XXII

Set Theory II: Model Theory and Forcing

Part XXII: Contents

92 Inner model theory

Model theory is *really* meta, so you will have to pay attention here.

Roughly, a "model of ZFC" is a set with a binary relation that satisfies the ZFC axioms, just as a group is a set with a binary operation that satisfies the group axioms. Unfortunately, unlike with groups, it is very hard for me to give interesting examples of models, for the simple reason that we are literally trying to model the entire universe.

§92.1 Models

Prototypical example for this section: $(\omega, \in)$ *obeys* PowerSet, V_κ *is a model for* κ *inaccessible (later).*

Definition 92.1.1. A **model** $\mathscr{M}$ consists of a set M and a binary relation $E \subseteq M \times M$. (The E relation is the "$\in$" for the model.)

Remark 92.1.2 — I'm only considering *set-sized* models where M is a set. Experts may be aware that I can actually play with M being a class, but that would require too much care for now.

If you have a model, you can ask certain things about it, such as "does it satisfy EmptySet?". Let me give you an example of what I mean and then make it rigorous.

Example 92.1.3 (A stupid model)

Let's take $\mathscr{M} = (M, E) = (\omega, \in)$. This is not a very good model of ZFC, but let's see if we can make sense of some of the first few axioms.

(a) $\mathscr{M}$ satisfies Extensionality, which is the sentence

$$\forall x \forall y \forall a : (a \in x \iff a \in y) \implies x = y.$$

This just follows from the fact that E is actually $\in$.

(b) $\mathscr{M}$ satisfies EmptySet, which is the sentence

$$\exists a : \forall x \neg(x \in a).$$

Namely, take $a = \varnothing \in \omega$.

(c) $\mathscr{M}$ does not satisfy Pairing, since $\{1, 3\}$ is not in ω, even though $1, 3 \in \omega$.

(d) Miraculously, $\mathscr{M}$ satisfies Union, since for any $n \in \omega$, $\cup n$ is $n - 1$ (unless $n = 0$). The Union axiom states that

$$\forall a \exists U \quad \forall x \ [(x \in U) \iff (\exists y : x \in y \in a)] .$$

An important thing to notice is that the "$\forall a$" ranges only over the sets in the model of the universe, $\mathscr{M}$.

Example 92.1.4 (Important: this stupid model satisfies PowerSet)

Most incredibly of all: $\mathscr{M} = (\omega, \in)$ satisfies PowerSet. This is a really important example.

You might think this is ridiculous. Look at $2 = \{0, 1\}$. The power set of this is $\{0, 1, 2, \{1\}\}$ which is not in the model, right?

Well, let's look more closely at PowerSet. It states that:

$$\forall x \exists a \forall y (y \in a \iff y \subseteq x).$$

What happens if we set $x = 2 = \{0, 1\}$? Well, actually, we claim that $a = 3 = \{0, 1, 2\}$ works. The key point is "for all y" – this *only ranges over the objects in* $\mathscr{M}$. In $\mathscr{M}$, the only subsets of 2 are $0 = \varnothing$, $1 = \{0\}$ and $2 = \{0, 1\}$. The "set" $\{1\}$ in the "real world" (in V) is not a set in the model $\mathscr{M}$.

In particular, you might say that in this strange new world, we have $2^n = n + 1$, since $n = \{0, 1, \dots, n-1\}$ really does have only $n + 1$ subsets.

Example 92.1.5 (Sentences with parameters)

The sentences we ask of our model are allowed to have "parameters" as well. For example, if $\mathscr{M} = (\omega, \in)$ as before then $\mathscr{M}$ satisfies the sentence

$$\forall x \in 3 (x \in 5).$$

§92.2 Sentences and satisfaction

With this intuitive notion, we can define what it means for a model to satisfy a sentence.

Definition 92.2.1. Note that any sentence ϕ can be written in one of five forms:

- $x \in y$
- $x = y$
- $\neg \psi$ ("not ψ") for some shorter sentence ψ
- $\psi_1 \vee \psi_2$ ("ψ_1 or ψ_2") for some shorter sentences ψ_1, ψ_2
- $\exists x \psi$ ("exists x") for some shorter sentence ψ.

> **Question 92.2.2.** What happened to $\wedge$ (and) and $\forall$ (for all)? (Hint: use $\neg$.)

Often (almost always, actually) we will proceed by so-called "induction on formula complexity", meaning that we define or prove something by induction using this. Note that we require all formulas to be finite.

Now suppose we have a sentence ϕ, like $a = b$ or $\exists a \forall x \neg (x \in a)$, plus a model $\mathscr{M} = (M, E)$. We want to ask whether $\mathscr{M}$ satisfies ϕ.

To give meaning to this, we have to designate certain variables as **parameters**. For example, if I asked you

"Does $a = b$?"

the first question you would ask is what a and b are. So a, b would be parameters: I have to give them values for this sentence to make sense.

On the other hand, if I asked you

"Does $\exists a \forall x \neg(x \in a)$?"

then you would just say "yes". In this case, x and a are *not* parameters. In general, parameters are those variables whose meaning is not given by some $\forall$ or $\exists$.

In what follows, we will let $\phi(x_1, \dots, x_n)$ denote a formula ϕ, whose parameters are $x_1, \dots, x_n$. Note that possibly $n = 0$, for example all ZFC axioms have no parameters.

> **Question 92.2.3.** Try to guess the definition of satisfaction before reading it below. (It's not very hard to guess!)

Definition 92.2.4. Let $\mathscr{M} = (M, E)$ be a model. Let $\phi(x_1, \dots, x_n)$ be a sentence, and let $b_1, \dots, b_n \in M$. We will define a relation

$$\mathscr{M} \vDash \phi[b_1, \dots, b_n]$$

and say $\mathscr{M}$ **satisfies** the sentence ϕ with parameters $b_1, \dots, b_n$.

The relationship is defined by induction on formula complexity as follows:

- If ϕ is "$x_1 = x_2$" then $\mathscr{M} \vDash \phi[b_1, b_2] \iff b_1 = b_2$.
- If ϕ is "$x_1 \in x_2$" then $\mathscr{M} \vDash \phi[b_1, b_2] \iff b_1 \; E \; b_2$.
  (This is what we mean by "E interprets $\in$".)
- If ϕ is "$\neg\psi$" then $\mathscr{M} \vDash \phi[b_1, \dots, b_n] \iff \mathscr{M} \nvDash \psi[b_1, \dots, b_n]$.
- If ϕ is "$\psi_1 \vee \psi_2$" then $\mathscr{M} \vDash \phi[b_1, \dots, b_n]$ means $\mathscr{M} \vDash \psi_i[b_1, \dots, b_n]$ for some $i = 1, 2$.
- Most important case: suppose ϕ is $\exists x \psi(x, x_1, \dots, x_n)$. Then $\mathscr{M} \vDash \phi[b_1, \dots, b_n]$ if and only if

  $$\exists b \in M \text{ such that } \mathscr{M} \vDash \psi[b, b_1, \dots, b_n].$$

  Note that ψ has one extra parameter.

Notice where the information of the model actually gets used. We only ever use E in interpreting $x_1 \in x_2$; unsurprising. But we only ever use the set M when we are running over $\exists$ (and hence $\forall$). That's well-worth keeping in mind:

> **The behavior of a model essentially comes from $\exists$ and $\forall$, which search through the entire model M.**

And finally,

Definition 92.2.5. A **model of ZFC** is a model $\mathscr{M} = (M, E)$ satisfying all ZFC axioms.

We are especially interested in models of the form $(M, \in)$, where M is a *transitive* set. (We want our universe to be transitive, otherwise we would have elements of sets which are not themselves in the universe, which is very strange.) Such a model is called a **transitive model**.

Abuse of Notation 92.2.6. If M is a transitive set, the model $(M, \in)$ will be abbreviated to just M.

Definition 92.2.7. An **inner model** of ZFC is a transitive model satisfying ZFC.

Remark 92.2.8 — The definition of a model of ZFC only uses $M \vDash \varphi$ where φ has no parameters; nevertheless, you can see that we define what $M \vDash \varphi$ means when φ has parameters because it's used in the definition of $M \vDash \exists x \psi(x)$.

The extension $\varphi(x_1, \dots, x_n)$ is written with round parentheses, but $M \vDash \varphi[b_1, \dots, b_n]$ is written with square brackets — you can think of it as "formally substitute" the parameters $b_1, \dots, b_n$ into φ, because if $b_1, \dots, b_n$ is "actually" substituted into φ, then $\varphi(b_1, \dots, b_n)$ is just a single boolean value.

§92.3 The Levy hierarchy

Prototypical example for this section: $\texttt{isSubset}(x, y)$ *is absolute. The axiom* EmptySet *is* Σ_1, $\texttt{isPowerSetOf}(X, x)$ *is* Π_1.

A key point to remember is that the behavior of a model is largely determined by $\exists$ and $\forall$. It turns out we can say even more than this.

Consider a formula such as

$$\texttt{isEmpty}(x) : \neg \exists a (a \in x)$$

which checks whether a given set x has an element in it. Technically, this has an "$\exists$" in it. But somehow this $\exists$ does not really search over the entire model, because it is *bounded* to search in x. That is, we might informally rewrite this as

$$\neg(\exists a \in x)$$

which doesn't fit into the strict form, but points out that we are only looking over $a \in x$. We call such a quantifier a **bounded quantifier**.

We like sentences with bounded quantifiers because they designate properties which are **absolute** over transitive models. It doesn't matter how strange your surrounding model M is. As long as M is transitive,

$$M \vDash \texttt{isEmpty}(\varnothing)$$

will always hold. Similarly, the sentence

$$\texttt{isSubset}(x, y) : x \subseteq y \text{ i.e. } \forall a \in x (a \in y)$$

is absolute. Sentences with this property are called Σ_0 or Π_0.

The situation is different with a sentence like

$$\texttt{isPowerSetOf}(y, x) : \forall z \left(z \subseteq x \iff z \in y \right)$$

which in English means "y is the power set of x", or just $y = \mathcal{P}(x)$. The $\forall z$ is *not* bounded here. This weirdness is what allows things like

$$\omega \vDash \text{“}\{0, 1, 2\} \text{ is the power set of } \{0, 1\}\text{”}$$

and hence

$$\omega \vDash \text{PowerSet}$$

which was our stupid example earlier. The sentence $\texttt{isPowerSetOf}$ consists of an unbounded $\forall$ followed by an absolute sentence, so we say it is Π_1.

More generally, the **Levy hierarchy** keeps track of how bounded our quantifiers are. Specifically,

- Formulas which have only bounded quantifiers are $\Delta_0 = \Sigma_0 = \Pi_0$.
- Formulas of the form $\exists x_1 \dots \exists x_k \psi$ where ψ is Π_n are considered Σ_{n+1}.
- Formulas of the form $\forall x_1 \dots \forall x_k \psi$ where ψ is Σ_n are considered Π_{n+1}.

(A formula which is both Σ_n and Π_n is called Δ_n, but we won't use this except for $n = 0$.)

Example 92.3.1 (Examples of Δ_0 sentences)

(a) The sentences $\texttt{isEmpty}(x)$, $x \subseteq y$, as discussed above.

(b) The formula "x is transitive" can be expanded as a Δ_0 sentence.

(c) The formula "x is an ordinal" can be expanded as a Δ_0 sentence.

Exercise 92.3.2. Write out the expansions for "x is transitive" and "x is an ordinal" in a Δ_0 form.

Example 92.3.3 (More complex formulas)

(a) The axiom EmptySet is Σ_1; it is $\exists a(\texttt{isEmpty}(a))$, and $\texttt{isEmpty}(a)$ is Δ_0.

(b) The formula "$y = \mathcal{P}(x)$" is Π_1, as discussed above.

(c) The formula "x is countable" is Σ_1. One way to phrase it is "$\exists f$ an injective map $x \hookrightarrow \omega$", which necessarily has an unbounded "$\exists f$".

(d) The axiom PowerSet is Π_3:

$$\forall y \exists P \forall x (x \subseteq y \iff x \in P).$$

Remark 92.3.4 (Why only alternating unbounded quantifier count?) — Note that a formula $\exists a \exists b \psi(a, b)$ can alternatively be written as $\exists c(c \text{ is an ordered pair } (a, b) \wedge \psi(a, b))$, which explains why we only want to consider the formula $\exists a \exists b \psi(a, b)$ as Σ_1.

§92.4 Substructures, and Tarski-Vaught

Let $\mathscr{M}_1 = (M_1, E_1)$ and $\mathscr{M}_2 = (M_2, E_2)$ be models.

Definition 92.4.1. We say that $\mathscr{M}_1 \subseteq \mathscr{M}_2$ if $M_1 \subseteq M_2$ and E_1 agrees with E_2; we say $\mathscr{M}_1$ is a **substructure** of $\mathscr{M}_2$.

That's boring. The good part is:

Definition 92.4.2. We say $\mathscr{M}_1 \prec \mathscr{M}_2$, or $\mathscr{M}_1$ is an **elementary substructure** of $\mathscr{M}_2$, if $\mathscr{M}_1 \subseteq \mathscr{M}_2$ and for *every* sentence $\phi(x_1, \dots, x_n)$ and parameters $b_1, \dots, b_n \in M_1$, we have

$$\mathscr{M}_1 \vDash \phi[b_1, \dots, b_n] \iff \mathscr{M}_2 \vDash \phi[b_1, \dots, b_n].$$

In other words, $\mathscr{M}_1$ and $\mathscr{M}_2$ agree on every sentence possible. Note that the b_i have to come from $\mathscr{M}_1$; if the b_i came from $\mathscr{M}_2$ then asking something of $\mathscr{M}_1$ wouldn't make sense.

Let's ask now: how would $\mathscr{M}_1 \prec \mathscr{M}_2$ fail to be true? If we look at the possible sentences, none of the atomic formulas, nor the "$\land$" and "$\neg$", are going to cause issues.

The intuition you should be getting by now is that things go wrong once we hit $\forall$ and $\exists$. They won't go wrong for bounded quantifiers. But unbounded quantifiers search the entire model, and that's where things go wrong.

To give a "concrete example": imagine $\mathscr{M}_1$ is MIT, and $\mathscr{M}_2$ is the state of Massachusetts. If $\mathscr{M}_1$ thinks there exist hackers at MIT, certainly there exist hackers in Massachusetts. Where things go wrong is something like:

$$\mathscr{M}_2 \vDash \text{“}\exists x : x \text{ is a course numbered } > 50\text{”}.$$

This is true for $\mathscr{M}_2$ because we can take the witness $x = \text{Math } 55$, say. But it's false for $\mathscr{M}_1$, because at MIT all courses are numbered 18.701 or something similar.

The issue is that the *witnesses* for statements in $\mathscr{M}_2$ do not necessarily propagate down to witnesses for $\mathscr{M}_1$.

The Tarski-Vaught test says this is the only impediment: if every witness in $\mathscr{M}_2$ can be replaced by one in $\mathscr{M}_1$ then $\mathscr{M}_1 \prec \mathscr{M}_2$.

Lemma 92.4.3 (Tarski-Vaught)

Let $\mathscr{M}_1 \subseteq \mathscr{M}_2$. Then $\mathscr{M}_1 \prec \mathscr{M}_2$ if and only if: For every sentence $\phi(x, x_1, \dots, x_n)$ and parameters $b_1, \dots, b_n \in M_1$: if there is a witness $\tilde{b} \in M_2$ to $\mathscr{M}_2 \vDash \phi(\tilde{b}, b_1 \dots, b_n)$ then there is a witness $b \in M_1$ to $\mathscr{M}_1 \vDash \phi(b, b_1, \dots, b_n)$.

Proof. Easy after the above discussion. To formalize it, use induction on formula complexity. □

§92.5 Obtaining the axioms of ZFC

We now want to write down conditions for M to satisfy ZFC axioms. The idea is that almost all the ZFC axioms are just Σ_1 claims about certain desired sets, and so verifying an axiom reduces to checking some appropriate "closure" condition: that the witness to the axiom is actually in the model.

For example, the EmptySet axiom is "$\exists a(\texttt{isEmpty}(a))$", and so we're happy as long as $\varnothing \in M$, which is of course true for any nonempty transitive set M.

Lemma 92.5.1 (Transitive sets inheriting ZFC)

Let M be a nonempty transitive set. Then

(i) M satisfies Extensionality, Foundation, EmptySet.

(ii) $M \vDash$ Pairing if $x, y \in M \implies \{x, y\} \in M$.

(iii) $M \vDash$ Union if $x \in M \implies \cup x \in M$.

(iv) $M \vDash$ PowerSet if $x \in M \implies \mathcal{P}(x) \cap M \in M$.

(v) $M \vDash$ Replacement if for every $x \in M$ and every function $F : x \to M$ which is M-definable with parameters, we have $F^{\text{img}}(x) \in M$ as well.

(vi) $M \vDash$ Infinity as long as $\omega \in M$.

Here, a set $X \subseteq M$ is **M-definable with parameters** if it can be realized as

$$X = \{x \in M \mid \phi[x, b_1, \dots, b_n]\}$$

for some (fixed) choice of parameters $b_1, \dots, b_n \in M$. We allow $n = 0$, in which case we say X is **M-definable without parameters**. Note that X need not itself be in M! As a trivial example, $X = M$ is M-definable without parameters (just take $\phi[x]$ to always be true), and certainly we do not have $X \in M$.

Exercise 92.5.2. Verify (i)-(iv) above.

Remark 92.5.3 — Converses to the statements of Lemma 92.5.1 are true for all claims other than (vi).

§92.6 Mostowski collapse

Up until now I have been only talking about transitive models, because they were easier to think about. Here's a second, better reason we might only care about transitive models.

Lemma 92.6.1 (Mostowski collapse lemma)

Let $X = (X, \in)$ be a model satisfying Extensionality, where X is a set (possibly not transitive). Then there exists an isomorphism $\pi : X \to M$ for a transitive model $M = (M, \in)$.

This is also called the *transitive collapse*. In fact, both π and M are unique.

Proof. The idea behind the proof is very simple. Since $\in$ is well-founded and extensional (satisfies Foundation and Extensionality, respectively), we can look at the $\in$-minimal element $x_\varnothing$ of X with respect to $\in$. Clearly, we want to send that to $0 = \varnothing$.

Then we take the next-smallest set under $\in$, and send it to $1 = \{\varnothing\}$. We "keep doing this"; it's not hard to see this does exactly what we want.

To formalize, define π by transfinite recursion:

$$\pi(x) := \{\pi(y) \mid y \in x\}.$$

This π, by construction, does the trick. ☐

Remark 92.6.2 (Digression for experts) — Earlier versions of Napkin claimed this was true for general models $\mathscr{X} = (X, E)$ with $\mathscr{X} \vDash \text{Foundation} + \text{Extensionality}$. This is false; it does not even imply E is well-founded, because there may be infinite descending chains of subsets of X which do not live in X itself. Another issue is that E may not be set-like.

The picture of this is "collapsing" the elements of M down to the bottom of V, hence the name.

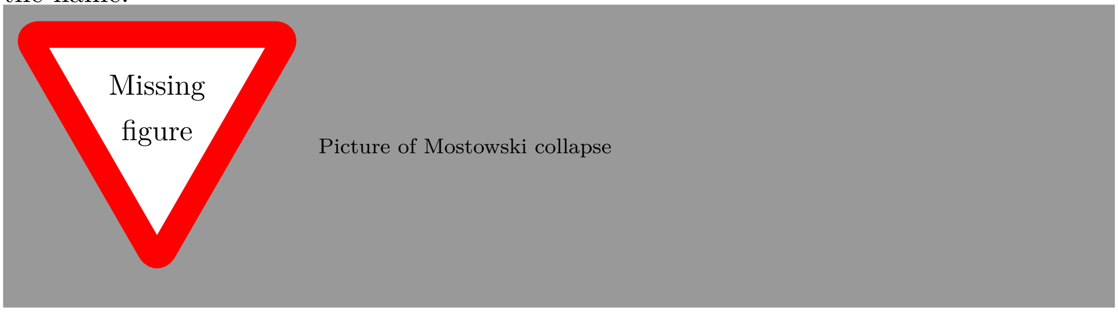

§92.7 Adding an inaccessible

Prototypical example for this section: V_κ

At this point you might be asking, well, where's my model of ZFC?

I unfortunately have to admit now: ZFC can never prove that there is a model of ZFC (unless ZFC is inconsistent, but that would be even worse). This is a result called Gödel's incompleteness theorem.

Nonetheless, with some very modest assumptions added, we can actually show that a model *does* exist: for example, assuming that there exists a strongly inaccessible cardinal κ would do the trick, V_κ will be such a model (Problem 92D*). Intuitively you can see why: κ is so big that any set of rank lower than it can't escape it even if we take their power sets, use the Replacement axiom, or any other method that ZFC lets us do.

More pessimistically, this shows that it's impossible to prove in ZFC that such a κ exists. Nonetheless, we now proceed under ZFC^+ for convenience, which adds the existence of such a κ as a final axiom. So we now have a model V_κ to play with. Joy!

Great. Now we do something *really* crazy.

Theorem 92.7.1 (Countable transitive model)
Assume ZFC^+. Then there exists a transitive model X of ZFC such that X is a *countable* set.

Proof. Fasten your seat belts.

First, since we assumed ZFC^+, we can take $V_\kappa = (V_\kappa, \in)$ as our model of ZFC. Start with the set $X_0 = \varnothing$. Then for every integer n, we do the following to get X_{n+1}.

- Start with X_{n+1} containing every element of X_n.
- Consider a formula $\phi(x, x_1, \dots, x_n)$ and $b_1, \dots, b_n$ in X_n. Suppose that V_κ thinks there is a $b \in V_\kappa$ for which
$$V_\kappa \vDash \phi[b, b_1, \dots, b_n].$$

We then add in the element b to X_{n+1}.

- We do this for *every possible formula in the language of set theory*. We also have to put in *every possible set of parameters* from the previous set X_n.

At every step X_n is countable. Reason: there are countably many possible finite sets of parameters in X_n, and countably many possible formulas, so in total we only ever add in countably many things at each step. This exhibits an infinite nested sequence of countable sets

$$X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$$

None of these is an elementary substructure of V_κ, because each X_n relies on witnesses in X_{n+1}. So we instead *take the union*:

$$X = \bigcup_n X_n.$$

This satisfies the Tarski-Vaught test, and is countable.

There is one minor caveat: X might not be transitive. We don't care, because we just take its Mostowski collapse. □

Please take a moment to admire how insane this is. It hinges irrevocably on the fact that there are countably many sentences we can write down.

Remark 92.7.2 — This proof relies heavily on the Axiom of Choice when we add in the element b to X_{n+1}. Without Choice, there is no way of making these decisions all at once.

Usually, the right way to formalize the Axiom of Choice usage is, for every formula $\phi(x, x_1, \dots, x_n)$, to pre-commit (at the very beginning) to a function $f_\phi(x_1, \dots, x_n)$, such that given any $b_1, \dots, b_n$, $f_\phi(b_1, \dots, b_n)$ will spit out the suitable value of b (if one exists). Personally, I think this is hiding the spirit of the proof, but it does make it clear how exactly Choice is being used.

These f_ϕ's have a name: **Skolem functions**.

The trick we used in the proof works in more general settings:

Theorem 92.7.3 (Downward Löwenheim-Skolem theorem)

Let $\mathscr{M} = (M, E)$ be a model, and $A \subseteq M$. Then there exists a set B (called the **Skolem hull** of A) with $A \subseteq B \subseteq M$, such that $(B, E) \prec \mathscr{M}$, and

$$|B| = \max\{\omega, |A|\}.$$

In our case, what we did was simply take A to be the empty set.

Question 92.7.4. Prove this. (Exactly the same proof as before.)

§92.8 FAQ's on countable models

The most common one is "how is this possible?", with runner-up "what just happened?".

Let me do my best to answer the first question. It seems like there are two things running up against each other:

(1) M is a transitive model of ZFC, but its universe is countable.

(2) ZFC tells us there are uncountable sets!

(This has confused so many people it has a name, **Skolem's paradox**.)

The reason this works I actually pointed out earlier: *countability is not absolute, it is a* Σ_1 *notion.*

Recall that a set x is countable if *there exists* an injective map $x \hookrightarrow \omega$. The first statement just says that *in the universe* V, there is an injective map $F \colon M \hookrightarrow \omega$. In particular, for any $x \in M$ (hence $x \subseteq M$, since M is transitive), x is countable *in* V. This is the content of the first statement.

But for M to be a model of ZFC, M only has to think statements in ZFC are true. More to the point, the fact that ZFC tells us there are uncountable sets means

$$M \vDash \exists x \text{ uncountable.}$$

In other words,

$$M \vDash \exists x \forall f \text{ If } f \colon x \to \omega \text{ then } f \text{ isn't injective.}$$

The key point is the $\forall f$ searches only functions in our tiny model M. It is true that in the "real world" V, there are injective functions $f \colon x \to \omega$.[1] But M has no idea they exist! It is a brain in a vat: M is oblivious to any information outside it.

So in fact, every ordinal which appears in M is countable in the real world. It is just not countable in M. Since $M \vDash \mathsf{ZFC}$, M is going to think there is some smallest uncountable cardinal, say $\aleph_1^M$. It will be the smallest (infinite) ordinal in M with the property that there is no bijection *in the model* M between $\aleph_1^M$ and ω. However, we necessarily know that such a bijection is going to exist in the real world V.

Put another way, cardinalities in M can look vastly different from those in the real world, because cardinality is measured by bijections, which I guess is inevitable, but leads to chaos.

§92.9 Picturing inner models

Here is a picture of a countable transitive model M.

[1] Since M is transitive and countable, for any $x \in M$, then x itself can have at most countably many elements. This isn't necessarily true if M is not transitive.

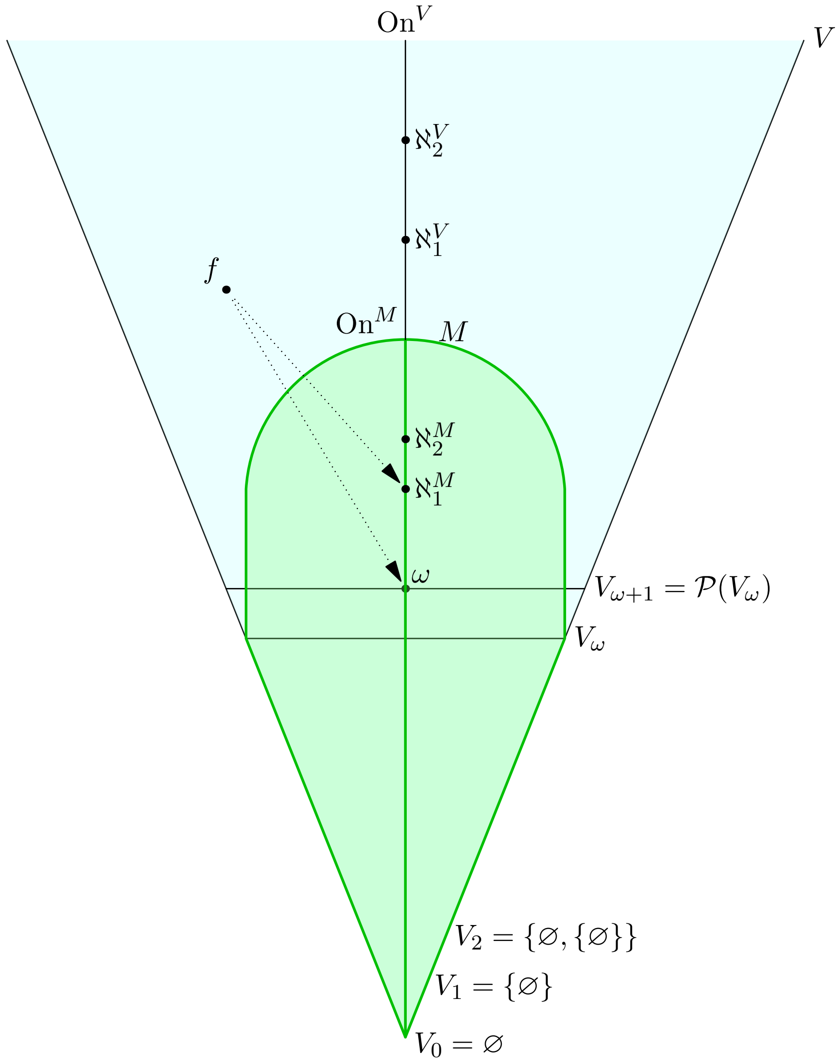

Note that M and V must agree on finite sets, since every finite set has a formula that can express it. However, past V_ω the model and the true universe start to diverge.

The entire model M is countable, so it only occupies a small portion of the universe, below the first uncountable cardinal $\aleph_1^V$ (where the superscript means "of the true universe V"). The ordinals in M are precisely the ordinals of V which happen to live inside the model, because the sentence "α is an ordinal" is absolute. On the other hand, M has only a portion of these ordinals, since it is only a lowly set, and a countable set at that. To denote the ordinals of M, we write On^M, where the superscript means "the ordinals as computed in M". Similarly, On^V will now denote the "set of true ordinals".

Nonetheless, the model M has its own version of the first uncountable cardinal $\aleph_1^M$. In the true universe, $\aleph_1^M$ is countable (below $\aleph_1^V$), but the necessary bijection witnessing this might not be inside M. That's why M can think $\aleph_1^M$ is uncountable, even if it is a countable cardinal in the original universe.

So our model M is a brain in a vat. It happens to believe all the axioms of ZFC, and so every statement that is true in M could conceivably be true in V as well. But M can't see the universe around it; it has no idea that what it believes is the uncountable $\aleph_1^M$ is really just an ordinary countable ordinal.

§92.10 A few harder problems to think about

Problem 92A*. Show that for any transitive model M, the set of ordinals in M is itself some ordinal.

Problem 92B†. Assume $\mathscr{M}_1 \subseteq \mathscr{M}_2$. Show that

(a) If ϕ is Δ_0, then $\mathscr{M}_1 \vDash \phi[b_1, \dots, b_n] \iff \mathscr{M}_2 \vDash \phi[b_1, \dots, b_n]$.

(b) If ϕ is Σ_1, then $\mathscr{M}_1 \vDash \phi[b_1, \dots, b_n] \implies \mathscr{M}_2 \vDash \phi[b_1, \dots, b_n]$.

(c) If ϕ is Π_1, then $\mathscr{M}_2 \vDash \phi[b_1, \dots, b_n] \implies \mathscr{M}_1 \vDash \phi[b_1, \dots, b_n]$.

(This should be easy if you've understood the chapter.)

Problem 92C† (Reflection)**.** Let κ be an inaccessible cardinal such that $|V_\alpha| < \kappa$ for all $\alpha < \kappa$. Prove that for any $\delta < \kappa$ there exists $\delta < \alpha < \kappa$ such that $V_\alpha \prec V_\kappa$; in other words, the set of α such that $V_\alpha \prec V_\kappa$ is *unbounded* in κ. This means that properties of V_κ reflect down to properties of V_α.

Problem 92D* (Strongly inaccessible cardinals produce models)**.** Let κ be a strongly inaccessible cardinal. Prove that V_κ is a model of ZFC.

93 Forcing

We are now going to introduce Paul Cohen's technique of **forcing**, which we then use to break the Continuum Hypothesis.

Here is how it works. Given a transitive model M and a poset $\mathbb{P}$ inside it, we can consider a "generic" subset $G \subseteq \mathbb{P}$, where G is not in M. Then, we are going to construct a bigger universe $M[G]$ which contains both M and G. (This notation is deliberately the same as $\mathbb{Z}[\sqrt{2}]$, for example – in the algebra case, we are taking $\mathbb{Z}$ and adding in a new element $\sqrt{2}$, plus everything that can be generated from it.) By choosing $\mathbb{P}$ well, we can cause $M[G]$ to have desirable properties.

Picture:

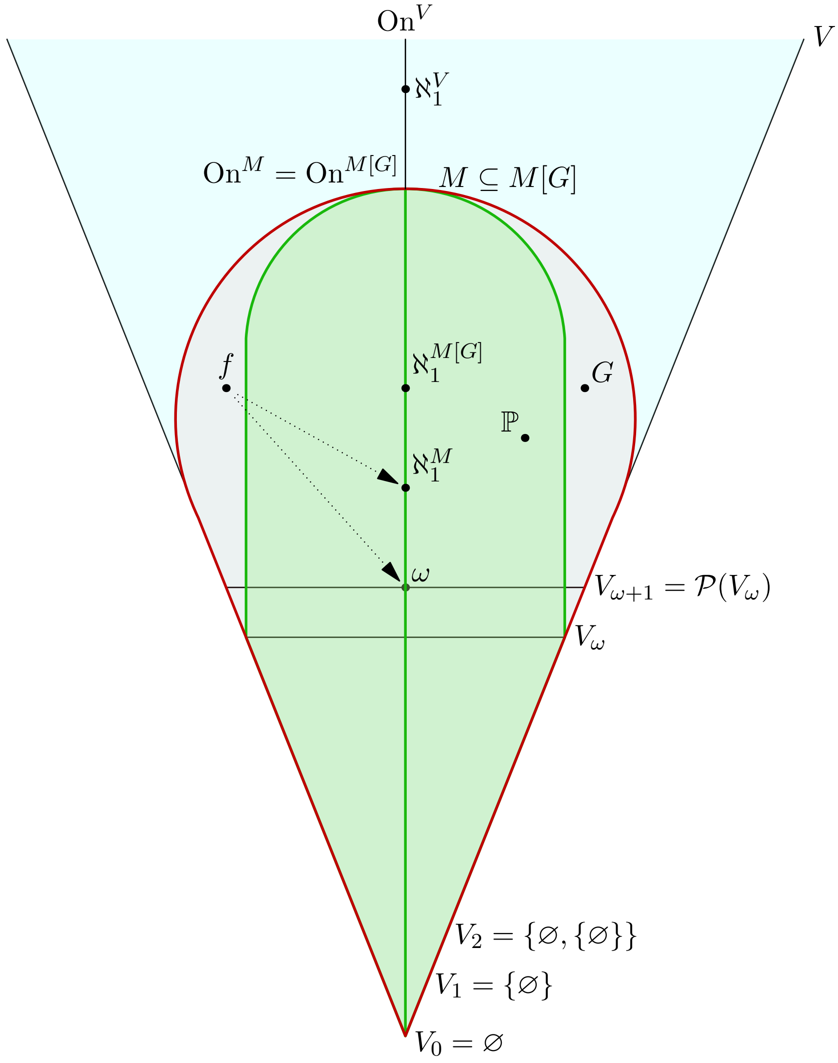

The model M is drawn in green, and its extension $M[G]$ is drawn in red.

The models M and $M[G]$ will share the same ordinals, which is represented here as M being no taller than $M[G]$. But one issue with this is that forcing may introduce some new bijections between cardinals of M that were not there originally; this leads to the phenomenon called **cardinal collapse**: quite literally, cardinals in M will no longer be cardinals in $M[G]$, and instead just an ordinal. This is because in the process of adjoining G, we may accidentally pick up some bijections which were not in the earlier universe. In the diagram drawn, this is the function f mapping ω to $\aleph_1^M$. Essentially, the difficulty is that "κ is a cardinal" is a Π_1 statement.

In the case of the Continuum Hypothesis, we'll introduce a $\mathbb{P}$ such that any generic subset G will "encode" $\aleph_2^M$ real numbers. We'll then show cardinal collapse does not occur, meaning $\aleph_2^{M[G]} = \aleph_2^M$. Thus $M[G]$ will have $\aleph_2^{M[G]}$ real numbers, as desired.

§93.1 Setting up posets

Prototypical example for this section: Infinite Binary Tree

Let M be a transitive model of ZFC. Let $\mathbb{P} = (\mathbb{P}, \leq) \in M$ be a poset with a maximum element $1_{\mathbb{P}}$ which lives inside a model M. The elements of $\mathbb{P}$ are called **conditions**; because they will force things to be true in $M[G]$.

Definition 93.1.1. A subset $D \subseteq \mathbb{P}$ is **dense** if for all $p \in \mathbb{P}$, there exists a $q \in D$ such that $q \leq p$.

Examples of dense subsets include the entire $\mathbb{P}$ as well as any downwards "slice".

Definition 93.1.2. For $p, q \in \mathbb{P}$ we write $p \parallel q$, saying "p is **compatible** with q", if there exists $r \in \mathbb{P}$ with $r \leq p$ and $r \leq q$. Otherwise, we say p and q are **incompatible** and write $p \perp q$.

Example 93.1.3 (Infinite binary tree)

Let $\mathbb{P} = 2^{<\omega}$ be the **infinite binary tree** shown below, extended to infinity in the obvious way:

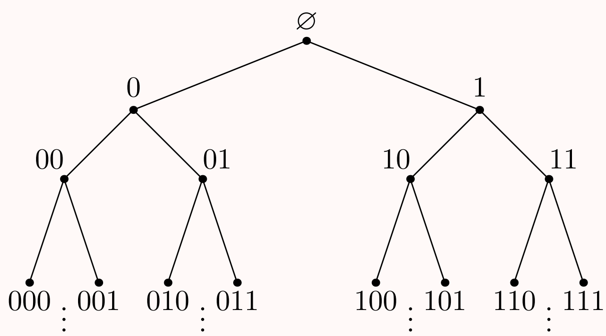

(a) The maximum element $1_{\mathbb{P}}$ is the empty string $\varnothing$.

(b) $D = \{\text{all strings ending in } 001\}$ is an example of a dense set.

(c) No two elements of $\mathbb{P}$ are compatible unless they are comparable.

Example 93.1.4 (Infinite chain)

Let $\mathbb{P} = (\mathbb{N}, \geq)$. This can be considered the "infinite unary tree".

- The maximum element $1_{\mathbb{P}}$ is 1.
- A set is dense if and only if it has infinitely many elements. For example, the

set of all positive even numbers and the set of all primes are dense.

Now, I can specify what it means to be "generic".

Definition 93.1.5. A nonempty set $G \subseteq \mathbb{P}$ is a **filter** if

(a) The set G is upwards-closed: $\forall p \in G(\forall q \geq p)(q \in G)$.

(b) Any pair of elements in G is compatible.

We say G is **M-generic** if for all D which are *in the model M*, if D is dense then $G \cap D \neq \varnothing$.

Question 93.1.6. Show that if G is a filter then $1_{\mathbb{P}} \in G$.

Note that the condition that $D \in M$ is important, because:

Question 93.1.7. On the infinite binary tree, show that:

- For every filter G, there's a dense subset D (not necessarily in M) such that $G \cap D = \varnothing$.
- Specifically, if $G \in M$, then such a set D can be chosen such that $D \in M$ — in particular, G is not M-generic.

We will formalize this later in Lemma 93.2.5.

Example 93.1.8 (Generic filters on the infinite binary tree)

Let $\mathbb{P} = 2^{<\omega}$. The generic filters on $\mathbb{P}$ are sets of the form

$$\{0,\ b_1,\ b_1b_2,\ b_1b_2b_3,\ \dots\}.$$

So every generic filter on $\mathbb{P}$ corresponds to a binary number $b = 0.b_1b_2b_3\dots$.[a]

It is harder to describe which reals correspond to generic filters, but they should really "look random". For example, the set of strings ending in 011 is dense, so one should expect "011" to appear inside b, and more generally that b should contain every binary string. So one would expect the binary expansion of $\pi - 3$ might correspond to a generic, but not something like $0.010101\dots$. That's why we call them "generic".

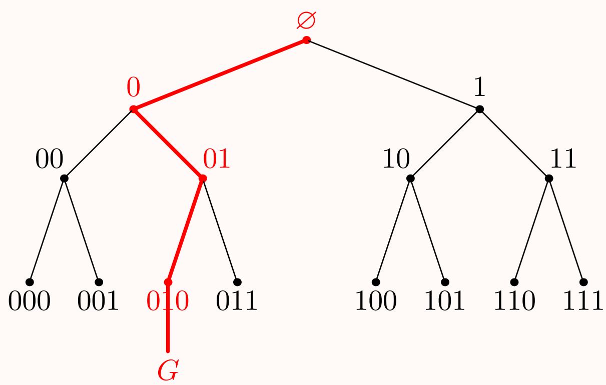

[a]Note that it may be the case that two distinct filters correspond to the same real number, such as $0.1000\dots$ and $0.0111\dots$, but such filters are necessarily not generic.

Example 93.1.9 (Generic filters on the infinite chain)

There's only one generic filter on $\mathbb{P} = (\mathbb{N}, \geq)$: $\mathbb{N}$ itself.

This is indeed generic — it hits every dense set. This doesn't "look random" by any measure, you may say — but if you think about it, if you start at the root $1_{\mathbb{P}} = 1$ and move down randomly at each step, there's only one choice where to go!

Exercise 93.1.10. Verify that every generic filter $2^{<\omega}$ has the form above. Show that conversely, a binary number gives a filter, but it need not be generic.

Notice that if $p \geq q$, then the sentence $q \in G$ tells us more information than the sentence $p \in G$. In that sense q is a *stronger* condition. In another sense $1_{\mathbb{P}}$ is the weakest possible condition, because it tells us nothing about G; we always have $1_{\mathbb{P}} \in G$ since G is upwards closed.

§93.2 More properties of posets

We had better make sure that generic filters exist.

Example 93.2.1

If $\mathbb{P}$ is the infinite binary tree and M contains every subset of $\mathbb{P}$, then generic filter does not exist — for every filter $G \subseteq \mathbb{P}$, $\mathbb{P} \setminus G$ is a dense set $D \in M$, and $G \cap M = \varnothing$.

In fact this is kind of tricky, but for countable models it works:

Lemma 93.2.2 (Rasiowa-Sikorski lemma)

Suppose M is a *countable* transitive model of ZFC and $\mathbb{P}$ is a partial order. Then there exists an M-generic filter G.

Proof. Essentially, hit them one by one. Problem 93B. □

Fortunately, for breaking CH we would want M to be countable anyways.

The other thing we want to do to make sure we're on the right track is guarantee that a generic set G is not actually in M. (Analogy: $\mathbb{Z}[3]$ is a really stupid extension.) The condition that guarantees this is:

Definition 93.2.3. A partial order $\mathbb{P}$ is **splitting** if for all $p \in \mathbb{P}$, there exists $q, r \leq p$ such that $q \perp r$.

Example 93.2.4 (Infinite binary tree is (very) splitting)

The infinite binary tree is about as splitting as you can get. Given $p \in 2^{<\omega}$, just consider the two elements right under it.

Lemma 93.2.5 (Splitting posets omit generic sets)

Suppose $\mathbb{P}$ is splitting. Then if $F \subseteq \mathbb{P}$ is a filter such that $F \in M$, then $\mathbb{P} \setminus F$ is dense. In particular, if $G \subseteq \mathbb{P}$ is generic, then $G \notin M$.

Proof. Consider $p \notin \mathbb{P} \setminus F \iff p \in F$. Since $\mathbb{P}$ is splitting, there exist $q, r \le p$ which are not compatible. Since F is a filter it cannot contain both; we must have one of them outside F, say q. Hence every element of $p \in \mathbb{P} \setminus (\mathbb{P} \setminus F)$ has an element $q \le p$ in $\mathbb{P} \setminus F$. That's enough to prove $\mathbb{P} \setminus F$ is dense.

Question 93.2.6. Deduce the last assertion of the lemma about generic G. □

§93.3 Names, and the generic extension

We now define the *names* associated to a poset $\mathbb{P}$.

Definition 93.3.1. Suppose M is a transitive model of ZFC, $\mathbb{P} = (\mathbb{P}, \le) \in M$ is a partial order. We define the hierarchy of **$\mathbb{P}$-names** recursively by

$$\begin{aligned} \mathrm{Name}_0 &= \varnothing \\ \mathrm{Name}_{\alpha+1} &= \mathcal{P}(\mathrm{Name}_\alpha \times \mathbb{P}) \\ \mathrm{Name}_\lambda &= \bigcup_{\alpha<\lambda} \mathrm{Name}_\alpha . \end{aligned}$$

Finally, $\mathrm{Name} = \bigcup_\alpha \mathrm{Name}_\alpha$ denote the class of all $\mathbb{P}$-names.

(These Name_α's are the analog of the V_α's: each Name_α is just the set of all names with rank $\le \alpha$.)

Definition 93.3.2. For a filter G, we define the **interpretation** of τ by G, denoted τ^G, using the transfinite recursion

$$\tau^G = \left\{ \sigma^G \mid \langle \sigma, p \rangle \in \tau \text{ and } p \in G \right\}.$$

We then define the model

$$M[G] = \left\{ \tau^G \mid \tau \in \mathrm{Name}^M \right\}$$

where Name^M are the elements of the class Name that belongs to M. Thus $M[G]$ is a set. In words, $M[G]$ is the interpretation of all the possible $\mathbb{P}$-names (as computed by M).

You should think of a $\mathbb{P}$-name as a "fuzzy set". Here's the idea. Ordinary sets are collections of ordinary sets, so fuzzy sets should be collections of fuzzy sets. These fuzzy sets can be thought of like the Ghosts of Christmases yet to come: they represent things that might be, rather than things that are certain. In other words, they represent the possible futures of $M[G]$ for various choices of G.

Every fuzzy set has an element $p \in \mathbb{P}$ pinned to it. When it comes time to pass judgment, we pick a generic G and filter through the universe of $\mathbb{P}$-names. The fuzzy sets with an element of G attached to it materialize into the real world, while the fuzzy sets with elements outside of G fade from existence. The result is $M[G]$.

Example 93.3.3 (First few levels of the name hierarchy)

Let us compute

$$\begin{aligned}\text{Name}_0 &= \varnothing \\ \text{Name}_1 &= \mathcal{P}(\varnothing \times \mathbb{P}) \\ &= \{\varnothing\} \\ \text{Name}_2 &= \mathcal{P}(\{\varnothing\} \times \mathbb{P}) \\ &= \mathcal{P}\left(\{\langle \varnothing, p\rangle \mid p \in \mathbb{P}\}\right).\end{aligned}$$

Compare the corresponding von Neumann universe.

$$V_0 = \varnothing,\ V_1 = \{\varnothing\},\ V_2 = \{\varnothing, \{\varnothing\}\}.$$

Example 93.3.4 (Example of an interpretation)

As we said earlier, $\text{Name}_1 = \{\varnothing\}$. Now suppose

$$\tau = \{\langle \varnothing, p_1\rangle, \langle \varnothing, p_2\rangle, \dots, \langle \varnothing, p_n\rangle\} \in \text{Name}_2.$$

Then

$$\tau^G = \{\varnothing \mid \langle \varnothing, p\rangle \in \tau \text{ and } p \in G\} = \begin{cases} \{\varnothing\} & \text{if some } p_i \in G \\ \varnothing & \text{otherwise.}\end{cases}$$

In particular, since $1_{\mathbb{P}} \in G$, then:

- when $n = 0$, $\tau = \varnothing$, so $\tau^G = \varnothing$.
- when $n = 1$ and $p_1 = 1_{\mathbb{P}}$, $\tau = \{\langle \varnothing, 1_{\mathbb{P}}\rangle\}$, so $\tau^G = \{\varnothing\}$.

So,

$$\left\{\tau^G \mid \tau \in \text{Name}_2\right\} = V_2.$$

In fact, this holds for any natural number n, not just 2.

So, $M[G]$ and M agree on finite sets.

Now, we want to make sure $M[G]$ contains the elements of M. The proof of $\left\{\tau^G \mid \tau \in \text{Name}_2\right\} = V_2$ above can be easily adapted: Since $1_{\mathbb{P}}$ must be in G, we define for every $x \in M$ the set

$$\check{x} = \{\langle \check{y}, 1_{\mathbb{P}}\rangle \mid y \in x\}$$

by transfinite recursion. Basically, $\check{x}$ is just a copy of x where we tag every element *at every nesting level* with $1_{\mathbb{P}}$.

Example 93.3.5

Compute $\check{0} = 0$ and $\check{1} = \{\langle \check{0}, 1_{\mathbb{P}}\rangle\}$. Thus

$$(\check{0})^G = 0 \quad \text{and} \quad (\check{1})^G = 1.$$

Question 93.3.6. Show that in general, $(\check{x})^G = x$. (Rank induction.)

However, we'd also like to cause G to be in $M[G]$. In fact, we can write down the name exactly: we define

$$\dot{\mathbb{P}} := \{\langle \check{p}, p\rangle \mid p \in \mathbb{P}\}.$$

Question 93.3.7. Show that $(\dot{\mathbb{P}}) \in \text{Name}^M$, and $(\dot{\mathbb{P}})^G = G$.

Question 93.3.8. Verify that $M[G]$ is transitive: that is, if $\sigma^G \in \tau^G \in M[G]$, show that $\sigma^G \in M[G]$. (This is offensively easy.)

In summary,

$M[G]$ is a transitive model extending M (it contains G).

Moreover, it is reasonably well-behaved even if G is just a filter. Let's see what we can get off the bat.

Lemma 93.3.9 (Properties obtained from filters)
Let M be a transitive model of ZFC. If G is a filter, then $M[G]$ is transitive and satisfies Extensionality, Foundation, EmptySet, Infinity, Pairing, and Union.

This leaves PowerSet, Replacement, and Choice.

Proof. We get Extensionality and Foundation for free. Then Infinity and EmptySet follows from $M \subseteq M[G]$.

For Pairing, suppose $\sigma_1^G, \sigma_2^G \in M[G]$. Then

$$\sigma = \{\langle \sigma_1, 1_{\mathbb{P}}\rangle, \langle \sigma_2, 1_{\mathbb{P}}\rangle\}$$

satisfies $\sigma^G = \{\sigma_1^G, \sigma_2^G\}$. (Note that we used $M \vDash \text{Pairing}$.) Union is left as a problem, which you are encouraged to try now. □

Up to here, we don't need to know anything about when a sentence is true in $M[G]$; all we had to do was contrive some names like $\check{x}$ or $\{\langle \sigma_1, 1_{\mathbb{P}}\rangle, \langle \sigma_2, 1_{\mathbb{P}}\rangle\}$ to get the facts we wanted. But for the remaining axioms, we *are* going to need this extra power. For this, we have to introduce the fundamental theorem of forcing.

§93.4 Fundamental theorem of forcing

The model M unfortunately has no idea what G might be, only that it is some generic filter.[1] Nonetheless, we are going to define a relation $\Vdash$, called the *forcing* relation. Roughly, we are going to write

$$p \Vdash \varphi(\sigma_1, \dots, \sigma_n)$$

where $p \in \mathbb{P}$, $\sigma_1, \dots, \sigma_n \in M[G]$, if and only if:

[1]You might say this is a good thing; here's why. We're trying to show that $\neg$CH is consistent with ZFC, and we've started with a model M of the real universe V. But for all we know CH might be true in V (what if $V = L$?), in which case it would also be true of M.

Nonetheless, we boldly construct $M[G]$ an extension of the model M. In order for it to behave differently from M, it has to be out of reach of M. Conversely, if M could compute everything about $M[G]$, then $M[G]$ would have to conform to M's beliefs.

That's why we worked so hard to make sure $G \in M[G]$ but $G \notin M$.

For *any* generic G, if $p \in G$, then $M[G] \vDash \varphi[\sigma_1^G, \dots, \sigma_n^G]$.

Note that $\Vdash$ is defined without reference to G: it is something that M can see. We say p **forces** the sentence $\varphi(\sigma_1, \dots, \sigma_n)$. And miraculously, we can define this relation in such a way that the converse is true: *a sentence holds if and only if some p forces it.*

Theorem 93.4.1 (Fundamental theorem of forcing)
Suppose M is a transitive model of ZF. Let $\mathbb{P} \in M$ be a poset, and $G \subseteq \mathbb{P}$ is an M-generic filter. Then,

(1) Consider $\sigma_1, \dots, \sigma_n \in \operatorname{Name}^M$. Then

$$M[G] \vDash \varphi[\sigma_1^G, \dots, \sigma_n^G]$$

if and only if there exists a condition $p \in G$ such that p *forces* the sentence $\varphi(\sigma_1, \dots, \sigma_n)$. We denote this by $p \Vdash \varphi(\sigma_1, \dots, \sigma_n)$.

(2) This forcing relation is (uniformly) definable in M.

I'll tell you how the definition works in the next section.

§93.5 (Optional) Defining the relation

Here's how we're going to go. We'll define the most generous condition possible such that the forcing works in one direction ($p \Vdash \varphi(\sigma_1, \dots, \sigma_n)$ means $M[G] \vDash \varphi[\sigma_1^G, \dots, \sigma_n^G]$). We will then cross our fingers that the converse also works.

We proceed by induction on the formula complexity. It turns out in this case that the atomic formulas (base cases) are hardest and themselves require induction on ranks.

For some motivation, let's consider how we should define $p \Vdash \tau_1 \in \tau_2$ assuming that we've already defined $p \Vdash \tau_1 = \tau_2$. We need to ensure this holds iff

$$\forall M\text{-generic } G \text{ with } p \in G: \ M[G] \vDash \tau_1^G \in \tau_2^G.$$

So it suffices to ensure that any generic $G \ni p$ hits a condition q which forces τ_1^G to *equal* a member τ^G of τ_2^G. In other words, we want to choose the definition of $p \Vdash \tau_1 \in \tau_2$ to hold if and only if

$$\{q \in \mathbb{P} \mid \exists \langle \tau, r \rangle \in \tau_2 \, (q \le r \land q \Vdash (\tau = \tau_1))\}$$

is dense below in p. In other words, if the set is dense, then the generic must hit q, so it must hit r (recall that a filter is upwards-closed), meaning that $\langle \tau_r \rangle \in \tau_2$ will get interpreted such that $\tau^G \in \tau_2^G$, and moreover the $q \in G$ will force $\tau_1 = \tau$.

Now let's write down the definition... In what follows, the $\Vdash$ omits the M and $\mathbb{P}$.

Definition 93.5.1. Let M be a countable transitive model of ZFC. Let $\mathbb{P} \in M$ be a partial order. For $p \in \mathbb{P}$ and $\varphi(\sigma_1, \dots, \sigma_n)$ a formula in the language of set theory, we write $\tau \Vdash \varphi(\sigma_1, \dots, \sigma_n)$ to mean the following, defined by induction on formula complexity plus rank.

(1) $p \Vdash \tau_1 = \tau_2$ means

(i) For all $\langle \sigma_1, q_1 \rangle \in \tau_1$ the set

$$D_{\sigma_1,q_1} := \{r \mid r \leq q_1 \to \exists \langle \sigma_2, q_2 \rangle \in \tau_2 \, (r \leq q_2 \wedge r \Vdash (\sigma_1 = \sigma_2))\} .$$

is dense in p. (This encodes "$\tau_1 \subseteq \tau_2$".)

(ii) For all $\langle \sigma_2, q_2 \rangle \in \tau_2$, the set D_{σ_2,q_2} defined similarly is dense below p.

(2) $p \Vdash \tau_1 \in \tau_2$ means

$$\{q \in \mathbb{P} \mid \exists \langle \tau, r \rangle \in \tau_2 \, (q \leq r \wedge q \Vdash (\tau = \tau_1))\}$$

is dense below p.

(3) $p \Vdash \varphi \wedge \psi$ means $p \Vdash \varphi$ and $p \Vdash \psi$.

(4) $p \Vdash \neg\varphi$ means $\forall q \leq p$, $q \nVdash \varphi$.

(5) $p \Vdash \exists x \varphi(x, \sigma_1, \dots, \sigma_n)$ means that the set

$$\{q \mid \exists \tau \, (q \Vdash \varphi(\tau, \sigma_1, \dots, \sigma_n))\}$$

is dense below p.

This is definable in M! All we've referred to is $\mathbb{P}$ and names, which are in M. (Note that being dense is definable.) Actually, in parts (3) through (5) of the definition above, we use induction on formula complexity. But in the atomic cases (1) and (2) we are doing induction on the ranks of the names.

So, the construction above gives us one direction (I've omitted tons of details, but...).

Now, how do we get the converse: that a sentence is true if and only if something forces it? Well, by induction, we can actually show:

Lemma 93.5.2 (Consistency and Persistence)

We have

(1) (Consistency) If $p \Vdash \varphi$ and $q \leq p$ then $q \Vdash \varphi$.

(2) (Persistence) If $\{q \mid q \Vdash \varphi\}$ is dense below p then $p \Vdash \varphi$.

You can prove both of these by induction on formula complexity. From this we get:

Corollary 93.5.3 (Completeness)

The set $\{p \mid p \Vdash \varphi \text{ or } p \Vdash \neg\varphi\}$ is dense.

Proof. We claim that whenever $p \nVdash \varphi$ then for some $\overline{p} \leq p$ we have $\overline{p} \Vdash \neg\varphi$; this will establish the corollary.

By the contrapositive of the previous lemma, $\{q \mid q \Vdash \varphi\}$ is not dense below p, meaning for some $\overline{p} \leq p$, every $q \leq \overline{p}$ gives $q \nVdash \varphi$. By the definition of $p \Vdash \neg\varphi$, we have $\overline{p} \Vdash \neg\varphi$. □

And this gives the converse: the M-generic G has to hit some condition that passes judgment, one way or the other. This completes the proof of the fundamental theorem.

§93.6 The remaining axioms

Theorem 93.6.1 (The generic extension satisfies ZFC)

Suppose M is a transitive model of ZFC. Let $\mathbb{P} \in M$ be a poset, and $G \subseteq \mathbb{P}$ is an M-generic filter. Then

$$M[G] \vDash \mathsf{ZFC}.$$

Proof. We'll just do Comprehension, as the other remaining axioms are similar.

Suppose $\sigma^G, \sigma_1^G, \dots, \sigma_n^G \in M[G]$ are a set and parameters, and $\varphi(x, x_1, \dots, x_n)$ is a formula in the language of set theory. We want to show that the set

$$A = \left\{ x \in \sigma^G \mid M[G] \vDash \varphi[x, \sigma_1^G, \dots, \sigma_n^G] \right\}$$

is in $M[G]$; i.e. it is the interpretation of some name.

Note that every element of σ^G is of the form ρ^G for some $\rho \in \operatorname{dom}(\sigma)$ (a bit of abuse of notation here, σ is a bunch of pairs of names and p's, and the domain $\operatorname{dom}(\sigma)$ is just the set of names). So by the fundamental theorem of forcing, we may write

$$A = \left\{ \rho^G \mid \rho \in \operatorname{dom}(\sigma) \text{ and } \exists p \in G \, (p \Vdash \rho \in \sigma \wedge \varphi(\rho, \sigma_1, \dots, \sigma_n)) \right\}.$$

To show $A \in M[G]$ we have to write down a τ such that the name τ^G coincides with A. We claim that

$$\tau = \{ \langle \rho, p \rangle \in \operatorname{dom}(\sigma) \times \mathbb{P} \mid p \Vdash \rho \in \sigma \wedge \varphi(\rho, \sigma_1, \dots, \sigma_n) \}$$

is the correct choice. It's actually clear that $\tau^G = A$ by construction; the "content" is showing that τ is in actually a name of M, which follows from $M \vDash$ Comprehension.

So really, the point of the fundamental theorem of forcing is just to let us write down this τ; it lets us show that τ is in Name^M without actually referencing G. $\square$

§93.7 A few harder problems to think about

Problem 93A. For a filter G and M a transitive model of ZFC, show that $M[G] \vDash$ Union.

Problem 93B (Rasiowa-Sikorski lemma)**.** Show that in a countable transitive model M of ZFC, one can find an M-generic filter on any partial order.

94 Breaking the continuum hypothesis

We now use the technique of forcing to break the Continuum Hypothesis by choosing a good poset $\mathbb{P}$. As I mentioned earlier, one can also build a model where the Continuum Hypothesis is true; this is called the *constructible universe*, (this model is often called "$V = L$"). However, I think it's more fun when things break...

§94.1 Adding in reals

Starting with a *countable* transitive model M.

We want to choose $\mathbb{P} \in M$ such that $(\aleph_2)^M$ many real numbers appear, and then worry about cardinal collapse later.

Recall the earlier situation where we set $\mathbb{P}$ to be the infinite complete binary tree; its nodes can be thought of as partial functions $n \to 2$ where $n < \omega$. Then G itself is a path down this tree; i.e. it can be encoded as a total function $G : \omega \to 2$, and corresponds to a real number.

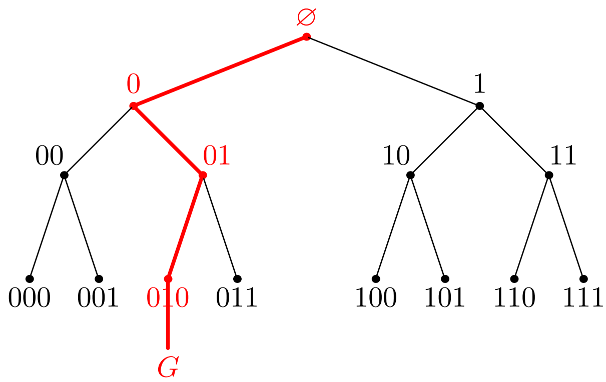

We want to do something similar, but with ω_2 many real numbers instead of just one. In light of this, consider in M the poset

$$\mathbb{P} = \operatorname{Add}(\omega, \omega_2) := (\{p : \omega_2 \times \omega \to 2, \operatorname{dom}(p) \text{ is finite}\}, \supseteq).$$

These elements p (conditions) are "partial functions": we take some finite subset of $\omega_2 \times \omega$ and map it into $2 = \{0, 1\}$. (Here $\operatorname{dom}(p)$ denotes the domain of p, which is the finite subset of $\omega_2 \times \omega$ mentioned.) Moreover, we say $p \le q$ if $\operatorname{dom}(p) \supseteq \operatorname{dom}(q)$ and the two functions agree over $\operatorname{dom}(q)$.

Question 94.1.1. What is the maximum element $1_{\mathbb{P}}$ here?

Exercise 94.1.2. Show that a generic G can be encoded as a function $\omega_2 \times \omega \to 2$.

Lemma 94.1.3 (G encodes distinct real numbers)
For $\alpha \in \omega_2$ define

$$G_\alpha = \{n \mid G(\alpha, n) = 0\} \in \mathcal{P}(\mathbb{N}).$$

Then $G_\alpha \neq G_\beta$ for any $\alpha \neq \beta$.

Proof. We claim that the set

$$D = \{q \mid \exists n \in \omega : q(\alpha, n) \neq q(\beta, n) \text{ are both defined}\}$$

is dense.

> **Question 94.1.4.** Check this. (Use the fact that the domains are all finite.)

Since G is an M-generic it hits this dense set D. Hence $G_\alpha \neq G_\beta$. $\square$

Since $G \in M[G]$ and $M[G] \vDash \mathsf{ZFC}$, it follows that each G_α is in $M[G]$. So there are at least $\aleph_2^M$ real numbers in $M[G]$. We are done once we can show there is no cardinal collapse.

§94.2 The countable chain condition

It remains to show that with $\mathbb{P} = \operatorname{Add}(\omega, \omega_2)$, we have that

$$\aleph_2^{M[G]} = \aleph_2^M.$$

In that case, since $M[G]$ will have $\aleph_2^M = \aleph_2^{M[G]}$ many reals, we will be done.

To do this, we'll rely on a combinatorial property of $\mathbb{P}$:

Definition 94.2.1. We say that $A \subseteq \mathbb{P}$ is a **strong antichain** if for any distinct p and q in A, we have $p \perp q$.

Example 94.2.2 (Example of an antichain)

In the infinite binary tree, the set $A = \{00, 01, 10, 11\}$ is a strong antichain (in fact maximal by inclusion).

This is stronger than the notion of "antichain" than you might be used to![1] We don't merely require that every two elements are incomparable, but that they are in fact *incompatible*.

> **Question 94.2.3.** Draw a finite poset and an antichain of it which is not strong.

> **Question 94.2.4.** Convince yourself that all antichains in the infinite binary tree are strong, but some antichains in the poset $\mathbb{P}$ defined above are not strong.

Definition 94.2.5. A poset $\mathbb{P}$ has the **κ-chain condition** (where κ is a cardinal) if all strong antichains in $\mathbb{P}$ have size less than κ. The special case $\kappa = \aleph_1$ is called the **countable chain condition**, because it implies that every strong antichain is countable.

Remark 94.2.6 (Notational digression: Why $<$ instead of $\le$?) — We could have defined that a poset $\mathbb{P}$ has the κ-chain condition if all strong antichains in $\mathbb{P}$ has size $\le \kappa$. Nevertheless, this alternative definition is less versatile — for instance, there would be no way to express that all strong antichains in $\mathbb{P}$ are finite!

[1] In the context of forcing, some authors use "antichain" to refer to "strong antichain". I think this is lame.

We are going to show that if the poset has the κ-chain condition then it preserves all cardinals $\geq \kappa$. In particular, the countable chain condition will show that $\mathbb{P}$ preserves all the cardinals. Then, we'll show that $\text{Add}(\omega, \omega_2)$ does indeed have this property. This will complete the proof.

We isolate a useful lemma:

Lemma 94.2.7 (Possible values argument)

Suppose M is a transitive model of ZFC and $\mathbb{P}$ is a partial order such that $\mathbb{P}$ has the κ-chain condition in M. Let $X, Y \in M$ and let $f: X \to Y$ be some function in $M[G]$, but $f \notin M$.

Then there exists a function $F \in M$, with $F: X \to \mathcal{P}(Y)$ and such that for any $x \in X$,

$$f(x) \in F(x) \quad \text{and} \quad |F(x)|^M < \kappa.$$

What this is saying is that if f is some new function that's generated, M is still able to pin down the values of f to less than κ many values.

Proof. The idea behind the proof is easy: any possible value of f gives us some condition in the poset $\mathbb{P}$ which forces it. Since distinct values must have incompatible conditions, the κ-chain condition guarantees there are at most κ such values.

Here are the details. Let $\dot{f}$, $\check{X}$, $\check{Y}$ be names for f, X, Y. Start with a condition p such that p forces the sentence

$$\text{“}\dot{f} \text{ is a function from } \check{X} \text{ to } \check{Y}\text{”}.$$

We'll work just below here.

For each $x \in X$, we can consider (using the Axiom of Choice) a maximal strong antichain $A(x)$ of incompatible conditions $q \leq p$ which forces $f(x)$ to equal some value $y \in Y$. Then, we let $F(x)$ collect all the resulting y-values. These are all possible values, and there are less than κ of them. $\square$

§94.3 Preserving cardinals

As we saw earlier, cardinal collapse can still occur. For the Continuum Hypothesis we want to avoid this possibility, so we can add in $\aleph_2^M$ many real numbers and have $\aleph_2^{M[G]} = \aleph_2^M$. It turns out that to verify this, one can check a weaker result.

Definition 94.3.1. For M a transitive model of ZFC and $\mathbb{P} \in M$ a poset, we say $\mathbb{P}$ **preserves cardinals** if $\forall G \subseteq \mathbb{P}$ an M-generic, the model M and $M[G]$ agree on the sentence "κ is a cardinal" for every κ. Similarly we say $\mathbb{P}$ **preserves regular cardinals** if M and $M[G]$ agree on the sentence "κ is a regular cardinal" for every κ.

Intuition: In a model M, it's possible that two cardinals which are in bijection in V are no longer in bijection in M. Similarly, it might be the case that some cardinal $\kappa \in M$ is regular, but stops being regular in V because some function $f: \overline{\kappa} \to \kappa$ is cofinal but happened to only exist in V. In still other words, "κ is a regular cardinal" turns out to be a Π_1 statement too.

Fortunately, each implies the other. We quote the following without proof.

Proposition 94.3.2 (Preserving cardinals $\iff$ preserving regular cardinals)

Let M be a transitive model of ZFC. Let $\mathbb{P} \in M$ be a poset. Then for any λ, $\mathbb{P}$ preserves cardinalities less than or equal to λ if and only if $\mathbb{P}$ preserves regular cardinals less than or equal to λ. Moreover the same holds if we replace "less than or equal to" by "greater than or equal to".

Thus, to show that $\mathbb{P}$ preserves cardinality and cofinalities it suffices to show that $\mathbb{P}$ preserves regularity. The following theorem lets us do this:

Theorem 94.3.3 (Chain conditions preserve regular cardinals)

Let M be a transitive model of ZFC, and let $\mathbb{P} \in M$ be a poset. Suppose M satisfies the sentence "$\mathbb{P}$ has the κ-chain condition and κ is regular". Then $\mathbb{P}$ preserves regularity greater than or equal to κ.

Proof. Use the Possible Values Argument. Problem 94A. □

In particular, if $\mathbb{P}$ has the countable chain condition then $\mathbb{P}$ preserves *all* the cardinals (and cofinalities). Therefore, it remains to show that $\operatorname{Add}(\omega, \omega_2)$ satisfies the countable chain condition.

§94.4 Infinite combinatorics

We now prove that $\operatorname{Add}(\omega, \omega_2)$ satisfies the countable chain condition. This is purely combinatorial, and so we work briefly.

Definition 94.4.1. Suppose C is an uncountable collection of finite sets. C is a **Δ-system** if there exists a **root** R with the condition that for any distinct X and Y in C, we have $X \cap Y = R$.

Lemma 94.4.2 (Δ-System lemma)

Suppose C is an uncountable collection of finite sets. Then $\exists \overline{C} \subseteq C$ such that $\overline{C}$ is an uncountable Δ-system.

Proof. There exists an integer n such that C has uncountably many guys of length n. So we can throw away all the other sets, and just assume that all sets in C have size n.

We now proceed by induction on n. The base case $n = 1$ is trivial, since we can just take $R = \varnothing$. For the inductive step we consider two cases.

First, assume there exists an $a \in C$ contained in uncountably many $F \in C$. Throw away all the other guys. Then we can just delete a, and apply the inductive hypothesis.

Now assume that for every a, only countably many members of C have a in them. We claim we can even get a $\overline{C}$ with $R = \varnothing$. First, pick $F_0 \in C$. It's straightforward to construct an F_1 such that $F_1 \cap F_0 = \varnothing$. And we can just construct $F_2, F_3, \dots$ □

Lemma 94.4.3

For all κ, $\operatorname{Add}(\omega, \kappa)$ satisfies the countable chain condition.

Proof. Assume not. Let

$$\{p_\alpha \mid \alpha < \omega_1\}$$

be a strong antichain. Let

$$C = \{\mathrm{dom}(p_\alpha) \mid \alpha < \omega_1\}.$$

Let $\overline{C} \subseteq C$ be such that $\overline{C}$ is uncountable, and $\overline{C}$ is a Δ-system with root R. Then let

$$B = \{p_\alpha \mid \mathrm{dom}(p_\alpha) \in R\}.$$

Each $p_\alpha \in B$ is a function $p_\alpha \colon R \to \{0, 1\}$, so there are two that are the same. □

Thus, we have proven that the Continuum Hypothesis cannot be proven in ZFC.

§94.5 A few harder problems to think about

Problem 94A. Let M be a transitive model of ZFC, and let $\mathbb{P} \in M$ be a poset. Suppose M satisfies the sentence "$\mathbb{P}$ has the κ-chain condition and κ is regular". Show that $\mathbb{P}$ preserves regularity greater than or equal to κ.