup \neg\Sigma(T,\phi)$$

(where these are as used in the proof of Theorem 4.7).

4.6 For each theory T let $\mathcal{B}(T)$ be the class of amalgamation bases for T. Show the following.

(a) $\mathcal{B}(T) = \mathcal{A}(T \cap \forall_1)$.
(b) If $\mathfrak{A} \prec_1 \mathfrak{A}'$ for some $\mathfrak{A}' \in \mathcal{B}(T)$, then $\mathfrak{A} \in \mathcal{A}(T)$.
(c) $\mathcal{B}(T)$ is closed under unions of directed families.

[*Held in 120-../B43-bit.. Last changed July 26, 2004*]

4.3 Submodel complete theories

In subsection 3.2 we saw the idea of moving from a theory T to a family of enriched theories $T[\mathfrak{A}]$ for various structures $\mathfrak{A}$. Model completeness is concerned with these enrichments for $\mathfrak{A} \in \mathcal{M}(T)$. We now go to the other extreme, and consider what happens when we adjoin the diagram of *submodels* of T.

4.8 DEFINITION. A theory T is submodel complete if the enriched theory $T[\mathfrak{A}]$ is complete for each $\mathfrak{A} \in \mathcal{S}(T)$. ■

Trivially, each submodel complete theory is model complete, but it should have further properties. The following result explains everything.

4.9 THEOREM. *For each theory T, the following are equivalent.*

(i) T is submodel complete.

(ii) T is model complete and (the theory axiomatized by) $T \cap \forall_1$ has AP.

(iii) The implication

$$(\mathfrak{A}, a) \equiv_0 (\mathfrak{B}, b) \Longrightarrow (\mathfrak{A}, a) \equiv (\mathfrak{B}, b)$$

holds for all models $\mathfrak{A}, \mathfrak{B}$ of T and points a from $\mathfrak{A}$ and b from $\mathfrak{B}$.

(iv) The implication

$$(\mathfrak{A}, a) \equiv_0 (\mathfrak{B}, b) \Longrightarrow (\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b)$$

holds for all models $\mathfrak{A}, \mathfrak{B}$ of T and points a from $\mathfrak{A}$ and b from $\mathfrak{B}$.

(v) For each $\forall_1$-formula ϕ, there is a quantifier-free formula δ such that

$$T \vdash \phi \leftrightarrow \delta$$

holds.

(vi) T has EQ.

Proof. (i)⇒(ii). This is straight forward.

(ii)⇒(iii). Assuming (ii), consider any situation

$$(\mathfrak{A}, a) \equiv_0 (\mathfrak{B}, b)$$

where $\mathfrak{A}, \mathfrak{B} \in \mathcal{M}(T)$. The point a of $\mathfrak{A}$ need not enumerate a substructure of $\mathfrak{A}$, but it certainly generates a substructure. (We merely take a and close off under the distinguished operations of $\mathfrak{A}$.) In the same way, b generates a substructure of $\mathfrak{B}$. The given relationship between a and b ensures that these two substructures are isomorphic. Let $\mathfrak{C}$ be an isomorphic copy of these isomorphic substructure. Thus we have a wedge

$$\begin{tikzcd}[column sep=small] & \mathfrak{B} \\ \mathfrak{C} \arrow[ur] \arrow[dr] & \\ & \mathfrak{A} \end{tikzcd} \qquad \begin{tikzcd}[column sep=small] & \mathfrak{B} \arrow[dr] & \\ \mathfrak{C} \arrow[ur] \arrow[dr] & & \mathfrak{D} \\ & \mathfrak{A} \arrow[ur] & \end{tikzcd}$$

of embeddings, as to the left, where $\mathfrak{C} \in \mathcal{S}(T)$ but $\mathfrak{A}, \mathfrak{B} \in \mathcal{M}(T)$. Since $T \cap \forall_1$ has AP, this gives us a commuting square of embeddings, as to the right, where $\mathfrak{D} \in \mathcal{S}(T)$. By a suitable enlargement we may suppose $\mathfrak{D} \models T$. Thus, since T is model complete, the two right hand embeddings are elementary.

Notice how the two points a of $\mathfrak{A}$ and b of $\mathfrak{B}$ are situated in this square. By construction, there is a single point c of $\mathfrak{C}$ which is sent to a and b by the left hand embeddings. But the square commutes, so there is a single point d of $\mathfrak{D}$ to which a and b are sent by the right hand embeddings. Thus

$$(\mathfrak{A}, a) \equiv (\mathfrak{D}, d) \equiv (\mathfrak{B}, b)$$

to verify (iii).

(iii)$\Rightarrow$(iv). This is trivial.

(iv)$\Rightarrow$(v). Consider any $\forall_1$-formula ϕ, and let Δ be the set of quantifier-free formulas δ with $\partial\delta \subseteq \partial\phi$ and such that $T \vdash \phi \to \delta$. Since Δ is close under conjunction, it suffices to show that

$$T \cup \Delta \vdash \phi$$

holds. Of course, to do this we need to assume (iv).

Consider any model $\mathfrak{A} \models T$, and consider any point a of $\mathfrak{A}$ with $\mathfrak{A} \models \Delta(a)$. We must show that $\mathfrak{A} \models \phi(a)$ holds. To this end consider

$$T \cup \mathit{Diag}(\mathfrak{A}, a) \cup \{\phi(a)\}$$

where, of of course, this diagram need not be the full diagram of $\mathfrak{A}$ (since a need not enumerate the whole of $\mathfrak{A}$). By a simple argument, we see that this set is consistent. Thus we obtain a model $\mathfrak{B}$ of T and a point b of $\mathfrak{B}$ such that

$$(\mathfrak{A}, a) \equiv_0 (\mathfrak{B}, b) \qquad \mathfrak{B} \models \phi(b)$$

hold. (We can't say that $\mathfrak{A} \subseteq \mathfrak{B}$ since we are not working with the full diagram of $\mathfrak{A}$.) By (iv) we have

$$(\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b) \qquad \mathfrak{B} \models \phi(b)$$

and hence $\mathfrak{A} \models \phi(a)$, as required.

(v)$\Rightarrow$(vi). This follows by induction on the complexity of ϕ.

(vi)$\Rightarrow$)(i). This is straight forward. ■

Exercises–*needed*

[*Held in 120../B50-bit... Last changed July 26, 2004*]

5 Companion theories and existentially closed structures

Almost trivially, there are theories which are not model complete. For any language L the smallest theory – the pure logic – (the theory of all L-structures) is not model complete.

A theory need not be model complete, but may be we can move, in a canonical fashion, from T to some model complete theory T^* (in the same language). If so, then perhaps T^* will be useful in the analysis of T. There are some theories for which a suitable T^* does exists, and some for which no suitable T^* exists. What do we do when T^* doesn't exists? There are two general possibilities. We can move to a theory $T^\bullet$ which has some, but not all, of the properties of T^*. We can move to a class $\mathcal{C}$ of structures which has some, but not all, of the properties of $\mathcal{M}(T^*)$. Let's start to develop these ideas.

[*Held in 120-../B51-bit.. Last changed July 26, 2004*]

5.1 Model companions

Given a theory T how can we select a model complete theory T^* which is closely attached to T? We need to make precise what 'closely attached' should mean.

5.1 DEFINITION. Two theories T_a, T_b (in the same language) are companions if each model of the one can be embedded in a model of the other. ■

Thus, two theories T_a, T_b are companions exactly when $\mathcal{S}(T_a) = \mathcal{S}(T_b)$ or, equivalently, when $T_a \cap \forall_1 = T_b \cap \forall_1$. For each language this puts an equivalence relation on the family of all theories in that language. We look for special members of each companion block.

5.2 LEMMA. *Suppose T_a and T_b are model complete companion theories. Then $T_a = T_b$.*

Proof. Consider any model $\mathfrak{A} \models T_a$. Since T_a and T_b are companions we have

$$\mathfrak{A} \subseteq \mathfrak{B} \models T_b$$

for some structure $\mathfrak{B}$, and then some structure $\mathfrak{A}'$ with

$$\mathfrak{B} \subseteq \mathfrak{A}' \models T_a$$

(by the same argument). Since T_a is model complete we have $\mathfrak{A} \prec \mathfrak{A}'$, and hence $\mathfrak{A} \prec_1 \mathfrak{B}$. Since T_b is $\forall_2$ axiomatizable this gives $\mathfrak{A} \models T_b$, and hence $\mathcal{M}(T_a) \subseteq \mathcal{M}(T_b)$, so that $T_b \subseteq T_a$. By symmetry we have $T_a \subseteq T_b$, to give $T_a = T_b$, as required. ■

This shows that each companion block contains at most one model complete theory. We select this nice member when it exists, and a good approximation when it doesn't.

5.3 DEFINITION. A model companion of a theory T is a companion T^* of T which is model complete. ■

Lemma 5.2 show that each theory has at most one model companion. Some theories do have a model companion and some don't. Let's look at a couple of non-trivial example.

We begin with an example of a theory which does have a model companion.

5.4 EXAMPLE. Let T be the theory of linearly ordered sets. Thus T is the theory of all structures $(A, \leq)$ where $\leq$ is a reflective linear comparison on the non-empty set A. As in subsection 2.3, we know that T is $\forall_1$-axiomatizable. Let T^* be the theory of lines. Thus $T \subseteq T^*$ and, by the results of subsection 2.3, we know that T^* is model complete (since it has EQ). To show that T^* is the model companion of T, it suffices to show that T and T^* are companions. For this, since $T \subseteq T^*$, it suffices to show that each linearly ordered set can be embedded in a line. This is done in Exercise 2.7. ■

Out next example, of a theory without a model companion, will take a little longer to develop, but illustrates many aspects of this topic.

5.5 EXAMPLE. Let T be the theory of commutative rings with 1. Such a ring is viewed as a structure $\mathfrak{A} = (A, +, \times, 0, 1)$ where $0 \neq 1$. We may follow the common practice and identify each ring $\mathfrak{A}$ with its carrier A. Also we do not need the qualifier 'commutative' since every ring we meet will be commutative.

We show that T does not have a model companion. We argue by contradiction. Thus we assume that T has a model companion T^* and eventually reach a contradiction.

The theory T is $\forall_2$-axiomatizable. Most of the axioms are $\forall_1$-sentences, but the existence of additive inverses requires a $\forall_2$-axiom. In particular $T \subseteq T^*$ holds. To obtain the contradiction we need a trick. ■

Recall that an element a of a ring A is nilpotent if $a^n = 0$ for some $n \in \mathbb{N}$. This can be expressed by an infinite disjunction

$$\bigvee\{v^n \doteq 0 \mid n \in \mathbb{N}\}$$

and we show that, modulo any model companion T^*, this can be reduced to a single formula. Let $\mathrm{nnil}(v)$ be the $\exists_1$-formula

$$(\exists w)[(vw \neq 0) \wedge ((vw)^2 = vw)]$$

where v is some selected variable. This $\exists_1$-formula has many quantifier-free consequences.

5.6 LEMMA. *For each $n \in \mathbb{N}$*

$$T \vdash \mathrm{nnil}(v) \rightarrow (v^n \neq 0)$$

holds.

This is proved by a simple calculation in an arbitrary ring. A similar calculation is used in the following crucial observation.

5.7 LEMMA. *For each ring A and element a of A, the following are equivalent.*

(i) The element a is not nilpotent.

(ii) The is some ring $A \subseteq B$ with $B \models \mathrm{nnil}(a)$.

Proof. $(i) \Rightarrow (ii)$. Suppose a is not nilpotent and consider the factor ring

$$B = A[X]/\langle aX - a^2X^2\rangle$$

of the polynomial extension $A[X]$ of A. A typical member of the ideal $\langle aX - a^2X^2\rangle$ has the form

$$p(X) = (b_0 + b_1X + b_2X^2 + \cdots)(aX - a^2X^2) = c_0 + c_1X + c_2X^2 + c_3 + X^3 + \cdots$$

where $b_0, b_1, b_2, \ldots$ is a list of elements of A with $b_n = 0$ for all sufficiently large n. On multiplying out we see that

$$\begin{aligned} c_0 &= 0 \\ c_1 &= ab_0 \\ c_2 &= ab_1 - a^2b_0 \\ &\vdots \\ c_{n+2} &= ab_{n+1} - a^2b_n \\ &\vdots \end{aligned}$$

are the coefficients of $p(X)$. In particular, the only member of A that belongs to the ideal is 0, and hence the canonical morphism

$$A \longrightarrow B$$

is an embedding.

It now suffices to show that $aX \notin \langle aX - a^2X^2\rangle$.

By way of contradiction, suppose $aX \in \langle aX - a^2X^2\rangle$. Then, using the case $p(X) = aX$ we have

$$\begin{array}{ll} a = ab_0 & a = ab_0 \\ 0 = ab_1 - a^2b_0 & a^2b_0 = ab_1 \\ \vdots & \vdots \\ 0 = ab_{n+1} - a^2b_n & a^2b_n = ab_{n+1} \\ \vdots & \vdots \end{array}$$

for some eventually constant list $b_0, b_1, b_2, \ldots$ of elements of A. A simple induction give

$$a^{n+1} = ab_n$$

(for $n \in \mathbb{N}$, and and hence $a^{n+1} = 0$ for all large n. This is the contradiction, since a is not nilpotent.

$(ii) \Rightarrow (i)$. This follows by a simple calculation. ■

This embedding result quickly leads to the required result.

5.8 THEOREM. *The theory T of commutative rings with 1 does not have a model companion.*

Proof. By way of contradiction, suppose T^* is the model companion of T. We show that

$$T^* \cup \{(v^n \neq 0) \mid n \in \mathbb{N}\} \vdash \text{nnil}(v)$$

holds.

To this end, consider any model (A, a) of the hypothesis set. Thus $A \models T^*$, and a is an element of A which is not nilpotent. By Lemma 5.7, there is some $A \subseteq B \models T$ with $b \models \mathrm{nnil}(a)$. There is some $B \subseteq C \models T^*$, and then $C \models \mathrm{nnil}(a)$ since $\mathrm{nnil}(v)$ is a $\exists_1$-formula. But T^* is model complete, so that $A \prec C$, and hence $A \models \mathrm{nnil}(a)$, as required.

We now apply compactness to get

$$T^* \vdash (v^n \neq 0) \rightarrow \mathrm{nnil}(v)$$

and hence

$$T^* \vdash (\forall v)[(v^n \neq 0) \rightarrow \mathrm{nnil}(v)]$$

for some $n \in \mathbb{N}$. But now, using Lemma 5.6, we have

$$T^* \vdash (\forall v)[(v^n \neq 0) \rightarrow (v^{n+1} \neq 0)]$$

and hence

$$T \vdash (\forall v)[(v^n \neq 0) \rightarrow (v^{n+1} \neq 0)]$$

since we are dealing with a $\forall_1$-sentence. This is the contradiction since, for each n, it is easy to construct a ring with an element a where $a^{n+1} = 0$ but $a^n \neq 0$. ■

This completes the development of Example 5.5.

These two examples and ideas pose several questions.

- When does a theory have a model companion? What special properties does this companion have?
- When a theory does not have a companion, is there any kind of replacement which is nearly as good?

We will investigate these and similar questions further in the following subsections.

The notion of a model companion of a theory was refined (by a rather interesting gentleman, Eli Bers) from an earlier notion of a model completion. It is worth looking at this ancestor.

5.9 DEFINITION. A model completion of a theory T is a companion T^* such that $T \subseteq T^*$ and such that $T^*[\mathfrak{A}]$ is complete for each model $\mathfrak{A} \models T$. ■

Exercise 5.1 indicates how this relates to the notion of a model companion and why model companions are more useful than completions.

Exercises

5.1 (a) Suppose T^* is a model completion of a theory T. Show that T^* is the model companion of T, and T has AP.

(b) Suppose T has AP and a model companion $T \subseteq T^*$. Show that T^* is a model completion of T.

(c) Find an example of a theory which has a model companion but no model completion.

5.2 Complete the proofs of Lemmas 5.6 and 5.7.

5.3 *Use Greg Cherlin's book pages 92 -- 95 to sort out an exercise on rings.*

[*Held in 120-../B52-bit.. Last changed July 26, 2004*]

5.2 Companion operators

What can we do when a theory T does not have a model companion? We need some kind of substitute. We can view the attempted selection of a model companion in two ways, both of which lead to different possible substitutes.

- When the theory T has a model companion, we attach to T a particularly nice theory, this companion. In general, we can look for ways of attaching to T fairly nice theories which have some of the properties of this missing model companion.
- When the theory T has a model companion, we attach to T a particularly nice elementary class of structures, the models of the companion. In general, we can look for ways of attaching to T fairly nice (but non-elementary) classes of structures theories which have some of the properties of this missing elementary class.

In this subsection we develop the first approach, The second approach will be developed later.

How can we attach to a theory a fairly nice companion theory. However we do it, the selection should be done in a uniform way.

5.10 DEFINITION. For a fixed language L, a companion operator is an assignment

$$T \longmapsto T^a$$

between theories (of the underlying language) such that

(i) T and T^a are companions

(ii) if T_1 and T_2 are companions, then $T_1^a = T_2^a$

(iii) $T \cap \forall_2 \subseteq T^a$

hold for all theories T, T_1, T_2. ■

Thus a companion operator selects from each companion block a particular member. Furthermore, it ensures that the $\forall_2$ part of the selected theory is a large as possible.

Shortly we will produce the first example of a companion operator. Before that let's see why these gadgets are useful.

5.11 THEOREM. *Let $(\cdot)^a$ be a companion operator (for some language L). If the theory T has a model companion T^*, then $T^a = T^*$.*

Proof. Suppose the theory T has a model companion T^*. Since T and T^* are companions, we have $T^a = T^{*a}$, by condition (ii). But T^* is $\forall_2$-axiomatizable, so $T^* \subseteq T^{*a} = T^a$, by condition (iii). The theory T^* is model complete, hence so it T^a. Finally, by condition (i), the two theories T^a, T^* are model complete companions, and hence $T^a = T^*$, by Lemma 5.2. ■

Where can we find an example of a companion operator? A simple exercise show that for each companion operator $(\cdot)^a$ the set $T^a \cap \forall_2$ depends only on the theory T and not on the particular operator $(\cdot)^a$. We can construct this set directly, and this leads to the minimum companion operator.

5.12 DEFINITION. For a theory T and $\forall_2$-sentence σ is **0-tame over** T if

$$T \cap \forall_1 \vdash (\sigma \to \alpha) \Longleftrightarrow T \vdash \alpha$$

holds for each $\forall_1$-sentence α. Let $0(T)$ be the set of all such $\forall_2$-sentences. ■

Thus the $\forall_2$-sentence σ is 0-tame over T if and only if $(T \cap \forall_1) \cup \{\sigma\}$ axiomatizes a companion of T. Notice also that $T \cap \forall_2 \subseteq 0(T)$.

5.13 LEMMA. *For each theory T the set $0(T)$ is consistent and axiomatizes a companion of T.*

Proof. To begin we show that

$$\sigma, \tau \in 0(T) \Longrightarrow \sigma \wedge \tau \in 0(T)$$

holds for all $\forall_2$ sentences σ, τ.

Thus, consider $\sigma, \tau \in 0(T)$. Trivially, $\sigma \wedge \tau$ is a $\forall_2$-sentence, and every model of $(T \cap \forall_1) \cup \{\sigma \wedge \tau\}$ is a submodel of T. Thus it is sufficient to show that each model of T is a submodel of $(T \cap \forall_1) \cup \{\sigma \wedge \tau\}$.

Consider any $\mathfrak{A} \in \mathcal{S}(T)$. Using first σ and then τ we see there are submodels $\mathfrak{B}, \mathfrak{C} \in \mathcal{S}(T)$ such that

$$\mathfrak{A} \subseteq \mathfrak{B} \subseteq \mathfrak{C} \qquad \mathfrak{B} \models \sigma \quad \mathfrak{C} \models \tau$$

hold. By iteration we generate two interlacing chains

$$\mathcal{B} = \{\mathfrak{B}_i \mid i < \omega\} \qquad \mathcal{C} = \{\mathfrak{C}_i \mid i < \omega\}$$

of submodels of T such that

$$\mathcal{B}_i \models \sigma \qquad \mathcal{C}_i \models \tau$$

hold for each $i < \omega$. Let $\mathfrak{U}$ be the common union of these two chains. Then

$$\mathfrak{A} \subseteq \mathfrak{U} \in \mathcal{S}(T) \qquad \mathfrak{U} \models \sigma \wedge \tau$$

to give the required result.

Trivially we have $T \cap \forall_1 \subseteq 0(T)$, so to complete the proof it suffices to show that

$$0(T) \vdash \alpha \Longrightarrow \alpha \in T$$

holds for each $\forall_1$-sentence α. To this end, consider any $\forall_1$-sentence α with $0(T) \vdash \alpha$. Since, by the implication above, the set $0(T)$ is closed under conjunction, there is some $\sigma \in 0(T)$ such that $\vdash \sigma \to \alpha$. Consider any model $\mathfrak{A}$ of T. Since $\sigma \in 0(T)$, there is some $\mathfrak{B} \in \mathcal{S}(T)$ with $\mathfrak{A} \subseteq \mathfrak{B} \models \sigma$. But now $\mathfrak{B} \models \alpha$, and hence $\mathfrak{A} \models \alpha$, as required. ■

Notice how useful this interlacing argument is.

5.14 DEFINITION. For each theory T let T^0 be the theory axiomatized by $0(T)$. ■

This shows how to construct the first example of a companion operator. In fact, it is the least companion operator.

5.15 THEOREM. *For each language L the assignment $(\cdot)^0$ is a companion operator. Furthermore, $T^0 \subseteq T^a$ for each theory T and each companion operator $(\cdot)^a$.*

Proof. By Lemma 5.13, T and T^0 are companions.

By construction, T^0 depends only on $T \cap \forall_1$, and hence $T_1^0 = T_2^0$ for companions T_1, T_2.

By construction, $T \cap \forall_2 \subseteq T^0$.

This shows that $(\cdot)^0$ is a companion operator.

Finally, for each companion operator $(\cdot)^a$, we have

$$T^0 \subseteq T^{0a} = T^a$$

using simple properties of such operator. ■

As we proceed we will see several more companion operators, all of which are more interesting than $(\cdot)^0$.

Exercises

5.4 Let $(\cdot)^a$ and $(\cdot)^b$ be companion operators. Show that

$$T^{ab} = T^b \qquad T^a \cap \forall_2 = T^b \cap \forall_2$$

hold for all theories T.

[*Held in 120../B53-bit.. Last changed July 26, 2004*]

5.3 Existentially closed structures

When a theory T has a model companion T^*, we have a particularly nice elementary subclass $\mathcal{M}(T^*)$ of $\mathcal{S}(T)$. What can we do when T does not have a model companion? We look for a subclass of $\mathcal{S}(T)$ which has most of the characteristic properties of $\mathcal{M}(T^*)$, except that it need not be elementary. There are several possible such subclasses. In this section we look at one of the simplest.

5.16 DEFINITION. A structure $\mathfrak{A}$ is existentially closed for a theory T if $\mathfrak{A} \in \mathcal{S}(T)$ and

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_1 \mathfrak{B}$$

holds for each model $\mathfrak{B} \models T$.

Let $\mathcal{E}(T)$ be the class of structures which are existentially closed for T. ■

The definition of $\mathcal{E}(T)$ makes specific reference to models $\mathfrak{B}$ of the theory T. In fact, this is not necessary. The proof of the following is a simple exercise.

5.17 LEMMA. *For each theory T and $\mathfrak{A} \in \mathcal{E}(T)$, the implication*

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_1 \mathfrak{B}$$

holds for all $\mathfrak{B} \in \mathcal{S}(T)$.

This shows that $\mathcal{E}(T)$ depends only on the class $\mathcal{S}(T)$ and not the smaller class $\mathcal{M}(T)$. In other words, $\mathcal{E}(T)$ depends only on the companion block of T.

5.18 THEOREM. *Suppose the theory T has a model companion T^*. Then $\mathcal{E}(T) = \mathcal{M}(T^*)$.*

Proof. Consider any $\mathfrak{A} \in \mathcal{M}(T^*)$ and any $\mathfrak{B} \models T$ with $\mathfrak{A} \subseteq \mathfrak{B}$. Since T and T^* are companions, we have $\mathfrak{B} \subseteq \mathfrak{C}$ for some $\mathfrak{C} \models T^*$. But T^* is model complete, so $\mathfrak{A} \prec \mathfrak{C}$, and hence $\mathfrak{A} \prec_1 \mathfrak{B}$. Thus $\mathfrak{A} \in \mathcal{E}(T)$. This shows that $\mathcal{M}(T^*) \subseteq \mathcal{E}(T)$.

Conversely, consider any $\mathfrak{A} \in \mathcal{E}(T)$. Since T and T^* are companions, we have $\mathfrak{A} \subseteq \mathfrak{B}$ for some model $\mathfrak{B} \models T^*$. Then $\mathfrak{A} \prec_1 \mathfrak{B}$, so that $\mathfrak{B} \equiv\rangle(\forall_2)\,\mathfrak{A}$, and hence $\mathfrak{A} \models T^*$ (since T^* is $\forall_2$-axiomatizable). Thus $\mathfrak{A} \in \mathcal{M}(T^*)$. This shows that $\mathcal{E}(T) \subseteq \mathcal{M}(T^*)$. ■

At this point we should show that $\mathcal{E}(T)$ is non-empty for every theory T. We should prove the following.

5.19 THEOREM. *For each theory T the class $\mathcal{E}(T)$ is cofinal in $\mathcal{S}(T)$. In other words, for each $\mathfrak{A} \in \mathcal{S}(T)$ there is some $\mathcal{B} \in \mathcal{E}(T)$ with $\mathfrak{A} \subseteq \mathfrak{B}$.*

The proof of this is not particularly instructive (and later we will look at more intricate version of the construction used). Thus, in order not to upset the flow, we will postpone the proof until later and continue to develop the properties of $\mathcal{E}(\cdot)$. To do this we will occasionally use this existence result.

5.20 LEMMA. *For each theory T*

$$\mathfrak{A} \prec_1 \mathfrak{B} \in \mathcal{E}(T) \Longrightarrow \mathfrak{A} \in \mathcal{E}(T)$$

holds (for all structures $\mathfrak{A}, \mathfrak{B}$).

Proof. Consider any situation

$$\mathfrak{A} \prec_1 \mathfrak{B} \in \mathcal{E}(T) \qquad \mathfrak{A} \subseteq \mathfrak{C} \models T$$

(for structure $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$). We must show that $\mathfrak{A} \prec_1 \mathfrak{C}$ holds.

We have a wedge, as to the left

$$\begin{tikzcd} & \mathfrak{B} \\ \mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & \\ & \mathfrak{C} \end{tikzcd} \qquad \begin{tikzcd} & \mathfrak{B} \arrow[dr, "h"] & \\ \mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & & \mathfrak{D} \\ & \mathfrak{C} \arrow[ur, "k"'] & \end{tikzcd}$$

where f is a $\prec_1$-embedding and g is an embedding. (In fact, both are insertions.) By Lemma 4.5 there is a commuting square of embeddings, as to the right, where k is elementary. In particular, $\mathfrak{C} \models T$, and hence h is a $\prec_1$-embedding (since $\mathfrak{B} \in \mathcal{E}(T)$.

Now consider any $\forall_1$-formula $\phi(v)$ and point a of $\mathfrak{A}$ (which matches v). Using the various kinds of embeddings we have

$$\mathfrak{A} \models \phi(a) \Longrightarrow \mathfrak{B} \models \phi(fa) \Longrightarrow \mathfrak{D} \models \phi((h \circ f)a) \Longrightarrow \mathfrak{D} \models \phi((k \circ g)a) \Longrightarrow \mathfrak{C} \models \phi(ga)$$

so that

$$\mathfrak{A} \models \phi(a) \Longrightarrow \mathfrak{C} \models \phi(a)$$

and hence $\mathfrak{A} \prec_1 \mathfrak{C}$, as required. ■

With this result we can produce an intrinsic characterization of $\mathcal{E}(\cdot)$.

5.21 THEOREM. *For each theory T the class $\mathcal{E}(T)$ is uniquely characterized by the following three properties.*

(i) $\mathcal{E}(T)$ is cofinal in $\mathcal{S}(T)$.

(ii) The implication

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_1 \mathfrak{B}$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{E}(T)$.

(iii) The implication

$$\mathfrak{A} \prec_1 \mathfrak{B} \in \mathcal{E}(T) \Longrightarrow \mathfrak{A} \in \mathcal{E}(T)$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{S}(T)$.

Proof. We must do two things. We must show that the class $\mathcal{E}(T)$ has these three properties, and we must show that $\mathcal{E}(T)$ is the only class with these three properties.

First we verify that $\mathcal{E}(T)$ has (i,ii,iii).

(i) This is the basic existence result, Theorem 5.19 (which, of course, we have not yet proved).

(ii) As in Lemma 5.17, this is a simple consequence of the definition of $\mathcal{E}(T)$.

(iii) This is just Lemma 5.20.

Secondly, let $\mathcal{E}'(T)$ be any class with properties (i′,ii′,iii′) corresponding to (i,ii,iii). We must show that $\mathcal{E}(T) = \mathcal{E}'(T)$.

Consider any $\mathfrak{A} \in \mathcal{E}(T)$. By (i′) and (i) there are structures

$$\mathfrak{A} \subseteq \mathfrak{B} \subseteq \mathfrak{C}$$

where $\mathfrak{B} \in \mathcal{E}'(T)$ and $\mathfrak{C} \in \mathcal{E}(T)$. By (ii) we have $\mathfrak{A} \prec_1 \mathfrak{C}$, and hence $\mathfrak{A} \prec_1 \mathfrak{B}$. But now (iii′) gives $\mathfrak{A} \in \mathcal{E}'(T)$.

This gives $\mathcal{E}(T) \subseteq \mathcal{E}'(T)$, and a symmetric argument gives the converse inclusion. ■

In general, the class $\mathcal{E}(T)$ is not elementary. When it is elementary it is the class of models of the model companion of T. However, even though the class need not be elementary, it still has a theory.

5.22 DEFINITION. For each theory T let $T^e = Th(\mathcal{E}(T))$. ■

When T has a model companion T^*, we have $\mathcal{E}(T) = \mathcal{M}(T^*)$, and then $T^e = T^*$. In fact, this equality is a consequence of the following (whose proof is left as an exercise).

5.23 THEOREM. *For each language L the assignment $(\cdot)^e$ is a companion operator.*

Existentially closed structures are an important tool in model theory. We will use them all the time, and develop some special classes of these structures.

Exercises

5.5 Show that for each theory T, if the class $\mathcal{E}(T)$ is elementary, then T has a model companion.

5.6 Show that for each theory T the class $\mathcal{E}(T)$ is closed under unions of directed systems.

5.7 Let T be an arbitrary theory.

(a) Show that

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\,(\mathfrak{B}, b) \Longrightarrow (\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b)$$

holds for each $\mathfrak{A} \in \mathcal{E}(T), \mathfrak{B} \in \mathcal{S}(T)$, and matching points a, b of these structures.

(b) Show that

$$\mathfrak{A} \equiv\rangle(\exists_1)\,\mathfrak{B} \Longrightarrow \mathfrak{A} \equiv_2 \mathfrak{B}$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{E}(T)$.

5.8 Prove Theorem 5.23.

5.9 Let T be the theory of commutative rings (as in Example 5.5), and let

$$\mathrm{NNil}(v) = \{(v^n \neq 0) \mid n \in \mathbb{N}\}$$

to obtain a quantfier-free type.

Show that

$$\mathcal{E}(T) \models (\forall v)[\bigwedge \mathrm{NNil}(v) \leftrightarrow \mathrm{nnil}(v)]$$

(where this infinite conjunction has the obvious semantics).

Does this show that

$$T^e \cup \mathrm{NNil}(v) \vdash \mathrm{nnil}(v)]$$

holds?

[*Held in 120../B54-bit.. Last changed July 26, 2004*]

5.4 Existence and characterization

We now turn to the characterization and construction of existentially closed structures. In fact, this can be done quite quickly. However, to prepare for later material, and to give a comparison, we will do more than is strictly necessary.

Recall that a type is a set Γ of formulas such that

$$\partial\Gamma = \bigcup\{\partial\phi \mid \phi \in \Gamma\}$$

is finite. We often write $\Gamma(v)$ to indicate that Γ is a type with v as its batch of free variables.

5.24 DEFINITION. Let T be a theory.

A **maximal-$\exists_1$ type over** T is a type $\Sigma(v)$ which is a set of $\exists_1$-formulas, such that $T \cup \Sigma(v)$ is consistent, and such that

$$T \cup \Sigma \cup \{\theta(v)\} \text{ is consistent} \Longrightarrow \theta \in \Sigma$$

holds for each $\exists_1$-formula θ.

For each list v of variable let $M(T, v)$ be the set of types $\Sigma(v)$, in the variables v, which are maximal-$\exists_1$ over T. ■

We often vary the phrase 'maximal-$\exists_1$ type over T' and refer to a $\exists_1$-type which is maximally consistent with T, or some similar phrase.

What has this notion got to do with existential closedness?

5.25 LEMMA. *Let T be a theory. For each $\exists_1$-type $\Sigma(v)$ the following are equivalent.*

(i) The type $\Sigma(v)$ is maximally consistent with T.

(ii) The type $\Sigma(v)$ is the $\exists_1$-type of some point a of some $\mathfrak{A} \in \mathcal{E}(T)$.

Proof. $(i) \Rightarrow (ii)$. Suppose $\Sigma(v)$ is a maximal-$\exists_1$ type over T. Since Σ is consistent with T, it is realized in some model $\mathfrak{B}$ of T. Consider any $\mathfrak{A} \in \mathcal{E}(T)$ with $\mathfrak{B} \subseteq \mathfrak{A}$. The $\exists_1$-type is realized in $\mathfrak{A}$ by some point a, say. We show that Σ is the $\exists_1$-type of a in $\mathfrak{A}$.

Consider any $\exists_1$-formula $\theta(v)$ such that $\mathfrak{A} \models \theta(a)$. Then

$$T \cup \Sigma \cup \{\theta(v)\}$$

is consistent, and hence (i) gives $\theta \in \Sigma$, as required.

$(ii) \Rightarrow (i)$. Consider any $\mathfrak{A} \in \mathcal{E}(T)$ and point a of $\mathfrak{A}$, and suppose $\Sigma(v)$ is the $\exists_1$-type of a in $\mathfrak{A}$. In particular, Σ is consistent with T. Consider any $\exists_1$-formula $\theta(v)$ such that

$$T \cup \Sigma \cup \{\theta(v)\}$$

is consistent. There is some model $\mathfrak{B}$ of T and some point b of $\mathfrak{B}$ such that

$$\mathfrak{B} \models \Sigma(b) \qquad \mathfrak{B} \models \theta(b)$$

hold. The first of these gives

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\,(\mathfrak{B}, b)$$

so that

$$(\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b)$$

(since $\mathfrak{A} \in \mathcal{E}(T)$). But now $\mathfrak{A} \models \theta(a)$, which leads to the required result. ■

Using a variant of this proof we obtain the following characterization.

5.26 LEMMA. *Let T be a theory. For each $\mathfrak{A} \in \mathcal{S}(T)$ the following are equivalent.*

(i) $\mathfrak{A} \in \mathcal{E}(T)$.

(ii) For each point a of $\mathfrak{A}$, the $\exists_1$-type of a in $\mathfrak{A}$ is maximally consistent with T.

This indicates that to construct an existentially closed structure we must somehow maximize the existential types involved. This can be done by a rudimentary saturation process (by the kind of construction we look at in section 12). However, it is more instructive to obtain existentially closed structures in a slightly different way.

We continue to use types, but not just $\exists_1$-types.

5.27 DEFINITION. For a theory T and formula ϕ let

$$\exists_1(T, \phi)$$

be the set of $\exists_1$-formulas θ such that

$$\partial\theta \subseteq \partial\phi \qquad T \cup \theta \text{ is consistent} \qquad T \vdash \theta \rightarrow \phi$$

hold. We call $\exists_1(T, \phi)$ the set of $\exists_1$-generators of ϕ over T. ■

What we have here is an $\exists_1$-type in the underlying language. However, it is not the one we want. We look at the negation set

$$\neg\Sigma = \{\neg\theta \mid \theta \in \Sigma\}$$

of the type $\Sigma = \exists_1(T, \phi)$. This gives a $\forall_1$-type which we augment.

5.28 DEFINITION. For a theory T and $\forall_1$-formula ϕ let

$$\Omega(T, \phi) = \{\phi\} \cup \neg\exists_1(T, \phi)$$

to produce an $\forall_1$-type. ■

The statement of the next result could be simplified by leaving out certain clauses. However, there are different times when each one is useful.

5.29 THEOREM. *Let T be a theory and let $\mathfrak{A} \in \mathcal{S}(T)$. The following are equivalent.*

(i) $\mathfrak{A} \in \mathcal{E}(T)$

(ii) For each $\forall_1$-formula $\phi(v)$ and point a of $\mathfrak{A}$ with $\mathfrak{A} \models \phi(a)$,

$$T \cup Diag(\mathfrak{A}, \mathsf{a}) \vdash \phi(a)$$

holds (where a is a full enumeration of $\mathfrak{A}$).

(iii) For each $\forall_1$-formula $\phi(v)$ and point a of $\mathfrak{A}$ with $\mathfrak{A} \models \phi(a)$, there is an $\exists_1$-formula $\theta(v)$ such that

$$\mathfrak{A} \models \theta(a) \qquad T \vdash \theta \rightarrow \phi$$

hold.

(iv) For each $\forall_1$-formula ϕ the structure $\mathfrak{A}$ omits $\Omega(T, \phi)$.

Proof. (i)⇒(ii). Assuming (i) and $\mathfrak{A} \models \phi(a)$, consider any model of $T \cup \mathit{Diag}(\mathfrak{A})$. In other words, consider a structure $\mathfrak{B}$ with $\mathfrak{A} \subseteq \mathfrak{B} \models T$. Then (i) gives $\mathfrak{A} \prec_1 \mathfrak{B}$ so that $\mathfrak{B} \models \phi(a)$ to give the required result.

(ii)⇒(iii). Assuming (ii) consider any $\forall_1$-formula $\phi(v)$ where $\mathfrak{A} \models \phi(a)$ for some point a of $\mathfrak{A}$. By (ii) we have

$$T \cup \mathit{Diag}(\mathfrak{A}) \vdash \phi(a)$$

and hence

$$T \vdash \delta(w, v) \to \phi(v)$$

for some quantifier-free formula $\delta(w, v)$ such that $\mathfrak{A} \models \delta(b, a)$ for some point b of $\mathfrak{A}$. Let

$$\theta(v) = (\exists w)\delta(w, v)$$

to produce the required $\exists_1$-formula.

(iii)⇒(iv). Condition (iv) is little more than a rephrasing of (iii).

(iv)⇒(i). Assuming (iv) consider structure $\mathfrak{B}$ with $\mathfrak{A} \subseteq \mathfrak{B} \models T$. Suppose $\mathfrak{A} \models \phi(a)$ where $\phi(v) \in \forall_1$ and a is a point of $\mathfrak{A}$. We know that $\mathfrak{A}$ omits $\Omega(T, \phi)$, so there is some $\theta(v) \in \exists_1$ with

$$T \vdash \theta \to \phi$$

and $\mathfrak{A} \models \theta(a)$. Since $\mathfrak{B} \models T$ we see that

$$\mathfrak{B} \models (\forall v)[\theta \to \phi]$$

holds. But $\mathfrak{A} \subseteq \mathfrak{B}$ and $\theta \in \exists_1$, so that $\mathfrak{B} \models \theta(a)$, which gives $\mathfrak{B} \models \phi(a)$, as required. ■

We are now almost in a position to prove Theorem 5.19. This is not a very difficult proof, but it does lead to more refined techniques. In later section we will look at the saturation technique and the omitting types technique. The proof of Theorem 5.19 is a rather feeble version of both of these. It is convenient to phrase it as an omitting types argument, but that is merely to allow the following definition.

5.30 DEFINITION. Given a pair $\mathfrak{A} \subseteq \mathfrak{B}$ of structures and a type Φ, we say $\mathfrak{B}$ omits Φ within $\mathfrak{A}$ is there is no point a of $\mathfrak{A}$ such that $\mathfrak{B} \models \Phi(a)$ holds. ■

Many of the refined constructions we use come in two parts; a 1-step construction, and then an accumulation construction, in which the 1-step construction is iterated to closure.

5.31 LEMMA. (The 1-step construction) *Let T be a theory. For each $\mathfrak{A} \in \mathcal{S}(T)$, there is a model $\mathfrak{B}$ of T with $\mathfrak{A} \subseteq \mathfrak{B}$, and such that for each $\forall_1$-formula ϕ, the structure $\mathfrak{B}$ omits $\Omega(T, \phi)$ within $\mathfrak{A}$.*

Proof. Given the structure $\mathfrak{A} \in \mathcal{S}(T)$, we extend the underlying language L to $L(\mathsf{a})$ by adding a parameter to name each member of $\mathfrak{A}$. Thus $\Delta = \mathit{Diag}(\mathfrak{A}, \mathsf{a})$ is a set of quantifier-free $L(\mathsf{a})$-sentences. Consider those sets Ψ of $\exists_1$-sentences of $L(\mathsf{a})$ such that

$$\Delta \subseteq \Psi \qquad T \cup \Psi \text{ is consistent}$$

hold. In particular, Δ is one such set. This set is partially ordered by inclusion, and is closed under unions of directed families.

By Zorn's lemma, there is a maximal such set Ψ.

Consider any model of this maximal set. This gives a structure $\mathfrak{A} \subseteq \mathfrak{B} \models T$ such that $(\mathfrak{B}, \mathsf{a}) \models \Psi$. Consider any $\forall_1$-formula $\phi(v)$ (of the parent language L) and consider any point a from $\mathfrak{A}$ which matches the free variables v. We must show that $\mathfrak{B}$ does not realize $\Omega(T, \phi)$ at a.

If

$$T \cup \Psi \cup \{\neg\phi(a)\}$$

is consistent, then (since $\neg\phi \in \exists_1$) the maximality of Ψ gives $\neg\phi(a) \in \Psi$, and hence $\mathfrak{B} \models \neg\phi(a)$, so that $\mathfrak{B}$ does not realize $\Omega(T, \phi)$ at a.

If

$$T \cup \Psi \cup \{\neg\phi(a)\}$$

is inconsistent then

$$T \cup \Psi \vdash \phi(a)$$

and hence

$$T \vdash \psi(w, v) \to \phi(v)$$

for some formula $\psi \in \exists_1$ where $\mathfrak{B} \models \psi(b, a)$ for some point b of $\mathfrak{B}$ taken from $\mathfrak{A}$. Let $\theta(v) = (\exists w)\psi(w, v)$, so that $\neg\theta \in \Omega(T, \phi)$ and $\mathfrak{B} \models \theta(a)$, so that $\mathfrak{B}$ does not realize $\Omega(T, \phi)$ at a. ■

The structure $\mathfrak{B}$ produced in this way omits each type $\Omega(T, \phi)$ only in a part of its carrier. It is possible that this type is realized in the remaining part. To ensure that doesn't happen, we must repeat the construction.

5.32 THEOREM. (The accumulation construction) *Let T be a theory. For each $\mathfrak{A} \in \mathcal{S}(T)$ there is some $\mathfrak{B} \in \mathcal{E}(T)$ with $\mathfrak{A} \subseteq \mathfrak{B}$.*

Proof. Given a structure $\mathfrak{A} \in \mathcal{S}(T)$ we may iterate the use of Lemma 5.31 to produce an ω-chain

$$\mathfrak{A} = \mathfrak{A}_0 \subseteq \mathfrak{A}_1 \subseteq \cdots \subseteq \mathfrak{A}_i \subseteq \cdots \qquad (i < \omega)$$

of submodels of T where, for each step i, the structure $\mathfrak{A}_{i+1}$ omits within $\mathfrak{A}_i$ the appropriate family of types. Let

$$\mathfrak{B} = \bigcup\{\mathfrak{A}_i \mid i < \omega\}$$

so that $\mathfrak{A} \subseteq \mathfrak{B} \in \mathcal{S}(T)$. We show that $\mathfrak{B} \in \mathcal{E}(T)$.

Consider any structure $\mathfrak{C}$ with $\mathfrak{B} \subseteq \mathfrak{C} \models T$, and any formula $\phi \in \forall_1$ with $\mathfrak{B} \models \phi(b)$ for some point b of $\mathfrak{B}$. Since b has just finitely many elements, there is some $i < \omega$ such that b comes from $\mathfrak{A}_i$, and hence from $\mathfrak{A}_j$ for each $i \leq j < \omega$. Since

$$\mathfrak{A}_i \subseteq \mathfrak{A}_{i+1} \subseteq \mathfrak{B} \models \phi(b)$$

and $\phi \in \forall_1$, we have $\mathfrak{A}_{i+1} \models \phi(b)$. But $\mathfrak{A}_{i+1}$ omits the type $\Omega(T, \phi)$ within $\mathfrak{A}_i$, and hence $\mathfrak{A}_{i+1} \models \theta(b)$ for some $\theta \in \exists_1$ where $T \vdash \theta \to \phi$. But now $\mathfrak{C} \models (\forall)[\theta \to \phi]$ and $\mathfrak{C} \models \theta(b)$, to give $\mathfrak{C} \models \phi(b)$, as required. ■

You should make sure you understand this last proof. The same idea is used many times in model theory.

5.10 Prove Lemma 5.26.

[*Held in 120../B55-bit.. Last changed July 26, 2004*]

5.5 Theories which are weakly complete

Theorem 4.9 gives us four different ways of characterizing the model completeness of a theory T. Three of these are concerned with with the equivalence (in T) of certain kinds of formulas. There are several weakenings of these notions, each obtained by replacing the equivalence (in T) by an implication (in T). Here is a selection of these weakenings.

5.33 DEFINITION. ([ω]) A theory T is f-complete if for each formula $\phi(v)$ which is consistent with T, there is an $\exists_1$-formula $\theta(v)$ which is consistent with T and such that $T \vdash \theta \rightarrow \phi$.

([n]) For $n \in \mathbb{N}$, a theory T is [n]-complete if for each $\forall_{n+1}$-formula $\phi(v)$ which is consistent with T, there is an $\exists_1$-formula $\theta(v)$ which is consistent with T and such that $T \vdash \theta \rightarrow \phi$.

((n)) For $n \in \mathbb{N}$, a theory T is (n)-complete if for each $\forall_{n+1}$-formula $\phi(v)$ which is consistent with T, there is an $\exists_{n+1}$-formula $\theta(v)$ which is consistent with T and such that $T \vdash \theta \rightarrow \phi$. ■

There are several observations that can be made immediately.

Because we are concerned with implications in T, at each step we consider only formulas that are consistent with T. This prevents some silly consequences.

Each model complete theory is f-complete. Also a theory is f-complete exactly when it is [n]-complete for each $n \in \mathbb{N}$. Furthermore, if a theory is [$n+1$]-complete then it is [n]-complete. Thus we have a chain of notions with increasing strength, and f-completeness is the obvious limit of the properties. (A more natural terminolgy is 'ω-complete' rather than 'f-complete'. However, we keep to the standard term. Originally the notion of f-completeness came from an analysis of finite forcing, and the name as stuck.)

Finally, note that for each $n \in \mathbb{N}$, if a theory is [n]-complete then it is (n)-complete. In fact, the two notions agree for $n = 0$, but thereafter [n]-completeness becomes much stronger than (n)-completeness as n increases.

We will use the notion of f-completeness to produce another companion operator $(\cdot)^f$, but this one is more delicate than $(\cdot)^e$.

5.34 LEMMA. *Let T_a, T_b be theories with $T_a \cap \forall_{n+1} = T_b \cap \forall_{n+1}$, and suppose T_b is (n)-complete. Then $T_a \cap \forall_{n+2} \subseteq T_b$.*

Proof. Consider any sentence $\sigma \in T_a \cap \forall_{n+2}$ and, by way of contradiction, suppose $\sigma \notin T_b$. Let $\sigma = (\forall v)\neg\phi(v)$ where ϕ is a $\forall_{n+1}$-formula. Then ϕ is conssistent with T_b, and hence

$$T_b \vdash \theta \rightarrow \phi$$

for some $\exists_{n+1}$-formula $\theta(v)$ which is consistent with T_b. Since $T_a \cap \forall_{n+1} = T_b \cap \forall_{n+1}$, we see that θ is consistent with T_a and

$$T_a \vdash \theta \rightarrow \phi$$

holds. This leads to the contradiction. ■

This result has several consequence.

5.35 COROLLARY. *Let T_a, T_b be companion theories where both are $[n]$-complete. Then $T_a \cap \forall_{n+2} = T_b \cap \forall_{n+2}$.*

Proof. We proceed by induction on n. The base case, $n = 0$, is trivial. The induction step, $n \mapsto n+1$, follows by Lemma 5.34. ■

Combining these results for each n we get the following.

5.36 COROLLARY. *If T_a, T_b are f-complete companions, then $T_a = T_b$.*

This result show that each theory has at most one f-complete companion. Our main job is to show that each theory has at least one f-complete companion.

As yet, the only examples of these weakly complete theories that we have recognized are the model complete ones. In fact, there are more interesting one.

For each theory T and companion operator $(\cdot)^a$ the set $T^a \cap \forall_2$ depends only on T and not on $(\cdot)^a$. In particular, $T^0 \cap \forall_2 = T^e \cap \forall_2$ holds. Thus we may use the class $\mathcal{E}(T)$ to extract some information about T^0.

5.37 THEOREM. *For each theory T the companion T^0 is (0)-complete.*

Proof. Consider any $\forall_1$-formula ϕ which is consistent with T^0. Since $T^0 \cap \forall_2 = T^e \cap \forall_2$, we see that the formula ϕ is also consistent with T^e. Thus it is realized in some $\mathfrak{A} \in \mathcal{E}(T)$. A use of Lemma 5.29(iii) now provides the required $\exists_1$-formula θ. ■

This gives us the first step in the construction of the f-companion T^f of a theory. In fact, Theorem 5.37 contains most of the work needed to produce T^f, provided we are prepared to relativize the result.

5.38 THEOREM. *For each theory T and $n \in \mathbb{N}$, there is a unique theory $T^{(n)}$ such that*

(n i) $T \cap \forall_{n+1} = T^{(n)} \cap \forall_{n+1}$

(n ii) *$T^{(n)}$ is $\forall_{n+2}$-axiomatizable*

(n iii) *$T^{(n)}$ is (n)-complete*

(n iv) $T \cap \forall_{n+2} \subseteq T^{(n)}$

hold.

We don't need to prove this here, but a few words won't go amiss.

First of all, the uniqueness is just Lemma 5.34, It is the existence of $T^{(n)}$ that requires a little more work.

Consider the case $n = 0$. We know that the 0-companion T^0 satisfies (0 i), (0 ii), (0 iii), and Theorem 5.37 gives (0 iv). For the general case we relativize the construction of T^0 using $\exists_{n+1}$-formulas in place of $\exists_1$-formulas, and $\forall_{n+1}$-formulas in place of $\forall_1$-formulas.

With this result the companion T^f is easy to construct.

5.39 DEFINITION. For each theory T let

$$T^{[0]} = T^0 \qquad T^{[n+1]} = (T^{[n]})^{(n+1)}$$

for each $n \in \mathbb{N}$. ■

In other words, the sequence $T^{[\cdot]}$ is generated by recursion along $\mathbb{N}$ using the appropriate version of Theorem 5.38 at each step.

5.40 LEMMA. *For each theory T*

[n i] $\quad T \cap \forall_1 = T^{[n]} \cap \forall_1$

[n ii] $\quad T^{[n]}$ *is $\forall_{n+2}$-axiomatizable*

[n iii] $\quad T^{[n]}$ *is $[n]$-complete*

hold for each $n \in \mathbb{N}$. Furthermore

[n iv] $\quad T^{[n]} \cap \forall_{n+2} = T^{[n+1]} \cap \forall_{n+2}$

hold for each $n \in \mathbb{N}$.

Proof. We verify these properties by induction on n.

For the base case, $n = 0$, we have $T^{[0]} = T^0$ and then [0 i], [0 ii], [0 iii] are immediate.

For the step $n \mapsto n + 1$ we have

$$T^{[n+1]} = (T^{[n]})^{(n+1)}$$

and then

($n+1$ i) $\quad T^{[n]} \cap \forall_{n+2} = T^{[n+1]} \cap \forall_{n+2}$

($n+1$ ii) $\quad T^{[n+1]}$ is $\forall_{n+3}$-axiomatizable

($n+1$ iii) $\quad T^{[n+1]}$ is $(n+1)$-complete

hold by Lemma 5.38. Notice that [n iv] is just ($n + 1$ i), so it remains to verify [$n + 1$ i], [$n + 1$ ii], [$n + 1$ iii]. This is where the induction hypothesis is used.

Using [n i] and ($n + 1$ i) we have

$$T \cap \forall_1 = T^[n] \cap \forall_1 = T^{[n+1]} \cap \forall - 1$$

to give [$n + 1$ i].

Property [$n + 1$ ii] is just ($n + 1$ ii).

To verify [$n + 1$, iii] consider any $\forall_{n+2}$-formula $\phi(v)$ which is consistent with $T^{[n+1]}$. By ($n + 1$ iii) there is some $\exists_{n+2}$-formula $\psi(v)$ which is consistent w3ith $T^{[n+1]}$ such that

$$T^{[n+1]} \vdash \psi(v) \to \phi(v)$$

holds. Let $\psi(v) = (\exists w)\chi(w, v)$ where $\chi(w, v)$ is a $\forall_{n+1}$-formula. This χ is consistent with $T^{[n+1]}$ and hence, by ($n + 1$ i), it is consistent with $T^{[n]}$. But now [n iv] gives us a $\exists_1$-formula $\xi(w, v)$ which is consistent with $T^{[n]}$ and such that

$$T^{[n]} \vdash \xi(w, v) \to \chi(w, v)$$

holds. A second use of $(n+1\text{ i})$ shows that $\xi(w,v)$ is consistent with $T^{[n+1]}$ and

$$T^{[n+1]} \vdash \xi(w,v) \to \chi(w,v)$$

holds. Finally, let $\theta(v) = (\exists w)\chi(w,v)$, so that θ is a $\exists_1$-formula which is consistent with $T^{[n+1]}$ and such that

$$T^{[n+1]} \vdash \theta(v) \to \psi(v)$$

holds. In particular, we have

$$T^{[n+1]} \vdash \theta(v) \to \phi(v)$$

to give the required result. ■

By $[n\text{ ii}]$ and $[n\text{ iii}]$ we have an ascending chain

$$T^{[0]} \subseteq T^{[1]} \subseteq \cdots \subseteq T^{[n]} \subseteq \cdots$$

of theories each of which is a companion of T (by $(n\text{ i})$). Thus we may set

$$T^f = \bigcup \{T^{[n]} \mid n \in \mathbb{N}$$

to obtain a companion of T. The conditions $[n\text{ iii}]$ ensure that T^f is f-complete.

To be finished off

Exercises

5.11 Show that a theory T is 0-complete if and only if $T^0 \subseteq T$.

5.12 *Some properties of T^f?*

[*Held in 120../B60-bit... Last changed July 26, 2004*]

6 Pert and Buxom structures

In this section we look at two kinds of structures which, in a sense, are at the opposite ends of the spectrum of submodels of a theory. They are the 'small' and the 'large' submodels, but not in the sense of size. The 'small' ones realize as few types as possible. Informally these are the pert structures, and formally they are the appropriately atomic structures (where the nature of the embeddings involved must be taken into account). The 'large' structures are the ones that realize as many types as possible. Informally these are the buxom structures, and formally they are the appropriately saturated structures (where again the nature of the embeddings involved must be taken into account).

We look at two versions of such structures. One version, (0), uses embeddings, and the other, (ω), uses elementary embedding. (There are also many versions (n) in between but, on the whole, these are not so interesting. Once we understand the 0-version, it is easy to generate the n-version.)

On the whole, we will concentrate on the 0-version. This is more complicated of the two. We will merely sketch the details of the ω-version.

[*Held in 120-../B61-bit.. Last changed July 26, 2004*]

6.1 Atomicity

In this subsection we consider structures that are rather pert in the sense that each such structure contains a point only when it has to. The type of each point is controlled by a single sentence, and so it is easy to say that such a point must exists (by existentially quantifying out the variables of this controlling formula). At a crucial stage, when we show that such structures exists, we use the omitting types technique. At that stage we need to assume that the underlying language is countable. However, the general notions can be set up for any language.

We will consider the two extreme versions. The 0-version is concerned with embeddings and the like, and the ω-version is concerned with elementary embeddings and the like.

For the 0-version we must analyse a theory T which, in practice, will have JEP. We use associated gadgets such as T^0 and $\mathcal{E}(T)$, and work within the class $\mathcal{S}(T)$ of submodels of T. For the ω-version we analyse a complete theory T and can work entirely within the class $\mathcal{M}(T)$ of models of T.

6.1 DEFINITION. Let T be a theory.

(ω) A formula θ is complete over T if it is consistent with T and exactly one of

$$T \vdash \theta \rightarrow \psi \qquad T \vdash \theta \rightarrow \neg\psi$$

holds for each formula ψ (with $\partial\psi \subseteq \partial\theta$).

(0) An $\exists_1$-formula θ is $\exists_1$-complete over T if it is consistent with T and

$$\left.\begin{array}{l} T \cup \{\theta, \psi_1\} \text{ is consistent} \\ T \cup \{\theta, \psi_2\} \text{ is consistent} \end{array}\right\} \Longrightarrow T \cup \{\psi_1, \psi_2\} \text{ is consistent}$$

holds for all $\exists_1$-formulas ψ_1, ψ_2 (with $\partial\psi_1 \cup \partial\psi_2 \subseteq \partial\theta$). ■

The ω-version of this notion can be rephrased to look more like the 0-version. However, the phrasing above is the standard way to describe this notion. The 0-version can *not* be rephrased to look like the ω-version. This is because the negation of a $\exists_1$-formula need not be a $\exists_1$-formula.

There are various characterizations of $\exists_1$-completeness, and we need some of these. We use the 0-companion T^0 of the parent theory T. Recall that an $\exists_2$-formula $\chi(v)$ is consistent with T^0 exactly when $\mathfrak{A} \models \chi(a)$ for some $\mathfrak{A} \in \mathcal{E}(T)$ and some point a of $\mathfrak{A}$.

[`Put this as an earlier exercises`]

6.2 LEMMA. *Let T be a theory, and let θ be an $\exists_1$-formula consistent with T. The following are equivalent.*

(i) The formula θ is $\exists_1$-complete over T.

(ii) The implication

$$T^0 \cup \{\theta, \phi\} \text{ is consistent } \Longrightarrow T \vdash \theta \to \phi$$

holds for each $\forall_1$-formula ϕ (with $\partial\phi \subseteq \partial\theta$).

(iii) The implication

$$T \cup \{\theta, \psi\} \text{ is consistent } \Longrightarrow T^0 \vdash \theta \to \psi$$

holds for each $\exists_1$-formula ψ (with $\partial\psi \subseteq \partial\theta$).

Proof. (i)$\Rightarrow$(ii). Assuming (i), consider any $\forall_1$-formula ϕ (with $\partial\phi \subseteq \partial\theta$) where $T^0 \cup \{\theta, \phi\}$ is consistent. Since $\theta \wedge \phi$ is a $\exists_2$-formula, we have

$$\mathfrak{A} \models \theta(a) \qquad \mathfrak{A} \models \phi(a)$$

for some $\mathfrak{A} \in \mathcal{E}(T)$ ands some point a from $\mathfrak{A}$. By Lemma 5.29 the second of these gives some $\exists_1$-formula ψ such that

$$\mathfrak{A} \models \psi(a) \qquad T \vdash \psi \to \phi$$

hold. In particular,

$$T \cup \{\theta, \psi\}$$

is consistent, but

$$T \cup \{\psi, \neg\phi\}$$

is not consistent. Thus, since $\neg\phi$ is a $\exists_1$-formula, the assumption (i) shows that

$$T \cup \{\theta, \neg\phi\}$$

is not consistent, and hence

$$T \vdash \theta \to \phi$$

as required.

(ii)⇒(iii). Assuming (ii) consider any $\exists_1$-formula ψ (with $\partial\psi \subseteq \partial\theta$) where $T \cup \{\theta, \psi\}$ is consistent. If $T^0 \nvdash \theta \rightarrow \psi$, then $T^0 \cup \{\theta, \neg\psi\}$ is consistent, and so (ii) gives

$$T \vdash \theta \rightarrow \neg\psi$$

(since $\neg\psi \in \forall_1$). This is not so,and hence we have $T^0 \vdash \theta \rightarrow \psi$, as required.

(iii)⇒(i). Since θ is consistent with T we have $\mathfrak{A} \models \theta(a)$ for some $\mathfrak{A} \models T$ and some point a from $\mathfrak{A}$. Consider

$$\mathfrak{A} \subseteq \mathfrak{B} \subseteq \mathfrak{C}$$

where $\mathfrak{B} \models T^0$ and $\mathfrak{C} \models T$. Note that $\mathfrak{B} \models \theta(a)$ and $\mathfrak{C} \models \theta(a)$. Let Ψ be the $\exists_1$-type of a in $\mathfrak{C}$. For each $\exists_1$-formula ψ, the assumption (iii) gives

$$T \cup \{\theta, \psi\} \text{ consistent} \Longrightarrow T^0 \vdash \theta \rightarrow \psi \Longrightarrow \mathfrak{B} \models \psi(a) \Longrightarrow \mathfrak{C} \models \psi(a) \Longrightarrow \psi \in \Psi$$

which leads to (i). ■

Almost invariably we use the notion of an $\exists_1$-complete formula only for a theory with JEP. We use the notion of a complete formula only for a complete theory.

6.3 DEFINITION. Let T be a theory.

(ω) A structure $\mathfrak{A}$ is atomic for T if $\mathfrak{A} \models T$ and for each point a of $\mathfrak{A}$ there is a formula θ which is complete over T with $\mathfrak{A} \models \theta(a)$.

(0) A structure $\mathfrak{A}$ is $\exists_1$-atomic for T if $\mathfrak{A} \models T^0$ and for each point a of $\mathfrak{A}$ there is a formula θ which is $\exists_1$-complete over T with $\mathfrak{A} \models \theta(a)$.

Let $\mathcal{A}(T)$ be the class of structures which are $\exists_1$-atomic for T. ■

Notice again how the companion T^0 is used in the 0-version. This ensures that an $\exists_1$-atomic structure is existentially closed. In fact, as we will see later, it is rather a special kind of e. c. structure.

Unlike most other classes we associate with a theory T, the class $\mathcal{A}(T)$ can be empty. Shortly we will see how to ensure that $\mathcal{A}(T)$ is non-empty.

6.4 LEMMA. *Let T be a theory. For each structure $\mathfrak{A}$ the following are equivalent.*

(i) $\mathfrak{A}$ is $\exists_1$-atomic for T.

(ii) $\mathfrak{A} \in \mathcal{E}(T)$ and for each point a of $\mathfrak{A}$ there is an $\exists_1$-formula θ such that

$$\mathfrak{A} \models \theta(a) \qquad T^0 \vdash \theta \rightarrow \bigwedge\Sigma$$

where Σ is the $\exists_1$-type of a in $\mathfrak{A}$.

(iii) $\mathfrak{A} \in \mathcal{S}(T)$ and for each point a of $\mathfrak{A}$ there is an $\exists_1$-formula θ such that

$$\mathfrak{A} \models \theta(a) \qquad T \vdash \theta \rightarrow \bigwedge\Pi$$

where Π is the $\forall_1$-type of a in $\mathfrak{A}$.

Proof. (i)⇒(ii). Assuming (i) consider any point a of $\mathfrak{A}$, and let θ be any formula which is $\exists_1$-complete of T and $\mathfrak{A} \models \theta(a)$. Let Σ and Π be, respectively, the $\exists_1$-type of a and the $\forall_1$-type of a in $\mathfrak{A}$. Since $\mathfrak{A} \models T^0$ we see that

$$T^0 \cup \{\theta\} \cup \Sigma \cup \Pi$$

is consistent. Thus

$$T \vdash \theta \to \bigwedge \Pi \qquad T^0 \vdash \theta \to \bigwedge \Sigma$$

hold by Lemma 6.2. The first of these shows that $\mathfrak{A} \in \mathcal{E}(T)$ and the second completes the proof of (ii).

(ii)⇒(iii). Assuming (ii) consider any point a of $\mathfrak{A}$, and let θ be the $\exists_1$-formula given by (ii). Let Σ and Π be, respectively, the $\exists_1$-type of a and the $\forall_1$-type of a in $\mathfrak{A}$. Since $\mathfrak{A} \in \mathcal{E}(T)$, we know that

$$T \vdash \bigwedge \Sigma \to \bigwedge \Pi$$

holds. By (ii) we have

$$T^0 \vdash \theta \to \bigwedge \Sigma$$

and hence

$$T \vdash \theta \to \bigwedge \Pi$$

to complete the proof of (iii).

(iii)⇒(i). Assuming (iii) observe first of all that $\mathfrak{A} \in \mathcal{E}(T)$ and hence $\mathfrak{A} \models T^0$. Consider any point a of $\mathfrak{A}$ and let θ be the $\exists_1$-formula given by (iii). We show that θ is $\exists_1$-complete over T. To do this we verify the property of Lemma 6.2(ii).

Consider any $\forall_1$-formula ϕ (with $\partial\phi \subseteq \partial\theta$ and) where $T^0 \cup \{\theta, \phi\}$ is consistent. There is some $\mathfrak{B} \in \mathcal{E}(T)$ and some point b of $\mathfrak{B}$ with $\mathfrak{B} \models \theta(b) \wedge \phi(b)$. The assumed property (iii) ensures that

$$(\mathfrak{A}, a) \equiv\rangle(\forall_1)\, (\mathfrak{B}, b)$$

and hence

$$(\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b)$$

(since $\mathfrak{B} \in \mathcal{E}(T)$). Thus we have $\mathfrak{A} \models \phi(a)$ (so that $\phi \in \Pi$), and hence $T \vdash \theta \to \phi$ by a second use of (iii). ■

In general, the existence of an $\exists_1$-atomic or an atomic structure structure for a theory requires some restrictions on the theory and the underlying language.

6.5 DEFINITION. (ω) A theory T is atomic is it complete and for each formula ψ which is consistent with T, there is a formula θ which is complete over T such that $T \vdash \theta \to \psi$ holds.

(0) A theory T is $\exists_1$-atomic is it has JEP and for each $\exists_1$-formula ψ which is consistent with T, there is a formula θ which is $\exists_1$-complete over T such that $T^0 \vdash \theta \to \psi$ holds. ■

We now come to the crucial result, the existence of atomic structures. This is where we use the omitting types technique. We do the 0-version, and leave the ω-version as an exercise.

To employ Theorem 9.4 we need to assume that the underlying language is countable. We also need to ensure that a certain set $\mathbf{\Pi}$ of types is countable. This will be automatic.

6.6 THEOREM. *Let T be a theory in a countable language, and suppose T has JEP. The following are equivalent.*

(i) There is a countable structure which is $\exists_1$-atomic for T.

(ii) The theory T is $\exists_1$-atomic.

Proof. (i)$\Rightarrow$(ii). Suppose the structure $\mathfrak{A}$ is $\exists_1$-atomic for T, and consider any $\exists_1$-formula ψ which is consistent with T. Since T has JEP we know that ψ is realized in $\mathfrak{A}$, so there is some point a of $\mathfrak{A}$ with $\mathfrak{A} \models \psi(a)$. By Lemma 6.4(ii) there is some formula θ which is $\exists_1$-complete over T with $\mathfrak{A} \models \theta(a)$ and $T^0 \vdash \theta \to \psi$. This gives (ii).

(Notice that this part does not use the countability of the underlying language.)

(ii)$\Rightarrow$(i). Suppose T is $\exists_1$-atomic. For each finite list v of variables let

$$\Sigma(v) = \{\theta \mid \partial\theta \subseteq v \text{ and } \theta \text{ is } \exists_1\text{-complete over } T\}$$

to produce an $\exists_1$-type. Let

$$\Pi(v) = \neg\Sigma(v)$$

to produce a $\forall_1$-type. Since there are only countably many finite lists v of variables, there are countably many such types Π. Thus these can be omitted in a single structure provided each is not $\exists_1$-principal over T.

By way of contradiction, suppose some such type $\Pi(v)$ is $\exists_1$-principal over T. Thus, there is some $\exists_1$-formula $\psi(v)$ which is consistent with T and such that

$$T \vdash \psi \to \bigwedge \Pi$$

holds. By (ii) there is formula $\theta(v)$ which is $\exists_1$-complete over T and such that $T^0 \vdash \theta \to \psi$ holds. Remembering how T and T^0 are related, we see that

$$T \vdash \theta \to \bigwedge \Pi$$

holds. But, by construction, we have $\theta(v) \in \Sigma(v)$ and hence $\neg\theta(v) \in \Pi(v)$, to give $T \vdash \theta \to \neg\theta$, which is a contradicts since $T \cup \{\theta\}$ is consistent.

We now apply the omitting types result Theorem 9.4, as modified by Exercise 9.1, to obtain some (countable) structure $\mathfrak{A} \in \mathcal{E}(T)$ which omits each of these types Π. By Lemma 6.4(ii), this structure is $\exists_1$-atomic for T. ■

This is the main existence result for the 0-version. The corresponding existence result for the ω-version is similar.

6.7 THEOREM. *Let T be a complete theory in a countable language. The following are equivalent.*

(i) There is a countable structure which is atomic for T.

(ii) The theory T is atomic.

The proof of this follows the same pattern as the proof of Theorem 6.6. The difference is that we no longer have any restrictions on the quantifier complexity of formulas and types. To help with this, all embeddings are elementary, and so formulas are preserved when we need them to be. At the heart of the proof is an application of the omitting types result for full types.

In the remainder of this section we look at some of the properties of $\exists_1$-atomic structures. These are concerned with the back-and-forth technique developed in section 10.

6.8 THEOREM. *Let T be a theory, and let $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}(T)$. The set of pairs (a, b) of points a from $\mathfrak{A}$ and b from $\mathfrak{B}$ for which*

$$(\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b)$$

is a back-and-forth system for $\mathfrak{A}, \mathfrak{B}$.

Proof. Consider any pair (a, b) with

$$(\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b)$$

and any element x of $\mathfrak{A}$. Let $\Sigma(u, v)$ be type of the extended point $a^\frown x$ of $\mathfrak{A}$. Thus u is a list of variable matching a, and v is a single variable. By Lemma 6.4(ii) there is an $\exists_1$-formula $\theta(u, v)$ such that

$$\mathfrak{A} \models \theta(a, x) \qquad T^0 \vdash \theta \to \bigwedge \Sigma$$

holds. In particular, we have $\mathfrak{A} \models (\exists v)\theta(a, v)$ and hence $\mathfrak{B} \models (\exists v)\theta(b, v)$ (by the relationship between a and b). This provides an element y of $\mathfrak{B}$ with $\mathfrak{B} \models \theta(b, y)$. Since $\mathfrak{B} \models T^0$, this gives $\mathfrak{B} \models \Sigma(b, y)$, so that

$$(\mathfrak{A}, a, x) \equiv\rangle(\exists_1)\, (\mathfrak{B}, b, y)$$

and hence

$$(\mathfrak{A}, a, x) \equiv_1 (\mathfrak{B}, b, y)$$

(since $\mathfrak{A}, \mathfrak{B} \in \mathcal{E}(T)$).

This, with a similar argument for the other direction, verifies the back-and-forth property. ■

This has an immediate consequence for the family of $\exists_1$-atomic structures.

6.9 COROLLARY. *Let T be a theory. The implication*

$$\mathfrak{A} \equiv_1 \mathfrak{B} \Longrightarrow \mathfrak{A} \equiv_p \mathfrak{B}$$

holds for each $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}(T)$.

In particular, each such structure is $(1, p)$-homogeneous.

We know that the two relations $\equiv_p$ and $\equiv_\omega$ agree on countable structures, and hence we have the following uniqueness property.

6.10 COROLLARY. *Let T be a theory with JEP. Then, up to isomorphisms, there is at most one countable structure which is $\exists_1$-atomic for T.*

Let T be a theory which is $\exists_1$-atomic (in a countable language). In particular, T has JEP. By combining this Corollary with Theorem 6.6 we see that T has a unique countable $\exists_1$-atomic submodel. This must hold some privileged position in the spectrum of all submodels. What is this? Let $T^a = Th(\mathcal{A}(T))$. By Corollary 6.9 we see that T^a is complete, and $T^e \subseteq T^a$ (since $\mathcal{A}(T) \subseteq \mathcal{E}(T)$). What is this theory T^a?

Exercises

6.1 Let θ be an $\exists_1$-formula consistent with a theory T. Show that θ is $\exists_1$-complete over T if and only if

$$T \vdash \phi_1 \vee \phi_2 \Longrightarrow T \vdash \theta \to \phi_1 \text{ or } T \vdash \theta \to \phi_2$$

holds for all $\forall_1$-formulas ϕ_1, ϕ_2.

6.2 By Theorem 5.34 the companion T^0 of a theory T is 0-complete. Sometimes this can be strengthened.

Show that if T is $\exists_1$-atomic, the for each $\forall_1$-formula $\phi(v)$ which is consistent with T, there is a formula $\theta(v)$ which is $\exists_1$-complete over T for which $T \vdash \theta \to \phi$.

6.3 (a) For an arbitrary $n < \omega$ write down the notion of a formula being $\exists_{n+1}$-complete over a theory T.

(b) Show a formula θ is complete over a theory T if and only if it is $\exists_{n+1}$-complete over T for all sufficiently large $n < \omega$.

6.4 Prove Theorem 6.7.

6.5 In Corollary 6.10, why is it necessary to assume that the theory has JEP?

6.6 Let T be a theory in a countable language, and suppose T has JEP. Let $\mathfrak{A}$ be the unique countable $\exists_1$-atomic structure for T. Show that $\mathfrak{A}$ is embeddable in each model of T^0.

6.7 Show that if a $\exists_1$-atomic theory T has a model companion T^*, then $T^a = T^*$.

[*Held in 120../B62-bit.. Last changed July 26, 2004*]

6.2 Existentially universal structures

In the remainder of this section we consider structures which are rather buxom in the sense that each such structure contains as many different kinds of points as is compatible with its environment. These are saturated structures, but there are many variants of this notion, and we concentrate on just one of them. We deal with what could be termed the $\aleph_0$-(0)-version. Later on we will sketch the possible variants of this notion.

We have met the idea of saturation all ready when we looked at e. c. structures. However, there the saturation is so feeble that we didn't bother to describe it in this way.

Saturation is concerned with the realization of types. In a way it is the obverse to omitting types.

So far most types that we have seen have been pure, that is they have been sets of formulas of the underlying language. We now begin to use types which may contain parameters from some selected structure. In other words, we use types in various enrichments of the underlying language.

Consider a structure $\mathfrak{A}$ for some language L. let a be a point of $\mathfrak{A}$. This gives a finite enrichment $L(a)$ of L. Let $\Phi(a, u)$ be a type in this enriched language. Thus we have a pure type $\Phi(u, v)$ in a larger list of variables, and some of these are instantiated by the selected parameters. In this section we are concerned with $\exists_1$-types in enriched languages. This gives us the notion of a $\exists_1$-type over a structure $\mathfrak{A}$. The two crucial restrictions are that all the parameters come from $\mathfrak{A}$ and the are only finitely many of these, and all the formulas are $\exists_1$.

How can we find examples of such $\exists_1$-types over $\mathfrak{A}$?

Fix the point a, and consider also another point b of $\mathfrak{A}$ (which may overlap a). Thus we have an extended point $a^\frown b$, but it is convenient to keep the two part separate. Let $\Theta(u,v)$ be the $\exists_1$-type in $\mathfrak{A}$ of this extended point. Thus $\Theta(u,v)$ consists of all $\exists_1$-formulas $\theta(u,v)$ such that $\mathfrak{A} \models \theta(a,b)$.

The set

$$\Theta(a,v)$$

is a $\exists_1$-types over $\mathfrak{A}$. In fact, this type is realized in $\mathfrak{A}$ since

$$\mathfrak{A} \models \Theta(a,b)$$

holds for some point b. We may write

$$\mathfrak{A} \models (\exists v)\bigwedge\Theta(a,v)$$

to indicate this without naming the point b.

There may be other types which are not realized in $\mathfrak{A}$. Consider any structure $\mathfrak{A} \subseteq \mathfrak{B}$. With the fixed point a from $\mathfrak{A}$, and any point b from $\mathfrak{B}$, let $\Theta(a,v)$ be the $\exists_1$-type of b in $\mathfrak{B}$. Thus $\mathfrak{B}$ realizes this type

$$\mathfrak{B} \models (\exists v)\bigwedge\Theta(a,v)$$

but it may be omitted by $\mathfrak{A}$.

6.11 DEFINITION. Let T be a theory, and let $\mathfrak{A} \in \mathcal{S}(T)$. A $\exists_1$-type $\Theta(a,v)$ over $\mathfrak{A}$ is T-consistent over $\mathfrak{A}$ if there is some $\mathfrak{A} \subseteq \mathfrak{B} \models T$ in which $\Theta(a,v)$ is realized. ■

Since we are dealing with $\exists_1$-types, the restriction $\mathfrak{B} \in \mathcal{M}(T)$ on the realizing structure can be weakened to $\mathfrak{B} \in \mathcal{S}(T)$.

With these preliminaries we can make precise the idea that a structure contains as many different kinds of points as possible.

6.12 DEFINITION. A structure $\mathfrak{A}$ is existentially universal for a theory T if $\mathfrak{A} \in \mathcal{S}(T)$ and if $\mathfrak{A}$ realizes each $\exists_1$-type over $\mathfrak{A}$ which is T-consistent over $\mathfrak{A}$.

Let $\mathcal{U}(T)$ be the class of structures which are existentially universal for T. ■

You should compare this with Definition 5.16. Although they are phrased rather differently, the two notions are quite similar. The crucial difference is that an e. c. structure is concerned only with the realization of a single formula, whereas an e. u. structure is concerned with the realization of a whole type. In particular, if we consider the case where the $\exists_1$-type is a single formula, then we obtain the following.

6.13 LEMMA. *For each theory T the inclusion $\mathcal{U}(T) \subseteq \mathcal{E}(T)$ holds.*

Eventually we will show that $\mathcal{U}(T)$ is cofinal in $\mathcal{S}(T)$, but the proof is a little complicated, so we postpone it for a while.

We often write $\mathfrak{M}, \mathfrak{N}, \ldots$ to indicate that we are dealing with an e. u. structure.

Existentially universal structures have rather strong back-and-forth properties. Here is the crucial result (whose proof you should compare with Solution 6.5).

6.14 LEMMA. *Let T be a theory and consider a situation*

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\,(\mathfrak{M}, b)$$

where $\mathfrak{A} \in \mathcal{S}(T)$, $\mathfrak{M} \in \mathcal{U}(T)$, and a, b are matching points. Then for each element x of $\mathfrak{A}$, there is a element y of $\mathfrak{M}$ such that

$$(\mathfrak{A}, a^\frown x) \equiv\rangle(\exists_1)\,(\mathfrak{M}, b^\frown y)$$

holds.

Proof. From the given situation we have

$$(\mathfrak{A}, a) \xrightarrow{f} (\mathfrak{C}, b) \qquad\qquad (\mathfrak{M}, b) \subseteq (\mathfrak{C}, b)$$

for some $\mathfrak{C} \in \mathcal{S}(T)$ and embedding f. (In fact, we can arrange that $\mathfrak{M} \prec \mathfrak{C}$, but that is not needed.)

Let $\Theta(a, v)$ be the $\exists_1$-type of x in $(\mathfrak{A}, a)$. Thus $\mathfrak{A} \models \Theta(a, x)$ so that $\mathfrak{C} \models (b, fx)$, to show that

$$\mathfrak{C} \models (\exists v)\bigwedge\Theta(b, v)$$

holds. But $\mathfrak{M} \in \mathcal{U}(T)$, so that

$$\mathfrak{M} \models (\exists v)\bigwedge\Theta(b, v)$$

to give the required element of $\mathfrak{M}$. ■

Using this one construction we can draw several conclusions.

6.15 THEOREM. *Let T be a theory.*

For each $\mathfrak{M}, \mathfrak{N} \in \mathcal{U}(T)$

$$(\mathfrak{M}, a) \equiv\rangle(\exists_1)\,(\mathfrak{N}, b) \Longrightarrow (\mathfrak{M}, b) \equiv_p (\mathfrak{N}, b)$$

holds for all matching points a from $\mathfrak{M}$ and b from $\mathfrak{N}$.

If T has JEP then

$$(\mathfrak{M}, a) \equiv_p (\mathfrak{N}, b)$$

holds for each $\mathfrak{M}, \mathfrak{N} \in \mathcal{U}(T)$ and matching points a and b.

Proof. Since $\mathfrak{M}, \mathfrak{N}$ are e. c. we have

$$(\mathfrak{M}, a) \equiv\rangle(\exists_1)\,(\mathfrak{N}, b) \Longrightarrow (\mathfrak{M}, a) \equiv_2 (\mathfrak{N}, b)$$

and Lemma 6.14 enables us to set up a back-and-forth system.

When T has JEP we have

$$(\mathfrak{M}, a) \equiv_0 (\mathfrak{N}, b) \Longrightarrow (\mathfrak{M}, a) \equiv_2 (\mathfrak{N}, b)$$

For each $\mathfrak{M}, \mathfrak{N} \in \mathcal{U}(T)$. ■

When we restrict attention to a single e. u. structure we see it has some homogeneity.

6.16 COROLLARY. *For each theory T, each $\mathfrak{M} \in \mathcal{U}(T)$ is $(1,p)$-homogeneous.*
If T has JEP, then each $\mathfrak{M} \in \mathcal{U}(T)$ is $(0,p)$-homogeneous.

We know that embeddings between e. c. structures have reasonable strong preservation properties. Theorem 6.15 shows these are much stronger for e. u. structures.

6.17 COROLLARY. *For each theory T*

$$\mathfrak{M} \subseteq \mathfrak{N} \Longrightarrow \mathfrak{M} \prec \mathfrak{N}$$

holds for each $\mathfrak{M}, \mathfrak{N} \in \mathcal{U}(T)$.

Of course, Lemma 6.14 doesn't have to be used in two directions.

6.18 THEOREM. *For each theory T, each $\mathfrak{M} \in \mathcal{U}(T)$, each countable $\mathfrak{A} \in \mathcal{S}(T)$, and each pair a, b of matching points from $\mathfrak{A}, \mathfrak{M}$, the implication*

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\, (\mathfrak{M}, b) \Longrightarrow (\mathfrak{A}, a) \textit{ is embeddable in } (\mathfrak{M}, b)$$

holds.

Proof. It suffices to iterate Lemma 6.14 along an enumeration of $\mathfrak{A}$. ∎

We say a structure $\mathfrak{M}$ is universal for countable submodels of a theory if each such structure is embeddable in $\mathfrak{M}$. The two uses of 'universal' here is no coincidence.

6.19 COROLLARY. *If the theory T has JEP, the each $\mathfrak{M} \in \mathcal{U}(T)$ is universal for countable submodels of T.*

Notice that we still haven't verified the existence of e. u. structures. This is the topic of subsection 6.4.

Exercises–*needed*

[*Held in 120../B63-bit.. Last changed July 26, 2004*]

6.3 A companion operator

We have just seen that the class $\mathcal{U}(T)$ has some nice properties. In particular, Corollary 6.17 suggest that $\mathcal{U}(T)$ is something like the class of models of a model complete theory, or at least can be used to generate a companion operator. We have not yet shown that $\mathcal{U}(T)$ is cofinal in $\mathcal{S}(T)$, but we will do so in the next subsection. The class $\mathcal{U}(T)$ is missing one useful property. It is not closed under taking elementary substructures. We correct this as follows.

6.20 DEFINITION. For each theory T let $\mathcal{G}(T)$ be given by

$$\mathfrak{A} \in \mathcal{G}(T) \Longleftrightarrow \mathfrak{A} \prec \mathfrak{M} \text{ for some } \mathfrak{M} \in \mathcal{U}(T)$$

to produce a subclass of $\mathcal{S}(T)$. ∎

Since each e. u. structure is e. c. we see that $\mathcal{G}(T) \subseteq \mathcal{E}(T)$. In fact, in some ways $\mathcal{G}(T)$ is a nicer version of $\mathcal{E}(T)$. The following result should be compared with Theorem 5.21.

6.21 THEOREM. *For each theory T the class $\mathcal{G}(T)$ is uniquely characterized by the following three properties.*

(i) $\mathcal{G}(T)$ is cofinal in $\mathcal{S}(T)$.

(ii) The implication

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec \mathfrak{B}$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{G}(T)$.

(iii) The implication

$$\mathfrak{A} \prec \mathfrak{B} \in \mathcal{G}(T) \Longrightarrow \mathfrak{A} \in \mathcal{G}(T)$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{S}(T)$.

Proof. We must do two things. We must show that the class $\mathcal{G}(T)$ has these three properties, and we must show that $\mathcal{G}(T)$ is the only class with these three properties.

First we verify that $\mathcal{G}(T)$ has (i,ii,iii).

(i) This is an immediate consequence of the existence result Theorem 6.27 (which, of course, we have not yet proved).

(ii) Consider $\mathfrak{A}, \mathfrak{B} \in \mathcal{G}(T)$ with $\mathfrak{A} \subseteq \mathfrak{B}$. There are $\mathfrak{M}, \mathfrak{N} \in \mathcal{U}(T)$ with $\mathfrak{A} \prec \mathfrak{M}$ and $\mathfrak{B} \prec \mathfrak{N}$. This gives a commuting square of embeddings

$$\begin{tikzcd} \mathfrak{M} \arrow[rr, "f"] & & \mathfrak{C} \\ \mathfrak{A} \arrow[u, hook, "\prec"] & \subseteq \mathfrak{B} \prec & \mathfrak{N} \arrow[u, "g"'] \end{tikzcd}$$

where g and the right hand insertion are elementary. In particular, $\mathfrak{C} \in \mathcal{S}(T)$. Using the existence result Theorem 6.27, this gives a commuting square of embeddings

$$\begin{tikzcd} \mathfrak{M} \arrow[rr, "f"] & & \mathfrak{P} \\ \mathfrak{A} \arrow[u, hook, "\prec"] & \subseteq \mathfrak{B} \prec & \mathfrak{N} \arrow[u, "g"'] \end{tikzcd}$$

where $\mathfrak{P} \in \mathcal{U}(T)$. By Corollary 6.17, both f and g (in this diagram) are elementary, and hence $\mathfrak{A} \prec \mathfrak{B}$, as required.

(iii) This is an immediate consequence of the definition of $\mathcal{G}(T)$.

For the second part let $\mathcal{G}'(T)$ be any class with properties (i$'$,ii$'$,iii$'$) corresponding to (i,ii,iii). We must show that $\mathcal{G}(T) = \mathcal{G}'(T)$. To do this we show

For each $\mathfrak{A} \in \mathcal{G}(T)$ and $\mathfrak{A}' \in \mathcal{G}'(T)$

$[n]$ $\quad \mathfrak{A} \subseteq \mathfrak{A}' \Longrightarrow \mathfrak{A} \prec_n \mathfrak{A}' \qquad \mathfrak{A}' \subseteq \mathfrak{A} \Longrightarrow \mathfrak{A}' \prec_n \mathfrak{A}$

by induction on n.

The base case, [0], is trivial.

For the induction step, $[n] \Rightarrow [n+1]$, consider $\mathfrak{A} \in \mathcal{G}(T)$ and $\mathfrak{A}' \in \mathcal{G}'(T)$ with $\mathfrak{A} \subseteq \mathfrak{A}'$. Since $\mathfrak{A}' \in \mathcal{G}'(T) \subseteq \mathcal{S}(T)$, property (i) gives

$$\mathfrak{A} \subseteq \mathfrak{A}' \subseteq \mathfrak{B}$$

for some $\mathfrak{B} \in \mathcal{G}(T)$. But now

$$\mathfrak{A} \prec \mathfrak{B} \qquad \mathfrak{A}' \prec_n \mathfrak{A}$$

by (ii) and the induction hypothesis $[n]$. This gives $\mathfrak{A} \prec_{n+1} \mathfrak{A}'$ to verify one half of $[n+1]$. The other half follows by a symmetric argument.

Using $[n]$ for each $n \in \mathbb{N}$ we have

$$\mathfrak{A} \subseteq \mathfrak{A}' \Longrightarrow \mathfrak{A} \prec \mathfrak{A}' \qquad \mathfrak{A}' \subseteq \mathfrak{A} \Longrightarrow \mathfrak{A}' \prec \mathfrak{A}$$

for each $\mathfrak{A} \in \mathcal{G}(T)$ and $\mathfrak{A}' \in \mathcal{G}'(T)$. The obvious argument, using (i,iii) and (i$'$,iii$'$), now shows that $\mathcal{G}(T) = \mathcal{G}'(T)$. ■

Given this characterization we must at least do one thing.

6.22 DEFINITION. For each theory T let $T^g = Th(\mathcal{G}(T))$. ■

The associated result is proved in the same way as Theorem 5.23.

6.23 THEOREM. *For each language L the assignment $(\cdot)^g$ is a companion operator.*

We know that structures in $\mathcal{U}(T)$ have some nice back-and-forth properties. Some of these transfer to $\mathcal{G}(T)$.

6.24 THEOREM. *Let T be a theory. Then*

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\, (\mathfrak{B}, b) \Longrightarrow (\mathfrak{A}, a) \equiv (\mathfrak{B}, b)$$

holds for each $\mathfrak{A}, \mathfrak{B} \in \mathcal{G}(T)$ and matching points a and b.

Proof. Suppose

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\, (\mathfrak{B}, b)$$

and consider $\mathfrak{A} \prec \mathfrak{M} \in \mathcal{U}(T)$ and $\mathfrak{B} \prec \mathfrak{N} \in \mathcal{U}(T)$. We have

$$(\mathfrak{M}, a) \equiv\rangle(\exists_1)\, (\mathfrak{N}, b)$$

and hence

$$(\mathfrak{M}, a) \equiv (\mathfrak{N}, b)$$

by Theorem 6.15, to give

$$(\mathfrak{A}, a) \equiv (\mathfrak{B}, b)$$

as required. ■

This back-and-forth property leads to a characterization of the members of $\mathcal{G}(T)$. This result should be compared with Theorem 5.29.

6.25 THEOREM. *Let T be a theory and let $\mathfrak{A} \in \mathcal{S}(T)$. The following are equivalent.*

(i) $\mathfrak{A} \in \mathcal{G}(T)$

(ii) For each formula $\phi(v)$ and point a of $\mathfrak{A}$ with $\mathfrak{A} \models \phi(a)$, there is an $\exists_1$-type $\Theta(v)$ such that

$$\mathfrak{A} \models \Theta(a) \qquad \mathcal{G}(T) \models (\forall v)[\bigwedge\Theta \to \phi]$$

hold.

Proof. (i)⇒(ii). suppose $\mathfrak{A} \in \mathcal{G}(T)$ and $\mathfrak{A} \models \phi(a)$ for some formula $\phi(v)$ and point a of $\mathfrak{A}$. Let $\Theta(v)$ be the $\exists_1$-type of a in $\mathfrak{A}$. It suffices to show that

$$\mathcal{G}(T) \models (\forall v)[\bigwedge\Theta \to \phi]$$

holds.

Consider any $\mathfrak{B} \in \mathcal{G}(T)$ and any point b of $\mathfrak{B}$ with $\mathfrak{B} \models \Theta(b)$. We have

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\,(\mathfrak{B}, b)$$

and hence

$$(\mathfrak{A}, a) \equiv (\mathfrak{B}, b)$$

by Theorem 6.24, to give $\mathfrak{B} \models \phi(b)$, as required.

(ii)⇒(i). Suppose $\mathfrak{A}$ has (ii), and consider any $\mathfrak{M} \in \mathcal{U}(T)$ with $\mathfrak{A} \subseteq \mathfrak{M}$. Then $\mathfrak{A} \prec \mathfrak{M}$, so that $\mathfrak{A} \in \mathcal{G}(T)$. ■

The members of the class $\mathcal{G}(T)$ were originally obtained in a different way. They were constructed using the method of infinite forcing, and they are the generic structures for that technique. This is where the '$\mathcal{G}(\cdot)$' and the '$(\cdot)^g$' comes from. Only later were these structures connected with the saturation technique.

Exercises

6.8 Show that for each theory T, the class $\mathcal{G}(T)$ is elementary if and only if T has a model companion, in which case T^g is the model companion.

6.9 Show that for each theory T, the class $\mathcal{G}(T)$ is closed under unions of directed systems.

[*Held in 120../B64-bit.. Last changed July 26, 2004*]

6.4 Existence of e. u. structures

We want to show that, for each theory T, the class $\mathcal{U}(T)$ is cofinal in $\mathcal{S}(T)$. That is, each $\mathfrak{A} \in \mathcal{S}(T)$ can be embedded in some $\mathfrak{M} \in \mathcal{U}(T)$. We produce $\mathfrak{M}$ from $\mathfrak{A}$ is stages. We first use a 1-step construction to embed $\mathfrak{A}$ into some $\mathfrak{A}' \in \mathcal{S}(T)$ which is only partly saturated (in an appropriate sense). We then iterate this construction to accumulates this partial saturation into a full saturation of the required kind. These stages are analogous to Lemma 5.31 and 5.32.

We need to gain some control over the size of the constructed $\mathfrak{M}$. In particular, we wish to keep $\mathfrak{M}$ as small as possible. With the construction of e. c. structures, this is not a problem. Once we have one such structure, we may take an elementary substructure to produce a smaller one. With e. u. structures this is not possible (since we need a preservation property that is stronger than the one provided by an elementary embedding). Thus, we must take more care with the construction $\mathfrak{A} \longmapsto \mathfrak{M}$.

In the 1-step construction $\mathfrak{A} \longmapsto \mathfrak{A}'$ we need to calculate the size of $\mathfrak{A}'$ in comparison with the size of $\mathfrak{A}$. Let's see how we do that.

We work relative to a theory T which we assume is formalized in a countable language. (Strictly speaking this countability is not essential here, but it will be for many applications. In any case, we are not attempting to describe the most general construction.)

We assume given some $\mathfrak{A} \in \mathcal{S}(T)$, and we want to produce some $\mathfrak{A} \subseteq \mathfrak{A}' \in \mathcal{S}(T)$ which realizes many $\exists_1$-types. Each such type has the form

$$\Theta(a, v)$$

where $\Theta(u, v)$ is a pure $\exists_1$-type, and the point a come from $\mathfrak{A}$ (not yet the whole of $\mathfrak{A}'$). How many such types are there?

Since the language is countable, there are no more than $2^{\aleph_0}$ pure types $\Theta(u, v)$. Each one of these may give many different types $\Theta(a, v)$ as a varies through $\mathfrak{A}$. Suppose $|\mathfrak{A}| \leq 2^\kappa$ for some infinite cardinal κ. Then there are no more than

$$2^{\aleph_0} \cdot 2^\kappa = 2^{\aleph_0 + \kappa} = 2^\kappa$$

types $\Theta(a, v)$. Using this we will arrange that $|\mathfrak{A}'| \leq 2^\kappa$ holds. In particular, as we accumulate

$$\mathfrak{A} \longmapsto \mathfrak{A}' \longmapsto \mathfrak{A}'' \longmapsto \mathfrak{A}''' \longmapsto \cdots$$

at each step the structure produced will have size no more that 2^κ.

It's time to fill in the details.

6.26 LEMMA. (The 1-step construction) *Let T be a theory in a countable language. Let $\mathfrak{A} \in \mathcal{S}(T)$ with $|\mathfrak{A}| \leq 2^\kappa$ for some infinite cardinal κ. Then there is a structure $\mathfrak{A}' \in \mathcal{S}(T)$ with*

$$\mathfrak{A} \subseteq \mathfrak{A}' \qquad |\mathfrak{A}'| \leq 2^\kappa$$

such that for each $\exists_1$-type $\Theta(a, v)$ with parameters from $\mathfrak{A}$, if this type is T-consistent over $\mathfrak{A}'$, then it is already realized in $\mathfrak{A}'$.

Proof. Consider all possible $\exists_1$-types with parameters from $\mathfrak{A}$. By the discussion above, we know there are no more than 2^κ such types. Let

$$\boldsymbol{\Theta} = \{\Theta_i \mid i < 2^\kappa\}$$

by an ordinal enumeration of these types. It doesn't matter if there are some repetitions in this enumeration. Also, each Θ_i will contain certain parameters and free variables, and these may change with i, but we do not need the details of this dependence.

We produce $\mathfrak{A}'$ as the union of a long ascending chain

$$\mathcal{A} = \{\mathfrak{A}_i \mid i < 2^\kappa\}$$

of submodels of T. We generate this chain by recursion along the ordinal 2^κ. For each $i < 2^\kappa$, the recursion step $\mathfrak{A}_i \longmapsto \mathfrak{A}_{i+1}$ deals with the type Θ_i.

Here is what we arrange.

(0) $\mathfrak{A}_0 = \mathfrak{A}$.

(1) $i < j < 2^\kappa \Longrightarrow \mathfrak{A}_i \subseteq \mathfrak{A}_j$.

(2) $|\mathfrak{A}_i| \leq 2^\kappa$ for each $i < 2^\kappa$.

(3) If $l < 2^\kappa$ is a limit ordinal, then $\mathfrak{A}_l = \bigcup\{\mathfrak{A}_i \mid i < l\}$.

(4) For each $i < 2^\kappa$, if the type Θ_i is T-consistent over $\mathfrak{A}_i$, then it is realized in $\mathfrak{A}_{i+1}$.

Before we generate this chain $\mathcal{A}$, let's see why

$$\mathfrak{A} = \bigcup \mathcal{A}$$

has the required properties.

Firstly, this $\mathfrak{A}'$ is the union of an ascending chain of submodels of T, and hence itself is a submodel of T.

Secondly, by (2) and since the length of the chain is 2^κ, we have $|\mathfrak{A}'| \leq 2^\kappa$.

Thirdly, it has the required partial saturation property. Consider any $\exists_1$-type $\Theta(a, v)$ with parameters from $\mathfrak{A}$. Suppose also that this type is T-consistent over $\mathfrak{A}'$. By the choice of $\mathbf{\Theta}$, we have $\Theta = \Theta_i$ for some $i < 2^\kappa$, and then Θ_i is T-consistent over $\mathfrak{A}_i$. By clause (4) this Θ_i is realized in $\mathfrak{A}_{i+1}$ and hence, since $\Theta = \Theta_i$ is a $\exists_1$-type, it is realized in $\mathfrak{A}'$.

It remains to generate the chain $\mathcal{A}$. We do this by recursion along the ordinal 2^κ.

For the base case, $i = 0$, we set $\mathfrak{A}_0 = \mathfrak{A}$, to obtain (0). Since we are given $|\mathfrak{A}| \leq 2^\kappa$, we have (2) (for this i).

For the recursion step, $i \mapsto i + 1$, we look at the type Θ_i. There are two sub-cases.

If Θ_i is *not* T-consistent over $\mathfrak{A}_i$, then we set $\mathfrak{A}_{i+1} = \mathfrak{A}_i$. In this sub-case, properties (1), (2), and (4) are immediate.

Suppose Θ_i is not T-consistent over $\mathfrak{A}_i$. There is some $\mathfrak{A}_i \subseteq \mathfrak{B} \in \mathcal{S}(T)$ in which Θ_i is realized. By taking a suitable elementary substructure, we may suppose $|\mathfrak{B}| = |\mathfrak{A}_i| \leq 2^\kappa$. We let $\mathfrak{A}_{i+1}$ be such a $\mathfrak{B}$, and so preserve properties (1), (2), and (4).

For the leap to a limit ordinal $l < 2^\kappa$, we must set $\mathfrak{A}_l = \bigcup\{\mathfrak{A}_i \mid i < l\}$ by (3). Property (1) is preserved, and (4) is vacuous, so we must check (3). But, using (2) we have

$$|\mathfrak{A}_l| \leq 2^\kappa \cdot |l| = 2^\kappa$$

(since $|l| < 2^\kappa$), as required.

This completes the construction, and the whole proof. ■

Before we continue it is worth looking at the restrictions on the size of the constructed $\mathfrak{A}'$. Let's look at the simplest case where the given structure $\mathfrak{A}$ is countable, so that $|\mathfrak{A}| \leq \aleph_0 < 2^{\aleph_0}$. Perhaps this strict comparison leads to a smaller $\mathfrak{M} \in \mathcal{U}(T)$. The problem. however, is not the size of $\mathfrak{A}$ but the potential number of types that have to be realized. This is still $2^{\aleph_0}$, so the size of the first step structure $\mathfrak{A}'$ still could be $2^{\aleph_0}$. At the next step we move from $\mathfrak{A}'$ to $\mathfrak{A}''$, so we are back in the situation of Lemma 6.26 with $\kappa = \aleph_0$. We will return to this discussion later.

6.27 THEOREM. (The accumulation construction) *Let T be a theory in a countable language. Let $\mathfrak{A} \in \mathcal{S}(T)$ with $|\mathfrak{A}| \leq 2^\kappa$ for some infinite cardinal κ. Then there is a some $\mathfrak{M} \in \mathcal{U}(T)$ with $\mathfrak{A} \subseteq \mathfrak{M}$ and $|\mathfrak{M}| \leq 2^\kappa$.*

Proof. Given a structure $\mathfrak{A} \in \mathcal{S}(T)$ we may iterate the 1-step construction of Lemma 6.26 to produce an ω-chain

$$\mathfrak{A} = \mathfrak{A}_0 \subseteq \mathfrak{A}_1 \subseteq \cdots \subseteq \mathfrak{A}_i \subseteq \cdots \qquad (i < \omega)$$

of submodels of T where, for each step i, the structure $\mathfrak{A}_{i+1}$ realizes certain types taken from $\mathfrak{A}_i$. Let

$$\mathfrak{M} = \bigcup\{\mathfrak{A}_i \mid i < \omega\}$$

so that $\mathfrak{A} \subseteq \mathfrak{M} \in \mathcal{S}(T)$. We show that $\mathfrak{M} \in \mathcal{U}(T)$.

Consider any $\exists_1$-type $\Theta(a, v)$ with parameters a from $\mathfrak{M}$ and which is T-consistent over $\mathfrak{M}$. This point a comprises just finitely many elements of $\mathfrak{M}$, and hence a is a point of $\mathfrak{A}_i$ for some $i < \omega$. Furthermore, the type Θ is T-consistent over $\mathfrak{A}_i$. By construction, the $\exists_1$-type Θ is realized in $\mathfrak{A}_{i+1}$, and hence is realized in $\mathfrak{M}$, as required. ■

Suppose T is a theory in a countable language, and suppose $\mathfrak{A} \in \mathcal{S}(T)$ is countable. Theorem 6.27 gives us some $\mathfrak{A} \subseteq \mathfrak{M} \in \mathcal{U}(T)$ with $|\mathfrak{M}| \leq 2^{\aleph_0}$. In terms of size, is this the best we can do? The following example shows that, in general, it is.

6.28 EXAMPLE. Consider the binary splitting tree as a structure

$$\mathfrak{T} = (\mathbf{\Psi}, \leq, \perp, S_0, S_1)$$

where $\mathbf{\Psi}$ is the set of all node; $\leq$ is the comparison of these nodes; and $\perp$ is the empty node, the bottom of the poset. The two 1-placed operation symbols S_0 and S_1 are the left and right successor operations. Thus

$$S_0 a = a0 \qquad S_1 a = a1$$

for each node a. Notice that each node a has a canonical name

$$a = S_{i(n)}(\cdots(S_{i(0)})\cdots)$$

where each $S_{i(\cdot)}$ is S_0 or S_1, as appropriate.

Let $T = Th(\mathfrak{T})$. This is the theory of two successor functions. We need not write down the axioms for this theory, but we should note that

$$(\forall u, v, w)[(u \leq w) \land (v \leq w) \rightarrow (u \leq v) \lor (v \leq u)]$$

holds in T.

The structure $\mathfrak{T}$ is countable, and $\mathfrak{T} \subseteq \mathfrak{M}$ for some $\mathfrak{M} \in \mathcal{U}(T)$. How big must $\mathfrak{M}$ be?

Consider any branch p of $\mathbf{\Psi}$. Let

$$\Theta_p(w) = \{(a \leq w) \mid a < p\}$$

to produce a quantifier-free pure type in the variable w. This type is finitely satisfiable in $\mathfrak{T}$, and hence in $\mathfrak{M}$, so it is realized in some elementary extension of $\mathfrak{M}$. But $\mathfrak{M} \in \mathcal{U}(T)$ and hence the type is realized in $\mathfrak{M}$.

For each branch p of $\mathbf{\Psi}$ the type Θ_p is realized by some element of $\mathfrak{M}$. We check that no two branches are realized by the same element, and hence $|\mathfrak{M}| \geq 2^{\aleph_0}$.

Consider disting branches p, q of $\mathbf{\Psi}$, and suppose

$$\mathfrak{M} \models \Theta_p(m) \qquad \mathfrak{M} \models \Theta_q(m)$$

for some $m \in \mathfrak{M}$. Since $p \neq q$, there are nodes $a, b \in \mathbf{\Psi}$ such that

$$a < p \quad b \not< q \quad b < q \quad a \not< p$$

hold. But

$$a \leq m \qquad b \leq m$$

in $\mathfrak{M}$, so that

$$a \leq b \quad \text{or} \quad b \leq a$$

and hence

$$a \leq q \quad \text{or} \quad b \leq p$$

neither of which can hold. ■

The problem with the size of an e. u. structure is the number of types it must realize. If we can keep this small, then we have a chance of keeping down the size of any e. u. structure we may construct.

Have another look at the calculations just before Lemma 6.26. Suppose we want to embed some countable $\mathfrak{A} \in \mathcal{S}(T)$ into some $\mathfrak{M} \in \mathcal{U}(T)$. The number of types to be considered could be $2^{\aleph_0}$, and then $|\mathfrak{M}| = 2^{\aleph_0}$ is the best we can achieve. However, sometimes we know there is a smaller number of types, and then we can do better.

Consider how we produce $\mathfrak{A}'$ from $\mathfrak{A}$ (as in the proof of Lemma 6.26). We have to realize in $\mathfrak{A}'$ each $\exists_1$-type $\Theta(a, v)$ which is T-consistent over $\mathfrak{A}'$ but has parameters from $\mathfrak{A}$. Each such type is realized in some e. c. extension of $\mathfrak{A}'$ and so is a subtype of some $\exists_1$-type which is maximal over T. Thus, in the construction, it suffices to consider only types Θ which arise by instantiation from maximal types. This in itself doesn't bring down the number of types, but we know a method that does.

Recall that for each list w of variable

$$M(T, w)$$

is the set of $\exists_1$-types (in W) each of which is maximally consistent over T. Let

$$m(T, w) = |M(T, w)|$$

so that $m(T, w) \leq 2^{\aleph_0}$, in general.

6.29 LEMMA. (The refined 1-step construction) *Let T be a theory in a countable language, and suppose $m(T, w) \leq \aleph_0$ for each list w of variable. Then for each countable $\mathfrak{A} \in \mathcal{S}(T)$ there is a countable $\mathfrak{A}' \in \mathcal{S}(T)$ with $\mathfrak{A} \subseteq \mathfrak{A}'$ and such that for each $\exists_1$-type $\Theta(a, v)$ with parameters from $\mathfrak{A}$, if this type is T-consistent over $\mathfrak{A}'$, then it is already realized in $\mathfrak{A}'$.*

Proof. We repeat the construction of the proof of Lemma 6.29. But now we have a countable enumeration

$$\boldsymbol{\Theta} = \{\Theta_i \mid i < \omega\}$$

of the types that have to be handle. Thus we can produce $\mathfrak{A}'$ as the union of an ascending ω-chain

$$\mathcal{A} = \{\mathfrak{A}_i \mid i < \omega\}$$

of structures each of which is countable. Hence $\mathfrak{A}'$ is countable. ■

The accumulation construction is always an ω-iteration iteration of the 1-step construction. Thus an iteration of Lemma 6.29 gives the following.

6.30 THEOREM. (The refined accumulation construction) *Let T be a theory in a countable language, and suppose $m(T, w) \leq \aleph_0$ for each list w of variable. Then for each countable $\mathfrak{A} \in \mathcal{S}(T)$ there is a countable $\mathfrak{M} \in \mathcal{U}(T)$ with $\mathfrak{A} \subseteq \mathfrak{M}$.*

What is the use of this? In other words, how can we keep down the size of $M(T, w)$? That is one of the topics of the next section.

Exercises–*needed*

7 A hierarchy of properties

In this section we look at a list of properties which each theory may or may not have. These properties form, in strength, a linear sequence, with the nicest (or most amenable) theories having the strongest property. Most of these properties are concerned with cardinalities, so we need to impose some global restrictions.

We analyse an arbitrary theory T under the following conditions.

- The underlying language is countable.
- The theory T has JEP.
- The theory T has no finite models (or, at least, each finite model is embeddable in an infinite model).
- *Sort out this last restriction*

Of course, at any one time we may not need all three restrictions, but each is used somewhere in the full analysis.

We consider seven properties such a theory T may or may not have.

(0) T is $\aleph_0$-categorical.

(1) T has an $\aleph_0$-categorical model companion.

(2) $\imath(\mathcal{E}(T)) = 1$.

(3) $\imath(\mathcal{E}(T)) < 2^{\aleph_0}$.

(4) There is some $\mathfrak{M} \in \mathcal{U}(T)$ with $|\mathfrak{M}| \leq \aleph_0$.

(5) T is $\exists_1$-atomic.

(6) For each $\forall_1$-formula ϕ consistent with T^0, there is a formula θ which is $\exists_1$-complete over T with $T \vdash \theta \to \phi$.

We will show that each of the six implications

$$(0) \Longrightarrow (1) \Longrightarrow (2) \Longrightarrow (3) \Longrightarrow (4) \Longrightarrow (5) \Longrightarrow (6)$$

holds (under the global restrictions given above). Along the way we will see various other properties that are equivalent to one or other of these seven properties.

The implications

$$(1) \Longrightarrow (2) \Longrightarrow (3)$$

are immediate (or even trivial). The implication

$$(5) \Longrightarrow (6)$$

is just Exercise 6.2. Thus is remains to show that

$$(0) \Longrightarrow (1) \qquad (3) \Longrightarrow (4) \Longrightarrow (5)$$

hold, and this will take us most of the rest of this section.

The crucial technique of this section is a splitting argument which enables us to produce $2^{\aleph_0}$ gadgets (of one sort or another). We will use this twice, once with Good formulas, and once with not-Bad formulas.

In the final subsection we look at a collection of examples of theories which show various of these implications are strict, that is, can not be reversed without extra restrictions.

[`Held in 120-../B71-bit..  Last changed July 26, 2004`]

7.1 Splitting with Good formulas

Recall that for each list v of variables

$$M(T, v)$$

is the set of $\exists_1$-types Σ (in v) each of which is maximally consistent with T. We let

$$m(T, v) = |M(T, v)|$$

so that (since the underlying language is countable)

$$m(T, v) \leq 2^{\aleph_0}$$

and, in general, this upper bound is achieved. In this subsection we are concerned with improving this bound (given the right circumstances). To this end we introduce a further property (4.3). Later we will see that this is equivalent to property (4), and this will explain its label.

(4.3) For each list v of variables, $m(T, v) < 2^{\aleph_0}$.

(Later it will be clear why this property is labelled in this way.)

In this subsection we show that (4.3) $\Rightarrow$ (5) holds.

For the time being fix the list v of variables

7.1 DEFINITION. For a theory T, an $\exists_1$-formula θ (in v) is Good (over T) if it is consistent with T and there is no formula ψ (in v) which is $\exists_1$-complete (over T) such that $T^0 \vdash \psi \to \theta$ holds. ■

You should compare this notion with that of an $\exists_1$-atomic theory. By inserting a negation at the appropriate places we have the following.

7.2 LEMMA. *Let T be a theory (with JEP) which is not $\exists_1$-atomic. Then there is a formula which is Good (over T).*

This shows how we can obtain at least one Good formula. The nice thing is, once we have one of them, we can generate many more.

7.3 LEMMA. (1-step splitting) *Let T be a theory and let v be a list of variables. For each formula θ (in v) which is Good (over T), there are formulas θ_0, θ_1 (in v) both of which are Good (over T), and for which*

$$T \vdash \theta_0 \to \theta \quad T \vdash \theta_1 \to \theta \quad T \vdash \neg\theta_0 \vee \neg\theta_1$$

hold.

Proof. This is little more than an unravelling of what the words mean.

Suppose θ is Good. Then there is no formula ψ with

$$T^0 \vdash \psi \to \theta$$

where ψ is $\exists_1$-complete (over T). In particular, θ itself is not $\exists_1$-complete. Thus, by Definition 6.1, there are $\exists_1$-formulas ψ_0, ψ_1 such that

$$T \cup \{\theta, \psi_0\} \text{ is consistent} \quad T \cup \{\theta, \psi_1\} \text{ is consistent} \quad T \cup \{\psi_0, \psi_1\} \text{ is not consistent}$$

(and $\partial\psi_0 \cup \partial\psi_1 \subseteq \partial\theta$). The last of these show that

$$T \vdash \neg\psi_0 \vee \neg\psi_1$$

holds. Let

$$\theta_i = \theta \wedge \psi_i$$

for $i = 0, 1$. Thus

$$T \vdash \theta_0 \to \theta \quad T \vdash \theta_1 \to \theta \quad T \vdash \neg\theta_0 \vee \neg\theta_1$$

and so it suffices to show that each of θ_0, θ_1 is Good.

The universal closure of $\theta_i \to \theta$ is an $\forall_2$-sentence, and hence

$$T^0 \vdash \theta_i \to \theta$$

holds. If θ_i is not Good, then

$$T^0 \vdash \psi \to \theta_i$$

for some formula ψ which is $\exists_1$-complete (over T), so that

$$T^0 \vdash \psi \to \theta$$

which is not so (since θ is Good). ■

This is the basic machinery we need. We are going to iterate this splitting to produce a largish family of Good formulas. Before you read the description of this iteration, you might want to re-read Solution 1.14 (which is a slightly simpler version of the argument).

We use the Cantor tree $\mathbf{\Psi}$, the full binary splitting tree. The nodes $\mu, \nu, \ldots$ of $\mathbf{\Psi}$ are the finite lists taken from $\{0, 1\}$ ordered by extension. The branches of $\mathbf{\Psi}$ are essentially the functions $p : \mathbb{N}* \longrightarrow *\{0, 1\}$. When appropriately topologized, the set $\mathcal{C}$ of all such branches is Cantor space. The crucial property is $|\mathcal{C}| = 2^{\aleph_0}$.

We use $\mathbf{\Psi}$ to index a tree of formulas.

7.4 LEMMA. (Iterated splitting) *Let T be a theory, and let θ be a formula which is Good (over T). There is a tree*

$$\Theta = (\theta_\nu \mid \nu \in \mathbf{\Psi})$$

of formulas, each of which is Good (over T) with $\theta_\perp = \theta$ and such that

$$\nu \leq \mu \Longrightarrow T \vdash \theta_\mu \to \theta_\nu \qquad \mu | \nu \Longrightarrow T \vdash \neg\theta_\mu \vee \neg\theta_\nu$$

and $\partial\theta_\nu \subseteq \partial\theta$ hold for all nodes μ, ν of $\mathbf{\Psi}$.

Proof. We construct Θ by recursion up $\mathbf{\Psi}$.

At the base we set

$$\theta_\perp = \theta$$

(the given Good formula).

Once we have θ_ν (for $\nu \in \mathbf{\Psi}$), Lemma 7.2 gives us Good formulas $\theta_{\nu 0}, \theta_{\nu 1}$ such that

$$T \vdash \theta_{\nu 0} \rightarrow \theta_\nu \quad T \vdash \theta_{\nu 1} \rightarrow \theta_\nu \quad T \vdash \neg\theta_{\nu 0} \vee \neg\theta_{\nu 1}$$

hold. This will ensure the required monotone and incomparability properties of Θ. ■

We now have all the ingredients to prove the main result of this subsection.

7.5 **THEOREM.** *Let T be a theory (in a countable language) with JEP . Then*

$$(4.3) \Longrightarrow (5)$$

holds.

Proof. In fact, we prove the contrapositive of this implication. Thus, suppose T is not $\exists_1$-atomic. Then, by Lemma 7.2, there is at least one formula θ which is Good (over T). Let $v = \partial\theta$. Lemma 7.4 gives us a tree Θ of formulas θ_ν with $\partial\theta_\nu = v$ for each $\nu \in \mathbf{\Psi}$.

Consider an $p \in \mathcal{C}$ (a branch of $\mathbf{\Psi}$). Let

$$\Sigma_p = \{\theta_\nu \mid \nu < p\}$$

to produce an $\exists_1$-type in v. We first check that this is consistent with T.

By way of contradiction, suppose that Σ_p is not consistent with T. Then there are nodes $\nu(1), \ldots, \nu(n)$ below p such that

$$T \cup \{\theta_{\nu(1)}, \ldots, \theta_{\nu(n)}\}$$

is inconsistent. Now consider and node $\nu < p$ which lies above each of $\nu(1), \ldots, \nu(n)$. By the monotone property of Θ we have $T \vdash \theta_\nu \rightarrow \theta_{\nu(j)}$ for each $1 \leq j \leq n$, and hence $T \cup \{\theta_\nu\}$ is not consistent. This contradicts the construction of Θ.

Next consider distinct branches $p, q \in \mathcal{C}$. The are incomparable nodes $\mu|\nu$ with $\mu < p$ and $\nu < q$. Thus

$$T \cup \Sigma_p \cup \Sigma_q$$

is inconsistent.

Finally, we may extend each Σ_p to a type Γ_p which is $\exists_1$-maximal over T, and then

$$\{\Gamma_p \mid p \in \mathcal{C}\}$$

is a subfamily of $M(T, v)$ of cardinality $2^{\aleph_0}$, to give $m(T, v) = 2^{\aleph_0}$. ■

This splitting technique is used many times in model theory (and in other parts of mathematics). It has many refinements, and we will see a couple of these.

7.1 The following strengthens the result of Exercise 1.14.

Let T be a consistent theory, Say a sentence σ is good over T if it is consistent with T and there is no sentence τ such that $T \cup \{\tau\}$ axiomatizes a complete theory with $T \vdash \tau \rightarrow \sigma$.

(a) Show that if no finite extension of T is complete, then each sentence which is consistent with T is good over T.

(b) Show that if there is at least one sentence that is good over T, then T has at least $2^{\aleph_0}$ complete extensions.

7.2 Say an $\exists_1$-sentence σ is not-jep over a theory T if $T \cup \{\sigma\}$ is consistent but there is no $\exists_1$-sentence τ with $T \vdash \tau \rightarrow \sigma$ and which $T \cup \{\tau\}$ is consistent and axiomatizes a theory with JEP.

Show that if there is at least one not-jep sentence, then T^g has $2^{\aleph_0}$ complete extensions each of which is the theory of some $\mathfrak{A} \in \mathcal{G}(T)$.

[*Held in 120-../B72-bit.. Last changed July 26, 2004*]

7.2 Splitting with not-Bad formulas

The notion of a Good formula is concerned with the existence (or not) of $\exists_1$-complete formulas. Such formulas can be used to produce many maximal $\exists_1$-types. In this subsection we carry out a similar splitting construction using a larger class of formulas. There are determined by certain cardinality conditions.

7.6 DEFINITION. Let T be a theory and let v be a list of variables. For each $\exists_1$-formula θ (in v) let

$$M(\theta) = \{\Sigma \in M(T, v) \mid \theta \in \Sigma\}$$

be the set of types that are $\exists_1$-maximal over T, and which contain θ.

The formula θ is Bad (over T) if $|M(\theta)| \leq \aleph_0$.

The formula θ is not-Bad (over T) if $\aleph_0 < |M(\theta)|$. ■

For instance, if θ is not consistent with T then $M(\theta) = \emptyset$, and so θ is Bad. However, θ may be Bad for other reasons.

We know some formulas which are not-Bad.

7.7 LEMMA. *Suppose the $\exists_1$-formula θ is Good (over T). Then*

$$|M(\theta)| = 2^{\aleph_0}$$

and hence θ is not-Bad.

Proof. Consider the tree of Good formulas given by Lemma. In particular, $\theta_\perp = \theta$ is the root of this tree. The branches of Θ produce $2^{\aleph_0}$ members of $M(\theta)$. ■

We work towards an analogue of the Good tree Θ using not-Bad formulas. For that we need an analogue of 1-step splitting given by Lemma 7.3.

The Bad formulas cause trouble, and we want to avoid them. Thus (for fixed T and v) let

$$M(\mathrm{Bad}) = \bigcup\{M(\theta) \mid \theta \text{ is Bad}\}$$

to collect together all the troublesome types. This is a countable union of countable sets, and hence $M(\mathrm{Bad})$ is countable. In some circumstances $M(\mathrm{Bad})$ may be the whole of $M(T, v)$, but there is a simple way to avoid that.

7.8 LEMMA. *Let T be a theory and let v be a list of variables such that $m(T, v) > \aleph_0$. Then*

$$|M(T, v) - M(\mathrm{Bad})| > \aleph_0$$

holds.

Notice that if $\Sigma \in (M(T, v) - M(\mathrm{Bad}))$ then each $\theta \in \Sigma$ is not-Bad. Thus a simple cardinality condition gives us the existence of not-Bad formulas. With this observation we can obtain the analogue of Lemma 7.3.

7.9 LEMMA. (1-step splitting) *Let T be a theory and let v be a list of variables. For each formula θ (in v) which is not-Bad (over T), there are formulas θ_0, θ_1 (in v) both of which are not-Bad (over T), for which*

$$T \vdash \theta_0 \to \theta \qquad T \vdash \theta_1 \to \theta \qquad T \vdash \neg\theta_0 \vee \neg\theta_1$$

hold.

Proof. Let θ be the given not-Bad formula. The set $M(\theta)$ is uncountable (for otherwise θ is Bad). There may be a type $\Sigma \in M(\theta)$ which has a bad member. In other words,

$$M(\theta) \cap M(\mathrm{Bad})$$

may be non-empty. However, this intersection is countable (since $M(\mathrm{Bad})$ is countable). In particular, the set

$$M(\theta) - M(\mathrm{Bad})$$

is uncountable, and so contains at least two types Σ_0, Σ_1.

Since Σ_0 and Σ_1 are $\exists_1$-maximal (over T) and distinct, we see that

$$T \cup \Sigma_0 \cup \Sigma_1$$

is not consistent. Thus there are $\psi_0 \in \Sigma_0$ and $\psi_1 \in \Sigma_1$ such that

$$T \vdash \neg\psi_0 \vee \neg\psi_1$$

holds. Let

$$\theta_1 = \neg\psi_1 \wedge \theta$$

for $i = 0, 1$. We have $\theta_1 \in \Sigma_i$, and hence θ_i is not-Bad (for otherwise $\Sigma_i \in M(\mathrm{Bad})$). The other required conditions are immediate. ■

This indicates that if we have one not-Bad formula then we can generate many more. By iterating this 1-step splitting we obtain the following analogue of Lemma 7.4.

7.10 LEMMA. (Iterated splitting) *Let T be a theory, and let θ be a formula which is not-Bad (over T). There is a tree*

$$\Theta = (\theta_\nu \mid \nu \in \mathbf{\Psi})$$

of formulas, each of which is not-Bad (over T) with $\theta_\perp = \theta$ and such that

$$\nu \le \mu \Longrightarrow T \vdash \theta_\mu \to \theta_\nu \qquad \mu|\nu \Longrightarrow T \vdash \neg\theta_\mu \vee \neg\theta_\nu$$

and $\partial\theta_\nu \subseteq \partial\theta$ hold for all nodes μ, ν of $\mathbf{\Psi}$.

Given such a tree the branches of $\mathbf{\Psi}$ will enable us to produce many types which are $\exists_1$-maximal over T. To state the appropriate result we need two variants of property (4).

(4.3) For each list v of variables, $m(T, v) < 2^{\aleph_0}$.

(4.4) For each list v of variables, $m(T, v) \le \aleph_0$.

Property (4.3) is just the same as in subsection 7.1. Notice also that (4.4) $\Rightarrow$ (4.3).

7.11 THEOREM. *For each theory T (in a countable language)*

$$(4.3) \Longleftrightarrow (4.4)$$

holds.

Proof. It suffices to show (4.3) $\Rightarrow$ (4.4), and for that we prove the contrapositive.

Thus suppose $m(T, v) > \aleph_0$ for some list v of variables. This gives us at least one $\exists_1$-formula θ (in v) which is not-Bad (over T). But now, by Lemma 7.10, we have a whole tree Θ of such formulas (with $\theta_\perp = \theta$). The branches of Θ show that $m(T, v) \ge 2^{\aleph_0}$. ■

You should understand how this splitting technique refines the one used in subsection 7.1. There is doesn't matter how we spilt a Good situation, for the global conditions ensure the we obtain two Good situations. Here we have to be more careful. There are certain splittings which don't work, and we may stay away from these. The global conditions give us enough room to do this.

Exercises

7.3 Show that for each theory T (in a countable language) and each $\exists_1$-formula θ, either $|M(\theta)| \le \aleph_0$ or $2^{\aleph_0} = |M(\theta)|$.

7.4 *Use one of the french paper results*

[*Held in 120-../B73-bit.. Last changed July 26, 2004*]

7.3 Countable existentially universal structures

By Theorem 6.27 we know that each $\mathfrak{A} \in \mathcal{S}(T)$ is embeddable in some $\mathfrak{M} \in \mathcal{U}(T)$. However, using that general construction we can not keep the cardinality of $\mathfrak{M}$ below $2^{\aleph_0}$. This is because, potentially, there are $2^{\aleph_0}$ possible types which may have to be realized in $\mathfrak{M}$. In general, as Example 6.29 shows, this is the best we can do. To obtain a countable existentially universal structure we need to impose certain restrictions on the parent theory (and the underlying language). Theorem 6.30 indicates one way to achieve this.

Is the axiomatized version of that theory complete or model complete?. Does it have JEP or AP?

At this point we produce a list of the variants of property (4). Two of these variants have occurred already.

7.12 THEOREM. *Let T be a theory in a countable language, and suppose T has JEP. The following are equivalent.*

(4) There is some countable $\mathfrak{M} \in \mathcal{U}(T)$.

(4.1) There is some countable structure $\mathfrak{M} \in \mathcal{S}(T)$ such that each countable structure $\mathfrak{A} \in \mathcal{S}(T)$ is embeddable in $\mathfrak{M}$.

(4.2) There is some structure $\mathfrak{M} \in \mathcal{S}(T)$ with $|\mathfrak{M}| < 2^{\aleph_0}$ such that each countable structure $\mathfrak{A} \in \mathcal{S}(T)$ is embeddable in $\mathfrak{M}$.

(4.3) For each list v of variables, $m(T, v) < 2^{\aleph_0}$.

(4.4) For each list v of variables, $m(T, v) \leq \aleph_0$.

Proof. (4)⇒(4.1). This is a consequence of Theorem 6.18.

(4.1)⇒(4.2). This is trivial.

(4.2)⇒(4.3). Let $\mathfrak{M} \in \mathcal{S}(T)$ be the structure given by (4.2). Let v be any list of variables. Consider any $\Sigma \in M(T, v)$. This type is realized in some countable $\mathfrak{A} \in \mathcal{E}(T)$, and this is embeddable in $\mathfrak{M}$, so that Σ is realized in $\mathfrak{M}$. Thus $m(T, v) \leq |\mathfrak{M}| < 2^{\aleph_0}$.

(4.3)⇒(4.4). This is Theorem 7.11.

(4.4)⇒(4). This follows by the refined existence result, Theorem 6.30. ■

With this we can obtain two of the remaining required implications.

7.13 THEOREM. *Let T be a theory in a countable language, and suppose T has JEP. Then two implications*

$$(3) \Longrightarrow (4) \Longrightarrow (5)$$

hold.

Proof. A simple counting argument gives (3) ⇒ (4.3), and we have (4.3) ⇒ (5) by Theorem 7.5. ■

End chat

[*Held in 120-../B74-bit.. Last changed July 26, 2004*]

7.4 Categoricity properties

In this subsection we prove the remaining implication

$$(0) \Longrightarrow (1)$$

and, along the way, we obtain several equivalent variants of these two properties. In this vein, we begin with some equivalents of (2).

Recall that $\mathcal{A}(T)$ is the class of $\exists_1$-atomic structures for T.

7.14 THEOREM. *Let T be a theory in a countable language, and suppose T has JEP. The following are equivalent.*

(2) $\imath(\mathcal{E}(T)) = 1$

(2.1) $\mathcal{A}(T) \cap \mathcal{U}(T) \neq \emptyset$

(2.2) Each $\exists_1$-type which is maximally consistent with T contains a formula that is $\exists_1$-complete (over T).

(2.3) $\mathcal{A}(T) = \mathcal{E}(T)$

Proof. (2)⇒(2.1). Since

$$(2) \Longrightarrow (4) \Longrightarrow (5)$$

the hypothesis (2) gives structures

$$\mathfrak{M} \in \mathcal{U}(T) \subseteq \mathcal{E}(T) \qquad \mathfrak{A} \in \mathcal{A}(T) \subseteq \mathcal{E}(T)$$

where both are countable. But then (2) gives $\mathfrak{M} \cong \mathfrak{A}$, and hence (2.1) holds.

(2.1)⇒(2.2). Consider any

$$\mathfrak{A} \in \mathcal{A}(T) \cap \mathcal{U}(Y)$$

(where we do not assume that $\mathfrak{A}$ is countable). Consider any type $\Sigma(v) \in M(T, v)$. Since T has JEP and $\mathfrak{A} \in \mathcal{U}(T)$, this type is realized by some point a of $\mathfrak{A}$. In fact, by Lemma 5.25, we have

$$\mathfrak{A} \models \psi(a) \Longleftrightarrow \psi(v) \in \Sigma(v)$$

for each $\exists_1$-formula $\psi(v)$. Since $\mathfrak{A} \in \mathcal{A}(T)$, there is some formula $\theta(v)$ which is $\exists_1$-complete (over T) and $\mathfrak{A} \models \theta(a)$. This is the required formula.

(2.2)⇒(2.3). Consider any $\mathfrak{A} \in \mathcal{E}(T)$. Consider any point a of $\mathfrak{A}$, and let $\Sigma(v)$ be the $\exists_1$-type of a in $\mathfrak{A}$. By Lemma 5.25 we have $\Sigma(v) \in M(T, v)$. By the assumption (2.2) there is some $\theta(v) \in \Sigma$ which is $\exists_1$-complete (over T). Hence $\mathfrak{A} \in \mathcal{A}(T)$.

(2.3)⇒(2). This is an immediate consequence of Corollary 6.9. ■

Next we analyse property (1) and provide a collection of several variants.

7.15 THEOREM. *Let T be a theory in a countable language, and suppose T has JEP. The following are equivalent.*

(1) T has an $\aleph_0$-categorical model companion.

(1.1) T^0 is $\aleph_0$-categorical.

(1.2) For each $\forall_1$-type Π which is consistent with T^0, there is an $\exists_1$-formula which is $\exists_1$-complete (over T) with $\partial(\theta) \subseteq \partial(\Pi)$ and $T \vdash \phi \rightarrow \bigwedge \Pi$.

(1.3) For each list v of variables, the are are formulas $\theta_1, \ldots, \theta_n$ in v (and where n may depend on v) where each is $\exists_1$-complete (over T) and

$$T^0 \vdash \theta_1 \vee \cdots \vee \theta_n$$

holds.

(1.4) For each list v of variables, $m(T, v) < \aleph_0$.

(1.5) For each list v of variables, there are up to T^0-equivalence only finitely many $\exists_1$-formulas in v.

Proof. (1)⇒(1.1). This is immediate.

(1.1)⇒(1.2). [*`This may need no finite models`*] Assuming (1.1), let $\mathfrak{A}$ be the unique countable model of T^0. Since (2)⇒(1), we see that $\mathfrak{A} \in \mathcal{A}(T)$. Consider any $\forall_1$-type $\Pi(v)$ which is consistent with T^0. This is realized in some countable model of T^0, and hence in $\mathfrak{A}$ by some point. Since $\mathfrak{A} \in \mathcal{A}(T)$, this point satisfies some formula $\theta(v)$ which is $\exists_1$-complete over T, which leads to the required result.

(1.2)⇒(1.3). Let v be a list of variables, and consider the $\forall_1$-type

$$\{\neg\theta \mid \partial\theta \subseteq v \text{ and } \theta \text{ is } \exists_1\text{-complete (over } T)\}$$

(as used in the proof of Theorem 6.6). The assumption (1.2) implies that Π is not consistent with T, and hence we obtain (1.3).

(1.3)⇒(1.4). For each v list of variables, each $\Sigma \in M(T, v)$ must contain one of the formula $\theta_1, \ldots, \theta_n$ provided by (1.3). Let θ be this formula. Since this is $\exists - 1$-complete (over T) we have

$$\psi \in \Sigma \Longrightarrow T^0 \vdash \theta \rightarrow \psi$$

and hence Σ is uniquely determined by θ. Thus $m(T, v) \leq n$.

(1.4)⇒(1.5). Two $\exists_1$-formulas, in the same list v of variable, are T^0-equivalent if and only if they are equivalent in each member of $\mathcal{E}(T)$, and hence belong to exactly the same members of $M(T, v)$.

(1.5)⇒(1). Suppose (2.5) holds. We first show that T^0 is model complete.

Consider any $\forall_1$-formula $\phi(v)$ which is consistent with T^0. Since T^0 is 0-complete, there is at least one $\exists_1$-formula $\theta(v)$ which is consistent with T and such that $T^0 \vdash \theta \rightarrow \phi$ holds. By (1.5) there are only finitely many such formulas

$$\theta_1, \ldots, \theta_n$$

up to T^0-equivalence. It suffices to show that

$$T^0 \vdash \phi \to (\theta_1 \vee \cdots \vee \theta_n)$$

holds.

By way of contradiction, suppose this is not so. Thus the $\forall_1$-formula

$$\psi = \phi \wedge \neg\theta_1 \wedge \cdots \wedge \neg\theta_n$$

is consistent with T^0, and hence

$$T^0 \vdash \theta \to \psi$$

for some $\exists_1$-formula θ which is consistent with T^0. In particular,

$$T^0 \vdash \theta \to \phi$$

so that θ is T^0-equivalent to one of $\theta_1, \ldots, \theta_n$. But then

$$T^0 \vdash \theta \to \neg\theta$$

which is the contradiction.

This show that T^0 is model complete. It remains to show that T^0 is $\aleph_0$-categorical.

Consider any countable model $\mathfrak{A} \models T^0$. Since T^0 is model complete, we have $\mathfrak{A} \in \mathcal{E}(T)$. Consider any point a of $\mathfrak{A}$, and let Σ be the $\exists_1$-type of a in $\mathfrak{A}$. By (1.5), this type is T^0-equivalent to a single $\exists_1$-formula. Thus, by Lemma 6.4(ii), we see that $\mathfrak{A} \in \mathcal{A}(T)$.

The required catogoricity now follows from Corollary 6.10. ■

Before we give a similar analysis let's look at a particular case of this last result. We will use this to prove (0) $\Rightarrow$ (1).

Consider the case of Theorem 7.15 where the given theory T is known to be model complete (and hence complete). For such a theory properties (1) and (1.1) coalesce to

(0) T is $\aleph_0$-categorical

(the top property of the original list). Since T is model complete, each formula is T-equivalent to both an $\exists_1$-formula and a $\forall_1$-formula. This, in (1.2, 1.3, 1.4, 1.5) all references to quantifier complexity can be removed. In this way we see that the model complete case of Theorem 7.15 is also a particular case of the following result. (In the statement of this result the variants are labelled to match the corresponding variants in Theorem 7.15. Thus there is no variant (0.1).)

7.16 THEOREM. *Let T be a complete theory in a countable language. The following are equivalent.*

(0) T is $\aleph_0$-categorical.

(0.2) For each type Γ which is consistent with T, there is a formula which is complete (over T) with $\partial(\theta) \subseteq \partial(\Gamma)$ and $T \vdash \phi \to \bigwedge \Gamma$.

(0.3) For each list v of variables, the are are formulas $\theta_1, \ldots, \theta_n$ in v (and where n may depend on v) where each is complete (over T) and

$$T \vdash \theta_1 \vee \cdots \vee \theta_n$$

holds.

(0.4) For each list v of variables, $f(T, v) < \aleph_0$.

[`Use` $F(T, v)$ `for the full types over` T]

(0.5) For each list v of variables, there are up to T-equivalence only finitely many formulas in v.

Proof. We sketch two different proofs.

(First) By a process of adding new relation symbols we can enrich the underlying language in such a way that T axiomatizes a model complete theory T^+ in the extended language. Furthermore, we can ensure that the models of T^+ are essentially the same as the models of T. We may then apply the model complete case of Theorem 7.15 to T^+.

(Second) We may rework the proof of Theorem 7.15, but this time dealing with arbitrary formulas and types (without any restrictions on the quantifier complexity). ■

Neither of these two sketched proofs is better than the other. Initially we may want just the bare bones of the characterization, as stated. In that the case the first proof will do. However, in more advanced work we may need more information, or more control on the various structures involved. In that case we need to refine the second proof.

With these two characterzations we can complete the chain of properties.

7.17 THEOREM. *Let T be a theory in a countable language, and suppose T has JEP. Then*

$$(0) \Longrightarrow (1)$$

holds.

Proof. Suppose T is $\aleph_0$-categorical. Then [*`no finite models`*] T is complete and satisfies (0.1 – 0.5). For each list v of variables we have

$$m(T, v) \leq f(T, v)$$

(since each $\Sigma \in M(T, v)$ can be extended to a full type $\Gamma \in F(T, v)$ in which there are no new $\exists_1$-formulas.) Thus

$$(0.4) \Longrightarrow (1.4)$$

holds, to show that T has (1). ■

`End chat`

Exercises

7.5 Let T be a theory in a countable language and suppose that T has JEP. Show that T has an $\aleph_0$-categorical model companion if and only if $\mathcal{M}(T^0) = \mathcal{A}(T)$.

7.6 [*`Check this`*] Let T be a theory in a countable language which is $\forall_2$-axiomatizable theory, has no finite models, and is $\aleph_0$-categorical. Show that T is model complete.

7.7 Let T be a complete theory in a countable language. Show that the following are equivalent.

(i) T is $\aleph_0$-categorical.

(ii) T has an atomic model which is countably saturated.

(iii) Each model of T is atomic.

7.5 Some particular examples

Do we need this

(0)$\not\Leftarrow$(1). Let T be the theory of linearly ordered sets and let T^* be the theory of lines. Then T^* is $\aleph_0$-categorical and is the model companions of T, but T is not α_0-categorical.

(1)$\not\Leftarrow$(2). *?????????*

(2)$\not\Leftarrow$(3). This is related to Vaught's conjecture, and preseumably is quite difficult.

(3)$\not\Leftarrow$(4). *?????????*

(4)$\not\Leftarrow$(5). Try the theory of the Cantor tree. [*Sort out the details*]

(5)$\not\Leftarrow$(6). Try the modified successor example which has $T^e = T^f = T^g = T^0$.

Further remarks

****

(3) .v. (4)

On the face of it it seems that

$$(4) \Longrightarrow (3)$$

might hold. Here is the bones of a proof.

Consider any countable $\mathfrak{M} \in \mathcal{U}(T)$. Each countable $\mathfrak{A} \in \mathcal{E}(T)$ is embeddable in $\mathfrak{M}$. So $\imath(\mathcal{E}(T)) \leq$ number of substructures of $\mathfrak{M}$. Why should this be small?

Perhaps some stability conditions will give this.

Part II

Construction techniques

8 The construction of canonical models

In this section we look at a method by which we can produce a structure out of a set of sentences. Of course, the aim is to make the structure a model of the set. We see how to achieve this provided we start from a suitable set.

The construction technique we use has several phases, each of which can be modified to produce extra properties of the end result. This is important in more advanced work. (Indeed, the account here is based on Chapter 3 of [9] which is about languages which are not even first order.)

We use the method to prove the compactness theorem. A slight modification gives a proof of the completeness theorem. As indicated, the method is also the basis of several more sophisticated results.

Let L be the underlying language. The crucial trick is to work in an enrichment of L. Let W be a set of new constant symbols with $|W| = |L|$. We work in the enriched language $L(W)$. Each constant symbol $a \in W$ does a special kind of job. We refer to each member $a \in W$ as a witness. Note that $|L| = |W| = |L(W)|$. This will be important.

Here is the strategy we employ.

- The structure is built directly out of the witnesses. Thus each element of the structure is a certain equivalence class of witnesses. The other attributes are formed using a modified term algebra (or Lindenbaum algebra) construction.

- In general, a term algebra is suitable for modelling equational properties, and positive combinations of these. Here we need to handle quantification and negation. We arrange that every existential quantification is witnessed by an element of W (which is why these are called witnesses). This trick is sometimes called a Henkin construction, but similar techniques occur in other parts of mathematical logic.

- To ensure everything goes through as we want it to, we work with a special kind of set of $L(W)$-sentences. This has certain closure properties which help us with the construction and the required verifications. It is a particular kind of Hintikka set.

- To obtain such a helpful set we use a maximizing argument. We want the witnesses to behave in an appropriate way, so this maximization must be done with some care. To guide this process we use a consistency property, a certain family of sets of $L(W)$-sentences.

In most parts of model theory we do not need to distinguish between logically equivalent sentences. However, here several verifications proceed by induction over the construction of the sentences involved. Furthermore, most of the verification are of implications, so the step across a use of the negation symbol can cause trouble. Luckily, there is an observation, due to Smullyan, which helps cut down the amount of work.

Each sentence of $L(W)$ is either atomic, or of one of the shapes

$$(\rho \wedge \sigma) \quad (\rho \vee \sigma) \quad (\rho \to \sigma) \quad (\forall v)\phi(v) \quad (\exists v)\phi(v)$$

are the negation of one of these, or a doubly negated sentences. Here ρ and σ are simpler sentences, and $\phi(v)$ is a formula with just the one free variable, as indicated. This gives us 13 cases which suggests that each induction step will be split into 13 part. However, we can group these shapes into blocks where the members of each block are dealt with in a similar way.

Each sentence of $L(W)$ has one of the shapes

$$
\begin{array}{llll}
(0) & \alpha & \neg\alpha & \\
(\neg) & \neg\neg\tau & & \\
(\wedge) & (\rho \wedge \sigma) & \neg(\rho \vee \sigma) & \neg(\rho \to \sigma) \\
(\vee) & (\rho \vee \sigma) & \neg(\rho \wedge \sigma) & (\rho \to \sigma) \\
(\forall) & (\forall v)\phi(v) & \neg(\exists v)\phi(v) & \\
(\exists) & (\exists v)\phi(v) & \neg(\forall v)\phi(v) &
\end{array}
$$

where α is atomic, where ρ, σ, τ are arbitrary, and where $\phi(v)$ is a formula with just the one free variable. Notice how these are grouped. For instance, the sentences of $(\wedge)$ are logically equivalent to

$$(\rho \wedge \sigma) \quad (\neg\rho \wedge \neg\sigma) \quad (\rho \wedge \neg\sigma)$$

respectively. We find that if we can handle the principal, left hand, shape of a group, then we can handle all the shapes of that group.

In the next two subsections we describe the details of this method of construction.

[`Held in 120../C11-bit..  Last changed July 26, 2004`]

8.1 Helpful set and its canonical model

In this subsection we show how to turn a suitable set Ξ of $L(W)$-sentences into a structure $\mathfrak{M}(\Xi)$ which is a model of Ξ. Furthermore, each element of $\mathfrak{M}(\Xi)$ is named by some witness. The trick is to build this structure on the set $A = W/\sim$ of blocks of some equivalence relation $\sim$ on the set W of witnesses.

Of course, the set Ξ needs to have some rather spacial properties to help us with the construction of $\mathfrak{M}(\Xi)$.

8.1 DEFINITION. A set Ξ of $L(W)$-sentences is helpful if it has the closure properties listed in Table 2. ■

In the next subsection we see how to generate helpful sets. For this subsection we assume we have some such set Ξ. We construct $\mathfrak{M}(\Xi)$ in several phases, some of which are quite finicky. Unfortunately, these occur at the beginning.

We start by constructing the carrier A.

8.2 DEFINITION. Let $\sim$ be the relation on W given by

$$a \sim b \iff (a \doteq b) \in \Xi$$

(for $a, b \in W$). ■

$(0)(i)$	$\mathsf{false} \notin \Xi$ and $\neg\mathsf{true} \notin \Xi$	
$(0)(ii)$	$\alpha \notin \Xi$ or $\neg\alpha \notin \Xi$	for each atomic sentence α
$(\neg)$	$\neg\neg\tau \in \Xi \Longrightarrow \tau \in \Xi$	
$(\wedge)(i)$	$(\rho \wedge \sigma) \in \Xi \Longrightarrow \rho, \sigma \in \Xi$	
$(\wedge)(ii)$	$\neg(\rho \vee \sigma) \in \Xi \Longrightarrow \neg\rho, \neg\sigma \in \Xi$	
$(\wedge)(iii)$	$\neg(\rho \to \sigma) \in \Xi \Longrightarrow \neg\rho, \sigma \in \Xi$	
$(\vee)(i)$	$(\rho \vee \sigma) \in \Xi \Longrightarrow \rho \in \Xi$ or $\sigma \in \Xi$	
$(\vee)(ii)$	$\neg(\rho \wedge \sigma) \in \Xi \Longrightarrow \neg\rho \in \Xi$ or $\neg\sigma \in \Xi$	
$(\vee)(iii)$	$(\rho \to \sigma) \in \Xi \Longrightarrow \neg\rho \in \Xi$ or $\sigma \in \Xi$	
$(\forall)(i)$	$(\forall v)\phi(v) \in \Xi \Longrightarrow \phi(a) \in \Xi$	for each $a \in W$
$(\forall)(ii)$	$\neg(\exists v)\phi(v) \in \Xi \Longrightarrow \neg\phi(a) \in \Xi$	for each $a \in W$
$(=)(i)$	$(a \doteq b) \in \Xi \Longrightarrow (b \doteq a) \in \Xi$	for each $a, b \in W$
$(=)(ii)$	$(a \doteq t), \phi(t) \in \Xi \Longrightarrow \phi(a) \in \Xi$	for each $a \in W$ and closed term t
$(W)(t)$	For each closed term t, there is some $a \in W$ with $(a \doteq t) \in \Xi$	
$(W)(\exists)$	$(\exists v)\phi(v) \in \Xi \Longrightarrow \phi(a) \in \Xi$	for some $a \in W$
$(W)(\forall)$	$\neg(\forall v)\phi(v) \in \Xi \Longrightarrow \neg\phi(a) \in \Xi$	for some $a \in W$

Table 2: The closure properties of a helpful set Ξ

This is the equivalence relation which generates the blocks of W which we make the elements of the canonical model. We need to check this first assertion.

8.3 LEMMA. *The constructed relation $\sim$ is an equivalence relation on W.*

Proof. We must show that $\sim$ is reflexive, symmetric, and transitive.

(R) Consider any $a \in W$. This is a closed term, so $(W)(t)$ gives $(b \doteq a) \in \Xi$ for some $b \in W$. By $(=)(i)$ we have $(a \doteq b) \in \Xi$. Let $\phi(v)$ be $(a \doteq v)$. Then $(a \doteq b), \phi(b) \in W$, so that $(=)(ii)$ give $\phi(a) \in W$, and hence $a \sim a$.

(S) This is an immediate consequence of $(=)(i)$.

(T) Suppose $a \sim b \sim c$ for $a, b, c \in \Xi$. Let $\phi(v)$ be $(v \doteq c)$. Thus $(a \doteq b), \phi(b) \in \Xi$, so that $(=)(ii)$ gives $\phi(a) \in \Xi$, and hence $a \sim c$, as required. ■

Since $\sim$ is an equivalence relation on W, we may form the set $A = W/\sim$ of $\sim$-blocks. For each $a \in W$ let

$$[a] = \{x \in W \mid a \sim x\}$$

to form a typical block $[a]$ with a representative a. As we build the structure carried by A we must check that each construction is well defined, that is it is independent of the choice of the representatives used. As usual, this can be a little tedious.

To produce a structure carried by A we must construct interpretations for each symbol of the signature of $L(W)$. Thus we require attributes

$$[\![K]\!] \quad [\![R]\!] \quad [\![O]\!] \quad [\![a]\!]$$

for each constant symbol, each relation symbol, and each operation symbol O of L, and each $a \in W$. Clearly, we take

$$[\![a]\!] = [a]$$

for each $a \in W$. The other attributes take just a little longer to produce.

As we go through the construction we will need several facts concerning the interaction of $\sim$ and the set Ξ. Here is the first of these.

8.4 LEMMA. *Suppose* $(a \doteq t) \in \Xi$ *for some* $a \in W$ *and closed term* t. *Then*

$$(b \doteq t) \in \Xi \Longleftrightarrow a \sim b$$

holds for each $b \in W$.

Proof. Suppose $(b \doteq t) \in \Xi$. Let $\phi(v)$ be $(b \doteq t)$, so that $(a \doteq t), \phi(t) \in \xi$. Thus $(=)(ii)$ give the required result.

Conversely, suppose $a \sim b$. Let $\phi(v)$ be $(v \doteq t)$, so that $(b \doteq a), \phi(a) \in \Xi$. Thus $(=)(ii)$ give the required result. ■

The constant symbol K is a close term. Thus, by $(W)(t)$, we have

$$(a \doteq K) \in \Xi$$

for at least on $a \in W$. There may be several such a. However, Lemma 8.4 gives

$$(b \doteq K) \in \Xi \Longleftrightarrow a \sim b$$

so we may take

$$[\![K]\!] = [a]$$

and use any $b \in [a]$ as a name of $[\![K]\!]$.

We use a similar method to produce $[\![R]\!]$ and $[\![O]\!]$. But before we look at the details, let's obtain a useful generalization of $(=)(ii)$.

8.5 LEMMA. *Let* $\phi(v_1, \ldots, v_n)$ *be a formula (of* $L(W)$*) with indicated free variables. Let* $a_1, \ldots, a_n \in W$ *and let* $t_1, \ldots, t_n$ *be closed terms such that*

$$(a_i \doteq t_i) \in \Xi$$

for each $1 \leq i \leq n$. *Then*

$$\phi(t_1, \ldots, t_n) \in \Xi \Longrightarrow \phi(a_1, \ldots, a_n) \in \Xi$$

holds.

Proof. For the time being fix i with $1 \le i \le n$. let $\phi_i(v)$ be

$$\phi(t_1, \ldots, t_{i-1}, v, t_{i+1}, \ldots, t_n)$$

where the free variable occurs in position i of ϕ and all the other free variables are instantiated. Then $(=)(ii)$ show that

$$(a_i \doteq t_i), \phi(t_i) \in \Xi \Longrightarrow \phi_i(a_i)$$

holds. We use this several times for different i and different selections of terms. Thus

$$\begin{aligned}
\phi(t_1, t_2, \ldots, t_{n-1}, t_n) \in \Xi &\Longrightarrow \phi(a_1, t_2, \ldots, t_{n-1}, t_n) \in \Xi \quad \text{since } (a_1 \doteq t_1) \in \Xi \\
&\Longrightarrow \phi(a_1, a_2, \ldots, t_{n-1}, t_n) \in \Xi \quad \text{since } (a_2 \doteq t_2) \in \Xi \\
&\quad \vdots \\
&\Longrightarrow \phi(a_1, a_2, \ldots, a_{n-1}, t_n) \in \Xi \quad \text{since } \ldots \\
&\Longrightarrow \phi(a_1, a_2, \ldots, a_{n-1}, a_n) \in \Xi \quad \text{since } (a_n \doteq t_n) \in \Xi
\end{aligned}$$

to give the required result. ■

It's time to produce $[\![R]\!]$ and $[\![O]\!]$.

Let R be a n-placed relation symbol of the signature of L. How do we produce the n-placed relation $[\![R]\!]$ on A? We must determine when

$$[\![R]\!][a_1] \cdots [a_n]$$

is true and when it is false for each $[a_1], \ldots, [a_n] \in A$.

We wish to define $[\![R]\!]$ by

$$[\![R]\!][a_1] \cdots [a_n] \Longleftrightarrow Ra_1 \cdots a_n \in \Xi$$

where $a_1, \ldots, a_n$ are representative of the blocks involved. However, this definition will make sense only if the equivalence is independent of the representatives.

Let $b_1, \ldots, b_n$ be a different set of representatives. Thus

$$a_i \sim b_i$$

for each $1 \le i \le n$. Using Lemma 8.5 we have

$$Ra_1 \cdots a_n \in \Xi \Longleftrightarrow Rb_1 \cdots b_n \in \Xi$$

so that, as required, the attempted construction of $[\![R]\!]$ is well defined.

Let O be a n-placed operation symbol of the signature of L. How do we produce the n-placed operation $[\![O]\!]$ on A? We must determine an output

$$[\![O]\!][a_1] \cdots [a_n]$$

in A for each selection $[a_1], \ldots, [a_n]$ of inputs in A.

Let $a_1, \ldots, a_n$ be representatives of the input blocks. Observe that

$$t = Oa_1 \cdots a_n$$

is a closed term of $L(W)$, and hence $(W)(t)$ gives

$$(a \simeq Oa_1 \cdots a_n) \in \Xi$$

for some $a \in W$. We wish to set

$$[\![O]\!][a_1] \cdots [a_n] = [a]$$

but before we can do that, we must check that it makes sense.

Let $b_1, \ldots, b_n \in W$ satisfy

$$a_i \sim b_i$$

for each $1 \leq i \leq n$. Thus we have a second set of representative for the input blocks. Let

$$(b \simeq Ob_1 \cdots b_n) \in \Xi$$

so that $[b]$ is an attempted output for these representatives. A use of Lemma 8.5 gives

$$(b \simeq Oa_1 \cdots a_n) \in \Xi$$

and then Lemma 8.4 gives $a \sim b$ so that, as required, the attempted construction of $[\![O]\!]$ is well defined.

This completes the construction of the structure $\mathfrak{M}(\Xi)$. We now begin to relate it properties to those of Ξ, and eventually we show that $\mathfrak{M}(\Xi)$ is a model of Ξ. We work through several phases following the overall definition of the satisfaction relation given in section 1.3.

Consider a closed term t of $L(W)$. This has some value $[\![t]\!]$ in A as generated by Definition 1.13. Here, however, we do not need an assignment, (since each term we meet is closed). There is also a more direct way to determine $[\![t]\!]$.

8.6 LEMMA. *The implication*

$$(a \simeq t) \in \Xi \Longrightarrow [\![t]\!] = [a]$$

holds for each closed term t and $a \in W$.

Proof. We proceed by induction on the construction of t.

The base case, when t is a constant symbol K, is immediate from the definition of $[\![K]\!]$.

For the induction step, suppose

$$t = Ot_1 \cdots t_n$$

where O is an n-placed operation symbol and $t_1, \ldots t_n$ a smaller terms. For each $1 \leq i \leq n$, $(W)(t)$ gives some $a_i \in W$ with $(a_i \simeq t_i) \in \Xi$, and then $[\![t_i]\!] = [a_i]$ follows by the induction hypothesis. Consider any $a \in W$ such that

$$(a \simeq t) \in \Xi$$

holds. Lemma 8.5 gives

$$(a \simeq Oa_1 \cdots a_n) \in \Xi$$

and then Definition 1.13 gives the required

$$[\![t]\!] = [\![O]\!][\![t_1]\!]\cdots[\![t_n]\!] = [\![O]\!][a_1]\cdots[a_n] = [a]$$

where the last step holds by the definition of $[\![O]\!]$. ■

This is most of the finicky work done. To show that $\mathfrak{M}(\Xi)$ is a model of Ξ we show that

$$\tau \in \Xi \Longrightarrow \mathfrak{M}(\Xi) \models \tau$$

holds for each sentence τ. Of course, we proceed by induction on the construction of τ.

8.7 LEMMA. *For each atomic sentence τ both*

$$\tau \in \Xi \Longrightarrow \mathfrak{M}(\Xi) \models \tau \qquad \neg\tau \in \Xi \Longrightarrow \mathfrak{M}(\Xi) \models \neg\tau$$

hold.

Proof. Each atomic sentence has one of the forms

$$\mathsf{true} \quad \mathsf{false} \quad (s \doteq t) \quad Rt_1\cdots t_n$$

where $s, t, t_1, \ldots, t_n$ are closed terms and R is a n-placed relation symbol. Let's look at the case where τ has the fourth form.

By $(W)(t)$ there are $a_i \in W$ such that

$$(a_i \doteq t_i) \in \Xi$$

for each $1 \leq i \leq n$, and then

$$[\![t_i]\!] = [a_i]$$

by Lemma 1.6. Thus

$$\mathfrak{M}(\Xi) \models \tau \Longleftrightarrow [\![R]\!][\![t_1]\!]\cdots[\![t_n]\!] \Longleftrightarrow [\![R]\!][a_1]\cdots[a_n] \Longleftrightarrow Ra_1\cdots a_n \in \Xi$$

using Definition 1.16 and the definition of $[\![R]\!]$. With this Lemma 8.5 gives

$$\tau \in \Xi \Longrightarrow Ra_1\cdots a_n \in \Xi \Longrightarrow \mathfrak{A} \models \tau$$

and a similar argument using $(0)(ii)$ gives

$$\neg\tau \in \Xi \Longrightarrow \neg(Ra_1\cdots a_n) \in \Xi \Longrightarrow Ra_1\cdots a_n \notin \Xi \Longrightarrow \mathrm{not}(\mathfrak{M}(\Xi) \models \tau) \Longrightarrow \mathfrak{M}(\Xi) \models \neg\tau$$

as required. ■

So far we haven't used any of the conditions $(\neg)$ — $(\forall)(ii)$, and neither $(W)(\exists)$ nor $(W)(\forall)$ These merely unravel the semantics of $\mathfrak{M}(\Xi)$ and lead to a proof of the following.

8.8 THEOREM. *The implication*

$$\tau \in \Xi \Longrightarrow \mathfrak{M}(\Xi) \models \tau$$

holds for each sentence τ. In particular, $\mathfrak{M}(\Xi)$ is a model of Ξ.

Naturally, we call $\mathfrak{M}(\Xi)$ the canonical model of Ξ.

Exercises

8.1 Complete the proof of Theorem 8.8.

[*Held in 120-../C17-bit.. Last changed July 26, 2004*]

8.2 Consistency property

By Theorem 8.8, to show that a set Σ of $L(W)$-sentences has a model, it is sufficient to embed Σ in a helpful set Ξ. We do this by maximizing in a certain way. The conditions $(0)(i)$ — $(=)(ii)$ are fairly easy to arrange, but the three witnessing conditions $(W)(t), (W)(\exists), (W)(\forall)$ are more delicate. We need a mechanism to help guide the maximizing process.

We work with certain families $\boldsymbol{Con}$ of sets Σ of $L(W)$-sentences. These families must have certain closure properties.

8.9 DEFINITION. Let $\boldsymbol{Con}$ be a family of sets $L(W)$-sentences.

$\boldsymbol{Con}$ is downward closed if

$$\Gamma \subseteq \Sigma \in \boldsymbol{Con} \Longrightarrow \Gamma \in \boldsymbol{Con}$$

holds.

$\boldsymbol{Con}$ has finite character if $\Sigma \in \boldsymbol{Con}$ whenever $\Gamma \in \boldsymbol{Con}$ for each finite $\Gamma \subseteq \Sigma$. ■

The finite character ensures that a consistency property $\boldsymbol{Con}$ is closed under union of chains (and directed families), and so Zorn's Lemma gives us maximal members of $\boldsymbol{Con}$. Unfortunately, this process is too crude to give us what we need.

8.10 DEFINITION. A consistency property is a family $\boldsymbol{Con}$ of sets of $L(W)$-sentences which is downward closed, has finite character, and for each $\Sigma \in \boldsymbol{Con}$ the conditions of Table 3 hold. ■

There are several examples of consistency properties. The one we need (for the compactness theorem) is based on finite satisfiability.

8.11 THEOREM. *The family* $\boldsymbol{FinSat}$ *of finitely satisfiable sets of* $L(W)$*-sentences is a consistency property.*

Proof. Almost trivially, $\boldsymbol{FinSat}$ is downward closed and has finite character. It is routine to check that $\boldsymbol{FinSat}$ has properties $(0)(i)$ — $(=)(ii)$, for these just unravel the semantics of sentences. The three witnessing conditions $(W)(t), (W(\exists)$, and $(W)(\forall)$ are not so immediate.

Consider any $\Sigma \in \boldsymbol{FinSat}$ and look at

$$\Sigma \cup \{(a \doteq t)\}$$

where t is a closed term and a is a witness that does not appear in Σ not in t. Any finite subset of this set is a subset of

$$\Gamma \cup \{(a \doteq t)\}$$

A family $\boldsymbol{Con}$ of sets of $L(W)$-sentences is a consistency property if it is downward closed, has finite character, and for each $\Sigma \in \boldsymbol{Con}$, the following hold.

$(0)(i)$	$\mathsf{false} \notin \Sigma$ and $\neg\mathsf{true} \notin \Sigma$	
$(0)(ii)$	$\alpha \notin \Sigma$ or $\neg\alpha \notin \Sigma$	for each atomic sentence α
$(\neg)$	$\neg\neg\tau \in \Sigma \Longrightarrow \Sigma \cup \{\tau\} \in \boldsymbol{Con}$	
$(\wedge)(i)$	$(\rho \wedge \sigma) \in \Sigma \Longrightarrow \Sigma \cup \{\rho, \sigma\} \in \boldsymbol{Con}$	
$(\wedge)(ii)$	$\neg(\rho \vee \sigma) \in \Sigma \Longrightarrow \Sigma \cup \{\neg\rho, \neg\sigma\} \in \boldsymbol{Con}$	
$(\wedge)(iii)$	$\neg(\rho \to \sigma) \in \Sigma \Longrightarrow \Sigma \cup \{\neg\rho, \sigma\} \in \boldsymbol{Con}$	
$(\vee)(i)$	$(\rho \vee \sigma) \in \Sigma \Longrightarrow \Sigma \cup \{\rho\} \in \boldsymbol{Con}$ or $\Sigma \cup \{\sigma\} \in \boldsymbol{Con}$	
$(\vee)(ii)$	$\neg(\rho \wedge \sigma) \in \Sigma \Longrightarrow \Sigma \cup \{\neg\rho\} \in \boldsymbol{Con}$ or $\Sigma \cup \{\neg\sigma\} \in \boldsymbol{Con}$	
$(\vee)(iii)$	$(\rho \to \sigma) \in \Sigma \Longrightarrow \Sigma \cup \{\neg\rho\} \in \boldsymbol{Con}$ or $\Sigma \cup \{\sigma\} \in \boldsymbol{Con}$	
$(\forall)(i)$	$(\forall v)\phi(v) \in \Sigma \Longrightarrow \Sigma \cup \{\phi(a)\} \in \boldsymbol{Con}$	for each $a \in W$
$(\forall)(ii)$	$\neg(\exists v)\phi(v) \in \Sigma \Longrightarrow \Sigma \cup \{\neg\phi(a)\} \in \boldsymbol{Con}$	for each $a \in W$
$(=)(i)$	$(a \doteq b) \in \Sigma \Longrightarrow \Sigma \cup \{(b \doteq a)\} \in \boldsymbol{Con}$	for each $a, b \in W$
$(=)(ii)$	$(a \doteq t), \phi(t) \in \Sigma \Longrightarrow \Sigma \cup \{\phi(a)\} \in \boldsymbol{Con}$	for each $a \in W$ and each close term t
$(W)(t)$	$\Sigma \cup \{(a \doteq t)\} \in \boldsymbol{Con}$	for each closed term t and each $a \in W - w(\Sigma \cup \{t\})$
$(W)(\exists)$	$(\exists v)\phi(v) \in \Sigma \Longrightarrow \Sigma \cup \{\phi(a)\} \in \boldsymbol{Con}$	for each $a \in W - w(\Sigma)$
$(W)(\forall)$	$\neg(\forall v)\phi(v) \in \Sigma \Longrightarrow \Sigma \cup \{\neg\phi(a)\} \in \boldsymbol{Con}$	for each $a \in W - w(\Sigma)$

Table 3: The closure conditions for a consistency property

for some finite $\Gamma \subseteq \Sigma$. We must show that this set has a model.

We know that Γ has a model. This has the form

$$(\mathfrak{A}, \mathsf{a})$$

where $\mathfrak{A}$ is a L-structure and a is an enumeration which interprets each witness. In particular, for the witness a there is some element $[\![a]\!]$ in the a-position.

Suppose we form a second enumeration b by replacing this element $[\![a]\!]$ by some other element. This gives us a second structure $(\mathfrak{A}, \mathsf{b})$. Notice that

$$(\mathfrak{A}, \mathsf{b}) \models \Gamma \qquad (\mathfrak{A}, \mathsf{a})[\![t]\!] = (\mathfrak{A}, \mathsf{b})[\![t]\!]$$

because the element in the a-position is not involved in either calculation. Thus, we form b by replacing the element in the a-position of a by the element $(\mathfrak{A}, \mathsf{a})[\![t]\!]$, and so

$$(\mathfrak{A}, \mathsf{b}) \models (a \doteq t)$$

to exhibit the required model.

The same trick can be used to verify $(W)(\exists)$ and $(W)(\forall)$.

For instance, consider any $(\exists v)\phi(v) \in \Sigma \in \boldsymbol{Con}$ and any witness a not occurring in Σ. Any finite subset of

$$\Sigma \cup \{\phi(a)\}$$

is a subset of

$$\Gamma \cup \{\phi(a)\}$$

for some finite $\Gamma \subseteq \Sigma$ with $(\exists v)\phi(v) \in \Gamma$. We know that Γ has a model $(\mathfrak{A}, \mathsf{a})$, and we modify this to obtain a model $(\mathfrak{A}, \mathsf{b})$ of $\phi(a)$. We do this by altering the element in the a-position of a. Since $(\mathfrak{A}, \mathsf{a}) \models (\exists v)\phi(v)$ there is some element b with $(\mathfrak{A}, \mathsf{a}) \models \phi(b)$. We form b by putting b into the a-position, so that $(\mathfrak{A}, \mathsf{b})[\![a]\!] = b$. This ensures that $(\mathfrak{A}, \mathsf{b})$ is a model of $\Gamma \cup \{\phi(a)\}$, as required. ■

You should compare Table 3 with Table 2. Notice how that conditions in the two tables match. In particular, conditions $(0)(i)$ — $(=)ii)$ of Table 3 ensure that each maximal member of $\boldsymbol{Con}$ will satifies the conditions $(0)(i)$ — $(=)ii)$ of Table 2. But the witnessing conditions are different. The conditions Table 3 give us enough room to manipulate the witnesses into the places we want them, *provided* we start from a position with enough room.

For each set Σ of $L(W)$-sentences, let $w(\Sigma)$ be the set of witnesses occurring in W. Thus $w(\Sigma) \subseteq W$ and could be the whole of W. However, we want to make sure there is enough witesses left in $W - w(\Sigma)$ to do any job that might turn up. The easiest way to achieve this is to impose a cardinality condition in $w(\Sigma)$.

8.12 THEOREM. *Let $\kappa = |L| = |W| = |L(W)|$.*
Let $\boldsymbol{Con}$ be a consistency property.
Let $\Sigma \in \boldsymbol{Con}$ and suppose $\nu = |w(\Sigma)| < \kappa$.
Then $\Sigma \subseteq \Xi \in \boldsymbol{Con}$ for some helpful set Ξ.

Proof. We obtain Ξ by a maximizing argument. A simple use of Zorn's Lemma will achieve $(0)(i)$ — $(=)(ii)$ of Table 2, but may not achieve the witnessing conditions $(W)(t), (W)(\exists)$, and $(W)(\forall)$. To do that we move from Σ to Ξ in a more controlled way. The initial condition $\nu < \kappa$ gives us enough room to organize the witnesses to do the jobs required.

Let

$$(\tau_i \,|\, i < \kappa) \qquad (t_i \,|, i < \kappa)$$

be ordinal indexed families of all $L(W)$-sentences τ and all closed $L(W)$-terms t. There are no restrictions on these enumerations, and can contain repetitions.

We produce an ascending chain

$$(\boldsymbol{\Sigma}) \quad \Sigma = \Sigma_0 \subseteq \cdots \subseteq \Sigma_i \subseteq \cdots \quad (i < \kappa)$$

of members of $\boldsymbol{Con}$ with

$$\nu_i = |w(\Sigma_i)| \leq \max\{\nu, |i|\}$$

for each $i < \kappa$. This condition will give us plenty of room to manoeuvre the witnesses at later stages.

We set

$$\Sigma_0 = \Sigma \qquad \Sigma_\ell = \bigcup\{\Sigma_i \,|\, i < \ell\}$$

for each limit ordinal $\ell < \kappa$. Note that

$$\nu_\ell = |w(\Sigma_\ell)| = \sup\{\nu_i \,|, i < \ell\} \leq \max\{\nu, |\ell|\} < \kappa$$

to preserve the cardinality condition.

At the step $\Sigma_i \mapsto \Sigma_{i+1}$ we deal with the sentence $\tau = \tau_i$ and the term $t = t_i$. We produce two intermediate stages

$$\Sigma_i \subseteq \Sigma_i' \subseteq \Sigma_i'' \subseteq \Sigma_{i+1}$$

each designed to do a particular job. At each stage we use no more than finitely many new witnesses, so that $\nu_{i+1} = \nu_i + n$ for some $n \in \mathbb{N}$. This ensure the cardinality condition is preserved.

We know that $\Sigma_i \in \boldsymbol{Con}$.

For the first stage we look at $\tau = \tau_i$ and set

$$\Sigma_i' = \begin{cases} \Sigma_i \cup \{\tau\} & \text{if this is in } \boldsymbol{Con} \\ \Sigma_i & \text{otherwise} \end{cases}$$

so that $\Sigma_i' \in \boldsymbol{Con}$ and the cardinality restriction is preserved.

We now inspect the shape of τ. If τ is of neither of the shapes

$$(\exists v)\phi(v) \qquad \neg(\forall v)\phi(v)$$

then we set $\Sigma_i'' = \Sigma_i'$ and proceed to the next stage. It τ is of one of these shapes, Then we arrange that $(W)(\exists)$ or $(W)(\forall)$ of Table 2 holds. We have $|w(\Sigma_i')| < \kappa = |W|$ so we may select a new witness $a \in W - w(\Sigma_i')$, and set

$$\Sigma_i'' = \Sigma_i' \cup \{\phi(a)\} \qquad \Sigma_i'' = \Sigma_i' \cup \{\neg\phi(a)\}$$

accordingly. The corresponding consistency property $(W)(\exists)$ or $(W)(\forall)$ of Table 3 ensures that $\Sigma_i'' \in \boldsymbol{Con}$, and, of course, the cardinality restriction is preserved.

Finally, we deal with the closed term $t = t_i$. We have $|w(\Sigma_i'' \cup \{t\})| < \kappa = |W|$ so we may select a new witness $a \in W - w(\Sigma_i'' \cup \{t\})$, and set

$$\Sigma_{i+1} = \Sigma_i'' \cup \{(a \doteq t)\}$$

to complete this step. Notice that the consistency property $(W)(t)$ of of Table 3 ensures that $\Sigma_{i+1} \in \boldsymbol{Con}$.

Consider the chain $\boldsymbol{\Sigma}$ generated in this way. We set

$$\Xi = \bigcup \boldsymbol{\Sigma}$$

so that $\Xi \in \boldsymbol{Con}$. We need to check that Ξ has the closure properties of Table 2.

Each of the conditions $(0)(i)$ — $(=)(ii)$ follows from the corresponding consistency property of Table 3. Let's look at a few of these verifications.

$(\wedge)(i)$ Suppose $\rho \wedge \sigma \in \Xi$ for some sentences ρ, σ. Note that $\rho \wedge \sigma \in \Sigma_i$ for all sufficiently large i, and hence the closure condition $(\wedge)(i)$ (of Table 3) gives

$$\Sigma_i \cup \{\rho, \sigma\} \in \boldsymbol{Con}$$

for all such i. Let $\rho = \tau_r$ and $\sigma = \tau_s$ to locate two particular indexes $r, s < \kappa$. We have arranged that

$$\Sigma_r \cup \{\rho\} \in \boldsymbol{Con} \Longrightarrow \rho \in \Sigma_{r+1} \subseteq \Xi \qquad \Sigma_s \cup \{\sigma\} \in \boldsymbol{Con} \Longrightarrow \sigma \in \Sigma_{s+1} \subseteq \Xi$$

hold. Now consider any sufficiently large $i > r, s$. Then

$$\Sigma_r \cup \Sigma_s \cup \{\rho, \sigma\} \subseteq \Sigma_i \cup \{\rho, \sigma\} \in \boldsymbol{Con}$$

so that

$$\Sigma_r \cup \{\rho\} \in \boldsymbol{Con} \qquad \Sigma_s r \cup \{\sigma\} \in \boldsymbol{Con}$$

and hence $\rho, \sigma \in \Xi$, as required.

$(\vee)(i)$ Suppose $\rho \vee \sigma \in \Xi$ for some sentences ρ, σ. Note that $\rho \vee \sigma \in \Sigma_i$ for all sufficiently large i, and hence the closure condition $(\vee)(i)$ (of Table 3) ensures that one of

$$\Sigma_i \cup \{\rho\} \in \boldsymbol{Con} \qquad \Sigma_i \cup \{\sigma\} \in \boldsymbol{Con}$$

holds for all such i. Let $\rho = \tau_r$ and $\sigma = \tau_s$ to locate two particular indexes $r, s < \kappa$. We have arranged that

$$\Sigma_r \cup \{\rho\} \in \boldsymbol{Con} \Longrightarrow \rho \in \Sigma_{r+1} \subseteq \Xi \qquad \Sigma_s \cup \{\sigma\} \in \boldsymbol{Con} \Longrightarrow \sigma \in \Sigma_{s+1} \subseteq \Xi$$

hold. Now consider any sufficiently large $i > r, s$. Then one of

$$\Sigma_r \cup \{\rho\} \subseteq \Sigma_i \cup \{\rho\} \in \boldsymbol{Con} \qquad \Sigma_s \cup \{\sigma\} \subseteq \Sigma_i \cup \{\sigma\} \in \boldsymbol{Con}$$

holds, so that one of

$$\Sigma_r \cup \{\rho\} \in \boldsymbol{Con} \qquad \Sigma_s \cup \{\sigma\} \in \boldsymbol{Con}$$

holds, and hence we have one of

$$\rho \in \Xi \qquad \sigma \in \Xi$$

as required.

Each of the conditions $(0)(i)$ — $(=)(ii)$ follows by a similar kind or argument. Let's look at the final last one.

$(=)(ii)$ Suppose $(a \doteq t), \phi(t) \in \Xi$ for some $a \in W$, some closed term t, and some formula $\phi(v)$. Note that

$$(a \doteq t), \phi(t) \in \Sigma_i$$

for all sufficiently large i. and hence the closure condition $(=)(ii)$ (of Table 3) ensures that

$$\Sigma_i \cup \{\phi(a)\} \in \boldsymbol{Con}$$

holds for all such i. Let $\phi(a) = \tau_r$ to locate a particular index. Consider any sufficiently large $i > r$. We have

$$\Sigma_r \cup \{\phi(a)\} \subseteq \Sigma_i \cup \{\phi(a)\} \in \boldsymbol{Con}$$

so that

$$\Sigma_r \cup \{\phi(a)\} \in \boldsymbol{Con}$$

and hence, by the construction, we have

$$\phi(a) \in \Sigma_{r+1} \subseteq \Xi$$

as required.

It remains to verify the witnessing conditions $(W)(t)$, $(W)(\exists)$, and $(W)(\forall)$. The arguments for these are similar to the above arguments, but let;s look at two of them.

$(W)(t)$ Consider any close term t, and let $t = t_i$ to locate an index $i < \kappa$. At the step $i \mapsto i+1$ we have arranged that

$$(a \doteq t) \in \Sigma_{i+1} \subseteq \Xi$$

for some $a \in W$, as required.

$(W)(\exists)$. Consider any sentence $\tau = (\exists v)\phi(v) \in \Xi$. Let $\tau = \tau_i$ to locate an index $i < \kappa$. Since $\Sigma_i \cup \{\tau\} \subseteq \Xi$, we have $\Sigma_i \cup \{\tau\} \in \boldsymbol{Con}$, and hence we have arranged that

$$\tau \in \Sigma_i' \qquad \phi(a) \in \Sigma_i'' \subseteq \Xi$$

for some $a \in W$, as required.

This completes the whole proof. ■

How does this prove the refined compactness theorem for a language L? Consider any set Σ is L-sentences which is finitely satisfiable. We may view Σ as a set of $L(W)$-sentences with $w(\Sigma) = \emptyset$. As such we have $\Sigma \in \boldsymbol{FinSat}$ where this is the consistency property of Theorem 8.11. Furthermore, there are no witnesses in Σ, so that $|w(\Sigma)| = 0 < |L|$, and hence Theorem 8.12 gives $\Sigma \subseteq \Xi \in \boldsymbol{FinSat}$ for some helpful set Ξ. By Theorem 8.8, Ξ has a canonical model and this, of course, provides a model of Σ. The construction of this structure ensures that it has cardinality no bigger than $|W| = |L|$.

The compleness theorem can be proved in the same way. We use

$$\Sigma \in \boldsymbol{ProofCon} \Longleftrightarrow \text{not}(\Sigma \dot{\vdash} \mathsf{false})$$

to produce from a formal proof system $\dot{\vdash}$ a family $\boldsymbol{ProofCon}$ of sets Σ of $L(W)$-sentences. The formal proof rules are designed to ensure that $\boldsymbol{ProofCon}$ is a consistencey property. A simple argument then gives the adequacy of the system.

Exercises

8.2 Complete the proof of Theorem 8.12

8.3 Use the notion of a consistency property to design a system of formal proofs which is complete.

9 Omitting types

Consider any set Γ of formulas (in some language L). For each $\phi \in \Gamma$ the set $\partial\phi$ is finite, but the set

$$\partial\Gamma = \bigcup\{\partial\phi \mid \phi \in \Gamma\}$$

may be infinite (for different formulas may use different variables).

9.1 DEFINITION. let L be a language.

(a) A type (in L) is a set Γ of formulas such that $\partial\Gamma$ is finite. Sometimes we write $\Gamma(v)$ to indicate that $\partial\Gamma = v$.

(b) A type Γ is realized in a structure $\mathfrak{A}$ (suitable for L) if $\mathfrak{A} \models \Gamma(a)$ for some point a of $\mathfrak{A}$ (matching the free variables of Γ).

(c) A type is omitted in a structure if it is not realized in that structure. ■

Compactness arguments can be used to produce a structure which realizes a given type, but how can we produce a structure which omits a type? We have seen already an example of such a construction. By Theorem 5.29, a submodel $\mathfrak{A}$ of a theory T is in $\mathcal{E}(T)$ if an only if $\mathfrak{A}$ omits a certain family of types. However, the construction of $\mathfrak{A}$, as in Lemma 5.31 and Theorem 5.32, is rather crude, and doesn't refine in a straight forward manner. In general, to omit types we need some restrictions on the situation.

9.2 DEFINITION. Let T be a theory.

(ω) Let Γ be a type. The type Γ is principal over T if there is a formula θ (with $\partial\theta = \partial\Gamma$) which is consistent with T and such that

$$T \vdash \theta \to \bigwedge\Gamma$$

holds.

(0) Let Π be a $\forall_1$-type, that is type of $\forall_1$-formulas. The type Π is $\exists_1$-principal over T if there is an $\exists_1$-formula θ (with $\partial\theta = \partial\Pi$) which is consistent with T and such that

$$T \vdash \theta \to \bigwedge\Pi$$

holds. ■

In the 0-version we put a restriction on the quantifier complexity of both the type and the generating formula. If this generating formula is realized in any model of the theory, then we automatically realize the type. However, that does not help us to omit the type. In fact, to ensure that a type can be omitted it should *not* be principal.

9.3 EXAMPLE. Let T be a theory. Recall that for each $\forall_1$-formula ϕ we set

$$\Omega(T, \phi) = \{\phi\} \cup \neg\exists_1(T, \phi)$$

where $\exists_1(T, \phi)$ is the set of $\exists_1$-formulas θ such that

$$T \cup \{\theta\} \text{ is consistent} \qquad T \vdash \theta \to \phi$$

hold. Thus $\Omega(T, \phi)$ is an $\forall_1$-type. Almost by construction, this type is not $\exists_1$-principal over T (for think what an $\exists_1$-generator could be).

The proof of the following result is rather involved. However, the technique used is important, so we will take it rather slowly.

9.4 THEOREM. *Let T be a theory in a countable language L. Let $\mathbf{\Pi}$ be countable collection of $\forall_1$-types each of which is not $\exists_1$-principal over T. Then there is a countable structure $\mathfrak{A} \in \mathcal{S}(T)$ which omits each type in $\mathbf{\Pi}$.*

Proof. We enrich the language L to a language $L(\mathsf{a})$ be adding countably many parameters a. We will produce a structure $(\mathfrak{A}, \mathsf{a})$ for $L(\mathsf{a})$ such that a enumerates the whole of $\mathfrak{A}$, and this is the required structure.

To obtain $(\mathfrak{A}, \mathsf{a})$ we first produce a set Ξ of $\exists_1$-sentence of $L(\mathsf{a})$ such that the following hold.

(i) $T \cup \Xi$ is consistent.

(ii) For each $\exists_1$-sentence τ of $L(\mathsf{a})$, if $T \cup \Xi \cup \{\tau\}$ is consistent then $\tau \in \Xi$.

(iii) For each quantifier-free formula $\delta(v)$ of $L(\mathsf{a})$, if $(\exists v)\delta(v) \in \Xi$ then $\delta(a) \in \Xi$ for some selection a of parameters (matching v).

(iv) For each $\Pi(v) \in \mathbf{\Pi}$ and selection a of parameters (matching v), there is some formula $\phi(v) \in \Pi(v)$ with $\neg\phi(a) \in \Xi$.

Suppose we have found such a set Ξ. By (i) there is a model $(\mathfrak{B}, \mathsf{a})$ of $T \cup \Xi$. Let $\mathfrak{A}$ be the substructure of $\mathfrak{B}$ generated by a. Thus $(\mathfrak{A}, \mathsf{a}) \subseteq (\mathfrak{B}, \mathsf{a})$. As yet, we do not know that a enumerates the whole of $\mathfrak{A}$, but we will prove this shortly.

Consider any $\tau \in \Xi$. Then

$$\tau = (\exists v)\delta(v)$$

for some quantifier-free formula $\delta(v)$ of $L(\mathsf{a})$. This δ may contain parameters not shown explicitly. By (iii) we have $\delta(a) \in \Xi$ for some selection a of parameters. This gives $(\mathfrak{B}, \mathsf{a}) \models \delta(a)$, and hence $(\mathfrak{A}, \mathsf{a}) \models \delta(a)$, so that $(\mathfrak{A}, \mathsf{a}) \models \tau$.

Using this argument we draw a couple of conclusions.

Firstly it shows that $(\mathfrak{A}, \mathsf{a}) \models \Xi$.

Now consider any $\exists_1$-sentence τ of $L(\mathsf{a})$ where $(\mathfrak{B}, \mathsf{a}) \models \tau$. Then $T \cup \Xi \cup \{\tau\}$ is consistent, so that (ii) gives $\tau \in \Xi$, and hence $(\mathfrak{A}, \mathsf{a}) \models \tau$. This shows that

$$\mathfrak{B} \models \theta(a) \Longrightarrow \mathfrak{A} \models \theta(a)$$

for each $\exists_1$-formula $\theta(v)$ of L and each point a from a.

Next we see that a enumerates the whole of $\mathfrak{A}$. For consider a sentence

$$\tau := (\exists v)[v = Oa_1, \dots, a_n]$$

where O is an n-placed operation symbol and $a_1, \dots, a_n$ are parameters. This $\exists_1$-sentence of $L(\mathsf{a})$ is true in $\mathfrak{B}$, and hence $\tau \in \Xi$ (as in the previous paragraph). But now (iii) give some elemement $a \in \mathsf{a}$ with

$$(a = Oa_1, \dots, a_n) \in \Xi$$

to show that a is closed under O.

From this, we see that

$$(\mathfrak{A}, \mathsf{a}) \prec_1 (\mathfrak{B}, \mathsf{a})$$

holds.

Finally, consider any type $\Pi(v) \in \mathbf{\Pi}$ and any selection a of parameters matching v. By (iv) we have $\neg\phi(a) \in \Xi$ for some $\phi(v) \in \Pi(v)$. But now $\mathfrak{A} \models \phi(a)$, and hence $\mathfrak{A}$ does not realize $\Pi(v)$ at a. Since a enumerates the whole of $\mathfrak{A}$, this shows that $\mathfrak{A}$ omits $\Pi(v)$.

Our job now is to produce such a set Ξ of sentence of $L(\mathsf{a})$. This is where the assumed countability is important.

Consider a triple

$$(\tau, \Pi(v), a)$$

where τ is an $\exists_1$-sentence of $L(\mathsf{a})$, where $\Pi(v) \in \mathbf{\Pi}$, and where a is a selection of parameters matching v. The language $L(\mathsf{a})$ is countable, and the collection $\mathbf{\Pi}$ is countable, so there are just countably many such triples. Let

$$(\tau, \Pi(v), a)_i \quad i < \omega$$

be an enumeration of all such triples.

We produce Ξ as the union of an ascending chain

$$\Xi_0 \subseteq \Xi_1 \subseteq \cdots \subseteq \Xi_i \subseteq \ldots \qquad i < \omega$$

of finite sets of $\exists_1$-sentences of $L(\mathsf{a})$. We generate this ω-chain by recursion on i. We ensure that each $T \cup \Xi_i$ is consistent, so that (i) will hold. The other requirements (ii,iii,iv) are dealt with during the course of the construction. The step $\Xi_i \mapsto \Xi_{i+1}$ deals with the triple $(\tau, \Pi(v), a)$ with index i.

At the base we set $\Xi_0 = \emptyset$.

Suppose we have obtain the finite set Ξ_i which is consistent with T. We look at the i-triple

$$(\tau, \Pi(v), a)$$

enumerated above. Using this we do several things to obtain Ξ_{i+1}.

(v) If $T \cup \Xi_i \cup \{\tau\}$ is consistent we put

$$\Xi_i' = \Xi_i \cup \{\tau\}$$

otherwise we put $\Xi_i' = \Xi_i$.

(vi) We have a finite set Ξ_i' of $\exists_1$-sentences of $L(\mathsf{a})$ which is consistent with T. Since $\Pi(v)$ is not $\exists_1$-principal over T we know that

$$T \cup \Xi_i' \vdash \bigwedge \Pi(a)$$

can not hold. (If it does hold, then $\bigwedge \Xi_i'$ will produce a generator of $\Pi(v)$ over T.) Thus there is some $\phi(v) \in \Pi(v)$ such that

$$T \cup \Xi_i' \cup \{\neg\phi(a)\}$$

is consistent. We put

$$\Xi_i'' = \Xi_i' \cup \{\neg\phi(a)\}$$

to obtain a further extension of Ξ_i'.

(vii) To get from Ξ_i to Ξ''_i we may have put in one or two new $\exists_1$-sentence of $L(\mathsf{a})$. Each of these has the form

$$(\exists v)\delta(v)$$

where $\delta(v)$ is a quantifier-free formula of $L(\mathsf{a})$ (where the occurring parameters have not been displayed). In total there are only finitely many parameters occurring in Ξ''_i, so there are plenty that have not yet been used. We may select a point a of new parameters matching v (the free variables of this δ). With this, $T \cup \Xi''_i \cup \{\delta(a)\}$ is consistent. We witness all the new $\exists_1$-sentences in the same way to obtain Ξ_{i+1}.

Once this construction has be carried out we set

$$\Xi = \bigcup\{\Xi_i \mid i < \omega\}$$

to obtain the final set Ξ.

Why does this set Ξ satisfy (i,ii,iii,iv)?

Property (i) holds since at each stage $T \cup \Xi_i$ is consistent.

Property (ii) holds since this sentence τ is consider at some stage i. If $T \cup \Xi \cup \{\tau\}$ is consistent, then so is $T \cup \Xi_i \cup \{\tau\}$, and then (v) ensures that $\tau \in \Xi_{i+1} \subseteq \Xi$.

Property (iii) holds by the final part of the step $\Xi_i \mapsto \Xi_{i+1}$ described in (vii).

Property (iv) holds by the arrangement described in (vi). The type $\Pi(v)$ and the selection a is part of a triple $(\tau, \Pi(v), a)$ considered at some stage i. The small step $\Xi'_i \mapsto \Xi''_i$ ensures that $\neg\phi(a) \in \Xi_{i+1} \subseteq \Xi$ for some $\phi(v) \in \Xi$.

With this the proof is completed. ■

This result is the 0-version of the omitting types result. There is an analogous, and more common, ω-version.

9.5 THEOREM. *Let T be a theory in a countable language L. Let $\mathbf{\Pi}$ be countable collection of types each of which is not principal over T. Then there is a countable model $\mathfrak{A}$ of T which omits each type in $\mathbf{\Pi}$.*

This is proved by almost the same proof as Theorem 9.4. The difference is there is no quantifier complexity on the set Ξ is sentences of the enriched language $L(\mathsf{a})$. From this we obtain a model $(\mathfrak{B}, \mathsf{a})$ of $T \cup \Xi$, and a generated substructure $(\mathfrak{A}, \mathsf{a})$ of $(\mathfrak{B}, \mathsf{a})$. However, we have to show that $\mathfrak{A} \prec \mathfrak{B}$. Conditions (i, ii, ii, iv) and the generating conditions (v, vi, vii) are modified by removing all the restrictions on quantifier complexity.

Exercises

9.1 Theorem 9.4 can be strengthened. Prove the following.

Let T be a theory in a countable language L. Let $\mathbf{\Pi}$ be countable collection of $\forall_1$-types each of which is not $\exists_1$-principal over T. Then there s a countable structure $\mathfrak{A} \in \mathcal{E}(T)$ which omits each type in $\mathbf{\Pi}$.

In other words, show that we can arrange that the omitting structure is e. c.

9.2 Prove Theorem 9.5.

10 The back and forth technique

Recall that, from section 2.3, a line is a dense linear ordering without end points. The rationals $\mathbb{Q}$ and the reals $\mathbb{R}$ are the two best known examples of lines. However, since the class of lines is elementary and its theory is $\forall_2$-axiomatizable (if fact, it has EQ), we see there are examples of lines in every (non-finite) cardinality. How can we compare these different lines.

Consider any two lines A and B , and let's compare configurations in A with those in B. Thus, let

$$a = a_1, \dots, a_n \qquad b = b_1, \dots, b_n$$

be, respectively, a point of A and a point of B. (Here, of course, we mean 'point' in the sense used throughout these notes, not in the sense of a geometrical point.) The elements of a and b need not be listed in the order that they occur in their respective lines, but we do want them to have the same configuration. Let us write

$$a \sim b$$

to indicate this similarity. Thus, in detail, this means that

$$a_i \leq a_j \iff b_i \leq b_j$$

holds for all $i, j \in \{1, \dots, n\}$.

Suppose now that we are given a new element x of A. We may attach x to the end of a to form an extended point ax. This will have a certain configuration in A. Can we find a element y of B such that

$$ax \sim by$$

holds?

We can, because both A and B are lines.

If $x = a_i$ for some i, then we simply take $y = b_i$. Thus we may suppose that $x \neq a_i$ for each i. We divide the a_i into two part, a left hand part $L(A)$ and a right hand part $R(A)$. Using x we let

$$a_i \in L(A) \iff a_i < x \qquad a_i < x \iff a_i \in R(A)$$

to generate these part. Notice that

$$a_l < x < a_r$$

holds for all $a_l \in L(A)$ and $a_r \in R(A)$. We pass these parts across to B. Thus

$$b_i \in L(B) \iff a_i \in L(A) \qquad a_i \in R(A) \iff b_i \in R(B)$$

gives us two finite subsets $L(B), R(B)$ of B. The similarity $a \sim b$ ensures that

$$b_l < b_r$$

holds for all $b_l \in L(B), b_r \in R(B)$. Now we make the crucial observation. Because B is a line, there is some $y \in B$ such that

$$b_l < y < b_r$$

holds for all $b_l \in L(B)$ and $b_r \in R(B)$. To see this simply consider the three possible cases; $L(B) = \emptyset, R(B) = \emptyset$, both $L(B)$ and $R(B)$ non-empty. Remember that B has no left hand end, has no right hand end, and is dense. This choice of y ensures that

$$ax \sim by$$

holds.

We want to extend this kind of argument to more general structures. To do that we need a notion of 'points in the same configuration'.

10.1 DEFINITION. Let $\mathfrak{A}, \mathfrak{B}$ be structures for some language L. A partial isomorphism between $\mathfrak{A}$ and $\mathfrak{B}$ is a pair (a, b) of points (a from $\mathfrak{A}$ and b from $\mathfrak{B}$) such that

$$(\mathfrak{A}, a) \equiv_0 (\mathfrak{B}, b)$$

holds. ■

For the case of lines a partial isomorphisms (a, b) is merely a pair of points in the same configuration. The description of this used above sets down precisely what $(A, a) \equiv_0 (B, b)$ means in this particlar case. In the same way we could set down the details of what $(\mathfrak{A}, a) \equiv_0 (\mathfrak{B}, b)$ means for the general case. We don't do that because these details are never needed here. Notice that we do not assume that the points a and b enumerate substructures of $\mathfrak{A}$ and $\mathfrak{B}$. However, the two points will generate substructures, and these substructures will be canonically isomorphic.

10.2 DEFINITION. Let $\mathfrak{A}, \mathfrak{B}$ be structures for some language L.

(a) A back-and-forth system for the pair $\mathfrak{A}, \mathfrak{B}$ is a non-empty set P of partial isomorphisms (a, b) with the following two properties.

(back) For each $(a, b) \in P$ and each element y of $\mathfrak{B}$, there is an element x of $\mathfrak{A}$ with $(ax, by) \in P$.

(forth) For each $(a, b) \in P$ and each element x of $\mathfrak{A}$, there is an element y of $\mathfrak{B}$ with $(ax, by) \in P$.

(b) We write

$$\mathfrak{A} \equiv_p \mathfrak{B}$$

and say $\mathfrak{A}$ and $\mathfrak{B}$ are potentially isomorphic if there is at least one back-and-forth system for the pair $\mathfrak{A}, \mathfrak{B}$. ■

For instance, we saw above that each pair of lines are potentially isomorphic. In particular, potential isomorphisms does not require that the two structures have the same cardinality.

The relation $\equiv_p$ should be compared with the semantic relations

$$\equiv_n \qquad \equiv_\omega$$

(for $n < \omega$) and the isomorphism relation

$$\equiv_\infty$$

defined in section 3.1. We have used a similar notation for these relations because of the following result.

10.3 THEOREM. *For each pair $\mathfrak{A}, \mathfrak{B}$ of structures (for some language), both the implications*

$$\mathfrak{A} \equiv_\infty \mathfrak{B} \Longrightarrow \mathfrak{A} \equiv_p \mathfrak{B} \Longrightarrow \mathfrak{A} \equiv_\omega \mathfrak{B}$$

hold.

Proof. For the first implication suppose f is an isomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. Let P be the set of finite restrictions of f. That is, P is the set of pair (a, b) where

$$a = a_1, \dots, a_n \qquad b = b_1 \dots, b_n$$

and where

$$fa_i = b_i$$

for each $i \le i \le n$. A simple calculation shows that this is a back-and-forth system.

For the second implication suppose we have a back-and-forth system P for $\mathfrak{A}, \mathfrak{B}$. We show that

$$[n] \quad (\mathfrak{A}, a) \equiv_n (\mathfrak{B}, n)$$

holds for each $n < \omega$ and each $(a, b) \in P$. We proceed by induction on n.

The base case, [0] is trivial.

For the induction step, $[n] \Rightarrow [n+1]$, suppose

$$\mathfrak{A} \models (\exists v_1, \dots, v_m)\phi(v_m, \dots, v_1, a)$$

where $(a, b) \in P$ and $\phi \in \forall_n$. There are elements $x_1, \dots, x_m$ of $\mathfrak{A}$ such that

$$\mathfrak{A} \models \phi(x_m, \dots, x_1, a)$$

holds. By repeated use of the 'forth' property there are elements $y_1, \dots, y_m$ of $\mathfrak{B}$ such that

$$(ax_1 \cdots x_m, by_1 \cdots y_m) \in P$$

holds, and hence

$$(\mathfrak{A}, ax_1 \cdots x_m) \equiv_n (\mathfrak{B}, by_1 \cdots y_m)$$

by the induction hypothesis $[n]$. This gives

$$\mathfrak{B} \models \phi(y_m, \dots, y_1, ba)$$

and hence

$$\mathfrak{B} \models (\exists v_1, \dots, v_m)\phi(v_m, \dots, v_1, b)$$

holds.

This, with a symmetric argument, gives $[n+1]$. ■

As a simple consequence of this we see that the theory of lines is complete. For any two lines A, B we have $A \equiv_p B$ and hence $A \equiv_\omega B$, in other words A and B are elementary equivalent.

Originally, back-and-forth systems were invented to generate full ismorphisms. Here is that first use.

10.4 THEOREM. *The implication*

$$\mathfrak{A} \equiv_p \mathfrak{B} \Longrightarrow \mathfrak{A} \equiv_\infty \mathfrak{B}$$

holds for all countable structures $\mathfrak{A}, \mathfrak{B}$.

Proof. Let P be a back-and-forth system for the two countable structures $\mathfrak{A}$ and $\mathfrak{B}$. We use P to generate a full isomorphisms f from $\mathfrak{A}$ to $\mathfrak{B}$.

Let A be the carrier of $\mathfrak{A}$ and let B be the carrier of $\mathfrak{B}$. Since both A and B are countable there are enumerations

$$x_i \quad y_i \qquad i < \omega$$

of these sets. (These enumerations can be quite arbitrary.) We use P to produce an ascending chain

$$f_0 \subseteq f_1 \subseteq \cdots \subseteq f_i \subseteq \qquad i < \omega$$

of partial isomorphism

$$f_i : A_i \longrightarrow B_i$$

between finite subsets of A and B. Thus, for each $i < \omega$ we have

$$(\mathfrak{A}, a_1, \dots, a_n) \equiv_0 (\mathfrak{B}, b_0, \dots, b_n)$$

where

$$a_1, \dots, a_n$$

is an enumeration of A_i, and

$$b_1 = f_i a_1, \ \dots \ , b_n = f_i a_n$$

is the corresponding enumeration of B_i. For each $i < \omega$ this constructed pair (a, b) of enumerations will be in P and we will ensure that

$$x_i \in A_i \qquad y_i \in B_i$$

hold. The union

$$f = \bigcup \{ f_i \mid i < \omega \}$$

is then the required isomorphism.

Suppose we have produced the partial isomorphism f_i with (a, b) as the corresponding pair. We look first at the element $x = x_{i+1}$ of A. By the forth property there is some element v of B with $(ax, bv) \in P$. We next look at the element $y = y_{i+1}$ of B. By the back property there is some element u of A with $(axu, bvy) \in P$. We take this as the extension f_{i+1} of f_i. ■

As a simple consquence of this we see that the theory of lines is $\aleph_0$-categorical.

Homogeneous structures

Exercises

10.1 Show that potential isomorphism, $\equiv_p$, is an equivalence on the class of all structures (for a given language).

11 Homogeneous-universal models

12 Saturation – *see earlier*

13 Forcing techniques

14 Ultraproducts–*to be done*

Part III

Some solutions to the exercises

[*Held in 120../B10-sols.. Last changed July 26, 2004*]

A For section 1

[*Held in 120../B11-sols.. Last changed July 26, 2004*]

A.1 For §1.1

1.1 The shape of the formulas can be seen by subscripting the different uses of the variables in ϕ_3.

$$\begin{array}{lcl} \phi_3 & := & (\exists u_2)[(u_2 < v_3) \wedge \\ & & \qquad (\exists v_2)[(u_2 \doteq v_2) \ \wedge \phi_2]] \text{ where} \\ \phi_2 & := & \qquad\qquad (\exists u_1)[(u_1 < v_2) \wedge \\ & & \qquad\qquad\qquad (\exists v_1)[(u_1 \doteq v_1) \ \wedge \phi_1]] \text{ where} \\ \phi_1 & := & \qquad\qquad\qquad\qquad (\exists u_0)[(u_0 < v_1) \wedge \\ & & \qquad\qquad\qquad\qquad\qquad (\exists v_0)[(u_0 \doteq v_0) \ \wedge \phi_0]] \text{ where} \\ \phi_0 & := & \qquad\qquad\qquad\qquad\qquad\qquad (v_0 \doteq v_0) \end{array}$$

Continuing this sequence in the obvious way, a simple induction gives $\partial\phi_r = \{v_r\}$ and $\partial\theta_r = \{u_r\}$ for each $r < \omega$. ■

A.2 For §1.2–*not yet done*

[*Held in 120../B13-sols.. Last changed July 26, 2004*]

A.3 For §1.3

1.4 (a) A simple induction shows that

$$\mathfrak{N} \models \phi_r(a) \Longleftrightarrow r \leq a$$

holds for each $a \in \mathfrak{N}$ and $r < \omega$.

(b) This is an immediate consequence of (a).

(c) Consider any such formula $\psi(v)$. If there is any $a \in \mathfrak{N}$ with $\mathfrak{N} \models \psi(a)$ then, by (a), we have $r \leq a$ for each $r < \omega$. This is impossible, and hence there is no such a. Thus $\mathfrak{N} \models \neg(\exists v)\psi(v)$ must hold. ■

[*Held in 120../B14-sols.. Last changed July 26, 2004*]

A.4 For §1.4 – *most not done*

1.13 (i)⇒(ii). Suppose $T \cup \{\sigma\}$ is consistent and consider any $\mathfrak{A} \models T \cup \{\sigma\}$. Then, using i, we have

$$\mathfrak{B} \models T \Longrightarrow \mathfrak{B} \equiv \mathfrak{A} \Longrightarrow \mathfrak{B} \models \sigma$$

so that $\sigma \in T$.

(ii)$\Rightarrow$(iii). Suppose $T \subseteq T'$ where T' is consistent. Then (ii) gives

$$\sigma \in T' \Longrightarrow T \cup \{\sigma\} \text{ consistent} \Longrightarrow \sigma \in T$$

and hence $T = T'$.

(iii)$\Rightarrow$(iv). We verify the contrapositive of (iv). Consider a pair of sentences σ, τ neither of which is in T. Then both

$$T \cup \{\neg\sigma\} \qquad T \cup \{\neg\tau\}$$

are consistent, and hence (iii) (or even (ii)) gives

$$T \vdash \neg\sigma \qquad T \vdash \neg\tau$$

so that $\neg\sigma \wedge \neg\tau \in T$, which leads to $\sigma \vee \tau \notin T$.

(iv)$\Rightarrow$(i). For each sentence σ, we have $\sigma \vee \neg\sigma \in T$, and hence (iv) ensures that one of

$$\sigma \in T \qquad \neg\sigma \in T$$

holds, which leads to the required result. ■

1.14 Let $\boldsymbol{\Psi}$ be the full binary splitting tree (sometimes called the Cantor tree). The nodes $\mu, \nu, \ldots$ of $\boldsymbol{\Psi}$ are the finite lists

$$i(0)i(1)\cdots i(n-1)$$

taken from $\{0,1\}$. This particular list has length n, and the empty list

$$\perp$$

of length 0 is allowed. These nodes are partially ordered by extension.

$$\dfrac{\dfrac{\dfrac{\overset{\vdots}{000}\quad \overset{\vdots}{001}}{00} \qquad \dfrac{\overset{\vdots}{010}\quad \overset{\vdots}{011}}{01}}{0} \qquad \dfrac{\dfrac{\overset{\vdots}{100}\quad \overset{\vdots}{101}}{10} \qquad \dfrac{\overset{\vdots}{110}\quad \overset{\vdots}{111}}{11}}{1}}{\perp}$$

In particular

$$\dfrac{\nu 0 \quad \nu 1}{\nu}$$

is a typical node ν and its two successors $\nu 0, \nu 1$.

We write

$$\nu \leq \mu \qquad \mu | \nu$$

to indicate, respectively, that μ extends ν, and that μ and ν are incomparable.

A branch of this tree is just a function

$$p : \mathbb{N} \longrightarrow \{0,1\}$$

and there are $2^{\aleph_0}$ such branches. This branch lies above the nodes

$$\perp,\ p(0),\ p(0)p(1),\ p(0)p(1)p(2),\ \ldots$$

the finite initial sections of p.

Assuming that no finite extension of the theory T is complete, we produce a family

$$\Sigma = \{\sigma_\nu \mid \nu \in \Psi\perp\}$$

of sentences, each consistent with T, with

$$\sigma_\perp = \mathsf{true}$$

and such that

$$(i)\quad \nu \le \mu \Longrightarrow T \vdash \sigma_\mu \to \sigma_\nu \qquad (ii)\quad \mu|\nu \Longrightarrow T \vdash \neg\sigma_\mu \vee \neg\sigma_\nu$$

hold for all nodes μ, ν of $\mathbf{\Psi}$.

We generate this family by recursion up the tree $\mathbf{\Psi}$. Suppose we have generated σ_ν for some $\nu \in \mathbf{\Psi}$. Thus $T \cup \{\sigma_\nu\}$ is consistent, and does not axiomatize a complete extension of T. Thus there is a sentences τ such that both

$$T \cup \{\sigma_\nu, \tau\} \qquad T \cup \{\sigma_\nu, \neg\tau\}$$

are consistent. We set

$$\sigma_{\nu 0} = \sigma_\nu \wedge \tau \qquad \sigma_{\nu 1} = \sigma_\nu \wedge \neg\tau$$

to produce the two immediate successors of σ_ν. Note that

$$T \vdash \sigma_{\nu i} \to \sigma_\nu \qquad T \vdash \neg\sigma_{\nu 0} \vee \neg\sigma_{\nu 1}$$

hold.

Once we have the family Σ we set

$$\Sigma_p = \{\sigma_\nu \mid \nu < p\}$$

for each branch p of $\mathbf{\Psi}$. By (i) we see that each such theory $T \cup \Sigma_p$ is consistent. Let

$$\mathfrak{A}_p \models T \cup \Sigma_p$$

to produce a complete extension $T(p) = Th(\mathfrak{A}_p)$ of T. We show that these are pairwise distinct.

Consider distinct branches p, q of $\mathbf{\Psi}$. There are nodes $\mu < p, \nu < q$ with $\mu|\nu$. Then (ii) gives

$$\mathfrak{A}_p \models \neg\sigma_\nu \qquad \mathfrak{A}_q \models \neg\sigma_\mu$$

so that $\mathfrak{A}_\mu \not\equiv \mathfrak{A}_q$, as required. ■

[*Held in 120../B15-sols.. Last changed July 26, 2004*]

A.5 For §1.5

1.16 (a) If $\Sigma \vdash \tau$, then $\Sigma \cup \{\neg\tau\}$ does not have a model, and hence, by compactness, some finite part of this set does not have a model. In other words, there is some finite $\Gamma \subseteq \Sigma$ such that $\Gamma \cup \{\neg\tau\}$ does not have a model. With this Γ we have $\Gamma \vdash \tau$, to give the required result.

(b) Suppose Σ has the closure properties. We show that

$$\Sigma \vdash \tau \Longrightarrow \tau \in \Sigma$$

holds, and hence Σ is a theory. (The other required implication is trivial.)

Suppose $\Sigma \vdash \tau$. By part (a) we have $\Gamma \vdash \tau$ for some finite part

$$\Gamma = \{\sigma_m, \dots, \sigma_1\}$$

of Σ. But now the sentence

$$\sigma_m \to \cdots \sigma_1 \to \tau$$

is logically valid, and hence belongs to Σ. With this we may strip of $\sigma_m, \dots, \sigma_1$ in turn, to get $\tau \in \Sigma$. ■

1.17 Since each $\mathcal{K}_r$ is strictly elementary we have $\mathcal{K}_r = \mathcal{M}(\sigma_r)$ for some family

$$\Sigma = \{\sigma_r \mid r < \omega\}$$

of sentences. Since $\mathcal{K}_r \supseteq \mathcal{K}_{r+1}$ we may assume that $\vdash \sigma_{r+1} \to \sigma_r$ holds. (This may be achieved by taking conjunctions of the original σ_r.)

We see that $\mathcal{K} = \mathcal{M}(\Sigma)$, so that $\mathcal{K}$ is elementary.

By way of contradiction, suppose that $\mathcal{K}$ is strictly elementary. Thus $\mathcal{K} = \mathcal{M}(\tau)$ for some sentence τ. In particular

$$\Sigma \vdash \tau$$

and

$$\vdash \tau \to \sigma_r$$

for each $r < \omega$. These give some particular $n < \omega$ such that $\vdash \tau \leftrightarrow \sigma_n$, and hence $\mathcal{K} = \mathcal{M}(\sigma_n) = \mathcal{K}_n$, which is not so. ■

1.18 Let

$$\mathcal{K} = \mathcal{M}(\sigma) \qquad \mathcal{L} = \mathcal{M}(L) \qquad \mathcal{R} = \mathcal{M}(R)$$

where σ is a sentence and L, R are sets of sentences. Since $\mathcal{L} \cap \mathcal{R} = \emptyset$, we see that $L \cup R$ does not have a model, and hence compactness gives sentences λ, ρ such that

$$\mathcal{L} \subseteq \mathcal{M}(\lambda \wedge \sigma) \qquad \mathcal{R} \subseteq \mathcal{M}(\rho \wedge \sigma) \qquad \vdash \neg\lambda \vee \neg\rho$$

hold. We show that, in fact,

$$\mathcal{L} = \mathcal{M}(\lambda \wedge \sigma) \qquad \mathcal{R} = \mathcal{M}(\rho \wedge \sigma)$$

hold.

Consider any $\mathfrak{A} \in \mathcal{M}(\lambda \wedge \sigma)$. Since $\mathfrak{A} \models \sigma$, we have $\mathfrak{A} \in \mathcal{K} = \mathcal{L} \cup \mathcal{R}$. Since $\mathfrak{A} \models \lambda$, we have $\mathfrak{A} \models \neg\rho$, and hence $\mathfrak{A} \notin \mathfrak{R}$, to give $\mathfrak{A} \in \mathcal{L}$, as required. ■

1.19 (a,b) Let α be conjunction of th eusual axioms for fields. This shows that $\mathcal{F}$ is strictly elementary. and axiomatizes T.

For each prime p let

$$(p \doteq 0) \quad \text{abbreviate} \quad 1 + \cdots + 1 \doteq 0$$

where there are p number of 1s in this sum. Then

$$\alpha \wedge (p \doteq 0)$$

axiomatizes T_p, to show that $\mathcal{F}_p$ is strictly elementary. Also

$$\{\alpha\} \cup \{(p \neq 0) \mid p \text{ a prime}\}$$

axiomatizes T_0, to show that $\mathcal{F}_0$ is elementary.

We have

$$\mathcal{F}_0 \cup \mathcal{F}_f = \mathcal{F} \qquad \mathcal{F}_0 \cap \mathcal{F}_f = \emptyset$$

so that, by Exercise 1.18, if $\mathcal{F}_f$ is elementary then $\mathcal{F}_0$ is strictly elementary. Part (c) show that this is not the case.

(c) For each $\tau \in T_0$ we have

$$T \cup \{(p \neq 0) \mid p \text{ a prime}\} \vdash \tau$$

and hence

$$T \cup \{(2 \neq 0), \ldots, (p \neq 0)\} \vdash \tau$$

for some sufficiently large prime p. Any field of larger prime characteristic show that this can not be so.

In particular, T_0 is not finitely axiomatizable, and $\mathcal{F}_0$ is not strictly elementary.

(d) The same method show that

$$T_f \cup \{(p \neq 0) \mid p \text{ a prime}\}$$

is finitely satisfiable, and hence has a model. ■

1.20 For the first part all we need to observe is that $Th(\mathfrak{R})$ can be formalised in a countable language.

The second part depends on which signature is used.

If we view $\mathfrak{R}$ as a line $(\mathbb{R}, \leq)$ (in the sense of subsection 2.3) then the rational line $(\mathbb{Q}, \leq)$ is the only possibility.

If we view $\mathfrak{R}$ in the most natural way as a linearly ordered real field, then the real part of the algebraic closure of $\mathbb{Q}$ will do (and there are other possibilities).

This second part requires some results not yet covered and some information about the structure of real fields. ■

[*Held in 120../B20-sols.. Last changed July 26, 2004*]

B For section 2

[*Held in 120../B21-sols.. Last changed July 26, 2004*]

B.1 For §2.1

2.1 Since there are no operation symbols in the signature, each atomic sentence has one of the shapes

$$\mathsf{true} \quad \mathsf{false} \quad (K_1 \doteq K_2) \quad RK_1, \cdots K_n$$

where $K_1, \dots, K_n$ are constant symbols and R is an n-placed relation symbol. Since the signature is finite, there are only finitely many such atomic sentences. Let $\alpha_1, \dots, \alpha_n$ be these atomic sentences.

Each quantifier-free sentence is T-equivalent to a conjunction of disjunction of literals. Each such disjunction can be put in the form

$$\pm\alpha_1 \vee \cdots \vee \pm\alpha_n$$

where each component $\pm\alpha$ is α or $\neg\alpha$, as approriate. There are only 2^n possible disjunctions of this kind. Each quantifier-free sentence is T-equivalent to a conjunction of such disjunctions, and hence there are only finitely many quantifier-free sentences up to T-equivalence.

But every sentence is T-equivalent to a quantifier-free sentence, and hence the boolean algebra of sentences modulo T is finite.

Each complete extension is determined by the set of sentences that need to be added to T. Each such set is finite and, in fact, bounded by the size of the boolean algebra of all sentences. Thus T has only finitely many complete extension. ■

[*Held in 120../B22-bit.. Last changed July 26, 2004*]

B.2 For §2.2

2.2 We are required to show

$$(a) \quad T^+ \vdash (2) \qquad (b) \quad T^+ \vdash (3) \qquad (c) \quad T \vdash (4)$$

where (4) is the sentence of Lemma 2.5.

Handling internal (formal) induction can be quite tricky, especially when it is combined with external induction.

Let $\phi(v)$ be any formula with just the free variable v, as indicated. Notice that

$$\text{(case)} \quad \phi(0) \wedge (\forall v)\phi(Sv) \mathrel{.\!\!\rightarrow\!.} (\forall v)\phi(v)$$

is a simple consequence of the induction axiom for ϕ. This is the parameter-free case induction.

(a) Let

$$\phi(v) := (v \doteq 0) \vee (\exists w)[Sw \doteq v]$$

so that

$$T^+ \vdash \phi(0) \qquad T^+ \vdash \phi(Sv)$$

(for we can witness the approriate disjunct of ϕ). The required result

$$T^+ \vdash (\forall v)\phi(v)$$

is an immediate application of (case).

(b) Let

$$\phi_k(v) := (S^{k+1}v \neq v)$$

(for each $k \in \mathbb{N}$). First we check that

$$(i) \quad T^+ \vdash \phi_0(v) \qquad (ii) \quad T^+ \vdash \phi_{k+1}(0) \qquad (iii) \quad T^+ \vdash \phi_k(v) \rightarrow \phi_{k+1}(Sv)$$

hold.

We have

$$T^+ \vdash \phi_0(0)$$

by axiom (0), and

$$T^+ \vdash (S(Sv) \doteq Sv) \rightarrow (Sv \doteq v)$$

by axiom (1), so that

$$T^+ \vdash \phi_0(v) \rightarrow \phi_0(Sv)$$

and hence a use of internal induction gives (i).

An instance of axiom (0) fives (ii).

By axiom (1) we have

$$T^+ \vdash (S^{k+2}v \doteq Sv) \rightarrow (S^{k+1}v \doteq v)$$

which gives (iii).

We now show

$$[k] \quad T^+ \vdash (\forall v)\phi_k(v)$$

by an external induction over k.

The base case [0], is just (i).

For the induction step, $[k] \Rightarrow [k+1]$, we have

$$T^+ \vdash (\forall v)\phi_k(v) \rightarrow (\forall v)\phi_{k+1}(Sv)$$

by (iii), so that $[k]$, (ii), and (case) give $[k+1]$.

(c) We prove the two implications of (4) separately.

For each $l \leq k$ we have a sequence of implications

$$T \vdash (S^{k+1}w \doteq v) \wedge (v \doteq S^l 0) \rightarrow (S^{k+1}w \doteq S^l 0) \rightarrow (S^{(k-l)+1}w \doteq 0) \rightarrow \mathsf{false}$$

by the properties of equality, axiom (2), and axiom (0). Thus

$$T \vdash (\exists w)[S^{k+1}w \doteq v) \rightarrow (v \neq \ulcorner l \urcorner)$$

to give

$$T \vdash (\exists w)[S^{k+1}w \doteq v) \to (v \neq \ulcorner 0 \urcorner) \wedge \cdots \wedge (v \neq \ulcorner k \urcorner)$$

which is the first implication.

For the converse let

$$\phi_k(v) := (v \doteq \ulcorner 0 \urcorner) \vee \cdots \vee (v \doteq \ulcorner k \urcorner) \vee (\exists w)[S^{k+1}w \doteq v]$$

(for each $k \in \mathbb{N}$). We first observe that

$$(iv) \quad T \vdash \phi_{k+1}(0) \qquad (v) \quad T \vdash \phi_{k+1}(Sv) \leftrightarrow \phi_k(v)$$

hold. The first, (iv), is trivial. For the second, (v), we have

$$\begin{aligned} T \vdash \phi_{k+1}(Sv) &\leftrightarrow (Sv \doteq \ulcorner 0 \urcorner) \vee (Sv \doteq \ulcorner 1 \urcorner) \vee \cdots \vee (Sv \doteq \ulcorner k+1 \urcorner) \vee (\exists w)[S^{k+2}w \doteq Sv] \\ &\leftrightarrow \qquad\qquad\qquad\quad (Sv \doteq \ulcorner 1 \urcorner) \vee \cdots \vee (Sv \doteq \ulcorner k+1 \urcorner) \vee (\exists w)[S^{k+2}w \doteq Sv] \\ &\leftrightarrow \phi_k(v) \end{aligned}$$

using first axiom (0), and then axiom (1). With these we show

$$[k] \quad T \vdash (\forall v)\phi_k(v)$$

by an external induction over k.

The base case [0], is axiom (2).

For the induction step, $[k] \Rightarrow [k+1]$, we have

$$T \vdash (\forall v)\phi_{k+1}(Sv)$$

by $[k]$ and (v), so that (iv) and (case) give $[k+1]$. ■

2.3 (a) Consider any element a of any model $\mathfrak{A}$ of T. We may generate a chain

$$\{a(r) \mid r \in \mathbb{N}\}$$

of elements of $\mathfrak{A}$ by setting $a(r) = S^r a$. By axiom (3) (as given in Solution 2.2), all these elements are distinct. Thus we generate a copy of $\mathfrak{N}$. In particular, when a is the distinguished element of $\mathfrak{A}$, we identify this copy with $\mathfrak{N}$. Recall that axiom (0) says that 0 does not have a predecessor.

Now consider any element a not in (this canonical copy of) $\mathfrak{N}$. Using axiom (2) we may regress to form a chain

$$\{a(-r) \mid r \in \mathbb{N}\}$$

of elements of $\mathfrak{A}$ where $a(0) = a$ (as before) and $Sa(-(r+1)) = a(-r)$ for each $r \in \mathbb{N}$. Using axiom (3) we see these two chains form a copy of $\mathfrak{Z} = (\mathbb{Z}, S, 0)$.

We now remove $\mathfrak{N}$ from $\mathfrak{A}$. What we have left can be partitioned into blocks each of which is a copy of $\mathfrak{Z}$. We count the number of such blocks, to see that

$$\mathfrak{A} = \mathfrak{N} \cup \mathfrak{Z}^{(\kappa)}$$

where the superscript indicates there are κ copies of $\mathfrak{Z}$. In particular, this κ completely determines the (ismorphism) type of $\mathfrak{A}$. Let us write $\mathfrak{A}(\kappa)$ for this model. Thus every model is isomorphic to $\mathfrak{A}(\kappa)$ for precisely one κ.

Notice that the cardinality of $\mathfrak{A}(\kappa)$ is κ or $\aleph_0$, which ever is the larger.

(b) For uncountable κ, each model of T of cardinality κ is isomorphic to $\mathfrak{A}(\kappa)$.

(c) Each countable model of T is isomorphic to precisely one of

$$\mathfrak{N} = \mathfrak{A}(0) \subseteq \mathfrak{A}(1) \subseteq \cdots \subseteq \mathfrak{A}(r) \subseteq \cdots \subseteq \mathfrak{A}(\aleph_0)$$

which is a chain of order type $\omega + 1$.

This spectrum of models is typical of a large class of theories.

[`What are this kind of theories called?`]. ■

2.4 (a) Since $\mathfrak{N} \models T$, we have

$$T \vdash \delta \Longrightarrow \mathfrak{N} \models \delta$$

(for each quantifier-free sentence), so it suffices to show that converse. There are two ways to do this.

Each atomic sentence α of this language has the shape

$$(\ulcorner m \urcorner \doteq \ulcorner n \urcorner)$$

for some $m, n \in \mathbb{N}$ (or is one of true or false). Then

$$\begin{aligned} \mathfrak{N} \models \alpha \quad &\Longrightarrow (m = n) \Longrightarrow T \vdash \alpha \\ \mathfrak{N} \models \neg\alpha &\Longrightarrow (m \neq n) \Longrightarrow T \vdash \neg\alpha \end{aligned}$$

where the last implication makes use of axioms (1) and (0). Using this we show

$$\mathfrak{N} \models \delta \Longrightarrow T \vdash \delta \qquad \mathfrak{N} \models \neg\delta \Longrightarrow T \vdash \neg\delta$$

for each quantifier-free sentence δ. We proceed by induction on the construction of δ. Notice that we verify the two implications in tandem.

The base case, where δ is atomic, is just the observation above. For the induction step we survey the possible shapes of δ. There are four cases

$$\begin{array}{ll} \delta = (\alpha \wedge \beta) & T \vdash \neg\delta \leftrightarrow (\neg\alpha \vee \neg\beta) \\ \delta = (\alpha \vee \beta) & T \vdash \neg\delta \leftrightarrow (\neg\alpha \wedge \neg\beta) \\ \delta = (\alpha \rightarrow \beta) & T \vdash \neg\delta \leftrightarrow (\alpha \wedge \neg\beta) \\ \delta = \neg\gamma & T \vdash \neg\delta \leftrightarrow \gamma \end{array}$$

where α, β, γ are simpler sentences.

For the first case we have

$$\begin{array}{lllllll} \mathfrak{N} \models \delta & \Longrightarrow & \mathfrak{N} \models \alpha \text{ and } \mathfrak{N} \models \beta & \Longrightarrow & T \vdash \alpha \text{ and } T \vdash \beta & \Longrightarrow & T \vdash \delta \\ \mathfrak{N} \models \neg\delta & \Longrightarrow & \mathfrak{N} \models \neg\alpha \text{ or } \mathfrak{N} \models \neg\beta & \Longrightarrow & T \vdash \neg\alpha \text{ or } T \vdash \neg\beta & \Longrightarrow & T \vdash \neg\delta \end{array}$$

where, for both parts, the central implication uses the induction hypothesis.

The second case is similar.

For the third case we have

$$\begin{array}{lllllll} \mathfrak{N} \models \delta & \Longrightarrow & \mathfrak{N} \models \neg\alpha \text{ or } \mathfrak{N} \models \beta & \Longrightarrow & T \vdash \neg\alpha \text{ or} T \vdash \beta & \Longrightarrow & T \vdash \delta \\ \mathfrak{N} \models \neg\delta & \Longrightarrow & \mathfrak{N} \models \alpha \text{ and } \mathfrak{N} \models \neg\beta & \Longrightarrow & T \vdash \alpha \text{ and } T \vdash \neg\beta & \Longrightarrow & T \vdash \neg\delta \end{array}$$

where again the central implications use the induction hypothesis.

For the fourth case we have

$$\begin{array}{ccccccc}\mathfrak{N} \models \delta & \Longrightarrow & \mathfrak{N} \models \neg\gamma & \Longrightarrow & T \vdash \neg\gamma & \Longrightarrow & T \vdash \delta \\ \mathfrak{N} \models \neg\delta & \Longrightarrow & \mathfrak{N} \models \gamma & \Longrightarrow & T \vdash \gamma & \Longrightarrow & T \vdash \neg\delta\end{array}$$

where the central implications use the induction hypothesis. This is why we verify the two implications in tandem.

Another way to prove

$$\mathfrak{N} \models \delta \Longrightarrow T \vdash \delta$$

to to rephrase δ as a conjunction of disjunctions of literals, and then unravel this shape.

(b) We know that $T \subseteq T^+ \subseteq Th(\mathfrak{N})$. For the converse, consider any sentence σ. We have

$$T \vdash \sigma \leftrightarrow \delta$$

for some quantifier-free sentence δ. Then

$$\mathfrak{N} \models \sigma \Longrightarrow \mathfrak{N} \models \delta \Longrightarrow T \vdash \delta \Longrightarrow T \vdash \sigma$$

using part (a), to give the required result. ■

[*Held in 120../B23-bit.. Last changed July 26, 2004*]

B.3 For §2.3

2.6 Both $\mathbb{Q}$ and $\mathbb{R}$ are models of the theory T (of lines), which is complete, so that $Th(\mathbb{Q}) = T = Th(\mathbb{R})$ and hence $\mathbb{Q} \equiv \mathbb{R}$.

This shows that the (Dedekind) completeness of the line $\mathbb{R}$ can not be captured in a first order way (at least within the languages of lines). ■

2.7 Let $(A, \leq)$ be a linearly ordered set. We form a new linear ordering by replacing each element of A by a copy of $\mathbb{Q}$. These blocks are ordered as the occur in A. This gives the ordinal product $\mathbb{Q} \cdot A$.

In more detail, we look at the set of pairs

$$(p, a)$$

where $p \in \mathbb{Q}$ and $a \in A$. We compare these by

$$(p, a) \leq (q, b) \Longleftrightarrow (a < b) \text{ or } [(a = b) \text{ and } (p \leq q)]$$

(for $p, q \in \mathbb{Q}$ and $a, b \in A$). It is routine to check that this gives a line, and $a \longmapsto (0, a)$ is an embedding. ■

2.8 The $\aleph_0$-categoricity is the simplest application of the back-and-forth method. This is discussed in section 10.

Consider any cardinal $\kappa > \aleph_0$. There are several methods of producing non-isomorphic lines of cardinality κ. Here is one of them.

Consider any limit ordinal λ of cardinality κ. Let λ^* be the reverse of λ, and consider the ordinal products

$$\mathbb{Q}\cdot\lambda \qquad \mathbb{Q}\cdot\lambda^*$$

both of which are lines of cardinality κ. We know that, as a linear order, λ is embedded in the first, and λ^* is embedded in the second. We show that λ^* can not be embedded in the first, and so the two are not isomorphic.

By way of contradiction suppose λ^* can be embedded in $\mathbb{Q}\cdot\lambda$. Thus we have an indexed family

$$a(i) = (p(i), \alpha(i)) \qquad (i < \lambda)$$

of elements of the line such that

$$j < i < \lambda \Longrightarrow a(i) < a(j)$$

holds. Remembering how these pairs are compared, we see that

$$j < i < \lambda \Longrightarrow \alpha(i) \leq \alpha(j)$$

holds. In other words, we have a descending chain through λ. But λ is an ordinal, so this chain must stabilize after a finite number of steps. In other words, there is some finite $k < \omega$ such that

$$k \leq j < i < \lambda \Longrightarrow \alpha(i) = \alpha(j)$$

and hence

$$k \leq j < i < \lambda \Longrightarrow p(i) < p(j)$$

holds. This gives a subset of $\mathbb{Q}$ of cardinality $|\lambda| = \kappa$, which is impossible. ■

B.4 For §2.4–*not yet done*

[*Held in 120../B30-sols.. Last changed July 26, 2004*]

C For section 3

[*Held in 120../B31-sols.. Last changed July 26, 2004*]

C.1 For §3.1

3.2 Suppose $\mathfrak{A} \prec_r \mathfrak{B}$ and $\mathfrak{A} \models \sigma$ where σ is a $\exists_{r+1}$-sentence. We have

$$\sigma = (\exists v)\phi(v)$$

for some $\forall_r$-formula $\phi(v)$, and there is some point a of $\mathfrak{A}$ with $\mathfrak{A} \models \phi(a)$. But now $\mathfrak{A} \prec_r \mathfrak{B}$ gives $\mathfrak{B} \models \phi(a)$ so that $\mathfrak{B} \models \sigma$, as required. ■

[*Held in 120../B32-sols.. Last changed July 26, 2004*]

C.2 For §3.2

3.3 Given such a pair of embeddings, for each $\forall_{n+1}$-sentence we have

$$\mathfrak{A} \models \sigma \Longrightarrow \mathfrak{C} \models \sigma \Longrightarrow \mathfrak{B} \models \sigma$$

using first f then g.

Conversely, suppose that $\mathfrak{A} \equiv\rangle(\forall_{n+1})\, \mathfrak{B}$ holds. Let a be an enumeration of $\mathfrak{A}$ and let b be an enumeration of $\mathfrak{B}$. We enrich the underlying language L to both $L(\mathsf{a})$ and $L(\mathsf{b})$ and combine these to obtain $L(\mathsf{a}, \mathsf{b})$. We arrange that the a-parameters and the b-parameters are disjoint. Look at

$$Th(\mathfrak{A}, \mathsf{a}) \qquad \Delta(\mathfrak{B}, \mathsf{b}) = Th(\mathfrak{B}, \mathsf{b}) \cap \forall_n$$

which are, respectively, a set of $L(\mathsf{a})$-sentences and a set of $L(\mathsf{b})$-sentences. These combine to give

$$Th(\mathfrak{A}, \mathsf{a}) \cup \Delta(\mathfrak{B}, \mathsf{b})$$

a set of $L(\mathsf{a}, \mathsf{b})$-sentences, and it suffices to show that this set is consistent.

By way of contradiction, suppose the set is not consistent. Then

$$Th(\mathfrak{A}, \mathsf{a}) \vdash \neg\delta(b)$$

for some $\forall_n$-formular $\delta(w)$ and a point b of $\mathfrak{B}$ with $\mathfrak{B} \models \delta(b)$. This gives first

$$Th(\mathfrak{A}, \mathsf{a}) \vdash (\forall w)\neg\delta(w)$$

and then

$$\mathfrak{A} \models (\forall w)\neg\delta(w)$$

(since a and b are disjoint). But $(\forall w)\neg\delta(w)$ is a $\forall_{n+1}$-sentence (of the original language), and hence

$$\mathfrak{B} \models (\forall w)\neg\delta(w)$$

which leads to the contradiction. ■

[*Held in 120../B33-sols.. Last changed July 26, 2004*]

C.3 For §3.3

3.5 This proof is very like that of Solution 3.3, and generalizes the proof of Lemma 3.19.

On of the required implications is immediate. For the other suppose $\mathfrak{A} \prec_{n+1} \mathfrak{B}$. Let a be an enumeration of $\mathfrak{A}$, and let b be an enumeration of $\mathfrak{B}$. It suffices to show that

$$Th(\mathfrak{A}, \mathsf{a}) \cup (Th(\mathfrak{B}, \mathsf{a}, \mathsf{b}) \cap \forall_n)$$

is consistent. The left hand component is a set of $L(\mathsf{a})$ sentences. The right hand component is a set of $L(\mathsf{a}, \mathsf{b})$-sentences, and will contains some sentences $(a \doteq b)$ where a and b are parameters that name the same element of $\mathfrak{B}$.

By way of contradiction suppose the displayed set is not consistent. Then

$$Th(\mathfrak{A}, \mathsf{a}) \vdash \neg\delta(b, a)$$

for some $\forall_n$-formula $\delta(v, u)$ of the parent language, an points a from a and b from b. This gives $\mathfrak{A} \models (\forall v)\neg\delta(v, a)$ and then $\mathfrak{B} \models (\forall v)\neg\delta(v, a)$ (since $\mathfrak{A} \prec_{n+1} \mathfrak{B}$), which leads to the contradiction. ■

[*Held in 120../B34-sols.. Last changed July 26, 2004*]

C.4 For §3.4

3.6 We verify this by induction on n.

The base case, $n = 0$, is trivial.

For the induction step, $n \mapsto n + 1$, suppose

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_{n+1} \mathfrak{B}$$

for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$. Consider any $\mathfrak{A} \in \mathcal{A}$, and suppose

$$\mathfrak{A} \models (\forall v)\phi(v, a)$$

where $\phi(v, u)$ is a $\exists_n$-formula, and a is a point from $\mathfrak{A}$. To show

$$\mathfrak{U} \models (\forall v)\phi(v, a)$$

consider any point b from $\mathfrak{U}$ which matches v. There is some $\mathfrak{B} \in \mathcal{A}$ with $\mathfrak{A} \subseteq \mathfrak{B}$ and such that b come from $\mathfrak{B}$. But now $\mathfrak{A} \prec_{n+1} \mathfrak{B}$, so that

$$\mathfrak{B} \models (\forall v)\phi(v, a)$$

and hence

$$\mathfrak{B} \models \phi(b, a)$$

holds.

The induction hypothesis gives $\mathfrak{B} \prec_n \mathfrak{U}$, and hence

$$\mathfrak{U} \models \phi(b, a)$$

as required. ■

3.7 Using induction on n we show that

$$[n] \quad \mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_n \mathfrak{B}$$

holds for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$ (quantified inside the induction hypothesis).

The base case, $[0]$, is trivial.

For the induction step, $[n] \Rightarrow [n+1]$, consider $\mathfrak{A}, \mathfrak{B} \in \mathcal{A}$ with $\mathfrak{A} \subseteq \mathfrak{B}$. The condition on $\mathcal{A}$ gives

$$\mathfrak{A} \prec \mathfrak{C} \qquad \mathfrak{B} \subseteq \mathfrak{C}$$

for some $\mathfrak{C} \in \mathcal{A}$. The induction hypothesis applied to $\mathfrak{B}, \mathfrak{C}$ gives $\mathfrak{B} \prec_n \mathfrak{C}$, and hence $\mathfrak{A} \prec_{n+1} \mathfrak{B}$, as required. ■

[*Held in 120.../B35-sols.. Last changed July 26, 2004*]

C.5 For §3.5

3.9 Under the given conditions we show

$$\mathfrak{A} \prec_n \mathfrak{B}$$

by induction on n. The base case, $n = 0$, is trivial. For the induction step, $n \mapsto n+1$, suppose $\mathfrak{B} \models (\exists v)\phi(v, a)$ where $\phi(v, u)$ is a $\forall_n$-formula, and a is a point from $\mathfrak{A}$. We do not assume that v is a single variable, since iterated use of the given condition gives $\mathfrak{B} \models \phi(b, a)$ for some point b also from $\mathfrak{A}$. The induction hypothesis, $\mathfrak{A} \prec_n \mathfrak{B}$, now gives $\mathfrak{A} \models \phi(b, a)$, so that $\mathfrak{A} \models (\exists v)\phi(v, a)$, as required. ■

[*Held in 120../B40-sols.. Last changed July 26, 2004*]

D For section 4

D.1 For §4.1 -*to be done*

[*Held in 120../B42-sols.. Last changed July 26, 2004*]

D.2 For §4.2

4.3 This has two different proofs, both of which are instructive.

For the first proof observe that the characterization of Theorem 4.7(ii) depends only $T \cap \forall_2$.

For the second proof we show that if T has AP then so does $T \cap \forall_2$. (The other required implication is immediate.)

Consider any wedge of embeddings between models of $T \cap \forall_2$, as in (1) of Table 4. By enlarging $\mathfrak{B}, \mathfrak{C}$ we may suppose that both are models of T. A similar trick gives us a diagram as in (2) of Table 4 where $\mathfrak{A}', \mathfrak{B}, \mathfrak{C} \in \mathcal{M}(T)$ and l is a $\prec_1$-embedding. (This is because $\mathfrak{A} \models T \cap \forall_2$.) Two uses of Lemma 4.5 gives us a commuting diagram of embeddings as in (3) of Table 4 where m and n are $\prec$-embeddings. In particular, $\mathfrak{A}', \mathfrak{B}', \mathfrak{C}' \in \mathcal{M}(T)$. But T has AP, so the right hand wedge can be closed to give a diagram as in (4) of Table 4 where $\mathfrak{D}' \models T$. ∎

4.4 Suppose $T \cap \forall_1$ has AP, and consider formulas $\theta(v) \in \exists_1$ and $\phi(v) \in \forall_1$ (in the indicated variables) such that

$$T \vdash \theta(v) \rightarrow \phi(v)$$

holds. Let $\Delta(v)$ be the set of quantifier-free formulas $\delta(v)$ such that

$$T \vdash \theta(v) \rightarrow \delta(v)$$

holds. If suffices to show that

$$T \cup \Delta(v) \vdash \phi(v)$$

holds.

Consider any model $(\mathfrak{B}, a)$ of the hypothesis set. Thus

$$\mathfrak{B} \models T \qquad \mathfrak{B} \models \Delta(a)$$

for some point a of $\mathfrak{B}$. We must show that $\mathfrak{B} \models \phi(a)$.

Let $\mathfrak{A}$ be the substructure of $\mathfrak{B}$ generated by a. Thus $\mathfrak{A} \models T \cap \forall_1$. A simple argument shows that

$$T \cup Diag(\mathfrak{A}, a) \cup \{\theta(a)\}$$

is consistent. Thus we have a wedge of embeddings as in (1) of Table 4 where $\mathfrak{C} \models T$ and $\mathfrak{C} \models \theta(ga)$. Here f is the insertion of $\mathfrak{A}$ into $\mathfrak{B}$. Since $T \cap \forall_1$ has AP, this closes to

(1) (2) (3) (4)

$$\begin{tikzcd}
& \mathfrak{B} \\
\mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & \\
& \mathfrak{C}
\end{tikzcd}
\qquad
\begin{tikzcd}
& \mathfrak{B} \\
\mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] \arrow[r, "l"] & \mathfrak{A}' \\
& \mathfrak{C}
\end{tikzcd}
\qquad
\begin{tikzcd}
& \mathfrak{B} \arrow[r, "m"] & \mathfrak{B}' \\
\mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] \arrow[r, "l"] & \mathfrak{A}' \arrow[ur, "f'"'] \arrow[dr, "g'"] & \\
& \mathfrak{C} \arrow[r, "n"'] & \mathfrak{C}'
\end{tikzcd}
\qquad
\begin{tikzcd}
& \mathfrak{B} \arrow[r, "m"] & \mathfrak{B}' \arrow[dr] & \\
\mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] \arrow[r, "l"] & \mathfrak{A}' \arrow[ur, "f'"'] \arrow[dr, "g'"] & & \mathfrak{D}' \\
& \mathfrak{C} \arrow[r, "n"'] & \mathfrak{C}' \arrow[ur] &
\end{tikzcd}$$

Table 4: Diagrams for Solution 4.3

---

a commuting square

$$\begin{tikzcd}
& \mathfrak{B} \arrow[dr, "h"] & \\
\mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & & \mathfrak{D} \\
& \mathfrak{C} \arrow[ur, "k"'] &
\end{tikzcd}$$

where $\mathfrak{D} \models T \cap \forall_1$. In fact, by a suitable enlargement we can arrange that $\mathfrak{D} \models T$.

We have

$$\mathfrak{C} \models \theta(ga) \qquad T \vdash \theta \to \phi$$

and $\theta \in \exists_1$, so that

$$\mathfrak{D} \models \theta((k \circ g)a) \qquad T \vdash \theta \to \phi$$

to give

$$\mathfrak{D} \models \theta((h \circ f)a)$$

since $h \circ f = k \circ g$. But $\phi \in \forall_1$, so that

$$\mathfrak{B} \models \phi(fa)$$

and hence

$$\mathfrak{B} \models \phi(a)$$

since f is an insertion. This is the required result. ■

4.5 Suppose first that $\mathfrak{A}$ is an amalgamation base for T, and suppose

$$T \vdash \psi \vee \phi$$

where $\psi(v)$ and $\phi(v)$ are $\forall_1$-formulas. By way of contradiction, suppose $\mathfrak{A}$ realizes the associated types at some point a. Let a be an enumeration of the whole of $\mathfrak{A}$. We check that both

$$T \cup Diag(\mathfrak{A}, \mathsf{a}) \cup \{\neg\psi(a)\} \qquad T \cup Diag(\mathfrak{A}, \mathsf{a}) \cup \{\neg\phi(a)\}$$

are consistent, and hence we obtain a wedge of embeddings as in (1) of Table 4 where

$$\mathfrak{B} \models \neg\psi(a) \qquad \mathfrak{C} \models \neg\phi(a)$$

and with $\mathfrak{B}, \mathfrak{C} \in \mathcal{M}(T)$. Since $\mathfrak{A}$ is an amalgamation base, we may close this wedge, and then argue as in the proof of Theorem 4.7, (i)⇒(ii).

Conversely, suppose $\mathfrak{A} \in \mathcal{S}(T)$ has the omitting property, and consider a wedge of embeddings as in (1) of Table 4 where $\mathfrak{B}, \mathfrak{C} \in \mathcal{M}(T)$. As in the proof of Theorem 4.7, (ii)⇒(i) we may suppose

$$\mathfrak{A} \subseteq \mathfrak{B} \quad \mathfrak{A} \subseteq \mathfrak{C} \quad \mathfrak{B} \cap \mathfrak{C} = \mathfrak{A}$$

hold. Let $\mathsf{a}, \mathsf{b}, \mathsf{c}$ be enumerations of $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$ respectively. It suffices to show that

$$T \cup \mathit{Diag}(\mathfrak{B}, \mathsf{b}, \mathsf{a}) \cup \mathit{Diag}(\mathfrak{C}, \mathsf{c}, \mathsf{a})$$

is consistent.

By way of contradiction, suppose this set is not consistent. Then we find $\forall_1$-formulas $\psi(u), \phi(u)$ such that

$$T \vdash \psi \vee \phi \quad \mathfrak{B} \models \neg\psi(a) \quad \mathfrak{A} \models \neg\phi(a)$$

for some point a of $\mathfrak{A}$. This point a can not realize

$$\neg\Sigma(T, \psi) \cup \neg\Sigma(T, \phi)$$

in $\mathfrak{A}$. Suppose it does not realize $\neg\Sigma(T, \psi)$, say. Then we have

$$\mathfrak{A} \models \lambda(a) \qquad T \vdash \lambda \to \psi$$

for some $\exists_1$-formula $\lambda(u)$, and hence $\mathfrak{B} \models \lambda(a)$, which leads to the contradiction. ■

4.6 (a) This is almost immediate.

(b) This is prove in the same way as the argument of Solution 4.3.

(c) This is an immediate consequence of the characterization given in Exercise 4.5. ■

D.3 For §4.3 -*no exercises yet*

[*Held in 120../B50-sols.. Last changed July 26, 2004*]

E For section 5

[*Held in 120../B51-sols.. Last changed July 26, 2004*]

E.1 For §5.1

5.1 (a) Suppose T^* is a model completion of T. Consider any model $\mathfrak{A} \models T^*$. Then $\mathfrak{A}T$, so that $T^*[\mathfrak{A}]$ is complete, to show that T^* is model complete.

Consider any wedge

$$\begin{tikzcd} & \mathfrak{B} \\ \mathfrak{A} \arrow[ur] \arrow[dr] & \\ & \mathfrak{C} \end{tikzcd}$$

of models of T. By enlargement we may suppose $\mathfrak{B}, \mathfrak{C}$ are models of T^*. In the usual way we may suppose $\mathfrak{B} \cap \mathfrak{C} = \mathfrak{A}$. Let L be the underlying language. The three theories $T^*[\mathfrak{A}], T^*[\mathfrak{B}], T^*[\mathfrak{C}]$ are complete in the respective enrichments of L. Let a be the enumeration of $\mathfrak{A}$, so that $L(\mathfrak{A}) = L(\mathsf{a})$. Then

$$T^*[\mathfrak{A}] \subseteq T^*[\mathfrak{B}] \cap L(\mathsf{a})$$

to give equality, with a simuilar observation for $\mathfrak{C}$. Thus

$$(\mathfrak{B}, \mathsf{a}) \equiv (\mathfrak{C}, \mathsf{a})$$

which, by a standrad argument, leads to an amalgam.

(Notice that this argument does *not* show $\mathfrak{A} \prec \mathfrak{B}$, unless $\mathfrak{A}$ is a model of T^*.)

(b) Given the considtions, it suffices to show that $T^*[\mathfrak{A}]$ is comlete for each $\mathfrak{A} \models T$. Consider any pair of models of this enriched theory. Thus we have a wedge of embeddings

$$\begin{tikzcd} & \mathfrak{B} \\ \mathfrak{A} \arrow[ur] \arrow[dr] & \\ & \mathfrak{C} \end{tikzcd}$$

where $\mathfrak{B}, \mathfrak{C} \in \mathcal{M}(T^*)$. Since T has AP, we may close this wedge

$$\begin{tikzcd} & \mathfrak{B} \arrow[dr] & \\ \mathfrak{A} \arrow[ur] \arrow[dr] & & \mathfrak{D} \\ & \mathfrak{C} \arrow[ur] & \end{tikzcd}$$

to a model $\mathfrak{D} \models T^*$. But T^* is model complete, so both the right hand embeddings are elementary, which leads to the required result.

(c) *???????* ■

5.2 Consider elements a, b of a commutative ring such that

$$ab \neq 0 \qquad (ab)^2 = ab$$

hold. A simple induction gives

$$a^{n+1}b^{n+1} = ab$$

for each $n \in \mathbb{N}$, and hence $a^n \neq 0$.

This is enough to prove both Lemma 5.6 and Lemma 5.7((ii)⇒(i)). ■

[*Held in 120../B52-sols.. Last changed July 26, 2004*]

E.2 For §5.2

5.4 By Definition 5.10(i) the two theories T, T^b are companions, and hence $T^b = T^{ab}$ by 5.10(ii).

By 5.10(iii) we have

$$T^a \cap \forall_2 \subseteq T^{ab} = T^b$$

and hence $T^a \cap \forall_2 \subseteq T^b \cap \forall_2$. The converse holds by symmetry. ■

[*Held in 120../B53-sols.. Last changed July 26, 2004*]

E.3 For §5.3

5.5 Suppose $\mathcal{E}(T)$ is elementary and let $T^* = Th(\mathcal{E}(T))$, so that $\mathcal{E}(T) = \mathcal{M}(T^*)$. Since

$$\mathfrak{A} \subseteq \mathfrak{B} \Longrightarrow \mathfrak{A} \prec_1 \mathfrak{B}$$

for $\mathfrak{A}, \mathfrak{B} \in \mathcal{M}(T^*)$. we see that T^* is model complete. Also $\mathcal{M}(T^*) = \mathcal{E}(T) \subseteq \mathcal{S}(T)$, so it suffices to show $\mathcal{M}(T) \subseteq \mathcal{S}(T^*)$. But this is an immediate consequence of the existence Theorem 5.19. ■

5.6 Let $\mathfrak{U}$ be the union of a directed family $\mathcal{A} \subseteq \mathcal{E}(T)$, consider $\mathfrak{U} \subseteq \mathfrak{B} \models T$, and suppose $\mathfrak{U} \models \phi(a)$ for some $\forall_1$-formula $\phi(v)$ and point a from $\mathfrak{U}$. There is some $\mathfrak{A} \in \mathcal{A}$ such that a comes from $\mathfrak{A}$. But then $\mathfrak{A} \models \phi(a)$ and $\mathfrak{A} \prec_1 \mathfrak{B}$, to give $\mathfrak{B} \models \phi(a)$, as required. ■

5.7 (a) Suppose

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\, (\mathfrak{B}, b)$$

so that we have a pair of embeddings

$$(\mathfrak{A}, a) \xrightarrow{f} (\mathfrak{C}, c) \qquad (\mathfrak{B}, a) \xrightarrow{g} (\mathfrak{C}, c)$$

where g is elementary. In particular, $\mathfrak{C} \in \mathcal{S}(T)$. But $\mathfrak{A} \in \mathcal{E}(T)$, so that f is a $\prec_1$-embedding, and hence

$$(\mathfrak{A}, a) \equiv_1 (\mathfrak{C}, c) \equiv_1 (\mathfrak{B}, b)$$

as required.

(b) Consider first $\mathfrak{A}, \mathfrak{B} \in \mathcal{E}(T)$ with $\mathfrak{A} \subseteq \mathfrak{B}$. Then $\mathfrak{A} \prec_1 \mathfrak{B}$ (since $\mathfrak{A} \in \mathcal{E}(T)$), and hence

$$\mathfrak{A} \subseteq \mathfrak{B} \subseteq \mathfrak{C} \qquad \mathfrak{A} \prec \mathfrak{C}$$

for some structure $\mathfrak{C}$. In particular, $\mathfrak{C} \in \mathcal{S}(T)$, so that $\mathfrak{B} \prec_1 \mathfrak{C}$, and hence $\mathfrak{A} \prec_2 \mathfrak{B}$.

Now suppose

$$\mathfrak{A} \equiv\rangle(\exists_1)\, \mathfrak{B}$$

where $\mathfrak{A}, \mathfrak{B} \in \mathcal{E}(T)$. Up to isomorphism there is some $\mathfrak{C} \in \mathcal{E}(T)$ with $\mathfrak{A} \subseteq \mathfrak{C}$ and $\mathfrak{B} \subseteq \mathfrak{C}$. But now, by the observation above, we have $\mathfrak{A} \prec_2 \mathfrak{C}$ and $\mathfrak{B} \prec_2 \mathfrak{C}$, to give

$$\mathfrak{A} \equiv_2 \mathfrak{C} \equiv_2 \mathfrak{B}$$

as required. ■

5.8 Let $T^e = Th(\mathcal{E}(T))$. Since $\mathcal{E}(T) \subseteq \mathcal{S}(T)$, we have $T \cap \forall_1 \subseteq T^e \cap \forall_1$. The converse follows by Theorem 5.19.

Lemma 5.17 shows that $\mathcal{E}(T) = \mathcal{E}(T \cap \forall_1)$, and hence T^e depends only on $T\forall_1$.

Finnaly, consider any $\mathfrak{A} \in \mathcal{E}(T)$. We have $\mathfrak{A} \subseteq \mathfrak{B}$ for some $\mathfrak{B} \models T$. But Then $\mathfrak{A} \prec_1 \mathfrak{B}$, and hence $\mathfrak{B} \equiv\rangle(\forall_2)\, \mathfrak{A}$. so that $\mathfrak{A} \models T \cap \forall_2$. This shows that $T \cap \forall_2 \subseteq T^e$. ■

5.9 By Lemma 5.6 it suffices to show

$$\mathcal{E}(T) \models (\forall v)[\bigwedge \mathrm{NNil}(v) \to \mathrm{nnil}(v)]$$

holds. To this end consider any $\mathfrak{A} \in \mathcal{E}(T)$ and any element a of $\mathfrak{A}$ with $\mathfrak{A} \models \mathrm{Nnil}(a)$. This element is not nilpotent, so that Lemma 5.7 gives some $\mathfrak{A} \subseteq \mathfrak{b} \models T$ with $\mathfrak{B} \models \mathrm{nnil}(a)$. But now $\mathfrak{A} \prec_1 \mathfrak{B}$ and hence $\mathfrak{A} \models \mathrm{nnil}(a)$, as required.

This does not show

$$T^e \cup \mathrm{NNil}(v) \vdash \mathrm{nnil}(v)]$$

for the argument above works only for $\mathfrak{A} \in \mathcal{E}(T)$, not for arbitrary $\mathfrak{A} \models T^e$. ■

E.4 For §5.4 *-to be done*

[*Held in 120../B55-sols.. Last changed July 26, 2004*]

E.5 For §5.5

5.11 Suppose $T^0 \subseteq T$ and consider any $\forall_1$-formula ϕ which is consistent with T. Then ϕ is consistent with T^0, so that

$$T^0 \vdash \theta \to \phi$$

for some $\exists_1$-formula which is consistent with T^0. But T and T^0 are companions, so θ is consistent with T, to show that T is 0-complete.

Conversely, suppose T is 0-complete, and consider any $\forall_2$-sentence σ which is 0-tame over T. We require $T \vdash \sigma$.

By way of contradiction, suppose $T \cup \{\neg\sigma\}$ is consistent. Let $\sigma = (\forall v)\neg\phi(v)$ where ϕ is a $\forall_1$-formula. Then $\phi(v)$ is consistent with T, so that, since T is 0-complete, we have

$$T \vdash \theta \to \phi$$

for some $\exists_1$-formula $\theta(v)$ which is consistent with T. Let $\alpha = (\forall v)\neg\theta(v)$, so that

$$T \vdash \sigma \to \alpha$$

and hence $\vdash \alpha$, since σ is 0-tame over T. This leads to the contradiction. ■

[*Held in 120../B60-sols.. Last changed July 26, 2004*]

F For section 6

[*Held in 120../B61-sols.. Last changed July 26, 2004*]

F.1 For §6.1

6.1 Observe that the formal Definition 6.1 and the implication of the exercise are contrapostives. ■

6.2 Consider any $\forall_1$-formula ϕ which is consistent with T^0. Since T^0 is 0-complete, there is a $\exists_1$-formula ψ which is consistent with T^0 and with $T^0 \vdash \psi \to \phi$. This formula ψ is consistent with T, which is given to be $\exists_1$-atomic, so there is some formula θ which is $\exists_1$-complete over T and with $T^0 \vdash \theta \to \psi$. THus we have $T^0\theta \to \phi$ and hence $T\theta \to \phi$, as required. ■

6.3 (a) An $\exists_{n+1}$-formula θ is $\exists_{n+1}$-complete over a theory T if it is consistent with T and

$$\left.\begin{array}{l} T \cup \{\theta, \psi_1\} \text{ is consistent} \\ T \cup \{\theta, \psi_2\} \text{ is consistent} \end{array}\right\} \Longrightarrow T \cup \{\psi_1, \psi_2\} \text{ is consistent}$$

holds for all $\exists_{n+1}$-formulas ψ_1, ψ_2 (with $\partial\psi_1 \cup \partial\psi_2 \subseteq \partial\theta$). Equivalently, using Exercise 6.1, an $\exists_{n+1}$-formula θ is $\exists_1$-complete over T if

$$T \vdash \phi_1 \vee \phi_2 \Longrightarrow T \vdash \theta \to \phi_1 \text{ or } T \vdash \theta \to \phi_2$$

holds for all $\forall_{n+1}$-formulas ϕ_1, ϕ_2.

(b) Consider a formula θ which is $\exists_{n+1}$-complete over T for all large n. In particular, θ is consistent with T. Consider any formula ψ (with $\partial\psi \subseteq \partial\theta$). Simce θ, is consistent with T, at most one of

$$T \vdash \theta \to \psi \qquad T \vdash \theta \to \neg\psi$$

holds, so it suffices to show that at least one holds. We may chose n large enough so that θ is a $\exists_{n+1}$-formula, and both $\psi, \neg\psi$ are $\forall_{n+1}$-formulas. Since

$$T \vdash \psi \vee \neg\psi$$

the required result follows by the variant of the official definition.

Conversely, suppose that θ is a $\exists_{n+1}$-formula which is complete over T. Suppose

$$T \vdash \phi_1 \vee \phi_2$$

where both ϕ_1 and ϕ_2 is a $\forall_{n+1}$-formula. Since θ is complete over T at least one of

$$T \vdash \theta \to \phi_1 \qquad T \vdash \theta \to \neg\phi_1$$

holds. If then second one holds, then so does

$$T \vdash \theta \to \phi_2$$

(since $T \vdash \phi_1 \vee \phi_2$) to show that θ is $\exists_{n+1}$-complete over T. ■

6.4 *To be done if necessary* ■

6.5 This is needed to ensure that the required b&f system is non-empty. ■

6.6 This is a refinement of the proof of Theorem 6.8. Let $\mathfrak{A}$ be the unique countable $\exists_1$-atomic structure for T. Consider and model $\mathfrak{B} \models T^0$.

For the time being suppose [*Check this is right*]

$$(\mathfrak{A}, a) \equiv\rangle(\exists_1)\,(\mathfrak{B}, b)$$

for points a from $\mathfrak{A}$ and b from $\mathfrak{B}$. Since $\mathfrak{A} \in \mathcal{E}(T)$, this ensures that

$$(\mathfrak{A}, a) \equiv_1 (\mathfrak{B}, b)$$

holds. Consider any element x of $\mathfrak{A}$. As in the proof of Theorem 6.8, there is some element y of $\mathfrak{B}$ such that

$$(\mathfrak{A}, a^\frown x) \equiv\rangle(\exists_1)\,(\mathfrak{B}, b^\frown y)$$

holds.

Now let a be a full enumeration of $\mathfrak{A}$. By iterated use of the above observation we produce a partial enumeration b of $\mathfrak{B}$ such that

$$(\mathfrak{A}, \mathsf{a}) \equiv\rangle(\exists_1)\,(\mathfrak{B}, \mathsf{b})$$

holds, and hence

$$(\mathfrak{A}, \mathsf{a}) \equiv_1 (\mathfrak{B}, \mathsf{b})$$

(since $\mathfrak{A} \in \mathcal{E}(T)$). This gives the required embedding. ■

6.7 Left $\mathfrak{A}$ be the unique countable structure which is $\exists_1$-atomic for T. Since $\mathcal{A}(T) \subseteq \mathcal{E}(T)$, we have $T^* = T^e \subseteq T^a$. Consider any $\mathfrak{B} \models T^*$. By Exercise 6.5 we have $\mathfrak{A} \subseteq \mathfrak{B}$, and hence $\mathfrak{A} \prec \mathfrak{B}$ (since T^* is model complete), and hence $\mathfrak{B} \models T^a$, to show that $T^* = T^a$. ■

F.2 For §6.2 -*no exercises yet*

F.3 For §6.3 -*to be done*

F.4 For §6.4 -*no exercises yet*

[*Held in 120../B70-sols.. Last changed July 26, 2004*]

G For section 7

[*Held in 120-../B71-sols.. Last changed July 26, 2004*]

G.1 For §7.1

7.1 (a) This is immediate.

(b) Suppose the sentence σ is good over T. In particular, $T \cup \{\sigma\}$ is consistent but does not axiomatize a comlete theory. Thus there is at least one sentence *tau* such that both

$$T \cup \{\sigma, \tau\} \qquad T \cup \{\sigma, \neg\tau\}$$

are consistent. Let

$$\sigma(0) = \sigma \wedge \tau \qquad \sigma(1) = \sigma \wedge \neg\tau$$

so that $T \cup \{\sigma(i)\}$ is consistent for $i = 0, 1$. Since $T \vdash \sigma(i) \rightarrow \sigma$, this sentence $\sigma(i)$ is good over T.

This gives us a 1-step splitting, which can be iterated to produce a suitable tree of good sentences. ■

7.2 Consider any $\exists_1$-sentence which is not-jep over T. Thus $T \cup \{\sigma\}$ is consistent and does not axiomatize a theory with JEP. (In fact, σ has a stronger property, but we don't need that just yet.) This gives a pair α, β of $\forall_1$-sentences such that

$$T \vdash \sigma \rightarrow \alpha \vee \beta \qquad T \cup \{\sigma, \neg\alpha\} \text{ is consistent} \qquad T \cup \{\sigma, \neg\beta\} \text{ is consistent}$$

hold. Let

$$\sigma(0) = \sigma \wedge \neg\alpha \qquad \sigma(1) = \sigma \wedge \neg\beta$$

so that

$$T \vdash \sigma(i) \rightarrow \sigma \qquad T \vdash \neg\sigma(0) \vee \neg\sigma(1)$$

and, furthermore, both $\sigma(i)$ are not-jep over T.

By iterating this splitting we produce a continuum of consistent extensions T_p of T each formed by adding a certain set Σ_p of $\exists_1$-sentences to T. Furthermore, for distinct indexes p, q, the set

$$T \cup \Sigma_p \cup \Sigma_q$$

is not consistent.

For each p let $\mathfrak{A} \models T_p$. This can be embedded in some $\mathfrak{A} \subseteq \mathfrak{B}_p \in \mathcal{G}(T)$, and $\mathfrak{B}_p \models \Sigma_p$ (since Σ_p consists of $\exists_1$-sentences). The incompatibility of the $\Sigma_\bullet$ ensure that $\mathfrak{B}_p \not\equiv \mathfrak{B}_q$ for distinct p, q. ■

[*In fact, the french paper shows that* $j(\mathcal{F}(T)) = 2^{\aleph_0}$]

G.2 For §7.2 -*to be done*

G.3 For §7.3 -*no exercises yet*

G.4 For §7.4 -*to be done*