1 Model theory; introduction and overview

Model theory, sometimes described as "algebra with quantifiers" is, roughly, mathematics done with attention to definability. The word can refer to definability of a whole class of abstract mathematical objects (like the way we can axiomatically define groups, or metric spaces, or ...) but more often it refers to how we define significant subsets of such objects (like solution-sets of polynomial equations, or ϵ-neighbourhoods of points, ...). Sometimes we will be looking at uniform definability within a class of objects (e.g. we can write down a definition of the centre of a group which will apply in all groups), more often we look at the definable subsets of a specific mathematical object (of, e.g., n-dimensional space over a field, or the complex numbers).

"Definability" means definability in some formal language which is set up to be appropriate to the kind of structure we are interested in. Describing such (finitary, first-order) languages is something that most (but probably not all) of you have seen in a course that includes some Predicate Logic. Having chosen a language $\mathcal{L}$, one describes how to build up the terms and formulas of $\mathcal{L}$. Then one says what is meant by an $\mathcal{L}$-structure and one makes the connection between syntax (formulas of $\mathcal{L}$) and semantics (the interpretation of formulas in $\mathcal{L}$-structures), in particular describing how to read $\mathcal{L}$-formulas in $\mathcal{L}$-structures. If you've not seen this before, or want to remind yourself about it, then see the Appendix, which also defines homomorphisms - the structure-preserving maps which connect $\mathcal{L}$-structures.

In this course, we will emphasise 'semantics' - structure and particular examples, taking the view that the more formal aspects are what we use for proving general results that can then be applied in many different contexts.

We will begin by describing the ultraproduct construction and giving some examples of the weird and wonderful things that we can get by using it and the associated theorem of Łos. We will also use it to give a proof the Compactness Theorem which does not depend on setting up a formal proof system.

The interplay of syntax (the formulas of a formal language) and semantics (the meanings of those formulas in structures) leads to the notion of the theory of a structure (the set of all sentences true in the structure) and the class of models of a theory (the class of structures in which all sentences in the theory are true). The latter is an extension of the usual way of specifying abstract classes by writing down a set of axioms. In this "logic narrative", elementary equivalence - having identical syntactically-expressible properties - is the key relation and the class of models of a theory is the natural context, with a central task being to classify the models.

In the more algebraic narrative, we can start with the observation that we often want to understand the solution sets of equations or inequalities and, if there is more than one variable, the projections of these solution sets to fewer variables. These are subsets which are definable in the kind of language $\mathcal{L}$ referred to above, so understanding the definable subsets becomes a central task.

We'll see ideas and results from both these strands which, in any case, cannot be disentangled. Model theory is abstract - it applies in many different contexts (algebraic, analytic, geometric) - and concept-heavy but we will use a variety of examples - some familiar, some less so - to illustrate and apply the results; this should help to make them more concrete (or 'concrete' - this is pure maths).

2 Ultraproducts and Łos' Theorem

2.1 Producing infinitesimals

Question 2.1. Is $0.\overline{9} = 1.0$?

Consider the difference $\epsilon = 1 - 0.\overline{9}$; what can we say about it? It's easy to see that we must have $\epsilon < 1/n$ for every positive integer n but why shouldn't we be able to add the condition $\epsilon > 0$? We can appeal to the Compactness Theorem, which you might have seen already (and which we will prove using ultraproducts).

Theorem 2.2. *(Compactness Theorem v.1) If you want something and there's no reason you can't have it, then you can get it.*

(There is some small print; we'll get to that.)

The Compactness Theorem applies to our question about ϵ: let's consider the set, $\{\epsilon < 1/n : n \in \mathbb{Z}^+\} \cup \{\epsilon > 0\}$, of conditions on ϵ. If we take any finitely many of these, then there is a solution in the reals $\mathbb{R}$, so "there's no reason we can't have" ϵ. Admittedly the Completeness Property for the reals does exclude there being a solution in $\mathbb{R}$ but, and this is part of the small print, *we might have to move* to get the "something" in the theorem.

What that means in this example is that there is a structure - a non-standard version, $\mathbb{R}^*$, of the reals - which has an **infinitesimal** (a solution to all those conditions). So, in $\mathbb{R}^*$, we will have $0.\overline{9} < 1 - \epsilon < 1$, giving an alternative to the standard answer to the original question. This structure $\mathbb{R}^*$ will share a great many properties with the reals and it will have a copy of the reals sitting nicely inside it. But it will contain an infinitesimal, with all that that implies. For example, since $\mathbb{R}$ is an ordered field, so also will be $\mathbb{R}^*$, therefore $\mathbb{R}^*$ will contain elements, such as ϵ^{-1}, greater than every integer and it will contain many infinitesimals, $\epsilon + \epsilon$, ϵ^2, ϵ^3, ...

The Compactness Theorem, properly stated, says that if we have a set of conditions (of a certain form: namely which can be expressed by formulas of a first order predicate language appropriate for the structure, $\mathbb{R}$ in this case) such that every finite subset has a solution in some structure M, then there will be an "elementary extension" M^* of M which contains a simultaneous solution to all the conditions. Of course it might be that the original structure M already contains a solution but it may be that, as in our example, the original structure contains no solution and we do have to move to a proper elementary extension.

The definition is: M^* is an **elementary extension** of M (equally, M is an **elementary substructure** of M^*) if, whenever $\varphi(\overline{x})$ is a formula (in n free variables and parameters from M) then the solution set $\varphi(M)$ in M is the intersection of M^n with the solution set $\varphi(M^*)$ in M^*. Even if you don't know what is meant by a formula (with parameters), perhaps this gives some flavour of the idea. Here's a specific example. Can -1 be a square in $\mathbb{R}^*$? No, because, if so, then it would be in the solution set in $\mathbb{R}^*$ of the formula $\exists y\,(y^2 = x)$. So it would be in the intersection of that solution set with $\mathbb{R}$, hence in the solution set of the same formula in $\mathbb{R}$ - which is not the case. Here, $\exists y\,(y^2 = x)$ is an example of a "formula"; an example of a formula with parameters (namely $\pi \in \mathbb{R}$) in this case) is $\quad x^3 = \pi \wedge \forall y\,(y^3 = \pi \to y = x) \quad$ (read $\wedge$ as "and").

Formulas with no free variables are just statements (or "sentences") and if M is an elementary substructure of M^* then M and M^* must satisfy the same sentences. For instance $\mathbb{R}^*$ will be densely ordered because $\mathbb{R}$ satisfies the sentence $\forall y, z\,(y < z \rightarrow \exists w\,(y < w \wedge w < z)$ which expresses the "densely ordered" property.

Precise definitions of "language", "term", "formula", "sentence", "structure" *etc.* are given in the Appendix sections on Predicate Logic. I will explain these briefly but, if they are new to you, you should refer to the Appendix for definitions, explanations and examples.

Rather than quote the Compactness Theorem, we will directly produce an elementary extension of the reals which contains an infinitesimal. We will use the ultraproduct construction.

First we form the direct product $\mathbb{R}^{\mathbb{P}}$ where $\mathbb{P}$ denotes the set of positive integers. This is the set of sequences $(r_i)_{i\in\mathbb{P}}$ with each $r_i \in \mathbb{R}$. We define addition and multiplication on $\mathbb{R}^{\mathbb{P}}$ pointwise: $(r_i)_i + (s_i)_i = (r_i + s_i)_i$ and $(r_i)_i \times (s_i)_i = (r_i \times s_i)_i$. Then these operations make $\mathbb{R}^{\mathbb{P}}$ into a commutative ring with multiplicative identity element $1 = (1_i)_i$ and additive zero element $0 = (0_i)_i$. Notice, *exercise*, that, unlike $\mathbb{R}$, $\mathbb{R}^{\mathbb{P}}$ is not a field.

We are going to factor out a (maximal) ideal so as to obtain a field. We need the following definitions.

Definition 2.3. *A set $\mathcal{F}$ of subsets of $\mathbb{P}$ is a* **filter on** $\mathbb{P}$ *if:*
- $\mathbb{P} \in \mathcal{F}$;
- $\emptyset \notin \mathcal{F}$;
- *if* $J \subseteq K \subseteq \mathbb{P}$ *and* $J \in \mathcal{F}$ *then* $K \in \mathcal{F}$;
- *if* $J, K \in \mathcal{F}$, *then* $J \cap K \in \mathcal{F}$.

Given any such filter $\mathcal{F}$, we define the set $Z_{\mathcal{F}} = \{(r_i)_{i\in\mathbb{P}} : \{i : r_i = 0\} \in \mathcal{F}\}$ - the set of elements of $\mathbb{R}^{\mathbb{P}}$ which are zero on a set of coordinates in $\mathcal{F}$. Then, *exercise*, $Z_{\mathcal{F}}$ is a (proper) ideal of $\mathbb{R}^{\mathbb{P}}$. In order that $\mathbb{R}^{\mathbb{P}}/I_{\mathcal{F}}$ be a field, it is necessary and sufficient that $Z_{\mathcal{F}}$ be a maximal ideal (*exercise* if you haven't seen/don't recall how to prove this fact). That will be the case iff $\mathcal{F}$ is a maximal filter - or "**ultrafilter**", meaning a filter $\mathcal{F}$ on $\mathbb{P}$ such that, if $\mathcal{F}'$ is any filter on $\mathbb{P}$ with $\mathcal{F} \subseteq \mathcal{F}'$, then $\mathcal{F} = \mathcal{F}'$. It is the case that, if $\mathcal{F}$ is any filter on $\mathbb{P}$, then there is an ultrafilter $\mathcal{U}$ on $\mathbb{P}$ with $\mathcal{F} \subseteq \mathcal{U}$. This follows from Zorn's Lemma, which we will discuss later when we look at the ultraproduct construction in the general context.

Here's how we can use this to produce infinitesimals in some extension/enriched version of the reals.

Define the filter $\mathcal{F}$ to consist of all the **cofinite** subsets I of $\mathbb{P}$ (those with finite complement); *exercise*: show that this is indeed a filter on $\mathbb{P}$. (*Another exercise*: show that $\mathcal{F}$ is the smallest filter containing all the subsets of $\mathbb{P}$ of the form $\{m : m \geq n\}$ for some $n \in \mathbb{P}$.) Let $\mathcal{U}$ be any ultrafilter containing $\mathcal{F}$ and consider the ideal $Z_{\mathcal{U}}$ of $\mathbb{R}^{\mathbb{P}}$: $Z_{\mathcal{U}} = \{(r_i)_i : \{i : r_i = 0\} \in \mathcal{U}\}$. Then the quotient ring $\mathbb{R}^{\mathbb{P}}/Z_{\mathcal{U}}$ is a field (*exercise*) - either show this directly or show that $Z_{\mathcal{U}}$ is a maximal ideal) denoted $\mathbb{R}^*$ say. We show that $\mathbb{R}^*$ contains a copy of $\mathbb{R}$.

Consider $\delta : \mathbb{R} \xrightarrow{\Delta} \mathbb{R}^{\mathbb{P}} \xrightarrow{\pi} \mathbb{R}^*$, where Δ is the diagonal embedding, given by taking an element $r \in \mathbb{R}$ to the constant sequence $(r)_i$, and where π is the canonical projection of the ring $\mathbb{R}^{\mathbb{P}}$ to its factor ring $\mathbb{R}^* = \mathbb{R}^{\mathbb{P}}/Z_{\mathcal{U}}$. Then, we claim, δ is an embedding of rings: we must show that $\delta(r) = \delta(s)$ implies that

$r = s$. That follows since $\delta(r) - \delta(s) = 0$ implies $\Delta(r) - \Delta(s) \in \ker(\pi) = Z_{\mathcal{U}}$, hence that $(\Delta(r))_i = (\Delta(s))_i$ for some (and hence for every) coordinate i, hence that $r = s$.

Finally, we show that $\mathbb{R}^*$ contains an infinitesimal: namely the element $\epsilon = \pi((1/i)_i)$. For that, we refer to the orderings on $\mathbb{R}$ and $\mathbb{R}^*$. These can be defined just using the arithmetic operations ($x \leq y$ iff $y - x$ is a square) but we can define the ordering on $\mathbb{R}^*$ directly by setting $\pi((r_i)_i) \leq \pi((s_i)_i)$ iff $\{i : r_i \leq s_i\} \in \mathcal{U}$. It's an *exercise* to show that this is well-defined, meaning independent of choices of representatives of equivalence classes and that it gives a total ordering on $\mathbb{R}^*$; it's a nice additional *exercise* to show then that the ordering defined this way can also be defined algebraically as stated above.

So, that done, let's check that ϵ is, indeed, an infinitesimal. We will identify $\mathbb{R}$ with the copy, $\delta(\mathbb{R})$, of it sitting inside $\mathbb{R}^*$.

Given $n \in \mathbb{P}$, we have that $1/i < 1/n$ for all $i \geq n+1$, hence $\{i \in \mathbb{P} : (1/i)_i < \Delta(1/n)\} \in \mathcal{U}$, so $\epsilon < 1/n$ (where the latter really means $\delta(1/n)$).

Finally $\epsilon \neq 0$ since clearly $\epsilon \notin Z_{\mathcal{U}}$.

That was all rather fast (and particular); in the next few lectures we will go through the ideas and constructions more slowly, and in great generality (but using many examples to illustrate the general ideas).

2.2 Products, filters and ultraproducts

Suppose that I is some index set, and that, for each $i \in I$ we have a structure M_i, where the M_i all are structures of the same kind (all groups, or rings, or partially ordered sets, or...), that is, all $\mathcal{L}$-structures for some language $\mathcal{L}$. Their **product** is, as a set, the product $\prod_{i\in I} M_i$ of their underlying sets. This has, for its elements, the sequences $(a_i)_{i\in I}$ with $a_i \in M_i$. So these are sequences, indexed by I and with the ith coordinate coming from the structure M_i.[1] We will also use the shorter notations $\overline{a} = (a_i) = (a_i)_i$ for such elements.

The $\mathcal{L}$-structure is defined on this set pointwise. For instance, if we have a binary operation, denoted $+$ say, in each structure (more preciseli, let $+_i$ denote the operation on M_i), then we define the operation $+$ on $\prod_i M_i$ by $(a_i)_i + (b_i)_i = (a_i + b_i)_i$ (more precisely, $(a_i +_i b_i)_i$). Here's the general definition.

Given $\mathcal{L}$-structures M_i ($i \in I$) we make the product $\prod_i M_i$ into an $\mathcal{L}$-structure as follows.

• For each constant symbol c in $\mathcal{L}$, we define the interpretation $c^{\prod M_i}$ of c in $\prod M_i$ to be the element $(c^{M_i})_i$.
• Given an n-ary function symbol f in $\mathcal{L}$, we define its interpretation $f^{\prod M_i}$ in $\prod M_i$ to be the function given by: if $a^1, \dots, a^n \in \prod M_i$, with $a^j = (a_i^j)_i$, then $f(a^1, \dots, a^n) = (f^{M_i}(a_i^1, \dots, a_i^n))_i$.[2]
• Given an n-ary relation symbol R in $\mathcal{L}$, we define its interpretation $R^{\prod M_i}$ in $\prod M_i$ to be set of all n-tuples $\overline{a} = (a^1, \dots, a^n) \in (\prod M_i)^n$ such that $R^{M_i}(a_i^1, \dots, a_i^n)$ holds for each $i \in I$ (where, as above, the ith coordinate of a^j is written a_i^j).

[1] Formally, the elements of such a product can be defined to be the functions a from I to $\bigcup_{i\in I} M_i$ such that, for every i, $a(i) \in M_i$.

[2] The notation is useful in that it lets us make a precise and general definition but it obscures the idea, and maybe you need to know the idea in order to make sense of the notation! To understand what is meant, take specific cases, like $n = 1$, $n = 2$ and maybe even a small index set I.

Exercise: to make sense of this, take $\mathcal{L}$ to be the language with one binary relation symbol, written $<$, take $I = \{1, 2\}$, take M_1 and M_2 to be respectively the sets $\{0, 1\}$ and $\{3, 4, 5\}$ both with their natural ordering. Figure out the ordering on the product (draw its Hasse diagram for example).

In this way we turn the product of any set of $\mathcal{L}$-structures into an $\mathcal{L}$-structure. If all the component structures M_i are copies of the same structure M, you can check (*exercise*) that the diagonal embedding $\delta' : M \to M^I$ is an embedding of $\mathcal{L}$-structures.

Let's use the following **running example** in this section: take the index set I to be the set of positive prime integers, and the structure indexed by $p \in I$ to be field $\mathbb{F}_p$ with exactly p elements, that is, the ring of integers modulo p, also written $\mathbb{Z}_p$ or $\mathbb{Z}/p$.

So what does the product construction give in this example? We get the structure $\prod_p \mathbb{Z}_p$ where p ranges over the primes. The *structure* on this is that of a ring: there's an addition and a multiplication, both defined coordinatewise, an identity $1 = (1_p)_p$ for the multiplication and an identity $0 = (0_p)_p$ for the addition, where 1_p denotes the congruence class of $1 \in \mathbb{Z}$ in $\mathbb{F}_p$ and similarly for 0. For example, $1+1 = (0_2, 2_3, 2_5, 2_7, 2_{11}, \dots)$, $1+1+1 = (1_2, 0_3, 3_5, 3_7, 3_{11}, \dots)$, *et cetera* (using a hopefully self-explanatory notation).

This makes the product into a commutative ring. For example, to check commutativity of the multiplication, we compute:
$(a_i)_i \times (b_i)_i = (a_i \times_i b_i)_i$ (by definition of the structure on the product)
$= (b_i \times_i a_i)_i$ (since each component structure is commutative)
$= (b_i)_i \times (a_i)_i$ (by definition).
It is not, however, a field: e.g. $(1, 0, 0, 0, \dots) \times (0, 1, 0, 0, \dots) = (0, 0, 0, 0, \dots)$.

We can, however, get a field from this product if we factor out by a maximal ideal, in other words, if we collapse elements appropriately. By which I mean that we will define an equivalence relation on the product, form the set of equivalence classes and induce a structure on that set (rather as we do in forming $\mathbb{F}_5 = \mathbb{Z}_5$ from $\mathbb{Z}$). Because the additive group structure is there, so we have cosets, it's actually enough to specify which elements get collapsed together with 0 - that is, to specify the ideal we factor out by - but, to better illustrate the general process, we'll not use that fact.

In forming an ultraproduct from a product, the idea is that we collapse (declare to be equivalent) elements which agree on a "large" set of coordinates.

What should we mean by a "large" set of coordinates? To recap: we have an index set I.[3] and, for each $i \in I$, we have an $\mathcal{L}$-structure M_i. We form the product $\mathcal{L}$-structure $\prod_{i \in I} M_i$ as above and we are going to identify/collapse elements which agree on a large set of coordinates. But how do we decide which subsets of I should count as "large"?

Certainly I itself should be a large subset (equal elements should be identified) and the empty set $\emptyset$ should not be large (collapsing all elements together would not give an interesting result). If $J \subseteq I$ is large and $J \subseteq K \subseteq I$ then surely K should also be large. If we're going to identify a and b and also identify b and c then we're going to have to identify a and c ("identification" will be an equivalence relation). If $J = \{i \in I : a_i = b_i\}$ and $K = \{i \in I : b_i = c_i\}$ are the, "large", sets where these pairs of elements agree, then all we can really say

[3] Think of I as being infinite; finite index sets won't give anything new.

about the set of coordinates where $a_i = c_i$ is that it contains $J \cap K$; so it looks as if we should require this set to be large. Let's extract those conditions.

Definition 2.4. *If I is a set then a* **filter** *on I is a collection $\mathcal{F}$ of subsets of I such that:*
- *$I \in \mathcal{F}$;*
- *$\emptyset \notin \mathcal{F}$;*
- *if $J \subseteq K \subseteq I$ and $J \in \mathcal{F}$ then $K \in \mathcal{F}$;*
- *if $J, K \in \mathcal{F}$, then $J \cap K \in \mathcal{F}$.*

Note (*exercise*) that, as a consequence of these clauses, if a subset $J \subseteq I$ is large then its complement $J^{\mathrm{c}} = I \setminus J$ cannot be large.

Given a filter $\mathcal{F}$ on I, we define the corresponding equivalence relation $\sim_{\mathcal{F}}$, or $\sim$ for short, on the product $\prod_{i \in I} M_i$ by $(a_i)_i \sim (b_i)_i$ iff $\{i \in I : a_i = b_i\} \in \mathcal{F}$. Denote by $\prod_{i \in I} M_i / \mathcal{F}$ the set of equivalence classes, writing $a/\sim$ for the equivalence class of an element $a \in \prod_{i \in I} M_i$. We can then turn $\prod_{i \in I} M_i / \mathcal{F}$ into an $\mathcal{L}$-structure, defining operations and relations pointwise but paying attention only to what happens on "large" sets of indices. This structure is called the **reduced product** of the M_i with respect to the filter $\mathcal{F}$. If all the structures M_i are copies of the same structure M then we use the notation $M^I/\mathcal{F}$ and refer to this as a **reduced power** of M.

Before doing the general case carefully, let's do this with the running example, using the filter $\mathcal{F}$ of cofinite subsets of the set I of primes. We define the algebraic operations by setting $\big((a_p)_p/\sim\big) + \big((b_p)_p/\sim\big) = \big((a_p + b_p)_p/\sim\big)$ and $\big((a_p)_p/\sim\big) \times \big((b_p)_p/\sim\big) = \big((a_p \times b_p)_p/\sim\big)$. It has to be checked that this is well-defined (e.g. that if $(a'_p)_p \sim (a_p)_p$ and $(b'_p)_p \sim (b_p)_p$ then $(a'_p + b'_p)_p \sim (a_p + b_p)_p$) but the conditions in the definition of a filter include what we need to do this (*exercise*). We can then check that $(0_p)_p/\sim$ is the zero for addition and $(1_p)_p/\sim$ is the identity for multiplication and, indeed, that all the axioms for a commutative ring are satisfied by $\prod_p \mathbb{F}_p/\mathcal{F}$ (more *exercises*).

Exercise 2.5. Prove that the map $\prod_p \mathbb{F}_p \to \prod_p \mathbb{F}_p/\mathcal{F}$ is a surjective homomorphism of rings and identify its kernel.

So let's do that for general structures M_i, $i \in I$. We're supposing that all these are $\mathcal{L}$-structures for some language $\mathcal{L}$ and we must turn the reduced product into an $\mathcal{L}$-structure. That means that we have to interpret every function, constant and relation symbol of the language. We can do it directly but it's quicker to define it now with reference to the $\mathcal{L}$-structure on $\prod_i M_i$. Recall that $\pi : \prod M_i \to M^* = \prod_{i \in I} M_i / \mathcal{F}$ is the projection map, which takes $a \in \prod M_i$ to its equivalence class $a/\sim$ where $\sim$ means $\sim_{\mathcal{F}}$
- For each constant symbol c in $\mathcal{L}$, we define $c^{M^*} = \pi(c^{\prod M_i})$, that is $(c^{M_i})_i/\sim$.
- Given an n-ary function symbol f in $\mathcal{L}$, we define f^{M^*} to be the n-ary function on M^* given by: if $b^1, \ldots, b^n \in M^*$ then choose, for each $j = 1, \ldots, n$ some $a^j \in \prod M_i$, $a^j = (a_i^j)_i$ say, with $\pi(a^j) = b^j$ and set $f^{M^*}(b^1, \ldots, b^n) = \pi(f^{\prod M_i}(a^1, \ldots, a^n))$. That is, choose a representative in $\prod M_i$ for each equivalence class $b^1, \ldots, b^n$, evaluate the function f on that n-tuple in $\prod M_i$ and then take the $\sim$-equivalence class of the result. Of course it has to be shown that the result is independent of choice of representatives. You should do that as an important *exercise* - important because managing to do it probably means that you have got behind the notation and understood the idea.

• Given an n-ary relation symbol R in $\mathcal{L}$, we define R^{M^*} to be set of all n-tuples $\bar{b} = (b^1, \dots, b^n) \in (\prod M_i)^n$ such that $\{i \in I : (b_i^1, \dots, b_i^n) \in R^{M_i}\} \in \mathcal{F}$.

In this way, we make any reduced product of $\mathcal{L}$-structures into an $\mathcal{L}$-structure. We already saw this construction in our running example and in the ultrapower of the ordered field $\mathbb{R}$, but this, notationally rather unwieldy, definition shows how to do it in general.

Coming back to our running example using the prime fields $\mathbb{F}_p$, we produced a ring $\prod_p \mathbb{F}_p/\mathcal{F}$, where $\mathcal{F}$ is the filter of cofinite sets, but this is still not a field. To see that, split the primes into two infinite disjoint subsets J and K. Define the element a to be 1 on the indices in J and 0 on those in K; define b *vice versa*. Their product is 0 but neither is 0, so this ring is not even an integral domain, let alone a field.

So we need to go further, and impose a further condition on a filter.

Definition 2.6. *An* **ultrafilter** $\mathcal{U}$ *on a set* I *is a filter on* I *which satisfies the further equivalent conditions (we will prove their equivalence):*
• *for each* $J \subseteq I$ *either* $J \in \mathcal{U}$ *or* $J^{\mathrm{c}} = I \setminus J \in \mathcal{U}$*;*
• *if* $J \cup K \in \mathcal{U}$ *then either* $J \in \mathcal{U}$ *or* $K \in \mathcal{U}$*;*
• $\mathcal{U}$ *is a maximal filter (meaning that no collection of subsets of* I *can be a filter and properly include all the sets in* $\mathcal{U}$*).*

So an ultrafilter splits the subsets of I into "large" ones (those in $\mathcal{U}$) and "small" ones (those not in $\mathcal{U}$, equivalently, those whose complement is in $\mathcal{U}$). "Small" does not really mean small (say in the sense of cardinality), just small according to $\mathcal{U}$. But we do have that the union of two "small" sets is still "small" (this is the second of the equivalent conditions above).

In the case where we collapse using an ultrafilter $\mathcal{U}$ rather than any old filter, we refer to the result $\prod_i M_i/\mathcal{U}$ as an **ultraproduct** or, in the case that the M_i all are equal (or isomorphic), an **ultrapower**.

There is one type of ultrafilter that is not interesting. Suppose that $i_0 \in I$. Set $\mathcal{U}(i_0) = \{J \subseteq I : i_0 \in J\}$. Then, *exercise*, $\mathcal{U}(i_0)$ is an ultrafilter, called the **principal ultrafilter generated by** i_0. You can check that, in the ultraproduct $\prod_{i \in I} M_i \,/\mathcal{U}(i_0)$, all that matters is what happens at the coordinate i_0 indeed, another *exercise*, this ultraproduct is isomorphic to M_{i_0}, so we got nothing new from the construction. Therefore we will consider only *non-principal* ultrafilters. But first - are there any?

If I is any infinite set then the collection of all cofinite sets is a filter, sometimes called the **Fréchet filter** $\mathcal{F}_0$. If $\mathcal{U}$ is any ultrafilter containing $\mathcal{F}_0$ then $\mathcal{U}$ cannot be principal, and conversely (quick *exercise*). We do need to call on Zorn's Lemma to give the existence of a maximal=ultra filter containing any given filter but that can be proved (and is an *exercise* for those who have seen Zorn's Lemma).

Let's continue our example using, in place of the cofinite filter $\mathcal{F} = \mathcal{F}_0$, an ultrafilter $\mathcal{U}$ containing $\mathcal{F}$. Let's check that the ultraproduct $\prod_p \mathbb{F}_p/\mathcal{U}$, which equals the quotient ring $\prod_p \mathbb{F}_p/Z_{\mathcal{U}}$, is a field. So take any non-zero element $a/\sim\, = (a_p)_p/\sim\, \neq 0$ in the ultraproduct. Since it is not in the same $\sim$ class as 0, it must be that $J = \{p : a_p = 0\} \notin \mathcal{U}$. Since we have an *ultra*filter, it follows that $J^{\mathrm{c}} = \{p : a_p \neq 0\} \in \mathcal{U}$. But whenever $a_p \neq 0$, a_p has an inverse, b_p say. Set $b = (b_p)_p \in \prod_p \mathbb{F}_p$ (where, if $p \in J$, set $b_p = 0$, say). Then $(a/\sim)\,(b/\sim) = (a_p b_p)/\sim\, = 1$, since the set of coordinates p where $a_p b_p = 1_p$ is in $\mathcal{U}$. So $a/\sim$ has a multiplicative inverse, $b/\sim$, as required.

Exercise 2.7. Consider the product $R = \prod_{i\in I} R_i$ where, for each i in the index set I, R_i is a ring (with 1, commutative if you like). Show that if $\mathcal{F}$ is a filter on I then $Z_{\mathcal{F}} = \{r = (r_i)_i \in R : \{i \in I : r_i = 0\} \in \mathcal{F}\}$ is a (2-sided) ideal of R, Prove that if $\mathcal{F}$, $\mathcal{G}$ are filters on I then $Z_{\mathcal{F}} \subseteq Z_{\mathcal{G}}$ iff $\mathcal{F} \subseteq \mathcal{G}$. Show that if $Z_{\mathcal{F}}$ is a maximal ideal of R then $\mathcal{F}$ must be an ultrafilter on I; is the converse true?

Exercise 2.8. Take the index set I to be the set of positive integers and, for each i, set M_i to be the ring $\mathbb{Z}$ of integers. Let $\mathcal{U}$ be any non-principal ultrafilter on I. Consider the element $p = (p_i)_i/\sim$ where p_i is the ith prime. Show that p is a prime element of $\mathbb{Z}^* = \prod_{i\in I} \mathbb{Z}/\mathcal{U}$ and is not equal to any standard prime (thinking of those as sitting inside the diagonal copy $\delta(\mathbb{Z})$ of $\mathbb{Z}$ in $\mathbb{Z}^*$).

In contrast, let b_i be the product of the first i primes. Show that $b = (b_i)_i/\sim$ has every standard prime as a factor - so this is an element which has infinitely many prime divisors. Show that b also has a prime divisor different from any standard prime.

You might wonder whether every element of $\mathbb{Z}^*$, apart from ± 1, must have at least one prime divisor. The, not so obvious, answer is "yes"; this will follow (*exercise*) directly from Łos' Theorem.

An informal statement of Łos' Theorem is as follows.

Theorem 2.9. *(Łos' Theorem v.1) If all the component structures (or even just a "large" set of them) have a certain property, then their ultraproduct also has that property.*

The small print is that the "property" must be "definable", that is, expressible in terms of the formal language $\mathcal{L}$, where all the component structures M_i are $\mathcal{L}$-structures. Of course, we need to be more precise in the formulation of Łos' Theorem, but let's proceed a little further with this in the specific examples we considered, since the properties clearly will be definable.

For example, this is why the ultraproduct of fields is a field, not just a ring. And also why our running example (the ultraproduct of the prime fields by a non-principal ultrafilter) produces an infinite field: because, given any number N, all but finitely many, hence a "large" set of, components have cardinality greater than N. By Łos' Theorem, the ultraproduct must have cardinality greater than N. That's true for every N, so the ultraproduct is infinite.

Finally, let's come back to infinitesimals. Recall that the requirement on an infinitesimal is that it should be a solution to the set of conditions:
$0 < x$ and, for every positive integer n, $nx < 1$.

In our construction we took our index set I to be the set of positive integers. For each $n \in I$, we took the structure M_n to be the reals $\mathbb{R}$. We called on Zorn's Lemma to get a non-principal ultrafilter $\mathcal{U}$ on I and formed the corresponding ultrapower, $\mathbb{R}^* = \mathbb{R}^I/\mathcal{U}$. We noted that this contains the diagonally-embedded copy of the reals. We then considered the element $\epsilon = \big((1/n)_n/\sim\big) \in \mathbb{R}^*$. Each component is > 0 so, from the way we define the ordering in the ultraproduct, $\epsilon > 0$. Also, given a positive integer n, for all but finitely many i, the ith component of ϵ is $< 1/n$; so, again by the definition of the ordering relation in the ultraproduct, $\epsilon < 1/n$. Before, we saw this by arguing directly but we can see that the fact that ϵ is an infinitesimal is an immediate consequence of Łos' Theorem.

2.3 Definable sets

We have seen that, in an ultraproduct, equations (like those expressing commutativity for the multiplication in a ring) and conditions built from equations (like having an inverse with respect to multiplication) somehow reflect those from the component structures. This is the key to making the precise statement of Łos' Theorem. So we will look at solution sets of equations - these are examples of definable sets. But we will also consider solution sets of more complicated conditions - conjunctions of equations, inequations, even projected (to some components) solution sets of equations. In fact, if we take a structure M and consider the collection of subsets of powers of M which are solution sets of equations (in any finite number of unknowns), and then close under the operations that we just mentioned (forming finite intersections, complements, and projections - but in general repeatedly, to get increasingly complicated sets), then we obtain the definable sets. Let's give a more formal definition.

Definition 2.10. *Let M be a structure. Let $x_1, \dots, x_n$ be variables ("unknowns") and let t_1, t_2 be two terms built up from these variables, using the algebraic operations and also allowing the constants, if there are any, to appear. We write $t(\overline{x})$ to emphasise the variables which may appear. We refer to the expression "$t_1 = t_2$" as an* **equation** *and we define its* **solution set** *to be $\{\overline{a} \in M^n : t_1(\overline{a}) = t_2(\overline{a})\}$.*

Example 2.11. Suppose that we take a field K for our structure (with the ring operations, $+$, $\times$ and constants 0, 1). Then a term built from variables $\overline{x} = (x_1, \dots, x_n)$ is essentially a polynomial, with integer coefficients, in those variables. So, if p, q are such terms/polynomials, then the solution set of the equation $p = q$ is the subset of K^n consisting of all $\overline{a} \in K^n$ where p and q take the same value. Another way of saying this is that the solution set is the zero-set of the polynomial $p - q$.[4]

Definition 2.12. *Suppose first that M is a purely algebraic structure (meaning the "structure" is given by operations and constants - that is, no relation symbols in the language $\mathcal{L}$). The* **definable subsets** *of (the various finite powers of) M are the sets obtained as follows:*
- *the solution set of every equation $t_1 = t_2$ between terms is a definable subset;*
- *the complement, $M^n \setminus D$ of any definable subset D of M^n is definable;*
- *the intersection of any two definable subsets of M^n is definable (therefore, in view of the previous clause, their union also is definable);*
- *if D is a definable subset of M^n and $i \in \{1, \dots, n\}$ then the image of D under projection along the ith axis, that is $\{(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) : \exists a \in M \text{ with } (a_1, \dots, a_{i-1}, a, a_{i+1}, \dots, a_n) \in D\}$, is a definable subset of M^{n-1}.*

If the structure M also has relations then we just add those in at the beginning, along with the solution sets of equations. This does make sense, since an n-ary relation on a set M is, formally, a subset of M^n. For instance, a partial

[4]You may know these already as subvarieties of affine space over K - so these are definable subsets of K or, to say it better, subsets of K^n definable in K. As I've said it, these would just be the subvarieties defined over - that is, zero-sets of polynomials with coefficients in - the prime field, $\mathbb{Q}$ or some $\mathbb{F}_p$. To get more general subvarieties, we should allow elements of K to appear as parameters in our formulas, so that those formulas can refer to polynomials with coefficients in K.

order "<" is treated formally as a set of pairs - exactly those pairs (a, b) with $a < b$. So $\{(a, b) \in M^2 : a < b\}$ would be one of the basic definable subsets.

Every definable set has a definition, namely it is the solution set $\varphi(M)$ of some formula φ of a language $\mathcal{L}$ such that M is an $\mathcal{L}$-structure. The atomic (=most basic) formulas are those of the form $t(\overline{x}) = t'(\overline{x})$ where t and t' are terms, together with those of the form $R(\overline{x})$ where R is some relation symbol of $\mathcal{L}$. The other formulas are built up using the boolean operations of negation and conjunction, together with the prefixing of existential quantifiers - these exactly correspond to the operations on definable sets of taking the complement, forming the intersection and projecting along a coordinate. (Recall that disjunction, implication and universal quantification can be defined using the other operations so, although we use them, we treat them formally as defined in terms of the others - this is useful when proving things by induction on complexity of formulas.)

Exercise: convince yourself, through a variety of examples, that if D is a definable subset of a structure M then there is a formula φ, with free = unquantified variables (among) $x_1, \ldots, x_n$, such that D is the solution set in M of φ. Part of the point of using formulas like φ is that they can be applied to any $\mathcal{L}$-structure, not just the structure we started with, and its meaning will be 'the same', in the sense that it expresses the same property (though, of course, its solutions will be very different).

So, from now on we will use formulas when we talk about definable sets in more than one structure. For instance, the centre of a group is defined by the formula $\forall x_2 \, (x_1 * x_2 = x_2 * x_1)$, where $*$ denotes the operation in the group. So the centre of any group G is a definable subset of G but the formula $\varphi(x_1)$ that we just wrote down can be read in any group H and its solution set, which we denote $\varphi(H)$, is exactly the centre of that group.

2.4 Łos' Theorem

Łos' Theorem is about what is true in an ultraproduct M^*; more precisely it tells us what the $\mathcal{L}$-theory of M^* is and what are its definable subsets. It says that an element $a = (a_i)_i/\sim$ of an ultraproduct belongs to a definable subset $\varphi(M^*)$ iff, for a "large" set of indices i, the component a_i belongs to the corresponding definable subset $\varphi(M_i)$ of M_i. It says, furthermore, that an ultraproduct has a property which can be expressed by a sentence σ of $\mathcal{L}$ iff "most" of its coordinate structures have that property (that is, satisfy σ). Here "large" and "most" mean with respect to the ultrafilter.

Theorem 2.13. *(Łos' Theorem v.2.1) Suppose that $M^* = \prod_{i \in I} M_i/\mathcal{U}$ is an ultraproduct. Suppose that φ is a formula (with free variables $x_1, \ldots, x_n$). Then $\overline{a} = (a^1, \ldots, a^n)$ is in the solution set, $\varphi(M^*)$, of φ in M^* iff $\{i \in I : \overline{a_i} \in \varphi(M_i)\}$ is in $\mathcal{U}$, where $a^j = (a_i^j)_i/\sim$ and $\overline{a_i} = (a_i^1, \ldots, a_i^n)$.*

Proof. So suppose that $M^* = \prod_{i \in I} M_i/\mathcal{U}$ is an ultraproduct.

The assertion is that, given a formula ψ,
$(*)$ for all $\overline{a}$, we have $\overline{a} \in \psi(M^*)$ iff $\{i \in I : \overline{a_i} \in \psi(M_i)\} \in \mathcal{U}$
where
$\overline{x} = (x_1, \ldots, x_n)$, $\overline{a} = (a^1, \ldots, a^n) \in (M^*)^n$, $a^j = (a_i^j)_i/\sim$ and $\overline{a_i} = (a_i^1, \ldots, a_i^n)$.

This is proved by induction on the complexity of ψ. This is "complexity" in the sense that atomic formulas are the least complex formulas (so the starting point of the induction) and then "complexity" is increased each time we apply a boolean operation ("and", "or", "not", "implies") or a quantifier ("there exists", or "for all"). Because "not", "and" and "there exists" are enough to define the others, we only need to consider those. So we have the base cases - where ψ is an equation or a relation - and three types of induction step. Here are the statements that, therefore, have to be proved.

If t and t' are terms built from variables $\overline{x} = (x_1, \dots, x_n)$ (and perhaps constants) and $\overline{a} = (a^1, \dots, a^n) \in (M^*)^n$ where $a^k = (a^k_i)_i/\sim$, then $t(\overline{a}) = t'(\overline{a})$ iff $\{i \in I : t(\overline{a^j_i}) = t'(\overline{a^j_i})\} \in \mathcal{U}$. This statement, in turn, has to be proved by induction on complexity of terms (how they are built up from the variables and constant symbols by successively applying function symbols). I will do some, maybe all, the details of this in class.

The other base case is that of a basic relation $R(x_1, \dots x_n)$ and we need the statement that, with notation $\overline{a}$ etc. as above, $R(a^1, \dots, a^n)$ holds in M^* iff $\{i \in I : R(\overline{a_i}) \text{ holds in } M_i\} \in \mathcal{U}$. But this is how we defined the $\mathcal{L}$-structure on M^* so (assuming we already proved the well-definedness of this), there is nothing to do here.

That's the base case; the induction steps have the following (three) forms.

If ψ and ψ' are formulas and if each of these satisfies $(*)$, then so does the conjunction $\psi \wedge \psi'$ [the proof uses the closure of $\mathcal{U}$ under intersections].

If ψ is a formula which satisfies $(*)$, then so does the negation $\neg\psi$ (this proof of this uses that $\mathcal{U}$ is actually an *ultra*filter).

These two cases are very straightforward. I'll do the third here.

If ψ is a formula which satisfies $(*)$ and y is a variable then $\exists y\, \psi$ satisfies $(*)$ (it doesn't matter whether or not y actually appears in ψ, though it's rather pointless to stick $\exists y$ in front if y doesn't appear in ψ). Let's look at this one more closely (you might guess that this will use the "closed upwards" property of filters; let's see). Suppose then that $\overline{a} \in (\exists y\, \psi(\overline{x}, y))(M^*)$. Then there is $b \in M^*$ such that $(\overline{a}, b) \in \psi(M^*)$. So, by the induction hypothesis, $\{i \in I : (\overline{a}_i, b_i) \in \psi(M_i)\} \in \mathcal{U}$. Now, this set is certainly contained in $\{i \in I : \overline{a} \in (\exists y\, \psi(\overline{x}, y))(M_i)\}$, so this set is in $\mathcal{U}$, proving one direction of $(*)$ (and, indeed, using the upwards-closed property). We must check the other direction.

So suppose $K = \{i \in I : \overline{a} \in (\exists y\, \psi(\overline{x}, y))(M_i)\} \in \mathcal{U}$. For each index i in this set, choose a "witness" to the existential quantifier. That is, choose some $c_i \in M_i$ such that $(\overline{a}_i, c_i) \in \psi(M_i)$. Define the element $c \in M^*$ to be $(c_i)_i/\sim$ where, for $i \in I \setminus K$, you can choose c_i to be any element in M_i (since $I \setminus K$ is a "small" set, it doesn't matter what happens on any of those components). Then, $\{i \in I : (\overline{a}_i, c_i) \in \psi(M_i)\} \in \mathcal{U}$ and so, by the inductive hypothesis, $(\overline{a}, c) \in \psi(M^*)$. Therefore $\overline{a} \in (\exists y\, \psi(\overline{x}, y))(M^*)$, as required. $\square$

Note that definable subsets of $\mathbb{R}$ include solution sets to equations and inequalities, so each of the requirements that characterise an infinitesimal can be expressed by a suitable formula - say in the language of ordered rings. And any finitely many of these requirements can be satisfied in $\mathbb{R}$ (that is, the intersection of the corresponding finitely many definable sets is nonempty). So the construction that we gave earlier is just a special case of Łos' Theorem.

Notice that, if $\mathcal{U}$ is a non-principal ultrafilter, hence every cofinite subset of the index set is in $\mathcal{U}$, an element $a = (a_i)_i/\sim$ of the ultraproduct will have a property which can be expressed by a formula $\varphi(x)$ if some $\sim_{\mathcal{U}}$-representative[5] of a has all but finitely many of its components satisfying that property. Although this condition is sufficient, it is by no means necessary for a to have the given property but I mention it explicitly because it often is used in particular cases.

We're not done with stating Łos' Theorem yet. The version above applies to properties of elements and n-tuples, but what about properties of the whole structure? For instance, the property that a commutative ring is a field. That particular property can be expressed by a formula, $\forall x\,(x \neq 0 \to \exists y\,(xy = 1))$, which has no free variables, that is, by a **sentence**. Sentences express properties of structures, rather than properties of elements. Here (in part (b)) is the version of Łos' Theorem that applies to them. Part (a) is a restatement of the previous version for the case of a single element rather than an n-tuple (less notation being necessary, perhaps the meaning is clearer).

Theorem 2.14. *(Łos' Theorem v.2.2) Suppose that $M^* = \prod_{i\in I} M_i/\mathcal{U}$ is an ultraproduct of $\mathcal{L}$-structures.*
(a) Suppose that $\varphi(x)$ is a formula of $\mathcal{L}$. Then $a = (a_i)_i/\sim$ is in the solution set, $\varphi(M^*)$, of φ in M^* iff $\{i \in I : a_i \in \varphi(M_i)\}$ is in $\mathcal{U}$.
(b) Suppose that σ is a sentence of $\mathcal{L}$. Then σ is true in M^* iff $\{i \in I : \sigma \text{ is true in } M_i\}$ is in $\mathcal{U}$.

Proof. Part (a) is the case $n = 1$ of the version above. But, for part (b) just notice that since a sentence is a special case of a formula, this is already covered by (a).
[If the argument for (b) seems unconvincing, then do the following *exercise*. Consider the case that σ has the form $\forall x\,\varphi(x)$ and suppose that $M^* \models \sigma$ (recall that $\models$ is the notation for the satisfaction/"true in" relation between structures and sentences/formulas with parameters). This means that the definable set $\varphi(M)$ is all of M. Assume the statement of part (a) for the formula φ and deduce that the set of indices i such that $M_i \models \forall x\,\varphi(x)$ is in $\mathcal{U}$. Be careful in your argument. You could also prove the converse, assuming $\{i \in I : M_i \models \sigma\} \in \mathcal{U}$, and deducing $M^* \models \sigma$.] $\square$

Here is the previous version written using the compact notation $\models$.

Theorem 2.15. *(Łos' Theorem v.2.2.5) Suppose that $M^* = \prod_{i\in I} M_i/\mathcal{U}$ is an ultraproduct of $\mathcal{L}$-structures.*
(a) Suppose that $\varphi(x)$ is a formula of $\mathcal{L}$ and let $a = (a_i)_i/\sim \in M^*$. Then $M^* \models \varphi(a)$ iff $\{i \in I : M_i \models \varphi(a_i)\} \in \mathcal{U}$.
(b) Suppose that σ is a sentence of $\mathcal{L}$. Then $M^* \models \sigma$ iff $\{i \in I : M_i \models \sigma\} \in \mathcal{U}$.

To emphasise: "formula" and "sentence" have very precise meanings here - they refer to formulas constructed from the basic algebraic relations, constants, and relations which give meaning to the phrase "type of structure" and where "constructed" means constructed using the boolean operations and quantifiers (the operations that we use when constructing definable sets). Of course, there has to be some kind of restriction: we know, for instance, that the condition "there are only finitely many elements" cannot be expressed by such a sentence, otherwise we could make an ultraproduct which would contradict Łos' Theorem.

[5] meaning an element $(a'_i)_i$ in the product such that $(a'_i)_i/\sim = a$

(On the other hand, saying "there are no more than N elements" is, for any particular natural number N, certainly expressible in any language - all we need is equality to say that.)

2.5 The Compactness Theorem

Now we derive the Compactness Theorem from Łos' Theorem.

Theorem 2.16. *(Compactness theorem, v.2a) Suppose that we have a set T of sentences in a language $\mathcal{L}$ appropriate for some specific type of structure. Suppose that, for every finite subset S of T, there is an $\mathcal{L}$-structure M_S which satisfies all the sentences in S. Then there is an ultraproduct of the M_S which satisfies all the sentences in T.*

Proof. We take the index set I to be the set of finite subsets S of T, with the structure being indexed by S being (some chosen) M_S. Any old ultrafilter won't do, so we first set up a filter (really a 'filter basis') by taking, for each $S \in I$, the subset $\langle S \rangle = \{S' \in I : S \subseteq S'\}$ of I. Note that if $S, S' \in I$ then $\langle S \rangle \cap \langle S' \rangle \supseteq \langle S \cup S' \rangle$.

It follows (see the *exercise* below) that the set $\mathcal{F} = \{J \subseteq I : J \supseteq \langle S \rangle$ for some $S \in I\}$ of subsets of I which contain some set of the form $\langle S \rangle$, is a filter. We then call on Zorn's Lemma to bring an ultrafilter $\mathcal{U}$ on I (necessarily non-principal) containing $\mathcal{F}$. Form the ultraproduct $M^* = \prod_{S \in I} M_S / \mathcal{U}$. I claim that this does the job.

So take any sentence $\sigma \in T$. Then $\{\sigma\} \in I$, so $\langle \{\sigma\} \rangle = \{S \in I : \sigma \in S\}$ is in the filter base, hence in $\mathcal{F}$, hence in $\mathcal{U}$. Note that if $S \in \langle \{\sigma\} \rangle$ then M_S satisfies σ. Therefore the set of indices S where the structure M_S satisfies σ is in $\mathcal{U}$ and hence, by Łos' Theorem, M^* satisfies σ. As required. $\square$

Exercise 2.17. If I is a(n index) set then a set $\mathcal{B}$ of subsets of I is a **filter basis** if the intersection of any finitely many members of $\mathcal{B}$ is nonempty - we say that $\mathcal{B}$ has the **finite intersection property** (**fip** for short). Show that $\mathcal{F}_{\mathcal{B}} = \{J \subseteq I : J \supseteq K_1 \cap \cdots \cap K_n$ for some $K_1, \ldots, K_n \in \mathcal{B}\}$ is a filter on I and is the smallest filter on I which contains every set in $\mathcal{B}$. It is called the filter **generated** by $\mathcal{B}$.

Theorem 2.18. *(Compactness theorem, v.2b) Suppose that M is an $\mathcal{L}$-structure and that Φ is a set of formulas of $\mathcal{L}$ with free variables (among) $\overline{x} = (x_1, \ldots, x_n)$. Suppose that, for every finite subset S of Φ, there is $\overline{a} \in M^n$ which satisfies all the formulas in S. Then there is an ultrapower M^* of M and $\overline{c} \in (M^*)^n$ which satisfies all the formulas in Φ.*

Proof. The proof is quite similar to that above. In fact, it can be made into a special case by introducing n new constant symbols to replace the variables $x_1, \ldots, x_n$, so that a formula $\varphi(\overline{x})$ can be replaced by a sentence in this slightly enriched language, thus replacing Φ by a set of sentences, which can then be fed into the version above. I'll go through the details in the lecture. $\square$

Definition 2.19. *A set T of sentences (of some language $\mathcal{L}$) is* **finitely satisfiable** *if every finite subset of T has a model; the set T is* **satisfiable** *if it has a model, that is, if there is an $\mathcal{L}$-structure $\mathcal{M}$ such that $\mathcal{M} \models \sigma$ for every $\sigma \in T$.*

Theorem 2.20. *(Compactness Theorem, v2.a.5) If a set of sentences of a language $\mathcal{L}$ is finitely satisfiable then it is satisfiable.*

Here is a corollary of the Compactness Theorem, stated more in the style of the next section, where we consider theories and their models.

Corollary 2.21. *Suppose that T is a set of sentences with arbitrarily large finite models. Then T has an infinite model*

Proof. Our assumption is that, for every n, there is a finite model of T with $\geq n$ elements. So consider the set $T \cup \{\sigma_{\geq n} : n \geq 1\}$ (where $\sigma_{\geq n}$ is a sentence saying there are at least n distinct elements). By our assumption this is finitely satisfiable, hence by Compactness it is satisfiable, meaning that there is an infinite model of T. $\square$

Exercise 2.22. Show that if a definable subset $\varphi(\mathcal{M}) \subseteq M$ of an $\mathcal{L}$-structure $\mathcal{M}$ is infinite then there is an elementary extension $\mathcal{M}' \succ \mathcal{M}$ of $\mathcal{M}$ which contains an element $a \in \varphi(\mathcal{M}') \setminus M$. Deduce that every infinite $\mathcal{L}$-structure has a proper elementary extension.

Exercise 2.23. Let p be a prime integer and let G be the multiplicative group of complex p^nth roots of unity (the union over all n). Note that every element of G is torsion. Show that there is an elementary extension of G which contains an element which is not torsion.

3 Theories and Models

Definition 3.1. *Suppose that $\mathcal{L}$ is some language of the kind that we have been considering. A* **theory** *in $\mathcal{L}$, or $\mathcal{L}$-***theory***, is a set, T, of sentences of $\mathcal{L}$. A* **model** *of an $\mathcal{L}$-theory T is an $\mathcal{L}$-structure $\mathcal{M}$ such that $\mathcal{M}$ satisfies every sentence in T: $\mathcal{M} \models \sigma$ for every $\sigma \in T$; we write $\mathcal{M} \models T$. If T is a set of sentences of $\mathcal{L}$, we let $\mathrm{Mod}(T)$ denote the collection of all models of T.*

Usually we include the requirement that, in order to be allowed as a theory, a set T of sentences must be **consistent**, in the sense that it has some model.

Definition 3.2. *The* **deductive closure** *of T is the set $\overrightarrow{T}$ of all sentences σ which are true in every model of T; we could write $\overrightarrow{T} = \mathrm{Th}(\mathrm{Mod}(T))$ (cf. 3.5 below). Note, exercise, that $\mathrm{Mod}(\overrightarrow{T}) = \mathrm{Mod}(T)$.*

Corollary 3.3. *(of the Compactness Theorem) If T is an $\mathcal{L}$-theory and σ is a sentence of $\mathcal{L}$ in the deductive closure of T, then there is a finite subset T' of T such that σ is in the deductive closure of T'.*

Proof. *Exercise* □

In practice we often blur the distinction between a theory and its deductive closure, so "the theory of groups" could mean just some choice of axioms for groups or it could mean everything that follows from those, that is, every sentence (in the chosen language) which is true in every group; officially, its the latter.

The terms "consistent" and "deductive closure" suggest the notion of formal deduction and you may know that the Completeness Theorem for Predicate Calculus implies that a theory is consistent iff there is no contradiction deducible from it and that $\overrightarrow{T}$ is the set of sentences of $\mathcal{L}$ formally deducible from T. We are not going to consider formal deductive systems in this course and, instead, we have defined these notions "semantically" - by reference to truth in models (the Completeness Theorem for Predicate Logic says that the two approaches give the same result).

Definition 3.4. *A theory T is* **complete** *if for every sentence $\sigma \in \mathcal{L}$ either $\sigma \in \overrightarrow{T}$ or $\neg\sigma \in \overrightarrow{T}$.*

Definition 3.5. *If $\mathcal{M}$ is an $\mathcal{L}$-structure then the (***complete***)* **theory of** *$\mathcal{M}$ is the set of all sentences of $\mathcal{L}$ which are true in $\mathcal{M}$:*

$$\mathrm{Th}(\mathcal{M}) = \{\sigma \in \mathcal{L} : \sigma \textit{ is a sentence and } \mathcal{M} \models \sigma\}.$$

Corollary 3.6. *If T is an $\mathcal{L}$-theory and σ is a sentence of $\mathcal{L}$ in the deductive closure of T, then there is a finite subset T' of T such that σ is in the deductive closure of T'.*

Proof. *Exercise* □

The theory of $\mathcal{M}$ is indeed complete because, if σ is any sentence then, either $\mathcal{M} \models \sigma$ (so $\sigma \in \mathrm{Th}(\mathcal{M})$) or, if not, that is if $\mathcal{M} \not\models \sigma$, then $\mathcal{M} \models \neg\sigma$ (so $\neg\sigma \in \mathrm{Th}(\mathcal{M})$).

Definition 3.7. *Two $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{M}'$ are* **elementarily equivalent** *if they satisfy exactly the same sentences of $\mathcal{L}$: we then write $\mathcal{M} \equiv \mathcal{M}'$.*

It is immediate from the definitions that $\mathcal{M} \equiv \mathcal{M}'$ iff $\mathrm{Th}(\mathcal{M}) = \mathrm{Th}(\mathcal{M}')$. Recall (3.14) that if structures are isomorphic then they are elementarily equivalent.

Lemma 3.8. *The following are equivalent for a consistent theory T:*
(i) T is complete;
(ii) $\mathcal{M}, \mathcal{M}' \models T$ implies $\mathcal{M} \equiv \mathcal{M}'$;
(iii) $\overrightarrow{T} = \mathrm{Th}(\mathcal{M})$ for some $\mathcal{L}$-structure $\mathcal{M}$.

Proof. Direct from the definitions. □

A complete theory is the theory of a single structure; this is not to say that a complete theory has only one model!! In fact, most complete theories have lots of models, for example, ultrapowers of any given infinite model. But there is one exception, as follows.

Proposition 3.9. *If T is complete and has a finite model then T has just one model up to isomorphism.*

Proof. (Outline) Say $\mathcal{M} \models T$ has exactly n elements. Then $\sigma_{=n}$, a sentence saying there are exactly n elements, is in T. Therefore every model of T has exactly n elements. But we have to show that $T = \mathrm{Th}(\mathcal{M})$ specifies everything about $\mathcal{M}$, not just how many elements it has. We do this by taking any $\mathcal{N} \models T$ (so N must have n elements) and then we show that one of the $n!$ bijections between M and N must be an isomorphism - otherwise we produce a sentence true in one of $\mathcal{M}$, $\mathcal{N}$ but false in the other - contradicting completeness of T. □

Let T be a theory in some language $\mathcal{L}$ and suppose that we want to understand the models of T. First we may reduce to the case where T is complete.

Definition 3.10. *Say that T' is a* **completion** *of T if $T \subseteq T'$ (so any model of T' is a model of T) and if T' is complete.*

Note that if $M \models T$ then $\mathrm{Th}(M)$ is a completion of T. Therefore we can express $\mathrm{Mod}(T)$ as the disjoint (why?) union of the classes $\mathrm{Mod}(T')$ where T' runs over the distinct completions of T.

Exercise 3.11. Prove that if T is not complete then it has at least two different completions.

Exercise 3.12. (if you know the relevant cardinal arithmetic from Set Theory) Suppose that T is a theory in a countable language. Show that T has at most $2^{\aleph_0}$ completions.

Exercise 3.13. For each $\kappa \in \{1, 2, \ldots, n, \ldots \aleph_0, 2^{\aleph_0}\}$ give an example of a theory T_κ in a countable language which has exactly κ completions (in fact, these are the only possibilities for the number of completions of a theory in a countable language).

Understanding the class of models of a theory is one theme that runs through model theory. We will prove the Downwards Löwenheim-Skolem Theorem, which produces "small" models, for example saying that if the language $\mathcal{L}$ is countable then every infinite $\mathcal{L}$-structure contains a countable (necessarily infinite) elementary substructure. Later we prove the Upwards Löwenheim-Skolem Theorem, which says that any infinite model (of cardinality at least that of the language) has an elementary extension of any prescribed larger cardinality.

It is a consequence of those theorems that if $\mathcal{L}$ is countable then any $\mathcal{L}$-theory with an infinite model has at least one model of every infinite cardinality; in particular, we can't completely specify an infinite structure using first order logic. So the most we can ask of a theory T is that there be, for each infinite cardinality κ, just one model of cardinality κ up to isomorphism. That property of T is referred to as *categoricity.* A weakening of the condition of categoricity is that there be just one countably infinite model (up to isomorphism) - that is $\aleph_0$-categoricity and we will prove some nice characterisations of such theories. Alternatively we could ask for categoricity in higher, uncountable, cardinalities. Morley (in the 60's) proved a remarkable and extremely influential theorem in model theory that a theory in a countable language which is categorical in one uncountable cardinality is categorical in *all* uncountable cardinalities. We won't get to proving that but the upshot is that, for theories in countable languages, categoricity comes in three varieties: countable categoricity; uncountable categoricity; total categoricity (i.e. both the others).

3.1 Building around a set; the Downwards Löwenheim-Skolem theorem

The simplest example of "building around a set" is how we generate a substructure, $\langle A\rangle$, of an $\mathcal{L}$-structure $\mathcal{M}$ from any subset A of M. It is more difficult to build an elementary substructure of $\mathcal{M}$ containing A (and we shouldn't expect the uniqueness/minimality that we got for $\langle A\rangle$) but it can be done.

Theorem 3.14. *(Downwards Löwenheim-Skolem) Suppose that $\mathcal{M}$ is an infinite $\mathcal{L}$-structure, where $\mathcal{L}$ is a countable language, and that A is a countable subset of M. Then there is a countable elementary substructure $\mathcal{N}$ of $\mathcal{M}$ with $A \subseteq N$.*

Proof. We have to produce a subset of M which contains A and which is (the underlying set of) an elementary substructure of $\mathcal{M}$ (we'll worry about its countability later). Certainly any such subset must be a substructure, so we have to include, in addition to A, any interpretations of constant symbols and then "close under the ($\mathcal{L}-$)functions". Doing that would give us a substructure of $\mathcal{M}$ (we proved this in 10.2) but not necessarily an elementary substructure of $\mathcal{M}$. But Tarski's Lemma 10.17 says that a substructure is an elementary one if it contains "witnesses for existential quantifiers". So we list all formulas $\varphi(x)$ with parameters from $A \cup \{\text{constants}\}$ and then, for each φ with $\mathcal{M} \models \exists x\varphi(x)$ we choose some element $b \in M$ with $\mathcal{M} \models \varphi(b)$ (choose just one for each such φ). Add in all these "witnesses" b to get a new set $A \cup \{\text{witnesses}\}$. Now we realise that we have to repeat the process (because we have added new parameters). In fact we have to repeat the process ω-many times but no more (since a formula is a finite object). We end up with a subset of M which satisfies the criterion of Tarski's Lemma for being an elementary substructure of $\mathcal{M}$. Cardinal arithmetic and the countability assumptions we made ensure that this set is indeed countable (and necessarily infinite since it is elementarily equivalent to $\mathcal{M}$).

That was the recipe; let's make the proof now.

Set $A_0 = A$. For each formula $\varphi(x)$, with free variable x and parameters in A_0, such that $\mathcal{M} \models \exists x\, \varphi(x)$ choose some $b_\varphi \in M$ such that $\mathcal{M} \models \varphi(b_\varphi)$. Set $A_1 = A_0 \cup \{b_\varphi\}_\varphi$ (φ running over such formulas). Note that A_1 contains the interpretations in $\mathcal{M}$ of each constant symbol c (take φ to be the formula $x = c$) and, if f is an n-ary function symbol and $a_1, \ldots, a_n \in A_0$ then $f^{\mathcal{M}}(a_1, \ldots, a_n) \in A_1$ (consider the formula with parameters in A_0, $x = f(a_1, \ldots, a_n)$).

We repeat this process: having defined A_i, for every formula $\varphi(x)$, with free variable x and parameters in A_i, such that $\mathcal{M} \models \exists x\, \varphi(x)$ choose some $b_\varphi \in M$ such that $\mathcal{M} \models \varphi(b_\varphi)$. Then set $A_{i+1} = A_i \cup \{b_\varphi\}_\varphi$ (φ running over such formulas); in particular $A_i \subseteq A_{i+1}$.

Set $A_\omega = \bigcup_{i \geq 0} A_i$. We will show that we can take $N = A_\omega$ - that A_ω is the underlying set of an elementary substructure of $\mathcal{M}$. First we show it's the underlying set of a substructure.

We've already seen that A_1, hence A_ω contains all the $c^{\mathcal{M}}$ - the first requirement in 10.1 - so we check the second requirement of that lemma. Let f be an n-ary function symbol of $\mathcal{L}$ and $a_1, \ldots, a_n \in A_\omega$. For each i, $a_i \in A_{m(i)}$ for some $m(i)$. Let $m = \max\{m(1), \ldots, m(n)\}$ - so, since $A_i \subseteq A_j$ if $i \leq j$, each a_i is in A_m. Then, by construction (and as already argued above), $f^{\mathcal{M}}(a_1, \ldots, a_n) \in A_{m+1}$ and hence $f^{\mathcal{M}}(a_1, \ldots, a_n) \in A_\omega$. Therefore A_ω is indeed the underlying set of a substructure, we label it $\mathcal{N}$, of $\mathcal{M}$. To show that it is an elementary substructure, we check the criterion of 10.17.

So let $a_1, \ldots, a_n \in A_\omega$ and suppose that $\psi(y_1, \ldots, y_n, x)$ is such that $\mathcal{M} \models \exists x\, \psi(a_1, \ldots, a_n, x)$. Choose the $m(i)$ and m as we did in the previous paragraph. Then, by construction (with $\varphi(x)$ being $\psi(a_1, \ldots, a_n, x)$), there is $b_\varphi \in A_{m+1}$, hence $b_\varphi \in A_\omega$, such that $\mathcal{M} \models \varphi(b_\varphi)$, that is, $\mathcal{M} \models \psi(a_1, \ldots, a_n, b_\varphi)$. So, by 10.17, $\mathcal{N} \prec \mathcal{M}$.

Clearly $A \subseteq N = A_\omega$ so it remains to count the elements of A_ω. We claim that, for each i, A_i is countable; we prove this by induction. The base case is the assumption that $A_0 = A$ is countable. So suppose inductively that A_i is countable. There are only countably many formulas $\psi(y_1, \ldots, y_n, x)$ (since $\mathcal{L}$ is countable) and, for each of these, only countably many choices for replacing y_i by an element of A_i (since A_i is countable). Hence there are only countably many formulas with parameters from A_i, hence (since we add at most one witness for each), only countably many elements b_φ added to A_i to get A_{i+1}. Finally, a countable union (A_ω) of countable sets is countable. And we're done. $\square$

Corollary 3.15. *Suppose that $\mathcal{L}$ is a countable language and that $\mathcal{M}$ is an infinite $\mathcal{L}$-structure. Then $\mathcal{M}$ has a countable elementary substructure.*

Corollary 3.16. *Suppose that $\mathcal{L}$ is a countable language and T is a consistent $\mathcal{L}$-theory. Then T has a countable model.*

A little cardinal arithmetic

If X and Y are sets and there is a bijection from X to Y then we write $|X| = |Y|$ and say that X and Y have **the same cardinality** (or **the same size**). This is an equivalence relation (and you might note that this is just isomorphism as $\mathcal{L}_0$-structures).

You might already have seen the proofs that $|2\mathbb{N}| = |\mathbb{N}|$, that $|\mathbb{N}| = |\mathbb{Z}|$ and that $|\mathbb{N}| = |\mathbb{Q}|$ - all these sets are **countably infinite** but $|\mathbb{N}| \neq |\mathbb{R}|$ (by Cantor's Diagonal Argument).

Just as in the finite case, it is useful to abstract actual numbers from classes of sets all of the same size. So we set $\aleph_0$ to be the "number of elements in any countably infinite set - this is the smallest infinite number - actually we call these **cardinal numbers** and refer to the **cardinality** of a set.

We can order cardinal numbers: if $\kappa = |X|$ and $\lambda = |Y|$, then we set $\kappa \leq \lambda$ if there is an injection from X to Y. The Cantor-Schröder-Bernstein Theorem says that if there is an injection from X to Y and an injection from Y to X, then there is a bijection between X and Y. That is, if $\kappa \leq \lambda$ and $\lambda \leq \kappa$, then $\lambda = \kappa$. So we get an ordering on the collection of cardinal numbers.

A generalisation of Cantor's Diagonal Argument shows that for every cardinal κ, we have $2^\kappa > \kappa$ (so there is no greatest cardinal number).

Also, assuming the Axiom of Choice, one can prove that the cardinal numbers are totally ordered: if κ and λ are cardinal numbers then either $\kappa \leq \lambda$ or $\lambda \leq \kappa$. Indeed, the collection of cardinal numbers is **well-ordered**, meaning that every non-empty set of cardinal numbers has a least element; it follows that, given κ, there is a least cardinal strictly greater than κ - called the **successor** of κ and written κ^+. We set $\aleph_1 = \aleph_0^+$, $\aleph_2 = \aleph_1^+$, ...

Cardinal arithmetic We define addition, multiplication and exponentiation of cardinals in terms of representative sets - of course one then has to show the definitions are independent of choice of representatives, but that's not hard.

Given cardinals κ, λ, choose sets X, Y with $|X| = \kappa$ and $|Y| = \lambda$. Define $\kappa + \lambda = |X \cup Y|$ provided X and Y are disjoint (which can be arranged by rechoosing one of them if necessary). Also define $\kappa \times \lambda = |X \times Y|$ (the cartesian product of X and Y) and $\kappa^\lambda = |X^Y|$ where by X^Y we mean the set of all functions from Y to X. (You should check that these do actually give the right numbers when κ and λ are finite.) Then we can prove the following arithmetic rules:

$\kappa + \lambda = \lambda + \kappa$, $\kappa \times \lambda = \lambda \times \kappa$, $\kappa \times (\lambda \times \mu) = (\kappa \times \lambda) \times \mu$, $\kappa \times (\lambda + \mu) = (\kappa \times \lambda) + (\kappa \times \mu)$, *et cetera.*

Also if κ is infinite then $\kappa + \kappa = \kappa$ and $\kappa \times \kappa = \kappa$.

And, more generally, if $\kappa \leq \lambda$ then $\kappa + \lambda = \lambda$ and $\kappa \times \lambda = \lambda$.
For example, the fact that $\aleph_0 \times \aleph_0 = \aleph_0$ implies that a countable union of countable sets is countable.

4 Building up maps: Back and Forth between Dense Linear Orders; the Random Graph

4.1 Back and Forth

The idea is nicely illustrated in the proof, below, of $\aleph_0$-categoricity of the theory of dense linear orders without endpoints: there one shows that any two countable models of that theory are isomorphic. The general shape is as follows.

Suppose that we start with two countably infinite structures $\mathcal{M}$, $\mathcal{N}$. We enumerate M as $a_1, \ldots, a_k, \ldots$ and enumerate N as $b_1, \ldots, b_k, \ldots$

The aim is to try to build up an isomorphism from $\mathcal{M}$ to $\mathcal{N}$, so we have to decide where this isomorphism is going to send a_1 - it has to be sent to an element of N with "the same properties". The way we find such an element in N depends on the particular situation, and it might or might not be possible. But, if we can do that, then we can look at a_2 and decide where to send that. As at the first stage, if were going to send a_1, a_2 to, say, b, b', then the pair (b, b') should have "the same properties" or as "look the same as" (a_1, a_2). And so on.

If we (can) continue this process, then we end up with an embedding of M into N. Whether this will be an embedding of $\mathcal{M}$ as a substructure or as an elementary substructure of $\mathcal{N}$ depends on how we interpreted "having the same properties" during the construction.

But that is just a "forth" construction, which certainly has its uses. If we want an isomorphism then we need to make sure that the eventual map is onto and there's no reason that the process we described will give a surjection. So, to make sure that every element $b_i \in N$ is in the image of the constructed map, we spend half the time on the "back" construction, essentially constructing what will be the inverse of the map from M to N.

That is, the pattern for a back-and-forth construction is:
• enumerate the elements of the two sets M and N as above:
• decide where to send a_1;
• if b_1 is not in the image of the partly-constructed map (which in this starting case just means that we didn't send a_1 to b_1) then decide which element of M is going to get mapped to b_1;
• return to a_2 and, if it's not already in the domain of the partly-constructed map, decide where to send it;
• then back to N - and decide which element of M is going to get mapped to b_2 if that's not already determined.
• *et cetera.*

The process is similar to the forth construction but the forth and back constructions are interleaved so that every element of N, as well as every element of M, gets dealt with at some stage.

Of course, this may not work - there are many theories which have non-isomorphic countable models - but, for a number of important theories, a back-and-forth construction can be carried through to establish $\aleph_0$-categoricity and/or other consequences.

4.2 An example: dense linear orders

Take $\mathcal{L}$ to be $\mathcal{L}_0$ augmented by a binary relation symbol, which we will write as "$\leq$". Let T_{dlo} be the theory of densely linearly ordered sets without endpoints:

(poset) $\forall x\,(x \leq x)$; $\quad \forall x, y, z\,(x \leq y \wedge y \leq z \to x \leq z)$;
$\forall x, y\,(x \leq y \wedge y \leq x \to x = y)$
(linear) $\forall x, y\,(x \leq y \vee y \leq x)$
(densely ordered) $\forall x, y\,(x < y \to \exists z (x < z \wedge z < y))$
(w/o endpoints) $\forall x \exists y \exists z\,(x < y \wedge z < x)$
where we use the symbol "$<$" as an abbreviation (so "$x < y$ stands for "$x \leq y \wedge \neg(x = y)$"); similarly for $>$ and $\geq$.
For instance, $(\mathbb{Q}; \leq)$ is a model of T_{dlo}, as is $(\mathbb{R}; \leq)$.

Theorem 4.1. *Let $\mathcal{M}, \mathcal{N} \models T_{\text{dlo}}$ be countable and let $a_1 < \cdots < a_k$ be elements of M and $b_1 < \cdots < b_k$ be elements of N. Then there is an isomorphism $\alpha : \mathcal{M} \longrightarrow \mathcal{N}$ taking a_i to b_i $(i = 1, \ldots, k)$.*

Proof. The proof is a back-and-forth argument as outlined in Section 4.1. We carry out the argument in all its detail here. It is a slight elaboration of the basic argument in that we start with part of the isomorphism, let's call it α_0, already determined (namely a_i has to be sent to b_i for $i \leq k$).

We start by enumerating
M as $a_1, \ldots, a_k, m_1, \ldots, m_t, \ldots$ and
N as $b_1, \ldots, b_k, n_1, \ldots, n_t, \ldots$.

Since the details really are quite tedious to write out, let's summarise the proof first to make sure the idea is clear. We build up the isomorphism, starting by sending each a_i to b_i. Then we have to decide where to send m_1: so we look for an element - which we will set to be $\alpha(m_1)$ - which bears the same relation to $b_1, \ldots, b_k$ that m_1 does to $a_1, \ldots, a_k$ (we can find such an element because the ordering is dense and because there are no end-points). We move on to m_2 - we find an element which bears the same relation to $b_1, \ldots, b_k, \alpha(m_1)$ that m_2 does to $a_1, \ldots, a_k, m_1$ and set this to be $\alpha(m_2)$, etc. If we continue in this way we build up an order-preserving injection from $\mathcal{M}$ to $\mathcal{N}$: but we want an isomorphism and there's no reason why n_1, say, should be in the image of α - that is just a "forth" construction. For a back-and-forth we modify the construction as follows: at odd-numbered stages we follow the above; at even-numbered stages we reverse the roles of $\mathcal{M}$ and $\mathcal{N}$ (so at stage 2 we will find an element which bears the same relation to $a_1, \ldots, a_k, m_1$ that n_1 does to $b_1, \ldots, b_k, \alpha(m_1)$ and set this to be $\alpha^{-1}(n_1), etc.$). We end up with an isomorphism as described.

Beyond that, the details are just an exercise in precision and notation. Here they are, in all their glory.

We define, inductively, a sequence of partial maps from M to N (a **partial map** from M to N is a subset α of $M \times N$ such that for each $a \in M$ there is at most one $b \in N$ with $(a, b) \in \alpha$ - we write $b = \alpha(a)$ if there is such a b). A partial map α is **injective** if $\alpha(a) = \alpha(a')$ implies $a = a'$. The **domain** of such a partial map α, $\text{dom}(\alpha)$, is $\{a \in M : \exists b \in N,\, (a, b) \in \alpha\}$ and its **image**, $\text{im}(\alpha)$ is $\{b \in N : \exists a \in M,\, (a, b) \in \alpha\}$.

We start by setting $\alpha_0 = \{(a_1, b_1), \ldots, (a_k, b_k)\}$. Note that α_0 is injective. We will inductively define partial maps α_i such that:
α_i is finite;
α_i is injective;
α_i is order-preserving;
$\alpha_i \subseteq \alpha_j$ if $i \leq j$;

m_t is in the domain of α_i for all $t \leq (i+1)/2$;
n_t is in the image of α_i for all $t \leq i/2$.
We have the base case $i = 0$, so assume inductively that we have defined α_i, that α_i satisfies these inductive hypotheses. We define α_{i+1}; there are two cases.

• $i + 1$ is even. If $m = m_{(i+1)/2} \in \mathrm{dom}(\alpha_i)$ then we set $\alpha_{i+1} = \alpha_i$; if not then we proceed as follows.
case (a) If $m < m'$ for every $m' \in \mathrm{dom}(\alpha_i)$, then choose $b \in N$ such that $b < \alpha(m')$ for every $m' \in \mathrm{dom}(\alpha_i)$ - we can do this since $\mathrm{dom}(\alpha_i)$ is finite and $\mathcal{N}$ has no least element. Set $\alpha_{i+1} = \alpha_i \cup \{(m_{(i+1)/2}, b)\}$. Clearly α_{i+1} is finite; it is injective since α_i is injective and $b \notin \mathrm{im}(\alpha_i)$; it is order-preserving since α_i is and since $m_{(i+1)/2} < m'$ for all $m' \in \mathrm{dom}(\alpha_i)$ and $b < n$ for all $n \in \mathrm{im}(\alpha_i)$; it is a superset of α_i and hence of all α_j for $j \leq i$; we have just put $m_{(i+1)/2}$ into the domain of α_{i+1} so, using the inductive hypothesis on α_i, the domain condition holds for α_{i+1} and, since $i+1$ is odd, the image condition holding for α_i implies that it holds for α_{i+1}.
case (b) If $m > m'$ for every $m' \in \mathrm{dom}(\alpha_i)$ then choose $b \in N$ such that $b > \alpha(m')$ for every $m' \in \mathrm{dom}(\alpha_i)$ - we can do this since $\mathrm{dom}(\alpha_i)$ is finite and $\mathcal{N}$ has no greatest element. Set $\alpha_{i+1} = \alpha_i \cup \{(m_{(i+1)/2}, b)\}$. Clearly α_{i+1} is finite; it is injective since α_i is injective and $b \notin \mathrm{im}(\alpha_i)$; it is order-preserving since α_i is and since $m_{(i+1)/2} > m'$ for all $m' \in \mathrm{dom}(\alpha_i)$ and $b > n$ for all $n \in \mathrm{im}(\alpha_i)$; it is a superset of α_i and hence of all α_j for $j \leq i$; we have just put $m_{(i+1)/2}$ into the domain of α_{i+1} so, using the inductive hypothesis on α_i, the domain condition holds for α_{i+1} and, since $i+1$ is odd, the image condition holding for α_i implies that it holds for α_{i+1}.
case (c) The only other possibility is that m lies strictly between two elements in $\mathrm{dom}(\alpha_i)$, say $c < m < c'$ with $c, c' \in \mathrm{dom}(\alpha_i)$ and no element of $\mathrm{dom}(\alpha_i)$ strictly between c and c'. Since α_i is injective and order-preserving, $\alpha_i(c) < \alpha_i(c')$ and there is no element of $\mathrm{im}(\alpha_i)$ strictly between $\alpha_i(c)$ and $\alpha_i(c')$. So choose $b \in N$ with $\alpha_i(c) < b < \alpha_i(c')$ - possible since $\mathcal{N}$ is densely ordered. Set $\alpha_{i+1} = \alpha_i \cup \{(m = m_{(i+1)/2}, b)\}$. Clearly α_{i+1} is finite; it is injective since α_i is injective and $b \notin \mathrm{im}(\alpha_i)$; it is order-preserving since α_i is and by construction; it is a superset of α_i and hence of all α_j for $j \leq i$; we have just put $m_{(i+1)/2}$ into the domain of α_{i+1} so, using the inductive hypothesis on α_i, the domain condition holds for α_{i+1} and, since $i+1$ is odd, the image condition holding for α_i implies that it holds for α_{i+1}.

• i is even. If $m = m_{i/2} \in \mathrm{im}(\alpha_i)$ then we set $\alpha_{i+1} = \alpha_i$; if not then can proceed as in the other case, interchanging the roles of elements of M and N. Or we can save ourselves a bit of tedium by noting that the hypotheses on α_i imply the same on α_i^{-1}, where $\alpha_i^{-1} = \{(b, a) : (a, b) \in \alpha_i\}$ is the inverse relation (indeed inverse partial map), except that the roles of domain and image get interchanged. So we can apply the above procedure to α_i^{-1} in order to get a partial map β, with one more pair in it, which also satisfies the reversed induction hypotheses. Then define α_{i+1} to be β^{-1} and notice that α_{i+1} will satisfy the induction hypothesis. (If you like, you can write out the details as an *exercise*.)

That describes the induction. We then set $\alpha = \bigcup_{i \geq 0} \alpha_i$ and we need to check that α is an isomorphism from $\mathcal{M}$ to $\mathcal{N}$ taking a_i to b_i ($i = 1, \ldots, k$). The last point is clear since $\alpha_0 \subseteq \alpha$. By construction, for each $a \in M$ there is some i such that $(a, b) \in \alpha_i$, hence $(a, b) \in \alpha$ and, also by construction, there is no other

$b' \in N$ such that $(a, b') \in \alpha$, so α is a total, well-defined map from M to N, with $\alpha(a) = b$ iff there is some i such that $(a, b) \in \alpha_i$. It is injective since each α_i is injective and since injectivity is a finitary property. Moreover, by construction, for each $b \in N$ there is some $a \in M$ and i such that $(a, b) \in \alpha_i$, so α is surjective. If $a < a'$ in $\mathcal{M}$ then there is i such that a and a' are both in the domain of α_i; since α_i is order-preserving, we have $\alpha(a) = \alpha_i(a) < \alpha_i(a') = \alpha(a')$. Since the ordering is total and α is injective, it follows that $a < a'$ iff $\alpha(a) < \alpha(a')$. As required. $\square$

Corollary 4.2. *If $\mathcal{M}$, $\mathcal{N}$ are countable models of T_{dlo} then $\mathcal{M} \simeq \mathcal{N}$.*

Proof. Take $\bar{a} = \emptyset = \bar{b}$ in 4.1. $\square$

Corollary 4.3. *If $\mathcal{M}$ is a countable densely linearly ordered set without endpoints then $\mathcal{M} \simeq (\mathbb{Q}, \leq)$.*

Definition 4.4. *A theory T is* **κ-categorical** *(κ a cardinal) if it has, up to isomorphism, just one model of cardinality κ.*

Corollary 4.5. *The theory of densely linearly ordered sets without endpoints is $\aleph_0$-categorical and hence complete.*

Proof. Suppose that $\mathcal{M}, \mathcal{N}$ are models of T_{dlo}. Since T_{dlo} has no finite model, these are infinite and so, by downwards Löwenheim-Skolem (3.14), there are countable $\mathcal{M}'$, $\mathcal{N}'$ with $\mathcal{M}' \prec \mathcal{M}$ and $\mathcal{N}' \prec \mathcal{N}$. By 4.2, $\mathcal{M}' \simeq \mathcal{N}'$ and hence (10.15) $\mathcal{M}' \equiv \mathcal{N}'$. So $\mathcal{M} \equiv \mathcal{N}$. $\square$

Exercise 4.6. Show that if T is a theory in a countable language $\mathcal{L}$, has no finite model, and is $\aleph_0$-categorical, then T is complete.

Exercise 4.7. Show that T_{dlo} is not κ-categorical where κ is the cardinal of the continuum.

Proposition 4.8. *Let $\mathcal{M} \models T_{\mathrm{dlo}}$ and let $\bar{a} = (a_1, \ldots, a_n)$ and $\bar{b} = (b_1, \ldots, b_n)$ be tuples from M with the same order-type in $\mathcal{M}$ - that is, $a_i < a_j$ iff $b_i < b_j$ for all i, j.*

Then for all $\varphi(\bar{x}) \in \mathcal{L}$ we have $\mathcal{M} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\bar{b})$.

Proof. By Downwards Löwenheim-Skolem we can replace $\mathcal{M}$ by a countable elementary substructure still containing $\bar{a}$ and $\bar{b}$, so wlog $\mathcal{M}$ is countable. Renumbering if necessary we may assume that we are in the situation of the hypothesis of Proposition 4.1 and so we may conclude that there is an automorphism of $\mathcal{M}$ taking $\bar{a}$ to $\bar{b}$. An appeal to 10.14 finishes the proof. $\square$

Definition 4.9. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure, $\bar{a}$ a tuple from M. The* **type** *of $\bar{a}$ in $\mathcal{M}$ is the set of formulas (in some fixed tuple, $\bar{x}$, of free variables) satisfied by $\bar{a}$ in $\mathcal{M}$:*

$\mathrm{tp}^{\mathcal{M}}(\bar{a}) = \{\varphi(\bar{x}) \in \mathcal{L} : \mathcal{M} \models \varphi(\bar{a})\}$.

Corollary 4.10. *Let $\mathcal{M} \models T_{\mathrm{dlo}}$ and let $\bar{a}$ be in M. Then the type of $\bar{a}$ in $\mathcal{M}$ is completely determined by the order-type of $\bar{a}$.*

Definition 4.11. *A complete $\mathcal{L}$-theory T has* **elimination of quantifiers** *if, for every formula $\varphi(\overline{x})$ in $\mathcal{L}$ there is a quantifier-free formula $\theta(\overline{x})$ in $\mathcal{L}$ such that φ is equivalent to θ in every model of T; that is,*

$\mathcal{M} \models T$ implies $\varphi(\mathcal{M}) = \theta(\mathcal{M})$,

equivalently,

$\mathcal{M} \models T$ implies $M \models \forall \overline{x}\,(\varphi(\overline{x}) \leftrightarrow \theta(\overline{x}))$.

Corollary 4.12. *T_{dlo} has elimination of quantifiers.*

Proof. Given $\varphi(\bar{x}) \in L$ we must show that there is a quantifier-free formula $\theta(\bar{x})$ equivalent to φ modulo T_{dlo}. By 4.10, whether or not a tuple $\bar{a}$ satisfies φ is determined by the order-type of $\bar{a}$. The order-type τ of a tuple can be described by a quantifier-free formula, ρ_τ say, so we let θ be the disjunction of the ρ_τ such that some (hence by 4.8 every) n-tuple (where $n = l(\bar{x})$) of order-type τ satisfies φ. $\square$

4.3 The Random Graph

By a **graph** we mean an $\mathcal{L}$-structure $\mathcal{M}$ where $\mathcal{L} = \mathcal{L}_0 \vee \{R\}$, with R a binary relation symbol, such that $R^{\mathcal{M}}$ is symmetric and $\mathcal{M} \models \forall x(\neg R(x,x))$ - so we are looking at (undirected) graphs without multiple edges or loops.

Theorem 4.13. *There is a graph $\mathcal{R}$ such that if any countably infinite set of vertices is made into a graph $\mathcal{G}$ by choosing, with probability $= \frac{1}{2}$, pairs of points to be edges or non-edges then, with probability=1, $\mathcal{G} \simeq \mathcal{R}$.*

The proof turns on the following key property which, we will show, characterises $\mathcal{R}$:

($\star$) for every pair (U, V) of disjoint finite sets of vertices there is a vertex, not in U or V, connected to every point in U and to no point in V.

For a particular pair (U, V), write $(\star_{U,V})$ for the statement "there is a vertex not in U or V which is connected to every point in U and to no point in V".

First, we prove that if $\mathcal{G}$ is constructed as above then, with probability=1, $\mathcal{G}$ satisfies ($\star$). Then we prove that any two countable graphs which satisfy ($\star$) must be isomorphic (this is proved *via* a "back-and-forth" construction). Putting those two results together will give us the above theorem 4.13.

Also note (*exercise!*), for future use, that ($\star$) is an elementary property (i.e. can be expressed by a set of sentences in the language $\mathcal{L}$). And that any graph satisfying ($\star$) must be infinite.

Exercise 4.14. Show that the property ($\star$) is equivalent to the apparently weaker property:

($\sharp$) for every pair (U, V) of disjoint finite sets of vertices there is a vertex connected to every point in U and to no point in V.

Theorem 4.15. *Let G be a countably infinite set. For each pair $(a, b) \in G^2$ with $a \neq b$ choose (a, b) to be an edge/non-edge with probability $= \frac{1}{2}$. Then, with probability=1, the resulting graph $\mathcal{G}$ satisfies property ($\star$).*

Proof. Well, sort of proof (since we're not going to set up the appropriate measure space explicitly). We'll prove the equivalent form $\sharp$. Let U, V be any pair of disjoint finite subsets of G. Suppose there are m elements in U and n elements in V. The probability that a random element of $G \setminus (U \cup V)$ satisfies $((\sharp)_{U,V})$ is the probability that a random point is connected to no point of V (that is $\frac{1}{2^n}$) and is unconnected to every point of U (that is $\frac{1}{2^m}$), hence is $\frac{1}{2^{m+n}}$. If we choose N distinct points from $G \setminus (U \cup V)$ at random then the probability that none of them satisfies $(\sharp_{U,V})$ is therefore $\left(1 - \frac{1}{2^{m+n}}\right)^N$, the limit of which as $N \to \infty$ is 0. There are countably many choices for the pair (U, V) so, assuming that the relevant measure satisfies countable additivity (which we don't justify here), the probability that there is some pair (U, V) not satisfying $(\sharp_{U,V})$ also is 0, as required. $\square$

Theorem 4.16. *Suppose that $\mathcal{G}$ and $\mathcal{H}$ are countable graphs both of which satisfy $(\star)$. Then $\mathcal{G} \simeq \mathcal{H}$.*

Proof. This is proved by the back-and-forth method, using $(\star)$ at the key stage. Namely, at each stage of the induction, we will have a partial isomorphism α_i between two finite subsets $G' \subseteq G$ and $H' = \mathrm{im}(\alpha_i) \subseteq H$ of the graphs. In particular, α_i will be an isomorphism between G' and H'. This partial isomorphism has to be extended to include one more element of G (or H) into its domain (or image), let's say the former, say $a \in G \setminus G'$ is the element to be included.

We set $U = \{g \in G' : a \text{ is connected to } g\}$ and $V = \{g \in G' : a \text{ is not connected to } g\}$. So $U \cap V = \emptyset$ and $U \cup V = G'$. Set $\alpha U = \{\alpha_i(g) : g \in U\}$ and $\alpha V = \{\alpha_i(g) : g \in V\}$; since α_i is injective we have $\alpha U \cap \alpha V = \emptyset$. Since $\mathcal{H}$ is assumed to satisfy the condition $(\star_{\alpha U, \alpha V})$ there is $b \in H$ connected to each element of αU and connected to no element of αV. Define $\alpha_{i+1} = \alpha_i \cup \{(a, b)\}$; then clearly α_{i+1} is a partial isomorphism extending α_i and has a in its domain. That's how we do the extension step and the rest is just the standard framework. From which we deduce that there is an isomorphism ($\alpha = \bigcup_i \alpha_i$) from $\mathcal{G}$ to $\mathcal{H}$ as stated. $\square$

Therefore 4.13 is proved: 4.15 gives us existence of a countable graph satisfying $(\star)$ and 4.16 says that any two countable graphs satisfying that condition are isomorphic - thus justifying calling the result *the* **Random Graph**, denoted $\mathcal{R}$. We will use Ran to denote the underlying set (that is, the set of vertices) of the Random Graph.

Corollary 4.17. *Any countable graph which satisfies $(\star)$ is isomorphic to the Random Graph $\mathcal{R}$.*

Proposition 4.18. *Let U, V be finite disjoint subsets of the Random Graph and define $Z = \{z \in \mathrm{Ran} \setminus V : \forall u \in U, R(z,u) \text{ holds and } \forall v \in V, \neg R(z,v)\}$. Let $\mathcal{Z}$ be the induced subgraph on Z. Then $\mathcal{Z} \simeq \mathcal{R}$.*

Proof. In view of 4.17, all we have to do is check property $(\star)$ for $\mathcal{Z}$. Notice that, since there are no loops in $\mathcal{R}$, we have $Z \subseteq \mathrm{Ran} \setminus (U \cup V)$. So let U', V' be disjoint finite subsets of Z; set $U'' = U \cup U'$ and $V'' = V \cup V'$. Then $U'' \cap V'' = \emptyset$. Since the Random Graph satisfies $(\star_{U'',V''})$ there is $g \in \mathrm{Ran} \setminus (U'' \cup V'')$ such that g is connected to every element of U'' and to no element of V''. In particular, g is an element of $Z \setminus (U' \cup V')$ which is connected to every element of U' and to no element of V' - existence of which is what had to be checked. $\square$

Proposition 4.19. *Let* $\mathrm{Ran} = X_1 \cup \cdots \cup X_k$ *be a partition of the Random Graph into finitely many disjoint pieces. For each* i *let* $\mathcal{X}_i$ *be the induced subgraph on* X_i. *Then for some* i *we have* $\mathcal{X}_i \simeq \mathcal{R}$.

Proof. Argue by contradiction: assume not and then produce sets U, V for which $(\star)$ fails in $\mathcal{R}$. $\square$

Exercise 4.20. Prove that if W is any finite subset of Ran then the induced subgraph on $\mathrm{Ran} \setminus W$ is isomorphic to the Random Graph.

Exercise 4.21. Does the conclusion of 4.19 still hold if we don't assume that the X_i are disjoint?

Exercise 4.22. Let $\mathcal{S}$ be a countable model of (Zermelo-Fraenkel) set theory. On S define a graph by joining x to y iff either $x \in y$ or $y \in x$. Prove that the resulting graph is isomorphic to $\mathcal{R}$.

Exercise 4.23. Prove that $\mathcal{R} \simeq \mathcal{R}^{\mathrm{c}}$ where, for any graph $\mathcal{G}$, we define $\mathcal{G}^{\mathrm{c}}$ by replacing the edges of $\mathcal{G}$ with non-edges and the non-edges of $\mathcal{G}$ with edges.

Proposition 4.24. *Every finite or countably infinite graph embeds as a substructure of* $\mathcal{R}$.

Proof. A "forth" construction (cf. proof of 4.1). That is, we enumerate the elements of the given finite or countably infinite graph and build up an embedding through a sequence of partial isomorphisms, one-by-one adding the enumerated elements to the domain, using that the Random Graph satisfies $(\star)$. We simply omit the stages of a back-and-forth argument where we make sure that every element of the second structure gets into the image of some partial isomorphism. $\square$

Exercise 4.25. A relational structure is said to be $\aleph_0$-**homogeneous** if every isomorphism between finite substructures extends to an automorphism of the structure. Prove that the Random Graph is $\aleph_0$-homogeneous.

A countably infinite structure is said to be $\aleph_0$-**categorical** if it is the only countable model of its theory (i.e. if it is characterised, among countable structures, by its theory).

Proposition 4.26. *The Random Graph is* $\aleph_0$*-categorical.*

Proof. Immediate from 4.17 and the observation that $(\star)$ is part of $\mathrm{Th}(\mathcal{R})$. $\square$

4.4 The 0/1-Law for Relational Structures

You can ignore all of of this section.

Suppose that $\mathcal{L}$ is a language with just finitely many constant, function and relation symbols; we say that $\mathcal{L}$ is a "finite" language. Consider a set, say $\{1, 2, \ldots, n\}$, with exactly n elements and let $S(\mathcal{L}, n)$ denote the set of all $\mathcal{L}$-structures on $\{1, \ldots, n\}$. So $\mid S(\mathcal{L}, n) \mid$ is the number of $\mathcal{L}$-structures on that set. Note that we are counting structures, not structures up to isomorphism.

Exercise 4.27. For some specific examples of $\mathcal{L}$, write down a formula for $\mid S(\mathcal{L}, n) \mid$.

Exercise 4.28. Let $\mathcal{L}$ be a language with no function symbols, one constant symbol c and one unary relation symbol R, show that $\mid S(R(c), n) \mid = \mid S(\neg R(c), n) \mid$.

If σ is a sentence of $\mathcal{L}$, let $S(\sigma, n)$ denote the set of those structures in $S(\mathcal{L}, n)$ which satisfy σ:

$$S(\sigma, n) = \{\mathcal{M} : \mathcal{M} \in S(\mathcal{L}, n) \text{ and } \mathcal{M} \models \sigma\}.$$

Definition 4.29. *Set* $p(\sigma, n) = \dfrac{\mid S(\sigma, n) \mid}{\mid S(\mathcal{L}, n) \mid}$ *to be the proportion of* $\mathcal{L}$*-structures with underlying set* $\{1, \ldots, n\}$ *which satisfy* σ*.*

Say that σ *is* **true almost surely** *if* $p(\sigma, n) \longrightarrow 1$ *as* $n \longrightarrow \infty$*.*

Let $T^{\mathrm{as}} = \{\sigma \in \mathcal{L} : \sigma$ *a sentence which is true almost surely*$\}$*.*

Theorem 4.30. *Suppose that* $\mathcal{L}$ *is a finite language. Then* T^{as} *has a model, so is a consistent theory. Every model of* T^{as} *is infinite.*

Proof. By the Compactness Theorem (which we will prove later), it is enough to show that every finite subset of T^{as} has a model. Given a finite subset, we make suitable estimates.

In more detail, if $\sigma_1, \ldots, \sigma_n \in T^{\mathrm{as}}$, then there is N such that, for each $i = 1, \ldots, n$, $p(\sigma_i, N) > 1 - \frac{1}{n+1}$. Therefore there is at least one structure on $\{1, \ldots, N\}$ which satisfies all of $\sigma_1, \ldots, \sigma_n$, as required. $\square$

Exercise 4.31. Prove that, for any sentences $\sigma_1, \ldots, \sigma_k \in T^{\mathrm{as}}$, we have $\bigwedge_{i=1}^{k} \sigma_i \in T^{\mathrm{as}}$. Show that T^{as} is **deductively closed**, that is $\overrightarrow{T^{\mathrm{as}}} = T^{\mathrm{as}}$.

Corollary 4.32. *If* T^{as} *is a complete theory then the 0/1-law holds for* $\mathcal{L}$*-structures, that is, for any sentence* $\sigma \in \mathcal{L}$*, either* $p(\sigma, n) \longrightarrow 1$ *as* $n \longrightarrow \infty$ *or* $p(\neg\sigma, n) \longrightarrow 1$ *as* $n \longrightarrow \infty$*.*

Proof. Note that $p(\neg\sigma, n) \longrightarrow 1$ as $n \longrightarrow \infty$ is equivalent to $p(\sigma, n) \longrightarrow 0$ as $n \longrightarrow \infty$ (justifying the term "0/1-law"). Completeness implies that, for each sentence σ, either $\sigma \in T^{\mathrm{as}}$ or $\neg\sigma \in T^{\mathrm{as}}$. If $\sigma \in T^{\mathrm{as}}$ then $p(\sigma, n) \longrightarrow 1$ as $n \longrightarrow \infty$. Otherwise $\neg\sigma \in T^{\mathrm{as}}$ so $p(\neg\sigma, n) \longrightarrow 1$ as $n \longrightarrow \infty$ and hence $p(\sigma, n) \longrightarrow 0$ as $n \longrightarrow \infty$. $\square$

In fact (a result of Fagin from 1976), if $\mathcal{L}$ is a finite language with no function or constant symbols, then T^{as} is a complete theory and hence the 0/1 law holds for $\mathcal{L}$-structures.

Exercise 4.33. Give an example to show that even for a language with just one function symbol the 0/1-law will fail.

We won't prove Fagin's result but we do prove the following special case. First we need to notice that we can do all the above if, instead of considering *all* $\mathcal{L}$-structures, we consider only those that satisfy some sentence ρ, say and we define the proportions $S(\sigma, n)$ and the notion of "almost surely" by counting only those structures which satisfy ρ.

Theorem 4.34. *Let $\mathcal{L} = \mathcal{L}_0 \vee \{R(-,-)\}$ be a language for finite graphs. Then, relative to the class of graphs (undirected and without loops), $T^{\mathrm{as}} = \mathrm{Th}(\mathcal{R})$ where $\mathcal{R}$ is the Random Graph. In particular T^{as} is complete and hence the 0/1-law holds for finite graphs.*

Proof. From what we have shown already it's enough to show that, for each $m =| U |$, $n =| V |$, the condition $(*)$ is in T^{as}. We count more or less as we did in the proof of 4.13, using that probability approximates proportion in very large sample spaces, and hence in the limit these are equal. Fix m, n and then consider graphs on a fixed set, G, of N vertices and the probability that such a graph satisfies the condition $(*)$ for these values. Fix disjoint subsets of G of sizes m and n. There are 2^{m+n} ways that an element of N can be connected to these $m + n$ vertices, exactly one of which is the "correct" (for $(*)$) way, so the probability that G contains no "correctly" connected vertex is, for $N >> m+n$, approximately $\big(1 - \frac{1}{2^{m+n}}\big)^N$. As $N \to \infty$ this $\to 0$, so the proportion of graphs of size N that have at least one correctly connected vertex, hence satisfy $(*)$, tends to 1 as $N \to \infty$, as required. $\square$

5 Building large models

5.1 The Method of Diagrams

You can ignore all of of this section.

The "method of diagrams" is a way of producing $\mathcal{L}$-structures which contain (an isomorphic copy of) a given $\mathcal{L}$-structure as either a substructure or as an elementary substructure.

Given an $\mathcal{L}$-structure $\mathcal{M}$ and a subset $A \subseteq M$, enrich $\mathcal{L}$ to $\mathcal{L}_A$ by adding new, distinct, constant symbols, one for every element of A: we will denote the constant symbol corresponding to $a \in A$ as c_a (another popular notation is $\ulcorner a \urcorner$). Now extend the $\mathcal{L}$-structure $\mathcal{M}$ to an $\mathcal{L}_A$-structure, denoted $(\mathcal{M}, A)$, in the obvious way i.e. by interpreting c_a as a.

Let $T_A = \mathrm{Th}(\mathcal{M}, A)$ - the (**full**) **diagram of** A; let $D_0(A)$ be the set of quantifier-free formulas in T_A - the **quantifier-free diagram of** A; and let D_0^+ denote the set of atomic formulas in T_A - the **atomic diagram of** A.

Proposition 5.1. *Let $A \subseteq M$, $\mathcal{M}$ an $\mathcal{L}$-structure.*
(0) Let $\mathcal{N}$ be an $\mathcal{L}_A$-structure with $\mathcal{N} \models D_0^+(A)$. If A is the underlying set of a a substructure, $\mathcal{A}$ say, of $\mathcal{M}$, then the map, α, from A to N which takes $a \in A$ to the interpretation, ${c_a}^{\mathcal{N}}$, of c_a in $\mathcal{N}$ is a homomorphism of $\mathcal{L}$-structures.
(1) Let $\mathcal{N}$ be an $\mathcal{L}_A$-structure with $\mathcal{N} \models D_0(A)$. Then the map from A to N which takes $a \in A$ to the interpretation, ${c_a}^{\mathcal{N}}$, of c_a in $\mathcal{N}$ is injective and, if A is the underlying set of a substructure of $\mathcal{M}$, is an embedding of $\mathcal{L}$-structures.
(2) In particular, taking $A = M$, if $\mathcal{N} \models D_0(M)$ then the map defined in (1) is an embedding of $\mathcal{L}$-structures from $\mathcal{M}$ to the reduct, $\mathcal{N} \restriction \mathcal{L}$, of $\mathcal{N}$ to $\mathcal{L}$.
(3) Let $\mathcal{N}$ be an $\mathcal{L}_A$-structure with $\mathcal{N} \models T_A$. Then, with α being the map from A to N which takes $a \in A$ to the interpretation, ${c_a}^{\mathcal{N}}$ of c_a in $\mathcal{N}$, we have, for all $\overline{a}$ from A and $\varphi(\overline{x}) \in \mathcal{L}$, that $\mathcal{M} \models \varphi(\overline{a})$ iff $\mathcal{N} \models \varphi(\alpha\overline{a})$.
(4) In particular, taking $A = M$, if $\mathcal{N} \models T_A$ then the map α defined in (3) is an elementary embedding of the $\mathcal{L}$-structure $\mathcal{M}$ into the reduct, $\mathcal{N} \restriction \mathcal{L}$, of $\mathcal{N}$ to $\mathcal{L}$.

Proof. The proof is rather direct from the construction of T_A and $D_0(A)$.
(0) We have to check the conditions of the definition of morphism of $\mathcal{L}$-structures. If c is a constant symbol of $\mathcal{L}$ then $a = c^{\mathcal{M}} \in A$ (since A is the underlying set of a substructure of $\mathcal{M}$), so $c = c_a$ will be in $D_0^+(A)$ and also $\alpha(c^{\mathcal{M}}) = \alpha(a) = c_a^{\mathcal{N}}$ and the latter therefore equals $c^{\mathcal{N}}$ since $\mathcal{N} \models c = c_a$.

If f is an n-ary function symbol of $\mathcal{L}$ and $a_1, \ldots, a_n$ are elements of A, then $b = f^{\mathcal{M}}(a_1, \ldots, a_n) \in A$, so $f(c_{a_1}, \ldots, c_{a_n}) = c_b$ is in $D_0^+(A)$. Therefore $\alpha(f^{\mathcal{M}}(a_1, \ldots, a_n)) = \alpha(b) = c_b^{\mathcal{N}} = (f(c_{a_1}, \ldots, c_{a_n}))^{\mathcal{N}}$ (since $\mathcal{N} \models D_0^+(A)$) $= f^{\mathcal{N}}(c_{a_1}^{\mathcal{N}}, \ldots, c_{a_n}^{\mathcal{N}})$ (by definition of interpretation) $= f(\alpha(a_1), \ldots, \alpha(c_n))$ (by definition of α).

The relation symbol case is very similar.
(1) The added ingredients give injectivity by considering formulas $c_a \neq c_b$ which are in $D_0(A)$ whenever $a \neq b$ are distinct elements of A; they give the "relation" part of being an embedding (condition (c') of the definition) *via* the formulas $\neg R(c_{a_1}, \ldots, c_{a_n})$ which are in $D_0(A)$ whenever $(a_1, \ldots, a_n) \in A^n \setminus R^{\mathcal{M}}$.
(3) We have $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ (where $(a_1, \ldots, a_n) \in A^n$) iff $\varphi(c_{a_1}, \ldots, c_{a_n}) \in T_A$ iff $\mathcal{N} \models \varphi(c_{a_1}^{\mathcal{N}}, \ldots, c_{a_n}^{\mathcal{N}})$ iff $\mathcal{N} \models \varphi(\alpha(a_1), \ldots, \alpha(a_n))$.

$\square$

5.2 Elementary Chains and Upwards Löwenheim-Skolem

Theorem 5.2. *(Elementary Chain Theorem) Suppose that we have a sequence of $\mathcal{L}$-structures, indexed by the natural numbers and each elementarily embedded in the next: for all $i \in \mathbb{N}$, $\mathcal{M}_i \prec \mathcal{M}_{i+1}$.*

Let $M_\omega = \bigcup_i M_i$. Then:
(1) there is an $\mathcal{L}$-structure on M_ω, induced by the $\mathcal{L}$-structures on the M_i, such that, for all i, $\mathcal{M}_i \leq \mathcal{M}_\omega$;
(2) for all i, $\mathcal{M}_i \prec \mathcal{M}_\omega$.

Proof. The $\mathcal{L}$-structure on M_ω is the "obvious" one (e.g. to determine the value of an n-ary function on $\overline{a}$, choose i large enough that all entries of $\overline{a}$ lie in M_i and assign the value $f^{\mathcal{M}_i}(\overline{a})$) - this process gives a well-defined structure on M_ω because for all $i \leq j$, $\mathcal{M}_i$ is a substructure of $\mathcal{M}_j$ and then (1) is an immediate consequence.

To show (2), first observe that, for all $i \leq j$, $\mathcal{M}_i \prec \mathcal{M}_j$ (by hypothesis, transitivity of $\prec$ (10.9(2)), and induction). Then we show that for all $\varphi(\overline{x}) \in \mathcal{L}$, for all i and for all $\overline{a}$ in M_i, $\mathcal{M}_i \models \varphi(\overline{a})$ iff $\mathcal{M}_\omega \models \varphi(\overline{a})$. This is proved by, you guessed it, induction on complexity of φ with the only (slightly) non-trivial case being that where $\varphi(\overline{x})$ is $\exists y\, \psi(\overline{x}, y)$. $\square$

Theorem 5.3. *(Upwards Löwenheim-Skolem, for countable languages) Suppose that $\mathcal{M}$ is an infinite $\mathcal{L}$-structure where $\mathcal{L}$ is a countable language and suppose that κ is a cardinal with $\kappa \geq \mid M \mid$. Then there is an elementary extension $\mathcal{M}'$ of $\mathcal{M}$ with $\mid M' \mid = \kappa$.*

Proof. You can ignore this proof for now; we'll come back to it once we have proved some things about realising types and give the proof from that point of view.)

The idea is: since M is infinite it is consistent to say that there is an element not equal to any member of M, so take an elementary extension which contains such an element. Of course this "adds" perhaps only one extra element, but the procedure can be continued by transfinite induction up to κ.

That's the basic idea but, when it comes to writing out the procedure carefully, it turns out to be just as easy to add the κ new elements all at once. So, we first extend the language to include a name for every element of M: replace $\mathcal{L}$ by $\mathcal{L}_M = \mathcal{L} \vee \{c_a : a \in M\}$, these being distinct and new constant symbols (see Section 5.1). Then choose a set C of κ new constant symbols and set $\mathcal{L}_{M\cup C} = \mathcal{L}_M \vee C$. Now just write down the set $T_{\neq}$ of sentences in $\mathcal{L}_{M\cup C}$ which says "the elements of C are all different from each other". Then we take a model of this theory. But we also want $\mathcal{M}$ to be an elementary substructure, so we actually need a model of $T_{\neq} \cup \mathrm{Th}(\mathcal{M}, M)$, where $(\mathcal{M}, M)$ denotes the $\mathcal{L}_M$-structure with underlying set M and with c_a interpreted as a for each $a \in M$ (see Section 5.1, "The Method of Diagrams" for more explanation of this).

To show that $T_{\neq} \cup \mathrm{Th}(\mathcal{M}, M)$ has a model it is enough, by the Compactness Theorem 2.20 to show that every finite subset has a model; in fact $\mathcal{M}$ already provides this if we interpret the constant symbols in C appropriately. Let's do that in detail.

Any finite subset T' of $T_{\neq} \cup \mathrm{Th}(\mathcal{M}, M)$ is contained in one consisting of a finite set $\{\sigma_1, \ldots, \sigma_n\}$ of sentences of $\mathrm{Th}(\mathcal{M}, M)$, together with a finite set of inequalities $c_i \neq c_j$ $(i \neq j)$ where $\{c_1, \ldots, c_k\}$ is a set of distinct elements of C. We make an $\mathcal{L}_{M \cup C}$-structure with underlying set M, by starting with the given $\mathcal{L}_M$-structure on M, choosing k distinct elements $a_1, \ldots, a_k$ say, of M (it doesn't matter which we choose), and then interpreting c_i as a_i and any $c \in C \setminus \{c_1, \ldots, c_k\}$ as (say) a_1. Then note that this structure is a model of T'. Thus every finite subset of $T_{\neq} \cup \mathrm{Th}(\mathcal{M}, M)$ has a model.

We conclude that $T_{\neq} \cup \mathrm{Th}(\mathcal{M}, M)$ has a model. But such a model contains a copy of $\mathcal{M}$ as an elementary substructure (5.1); also such a model contains at least κ many distinct elements. You might notice that we haven't used the assumption $\kappa \geq \mid \mathcal{L} \mid$. The point is that this model of $T_{\neq} \cup \mathrm{Th}(\mathcal{M}, M)$ may be of cardinality larger than κ and so we might have to cut it down a bit. For this we use the Downwards Löwenheim-Skolem Theorem 3.14 (which needs the assumption on the size of the language compared with κ). $\square$

Corollary 5.4. *Let T be a theory in the countable language $\mathcal{L}$. If T has an infinite model then, for every infinite cardinal κ, T has a model of cardinality exactly κ.*

Proof. Use the upwards, then the downwards, Löwenheim-Skolem theorems. $\square$

Corollary 5.5. *Suppose that $\mathcal{L}$ is a countable language and that T is an $\mathcal{L}$-theory. Then either there is a finite bound on the cardinalities of models of T or for every infinite cardinal κ there is a model of T of cardinality κ.*

Proof. Combine 2.21 with 5.4. $\square$

If the language $\mathcal{L}$ is uncountable then the statement of the theorem is as follows (and the proof is essentially as above).

Theorem 5.6. *(Upwards Löwenheim-Skolem, for arbitrary languages) Suppose that $\mathcal{M}$ is an infinite $\mathcal{L}$-structure and suppose that κ is a cardinal with $\kappa \geq \mid M \mid$ and $\kappa \geq \mid \mathcal{L} \mid$. Then there is an elementary extension $\mathcal{M}'$ of $\mathcal{M}$ with $\mid M' \mid = \kappa$.*

6 Types

Suppose that $\mathcal{M}$ is an $\mathcal{L}$-structure and $a \in M$. The idea is that the type of a in $\mathcal{M}$ is everything we can say about a using the language $\mathcal{L}$. More generally, if $B \subseteq M$ then the type of a in $\mathcal{M}$ over B is everything we can say about a and how it sits in $\mathcal{M}$ using formulas of $\mathcal{L}$ with parameters from B. The formal definition (which applies to n-tuples of elements as well as single elements) is as follows.

Definition 6.1. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure, $B \subseteq M$ and $\overline{a}$ an n-tuple of elements of M. The* **type of $\overline{a}$ in $\mathcal{M}$ over** *B is the set of formulas with parameters from B satisfied by $\overline{a}$: fix a matching sequence $\overline{x}$ of variables:*
$\mathrm{tp}^{\mathcal{M}}(\overline{a}/B) = \{\varphi(\overline{x}, \overline{b}) : \varphi(\overline{x}, \overline{y}) \in \mathcal{L}, \overline{b}$ *from B and* $\mathcal{M} \models \varphi(\overline{a}, \overline{b})\}$.
In the case that $B = \emptyset$, we just refer to the **type of $\overline{a}$ in $\mathcal{M}$** *and can write $\mathrm{tp}^{\mathcal{M}}(\overline{a})$ as an alternative to $\mathrm{tp}^{\mathcal{M}}(\overline{a}/\emptyset)$.*

We want to allow for the possibility of a type - the collection of all information about an element - not actually being realised in $\mathcal{M}$ - that is, a type could be a description of a *potential* element (think of the description of an infinitesimal in the reals).

Definition 6.2. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure, $B \subseteq M$. An n-***type** *for $\mathcal{M}$* **over** *B is a set, $p(\overline{x})$ ($\overline{x} = (x_1, \ldots, x_n)$), of formulas $\varphi(\overline{x})$ with parameters from B (not shown here in the notation) which is finitely realised in $\mathcal{M}$ and which is such that, for every formula $\psi(\overline{x})$ with parameters from B, either $\psi \in p$ or $\neg\psi \in p$. Here we say that a set $\Phi(\overline{x})$ of formulas is* **finitely realised in** *$\mathcal{M}$ if, whenever $\{\varphi_1, \ldots, \varphi_k\}$ is a finite subset of Φ, $\mathcal{M} \models \exists \overline{x}\, (\varphi_1(\overline{x}) \wedge \cdots \wedge \varphi_n(\overline{x}))$.*
By a **type** *over B we mean an n-type over B for some n.*
We say that a type p over B is **realised** *in $\mathcal{M}$ if there is $\overline{a}$ from M such that $\mathrm{tp}^{\mathcal{M}}(\overline{a}/B) = p$.*

Remark 6.3. The type of an element (or n-tuple of elements) in $\mathcal{M}$ is a type: that is, it is finitely realised in $\mathcal{M}$ (in fact, it is realised in $\mathcal{M}$) and it satisfies the condition that, for every formula over B, it contains that formula or its negation.

Exercise 6.4. Show that if $\Phi(\overline{x}) = \overrightarrow{\Phi}$ is a deductively closed set of formulas with parameters from B which is finitely realised in $\mathcal{M}$, then Φ is a type over B iff Φ is a maximal such set (meaning that if $\Phi'(\overline{x})$ is a set of formulas with parameters from B which is finitely realised in $\mathcal{M}$ and if $\Phi \subseteq \Phi'$ then $\Phi = \Phi'$).

(We say that any set Φ of sentences, in this case in the language $\mathcal{L}_B$, is **deductively closed** if $\Phi = \overrightarrow{\Phi}$. To extend this to formulas with free variables $\overline{x} = (x_1, \ldots, x_n)$, we can define deductive closure using "finitely realised" in place of "model" in the definition of deductive closure, or replace the free occurrences of each x_i by a new constant symbol c_i, then treating a type as a set of sentences in the extended language $\mathcal{L}_B \cup \{c_1, \ldots, c_n\}$, as in the proof of 6.20.)

Exercise 6.5. Replacing formulas by the sets that they define, give an alternative definition of type in terms of B-definable subsets of $\mathcal{M}$.

Lemma 6.6. *If p is a type for $\mathcal{M}$ over B and if $\mathcal{M} \prec \mathcal{N}$, then p is a type for $\mathcal{N}$ over B.*

Proof. The proof follows directly from the definitions. First, note that the set of formulas in $\bar{x}$ with parameters from B is the same, whether we are working in $\mathcal{M}$ or in $\mathcal{N}$. So we just have to check that a set $\{\varphi_1, \ldots, \varphi_k\}$ of such formulas is realised in $\mathcal{M}$ iff it is realised in $\mathcal{N}$. The first means $\mathcal{M} \models \exists \bar{x}\, (\varphi_1(\bar{x}) \wedge \cdots \wedge \varphi_n(\bar{x}))$, the second means $\mathcal{N} \models \exists \bar{x}\, (\varphi_1(\bar{x}) \wedge \cdots \wedge \varphi_n(\bar{x}))$ so, since $\mathcal{M} \prec \mathcal{N}$ and $B \subseteq M$, this follows. $\square$

Remark 6.7. Formulas of $\mathcal{L}$ with parameters from B can alternatively be seen as formulas in the language $\mathcal{L}_B = \mathcal{L} \vee \{c_b : b \in B\}$ which is obtained from $\mathcal{L}$ by adding a new constant symbol, c_b, for each element $b \in B$, then extending $\mathcal{M}$ in the natural way to an $\mathcal{L}_B$-structure (by interpreting c_b as b for each $b \in B$, see Section 9.3) - we write this enriched structure as $(\mathcal{M}, B)$. Then a type over B becomes a set of $\mathcal{L}_B$-formulas. See Sections 5.1, 5.2, 6.5, for more development and applications of this idea.

Exercise 6.8. For $\mathcal{M} = (\mathbb{Q}, <)$ find the types of the form $\text{tp}^{\mathcal{M}}(a_1, \ldots, a_n)$ where $n = 1, 2, 3$. Fix $a_0 \in \mathbb{Q}$: find the types of the form $\text{tp}^{\mathcal{M}}(a_1, \ldots, a_n/a_0)$ for the same values of n. Hint: you should use what we proved in Section 4.2 about T_{dlo} and the exercise becomes easier if you also use the result 6.13 in the next subsection.

6.1 Automorphisms and types

Recall that $\text{Aut}(\mathcal{M})$ denotes the set of automorphisms of the $\mathcal{L}$-structure $\mathcal{M}$.

Exercise 6.9. $\text{Aut}(\mathcal{M})$, under composition of functions, forms a group.

Exercise 6.10. Let $G = \text{Aut}(\mathbb{Q}, <)$.
(i) Show that G is infinite.
(ii) Show that G is uncountable.
(iii) Does G have elements of finite order (other than the identity)?
(iv) Describe some elements of G of infinite order.

Definition 6.11. *Suppose that $\mathcal{M}$ is an $\mathcal{L}$-structure and that $B \subseteq M$. Define $\text{Aut}_B(\mathcal{M})$ to be the set of automorphisms of $\mathcal{M}$ which fix B pointwise (the B-***automorphisms** *of $\mathcal{M}$): $\text{Aut}_B(\mathcal{M}) = \{\alpha \in \text{Aut}(\mathcal{M}) : \alpha(b) = b$ for all $b \in B\}$.*

Exercise 6.12. $\text{Aut}_B(\mathcal{M})$ is a subgroup of $\text{Aut}(\mathcal{M})$.

Proposition 6.13. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure, $B \subseteq M$, $\bar{a}$ in M and let $\alpha \in \text{Aut}_B(\mathcal{M})$.*
Then $\text{tp}^{\mathcal{M}}(\bar{a}/B) = \text{tp}^{\mathcal{M}}(\alpha(\bar{a})/B)$.

Proof. Direct from 10.14 in the case that $B = \emptyset$; the extension to general B can be made by applying 10.14 to the $\mathcal{L}_B$-structure $(\mathcal{M}, B)$. $\square$

6.2 The Space of Types

Given an $\mathcal{L}$-structure $\mathcal{M}$, a subset B of M and an integer n, we let $S_n(B)$ be the set of all n-types (in a fixed sequence of n free variables) over B. We call this the **set of n-types** (of $\mathcal{M}$) **over** B. We also write $S_n^T(B)$ instead of $S_n(B)$, where $T = \text{Th}(\mathcal{M})$ is the complete theory of M. This is justified since if $\mathcal{M} \prec \mathcal{N}$ then we get the same set of types whether constructed with reference to $\mathcal{M}$ or $\mathcal{N}$

(you should check that you see why). We may also write $S^T(B)$ if there is no need to specify n.

We turn the set $S_n(B)$ into a topological space. In order to do so, we must specify the open sets. Given any formula $\phi(x_1, \ldots, x_n)$ with parameters from B we consider the set of types which contain ϕ: $\mathcal{O}_\phi = \{p \in S_n(B) : \phi \in p\}$. We take these sets to form a basis of open sets; therefore the open sets are exactly the unions (possibly infinite) of these. To check that this does give a topology it's enough to show that the intersection of any two basic open sets is open. Which is true: $\mathcal{O}_\phi \cap \mathcal{O}_\psi = \mathcal{O}_{\phi\wedge\psi}$ (*exercise*). (Another exercise: show that $\mathcal{O}_\phi \cup \mathcal{O}_\psi$ has the form $\mathcal{O}_\theta$ for some formula θ.)

Exercise 6.14. (if you know the required topology) Show that every point of $S_n(B)$ is closed, indeed, that this space is Hausdorff.

Show that every basic open set is compact. Deduce that the whole space is compact.

Suppose that $A \subset B$ are subsets of some model $\mathcal{M}$. Define the restriction map $f : S_n(B) \to S_n(A)$ by $f(q) = \{\phi \in q : \text{ all parameters in } \phi \text{ are in A}\}$. Show that $f(q)$ is indeed an n-type over A and hence the map is well-defined. Show that it is continuous.

6.3 Principal types

Given an n-type p over B we say that $\overline{a} \in M^n$ is a **realisation of** p if $\mathrm{tp}(\overline{a}/B) = p$; if there is such an $\overline{a}$, then we say that p is **realised in** $\mathcal{M}$. Thus p is realised in $\mathcal{M}$ iff there is $\overline{a} \in M^n$ with $\mathcal{M} \models \varphi(\overline{a}, \overline{b})$ for every $\varphi(\overline{x}, \overline{b}) \in p$.

Exercise 6.15. Let p be a type over B. Then p is realised in $\mathcal{M}$ iff $\bigcap\{\varphi(\mathcal{M}, \overline{b}) : \varphi(\overline{x}, \overline{b}) \in p\} \neq \emptyset$.

Definition 6.16. *Given $\mathcal{M}$ and $B \subseteq M$, a type $p = p(\overline{x}) \in S(B)$ is* **principal** *if there is a formula $\varphi(\overline{x})$ in p such that for every $\psi(\overline{x})$ in p we have $\mathcal{M} \models \forall\overline{x}\,(\varphi(\overline{x}) \to \psi(\overline{x}))$, that is, if $\varphi(\mathcal{M}) \subseteq \psi(\mathcal{M})$ for every $\psi \in p$. Then we say that φ* **generates** *p.*

Principal types are also referred to as **isolated** types. Because they are exactly the isolated points of the topological space $S_(B)$ (a point p of a topological space is said to be **isolated** iff the set $\{p\}$ is open). Exercise: prove this.

Lemma 6.17. *Let $T = \mathrm{Th}(\mathcal{M})$ and suppose that $p \in S(\emptyset)$ is principal, generated by φ. If $\mathcal{N} \models T$ and if $\overline{a} \in \varphi(\mathcal{N})$, then $\mathrm{tp}^{\mathcal{N}}(\overline{a}) = p$.*

More generally, suppose also $\overline{b}$ is a finite tuple of elements of M and $p \in S(\overline{b})$ is principal, generated by $\varphi(\overline{x}, \overline{b})$. Suppose that $\mathcal{N} \models T_{\overline{b}}$. Define the partial map α from M to N by sending b_j to the interpretation in $\mathcal{N}$ of the constant symbol of $\mathcal{L}_{\overline{b}}$ corresponding to b_j. Then, if $\overline{a} \in \varphi(\mathcal{N}, \alpha(\overline{b}))$, then $\mathrm{tp}^{\mathcal{N}}(\overline{a}/\alpha(\overline{b})) = p$.

Proof. The first statement is direct from the definitions. If $\overline{a} \in \varphi(\mathcal{N})$ then, for every $\psi \in p$, $\overline{a} \in \psi(\mathcal{N})$ (because $\mathcal{N} \models \forall\overline{x}\,(\varphi(\overline{x}) \to \psi(\overline{x}))$), so $\mathrm{tp}^{\mathcal{N}}(\overline{a}) = p$. The second statement is obtained by applying this to the language $\mathcal{L}_{\overline{b}}$ (see Section 5.1). □

Proposition 6.18. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure, let $B \subseteq M$ and set $T = \mathrm{Th}(\mathcal{M})$.*
(a) If $p \in S(B)$ is principal then p is realised in $\mathcal{M}$.
(b) Every principal type in $S^T(\emptyset)$ is realised in every model of T.
(c) If $\mathcal{M}$ is finite then every type in $S(B)$ is principal.

Proof. (a) Suppose that $\varphi(\overline{x})$ generates $p(\overline{x})$. Since p is finitely realised in $\mathcal{M}$ we have $\mathcal{M} \models \exists \overline{x}\, \varphi(\overline{x})$, say $\mathcal{M} \models \varphi(\overline{a})$. Then $\mathrm{tp}(\overline{a}/B) = p$ (since φ generates p).

(b) Immediate from (a).

(c) Any type over B is, by 6.20, realised in an elementary extension of $\mathcal{M}$. But $\mathcal{M}$ is finite so, by 3.9, has no proper elementary extension. That is, every type over B is already realised in $\mathcal{M}$. Given n, there are only finitely many n-tuples in M and hence, by the first observation, only finitely many types in $S_n(B)$. Let p be one of these. For each $q \in S_n(B)$ with $q \neq p$, choose a formula φ_q such that $\varphi_q \in p$ but $\varphi_q \notin q$ (possible since types are complete). Let $\varphi = \bigwedge_q \varphi_q$ - so $\varphi \in p$ and, if there were $\psi \in p$ such $\mathcal{M} \not\models \forall \overline{x}(\varphi(\overline{x}) \to \psi(x))$ then $\varphi(\overline{x}) \wedge \neg\psi(\overline{x})$ would be finitely satisfied and so, by 6.22, extend to some $q \in S_n(B)$. But then we would have $q \neq p$ and hence a contradiction (since we would obtain both $\varphi_q \in q$ and $\varphi_q \notin q$). $\square$

6.4 On the definition of types

Quibbles When we defined the type of a tuple we considered formulas $\varphi(\overline{x})$ in a chosen and fixed tuple of free variables $\overline{x}$. On the other hand, the choice of those free variables is rather immaterial: if we had chosen our free variable sequence to be $\overline{y}$, then the resulting type is still equivalent to the first one if we regard types as descriptions of elements. Often it is convenient to replace the chosen tuple $\overline{x}$ by a different tuple of free variables and, in practice, the notation and terminology used for types is rather lax about all this (if you get confused by the implicit conventions then use the definition from the remark below).

Remark 6.19. An alternative way of introducing types (which is, perhaps, less intuitive but does avoid some quibbles referred to above): given an $\mathcal{L}$-structure $\mathcal{M}$, a subset B of M and an integer n, an n-**type over** (or **with parameters from**) B is a completion of the theory $T_B = \mathrm{Th}(\mathcal{M}, B)$ in the language $\mathcal{L}_{B,\overline{c}}$ where $\overline{c}$ is a fixed n-tuple of new constant symbols. Then, given an n-tuple $\overline{a}$ of elements of M, the **type of** $\overline{a}$ **in** M **over** B is $\mathrm{tp}^M(\overline{a}/B) = \{\varphi(\overline{c}, \overline{b}) : M \models \varphi(\overline{a}, \overline{b}) \text{ where } \varphi \in \mathcal{L}, \overline{b} \in B\}$ - note that this is just the theory of the $\mathcal{L}_{B,\overline{c}}$-structure $(\mathcal{M}, B, \overline{a})$ which is obtained from $(\mathcal{M}, B)$ by extending it to the $\mathcal{L}_{B,\overline{c}}$-structure which has the constant symbol c_i interpreted as a_i for $i = 1, \ldots, n$. You should think about how this relates to the first definition given.

6.5 Realising Types

Recall that if $\mathcal{M}$ is an $\mathcal{L}$-structure, $B \subseteq M$ and $\overline{a}$ is in M then

$$\mathrm{tp}(\overline{a}/B) = \{\varphi(\overline{x}, \overline{b}) : \mathcal{M} \models \varphi(\overline{a}, \overline{b}), \varphi \in \mathcal{L}, \overline{b} \text{ in } B\}$$

is the the set of all formulas with parameters from B satisfied by $\overline{a}$ (in $\mathcal{M}$, equivalently, cf.6.6, in every elementary extension of $\mathcal{M}$).

These types are descriptions (but limited by the expressive power of $\mathcal{L}$) of the properties (including properties which make reference to the elements of B) of tuples $\overline{a}$ of elements of M. These are the "realised (in $\mathcal{M}$)" types over B: but, as discussed in Section 6, one can also have descriptions of elements (and tuples but let's just say "elements") which are consistent but are not actually realised in the given model (think of infinitesimals for the reals). So we defined the more general idea that a type is a consistent and complete description of an element that could be, but might not be, there in $\mathcal{M}$: some of these types will be realised in a given model but others are only descriptions of "potential elements". The next result says that those "potential elements" can indeed be found as actual elements in elementary extensions of $\mathcal{M}$.

Theorem 6.20. *(Every type is realised in an elementary extension) Suppose that $\mathcal{M}$ is an $\mathcal{L}$-structure, that $B \subseteq M$ and that $p = p(\overline{x}) \in S(B)$ is a type over B. Then p is realised in an elementary extension of $\mathcal{M}$.*

Proof. We form an appropriate ultrapower of $\mathcal{M}$. We take the index set I to be the set of finite subsets of p. Given $\phi \in p$, let $J_\phi = \{S \in I : p \in S\}$ be the set of finite subsets which contain ϕ. Since $J_\phi \cap J_\psi = J_{\phi\wedge\psi}$, these sets form a filter basis, so we choose any ultrafilter $\mathcal{U}$ containing I and form the ultraproduct $M^* = M^I/\mathcal{U}$.

Given $S \in I$, choose an n-tuple (assuming that p is an n-type) $a_S \in M$ such that a_S satisfies each formula in S - such elements exist by definition of type (finite realisability). Set $a = (a_S)_S/\sim$; we claim that a realises p. To see that, take $\phi \in p$. Then $\{S : a_S \in \phi(M)\} \supseteq J_{\{\phi\}}$ by definition of $J_{\{\phi\}}$. But $J_\phi \in \mathcal{U}$ by choice of $\mathcal{U}$. So, by Łos' Theorem, $a \in \phi(M^*)$, as required. $\square$

Corollary 6.21. *Suppose that $\mathcal{L}$ is a countable language, that $\mathcal{M}$ is an $\mathcal{L}$-structure, that $B \subseteq M$ is countable and that $p = p(\overline{x}) \in S(B)$ is a type over B. Then p is realised in a countable model of* $\mathrm{Th}(\mathcal{M})$.

Proof. By 6.20, p is realised in some elementary extension, $\mathcal{N}$, of $\mathcal{M}$, say by $\overline{a} = (a_1, \ldots, a_n)$. By Downwards Löwenheim-Skolem 3.14, there is a countable elementary submodel $\mathcal{N}'$ of $\mathcal{N}$ which contains $B \cup \{a_1, \ldots, a_n\}$. Noting that $\mathrm{tp}^{\mathcal{N}'}(\overline{a}) = \mathrm{tp}^{\mathcal{N}}(\overline{a}) = p$, we have the result. $\square$

Extending partial types to types

Proposition 6.22. *(Extending finitely satisfied sets of formulas to types) Let $\mathcal{M}$ be an $\mathcal{L}$-structure, let $B \subseteq M$ and let $\Phi = \Phi(\overline{x})$ be a set of formulas in $\mathcal{L}_B$ (equivalently, formulas of $\mathcal{L}$ with parameters from B) which is finitely satisfied in $\mathcal{M}$. Then Φ is contained in an n-type over B where $n = l(\overline{x})$.*

Proof. We have to extend Φ to a maximal consistent set of formulas. We will give the proof for the case that $\mathcal{L}$ is countable and indicate what has to be done in general.

If $\mathcal{L}$ is countable then we can do this by enumerating the formulas of $\mathcal{L}$ with free variables contained in $\overline{x}$, as $\psi_0, \psi_1, \ldots, \psi_n, \ldots$ and then going through this list, starting with $\Phi = \Phi_0$ and at the n-th stage setting $\Phi_{n+1} = \Phi_n \cup \{\psi_n\}$ if this set is consistent and setting $\Phi_{n+1} = \Phi_n$ otherwise. Then $p = \bigcup\{\Phi_n : n \in \omega\}$ is easily seen to be a complete type containing Φ.

(If $\mathcal{L}$ is uncountable we need to use Zorn's Lemma. We consider the collection of all those finitely satisfiable subsets $\mathcal{K} \subseteq \mathcal{L}$ which consist of formulas with free variables contained in $\overline{x}$ and with $\Phi \subseteq \mathcal{K}$. Regard this collection as a poset, ordered by set inclusion. The conditions of Zorn's Lemma are easily checked (the point being that finite satisfiability is a finitary property) and so, by Zorn's

Lemma, we deduce the existence of a maximal element in this set of extensions of Φ and any such maximal set is easily checked to be a type - we can say a "complete type" to emphasise that for every formula $\psi(\overline{x})$ either ψ or $\neg\psi$ is in p.) $\square$

Corollary 6.23. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure and let $B \subseteq M$. Suppose that $\Phi(\overline{x})$ is a set of $\mathcal{L}$-formulas which is finitely realised in $\mathcal{M}$. Then there is $\mathcal{N} \succ \mathcal{M}$ and $\overline{a}$ in N with $\mathcal{N} \models \Phi(\overline{a})$ (meaning $\mathcal{N} \models \varphi(\overline{a})$ for every $\varphi \in \Phi$).*

Proof. Immediate from 6.22 and 6.20. $\square$

6.6 Omitting types

Definition 6.24. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure and set $T = \mathrm{Th}(\mathcal{M})$. Let $B \subseteq M$, $p \in S^T(B)$ and let $\mathcal{N} \models T_B$. We say that $\mathcal{N}$* **omits** *p if there is no tuple $\overline{a}$ from $\mathcal{N}$ with $\mathrm{tp}^{\mathcal{N}}(\overline{a}/B) = p$.*

The Omitting Types theorem says that, for a countable theory, any non-principal type may be omitted.

Theorem 6.25. *(Omitting Types Theorem) Let T be a complete theory in a countable language $\mathcal{L}$ and let $p \in S(\emptyset)$. Then there is a model $\mathcal{M}$ of T which omits p iff p is non-principal.*
(The statement can be extended to types $p \in S(B)$ by using the enriched language $\mathcal{L}_B$.)

Proof. This is also the Omitted Proof Theorem but you can find the proof (which is quite hard) in Chang and Keisler or in Hodges (for instance). There is a sketch proof in my lecture notes on Intermediate Model Theory at http://www.ma.man.ac.uk/~mprest/leedslecs5.pdf - those are also a source for a little on topics (ultraproducts, topology on the set of types, saturation, etc.) that may be mentioned but not developed in this first course on Model Theory. $\square$

7 Characterisation of $\aleph_0$-categorical theories

Recall that the complete theory T is $\aleph_0$-categorical if there is, up to isomorphism, only one countably infinite model of T. The following theorem characterises these theories.

Theorem 7.1. *(Ryll-Nardzewski) Suppose that T is a complete theory, in a countable language $\mathcal{L}$, without a finite model. Then the following are equivalent.*
(i) T is $\aleph_0$-categorical.
(ii) For every integer n, the set $S_n^T(\emptyset)$ of n-types is finite.
(iii) For every integer n, every type in $S_n^T(\emptyset)$ is principal.
(iv) For every tuple $\overline{x}$ of variables, there are only finitely many formulas with free variables $\overline{x}$ up to T-equivalence (where we say that $\varphi(\overline{x})$ is T-equivalent to $\psi(\overline{x})$ if $\varphi(\mathcal{M}) = \psi(\mathcal{M})$ for every, equivalently any, model $\mathcal{M}$ of T).
(v) For every countable model $\mathcal{M}$ of T and every integer n, there are only finitely many orbits of $\mathrm{Aut}(\mathcal{M})$ *on n-tuples of elements of M.*

Proof. (i)$\Rightarrow$(iii) If there were some non-principal type then, by 6.21, there would be a countable model which realises that type and, by the Omitting Types Theorem 6.25, a countable model which omits that type. These two models cannot be isomorphic, by 6.13 (or 10.14).

(iii)$\Rightarrow$(ii) For each $p \in S_n^T(\emptyset)$ choose (assuming (iii)) some formula φ_p which generates that type. Then consider the set of formulas $\Phi(\overline{x}) = \{\neg\varphi_p(\overline{x}) : p \in S_n^T(\emptyset)\}$. If $S_n^T(\emptyset)$ is infinite then it follows easily that $\Phi(\overline{x})$ is finitely satisfiable in $\mathcal{M}$, where $\mathcal{M}$ is any chosen model of T. Therefore, by 6.22, $\Phi(\overline{x})$ extends to some member, say q, of $S_n^T(\emptyset)$. But then q would contain both φ_q and, since it contains Φ, $\neg\varphi_q$ - contradiction, as required.

(ii)$\Rightarrow$(iv) Fix n and let $p_1, \dots, p_k$ be the types in the finite set $S_n^T(\emptyset)$. Let $\psi(\overline{x})$ be a formula in n free variables $\overline{x}$. Define $S_\psi = \{i : \psi \in p_i\}$. It is easy to show that for formulas ψ and ψ', $\psi(\overline{x})$ and $\psi'(\overline{x})$ are T-equivalent iff $S_\psi = S_{\psi'}$. Since there are only finitely many possibilities for S_ψ there are, therefore, only finitely many formulas in n free variables up to T-equivalence.

(iv)$\Rightarrow$(iii) Let $\psi_1, \dots, \psi_k$ be representatives of the finitely many T-equivalence classes of formulas. Given $p \in S_n^T(\emptyset)$ set $\varphi_p = \bigwedge\{\psi_i : \psi_i \in p\}$. So $\varphi_p \in p$ and any formula in p is T-equivalent to one of the ψ_i and hence is a consequence (modulo T) of φ_p. Thus φ_p is a generator of p and p is principal, as required.

(iii)$\Rightarrow$(i) This is the longest part of the proof; it is a back-and-forth argument where we use the fact that types are principal. Here's the start of the argument. Take countable models, $\mathcal{M}$ and $\mathcal{N}$ of T; we have to show that they are isomorphic. So list the elements of M as $a_1, a_2, \dots$ and similarly list the elements of N. Consider a_1; we must find an element b of N such that $\mathrm{tp}^{\mathcal{N}}(b) = \mathrm{tp}^{\mathcal{M}}(a_1)$. By assumption, there is a formula $\varphi(x)$ which generates $\mathrm{tp}^{\mathcal{M}}(a_1)$. Certainly $\mathcal{M} \models \exists x\, \varphi(x)$ so, since $\mathcal{M} \equiv \mathcal{N}$, $\mathcal{N} \models \exists x\, \varphi(x)$, say $b \in N$ is such that $\mathcal{N} \models \varphi(b)$. By 6.17, $\mathrm{tp}^{\mathcal{N}}(b) = \mathrm{tp}^{\mathcal{M}}(a_1)$, as required. We set $\alpha_1(a_1) = b$.

For the induction step, in the "forth" direction, suppose that we have defined the partial map α_n, sending $a_1, \dots, a_k$ in M to $b_1', \dots, b_k'$ in N and such that $\mathrm{tp}^{\mathcal{M}}(a_1, \dots, a_k) = \mathrm{tp}^{\mathcal{N}}(b_1', \dots, b_k')$. Where do we send a_{k+1}? We need to find an element b_{k+1}' such that $\mathrm{tp}^{\mathcal{M}}(a_1, \dots, a_k, a_{k+1}) = \mathrm{tp}^{\mathcal{N}}(b_1', \dots, b_k', b_{k+1}')$. Let $\psi(x_1, \dots, x_n, y)$ be a generator for $\mathrm{tp}^{\mathcal{M}}(a_1, \dots, a_k, a_{k+1})$. Then $\exists y\, \psi(\overline{x}, y) \in \mathrm{tp}^{\mathcal{M}}(a_1, \dots, a_k)$ which, by the inductive assumption, equals $\mathrm{tp}^{\mathcal{N}}(b_1', \dots, b_k')$. So $\mathcal{N} \models \exists y\, \psi(b_1', \dots, b_k', y)$, so $\mathcal{N} \models \psi(b_1', \dots, b_k', b)$ for some b; set $b_{k+1}' = b$.

Then, by construction and 6.17, $\mathrm{tp}^{\mathcal{N}}(b'_1, \ldots, b'_k, b'_{k+1}) = \mathrm{tp}^{\mathcal{M}}(a_1, \ldots, a_k, a_{k+1})$, as required.

As usual, half the time we use the "forth" argument and the other half, use the "back" argument (in other words, the "forth" argument using α_n^{-1}) and build up an isomorphism between $\mathcal{M}$ and $\mathcal{N}$.

(v)$\Rightarrow$(ii) If there are infinitely many types then choose countably many of them - $p_1, p_2, \ldots, p_i, \ldots$. By 6.20 used repeatedly, there is some model $\mathcal{N}$ of T which realises each of these, say $\mathrm{tp}^{\mathcal{N}}(\overline{a_i}) = p_i$. By the downwards Löwenheim-Skolem Theorem, there is a countable model of T which contains all (the entries of) the $\overline{a_i}$ and hence which realises all of the infinitely many types p_i. But, by 6.13, these must lie in different orbits of $\mathrm{Aut}(\mathcal{M})$.

(ii),(iii)$\Rightarrow$(v) We have to show that if $\overline{a}, \overline{b}$ are n-tuples in $\mathcal{M}$ with the same type then there is $\alpha \in \mathrm{Aut}(\mathcal{M})$ with $\alpha(\overline{a}) = \overline{b}$. To get such an α we just apply the proof of part (iii)$\Rightarrow$(i) but, instead of starting at the empty function, we begin with the partial automorphism which takes $\overline{a}$ to $\overline{b}$ and extend that by a back-and-forth argument. $\square$

8 Appendix (this section to the end of the notes): Languages and Structures done carefully

If you have done some predicate logic then this will all/mostly be material you have seen already, perhaps with somewhat different notation. We will cover some of it in lectures, though quickly and less formally, but you might sometimes need to refer back to this more detailed presentation. Throughout, "language" means "first order, finitary, 1-sorted, predicate language" (in case you were wondering; if you weren't, then ignore this comment).

One important notational difference between this appendix and the main body of notes is that, in the latter, we usually don't distinguish notationally between the structure, in this appendix denoted $\mathcal{M}$, and its underlying set, M (as is normal in algebra: usually we say, "let G be a group", not "let $(G, *)$ be a group").

8.1 Languages: the ingredients

What goes into the language is, to some extent, up to us: it depends on the kind of structure in which we want to interpret the terms and formulas of the language. Here's the stuff that always has to be there.

(i) all the **propositional** (or **Boolean**) **connectives** $\wedge$, $\vee$, $\neg$, $\rightarrow$, $\leftrightarrow$
(ii) countably many **variables** x, y, u, v, v_0, v_1, ...
(iii) the **existential quantifier** $\exists$
(strictly speaking, for each variable, say x, a symbol $\exists x$)
(iv) the **universal quantifier** $\forall$
(strictly speaking, for each variable, say x, a symbol $\forall x$)
(v) a symbol, $=$, for equality[6].

And here's the optional stuff: it consists of **function**, **relation** and **constant symbols**, as many (maybe none) of each as we want, and the function and relation symbols can be of whatever arities we want (the **arity** of a function or relation symbol is the number of arguments, i.e. inputs, it takes). The collection of all these symbols (or, rather, a record of how many of each type there are) is often called the **signature** of the language. We choose these extra symbols depending on what kinds of structures we want to interpret the language in.

Notation: We write $\mathcal{L}_0$ for the basic language - the one with empty signature, meaning no function, relation or constant symbols - and, if S is the set of "extra" symbols we have added, then we will write $\mathcal{L} = \mathcal{L}_0 \vee S$ for the language with these extra symbols.

(It is notationally convenient to regard $\mathcal{L}$ as being, literally, the set of all formulas that can be built from the symbols we have chosen, so then writing, for example, $\varphi \in \mathcal{L}$ makes literal sense. Thus the "$\vee$" should be understood as some sort of "join", not union of sets.)

[6]Some people prefer to use a different, modified equality, symbol for the equality sign that appears in formulas, because equality already has a use in mathematics (as part of the 'metalanguage' as opposed to the formal language we're setting up).

8.2 Building terms

The **terms** of a language $\mathcal{L}$ and, alongside these, their sets, fv$(-)$, of **free variables**, are defined inductively by:

(i) each variable x is a term, $\text{fv}(x) = \{x\}$;

(ii) each constant symbol c is a term, $\text{fv}(c) = \emptyset$;

(iii) if f is an n-ary function symbol and if $t_1, \dots, t_n$ are terms, then $f(t_1, \dots, t_n)$ is a term, $\text{fv}(f(t_1, \dots, t_n)) = \text{fv}(t_1) \cup \cdots \cup \text{fv}(t_n)$;

(iv) that's it (meaning that, in order to be a term, an expression must be obtained by applying the above clauses of the definition).

Also, every occurrence of a variable in a term is a **free occurrence**.

Notice that this is a generalised inductive definition (see the end of this section for a discussion), with the variables and constants being the "base cases" and each function symbol providing one type of "inductive construction step".

Comments: The propositional connectives, $\wedge$, $\vee$, $\neg$, $\rightarrow$, $\leftrightarrow$, are used for combining already-formed formulas/statements into more complex ones. The statements that we make with these languages will be about things happening inside specific mathematical structures, where we take the standard view in pure mathematics that the main objects of study are sets-with-structure. So we need to be able to write formulas/statements which refer to elements of sets - this is reflected in our languages having variables, x, y, $\dots$, which range over the elements of a given set. We also have the universal quantifier $\forall$ ("for all") and the existential quantifier $\exists$ ("there is") prefixing variables, allowing us to write sweeping statements about what happens in a structure: so a formula in this language can begin $\forall x \exists y \dots$ ("for all x there is a y such that .."). Of course, at "..." we want to be able to insert something about x and y and that's where the function, relation and constant symbols come in.

In brief, using a pick-and-mix approach, we set up a language by choosing a certain collection of symbols which can stand for **constants** (specific and fixed elements of structures), for **functions** and for **relations**. Here are some examples.

Example 8.1. One piece of structure that is always there in a set-with-structure is equality - so we will (in this course) always have the relation $=$ and use it to express equality between elements of a set. This is a **binary** (=**2-ary**) relation, meaning that it specifies certain ordered pairs of elements as being related (in a specific sense, in this case as being equal.

Example 8.2. Part of the structure on a set might be an ordering - for example the integers, $\mathbb{Z}$, or reals, $\mathbb{R}$, with the usual ordering. If so, then we would also include a binary relation symbol, say $\leq$, different from equality, in our language.

Example 8.3. Continuing with the examples $\mathbb{Z}$ and $\mathbb{R}$, we might want to express the arithmetic operations of addition, multiplication and taking-the-negative in our language: so we would add two binary function (or "**operation**") symbols (i.e. function symbols taking two arguments), $+$ and $\times$, and also a **unary** (=**1-ary**) function symbol $-$ (which is meant to be used for the function $a \mapsto -a$, not the binary function subtraction). We might also add symbols 0 and 1 as constants.

Example 8.4. Functions with more than two arguments are not uncommon; for instance we might want to have some polynomial functions (or, if we were dealing with the complex numbers, $\mathbb{C}$, perhaps some analytic functions) built into our language, say a 3-ary function symbol f with which we could express the function given by $F(x, y, z) = x^2 + y^2 + xz + 1$ (though clearly we could define this one in terms of the binary operations $+$ and $\times$ and the constant symbol 1, as opposed to, say, $G(x, y, z) = x^y e^z$).

Example 8.5. Relation symbols with more than two arguments are not so common but here's an example. Consider the real line and define the relation $B(x, y, z)$ to mean "y lies (strictly) between x and z".

8.3 Building formulas

The **atomic formulas** of $\mathcal{L}$ (and their **free** (**occurrences of**) **variables**) are defined as follows:

(i) if s, t are terms then $s = t$ is an atomic formula, $\mathrm{fv}(s = t) = \mathrm{fv}(s) \cup \mathrm{fv}(t)$,[7] every occurrence of a variable is free;

(ii) if R is an n-ary relation symbol and if $t_1, \ldots, t_n$ are terms, then $R(t_1, \ldots, t_n)$ is an atomic formula, $\mathrm{fv}(R(t_1, \ldots, t_n)) = \mathrm{fv}(t_1) \cup \cdots \cup \mathrm{fv}(t_n)$, every occurrence of a variable is free.

The **formulas** of $\mathcal{L}$ (and their **free** (**occurrences of**) **variables**) are defined as follows:

(0) every atomic formula is a formula;

(i) if φ is a formula then so is $\neg\varphi$ ("not φ"), $\mathrm{fv}(\neg\varphi) = \mathrm{fv}(\varphi)$, the free occurrences of variables are those of φ;

(ii) if φ and ψ are formulas then so are $\varphi \wedge \psi$ ("φ and ψ"), $\varphi \vee \psi$ ("φ or ψ"), $\varphi \to \psi$ ("φ implies ψ") and $\varphi \leftrightarrow \psi$ ("φ iff ψ"), $\mathrm{fv}(\varphi \wedge \psi) = \mathrm{fv}(\varphi \vee \psi) = \mathrm{fv}(\varphi \to \psi) = \mathrm{fv}(\varphi \leftrightarrow \psi) = \mathrm{fv}(\varphi) \cup \mathrm{fv}(\psi)$, and the free occurrences of variables are those of φ together with those of ψ;

(iii) if φ is a formula and x is any variable then $\exists x\varphi$ and $\forall x\varphi$ are formulas, $\mathrm{fv}(\exists x\varphi) = \mathrm{fv}(\forall x\varphi) = \mathrm{fv}(\varphi) \setminus \{x\}$, the free occurrences of variables are those of φ except for all occurrences of x, those which *were* free in φ, *now* being **bound by**, or **within the scope of**, the **leading** (= leftmost) **quantifier** of the new formula.

Again, this is an inductive definition with (0) giving the base cases and (i), (ii), (iii) giving the inductive processes.

A **sentence** of $\mathcal{L}$ is a formula, σ, of $\mathcal{L}$ with no free variables (i.e. $\mathrm{fv}(\sigma) = \emptyset$).

Since formulas were *constructed* by induction we can *prove* things about them by induction ("on complexity"). One is the fact ("unique readability") that each formula can be constructed in just one way. That is, given a formula, we can form a "(de)construction tree" which shows exactly how it was built up, subformula by subformula. This is important since, when we use induction on complexity to prove things about formulas, we need to know that there's no potential problem about different "construction routes" giving incompatible properties of a particular formula. Unique readability does hold for formulas, and also for terms but we won't prove those here. Both proofs are done by induction (on complexity) and are not very difficult (but not that easy either).

If φ is a formula then, by a **subformula** of φ, we mean any formula which occurs in the "(de)construction tree" of φ.

A comment on use of variables when you are constructing formulas. Note that bound variables are "dummy variables": the formula $\exists x\,(f(x) = y)$ and $\exists z\,(f(z) = y)$ are equivalent (in the sense that they express the same condition, equivalently, the first is true in a given structure whenever the second is, and conversely). A formula with nested occurrences of the same variable being bound by distinct quantifiers can be confusing to read: $\exists x(\forall x(f(x) = x) \to f(x) = x)$ could be written less confusingly as $\exists x(\forall y(f(y) = y) \to f(x) = x)$. Of

[7] This line illustrates why some people prefer to use a different symbol in the language for equality.

course these are not the same formula but one can prove that they are logically equivalent and the second is preferable.

You'll notice that we also use pairs of parentheses "(...)" in the construction of formulas - this is necessary, both in the formal definition and in practice but I left them out earlier because often one can drop some of them without losing unique (or easy) readability.

Another informal notation that we will sometimes use is to "collapse repeated quantifiers", for example writing $\forall x, y\,(x = y \to y = x)$ instead of $\forall x \forall y\,(x = y \to y = x)$. Sometimes the abbreviations $\exists^! x$, $\exists^{\leq n} x$, $\exists^{=n} x$ are useful abbreviations for the quite long formulas which say "there is exactly one x such that ...", "there are less than or equal to n x such that...", "there are exactly n x such that..." respectively.

8.4 Definition and Proof by Induction on Complexity

This is used all over the place in mathematical logic (and is used elsewhere). Recall the usual structure of a definition or proof by induction: we start at a "base case" and have way of getting from one stage to the "next" stage. So the "shape" of the proof (as opposed to the details of it) follows that of the natural numbers $\mathbb{N} = \{0, 1, 2, \dots\}$, where we start with 0 (the base) and have a way (adding 1) of moving from the current thing to the next one. Then, if we have defined something by induction, typically we prove the basic things about it using proof by induction. The validity of proof by induction is stated in the the "Principle of Induction" which, if $\mathbb{N}$ is the base of the induction, takes the form "Given a statement $P(n)$, depending on $n \in \mathbb{N}$, if $P(0)$ is true and if from $P(n)$ we can prove $P(n+1)$, then we may conclude that $P(n)$ is true for every $n \in \mathbb{N}$."

In some contexts, such as constructing the terms and formulas of a language, there may be many starting points and there may be more than one way of going from what we currently have to the "next" stage. But, still, the Principle of Induction is valid, though we would have to phrase it so as to refer to the new "shape" underlying our construction process - roughly said, if something is true for all the "base cases" and is "preserved by each construction process", then it's true of everything we can construct. We usually say that the resulting statement has been proved "by induction on complexity".

9 Languages and structures

9.1 Specifying structures

Fix a language $\mathcal{L}$, equivalently a "signature", which was probably chosen with some specific kind of structure in mind. An $\mathcal{L}$-structure is, roughly, a place where the formulas of the language can make sense. Typically, the class of $\mathcal{L}$-structures will include many more structures than the ones we first thought of in choosing the language $\mathcal{L}$, but we can cut down the class of structures by adding requirements/axioms which are themselves sentences of the language $\mathcal{L}$: see some of the examples below.

The idea: Formulas and sentences do not take on meaning until they are interpreted in a particular structure. Having fixed a language, a structure for that language provides:

a set for the variables to range over (so, if M is the set then, "$\forall x$" will mean "for all x in M" and "$\exists x$" will mean "there is an x in M (such that ...)");

for each constant symbol a specific (and fixed) element of M (which is "named" by the constant symbol in this interpretation);

for each function symbol of the language an actual function (of the correct arity) on M;

for each relation symbol of the language an actual relation (of the correct arity) on M.

Typically we use a script version of the letter used for the underlying set to refer to this whole package of information. Here's the precise definition.

An **$\mathcal{L}$-structure** $\mathcal{M}$ (or **structure for** the language $\mathcal{L}$) consists of a non-empty set M, called the **domain** or **underlying set** of $\mathcal{M}$ (we can write $M = |\mathcal{M}|$) together with an interpretation in M of each of the function, relation and constant symbols of $\mathcal{L}$. By an **interpretation**, $\mathcal{M}$ in M, of one of these symbols we mean the following (and we also insist that the symbol "$=$" for equality be interpreted as actual equality between elements of M):

(i) if c is a constant symbol, then the interpretation of c, which is denoted $c^{\mathcal{M}}$, must be an element of M;

(ii) if f is an n-ary function symbol, then the interpretation of f, which is denoted $f^{\mathcal{M}}$, must be a function from M^n to M;

(iii) if R is an n-ary relation symbol, then the interpretation of R, which is denoted $R^{\mathcal{M}}$, must be a subset of M^n (in particular, the interpretation of a 1-ary predicate symbol is a subset of M - a "property" of elements of M).

To show, or at least indicate, this structure, we can write

$$\mathcal{M} = (M; c^{\mathcal{M}}, d^{\mathcal{M}}, \ldots, f^{\mathcal{M}}, g^{\mathcal{M}}, \ldots, R^{\mathcal{M}}, S^{\mathcal{M}}, \ldots).$$

If no confusion should arise from doing so, the superscript "$\mathcal{M}$" may be dropped (thus the same symbol "f" might used for the function symbol and for the particular interpretation of this symbol in a given $\mathcal{L}$-structure).

Important Note: In the main part of the notes we will often denote a structure $\mathcal{M}$, meaning $\mathcal{M} = (M; c^{\mathcal{M}}, d^{\mathcal{M}}, \ldots, f^{\mathcal{M}}, g^{\mathcal{M}}, \ldots, R^{\mathcal{M}}, S^{\mathcal{M}}, \ldots)$, just by M since, notationally, this is a bit simpler and more in line with the usual practice in algebra (where a group or ring is usually just denoted by the same letter that we use for its underlying set). In this appendix, however, and in some other places in the notes where we need to take a bit more care, we will use the script form of the notation in order to emphasise the distinction between the structure and its underlying set.

Some basic examples

Example 9.1. An $\mathcal{L}_0$-structure is simply a set, so $\mathcal{L}_0$-structures have rather limited value as illustrations of definitions and results.

Example 9.2. (**Directed graphs**) An $\mathcal{L} = \mathcal{L}_0 \vee \{R(-,-)\}$-structure $\mathcal{M}$ consists of a set M together with an interpretation of the binary relation symbol R as a particular subset, $R^{\mathcal{M}}$, of $M \times M$: $\mathcal{M} = (M; R^{\mathcal{M}})$. That is, an $\mathcal{L}$-structure consists of a set together with a specified binary relation on that set.
Given such a structure, its **directed graph**, or **digraph** for short, has for its vertices the elements of M and has an arrow going from vertex a to vertex b iff $(a, b) \in R^{\mathcal{M}}$. This gives an often useful graphical way of picturing or even defining a relation $R^{\mathcal{M}}$ (note that the digraph of a relation specifies the relation completely).

Certain types of binary relation are of particular importance in that they occur frequently in mathematics (and elsewhere).

Example 9.3. (**Posets**) A **partially ordered set** (**poset** for short) consists of a set P and a binary relation on it, usually written $\leq$, which satisfies:
for all $a \in P$, $a \leq a$ ($\leq$ is **reflexive**);
for all $a, b, c \in P$, $a \leq b$ and $b \leq c$ implies $a \leq c$ ($\leq$ is **transitive**);
for all $a, b \in P$, if $a \leq b$ and $b \leq a$ then $a = b$ ($\leq$ is **weakly antisymmetric**).
(short) *Exercise:* check that all these conditions may be expressed by a sentence of the language $\mathcal{L}$.
The **Hasse diagram** of a poset is a diagrammatic means of representing a poset. It is obtained by, for each element a of the poset, placing a point on the plane to represent a and connecting it to each of the immediate successors of a (if there are any) by a line which goes in an upwards direction from that point. We say that b is an **immediate successor** of a if $a < b$ (i.e. $a \leq b$ and $a \neq b$) and if $a \leq c \leq b$ implies $a = c$ or $c = b$: we also then say that a is an **immediate predecessor** of b.

Example 9.4. (**Equivalence relations**) An **equivalence relation**, $\equiv$, on a set X is a binary relation which satisfies:
for all $a \in X$, $a \equiv a$ ($\equiv$ is reflexive);
for all $a, b \in X$, $a \equiv b$ implies $b \equiv a$ ($\equiv$ is **symmetric**);
for all $a, b, c \in X$, $a \equiv b$ and $b \equiv c$ implies $a \equiv c$ ($\equiv$ is transitive).
The ($\equiv$-)**equivalence class** of an element $a \in X$ is denoted, for instance, $a/\equiv$, or $[a]_{\equiv}$, or just $[a]$, and is the set $\{b \in X : b \equiv a\}$ of elements which are $\equiv$-equivalent to a. The key point is that equivalence classes are equal or disjoint: if $a, b \in X$, then either $[a] = [b]$ or $[a] \cap [b] = \emptyset$. Thus the distinct $\equiv$-equivalence classes partition X into disjoint subsets.

The above are sometimes referred to as "toy examples" of structures because they are easily specified, quite malleable, not so complicated to understand, but often rather bare of interesting structure, though they are important and do underly many interesting, "richer" structures.

9.2 Interpreting terms

This is the process of substituting in values for variables - something we are all familiar with from contexts such as evaluating polynomials or other functions on particular domains.

The general context is that we have a language $\mathcal{L}$ and an $\mathcal{L}$-structure $\mathcal{M}$. Take some term t of $\mathcal{L}$ with (free) variables $x_1, \ldots, x_n$ (say). We evaluate this on $\mathcal{M}$ by choosing an n-tuple $(a_1, \ldots, a_n) \in M^n$ and then replacing each occurrence of x_i by a_i, to get an element of M, which we write variously as $t^{\mathcal{M}}(a_1, \ldots, a_n)$, $t(a_1, \ldots, a_n)$, $t(\overline{x}/\overline{a})$, or $t(\overline{a})$.

Defining this process precisely is done by induction on complexity of the term t: suppose the (free) variables of t are among $x_1, \ldots, x_n$, and let $\overline{a} = (a_1, \ldots, a_n) \in M^n$; then the **value**, $t(\overline{a})$, of t at $\overline{a}$ (or the **evaluation** of t at $\overline{a}$) is defined as follows:

(i) if t is a constant symbol c, then $t(\overline{a})$ is the element $c^{\mathcal{M}} \in M$; if t is the variable x_i then $t(\overline{a})$ is a_i;

(ii) if t is $f(t_1, \ldots, t_m)$ where $t_1, \ldots, t_m$ are terms then, assuming that we have inductively defined the elements $t_1(\overline{a}), \ldots, t_m(\overline{a})$ of M, we define $t(\overline{a})$ to be $f^{\mathcal{M}}(t_1(\overline{a}), \ldots, t_m(\overline{a}))$ - the value of the function $f^{\mathcal{M}}$ on the element $(t_1(\overline{a}), \ldots, t_m(\overline{a}))$ of M^m.

You should work through a few examples (say, polynomial expressions) to see how this works and to realise that it is a process already familiar from particular contexts.

Important! Notice that a term t is a formal expression of the language $\mathcal{L}$ whereas the result, $t(\overline{a})$, of substitution is an element of M. (That said, by adding new constant symbols to the language $\mathcal{L}$ to "name" $a_1, \ldots, a_n$, we can narrow the gap between these formal expressions and elements of M, see the next subsection.)

9.3 Interpreting formulas

Suppose that σ is a sentence of the language $\mathcal{L}$ and that $\mathcal{M}$ is an $\mathcal{L}$-structure. Then σ may be "interpreted" in $\mathcal{M}$ and, so interpreted, it will make a definite assertion "about $\mathcal{M}$" which will be either true or false. We will give a precise definition of what we mean by "interpreting σ in $\mathcal{M}$" but, in order to do so, we must consider the interpretation of general formulas, including those with free variables, in $\mathcal{M}$. The reason for this is that we will define the notion of interpretation by induction on the complexity of a formula and the inductive construction of a sentence goes via formulas which may (and generally will) have free variables.

But then we have the following issue: a formula, when interpreted in a structure, need not have a truth value. For instance, the formula $\exists y R(x, y)$ has free variable x and its interpretation in $\mathcal{M}$ asserts the existence of an element $b \in M$ (say) such that x is R-related to b. This "assertion" does not have a truth value, since its truth or falsity depends on the value we assign to x: in general, for some values of x the "assertion" will be true and for others it will be false. One way to deal with this is to introduce formulas with (some) free variables replaced by values in M. The idea is familiar from, say, polynomials, where we might replace one variable by a value and regard the result as a polynomial with one fewer variable (and possibly over a larger ring of coefficients). We define this formally now.

Substituting values into formulas: Suppose that $\varphi(x_1, \ldots, x_n)$ is a formula of the language $\mathcal{L}$ and has free variables among $x_1, \ldots, x_n$ (it is technically convenient not to insist that each of $x_1, \ldots, x_n$ actually occurs in φ). Let $a_1, \ldots, a_n$ be elements of M (they need not be distinct). Then $\varphi(a_1, \ldots, a_n)$ is the "formula" which results when each free occurrence of x_i in φ is replaced by a_i.

At least, that is the idea: but an element cannot literally be substituted into a formula, so what we can do is the following. Extend, temporarily, the language $\mathcal{L}$ by adding a "name" for each element $a_1, \ldots, a_n$. That is, add to $\mathcal{L}$ new constant symbols $c_1, \ldots, c_n$ and regard $\mathcal{M}$ as a structure for this enriched language, $\mathcal{L}'$, by interpreting each constant symbol c_i as the corresponding element a_i of M: we could write $\mathcal{M}' = (\mathcal{M}, c_1^{\mathcal{M}'} = a_1, \ldots, c_n^{\mathcal{M}'} = a_n)$ for this enriched (or "extended") structure. Then we can take "$\varphi(a_1, \ldots, a_n)$" literally to mean the formula, $\varphi(c_1, \ldots, c_n)$, of this enriched language which is obtained by replacing each free occurrence of x_i by the constant symbol c_i and which is then interpreted in the structure $\mathcal{M}'$ for the enriched language. When we use intuitively understandable expressions like $\varphi(a_1, \ldots, a_n)$ this what we literally mean (but it would be cumbersome to refer to this extension process every time we use such an expression, so we don't).

By the way, if we wish to replace only some of the free variables then we use notation such as $\varphi(a_1, x_2, x_3, a_4)$: x_1 and x_4 have been substituted for, but x_2 and x_3 have not.

It is important to note that it is only the *free* occurrences of variables that are replaced. For instance if $\varphi(x)$ is the formula $x + 1 = 0 \wedge \forall x(x = x)$ then it is only the first occurrence of x that is free, so the formula $\varphi(-1)$ is $-1 + 1 = 0 \wedge \forall x(x = x)$.

Satisfaction of formulas: If $\varphi(x_1, \ldots, x_n)$ is a formula (with, by our convention, all its free variables among $x_1, \ldots, x_n$) and if $a_1, \ldots, a_n$ are elements of the structure $\mathcal{M}$, then the expression $\varphi(a_1, \ldots, a_n)$, defined above, is an assertion with a definite truth value. We will refer to it as a **formula with parameters**: it is (notation for) a sentence ($\varphi(c_1, \ldots, c_n)$) of a language $\mathcal{L}'$ which is $\mathcal{L}$ extended as above. We now define what we mean by such an expression being satisfied in a structure $\mathcal{M}$. The notation we use is $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$, which is read as "$\mathcal{M}$ satisfies $\varphi(a_1, \ldots, a_n)$" or "$\varphi(a_1, \ldots, a_n)$ is true in $\mathcal{M}$". (An equivalent notation is $(a_1, \ldots, a_n) \in \varphi(\mathcal{M})$, where $\varphi(\mathcal{M})$ is the set of "solutions" of φ in $\mathcal{M}$ - see "definable sets" below.)

The definition is by induction on complexity; the base case is that $\varphi = \varphi(x_1, \ldots, x_n)$ is an atomic formula - hence of the form

(i) $t_1 = t_2$ where t_1, t_2 are terms with $\mathrm{fv}(t_1), \mathrm{fv}(t_2) \subseteq \{x_1, \ldots, x_n\}$, or

(ii) $R(t_1, \ldots, t_m)$ where $t_1, \ldots, t_m$ are terms with $\mathrm{fv}(t_j) \subseteq \{x_1, \ldots, x_n\}$ for each j.

In case (i) we set $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ iff $t_1(a_1, \ldots, a_n) = t_2(a_1, \ldots, a_n)$. In case (ii) we set $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ iff $(t_1(a_1, \ldots, a_n), \ldots, t_m(a_1, \ldots, a_n)) \in R^{\mathcal{M}}$.

We now proceed by induction on complexity - that is, following the construction of formulas from simpler ones. So suppose, inductively, that we have defined the meaning of $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ and $\mathcal{M} \models \psi(a_1, \ldots, a_n)$ for certain formulas $\varphi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ and for *every* $(a_1, \ldots, a_n) \in M^n$. We're taking advantage here of our convention that the variables listed beside a formula can contain variables that don't actually appear in the formula - so we can take a common list of variables for both φ and ψ. (Otherwise we'd have to combine the lists explicitly, which would be messy.)

($\wedge$) We set $\mathcal{M} \models (\varphi \wedge \psi)(a_1, \ldots, a_n)$ iff $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ and $\mathcal{M} \models \psi(a_1, \ldots, a_n)$.

($\neg$) We set $\mathcal{M} \models \neg\varphi(a_1, \ldots, a_n)$ iff $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ does not hold.

($\exists$) Consider a formula of the form $\exists y\, \varphi(x_1, \ldots, x_n)$. If y is not one of the free variables $x_1, \ldots, x_n$ of φ then we set $\mathcal{M} \models \exists y\, \varphi(a_1, \ldots, a_n)$ exactly if $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$. The non-trivial case is that y is one of the free variables of φ : let's say it is x_1 (the other cases being similar) - so we are considering the formula $\exists x_1\, \varphi(x_1, \ldots, x_n)$. Then we set $\mathcal{M} \models \exists x_1\, \varphi(a_2, \ldots, a_n)$ (we can drop x_1 now from the list of free variables if we want to) iff there is *some* element c of M such that $\mathcal{M} \models \varphi(c, a_2, \ldots, a_n)$.

We don't have to consider the other connectives ($\vee$, $\rightarrow$, $\leftrightarrow$) because they can be written in terms of $\wedge$ and $\neg$ and, since "$\forall = \neg\exists\neg$", we don't have to have a "$\forall$" case. That just makes definitions, and proofs by induction following those definitions, shorter.

Other methods. There are other (equivalent) ways of setting up a notion of interpretation of a formula in a structure $\mathcal{M}$. One that is commonly used is to work with functions ("assignments" or "valuations") from the set of all variables of the language to the underlying set M and then to define the truth value of every formula relative to any given such function. But I find that very contrived and feel that the above is more natural.

9.4 Definable subsets

Suppose that $\varphi = \varphi(x)$ is a formula (of the language $\mathcal{L}$) with free variable x and let $\mathcal{M}$ be an $\mathcal{L}$-structure. Then the set of all elements of $\mathcal{M}$ which satisfy φ - the "solution set" of φ in $\mathcal{M}$ - is denoted $\varphi(\mathcal{M})$ and is usually termed the **subset of $\mathcal{M}$ defined by φ**:

$$\varphi(\mathcal{M}) = \{a \in M : \mathcal{M} \models \varphi(a)\}.$$

More generally, if $\varphi = \varphi(x_1, \ldots, x_n)$, then this defines a set of n-tuples - a **subset of M^n definable in M**:

$$\varphi(\mathcal{M}) = \{(a_1, \ldots, a_n) \in M^n : \mathcal{M} \models \varphi(a_1, \ldots, a_n)\}.$$

Even more generally, if $b_1, \ldots, b_m$ are elements of M and if $\varphi(x_1, \ldots, x_n, b_1, \ldots, b_m)$ is a formula with parameters (that is, a formula of the language enriched by constant symbols for these elements), then

$$\varphi(\mathcal{M}, b_1, \ldots, b_m) = \{(a_1, \ldots, a_n) \in M^n : \mathcal{M} \models \varphi(a_1, \ldots, a_n, b_1, \ldots, b_m)\}$$

is a subset of M^n **definable with parameters** $b_1, \ldots, b_n$. Generally we use compact notation, writing $\overline{a}$ for $(a_1, \ldots, a_n)$ etc., so would set

$$\varphi(\mathcal{M}, \overline{b}) = \{\overline{a} \in M^n : \mathcal{M} \models \varphi(\overline{a}, \overline{b})\}.$$

Subsets of the above forms are generally referred to as **definable subsets** (of the given structure $\mathcal{M}$). We might want to restrict the parameters allowed in formulas to some subset, A, of the underlying set, M, of the $\mathcal{L}$-structure $\mathcal{M}$ - in which case we refer to "A-**definable**" subsets of $\mathcal{M}$ (or, indeed, of any structure containing that set A).

Exercise 9.5. Show that the set of (A-)definable subsets of $\mathcal{M}$ is closed under finite intersection and union, complement and projections (the last referring to the canonical projections from one power of M to another). In particular, note that existential quantification corresponds to projection.

Exercise 9.6. Give examples to show that the definable subsets of a structure are not in general closed under infinite unions and intersections.

10 Substructures and Morphisms

10.1 Substructures

Let $\mathcal{M}$ be an $\mathcal{L}$-structure and take any nonempty subset, N, of M. The $\mathcal{L}$-structure **induced on** N - if it exists - is given by:

$c^{\mathcal{N}} = c^{\mathcal{M}}$ for every constant symbol c (so we need $c^{\mathcal{N}} \in N$);

$f^{\mathcal{N}}(\bar{a}) = f^{\mathcal{M}}(\bar{a})$ for every n-ary function symbol f and $\bar{a} \in N^n$ (so we need $f^{\mathcal{N}}(N^n) \subseteq N$);

$R^{\mathcal{N}}(\bar{a})$ iff $R^{\mathcal{M}}(\bar{a})$ for all $\bar{a} \in N^n$ for every n-ary relation symbol R (no requirement on N).

If this structure is actually defined - that is, if N "contains the constants" and is closed under all the functions - then we write $\mathcal{N} = \mathcal{M} \upharpoonright N$ (read "$\upharpoonright$" as "restricted to") and say that $\mathcal{N}$ is the **substructure of** $\mathcal{M}$ **based on** N. We write $\mathcal{N} \leq \mathcal{M}$ to mean that $\mathcal{N}$ is a substructure of $\mathcal{M}$.

The next statement is immediate from the definition.

Lemma 10.1. *If $\mathcal{M}$ is an $\mathcal{L}$-structure and N is a subset of M, the underlying set of $\mathcal{M}$, then N is the underlying set of a substructure of $\mathcal{M}$ iff:*
for every constant symbol c of $\mathcal{L}$, $c^{\mathcal{M}} \in N$ and;
for every (n-ary) function symbol f of $\mathcal{L}$ and $a_1, \dots, a_n \in N$, $f^{\mathcal{M}}(a_1, \dots, a_n) \in N$.

Proof. If N satisfies these conditions then we define the $\mathcal{L}$-structure $\mathcal{N}$ with underlying set N to be given by $c^{\mathcal{N}} = c^{\mathcal{M}}$, $f^{\mathcal{N}} = f^{\mathcal{M}} \upharpoonright N^n$ if f is an n-ary function symbol and $R^{\mathcal{N}} = R^{\mathcal{M}} \cap N^n$ if R is an n-ary relation symbol. Then $\mathcal{N} \leq \mathcal{M}$. In the other direction, if there is a substructure of $\mathcal{M}$ with underlying set N then this is how its structure must be defined - which will be possible only if N satisfies the conditions stated in the lemma. $\square$

The notion of substructure is sensitive to choice of language. For example, if we want a language suitable for groups then we can use just a binary operation symbol (for the group operation) or we can also choose to add a constant symbol (for the identity of the group) and a 1-ary function symbol for taking inverses of elements. In the first case, a subset of a group will give a subgroup provided it is closed under the group operation - it need not contain the identity nor be closed under inverses, so need not be a (sub)group; whereas a substructure with respect to the second language will contain the identity and be closed under taking inverses, so will be a subgroup.

Proposition 10.2. *Let $\mathcal{L}$ be a language and $\mathcal{M} = (M; \dots)$ be an $\mathcal{L}$-structure. Given any subset A of M there is a smallest $\mathcal{L}$-substructure of $\mathcal{M}$ which contains A - this substructure is denoted $\langle A \rangle$ and called the* **substructure generated by** *A (since it can also be obtained by taking the set A together with the set of elements of M which are the interpretations of the constant symbols of $\mathcal{L}$ and then closing this set under repeated application of the functions $f^{\mathcal{M}}$).*

Proof. One proof takes the parenthetical description of $\langle A \rangle$, gives a careful definition of that, shows that the result is a substructure containing A and hence must be the smallest such substructure. Another, cheaper, proof establishes just the first statement (not the description). We give the cheap proof first.

Let $\mathbb{S}$ be the set of all substructures of $\mathcal{M}$ containing A. This is non-empty since $\mathcal{M} \in \mathbb{S}$. Set $A' = \bigcap_{\mathcal{S}\in\mathbb{S}} S$ be the intersection of the underlying sets of all the substructures in $\mathbb{S}$. Since $A \subseteq S$ for every $\mathcal{S} \in \mathbb{S}$, we have $A \subseteq A'$. Since $c^{\mathcal{M}} \in S$ for every $\mathcal{S} \in \mathbb{S}$, we have $c^{\mathcal{M}} \in A'$. Let f be an n-ary function symbol and $a_1, \dots, a_n \in A'$. Then, for every $\mathcal{S} \in \mathbb{S}$, $a_1, \dots, a_n \in S$ so, since $\mathcal{S}$ is a substructure of $\mathcal{M}$, $f^{\mathcal{M}}(a_1, \dots, a_n) \in S$. This being true for every $\mathcal{S} \in \mathbb{S}$, we have $f^{\mathcal{M}}(a_1, \dots, a_n) \in A'$. Thus A' contains all the constants and is closed under the functions. That is, A' satisfies the conditions of 10.1 for being the underlying set of a substructure, $\mathcal{A}'$ say, of $\mathcal{M}$. This substructure contains A, so $\mathcal{A}' \in \mathbb{S}$; hence is the (unique) smallest member of $\mathcal{S}$, that is, it is the smallest substructure of $\mathcal{M}$ containing A.

Here's the more informative proof which, instead of intersecting down above A, builds up from A (there will be more examples of this style of construction later). Set $A_0 = A \cup \{c^{\mathcal{M}} : c \text{ a constant symbol}\}$. Having inductively defined A_i, set

$$A_{i+1} = A_i \cup \{f^{\mathcal{M}}(a_1, \dots, a_n) : f \text{ an } n\text{-ary function symbol}, a_1, \dots, a_n \in A_i\}.$$

Note that $A_i \subseteq A_j$ whenever $i \leq j$. Set $A_\omega = \bigcup_{i\geq 0} A_i$. We claim that A_ω is the underlying subset of a substructure of $\mathcal{M}$ (in fact, it will turn out to be A' in the paragraph above); we check the conditions of 10.1. The condition on constants was satisfied at the first stage (by definition of A_0 and since $A_0 \subseteq A_\omega$). For the other condition, suppose $a_1, \dots, a_n \in A_\omega$. For each i there is $m(i)$ such that $a_i \in A_{m(i)}$. Let $m = \max\{m(1), \dots, m(n)\}$; then $a_i \in A_m$ for every i (since the A_i form an increasing sequence of subsets). So $f^{\mathcal{M}}(a_1, \dots, a_n) \in A_{m+1}$ by construction. Since $A_{m+1} \subseteq A_\omega$, we deduce $f^{\mathcal{M}}(a_1, \dots, a_n) \in A_\omega$, which is what we wanted. Therefore, by 10.1 we deduce that A_ω is the underlying set of a substructure of $\mathcal{M}$ which contains A. To see that it's the smallest, we note that if $\mathcal{S}$ is any substructure of $\mathcal{M}$ which contains A then it must contain A_0 and then, inductively on i, since it contains A_i it must also contain A_{i+1}, hence it contains A_ω, as required. (And therefore, both being the smallest substructure containing A, the substructure with underlying set A_ω must equal the substructure $\mathcal{A}'$ that was produced in the first paragraph, for which we use the notation $\langle A \rangle$.) $\square$

10.2 Morphisms between $\mathcal{L}$-structures

These are the "structure-preserving" maps between $\mathcal{L}$-structures. You likely have seen special cases before, particularly algebraic ones such as groups, fields, vector space, boolean algebras, rings, Lie algebras, ...; possibly also relational examples such as partial orders.

Definition 10.3. *If $\mathcal{M}$ and $\mathcal{N}$ are $\mathcal{L}$-structures then a* **homomorphism***, or just* **morphism***, from $\mathcal{M}$ to $\mathcal{N}$ is a map $\alpha : M \longrightarrow N$ such that:*
(a) for every constant symbol c of $\mathcal{L}$, $\alpha(c^{\mathcal{M}}) = c^{\mathcal{N}}$;
(b) for every n-ary function symbol f of $\mathcal{L}$ and every n-tuple $a_1, \dots, a_n$ of elements of M, $\alpha(f^{\mathcal{M}}(a_1, \dots, a_n)) = f^{\mathcal{N}}(\alpha(a_1), \dots, \alpha(a_n))$;
(c) for every n-ary relation symbol R of $\mathcal{L}$ and every n-tuple $a_1, \dots, a_n$ of elements of M, if $R^{\mathcal{M}}(a_1, \dots, a_n)$ holds then $R^{\mathcal{N}}(\alpha(a_1), \dots, \alpha(a_n))$ holds.

Such a morphism is an **embedding** *if it satisfies, in place of (c), the stronger condition*
(c′) for every n-ary relation symbol R of $\mathcal{L}$ and every n-tuple $a_1, \ldots, a_n$ of elements of M, $R^{\mathcal{M}}(a_1, \ldots, a_n)$ holds iff $R^{\mathcal{N}}(\alpha(a_1), \ldots, \alpha(a_n))$ holds.

An embedding has to be injective (one-to-one) - we mean condition (c′) to apply also to the symbol for equality. But an injective morphism need not be an embedding if there are relation symbols in the language.

An embedding which is also surjective is an **isomorphism**. *The idea is that an isomorphism between structures is a bijection between their underlying sets which preserves and reflects all the structure, so the structures are "copies" of each other.*

Lemma 10.4. *If $\mathcal{M}, \mathcal{N}, \mathcal{N}'$ are $\mathcal{L}$-structures and $\alpha : \mathcal{M} \to \mathcal{N}$ and $\beta : \mathcal{N} \to \mathcal{N}'$ are homomorphisms, then the composition, $\beta\alpha : \mathcal{M} \to \mathcal{N}'$ is a homomorphism.*

Proof. Repeatedly use that $\beta\alpha(a) = \beta(\alpha(a))$ to check this.

The key lines, with most of the verbiage missing, are:
- $(\beta\alpha)(c^{\mathcal{M}}) = \beta(\alpha(c^{\mathcal{M}})) = \beta(c^{\mathcal{N}})$ (since α is a morphism) $= c^{\mathcal{N}'}$ (since β is a morphism);
- $(\beta\alpha)(f^{\mathcal{M}}(a_1, \ldots, a_n)) = \beta(\alpha(f^{\mathcal{M}}(a_1, \ldots, a_n)) = \beta(f^{\mathcal{N}}(\alpha(a_1), \ldots, \alpha(a_n)))$ (since α is a morphism) $= f^{\mathcal{N}'}(\beta(\alpha(a_1)), \ldots, \beta(\alpha(a_n)))$ (since β is a morphism) $= f^{\mathcal{N}}((\beta\alpha)(a_1), \ldots, (\beta\alpha)(a_n))$;
- and a similar line for relation symbols. □

Lemma 10.5. *(1) If $\mathcal{M}_0$ is a substructure of $\mathcal{M}$ then the inclusion map from $\mathcal{M}_0$ to $\mathcal{M}$ is an embedding.*
(2) If α is an embedding from $\mathcal{N}$ to $\mathcal{M}$ then the image, $\mathrm{im}(\alpha)$, equipped with the $\mathcal{L}$-structure copied from $\mathcal{N}$, is the substructure of $\mathcal{M}$ based on $\mathrm{im}(\alpha)$.

Proof. The sort of proof where the only difficulty is in seeing what needs to be proved.

(1) The inclusion map, $i : M_0 \to M$ is defined by $i(a) = a$ for $a \in M_0$. Clearly this is injective. For condition (a) of the definition of homomorphism, we have $i(c^{\mathcal{M}_0}) = c^{\mathcal{M}_0} = c^{\mathcal{M}}$ (since $\mathcal{M}_0 \leq \mathcal{M}$). For (b), $i(f^{\mathcal{M}_0}(a_1, \ldots, a_n)) = f^{\mathcal{M}_0}(a_1, \ldots, a_n) = f^{\mathcal{M}}(a_1, \ldots, a_n)$ (since $\mathcal{M}_0 \leq \mathcal{M}$) $= f^{\mathcal{M}_0}(i(a_1), \ldots, i(a_n))$. For (c), $R^{\mathcal{M}}(i(a_1), \ldots, i(a_n))$ holds iff $R^{\mathcal{M}}(a_1, \ldots, a_n)$ holds iff $R^{\mathcal{M}_0}(a_1, \ldots, a_n)$ holds (since $\mathcal{M}_0 \leq \mathcal{M}$).

(2) First, we should say what we mean by copying the structure. Set $A = \mathrm{im}(\alpha)$ and define the "copied" $\mathcal{L}$-structure $\mathcal{A}$ on A by: $c^{\mathcal{A}} = \alpha(c^{\mathcal{N}})$; $f^{\mathcal{A}}(b_1, \ldots, b_n) = \alpha(f^{\mathcal{N}}(a_1, \ldots, a_n))$ where $b_i \in A$ and a_i is the (unique!) element of N with $\alpha(a_i) = b_i$; $R^{\mathcal{A}}(b_1, \ldots, b_n)$ holds iff $R^{\mathcal{N}}(a_1, \ldots, a_n)$ holds in $\mathcal{N}$. Simply from this definition, the corestriction of α to a map from N to A becomes an isomorphism of $\mathcal{N}$ with $\mathcal{A}$, so the latter is, indeed, a "copy" of $\mathcal{N}$.

Then we note that $\mathcal{A}$ is the substructure of $\mathcal{M}$ based on $A = \mathrm{im}(\alpha)$. For (a), $c^{\mathcal{A}} = \alpha(c^{\mathcal{N}})$ (by definition) $= c^{\mathcal{M}}$ (since α is a morphism). For (b), $f^{\mathcal{A}}(\alpha(a_1), \ldots, \alpha(a_n)) = \alpha(f^{\mathcal{N}}(a_1, \ldots, a_n))$ (by definition) $= f^{\mathcal{M}}(\alpha(a_1), \ldots, \alpha(a_n))$ (since α is a morphism). Similarly for (c). □

Part (2) of the proof shows that the map α factors as an isomorphism (of $\mathcal{N}$ with the "copied" structure $\mathcal{A}$) followed by an embedding (of the copied structure into $\mathcal{M}$).

The notation $\alpha : \mathcal{M} \simeq \mathcal{N}$ means that α is an isomorphism from $\mathcal{M}$ to $\mathcal{N}$. If there is an isomorphism from $\mathcal{M}$ to $\mathcal{N}$ then we say that $\mathcal{M}$ and $\mathcal{N}$ are **isomorphic** and write $\mathcal{M} \simeq \mathcal{N}$. The conclusion of the exercise below is that this is an equivalence relation on $\mathcal{L}$-structures.

Exercise 10.6. (i) For every structure $\mathcal{M}$ the **identity map** $id_{\mathcal{M}} : \mathcal{M} \longrightarrow \mathcal{M}$, given by $id_{\mathcal{M}}(a) = a$ for every $a \in M$, is an isomorphism from $\mathcal{M}$ to itself.
(ii) If $\alpha : \mathcal{M} \longrightarrow \mathcal{M}'$ and $\alpha' : \mathcal{M}' \longrightarrow \mathcal{M}''$ are isomorphisms then so is their composition $\alpha'\alpha : \mathcal{M} \longrightarrow \mathcal{M}''$ (which is given by $\alpha'\alpha(a) = \alpha'(\alpha(a))$ for every $a \in M$).
(iii) If $\alpha : \mathcal{M} \longrightarrow \mathcal{M}'$ is an isomorphism then the inverse $\alpha^{-1} : \mathcal{M}' \longrightarrow \mathcal{M}$, defined by $\alpha^{-1}(b) = a$ iff $\alpha(a) = b$ ($a \in M, b \in M'$), is an isomorphism.

Proposition 10.7. *Suppose that $\alpha : \mathcal{N} \longrightarrow \mathcal{M}$ is morphism of $\mathcal{L}$-structures. Then, for any term $t(x_1, \ldots, x_n)$ of $\mathcal{L}$ and any $a_1, \ldots, a_n \in N$, we have*

$$\alpha(t^{\mathcal{N}}(a_1, \ldots, a_n)) = t^{\mathcal{M}}\big(\alpha(a_1), \ldots, \alpha(a_n)\big).$$

Proof. This is proved by induction on complexity of terms.

So here goes. Let $t = t(x_1, \ldots, x_n)$ be a term (with its (free) variables *among* $x_1, \ldots, x_n$) and let $a_1, \ldots, a_n \in N$.

Base cases:
(a) t is a constant symbol, c say. Then $t^{\mathcal{N}}$ is $c^{\mathcal{N}}$ and $t^{\mathcal{M}}$ is $c^{\mathcal{M}}$ and, since α is a morphism, $\alpha(c^{\mathcal{N}}) = c^{\mathcal{M}}$, as required.
(b) t is a variable x_i say. Then $t^{\mathcal{N}}(a_1, \ldots, a_n)$ is a_i and $t^{\mathcal{M}}(\alpha(a_1), \ldots, \alpha(a_n))$ is $\alpha(a_i)$ as required.

Induction cases:
(c) Suppose that t is $f(t_1, \ldots, t_m)$ where f is an m-ary function symbol and the terms $t_1, \ldots, t_m$ have (free) variables among $x_1, \ldots, x_n$. Write $\overline{x}$ for $(x_1, \ldots, x_n)$ and $\overline{a}$ for $(a_1, \ldots, a_n)$. Then $t^{\mathcal{N}}(\overline{a})$ is $\big(f(t_1, \ldots, t_m)\big)^{\mathcal{N}}(\overline{a})$ which, by the definition of interpretation of terms, is $f^{\mathcal{N}}\big(t_1^{\mathcal{N}}(\overline{a}), \ldots, t_m^{\mathcal{N}}(\overline{a})\big)$. So $\alpha\big(t^{\mathcal{N}}(\overline{a})\big) = \alpha\big(f^{\mathcal{N}}(t_1^{\mathcal{N}}(\overline{a}), \ldots, t_m^{\mathcal{N}}(\overline{a}))\big) = f^{\mathcal{M}}(\alpha(t_1^{\mathcal{N}}(\overline{a})), \ldots, \alpha(t_m^{\mathcal{N}}(\overline{a})))$ (since α is a morphism) $= f^{\mathcal{M}}\big(t_1^{\mathcal{M}}(\alpha(\overline{a})), \ldots, t_m^{\mathcal{M}}(\alpha(\overline{a}))\big)$ (by the induction hypothesis) $= \big(f(t_1, \ldots, t_m)\big)^{\mathcal{M}}(\alpha(\overline{a}))$ which is just $t^{\mathcal{M}}\big(\alpha(\overline{a})\big)$, as required.

Therefore, by induction on complexity/construction of terms, we have the statement of the proposition. □

10.3 Elementary substructures and elementary embeddings

Consider the structure $(\mathbb{Z}; +, 0)$ (the group of integers under addition). It's easy to check that the set, $2\mathbb{Z}$, of even integers is the basis for a substructure (the subgroup of even integers) of $(\mathbb{Z}, +, 0)$. But notice that if $\varphi(x)$ is the formula $\exists y\, (x = y + y)$ then we have $\mathbb{Z} \models \varphi(2)$ but $2\mathbb{Z} \models \neg\varphi(2)$. Thus, the properties of an element of $2\mathbb{Z}$ may be different depending on whether we regard it as an element of $2\mathbb{Z}$ or of the larger set $\mathbb{Z}$.

If we have a structure $\mathcal{M}$ and want to embed it into a larger structure $\mathcal{M}'$ we may well want the properties of elements of $\mathcal{M}$ to be the same whether we think of them as elements of $\mathcal{M}$ or of $\mathcal{M}'$. For instance, if we extend the reals $\mathbb{R}$ to a structure with infinitesimals then we wouldn't want the properties of the various real numbers to have changed in doing so. The conclusion is that the property of being a substructure is not strong enough for some purposes.

Definition 10.8. *Suppose that $\mathcal{N}$ and $\mathcal{M}$ are $\mathcal{L}$-structures with $\mathcal{N} \leq \mathcal{M}$, that is $\mathcal{N}$ a substructure of $\mathcal{M}$. We say that $\mathcal{N}$ is an* **elementary substructure** *of $\mathcal{M}$ (and that $\mathcal{M}$ is an* **elementary extension** *of $\mathcal{N}$), writing $\mathcal{N} \prec \mathcal{M}$ if:*
$(*)$ for every $\mathcal{L}$-formula $\varphi(\bar{x})$ and tuple $\bar{a}$ of elements of N we have $\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\bar{a})$.
Convention: *we do not usually explicitly state the condition that tuples should match - so in the above, were assuming that $l(\bar{a}) = l(\bar{x})$ where "l" denotes the length of a tuple*

Lemma 10.9. *(1) $\mathcal{N} \prec \mathcal{M}$ implies that for every sentence σ of $\mathcal{L}$ we have $\mathcal{N} \models \sigma$ iff $\mathcal{M} \models \sigma$ (then we say that $\mathcal{N}$ and $\mathcal{M}$ are elementarily equivalent, and write $\mathcal{N} \equiv \mathcal{M}$).*
(2) $\mathcal{N} \prec \mathcal{M}$ and $\mathcal{M} \prec \mathcal{M}'$ together imply $\mathcal{N} \prec \mathcal{M}'$.

Proof. Immediate from the definitions. □

Exercise 10.10. (i) If $\mathcal{N} \leq \mathcal{M}$ then $\mathcal{N} \prec \mathcal{M}$ iff for every formula $\varphi(\bar{x})$ of $\mathcal{L}$ and every (matching) tuple $\bar{a}$ from N, if $\mathcal{N} \models \varphi(\bar{a})$ then $\mathcal{M} \models \varphi(\bar{a})$ (and similarly the other way round: it's enough to test that if $\mathcal{M} \models \varphi(\bar{a})$ then $\mathcal{N} \models \varphi(\bar{a})$, provided of course that $\overline{a}$ comes from N).

(ii) Suppose that $\mathcal{M} \leq \mathcal{N}$ are $\mathcal{L}$-structures with $\mathcal{M} \equiv \mathcal{N}$. Prove that $\mathcal{M} \prec \mathcal{N}$ iff for every $\varphi(\bar{x}) \in \mathcal{L}$ we have $\varphi(\mathcal{M}) = M^{l(\bar{x})} \cap \varphi(\mathcal{N})$.

Exercise 10.11. Let $\mathcal{L} = \mathcal{L}_0 \vee \{P(-)\}$ where P is a unary predicate symbol. Let $\mathcal{M} = (M; P^{\mathcal{M}})$ with M infinite be such that $P^{\mathcal{M}}$ is an infinite, coinfinite (i.e. $M \setminus P^{\mathcal{M}}$ is infinite) subset of M. Let N be any subset of M such that both $N \cap P^{\mathcal{M}}$ and $N \cap (M \setminus P^{\mathcal{M}})$ are infinite. Prove that the substructure $\mathcal{N} = (N; P^{\mathcal{N}} = N \cap P^{\mathcal{M}})$ based on N is an elementary substructure of $\mathcal{M}$. You can use 10.17, although there may be more direct arguments.

Definition 10.12. *An embedding $\alpha : \mathcal{N} \longrightarrow \mathcal{M}$ of $\mathcal{L}$-structures is an* **elementary embedding** *if for every $\mathcal{L}$-formula $\varphi(\bar{x})$ and tuple $\bar{a}$ from N we have $\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\alpha\bar{a})$. Here $\alpha\bar{a}$ is short for $(\alpha(a_1), \dots, \alpha(a_n))$ if $\overline{a} = (a_1, \dots, a_n)$.*

Lemma 10.13. *(cf. 10.5) (1) If $\mathcal{M}_0$ is an elementary substructure of $\mathcal{M}$ then the inclusion map from $\mathcal{M}_0$ to $\mathcal{M}$ is an elementary embedding.*
(2) If α is an elementary embedding from $\mathcal{N}$ to $\mathcal{M}$ then the image, $\mathrm{im}(\alpha)$, equipped with the $\mathcal{L}$-structure copied from $\mathcal{N}$ (= the $\mathcal{L}$-structure induced from $\mathcal{M}$ by 10.5(2)) is an elementary substructure of $\mathcal{M}$.

Proof. (Compare with 10.5.)
(1) If $\varphi(\overline{x})$ is a formula and $\overline{a}$ is from M_0 then we have $\mathcal{M} \models \varphi(\overline{a})$ iff $\mathcal{M}_0 \models \varphi(\overline{a})$ but, since $i(\overline{a}) = \overline{a}$, this is iff $\mathcal{M}_0 \models \varphi(i(\overline{a}))$, hence i is an elementary embedding.
(2) Write $\mathcal{A}$ for the structure copied from $\mathcal{N}$ on underlying set $\mathrm{im}(\alpha)$. From 10.5 we already have that the inclusion of $\mathcal{A}$ in $\mathcal{M}$ is an embedding of $\mathcal{L}$-structures. So, let $\varphi(\overline{x})$ be a formula and $\overline{b} = \alpha(\overline{a})$ be from $A = \mathrm{im}(\alpha)$. Then $\mathcal{M} \models \varphi(\overline{b})$, that is, $\mathcal{M} \models \varphi(\alpha(\overline{a}))$ iff $\mathcal{N} \models \varphi(\overline{a})$ (since α is an elementary embedding), iff $\mathcal{A} \models \varphi(\overline{b})$ (since the $\mathcal{L}$-structure on $\mathcal{A}$ is copied from $\mathcal{N}$ plus 10.14 below), as required. □

Proposition 10.14. *Suppose that $\alpha : \mathcal{N} \longrightarrow \mathcal{M}$ is an isomorphism. Let $\varphi(\bar{x}) \in \mathcal{L}$ and let $\bar{a}$ be from N (that is, it is a tuple of elements of N). Then $\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\alpha\bar{a})$.*

Proof. The proof is by induction on complexity of formulas. First, one needs the relevant statement for terms, which is that, for any term $t(x_1, \ldots, x_n)$ of $\mathcal{L}$ and any $a_1, \ldots, a_n \in N$ we have $\alpha(t^{\mathcal{N}}(a_1, \ldots, a_n)) = t^{\mathcal{M}}(\alpha(a_1), \ldots, \alpha(a_n))$. That was proved in 10.7 above (and was shown by induction on the construction of terms).

So let's begin; let $\varphi = \varphi(\overline{x})$ be a formula. We prove the statement "$\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\alpha\bar{a})$, for all tuples $\overline{a}$ of elements of N" by induction on complexity of φ (you'll see why we need to say it this way - with non-fixed $\overline{a}$ - when we deal with the $\exists$ induction step).

Base cases: φ is an atomic formula.
(a) $\varphi(\overline{x})$ has the form $t_1(\overline{x}) = t_2(\overline{x})$ for some terms t_1, t_2. Then $\mathcal{N} \models \varphi(\overline{a})$ iff $t_1^{\mathcal{N}}(\overline{a}) = t_2^{\mathcal{N}}(\overline{a})$ (by definition of satisfaction of formulas) iff $\alpha({t_1}^{\mathcal{N}}(\overline{a})) = \alpha(t_2^{\mathcal{N}}(\overline{a}))$ (one direction since α is a map, the other since α is injective), iff $t_1^{\mathcal{M}}(\alpha(\overline{a})) = t_2^{\mathcal{M}}(\alpha(\overline{a}))$ (by 10.7) iff $\mathcal{M} \models \varphi(\alpha(\overline{a}))$ (again by definition of satisfaction of formulas).
(b) φ has the form $R(\overline{x})$ where R is a relation symbol. Then $\mathcal{N} \models \varphi(\overline{a})$ iff $R^{\mathcal{N}}(\overline{a})$ holds that is, iff $\overline{a} \in R^{\mathcal{N}}$ (definition of satisfaction of formulas) iff $R^{\mathcal{M}}(\alpha(\overline{a}))$ holds (since α is an embedding) iff $\mathcal{M} \models \varphi(\alpha(\overline{a}))$.

Induction cases (again, each case refers back to the definition of satisfaction of formulas in structures, but I won't keep writing it in):
(c) φ is $\psi_1 \wedge \psi_2$, where we assume, inductively, that the statement holds for ψ_1 and ψ_2. Then $\mathcal{N} \models \varphi(\overline{a})$ iff $\mathcal{N} \models \psi_1(\overline{a})$ and $\mathcal{N} \models \psi_1(\overline{a})$ iff $\mathcal{M} \models \psi_1(\alpha(\overline{a}))$ and $\mathcal{M} \models \psi_2(\alpha(\overline{a}))$ (induction hypothesis) iff $\mathcal{M} \models \varphi(\alpha(\overline{a}))$.
(d) φ is $\neg\psi$ where we assume, inductively, that the statement holds for ψ. Then $\mathcal{N} \models \varphi(\overline{a})$ iff $\mathcal{N} \not\models \psi(\overline{a})$ iff $\mathcal{M} \not\models \psi(\alpha(\overline{a}))$ iff $\mathcal{M} \models \varphi(\alpha(\overline{a}))$.
(e) φ has the form $\exists x_0 \, \psi(x_0, x_1, \ldots x_n)$ where we assume, inductively, that the statement holds for ψ. Then $\mathcal{N} \models \varphi(\overline{a})$ iff there is $a_0 \in N$ such that $\mathcal{N} \models \psi(a_0, a_1, \ldots, a_n)$. This implies that $\mathcal{M} \models \psi(\alpha(a_0), \alpha(a_1), \ldots, \alpha(a_n))$ by our inductive hypothesis on ψ and hence that $\mathcal{M} \models \varphi(\alpha(\overline{a}))$. Conversely, if $\mathcal{M} \models \varphi(\alpha(\overline{a}))$, then there is some $b_0 \in M$ such that $\mathcal{M} \models \psi(b_0, \alpha(a_1), \ldots, \alpha(a_n))$ and then, because α is surjective, there is $a' \in N$ such that $b_0 = \alpha(a')$ and so, from $\mathcal{M} \models \psi(\alpha(a'), \alpha(a_1), \ldots, \alpha(a_n))$ we deduce, by the inductive hypothesis on ψ, that $\mathcal{N} \models \psi(a', a_1, \ldots, a_n)$, hence that $\mathcal{N} \models \varphi(\overline{a})$, as required. $\square$

Corollary 10.15. *$\mathcal{M} \simeq \mathcal{N}$ implies $\mathcal{M} \equiv \mathcal{N}$.*

Proof. Apply 10.14 with φ an arbitrary sentence. $\square$

Notice, with reference to the example at the beginning of this section, that the groups of integers and of even integers are isomorphic, hence elementarily equivalent, but the latter is not an elementary substructure of the former. By a **quantifier-free** formula we mean one which has no occurrences of any quantifier.

Exercise 10.16. These are variations on the proof of 10.14.
(i) Suppose $\alpha : \mathcal{N} \to \mathcal{M}$ is a morphism of $\mathcal{L}$-structures, $\varphi(\overline{x})$ is a positive quantifier-free formula (quantifier-free and built using only $\wedge$ and $\vee$) and $\overline{a}$ is

from N. Show that $\mathcal{N} \models \varphi(\overline{a})$ implies $\mathcal{M} \models \varphi(\alpha\overline{a})$. Give an example to show that the converse implication is false.
(ii) Suppose $\alpha : \mathcal{N} \to \mathcal{M}$ is an embedding of $\mathcal{L}$-structures, $\varphi(\overline{x})$ is an existential formula (a string of existential quantifiers prefixing a quantifier-free formula) and $\overline{a}$ is from N. Show that $\mathcal{N} \models \varphi(\overline{a})$ implies $\mathcal{N} \models \varphi(\alpha\overline{a})$. Give an example to show that the converse implication is false.
(iii) Suppose $\alpha : \mathcal{N} \longrightarrow \mathcal{M}$ is an embedding and $\varphi(\bar{x}) \in \mathcal{L}$ is quantifier-free and $\bar{a}$ is from N; show that $\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\alpha(\bar{a}))$.
(iv) Deduce that, in particular, if $\mathcal{N}$ is a substructure of $\mathcal{M}$ and if $\varphi(\bar{x}) \in \mathcal{L}$ is quantifier-free and $\bar{a}$ is from N, then $\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\alpha(\bar{a}))$.

10.4 A test for elementary substructure

Theorem 10.17. *(Tarski's Lemma) Suppose that $\mathcal{N} \leq \mathcal{M}$. Then $\mathcal{N} \prec \mathcal{M}$ iff for all $\psi(x_1, \ldots, x_n, y) \in \mathcal{L}$ and $a_1, \ldots, a_n \in N$ if $\mathcal{M} \models \exists y\, \psi(\bar{a}, y)$ then there is $b \in N$ with $\mathcal{M} \models \psi(\bar{a}, b)$.*

Proof. The direction "$\Rightarrow$" is direct from the definition of elementary embedding. For the direction "$\Leftarrow$" we show, by induction on the complexity of φ, that for all formulas $\varphi \in \mathcal{L}$ and tuples $\bar{a}$ from N we have $\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\bar{a})$.

In detail:
($\Rightarrow$) With notation as in the statement, suppose that $\mathcal{N} \prec \mathcal{M}$ and that we have $\mathcal{M} \models \exists y\, \psi(\bar{a}, y)$. Since $\overline{a}$ is from N, by definition of $\prec$ we have $\mathcal{N} \models \exists y\, \psi(\bar{a}, y)$, so $\mathcal{N} \models \psi(\bar{a}, b)$ for some $b \in N$. Then, again by definition of $\prec$, it follows that $\mathcal{M} \models \psi(\bar{a}, b)$, as required.
($\Leftarrow$) Suppose that the right-hand condition is satisfied; we have to show $\mathcal{N} \prec \mathcal{M}$, that is, that for all formulas $\varphi \in \mathcal{L}$ and tuples $\bar{a}$ from N we have $\mathcal{N} \models \varphi(\bar{a})$ iff $\mathcal{M} \models \varphi(\bar{a})$. This is proved by induction on complexity of φ, though a lot of the argument has been seen already. For, since $\mathcal{N}$ is a substructure of $\mathcal{M}$, the argument for 10.16(iii) already covers the base cases and the induction steps involving the propositional connectives $\neg$ and $\wedge$ (we can't use the statement of that exercise because, in building formulas, these connectives can be applied to formulas which already contain quantifiers). So it's only the $\exists$ case that we need to detail here (and remember that we only need to show one direction of the implication in the definition of $\prec$ - that's Exercise 10.10(i)).

Suppose, then, that φ has the form $\exists y\, \psi(\overline{x}, y)$ where we already have the induction statement for ψ. Suppose that $\mathcal{M} \models \varphi(\overline{a})$, that is $\mathcal{M} \models \exists y\, \psi(\overline{a}, y)$. By assumption there is $b \in N$ such that $\mathcal{M} \models \psi(\overline{a}, b)$. By the inductive assumption, it follows that $\mathcal{N} \models \psi(\overline{a}, b)$, hence that $\mathcal{N} \models \varphi(\overline{a})$, as required. $\square$