2009 lectures on Set Theory

Thomas Forster

April 3, 2010

Contents

How to use this file:

Everything that is examinable in 2009/10 is to be found either in this file or in files linked from my home page or in Logic, Induction and Sets [15]. (This last contains mainly the Part II prerequisites.) There is much in this file that is not examinable, and that is generally in small type. The files `fundamentalsequence.pdf` and `axiomsofsettheory.pdf` contain useful material, and they are tutorials rather than research articles. However it is not clear that there is anything in either of them that is hard enough to be examinable.

0.1 Foreword

The Part III set theory course has been slimmed down to make room for extra material. One of the new topics is the theory of computable functions. This used to be lectured at Part II, and as a result the two blue CUP books—[15] and [23]—based on the Part II syllabus both cover it. There is nothing in these notes on the theory of computable functions: I shall lecture it from [15]. I might also include a very high-speed treatment of Basic proof theory. Like the theory of computable functions this is the kind of thing that really belongs at undergraduate level, so my treatment really will be very high speed indeed. The notes from which I lecture it will not be in this file but in [16]. If students prefer it, I am open to the suggestion that this material be not lectured.

There are plenty of mainstream books on set theory at this level and above: Van Dalen/Doets/de Swart is still in many college libraries (used to be recommended for Part II Logic). Jech and Kunen both have books called 'Set theory'. (Jech [28] is Academic Press, Kunen [26] is North Holland). Drake and Singh Intermediate set theory [10] looks the right level; Hajnal/Hamburger [19] "Set Theory', looks good too, tho' it does have a strongly combinatorial emphasis quite unlike the drift of these notes. I know it sounds greedy and naff to recommend one's own book, but there is an excuse! This course is designed to be a sequel to the Part II course that most of you have sat through. 'Logic, Induction and Sets' [15] arose from lecture notes written to the specifications set down by the faculty board for that course. It is more elementary than any of the books mentioned above. Recently Bell and Slomson [3] has been reissued by Dover. I recommend it strongly.

There are no textbooks on BQO theory—yet! Monika Seisenberger and i are thinking of writing one.

I think the course is entitled "Set theory". I think however that it should really be entitled "Set Theory and Logic" since it casts its net wider than just to encompass set theory.

Further stuff to fit in: reflection and non-finite axiomatisability of any theory extending ZF.

Stuff to fit in on finite axiomatisablity

Can anyone on this list state—and cite!—a theorem to the effect that to any recursively axiomatisable theory T in a language L there is a finitely axiomatisable theory T' in a suitable language L' with T' equivalent to T' in some very strong sense. Somebody must have proved a rigorous version of this, and I am hoping that listmembers will know who and how.

The basic references here are [25] and [5]. There is a very clear review of both papers by Makkai in the JSL, **36** (1971), pp. 334-335 that also provides a sketch of the proof. Kleene's main result is:

> If T is a recursively axiomatizable theory that has only infinite models, then it has a finitely axiomatized conservative extension.

The fundamental idea of the proof is to formalize the inductive clauses of the truth definition for T.

Thanks to Alasdair Urquhart

0.2 Notation and definitions

I seem to have got into the habit of using `teletype` font for functions with a computer science flavour. `fst` and `snd` extract the two components of an ordered pair. `hd` gives you the first member of a list, and `tl` gives you the tail; `last` gives you the last element of a sequence, `len` gives you the length of a sequence, and `butlast` gives you everything but the last element..

DC and DC_λ

DC_λ says that whenever $\langle X, \leq \rangle$ is a poset in which every chain of length $< \lambda$ has a proper end-extension then it has a path of length λ. DC is DC_ω. It's more usually expressed in the form: If $\langle X, R \rangle$ is a binary structure such that for every $x \in X$ there is y in X with $R(x, y)$ then there is an infinite R-chain. (And yes, X might be finite and the chain might have repeats!)

Definition 1. *The* **relativisation** ϕ^X *of a formula* ϕ *to a term* X *is the result of replacing every universal quantifier '*$(\forall x)($*' in* ϕ *by '*$(\forall x \in X)($*' and every existential quantifier '*$(\exists x)($*' in* ϕ *by '*$(\exists x \in X)($*'.)*

1. Rectypes and Wellfoundedness. Does every wellfounded relation arise from the engendering relation of a rectype? Not clear. Finite antichain condition seems to be relevant here: $\leq^+$ is wellfounded if $\leq$ is WQO.

2. recursion theory and λ-calculus

   `http://ellemose.dina.kvl.dk/~sestoft/lamreduce/`

Chapter 1

Proof theory and recursive function theory

1.1 Completeness of LPC

Theorem 2. *Every theory in a countable language can be extended to a complete theory.*

Proof: Suppose T is a theory in a language $\mathcal{L}(T)$ which is countable. Then we count the formulæ in $\mathcal{L}(T)$ as $\phi_1, \phi_2 \ldots$ and define a sequence of theories T_i as follows.

$T_0 = T$; and thereafter

T_{i+1} is to be T_i if T_i decides ϕ_i and is $T_i \cup \{\phi_i\}$ otherwise.

$\in$-terms

For any theory T we can always add constants to $\mathcal{L}(T)$ to denote witnesses to $\exists^n$ sentences in T.

Suppose $T \vdash (\exists x)(F(x))$. There is nothing to stop us adding to $\mathcal{L}(T)$ a new constant symbol 'a' and adding to T an axiom $F(a)$. Clearly the new theory will be consistent if T was. Why is this? Suppose it weren't, then we would have a deduction of the `false` from $F(c)$. Therefore, by $\rightarrow$-introduction, we would have a proof of $\neg F(c)$. But then, by UG, we would have a proof of $\forall x \neg F(x)$ in T. But T also proves $(\exists x)(F(x))$, so T was inconsistent.

Notice that nothing about the letter 'a' that we are using as this constant tells us that a a thing which is F. We could have written the constant 'a_F' or something like that. Strictly it shouldn't matter: variables and constant symbols do not have any internal structure that is visible to the language, and the 'F' subscript provides a kind of spy-window available to anyone *mentioning* the language, but not to anyone merely *using* it. The possibility of writing out novel constants in suggestive ways like this will be useful later.

That is to say that for any F with one free variable we can invent a constant whose job it is to denote an object which has property F as long as anything does. If there is indeed a thing which has F then this constant can denote one of them. If there isn't it doesn't matter what it denotes. The appeal to the law of excluded middle in this patter is actually necessary.

This constant is often written $(\epsilon x)F(x)$. Since it points to something that has F as long as there is something that has F, we can see that

$$(\exists x)(F(x)) \qquad \text{and} \qquad F((\epsilon x)F(x))$$

are equivalent.

Notice that this gives us an equivalence between a formula that definitely belongs to LPC (co's it has a quantifier in it) and something that appears not to. Hilbert was very struck by this fact, and thought he had stumbled on an important breakthrough: a way of reducing predicate logic to propositional logic. Sadly he hadn't, but the ϵ-terms are useful gadgets all the same, as we are about to see.

THEOREM 3. *Every consistent theory in a countable language has a model.*

Proof:

Let T_1 be a consistent theory in a countable language $\mathcal{L}(T_1)$.

We do the following things

1. Add axioms to T_1 to obtain a complete extension;
2. Add ϵ terms to the language.

Notice that when we add ϵ-terms to the language we add new formulæ: if '$(\epsilon x)F(x))$' is a new ϵ-term we have just added then '$G((\epsilon x)F(x)))$' is a new formula, and T_1 doesn't tell us whether it is to be true or to be false. That is to say $\mathcal{L}(T_1)$ doesn't contain '$(\epsilon x)F(x)$' or '$G((\epsilon x)F(x)))$'. Let $\mathcal{L}(T_2)$ be the language obtained by adding to $\mathcal{L}(T_1)$ the expressions like '$(\epsilon x)F(x)$' and '$G((\epsilon x)F(x)))$'.

We extend T_1 to a new theory in $\mathcal{L}(T_2)$ that decides all these new formulæ we have added. This gives us a new theory, which we will—of course—call T_2. Repeat and take the union of all the theories T_i we obtain in this way: call it T_∞. (Easy to see that all the T_i are consistent—we prove this by induction).

It's worth thinking about what sort of formulæ we generate. We added terms like $(\epsilon x)(F(x))$ to the language of T_1. Notice that if H is a two-place predicate in $\mathcal{L}(T)$ then we will find ourselves inventing the term $(\epsilon y)H(y, (\epsilon x)F(x))$ which is a term of—one might say—*depth* 2. And there will be terms of depth 3, 4 and so on as we persist with this process. All atomic questions about ϵ terms of depth n are answered in T_{n+1}.

T_∞ is a theory in a language $\mathcal{L}_\infty$, and it will be complete. The model $\mathfrak{M}$ for T_∞ will be the structure whose carrier set is the set of ϵ terms we have generated *en route*. All questions about relations between the terms in the domain are answered by T_∞. Does this make $\mathfrak{M}$ into a model of T? We will establish the following:

LEMMA 4. $\mathfrak{M} \models \phi(t_1, \ldots t_n)$ *iff* $T_\infty \vdash \phi(t_1, \ldots t_n)$

Proof: We do this by induction on the logical complexity of ϕ. When ϕ is atomic this is achieved by stipulation. The induction step for propositional connectives is straightforward. (Tho' for one direction of the '$\vee$' case we need to exploit the fact that T_∞ is complete, so that if it proves $A \vee B$ then it proves A or it proves B.)

The remaining step is the induction step for the quantifiers. They are dual, so we need consider only one. We consider only the hard direction.

Suppose $\mathfrak{M} \models (\forall x)\phi(x, t_1, \ldots t_n)$. Then $\mathfrak{M} \models \phi(t_0, t_1, \ldots t_n)$ for all terms t_0. In particular it must satisfy it even when $t_0 = (\epsilon x)(\neg\phi(x, t_1, \ldots t_n))$, which is to say

$$\mathfrak{M} \models \phi((\epsilon x)(\neg\phi(x, t_1, \ldots t_n)), t_1, \ldots t_n)$$

So, by induction hypothesis we must have

$$T_\infty \vdash \phi((\epsilon x)(\neg\phi(x, t_1, \ldots t_n)), t_1, \ldots t_n)$$

whence of course $T_\infty \vdash (\forall x)\phi(x, t_1, \ldots t_n)$.

■

This completes the proof of theorem 3.

Observe the essential rôle played in this proof by the ϵ terms.

Notice that the model is countable as long as the original language was. If we had uncountably many constants to start with then of course the model will be uncountable. It will also be uncountable if we have uncountably many predicate letters or function letters. But not otherwise.

This is a result of fundamental importance. Any theory that is not actually self-contradictory is a description of *something*. It's important that this holds only for first-order logic. It does not work for second-order logic, and this fact is often overlooked.

Fit this in somewhere: take a formula of the Σ^1 fragment of second-order logic. Delete the existential quantifiers. The result is a formula in 1st order logic with function letters. If it is refutable then so was the Σ^1 formula we started with. So there is a refutation procedure for the Σ^1 fragment of second-order logic.

Similarly there is a refutation procedure for the set of formulæ true in alll finite structures.

1.2 Do some ordinals here

Chapter 2

Some Model Theory

I will try to adhere to the habit of using $\mathfrak{FRAKTUR}$ font[1] for letters denoting structures and the corresponding upper-case Roman letter for the carrier set. I probably won't manage it!

Definition 5. *The* **Skolem Hull** *of a structure* $\mathfrak{M}$ *is what one obtains as follows. For each sentence* $\exists x \phi(x)$ *true in* $\mathfrak{M}$ *pick the first such* x*. For each sentence* $\forall x \exists y \psi(x, y)$ *true in* $\mathfrak{M}$ *let* f_ψ *send each* x *to the first* y *such that* $\psi(x, y)$*. Close under these operations. The result is the Skolem Hull.*

Of course we can generalise this by requiring that the Skolem hull should contain some specified things to start with. It's another recursive datatype.

Definition 6. *An embedding* $i : \mathfrak{M} \to \mathfrak{N}$ *is* Γ-**elementary** *iff for all* $\phi \in \Gamma$ $\mathfrak{M} \models \phi(x_1 \ldots x_n)$ *iff* $\mathfrak{N} \models \phi(i(x_1) \ldots i(x_n))$

Very important notion, this. For example, end-extensions are elementary for formulæ in which all quantifiers are restricted!

If i is Γ-elementary where Γ is the set of all formulæ we say i is just plain elementary.

Remember elementary equivalence?!

Some examples:

inclusion embedding from the rationals-as-an-ordered set into the reals (ditto) is elementary. Not as an ordered field.

That example (rats into reals) is one where the embedding is in some sense canonical. There is no suggestion that this is so in general.

The simplest application of elementary embeddings known to me is the usual proof that classical monadic predicate logic is decidable.

Remark 7. *Classical monadic predicate logic is decidable.*

Proof: Suppose we have a monadic formula Φ, and let $\mathfrak{M}$ be a model. Φ contains only finitely many monadic predicate letters, say $\psi_1 \ldots \psi_i$. Let $\mathcal{L}_\Psi$ be

[1] Often called 'Gothic' by the ignorant. The Goths had a different alphabet (and a different language!) not just a different font.

the language with these monadic predicates and no other predicates or function letters. The various ψ divide the M into 2^i classes in the obvious way: a typical class looks like $\{x : \psi_1(x) \wedge \neg\psi_2(x) \wedge \ldots\}$. Any selection set for this partition gives a submodel of $\mathfrak{M}$ for which (we prove by induction on the recursive datatype of $\mathcal{L}_\Psi$-formulæ) the inclusion embedding is $\mathcal{L}_\Psi$-elementary. The submodel is finite and it will only take a finite time to check the truth value in it of any formula. ■

There is another proof of this using elementary formulæ which i probably won't get round to lecturing but which you can probably find yourselves. See the appendix on p. 143. Notice that predicate calculus itself is undecidable. It's *semidecidable* because of the completeness theorem: every valid formula has a proof, and a proof is a finite object which can be found by systematic exhaustive search. If every falsifiable formula of LPC had a finite countermodel we would have an analogous demonstration that the set of falsifiable sentences is semidecidable, and thereby a decision procedure. However altho' the negation of DLO is a falsifiable formula it has no finite countermodel (all dense linear orders are infinite!) so this strategy doesn't work.

2.1 Products

2.1.1 Direct Products and Reduced Products

If $\{\mathfrak{A}_i : i \in I\}$ is a family of structures, we define the product

$$\prod_{i \in I} \mathfrak{A}_i$$

to be the structure whose carrier set is the set of all functions f defined on the index set I such that $(\forall i \in I)(f(i) \in A_i)$ and the relations of the language are interpreted *"pointwise"*: the product believes $f\ R\ g$ iff $(\forall i \in I)(f(i)\ R\ g(i))$.

The $\{\mathfrak{A}_i : i \in I\}$ are said to be the *factors* of the product $\prod_{i \in I} \mathfrak{A}_i$

For this operation to make sense it is of course necessary that all the $\mathfrak{A}_i$ should have the same signature!

Products are nice in various ways. They preserve horn sentences. What do we mean by "preserve"?

Definition 8. *Let Γ be a class of formulæ. Products* **preserve** Γ *if whenever $\prod_{i \in I} A_i$ is a product of a family $\{\mathfrak{A}_i : i \in I\}$ and $\phi \in \Gamma$ then*

$$\prod_{i \in I} \mathfrak{A}_i \models \phi \text{ as long as } (\forall i \in I)(\mathfrak{A}_i \models \phi).$$

In these circumstances we also say that ϕ is **preserved**.

By definition of product, products preserve atomic formulæ. Clearly they also preserve conjunctions of anything they preserve, and similarly universal quantifications over things they preserve.

DEFINITION 9. *Horn formulæ Horn clauses: disjunctions of atomics and negatomics at most one of which is atomic, closed under* $\forall$ *and* $\wedge$.

EXERCISE 1. *Verify that products preserve Horn formulæ*

(This was proved by a man named 'Horn'!) However they do not always preserve formulæ containing $\vee$ or $\neg$. How so? If ϕ is preserved, then the product will fail to satisfy it if even *one* of the factors does not satisfy it (c.f. Genesis [19:23-33] where not even one righteous man is enough to save the city) but all the rest do. (The product is not righteous unless *all* its factors are). In these circumstances the product $\models \neg\phi$ but it is not the case that all the factors $\models \neg\phi$. As for $\vee$, if ϕ and ψ are preserved, it can happen that $\phi \vee \psi$ is not, as follows. If half the factors satisfy ϕ and half satisfy ψ, then they all satisfy $\psi \vee \phi$. Now the product will satisfy $\psi \vee \psi$ iff it satisfies one of them. But in order to satisfy one of them, that one must be true at *all* the factors, and by hypothesis it is not. Something similar happens with the existential quantifier.

Intersection-closed properties

We say that a property ϕ of $R_1, \ldots R_n$ is Horn if $\phi(\vec{R})$ is a Horn clause. The properties *symmetric reflexive*, and *transitive* are all Horn. The idea is an important one because for Horn properties one has an idea of **closure**. Take transitivity for example. R is transitive iff

$$(\forall xyz)((R(x,y) \wedge R(y,z)) \to R(x,z))$$

That is to say, the graph of R (R thought of as a set of ordered pairs) is closed under the operation that accepts the pair $\langle x,y \rangle$ and the pair $\langle y,z \rangle$ as inputs and outputs the pair $\langle x,z \rangle$.

As usual, when one has a set and an operation that can be applied to its members, one has a notion of canonical unique closure of that set under that operation. The point about horn properties is that in the horn clause there are lots of premisses (all positive, each saying—as it were—that the relation contains a certain tuple) and one conclusion, saying that in that case the relation contains this other tuple.

So each transitive relation is an inductively defined set!

Reduced products

Given a filter F over the index set, we can define $f \sim_F g$ on elements of the product if $\{i \in I : f(i) = g(i)\} \in F$. Then we *either* take this $\sim_F$ to be the interpretation of '=' in the new product we are defining, keeping the elements of the carrier set of the new product the same as the elements of the old *or* we take the elements of the new structure to be equivalence classes of functions under $\sim$. These we will write $[g]_{\sim_F}$ or $[g]$ if there is no ambiguity.

This new object is denoted by the following expression:

$$(\prod_{i \in I} \mathfrak{A}_i)/F$$

Similarly we have to revise our interpretation of atomic formulæ so that

$$(\prod_{i \in I} \mathfrak{A}_i)/F \models \phi(f_1, \ldots f_n) \text{ iff } \{i : \phi(f_1(i), \ldots f_n(i)\} \in F.$$

It may be worth bearing in mind that to a certain extent the choice between thinking of elements of the carrier set of the reduced product as the set of $\sim_F$-equivalence classes *versus* thinking of them as the functions is a real one and might matter. I have proceeded here on the basis that the carrier set is the set of $\sim_F$-equivalence classes because that seems more natural. However, in principle there are set-existence issues involved in thinking of a product this way—how do we know that the $\sim_F$-equivalence classes are sets?—so we want to keep alive in our minds the possibility of doing things the second way. This will matter when we come to consider reduced products where the factor structures are proper classes (= have carrier sets that are proper classes). In practice these issues are usually swept under the carpet; this is a safe strategy only because it is in fact possible to sort things out properly! There is of course also the possibility of picking representatives from the equivalence classes, possibly by means of AC.

(For those of a philosophical turn of mind, there is an interesting contrast here with the case of quotient structures like, say, integers mod p. I have the impression that on the whole mathematicians do not think of integers mod p as sets of integers, nor as integers equipped with a nonstandard equality relation, but rather think of them as objects of a new kind. These reflections may have significance despite not really belonging to the study of *mathematics*: the study of *how we think about mathematics* is important too.)

The reason for proceeding from products to reduced products was to complicate the construction and hope to get more things preserved. In fact nothing exciting happens (we still have the same trouble with $\vee$ and $\neg$) unless the filter we use is ultra. Then everything comes right.

At this stage we assume ultrafilters, why they exist etc. This is all in Logic induction and sets. If you aren't happy about it see me.

Worth making the point that the collection of filters on X is a complete poset and the collection of proper filters is merely a chain-complete poset.

2.1.2 Ultraproducts and Łoś's theorem

THEOREM 10. *(Łoś's theorem)*
Let $\mathcal{U}$ be an ultrafilter on I. For all expressions $\phi(f, g, h \ldots)$,

$$(\prod_{i \in I} \mathfrak{A}_i)/\mathcal{U} \models \phi(f, g, h \ldots) \text{ iff } \{i : \mathfrak{A}_i \models \phi(f(i), g(i), h(i) \ldots)\} \in \mathcal{U}.$$

This is an exercise: you should be able to do it!

You will notice that in the induction step for the existential quantifier you use the axiom of choice to pick a witness from each factor, and this use of AC seems unavoidable This might lead you to suppose that Łoś's theorem is actually equivalent to AC, but this seems not to be the case. Try it! I am endebted to Phil Freeman for drawing my attention to

`http://www.jstor.org/sici?sici=0002-9939(197506)49%3A2%3C426%3ALTATBP%3E2.0.CO%3B2-F`
[21]

This has the incredibly useful corollary (which we shall not prove) that

COROLLARY 11. *A formula is equivalent to a first-order formula iff the class of its models is closed under elementary equivalence and taking ultraproducts.*

Theorem 16 enables us to show that a lot of things are not expressible in any first order language. Since, for example, an ultraproduct of finite p-groups (which are all simple) is not simple, it follows that the property of being a simple group is not capturable by a language in which you are allowed to quantify only over elements of the object in question.

Miniexercise: If the ultrafilter is principal ($\{J \subseteq I : i \in J\}$), then the ultraproduct is isomorphic to the ith factor. So principal ultrafilters are no use.

In contrast if the ultrafilter is nonprincipal you can make good use of the construction even if all the models you feed into it are the *same*. In this case you speak of the output as an **ultrapower**, and write it as $\mathfrak{M}^\kappa/\mathcal{U}$. One of the reasons why this process is useful is that there is an elementary embedding $\mathfrak{M} \hookrightarrow \mathfrak{M}^\kappa/\mathcal{U}$.

LEMMA 12. *The embedding $i = \lambda m_{\mathfrak{M}} \lambda f_{\mathfrak{M}^\kappa/\mathcal{U}}.m$ is elementary.*

(This embedding i is just a typed version of the K combinator!)
Proof:
It will be sufficient to show that, for any $m \in \mathfrak{M}$, if there is an $x \in \mathfrak{M}^\kappa/\mathcal{U}$ such that $\mathfrak{M}^\kappa/\mathcal{U} \models \phi(x, i(m))$ then there is $x \in \mathfrak{M}$ s.t. $\mathfrak{M} \models \phi(x, m)$. Consider such an $x \in \mathfrak{M}^\kappa/\mathcal{U}$. It is the equivalence class of a family of functions which almost everywhere (in the sense of $\mathcal{U}$) are related to m by ϕ, so by Łoś's theorem , there must be something x in $\mathfrak{M}$ such that $\mathfrak{M} \models \phi(x, m)$. Then $\lambda i < \kappa.x$ will do.

■

We shall see the utility of this later in connection with measurable cardinals.

Ultraproducts enable us to give a particularly slick proof of the compactness theorem for predicate calculus.

THEOREM 13. *(Compactness theorem for predicate logic)*
If Δ is a set of sentences of predicate calculus such that every finite $\Delta' \subseteq \Delta$ has a model, (we say Δ is "finitely satisfiable") then Δ has model.

Proof: Let Δ be a set of wffs that is finitely satisfiable. Let $\mathcal{S}$ be the set of finite subsets of Δ (elsewhere in these notes notated $\mathcal{P}_{<\aleph_0}(\Delta)$), and let $X_s = \{t \in \mathcal{S} : s \subseteq t\}$. Pick $\mathfrak{M}_s \models s$ for each $s \in \mathcal{S}$. Notice that $\{X_s : s \in \mathcal{S}\}$ generates a proper filter. Extend this to an ultrafilter $\mathcal{U}$ on $\mathcal{S}$. Then

$$(\prod_{s \in \mathcal{S}} \mathfrak{M}_s)/\mathcal{U} \models \Delta.$$

This is because, for any $\phi \in \Delta$, $X_{\{\phi\}}$ is one of the sets that generated the filter that was extended to $\mathcal{U}$. For any $s \in X_{\{\phi\}}$, $\mathfrak{M}_s \models \phi$, so $\{s : \mathfrak{M}_s \models \phi\} \in \mathcal{U}$. ■

Notice we are not making any assumption that the language is countable. We will have cause to refer to this construction again when we reach definition

117. The ultrafilter $\mathcal{U}$ on $\mathcal{S}$ mentioned in this proof is *fine* in the sense of definition 140.

Notice the relation between Arrow's paradox and the nonexistence of non-principal ultrafilters on finite sets. Consider an ultraproduct of finitely many linear orders: it must be isomorphic to one of the quotients. This is Arrow's "dictatorship" condition.

Exercise 2. *Let $\{A_i : i \in \mathbb{N}\}$ be a family of finite structures, and $\mathcal{U}$ a nonprincipal ultrafilter on $\mathbb{N}$. Show that the ultraproduct is finite if there is a finite bound on the size of A_i and is of size $2^{\aleph_0}$ if every infinite subset of $\{A_i : i \in \mathbb{N}\}$ contains arbitrarily large elements.*

2.2 Saturated Models

(Review countable categoricity. See a property S so that any two models of a complete theory that are both S are not only elementarily equiv but iso)

There is a very beautiful theorem of Ryll-Nardzewski concerning countably categorical structures which i shall not prove, though i shall throw out a couple of hints.

Theorem 14. *Let $\mathfrak{M}$ be a countable structure. Then the following are equivalent*

- *For all $n \in \mathbb{N}$ $Aut(\mathfrak{M})$ has only finitely many orbits on n-tuples from $\mathfrak{M}$;*
- *$Th(\mathfrak{M})$ is countably categorical.*

Proof:

Sketch: One way we use a back-and-forth construction. The other way we use the omitting types theorem, theorem 31. ■

Total orderings and η_α sets

Definition 15. *A κ-**like** total order is one of power κ all of whose proper initial segments are of power $< \kappa$.*

We don't really need this definition yet but it has to go in somewhere!

This next concept arises naturally if we seek a condition on uncountable dense linear orders that enables us to prove an analogue of Cantor's isomorphism theorem.

If you were one of those children who opened the windows on the Advent calendar before the day you were supposed to, you might find this exercise tries your patience.

Exercise 3. *How do you have to modify the definition of denseness to obtain a definition for which the back-and-forth argument can be continued transfinitely?*

Definition 16. *An η_α set is a total order $\langle X, \leq \rangle$ in which, for all x, if A is a subset of X bounded above by x, and $|A| < \aleph_\alpha$, then A has an upper bound strictly below x.*

Check here that you know what a regular ordinal is!

Other than the rationals (which form a η_0 set), these are rather hard to visualise. $\mathbb{R}$ is not an η_α set. Perhaps an example would help. Let κ be an infinite regular initial ordinal and let us consider the set of κ-sequences of 0s and 1s. (Free associate to: transfinite generalisation of $\mathbb{R}$). Consider also the subordering containing only those sequences that are eventually constantly 0. (Free associate to: transfinite generalisation of the rationals). Order all these things lexicographically. Clearly the eventually-constant things are dense in the whole thing. Consider now the sequences that are *not* eventually constant. (The "irrationals"). If i have one of these chaps, s, say, and an increasing α-sequence $\langle s_\zeta : \zeta < \alpha \rangle$ of things all below s in the lexicographic order, there must be something above all the s_ζ that is below s, for consider the set $\{i < \kappa : \exists \zeta_1\ \zeta_2$ the first coordinate at which $s_{\zeta_1} \neq s_{\zeta_2}$ is $i\}$. This is a set of size α, all of whose members are below κ, so its sup is below κ, by regularity of κ. Then construct a bound for the s_ζ in the obvious way, and there is still room to squeeze another point in.

This is an exact transfinite analogue of the presentation of the (dyadic) rationals as finite sequences of 0s and 1s. Also, like the rationals, every total order of size κ embeds in this total order.[2]

Q: How big are these analogues of the rationals?

A: 2 to the "weak power κ". This is the sum of 2^β for $\beta < \kappa$.

The reals (Baire space: ω^ω with the product topology, each factor having the discrete topology) have a countable dense subset. This is true in general and we will need it later:

Exercise 4. *The lexicographic order of 2^κ has a dense subset of size 2 to the weak power κ.*

Exercise 5. *For α regular, any two η_α sets are isomorphic tosets.*

Exercise 6. *If you understand arrow notation for partition theorems have a bash at this one. $\kappa^+ >$ 2-to-the-weak-power-κ $\rightarrow 2^\kappa \not\rightarrow (\kappa^+)^2$.*

Here is as good a place as any to introduce the idea of a saturated model. A **type** (for model theorists, not for set theorists or type theorists!!) is a set of formulæ. If Σ is a type with free variables $\vec{x}$ we say that a tuple $\vec{a}$ (in a structure $\mathfrak{M}$) **realizes** Σ if $\mathfrak{M} \models \sigma(\vec{a})$ for every $\sigma \in \Sigma$.[3]

Definition 17. *A type is finitely satisfiable if every finite subset of it can be realized. A model is* **saturated** *iff every finitely satisfiable type is realized.*

---

[2]Let X be a countable linear order, and replace every point of X by a copy of the rationals. X obviously embeds into this, and it is also dense and countable, therefore isomorphic to the rationals.

[3]Model theorists tend to use capital Greek letters for types (in this sense of 'type') and corresponding lower-case Greek letters for formulæ in them.

We can use ultraproducts to prove the existence of saturated models.

THEOREM 18. *(The existence of saturated models)*
Let $\mathcal{L}$ be a countable language, and let $\mathcal{U}$ be an ultrafilter over an index set I, where $\mathcal{U}$ is not countably complete. Then for every family $\{A_i : i \in I\}$ the ultraproduct $(\prod_{i\in I} \mathfrak{A}_i)/\mathcal{U}$ is $\aleph_1$-saturated.

Proof:
We must show that for every countable set $\{f_i : i \in \mathbb{N}\}$ of elements of $(\prod_{i\in I} \mathfrak{A}_i)/\mathcal{U}$ and every set $\Sigma(x)$ of formulæ from $\mathcal{L}$ (with countably many new constants $c_i \ldots$), if each finite subset of $\Sigma(x)$ is satisfiable in $(\prod_{i\in I} \mathfrak{A}_i)/\mathcal{U}$ (with names for the f_i) then so is $\Sigma(x)$ itself.

Since $\mathcal{L}$ with the new constants is also a countable language it will be sufficient to prove it without the constants.

Suppose every finite subset of $\Sigma(x)$ is satisfiable in $(\prod_{i\in I} \mathfrak{A}_i)/\mathcal{U}$. $\Sigma(x)$ is countable, so we can think of it as $\{\sigma_i : i \in \mathbb{N}\}$. Since $\mathcal{U}$ is countably incomplete, we find a $\subseteq$-descending ω sequence $\langle I_i : i \in \mathbb{N}\rangle$ of $\mathcal{U}$-large subsets of I whose intersection is empty.

Define a new sequence $\langle X_i : i \in \mathbb{N}\rangle$ by

$$X_0 =: I$$

and thereafter

$$X_n =: I_n \cap \{i \in I : \mathfrak{A}_i \models (\exists x)(\bigwedge_{j<n} \sigma_j(x))\}.$$

$(\prod_{i\in I} \mathfrak{A}_i)/\mathcal{U}$ satisfies every finite subset of Σ so, by Łoś's theorem ,

$$\{i \in I : \mathfrak{A}_i \models (\exists x)(\bigwedge_{j<n} \sigma_j(x))\} \in \mathcal{U}.$$

This ensures that (i) each X_n is in $\mathcal{U}$, (ii) the X_i are nested and (iii) their intersection is empty. From (iii) it follows that for each $i \in I$ there is a last $n \in \mathbb{N}$ s.t. $i \in X_n$. Let this last n be $n(i)$. We are now going to construct an $f \in \prod_{i\in I} \mathfrak{A}_i$ such that $[f]_{\mathcal{U}}$ realizes Σ. If $n(i) = 0$ then $f(i)$ can be anything. Otherwise set $f(i)$ to be any x such that $\mathfrak{A}_i \models \bigwedge_{j<n(i)} \sigma_j(x)$.

■

EXERCISE 7. *Show that any two countably saturated countable elementarily equivalent structures are isomorphic.*

The effect of the ultraproduct construction is to add lots of things whose presence cannot be detected by finitistic first-order methods. Thus we can add infinitesimals to the reals. Hence Nonstandard Analysis: an ultraproduct of $\mathbb{R}$ (modulo a countably incomplete ultrafilter at least) is saturated (theorem 24) and therefore contains infinitesimals. This means we can reconstruct the eighteenth century theory of differentiation and integration!

2.2.1 The Ehrenfeucht-Mostowski theorem

Ultraproducts contain lots of nonstandard funny stuff, but they don't obviously admit automorphisms. However we can use them to create models that do. The theorem of this section was proved in the 1950's by Ehrenfeucht and Mostowski, using methods of Ramsey theory and compactness. Our proof here uses ultrapowers and is due to Gaifman.

DEFINITION 19. *$\mathcal{I} = \langle I, \leq_{\mathcal{I}} \rangle$ is a* **set of indiscernibles** *for a model $\mathfrak{M}$ for a language $\mathcal{L}$ iff for all $\phi \in \mathcal{L}$, if ϕ is a formula with n free variables in it then for all distinct n-tuples $\vec{x}$ and $\vec{y}$ from $\mathcal{I}$* **taken in increasing order** *we have $\mathfrak{M} \models \phi(\vec{x}) \longleftrightarrow \phi(\vec{y})$.*

We can create a nonstandard model of the reals by adding to our language a constant symbol c_α for each countable ordinal α, and—whenever $\alpha < \beta$—an axiom $c_\alpha < c_\beta$. By compactness this gives a consistent theory, so there is definitely going to be a nonstandard model of the reals containing a copy of the countable ordinals—or any other total order we want, come to think of it! However nothing in this construction will ensure that constants $\{c_\alpha : \alpha < \omega_1\}$ are embedded as a set of indiscernibles. The fact that this apparently much more difficult feat can be achieved is the content of

THEOREM 20. *(Ehrenfeucht-Mostowski theorem)*

Let I be a total order, T a theory with infinite models and a formula $P()$ with one free variable s.t. T thinks that the extension of P is an infinite total order. Then T has a model $\mathfrak{M}$ in which I is embedded in (the interpretation of) P and in which every automorphism of I extends to an automorphism of $\mathfrak{M}$. Finally the copy of I in $\mathfrak{M}$ is a set of indiscernibles.

Notice that there is no suggestion that the copy of I in the model we build is a set of that model, or is in any way definable.

We give first an outline of the original proof, due to Ehrenfeucht and Mostowski.

Proof:

Add to the language of T names c_i for every element of I, and axioms to say the c_i are all distinct. Next we add axioms providing correct order information about the c_i. Let this theory be T^*. By compactness we know that this theory is consistent, since T has an infinite model, and believes that the domain of $<$ is infinite. That is to say, if T believes that the domain of $<$ is infinite, we can find a model of T in which I is embedded in the domain of $<$. This much is a straightforward application of completeness and compactness.

Now we add axioms saying that these constants are a set of indiscernibles; there will be infinitely many of these axioms. This gives us a theory T^I.

We have to show that every finite fragment of T^I is consistent. For the moment let T' be one such finite fragment. It mentions only finitely many constants—c_1, c_3, c_4 and c_5, say—and it says that they form a set of indiscernibles for finitely many predicates—ϕ_1, ϕ_2, ϕ_3 and ϕ_4, for the sake of argument.

Now the task of proving this theory consistent is precisely the same task as proving consistent the theory T'' obtained from T' by replacing c_1, c_3, c_4 and c_5 by any other sequence of constants of length 4. So if we liked we could drop the names c_1, c_3 c_4 and c_5 and call them something noncommittal like a b c and d. Once we have the model for the noncommittal version of T'' we can restore the c labels.

Now let $\mathfrak{M}$ be a model of T^*. These predicates—ϕ_1, ϕ_2, ϕ_3 and ϕ_4, together with the order relation $<$ that T^* imposes on the c_i—divide up $[\{c_i : i \in I\}^{\mathfrak{M}}]^m$ (where m is the supremum of the arities of the ϕ's) into finitely many pieces. How do they do this? If ϕ_4 is of arity m then it obvious splits $[\{c_i : i \in I\}^{\mathfrak{M}}]^m$ into two bits. But what if $m = 3$ and ϕ_4 is of arity two? What is it to do with a triple from $\{c_i : i \in I\}^{\mathfrak{M}}$. Well, any such triple gives rise to three pairs, and we feed each pair into ϕ_4 in increasing order. So ϕ_3 splits $[\{c_i : i \in I\}^{\mathfrak{M}}]^m$ into eight pieces. So the number of pieces into which we split $[\{c_i : i \in I\}^{\mathfrak{M}}]^m$ is the product of the numbers of pieces mandated by each ϕ_i mentioned in T'. Call this partition Π.

Now $\{c_i : i \in I\}^{\mathfrak{M}}$ is infinite, so there must be a monochromatic set for Π of size 4, and we take its elements, read in increasing order, to be a, b, c and d.

Now we invoke compactness to conclude that T^I has a model. Any model of T^I has I embedded as a set of indiscernibles. Then consider the Skolem hull of the indiscernibles. ∎

Gaifman's proof

We start with some standard observations about direct limits. Given a directed family of structures with embeddings (satisfying commutation conditions) there is a well-defined notion of **direct limit** which i won't explain further.

Check that you know what a direct limit is!

THEOREM 21.

1. *A direct limit of structures preserves* Π_2 *sentences;*
2. *A direct limit of an elementary family (one where the embeddings are elementary) preserves everything.*

I think both these can be safely left as an exercise. Email me in case of doubt

The idea underlying the proof of Gaifman's is that one can recover any structure from the embedding relations between its finitely generated substructures: it's a direct limit of them (co-limit if you're a categorist).

Specifically if $\langle I, \leq_{\mathcal{I}} \rangle$ is an ordered set then it is the direct limit of its finite substructures where the embedding relations are the obvious inclusion embeddings. Remarkably, this banal fact is almost all we need!

The Construction

We start with an infinite model $\mathfrak{M}$ of T. We are going to create a directed family of elementary embeddings and iterated ultrapowers of $\mathfrak{M}$ indexed by the set of finite substructures of $\langle I, \leq_{\mathcal{I}} \rangle$, and the desired model will be a substructure of the direct limit, which we notate '$\mathfrak{M}_\infty$'. We will use the letters 's' and 't' to range over these finite substructures and we will notate the corresponding models $\mathfrak{M}_s$.

Let P be $\{x : P(x)\}^{\mathfrak{M}}$. We will assume that P has no last element in the sense of the ordering of P according to $\mathfrak{M}$. $\mathcal{U}$ will be an ultrafilter on P that contains all terminal segments of P. (So P had better not have a last element!)

Now to define the models in the family. $\mathfrak{M}_s$ will simply be the result of doing the ultrapower construction $|s|$ times to $\mathfrak{M}$, so that what $\mathfrak{M}_s$ actually *is* depends only on the *length* of s and in no way on what the members of s are. $\mathfrak{M}_\emptyset$ is just the $\mathfrak{M}$ we started with.

Now we have to define a family of embeddings and establish that they commute. We need to recall some notation: $\texttt{last}(s)$ is the last member of s (remember s is thought of as an increasing sequence) and $\texttt{butlast}(s)$ is s minus its last element. We now define by recursion a family $\{I(s,t) : s \subseteq t \in I^{<\omega}\}$ of embeddings: $I(s,t)$ will be an elementary embedding from $\mathfrak{M}_s$ into M_t. The recursion needs two constructions. (i) K is the standard elementary embedding by constant functions from a structure into its ultrapower. (ii) If i is an embedding from $\mathfrak{M}$ to $\mathfrak{N}$ then there is an injection from $\mathfrak{M}^\kappa/\mathcal{U}$ into $\mathfrak{N}^\kappa/\mathcal{U}$ "compose with i on the right". Perhaps a picture will help.

$$\begin{tikzcd}
\mathfrak{M}_s \arrow[r, "i"] \arrow[d, "K"'] & \mathfrak{M}_t \arrow[d, "K"] \\
{\mathfrak{M}_s}^P/\mathcal{U} \arrow[r, "{\lambda[f].[\lambda\alpha.i(f(\alpha))]}"'] & {\mathfrak{M}_t}^P/\mathcal{U}
\end{tikzcd}$$

Let us call this operation L, so that $L =: \lambda i.\lambda f.i \circ f$.[4]

Now we can give the recursive definition of $I(s,t)$ when $s \subseteq t$.

$\texttt{If}$ $s = t$ $\texttt{then}$ $I(s,t)$ is the identity $\texttt{else}$
$\texttt{If}$ $\texttt{last}(s) = \texttt{last}(t)$ $\texttt{then}$ $I(s,t) =: L(I(\texttt{butlast}(s), \texttt{butlast}(t)))$
$\texttt{else}$ $I(s,t) =: K \circ I(s, \texttt{butlast}(t))$.

[4]A Curry-Howard point. The constructor L explains why '$M \to N. \to (K \to M) \to (K \to N)$' is intuitionistically correct. It is also the embedding underlying the cardinal arithmétic banality that $\alpha \leq \beta. \to \alpha^\zeta \leq \beta^\zeta$.

Notice that $\lambda[f].[\lambda\alpha.i(f(\alpha))]$ sends "new stuff" in $\mathfrak{M}_s^P/\mathcal{U}$ (by which I mean $(\mathfrak{M}_s^P/\mathcal{U}) \setminus K``\mathfrak{M}_s)$ to to "new stuff" in $\mathfrak{M}_t^P/\mathcal{U}$ (by which I mean $(\mathfrak{M}_t^P/\mathcal{U}) \setminus K``\mathfrak{M}_t)$. This will be essential later.

To check that this system of models and embeddings is genuinely a directed system it remains only to show that the embeddings are elementary and that they commute.

K is elementary by lemma 18. L of an elementary embedding is elementary as follows. Suppose $i : \mathfrak{M}_s \hookrightarrow \mathfrak{M}_t$ is elementary, and that $\mathfrak{M}_s^P/\mathcal{U} \models \phi(f_1 \ldots f_n)$. That is to say, $\{p : \mathfrak{M}_s \models \phi(f_1(p) \ldots f_n(p)\} \in \mathcal{U}$. Now $i : \mathfrak{M}_s \hookrightarrow \mathfrak{M}_t$ is elementary so this is equivalent to $\{p : \mathfrak{M}_t \models \phi(i(f_1(p)) \ldots i(f_n(p))\} \in \mathcal{U}$ which is equivalent to $\mathfrak{M}_s^P/\mathcal{U} \models \phi((i \circ f_1) \ldots (i \circ f_n))$ as desired.

To check that the family is commutative it is sufficient to check that the representative diagram below is commutative.

$$\begin{tikzcd}
\mathfrak{M}_1 \arrow[r, "i"] \arrow[d, "K"'] & \mathfrak{M}_{1,2} \arrow[d, "K"] \\
\mathfrak{M}_{1,3} \arrow[r, "L(i)"'] & \mathfrak{M}_{1,2,3}
\end{tikzcd}$$

That is to say that, for any i, $K \circ i = L(i) \circ K$.

- $K \circ i$ gives: $x \mapsto i(x) \mapsto \lambda p.i(x)$;
- $L(i) \circ K$ gives: $x \mapsto K(x) \mapsto L(i)(K(x)) = \lambda p.(i \circ p)(K(x)) = i \circ K(x) = \lambda p.(i \circ K(x))p = \lambda p.(i(K(x)p)) = \lambda p.i(x)$.

This will show that all paths from $\mathfrak{M}_s$ to $\mathfrak{M}_t$ (and the number of such paths is presumably $(|t| - |s|)!$) correspond to the same injection.

Embedding I in the direct limit

The point of this direct limit construction was to obtain a structure $\mathfrak{M}_\infty$ in which I was embedded. To achieve this we ensure that each finite subset $s \subset I$ is embedded in $\mathfrak{M}_s$ in such a way that the manifestations of the elements of I in the $\mathfrak{M}_s$ get stitched together properly. That means that inside $\mathfrak{M}_s$ we must be able to point to $|s|$ distinct **things**. We will find these **things**[5] by a recursive construction, and we will prove by induction on n that the construction works for s of length n. We can think of $\mathfrak{M}_s$ as a segmented structure: it has $|s|$ segments, and each new segment consists of the junk added by the ultraproduct construction applied to the object so far, and each segment contains a **thing**—each time we zap the model with our ultraproduct wand we add a new **thing**.

What is the $|t|$th **thing** in $\mathfrak{M}_t$ to be? The following train of tho'rt gives us a fix on what it must be, and tells us how we might find the $|t|$th thing by recursion on $|t|$. Suppose we know how to find the n **thing**s in $\mathfrak{M}_s$ when $|s| = n$,

[5]'Things'?? I've got to call them *something*!!

and let $|t| = n + 1$. Now let t' be t with its *penultimate* element deleted (so we are assuming that $n \geq 2$). By induction hypothesis we know already what the last **thing** in $\mathfrak{M}_{t'}$ is. But we also know the embedding $I(t', t)$ from $\mathfrak{M}_{t'} \hookrightarrow \mathfrak{M}_t$. This tells us that the last **thing** in $\mathfrak{M}_t$ is that object to which $I(t', t)$ sends the last **thing** in $\mathfrak{M}_{t'}$. (We flagged this earlier.)

It's worth checking that we have got the base case right. This construction giving us the $n+1$th **thing** is guaranteed to work only as long as $\mathfrak{M}_s$ (where $|s| = n$) is already an ultrapower, since any embedding created by the recursion expects its domain to be an ultrapower. But this is all right; $\mathfrak{M}_\emptyset$ is not an ultrapower (it is $\mathfrak{M}$) but then it isn't expected to have a **thing** in it.

So how do we decide what the first **thing** is? We know that as the various **thing**s get added they must form an increasing sequence according to the (extension of the) order (in $\mathfrak{M}$) to which I will belong. It would be helpful to ensure that the first **thing**—which we see for the first time in $\mathfrak{M}^I/\mathcal{U}$—is later than everything in $\mathfrak{M}$. This will be the case if (i) the total order in $\mathfrak{M}$ has a last element and (ii) the index set I is precisely the extension of the order in $\mathfrak{M}$ and the ultrafilter $\mathcal{U}$ contains all terminal segments of I.

So we let the first **thing** be an arbitrary object in $\mathfrak{M}^P\mathcal{U} \setminus \mathfrak{M}$.

A last thought: how do we identify the nth **thing** in $\mathfrak{M}_s$ with the nth **thing** in $\mathfrak{M}_t$ where $|t| > |s| \geq n$? Well, K is an embedding from $\mathfrak{M}_s$ into $\mathfrak{M}_{s \cup \{i\}}$ and obviously we want the nth **thing** in $\mathfrak{M}_{s \cup \{i\}}$ to be K of the nth **thing** in $\mathfrak{M}_s$. Repeat as necessary.

The **things** end up in $\mathfrak{M}_\infty$ as a subclass $\{c_i : i \in I\}$. We want to show that the c_i form a set of indiscernibles in $\mathfrak{M}_\infty$. Let $\vec{s}$ and $\vec{t}$ be two finite subsets of I (tho'rt of as increasing sequences). We want to show that $\phi(\vec{s})$ iff $\phi(\vec{t})$ (identifying each c_i with i for the moment). Now the embedding from $\mathfrak{M}_s$ into $\mathfrak{M}_\infty$ is elementary, so $\mathfrak{M}_\infty \models \phi(\vec{s})$ iff $\mathfrak{M}_s \models \phi(\vec{s})$; similarly the embedding from $\mathfrak{M}_t$ into $\mathfrak{M}_\infty$ is elementary, so $\mathfrak{M}_\infty \models \phi(\vec{t})$ iff $\mathfrak{M}_t \models \phi(\vec{t})$. Now comes the step at which we exploit the fact that $\mathfrak{M}_s = \mathfrak{M}_t$ as long as $|s| = |t|$. This fact tells us that $\mathfrak{M}_\infty \models \phi(\vec{s})$ iff $\mathfrak{M}_\infty \models \phi(\vec{t})$.

■

We can now do various other clever things. We can consider the Skolem hull of the indiscernibles. We then find that any order-automorphism of $\mathcal{I} = \langle I, \leq_\mathcal{I} \rangle$ extends to an automorphism of the Skolem hull.

2.3 Unsaturated models

Any fool can realize a type: it takes a model theorist to omit one.

Gerald Sacks.

Sacks is right—omitting types is hard!

First some definitions

DEFINITION 22.

1. *An* n-**type** *is a set of formulæ all with at most* n *free variables*

2. *A model* $\mathfrak{M}$ **realises** *an n-type* Σ *if there is a tuple* $\vec{x}$ *s.t* $\mathfrak{M} \models \phi(\vec{x})$ *for every* $\phi \in \Sigma$*;*

3. T **locally omits** *a* 1*-type* Σ *if, whenever* ϕ *is a formula s.t.* $T \vdash (\forall x)(\phi(x) \to \sigma(x))$ *for all* σ *in* Σ*, then* $T \vdash (\forall x)(\neg\phi(x))$*.*

(I state this only for 1-types. The existence of a universal pairing function means that the treatment of n-types for higher n follows automatically, but it might be worth just checking this at some point.)

The property of theories of locally omitting a particular type is universal horn:

$$(\forall\phi)(\bigwedge_{\sigma\in\Sigma}(T \vdash (\forall x)(\phi(x) \to \sigma(x))) \ \to \ T \vdash (\forall x)(\neg\phi(x)))$$

(It looks like $(\forall x)((\bigwedge_{i\in I} p_i) \to q)$.) The following is a consequence of this observation:

Remark 23. *An intersection of an arbitrary family of theories each locally omitting a type* Σ *also locally omits* Σ*.*

Proof: Obvious... ■

...but worth noting, since it means that we have a good notion of closure: keep adding axioms to T until you obtain something that locally omits Σ.

Definition 24. *Let* T^0 *be* T*. Obtain* $T^{\alpha+1}$ *from* T^α *as follows. Whenever* ϕ *is a formula s.t.* $T^\alpha \vdash (\forall x)(\phi(x) \to \sigma(x))$ *for all* σ *in* Σ*, then add to* T^α *the new axiom* $(\forall x)(\neg\phi(x))$*. The result of doing this for all* $\sigma \in \Sigma$ *is* $T^{\alpha+1}$*.*

I can't think of any reason why this process should close up at ω so we iterate transfinitely until it closes or becomes inconsistent. Let the result be T_∞ where ∞ is the closure ordinal (countable if $\mathcal{L}_T$ is countable). I suppose the 'Σ' should appear in this notation somewhere!

Theorem 25. *If* T *locally omits* Σ *then it has a model omitting* Σ*.*

Proof:
Let T be a theory locally omitting a type Σ and let $C = \langle c_i : i \in \mathbb{N}\rangle$ be a countable set of new constant letters. Let $\langle \phi_i : i \in \mathbb{N}\rangle$ be an enumeration of the sentences of $\mathcal{L}_T$.

We will construct recursively a $\subseteq$-increasing sequence $\langle T_i : i \in \mathbb{N}\rangle$ of finite extensions of T with the property that, for each $m \in \mathbb{N}$,

1. T_{m+1} decides ϕ_n for all $n \leq m$;

2. If ϕ_m is $(\exists x)\psi(x)$ and $\phi_m \in T_{m+1}$, then $\psi(c_p) \in T_{m+1}$ where c_p is the first constant not occurring in T_m or ϕ_m;

3. There is a formula $\sigma(x) \in \Sigma$ such that $(\neg\sigma(c_m)) \in T_{m+1}$.

Given T_m, we construct T_{m+1} as follows. Think of T_m as $T \cup \{\theta_1 \ldots \theta_r\}$, and the conjunction of the theta's as Θ. Let $\{c_1 \ldots c_n\}$ be the constants from C that have appeared in Θ, and let $\Theta(\vec{x})$ be the result of replacing 'c_i' by 'x_i' in Θ. Then (obviously!) $\Theta(\vec{x})$ is consistent with T. Therefore, for some $\sigma(x) \in \Sigma$, $\Theta \wedge \neg\sigma(x_m)$ is consistent with T. Put '$\neg\sigma(c_m)$' into T_{m+1}. This makes (3) hold.

If ϕ_m is consistent with $T_m \cup \{\neg\sigma(c_m)\}$, put it into T_{m+1}. Otherwise put in $\neg\phi_m$. This takes care of (1). If ϕ_m is $(\exists x)\psi(x)$ and is consistent with $T_m \cup \{\neg\sigma(c_m)\}$, put $\psi(c_p)$ into T_{m+1}. This takes care of (2). This ensures that (1-3) hold for T_{m+1}.

Now consider $T^* = \bigcup\limits_{i \in \mathbb{N}} T_i$. T^* is complete by construction. Consider an arbitrary countable model of T^* and the submodel of that model generated by the constants in C. This will be a model of T^*, and condition 3 ensures that it omits Σ. ■

Lemma 26. *If T locally omits a type Σ, then so does any finite extension of T.*

Proof: Suppose T locally omits Σ; we will show that $T \cup \{p\}$ locally omits Σ too. Suppose for each $\sigma \in \Sigma$, $T \cup \{p\} \vdash (\forall x)(\phi(x) \to \sigma(x))$. Then, for each $\sigma \in \Sigma$,

$$T \vdash p \to (\forall x)(\phi(x) \to \sigma(x))$$

so

$$T \vdash (\forall x)(p \to (\phi(x) \to \sigma(x)))$$

and

$$T \vdash (\forall x)((p \wedge \phi(x)) \to \sigma(x))$$

so, since T locally omits Σ,

$$T \vdash (\forall x)(\neg(p \wedge \phi(x))$$

whence

$$T \cup \{p\} \vdash (\forall x)(\neg\phi(x))$$

as desired. ■

This seems to work for finite extensions only. One might think that one can do it for arbitrary extensions by using compactness but the proof has a hole. (Check it!) In particular i can see no reason why the union of a $\subseteq$-chain of theories each locally omitting Σ should omit Σ—unless of course the chain has uncountable cofinality. However it does enable us to prove the following

Corollary 27. *Let T be a theory, Σ a type $\subseteq \mathcal{L}_T$ and T^∞ the least theory $\supseteq T$ that locally omits Σ. Let T^* be the theory of all models of T that omit Σ. (That is to say $T^* = \bigcap\{Th(\mathfrak{M}) : \mathfrak{M} \models T$ and $\mathfrak{M}$ omits $\Sigma\}$)*

Then $T^\infty = T^*$.

Proof:

Clearly T^* locally omits Σ, so $T^\infty \subseteq T^*$. Suppose the inclusion is proper, so that there is $p \in T^* \setminus T^\infty$. But then, by lemma 32, $T^\infty \cup \{\neg p\}$ locally omits Σ. Therefore, by theorem 31, there will be a model of $T^\infty \cup \{\neg p\}$ that omits Σ. But then p cannot be in T^* which—after all—is $\bigcap\{Th(\mathfrak{M}) : \mathfrak{M} \models T$ and $\mathfrak{M}$ omits $\Sigma\})$. ■

2.3.1 Rich Theories

(not examinable)

Definition 28. *T* **has SEP** *(the "strong existence property") iff whenever $T \vdash (\exists x)\phi(x)$ there is some finite family $\langle t_i : i \in I\rangle$ of terms so that $T \vdash \bigvee_{i \in I} \phi(t_i)$. Alternatively we can say that T is* **rich**.

An alternative way of thinking of rich theories is given by their characterisation as theories extending finite intersections of complete theories locally omitting the type "I am not definable".

Lemma 29. *If T is rich, so is any extension of T.*

Proof: Suppose T is rich, and let Γ be a set of formulæ. If $T \cup \Gamma \vdash (\exists x)\Psi(x)$ then for some finite $\Phi \subseteq \Gamma$ we must have $T \vdash \Phi \to (\exists x)\Psi(x)$ which is to say $T \vdash (\exists x)(\Phi \to \Psi(x))$ whence (since T is rich) $T \vdash (\Phi \to \Psi(t_1)) \vee (\Phi \to \Psi(t_2)) \vee \ldots (\Phi \to \Psi(t_n))$ for some list of terms $\vec{t}$ which is to say $T \cup \{\Phi\} \vdash \Psi(t_1) \vee \Psi(t_2) \vee \ldots \Psi(t_n)$, so certainly $T \cup \Gamma \vdash \Psi(t_1) \vee \Psi(t_2) \vee \ldots \Psi(t_n)$ as well. ■

Remark 30. *The intersection of two rich theories is rich.*

Proof: Suppose $S \cap T \vdash (\exists x)\Psi(x)$.

Then $S \vdash (\exists x)\Psi(x)$ and $T \vdash (\exists x)\Psi(x)$.

So $S \vdash \bigvee_{i \in I} \phi(t_i)$ and $T \vdash \bigvee_{i \in J} \phi(t_i)$;

So both S and T prove $\bigvee_{i \in I \cup J} \phi(t_i)$, so $S \cap T \vdash \bigvee_{i \in I \cup J} \phi(t_i)$. ■

Proposition 31. *For the nonce let T_1 be the theory of all term models of T, and let T_2 be the intersection of all rich theories extending T. T_3 is the canonical inductively defined extension of T that locally omits the type "I am not a term" (as on page 32).*

Then:

T_1, T_2 and T_3 are all the same theory.

Proof:

$T_2 \subseteq T_1$:

This holds because T_2 is the intersection of *all* rich theories extending T and T_1 is the intersection of only some of them.

$T_3 \subseteq T_2$:

Let T_4 be a rich theory extending T: we want to show $T_3 \subseteq T_4$. Recall definition 30 on page 32. We show by induction on α that ${T_3}^{\alpha} \subseteq T_4$. This is certainly true for $\alpha = 0$. Now suppose ${T_3}^{\alpha} \vdash (\forall x)(\phi(x) \to x \neq t)$ for all terms t, then $T_4 \vdash (\forall x)(\phi(x) \to x \neq t)$ (since, by induction hypothesis, ${T_3}^{\alpha} \subseteq T_4$). Now if $(\exists x)(\phi(x))$ is consistent with T_4 there is a term model $\mathfrak{M} \models T + (\exists x)(\phi(x))$. But this is impossible. Therefore $(\forall x)(\neg\phi(x)) \in T_4$. But ϕ was arbitrary, so ${T_3}^{\alpha+1} \subseteq T_4$ as desired. The limit case is easy.

$T_1 \subseteq T_3$:

Suppose $\phi \in T_1 \setminus T_3$. Consider $T_3 \cup \{\neg\phi\}$. It is consistent ($\phi \notin T_3$ by hypothesis), it locally omits the type $x \neq t$ by lemma 32 and so it has a term model. But ϕ is true in all term models.

■

It might be worth noting that if T has no term models then all these three theories are axiomatised by $\perp$.

2.4 Preservation Theorems

(This lemma is probably not going to be lectured. It's here beco's, well ... it's the key lemma one uses for proving preservation theorems)

Lemma 32. *Let T be a consistent theory in $\mathcal{L}$ and let Δ be a set of sentences of $\mathcal{L}$ which is closed under $\vee$. Then the following are equivalent*

1. *T has a set Γ of axioms where $\Gamma \subseteq \Delta$;*
2. *If $\mathfrak{A}$ is a model of T and $(\forall\delta \in \Delta)(A \models \delta \to \mathcal{B} \models \delta)$ then $\mathcal{B} \models T$.*

Proof:

It is obvious that 1 implies 2. For the converse, assume (2), and suppose Δ and T given. Let $\Gamma = \{\phi \in \Delta : T \vdash \phi\}$. Then $T \vdash \Gamma$. We will show that Γ entails the whole of T. Let $\mathcal{B}$ be a model of T. Let

$$\Sigma = \{\neg\delta : \delta \in \Gamma \wedge \mathcal{B} \models \neg\delta\}$$

We show that $T \cup \Sigma$ is consistent. T is consistent by hypothesis; Suppose $T \cup \Sigma$ is inconsistent. Then there are $\neg\delta_1 \ldots \neg\delta_n$ all in Σ such that $T \vdash \neg(\neg\delta_1 \wedge \ldots \wedge \neg\delta_n)$ which is to say $T \vdash \delta_1 \vee \ldots \vee \delta_n$. Since Δ is closed under $\vee$ this theorem belongs to Δ, and therefore to Γ and therefore holds in $\mathcal{B}$. But this contradicts the fact that these δ_i are *false* in $\mathcal{B}$. So $\Sigma \cup T$ must have been consistent, and has a model $\mathfrak{A}$. Then every sentence $\delta \in \Delta$ which holds in $\mathfrak{A}$ holds also in $\mathcal{B}$ (by (2)). So Γ is an axiomatisation of T as desired.

Definition 33.

The triple $\langle \mathfrak{A}, \mathcal{B}, \mathcal{C} \rangle$ form a **sandwich** *if $\mathfrak{A} \subseteq \mathcal{B} \subseteq \mathcal{C}$ and $\mathfrak{A} \prec \mathcal{C}$;*

$\mathfrak{A}$ *is* **sandwiched by** $\mathcal{B}$ *if there are elementary extensions $\mathfrak{A}'$ of $\mathfrak{A}$ and $\mathcal{B}'$ of $\mathcal{B}$ such that $\mathcal{B} \subseteq \mathfrak{A}' \subseteq \mathcal{B}'$.*

Theorem 34. *the following are equivalent*

1. *T has a set of universal-existential axioms*
2. *T is preserved under unions of chains of models*
3. *if whenever $\mathfrak{A} \models T$ and $\mathfrak{A}$ is sandwiched by $\mathcal{B}$ then $\mathcal{B} \models T$.*

Proof:
We prove $1 \to 2 \to 3 \to 1$.
$1 \to 2$ is easy; $3 \to 1$ follows from lemma 38; we now prove $2 \to 3$.
Suppose $\mathfrak{A}$ is sandwiched by $\mathcal{B}$. We shall construct a chain of models

$$\mathcal{B}_0 \subseteq \mathfrak{A}_0 \subseteq \mathcal{B}_1 \subseteq \mathfrak{A}_1 \ldots \mathcal{B}_n \subseteq A_n \ldots$$

where $\mathcal{B}_0 = \mathcal{B}$; each triple $\mathcal{B}_n, \mathfrak{A}_n, \mathcal{B}_{n+1}$ forms a sandwich; $\mathfrak{A} \prec \mathfrak{A}_0$ and each $\mathfrak{A}_n$ is elementarily equivalent to $\mathfrak{A}$. We will attempt to construct this sequence by recursion, and to do this we will need to be able—on being presented with $\mathcal{B}_n, \mathfrak{A}_n$ and $\mathcal{B}_{n+1}$ forming a sandwich—to find $\mathfrak{A}_{n+1}$ elementarily equivalent to $\mathfrak{A}_n$ and an elementary extension $\mathcal{B}_{n+2}$ of $\mathcal{B}_{n+1}$ so that $\mathcal{B}_{n+1}, \mathfrak{A}_{n+1}$ and $\mathcal{B}_{n+2}$ form a sandwich.

How do we do this? We extend the language-in-hand by adding a new one-place predicate U and a constant name c_b for every element $b \in \mathcal{B}_{n+1}$. Let us call this new language Ł$'$. Note let T' be the theory

(elementary diagram of $\mathcal{B}_{n+1}$) $\cup \{\phi^U : \mathfrak{A}_n \models \phi\} \cup \{U(c_b) : b \in B_{n+1}\}$
where ϕ^U is the relativisation of ϕ.

Thus any model of T' will be an elementary extension $\mathcal{B}_{n+2}$ of $\mathcal{B}_{n+1}$ which contains a subset U that includes at least all the elements of $\mathcal{B}_{n+1}$. Also the submodel determined by the extension of U (in $\mathcal{B}_{n+2}$) is elementarily equivalent to $\mathfrak{A}_n$. That T' is consistent can be shown as follows:

- A sentence $F(c_{b_1} \ldots c_{b_n})$ such that $\mathcal{B}_{n+1} \models S(b_1 \ldots b_n)$
- A sentence ϕ^U such that $\mathfrak{A}_n \models \phi$
- The sentence $U(c_{b_1}) \wedge U(c_{b_2}) \wedge \ldots U(c_{b_n})$.

Since $\mathcal{B}_n \prec \mathcal{B}_{n+1}$ there will be $d_1 \ldots d_n \in B_n$ such that

$$\mathcal{B}_n \models S(d_1 \ldots d_n), \quad \mathcal{B}_{n+1} \models S(d_1 \ldots d_n)$$

Now we find that B_{n+1} is a model of T' if we interpret U as membership in A_n, and the constant c_{d_i} as d_i.

Now consider the sequence of models in this chain we are building. Clearly all the $\mathfrak{A}_n$ are models of T and so (since we are assuming 2) the union is also a model of T. But this union is also the union of the $\mathcal{B}_n$, which are all elementarily equivalent, and is therefore elementarily equivalent to them too, and is a model of T, so $\mathcal{B}$ was a model of T too.

We have used the fact that a direct limit of a family of elementary embeddings is an elementary extension. This was theorem 27.

This lemma is the crucial lemma in the proof of lots of completeness theorems: a formula is equivalent to a [syntactic property] formula iff the class of its models is closed under [some operations].

2.4.1 Ultralimits and Frayne's Lemma

(A reminder of two bits of jargon: an **expansion** of a structure $\mathcal{B}$ is a structure with the same carrier set and more gadgets. e.g. the rationals as a field are an expansion of the rationals as an additive group. The converse relation is a **reduction**: the rationals as an additive group are a reduction of the rationals as a field.) This is no longer in small letters co's it's examinable and dead cute.

Lemma 35. *Suppose $\mathfrak{A}$ and $\mathcal{B}$ are elementarily equivalent. Then there is an ultrapower $\mathfrak{A}^I/\mathcal{U}$ of A and an elementary embedding from $\mathcal{B}$ into it.*

Proof:

Supply names $\mathtt{b}$ for every member b of B. Let $\mathcal{L}$ be the language with the new constants. There is an obvious way of expanding $\mathcal{B}$ to a structure for this new language, namely to let each constant $\mathtt{b}$ denote that element b of B which gave rise to it. (Of course this is not the only way of doing it: *any* map $B \to B$ will give rise to an expansion of $\mathcal{B}$ of this kind—and later we will have to consider some of those ways). Let us write '$\mathcal{B}'$' to denote this obvious expansion of $\mathcal{B}$, and let I be the set of sentences of $\mathcal{L}$ true in $\mathcal{B}'$. (Use of the letter 'I' for this is a bit of a give-away!)

Consider ϕ a formula in I. It will mention finitely many constants—let us say two, for the sake of illustration. Replace these two constants by new variables 'v_1' and 'v_2' (*not* mentioned in ϕ!) to obtain ϕ'' and bind them both with '$\exists$' to obtain $(\exists v_1)(\exists v_2)\phi''$ which we will call 'ϕ'' for short. This new formula is a formula of the original language which is true in $\mathcal{B}$ and is therefore also true in $\mathfrak{A}$ (since $\mathfrak{A}$ and $\mathcal{B}$ are elementarily equivalent).

The next step is to expand $\mathfrak{A}$ to a structure for the language $\mathcal{L}$ by decorating it with the with the extra constants $\mathtt{b}$ etc that we used to denote members of B. Of course any function $B \to A$ gives us a way of decorating $\mathfrak{A}$ but with ϕ in mind we are interested only in those decorations which give us a structure that satisfies ϕ. If ϕ contained the constants $\mathtt{b}$ and $\mathtt{b}'''$ for example then the obvious way to expand A involves using those two constants to denote the witnesses in $\mathfrak{A}$ for the two existential quantifiers in ϕ'. Since ϕ contains only finitely many constants this nails down denotations for only finitely many of the constant-names-for-members-of-B. However any finite map from B to A can be extended to a total function $B \to A$ so we can extend this to a way of labelling members of A with these constants in such a way that the decorated version of $\mathfrak{A}$ satisfies the original formula ϕ.

Pick one such labelling and call it $a(\phi)$. (Thus $a(\phi)$ is merely an element of $B \to A$ satisfying an extra condition parametrised by ϕ. We can think of a as a function $\mathcal{L} \to (B \to A)$ or as a function $(\mathcal{L} \times B) \to A$ *ad libitum*). $\mathfrak{A}$ expanded with this decoration we call $\langle \mathfrak{A}, a(\phi) \rangle$. Now consider the set

$J(\phi) =: \{\psi \in I : \langle \mathfrak{A}, a(\psi) \rangle \models \phi\}$

It is easy to check that the family $\{J(\phi) : \phi \in I\}$ of subsets of I has the finite intersection property and so gives rise to a ultrafilter[6] $\mathcal{U}$ on I and thence to an ultrapower $\mathfrak{A}^I/\mathcal{U}$. Evidently if $\phi \in I$ then $J(\phi) \in \mathcal{U}$ and the ultrapower will believe ϕ.

We have to find an elementary embedding from $\mathcal{B}$ into this ultrapower. Given $b \in B$ whither do we send it? The obvious destination for b is the equivalence class of the function $\lambda\phi.a(\phi)b$ that sends ϕ to $a(\phi)b$. The function that sends b to $[\lambda\phi.a(\phi)b]$ is $\lambda b.[\lambda\phi.a(\phi)b]$—which we will write '$h$' for short. We must show that h is elementary.

The best way to understand what h does and why it is elementary is to think of the ultrapower as a reduction of the ultraproduct

$$\prod_{\psi \in I} \langle \mathfrak{A}, a(\psi) \rangle / \mathcal{U}.$$

("**expand** the factors; take an ultraproduct; **reduce** the ultraproduct—to obtain a ultrapower of the factors ...")

Each of the factors $\langle \mathfrak{A}, a(\psi) \rangle$ is a structure for $\mathcal{L}$ and therefore the ultraproduct is too. By the same token, for each $b \in B$, each of the factors has an element which is pointed to by $\mathtt{b}$-the-constant-name-of-b, and therefore the ultraproduct will too. The key fact is that h is the function that sends each $b \in B$ to the thing in the ultraproduct that is pointed to by $\mathtt{b}$ the constant-name-of-b.

As for the elementarity of h, suppose $\mathcal{B} \models \phi(\vec{v})$. Then, for some choice of constants $\vec{\mathtt{b}}$, $\mathcal{B} \models \phi(\vec{\mathtt{b}})$, and $\mathcal{B}' \models \phi'$. But now $J(\phi)$ is $\mathcal{U}$-large, so the ultrapower believes ϕ. ■

(I lifted this proof from Bell-and-Slomson: Models and Ultraproducts. [3] A truly lovely book.)

But what we really need is Scott's lemma:

Lemma 36. *Suppose $g : \mathfrak{A} \hookrightarrow \mathcal{B}$ is an elementary embedding. Then there is an ultrapower $\mathfrak{A}^I/\mathcal{U}$ of A and an elementary embedding from $\mathcal{B}$ into it making the triangle commute.*

Proof:

[6] There doesn't seem to be any reason to conclude that this ultrafilter will be nonprincipal, but then nor does it seem to matter if it *isn't*. Bell and Slomson [3] don't say that it will be nonprincipal. Thanks to Phil Ellison for drawing my attention to this point.

The ideas are the same, but we need to be slightly more careful in the definition of $a(\phi)$. Fix once for all a member a of A. As before, we extend the language by adding names for every member of B, thus obtaining the language $\mathcal{L}$ as before. Now we expand $\mathcal{B}$ by decorating B with these names, but not in the obvious way. If b is in the range of g we allow $\mathtt{b}$ the constant-name-of-b to denote b; if b is not in the range of g, then $\mathtt{b}$ will denote $g(a)$. Let's call this expanded structure $\mathcal{B}'$.

If we are to expand $\mathfrak{A}$ to obtain a structure for $\mathcal{L}$ then we must ensure that, for each $b \in B$, the constant-name-$\mathtt{b}$-of-b points to something in A. The obvious way to do this is to ordain that $\mathtt{b}$ point to g^{-1} of the thing that that $\mathtt{b}$ points to in the expansion $\mathcal{B}'$ of $\mathcal{B}$. This decorated version of $\mathfrak{A}$ and the decorated version $\mathcal{B}'$ of $\mathcal{B}$ are elementarily equivalent (with respect to the extended language with the names) (*)

As before, let I be the set of sentences of $\mathcal{L}$ true in $\mathcal{B}'$. Consider a formula $\phi \in I$. Recall what we did at the same stage in the proof of Frayne's Lemma. This time we replace with existentially-quantified variables only those constants denoting elements of B not in the range of g. Let's call this formula ϕ' like last time. Evidently $\mathcal{B}' \models \phi'$ and so, by the remark (*) at the end of the last paragraph, the decorated version of $\mathfrak{A}$ also satisfies ϕ'. So, as before, there is another decoration of $\mathfrak{A}$ which actually satisfies the original ϕ. Pick one such decoration and call it $a(\phi)$, and call the structure thus decorated $\langle \mathfrak{A}, a(\phi)\rangle$. We define

$J(\phi) =: \{\psi \in I : \langle \mathfrak{A}, a(\phi)\rangle \models \phi\}$

as before, and it has the finite intersection property as before and gives us an ultrafilter $\mathcal{U}$ as before, and we have the same elementary embedding h from $\mathcal{B}$ into the ultrapower as before. It remains only to check that the diagram is commutative. I think this can safely be left as an exercise to the reader. ■

Now comes the fun part.

$$\begin{tikzcd}
A \arrow[r, hook] \arrow[d, phantom, "\equiv"] & A_1 \arrow[r, hook] \arrow[d, hook] & A_2 \arrow[r, hook] \arrow[d, hook] & {} & \cdots & A_\infty \\
B \arrow[r, hook] \arrow[ur, hook] & B_1 \arrow[r, hook] \arrow[ur, hook] & B_2 \arrow[r, hook] \arrow[ur, hook] & {} & \cdots & B_\infty
\end{tikzcd}$$

where A_n is an n-times ultrapower of A and B_n similarly.

Evidently A_∞ and B_∞ are isomorphic.

Chapter 3

Set theory

3.1 Cantor and di Giorgi

I have come to the conclusion that the best point of departure for set theory is an *aperçu* of di Giorgi's.

Di Giorgi thinks of a model of Set Theory (better: a *structure for the language of set theory*) as a set A of atoms (things with no internal structure) together with an injective map $i : A \to \mathcal{P}(A)$. We then define a membership relation between the members of A by $a_1 \in_i a_2$ iff $a_1 \in i(a_2)$. Thus an atom a "codes" the set of things that are in $i(a)$, and sets of the model are those elements of $\mathcal{P}(A)$ that are coded by atoms.

A theorem of Cantor's tells us that this injection cannot be surjective.

The proof is titillating and short: Suppose every subset of M is coded by an element of M. Consider the set of all those elements of M that code subsets of M and are not members of the subsets they code. Call it X. If every subset of M is coded by a member of M then X is coded by some element x. We then get a contradiction by asking whether x is a member of X.

This is not the usual way in which Cantor's theorem is presented: I present it this way because i prefer to think of Cantor's theorem as a constraint on our ability to code things naïve set theory is inconsistent.

The inconsistency of naïve set theory is one of those fundamental metaphysical disasters that befall intelligent life, like original sin except—if anything—worse. The endeavour to recover from it colours all our experience, rather in the way in which everything has to be seen in the context of our endeavours to rebuild the Tower of Babel, or to reconstitute the two-backed beasts.

Cantor's theorem tells us that not every collection of sets can be coded by a set but that seems to be the limit of the really obvious elementary facts one can state about which collections of sets can be coded as sets. In particular, for any subset A' of A we can devise i so that A' is in the range of i, that is to say, A' is coded by an atom.

Now evidently we can make a decision about *which subsets of A are to* *be coded by atoms* while leaving open *which atoms are going to code those subsets*. After all, if we compose an injection $i : A \to \mathcal{P}(A)$ on the right with a permutation π of A then we have a different injection, but one that makes the same decision about which subsets of A are to be sets of the model. These two injections will be different implementations of the same decision about which sets are to be coded.

How do the model corresponding to an injective map $i : A \to \mathcal{P}(A)$ and the model corresponding to $i \cdot \pi$ differ? It is natural to seek to ascribe a special status to the formulæ which are invariant under this change. You may like to think about what this class of formulæ might be.

Last year I deleted altogether the chapters on Antifoundation axioms and Positive set theory. This was partly because they were essay topics and I therefore couldn't lecture them properly; I have now put them back. The need to cover topics outside set theory has meant that the coverage of set theory is thinner than it otherwise would be—and thinner in some areas than others. There is no doubt that one of the drivers for set theory—historically—has been the study of the continuum, and other small infinite sets: *Analysis* in short.[1] There are to this day set theorists who see the study of the continuum and other small infinite sets as their focus in studying set theory, and who hang around large cardinals for the information that large cardinals let fall about small infinite sets. Although I have no quarrel with them I am not one of their number, and this fact is reflected in the coverage that the study of small infinite sets gets in these notes: thin! My reason for being interested in large cardinals is that they delight us with some interesting mathematics.

I shall have something later to say about the historical roots of set theory as opposed to its conceptual roots as immediately above.

Conceptual roots? What is set theory a theory *of*? Accessible pointed graphs? What is graph theory a theory of?

It is a matter of record that nowadays most people who call themselves set theorists or who study what they call 'set theory' in fact study *wellfounded sets* exclusively. The axioms that the universe of wellfounded sets is believed to obey comprise a list that includes all the axioms of ZFC and it gets added to gradually over the decades. This is a rich and fascinating topic, and i shall cover a lot of it in what follows.

Why do they study wellfounded sets only? It's not entirely clear. Marco Forti says that it's pure historical accident that people interested in set theory about a century ago (when it all started, and when they had a chance to define the area) opted for the axiom scheme of foundation rather than antifoundation. Is it a mistake? Almost certainly, but not a particularly grave one. Since ZF + antifoundation can be smoothly interpreted in ZF then if (and it's a big if) you think the job of set theory is to provide a basis for mathematics then there is no cost attached to believing that there are no illfounded sets. If, on the other hand, you think the job of set theory is to study sets then you might be missing out on some mathematics by closing your eyes to the possibility of illfounded

[1] "small" is slang of course, but it means something like "of size $< \beth_\omega$".

sets.

Chapter 4

ZF and beyond: Large Cardinals

This should really be subtitled *Safe Sets*.

> Independence of the axioms. Ramsey's theorem and Ramsey ultrafilters. Games and other topics concerning small infinite sets. Then Erdős-Rado and off into large cardinal land!

4.1 Hereditarily something-or-other

Much of this is in [15] pp 197 et seq.

A device which turns up in many of these independence proofs is the idea of the set of things that are hereditarily ϕ, where ϕ is a one-place predicate. The intuition is that x is hereditarily ϕ if everything in $TC(x)$ is ϕ. If you are ϕ and all your members are hereditarily ϕ then you are hereditarily ϕ, whence the rectype definitions:

Definition 37.

$$\mathcal{P}_\kappa(x) := \{y \subseteq x : |y| < \kappa\};\ H_\kappa := \bigcap\{y : \mathcal{P}_\kappa(y) \subseteq y\};$$
$$\mathcal{P}_\phi(x) := \{y \subseteq x : \phi(y)\};\quad H_\phi := \bigcap\{y : \mathcal{P}_\phi(y) \subseteq y\}$$

A word is in order on the definition and the notation involved. The use of the set-forming bracket inside the '$\bigcap$' is naughty: in general there is no reason to suppose that the collection of all y such that $\mathcal{P}_\phi(y) \subseteq y$ is a set. However its intersection will be a set—as long as it's nonempty! And if there is even one x such that $\mathcal{P}_\phi(x) \subseteq x$ then $\{y \subseteq x : \mathcal{P}_\phi(y) \subseteq y\}$ will have the same intersection as $\{y : \mathcal{P}_\phi(y) \subseteq y\}$ and so no harm is done. But this depends on there being such an x. If there is, we are in the same situation we were with the implementation of $\mathbb{N}$. If not, then the collection H_ϕ will be a proper class and we have to define it as the collection of those x with the property that everything in $TC(x)$ is of size $< \kappa$. If H_ϕ is a set then the two definitions are of course equivalent, but if it isn't, it is only the definition in terms of TC that works. The definition in terms of TC is the standard one, but I find that my definition is more helpful to people who are used to thinking in terms of inductive definitions. After all, H_ϕ is a rectype. It has an empty set of founders and one (infinitary!) constructor that says that a subset of H_ϕ that is itself ϕ is also in H_ϕ.

WF is just $H_{x=x}$.

> A Warning about deviant definitions of things like $H_{\beth_\omega}$. If you define it in the usual (bad) way as the set of things whose transitive closure is of size less than $\beth_\omega$ then the result contains fewer sets that the result of defining it as the intersection of all sets that contain all their smaller-than-$\beth_\omega$. This problem tends to always arise with H_κ where κ is singular, so beware. This matters because if you define $H_{\beth_\omega}$ in the more restrictive way, with fewer members, then $H_{\beth_\omega}$ and $V_{\omega+\omega}$ turn out to be elementarily equivalent wrt stratified formulæ. This is clearly not true for $V_{\omega+\omega}$ and $H_{\beth_\omega}$ the way we have defined $H_{\beth_\omega}$ because $V_{\omega+\omega}$ is a model of sumset and $H_{\beth_\omega}$ isn't. And Sumset is stratified!

Exercise 8. *Establish that if $cf(\kappa) = \kappa$ then $H_\kappa = \{x : |TC(x)| < \kappa\}$.*

Exercise 9.

1. *Show that if κ is regular and we have AC then we can take H_κ to be the set of x s.t. $|TC(x)| < \kappa$.*
2. *Show that the collection of hereditarily wellordered sets isn't a set.*

Remark 38. *If $\phi(x) \to \phi(f``x)$ for all x and f then H_ϕ is a model for replacement.*

Proof:

For H_ϕ to be a model of replacement it is sufficient that if $x \in H_\phi$ and $f : H_\phi \to H_\phi$ is defined by a formula *with parameters from H_ϕ only, all of whose quantifiers are restricted to H_ϕ* then $f``x$ is also in H_ϕ. But this condition is met because by assumption a surjective image of a set that is ϕ is also ϕ: indeed, we didn't even need the italicised condition. ■

The class of hereditarily countable sets is a set and is the size of the continuum.

$H_{\aleph_1}$, the collection of hereditarily countable sets, is a pedagogically useful object, being an essential prop in an elementary proof of the independence from the other axioms of ZF of the axiom of power set, seen in part II

Although it is not obvious, H_κ is always a set according to *ZF*. The proof does not need (much) AC, but the axiom of foundation is essential.[1]

[1]We will see in due course that if we do not assume the axiom of foundation we can easily construct models containing as many Quine atoms (sets $x = \{x\}$) as we want. Since these objects are clearly hereditarily of size less than κ^+ there is no point in asking about the size or sethood of $H_{\aleph_1}$ unless we assume some form of foundation.

THEOREM 39. *If κ is an aleph, $|H_{\kappa^+}| \leq 2^\kappa$.*

Proof:

Let X be a set of size 2^κ. Assume enough choice to be sure that $|X| = |\mathcal{P}_{\kappa^+}(X)|$ beco's κ^+ is well-behaved. We need a bit of choice to do this beco's $\kappa^2 = \kappa$ is not enough. For example there are $2^{\aleph_0}$ ω-sequences of reals, since there are $(2^{\aleph_0})^{\aleph_0} = 2^{\aleph_0 \cdot \aleph_0} = 2^{\aleph_0}$ but that doesn't tell us there are $2^{\aleph_0}$ countable sets of reals. To infer that from the fact that there are $2^{\aleph_0}$ ω-sequences of reals we would need to be able to pick, for each countable set of reals, a wellordering of it to length ω.

Let us fix an injection $\pi : \mathcal{P}_{\kappa^+}(X) \hookrightarrow X$.

We construct an injection $h : H_{\kappa^+} \hookrightarrow X$ by recursion thus: $h(x) =: \pi(h``x)$. By considering a member of H_{κ^+} of minimal rank not in the range of h we show easily that h is total. It is injective because π is one-one. The range of h is a set by comprehension, and so its domain (which is H_{κ^+}) is also a set, by replacement. ■

REMARK 40. $|H_{\aleph_1}| = 2^{\aleph_0}$

Proof: We have just seen $|H_{\aleph_1}| \leq 2^{\aleph_0}$. The other direction follows immediately from the fact that $\mathcal{P}(\omega)$ (the Von Neumann $\mathcal{P}(\omega)$) is a subset of $H_{\aleph_1}$ of size $2^{\aleph_0}$. ■

EXERCISE 10. *Show that, for any set x, $H_{|x|}$ is a set.*

EXERCISE 11. *In this context, the di Giorgi take on what a model of set theory is turns out to be quite helpful and natural. If $|X| = |\mathcal{P}_\kappa(X)|$ and i is a bijection $X \longleftrightarrow \mathcal{P}_\kappa(X)$ then this i is an injection $X \hookrightarrow \mathcal{P}(X)$ and the resulting model of set theory is a model for ... well, what is it a model for?*

THEOREM 41. *(Jech: JSL 1982 op cit) Everything in H_κ is of rank $< \kappa^+$.*

Proof: omitted.

The study of the H_κ is looming gradually larger in modern ZF studies. For example CH is true iff $H_{\aleph_2}$ satisfies "Every countable set has a power set"

The H_κs have other nice properties. The universe is not only an end-extension of H_κ ("no new members of old sets in the extension"); it is also a $\mathcal{P}$-**extension** of any H_κ. ("No new subsets of old sets in the extension.")

4.2 Independence proofs

Although clearly some instances of the axiom schemes of separation and replacement can be derived from others, it is standard that the remaining axioms of ZF are independent from each other. For any other axiom A we can show that $ZF \setminus A \not\vdash A$. And for replacement we can show that $ZF\setminus$ replacement does not imply all instances of replacement, though it does prove some.

The relatively straightforward way in which almost every axiom of ZF allows itself to be deduced from the remaining axioms provides a rich source of elementary exercises for beginners in set theory, but it also annoys those beginners. Why not cut down the set of axioms to something irredundant, as we did with the axioms of propositional logic?

The answer is that we want to preserve *graceful downward compatibility*. On the whole, separate axioms of ZF correspond to separate principles of set existence—to distinct reasons why collections might be genuine sets. That being the case, whenever one is interested in executing a particular construction—be it of Gödel's L to take a pertinant example, or something else—one has a motive for isolating those principles that are needed for its execution, and discarding—for the moment—all those that are *not* thus needed. Thus one comes to retain some axioms that follow from stronger axioms that are present, because one wants them to be there when one temporarily discards the stronger axiom beco's it is not needed for the construction one has in mind, or perhaps beco's one is interested in seeing how one manages without it. The idea is that if we have axioms A and B, which answer to different set-theoretic existence principles, then one wants to retain B even if it follows from A, as one might want to discard A in order to see what B can do unaided.

For this reason it's not terribly important that the axioms be independent of each other, but the exercise of establishing the independence exposes one to various nice idiomatic constructions and—as we all know—constructions are more important than theorems.

4.2.1 Extensionality

Copy this section from axiomsofsettheory.tex

4.2.2 Replacement

The proof of the independence of Replacement is the only one to use a V_α.

$V_{\omega+\omega}$ is a model for all the axioms except replacement. It contains wellorderings of length ω, so pick on of them and call it $\mathbb{N}$. Consider now the function $\lambda n.V_{\omega+n}$. The image of $\mathbb{N}$ in this function is $\{V_{\omega+n} : n \in \mathbb{N}\}$ and this cannot be a set because we can use the axiom of sumset (and $V_{\omega+\omega}$ is clearly a model for the axiom of sumset!) to get $V_{\omega+\omega}$.

Readers are encouraged to check the details for themselves to gain familiarity with the techniques involved.

4.2.3 Power set

H_κ is never a model for power set except when κ is strong limit.

There is another way of proving that $H_{\aleph_1}$ is a set. Recall that $\lambda x.\mathcal{P}_{\aleph_1}(x)$ is not ω-continuous. If you think about this for a while you will realise that this function is α-continuous for those α such that $cf(\alpha) > \omega$. The first such ordinal is ω_1. (Look back at remark 61). So all we have to do is iterate this function ω_1

times and we will reach a fixed point. $H_{\aleph_1}$ will be a subset of this fixed point and will be a set by comprehension.

$H_{\aleph_1}$ gives us a model of ZF minus the power set axiom. The axiom of infinity will hold because there are genuinely infinite sets in $H_{\aleph_1}$. This is not sufficient by itself as "is infinite" is not Δ_0, but whenever X is such a set then there will be a bijection from X onto a proper subset of itself, and this bijection (at least if our ordered pairs are Wiener-Kuratowski) will be a hereditarily countable set. So any actually infinite member of $H_{\aleph_1}$ will be believed by $H_{\aleph_1}$ to be actually infinite. We have been assuming the axiom of choice so the union of countable many elements of $H_{\aleph_1}$ is also an element of $H_{\aleph_1}$, so it is a model of the axiom of sumset. Everything in $H_{\aleph_1}$ is countable and therefore wellordered, and under most implementations of pairing functions, the wellorderings will be in $H_{\aleph_1}$ too, so $H_{\aleph_1}$ is a model of AC, even if AC was not true in the model in which we start.

Make sure you understand Δ_0 ...

4.2.4 Independence of the axiom of infinity

$H_{\aleph_0}$ provides a model for all the axioms of ZF except infinity and thereby proves the independence of the axiom of infinity. (We constructed a copy of $H_{\aleph_0}$ on page **??**).

That status of AC in $H_{\aleph_0}$ is like its status in $H_{\aleph_1}$. Everything in $H_{\aleph_0}$ is finite and therefore wellordered, and under most implementation of pairing functions, the wellorderings will be in $H_{\aleph_0}$ too, so $H_{\aleph_0}$ is a model of AC, even if AC was not true in the model in which we start. This is in contrast to the situation obtaining with the countermodels to sumset and foundation: the truth-value of AC in those models is the same as its truth-value in the model in which we start.

4.2.5 Sumset

Definition 42. *The Hebrew letter $\beth$ is called 'beth'. Beth numbers are defined by setting $\beth_\alpha := |V_{\omega+\alpha}|$, or recursively by $\beth_0 := \aleph_0; \beth_{\alpha+1} := 2^{\beth_\alpha}$, taking sups at limits.*

Let us for the moment say that a set of size less that $\beth_\omega$ is **small**.

Then $H_{\beth_\omega}$, the collection of hereditarily small sets, proves the independence of the axiom of sumset. This is because there are wellorderings of length $\omega+\omega$ inside $V_{\omega+n}$ for n small, so by replacement $\{V_\alpha : \alpha < \omega+\omega\}$ is a set. Indeed it is a hereditarily small set. But $\bigcup\{V_\alpha : \alpha < \omega+\omega\}$ is not hereditarily small, since it is $V_{\omega+\omega}$ and accordingly of size $\beth_\omega$.

Exercise 12. *Establish that the collection of hereditarily small sets is a set.*

4.2.6 Foundation

Look at the first few pages of `www.dpmms.cam.ac.uk/~tf/churchlatest.ps`

Antifoundation

There is another proof of the independence of the axiom of foundation, which goes back to work of Forti and Honsell *op. cit.*. Under this approach a set is regarded as an isomorphism class of accessible pointed digraphs ("APG"s). An APG is a digraph with a designated vertex v such that every vertex has a dipath reaching v.

The best-known exposition of this material is the eminently readable Aczel *op. cit.*. I shall not treat it further here, since—although attractive—it is recondite, and the proof of independence of foundation that it gives does not (unlike the previous one) naturally give rise to a proof of the independence of the axiom of choice. This is our next chore.

4.2.7 Choice

We use Fraenkel-Mostowski models. 'F' is for Abraham Fraenkel (He who put the 'F' in 'ZF') and 'M' for Mostowski. Do not confuse these with Rieger-Bernays permutation models which we saw earlier, in the independence of foundation. The FM construction is intended to prove the independence of AC.

We start with a model of ZF with urelemente. In the original treatment these urelemente are taken to be empty. For technical reasons it's easier to take them to be Quine atoms, sets identical to their singletons. The effect is that one drops foundation rather than extensionality, but the two constructions have the same *feel*.

We start with a model of ZF + foundation, and use Rieger-Bernays model methods to obtain a permutation model with a countable set A of Quine atoms. The permutation we use to achieve this is the product of all transpositions $(n, \{n\})$ for $n \in \mathbb{N}^+$. A will be a **basis** for the illfounded sets in the sense that any class X lacking an $\in$-minimal element contains a member of A. Since the elements of A are Quine atoms every permutation of A is an $\in$-automorphism of A, and since they form a basis we can extend any permutation σ of A to a unique $\in$-automorphism of V in the obvious way: set $\sigma(x) := \sigma``x$. Notice that the collection of sets that this definition does not reach has no $\in$-minimal member if nonempty, and so it must contain a Quine atom. But σ by hypothesis is defined on Quine atoms. (a, b) is of course the transposition swapping a and b, and we will write '(a, b)' also for the unique automorphism to which the transposition (a, b) extends. Every set x gives rise to an equivalence relation on atoms. Say $a \sim_x b$ if (a, b) fixes x. We say x is of (or has) **finite support** if $\sim_x$ has a cofinite equivalence class. (If it has a cofinite equivalence class it can have only one, and those remaining will all be finite). (Equivalently, using the cofinite quantifier $\forall_\infty$ "For all but finitely many ..." we can say that x is of finite support if $(\forall_\infty a \in A)(\forall_\infty b \in A)(a \sim_x b)$.)

However, below i write this automorphism $\tau_{(a,b)} \cdots$

The union of the (finitely many) remaining (finite) equivalence classes is the **support** of x. Does that mean that x is of finite support iff the transitive closure $TC(x)$ contains finitely many atoms? Well, if $TC(x)$ contains only finitely many atoms then x is of finite support (x clearly can't tell apart the cofinitely many atoms not in $TC(x)$) but the converse is not true: x can be of finite support if $TC(x)$ contains cofinitely many atoms. (Though that isn't a sufficient condition for x to be of finite support!!)[2]

It would be nice if the class of sets of finite support gave us a model of something sensible, but extensionality fails: if X is of finite support then $\mathcal{P}(X)$ and the set $\{Y \subseteq X : Y \text{ is of finite support}\}$ are both of finite support and have the same members with finite support. We have to consider the class of elements hereditarily of finite support. Let's call it HF. This time we *do* get a model of ZF.

Lemma 43. *The class of sets of finite support is closed under all the definable operations that the universe is closed under.*

Proof:

When x is of finite support let us write '$A(x)$' for the cofinite equivalence class of atoms under $\sim_x$. For any two atoms a and b the the transposition (a, b) of two atoms induces an $\in$-automorphism which for the moment we will write $\tau_{(a,b)}$.

Now suppose that $x_1 \ldots x_n$ are all of finite support, and that f is a definable function of n arguments. $x_1 \ldots x_n$ are of finite support, and any intersection of finitely many cofinite sets is cofinite, so the intersection $A(x_1) \cap \ldots A(x_n)$ is cofinite. For any a, b we have

$$\tau_{(a,b)}(f(x_1 \ldots x_n)) = f(\tau_{(a,b)}(x_1) \ldots \tau_{(a,b)}(x_n))$$

since $\tau_{(a,b)}$ is an automorphism. In particular, if $a, b \in A(x_1) \cap \ldots A(x_n)$ we know in addition that $\tau_{(a,b)}$ fixes all the $x_1 \ldots x_n$ so

$$\tau_{(a,b)}(f(x_1 \ldots x_n)) = f(x_1 \ldots x_n).$$

So the equivalence relation $\sim_{f(x_1 \ldots x_n)}$ induced on atoms by $f(x_1 \ldots x_n)$ has an equivalence class which is a superset of the intersection $A(x_1) \cap \ldots A(x_n)$, which is cofinite, so $f(x_1 \ldots x_n)$ is of finite support. ■

This takes care of the axioms of empty set, pairing, sumset and power set. To verify the axiom scheme of replacement we have to check that the image of a set hereditarily of finite support in a definable function (with parameters among the sets hereditarily of finite support and all its internal variables restricted to sets hereditarily of finite support) is hereditarily of finite support too. The operation of translating a set under a definable function (with parameters among the sets hereditarily of finite support and all its internal variables restricted to sets hereditarily of finite support) is definable and will (by lemma 49) take sets of finite support to sets of finite support.

So if X is in HF and f a definable operation as above, $f``X$ is of finite support. And since we are interpreting this in $HF$, all members of $f``X$ are in HF, so $f``X$ is in HF too, as desired.

[2] A counterexample: wellorder cofinitely many atoms. The graph of the wellorder has cofinitely many atoms in its transitive closure, but they are all inequivalent.

To verify the axiom of infinity we reason as follows. Every wellfounded set x is fixed under all automorphisms, and is therefore of finite support. Since all members of any wellfounded set are wellfounded they will all be of finite support as well, so every wellfounded set is hereditarily of finite support. So HF will contain all wellfounded sets that were present in the model we started with. In particular it will contain the von Neumann ω.

It remains only to show that AC fails in HF. Consider the set of (unordered) pairs of atoms. This set is in HF. However no selection function for it can be. Suppose f is a selection function. It picks a (say) from $\{a, b\}$. Then f is not fixed by $\tau_{(a,b)}$. Since f picks one element from every pair $\{a, b\}$ of atoms, it must be able to tell all atoms apart; so the equivalence classes of $\sim_f$ are going to be singletons, $\sim_f$ is going to be of infinite index, and f is not of finite support.

So the axiom of choice for countable sets of pairs fails. Since this axiom is about the weakest version of AC known to man, this is pretty good. The slight drawback is that we have had to drop foundation to achieve it. On the other hand the failure of foundation is not terribly grave: the only illfounded sets are those with a Quine atom in their transitive closures, so there are no sets that are gratuitously illfounded: there is a basis of countably many Quine atoms.

rest of this section probably not to be lectured

Now let's stand back and have a look at what features we have used.

It turns out that stabilisers of sets are useful. How can one characterise the sets of finite support by talking about stabilisers? Well, if x is of finite support, the stabiliser of x—or at least that subgroup of it consisting of permutations of finite support—is a product of finitely many groups all but one of which is a finite symmetric group, and this last is the group of permutations of finite support on a cofinite set of atoms. (I haven't managed to think through whether the stabiliser of a set of finite support is simply the product of the symmetric groups on the pieces of the partition corresponding to the equivalence relation induced on the atoms, or whether it is only the subgroup of that consisting of elements of finite support).

Sort this out

The collection of such groups is closed under finite intersections (which is the piece of finite character that we need to prove that the class of hereditarily nice sets is a model of ZF) and also closed under conjugation. This gives us the idea of a **normal filter**. So this is the more general picture that we are after.

The gadgets are as follows. (i) A set U of *Urelemente*; (ii) A group $\mathcal{G}$ of permutations of U; (iii) A normal filter $\mathcal{F}$ over $\mathcal{G}$. A normal filter is a collection of subgroups closed under intersection, superset and conjugation. We then say that a **stable** set is one whose stabiliser is in $\mathcal{F}$. We want the collection of hereditarily stable sets. Judicious choice of U, $\mathcal{G}$ and $\mathcal{F}$ can produce models of ZFU with very specific properties.

Other approaches

Readers should not form the impression that the "what do we trust?" approach that leads to the cumulative hierarchy is the only sensible response to the paradoxes: we might decide to trust syntax rather than creation. What does this mean? For most readers the most striking feature of the Russell class is not the fact that the cumulative hierarchy doesn't construct it, but rather the dangerous paradoxical reasoning involved and the obvious parallels with the other paradoxes. Perhaps there are other solutions to the paradoxes to be found by following up this syntactic insight? One might feel that sets ought to be naturally regimented into layers in such a way that the questions like $x \in x$? are somehow illegitimate. One way of doing this would be to restrict the naïve set existence scheme

$$(\forall \vec{x})(\exists y)(\forall z)(z \in y \longleftrightarrow \Phi)$$

to those where Φ is *stratified* in the sense of section 4.2.6. It is now known that this approach successfully skirts the paradoxes—at least if we weaken extensionality to the extent of allowing distinct empty sets. (Two nonempty sets with the same members must still be the same set). It is not known if this weakening of extensionality is needed. Why might anyone want to try this approach anyway? We noted on p. **??** that the problem with paradoxical sets might not be their size but the descriptions they answer to. Stratified descriptions, with their flavour of avoiding self-reference, have a certain appeal. However there is another reason for suspecting that this approach might have merit. We can think of Cantor's theorem as telling us something about models of Set Theory thought of in the di Giorgi way. In the di Giorgi view, the injection $i : A \hookrightarrow \mathcal{P}(A)$ can be thought of as a coding of subsets of A by elements of A. Cantor's theorem tells us that there are always subsets that are not coded. However, through i, one can make a choice of which subsets of A are to be coded and which not. Suppose we have a di Giorgi structure $\langle A, i\rangle$. If we now compose i on the right with a permutation π of A we have a new di Giorgi structure, but it has made the same choices as the original structure did about which subsets of A are coded by elements of A. What properties of the model are invariant under this kind of monkeying around with permutations? The reader of section 4.2.6 will be prepared for the news that it is precisely the stratified properties that are preserved.

So stratification is at least a sensible mathematical notion. Whether a Set Theory based on restricting naïve comprehension to accomodate it is going to give rise to a fruitful alternative to ZF remains to be seen. The question has been open for more than sixty years. In the meantime there are other syntactic tricks one can try. It is known that if one restricts naïve comprehension by not allowing negation signs inside Φ then no contradiction can be proved. In fact this constraint can be slightly relaxed, and the result is a Set Theory called GPC ("Generalised positive comprehension") but although this relaxation can be explained, the explanation is neither as elementary nor as satisfying as the story behind stratification.

I am not trying to persuade the reader that they should drop ZF and take up the study of stratified Set Theory or generalised positive comprehension. Even if you do not believe that set membership is a wellfounded relation, the rectype *WF* remains an object worthy of study. I merely wish to make the point that alternatives to ZF are available and are legitimate objects of study too.

Exercises

EXERCISE 13. *For P a poset, let P^* be the poset of chains-in-P partially ordered by end-extension. Show that there is no injective homomorphism $P^* \hookrightarrow P$.*

EXERCISE 14. *Take only the following ZF axioms: Extensionality, Infinity, Union, Power set, Replacement, Foundation and Choice. Then derive: Null, Pairing and Separation. The only hard bit is to derive null and separation together.*

EXERCISE 15.

Define E on $\mathbb{N}$ by: $n\ E\ m$ iff the n^{th} bit in the binary expansion of m is 1 (Remember to start counting at the 0th bit!!) Do you recognise this structure?

EXERCISE 16. *If you got that easily consider the following more complicated version: $n\ E_{\mathcal{O}}\ m$ iff either m is even and the n^{th} bit in the binary expansion of $m/2$ is 1 or m is odd and the n^{th} bit in the binary expansion of $(m-1)/2$ is 0. You have almost certainly never seen this structure before: what can you say about it?*

EXERCISE 17. *An* antimorphism *is a permutation π of V so that $\forall x\ y\ x \in y \iff \pi(x) \notin \pi(y)$. Prove (*without *using the axiom of foundation) that no model of ZF has an antimorphism.*

(i) Find an antimorphism of the second structure in exercise 17.

(ii) Is it unique? (hint: Consider the dual of the preceding structure, i.e., the natural numbers with the relation $n\ E_{\mathcal{O}^*}\ m$ iff either m is even and the n^{th} bit in the binary expansion of $m/2$ is 0 or m is odd and the n^{th} bit in the binary expansion of $(m-1)/2$ is 1. Prove that this is isomorphic to the naturals with $E_{\mathcal{O}}$)

EXERCISE 18. *Let X be a transitive set. If R is an equivalence relation on X and Y, Z are subsets of X we can define $R'(Y,Z)$ iff $(\forall y \in Y)(\exists z \in Z)(R(y,z)) \wedge (\forall z \in Z)(\exists y \in Y)(R(y,z))$. Check that the restriction of R' to X is also an equivalence relation on X.*

Show that this operation on equivalence relations has a fixed point. For any fixed point, one can take a quotient . Show how to define a membership relation on the quotient in a natural way, and that the result is a model of extensionality as well.

This construction is of particular interest if X is a V_α and the fixed point is the greatest fixed point. What can you say about the quotient in this case?

EXERCISE 19. *Use AC to show that if every chain in $\langle P, \leq_P \rangle$ has a sup then every directed subset does too.*

EXERCISE 20. *Von Neumann had an axiom which makes sense in the context of Set-Theory-with-classes.* **A class is a set iff it is not the same size as** V

Prove that Von Neumann's axiom is equivalent to replacement plus choice.

4.3 Some ordinal arithmetic

Definition 44. *A* **clubset** *is a* **CL***osed and* U*n***B***ounded set. Closed = contains all its limit points. Alternatively: the range of a total continuous function. (Sometimes called a* **normal** *function).*

Thus a normal function is strictly increasing and continuous. It's obvious that every normal function has a fixed point. If f is normal, then $\mathrm{Sup}\{f^n\alpha : n \in \mathbb{N}\}$ is the least fixed point for f above α. In fact:

Lemma 45. *The function enumerating the set of fixed points of a normal function is also normal.*

Definition 46. *A set X of ordinals is α-***closed** *if the supremum of any α-sequence of members of X is also in X. A set that is α-closed and unbounded is α***-club***.*

Notice that a closed set is the range of a continuous function and an α-closed set is the range of an α-continuous function.

α-cts was explained in [15].

Exercise 21. *Give examples of functions that are ω-continuous but not α-cts for other α.*

Definition 47. *Let $\langle C_\alpha : \alpha < \kappa \rangle$ be a sequence of subsets of (the ordinals below) κ. $C = \{\alpha : \alpha \in \bigcap_{\beta<\alpha} C_\beta\}$ is the* **diagonal intersection** *of the sequence.*

(Notice that the definition is **not** "$C = \{\alpha : \alpha \in C_\alpha\}$", which is how one is liable to misremember it.)

Definition 48. *A* **stationary** *subset of (the ordinals below) κ is one that meets every club subset of (the ordinals below) κ.*

("meets" means "has nonempty intersection with")

The intersection of a stationary and a club is stationary. Suppose S is a stationary set, and C a clubset. We want $S \cap C$ to be stationary. That is to say, we want $(S \cap C) \cap C'$ to be nonempty for any old clubset C'. But $(S \cap C) \cap C' = S \cap (C \cap C')$ so it will be sufficient to show that the intersection of two clubsets is club.

This last is not hard. Given two clubsets A and B build a set C as follows. C_0 =: first member of A, C_1 = first member of $B > C_0$, and keep on swapping. By closedness the sup must belong to both A and B. Keep on swapping transfinitely. The result is a club subset of $A \cap B$. So $A \cap B$ is at least unbounded. But it's obviously closed!

(Forgettable but important fact: this works only for ordinals of cofinality $> \omega$!)

We can in fact generalise this in two ways: to α-club sets and to intersections of infinitely many of them.

Club and **stationary** are dual to each other in the same was that $\exists$ and $\forall$ are dual to each other. (dense and open aren't quite dual: dense and has-a-nonempty-open-subset are dual) One can think of a quantifier as a family of subsets of the universe, so that a formula $(Qx)\Phi(x)$ says that the extension of Φ belongs to Q. Thus $\forall$ corresponds to the singleton of the universe, and $\exists$ corresponds to the set of all nonempty sets. A set is in $\exists$ iff it meets every set in $\forall$.

Some kinky quantifiers are quite useful. The first-order theory of Turing degrees, expressed in a language with the "measure zero" quantifier is apparently decidable. But beware: such quantifiers can be ill-behaved. Even nice monotone quantifiers like the cofinite quantifier have the nasty feature that adjacent like quantifiers do not commute (contrast: two adjacent $\exists$ or two adjacent $\forall$ *always* commute.)—for cofinitely many $x \in \mathbb{N}$ cofinitely many $y \in \mathbb{N}$ $y > x \ldots$!

(Sometimes i wonder if it might be easier for first-years if one were to explain convergence of sequences by means of the cofinite quantifier: $\{a_n\}$ converges iff for all $x > 0$, for cofinitely many n and for cofinitely many m, $|a_n - a_m| < x$. One particularly nice feature is that use of the cofinite quantifier allows one to express things formally using the same number of quantifiers as English does.)

4.3.1 A little factoid of Friedman

Although this next result is not terribly important, it is rather pretty, and makes very effectively the point that two apparently indistinguishable total orders can be nonisomorphic. Let A be a cofinal subset of ω_1 and define X_A to be the result of replacing (in ω_1) every element α by Q (if $\alpha \in A$) or by $1 + \mathbf{Q}$ (if not). The result is a total order of power $\aleph_1$ all of whose initial segments are of power $\aleph_0$. (Thus it is $\aleph_1$-**like**). Indeed all initial segments of X_A are isomorphic to the rationals. Despite this, there are oodles of nonisomorphic X_A. Let us think of X_A as being a concatenation of ω_1 *lumps* where each lump is either $\mathbf{Q}$ or $1 + \mathbf{Q}$.

$X_B{\restriction}\alpha =$: the initial segment consisting of the first α lumps of X_B. Suppose $X_A \simeq X_B$. Fix an isomorphism π. We are interested in ordinals α such that $X_A{\restriction}\alpha \simeq_\pi X_B{\restriction}\alpha$. Let this be the set of **concordant** ordinals.

The set of concordant ordinals is clearly closed, but it is not obviously nonempty. Let β be any ordinal. Now consider the first β lumps of X_B and the first β lumps of X_A. If they are π-isomorphic then β is concordant. If not, then wlog $X_B{\restriction}\beta$ is π-isomorphic to some end-extension of $X_A{\restriction}\alpha$. Take the union of the lumps that meet this end-extension. This will be $X_B{\restriction}\beta_1$. Repeat **sort this out**

This is isomorphic to some end-extension burble **closure points**.

Now suppose A and B are stationary sets whose intersection is countable. Let $C \subseteq \omega_1$ be the set of closure points. $C \cap A$ and $C \cap B$ are (i) nonempty (since A and B are stationary and C is club) and (ii) distinct. If $C \cap A = C \cap B$ then both these sets—which are uncountable—are $C \cap (A \cap B)$ which is a subset of $A \cap B$ which is countable by assumption. So there is one point $\alpha \in C$ which belongs to A but not to B (wlog). So $X_A{\restriction}\alpha \simeq_\pi X_B{\restriction}\beta$ but the terminal segment of X_B at this point has a first element (by construction, since $\alpha \notin B$ so that lump was $1 + Q$, and the corresponding terminal segment of X_A doesn't. So they are non-isomorphic!! ∎

DEFINITION 49. *An ordinal (or a cardinal) is* **indecomposible** *iff it is not the sum of two smaller ordinals (cardinals).*

DEFINITION 50. *If $f : I \to P$ and $g : I \to P$ are I-sequences of elements of a poset $\langle P, \leq \rangle$, then f is* **cofinal** *in g if $(\forall x)(\exists y)(f(y) \geq g(x))$.*

This definition makes sense in a context—namely the ordinals—much more general than the one in which we shall use it. f and g do not have to be wellordered sequences for definition 56 to make sense. I doesn't have to be a total order, tho' in all cases of interest to us here it will be. The sequences alluded to in the next definition are (wellordered) increasing sequences of ordinals.

DEFINITION 51. *The cofinality of α, written '$cf(\alpha)$' is the least ordinal the length of a cofinal subsequence of something of length α.*

Notice that the relation 'f is cofinal in g' is *transitive*.

DEFINITION 52. *An ordinal α is* **regular** *if $\alpha = cf(\alpha)$. Otherwise it is* **singular**.

Miniexercise: cf is idempotent ($cf(cf(\alpha)) = cf(\alpha)$) because of transitivity, so all cofinalities are regular.

> I mentioned earlier the important triviality that every normal function has a fixed point. This is true because we can always obtain a fixed point by iterating ω times. This gives us fixed points of cofinality ω. The assertion that normal functions have **regular** fixed points is a large cardinal axiom.

EXERCISE 22. *Prove that ω_1 (the first uncountable ordinal) is regular. You may use the axiom of countable choice.*

(Without AC the only (infinite) ordinal we can prove to be regular is ω, tho' finding models of ZF where all infinite ordinals are singular is very difficult indeed. Come to think of it, you may well wonder how can you can be sure that there are any uncountable ordinals in the first place. This is a consequence of a deeply mysterious theorem called **Hartogs' theorem**, which states that for every set x there is a wellorderable set y which cannot be injected into x. The proof involves a difficult discussion that we haven't got time for, so the following arm-waving proof will have to do. First prove that every ordinal is the order type of the set of all ordinals below it in their natural ordering. Then think about the set of countable ordinals: its order type must be uncountable.)

The fact that ω_1 is regular means that we cannot reach it by any countable iteration of a continuous function from On into itself: Think of any operation that takes countable ordinals to countable ordinals, iterate it ω times and take the `sup`, the result is never ω_1—because the way it is generated ensures that it is of cofinality ω!

EXERCISE 23. *Prove that if f is $cf(\alpha)$-continuous then it is α-continuous.*

Cardinals pertaining to ordinals

Definition 53. *An* **initial** *ordinal is one such that the carrier set of any wellordering of that length is cardinally larger than the carrier set of any wellordering of any shorter length.*

Thus initial ordinals are precisely those ordinals that are α-like for some α.

Exercise 24. *Assuming AC we can generalise exercise 23 to show that for every ordinal α the $\alpha + 1$th initial ordinal is regular. (This is standard set theory but we won't need it). Why does this not show that every initial ordinal is regular?*

Assuming *full* AC (as is common in the study of wellfounded sets) every cardinal corresponds to a unique initial ordinal. The αth (infinite) initial ordinal is ω_α (We start counting at '0' and we **always** omit the subscript '0' in 'ω_0'!) and the corresponding cardinal number is $\aleph_\alpha$.

This notation makes sense even without AC. A cardinal of a wellorderable set is called an **aleph** and the collection of alephs is naturally wellordered. The αth aleph is notated '$\aleph_\alpha$'.

Usual dire warning about the difference between ω^ω and ${\aleph_0}^{\aleph_0}$.

Next we show

Theorem 54. *Cofinalities are initial ordinals.*

Proof: Fix $\langle X, \leq_X \rangle$ a wellordering of length ζ, with ζ regular. Suppose further that κ is the initial ordinal corresponding to ζ and $\kappa < \zeta$. We will obtain a contradiction. We enumerate X (in a different order) as a κ-sequence: $\langle X, \leq_\kappa \rangle$. Delete from X any element which is $\leq_X$ something which is $\leq_\kappa$ of it. What is left is a subset of X cofinal in X in the sense of either ordering and which is of length κ at most, contradicting regularity of ζ ■

So every regular ordinal is initial. So every countable ordinal $> \omega$ is singular. So it has smaller cofinality. This cofinality cannot be a smaller countable ordinal $> \omega$ because cofinality is idempotent. So

Remark 55. *Every countable limit ordinal has cofinality ω.*

■

Lemma 56. *If $\kappa > \omega$ is a regular initial ordinal, then the collection of α-club subsets of the ordinals below κ is closed under $< \kappa$ intersections.*

Proof:

Let $\{B_\beta : \beta < \lambda < \kappa\}$ be a family of α-club subsets of (the ordinals below) κ. The intersection $\bigcap_{\beta<\lambda} B_\beta$ is clearly α-closed, so we have to show that it is unbounded.

Pick $\nu < \kappa$. Define by recursion

$\nu_{0,0}$ = least thing in B_0 bigger than ν;

$\nu_{0,\beta}$ = least thing in B_β bigger than all $\nu_{0,\gamma}$ with $\gamma < \beta$.

Remember, each $\nu_{0,\beta} \in B_\beta$. Since κ is regular and $\lambda < \kappa$ the sup of all the $\nu_{0,\beta}$ is less than κ. Therefore we can repeat the process starting with this sup, which we call $\nu_{1,0}$. This gives us a family $\langle \nu_{\beta,\gamma} : \beta, \gamma < \lambda \rangle$ of ordinals where $\nu_{\beta,\gamma} \in B_\gamma$ and if $\beta < \beta'$ then $\nu_{\beta,\gamma} < \nu_{\beta',\gamma}$

Now all the sequences $\langle \nu_{\eta,\beta} : \eta < \alpha \rangle$ have the same sup. Therefore this common sup belongs to all the B_β since they are all α-closed. Therefore it is an element of the intersection of them all which is bigger than ν. ν was arbitrary, so the intersection is unbounded as desired.

■

Lemma 57. *Clubsets of $\{\alpha : \alpha < \kappa\}$ are closed under diagonal intersection. (Probably need $cf(\kappa) > \omega$).*

Let $\langle C_\alpha : \alpha < \kappa \rangle$ be a sequence of clubsets of ordinals below κ. Let C be the diagonal intersection.

(i) C is closed. Let λ be a sup of elements of C. For β an element of a given C-sequence S tending to λ we have $\beta \in \bigcap_{\alpha<\beta} C_\alpha$. Let's write this as '$\beta \in A_\beta$' to save trees. The A_β form a nested $\subseteq$-sequence of closed sets, and we want to know that λ belongs to the intersection of them. let β_0 be an arbitrary member of S: we want $\lambda \in A_{\beta_0}$. Now λ is the sup of a lot of later βs, all of which are in A_{β_0}. After all, $\beta \in C$ whence $\beta \in \bigcap_{\gamma<\beta} C_\gamma$ so in particular $\beta \in A_{\beta_0}$ But A_{β_0} is closed, so it will contain the sup of all those β, namely λ.

(ii) C is unbounded. We will show that C has a subset that is club and so must be unbounded. For each α the set $\bigcap_{\beta<\alpha} C_\beta$ is club. It is closed because any intersection of closed sets is closed, and by lemma 62 it's unbounded as long as κ is regular. So the function sending each $\alpha < \kappa$ to the αth member of $\bigcap_{\beta<\alpha} C_\beta$ is a normal function and has a closed unbounded set of fixed points. Any fixed point for this function is in C.

■

Theorem 58. *"Fodor's theorem"*[3]

(AC). Let $\kappa > \omega$ be regular, and $S \subseteq \kappa$ stationary. Let $f : S \to \kappa$ be regressive[4] *($(\forall \alpha \in S)(f(\alpha) < \alpha)$). Then $\exists \alpha < \kappa$ f^{-1}"$\{\alpha\}$ is stationary.*

Proof: Suppose not. For each $\alpha < \kappa$ pick C_α club and disjoint from f^{-1}"$\{\alpha\}$. Let C be the diagonal intersection of the C_α's. C is club so $C \cap S \neq \emptyset$. Consider $\alpha \in C \cap S$. Since $\alpha \in S$ we have $f(\alpha) < \alpha$ so and $\alpha \in f^{-1}$"$\{f(\alpha)\}$. Now $\alpha \in C$, which is to say $\alpha \in \bigcap\{C_\beta : \beta < \alpha\}$ so in particular $\alpha \in C_{f(\alpha)}$, since $f(\alpha) < \alpha$, contradicting $C_{f(\alpha)} \cap f^{-1}$"$\{f(\alpha)\} = \emptyset$.

■

Think of this as saying something like "ordinals are wellfounded" only more so.

Imre sez that the core of the proof of Neumer is in the following factoid. Sse f is a pressing down 1-1 function from countable limit ordinals to countable ordinals.

[3] Imre says it was actually proved by Walter Neumer.

[4] The slang expression is "pressing down".

We will show it is not 1-1. Without loss of generality we can suppose that f is onto an initial segment of the countable ordinals, since if it isn't we can just "trickle" the values down. Set $t_1 =: \omega$ and thereafter t_{n+1} is $\sup\{\alpha : f(\alpha) \leq t_n\}$. There are only countably many α s.t. $f(\alpha) \leq t_n$ (can't map an uncountable set onto a countable set in a 1-1 way!) Since f is pressing down we will have $t_n < t_{n+1}$ for all n. So $t_\omega =: \sup\{t_n : n \in \mathbb{N}\}$ is limit. But then $f(t_\omega) < t_\omega$, whence $f(t_\omega) < t_n$ for some n. But then $t_\omega \leq t_{n+1}$ by definition of t_{n+1}.

We can give a particularly cute proof of lemma 63 using Neumer's theorem: If C is not closed unbounded then it is disjoint from a stationary set. That is to say there are stationarily many $\alpha < \kappa$ s.t. $a \notin \bigcap\{C_\beta : \beta < \alpha\}$. Send each such miscreant α to the least β s.t. $\alpha \notin C_\beta$. This is pressing down and so, by Neumer's theorem (thm 64), is constant on a stationary set. So for some β the set of α that aren't in C_β is stationary, contradicting the clubness (clubhood?) of C_β.

Of course this is no good unless we have a proof of Neumer's theorem!

Take a detour at this point to the file on countable ordinals.

Exercise 25. *Use Cantor Normal Forms to show that every ordinal can be expressed as a sum of powers of 2.*

Exercise 26. *The game of Sylver Coinage was invented by Conway, Berlekamp and Guy [4]. It is played by two players,* `I` *and* `II`*, who move alternately, with* `I` *starting. They choose natural numbers greater than 1 and at each stage the player whose turn it is to play must play a number that is not a sum of positive multiples of any of the numbers chosen so far. The last player loses.*

Notice that by 'sum of multiples' we mean 'sum of positive *multiples'. The give-away is in the name: 'Sylver Coinage'. What the players are doing is trying at each stage to invent a new denomination of coin, one that is of a value that cannot be represented by assembling coins of the denominations invented so far. (There is a significance to the spelling of 'silver', but I do not think we need to concern ourselves with that.)*

Prove that no play of this game can go on forever.

The way to do this is to identify a parameter which is altered somehow by each move. The set of values that this parameter can take is to have a well-founded relation defined on it, and each move changes the value of the parameter to a new value related to the old by the well-founded relation. The question for you is, what is this parameter? and what is the well-founded relation?

(You should give a much more rigorous proof of this than of your answer to exercise **??** *below: it is quite easy to persuade oneself that all plays are indeed finite as claimed, but rather harder to present this intuition as reasoning about a well-founded relation.)*

Exercise 27. *What is the ordinal of the game of Sylver coinage?*

Exercise 28. *Verify that the class of wellorderings is closed under substructure and cartesian product.*

Exercise 29. *Verify that the end-extension relation between wellfounded binary structures is wellfounded.*

EXERCISE 30. *Verify that the transitive closure of a wellfounded relation is wellfounded.*

EXERCISE 31. *Use Cantor normal form on the ordinals below* ϵ_0 *to give a model of a multiset version of ZF minus infinity but with foundation.*

This is fun, but it's hardly important!

DEFINITION 59. *Limit cardinal, strong limit cardinal. 0-mahlo is wk inacc and (strong?) you are* $n+1$ *Mahlo iff the* n*-Mahlo's below you are stationary.*

4.4 Some elementary cardinal arithmetic

LEMMA 60. *Bernstein's lemma.*

COROLLARY 61. $(\forall$ *cardinals* $\alpha)(\alpha = \alpha^2) \to \mathrm{AC}$

We did this in Part II

Notation: κ^+.

PROPOSITION 62. $(\forall\aleph)(\aleph = \aleph^2)$

Proof: We did this in part II

■

THEOREM 63. *The Jordan-König theorem* (AC)*:*

If $\langle A_i : i \in I\rangle$ *and* $\langle B_i : i \in I\rangle$ *are families of sets such that* $(\forall i \in I)(|A_i| < |B_i|)$ *then*

$$|\bigcup_{i\in I} A_i| < |\prod_{i\in I} B_i|$$

Proof:

We did this in part II

■

This is a sort of infinite-dimensional version of Bernstein's lemma.

COROLLARY 64. $\alpha < \alpha^{cf(\alpha)}$

■

PROPOSITION 65. *(*AC*). If* λ *is singular strong limit then* $2^\lambda = \lambda^{cf(\lambda)}$

Proof:

$\lambda < 2^\lambda$ so $\lambda^{cf(\lambda)} \leq (2^\lambda)^{cf(\lambda)} = 2^{\lambda\cdot cf(\lambda)} \leq 2^{\lambda^2} = 2^\lambda$. The other direction is a bit harder.

Express λ as $\bigcup_{i<cf(\lambda)} A_i$. Want to code any subset of λ as a $cf(\lambda)$-sequence of elements of λ. Let $X \subseteq \lambda$. Consider the bits $X_i = X \cap A_i$. If each X_i is coded by a member of λ we are ok.

Now, since λ is strong limit, $\bigcup_{i<cf(\lambda)} \mathcal{P}(A_i)$ is also of size λ and so X_i is coded by something in λ as desired.

■

We will need this in the case of λ of cofinality ω when we come to prove corollary 128.

DEFINITION 66. *We saw earlier than an* **aleph** *is a cardinal of a wellordered set.* $\aleph(\alpha)$ *is the least aleph* $\not\leq \alpha$*. This is* **Hartogs' aleph function**. *Notice that there is no notation for the first ordinal that is not the length of a wellordering of any set of size* $\leq \alpha$*: we really do have to exploit the nasty hacky identification of cardinals with initial ordinals.* $\aleph_\beta$ *(β an ordinal) is the βth* **infinite** *aleph.*

We'd better show that this is defined!

THEOREM 67. *Hartogs' theorem.* $\aleph(\alpha) < 2^{2^{2^\alpha}}$

We did this in part II

EXERCISE 32. *Prove without using the axiom of choice that for any two cardinals α and β, if $2 \cdot \alpha = 2 \cdot \beta$ then $\alpha = \beta$.*

EXERCISE 33. $\aleph_\alpha \leq_* \beth_\alpha$

EXERCISE 34. *Not true without AC that every infinite set has a countable subset, but every infinite subset of* $\mathbb{R}$ *has a countable partition. (Fairly easy)*

THEOREM 68. AC_{wo} *All values of the Hartogs' aleph function are regular successor initial ordinals.*

Proof:

First we establish that $\aleph(|X|)$ is a successor aleph.

Suppose $\kappa = \aleph(|X|)$ were a limit cardinal. Then for all cardinals $\alpha < \kappa$ ($\exists$ subset $X_\alpha \subset X)(\exists R_\alpha \subseteq X_\alpha \times X_\alpha)(\langle X_\alpha, R_\alpha\rangle$ is a wellordering of length α). Concatenate them to get a wellordering of length κ which is embedded in X contradicting $\kappa = \aleph(|X|)$. How?! Might they not overlap? There are various ways out of this.

1. We could use DC to pick the X_α so that they are all disjoint. This involves a bit of fiddling around with Bernstein's lemma (lemma 66) to show that the complements are big enough. But DC is a bit strong.

2. We rule that each x in X is counted only at its first appearance, so that we are concatenating on the end not the *whole* of X_α, but only those things that have not already appeared. That way we get a wellordering and we know that after α steps the wellordering is of size at least α, so its length is the sup of all these, which is κ.

So we now know that κ is successor, so reletter it as κ^+. We must show that a thing K^+ of size κ^+ is not a union of at most κ things of size at most κ. If $K^+ = \bigcup_{i<\kappa} K_i$, then, using AC_{wo}, pick for each i an injection $K_i \hookrightarrow K$ where

K is some fixed set of size κ. Then $\bigcup_{i<\kappa} K_i \hookrightarrow K \times \{i : i < \kappa\}$ which is of size $\kappa^2 = \kappa$.

■

We write α adj β iff $\alpha < \beta$ and there is no cardinal γ with $\alpha < \gamma < \beta$. The **generalised continuum hypothesis** (**GCH**) is usually taken to be $2^{\aleph_\alpha} = \aleph_{\alpha+1}$. ($\alpha$ is an ordinal here). This presupposes the axiom of choice. There is a version that doesn't, namely $(\forall\alpha)(\alpha \text{ adj } 2^\alpha)$.

Exercise 35. *Assume GCH in the form $(\forall\alpha)(\alpha \text{ adj } 2^\alpha)$. Use Hartogs'-Sierpinski to deduce AC from it. (Hint: If α adj 2^α adj 2^{2^α} adj $2^{2^{2^\alpha}}$ think about $\aleph(\alpha)$. It is possible (for example using Exercise* **??***) to prove that if α adj 2^α adj 2^{2^α} then 2^α is an aleph. It seems to be still an open question whether or not α adj 2^α implies that α is an aleph.*

Definition 69. *For α an aleph, $\alpha^{(\mu)}$ is the μth aleph $> \alpha$.*
Let us write $T{\restriction}\alpha$ for $\{a \cap \alpha : a \in T\}$.

The following proof comes from notes taken by Frank Drake of a lecture by Prikry on Silver's theorem.

Theorem 70. *(Prikry)*
Let κ be a singular cardinal such that $cf(\kappa) = \lambda > \omega$ and $\alpha < \kappa \to \alpha^\lambda < \kappa$. Suppose $\mu < \lambda$ and $T \subseteq \mathcal{P}(\kappa)$. Then if $\{\alpha < \kappa : |(T{\restriction}\alpha)| \le \alpha^{(\mu)}\}$ is stationary in κ, then $|T| \le \kappa^{(\mu)}$.

Proof:
By induction on μ. Take $C = \{\alpha_\zeta : \zeta < \lambda\}$ a strictly increasing continuous sequence of cardinals with limit κ. Set $S = \{\zeta < \lambda : |T{\restriction}\alpha_\zeta| \le (\alpha_\zeta)^{(\mu)}\}$. S is stationary in λ by the hypothesis (since in general A stationary in κ iff $A \cap C$ stationary in κ iff $\{\zeta : \alpha_\zeta \in A\}$ stationary in λ).

Given $\zeta \in S$, let $f_\zeta : T{\restriction}\alpha_\zeta \to (\alpha_\zeta)^{(\mu)}$ be $1-1$; and given $a \in T$, let $g_a(\zeta) = f_\zeta(a \cap \alpha_\zeta)$ for $\zeta \in S$.

Case $\mu = 0$

(Drake comments that this case is dealt with by Erdős, Hajnal and Milner [13]

Then $g_a(\zeta) < \alpha_\zeta$ for $\zeta \in S$. So if ζ is a limit, $g_a(\zeta) < \alpha_\eta$ for some $\eta < \zeta$ (since C is cts); let $h_a(\zeta)$ be the least such. If $S_0 = \{\zeta \in S : \zeta \text{ is a limit}\}$, then S_0 is stationary too, and $h_a : S_0 \to \lambda$ is regressive. Therefore, by theorem 64, it is constant on a stationary set.

So for each $a \in T$, fix $S_a \subseteq S_0$ and $\eta(a) < \lambda$ s.t. S_a is stationary in λ and $h_a(\zeta) = \eta(a)$ constantly for $\zeta \in S_a$.

Now there are at most 2^λ pairs $\langle S_a, \eta(a)\rangle$ possible, and $2^\lambda < \kappa$ by assumption. So, given $\langle S', \eta'\rangle$, S' stationary in λ and $\eta' < \lambda$, let

$$T' = \{a \in T : S_a = S' \wedge \eta(a) = \eta'\}$$

We show $|T'| \le \kappa$, and the result follows because in that case $|T| \le 2^\lambda \cdot \kappa = \kappa$.

Since $a \in T' \to g_a(\zeta) < \alpha_{\eta'}$ for any $\eta \in S'$, we must have $|T' \restriction \alpha_\zeta| \leq \alpha_{\eta'}$ for $\zeta \in S'$; and if $a, b \in T$ with $a \neq b$ then $\langle a \cap \alpha_\zeta : \zeta \in S' \rangle$ and $\langle b \cap \alpha_\zeta : \zeta \in S' \rangle$ are distinct sequences.

So $|T'| \leq (\alpha_{\eta'})^\lambda < \kappa$ by hypothesis and case $\mu = 0$ follows.

Case $\mu > 0$, μ **limit**

Fix $a \in T$, then $g_a(\zeta) < (\alpha_\zeta)^{(\mu)}$, so $g_a(\zeta) < (\alpha_\zeta)^{(\nu)}$ for some $\nu < \mu$, for each $\zeta \in S$.

So this splits S into $\leq \mu$ pieces, and $\mu < \lambda$, so at least one is stationary. Say $S_a \subseteq S$ is stationary, and $\nu(a) < \nu$, and $\zeta \in S_a \to g_a(\zeta') < (\alpha_\zeta)^{(\nu(a))}$.

Again, fix $S' \subseteq S$ stationary, and $\nu' < \mu$; by the induction hypothesis for case ν', $|\{a \in T : S_a = S' \wedge \nu_a = \nu'\}| \leq \kappa^{(\nu)}$. But again there are $\leq 2^\lambda$ pairs $\langle S', \nu' \rangle$ so $|T| \leq \kappa^{(\mu)}$.

Case $\mu = \nu + 1$ (main case)

First pick an ultrafilter $\mathcal{U}$ on S so that $X \in \mathcal{U} \to X$ is stationary on λ. (The usual proof that a filter extends to an ultrafilter adapts to show that the closed unbounded filter on S extends to such a $\mathcal{U}$: at each stage, make sure that the set added is stationary, using the fact that if $X \sqcup X' = S$ is stationary then $\to$ X' or X are stationary.) Now think about the ultraproduct

$$(\prod_{\zeta \in S} (\alpha_\zeta)^{(\mu)})/\mathcal{U}.$$

This ultraproduct has a canonical linear order (by Łoś's theorem) and this induces an order $\prec$ on T by:

$$a \prec b \text{ iff } \{\zeta \in S : g_a(\zeta) < g_b(\zeta)\} \in \mathcal{U}.$$

This will be a **linear** order on T, since if $a \neq b$ then $a \cap \alpha_{\zeta_0} \neq b \cap \alpha_{\zeta_0}$ for some $\zeta_0 \in S$, and hence $g_a(\zeta) \neq g_b(\zeta)$ for $\zeta > \zeta_0$; i.e., $\{\zeta : g_a(\zeta) \neq g_b(\zeta)\} \in \mathcal{U}$.

Claim: The order-type of T under $\prec$ is $\kappa^{(\mu)}$-like.

Proof of claim: Fix $b \in T$ and look at $T_b = \{a \in T : a \prec b\}$. For $a \in T_b$, let $S_a = \{\zeta \in S : g_a(\zeta) < g_b(\zeta)\}$. Since $S_a \in \mathcal{U}$, S_a is stationary on λ. So fix S' stationary in λ, and let $T' = \{a \in T_b : S_a = S'\}$. Then for $\zeta \in S'$, $g_a(\zeta) < g_b(\zeta) < (\alpha_\zeta)^{(\mu)}$ for each $a \in T'$. But $g_\beta(\zeta)$ has cardinal $\leq (\alpha_\zeta)^{(\mu)}$, since $(\alpha_\zeta)^{(\mu)} = ((\alpha_\zeta)^{(\nu)})^+$, and so $|T' \restriction \alpha_\zeta| \leq (\alpha_\zeta)^{(\nu)}$, for each $\zeta \in S'$. So we can apply the induction hypothesis (case ν) to T', and get $|T'| \leq \kappa^{(\nu)}$. Since (again) there are only 2^λ such S', $|T'| \leq \kappa^{(\nu)}$ too, which is the claim. But any $\kappa^{(\mu)}$-like ordering has cardinal $\leq \kappa^{(\mu)}$. That is to say $|T| \leq \kappa^{(\mu)}$.

■

Corollary 71. *Suppose* $\omega < cf(\kappa) = \lambda < \kappa$, $\mu < \lambda$ *and* $\{\alpha < \kappa : 2^\alpha \leq \alpha^{(\mu)}\}$ *is stationary in* κ. *Then* $2^\kappa \leq \kappa^{(\mu)}$

In particular, if the GCH holds for a stationary set of cardinals below κ, it holds at κ.

■

4.5 Games

There seems to be a deep question about what the general form is of a discrete game. Until further notice a discrete game is defined in the first instance by an **arena**, A, which is the set of elements that the two players, `I` and `II`, can play. (The first person pronoun is 'i'.[5] `I` is male and `II` is female. Altho' there are deep and obscure game-theoretic reasons for this which we might yet get round to, my motive here is simply to use the gender structure of English pronouns to help distinguish the players.) Elements of A can be reused, so some readers might prefer to think of A as a set of *streams*. `I` and `II` play alternately, `I` starting, thereby building a member of A^ω. A member of A^ω thus constructed is a **play**. A finite initial segment of a play (i.e., a member of $A^{<\omega}$) is a **position**. $A^{<\omega}$, the set of positions has an obvious tree structure. If there is an upper bound k on the length of paths thru' the tree, then the tree is said to be **of height** k. If T is a tree, $[T]$ is the set of all paths through T. Thus $A^\omega = [A^{<\omega}]$. (In general if T is infinite we need DC to prove that $[T] \neq \emptyset$: there will be more on this later).

For any position p, the set of end-extensions of p is a subtree. If t is a subtree and $\{t_i : i \in I\}$ are subtrees, then $t \setminus \bigcup_{i\in I} t_i$ is a subtree. (Well, almost: when deleting a position one must also delete all its children) A **`I`-imposed subgame** is a subtree containing all children of all its even positions, a **`II`-imposed subgame** dually.

We call positions **even** or **odd** according to the parity of their distance from the root of the tree: `I` plays when the game is at an even position; `II` plays when the game is at an odd position. We will need the notion of a **strategy**. A strategy for `I` is a function defined on all *even* positions that returns positions of length one greater than the length of its argument; a strategy for `II` is a function defined on all *odd* positions that returns positions of length one greater than the length of its argument. Thus the distinction good English makes between *strategy* and *tactic* is not respected here.

If T is a subtree of $A^{<\omega}$ then a **game** is a function G from $[T]$ to $\{\mathtt{I},\mathtt{II}\}$, namely a function that says which of $\{\mathtt{I},\mathtt{II}\}$ has won any play of the game corresponding to v. $\{a \in A^\omega : G(a) = \mathtt{I}\}$ (or $G^{(-1)}$"$\{\mathtt{I}\}$ if you prefer), the set of plays won by `I`, is sometimes called the **payoff set** of the game. If B is the payoff set of G we will customarily denote with 'G_B' the game G whose payoff set is B.

It is natural to look for descriptions of the payoff set that tell us significant things about the game, such as whether or not it has a winning strategy for one player or the other. How can we describe subsets of A^ω? If A has algebraic structure of any kind, A^ω will have that structure as well, but even if A has no structure at all A^ω will have the structure of a product space, since A can be given the structure of a discrete topological space—which is no structure! We give A the discrete topology and A^ω the product topology. For any topological

[5] "...one whom your present interlocutor is accustomed to denote by means of the perpendicular pronoun"—Sir Humphrey Appleby

property (“open”, “closed”, “Borel” ...) we will speak of a game as having that property when what we really mean is that the payoff set has that property.

The reason why this is a sensible approach is that we will eventually be able to prove that if the payoff set is a Borel subset of A^ω then the corresponding game is **determined**—one of the two players has a winning strategy. This is a hard theorem, but its naturalness is underlined by the fact that stronger conditions than Borel-ness make for easier proofs of determinacy.

In all the games that follow I shall use the convention that ‘Wins’ with a capital ‘W’ means ‘has a winning strategy for’. A winning strategy for `I` (resp. `II`) can be tho’rt of as a `I`-imposed (resp. `II`-imposed) subtree T of $A^{<\omega}$ such that $[T] \subseteq P$ (resp. $[T] \cap P = \emptyset$) where P is the payoff set. Notice that a strategy can be winning without being total: it doesn’t need to tell you how to get out of positions that it never led you into in the first place.

A **nondeterministic strategy** or *policy* or **restraint** or **nds** is simply a many-valued strategy. To be rigorous, a nds for `I` is a function defined on even positions sending each argument to a set of its children (a subset of its litter); dually for `II` and odd positions. I shall probably stick to ‘nds’ co’s it’s shortest. One can think of the value of the nds at a position p as the set of things that the nds allows/recommends you to do. Of course it can also be tho’rt of as a binary relation between positions rather than a function from positions to sets of positions, and this is often more convenient.

define litter somewhere

In some games on A^ω you are forbidden to make certain moves. This can be handled by the theory-maker in one of two ways. You can define the game as play on a closed subset, or you can rule that any player who plays outside that closed subset loses instantly. These two descriptions are equivalent: the end result is the same, and the only difference is the time you have to wait for this to become manifest. However this equivalence relies on the axiom of dependent choice, in a way we had better now illustrate.

Consider the Dedekind-finite set S of socks which we first met in Russell’s discussion of the axiom of choice in *Introduction to Mathematical Philosophy*. It is divided into countably many pairs $\{S_i : i \in \mathbb{N}\}$ and is located in the infinite attic of the millionaire. The millionaire’s Valet (`I`) and Maid (`II`) play a game. They pick socks from the collection. (Perhaps “indicate” would be better than ‘pick’.) The first servant to indicate a sock already indicated earlier in the game loses. Draws are impossible because the game cannot go on for ever: an infinite play would be a countable subset of S, and there are none. The Valet has a nondeterministic winning strategy σ_v: indicate a sock not yet indicated: there are always plenty. Indeed his strategy is to indicate a sock —either will do— from the first pair not yet used. Incredibly the Maid has a winning strategy too, and hers—call it σ_m—is actually *deterministic*! All she has to do is indicate the sock that is the mate of the sock pointed to by `I` in the move to which she is replying.

The moral is as follows. We made a decision to think of discrete games as subsets of A^ω. That subsumes games in which it has become clear at some finite stage who has won. (And i don’t just mean “becomes clear at a finite stage who has a winning strategy from that point”; i refer to the situation where there is a player who wins from that point *whatever happens subsequently*). It's more natural to think of games like that as games of finite length, but there is a generality to be gained by making all games look the same, and therefore to think of this as a game that goes on for ever even tho' the outcome of an play is known after some finite stage of that play. However this endeavour to coerce all games into this form works only if one has countable choice. If choice fails, the millionaire's sock provide us with an example of a discrete game which is well-defined but where the payoff set cannot be thought of as a subset of A^ω for any arena A. The Maid-Valet game—if coerced into the form of a subset of A^ω—has no plays. If we don't attempt to coerce it into this form, but leave it as a game where all plays are finite, we find that any attempt on the part of the two players—both using their optimal strategies σ_v and σ_m—to play a *partie* will result in the death of one or both players. *There is no play resulting from the interaction of σ_v with σ_m!!*

There is another moral to this story (there is always more than one moral!) The tree of the Maid-Valet game is an example of a structure that is wellfounded in the no-descending sequences sense, but not in the true sense—and certainly not in any useful sense. It may lack infinite paths, but then it has no endpoints either, so there are no initial labellings to which we can do an E as in definition 78 below.

Now it is well-known that *DC* is weaker than full AC, and there are plenty of occasions when it is sensible to assume the first axiom but not the second. Game theory is probably not one of them. Borel determinacy, which we will prove later, really does need full AC, and if we assume AC earlier we can give smoother, more abstract proofs of assertions like Gale-Stewart which don't actually need AC at all. AC is provable (well, classically) in finite universes so the question of whether or not it is wise to adopt AC only ever arises when we are handling infinite objects. The theory of finite games is pretty straightforward, and this point is being made implicitly by our characteristion of a game as a subset of A^ω. Although there are no doubt positive results about subsets of A^ω that one can prove without any AC at all, we will be assuming AC anyway, not least beco's of the moral pointed by the Valet-Maid game.

We have defined a game to be a payoff function $G : A^\omega \to \{\mathtt{I}, \mathtt{II}\}$. However, in order to win a play of a game we need a *strategy*, namely a function defined on *positions* rather than a function—like G—defined on *plays*. Fortunately in certain circumstances we can process a payoff function G into a function defined on at least some positions, and we can do this when there are positions such that every play through them is won by the same player. This certainly happens if the payoff set $G``[A]$ is closed or open, but it will also happen even if it merely has nonempty closed or open subsets. Let us use the word **valuation** for (possibly partial) functions sending positions to $\{\mathtt{I}, \mathtt{II}\}$.

If $G(\pi) = \mathtt{I}$ for every play π that is an end-extension of p then we can sensibly process G into a valuation sending p to $\mathtt{I}$. For a fixed arena and payoff set let us call the valuation v obtained from G in this way the **base valuation**.

Now there is an obvious way of extending valuations. I call it E (for "extend label"), and it is defined as follows. If v : positions $\to \{\mathtt{I}, \mathtt{II}\}$ then

Definition 72.

$E\ v\ p =$:
`if` p *is even* `then (if` *there is a child* p' *of* p *with* $v(p') =$ `I`*)*
`then I`*;*
`else if` $v(p') =$ `II` *for all children* p' *of* p `then II`*;*
`else if` p *is odd* `then (if` *there is a child* p' *of* p *with* $v(p') =$ `II`*)*
`then II`*;*
`else if` $v(p') =$ `I` *for all children* p' *of* p `then I else fail`.

Clearly for any fixed arena and payoff set, the collection of valuations forms a chain-complete poset under inclusion (valuations tho'rt of as sets of ordered pairs) and the base valuation is the bottom element of this poset. E is clearly a monotone function from this poset into itself, so there will certainly be fixed points for E. The fixed points form a chain-complete poset, so there will even be fixed points by Bourbaki-Witt.

Any fixed point v for E will give rise to a pair of canonical nondeterministic strategies. I call them **soot** strategies. It is the stay out of trouble strategy, which, for player i, is to play nodes labelled i wherever possible and please yourself otherwise.

Now suppose G is an open (if player `I` is to win this has become apparent by some finite position) or closed (if player `II` is to win this has become apparent by some finite position) game these maximal fixed points become interesting, and for two reasons. (i) In an open or closed game a soot strategy defined at the empty position is winning; (ii) In an open or closed game a maximal valuation must be total.

Proof of (ii) Suppose not, and let $v = E(v)$ be a maximal fixed point that is not a total function. If i is the player that wins every play whose fate is not determined by a finite initial segment, then add to v all ordered pairs $\langle p, i\rangle$ for all positions p at which v is undefined. The result is a total function, and is still a fixed point for E.

This has just proved

Theorem 73. *(Gale-stewart) Every open or closed game is determined.*

■

We can even show that

Remark 74. *G_B is determined if B is an intersection of countably many open sets.*

Proof: Let us write B for the payoff set: G^{-1}"{`I`}. Suppose $B = \bigcap_{i\in\mathbb{N}} B_i$ is an intersection of countably many open sets. (I'm not making any assumptions about the arena.) If `I` lacks a winning strategy for even one of the G_{B_i} then `II` has a winning strategy for G_B: the A_i are determinate, so if `I` has no winning strategy for G_{B_i} then `II` must have, and any winning strategy for `II` will also be winning for her in G_B. So for the rest of this discussion we assume that $\mathtt{I}$ has a winning strategy for each B_i.

Each B_i gives rise to a labelling v_i of the nodes as $\mathtt{I}$-good or $\mathtt{II}$-good. We now cook up a composite labelling v defined so that $v(p) = \mathtt{I}$ iff $(\forall i \in \mathbb{N})(v_i(p) = \mathtt{I})$.

However, this is not the end of the matter, since some of the positions we have labelled $\mathtt{I}$-good might not really be very good for $\mathtt{I}$ at all, since altho' he has winning strategies for all the B_i he has no move that keeps *all* his options open. That is to say there may be an even position p such that $v(p) = \mathtt{I}$ but for no child p' of p do we have $v(p') = \mathtt{I}$ (or an odd position p such that $v(p) = \mathtt{I}$ but for some child p' of p do have $v(p') = \mathtt{II}$. Think of these as *positions at which* $\mathtt{I}$*'s juggling fails*. He can choose to win any of them, but he must drop at least one (ball) B_i and $\mathtt{II}$ can catch it an win from there. Such positions should really be labelled $\mathtt{II}$-good. This reallocation of $\mathtt{I}$-good positions to $\mathtt{II}$-good must be allowed to propagate up the tree as before by repeated application of E. Once this cascade has run its course, we have the final labelling we need. A position that remains $\mathtt{I}$-good even after all this is a position from which $\mathtt{I}$ really does have a winning strategy. If the root node is labelled $\mathtt{I}$ then $\mathtt{I}$ can play always to nodes labelled $\mathtt{I}$-good. If he does this then *because each of the countably many sets he's trying to get into are open*, he will—for each one of them—have succeeded in getting to that one by some finite stage. So by the end of the play, he will have won every B_i. ■

The italicised passage in the last paragraph of the proof captures that part of the argument which *won't* generalise to show that an intersection of countably many determinate sets is determinate. This is worth noting because if we had a proof that an intersection of countably many determinate sets is determinate then we would have a chance of proving that all Borel games are determined by an induction on the Borel hierarchy. (Pierre de Fourcade has pointed out to me that nowhere in this proof do we seem to have used the assumption that the payoff set is an intersection of *countably* many open sets.

Borel Determinacy

But this kind of combinatorial monkeying-around does not get us very far. We need a new idea. Enter Tony Martin, 1976, with a new construction. Extra handouts here

4.6 James' and my notes on Martin's proof of Borel determinacy

In this section we have a bit more structure: A game is not a function $v : [A^\omega] \to \{\mathtt{I}, \mathtt{II}\}$ but a function $v : [T] \to \{\mathtt{I}, \mathtt{II}\}$ where T is a closed subset of A^ω. $G(A, T)$ is the game played thru' $[T]$ with payoff set (the set of plays that are wins for $\mathtt{I}$) $A \subseteq [T]$.

We will use the notation that for S a tree and s a sequence in S, $S \uparrow s$ is the subtree of sequences comparable with s.

Central to the proof will be the observation that if A is a closed payoff set then the family of winning nds's for `I` is closed under union. Similarly winning nds's for `II` in an open game.

Definition 75. *A* **covering** *of a tree T is a triple $\langle T^+, \pi, \Phi\rangle$ where*

1. *T^+ is a tree;*
2. *$\pi : [T^+] \to [T]$;*
3. *$\Phi : S(T^+) \to S(T)$, taking strategies for `I`(`II`) in T^+ to strategies for `I`(`II`) in T;*
4. *If $x \in [T]$ is a play consistent with $\Phi(s^+)$, then there is $x^+ \in [T^+]$ consistent with s^+ such that $\pi(x^+) = x$.*

A covering as above **unravels** *A, a payoff set for T, if and only if π^{-1}"A is a clopen subset of $[T^+]$.*

Remark 76. *If there is a covering that unravels A, then $G(A, T)$ is determined.*

Proof: it is actually not necessary that the covering should unravel A, but merely that $G(\pi^{-1}\text{“}A, T^+)$ be determined. This is where we need clause 4) in the definition of a covering. We will show that if s^+ is winning for player `I` (wlog) in $G(\pi^{-1}\text{“}A, T^+)$, then $\Phi(s^+)$ is winning for `I` in $G(A, T)$. Let s^+ be winning for `I` in $G(\pi^{-1}\text{“}A, T^+)$. By 4), for any play $x \in [T]$ which is a result of following $\Phi(s^+)$, we can find a play y in $[T^+]$ such that $\pi(y) = x$, and y is a result of following s^+. But s^+ is winning for `I`, so y was a win for `I`, which is to say $y \in \pi^{-1}\text{“}A$, but then $x = \pi(y)$ so $x \in A$ so x was a win for `I`.

■

$G(\pi^{-1}\text{“}A, T^+)$. will sometimes be abbreviated to G^+

Remark 77. *Coverings can be composed: if $\langle T_1, \pi_1, \Phi_1\rangle$ covers T_0 and $\langle T_2, \pi_2, \Phi_2\rangle$ covers T_1, then $\langle T_2, \pi_1 \cdot \pi_2, \Phi_1 \cdot \Phi_2\rangle$ covers T_0. Better still, if π_2 is continuous then sets unravelled by the first covering remain unravelled in the composed covering.*

Unfortunately "every Borel set of rank α is unravelled by some covering" is too weak a hypothesis to power an induction. So we make the definition:

Definition 78. *A covering $\langle T^+, \pi, \Phi\rangle$ of a tree T is a k-***covering** *(for some $k \in \mathbb{N}$) if*

1. *$\pi(x^+)\restriction n$ depends only on $x^+\restriction n$; (Henceforth in the context of k-coverings we shall regard π as a length preserving map from T^+ to T in the natural way.)*
2. *$\Phi\restriction(s^+)$ on positions of length less than n depends only on s^+ on positions of length less than n, so it can be regarded as a length and inclusion-preserving map on fragments of strategies.*

3. *$\pi\restriction(T^+)^k$ is a bijection between $(T^+)^k$ and $(T)^k$, where given some tree S we define $(S)^k$ as $\{s \in S : \mathrm{lh}(S) \leq k\}$.*

4. *$\Phi\restriction(s^+)$ is the same as s^+ for the first k moves (which is possible because of the third item above)*

What we are demanding here is a strong form of uniform continuity plus some kind of finite approximability of play on T by play on T^+.

Notice that if $k \leq m$ then an m-covering of a k-covering is a k-covering. The induction hypothesis will be "every Borel set of rank α is, for every k, unravelled by some k-covering". Most of the work comes in starting things going, i.e. in proving

LEMMA 79. *Let A be a closed payoff set for T. For every k there is a k-covering of T which unravels A.*

Proof: Without loss of generality k is even (clearly an n-covering is an m-covering for $m \leq n$). Instead of describing T^+ explicitly we give the rules of a certain infinite game.

$$\begin{array}{lllllll} \mathtt{I} & a_0, & a_2, & \dots & (a_k, T_{\mathtt{I}}) & & a_{k+2} \\ \mathtt{II} & a_1 & \dots & a_{k-1} & & (a_{k+1}, T_{\mathtt{II}}, \text{Answer}) & a_{k+3} \end{array}$$

For every j we demand $\langle a_0 \dots a_j \rangle \in T$. $T_{\mathtt{I}}$ must be an nds for `I` in games on $T \uparrow \langle a_0 \dots a_k \rangle$. `I` in effect says to `II` "i can win with $T_{\mathtt{I}}$". Answer must be "yes" or "no". If Answer is "yes" then $T_{\mathtt{II}}$ must be an nds for `II` in $T_{\mathtt{I}}$, such that $[T_{\mathtt{II}}] \subset a$. `II` in effect says to `I` "i resign". If "answer" is "no" then $T_{\mathtt{II}}$ must be $T \uparrow \tau$ for $\tau \in T_{\mathtt{I}}$ which extends $\langle a_0 \dots a_k \rangle$ and is such that $[T \uparrow \tau] \subset [T] \setminus A$. In this case `II` in effect says to `I` "You have envisaged the possibility of getting to position τ in $T_{\mathtt{I}}$: let's play it from there and i'll win". (Note that if there is no way of answering "no", then as A is closed it follows that $[T_{\mathtt{I}}] \subset A$ and `II` can surely answer "yes")

At every move j after the kth the only restriction imposed on the players is $\langle a_0 \dots a_j \rangle \in T_{\mathtt{II}}$. π is just the function which takes a position in T^+ and forgets everything but the a_i's. So 1) and 2) in the definition of a covering and 1) and 3) in the definition of a k-covering are clearly true. Also we can see already that this will unravel A because a play x^+ in $[T^+]$ such that $\pi' x^+$ is a win for `I` can only be a result of `II` saying "yes". This is because if `II` says "no" and they stay inside $T_{\mathtt{II}}$ (as they must), then `II` wins. Similarly if `II` says "yes" then `II` loses.

We now define Φ so as to make (4) in the definition of a covering hold, and this is much the trickiest part of the construction.

Construction of Φ (part I)

From s^+ we construct a function f that eats a play $x_{\mathtt{II}}$ from `II` and returns a play x and a *lift* $\pi^{-1}(x)$ which is consistent with s^+. f is strongly continuous on the space of `II`'s plays, in that we can construct `I`'s $n+1$th move only once we have `II`'s *nth*. In the game $G(A, T)$ having `II`'s nth move to hand is sufficient for `I` to know what his $n+1$th move is, but not always sufficient for him to know what he should put in the $n+1$th position in the lift. We might have to change our minds about—for example—what goes into the $k+1$th position in the lift. However we never have to change our minds about the contents of any position in the lift more than once. This construction is uniform in s^+ and so implicitly defines Φ. For the first k moves we just copy `II`'s moves into the lift, consult s^+ and tell `I` to make the move which it advises. At move k, s^+ produces a pair $\langle a_k, T_{\mathtt{I}}\rangle$ and `I` plays a_k. If by any chance $T_{\mathtt{I}}$ was winning, then life is very easy, for we then know that at position $k+1$ in the lift must be the answer "yes". For the time being we assign "yes" to position $k+1$ in the lift. There are various possibilities for $T_{\mathtt{II}} \subseteq T_{\mathtt{I}}$, and if we assign some arbitrary nds for losing $G([T_{\mathtt{I}}] \cap A, T_{\mathtt{I}})$ we might get a shock when she actually plays outside it. We can forestall these shocks by assigning the *union* of all such strategies, for then she can play outside her (provisionally assigned) strategy only by committing us to making the new assignment "no". The union of all these strategies is an nds of the kind `II` is allowed to play when saying "yes", because, as noted above, (i) the nds's available to a player trying to get into a closed set are closed under unions, and (ii) $[T_{\mathtt{I}}] \cap A$ is a closed set. If $T_{\mathtt{I}}$ was winning we thus never has to change our minds about what $T_{\mathtt{II}}$ was. However $T_{\mathtt{I}}$ might not be winning, so `II` might play outside the maximal winning nds for misère $G(T_{\mathtt{I}} \cap A, T_{\mathtt{I}})$. We now have to change our minds. If we can tell that `II` is not playing to lose $G(T_{\mathtt{I}} \cap A, T_{\mathtt{I}})$ then `I` can play to put `II` in a position τ such that $N_\tau \subseteq A$. So we rule that $\Phi(s^+)$ tells `I` to do this, and when he has done it, to continue to play on the assumption that $T_{\mathtt{II}}$ was $T \upharpoonright \tau$. It might seem odd that he is allowed to use *any* strategy to get to τ once `II` has surprised him, but any strategy will be retrospectively justified by the decision that $T_{\mathtt{II}}$ was $T \upharpoonright \tau$, and that is the set of positions comparable with τ, and they had to play to stay inside $T_{\mathtt{II}}$ anyway.

Construction of Φ (part II)

This is quite similar to part I, in that `II` is trying to make a run of play that will lift and finds this much easier when she has on the pad a play in T^+ in which she takes the "no" option at $k+1$.

To be precise let t^+ be a strategy for `II` in T^+. At positions of length $\leq k$ we copy moves into the lift as in part I. Then player `I` puts down a_k, and `II` considers the **closed** auxiliary game $G(B,T)$, where B is the set

> $\{x \in [T]$: for no $\tau \subset x$ is there $T_{\mathtt{I}}$ such that the response of t^+ in the position I $a_0(a_k, T_{\mathtt{I}})$ ... `II` $a_1 a_{k-1}$ is to say "no" and put $T_{\mathtt{II}} = T \upharpoonright \tau\}$.

If $\langle a_0 \ldots a_k\rangle$ is a won position for `I` in $G(B,T)$ then `II` lets $T_{\mathtt{I}}$ be `I`'s maximal winning nds for the game subsequent to this position. Clearly t^+ now produces a move $\langle a_{k+1}, T_{\mathtt{II}}, \text{“yes”}\rangle$ where $T_{\mathtt{II}}$ is an nds for `II` in $T_{\mathtt{I}}$ and is such that $[T_{\mathtt{II}}] \subset A$. Now—as long as play stays inside $T_{\mathtt{I}}$—`II` copies moves in the obvious way (note that as `II` is playing a strategy for T^+ and $T_{\mathtt{II}}$ is an nds for `II` **in** $\mathbf{T}_{\mathtt{I}}$ play stays inside $T_{\mathtt{II}}$ as long as it stays in $T_{\mathtt{I}}$). If `I` ever blunders out of

$T_{\mathtt{I}}$—or immediately if $G(B,T)$ is lost for him at move k—then II just plays to win $G(B,T)$ and, when a won position τ is reached, she sets $T_{\mathtt{I}}$ to be something appropriate. Subsequently she copies moves in the usual fashion. This concludes the construction.

It is obvious (at least to us) that (3) and (4) in the definition of a covering and (2) and (4) in the definition (definition 84) of a k-covering have been satisfied.

Before attacking Borel determinacy we need an essentially trivial lemma asserting the existence of a certain inverse limit.

Lemma 80. II *: For each $i \in \omega$ let $\langle T_{i+1}, \pi_{i+1}, \Phi_{i+1}\rangle$ be a $k+i$-covering of T_i. There are (i) a tree T^∞ and (ii) maps ρ_i, χ_i such that, for each $i \in \omega$,*

- *$\langle T^\infty, \rho_i, \chi_i\rangle$ is a $k+1$-covering of T_i;*
- *$\rho_i = \pi_{i+1} \cdot \rho_{i+1}$;*
- *$\chi_i = \Phi_{i+1} \cdot \chi_{i+1}$*

Proof: To make life easier we use the πs to identify $(T_i)^{k+i}$ and $(T_{i+1})^{k+i}$ for each i, so that $\pi_{i+1}\restriction(T_{i+1})^{k+i}$ is the identity map. The standard construction of an inverse limit tells us to consider "towers" as in Diagram two; we note that a tower will be ultimately constant and we will systematically confuse it with the eventual constant value. Observe that σ represents a tower iff σ is in T_i for all large i iff σ is in T_i where i is minimal with $\mathrm{lh}(\sigma) \leq\ k+i$. T^∞ will be the set of such σ. To see that it is a tree note that closure under initial segment is immediate, and σ in T^∞ implies σ in T_i with $\mathrm{lh}(\sigma)$ less than $k+i$ implies there is a one-point extension τ of σ in T_i, and $lh(\tau)$ is less than or equal to $k+i$ implies there is a one-point extension of σ in T^∞ The projection maps ρ_i will just take σ to the ith coordinate of the associated tower. Now let s^∞ be a strategy for T^∞. We claim that given n, s^∞ restricted to positions of length less than n (which we written $s^\infty n$) is a fragment of a strategy for T_i if i is sufficiently large. To be precise let σ be in T_i and let $\mathrm{lh}(\sigma)$ be less than n. Let i be such that $n \leq\ k+i$. $\sigma \in T^\infty$ so s^∞ gives a one-point extension τ of s, $\tau \in T^\infty$. $\mathrm{lh}(\tau) \leq\ k+i$ so $\tau \in T_i$. To construct the maps χ_i just contemplate diagram **??** (missing) until you realise that the columns are compatible with the maps Φ_i, and that row i is an increasing chain of fragments of strategies whose union is a strategy for T_i. The only thing which may not be immediately clear is that each projection has 4) in the definition of a covering. So let α_i be some run of play in $[T_i]$ following $\chi_i(s^\infty)$. By constructing successive liftings using 4) for the Φ_i, we get for $j \geq\ i$ runs α_j in $[T_j]$ following $\chi_j(s^\infty)$. It is now immediate that if we define α^∞ as the union of $\{\sigma : \sigma$ is an initial segment of α_j for all large $j\}$ then α^∞ is the sought-after lift of α_i. For α_j is played according to s^∞ for the first $k+j$ moves, so that in the limit α^∞ follows s^∞ for ever.

To get Borel determinacy, we just do an induction using the hypothesis "every set in the class $B_{\Sigma\nu}$ is nice", where we call a set **nice** if for every k it is unravelled by a k-covering. We have this already for $\nu = 1$ by Lemma 85. So let ν be less than ω_1 and suppose that the hypothesis holds for γ less than ν. Let A be a Borel subset of some $[T]$ of class $B_{\Sigma\nu}$ and let A be written as

$\bigcup\{A_i : i \in \omega\}$, where each A_i is a Borel subset of $[T]$ of class $B_{\Pi\gamma}$ for γ less than ν. Notice that niceness is preserved by taking complements, so by hypothesis all the A_i are nice. Fix k, and construct a tower of coverings as in Lemma 86, such that the first covering unravels A_0, the second unravels $\pi_1^{-1}(A_1)$ and so on. Take the inverse limit of this system, and notice that $\rho_0^{-1}(A)$ is the union of $\{\rho_0^{-1}(A_i) : i \in \omega\}$ and is open, because ρ_0 factors through an unravelling of each A_i, and so unravels each of them. So we have a k-covering of T which reduces A to an open set, and we just invoke Lemma I again to find a k-covering of this which reduces this open set to a clopen one. This establishes Borel determinacy.

I think this should be ρ_0^{-1} "$A\ldots$

AD and AC incompatible

THEOREM 81. $\neg(AC \wedge AD)$

Proof:

Consider the equivalence relation on $\mathbb{N}^\omega$ defined by:

$$f \sim g \text{ if } (\exists n_0 m_0 \in \mathbb{N})(\forall x \in \mathbb{N})(f(n_0 + x) = g(m_0 + x)).$$

If $f \sim g$ then there are minimal naturals we can take as witnesses to the '$\exists n_0, m_0 \in \mathbb{N}$' and the difference between the two things in this minimal pair will be either even or odd. Thus if $f \sim g$ they are either "an odd distance apart" or "an even distance apart" *but not both*.

Notice that there are $2^{\aleph_0}$ equivalence classes. (Consider a perfect binary tree with a bijection between its nodes and $\mathbb{N}$. Any two paths thru' this tree—and there are $2^{\aleph_0}$ of them—will be nonequivalent.) Now suppose the set of equivalence classes has a set of representatives. That enables us to pronounce each $f \in \mathbb{N}^\omega$ to be either "even" (if it is an even distance from the chosen representative of $[f]_\sim$) or "odd" (if it is an odd distance from the chosen representative of $[f]_\sim$).

Now let $A \subseteq \mathbb{N}^\mathbb{N}$ be the set of odd sequences, and consider G_A. We will show that G_A has no winning strategy for either player.

Suppose `II` has a winning strategy τ. (Here τ is the kind of strategy that looks at the sequence of all moves by either player so far). The game starts with player `I` making the 0th move. Consider a play where `II` uses τ (so we know she is going to win) and player `I` uses τ as well. He cannot do this straightforwardly, since τ eats sequences of *odd* length and `I` is confronted with sequences of *even* length. To use τ he turns the even sequence he is confronted with into an odd one by hanging a 0 on the front and *then* using τ. In particular his first move will be to play whatever τ ordains as the response to the sequence $\langle 0 \rangle$.

The result of this play must be a win for `II`, since τ is winning for her. Let us call it X. Now consider what happens if `I` starts, not by playing $\tau\langle 0 \rangle$, but 0 itself. `II` is still going to use τ, so `II`'s response will then be $\tau\langle 0 \rangle$ (which was `I`'s first move in the first play we considered), and `I`'s response to that will be whatever `II`'s second move in the original play would have been. Clearly this infinite sequence generated by this play will be just like the infinite sequence X generated by the original play but with a 0 hung on the front. But if X was odd this sequence must be even and *vice versa* so they can't possibly *both* be wins for `II` as they should. Therefore τ was not Winning.

Now suppose player `I` has a winning strategy σ. Consider the play where he uses σ and player `II` uses σ too, rather the way in which `I` used τ in the previous thought-experiment. σ only works on sequences of *even* length whereas `II` wants answers to problems of *odd* length, and she gets round this by inserting $x_0 + 1$ after x_0 in each sequence she is confronted with. (x_0 is `I`'s first move, ordained by σ). This play—call it Y—is a win for `I`. Now consider the play where `I` uses σ as before (and therefore starts with x_0) and `II` instead of playing $\sigma\langle x_0, x_0 + 1\rangle$ plays $x_0 + 1$. `I`'s next move will be $\sigma\langle x_0, x_0 + 1\rangle$. Since `I` is continuing to use σ the remaining moves will be the same as in the first play with σ we considered, except that they are now being made by the "other" player. Clearly this play cannot be a win for `I` if the last one was, for each is odd iff the other is even.

■

There is also the following—perhaps more natural—proof.

Proof:

Let $\mathcal{B}$ be Baire space, the set of all infinite sequences of natural numbers.

By AC, $\mathcal{B}$ is wellorderable, so enumerate it as as $\langle b_\nu : \nu < \kappa\rangle$ for some initial ordinal κ.

Strategies are functions defined on finite sequences of natural numbers returning a natural number. There are κ strategies for `I`, so list them as $\langle \sigma_\nu : \nu < \kappa\rangle$ and κ strategies for `II`, and list them as $\langle \tau_\nu : \nu < \kappa\rangle$.

If `I` uses strategy σ and `II` plays the sequence b, write $oc(\sigma, [b])$ for the outcome of the game.

Similarly write $oc([a], \tau)$ for the outcome when `I` writes a and `II` uses strategy τ.

(One can regard $[a]$ as a notation for the strategy for the first player of writing a, regardless of the conduct of the second player).

Note that for given σ, the map $b \mapsto oc(\sigma, [b])$ is 1-1, since b is every other point of the outcome play. Similarly for given τ.

We are going to build a subset X of $\mathcal{B}$ in κ stages.

At any stage before κ we have put fewer than κ points of B into X and fewer than κ points of $\mathcal{B}$ into Y, the complement of X. The rest of the points are called "undecided" at the given stage.

At stage ν, consider the strategy σ_ν, and find a point p such that $oc(\sigma_\nu, [p])$ is undecided. Place that latter point in Y. [That ensures that σ_ν cannot be a winning strategy for the first player in $G(X)$].

Then consider the strategy τ_ν and find a point q such that $oc([q], \tau_\nu)$ is still undecided (i.e. undecided and distinct from $oc(\sigma_\nu, [p])$) and place $oc([q], \tau_\nu)$ in X. [That ensures that τ_ν cannot be a winning strategy for the second player in $G(X)$].

■

DEFINITION 82. *AD_X is the assertion that games of length ω over the arena X are determined. AD tout court is $AD_{\mathbb{N}}$*

Remark 83. *If all Games (even of length 2!) over $\mathcal{P}(x)$ are determined then $\mathcal{P}(x)$ has a choice function.*

Proof: $\mathtt{I}$ picks $x' \subseteq x$ and $\mathtt{II}$ replies with a singleton $x_{\mathtt{II}}$. She wins if $x_{\mathtt{II}} \in x'$ or x' is empty. Clearly $\mathtt{I}$ cannot Win, and any winning strategy for $\mathtt{II}$ is a choice function on $\mathcal{P}(x)$.

■

More specifically (and more significantly)

Remark 84. *Suppose $f : \mathbb{R} \twoheadrightarrow X$. Assume AD_X. Then f has a right inverse.*

Proof: Let G_f be the game in which $\mathtt{I}$ picks $x_0 \in X$ for his first move, and thereafter plays anything he/we like: it won't matter; and $\mathtt{II}$ elaborates a real r. $\mathtt{II}$ wins if $f(r) = x_0$. $\mathtt{I}$ cannot Win, and a winning strategy for $\mathtt{II}$ is a right-inverse to f. ■

Remark 85. *AD implies that $\aleph_1$ is measurable.*

To prove this we need a lemma, which is of independent interest.

Let D be the set of degrees of unsolvability and let $E \subseteq D$. Then there is a degree d_E such that the "Turing cone" $\{d' \in D : d' \geq d_E\}$ is either included in E or is disjoint from it.

Proof of lemma: $\bigcup E$ is the set of those sequences in $\mathbb{N}^\omega$ whose degrees belong to E, and consider $G_{\bigcup E}$. Suppose $\mathtt{I}$ has a winning strategy σ and let $d(\sigma)$ be its degree. Let $d \geq d(\sigma)$ and let a be a sequence of degree d. If $\mathtt{I}$ plays according to his winning strategy σ and $\mathtt{II}$ plays α then $oc(\sigma, [a])$ will have degree d. Hence $d \in E$. So the Turing cone $\{d' \in D : d' \geq d(\sigma)\}$ is included in E.

On the other hand if $\mathtt{II}$ has a winning strategy τ in this game, let $d(\tau)$ be its degree. Let $d \geq d(\tau)$ and let b be a sequence of degree d. If $\mathtt{II}$ plays according to her winning strategy τ and $\mathtt{I}$ plays b then $oc([b], \tau)$ will have degree d. Hence $d \notin E$. So the "Turing cone" $\{d' \in D : d' \geq d(\tau)\}$ is disjoint from E. ■

This gives us a countable additive two-valued measure on the set of degrees. A set of degrees has measure 1 if it extends a Turing cone and 0 if it is disjoint from one.

Countable ordinals have degrees in the sense that every countable ordinal is the length of a wellordering of $\mathbb{N}$, and the graphs of these wellorderings are sets of ordered pairs of naturals and therefore sets of naturals. For d a degree let $f(d) =$: the least ordinal α such that there is no wellordering of $\mathbb{N}$ of length α of degree $\leq d$. Then for A a set of countable ordinals the measure of A is just the measure of f^{-1}"A. The measure is countably additive because every increasing ω-sequence in D has an upper bound. (So a countable intersection of Turing cones extends a Turing cone).

■

This tells us that AD is strong. The reason for this is that we can now run our construction of Gödel's L with the countably complete ultrafilter $\mathcal{U}$ on ω_1 as a parameter. We add intersection-with-$\mathcal{U}$ to the list of basic operations. The result is a structure $L[\mathcal{U}]$ which is a model of AC for the same reasons that L is, and the ordinal that was ω_1 will still be measurable. Since $L[\mathcal{U}]$ is a model of AC we can run all the usual arguments to show that there are lots of inaccessibles.

Remark 86. *$AD_{\mathbb{N}}$ →Choice for countable sets of sets of reals.*

Proof: .

Let $\{X_i : i \in \mathbb{N}\}$ be a countable family of nonempty sets of reals. The game whose determinacy we will exploit is defined as follows. `I` picks i for his first move and his subsequent moves do not matter. `II`'s moves define a real, and she wins iff it is in X_i. Clearly `I` cannot have a winning strategy, and any winning strategy for `II` is a choice function for $\{X_i : i \in \mathbb{N}\}$. ■

This is satisfactory. Countable choice for sets of reals is about the smallest amount of choice needed to make analysis tolerable. Notice that this does not suffice to carry out Vitali's "construction" of a nonmeasurable set of reals.

Exercise 36. *In attempting to prove this result from memory i considered the wrong game:* `I` *chops subtrees off the tree $2^{<\omega}$ to leave a closed subset.* `II` *wins if she codes a point in it. What does this game prove?*

Now during the "odd-even" trick of theorem 87 we saw a set (the set of equivalence classes under "x and y have the same tail") that is the image of a map from $\mathbb{R}$ and—assuming AD—that has no right inverse. This shows that AD_X must fail for some $|X| \leq_* 2^{\aleph_0}$.

4.7 Ramsey's theorem and the Erdős-Rado theorem

In the remainder of this chapter I am going to deal with large cardinals, and I am going to start by revisiting Ramsey's theorem. This is because a lot of the motivation for study of large cardinals (which is what passed for Set Theory for many years and in some quarters still does) arose from infinitary combinatorial questions. There are other reasons for logicians to be interested in it. One is that there is a finitary version which is used to prove that the decision problem for universal sentences is solvable, and this involves a use of elementary embeddings.

Notation: $[X]^n$ is the set of unordered n-tuples from X. $\alpha \to (\beta)^\gamma_\delta$. Take a set A of size α, partition the unordered γ-tuples of it into δ bits. Then there is a subset $B \subseteq A$ of size β such that all the unordered γ-tuples from it are in the same piece of the partition.

The proof of Ramsey's theorem I shall give will be oriented towards transfinite generalisations, so it won't be the cutest. First we prove $\omega \to (\omega)^2_2$.

We are given a two-colouring (red and blue) of all the edges in the complete undirected graph on $\aleph_0$ vertices. We are going to form an infinite finite-branching tree whose nodes are labelled with natural numbers. Below[6] 0, to the left (and to the right, respectively), we place the first natural number z such that there are infinitely many numbers greater than z to which z is connected by a blue edge (red edge respectively) and—strictly temporarily, we associate with it that set of greater numbers. We now build the tree recursively. Below each growing point—bud (the nice thing about trees is the plethora of nice imagery that comes free)—which is a number with a set of greater numbers temporarily associated with it, we place, to the left (and to the right) the smallest member of the set-temporarily-associated-to-the-bud such that there are infinitely many larger members of that set to which it is connected by a blue (resp. red) edge.

As we deal with each node we throw away the set that has been temporarily associated with it. When we have finished we have a tree in which every node has either one or two children. It cannot have no children at all since whenever you split an infinite set into two bits, one of the two is infinite. This is a finite-branching infinite tree, and so by König's infinity lemma must have an infinite branch. This infinite branch either has infinitely many left turns in it, or infinitely many right turns. ■

Notice (this will matter in the generalisations) that not every vertex eventually ends up at a node. In the transfinite generalisations every vertex will have its own node.

Now the version of König's infinity lemma that we need for this actually comes free, without AC. The usual proof is as follows: if a node has infinitely many descendents but only finitely many children, then at least one of those children has infinitely many descendents: choose one. In this case we can do this without any use of AC, since there is a canonical ordering of all children of all nodes, so we simply pick the first in the sense of this ordering.

The really beautiful thing about this theorem is the spectacular variety of ways in which it asks to be generalised

1. Can we do this for more than two colours?
2. Can we do this for unordered n-tuples, $n > 2$?
3. Can we construct monochromatic sets that are not just infinite but are infinite in particular ways (remember we started with a countable set)?
4. If we start with partitions of unordered tuples of uncountable sets can we get uncountable monochromatic sets?
5. Can we do this for partitions of infinite subsets (as opposed to unordered n-tuples?)
6. How constructive is this proof? Do recursive partitions have recursive monochromatic sets?

[6] For some reason these trees always grow downwards.

7. Are there finitary versions?

The next few subsections revolve around some of the ways just enumerated of potentially generalising this result.

4.7.1 Finitary versions

The answer to 7 is "yes": there are finitary versions. Indeed it was a finitary version that Ramsey needed to prove his theorem about decidability of universal formulæ—he proved the infinitary version only because it was easier. (Nowadays there are much easier proofs of both finitary and infinitary versions) Actually the finitary version —$(\forall mnk)(\exists j)(j \to (m)^n_k)$—follows from the infinitary one by compactness. (Ramsey didn't know this: the first appearance of compactness for predicate logic appeared the following year (1930) in a paper of Gödel[7].) If we spice up the proof of this a bit we can prove something that turns out to be much stronger. We need the notion of a relatively large subset of $\mathbb{N}$. $x \subseteq \mathbb{N}$ is **relatively large** if $|x| > inf(x)$.

The finite version we want is

Theorem 87. $(\forall mnk)(\exists j)(j \to (m)^n_k)$.

This is not the version Ramsey needed for his proof of his decidability result, but is of more interest to us here, since it is this version (not the version Ramsey needed for the decidability result) that can be spiced up to give Paris-Harrington.

It's not hard to see how one can prove $(\forall mnk)(\exists j)(j \to (m)^n_k)$ directly by careful applications of Rado's method; this method will prove $2^{2n} \to (n)^2_2$ for example—though this is far from best possible: e.g., we know $6 \to (3)^2_2$.

However it is also possible to deduce this finite version of Ramsey's theorem from the infinite version by a reverse compactness argument. There are several reasons why Ramsey didn't do it that way. For one thing, as we have observed, the first appearance of compactness for predicate logic did not appear until the following year. For another, the reverse compactness proof is highly ineffective in that no bounds can be recovered from it; nobody in their right mind would try to do it that way unless they had an ulterior motive. However we do have such an ulterior motive: since the proof of Paris-Harrington (at least the only proof known to me) proceeds by a reverse compactness argument it is very useful to run through the reverse-compactness proof of finite Ramsey by way of a rehearsal for it.[8]

So here[9] is the reverse compactness proof of theorem 93, $(\forall mnk)(\exists p)(p \to (m)^n_k)$.

Proof: Suppose that claim is false, and that there are n, m, k in $\mathbb{N}$ such that for all $p \in \mathbb{N}$ there is a set P with $|P| = p$ and a colouring $f : [P]^m \to \{1, 2, \ldots k\}$

[7] Theorem X: thanks to Torkel Franzen for the citation.

[8] There are other reverse compactness arguments to be found in the literature: for example Friedman's proof of FFF.

[9] (Thanks to Dave Turner for finding (some of the) mistakes in my transcription from Simpson's article!)

such that there is no set $X \subseteq P$ with $|X| = n$ and $|f``[X]^m| = 1$. Fix n, m, k, and for each p let Y_p be the set

$$\{f : f : [[1,p]]^m \to [1,k] \wedge \neg(\exists X) \bigwedge \left\{ \begin{array}{l} X \subseteq [1,p] \\ |X| \geq n \\ |f``[X]^m| = 1 \end{array} \right\}\}.$$

of bad k-colourings of the m-tuples of the naturals below p. (k and m are fixed.)

(Beware: square brackets are here being used *both* to denote intervals in $\mathbb{N}$—as in $[1,k]$—*and* to denote the set of m-sized subsets of things—as in $[X]^m$.)

For any k, the set F_k of all k-colourings of m-tuples of initial segments of $\mathbb{N}$ is countable. (Each initial segment $[1,p]$ has only a finite set of m-membered subsets and there are only finitely many ways of colouring the set of those subsets). So we can uniformly wellorder F_k. Suppose this to be done, somehow. Then, for each p, we set f_p to be the first element of Y_p in the sense of that ordering.

We are now going to define a (bad) partition π of $[\mathbb{N}]^{m+1}$ into k pieces. You are given a set $x \subseteq \mathbb{N}$ of size $m+1$ and have to decide which piece to put it into. Its last member is $p+1$ for some natural number p. $x \setminus \{p+1\}$ is now a subset of $[1,p]$ and is therefore a suitable input for f_p. $f_p(x \setminus \{p+1\})$ is now a number $< k$, and that tells you which piece to put x into. (Slightly more formally, put x into the $f_{(sup(x)-1)}(x \setminus \{sup(x)\})$th piece.) So π partitions $[\mathbb{N}]^{m+1}$ into k pieces. We will show that π is bad.

With a view to obtaining a contradiction suppose X to be an infinite set monochromatic for π. Let $p+1$ be a member of X (and we will want to be able to find arbitrarily large such $p+1$). Consider those $(m+1)$-tuples from $X \cap [1,p+1]$ whose last element is $p+1$. What does π do to them? It sends every such $(m+1)$-tuple x to $f_p(\texttt{butlast}(x))$, and—because X is monochromatic—all these $f_p(\texttt{butlast}(x))$ are the same, whatever x we pick up. Now *every* m-sized subset of $X \cap [1,p]$ can be turned into such an $(m+1)$-tuple by the simple expedient of sticking $p+1$ on the end, so f_p sends every m-tuple from $X \cap [1,p]$ to the same number $< k$. But that is simply to say that $X \cap [1,p]$ is a subset of $[1,p]$ that is monochromatic for f_p. Now f_p was chosen so that any set monochromatic for it was of size less than n. So $X \cap [1,p]$ is of size less than n. So—no matter how large we pick $(p+1) \in X$—we find that $X \cap [1,p]$ has at most n members. So $|X| \leq n+1$ and X was not infinite, contradicting the Infinite Ramsey theorem. ■

We need to make a note here of the way in which this proof is less effective than the proof of Rado's given in the previous section. It is true that Rado's proof uses excluded middle—and is therefore beyond the pale for the extremely-squeamish—but is is effective in the weak sense that, by close examination of it, we can quite straightforwardly recover bounds for witnesses to the existential quantifier. In contrast the proof we have just given does not divulge bounds in this way. The reader will not be surprised to be told that the proof we are about to give of Paris-Harrington will be similarly tight-lipped.

4.7.2 The Paris-Harrington Theorem

THEOREM 88. *(Paris-Harrington) For every n, m, k in $\mathbb{N}$, there is p so large that whenever $f : [\{1, 2, \ldots p\}]^m \to \{1, 2, \ldots k\}$ there is a relatively large $X \subseteq \{1, 2, \ldots p\}$ such that $|X| \geq n$ and $|f``[X]^m| = 1$.*

Proof:

We argue by reverse compactness as before.

Suppose there are n, m, k in $\mathbb{N}$ such that for all $p \in \mathbb{N}$ there is $f : [\{1, 2, \ldots p\}]^m \to \{1, 2, \ldots k\}$ such that there is no relatively large $X \subseteq \{1, 2, \ldots p\}$ such that $|X| = n$ and $|f``[X]^m| = 1$. Fix n, m, k and p and let Y be the set

$$\{f : f : [\{1, 2, \ldots p\}]^m \to \{1, 2, \ldots k\} \wedge \neg(\exists X) \bigwedge \left\{ \begin{array}{l} X \subseteq \{1, 2, \ldots p\} \\ |X| > min(X) \\ |X| \geq n \\ |f``[X]^m| = 1 \end{array} \right\} \}.$$

This time let Y_p be—not the set of

> colourings-that-are-bad-in-the-sense-of-lacking-a-monochromatic-set-of-size-n

but the set of

> colourings-that-are-bad-in-the-sense-of-not-having-any-monochromatic-sets-of-size-n-that-are-relatively-large.

As before, initial segments of the monochromatic set X will be monochromatic for the colourings f_p. Now sets that are monochromatic for f_p are either smaller than n or are not relatively large. By considering initial segments of X that are long enough we can take care of the first condition, so the only way they can manage to be monochromatic for f_p will be by failing to be relatively large. So, for some large j, consider the initial segment consisting of the first j elements of X. We now know that this is not relatively large, so its first element must be bigger than j. So the first element of X is at least j. But j could have been taken to be arbitrarily large.

■

The Quantifier Prefix of Paris-Harrington

Paris-Harrington is dramatically stronger than finite Ramsey (see [**?**] for example) and one might well wonder whether or not there are any syntactic clues to the source of this extra strength. The feature that chiefly caught my interest in this connection is the unstratified/ill-typed nature of the property of *relative largeness*, and we will get onto that in due course.

One obvious difference between Finite Ramsey and P-H is that P-H doesn't talk about colourings of tuples from arbitrary finite *sets* but of colourings of tuples quite specifically from *initial segments of the naturals*.

However we will first get out of the way a simple observation about the quantifier prefix. First we rephrase Finite Ramsey as an assertion about colourings of tuples from a finite set:

For all n, m, j in $\mathbb{N}$
There is k in $\mathbb{N}$ so large that
For every set X of size k and
For every m-colouring χ of $[X]^j$
there is $X' \subseteq X$ with $|X'| = n$
and X' monochromatic wrt χ.

Next we rephrase Finite Ramsey as an assertion about colourings of tuples of naturals. An *enumeration* of a set X is a bijection between X and an initial segment of $\mathbb{N}$. This enables us to extend the notion of relative largeness from sets of natural numbers to subsets of arbitrary sets Y once Y has been equipped with an enumeration: $Y' \subseteq Y$ is relatively large with respect to an enumeration e of Y iff $e``Y'$ is relatively large *tout court*. Given $X$ (as in the statement of the theorem) it is clear that once we have found $X' \subseteq X$ (as in the statement of the theorem) we can pick an enumeration $e$ of $X$ so that $e``X' = [0, n]$. X' is now relatively large wrt e. So here is Finite Ramsey phrased as an assertion about relatively large monochromatic sets.

For all n, m, j in $\mathbb{N}$
There is k in $\mathbb{N}$ so large that
For every set X of size k and
For every m-colouring χ of $[X]^j$
there is an enumeration e of X and
there is $X' \subseteq X$ with $|X'| = n$
with X' monochromatic wrt χ and relatively large wrt e.

Now that we have expressed Finite Ramsey in a syntax that is the same as that used to express P-H we are in a better position to compare them. Here is P-H.

For all n, m, j in $\mathbb{N}$
There is k in $\mathbb{N}$ so large that
For every set X of size k and
For every m-colouring χ of $[X]^j$ and
For every enumeration e of X
there is $X' \subseteq X$
with X' monochromatic wrt χ and relatively large wrt e.

We have set out Paris-Harrington and Finite Ramsey above in something very like Prenex Normal form. The two of them have the same matrix (the stuff after the prefix) and the prefixes

$$(\forall mnj)(\exists k)(\forall X)(\forall \chi)(\exists e)(\exists X') \quad \text{and} \quad (\forall mnj)(\exists k)(\forall X)(\forall \chi)(\forall e)(\exists X')$$

are almost the same except that the innermost part of the quantifier prefix of Finite Ramsey is $(\exists e)(\exists X')$ whereas the innermost part of the quantifier prefix of Paris-Harrington is $(\forall e)(\exists X')$. That is to say we have replaced an existential quantifier with a universal quantifier: clearly we must expect P-H to be stronger than Finite Ramsey.

Paris-Harrington is an assertion solely about natural numbers, not even sets of natural numbers, since assertions about finite sets of natural numbers can be coded as assertions about numbers. However the proof we have given involves reasoning about infinite sets of natural numbers. Net result: we have an assertion of elementary arithmetic who does not—on the face of it—have a proof from arithmetical axioms. **This one will run and run**. (Compare FFF)

4.7.3 Recursive partitions without recursive monochromatic sets

It is not hard to persuade oneself that the proof that there is an infinite monochromatic set for a recursive partition will not ensure that the infinite monochromatic set is going to be recursive. One way to see this is to reflect on how one would decide what the first element of the infinite monochromatic set would be. It is perfectly true that—since $\mathbb{N}$ is wellordered—one knows which way to jump when a balloon splits into two infinite balloons, so the instance of König's lemma that one has to have in order to obtain an infinite path through the tree comes free. However one does not know what the first element of the infinite monochromatic set is until one has found the infinite path and determined whether it has infinitely many blue points or infinitely many pink points. Of course this rumination doesn't *prove* that a recursive partition doesn't have an infinite recursive monochromatic sets, but it should nprepare us for this discover.

Remark 89. *(Jockusch) There are recursive partitions without recursive monochromatic sets.*

Proof:

Let's have some definitions. We know that $\{p\}(d)\downarrow$ says that program with gnumber p applied to data d halts; $\{p\}_z(d)\downarrow$ will mean that the program with gnumber p applied to data d halts in $\leq z$ steps.

Consider ρ: $[\mathbb{N}]^3 \to \{0,1\}$ by $\rho\{x,y,z\} = 0$ iff $x < y < z \to (\forall p, d < x)(\{p\}_y(d)\downarrow \longleftrightarrow \{p\}_z(d)\downarrow)$

Suppose X is an infinite subset of $\mathbb{N}$ monochromatic for ρ. We must have $\rho``[X]^3 = \{0\}$ since $\rho``[X]^3 = \{1\}$ is obviously impossible ("too few truth values"). We will show that if X is recursive then we can solve the halting problem.

To determine whether or not $\{p\}(d)\downarrow$ first find a member n of X larger than p and d. Then for any y, z in X bigger than n we have $\{p\}_y(d)\downarrow$ iff $\{p\}_z(d)\downarrow$. Since X is infinite it has arbitrarily large members and so if $\{p\}(d)$ ever halts at all there is $z \in X$ large enuff to ensure that $\{p\}_z(d)\downarrow$. But then, by monochromaticity of X it will be sufficient to check $\{p\}_z(d)\downarrow$ for even *one* $z \in X$ bigger than n. ■

On the other hand there is a theorem of Seetapun's that every recursive partition of $[\mathbb{N}]^2$ has a monochromatic set in which the halting problem is not recursive. (See Hummell: Journal of Symbolic Logic December 1994)

4.7.4 Ramsey Ultrafilters and $\beta\mathbb{N}$

not to be lectured, and accordingly supplied in very small print

For the answer to the question "Can we have monochromatic sets that are large in particular ways?" we examine the proof we have been through. What have we used? At the building-the-tree stage we need to know that every time we split an infinite set into two bits one of the two bits is infinite. The other feature we needed was that if you have an infinite sequence of infinite sets $\langle X_i : i \in \mathbb{N}\rangle$ where X_n contains no numbers smaller than n, then there is an infinite set that contains at least one member of each. It turns out that the condition we need is something like this:

> Say A is a **set-of-big-sets** (subsets of $\mathbb{N}$) if (1) whenever you split a thing in A into two, one fragment is in A, and (2) Whenever $\langle X_i : i \in \mathbb{N}\rangle$ is a $\subseteq$-descending sequence of things in A which generate a nonprincipal filter on $\mathbb{N}$, then there is an X in A such that, for all $i \in \mathbb{N}$, $X \setminus X_i$ is finite—and indeed is included in the first i members of X.
>
> Then we can rejig the proof of Ramsey's theorem to show that any two-colouring of the complete graph on $\mathbb{N}$ has a monochromatic set in A. We shall see conditions like this cropping up later.

Intuitively this is a rather odd situation: one "large" comes from the countable set we started with, the other comes from the countable set that indexes the construction. (It's probably unstratified.)

This rather odd condition (2) gives rise to a new definition:

DEFINITION 90.

1. *X **diagonalises** the family $\langle X_s : s \in \mathbb{N}^{<\omega}\rangle$ iff $X \subseteq X_\emptyset$ and for $s \in \mathbb{N}^{<\omega}$ if $sup(s) \in X$ then X minus all integers below $|s|$ is a subset of X_s.*
2. *$A \subseteq \mathcal{P}(\mathbb{N})$ is a **happy family** iff $\mathcal{P}(\mathbb{N}) \setminus A$ is a nonprincipal ideal over $\mathbb{N}$, and whenever $\langle X_s : s \in \mathbb{N}^{<\omega}\rangle$ generates a filter $\subseteq A$ there is $X \in A$ which diagonalises the $\vec{X}_s$*

DEFINITION 91. *A **filter on** X is a subset of $\mathcal{P}(X)$ closed under (finite) intersections and superset; it is **proper** if it does not contain the empty set; an **ultrafilter** is a $\subseteq$-maximal filter.*

DEFINITION 92. *A **Ramsey** ultrafilter $\mathcal{U}$ on $\mathbb{N}$ is an ultrafilter such that every partition of $[\mathbb{N}]^2$ has a monochromatic set in $\mathcal{U}$.*

we had better show that this is the same as the above with "member of $\mathcal{U}$" substituted for "$\mathbb{N}$".

In a topological space a point is a p-point iff every countable intersection of neighbourhoods of it is another neighbourhood of it. (Some of you may have done Banach algebras and encountered this notion before.)[10]

[10] Andreas Blass writes on FOM:

The terminology comes from general topology, where a point x in a space X is called a P-point iff every intersection of countably many neighbourhoods of x is a neighbourhood (not

Given the topology on ultrafilters as the Stone-Čech compactification of $\mathbb{N}$ we topologise the family of ultrafilters on $\mathbb{N}$—which is written $\beta\mathbb{N}$—by (for each $X \subseteq \mathbb{N}$) taking an open set to be the set of ultrafilters containing X. There is an analogous definition for topologising the nonprincipal ultrafilters on $\mathbb{N}$, which is simply the family of ultrafilters in the quotient algebra $\mathcal{P}(\mathbb{N})/Fin$ (Fin is the ideal of finite sets). This means that a p-point ultrafilter over $\mathbb{N}$ must satisfy $\forall$ seq $\langle A_i : i \in \mathbb{N}\rangle \in (\mathcal{P}(\mathbb{N}) \setminus \mathcal{U})^\omega$ $\exists u \in \mathcal{U}\ \forall i\ A_i \cap u$ **is finite**. Equivalently, whenever X is a countable subset of $\mathcal{U}$ there is $y \in \mathcal{U}$ such that $y \subseteq' x$ for all $x \in X$ where $y \subseteq' x$ means that $y \setminus x$ is finite. Again, this is simply to say that—considered as an ultrafilter on $\mathcal{P}(\mathbb{N})/Fin$—the ultrafilter is countably complete.

DEFINITION 93. *If $\mathcal{U}$ and $\mathcal{V}$ are ultrafilters over a set X, then $\mathcal{V} \leq_{RK} \mathcal{U}$ iff $\exists f : X \twoheadrightarrow X$ and $\mathcal{V} = \{f``Y : Y \in \mathcal{U}\}$*

(Check that if $f : \mathbb{N} \twoheadrightarrow \mathbb{N}$ then $\{f``X : X \in \mathcal{U}\}$ is an ultrafilter if $\mathcal{U}$ is.)

It's easy to see that $\leq_{RK}$ is transitive and reflexive (so it's a **preorder** or **quasiorder**) but there is no reason to expect it to be antisymmetrical. Writing '$\mathcal{U} \sim \mathcal{V}$' for $\mathcal{U} \leq_{RK} \mathcal{V} \wedge \mathcal{V} \leq_{RK} \mathcal{U}$, it turns out we have

PROPOSITION 94. *If $\mathcal{U}$ and $\mathcal{V}$ are ultrafilters over a set X then $\mathcal{U} \sim \mathcal{V}$ iff there is a permutation π of X such that $j^2(\pi)(\mathcal{U}) = \mathcal{V}$.*

Doesn't X have to be $\mathbb{N}$ for this to be true?

Proof:

We will need the following fact:

REMARK 95. *If $(j^2 f)(\mathcal{U}) = \mathcal{U}$ (or $f_*(\mathcal{U}) = \mathcal{U}$) if you prefer) then $\{n : f(n) = n\} \in \mathcal{U}$.*

Proof:

Let $F_= = \{n : f(n) = n\}$, $F_> = \{n : f(n) > n\}$ and $F_< = \{n : f(n) < n\}$. We will show $F_> \cup F_< \notin \mathcal{U}$, so $F_= \in \mathcal{U}$ by ultra-ness of $\mathcal{U}$.

Given $T \in \mathcal{U}$ let T_n be $\{m : n$ is the least integer such that $f^n(m) \notin T\}$. Thus $T = \bigcup_{i\in\mathbb{N}} T_i$. Suppose $F_< \in \mathcal{U}$. Since each of the disjoint sets $\bigcup\{(F_<)_{2n} : n \in \mathbb{N}\}$ and $\bigcup\{(F_<)_{2n+1} : n \in \mathbb{N}\}$ must be in $\mathcal{U}$ iff the other is as well we have a contradiction. So $F_< \notin \mathcal{U}$.

As in the previous case we can see that neither $\bigcup\{(F_>)_{2n} : n \in \mathbb{N}\}$ nor $\bigcup\{(F_>)_{2n+1} : n \in \mathbb{N}\}$ can be in $\mathcal{U}$. The set $\mathbb{N} \setminus \bigcup\{(F_>)_n : n \in \mathbb{N}\}$ can be partitioned into two bits in the same way, such that when one is in $\mathcal{U}$ the other is in $j^2 f(\mathcal{U})$. Thus $F_> \notin \mathcal{U}$.

■

We now apply this to the situation where f maps $\mathcal{U}$ onto $\mathcal{V}$ and g maps $\mathcal{V}$ onto $\mathcal{U}$. Remark 101 tells us that gf is the identity on a set in $\mathcal{U}$, and that fg is the identity on a set in $\mathcal{V}$.

...

---

necessarily an open neighbourhood) of x. There remains the question of why the letter 'p' is used for this purpose in topology. Half of the answer is that 'P' is the first letter of "prime ideal". The other half is the following connection between prime ideals and p-points.

For any topological space X, the continuous real-valued functions on X form a commutative ring $C(X)$, with the operations of pointwise addition and multiplication. For any point $x \in X$, the continuous functions that vanish at x form a prime (in fact maximal) ideal P_x in $C(X)$. In a sufficiently nice space X (I believe complete regularity suffices, but I haven't checked this), x is a p-point iff P_x does not properly include any other prime ideal.

I'm trying to reconstruct this. Let's start with $F_>$. Call it T for short. Since $<$ is wellfounded, if $n \in T$ it follows that $f^k(n) \notin T$ for some k, lest $\langle n, f(n), f^2(n) \ldots \rangle$ be a descending chain. So T is the union of all the T_n where T_n is $\{m : n$ is the least integer such that $f^n(m) \notin T\}$. (The first thing in the sequence is T_1.) Let T_{even} be the union of T_n for n even and T_{odd} be the union of T_n for n odd. Now $f``T_{\text{even}} \subseteq T_{\text{odd}}$ (so if $T_{\text{even}}$ is big so is $T_{\text{odd}}$) and $f``T_{\text{odd}} \subseteq T_{\text{even}}$ (so if T_{odd} is big so is T_{even}) They are disjoint so they can't both be big, so neither of them can be big.

I don't see how to work the same trick for $F_<$ because we cannot be sure that if $x \in F_<$ then there is n such that $f^n(x) \notin F_<$. Wellfoundedness is no use here.

Let's try anyway. $F_<$ splits naturally into those things x s.t. $\forall n \in \mathbb{N}\ f^n(x) \in F_<$ and the rest. The rest we can deal with as before, but what about $\{x : (\forall n \in \mathbb{N})(f^n(x) \in F_<)\}$? Why should this not be big?

The $\leq_{RK}$ ordering has something to tell us about embeddings between ultrapowers.

REMARK 96. *Suppose $\mathcal{V} \leq_{RK} \mathcal{U}$, and that $f : X \twoheadrightarrow X$ such that $\mathcal{V} = \{f``Y : Y \in \mathcal{U}\}$. Let $\mathfrak{M}$ be any structure and consider the ultraproducts $\mathfrak{M}^X/\mathcal{U}$ and $\mathfrak{M}^X/\mathcal{V}$. Then the embedding $\mathfrak{M}^X/\mathcal{V} \hookrightarrow \mathfrak{M}^X/\mathcal{U}$ sending $[g]_\mathcal{V}$ to $[g \circ f]_\mathcal{U}$ is elementary.*

Proof: let g, $h \in \mathfrak{M}^X$. Then $g \sim_\mathcal{V} h$ iff

$$\{i : g(i) = h(i)\} \in \mathcal{V}$$

$$f^{-1}``\{i : g(i) = h(i)\} \in \mathcal{U}$$

$$\{f^{-1}(i) : g(i) = h(i)\} \in \mathcal{U}$$

$$\{i : g \circ f(i) = h \circ f(i)\} \in \mathcal{U}$$

$$g \circ f \sim_\mathcal{U} h \circ f$$

. . . and similarly for other predicate letters. ■

THEOREM 97. *The continuum hypothesis implies that every happy family extends an ultrafilter.*

Proof: Let A be a happy family. We are interested in sequences $\{X_s : s \in [\mathbb{N}]^{<\omega}\}$ which generate filters on $\mathbb{N}$. Enumerate them as $\langle \{X^\zeta{}_s : s \in [\mathbb{N}]^{<\omega}\} : \zeta < \omega_1 \rangle$. Construct a sequence $\langle F_\zeta : \zeta < \omega_1 \rangle$ of countably generated filters such that $F_\zeta \subseteq F_{\zeta+1} \subseteq A$. For each ζ either there is an s with $X^\zeta{}_s \in F_{\zeta+1}$ or $\{X^\zeta{}_s : s \in [\mathbb{N}]^{<\omega}\} \subseteq F^\zeta$ and there is a Y in $F_{\zeta+1}$ which diagonalises $\{X^\zeta{}_s : s \in [\mathbb{N}]^{<\omega}\}$ and such that $(\forall Z \subseteq \mathbb{N})(\exists \zeta)(Z \in F_\zeta \ \vee \ (\mathbb{N} \setminus Z) \in F_{\zeta+1})$. Then $\bigcup\{F_\zeta : \zeta < \omega_1\}$ is the desired ultrafilter.

The only difficult part of the construction is the following. Given $F_\zeta \subseteq A$ and $\{X^\zeta{}_s : s \in [\mathbb{N}]^{<\omega}\}$, if there is s such that the filter generated by F_ζ and $\mathbb{N} \setminus X^\zeta{}_s$ is included in A, then let $F_{\zeta+1}$ be that filter so generated by s. Otherwise $(\forall s)(\forall X \in F_\zeta)(X^\zeta{}_s \cap Y \in A)$, and so the filter generated by F_ζ and $\{X^\zeta{}_s : s \in [\mathbb{N}]^{<\omega}\}$ is included in A, and is countably generated by $\{Y_i : i \in \mathbb{N}\}$ (say). Let $Z_i = X^\zeta{}_i \cap \bigcap\{Y_s : |s| \leq i\}$. Then the filter generated by the Z_i is included in A. Since A is happy, there is some

W in A which diagonalises the Z_s, and set $F_{\zeta+1}$ to be the filter generated by F_ζ and W.

■

THEOREM 98. *the following are equivalent*

1. *("The minimality property") $\mathcal{U}$ is $\leq_{RK}$ minimal.*
2. *("The partition property") Every finite partition of $[\mathbb{N}]^n$ has a monochromatic set in $\mathcal{U}$.*
3. *("the selection property") $\forall$ partitions $\langle A_i : i \in \mathbb{N}\rangle$ of a $\mathcal{U}$-big set into $\mathcal{U}$-small sets $\exists u \in \mathcal{U}\ \forall i\ A_i \cap u$ is a singleton.*
4. *("The DC property"). Given $\langle A_i : i \in \mathbb{N}\rangle \in \mathcal{U}^\omega$ there is f increasing $\mathbb{N} \to \mathbb{N}$ with f"$\mathbb{N} \in \mathcal{U}$ and $(\forall i)(f(i+1) \in A_{f(i)})$*

[Notice that it is easy to show that anything $\leq_{RK}$ a Ramsey ultrafilter is also Ramsey (notice that *principal* ultrafilters are ramsey!).] Also: an ultrafilter is Ramsey iff it is a happy family.

Proof:

We will prove the following chain of implications:

Selection property → DC property → partition property → minimality property → selection property.

The selection property implies the DC property

Given $A_0 \supseteq A_1 \supseteq A_2 \supseteq \ldots \bigcap A_i = \Lambda$, find $X \in \mathcal{U}$ s.t. $|X \cap (A_i - A_{i+1})| \leq 1$. We want the transversal to be increasing.

Define g by $g(0) = min(X)$, thereafter choose $g(k+1)$ to be minimal so that $g(k+1) \in X$; $g(k+1) > g(k)$; $g(k+1) \geq \max(X \setminus A_{g(k)})$. Consider the sequence $A_{g(i)}$ .... (I think the point of this fiddle is to ensure that the sequence of bottom elements of members of this sequence is strictly increasing.) Let $T \in \mathcal{U}$ be a transversal. Enumerate $T \cap X$ by an increasing function h.

Claim: $h(j+2) \in A_{h(j)}$, so half of $T \cap X$ (the half in $\mathcal{U}$) is what we want.

Proof of claim: (also, what do we mean by "half of ..."?)

The DC property implies the partition property.

Given Δ a partition of $[\mathbb{N}]^2$ let $A_i = \{j > i : \Delta(i,j) = \text{whichever makes it big}\}$. Apply DC property to this. (This is just like the *last* stage in Rado's proof of Ramsey's theorem).

The partition property implies the minimality property.

We want to show that if $\mathcal{U}$ has the partition property then any ultrafilter $\mathcal{V} \leq_{RK} \mathcal{U}$ is either principal or $\geq_{RK} \mathcal{U}$. Fix $f : \mathbb{N} \twoheadrightarrow \mathbb{N}$. Suppose $X \in \mathcal{U}$. Define a partition of $[X]^2$ by $x \sim y$ iff $f(x) = f(y)$. There is a monochromatic set in $\mathcal{U}$. This will be either a preimage of a singleton (in which case $f_*\mathcal{U}$ is principal) or a large subset of X on which f is 1-1. We extend this 1-1 subset of f to a permutation of $\mathbb{N}$ mapping $\mathcal{U}$ onto $\mathcal{V}$.

The minimality property implies the selection property

Partition a big set into small sets. This partition is coded by a function $f(n)$ that $= i$ when $n \in B_i$. This f cannot be constant on a big set, so it is 1-1 on a big set, which gives us the transversal.

■

EXERCISE 37. *Principal ultrafilters are κ complete and every cardinal has one, so why can't we use one of them to show that every cardinal has the tree property?*

4.7.5 Increasing the exponent

The answer to 1 is "yes". Clearly the choice of a two-colouring was unneccessary: the same construction would have worked for any finite n. So we have proved $\omega \to (\omega)^2_n$. The answer to (2) is also "yes".

Now we must see how to extend this to $\omega \to (\omega)^m_n$, since it is this generalisation of the exponents and the subscripts that is usually known as Ramsey's theorem. This we do by induction on the exponent. Assume $\omega \to (\omega)^m_2$ and try to prove $\omega \to (\omega)^{m+1}_2$. (increasing the subscript is trivial)

We build a tree somewhat as before, this time starting off with the first m elements of $\mathbb{N}$. We then partition the naturals above m into two bits according to the colour an element x makes when made up into a m-tuple $\{1 \ldots m, x\}$. At this point the tree splits into two (c).

At the next stage in the construction of the tree we have (at each node) $m + 1$ elements above it, and a set associated with it. Each element of the set can be joined into a $m + 1$-tuple with any of the m-tuples we can form from the $m + 1$ things above it and there are of course m such tuples. Each of these tuples gives a colouring of x, so the elements of this set come in 2^m colours.

Every nth stage node has $m+n$ elements above it and a set associated with it; each member of this set—before it can be coloured—must be packed into a $m+1$-tuple with m things chosen from the $m+n$ elements above it; of course this can be done in $\binom{m+n}{m}$ ways. So each element of this set has $\binom{m+n}{m}$ colourings, and so the node will have $2^{\binom{m+n}{m}}$ children.

By this process we have assigned every element of $\mathbb{N}$ to a node of a finitely-branching tree so there must be an infinite path thru it by König's Infinity lemma. (We don't actually need König's Infinity lemma in this case beco's $\mathbb{N}$ is naturally wellordered.) Then we consider an infinite branch through this tree.

One thing we know about this branch is that for each increasing sequence $\{x_1 \ldots x_m\}$ through it, all $m+1$-tuples $\{x_1 \ldots x_m, y\}$, with $y > x_m$ are the same colour. We can think of this colour as the colour of the tuple $\{x_1 \ldots x_m\}$. This means we have a colouring of $[\omega]^m$. Now we invoke the induction hypothesis (that $\omega \to (\omega)^m_2$) to get a set monochromatic for this colouring. But this set will also be monochromatic for the colouring of the $m + 1$-tuples.

For higher exponents we proceed like this, using the induction hypothesis only at the tree stage.

Notice that for finite n this argument will prove $2^{2n} \to (n)^2_2$. This is not best possible: e.g., we know $6 \to (3)^2_2$. Ramsey actually used the finite version of this proof to construct an algorithm to tell when a universal sentence of the predicate calculus has arbitrarily large models. His paper is worth reading, for two reasons. (i) The decidability result is interesting. (ii) It is a sobering experience to realise how difficult things can be if you do not have the right notation (which Ramsey didn't). His proof of the combinatorial result is much less attractive than the proof here which is due (I believe) to R. Rado.

> When i first lectured this, a part 3 student called Dieter Olp showed me a nice way of dealing with the induction step:
>
> We have a colouring of the unordered n-tuples from X where X is countable. Pick up the first thing in X and call it x_0. Look at $X \setminus \{x_0\}$. We can colour the unordered $n-1$-tuples from $X \setminus \{x_0\}$ according to the colour of the n tuple they make when we add x_0. By induction hypothesis this partition of $n-1$-tuples has a monochromatic set. Let x_1 be the first thing in this monochromatic set and throw away everything that isn't in it. Repeat the process. The desired monochromatic set is the sequence $\{x_0, x_1, \ldots\}$. This is certainly a shorter proof than the one i gave, but from our point of view it has two fatal drawbacks. We have to *choose* a monochromatic set at each stage and there are countably many stages so we are using DC. The other is that we are using the induction hypothesis (that partitions of $n-1$-tuples have large monochromatic sets) *infinitely* often instead of just once. The difficulty this would make for us if we tried to use this proof in a context where the monochromatic set might be smaller than the set whose $n-1$-tuples we are partitioning is that there is no longer any guarantee that it will give us a large set monochromatic for the partition of n-tuples.

4.7.6 Transfinite Generalisations of Ramsey's theorem

The Erdős-Rado theorem

on p. 77 I carefully chose a proof of Ramsey's theorem that generalises sweetly to the proof of the version of the Erdős-Rado theorem that i will be giving below. Given this fact, and a statement of the Erdős-Rado theorem, you may wish to reconstruct the proof yourself. That is how I did it when first preparing these lectures, since I had by then long since forgotten any proof i may ever have had of it.

THEOREM 99. *(Erdős-Rado)* $(\beth_n(\kappa))^+ \to (\kappa^+)^{n+1}$.

That is the usual version, but it is not what i am going to prove. The usual version makes heavy use of AC, which is a shame since the *construction* doesn't really need AC at all. If one drops AC one gets a more instructive proof—albeit of something much weaker. One can recover the full AC-intoxicated version by snorting as much AC as is needed.

The proof I shall give exploits only the core construction behind the Erdős-Rado theorem, and makes no use of the cardinal equalities used to get the stronger, published result.

Proof is in erdosrado.pdf on my homepage, and the file has some helpful comments from the referee of the article

Now we want to start thinking a bit about the function $\lambda\alpha.$(least β such that $\beta \to (\alpha)^2$). The Erdős-Rado theorem does at least tell us that this is everywhere defined. Its values are called **Erdős numbers**. (Not to be confused with Erdős numbers!) Let's call it $\mathbf{E}$ for short.

Then we start wondering about fixed points for $\mathbf{E}$, which would be cardinals that satisfy a generalisation of Ramsey's theorem. The obvious thing to do is to iterate $\mathbf{E}$ but when you start with β and reach $sup\{\mathbf{E}^n(\beta) : n \in \mathbb{N}\}$ all that happens is that you get an α such that $(\forall \beta < \alpha)(\alpha \to (\beta)^2)$. Obvious question: for what ordinals α is the function $\mathbf{E}$ α-continuous? What can we infer about α if given that $\mathbf{E}$ is α-continuous?

This will give rise to an axiom of infinity. Many an axiom of infinity arises from an assertion that is true of $\aleph_0$ (or ω_0) and then asserting that there is an uncountable cardinal (ordinal) with that property. The property of interest here is $\kappa > \aleph_0 \wedge (\forall n, m \in \mathbb{N})(\kappa \to \kappa^n_m)$. (We'll get round later to considering replacing "$\in \mathbb{N}$" by "$< \kappa$") But let's put it in context. Axioms of infinity are simply strong assertions. 'Continuous functions have fixed points' is not a strong assertion: it follows from replacement. On the face of it the assertion that functions that are α-cts for only some α nevertheless still have fixed points is a bit stronger, but in fact it, too, follows from replacement. If we try asserting that some function like $\mathcal{P}$ (which is not α-cts for any α) has a fixed point then the sky falls in (that's Cantor's theorem) so the way to flirt profitably with danger is to assert the existence of fixed points for functions that are α-discontinuous for almost all α. As things stand at this stage in this development we have no evidence that $\mathbf{E}$ is α-cts for any known α and so it could be a good function to postulate fixed points for.

Why might the supremum of an γ-sequence of values of $\mathbf{E}$ not be a fixed point? The problem is: if α is such a limit we know only that $(\forall \beta < \alpha)(\alpha \to (\beta)^2)$. Can we glue all the little (β-sized) monochromatic sets together to get a big (α-sized) monochromatic set? This seems to need a kind of compactness property which we will soon discover.

Let us look at properties that these fixed points must have.

Remark 100. *If $\kappa \to (\kappa)^2_2$ then κ is regular.*

Proof: Let K be a set of size κ, and consider an arbitrary partition. Assign a pair $\{k_1, k_2\}$ to one of two bins depending on whether or not the k_i belong to the same element of the partition. Sets monochromatic in one sense give us a κ-sized element of the partition, and sets monochromatic in the other sense tell us there are at least κ elements of the partition. This says that K is not a union of fewer than κ smaller sets, so κ is regular.

■

Remark 101. *If $\kappa \to (\kappa)^2_2$ then κ is strong limit.*

Proof:

We have proved already that if $\kappa \to (\kappa)^2_2$ then κ is regular.

Suppose there is $\beta < \kappa$ such that $2^\beta \geq \kappa$. Let B be a set of size β and K a set of size κ, both equipped with worders of the length of the corresponding initial ordinal. Let $g : K \hookrightarrow (B \to \{0,1\})$ be injective. ($B \to \{0,1\}$ is the set of maps from B to the pair $\{0,1\}$.) We now two-colour the unordered pairs from κ by $\chi(\{\gamma,\delta\}) =$ blue if $(\gamma < \delta) \longleftrightarrow g(\gamma) < g(\delta)$ in the lexicographic order on $B \to \{0,1\}$, and pink otherwise. There will be a monochromatic set $K' \subseteq K$ of size κ, and g"K' will be a subordering of $B \to \{0,1\}$ of length κ (or κ^*—it won't matter which: without loss of generality take it to be κ). There is nothing wrong with the total order $B \to \{0,1\}$ having large subsets. The point we are going to exploit is not that it has a subset of *size* κ but that it has a subset of *length* κ in the inherited order. That is strong.

These sequences in g"K' might all have the same first entry, and second entry too for that matter: they might all begin $1,0,0\ldots$ for example. Look for the first address at which they differ, and record it as β_0. This chops g"K' into two pieces: those whose β_0th entry is 0 and those whose β_0th entry is 1. This second piece will be of length κ and we retain it to work on. We record the length of the first piece as κ_0.

Now we look for the first address at which the sequences in this second piece differ, and record it as β_1. As before, this chops the (tail of) g"K' into two pieces. We record β_1 and we retain the tail for the same treatment as before. We also record the length of the concatenation of the two initial segments as κ_1.

We might not run out of β_is (since g"K' might in principle be included in a bounded interval of $B \to \{0,1\}$) but we are certainly going to run out of κ_is. This means that $cf(\kappa) \leq \beta$, contradicting regularity of κ. ■

DEFINITION 102. *Regular strong limit cardinals are said to be* **strongly inaccessible**.

So

COROLLARY 103. *Any $\kappa \to (\kappa)^2$ will be strongly inaccessible.*

EXERCISE 38.

1. *If κ is strongly inaccessible then $H_\kappa = V_\kappa$;*

2. *H_κ is a model of replacement if κ is regular;*

3. *V_κ is a model of power set if κ is limit.*

H_κ tends to be a model of everything except possibly power set (recall that $H_{\aleph_1}$ is the usual example of a structure that models all the ZF axioms except power set); V_κ tends to be a model of all the axioms except replacement (recall that $V_{\omega+\omega}$ is the usual example of a structure that models all the ZF axioms except replacement), so if $H_\kappa = V_\kappa$ we expect that structure to be a model of ZF. It therefore follows from Gödel's incompleteness theorem that the existence of such cardinals is unprovable in *ZF*.

We now need to consider a generalisation of König's Infinity lemma.

Definition 104. *A cardinal κ has the* **tree property** *if every tree of size κ all of whose levels have fewer than κ elements has a branch of length κ.*

Evidently $\aleph_0$ has the tree property, and it is what we need to prove Ramsey's theorem.[11]

Theorem 105. *If $\kappa \to (\kappa)^2_2$ and κ has the tree property, then $(\forall m, n \in \mathbb{N})(\kappa \to (\kappa)^m_n)$*

Theorem 106. *If κ is strong limit and has the tree property, then $\kappa \to (\kappa)^2_2$.*

Proof:

Consider the Erdős-Rado theorem and what it tells us about κ when κ is strong limit. Evidently $\kappa \to (\alpha)^n_2$ for all $\alpha < \kappa$. With the tree property we can "glue" together the growing small monochromatic sets to get a monochromatic set of size κ.

These small monochromatic sets are all of size $< \kappa$ by hypothesis, so there are κ of them. Organize them into a tree by partially ordering them by end-extension. But not exactly! The point is that this tree is not automatically $< \kappa$ branching. The way out of this is to put $s \frown \{\alpha\}$ not immediately below s, but to keep on repeating s until the αth level of the tree is reached. *Then* you put in $s \frown \{\alpha\}$. (So a tree here is not a kind of partially ordered set but a kind of partially ordered *multiset*).

■

This property we have just seen—$\kappa \to (\kappa)^2_2$—is equivalent to the conjunction of (i) strong inaccessibility (as we have just seen in corollary 109, strong inaccessibility follows from this) and (ii) a property called "Weak compactness". This is the following:

Definition 107. *κ is* **weakly compact** *if $\kappa > \aleph_0$ and for all $X \subseteq \mathcal{P}(\kappa)$ with $|X| = \kappa$ there is a measure $\mu : X \to \{0, 1\}$ satisfying:*

1. *if a and $\kappa \setminus a$ both $\in X$ then $\mu(a) + \mu(\kappa \setminus a) = 1$;*

2. *if $\{\alpha\} \in X$ then $\mu(\{\alpha\}) = 0$;*

3. *μ is $< \kappa$-additive.*

In some versions of this definition X is required to be a subalgebra of $\langle \mathcal{P}(\kappa), \subseteq \rangle$.

[11] At least in the strong form that every partition of an *infinite* set has an infinite monochromatic set: if we weaken 'infinite' to countably infinite' we don't.

Definition 108. *$\mathcal{L}_{\beta\alpha}$ is the language (or family of languages) declared recursively like the predicate calculus with two extra constructors: strings of quantifiers of length less than α may be prefixed to any formula, and a conjunction (or disjunction) of fewer than β formulæ in $\mathcal{L}_{\beta\alpha}$ is also in $\mathcal{L}_{\beta\alpha}$.*

Thus ordinary predicate calculus that you all know and love is $\mathcal{L}_{\omega\omega}$. Notice that these languages are all first-order. Despite that, they can express a variety of second-order concepts because an infinite quantifier prefix behaves a bit like a single second-order quantifier. For example, we can capture wellfoundedness in $\mathcal{L}_{\omega_1\omega_1}$ by

$$(\forall x_1, x_2 \ldots)(\bigvee_{i \in \mathbb{N}} \neg R(x_i, x_{i+1}))$$

... which says that R is wellfounded. However wellfoundedness cannot be captured in $\mathcal{L}_{\omega_1\omega}$.

F is **free** on $X \subseteq F$ if for every map $f : X \to A$ there is exactly one morphism $t : F \to A$ making the diagram below commute. *Prima facie* this assertion is not nth-order for any n and is not definable in $\mathcal{L}_{\kappa\kappa}$; see Eklof, [11].

$$X f A_i t F$$

$i : X \hookrightarrow F$ is of course the identity: the inclusion map

There is also

Theorem 109. *Scott's Isomorphism theorem*

Every countable structure can be characterised up to isomorphism by a single sentence of $\mathcal{L}_{\omega_1\omega}$.

Proof:

We can obviously do this by cheating: if we want to characterise $\mathfrak{A}$ up to isomorphism by providing a name a for every element of A, the carrier set of $\mathfrak{A}$. However we want to do it without cheating!

For every tuple $a_1 \ldots a_n$ and every ordinal $\beta < \omega_1$, the formula $\phi^\beta_{a_1 \ldots a_n}$ is defined by recursion as follows:

- $\phi^0_{a_1 \ldots a_n}$ is $\bigwedge\{\theta(x_1 \ldots x_n) : \mathfrak{A} \models \theta[a_1 \ldots a_n]\}$ where θ is atomic or negatomic.
- if β is limit $\phi^\beta_{a_1 \ldots a_n}$ is $\bigwedge\limits_{\gamma<\beta} \phi^\beta_{a_1 \ldots a_n}$
- $\phi^{\beta+1}_{a_1 \ldots a_n}$ is $\phi^\beta_{a_1 \ldots a_n} \wedge \bigwedge\limits_{a_{n+1} \in A} (\exists x_{n+1})(\phi^\beta_{a_1 \ldots a_{n+1}}) \wedge (\forall x_{n+1})(\bigvee\limits_{a_{n+1} \in A} (\phi^\beta_{a_1 \ldots a_{n+1}}))$.

Observe that, for all tuples $a_1 \ldots a_n$ and all $\beta < \omega_1$,

- the formula $\phi^\beta_{a_1 \ldots a_n}$ has at most the free variables 'x_1' ... 'x_n'; and
- $\mathfrak{A} \models \phi^\beta_{a_1 \ldots a_n}[a_1 \ldots a_n]$;

- $\mathfrak{A} \models (\forall x_1 \ldots x_n)\phi^{\beta}_{a_1 \ldots a_n} \rightarrow \phi^{\gamma}_{a_1 \ldots a_n}$ whenever $\gamma < \beta < \omega_1$.

We are now in a position to prove theorem 115.

Observe that—by the third bullet—for each tuple $a_1 \ldots a_n$ from A and for every tuple $x_1 \ldots x_n$ the truth-value of $\phi^{\beta}_{a_1 \ldots a_n}(x_1 \ldots x_n)$ decreases monotonically as β increases. (If it ever becomes false it remains false). So the truth value is eventually constant. So to each $2n$-tuple $a_1 \ldots a_n$ with tuple $x_1 \ldots x_n$ we can associate the ordinal at which the truth-value of $\phi^{\beta}_{a_1 \ldots a_n}(x_1 \ldots x_n)$ *settles down.* Fix $a_1 \ldots a_n$. There are only countably many tuples $x_1 \ldots x_n$ so there are only countably many such ordinals. ω_1 is regular, so, for each tuple $a_1 \ldots a_n$, there will come a stage by which the truth-values of $\phi^{\beta}_{a_1 \ldots a_n}(x_1 \ldots x_n)$ have settled down for all $x_1 \ldots x_n$. Again, there are only countably many tuples $a_1 \ldots a_n$, so (by regularity of ω_1 again) there is a countable sup of all the settling-down ordinals.

[stuff missing]

there is $\alpha < \omega_1$ such that for all $\beta \geq \alpha$

$$\mathfrak{A} \models (\forall x_1 \ldots x_n)(\phi^{\alpha}_{a_1 \ldots a_n} \longleftrightarrow \phi^{\beta}_{a_1 \ldots a_n})$$

It follows that there is $\alpha < \omega_1$ such that, for all tuples $a_1 \ldots a_n$ from A and all $\beta \geq \alpha$,

$$\mathfrak{A} \models (\forall x_1 \ldots x_n)(\phi^{\alpha}_{a_1 \ldots a_n} \longleftrightarrow \phi^{\beta}_{a_1 \ldots a_n})$$

Now let ϕ be the sentence

$$\phi^{\alpha}_{\emptyset} \wedge \bigwedge_{n<\omega} \bigwedge_{a_1 \ldots a_n} (\forall x_1 \ldots x_n)(\phi^{\alpha}_{a_1 \ldots a_n} \rightarrow \phi^{\alpha+1}_{a_1 \ldots a_n})$$

duplication

We use a back-and-forth construction to show that any two countable models of the Scott sentence are isomorphic.

Proposition 110. *If κ is weakly compact and X is a set of sentences of $\mathcal{L}_{\kappa\kappa}$ such that $|X| \leq \kappa$ and every $X' \subseteq X$ with $|X'| < \kappa$ has a model then X has a model.*

Proof: omitted ■

This is sometimes taken as a *definition* of weak compactness. If we drop the clause "$|X| \leq \kappa$" then we obtain the following definition of strong compactness:

Definition 111. *κ is* **strongly compact** *iff whenever X is a set of sentences of $\mathcal{L}_{\kappa\kappa}$ such that every $X' \subseteq X$ with $|X'| < \kappa$ has a model then X has a model.*

(If κ is strongly compact then the theory of free groups is $\mathcal{L}_{\infty,\kappa}$ axiomatisable. see Eklof [12])

Although we have aproached weak compactness thru' partition relations it could have been equally easily motivated as a generalisation of a property of $\aleph_0$. The next property we consider also arises from a simple generalisation of a property of $\aleph_0$.

4.8 Measurable cardinals

DEFINITION 112. *A filter F is κ-complete if $X \subseteq F \wedge |X| < \kappa \to \bigcap X \in F$. (it is "closed under fewer than κ intersections".)*

...except that when people say a filter is countably complete they always mean that it is $\aleph_1$-complete. Notice that the definition of a filter *tout court* says that a filter is $\aleph_0$-complete!

We can use Zorn's lemma to prove that there is a nonprincipal $\aleph_0$-complete ultrafilter on any countable set. Can we find uncountable cardinals with the analogous property?

DEFINITION 113. *A cardinal κ is* **measurable** *iff there is a nonprincipal ultrafilter $\mathcal{U}$ on $\{\alpha : \alpha < \kappa\}$ which is κ-complete.*

REMARK 114. *Every measurable cardinal is weakly compact and therefore strongly inaccessible.*

We appeal to this later, in part 3 of proposition 125.

4.8.1 Measurable cardinals and elementary embeddings

Let us consider the ultrapower $V^\kappa/\mathcal{U}$. There is a problem here because the equivalence classes (of which the model is composed) are proper classes. The functions themselves aren't (because of replacement). The obvious solution is to pick a representative from each class by Global AC but if you prefer not to use global AC you may replace each equivalence class by the set of those of its members that are of minimal rank. (This assumes foundation! It is *Scott's trick* of which I have told you earlier)

In general there is no reason to suppose that the ultrapower is wellfounded. Suppose we take an ultrapower $V^\omega/\mathcal{U}$ where $\mathcal{U}$ is nonprincipal over ω. Consider the (Von Neumann) integers of the ultrapower. We want to show that they are not wellfounded. Since $\leq_{\mathbb{N}}$ is the same as $\in$ where von Neumann integers is concerned, this will show that the ultrapower is not really well founded (though it will be a model of the (first-order) axiom of foundation—cf Peano arithmetic). To do this it will be sufficient to find a countable family $\langle u_i : i \in \mathbb{N}\rangle$ of elements of $\mathcal{U}$ such that each integer belongs to only finitely many u_i. That way we can form an infinitely descending sequence of functions $\mathbb{N} \to \mathbb{N}$ (i.e., natural numbers in the sense of the ultrapower) where the values start off sufficiently big, and we create a smaller natural number by decreasing the values at all coordinates in u_1, then at all coordinates in u_2, then $u_3 \ldots u_n$ so that we get an infinitely descending sequence without any coordinate being decreased more than finitely many times. This is easy: just take u_i to be $\{n \in \mathbb{N} : n > i\}$.

But this follows immediately from the fact that if the ultrafilter is countably incomplete then the ultraproduct is saturated. Clearly no saturated model can be wellfounded: if there is an x that has $\in$-chains of all finite lengths hanging off it ...

However if the ultrafilter is countably complete the ultrapower will be wellfounded.

Proposition 115. *If V is wellfounded and $\mathcal{U}$ is countably complete over κ then $V^\kappa/\mathcal{U}$ is wellfounded.*

Proof. Suppose $\langle [f_i] : i \in \mathbb{N}\rangle$ satisfies $(V^\kappa/\mathcal{U}) \models [f_{i+1}] \in [f_i]$ for each $i \in \mathbb{N}$. Use AC_ω to pick f_i for each i. Then, for each $i \in \mathbb{N}$, let A_i be $\{\alpha < \kappa : f_{i+1}(\alpha) \in f_i(\alpha)\}$. All A_i are in $\mathcal{U}$ by hypothesis. But then, by countable completeness of $\mathcal{U}$, the intersection $\bigcap_{i\in\mathbb{N}} A_i$ is nonempty (it is actually in $\mathcal{U}$), and for any address β in it, it is the case that $(\forall i)(f_{i+1}(\beta) \in f_i(\beta))$, contradicting the assumption that there are no ω-descending $\in$-chains in V.

■

We could have stated something more general:

Exercise 39. *If $\mathcal{U}$ is a κ-complete ultrafilter over a measurable cardinal κ, then Łoś's theorem holds for sentences of $\mathcal{L}_{\kappa\kappa}$ and ultraproducts modulo $\mathcal{U}$*

We then obtain proposition 121 as a consequence of the fact that wellfoundedness can be captured by an expression of $\mathcal{L}_{\omega_1,\omega_1}$.

This does make it sound as if what is really important is the countable completeness of $\mathcal{U}$, since that it what enables us to show that the ultrapower is wellfounded. So why the κ-completeness in the definition of "measurable"? Suppose there is a countably complete ultrafilter $\mathcal{U}$ on κ, and $f : A \twoheadrightarrow \kappa$. Then $\{f^{-1}``X : X \in \mathcal{U}\}$ is a countably complete ultrafilter on A, so that if we had taken the posession of a countably complete nonprincipal ultrafilter to be our definition of measurable, then any cardinal (surjectively) larger than a measurable one would be measurable.

Next we take the Mostowski collapse of the ultrapower to obtain a new model, which it is customary to denote '$\mathfrak{M}$'. We will do this sufficiently often to have a notation. Let us use 'ψ' for this. We have a picture:

$$V \hookrightarrow_{i_\mathcal{U}} V^\kappa/\mathcal{U} \simeq_\psi \mathfrak{M}$$

giving an elementary embedding $V \hookrightarrow \mathfrak{M}$—namely $\psi \cdot i_\mathcal{U}$—which tradition in this area requires us to write 'j'.[12] Tradition also requires that the measurable cardinal be 'κ', and the Mostowski collapse of the ultrapower be '$\mathfrak{M}$'.

Lemma 116. *j is not the identity.*

Proof:

Think about $\psi[\lambda\alpha.\kappa]$ and $\psi[id]$. They are both ordinals of $\mathfrak{M}$. Clearly we have $\psi[\lambda\alpha.\kappa] > \psi[id]$. Clearly j sends κ to $[\lambda\alpha.\kappa]$ and thence to $\psi[\lambda\alpha.\kappa]$. Now, for any $\beta < \kappa$, Łoś's theorem tells us that $\psi[\lambda\alpha.\beta]$ is an ordinal and, since $\lambda\beta.\alpha$

[12] I have earlier in these notes used the letter 'j' as a notation for $\lambda f.\lambda x.(f``x)$. This use of '$j$' (for '$j$ump' is a Forsterism and is not standard: in contrast the use here of 'j' as a variable for elementary embeddings *is* standard—in the same sense that the use of 'x axis' and 'y axis' for 'abscissa' and 'ordinate' is standard.

is almost everywhere less than the identity, it is a smaller ordinal than $\psi[id]$. (Again by Łoś's theorem).

So the situation is this. Ordinals below κ get sent to ordinals less than $\psi[id]$. On the other hand κ itself gets sent to $\psi[\lambda\alpha.\kappa]$ which is strictly bigger. This means that j is not continuous at κ and cannot be the identity. ■

This isn't the same as saying that $\mathfrak{M} \neq V$. For all we know j might be a nontrivial elementary embedding from the universe into itself. As it happens, it isn't—there are none!—but we don't know that until corollary 128.

Next we show

LEMMA 117. *A non-trivial elementary embedding must send some ordinal κ to something $j(\kappa) > \kappa$, and the least such ordinal is initial.*

Proof: (*Idea: if all ordinals are fixed then by induction on rank all sets are fixed too. No ordinal can be sent to anything smaller so some ordinal is moved to something bigger*)

We observe that by elementarity of j we have

$$(\forall x)(\forall y)(y \in x \longleftrightarrow j(y) \in j(x))$$

Now we also have (for any injective j whatever!)

$$(\forall x)(\forall y)(y \in x \longleftrightarrow j(y) \in j``(x))$$

This implies $j``x \subseteq j(x)$. (Q: Why don't we get equality? A: $j(x)$ might have members not in the range of j!). This enables us to prove by induction that j can never move anything to a set of lower rank. If j is not the identity, consider an object x of minimal rank that is moved by j. Suppose x is moved but $\rho(x)$ fixed. Then

Suppose $y \in j(x)$. Then $\rho(y) < \rho(j(x)) = \rho(x)$ so $j(y) = y$. So $j(x) \subseteq$ range of j.

$$j(y) \in j(x)$$

iff

$$y \in x$$

by elementarity, so $j(x) = x$ by extensionality and x is not moved. So if x is a thing of minimal rank, $\rho(x)$ is also moved.

It's probably worth noting here that we are exploiting the fact that (the graph of) j is locally a set, in the sense that its intersection with any set is a set. If j had been merely some random external automorphism (in which case—admittedly—we would be working in a nonstandard model) we would not have been able to say "consider an object x of minimal rank that is moved by j ..."

So if anything is moved, an ordinal is moved. Now let κ be the first ordinal moved by j, we must show that κ is an initial ordinal. First we notice that it must be limit. If it is not initial, we have $\langle \kappa, \leq_{On} \rangle \simeq \langle \mu, R \rangle$ for some relation R on some ordinal $\mu < \kappa$. But then $j\langle \kappa, \leq_{On} \rangle \simeq j\langle \mu, R \rangle$ by elementarity, and $j\langle \mu, R \rangle = \langle \mu, R \rangle$ by minimality of κ (this is where we need κ to be limit, so that $\rho\langle\mu, R\rangle < \kappa$) so $j\langle\kappa, \leq_{On}\rangle \simeq \langle\kappa, \leq_{On}\rangle$ contradicting assumption that κ is moved.

■

This has a nice side-effect which we will need:

COROLLARY 118. *If $j : V \hookrightarrow \mathfrak{M}$ with κ the first ordinal moved, then things of rank less than κ are fixed by j, so $\mathfrak{M} \cap V_\kappa = V_\kappa$.*

Proof: One direction of the inclusion we have just proved, and for the other direction, remember $\mathfrak{M} \subseteq V$!

■

PROPOSITION 119. *Fix a κ-complete nonprincipal ultrafilter $\mathcal{U}$ on κ and consider $\mathfrak{M}$, the transitive collapse of the ultrapower. Then*

1. *The elementary embedding $j : V \hookrightarrow M$ is not the identity and κ is the first ordinal moved;*
2. *$\mathfrak{M}$ is closed under κ sequences but not κ^+-sequences;*
3. *$\kappa < 2^\kappa < j(\kappa) < (2^\kappa)^+$.*

We write "$\mathfrak{M}$ is closed under κ sequences" as '$\mathfrak{M}^\kappa \subseteq \mathfrak{M}$' (even tho' we could have written it as '$\mathcal{P}_{\kappa^+}(\mathfrak{M}) \subseteq \mathfrak{M}$' using a notation we already have) because this is the notation most commonly used in the literature.

Proof:

1. Let β be an ordinal below κ. The elementary embedding into the ultrapower will send β to $[\lambda x.\beta]_\mathcal{U}$, and this must get sent to β in the Mostowski collapse. This is because anything below it in the ultrapower is $[f]$ where f is almost everywhere less than β. Consider the preimages in f of the ordinals below β. There are only β of them and they add up to a set of measure 1. Therefore one of them is a set of measure 1, which is to say that $f \sim_\mathcal{U} \lambda x.\alpha$ for some $\alpha < \beta$. Thus there are precisely β things in $\mathfrak{M}$ below $j(\beta)$, so $j(\beta) = \beta$.

   For any $\alpha < \kappa$ we have $[\lambda x.\kappa] >_\mathcal{U} [id] >_\mathcal{U} [\lambda x.\alpha]$, so κ is certainly moved. Everything below κ is fixed, so κ is the first thing moved.

2. $\mathfrak{M}$ is closed under κ-sequences but not κ^+-sequences.

   We will show that if $j``x \in \mathfrak{M}$, $y \subseteq M$ and $|y| \leq |x|$ then $y \in \mathfrak{M}$. (This will be sufficient to show that $\mathfrak{M}$ is closed under $\kappa$-sequences because $j``\kappa = \kappa \in \mathfrak{M}$.) Represent y as $\{\psi([t_a]) : a \in x\}$ and define $T : j``x \to y$ by $T(j(a)) = \psi([t_a])$. Since $T$ enumerates $y$, it suffices to show that $T \in \mathfrak{M}$. So, we need a $g$ so that $\psi([\lambda\alpha.g]) = T$, that is, $dom\psi([\lambda\alpha.g]) = j``x$ and for all $a \in x$, $\psi([\lambda\alpha.g])(j(a)) = \psi([t_a])$. Let $[f] = \psi^{-1}(j``x)$. So there is a thing in the ultrapower with the properties we need. By Łoś's theorem if for each $i \in \kappa$ we set $dom(g(i)) = f(i)$, and $(g(i))(a) = t_a(i)$ for each $a \in dom(g(i))$, then clearly g is a thing in V with the required properties.[13]

To show that $\mathfrak{M}$ is not closed under κ^+-sequences it will be sufficient to show that the particular κ^+-sequence $j``\kappa^+$ is not in $\mathfrak{M}$. Suppose $j``\kappa^+(= \psi([f]_\mathcal{U})) \in \mathfrak{M}$. Let $A = \{i < \kappa : |f(i)| \leq \kappa\}$. There are two cases to consider:

(i) If $A \in \mathcal{U}$ then—since κ^+ is regular—there is $\alpha \in \kappa^+ \setminus \bigcup f``A$. But then $j(\alpha) \notin \psi([f])$.
(ii) If $A \notin \mathcal{U}$ so $B = \{i < \kappa : |f(i)| > \kappa\} \in \mathcal{U}$, define h by induction on ordinals so that $h(i) \in f(i) \setminus \{h(j) : j < i \wedge j \in B\}$. Then $[h] \in_\mathcal{U} [f]$—yet h is not constant on any set in $\mathcal{U}$, since $\mathcal{U}$ is nonprincipal.

Hence either way we get a contradiction from the assumption that $j``\kappa^+ = \psi[f]$.

3. Notice that the set of $<_{V^\kappa/\mathcal{U}}$-predecessors of $[\lambda x.\kappa]$ is a quotient of κ^κ (Here κ^κ is of course not an ordinal but is the set of maps from ordinals-below-κ to ordinals-below-κ! O the joys of overloading). κ^κ is the same size as $\mathcal{P}(\kappa)$ so the wellorder they form must be shorter than $(2^\kappa)^+$. Also, (in $\mathfrak{M}$) by elementarity, $j(\kappa)$ is measurable and hence strongly inaccessible (by remark 120) so that $(2^\kappa)^\mathfrak{M} < j(\kappa)$. But $\mathfrak{M}$ is closed under κ-sequences by (2) so $\mathcal{P}(\kappa) = (\mathcal{P}(\kappa))^\mathfrak{M}$, so $2^\kappa \leq (2^\kappa)^\mathfrak{M}$.

Notice that the fact that $\mathfrak{M}$ is not closed under κ^+-sequences entails that $V \neq \mathfrak{M}$. However we can give slightly more information than this.

■

THEOREM 120. $\mathcal{U} \notin \mathfrak{M}$.

Proof: Assume $\mathcal{U} \in \mathfrak{M}$. [14]. If $\mathcal{U} \in \mathfrak{M}$ then the whole of $\mathcal{P}(\kappa)$ is in $\mathfrak{M}$ too, as is κ^κ. So all the equivalence classes $[f]$ are also in $\mathfrak{M}$, in particular $[\lambda x.\kappa]$. Therefore we can reproduce in $\mathfrak{M}$ the proof that $j(\kappa) \leq (2^\kappa)^+$, but $\kappa < j(\kappa) \leq (2^\kappa)^+$ means that $j(\kappa)$ is not strong limit, contradicting the fact that $j(\kappa)$ is measurable in $\mathfrak{M}$.

■

Notice that we have here used that $\mathcal{U}$ is κ-complete, not just countably complete.

Next we need the failure of a particular infinite exponent partition relation. (NB common notation here: "${}^\omega\lambda$" for the set of ω-sequences of ordinals below λ. This is written instead of "λ^ω" because the latter could also denote the ordinal λ raised to the power of the ordinal ω. This is the price one pays for identifying ordinals with the set of their predecessors! On the whole it's a good bargain!)

[13]Thanks to Nathan Bowler and Phil Ellison for tidying up some infelicities in this proof.

[14]It would be nice to be able to argue: "Then $\mathfrak{M} \models$ "κ is measurable". But $\mathfrak{M} \models$ "$j(\kappa)$ is the first measurable" and $j(\kappa) > \kappa$" but presumably this works only when κ is the first measurable!

Proposition 121. *(Erdős-Hajnal) Let $\lambda^{\aleph_0} = 2^\lambda$. Then there is $F : {}^\omega\lambda \to \lambda$ s.t. if $X \subseteq \lambda$, $|X| = \lambda$, then $F``{}^\omega X = \lambda$.*

Proof:

Let $\langle\langle A_\alpha, \zeta_\alpha\rangle : \alpha < 2^\lambda\rangle$ enumerate $\{X \subseteq \lambda : |X| = \lambda\} \times \lambda$ and let $\langle f_\alpha : \alpha < 2^\lambda\rangle$ enumerate ${}^\omega\lambda$.

Define a sequence S_α as follows.

S_α is the first f_β (in the ordering $\langle f_\alpha : \alpha < 2^\lambda\rangle$) such that $f_\beta \in (A_\alpha)^\omega$ and f_β is not already used.

Now set $F(S_\alpha) = \zeta_\alpha$.

Is F defined on the whole of ${}^\omega\lambda$? That is, does every f_β become an S_α at some stage?

Yes. Let f_β be the first one that does not become an S_α. This must be because $f_\beta \notin {}^\omega A_\alpha$ for any α which is impossible.

Now we must show $F``{}^\omega X = \lambda$ for $X \subseteq \lambda, |X| = \lambda$.

Fix X with $X \subseteq \lambda$ and $|X| = \lambda$. Suppose $F``(X)^\omega \neq \lambda$. Let α be minimal such that ζ_α is not in the range of $F \restriction X^\omega$. Now, by construction of F, $F(S_\alpha) = \zeta_\alpha$. Contradiction.

■

We really wanted this only for the following corollary:

Corollary 122. *(Kunen)*

ZFC $\vdash$ There is no non-trivial elementary embedding $V \hookrightarrow V$.

Proof:

Let $j : V \hookrightarrow V$ be elementary; by lemma 123 there is a first ordinal moved, and we will call it κ. Set $\lambda = sup\{j^n(\kappa) : n \in \mathbb{N}\}$. We already know (by lemma 123) that κ is an initial ordinal, and λ, being a limit of initial ordinals, is also an initial ordinal.

It's obvious that $cf(\lambda) = \omega$, but to invoke proposition 127 we need it to be strong limit as well. But we have already shown that $j(\kappa) \geq 2^\kappa$—this was proposition 125. Therefore $\lambda^{\aleph_0} = 2^\lambda$ (by AC). (prop 71). Also $\lambda = j(\lambda)$.

Now let F be as in the conclusion of proposition 127. Set $X = j``\lambda$. Then $|X| = \lambda$. F satisfies

$$(\forall Y \subseteq \lambda)(|Y| = \lambda \to F``Y^\omega = \lambda)$$

abbreviates to $\Phi(F, \lambda)$.

$$(\forall Y \subseteq j(\lambda))(|Y| = j(\lambda) \to (jF)``Y^\omega = j(\lambda))$$

(which is to say $\Phi(j(F), j(\lambda))$). Now by elementarity (since $j(\lambda) = \lambda$) we have

$$(\forall Y \subseteq \lambda)(|Y| = \lambda \to (j(F))``Y^\omega = \lambda)$$

So $\exists x \in {}^\omega X$ such that $(j(F))(x) = \kappa$ with $X = j``\lambda$ ($X \subseteq \lambda$ and $|X| = \lambda$). But ${}^\omega X = j``{}^\omega\lambda$ so x is $j(y)$ for some y. Therefore $(j(F))(j(y)) = \kappa$. Therefore κ is in the range of j. But κ was the first thing moved, and it certainly cannot be $j(\beta)$ for $\beta < \kappa$. Contradiction.

■

We could have used this to prove that $V \neq \mathfrak{M}$ earlier than we did ... The reader might say: "why are we bothering to prove this? We've already established that $V \neq \mathfrak{M}$!" But all we proved there was that the elementary embedding from V arising from a countably complete ultrafilter on a measurable cardinal is not an elementary embedding into V itself but rather into a submodel. We hadn't proved that there are no elementary embeddings $V \hookrightarrow V$ at all, merely that none can arise from a fancy ultrafilter on a measurable cardinal.

In contrast this last proof doesn't make the assumption that the elementary embedding arose from a measurable cardinal. However it does use AC. It is an open question whether or not there can be elementary embeddings from the universe into itself if AC fails. It's a hard problem beco's it's very difficult to do much with elementary embeddings without AC: Łoś's theorem is equivalent to *AC* after all.

THEOREM 123. *If κ is measurable then it has the tree property.*

Consider a tree which satisfies the conditions in the antecedent of the definition. The root has κ descendents but $< \kappa$ children so at least one child has κ descendents. Better: κ has a measure-one set of descendents but fewer than κ children, so *precisely one* child has measure-one children. Thereafter we rely on the fact that if you chop a measure-one set into fewer than κ bits then precisely one of them is of measure one.

■

4.8.2 Normal ultrafilters

Wwe have seen how a nonprincipal κ-complete ultrafilter on a cardinal κ gives us a submodel $\mathfrak{M}$ and an elementary embedding $j : V \hookrightarrow \mathfrak{M}$ where κ is the first ordinal moved. In this section we shall see how one can come back the other way: if we have a submodel $\mathfrak{M}$ and an elementary embedding $j : V \hookrightarrow \mathfrak{M}$ where κ is the first ordinal moved then there is a nonprincipal κ-complete ultrafilter on κ. We will see how to recover this ultrafilter from j in proposition 136. However, if we start from a nonprincipal κ-complete ultrafilter on a cardinal κ, obtain thereby a submodel $\mathfrak{M}$ and an elementary embedding $j : V \hookrightarrow \mathfrak{M}$ where κ is the first ordinal moved, and then try to recover this ultrafilter from j, it's natural to ask if we get back the ultrafilter we started with. Annoyingly we need the ultrafilter to satisfy an extra condition if this is to work.

DEFINITION 124. *An ultrafilter is* **normal** *if an $f : \kappa \to \kappa$ which is less than the identity almost everywhere is constant almost everywhere.*

PROPOSITION 125. *Every normal ultrafilter is countably (κ) complete.*

Proof: Given $X \in \mathcal{U}$, split it into $\aleph_0$ pieces $\langle X_i : i \in \mathbb{N}\rangle$. Then the function that sends each $\alpha < \kappa$ to the least member of the unique X_i s.t. $\alpha \in X_i$ is pressing down and therefore constant on a set in $\mathcal{U}$. So some X_i extends a large set and is large.

■

Should really say something about why not every κ-complete filtre is normal

PROPOSITION 126. *Normal ultrafilters are closed under diagonal intersection.*

Let $\mathcal{U}$ be normal and let $\langle C_\alpha : \alpha < \kappa\rangle$ be a sequence of things in $\mathcal{U}$. Set $C = \{\alpha : \alpha \in \bigcap_{\beta<\alpha} C_\beta\}$.

Consider $f = \lambda\alpha.(\mu\beta)(\alpha \notin C_\beta)$. If $C \notin \mathcal{U}$ then f is pressing-down. (why?) But if f is pressing-down it is constant on a set of measure 1, that is to say, there is α_0 st $\{\alpha : f(\alpha) = \alpha_0\} \in \mathcal{U}$. Now $f(\alpha) = \alpha_0$ implies $\alpha \notin C_{\alpha_0}$ so $\{\alpha : f(\alpha) = \alpha_0\} \cap C_{\alpha_0}$ is empty, which is impossible since they are both of measure 1.

■

This diagonal intersection business ought to remind you of the definition of p-point and Ramsey ultrafilters, (not lectured but in section 4.7.4) and suggest that we should expect results like the following.

PROPOSITION 127. *If $\mathcal{U}$ is a normal ultrafilter on κ and $f : [\kappa]^2 \to \{0,1\}$ then f has a monochromatic set in $\mathcal{U}$.*

Proof: Exercise. (use proposition 132.)

There are other similarities. (i) Normal ultrafilters are RK-minimal (tho' the converse no longer holds) (ii) If we arm ourselves with the notion of a **uniform ultrafilter on** κ (one all of whose elements are of size κ) we find that "uniform" corresponds to "nonprincipal" in the countable case. A nonuniform ultrafilter on a countable set is simply a principal ultrafilter. Any ultrafilter on a set X gives rise to an ultrafilter on any superset of X: a nonuniform ultrafilter on Y is simply one that arose from a (uniform) ultrafilter on a smaller set. (If Y is countable the only ultrafilters on smaller sets are principal).

The contrast with the countable case is that $\mathcal{U}$ is countably complete, and we can make use of this fact.

THEOREM 128. *Every measurable cardinal admits a normal measure.*

Consider the transitive collapse of the ultrapower, and let $j : V \hookrightarrow M$. κ is the first thing moved (by lemma 123). Then $\{K \subseteq \kappa : \kappa \in j(K)\}$ is a normal ultrafilter, $\mathcal{U}$.

Suppose $f : \kappa \to \kappa$ is pressing-down. Then $\{\alpha < \kappa : f(\alpha) < \alpha\} \in \mathcal{U}$ which is to say $\kappa \in j(\{\alpha < \kappa : f(\alpha) < \alpha\}) = \{\alpha < j(\kappa) : (j(f))(\alpha) < \alpha\}$. Therefore $(j(f))(\kappa) < \kappa$. Abbreviate $(j(f))(\kappa)$ to α_0. α_0 is fixed because it is below κ. Now consider $\{\alpha < j(\kappa) : (j(f))(\alpha) = \alpha_0\}$. κ is a member of it, so if it is j of some X we infer $X \in \mathcal{U}$. X is $\{\alpha < \kappa : f(\alpha) = \alpha_0\}$ (since α_0 is fixed). Therefore $\{\alpha < \kappa : f(\alpha) = \alpha_0\} \in \mathcal{U}$ as desired.

■

Notice that we can do the same construction of an ultrapower (and discover it to be wellfounded) as long as we have a countably complete ultrafilter on κ. As remarked earlier, if κ is measurable, then for any $\lambda > \kappa$ there is a countably complete ultrafilter on λ, so we can form the ultrapower using that. However, there is no reason to suppose that the first ordinal moved will be the λ we had in mind. If the ultrafilter we used arose from an ultrafilter on some $\kappa < \lambda$, then it is that κ which will be the first thing moved. Take this as an exercise after going through the proof of proposition 125.

Notice that we have proved on the fly that if there is an elementary embedding $V \hookrightarrow \mathfrak{M}$ then the first thing moved is measurable.

If $\mathcal{U}$ is normal we can say a bit more about proposition 125.

Remark 129. *Suppose $\mathcal{U}$ is normal. If $g : \kappa \to \kappa$ satisfies $[g] <_{\mathcal{U}} [id|\kappa]$ (so that it corresponds to an ordinal $\leq$ the ordinal corresponding to the identity) then it is almost constant$^{\mathcal{U}}$, i.e., it corresponds to an ordinal below κ. That is to say, $\psi([id|\kappa]) = \kappa$.*

Proof. If $\mathcal{U}$ is normal, then $j(\kappa) = \psi([id]_{\mathcal{U}})$. For $\psi([f])$ to be an ordinal $\leq \kappa$ we need $\{\alpha : f(\alpha) \leq \kappa\} \in \mathcal{U}$. By normality, if $f(\alpha) < \alpha$ on a set in $\mathcal{U}$ then f is almost constant, so gets sent to an ordinal below κ. Therefore $\psi([id])$ is the first thing bigger than all these, namely κ. (o/w $j(\kappa)$ might be something $< [id]$).

■

Let us continue thinking about the submodel $\mathfrak{M}$ and the elementary embedding j when $\mathcal{U}$ is normal (so that $\psi([id]) = \kappa$). By Łoś's theorem we conclude that $\mathfrak{M} \models \phi(\kappa)$ if and only if $V \models \{\beta < \kappa : \phi(\beta)\} \in \mathcal{U}$. So, at least if we consider properties of κ which are preserved between V and M, we find that $\phi(\kappa)$ iff $\{\beta < \kappa : \phi(\beta)\} \in \mathcal{U}$. For which ϕ is this going to hold?

Well, by corollary 124, V and $\mathfrak{M}$ have the same sets of rank $< \kappa$, so any property of κ that can be characterised by quantifying solely over things of rank $< \kappa$ will satisfy this. ? This tells us that if κ is measurable it is the κth weak compact, etc. So it is very big indeed! weak compactness

Proposition 130. *$\mathcal{U}$ is normal iff it is the filter $\{X \subseteq \kappa : \kappa \in j(X)\}$.*

Proof:

The $R \to L$ implication we have seen already (it was theorem 134). For the other direction let $\mathcal{U}$ be normal. We want: $(\forall X \subseteq \kappa)(X \in \mathcal{U} \longleftrightarrow \kappa \in j(X))$. Now $\kappa = \psi([id]_{\mathcal{U}})$ by normality of $\mathcal{U}$ and 135 and $j(X)$ (for $X \subseteq \kappa$) is $\psi([\lambda\alpha.X])$. So $\kappa \in j(X)$ is

$$[id]_{\mathcal{U}} \in_{\mathcal{U}} [\lambda\alpha.X]$$

which is

$$\{\alpha < \kappa : \alpha \in X\} \in \mathcal{U}$$

which is $X \in \mathcal{U}$.

■

Definition 131. *The notation*

$$\kappa \to (\alpha)^{<\beta}_{\gamma}$$

means that given a family $\{\pi_\zeta : \zeta < \beta\}$ of partitions of K^ζ into γ pieces there is a set which is monochromatic for all these partitions **simultaneously**.

In fact we will use this notation only when $\beta = \omega$.

Definition 132. *if $\kappa \to (\kappa)^{<\omega}_{\gamma}$ for all $\gamma < \kappa$ then κ is* **Ramsey**.

Proposition 133. *Every measurable cardinal is Ramsey.*

Measurable cardinals have the tree property and are strongly inaccessible. Inspection of the proof of the Erdoš-Rado theorem (theorem 105) reminds us that every strong inaccessible is a limit of values of **E** and that the tree property for α is what is needed to show that α is a fixed point for **E** given that α is a limit of values of **E**. So if κ is measurable we have $\kappa \to (\kappa)^n_m$. Normal ultrafilters generalise the concept of Ramsey ultrafilters on $\mathbb{N}$, and so we know—given a normal ultrafilter $\mathcal{U}$ on κ—that any partition of n-tuples of ordinals below κ will have a monochromatic set in $\mathcal{U}$. This was proposition 133. But $\mathcal{U}$ is countably complete! Therefore, given a sequence $\langle \Pi_i : i \in \mathbb{N}\rangle$ where each Π_i is a partition of κ^i, not only will $\mathcal{U}$ contain—for each i—a set monochromatic for Π_i, but it will contain the intersection of them all. This is a set that is monochromatic for all the $\langle \Pi_i : i \in \mathbb{N}\rangle$ *simultaneously*. In virtue of this, κ is Ramsey. ■

Ramsey cardinals give indiscernible sets of ordinals as follows. Let κ be Ramsey, and for each $k \in \mathbb{N}$ partition the increasing k-tuples of ordinals below κ into $2^{\aleph_0}$ pieces depending on which formulæ in the language of set theory (with k free variables) they satisfy. Since κ is Ramsey, there is a subset of (the ordinals below) κ that is monochromatic for all these partitions simultaneously.

We have seen how countably complete ultrafilters give elementary embeddigns, and how you can get countably complete ultrafilters from elementary embeddings. It turns out that you can get elementary embeddings (but not defined on the whole universe) merely from the existence of Ramsey cardinals. Suppose S is an indiscernible set of ordinals below κ. Then, if f is a map $\kappa \hookrightarrow \kappa$ with $f``S \subseteq S$, f will give an elementary embedding from L_κ into itself, but we don't know about L_κ just yet.

4.9 Even larger cardinals

Proposition 125 tells us that when j is the elementary embedding arising from an ultrafilter over a measurable cardinal $j(\kappa)$ isn't much bigger than κ, because all its predecessors arose from κ-sequences of ordinals almost all below κ and there aren't that many of them. If we want $j(\kappa)$ to be moved to anything bigger we have to arrange matters so that if κ is the first thing moved then j didn't arise from an ultrafilter on κ. It turns out we need to consider ultrafilters on $\mathcal{P}_\kappa(\lambda)$, but we learn this by trying to generalise the compactness theorem.

We have seen the definition of weak compactness, and the corresponding definition of strong compactness. A technical lemma and a definition.

Definition 134. *Recall that $\mathcal{P}_\kappa(\lambda)$ is the set of subsets of λ of size $< \kappa$.*[15] *Let us also write "$B_\kappa(\alpha)$" for "$\{z \in \mathcal{P}_\kappa(\lambda) : \alpha \in z\}$".*

Then if $\lambda \geq \kappa$ we say $\mathcal{U}$ over $\mathcal{P}_\kappa(\lambda)$ is **fine** *if*

$$(\forall \alpha < \lambda)(B_\kappa(\alpha) \in \mathcal{U})$$

This '$B_\kappa(\alpha)$' is my notation, and i use it because of the use of '$B(x)$' to denote $\{y : x \in y\}$ (as we have seen it's an important operation in set theory with a universal set and this is a local version of it.)

Proposition 135. *the following are equivalent:*

1. *κ is strongly compact; (see definition 117: iff X is a set of sentences of $\mathcal{L}_{\kappa\kappa}$ such that every $X' \subseteq X$ with $|X'| < \kappa$ has a model then X has a model)*
2. *$\forall \lambda \geq \kappa$ there is a fine ultrafilter on $\mathcal{P}_\kappa(\lambda)$;*
3. *For all λ, every κ-complete filter on λ can be extended to a κ-complete ultrafilter.*

Proof:

The proof will be omitted, co's strong compacts are not central.

We will procede to supercompacts and leave bits of the proof lying around for enthusiasts to reconstruct.

$1 \to 2$

Suppose $\lambda \geq \kappa$. Let $\langle A_\nu : \nu < \eta \rangle$ enumerate $\mathcal{P}(\mathcal{P}_\kappa(\lambda))$ and let X be the set of sentences

- $\neg \bigwedge_{\nu \in b} F(A_\nu)$ if $b \subseteq \eta$, $|b| < \kappa$ and $\bigcap_{\nu \in b} A_\nu = \emptyset$;
- $F(A_\nu)$ if $A_\nu = B_\kappa(\alpha)$ for some $\alpha < \lambda$;
- $F(A_\nu) \vee F(A_\beta)$ if $A_\nu = \mathcal{P}_\kappa(\lambda) \setminus A_\beta$.

If $Y \subseteq X$ with $|Y| < \kappa$ then $a = \{\alpha < \lambda : (\exists \nu)((B_\kappa(\alpha) = A_\nu) \wedge A_\nu$ occurs in $Y)\}$ has power $< \kappa$ so interpreting A_ν as A_ν and $F()$ as $\{A_\nu : a \in A_\nu\}$ yields a model of Y. By compactness X has a model and this is the desired filter.

$2 \to 1$.

Now is the moment to recall the slick proof of the completeness theorem: theorem 19. Notice that it depends on the existence, for each $\lambda > \aleph_0$, of a fine ultrafilter on $\mathcal{P}_{\aleph_0}(\lambda)$, tho' we did not have that terminology at that point. The proof we need here is precisely the result of replacing $\aleph_0$ by κ.

The way to fit 3 into this picture is to generalise the result for $\aleph_0$ that every non-principal filter can be extended to a non-principal ultrafilter. ■

Miniexercise. Every non-principal ultrafilter on $\mathcal{P}_{\aleph_0}(\kappa)$ is fine.

[15] Check this: i'm coming to the conclusion that i meant: of size $< \kappa$ tho' i originally intended: of size $\leq \kappa$.

Definition 136. *An ultrafilter $\mathcal{U}$ is* (κ, λ)**-regular** *if it has a subset of λ elements, the intersection of any κ of which is empty.*

Fine ultrafilters on $\mathcal{P}_\kappa(\lambda)$ are (κ, λ)-regular.

Strong compact not as smooth a notion as we would like. Must go to supercompact.

Definition 137.
κ is λ-supercompact iff there is an elementary embedding $j : V \hookrightarrow \mathfrak{M}$ with $crit(j) = \kappa$, $j(\kappa) \geq \lambda$ and $\mathfrak{M}^\lambda \subseteq \mathfrak{M}$;

κ is **supercompact** *iff it is λ-supercompact for all $\lambda \geq \kappa$.*

Two points to note: (i) $\mathfrak{M}^\lambda \subseteq \mathfrak{M}$ implies of course that $\mathfrak{M}$ extends H_{λ^+}; (ii) the definition of supercompact cardinals involves variables ranging over things of arbitrarily high rank.

We saw in proposition 136 how, on being given $j : V \hookrightarrow \mathfrak{M}$, we can recover a nice ("normal") ultrafilter on $crit(j)$. We can do the same here. Since $\mathfrak{M}^\lambda \subseteq \mathfrak{M}$ we look at $\{X : j``\lambda \in j(X)\}$. $|j``\lambda| = \lambda < j(\kappa)$, so $j``\lambda \in \mathcal{P}_{j(\kappa)}(j(\lambda))$ so $\mathcal{P}_\lambda(\lambda)$ is "big" so (since any ultrafilter cuts down to an ultrafilter on any of its members) it is permissible to think of this as an ultrafilter on $\mathcal{P}_\kappa(\lambda)$.

We will take this to be our paradigm normal ultrafilter on $\mathcal{P}_\kappa(\lambda)$. The following definition generalises definition 130 from the measurable case to the supercompact case.

Definition 138. *A* **normal** *ultrafilter on $\mathcal{P}_\kappa(\lambda)$ satisfies the extra condition that for every $f : S \to \lambda$ with $S \in \mathcal{U}$ and $(\forall x \in S)(f(x) \in x)$ there is $\nu < \lambda$ s.t. $f^{-1}``\{\nu\} \in \mathcal{U}$.*

Proposition 139. *For $\mathcal{U}$ a fine ultrafilter on $\mathcal{P}_\kappa(\lambda)$ the following are equivalent:*

1. *$\mathcal{U}$ is normal;*
2. *If $A_\nu \in \mathcal{U}$ for $\nu < \lambda$ then $\{x : x \in \bigcap_{\nu \in x} A_\nu\} \in \mathcal{U}$;*
3. *If j is the obvious embedding then $\psi[id] = j``\lambda$.*

Proof
$1 \to 3$.
$\psi([id]) = j``\lambda$
$(\forall x)(x \in \psi[id] \longleftrightarrow x \in j``\lambda)$
$(\forall x)(x \in \psi[id] \longleftrightarrow (\exists \alpha < \lambda)(x = j(\alpha)))$
$\psi[id] = \{\psi[f] : [f] \in_{\mathcal{U}} [id]\}$ and $[f] \in_{\mathcal{U}} [id]$ is $\{x \in \mathcal{P}_\kappa(\lambda) : f(x) \in x\} \in \mathcal{U}$ so LHS is
$(\exists f)(x = \psi[f] \wedge \{x \in \mathcal{P}_\kappa(\lambda) : f(x) \in x\} \in \mathcal{U}$, so, by normality f is almost constant.

what about the rest of the proof?

■

Proposition 140. *Suppose $\lambda \geq \kappa$. Then the following are equivalent*

1. *$\exists$ normal ultrafilter on $\mathcal{P}_\kappa(\lambda)$;*

2. *There is an elementary embedding $j : V \hookrightarrow \mathfrak{M}$ with κ the first ordinal moved, $j(\kappa) \geq \lambda$ and $M^\lambda \subseteq M$; (i.e., κ is λ-supercompact).*

Proof:

$(1 \to 2)$

The $\mathfrak{M}$ we want will be the Mostowski collapse of $V^{\mathcal{P}_\kappa(\lambda)}/\mathcal{U}$. Let ψ be the Mostowski collapsing function as usual. We want to show that $\mathfrak{M}$ is closed under λ-sequences.

Given $a = \{\psi([g_\nu]) : \nu < \lambda\} \subseteq \mathfrak{M}$, we want to show that $a \in \mathfrak{M}$. Define $h : \mathcal{P}_\kappa(\lambda) \to V$ by $h(x) = \{g_\nu(x) : \nu \in x\}$. Then $a \subseteq \psi([h])$. This is because:

$a \subseteq \psi([h])$ iff

$$\{\psi[g_\nu] : \nu < \lambda\} \subseteq \psi[h]$$

iff

$$(\forall \nu < \lambda)(\psi[g_\nu] \in \psi[h])$$

iff

$$(\forall \nu < \lambda)(g_\nu \in_{\mathcal{U}} h)$$

iff

$$(\forall \nu < \lambda)(\{x \in \mathcal{P}_\kappa(\lambda) : g_\nu(x) \in h(x)\} \in \mathcal{U})$$

and $h(x) = \{g_\nu(x) : \nu \in x\}$ so this is

$$(\forall \nu < \lambda)(\{x \in \mathcal{P}_\kappa(\lambda) : \nu \in x\} \in \mathcal{U}$$

and this is ok, since $\mathcal{U}$ is fine.

Now we want to show $\psi[h] \subseteq a$. So suppose $\psi[g] \in \psi[h]$. This is equivalent to

$$g \in_{\mathcal{U}} h$$

iff

$$\{x : g(x) \in h(x)\} \in \mathcal{U}$$

which is

$$\{x : (\exists \nu)((g(x) = g_\nu(x)) \wedge \nu \in x\} \in \mathcal{U}$$

Call this set S and let f be the function that takes an x in $\mathcal{P}_\kappa(\lambda)$ are returns $(\mu\nu)(g(x) = g_\nu(x) \wedge \nu \in x)$. Then $f : S \to \lambda$ and f satisfies $(\forall x \in S)(f(x) \in x)$ so by normality there is $\alpha < \lambda$ f^{-1}“$\{\alpha\} \in \mathcal{U}$. For this α we have

$$\{x \in \mathcal{P}_\kappa(\lambda) : g(x) = g_\nu(x) \wedge \alpha \in x\} \in \mathcal{U}$$

so $g \sim_{\mathcal{U}} g_\alpha$ whence $g \in_{\mathcal{U}} h$.

We also have to show that if $\mathcal{U}$ is normal then $j(\kappa) \geq \lambda$.

If α is an ordinal below $j(\kappa)$ it is $\psi[f]$ for some $f : \mathcal{P}_\kappa(\lambda) \to \lambda$. Also, if it is strictly below $\lambda x.\kappa$ it satisfies the clause “$(\forall x \in S)(f(x) \in x)$” in the definition of normality of $\mathcal{U}$, so we conclude that there is $\nu < \lambda$ s.t. f^{-1}“$\{\nu\} \in \mathcal{U}$, and $j(f)$ will in fact be this ν.

$(2 \rightarrow 1)$

We proved this direction in the discussion immediately preceding definition 144

■

If $\mathcal{U}$ is merely fine on $\mathcal{P}_\kappa(\lambda)$—but not actually normal—then $\mathfrak{M}^\lambda \not\subseteq \mathfrak{M}$ (for example: $j``\lambda \notin \mathfrak{M}$.) but at any rate each λ-sequence is *covered* by a thing of size λ. That is to say, for every λ-sequence x of elements of $\mathfrak{M}$ there is $y \in \mathfrak{M}$ with $|y| = \lambda$ and $x \subseteq y$.

THEOREM 141. *If κ is 2^κ-supercompact then κ is the κth measurable cardinal.*

Proof: Let $j : V \hookrightarrow \mathfrak{M}$ with critical point κ, so that $\mathfrak{M}$ is closed under 2^κ-sequences. Define $\mathcal{U}$ over κ by

$$X \in \mathcal{U} \longleftrightarrow X \subseteq \kappa \wedge \kappa \in j(X)$$

$\mathcal{U}$ is normal over κ as before. But since $\mathfrak{M}$ is closed under 2^κ-sequences, every ultrafilter over κ is in $\mathfrak{M}$. Hence $(\kappa \text{ is measurable})^{\mathfrak{M}}$, so, by Łoś's theorem , $\{\alpha < \kappa : \alpha \text{ is measurable } \} \in \mathcal{U}$, so κ must be the κth measurable.

■

Woodin has shown that if there is a supercompact cardinal then there is an inner model of AD.

Before we leave supercompacts, a factoid with a logical flavour.

REMARK 142. *If κ is supercompact, Σ_2 (and indeed Π_3) sentences generalise downward to V_κ*

Proof: Suppose $(\exists y)Q(y)$ where Q is Π_1. There will be a b such that $Q(b)$ and, for α sufficiently large, b is in both V_α and H_α. Fix such an α. Now κ is supercompact so there is j and $\mathfrak{M}$ with $\kappa = crit(j)$ such that $j(\kappa) > \alpha$ and $\mathfrak{M}$ is closed under α-sequences so that $H_\alpha \subseteq \mathfrak{M}$. So b is in $V_{j(\kappa)}$ and $V_{j(\kappa)} \models Q(b)$. b is in $\mathfrak{M}$ beco's $b \in H_\alpha \subseteq \mathfrak{M}$. $Q(b)$ is Π_1 and therefore generalises downward to $\mathfrak{M}$.

$\mathfrak{M} \models Q(b)$ therefore, in $\mathfrak{M}$, $V_{j(\kappa)} \models Q(b)$, therefore, in $\mathfrak{M}$, $V_{j(\kappa)} \models \exists y Q(y)$. Therefore $V_\kappa \models (\exists y)Q(y)$ by elementarity. ■

This is best possible beco's "I am a supercompact cardinal" is Σ_3 and cannot generalise down to V_κ if κ is the least supercompact.

4.9.1 coda

No more on supercompacts here, I'm afraid. There are concepts of cardinals with more extreme properties than supercompactness but I don't know a great deal about them. There is a theorem to the effect that if there is a supercompact then there are substructures of the universe in which AD holds. Later versions of these notes will contain a proof.

Beyond the scope of this course!

4.10 Infinite exponent partition relations

First we have to make clear what we mean: $\kappa \to (\kappa)^\alpha$ could mean partitions of *subsets* of cardinality α, or sequences of length α. (There is no ambiguity in the finite case). We mean the second.

THEOREM 143. *If $\kappa \to (\alpha)^\omega_2$ then ${}^\omega\kappa$ cannot be wellordered.*

Proof:

Let K be a set of size κ and consider the equivalence relation on ω-sequences from K of having identical tails. (that is, two sequences are equivalent if one can delete finite initial segments from them to leave identical terminal segments) Suppose that ${}^\alpha\kappa$ can be wellordered. Use such a wellordering to pick one representative from each class. An ω-sequence from K is **even** if the shortest initial segment you can saw off (to get something identical with a terminal segment of the canonical representative) is even, and odd o/w. Now partition the ω-sequences from K into odd and even. There can be no infinite monochromatic set.

■

(Compare refutation of AC from AD, theorem 87)

What happens if we postulate the existence of a cardinal κ such that $\kappa \to (\kappa)^\omega$ and drop AC? It turns out that κ is measurable.

In virtue of this, the α-closed subsets generate a κ-complete filter "on κ" (as we say, identifying κ with $\{\beta : \beta < \kappa\}$).

THEOREM 144. *If $\kappa \to (\kappa)^\alpha$ then the filter generated by the α-club sets is an ultrafilter.*

Proof:

(In this case the exponent α means we are talking about subsets of κ which are of length α in the restricted ordering)

Let X be a subset of κ. Define $F : [\kappa]^\alpha \to 2$ according to whether or not the sup of the α-sequence belongs to X. There is then a set $A \subseteq \kappa$ s.t. *either* every α-sequence of elements of A has a supremum in X *or* no α-sequence of elements of A has a supremum in X. In the first case X extends an α-club subset of κ and belongs to the filter generated by the α-club subsets, and in the second case the complement of X extends an α-club subset of κ and belongs to the filter generated by the α-club subsets.

■

We know from lemma 62 that this filter is α-complete; we now know it's ultra, so κ is measurable.

Chapter 5

Inner models

An inner model is a definable class—the class of all x such that ϕ (for some ϕ) that is transitive and supercomplete (every subset is a subset of a member) $\mathfrak{M} \subseteq \mathcal{P}(\mathfrak{M}) \subseteq \bigcup_{x \in \mathfrak{M}} \mathcal{P}(x)$ and is also a model of ZF (by which we mean that the relativisation to $\{x : \phi(x)\}$ of any axiom of ZF is a theorem of ZF.) and contains all ordinals.

The earliest example of an inner model is the proper class of all wellfounded sets. This proves the relative consistency of the axiom of foundation. There are other inner models, and time permitting we will consider three of them: HOD, the hereditarily ordinal-definable sets, L the inner model of constructible sets, and Powell's model that shows that Intuitionistic ZF is as strong as classical ZF. (We customarily exclude from the concept of "inner model" the bounded Zermelo cones—V_κ—though these form models of ZF very naturally.) An obvious strategy for getting definable models is to take H_ϕ—the class of things that are hereditarily ϕ, so let's start there.

5.1 Hereditarily ordinal-definable sets

After the use of the inner model of the wellfounded sets to prove the relative consistency of the axiom of foundation the idea was abroard that one might be able to find an inner model that was a model of the axiom of choice (whose status was unclear at that stage).

So how about $H_{\text{wellordered}}$? Altho' theorem **??** tells us that the collection of hereditarily wellordered sets is a proper class there is the problem that the power set of a wellordered set cannot be relied upon to be wellordered, and so this structure cannot be relied upon to be a model of the power set axiom.

We shall see that the method of inner models is of limited application and that in particular, there is no inner model of $\neg$AC. These are intended to prove the *consistency* of AC.

Some of you may have heard me on the subject of AC before: we come to believe it because the situations in which we cannot prove the instances we want are situations in which we have unrealistically little information (example of perfect binary trees). This relative consistency proof works by constructing a model in which we know so much about all sets that AC ends up being true: if everything is definable we can order things according to the formulæ that define them.

A set is ordinal-definable if for some formula ϕ and some list of parameters $\vec{\alpha}$ taken from the ordinals, it is the unique x such that $\phi(x, \vec{\alpha})$. This seems a perfectly reasonable concept in principle, but we have to be sure we can formalise it inside *ZF*. If we can't we can't use it. To do this properly we need to be able to do semantics for *ZF* inside *ZF*, and altho' this can be done the details are messy and time-consuming. If we take it as read that the concept of ordinal-definable really is capturable by a formula in the language of set-theory then the rest is comparatively straightforward. The ordinals are wellordered, so (the class of) finite sequences of ordinals can be wellordered; the set of formulæ is countable, so the set of unique definitions is wellorderable. So the proper class of ordinal-definable sets is a surjective image of a wellorderable class (a set might have more than one definition) and so is wellorderable. *A fortiori* the proper class of hereditarily ordinal-definable sets can be wellordered too. More to the point, the restriction of this wellordering to any given (hereditarily ordinal-definable) set is itself another hereditarily ordinal definable set and so not only does the proper class of hereditarily ordinal definable sets consist entirely of things that have wellorderings, it actually contains those wellorderings too, so it will be a model of the axiom of choice. Checking that it is a model of the other axioms of of *ZF*—at least to the standard of rigour in this paragraph!—is comparatively straightforward.

5.2 Constructible sets

There are two drawbacks to HOD. One it shares with the hereditarily wellordered sets, namely that the class you get depends on the model in which you are working, and the other is the need to provide a definition of satisfaction in a structure. Gödel's model of constructible sets (which he called 'L') has neither of these disadvantages. Study of L is a full-time job for many people but we are going to do only enough to prove the consistency of AC and the Generalised Continuum Hypothesis.

There are two approaches to Gödel's model—both of them due to Gödel himself. One way is to do something like HOD. (HOD is actually a later development). Define a transfinite hierarchy where the operation at limit stages is union as usual and at successor stages one takes the set of all first-order definable subsets. This way one has to do the grubby bit with the truth-definitions.

The other way is to find a finite collection of Δ_0 operations such that any class that is closed under those operations is a model for Δ_0 separation. We want the operations to be Δ_0 so that the construction is predicative in the sense of Russell. So let's go looking for such a finite collection. What we want to prove is that for all A and $\vec{z}$ in L and all $\phi \in \Delta_0$, $\{\langle y_0 \ldots y_n \rangle \in A : \phi(\vec{y}, \vec{z})\}$ exists.

The obvious thing to do is to set out to prove this by induction on ϕ, since the various steps will tell us what operations we need the universe to be closed under in order for the inductive steps to succeed.

Of course the interesting steps are going to be those involving restricted quantifiers, and of course we might want to bind a variable from the middle of the list. If we want a construction that binds the kth variable then we will need infinitely many constructors, so the obvious thing to do is to have a single rule that binds the first variable, and then have rules that take a set A of ordered n-tuples and permute them all simultaneously in the same way, so we can get $\{\langle x, z, y\rangle : \langle x, y, z\rangle \in A\}$ from A. We want to do this with only finitely many operations. So let us start by having—in addition to the rule we have just considered, a rule that says we can get $\{\langle y, z, x\rangle : \langle x, y, z\rangle \in A\}$ from A. This generates S_3 so we can permute triples *ad libitum*. Now let's try to prove by induction on n that we can do this for n-tuples. We will see what machinery we need to make this feasible.

In this development we will trade heavily on the fact that we have defined ordered n-tuples in such a way that the the $n+1$-tuple $\langle x_0 \ldots x_n\rangle$ is the ordered pair $\langle x_0, \langle x_1 \ldots x_n\rangle\rangle$.

Assume true for n, as let us suppose we wish to rearrange the $n + 1$-tuple

$$\langle x_1 \ldots x_{n+1}\rangle \tag{1}$$

into

$$\langle y_1 \ldots y_{n+1}\rangle \tag{2}$$

The $n + 1$-tuple (1) is the same object as the n-tuple

$$\langle x_1 \ldots \langle x_n, x_{n+1}\rangle\rangle \tag{3}$$

By induction hypothesis we can rearrange this to the n-tuple

$$\langle\langle x_n, x_{n+1}\rangle, \ldots y_1, y_2\rangle \tag{4}$$

while at the same time rearranging the tail (which of course contains x_1 up to x_{n-1}) into any order we chose, so we took this opportunity to put—in the last two places—the components that were to be the first two in the rearrangement, namely y_1 and y_2. Notice that the tuple (4) is in fact an ordered pair whose first component is an ordered pair, so if we have an operation that takes $\langle\langle x, y\rangle, z\rangle$ and returns $\langle x, \langle y, z\rangle\rangle$ we can invoke it to obtain the $n + 1$-tuple:

$$\langle x_n, x_{n+1} \ldots y_1, y_2\rangle \tag{5}$$

which as before (see lines 1 and 2) is the same as the n-tuple

$$\langle x_n, x_{n+1} \ldots \langle y_1, y_2\rangle\rangle \tag{6}$$

Now we repeat the trick that took us from line (3) to line (4) to get

$$\langle\langle y_1, y_2\rangle, x_n, x_{n+1} \ldots\rangle \tag{7}$$

and at the same time we can rearrange the tail into the right shape so that when we invoke again the operation that takes $\langle\langle x, y\rangle, z\rangle$ and returns $\langle x, \langle y, z\rangle\rangle$ we get

$$\langle y_1, y_2, \ldots y_n\rangle \tag{8}$$

Conclusion: we need an operation that, on being given A, returns $\{\langle x, \langle y, z\rangle\rangle : \langle\langle x, y\rangle, z\rangle \in A\}$

It turns out that the following functions eventually generate Δ_0-separation. $\mathcal{F}_0(x, y) =: \{x, y\}$; $\mathcal{F}_1(x, y) =: x \setminus y$; $\mathcal{F}_2(x, y) =: x \times y$; $\mathcal{F}_3(x, y) = \{\langle u, z, v\rangle : z \in x \wedge \langle u, v\rangle \in y\}$; $\mathcal{F}_4(x, y) =: \{\langle u, v, z\rangle : z \in x \wedge \langle u, v\rangle \in y\}$; $\mathcal{F}_5(x, y) =: \bigcup x$; $\mathcal{F}_6(x, y) =: \mathrm{dom}(x)$; $\mathcal{F}_7(x, y) =: \in\!\!\!\restriction x$; $\mathcal{F}_8(x, y) =: \{x``\{z\} : z \in y\}$. The curly '$\mathcal{F}$'s are traditional.

We now envisage a transfinite construction where we start with the empty set, and apply the various operations to what we have, every now and then stopping at limit stages to add—as a new item—the set of of all things we have constructed so far. We could describe this by saying something like

$$L_0 =: \emptyset; \quad L_{\alpha+1} =: \text{closure of } L_\alpha \cup \{L_\alpha\} \text{ under the operations,}$$

but that way it is less than blindingly obvious that the result is canonically wellordered. To make it obvious that it is wellordered we have to apply the operations one-by-one in a predetermined order to things as they are generated, so that at each ordinal stage we construct a single new element of L rather than a new subset (as in the declaration i've just written: *prima facie* that would give us a prewellorder rather than a wellorder) That way we can actually establish that the wellordering is Σ_1. (Actually we can see anyway that L must be wellordered, because if X is wellordered so is the closure of X under finitely many finitary operations, and we then do an induction on α to show that L_α is wellordered.)

So the result is a proper class that is closed under these operations (always called 'rudimentary', and pronounced 'rud' as in 'rudder' not as in 'ruder') and therefore a model for Δ_0 comprehension (aka separation). But Δ_0 comprehension isn't a great deal of use. We want **full** comprehension: for example $\{\langle\vec{y}\rangle \in A : (\exists x)\phi(x, \vec{y})\}$. where $\phi \in \Delta_0$. We want there to be a set B in L so large that $(\forall \vec{y} \in A)((\exists x)(\phi(x, \vec{y}) \to (\exists x \in B)(\phi(x, \vec{y}))$. Suppose that, for some A, there were no such B. Then consider the function f that takes $\vec{y} \in A$ and returns $\{\beta \in On : \neg(\exists x \in_\beta)(\phi(x, \vec{y}))\}$ (*i.e.*, that initial segment of the ordinals that fails to supply a witness). Then, by replacement, $f``A$ exists and is an unbounded *set* of ordinals, so $\bigcup f``A$ is a set and is equal to On. So there is a big enuff B and then the set we want is $\{\langle\vec{y}\rangle \in A : (\exists x \in B)\phi(x, \vec{y})\}$, which is now Δ_0. This trick gives us the inductive step we need to induct through arbitrary numbers of quantifiers.

The slow way is called the J hierarchy and the quick way the L hierarchy. The L hierarchy is declared thus:

DEFINITION 145.

> $L_0 =: \emptyset$; $\quad L_{\alpha+1} =: rudclos(L_\alpha \cup \{L_\alpha\})$; $\quad L_\lambda = \bigcup_{\alpha<\lambda} L_\alpha$ *for* λ *limit.*

...where $rudclos(x)$ *is of course the closure of* x *under all the rudimentary operations.*

Let L be the class of all the things we construct in this way. Clearly L has a definable (external) wellordering, because the closure of a wellordered set of size κ under finitely many finitary operations is the same size as the set of finite sequences from that set and will turn out to be of size κ. A little thought will reveal that this wellorder is also definable internally, and any anal-retentive will be able to check that the wellordering is Σ_1: x is before y if there is an ordinal α such that $F(x, y, \alpha)$ where F is Δ_0.

By induction on α we show that L is also transitive. Now we have to show that it is a model of all the axioms of *ZFC*. Do we need a proof of this?

1. Extensionality. Easy, given that L is transitive. Every transitive set X is a model of extensionality, since if x and y are distinct, any witness to this fact must belong to one of x and y and so be in X.

2. Pairing. This is $\mathcal{F}_0$.

3. Comprehension. First we show that, for any ϕ whatever,

$$(\forall \vec{z})(\forall X)(\exists \alpha)(\forall x \in X)[(\exists y)(\phi(x, y, X, \vec{z})) \longleftrightarrow (\exists y \in L_\alpha)(\phi(x, y, X, \vec{z}))]$$

   Fix X and $\vec{z}$ and consider, for each $x \in X$, the first ζ such that there is a $y \in L_\zeta$ such that $\phi(x, y, \vec{z})$. Let's call this ordinal $\zeta(x)$. The key idea is that $\{\zeta(x) : x \in X\}$ must be bounded. If it weren't, X could be mapped surjectively to an unbounded subset of On, which would be a class, and this is impossible, by replacement. So there is a bound of all these $\zeta(x)$. Let's call it $\zeta(X)$. Now the set of all $x \in X$ such that $(\exists y)(\phi(x, y, \vec{z}))$ is the same as the set of all $x \in X$ such that $(\exists y \in L_{\zeta(X)})(\phi(x, y, \vec{z}))$.

4. Sumset. If x is in L_α then so are all its members and all their members, by transitivity of L. So $\bigcup x$ is $\{y \in L_\alpha : (\exists z \in x)(y \in z)\}$ which makes it a subset of L_α defined by a Δ_0 formula;

5. Power set. Fix a set x. There must come some ordinal ζ such that all constructible subsets of x have been constructed by stage ζ, for otherwise the collection of such ζ (which is a *set*) would be cofinal in On, and its sumset (which is On) would also be a set. $\mathcal{P}(x)$ is now $\{y \in L_\zeta : y \subseteq x\}$ which makes it a subset of L_ζ defined by a Δ_0 formula;

6. Collection. This is easy, and is very like the proof of power set.

What we have established is that if A is an axiom of *ZF* and $\phi(x)$ is "x is constructible" then *ZF* proves A^ϕ.

Finally we have to verify that the axiom of choice is true in this structure. Every $x \in L$ corresponds to an ordinal—at least if we have constructed L the slow way, one member at a time—since it appears for the first time at some stage α. The two-place relation "the first stage at which x appears is earlier than the first stage at which y appears" is definable, and—since $L \models$ comprehension, the restriction of it to any set in L must be coded by some set of ordered pairs in L. This means that anything in L has a wellordering in L, so $L \models$ AC.

Notice that

Lemma 146. $|L_\alpha| = |\alpha|$ *(and this bijection is in L).*

Proof: An easy induction on α.

Lemma 147. The Condensation Lemma
If $M \prec L_\alpha$ then there is $\beta < \alpha$ $M \simeq L_\beta$

Sketch of proof: The first weapon we use will of course be a Mostowski collapse, which we shall write π. An essential fact is that π preserves all sufficiently simple formulæ. In particular, if f is one of the rudimentary functions then we have

$$f(\pi(\vec{x})) = \pi(f(\vec{x}))$$

This is because the rud functions are all Δ_0.

Now there is a formula (never mind what it is exactly) that says "x is constructible". In L_α it is the case that everything is constructible, this will be true in M as well. Since π is an isomorphism, it will be true in π"M as well.

The β we want is in fact $sup(On \cap \pi``M)$, (so $\pi``M \subseteq L_\beta$) ■

Because of the uniform wellordering of L by birthdays, each L_α is equipped with a family of definable Skolem functions. Now consider the closure of $L_\kappa \cup \{X\}$ under these Skolem functions. Call the result M. Clearly $|M| = |\alpha|$ and M must be an elementary submodel of L_α. We can then use the condensation lemma to make M iso to an L_β. This technique is very useful in the study of L.

5.2.1 "Diamond"

Definition 148. *$\diamond_\kappa$ says: $\exists$ sequence $\langle s_\nu : \nu < \kappa\rangle$ where $s_\nu \subseteq \nu$ and $\forall A \subseteq \kappa$ $\{\nu < \kappa : A \cap \nu = s_\nu\}$ is stationary in κ.*

$\diamond$ with no subscript is $\diamond_{\omega_1}$. What does this mean? *There is a single sequence of countable sets which almost-correctly predicts* **all** $A \subseteq \omega_1$.

Theorem 149. $V = L \to \diamond$.

Proof:

Assume $V = L$. We will define two transfinite sequences of S's ($S_\nu \subseteq \nu$) and C's (C_ν a clubset of ν) by mutual recursion. $S_0 =: C_0 =: \emptyset$. Thereafter at stage ζ ask:

Is there a pair S, C such that C is club in ζ and $(\forall \nu \in C)(S \cap \nu \neq S_\nu)$? If there is such a pair, grab the $\leq_L$-first such and set them to be S_ζ, C_ζ. o/w (if no such pair then $S_\zeta = C_\zeta = \zeta$).

Now we are going to use $V = L$ to show that $\langle S_\zeta : \zeta < \omega_1 \rangle$ is a diamond sequence.

Suppose not. Then $\exists S, C$ C club $\subseteq \omega_1$, $S \subseteq \omega_1$ and $(\forall \nu \in C)(S \cap \nu \neq S_\nu)$. Let $\{S, C\}$ be the $\leq_L$-first such pair. This whole construction can be described in (say) L_{ω_2}, so $\{S, C\} \in L_{\omega_2}$. Let H be the Skolem hull of L_{ω_2} and collapse to get π"$H = \mathfrak{M}$. $\{S, C\}$ was the $\leq_L$-first pair in L so $\{\pi(S), \pi(C)\}$ must be the $\leq_L$-first pair in $\mathfrak{M}$. (Have to do a bit of work to check that $\pi(S) = S \cap ({\omega_1}^{\mathfrak{M}})$ and $\pi(C) = C \cap ({\omega_1}^{\mathfrak{M}})$). But then at stage $\omega_1^{\mathfrak{M}}$ we should have picked $S \cap ({\omega_1}^{\mathfrak{M}})$ and $C \cap ({\omega_1}^{\mathfrak{M}})$ which we didn't.

■

Notice that $\diamondsuit$ (which is $\diamondsuit_{\omega_1}$) easily implies CH. This is beco's if we let A be a subset of $\mathbb{N}$, then $A \cap \nu = A$ for $\nu \geq \omega$, so every subset of $\mathbb{N}$ is an S_ν.

There is also a generalisation:

Definition 150.
$\diamondsuit_E$ (with E a stationary subset of κ understood) says: $\exists$ sequence $\langle s_\nu : \nu \in E \rangle$ where $s_\nu \subseteq \nu$ and $\forall A \subseteq \kappa$ $\{\nu \in E : A \cap \nu = S_\nu\}$ is stationary in κ.

We then find that $V = L \vdash (\forall E \text{ stationary} \subseteq \omega_1)(\diamondsuit_E)$

> This in turn implies that every Whitehead[1] group is free. (G is Whitehead iff whenever there is $f : H \twoheadrightarrow G$ with kernel $\mathbf{Z}$ then there is $g : G \hookrightarrow H$ splitting f.)
>
> We need the concept of a **pure** subgroup. A subgroup H is pure in G iff for any $h \in H$, and any natural number n, if for some $g \in G$, $ng = h$, then there is an $h' \in H$ such that $nh' = h$. If G and H are elementarily equivalent it actually means H is an elementary submodel of G.
>
> "...Shelah proved that the result is independent of ZFC for groups of power $\aleph_1$. Let G be a group of size $\aleph_1$. Write it as a union of a $\subseteq$-increasing chain of countable subgroups G_i and let E be the set of limit ordinals such that G_δ is not $\aleph_1$ pure in G. Then (ZFC) G is free iff and only if E is not stationary."
>
> see [11]
>
> Some of the flavor is given by the example at around pages 26-29 of http://www2.math.uic.edu/ jbaldwin/aecabgrpbeamgood.pdf
>
> But the Eklof paper is the best place for Whitehead problem and Prest's book is best intro to this kind of abelian group/model theory.

[1]Not the Whitehead of Russell-and-Whitehead and type theory, who is A.N. Whitehead, but Whitehead the algebraist, who is J.H.C. ("Jesus, He's Confusing!") Whitehead.

5.2.2 Souslin

The reals have a countable dense subset. From this it follows that every set of disjoint open intervals is countable. Might there be a converse? In other words, must a dense total order of size (details!!) without an uncountable set of disjoint open intervals be isomorphic to the reals? We say that such a space satisfies the **countable chain condition**. (Spaces are assumed to be Hausdorff) Souslin's hypothesis is that the answer is 'yes'.

We will show that the answer is independent of ZFC. (Well, we'll get at least part of the way)

Two facts i hope to get round to proving:

(i) $\diamond \to \neg SH$;

(ii) $MA + \neg CH \to SH$.

Martin's axiom arose from an attempt to separate the combinatorial content of CH from the cardinal arithmétic content.

Recall that a space is *Baire* iff the intersection of countably many dense open sets is dense. Baire's theorem asserts that every compact space is Baire. CH obviously enables us to strengthen this to "the intersection of fewer than $2^{\aleph_0}$ dense open sets is dense". This might look like a candidate for the combinatorial content of CH, but actually it's equivalent to CH. To get something weaker we have to restrict the spaces for which this assertion is made. Let MA be

> Let $\mathcal{X}$ be a compact Hausdorff space satisfying ccc. Then every intersection of fewer than $2^{\aleph_0}$ dense open sets is dense.

This is not the way ZF-istes usually express it. It can be phrased as a fact about boolean algebras. This is important beco's of the appearance later of boolean valued models for set theory.

First we need a deviant concept of antichain. We say p and q are **incompatible** if there is no $r \leq p$ and $r \leq q$. An **antichain** is now a set of pairwise incompatible elements. (This concept is in play in the idea of a set of forcing conditions). We will say a poset satisfies the **countable chain condition** iff every antichain (in the new sense) is countable.

We can topologise the carrier set of any poset by taking basis sets to be $\{p : p \leq q\}$ for each q in the carrier set. The regular open algebra of this poset is a boolean algebra.

MA is now equivalent to the assertion that if P is a ccc poset and $D \subseteq \mathcal{P}(P)$ with $|D| < 2^{\aleph_0}$. Then there is a "D-generic" filter

definition in Dales and Woodin not consistent

We will now deduce SH from MA + $\neg$CH. Let $\langle P, < \rangle$ be a Souslin line. Let M be a maximal family of pairwise disjoint open intervals. Since P is ccc, M is countable, so $\bigcup M$ is separable. But P is not separable, so $\bigcup M$ is not dense in P, so there is an interval J disjoint from $\bigcup M$. Now $\langle J, < \rangle$ is a Souslin line too, so wlog we could have started with a Souslin line P whose every interval is nonseparable. Let us suppose we did this. So every open subset of P is a union of countably many pairwise disjoint nonseparable open intervals. Let us take these to be its **components**.

We will now construct a sequence $\langle U_\alpha : \alpha < \omega_1 \rangle$ of dense open subsets of P s.t. for all $\alpha < \beta < \omega_1$,

1. $U_\beta \subset U_\alpha$
2. if I is a component of U_α, $I \not\subset U_\beta$

Given U_α, we form $U_{\alpha+1}$ by deleting one point from each component of U_α. Since no component of U_α is separable, the deleted point was not isolated in U_α so $U_{\alpha+1}$ is dense. (It's open beco's it's a union of things that are open-sets-minus-a-singleton—but such things are open) why?

Now for the limit case, λ.

For each $\alpha < \lambda$ we can find a countable S_α such that if $a < b$ belong to different components of U_α then there is $c \in S_\alpha$ with $a < c < b$. Let $S =: \bigcup\{S_\alpha : \alpha < \lambda\}$. Then set U_β to be the union of those intervals (a, b) which for all $\alpha < \lambda$ are included in a component of U_α. why?

To show U_λ is dense, let J be an open interval. J isn't separable so there is an interval $(x, y) \subset J$ which meets S and $(x, y) \subseteq U_\lambda$ so U_λ meets J. why?

For $\alpha < \omega_1$ let Q_α be the family of components of U_α, and Q be the union of all the Q_α. Set $D_\alpha =: \bigcup\{Q_\beta : \alpha < \beta\}$. Now consider the poset of Q under subset. We claim this is ccc, and each D_α is dense in it. By MA there is a filter G in (*sic*) Q which meets every D_α. For each $\beta < \omega_1$ pick $J_\beta \in Q$ so that $G \cap Q_\beta = \{J_\beta\}$. Then for $\beta < \alpha < \omega_1$, J_α is a strict subset of J_β. So for every $\beta < \omega_1$ there is an open interval $U_\beta \subseteq J_\beta \setminus J_{\beta+1}$ which makes the U_βs an uncountable antichain, which is impossible.

Dales and Woodin p 100

Diamond implies not-SH

Definition 151. *A (κ, λ)-tree is a tree with*

> *All levels below κ nonempty but level κ empty;*
> *All lower levels are of size $< \lambda$;*
> *It's* **normal** *if it has only one root; and*
> *every node has at least two children; and*
> *every node has descendents at all later levels; and*
> *chain-complete (i.e., chains have unique minimal upper bounds).*

In this notation König's Lemma says that every (ω, ω) tree has an ω-branch. A (ω_1, ω_1)-tree with no ω_1-branch is said to be an **Aronszajn** *tree.*

A Souslin tree is a normal (ω_1, ω_1)-tree with no uncountable antichain (not just no uncountable levels).

Remark 152. *SH iff there is no Souslin tree.*

Proof:

If there is a Souslin tree then the set of its maximal branches ordered lexicographically is a Souslin line.

For the other direction let $\langle X, \leq\rangle$ be a Souslin line. Let us define a sequence $\langle I_\alpha : \alpha < \omega_1\rangle$ as follows. At stage α consider the set A_α of endpoints of the intervals so far constructed. It's not dense so we can pick I_α whose closure is disjoint from it. The construction ensures that $I_\alpha \supset I_\beta$ implies $\alpha < \beta$ so the I_α form a tree under $\supset$. ■

Remark 153. $\Diamond$ *implies the existence of a Souslin tree.*

Proof: The tree we construct will have the countable ordinals as its carrier set, and the tree ordering will be a subset of the order relation on the countable ordinals. We will also ensure that for every countable ordinal α, $T \restriction \alpha$ is a normal (α, ω_1)-tree.

The construction proceeds by induction. Let's take successor stages first. To each leaf of T_α give the first two ordinals not used so far. That's easy.

The difficulty comes at limit stages, because altho' each level is countable, there will be continuum many branches and we cannot extend them all—yet we have to ensure that every node has descendents at all subsequent levels. This is where we use $\Diamond$. Let $\langle S_\alpha : \alpha < \omega_1\rangle$ be a $\Diamond$-sequence. At stage α we use S_α. If S_α is not a maximal antichain of $T \restriction \alpha$ then for each $x \in \alpha$ pick one of the cofinal branches containing x and give it an ordinal. On the other hand if S_α is a maximal antichain we procede as follows.

Each x in $T\restriction\alpha$ is comparable with a unique member x' of S_α. Let b_x be an arbitrary cofinal branch containing both of these.

It remains to be shown that the tree this constructs is Souslin. This is where we invoke $\Diamond$. Suppose it isn't, and that A is an uncountable maximal antichain. It doesn't take long to check that

$C =: \{\alpha < \omega_1 : lim(\alpha) \wedge T\restriction\alpha = (T \cap \alpha) \wedge (A \cap \alpha$ is a maximal antichain of $T\restriction\alpha)\}$

is club in ω_1. By $\Diamond$ we can now find $\alpha \in C$ s.t. $\alpha \cap A = S_\alpha$. Since α is in C, we have

$\lim(\alpha) \wedge T\restriction\alpha = (T \cap \alpha) \wedge (A \cap \alpha = S_\alpha$ is a maximal antichain of $T\restriction\alpha)$.

Thus by construction every element of T_α lies above some element of S_α. Since A is uncountable, it must have an element x of height $> \alpha$. The path this lies on contains some $y \in T_\alpha$, and we also know that there is $z \in S_\alpha$ below y. But then $z \in A \cap \alpha$ and $x \in A$ is above z, contradicting the assumption that A was an antichain.

5.3 The Negative Interpretation

Rewrite this section completely

5.4 Hereditarily real ordinal definable sets

Shepherdson's theorem (theorem 161) tells us that no inner model is provably a model of $\neg AC$, but it doesn't tell us that every inner model is provably *not* a model of $\neg AC$.

Consider the inner model of sets that are hereditarily definable with parameters from the ordinals **and the reals**. If the reals lack a definable wellordering—and they may lack a definable wellordering even if AC is true—then this class (which we call $HROD$) cannot be proved to be wellordered in the way that HOD can. Inner models of $\neg AC$ can arise in this way. If we are to explore this chink we have to establish at all events that $HROD$ is in fact an inner model (which isn't difficult) but we would also need to establish that there needn't necessarily a definable wellordering of the real line, and that needs forcing, which doesn't belong in this chapter.

Solovay famously used $HROD$ to supply a model in which DC holds and every set of reals is Lebesgue-measurable.

5.5 V not a $\mathcal{P}$-extension of L

In a $\mathcal{P}$-extension not only do old objects not acquire new members they don't even acquire new subsets.

REMARK 154. *V is not a $\mathcal{P}$-extension of L unless $V = L$*

Proof:

With a view to obtaining a contradiction assume $V \neq L$ but that every subset of a constructible set is constructible.

Let x be a set and let y be the set of nonconstructible elements of $TCl(\{x\})$. Suppose y is nonempty, and let z be an element of y of minimal rank. Since $TCl(z)$ contains no nonconstructible sets, z must be a subset of L_α for some α. But L_α is a constructible set, so by the hypothesis it follows that z is constructible. This is a contradiction, and so the transitive closure of $\{x\}$ contains no nonconstructible elements. Hence x is constructible. ■

5.6 Shepherdson's wall

We conclude this section with a negative result of Shepherdson.

THEOREM 155. *Let $\phi(x)$ be any formula in the language of set theory. Then ZF does not prove that the extension of ϕ is a model of ZF + $V \neq L$.*

Proof:

Suppose *per impossibile* that *ZF* proved that the extension of ϕ is a model of *ZF* + $V \neq L$. We had better be precise about what we mean. We might mean that the extension of ϕ is a set and that the structure $\langle \{x : \phi\}, \in \rangle$ is a model of *ZF* + $V \neq L$. What we actually have in mind is something weaker. If A is a formula, recall that the relativisation A^ϕ of A to ϕ is the formula obtained from A by replacing every quantifier $(\forall x) \ldots$" by $(\forall x)(\phi(x) \to \ldots$ and every quantifier $(\exists x) \ldots$" by $(\exists x)(\phi(x) \wedge \ldots$". We mean that if A is an axiom of *ZF* + $V \neq L$ then *ZF* proves A^ϕ.

Even this weaker thing cannot happen. If it did, we would be able to reproduce in $ZF + V = L$ the proofs of the various A^{ϕ} for A a theorem of $ZF + V \neq L$. But what is the extension of ϕ in a model of $V = L$? It's a model of ZF and so must be closed under all the rud functions and contain everything in L: since it is also a subset of L it must be L itself! But it is also a model of $V \neq L$ which is impossible. ■

. . . the point being that there is no way of proving that any of these inner models are distinct from L. This tells us that the method of inner models will never give us a proof of the independence of $V = L$, let alone AC or GCH.

Chapter 6

Forcing

curtaining, cladding, drying, flooring, guttering, ironing, icing, knitting, lagging, roofing, sewing, padding, paving, piping, plumbing, scaffolding, sheeting, stuffing wadding, wainscoting, washing, webbing, wiring ...

Have a look at `http://www-math.mit.edu/~tchow/mathstuff/forcingdum`

Shepherdson's Wall tells us that we cannot use inner models to show the independence of $V = L$, nor *a fortiori* the independence of AC or CH. Let us not worry too much at this stage about these last two, they're too specific: let's think about developing a general method which will establish the independence of $V = L$ in the first instance and then CH and AC later with a bit of tweaking.

Since $V = L$ says that every set has a certain property (constructibility), the natural thing to do is to look for a method that adjoins extra elements to models and hope that it can add extra elements that aren't constructible, and so that the old model is simply the constructible part of the new.

Presumably we want the extension to be an end-extension, and it might be worth saying something about why this is obvious.

On the face of it there are two ways in which you could add novel elements to a widget W and close to get a new widget.

1. The widget might already be a subwidget of a larger widget U, with the novel elements already in U, so all you have to do is consider the subwidget of U generated by W and the novel elements. This construction works properly only when the theory of widgets is algebraic, tho' one can sometimes finesse around the difficulties even when the theory isn't completely algebraic. The construction of field extensions (due i think to Artin?) by means of quotients of rings of polynomials is a nice example.[1]

2. The other way is to add the new objects as virtual objects, rather in the way one can add cardinals and ordinals.

The difference is that in the first case the thing you are trying to extend is an object *in* the universe, and in the second case it *is* the universe.

[1] Indeed Cohen, who invented forcing, used this as an illustration of what he was trying to do. So brush up your Galois theory!

Forcing is rather odd, in that it exhibits features of both these approaches. This is beco's there are two ways of thinking of forcing, and they correspond to the two courses of action i have just outlined. In version (i), the original Cohen presentation, you have a model of ZFC sitting around in your workshop and and you carefully and cleverly add bits of plasticine to it to obtain a new model. Now one might say that approach (i) is not in the spirit of forcing beco's the universe is supposed to be *everything* after all, so what one should really be doing is placing oneself inside the model, reasoning inside the model (so that there is no workshop—il n'y a pas de hors-modèle) and one adds the new object like fictions in the way one can add cardinal or ordinals or relational types to ZF, by interpreting in the old language all the assertions that we want to make in a new language.

Let us briefly review the possible ways of adding stuff to models of ZF to get new models, and see why we need some new ideas.

1. One can fake sets by isomorphism types of apgs, that sort of thing. This certainly enables us to fake illfounded sets, in exactly the way one can fake cardinals and ordinals. However it won't enable us to fake new wellfounded sets because if we can fake them then we can make transitive copies of them by Mostowski collapse, and then they are real and were there to start with.

2. Ultrapowers sound as if they might be a useful way in via approach (i). They add new sets (and this is useful in nonstandard analysis) but the new sets don't have any useful properties. Sadly, nothing like this can work:

   Let $\mathcal{U}$ be a nonprincipal ultrafilter on κ, X a set, and $i : V \hookrightarrow V^\kappa/\mathcal{U}$ be the elementary embedding. Let $V^\kappa/\mathcal{U} \models X' \subseteq X$. Then $\{x \in X : \{k \in \kappa : V^\kappa/\mathcal{U} \models x \in X\}$ is $\{x \in X : \{k \in \kappa : x \in X(k)\} \in \mathcal{U}\}$.

   In other words, $V^\kappa/\mathcal{U} \models$ "$V^\kappa/\mathcal{U} \setminus V$ has no $\in$-least member." What this means is that if we take a substructure of the ultrapower that is an extension of the original model it won't be wellfounded.

Exercise 40. *Let V be a wellfounded model, κ an infinite cardinal and $\mathcal{U}$ a countably incomplete ultrafilter. What is the wellfounded part of the ultraproduct?*

We are going to start off the easy way by taking up the easy challenge of extending a model that is a set of whatever model we are working in, so we *know* there is stuff outside the model, (so we're in case (1) on page 123) and it's just a question of gathering the right moss and twigs to make a new larger model. Once we've seen how to do that we will tackle the much harder task of trying to describe, within the model that is to be extended, the process of adding what—from the point of view of the model—are virtual objects.

So we have a model $\mathfrak{M}$. While we are about it we may as well assume it is countable. (Skolemheim tells us that if there are models of ZF at all then there are countable models. Gödel's theorem tells us that we cannot prove the existence of a model of ZF at all—in ZF itself, so to start down this particular path we will have to work not in ZF but in ZF + Con(ZF).) One of the elements of $\mathfrak{M}$ is a poset $\mathbb{P}$ of **conditions**. $\mathbb{P}$ typically has a top element—1—which is the **empty condition**. No useful $\mathbb{P}$ ever has a bottom element. The lower an element is, the more information it contains. The parallel is with Boolean algebras, where the top element corresponds to **true** and contains no information, and the bottom element—0—contains *too much* information.

Obvious question: where do we get these conditions from? The short answer is that they are always approximations to the object you are trying to add. (Usually when you are doing forcing you are trying to *add* something: in the original instance we were trying to add something that isn't constructible.) Let's take a simple illustration. Suppose we are trying to build a new model in which the old ω_1 has become countable. One wants a bijection between ω and ω_1. Clearly one approximates this bijection by finite partial maps betwen ω and ω_1, and one orders them by *reverse inclusion* since approximants containing more ordered pairs contain more information and come lower. This partial ordering for collapsing ω_1 to ω exhibits lots of incompatible ways of extending partial maps.

> Specifically if f and g are two partial maps (conditions) with $g \not\leq f$ then there is $h \leq g$ such that f and h have no common lower bound.

... the idea being: unless g contains the same info as f and possibly more, there will be h that refines g and is incompatible with f. All realistic sets of conditions have this feature. This captures the intuition that there should be no preferred way of stitching the approximants together.

Set theorists' concept of an antichain.

However the short answer should be flagged by saying that what determines what the new model contains is not what the conditions *are* but how they are *ordered*—if $\mathbb{P}$ and $\mathbb{P}'$ are $\mathfrak{M}$-isomorphic posets you get the same result (in the sense that if π is an isomorphism between $\mathbb{P}$ and $\mathbb{P}'$ and F is a generic filter[2] in $\mathbb{P}$ then $\mathfrak{M}[F]$ and $\mathfrak{M}[\pi``F]$ are the same model).

By starting off armed with a countable model we are putting ourselves in a situation like case (i) above, where the things that we want actually exist (somewhere) and the problem is how to get our hands on them.

The idea is that eventually we will be able to describe the new model entirely from within the old. Every element of the old model will point to, sponsor, or evaluate to, an element of the new model.

Definition 156.

1. *A set $D \subseteq \mathbb{P}$ of conditions is* **dense** *iff for every condition $p \in \mathbb{P}$ there is $q \leq p$ with $q \in D$;*
2. *A* **filter** *is a upward-closed subset $F \subseteq \mathbb{P}$ with $(\forall p, q \in F)(\exists r \in F)(r \leq p \wedge r \leq q)$;*

[2] Be patient: you will be told what a generic filter is soon!

3. *A filter G is* **generic** *if it meets every dense subset.*

4. $a^F =: \{b^F : (\exists p \in F)(\langle b, p\rangle \in a)\}$;

   (So a^F really depends on F, as the notation suggests)

5. *For each generic filter F there is a structure $\mathfrak{M}[F]$ whose carrier set is $\{a^F : a \in M\}$.*

   Clearly $\mathfrak{M}[F]$ is a transitive set, by construction. What sets are in it? Now, for each $a \in M$ there is an $\breve{a}$ such that for all generic F, $\breve{a}^F = a$, and it is defined by $\in$-recursion:

6. $\breve{a} =: \{\langle \breve{b}, 1\rangle : b \in a\}$;

   This of course has the effect that $\mathfrak{M} \subseteq \mathfrak{M}[F]$ for all F. That's not to say that all the $\mathfrak{M}[F]$ are the same. Consider the following object:

7. $g =: \{\langle \breve{p}, p\rangle : p \in \mathbb{P}\}$

We check easily that $g^F = F$, so that $F \in \mathfrak{M}[F]$.

Let's return briefly to the poset of finite partial maps from ω to ω_1 partially ordered by $\supseteq$. Notice that, for any natural number $n \in \omega$, the set of partial maps defined at n is dense, and for any countable ordinal α, the set of partial maps taking α as a value is likewise dense. Now it is easy to check that if F is a generic filter then it is a family of pairwise compatible finite partial maps whose union is defined on the whole of ω and whose range is the whole of ω_1.

(miniexercise)

If F is an $\mathfrak{M}$-generic filter on $\mathbb{P}$ then $\mathfrak{M}[F]$ is not only transitive (as we have seen) but is also a model of all sorts of formulæ true in $\mathfrak{M}$. Specifically if $\mathfrak{M} \models ZFC$ then $\mathfrak{M}[F] \models ZFC$.

Let us check that $\mathfrak{M}[F] \models$ pairing. Let a^F and b^F be two things in $\mathfrak{M}[F]$. If $c = \{\langle \breve{a}, 1\rangle, \langle \breve{b}, 1\rangle\}$ then $\mathfrak{M}[F]$ certainly believes that c is the unordered pair of a and b. Similarly we can check that $(a \cup b)^F = a^F \cup b^F$ so $\mathfrak{M}[F] \models$ axiom of binary union. We haven't used the fact that F is a filter, let alone generic—merely the fact that $1 \in F$. If we want to satisfy ourselves that $\mathfrak{M}[F] \models$ other things we will have to exploit the fact that F is a generic filter. And that needs the truth lemma.

6.0.1 Separative posets

Write '$p||q$' for '$p \not\leq q \not\leq p$'

Let us say a poset is **separative** if $(\forall p)(\forall q)(p \not\leq q \to (\exists p' \leq p)(p'||q))$.

For every poset there is a separative poset that "does the same thing".

check this

Notice that $p \Vdash \breve{q} \in G$ iff $(\forall p' \leq p)(\exists p'' \leq p)(p'' \leq q)$. Now given a (possibly non-separative) poset $\leq$ define $p \leq^* q$ by $(\forall p' \leq p)(\exists p'' \leq p)(p'' \leq q)$. This is of course a quasiorder and might not be antisymmetric. However the quotient over the equivalence relation $p \leq^* q \leq^* p$ *is* antisymmetric—and is separative to boot.

6.1 Internalising Forcing

We want to define a relation $\Vdash$ between conditions and propositions so that

$$p \Vdash \phi \text{ iff } (\forall \text{ generic } F)(p \in F \to \mathfrak{M}[F] \models \phi)$$

The major achievement is going to be reducing '$p \Vdash \phi$' to something that doesn't mention any Fs but talks only about the poset $\mathbb{P}$ and the syntax of ϕ. It turns out that the correct definition is:

Definition 157.

($\Vdash$ for atomics, by recursion on $\in$)

$$p \Vdash a \in b \longleftrightarrow (\forall q \leq p)(\exists \langle c, r\rangle \in b)(\exists s \leq q, r)(s \Vdash a = c)$$

and

$$p \Vdash a = b \longleftrightarrow \bigwedge \begin{pmatrix} (\forall \langle c, r\rangle \in a)(\forall s \leq p, r)(\exists \langle d, t\rangle \in b)(\exists u \leq s, t)(u \Vdash c = d) \\ (\forall \langle c, r\rangle \in b)(\forall s \leq p, r)(\exists \langle d, t\rangle \in a)(\exists u \leq s, t)(u \Vdash c = d) \end{pmatrix}$$

Then for molecular formulæ by recursion on quantifiers and connectives:

- $p \Vdash A \vee B$ *iff* $p \Vdash A \vee p \Vdash B$;
- $p \Vdash A \wedge B$ *iff* $p \Vdash A \wedge p \Vdash B$;
- $p \Vdash (\exists x)\phi(x)$ *iff* $(\exists x)(p \Vdash \phi(x))$;
- $p \Vdash \neg A$ *iff* $(\forall q \geq p)(q \nVdash A)$.

Then we can prove the truth lemma, which makes explicit the connection between the internal relation $\Vdash$ of forcing (internal to the model in which we started, that is) and truth in the models $\mathfrak{M}[F]$ corresponding to the generic filters.

Lemma 158. The Truth Lemma

$$\mathfrak{M}[F] \models \phi \longleftrightarrow (\exists p \in F)(p \Vdash \phi)$$

We will need the following lemma

Say that E is predense below p iff every extension of p is compatible with some $s \in E$.

Lemma 159. *if $E \in V$ and E is predense below $p \in F$ which is V-generic then F meets E.*

Proof: : let D be the set of t which are either incompatible with p or extend some s in E

claim 1 : D is in V

Proof of claim: V models (enough) set theory

claim 2 : D is dense

Let a be in P. if a is incompatible with p OK. else let b be a common extension of a and p. Since b extends p, there is $s \in E$ compatible with b. choose c extending both b and s, then $c \in D$ and c extends a.

claim 3 : E meets F

Proof: Let c be in $D \cap F$. Then c is compatible with p since $p \in F$ a filter, so c must refine some $s \in E$. F is upward-closed so $s \in E \cap F$!

The effect of the truth lemma is that, when we are working in a model $\mathfrak{M}$ and are given a set of conditions $\mathbb{P}$, although $\mathfrak{M}$ contains no generic filters for $\mathbb{P}$, $\mathfrak{M}$ can say a lot about what is true in the models $\mathfrak{M}[F]$ corresponding to those filters. In general of course we should not expect $\mathfrak{M}$ to be able to say anything (we should not expect to be able to code in $\mathfrak{M}$ anything) *at all* about an arbitrary extension of $\mathfrak{M}$. It's the extensions arising from generic filters that we can describe quite well.

As Shoenfield says, this means that, for each molecular formula ϕ with n free variables, $p \Vdash \phi(x_1 \ldots x_n)$ is an $n+1$-place relation.

ccc: every dense set of forcing conditions has a countable dense subset. (Every antichain is minimal dense)

6.1.1 Boolean-valued models

Somewhere in LIS I set an exercise that says that if—when building truth-tables—you try to use a lattice of truth-values—other than the usual two-element boolean algebra $\mathbb{2}$—then the set of formulæ recognised as valid by the lattice is the set of propositional tautologies iff the lattice is a boolean algebra.

Easy important fact, makes a nice exercise.

Think now of the recursive definition of the V_αs as saying that $V_{\alpha+1}$ is $V_\alpha \to \mathbb{2}$ rather than $\mathcal{P}(V_\alpha)$. Then it's natural to wonder about the variant construction where we take $V_{\alpha+1}$ to be $V_\alpha \to \mathbb{B}$ for some boolean algebra $\mathbb{B}$.

We can do this, and we obtain a structure that we might call $V^{\mathbb{B}}$. There is an obvious notion of hereditarily two-valued element of $V^{\mathbb{B}}$. It's probably pretty obvious that the hereditarily two-valued functions will be an isomorphic copy of the V we started with. However any ultrafilter in $\mathbb{B}$ corresponds to a boolean homomorphism $\mathbb{B} \twoheadrightarrow \mathbb{2}$ and so will give rise to a two-valued quotient...

Is there any mileage in this? There is a simple line of thought that says that there shouldn't be. After all, we can describe all this in the model we start in, and so all the hereditarily two-valued chaps—and all the clever quotient objects over the ultrafilter already existed in the ground model. Of course this depends on the ultrafilter in $\mathbb{B}$ being a real object of the model. It's such a good idea there ought to be a way of saving it!

Well, there is. A **regular open** set is one that is the interior of its closure. ("No cracks or pinholes!"!) Dormant fact that you probably all knew but didn't care about: it's trivial to check that in any topological space the collection of regular open sets forms a boolean algebra. (Actually a *complete* boolean algebra, and this will matter). Another fact from off the shelf is that any poset has an order topology .... Poset? Ah! What we wanted, when trying to ensure that the two-valued quotient of the $V^{\mathbb{B}}$ was not trivial, was an ultrafilter that wasn't in the model we started with. If we take $\mathbb{B}$ to be the regular open algebra of the poset $\mathbb{P}$, then perhaps generic filters $\subseteq \mathbb{P}$ will correspond to suitably shadowy ultrafilters in $\mathbb{B}$!

Chapter 7

Positive Set Theory

There are three ideas, and i don't yet see how they fit together. From an expository point of view it might be a good idea to say what they are, and invite the reader to join in the author's voyage of discovery!

1. There is the idea that if we restrict the naïve comprehension scheme to formulæ without negation in them then we should evade paradox.

2. Take the di Giorgi point of departure for set theory: $i : A \hookrightarrow \mathcal{P}(A)$. Topological set theory is the kind of thing that happens if i"A is the set of closed sets of a topology. That is to say: if i"A is closed under finite unions and arbitrary intersections then we say we have a *topological universe*.

3. Finally there is Malitz's explicit construction of a model.

We will have to explain how these ideas come to have a common core.

7.1 Three ideas

7.1.1 Positive formulæ and syntax

Definition 160.

1. *The class of* **Almost Generalised Positive** *formulæ is inductively defined as atomics closed under* $\wedge$, $\vee$, *universal quantification, bounded quantification, and a formation rule:*

   > *If F is GPF and G has only 'x' free, then $(\forall x)(G(x) \to F)$ is also GPF.*

2. *The* **Generalised Positive** *formulæ are additionally closed under* $\exists$.

Weydert proves that

1. Any topological Universe is a model of Extensionality + comprehension for almost generalised positive formulæ;

2. Any *compact* topological Universe is a model of Extensionality + comprehension for generalised positive formulæ.

7.1.2 Topological di Giorgi models

In these circumstances we have a Galois connection. Partially order A by: $a_1 \leq a_2$ iff $i(a_1) \subseteq i(a_2)$. Then the pair of i and $\lambda X \subseteq \mathcal{P}(A).i^{-1}(\overline{X})$ form a Galois connection. (There is a good explanation of Galois connection in Wikipædia.)

(Malitz tells me that) every set is approximated by a class of wellfounded sets in the sese that $\forall x$ there is a class $X \subseteq WF$ with $x = \bigcap X$.

7.1.3 Malitz's construction

This is the construction of Malitz [1976]. (For him it was the point of departure for positive set theory: the construction comes before the syntactic characterisation in terms of positive formulæ.)

Definition 161.

$\sim_0 = V \times V$;
$\sim_{n+1} = \sim_n^+$;
$\sim_\lambda = \bigcap_{\beta<\lambda} \sim_\beta$ *at limits.*
where $\sim^+$ *is defined by*

$$X \sim^+ Y \text{ iff } (\forall x \in X)(\exists y \in Y)(x \sim y) \wedge (\forall y \in Y)(\exists x \in X)(x \sim y)$$

taking intersections at limits.

We iterate up to κ but will say nothing at this stage about what conditions we want κ to satisfy. We can topologize V by taking the basic closed sets to be the equivalence classes. No reason to expect the space to be complete so we complete it by taking sequences: specifically nested κ-sequences whose αth elements are α-equivalence classes. Notice that we can pick representatives in a canonical way as long as we have foundation, beco's you take the union of the things of minimal rank. However (and this is in Malitz) if κ is weakly compact then the space is Cauchy-complete.

The universe of our model of set theory will be the completed space. (Say something about how to complete metric spaces)

What is the membership relation? We say f is a member of g iff $(\forall\alpha)(\exists x \in f(\alpha))(\exists y \in g(\alpha))(x \in y)$. We can ask for this to happen at all α beco's the x and y that make it true at α also make it true at all earlier ordinals.

Think about how extensionality might be provable. Suppose f and g are distinct sequences. Then for all sufficiently late arguments there is a witness to the symmetric difference of f and g. Lots of witnesses. Pick those of minimal rank and they form an oligobranching tree. This is the point at which we have to impose conditions on κ. If κ is weakly compact it will have the tree property so we can pick a κ-path through this tree and this is the witness we want.

Topologise the sequences in the obvious way: for each f, the set of things agreeing with f on some initial segment is an open ball.

The "sets" of the model will be the closed subsets in this topology, beco's each Cauchy sequence codes a closed set and each closed set is coded by a Cauchy sequence.

7.1.4 he Vietoris topology

DEFINITION 162. *The* **Vietoris topology**

$V(X)$ is a topology on the set of closed sets of a pre-existing topology X. For each point C of $V(X)$, the set of (X-)closed subsets of C is a basic closed subset of $V(X)$.

Since every set is coded by a closed subset it is natural to ask if the coding map is a homeomorphism between the obvious topology and the Vietoris topology. We can at least show that every closed subset of the κ-Vietoris topology is a set. Converse is not clear: it's true iff the Vietoris topology is κ-compact.

Notice that this means that the intersection of an arbitrary class is a set! So V really is a complete lower semilattice. In particular the transitive closure of any proper class is a set. Also if $X = \bigcap\{x : \mathcal{P}(x) \subseteq x\}$ then X is a set. This looks as if X ought to be an inner model of something (though presumably not P itself since the relativisation of a GPF with a bounded universal quantifier is liable not to be GPF). X will be a model in which there is no set other than V extending its power set. Thus in particular, in this model if $x \in x$ then $\{x\}$ lacks a complement, for we can show in EST that if $x \in x$ then $V \setminus \{x\}$ extends its own power set and is therefore the universe. Also we have $\in$-induction in this inner model for whatever formulæ we have comprehension for!

Remember that "$x = t$" for t a GPF set abstract is GPF as long as t is CLOSED.

A permutation model of GPC is a model of GPC.

Is any of this a model for the λ-calculus?

A good place to start would be the construction over V_ω.

Randall is doing a good job of explaining this to me.

Vexing questions:

Are the elements of an isolated set isolated? (Converse is true)

Is the graph of the power set function a set?

Is the product topology on $V \times V$ the same as the restriction of the topology to $V \times V$?

Is the closure of the wellfounded sets a transitive model of GPC?

A Conversation with Isaac

Isaac wants his sets to obey strong extensionality. His strong extensionality is defined by a game: $G_{x=y}$. Player $\neq$ picks a member x' of x or a member y' of y. Player $=$ replies with a member y' or x' of the set that $\neq$ did *not* pick from (and loses if she can't. They then play $G_{x'=y'}$. $=$ wins if the game goes on for ever. Strong extensionality says that $x = y$ iff $=$ has a winning strategy.

This means that if $X \in X$ then $X \setminus \{X\}$ cannot be a set, because $=$ has a winning strategy in $G_{X=X\setminus\{X\}}$. ($\neq$ has to pick X from X, so $=$ picks $X\setminus\{X\}$ from $X \setminus \{X\}$ and they are back where they started.)

But that means that $V \setminus \{X\}$ cannot be a set, beco's if it were then $X \cap (V \setminus \{X\}) = X \setminus \{X\}$ would be a set too. So $V \setminus \{x\}$ can be a set only if $x \notin x$.

But we knew (something like) this anyway. If complements of singletons were reliably sets, then $\{V\setminus\{x\} : x \notin X\}$ would be a proper class and $\bigcap\{V\setminus\{x\} : x \notin X\}$ (which would then be a set, since arbitrary intersections of proper classes are sets) would be X. But X was arbitrary.

Chapter 8

Set Theory with Antifoundation axioms

Abstract:

Mostowski collapse → Antifoundation axioms → permutation models → virtual entities, featuring equivalence classes as virtual entities. Then virtual sets of equivalence classes and Hinnion's + operation. Virtual cardinal arithmetic. Uniterated virtual ordinal arithmetic. (Here we treat ordinal arithmetic as the theory of relations for which orderisomorphism is a congruence relation. Ordinals as things you index constructions by appeared earlier). Then iterated virtuality: Paris-Harrington and Burali-Forti. Virtual arithmetic of various kinds of extensional relations. The "arithmetic" of these relations turns out to be set theory with antifoundation axioms.

8.1 Mostowski collapse

There are two spurs to developing Set theory without the axiom of foundation. One is the desire to retain (from naïve set theory) large collections like the collection of all ordinals or all cardinals or all sets, as *sets*. The existence of sets like these contradicts foundation with varying degrees of obviousness. But there is another source of axioms contradicting the axiom scheme of foundation. That is Mostowski's collapse lemma.

Lemma 163.

1. *If $\langle X, R\rangle$ is a wellfounded extensional structure then there is a* **unique** *transitive set Y and a unique isomorphism between $\langle X, R\rangle$ and $\langle Y, \in\rangle$.*
2. *If $\langle X, R\rangle$ is a wellfounded structure then there is a transitive set Y and a homomorphism $f{:}\langle X, R\rangle \to \langle Y, \in\rangle$.*

The requisite bijection exists because of the recursion theorem. It needs the axiom of replacement, but the function one needs has so few quantifiers, and the result is so central to the study of wellfounded sets, that various minimalist systems of set theory are formulated with precisely enough replacement to prove this result (as well as transitive closures, as discussed already!). Typically one invokes only the restriction to Π_1 formulæ. For example, one needs it in the standard proof that the continuum hypothesis is true in Gödel's model L and—of more relevance to us here—plays a central rôle in the theory of measurable cardinals.

Beyond that Mostowski's collapse lemma has a very distinctive character which makes it give rise to new set-theoretic axioms. It asserts that certain kinds of binary structures have $\in$-copies. Such assertions are likely to be useful in proving things about sets (as opposed merely to proving things about implementations of things other than sets, likes reals or ordinals).

New axioms for set theory always arise from loose ends. We saw in the discussion round p. 90 how various axioms of infinity arise from assertions that there are uncountable cardinals which have some of the properties of $\aleph_0$. The point of departure for antifoundation axioms is Mostowski collapse: strengthen it by weakening the wellfounded-and-extensional condition in the antecedent, or weaken it by weakening 'iso' to something else. (A bit like trying to get yourself into the appropriate gear on a mountain bike by tweaking both the gear on the crown wheel and on the rear wheel)

Two brief points:

(i) Notice that the proof of Mostowski relies on a definition by recursion, and the legitimacy of this definition relies on wellfoundedness of the relation in question. If it isn't wellfounded we simply don't know how to begin. Indeed, if we are assuming the axiom of foundation no deviant version of Mostowski which lacks the wellfoundedness condition can possibly be consistent. If we are going to find generalisations of Mostowski we will have to repudiate the axiom of foundation.

(ii) Even if $\in$ isn't wellfounded, it is extensional, so if we drop extensionality we will have to settle for a homomorphism instead of an isomorphism.

So we should be looking for something like: every extensional relation is iso to $\in$ on a (unique?) transitive set, or every binary relation has a (unique) homomorphism onto a transitive set.

There was for a while a little cottage industry of people spinning generalisations like this. One difficulty with "every extensional relation is iso to $\in$ on a unique transitive set" is that it tells us that there is a set $x = \{x\}$, and also that there are y and z such that $z = \{y\} \wedge y = \{z\}$. However it doesn't tell us that x, y and z are all the same set, and people have different views about whether or not all these things should be the same. Not being a Platonist i do not believe there is any fact of the matter here to have views about and accordingly i do not intend to express any. At all events, at this point one has to decide whether or not to go for a kind of **strong extensionality**: two sets are the same if they cannot be told apart by stepping downwards using $\in$. If you want that axiom, then you cannot say that every extensional binary relation is isomorphic to $\in$ on a transitive set, beco's of x, y and z above. You have to say something like: every extensional binary relation has a unique homomorphism onto $\in$ on a transitive set. In fact, once you've weakened the consequent to allow for a homomorphism not an isomorphism, you can strengthen the antecedent by dropping 'extensional' since every binary relation has an extensional quotient. Miniexercise: use one of the fixed-point theorems from Part II to prove that every binary relation has a (largest) extensional quotient.

This sketch is intended to prepare you for the Forti-Honsell Antifoundation Axiom. This is a combination of two ideas, as above: (i) The idea that every extensional relation has a homomorphism onto a transitive set, (ii) A strong Extensionality component to ensure that the homomorphism is onto a **unique** transitive set.

DEFINITION 164. *Forti-Honsell Antifoundation.*

For every X and f there is a unique g making the following diagram commute:

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d, "g"'] & \mathcal{P}(X) \arrow[d, "j(g)"] \\ V \arrow[r, "id"'] & V \end{tikzcd}$$

Try $f = \iota$. Then the unique g satisfies

$$g(x) = g``\iota(x) = \{g(x)\}$$

so for all x, $g(x) = \{g(x)\}$. (Such an object is a **Quine atom**.) Thus Forti-Honsell antifoundation certainly contradicts foundation!

Another kind of reflection

There is a newer and different use of the word 'reflection'. Suppose we have a set theory T, and we have a notion of widget (be it irredundant trees, labelled extensional relations or whatever). The notion of widget and widget-embedding in T gives rise to an interpretation of a set theory in T. This is something rather like the idea of **the arithmetic of** T: which is what we get when widgets are finite wellorderings with the obvious embedding. (People say things like "the consistency of Zermelo set theory is provable in the arithmetic of *ZF*" for example.) We could, by a daring neologism, call the result of doing this to irredundant trees, labelled extensional relations or whatever, the **set theory of** T. This brings us to the question: is the set theory of T the same as T itself? If it is, we say that T **reflects** itself.[1] Holmes (see *op. cit.*) has a consistent strengthing of *NF*U (*NF* with urelemente) which reflects itself.

[1]This is not a numbered definition because this notation is not in any sense standard—yet!

Does this matter? I don't know yet. If we think of the arithmetic of T as what T believes about numbers, then the set theory of T should be what T thinks about sets. If T is a set theory then what it *says* about sets should be what it *thinks* about sets, so perhaps this kind of reflection will turn out to be a Good Thing.

Chapter 9

NF and NFU

I said earlier that set theory is the first-order theory of one extensional binary relation. Simple type theory is the nth order theory of equality. (Contrast with Church's type theory).

Theorem 165. Specker [1962]. *Given a model* $\mathfrak{M}$ *of* TST *plus full ambiguity, there is* $\mathfrak{M}'$ *elementarily equivalent to* $\mathfrak{M}$ *with a tsau* σ*.*

Proof:

Suppose typical ambiguity is consistent. Using theorem 24 on the existence of saturated models there will be a saturated typically ambiguous model, and it will have a tsau. ■

Corollary 166. Specker's equiconsistency lemma.

NF is equiconsistent with TST *plus full ambiguity.*

Extensionality and stratified comprehension. Stratified formulæ are a way of avoiding the paradoxes. I shall have more to say about the genesis of typing in the next section.

If $\mathfrak{M} \models$ TST, with a tsau σ, we can construct a new structure with carrier set the set of things in M of type 0, equality in the sense of $\mathfrak{M}$, and $x \in y$ iff $\mathfrak{M} \models x \in \sigma(y)$. The proof that this resulting structure is a model of *NF* is a simple exercise along the lines of lemma **??** in section **??**, and will be left as an exercise for the reader.

This tells us that if *TNT* plus complete ambiguity is consistent, then so is *NF*. The converse is even easier, since if $\mathfrak{M}$ is a model of *NF* we obtain a model of *TNT* plus complete ambiguity by making Z copies of it ($\langle M \times \{z\} : z \in \mathrm{Z}\rangle$) and saying $\langle x, n\rangle$ is "in" $\langle y, n+1\rangle$ iff $\mathfrak{M} \models x \in y$.

9.0.1 NFU

Theorem 167. (Jensen [1969]). *NFU is consistent.*

Proof: (Boffa–Jensen.)

Let $\mathfrak{M} = \langle \mathfrak{M}_i : i \in \mathbb{N} \rangle$ be a model of TST. For $I \subseteq \mathbb{N}$, let the *extracted model* $\mathfrak{M}_I$ be $\langle \mathfrak{M}_i : i \in I \rangle$ with a new $\in$ relation. We say $x_{i_n} \in y_{i_{n+1}}$ iff y is a set of singletons$^{i_{n+1}-i_n-1}$ (otherwise y is an *urelement*)[1] and $\iota^{i_{n+1}-i_n-1}(x)$ is a member of y in the sense of M. We check that $\mathfrak{M}_I$ is a model of *TSTU*, that is TST with *urelemente*.

This extracted model construction is the same as the construction of the first lecture which gives a model of ZF with urelemente.

Now let $\mathfrak{M} \models$ *TSTU*, and let Φ be an arbitrary expression in $\mathcal{L}_{\text{TST}}$. Φ speaks of, say, five types. Let us partition $[\mathbb{N}]^5$. Let $I = \{i_1, \ldots, i_5\}$ and send $\{i_1, \ldots, i_5\}$ to 1 if $\mathfrak{M}_I \models \Phi$ and to 0 otherwise. We now invoke Ramsey's theorem to find an infinite $J \subseteq \mathbb{N}$ monochromatic for this partition and consider $\mathfrak{M}_J$. By monochromaticity, either every model extracted from $\mathfrak{M}_J \models \Phi$ or every model extracted from $\mathfrak{M}_J \models \neg\Phi$.

Notice that if every model extracted from $\mathfrak{M}_J$ satisfies ϕ it is also true that every model extracted from $\mathfrak{M}_J$ satisfies ϕ^+. (this is beco's, for any model, it satisifies ϕ^+ iff the result of chopping off its bottom type (which is a special case of extraction) satisfies ϕ). So certainly every model extracted from $\mathfrak{M}_J$ satisfies $\Phi \longleftrightarrow \Phi^+$.

So far so good. We are going to iterate, to achieve the same end for a second formula. However, before we do, we must just check something elementary but important. *Extraction is transitive* Or do i mean that in a system of extracted models the extractions commute? The point is that something obtained by two consecutive extractions can be obtained by a single extraction: if i start with a model $\mathfrak{M}_1$ and extract a model $\mathfrak{M}_2$ from $\mathfrak{M}_1$ and then extract further a model $\mathfrak{M}_3$ from $\mathfrak{M}_2$, then i could have extracted $\mathfrak{M}_3$ from $\mathfrak{M}_1$ in one go. The reason why this matters is because any model we now extract while processing a second ϕ will still satisfy ambiguity for the original ϕ.

We now repeat the process for a different Φ, this time starting with $\mathfrak{M}_J$.

This shows that for any finite collection of formulae $\langle \Phi_i : i \in I \rangle$, we can find a model of $TSTU + \bigwedge_{i \in I} \Phi_i \longleftrightarrow \Phi_i^+$. By compactness, we have a model of *TNTU* plus complete ambiguity.

By use of ultraproducts (theorem 24) we can take this model to be saturated. Then we can use a back-and-forth construction to show that this saturated model $\mathfrak{M}$ has a shifting automorphism σ. The structure $\langle (T_0)^{\mathfrak{M}}, \in \cdot\sigma \rangle$ is now a model of NFU. ■

[1]The superscript is the number of times that ι is to be iterated.

Chapter 10

Mostowski and Shoenfield absoluteness

Reconstructed from a handout of Martin Hyland's. Not lectured and not examinable

The Universal Π^1_1 tree and the absoluteness lemma.

Theorem 168.

1. *(Mostowski)*

   Let $\mathfrak{M}$ be a transitive model of ZF. (KP will probably do: we need enuff replacement to be able to prove the collapsing lemma). Let $P(\alpha)$ be Π^1_1. Then

   $$\mathfrak{M} \models P(\alpha) \text{ iff } P(\alpha)$$

2. *(Shoenfield) Ditto for Σ^2_1 as long as every actually countably ordinal is in $\mathfrak{M}$.*

We need the fact that any poset $\langle X, \leq \rangle$ which $\mathfrak{M}$ believes to be a wellorder is in fact a wellorder. By the Mostowski collapse lemma, every wellfounded poset has a homomorphism onto an initial segment of the von Neumann ordinals. If $\mathfrak{M}$ believes $\langle X, \leq \rangle$ to be wellfounded, there will be such a rank function in $\mathfrak{M}$. The von Neumann ordinals in $\mathfrak{M}$ are the same as von Neumann ordinals in the real world, since being a von Neumann ordinal is absolute. Nothing can prevent this rank function from remaining a rank function in the real world. All the real world can do is add subsets of X and the rank function doesn't look at any subsets that aren't preimages of singletons, and they are all already in $\mathfrak{M}$.

Proof of (i):

WLOG we can take P to be of the form $(\forall \beta)(\exists n)R(\alpha(n), \beta(n))$, where we are thinking of α and β as functions $\mathbb{N} \to \mathbb{N}$. This involves quantifier pushing and squashing. We now need the notion of tree given by R and α. $T\alpha$[1] is

$$\{s \in \mathbb{N}^{<\omega} : (\forall n \leq \texttt{len}(s))(\neg R(\alpha(n), s(n)))\}$$

We partially order $T\alpha$ by $s \prec t$ if s is an end-extension of t.

Evidently $T\alpha$ cannot have an infinite path, since such an infinite path would be a counterexample to the '$\forall\beta$'. By DC we infer from the lack of an infinite path to wellfoundedness in the sense required for the existence of a rank function.

This establishes that if $(\forall\beta)(\exists n)R(\alpha(n), \beta(n))$ holds in $\mathfrak{M}$ then it holds in V. The other direction is easy, because universal sentences generalise downwards and the matrix is Δ_0 in the sense of Levy. (Here again we exploit the fact that von Neumann naturals are absolute.)

Proof of (ii):

$P(\alpha) \longleftrightarrow \exists\beta\forall\gamma\exists n R(\alpha(n), \beta(n), \gamma(n))$

iff $\exists\beta$ $T\alpha\beta$ is wellfounded;

iff $(\exists\beta)(\exists f)(f : T\alpha\beta \to On$ f a homomorphism);

now let Ω be an uncountable ordinal ...

[1]It should really have an 'R' as a parameter but for some reason it doesn't...

Appendices

Appendix 1: Monadic Logic

Suppose we have a language $\mathcal{L}$ with various predicate and function letters. Let ϕ be a formula of $\mathcal{L}$. To ϕ we associate a **graph** G_ϕ whose vertex set is the set of variables appearing in ϕ. We join two vertices by an edge if there is an atomic subformula of ϕ that they both appear in. We say two variables are **connected** if there is some connected subgraph of G_ϕ to which they both belong. We say a formula is **elementary** if all variables in it are connected.

We now have the following theorem of classical predicate logic

THEOREM 169. *Every wff Φ is equivalent to a boolean combination of elementary formulæ.*

Proof:

We can manipulate Φ into a form where the only variables occurring within the scope of a quantifier 'Qx' are those connected to 'x'. (In fact not only are they connected to 'x' but the graph of variables occurring within its scope is connected.) We do this by induction on the quantifier structure, working from the inside out. Let us say that a quantifier Qx is *bad* if there are variables in its scope that are not connected to 'x'. We concern ourselves with the innermost quantifier which is bad. By the usual devices we can assume that everything within the scope of this quantifier has been put into disjunctive normal form, where we can think of quantified elementary formulæ as propositional letters (if they are closed) or as atomic predicate formulæ (otherwise).

Suppose are looking at

$$(\forall x)((F(x) \wedge p) \vee (G(x) \wedge q) \vee r)$$

where 'x' does not occur free in p, q or r, but where other variables may occur free in '$F(x)$'. The induction hypothesis is that all of p, q, r, $F(x)$ and $G(x)$ are elementary. (There will be free variables around because we are considering the innermost bad quantifier). This formula is of course equivalent to

$$(\forall x)((F(x) \wedge p) \vee (G(x) \wedge q)) \vee r$$

This distributes to

$$(\forall x)((F(x) \vee G(x) \wedge (F(x) \vee q) \wedge (p \vee G(x) \wedge (p \vee q)) \vee r$$

Now we can pull out $(p \vee q)$

$$((\forall x)((F(x) \vee G(x) \wedge (F(x) \vee q) \wedge (p \vee G(x)) \wedge (p \vee q)) \vee r$$

Now we can pull the $\forall$ inside the $\wedge$ to get

$$((\forall x)(F(x) \vee G(x) \wedge (\forall x)(F(x) \vee q) \wedge (\forall x)(p \vee G(x)) \wedge (p \vee q)) \vee r$$

and we can pull the $\forall$ past the $\vee$ to get

$$((\forall x)(F(x) \vee G(x) \wedge ((\forall x F(x) \vee q) \wedge (p \vee ((\forall x)G(x))) \wedge (p \vee q)) \vee r$$

and now the only things occurring within the scope of a '$\forall x$' are atomic formulæ containing 'x'. *id est* the only variables in the scope of '$\forall x$' are those connected to 'x'.

■

Note what we are using:
A normal form theorem.

$$(\forall x)(F(x) \vee p) \longleftrightarrow ((\forall x)(F(x) \vee p)$$
$$(\forall x)(F(x) \wedge p) \longleftrightarrow ((\forall x)(F(x)) \wedge p)$$

$$(\exists x)(F(x) \vee p) \longleftrightarrow ((\exists x)(F(x) \vee p)$$
$$(\exists x)(F(x) \wedge p) \longleftrightarrow ((\exists x)(F(x) \wedge p)$$

...and probably a lot more besides!
This gives us another proof of remark 13.

Corollary 170. *Classical monadic predicate logic is decidable.*

Proof: We can assume that the only quantifer we use is '$\forall$' say. Take a closed monadic formula Φ and import all the quantifiers. Now the only variable within the scope of 'Qx' is 'x' itself: since all predicate letters are monadic, no variable can be connected to 'x'! By distributing '$\forall$' over '$\wedge$' any subformula of Φ that begins with a quantifier can be taken to be $(\forall x) \bigvee_{i \in I} \phi_i(x)$ where ϕ_i is atomic or negatomic. It is easy to see that we can replace this formula by a disjunction of literals, one for each ϕ_i and thereby the problem reduces to one in propositional logic in the sense that Φ is valid *(resp.* satisfiable) iff the resulting propositional formula is valid *(resp.* satisfiable). ■

Bibliography

[1] P. Aczel *Non-well-founded sets*, Lecture Notes, Number 14, (Center for the Study of Language and Information, 1988).

[2] Jon Barwise and Laurence Moss *Vicious Circles* CLSI. 1996. ISBN.1-57586-009-0 (paper) 202pp. The book is much better than its title.

[3] J. L. Bell and A. Slomson. Models and Ultraproducts. North Holland 1969. Recently reprinted by Dover.

[4] Conway Berlekamp and Guy, Winning ways 2 vols Academic Preess 1982

[5] Craig and Vaught, Finite axiomatizability using additional predicates, JSL, **23**, pp. 289–308.

[6] Makkai. Review of [25] and [5] in JSL **36** (1971), pp. 334-335.

[7] Dales and Woodin. An introduction to independence for analysts. LMS lecture notes **115**.

[8] J. W. Dauben. Georg Cantor: His mathematics and the Philosophy of the Infinite. Princeton 1990

[9] A.J. Dodd. The Core model. LMS lecture notes **61**.

[10] Drake and Singh *Intermediate Set Theory* Wiley 1996 is a paperback that may be of use.

[11] Paul Eklof. Applications to algebra. *in* Model-theoretic logics, Barwise, J., Feferman, S. (eds.), Perspectives in Mathematical Logic, New York etc., Springer-Verlag, 1985, pp. 423-441.

[12] Paul C. Eklof Whitehead's Problem is Undecidable The American Mathematical Monthly, Vol. 83, No. 10. (Dec., 1976), pp. 775-788.)

[13] Erdős Hajnal and Milner. Acta Math. Hung **19** (1968) pp208–218

[14] Forster, T.E. Set Theory with a Universal Set. OUP 1994

[15] Forster, T. E. Logic, Induction and Sets, London Mathemtical Society Undergraduate Texts in Mathematics **56** Cambridge University Press.

[16] Forster, T. E. Canterbury notes on logic. Linked from my teaching page.

[17] Forti, M. and Honsell, F., "Set theory with free construction principles", Annali della Scuola Normale Superiore di Pisa, Scienze fisiche e matematiche 10 (1983), pp. 493-522.

[18] Forti, M. and Hinnion, R. [1989] The consistency problem for positive comprehension principles. *Journal of Symbolic Logic* **54** pp. 1401−18.

[19] Hajnal/Hamburger "Set Theory', LMS Undergraduate Texts in Mathematics **48** Cambridge University Press.

[20] M. Randall Holmes, *Elementary Set Theory with a Universal Set*, volume 10 of the Cahiers du Centre de logique, Academia, Louvain-la-Neuve (Belgium), 241 pages, ISBN 2-87209-488-1.

The book is the volume 10 of the series *Cahiers du Centre de Logique* edited by the Centre de Logique of the Universite de Louvain at Louvain-la-Neuve. The price of the book is 1150 BF. A list of known errata is maintained at http://math.idbsu.edu/ holmes/errata.txt

It can be obtained by ordering from the publisher:

BRUYLANT-ACADEMIA

Grand'Place, 29

1348 Louvain-la-Neuve

Belgium

Fax (0)10-45.44.80 email: academia.bruylant@skynet.be

[21] Paul Howard, Łoś's theorem and the Boolean Prime Ideal Theorem imply the Axiom of Choice. Proc AMS **49** (june 1975) pp 426-8.

[22] A.C.J Hurkens. The relation $\mathcal{P}(x \cap y) \subseteq y$ is wellfounded. Collegium Logicum, Annals of the Kurt Gödel Soc. **4** (1999) pp 244−6.

[23] Notes on Set Theory CUP.

[24] Solovay Reinhard and Kanamori. Strong Axioms of infinity and elementary embeddings. Annals of Mathematical Logic **13** feb 1978 pp 73–116.

[25] S.C. Kleene, Finite Axiomatizability of theories in the predicate calculus using additional predicate symbols, Memoirs of the AMS, **10**

[26] Ken Kunen: Set theory. North Holland

[27] Tomas Jech: Hereditarily countable sets. The Journal of Symbolic Logic 1982 pp. 43-

[28] Tomas Jech: Set Theory Academic press

[29] Shelah, S. *Cardinal Arithmetic for Sceptics*. Bull. Am. Math. Soc. **26** no 2 (1992) pp 197-210.

[30] Sidney Smith's Part III essay: at `http://ucsu.colorado.edu/~bsid/`

[31] http://www.cambridge.org/uk/catalogue/catalogue.asp?isbn=9780521674539

[32] J. C. Shepherdson Inner Models for Set Theory–Part I. The Journal of Symbolic Logic, Vol. 16, No. 3. (Sep., 1951), pp. 161-190.

[33] J. C. Shepherdson. Inner Models for Set Theory–Part II. The Journal of Symbolic Logic, Vol. 17, No. 4. (Dec., 1952), pp. 225-237.

[34] J. C. Shepherdson. Inner Models for Set Theory–Part III. The Journal of Symbolic Logic, Vol. 18, No. 2. (Jun., 1953), pp. 145-167.

[35] Smullyan and Fitting. Set theory and the continuum problem. Oxford Logic guides **34**

[36] van Dalen, Doets and de Swart. Set: naïve, axiomatic and applied. (*formerly a recommended text for Part II: possibly helpful. Copies in most college libraries*)

[37] W. Hugh Woodin The Continuum Hypothesis Part I Notices A.M.S june/july 2001 pp 567–575.

[38] W. Hugh Woodin The Continuum Hypothesis Part I Notices A.M.S august 2001 pp 681–690.

The journal of symbolic logic is available on JSTOR if you request from a .cam.ac.uk address.