Recall that informally, a *class* is a subcollection of the universe. To be more formal, each class is identified with a formula $\phi(x)$ defining the class (so that $C = \{x : \phi(x)\}$).

Definition 3.25. *A class $C \subseteq ON$ of ordinals is a club if it is closed and unbounded in the order topology on the ordinals. To say that it is unbounded means that for every ordinal α there exists some larger ordinal β with $\beta \in C$. To say that it is closed means that if $\langle \alpha_\xi : \xi < \lambda \rangle$ is an increasing enumeration of a subset of C, then $\sup_{\xi<\lambda} \alpha_x i$ also belongs to C.*

For example, the class of limit ordinals is a club. The class of cardinals is a club. The class of regular cardinals is not a club; this is because (for example) $\aleph_\omega$ is a limit point of this class, being the limit of the $\aleph_n$ where $n < \omega$. Since the following proposition refers to clubs, formally speaking it is a schema of propositions rather than a single theorem proven in ZFC.

Proposition 3.26. *The intersection of two clubs is a club.*

Proof. Let C_0 and C_1 be clubs. Let us show that $C = C_0 \cap C_1$ is a club. It is not hard to see that C is closed for if $\langle \alpha_\xi : \xi < \lambda \rangle$ is an increasing enumeration of a subset of C with limit α, then each $\alpha_x i$ belongs to C_0 and so $\alpha \in C_0$ by virtue of C_0 being a club. Similarly $\alpha \in C_1$ and so $\alpha \in C$ and C is closed.

For the unbounded part, let α be some given ordinal. We seek a β with $\alpha < \beta$ and $\beta \in C$. Since C_0 is a club we can find $\beta_0 \in C_0$ larger than α. Since C_1 is a club we can find $\beta_1 \in C_1$ larger than β_0. Proceed inductively so that we get an increasing sequence $\langle \beta_n : n \in \omega \rangle$ with $\beta_{2n} \in C_0$ and $\beta_{2n+1} \in C_1$. Set $\beta = \lim_{n<\omega} \beta_n$. Then $\beta = \lim_{n<\omega} \beta_{2n}$ and so β belongs to C_0. But also $\beta = \lim_{n<\omega} \beta_{2n+1}$ and so β belongs to C_1. And so $\beta \in C$. □

Proposition 3.27. *Let $G : ON \to ON$ be an increasing function. Then the set $C = \{\beta : (\forall \alpha < \beta) G(\alpha) < \beta\}$ is a club.*

Proof. As usual with these kinds of arguments, checking the closure is the easy part. For suppose $\beta = \lim_{\xi<\lambda} \beta_\xi$ where each β_ξ belongs to C. Let us check that $\beta \in C$. Take $\alpha < \beta$. Then there is some ξ such that $\alpha < \beta_\xi$. Since $\beta_\xi \in C$ it follows that $G(\alpha) < G(\beta_\xi)$. Since $G(\beta_\xi) < G(\beta)$ (since G is an increasing function) it follows that $G(\alpha) < G(\beta)$.

Let us next check that C is unbounded. Fix β_0, an ordinal. We want a $\beta > \beta_0$ with $\beta \in C$. Assume inductively that β_n is given. Set β_{n+1} to be the supremum of $\{G(\alpha) : \alpha < \beta_n\}$. Then let $\beta = \lim_{n<\omega} \beta_n$. We verify that $\beta \in C$. Suppose $\alpha_0 < \beta$. Then $\alpha_0 < \beta_n$ for some n. So by definition of β_{n+1}, $G(\alpha_0) < \sup\{G(\alpha) : \alpha < \beta_n\} = \beta_{n+1} < \beta$. So indeed $\beta \in C$. □

Now we can state the following, which is somewhat stronger than what we will need to get our finite fragments. As before, formally speaking the following theorem is a schema for a collection of theorems each provable in ZFC, rather than a single theorem itself.

Theorem 3.28. *Let $\phi(x_1, \ldots x_n)$ be a formula in the language of set theory with n free variables. There is a club C_ϕ such that whenever $\alpha \in C_\phi$, and $a_1, \ldots a_n$ belong to V_α, then*

$$V_\alpha \vDash \phi(a_1, \ldots a_n) \text{ if and only } \phi(a_1, \ldots a_n).$$

Proof. We do this by induction on the complexity of ϕ. If ϕ is an atomic formula (or in fact any quantifier free formula) then ϕ is absolute and so the condition above will hold regardless of α. In this case we can take C_ϕ to be the club of all ordinals.

Now for our inductive steps. Say $\phi = \neg\psi$, and inductively C_ψ has been defined. Since

$$V_\alpha \vDash \neg\phi(a_1, \ldots a_n) \text{ if and only } V_\alpha \nvDash \phi(a_1, \ldots a_n) \text{ if and only } \neg\phi(a_1, \ldots a_n)$$

we can in fact just take $C_\phi = C_\psi$.

Suppose $\phi = \chi \wedge \psi$. Then we may take $C_\phi = C_\chi \cap C_\psi$, which is also a club.

The last case to consider is $\phi = \exists x\psi(x)$. Define a class function whose domain is the set of *all* n-tuples by letting

$$F(a_1, \dots a_n) = \begin{cases} \beta : & \text{where } \beta \text{ is minimal such that } V_\beta \vDash \phi(a_1, \dots a_n) \\ 0 : & \text{if no such } \beta \text{ exists.} \end{cases}$$

Then define $G : \mathrm{ON} \to \mathrm{ON}$ by setting $G(\alpha) = \sup_{a_1,\dots a_n \in V_\alpha} F(a_1, \dots a_n)$. Clearly G is an increasing function, and hence $C = \{\beta : (\forall \alpha < \beta) G(\alpha) < \beta\}$ is a club. We claim that $C_\phi = C \cap C_\psi$ is the club we want.

Suppose $\gamma \in C_\phi$ and $a_1, \dots a_n \in V_\gamma$. Then

$$V_\gamma \vDash \phi(a_1, \dots a_n) \leftrightarrow V_\gamma \vDash \exists\, x\psi(x, a_1, \dots a_n)$$

$$\leftrightarrow \text{ there is a } b \in V_\gamma \text{ such that } V_\gamma \vDash \psi(b, a_1, \dots a_n)$$

$$\leftrightarrow \text{ there is a } b \in V_\gamma \text{ such that } \psi(b, a_1, \dots a_n) (\text{ since } \gamma \in C_\psi)$$

$$\leftarrow \text{ there is a } b \in V_{F(a_1,\dots a_n)} \text{ such that } \psi(b, a_1, \dots a_n)$$

$$\leftrightarrow \text{ there is a } b \text{ such that } \psi(b, a_1, \dots a_n) \text{ (by definition of } F)$$

$$\leftrightarrow \phi(a_1, \dots a_n).$$

The leftwards implication from the third-to-last line to the fourth-to-last line holds because $F(a_1, \dots a_n) < \gamma$, since $a_1, \dots a_n \in V_\beta$ for some $\beta < \gamma$ and $\gamma \in C$. In fact, these these lines give an if and only if, since the second to last line is implied by the fourth to last. □

Corollary 3.29. *If ϕ is a finite fragment of ZFC, ie ϕ is a finite conjunction of some of the axioms of ZFC, then there is an α such that $V_\alpha \vDash \phi$.*

Proof. Notice first of all that ϕ has no free variables, and secondly that ϕ holds (in the universe) since we always work inside ZFC. Let C_ϕ be the club as described in the statement of Theorem 3.28. If $\alpha \in C_\phi$, then $V_\alpha \vDash \phi$ if and only if ϕ holds. Since ϕ holds, $V_\alpha \vDash \phi$. □

One consequence of this that one can prove immediately is that ZFC is not finitely axiomatizable (unless it is inconsistent); that is there is no finite collection of axioms which is equivalent to the ZFC axioms. Suppose for contradiction that ϕ were the conjunction of such axioms; then ZFC $\vdash \phi$ and $\phi \vdash$ ZFC. Now we have seen that within ZFC there is a model of ϕ. Since one can prove within ZFC that satisfiable sentences are consistent, we have that ZFC $\vdash$ CON(ϕ). But ZFC also knows that $\phi \vdash$ ZFC, and that the logical consequences of consistent sentences are themselves consistent. So ZFC $\vdash$ CON(ZFC). But that contradicts Gödel's Second Incompleteness Theorem, unless ZFC is inconsistent.

That is a nice fact, but it isn't the one we were looking for. What we really care about is finding countable transitive models of fragments of ZFC. Corollary 3.29 gives us transitive models (as V_α is always transitive), but unfortunately we do not quite yet have countable ones. So we need a little bit more technology.

We earlier gave several examples of models using the actual $\in$ relation which were not transitive. It turns out that this not such a barrier. Recall that two models M and N are isomorphic if there is a bijection $f : M \to N$ such that $m_1 \in m_2$ if and only if $f(m_1) \in f(m_2)$.

Theorem 3.30. *Let M be a given model such that M satisfies extensionality. There is a transitive model N which is isomorphic to M.*

Proof. The technique we use here is what is known as the *Mostowski collapse*. In general it does not matter that our M used the actual $\in$ relation as its interpretation of membership; the only thing actually required of M is that its interpretation of membership be a well-founded relation.

We define a map $\pi : M \to V$ by ϵ-induction; this is possible because ϵ is a well-founded relation. Set $\pi(u) = \{\pi(v) : v \in M \cap u\}$ and then set $N = \{\pi(u) : u \in M\}$.

First of all we claim that N is a transitive set. Suppose $y \in N$ and that $x \in y$. By definition of N, $y = \pi(u)$ for some $u \in M$. Then, by definition of π it must be that $x = \pi(v)$ for some $v \in M \cap u$. Since $\pi(v)$ belongs to N by definition of M, we have $x \in N$ and N is indeed transitive.

Next we claim that π is an isomorphism. Certainly it is surjective, as we literally chose N to be the image of π under M. If $u \in v$ both belong to M, then the definition of π makes it clear that $\pi(u) \in \pi(v)$.

To finish we must show that π is injective. Notice that we have not yet made use of the fact that M satisfies extensionality. We do so now. We will show that $\pi(u_1) = \pi(u_2)$ implies $u_1 = u_2$, by ϵ-induction. So assume that $\pi(u_1) = \pi(u_2)$, and assume inductively that if $\pi(v_1) = \pi(v_2)$ and $\pi(v_1) \in \pi(u_1)$ then $v_1 = v_2$. We must show that $u_1 = u_2$. By definition of π we have that $\{\pi(v) : v \in M \cap u_1\} = \{\pi(v) : v \in M \cap u_2\}$.

Now, since M satisfies extensionality it is enough to show that $M \vDash (\forall z) z \in u_1 \leftrightarrow z \in u_2$; for then we will have $M \vDash u_1 = u_2$ which is the same as having $u_1 = u_2$. We must show $(\forall z \in M) z \in u_1 \to z \in u_2$ and $(\forall z \in M) z \in u_2 \to z \in u_1$. Appealing to symmetry we just show the first. Suppose $z \in M \cap u_1$. Then $\pi(z) \in \pi(u_1)$. It follows that $\pi(z) \in \pi(u_2)$. By definition of π there is some $y \in u_2 \cap M$ such that $\pi(z) = \pi(y)$. By our ϵ induction we have $z = y$. So $z \in u_2$ as desired.

So π is injective and the proof is complete. □

Now it is a basic fact of first order logic that isomorphic structures satisfy the same sentences. So, by the above theorem, in order to get a countable transitive N which satisfies ϕ, we just to find a countable M which satisfies ϕ as then its Mostowski collapse will be the desired N. We know by Corollary 3.29 that for any finite fragment ϕ of ZFC there is some α such that $V_\alpha \vDash \phi$. If we take a countable elementary substructure $M \subseteq V_\alpha$ then we have a countable M such that $M \vDash \phi$.

Thus we have established the following.

Theorem 3.31. *For any $\phi_1, \ldots \phi_n$ a finite fragment of ZFC, there is a countable transitive model M such that $M \vDash \phi_1 \wedge \ldots \phi_n$.*

4 Forcing

4.1 $M[G]$

Definition 4.1. *Let M be a countable transitive model of ZFC. Let $\mathbb{P}$ be a poset with $\mathbb{P} \in M$. A filter G is $\mathbb{P}$-generic over M (or just $\mathbb{P}$-generic when M is understood from context, as will usually be the case) if for every set $D \in M$ which is dense in $\mathbb{P}$ we have that $G \cap D \neq \varnothing$.*

Lemma 4.2. *Let M be a countable transitive model of ZFC with $\mathbb{P} \in M$. Then there is a $\mathbb{P}$-generic filter G. In fact, for any $p \in \mathbb{P}$ there is a $\mathbb{P}$-generic filter G which contains p.*

Proof. Since M is countable, getting a $\mathbb{P}$-generic filter G is the same as finding a $\mathcal{D}$-generic filter G where

$$\mathcal{D} = \{D \in M : D \text{ is dense }\}.$$

Since MA(ω) always holds such a filter exists. If we want to ensure that $p \in \mathbb{G}$ we use the same proof as that of MA(ω), starting our construction at p. □

Let us give a few motivating words.

Suppose we wanted to construct a model of CH, and we had given to us a countable transitive M, a model of ZFC. Now M satisfies ZFC, so within M one may define the partial order $\mathbb{P}$ consisting of all countable approximations to a function $f : \omega_1 \to \mathcal{P}(\omega)$. Of course M is countable, so the things that M believes are ω_1 and $\mathcal{P}(\omega)$ are not actually the real objects. But for each $X \in \mathcal{P}(\omega)^M$ the set $D_X = \{p \in \mathcal{P} : X \in \mathrm{ran}(p)\}$ is dense, as is the set $E_\alpha = \{p \in \mathbb{P} : \alpha \in \mathrm{dom}(p)\}$ for each $\alpha < \omega_1^M$. So a $\mathbb{P}$-generic filter G will intersect each of those sets, and will by the usual arguments yield a surjection $g : \omega_1^M \to \mathcal{P}(\omega)^M$. Thankfully, by the previous lemma, such a G exists. Unfortunately there is no reason for us to believe that this G belongs to M. What we now learn is how to force the generic into model M without doing too much damage to the universe of M.

Given any poset $\mathbb{P}$ in M, and a $\mathbb{P}$-generic filter G, the method of forcing will give us a way of creating a new countable transitive model $M[G]$ satisfying ZFC that extends M and contains G. Now just getting such a model is not enough. For in the example above the surjection $g : \omega_1 \to \mathcal{P}(\omega)$ defined from G was a mapping between the objects in M. But a priori it may well be that the model $M[G]$ has a different version of ω_1 and a different version of $\mathcal{P}(\omega)$ and so the CH still would not be satisfied. It turns out that in this (and many other cases) the forcing machinery will work out in our favor, and these things will not be disturbed.

It is worth pointing out that when $\mathbb{P} \in M$ then the notion of being a partial order, or being dense in $\mathbb{P}$ are absolute (written out the formulas just involve bounded quantifiers over $\mathbb{P}$). So if $D \in M$ then $M \vDash$ "D is dense"' exactly when D really is dense. Thus the countable set $\{D \in M : \text{D is dense}\}$, is exactly the same collection defined in M to be the collection of *all* dense subsets of $\mathbb{P}$. Unless $\mathbb{P}$ is something silly this will not actually be all the dense subsets, since M will be missing some. Let us isolate a class of not-silly posets.

Definition 4.3. *A poset $\mathbb{P}$ is* separative *if (1) for every p there is a q which properly extends p (ie $q < p$) and (2) whenever $p \not\leq q$ then there is an $r \leq p$ with $q \perp r$.*

Definition 4.4. *A poset $\mathbb{P}$ is* nonatomic *if for any $p \in \mathbb{P}$ there exist $q, r \leq p$ which are incomparable.*

Essentially every example of a poset that we have used thus far is separative. Notice that every separative poset is nonatomic.

Proposition 4.5. *Say $\mathbb{P}$ is* nonatomic *and $\mathbb{P} \in M$. Let G be $\mathbb{P}$-generic. Then $G \notin M$.*

Proof. Consider the set $D = \mathbb{P} \setminus G$. Then D belongs to M. Let us see that D is dense. Let $p \in \mathbb{P}$ be arbitrary. Since $\mathbb{P}$ is separative there are $q, r \leq p$ which are incomparable. Since G is a filter, at most one of them can belong to G and whichever one does not belongs to D.

Since D is dense and G is $\mathbb{P}$-generic, G should intersect D. But that is ridiculous. □

Now we will show how, given G and M, to construct $M[G]$. Clearly the model M will not know about the model $M[G]$, since G can not be defined within M. But it will be the case that this is the only barrier. All of the tools to create $M[G]$ can assembled within M itself; only a generic filter G is needed to get them to run.

Definition 4.6. *We define the class of $\mathbb{P}$-names by defining for each α the $\mathbb{P}$-names of name-rank α. (For a $\mathbb{P}$-name τ we will use $\rho(\tau)$ to denote the name-rank of τ). The only $\mathbb{P}$-name of name-rank 0 is the empty set $\varnothing$. And recursively, if all the $\mathbb{P}$-names of name-rank strictly less than α have been defined, we say that τ is a $\mathbb{P}$-name of name-rank α if every $x \in \tau$ is of the form $x = \langle \tau, p \rangle$ where τ is a $\mathbb{P}$-name and $p \in \mathbb{P}$.*

Another way of stating the definition is just to say that a set τ of ordered pairs is called a $\mathbb{P}$-name if it satisfies (recursively) the following property: every element of τ has the form $\langle \sigma, p \rangle$ where σ is itself a $\mathbb{P}$-name and p is an element of $\mathbb{P}$.

It is not hard to see that the notion of being a $\mathbb{P}$-name is absolute; that is, $M \models$ "τ is a $\mathbb{P}$-name$''$ exactly when τ is a $\mathbb{P}$-name. This is because the concept is defined by transfinite recursion from absolute concepts. As another piece of notation, since τ is a set of ordered pairs, it makes sense to use $\mathrm{dom}(\tau)$ as notation for all the σ occurring in the first coordinate of an element of τ.

Definition 4.7. *If M is a countable transitive model of ZFC, then $M^{\mathbb{P}}$ denotes the collection of all the $\mathbb{P}$-names that belong to M.*

Alone the $\mathbb{P}$-names are just words without any meaning. The people living in M have the names but they do not know anyway of giving them a coherent meaning. But once we are have a $\mathbb{P}$-generic filter G at hand, they can be given values.

Definition 4.8. *Let τ be a $\mathbb{P}$-name and G a filter on $\mathbb{P}$. Then the value of τ under G, denote $\tau[G]$, is by recursive definition the set*

$$\{\sigma[G] : \langle \sigma, p\rangle \in \sigma \text{ and } p \in G\}.$$

With this definition in mind, one can think of as an element $\langle \sigma, p\rangle$ of a $\mathbb{P}$-name τ as saying that $\sigma[G]$ has probability p of belonging to $\tau[G]$. The fact that we are calling the maximal element of our posets $\mathbb{1}$ makes this all the more suggestive, for $\mathbb{1}$ belongs to every filter G. So in particular, whatever G is, if we have $\tau = \{\langle \varnothing, \mathbb{1}\rangle\}$ then $\tau[G] = \{\varnothing\}$. On the other hand if $\tau = \{\langle \varnothing, p\rangle\}$ for some p that does not belong to G then $\tau[G] = \varnothing$.