FORCING AND THE CONTINUUM HYPOTHESIS

SPENCER DEMBNER

ABSTRACT. We rigorously develop the technique of forcing, invented by Paul Cohen to generate new models of set theory. First we state the required logic, model theory, set theory and order theory. Then we introduce the concept of a forcing poset and a generic filter over a poset, and explain how to construct the generic extension of a model. After verifying that generic extensions are models of set theory, we use the technique to verify both directions of the independence of the continuum hypothesis.

CONTENTS

1. INTRODUCTION

The Continuum Hypothesis (CH) is the statement that $2^{\aleph_0} = \aleph_1$ — in other words, that there is no cardinality strictly between the cardinality of the natural numbers and that of the real numbers, which is the same as that of the powerset of the natural numbers. Godel showed in 1940 that the Continuum Hypothesis is consistent with the axioms of Zermelo-Fraenkel Set Theory with the Axiom of Choice (ZFC). His proof involves constructing an inner model L of ZFC, known as the Constructible Universe, and showing that the Continuum Hypothesis holds in L. In 1963, Paul Cohen demonstrated that the negation of CH is also consistent with ZFC, completing the proof that CH is independent from ZFC.

Date: August 26, 2018.

Cohen's basic technique, known as forcing, starts with a ground model M of ZFC, then augments it with a so-called "generic" set G. The resulting generic extension, $M[G]$, will be another model of ZFC, containing every element of the model M, but possessing new properties that depend on the choice of generic set. In the case of the Continuum Hypothesis, we are effectively able to "add" a large number of additional subsets of $\mathbb{N}$, creating a model in which the powerset $2^{\aleph_0}$ has, for example, a cardinality of at least $\aleph_2$, and thus modeling $\neg CH$. By choosing a different generic set, we can effectively generate bijections between distinct cardinalities, creating a model that "collapses" the larger cardinal onto the smaller one. In fact, we can use forcing to collapse the real numbers onto $\aleph_1$, obtaining an easy proof in the other direction that $ZFC + CH$ is consistent.

In this paper, we develop the necessary theory in order to state and prove the independence of the continuum hypothesis, culminating in theorems 4.22 and 4.25. First, we develop enough basic set theory to justify the presentation of forcing. Then we define the concept of a generic set and describe how to construct generic extensions of a model. After proving several key results about the behavior of forcing extensions, we verify that they satisfy the axioms of ZFC and prove the key results about the cardinality of the continuum. Our exposition largely follows Kunen's approach in [7]. However, we also draw on [5] and [3] for the framing of some results.

2. Preliminaries

We assume a basic familiarity with first-order predicate logic and basic set theory, including the construction of ordinals and basic notions of cardinality (see [1] for details on any of these). For the purpose of this paper we always assume that our language includes a notion of equality.

Aside from logical connectives, quantifiers and the $=$ sign, the only other symbol in the language of set theory is a two-place membership relation, $\in$. In particular, the subset relation is a derived notion and not a primitive one. Because it only contains finitely many symbols, the language of set theory is countable.

2.1. Basics of ZFC.

Definition 2.1. ZF is the first-order theory with the following axioms (we state them in natural language for simplicity):

- **Extensionality:** For any two sets X and Y, $X = Y$ if and only if X and Y have the same elements.
- **Pairing:** Given any sets a and b, there exists a set $\{a, b\}$ which contains exactly a and b as elements.
- **Axiom Schema of Comprehension:** Consider any first-order formula $\varphi(x, y)$, where x occurs freely and y, a parameter, may or may not occur freely. Then for any set X and any set p, there exists the set $Y = \{x \in X \mid \varphi(x, p)\}$. This is the subset of X defined by the formula φ and the parameter p.
- **Union:** For any set X there exists a set $Y = \cup X$, which contains exactly the elements which are a member of some set x, where $x \in X$.
- **Power Set:** For any set X, there exists $Y = \mathcal{P}(X)$, the power set of X. The elements of Y are exactly the subsets of X (where $A \subset B$ means that $x \in A \Longrightarrow x \in B$).

Infinity: There exists an inductive set E. This means that E contains the empty set $\emptyset$, and that for any $x \in E$, the set $x \cup \{x\}$ is a member of E.

Axiom Schema of Replacement: Let $\varphi(x, y)$ define a function F on the set A, in the sense that for every $x \in A$, exactly one set y satisfies $\varphi(x, y)$. Writing $F(x) = y$ for $\varphi(x, y)$, there exists a set $F(A)$, the image of A, which contains every $f(x)$ for $x \in A$ and has no other elements.

Regularity: For any set X, there exists some $y \in X$ such that y and X are disjoint.

The final axiom is the Axiom of Choice, which is usually dealt with separately in the context of independence proofs.

Axiom 2.2 (Axiom of Choice)**.** Given a set X all of whose elements are nonempty, a choice function F exists that picks one element from each member of X. ZF with this final axiom is referred to as *ZFC*.

Note that this set of axioms is countably infinite, since it contains one instance of comprehension for every formula and one instance of replacement for every function. In fact, no finite axiomatization of ZF is possible, because of results related to Godel's Incompleteness Theorem that we will state below.

We should make some observations regarding the Axiom of Infinity. The natural numbers in ZFC are identified with the finite ordinals, according to the following construction:

$$0 = \emptyset, 1 = \{\emptyset\}, 2 = \{\emptyset, \{\emptyset\}\}, \ldots, n = n - 1 \cup \{n - 1\}, \ldots$$

It's easy to see that every such natural number n must be a member of the inductive set E which Infinity asserts to exist — otherwise, consider the first natural number not contained in E. The set of all natural numbers, ω, is itself an ordinal, the first transfinite ordinal. However, strictly speaking, the Axiom of Infinity only asserts that E is a superset of ω, not that $\omega = E$. We can recover ω from E by a suitable example of the Comprehension Axiom.[1] For specifics, see section I.12 of [5].

Finally, note that Infinity justifies the existence of the empty set.

2.2. Lévy Reflection and "Nice" Models of ZFC.

A *model of ZFC* is simply a first-order structure satisfying the nine axioms above. However, forcing is much easier to develop if we make certain simplifying assumptions about our model — namely, that it is a so-called countable transitive model. Our next goal is to develop and justify these assumptions.

Definition 2.3. Let M be some class which may or may not be a set. We say M is *transitive* if for any $m \in M$, m is a subset of M, or equivalently if $\cup M \subset M$. In particular, M is a countable transitive model of ZFC if M is a model of ZFC, M is countable and M is transitive.

Working within ZFC itself, we can't prove the existence of a countable transitive model. However the Reflection Principle, explained below, will effectively allow us to use such a model without assuming its existence. First, using the following basic results about logic, we can verify that ZFC has a countable model if its axioms are consistent. They can be found in any standard reference, such as [1].

---

[1] Note that some models of ZFC have nonstandard natural numbers, but as we will verify later, if a model is transitive then its ω must be the "real" ω.

Theorem 2.4 (Completeness of First-Order Logic)**.** *Consider any collection* Γ *of first-order sentences in a language* $\mathscr{L}$*.* Γ *is consistent if and only if it has a model whose domain is a set (a set model).*

Corollary 2.5 (Compactness)**.** *A collection* Γ *of first-order sentences in a language* $\mathscr{L}$ *has a set model if and only if every finite* $\gamma \subset \Gamma$ *has a set model.*

Corollary 2.6 (Downward Lowenheim-Skolem Theorem)**.** *Let* Γ *be a collection of first-order sentences in a language* $\mathscr{L}$*. If* Γ *has any infinite models, then it has a model with cardinality* $|\mathscr{L}|$*.*

Here we follow the convention that if $\mathscr{L}$ is finite, then $|\mathscr{L}| = \aleph_0$, the cardinality of the natural numbers. In particular, if ZFC is consistent, then it must have a model which is a set. By the Downward Lowenheim-Skolem Theorem, there must be some such model which is countable.

On the other hand, it turns out that the existence of a transitive model of ZFC is stronger than the consistency of ZFC, which only implies the existence of some (not necessarily transitive) model.[2]

Another obstacle is provided by Godel's Second Incompleteness Theorem, which prevents us from proving the existence of a model of ZFC within ZFC itself.

Theorem 2.7 (Godel)**.** *Consider any formal system* $\mathscr{F}$ *capable of formalizing Peano Arithmetic which has a recursively enumerable set of axioms. If* $\mathscr{F}$ *is consistent, then* $\mathscr{F}$ *cannot prove its own consistency.*

It's a straightforward exercise to verify that ZFC is a formal system satisfying the requirements of Theorem 2.7 — its set of axioms is recursive, and the finite ordinals as defined above can be proven to satisfy the relevant axioms. By the Completeness Theorem, consistency and the existence of a model are equivalent. It follows that ZFC *cannot prove that it has a model*, or equivalently, that not every model of ZFC contains a set which it believes to be a model of ZFC.

However, ZFC *can* prove the existence of sets that satisfy arbitrarily large finite fragments of its axioms, and for our purposes this is good enough. In order to make this rigorous, we need to decide what it means for a sentence to be "true within a set."

Definition 2.8. Let $\varphi(x_1, x_2, \ldots, x_n)$ be a formula in the language of set theory, and let M be a set in (or a subclass of) a given model of ZFC. The *relativization*

[2] One proof of this is as follows (thanks to Henning Makholm for pointing this out on Stackexchange). Let T be the formal theory $\mathrm{ZFC} + \mathrm{Con}(\mathrm{ZFC})$, and suppose T is consistent, which amounts to assuming that ZFC has a model. By the incompleteness theorem, T must not prove its own consistency: therefore, there must be a model of the theory $T + \neg\,\mathrm{Con}(T)$, or in other words a model which believes T is true but believes that no model of T exists. By the definition of T, this model must believe that ZFC has a model, but it also believes that no model of ZFC can believe that ZFC is consistent. Therefore it must believe that all of the models of ZFC that it includes contain nonstandard integers, encoding fake "proofs" that ZFC is inconsistent. Furthermore, nonstandard natural numbers are never well-founded, so the model must believe the "internal" model is not transitive (since the model believes itself to be well-founded). So this is a universe in which ZFC has a model, but none of its models are transitive.

of φ *to* M, denoted φ^M, is defined by the following recursion:

$$(x = y)^M \Leftrightarrow x = y$$
$$(x \in y)^M \Leftrightarrow x \in y$$
$$(\neg\varphi)^M \Leftrightarrow \neg(\varphi^M)$$
$$(\varphi \wedge \psi)^M \Leftrightarrow \varphi^M \wedge \psi^M$$
$$(\exists x\varphi)^M \Leftrightarrow \exists x \in M(\varphi^M)$$

Here and in the rest of the paper, we limit our connectives and quantifiers to $\exists, \neg$, and the conjunction $\wedge$. All others can be written in terms of those three.

Definition 2.9. Consider some set or subclass M, as above. We say that the formula φ is *absolute* for M if for any $x_1, x_2, \ldots, x_n \in M$, we have the equivalence $\varphi(x_1, x_2, \ldots, x_n) \Leftrightarrow \varphi^M(x_1, x_2, \ldots, x_n)$.[3]

A *bounded quantifier* is a quantifier of the form $\exists x \in y$. We say that a formula is Δ_0 if its only quantifiers are bounded. The following lemma is standard, but see for example section I.12 of [5].

Lemma 2.10. *Let* φ *be a* Δ_0 *formula. Then* φ *is absolute for any transitive class* M.

Proof. We work by induction on the complexity of φ. It follows from the definition of relativization that atomic formulas are absolute, and if two formulas φ and ψ are absolute, we can verify that their conjunctions and negations are also absolute, again working directly from the definition. Therefore, consider the case where φ is of the form $(\exists x \in y)\psi$ and ψ is absolute for M.

To be explicit: we write the relativization φ^M as $(\exists x \in M)[x \in y \wedge \psi^M]$, moving the condition $x \in y$ within the quantifier. First, suppose φ holds: then since y is free, we have $y \in M$, and by transitivity, $x \in M$ for any $x \in y$. This proves that φ^M holds as well, since ψ is already assumed to be absolute. On the other hand, suppose φ^M holds: then there exists some $x \in M$ such that $x \in y$ and ψ^M holds for x. Since ψ is absolute, ψ must also hold for x, and this gives us φ. □

Example 2.11. There exists a Δ_0 formula expressing the property "x is an ordinal" — the details are unenlightening and we suppress them. This means that all transitive models will agree about whether a given set is an ordinal.[4] In particular, the ordinals in any transitive model are well-founded.

We are now in a position to state the key result of this section.

Theorem 2.12 (Lévy Reflection)**.** *Consider any finite collection of formulas* $\varphi_1, \varphi_2, \ldots, \varphi_n$. *Within any model of ZFC, there exists a transitive set* M *such that* M *is absolute for every member of the collection — in other words,* $\varphi_j^M \Leftrightarrow \varphi_j$*, for any* $1 \leq j \leq n$.

---

[3]The usual convention is to list free variables in a formula in parentheses. So the formula $\varphi(x_1, x_2, \ldots, x_n)$ has the variables $x_1, x_2, \ldots, x_n$ occurring freely. As always, a sentence is a formula in which no variable occurs freely.

[4]Of course, different transitive models may have different ordinals: a countable model will have countably many ordinals and so on. But given two transitive models M and N, either the ordinals of M are an initial segment of the ordinals of N, or vice versa.

In particular, this means that for any finite subset of the axioms of ZFC, some transitive set M models the truth of those axioms.

To prove Theorem 2.12, we will need to introduce a new construction, effectively verifying that our set universe is well-founded.

Definition 2.13. Let γ be some ordinal. We define the Von Neumann Universe up to γ, denoted V_γ, by the following recursion:

$$V_0 = \emptyset$$
$$V_{\gamma+1} = \mathcal{P}(V_\gamma), \text{ where } \mathcal{P} \text{ denotes the power set}$$
$$V_\gamma = \cup_{\lambda<\gamma} V_\lambda, \text{ if } \gamma \text{ is a limit ordinal}$$

We can informally write $V = \cup_{\lambda \in Ord} V_\lambda$ for the entire Von Neumann Universe (also called the Cumulative Hierarchy), but note that V is a proper class and not a set. Also note that each stage V_γ is transitive, and hence that $V_\gamma \subset V_{\gamma+1}$ for any γ.

Lemma 2.14. *Any set p has a transitive closure, that is, a minimal set $TC(p)$ such that $TC(p) \supset p$ and $\cup TC(p) \subset TC(p)$. Furthermore, the transitive closure is unique.*

Proof. Let $p_0 = p$, then define $p_{n+1} = \cup p_n$ for each n. In other words, at each stage we take the union of the set constructed at the previous stage. Defining $s = \bigcup_{i=0}^{\infty} p_n$, it's immediate that $s \supset p$. Also, consider any $x \in s$, and any $y \in x$. There must be some n such that $x \in p_n$, and then it follows that $y \in p_{n+1} \subset s$. Therefore $x \subset s$, and the set s is transitive.

We can prove that this is the minimal such transitive set, and hence that it must be unique, by applying induction. Let C be any transitive set containing p. Then by definition $p_0 \subset C$. Suppose that $p_n \subset C$: then since C is transitive, we must have that $\cup p_n = p_{n+1} \subset C$. It follows that $s \subset C$, and hence that the transitive closure is unique. □

Lemma 2.15. *Given any nonempty* class *C, C has an $\in$-minimal element x — that is, an element $x \in C$ such that $x \cap C = \emptyset$.*

This extends the Axiom of Regularity to proper classes as well as sets.

Proof. Suppose not, and pick some $x \in C$. By our assumption, $x \cap C \neq \emptyset$. Consider the transitive closure $TC(x)$: since $TC(x) \supset x$, we likewise have that $TC(x) \cap C$ is nonempty. Label this set Y (the fact that Y is a set follows from the axiom schema of comprehension). By the axiom of regularity, Y has a minimal element y, such that $y \cap Y = \emptyset$. But then by the definition of Y, we have that $y \in C$. Also, $y \cap C = \emptyset$: otherwise if $z \in y$, $z \in C$, then $z \in TC(x)$ by transitivity, so $z \in TC(x) \cap C = Y$, contradicting our assumption that y is minimal in Y. So y must be an $\in$-minimal element of the class C. □

Theorem 2.16. *Given any model M of ZFC, $M = V$. In other words, for any set X there exists an ordinal γ with $X \in V_\gamma$.*

Proof. If this result fails, then the class C of sets not contained in V is nonempty. It follows that there must be some $\in$-minimal element $c \in C$: for any $x \in c$, $x \notin C$, so we have $x \in V$. In other words, every element of the set c is contained in some level of the hierarchy V.

For any set x, we can define the *rank* of x to be the minimal γ such that $x \in V_\gamma$. Then, let $P = \{\text{rank}(x) \mid x \in c\}$: c is a set and P is the image of c under a function, so by Replacement, P must be a set as well.

We want to show that there exists some ordinal ζ which bounds the set P from above, demonstrating that $c \subset V_\zeta$. If no such ζ existed, then P would be *cofinal* with the class of all ordinals, ON, in the sense that for any $\xi \in ON$, some $\psi \in P$ exists where $\psi \geq \xi$. But by the Burali-Forti Paradox, ON must be a proper class: otherwise (by the definition of an ordinal), ON would itself be an ordinal, and $ON \in ON$. If P, which is a set, were cofinal, then its union would also be a set, and that union is the class of all ordinals. Contradiction.

Therefore, for some sufficiently large ζ, $c \subset V_\zeta$. But then $c \in V_{\zeta+1}$, contradicting the assumption that $c \in C$. So the class C must be empty, and $M = V$. □

Now we are ready to prove the section's main result:

Proof of Theorem 2.12. Using a process of recursion, we will pick a sufficiently large V_γ such that V_γ reflects the satisfaction of our finite set of formulas. At each stage, we will induct on the complexity of the formula, expanding our set by adding "witnesses" to existentially quantified sentences. Then, by recursion, we can construct the extension we want.

For simplicity, we can assume that our set $\Sigma = \varphi_1, \varphi_2, \ldots, \varphi_n$ of formulas is *closed under subformulas*, which means that any valid formula contained within another is also a member of the set. Specifically, if $\varphi = \neg\psi$ or $\varphi = \exists x(\psi)$ and $\varphi \in \Sigma$, then $\psi \in \Sigma$, and if $\varphi = \psi \wedge \sigma$, then $\psi, \sigma \in \Sigma$.

Note again that, from the definition of relativization, that the base case of this inductive process is always trivial. For any $x, y \in M$, $(x = y)^M \Leftrightarrow x = y$ and $(x \in y)^M \Leftrightarrow x \in y$, so these formulas are trivially absolute between M and the whole model. The same holds for the propositional connectives: if $\varphi^M \Leftrightarrow \varphi$ and $\psi^M \Leftrightarrow \psi$ both hold, it follows immediately that $\neg\varphi$ and $\varphi \wedge \psi$ are both absolute as well. Therefore, the only formulas we need to deal with are of the form $\exists x \varphi$, where φ is some other formula.

At each stage j in our recursive process, we start with some set $P_j = V_{\gamma_j}$ for an ordinal γ, and we expand it to a set $P_{j+1} = V_{\gamma_{j+1}}$. Consider any formula $\varphi_p \in \Sigma$, such that $\varphi_p = \exists x \varphi_q$ for some $\varphi_q \in \Sigma$ (we know $\varphi_q \in \Sigma$ because we assumed that Σ is closed under subformulas). Then φ_q has some number r of free variables aside from x, which we write $a_1, a_2, \ldots, a_r$ and denote collectively as $\vec{a}$.

We define a function on our list of formulas in the following way: if φ_p is not of the form $\exists x \varphi_q$, then $F(p) = 0$. On the other hand, suppose $\varphi_p(\vec{a})$ is of the form $\exists x \varphi_q(\vec{a}, x)$ for some φ_q. If there does not exist any b such that for some $\vec{a} \in P_j$, $\varphi_q(\vec{a}, b)$ is satisfied, then once again $F(p) = 0$. Otherwise, let $F(p) = \gamma$, where γ is the minimal rank of an element b satisfying $\varphi_q(\vec{a}, b)$. Let $G = \text{Max}\{F(i) \mid 1 < i < n\}$, let $G' = \text{Max}\{G, \gamma_j\}$, and set $P_{j+1} = V_{G'}$, effectively picking the smallest Von Neumann stage that contains each of the elements b we chose.

Finally, let $M = \bigcup_{i=0}^{\infty} P_i$. It follows immediately that $M = V_\gamma$ for some ordinal γ: a countable set of ordinals must be bounded above as we proved earlier. Then since the ordinals are well-ordered, the set must have a least upper bound γ, and $M = V_\gamma$.

By a final induction argument, we demonstrate that every formula in our list is absolute for M. As above, the base case and the inductive step for propositional connectives are trivial, so assume that $\varphi_p(\vec{a}) = \exists x \varphi_q(\vec{a}, x)$, where $\varphi_q(\vec{a}, x)$ is absolute for M. If $\varphi_p(\vec{a})$ is satisfied by some $a_1, a_2, \dots, a_r \in M$, then some b (not necessarily in M) exists satisfying $\varphi_q(\vec{a}, b)$. There must be some stage P_n containing all of $a_1, a_2, \dots, a_r$, and it follows that for some $c \in P_{n+1} \subset M$, $\varphi_q(\vec{a}, c)$ holds. But this means $\varphi_p(\vec{a})$ holds. On the other hand, suppose no $a_1, a_2, \dots, a_r \in M$ satisfies $\varphi_p(\vec{a})$. Then since φ_q is absolute, no $r+1$-tuple of elements in M can satisfy it. In short:

$$\exists x \varphi_q \Leftrightarrow (\exists x \in M)\varphi_q$$

This completes the proof. □

Note that by a slight modification of this process, we can guarantee that a given set is contained in the set generated through reflection: simply start with the given set at the first step of the recursion.

Corollary 2.17. *Let Δ be any finite fragment of the axioms of ZFC. If ZFC is consistent, then there exists a countable transitive set model of Δ.*

Proof. By Theorem 2.12, there exists some $V_\gamma = M$ such that for every $\varphi \in \Delta$, M satisfies the relevant φ^M. The set M is transitive, and this is expressible in the metatheory:

$$\text{“M is transitive”} \Leftrightarrow [\forall x \forall y (x \in M \wedge y \in x) \Longrightarrow y \in M]$$

Therefore, the first-order sentence expressing “M is transitive and satisfies all of Δ” has a model. By the Downward Lowenheim-Skolem Theorem, it must have a countable model N, and N is a countable transitive model of Δ. □

2.3. **A Word on Metamathematics.** The approach we will take in this paper, forcing over countable transitive models of ZFC, is only one possible way of formalizing the process. The advantage is that since our model is countable, we can guarantee the existence of our generic set.

Formally, the structure of our proof is as follows. We want to establish relative consistency, which is a claim of the form $\text{Con}(ZFC) \Longrightarrow \text{Con}(ZFC + \neg CH)$. Suppose that ZFC and (for instance) $\neg CH$ are inconsistent. Then by compactness, some finite fragment of ZFC, Δ, provably contradicts $\neg CH$. On the other hand, once we construct the forcing extension $M[G]$, we will be able to prove it satisfies any given axiom of ZFC, and that proof will itself only use finitely many axioms of ZFC. So consider a *larger* fragment, Γ, containing all the axioms we need to prove $M[G]$ satisfies $\Delta + \neg CH$. By the reflection principle, Γ has a countable transitive model. So Γ, by our assumption, must itself be inconsistent, making ZFC as a whole inconsistent. By contrapositive, this will prove our claim.

For the rest of the paper, “countable transitive model of ZFC” will usually mean any countable transitive model of a sufficiently large fragment of ZFC.

3. Basics of Forcing

Our goal, as we said, is to expand the model M (called the ground model) in such a way that it violates the Continuum Hypothesis, by adding $\aleph_2^M$ many additional subsets of the natural numbers.[5] From a certain perspective, we should *expect*

[5] We use superscripts to indicate that this is the second uncountable cardinal *within* M — the outer model, of course, thinks of it as countable.

this to be possible: after all, since our set is countable, the larger model believes that all its infinite sets have the same cardinality. But there are several immediate limitations on how we go about adding new subsets. The following two observations are due to Cohen, and appear in section IV.2 of [3].

Example 3.1. Let M be a countable transitive model of ZFC, and let $A = \sup\{\alpha \mid \alpha \in M\}$ be the least upper bound of the ordinals contained in M. Then A is a countable ordinal.

Suppose that for some ordinal $\beta > A$, V_β is a model of ZFC — if so, we say β is "worldly." There must be a least ordinal, γ ,with this property, and if we work within V_γ, we have a model where no such ordinal β exists. So it's consistent with ZFC that no ordinal greater than A is worldly.

A well-known result (which we won't prove here) is that given a model M of ZF, the sets "constructible" in a specific way from the ordinals of M themselves form a model of ZFC. In fact, this model, called L_M or the constructible universe of M, is the smallest transitive model containing all the ordinals in M. This demonstrates that $V = L$, the axiom that every set is constructible, is consistent with ZFC. Therefore, we can consistently assume that every set in our universe is constructible.[6]

We claim, given these assumptions, that there are *no* transitive set models of ZFC containing ordinals larger than A (and hence no transitive models containing A: it should be obvious that no transitive model can have a last ordinal). Otherwise, suppose that a transitive model N existed, such that $\sup\{\alpha \mid \alpha \in N\} = B > A$. We could then form the constructible universe $L_N = L^B$, containing all the same ordinals. But since we assumed $V = L$, $L^B = V^B$, and B is worldly. This violates our result above.

In effect, what this means is that we may have terrible luck while trying to use forcing. We might find ourselves in a universe where all sets are constructible, and in which no ordinals are worldly (all of which is consistent with ZFC). In this case we can't hope to add new ordinals to a transitive model, setting a limit on how we go about constructing new models. In fact, it need not be the case that all sets are constructible, but the favored way to prove this is through forcing.

The next result, also due directly to Cohen, demonstrates that we can't include sets that "encode" the fact that our model is countable.

Example 3.2. Again, let A be the supremum of the ordinals in our model M. Since A is countable, there exists a bijection, not contained in M, between A and ω. This amounts to a well-ordering of ω, and since ordering is a binary relation, to a subset of $\omega \times \omega$. But it's not hard to define a bijection between $\omega \times \omega$ and ω within ZFC, so A in turn corresponds to some subset of the natural numbers, $\alpha \subset \omega$.

We claim that no method of forcing can add the subset α to our model. Suppose it did: then by reversing our bijection, we could obtain the subset of $\omega \times \omega$. In turn, this well-ordering allows us to obtain the ordinal A. But A is an ordinal not contained in the original model. So we can't add α to our model if we want the model to stay transitive.

[6]Technically, we'd need to verify that the absence of worldly ordinals is consistent with constructibility, but this is routine: given a universe where all sets are constructible, we can repeat the same trick and pick the "first" worldly ordinal.

3.1. **Partial Orders and Generic Sets.**

Definition 3.3. Let M be a countable transitive model of ZFC. A forcing poset P, also called a notion of forcing, is a partially ordered set contained in M, containing a greatest element 1.

Definition 3.4. Let P be any poset. A subset $D \subset P$ is *dense* in P if, for any $p \in P$, there exists some $q \in D$ such that $q \leq p$. A set D is *dense below* the element a if for any $p \leq a$, some $q \in D$ exists where $q \leq p$.

Definition 3.5. Let P be any poset. A subset $F \subset P$ of the poset is a *filter* if it satisfies the following conditions:

(1) F is nonempty.
(2) F is upward-closed: if $a \in F$ and $b \geq a$, then $b \in F$.
(3) F is downward-directed: if $a, b \in F$, then there exists some $c \in F$ where $c \leq a$ and $c \leq b$.

Definition 3.6. Let P be a forcing poset. A subset $G \subset P$ is a *generic set* if:

(1) G is a filter on P
(2) For any dense subset D such that $D \in M$, $G \cap D$ is nonempty.

In [5] and other sources, the elements of our poset P are suggestively named "forcing conditions." In general, each element $p \in P$ will represent a small portion of the object we want to construct, such as a function between two cardinal numbers only defined for a finite subset of the domain. A generic set G can be thought of as a specific way of combining the elements: the requirement that G be a filter forces its elements to be"compatible" in a way we'll define shortly. And by requiring G to intersect certain dense sets from the ground model M, we can control the generic set's properties.

Definition 3.7. Let P be a poset. $p, q \in P$ are incompatible, written $p \perp q$, if no $a \in P$ exists where $a \leq p$, $a \leq q$. Otherwise, they're compatible, written $p \not\perp q$.

Lemma 3.8. *Let F be any filter over a poset P (in particular a generic set). Then every two elements in F are compatible.*

Proof. Follows directly from the downward-directedness condition. □

The following example, with slight modification, is due to [5]. We say that a forcing condition p is stronger than q if $p < q$.

Example 3.9. For any sets $A, B \in M$, let $Fn(A, B)$ represent the set of finite partial functions from A to B. That is, each element of $Fn(A, B)$ is a function into B, whose domain is a finite subset of A. We define the forcing poset P in the following way: its set of elements is $Fn(\omega, \{0, 1\})$. To define the ordering on $Fn(A, B)$, we stipulate that $a \leq b$ if $a \supset b$. In other words, a forcing condition p is stronger than q if it extends q to a larger domain.

Suppose G is a generic set over P. Since any two elements $a, b \in G$ are compatible, there exists a set c that extends both of them — so they must agree on their shared domain. Therefore, $\cup G$ is a function, since it contains at most one value for each element in ω. Let $f = \cup G$.

For any $n \in \omega$, define $D_n = \{p \in P \mid n \in \text{dom}(\omega)\}$, where dom identifies the domain of a function. Each set D_n is dense in P, because for any $p \in P$, we can extend it by defining the value of the function at n. Therefore, given G is generic, it intersects every D_n, and thus the domain of f is all of ω.

The function f has domain ω and range $\{0, 1\}$, so it defines a subset of the natural numbers. Identifying the real numbers with subsets of ω, we call f a *Cohen generic real*.

In fact, we can prove that the real number f defined via this process is not a member of M. The following condition, found in [7], demonstrates that almost any interesting generic set will not be contained in the ground model M.

Definition 3.10. Let P be a forcing poset. An atom is an element $p \in P$ such that for any $a, b \leq p$, $a \not\perp b$. We say that P is atomless if it has no atoms.

Lemma 3.11. *Let $P \in M$ be an atomless forcing poset, and let G be a generic set over P. Then $G \notin M$.*

Proof. Let $D = P \setminus G$. Evidently $D \in M$ if and only if $G \in M$, because either one can be constructed from the other as its complement. Consider any $p \in P$: since P is atomless there exist $a, b \leq p$ where $a \perp b$. We already proved that all elements in a generic set are compatible, so at most one of a and b can be contained in G. Thus, either $a \in D$ or $b \in D$, and either way there exists an element $x \in D$ where $x \leq p$. This proves that the set D is dense in P.

If D were contained in M, then it would have nonempty intersection with G, since G is generic and intersects all dense sets. But trivially, $(P \setminus G) \cap G = \emptyset$, so $D \notin M$. By our argument above, this implies $G \notin M$. □

On the other hand, we might wonder whether a generic set exists at all. Given an arbitrary model of ZFC and an arbitrary poset P, we *can't* guarantee that generic sets exist. But because our ground model M is countable, we can explicitly construct a generic set for any poset $P \in M$.

Lemma 3.12 (Rasiowa-Sikorski Lemma)**.** *Let P be a partial order and let D be a countable family of dense subsets of P. Then for any $p_0 \in P$ there exists a filter G containing p which intersects every member of D. In particular, for any countable transitive model M of ZFC and any forcing poset $P \in M$, there exists a generic set G over M.*

Proof. Since D is countable, we have some enumeration $D_1, D_2, \ldots$ of the dense subsets. Pick some element p_0. By the definition of density there exists some $p_1 \leq p_0$, not necessarily distinct, where $p_1 \in D_1$. In turn, there exists a $p_2 \leq p_1$ such that $p_2 \in D_2$. By recursion, we can define a sequence $p_0 \geq p_1 \geq p_2 \geq \cdots$ that intersects every dense subset in D. Since these elements form a descending sequence, all of them are compatible and we can define $G = \{q \in P \mid q \leq p_n, \text{ for some n}\}$, the filter generated by this sequence. Then G is a filter intersecting every dense set, and it contains the element p_0. □

3.2. **Construction of M[G].** We will construct the forcing extension $M[G]$ inductively, by defining a class of "names" for elements in the new model. It turns out that every element which we name in this way can be constructed using the generic set G, which proves that $M[G]$ is the minimal model containing G. But although the names are themselves elements of M, M doesn't have access to the generic set G and can't use the names to construct new sets. Here, our argument follows [7].

Definition 3.13. Let P be a forcing poset. The set τ is a P-name if every element of τ is an ordered pair of the form $\langle \sigma, p \rangle$, where $p \in P$ and σ is a P-name.

For any P-name τ, define $\text{dom}(\tau) = \{\sigma \mid \exists p \in P \ \langle \sigma, p \rangle \in \tau\}$.

Note that this definition is recursive: it's trivial that $\emptyset$ is a P-name, and other P-names are built up as sets of ordered pairs each containing a P-name.[7] We effectively use elements of P as "labels" for sets that may or may not be included depending on our choice of generic set. Using the generic set, we use another recursive process to construct $M[G]$:

Definition 3.14. Let τ be a P-name. We define:

$$\text{val}(\tau, G) = \tau_G = \{\text{val}(\sigma, G) \mid \exists p \in G[\langle \sigma, p \rangle \in \tau]\}$$

That is, the valuation of τ is the collection of valuations of names in τ which were paired with an element of G. Any ordered pair $\langle \sigma, p \rangle$ for $p \notin G$ is ignored.

Finally, we define $M[G] = \{\tau_G \mid \tau \in M^P\}$, the set of all valuations of P-names.

We can immediately present two essential examples of P-names. First, given a set $x \in M$, define $\check{x} = \{\langle \check{y}, 1 \rangle \mid y \in x\}$. This is *also* an inductive definition: each set $\check{x}$ is built up out of the P-names representing its elements. The process is exactly analogous to the construction of the class of all sets as the Von Neumann Universe (described above): but instead of simply taking sets of sets, we take sets of ordered pairs, all of which are indexed to the element 1.

Second, define $\Gamma = \{\langle \check{p}, p \rangle \mid p \in P\}$. With the suitable framework in place, Γ will be our "name" for the generic set G, but note that Γ does not name a *specific* generic set. If we chose a different generic set H, then Γ_H would evaluate to H, not to G. This illustrates a key principle: the ground model M should be able to talk about elements of $M[G]$ in the abstract, but because it doesn't have the set G, it doesn't know what the elements of $M[G]$ actually are.

Theorem 3.15. *Assume M is a countable transitive model of ZFC, $P \in M$ is a forcing poset, and $G \subset P$ is generic over M. Then:*

(1) $M[G] \supset M$
(2) $G \in M[G]$
(3) $|M[G]| = |M|$*, and in particular* $M[G]$ *is countable.*
(4) $M[G]$ *is minimal in the following sense: Suppose N is a transitive model of ZFC such that $N \supset M$ and $G \in N$. Then $N \supset M[G]$*

Proof. First, we verify by induction that $\check{x}_G = x$ for any $x \in M$. This is relatively trivial: if $\check{y}_G = y$ for all $y \in x$, then $\check{x}_G = \{\check{y}_G \mid y \in x\} = \{y \mid y \in x\} = x$. Then, since every element x has a canonical name $\check{x}$ and $\check{x}_G = x$, every element x is contained in $M[G]$.

Next, we prove $G \in M[G]$. Recall that $\Gamma = \{\langle \check{p}, p \rangle \mid p \in P\}$. If $p \in G$, then $\check{p}_G = p \in \Gamma_G$. On the other hand, if $p \notin G$, then the term $\langle \check{p}, p \rangle$ is not evaluated. So Γ_G is exactly the generic set G.

Next: since $M \subset M[G]$, we immediately have $|M| \leq |M[G]|$. On the other hand, there exists a surjective mapping from M^P, the set of P-names to the elements they represent in $M[G]$. So $|M[G]| \leq |M|$, and the two cardinalities must be equal.

Finally, we prove that G is minimal. Consider any model N as described. If τ is a P-name, then $\tau \in N$ since $M \subset N$. But also $G \in N$, so we can repeat the

[7] For the fact that definition by transfinite recursion is valid, see Theorem 2.14 of [5].

transfinite recursion defining τ_G within N. Because N is transitive, its ordinals are standard, and this process yields the same set as it does in $M[G]$. So $M[G] \subset N$. $\square$

Lemma 3.16. *Let τ be a P-name: then rank$(\tau_G) \leq$ rank(τ).*

Proof. By induction: Suppose that for every $\sigma_G \in \tau_G$, $\text{rank}(\sigma_G) \leq \text{rank}(\sigma)$. Note that for any σ, the minimal rank of a set containing σ is $\text{rank}(\sigma)+1$. In particular, $\text{rank}(\tau_G) = \sup\{\text{rank}(\sigma_G)+1 \mid \sigma_G \in \tau_G\} = \sup\{\text{rank}(\sigma_G)+1 \mid \sigma \in \tau\}$. But by our assumption, $\sup\{\text{rank}(\sigma_G)+1 \mid \sigma \in \tau\} \leq \sup\{\text{rank}(\sigma)+1 \mid \sigma \in \tau\} = \text{rank}(\tau)$. $\square$

Corollary 3.17. *$M[G]$ has the same ordinals as M.*

Proof. $M \subset M[G]$, so $M[G]$ contains all the ordinals of M. On the other hand, suppose some ordinal α is contained in $M[G]$. Trivially, α has rank α. But there exists some name $\tau \in M$ such that $\tau_G = \alpha$, and by our lemma, $\text{rank}(\tau) \geq \text{rank}(\alpha) = \alpha$. The concept of rank, because it's defined by transfinite recursion, is absolute between transitive models: so the rank of τ in M is at least α, implying that M must contain the ordinal α. $\square$

Lemma 3.18. *$M[G]$ is transitive.*

Proof. Consider any $x \in M[G]$, and suppose $y \in x$. There must exist some P-name τ such that $\tau_G = x$, and likewise there must exist some ordered pair $\langle \sigma, p\rangle \in \tau$ such that σ is a P-name, $\sigma_G = y$ and $p \in G$. But since $\sigma_G = y$, it follows immediately that $y \in M[G]$. $\square$

3.3. **Verifying Some Axioms.**

Theorem 3.19. *$M[G]$ is a model of extensionality, pairing, infinity and regularity.*

Proof. Extensionality: One direction is trivial, since $x = y$ always implies that x and y have the same elements. In the other direction, suppose $\forall z(z \in y \Leftrightarrow z \in x)^{M[G]}$ holds. This means that for any $z \in M[G]$, $z \in y \Leftrightarrow z \in x$. But since $M[G]$ is transitive, every element of either x or y is an element of $M[G]$. So $\forall z(z \in y \Leftrightarrow z \in x)^{M[G]}$ holds if and only if $\forall z(z \in y \Leftrightarrow z \in x)$ does, and $M[G]$ is extensional.

Pairing: Consider any two elements $a, b \in M[G]$. As usual, we have $a = \alpha_G$, $b = \beta_G$, for names α and β. Let $C = \{\langle \alpha, 1\rangle, \langle \beta, 1\rangle\}$: then C is a P-name, and it's trivial that $C_G = \{a, b\}$.

Infinity: Since $M[G] \supset M$, $\omega \in M[G]$ and so $M[G]$ immediately satisfies the axiom.

Regularity: By induction on P-names. Let τ be a P-name, and suppose that every σ_G such that $\langle \sigma, p\rangle \in \tau$ is well-founded. This means that of the elements contained in τ_G, at least one must be of minimal rank: pick some such element and call it β_G. It follows that for any γ_G in τ_G, $\text{rank}(\gamma_G) \geq \text{rank}(\beta_G)$, and hence that $\gamma_G \notin \beta_G$, $\tau_G \cap \beta_G = \emptyset$. $\square$

Theorem 3.20. *$M[G]$ is a model of the axiom of union.*

Proof. The following construction is suggested on page 249 of [7]. Consider some $a \in M[G]$: then we have $a = \tau_G$ for some P-name τ. Define:

$$\pi = \{\langle v, p\rangle \mid \exists \langle \sigma, q\rangle \in \tau\, \exists r[\langle v, r\rangle \in \sigma \wedge p \leq r \wedge p \leq q]\}$$

Then we claim $\pi_G = \cup a$. In the first, direction, suppose $x \in \cup a$. Then $x \in b$ for some $b \in a$, and we have $b = \beta_G$ for some P-name β. Similarly, we have $x = \xi_G$

for some P-name ξ, and also $\langle \xi, r \rangle \in \beta$ for some $r \in G$. Finally, for some $q \in G$, $\langle \beta, q \rangle \in \tau$.

Apply the downward-directedness of G: since $q, r \in G$, there exists some $p \in G$ where $p \leq r$ and $p \leq q$. Therefore, by the definition of π, $\langle \xi, p \rangle \in \pi$. Since $p \in G$, we have $\xi_G = x \in \pi_G$.

In the other direction, suppose $x \in \pi_G$. Then $x = \xi_G$, and $\langle \xi, p \rangle \in \pi$ for some $p \in G$. By the definition of π, it immediately follows that for some P-name σ and some $q \geq p$, $\langle \sigma, q \rangle \in \tau$, and likewise that for some $r \geq p$, $\langle \xi, r \rangle \in \sigma$. But since G is a filter, $q, r \in G$: this yields that $a = \tau_G \ni \sigma_G \ni \xi_G = x$, and hence that $x \in \cup a$. □

4. Completing the Proof

The remaining work to do is as follows: First, we need to demonstrate that $M[G] \vDash ZFC$, by verifying the axioms of powerset, replacement, comprehension and choice. Then, we can construct forcing posets to verify both directions of the continuum hypothesis.

First, however, we need to develop a way of defining and proving results about forcing within ZFC itself. We've already seen that all elements of the forcing extension have names in the ground model M. However, we can strengthen this result: for a given $p \in P$, M "knows" which sentences must be true if $p \in G$. This allows us to prove the remaining axioms, and provides a language to discuss which properties we can or can't force to be true.

4.1. The Forcing Relation.

For the rest of the paper, we enrich our logical vocabulary by adding P-names as constants:

Definition 4.1. Let P be a forcing poset. The forcing language of P, written $\mathcal{FL}_P$, is the language of set theory with every P-name added as a constant.

The following definition achieves the purpose we outlined above, by introducing what is known as the *forcing relation*.[8]

Definition 4.2. Let $p \in P$, and let α, β, τ be P-names. Then the relation $p \Vdash \varphi$ for some sentence $\varphi \in \mathcal{FL}_P$, spoken as "$p$ forces φ," is defined by the following recursion:

1. $p \Vdash \alpha = \beta$ if for any $\sigma \in \mathrm{dom}(\alpha) \cup \mathrm{dom}(\beta)$ and any $q \leq p$, we have $q \Vdash \sigma \in \alpha \Leftrightarrow q \Vdash \sigma \in \beta$.
2. Define $X = \{q \leq p : \exists \langle \sigma, r \rangle \in \tau \, [q \leq r \wedge q \Vdash \alpha = \sigma]\}$. Then $p \Vdash \alpha \in \tau$ if the set X is dense below p.
3. $p \Vdash \varphi \wedge \psi$ if $p \Vdash \varphi$ and $p \Vdash \psi$.
4. $p \Vdash \neg\varphi$ if there does not exist a $q \leq p$ where $q \Vdash \varphi$.
5. $p \Vdash \exists x \varphi(x)$ if the set of elements q where $q \Vdash \varphi(\tau)$ for some τ is dense below p.

In addition to recursion on sentence complexity, this definition employs the recursive contruction of $M[G]$ itself. Notice that the criteria for atomic sentences are defined in terms of each other. Both become trivial for $\emptyset$, such that any p

[8] The entirety of this section substantially follows Kunen in [7]. However, here, our terminology differs slightly from Kunen. The relation which we label $\Vdash$, in Kunen, is called $\Vdash^*$. Kunen defines $p \Vdash \varphi \Leftrightarrow M[G] \vDash \varphi$ for any $G \ni p$, and then proves the two notions are equivalent. Here, we introduce the second notion first, but achieve the same result: as Kunen does, we demonstrate the two notions of $\Vdash$ are entirely equivalent

forces $\emptyset = \emptyset$ and no p forces that $\emptyset$ contains a member. Then at each stage in the recursion, atomic statements about more complex P-names are defined in terms of statements about their members. Finally, note that this is a relation with *names*, not elements of $M[G]$.

We can immediately verify several basic properties of the forcing relation:

Lemma 4.3. *(1) If $p \Vdash \varphi$ and $q \leq p$, then $q \Vdash \varphi$*
(2) $p \Vdash \varphi$ if and only if $\{q \leq p \mid q \Vdash \varphi\}$ is dense below p.
(3) $p \Vdash \varphi$ if and only no $q \leq p$ exists where $q \Vdash \neg\varphi$.

Proof. (1): By induction on φ: The atomic cases and propositional connectives follow immediately from the definition and the fact that $q \leq p$. The case for $\exists x \varphi(x)$ is similarly trivial, since any set dense below p will also be dense below q.

(2): The first direction is obvious, since if $p \Vdash \varphi$, any $q \leq p$ will also force φ and hence the set of such q is every element below p. In the other direction, suppose $\{q \leq p \mid q \Vdash \varphi\}$ is dense below p. We employ another induction:

- First, consider the sentence $\alpha \in \tau$. The set q for which $q \Vdash \alpha \in \tau$ is dense below p. But for any q, $q \Vdash \alpha \in \tau$ simply means that X is dense below q for the set X defined above. Then for an arbitrary $r \leq p$, there exists a $q \leq r$ such that $q \Vdash \alpha \in \tau$, and a $s \leq q$ such that $s \in X$. It follows that X is dense below p.
- Next, consider the sentence $\alpha = \beta$. Given any σ, suppose $\sigma \in \alpha$. Then the set $\{q \mid q \Vdash \sigma \in \alpha\}$ is dense below p. But also the set $\{r \leq p \mid r \Vdash \sigma \in \alpha \Leftrightarrow r \Vdash \sigma \in \beta\}$ is dense below p. Consider any $s \leq p$: then there exists a $q \leq s$ such that $q \Vdash \sigma \in \alpha$, and similarly there exists an $r \leq q$ such that $r \Vdash \sigma \in \alpha \Leftrightarrow r \Vdash \sigma \in \beta$. But since $r \leq q$, $r \Vdash \sigma \in \alpha$, and it follows that $r \Vdash \sigma \in \beta$, so the relevant set is dense below p.
- The propositional connectives follow immediately from definitions. In particular, for conjunction, note that density below allows us to generate an element below any q that forces both φ and ψ.
- Finally, consider the sentence $\exists x \varphi(x)$. Suppose the set of elements q where $q \Vdash \exists x \varphi(x)$ is dense below p. Then consider any $s \leq p$. It follows that some $q \leq s$ exists such that $q \Vdash \exists x \varphi(x)$. But by the definition of the forcing relation, we can then pick some $r \leq q$ and some τ such that $r \Vdash \varphi(\tau)$. This means the relevant set is dense below p, and that $p \Vdash \exists x \varphi(x)$.

(3): This follows from (2). Suppose $p \Vdash \varphi$: then any element $q \leq p$ has $q \Vdash \varphi$, and (by our definition), it is not the case that $q \Vdash \neg\varphi$. On the other hand, suppose p does not force φ: then $\{q \leq p \mid q \Vdash \varphi\}$ is not dense below p, so for some $r \leq p$, no $s \leq r$ has $s \Vdash \varphi$. This gives us that $r \Vdash \neg\varphi$, violating the condition. □

Lemma 4.4. *Assume G is generic, $p \in G$, and suppose that $D \subset P$ is dense below p. Then $G \cap D$ is nonempty.*

Proof. Define $D' = D \cup \{q \in P \mid q \perp p\}$. We claim D' is dense. Pick any $q \in P$: if q is compatible with p, then there exists some r with $r \leq p$, $r \leq q$, and since D is dense below, there exists some $r' \leq r \leq q$ where $r \in D \subset D'$. Otherwise, q is incompatible with p and by definition q is already a member of D'. Therefore, there exists some $r \in D' \cap G$. But since G is a filter, r and p must be compatible. So $r \in D$. □

The key result of this section is the following lemma, which asserts that the sentences true in a generic extension are exactly those forced by some element of the extension. Our proof is substantially due to Kunen, though presented slightly differently.

Lemma 4.5 (The Truth Lemma)**.** *Let $\varphi \in \mathcal{FL}_P$ be a sentence. Then:*
(a) If $p \in G$ and $p \Vdash \varphi$, then $M[G] \vDash \varphi$.[9]
(b) If $M[G] \vDash \varphi$ then for some p, we have $p \Vdash \varphi$.

Proof. We prove this in two steps. First we verify for atomic sentences, proving (a) and (b) separately. Then, we use induction on sentence complexity to extend both (a) and (b) simultaneously to any sentence.

First, suppose $p \Vdash \pi \in \tau$ for some P-names π and τ. Recall: this means that the set $D = \{q \leq p \mid \exists \langle \sigma, r \rangle \in \tau \, [q \leq r \wedge q \Vdash \pi = \sigma]\}$ is dense below p. But then by the lemma we proved, $G \cap D$ is nonempty. Pick some $q \in G \cap D$: then using the definition of D, pick $\langle \sigma, r \rangle \in \tau$ such that $q \leq r$ and $q \Vdash \pi = \sigma$. By the inductive hypothesis, we have that $M[G] \vDash \pi = \sigma$, meaning $\pi_G = \sigma_G$. But also, since $q \in G$, we have $r \in G$, $\sigma_G \in \tau_G$. Since $\sigma_G = \pi_G$, this proves the result.

Next, suppose that $p \Vdash \pi = \tau$. This, by definition, means that for any P-name $\sigma \in \operatorname{dom}(\pi) \cup \operatorname{dom}(\tau)$, we have $p \Vdash \sigma \in \pi \Leftrightarrow p \Vdash \sigma \in \tau$. We will show that $\pi_G \subset \tau_G$: the result in the other direction is identical. To that end, pick some $x \in \pi_G$. We have $x = \sigma_G$ for some P-name σ, and that for some $r \in G$, $\langle \sigma, r \rangle \in \pi$. Since p, r are elements of G and G is a filter, we can pick some $q \in G$ with $q \leq p, q \leq r$.

We claim that $q \Vdash \sigma \in \pi$, and this will follow easily from the definition. For any $q_1 \leq q$, we have $q_1 \leq r$, and $\langle \sigma, r \rangle \in \pi$. The only remaining condition is that $q \Vdash \sigma = \sigma$, but this is once again an immediate result of the definition. However, since $q \leq p$, we also have that $q \Vdash \pi = \tau$, so $q \Vdash \sigma \in \pi \Leftrightarrow q \Vdash \sigma \in \tau$. Therefore, $q \Vdash \sigma \in \tau$. By the inductive hypothesis, this means $M[G] \vDash \sigma \in \tau$, and that $x = \sigma_G \in \tau$.

Next, we verify claim (b) for atomic sentences. First, suppose that $M[G] \vDash \pi \in \tau$. Then for some $r \in G$, and some P-name σ, we have $\langle \sigma, r \rangle \in \tau$ and $\sigma_G = \pi_G$. Applying the inductive hypothesis, there exists some $p \in G$ such that $p \Vdash \sigma = \pi$. Applying the fact that G is a filter, pick some $r \in G$ such that $r \leq q$, $r \leq p$. Then since $r \leq p$, $r \Vdash \sigma = \pi$. So it follows directly from the definition of the forcing relation that $r \Vdash \pi \in \tau$.

Finally, for atomic sentences, suppose $M[G] \vDash \pi = \tau$. Define $D \subset P$ to be the set of elements q for which one of the following holds:

(1) $q \Vdash \pi = \tau$
(2) For some P-name $\sigma \in \operatorname{dom}(\pi) \cup \operatorname{dom}(\tau)$, $q \Vdash \sigma \in \pi$ and $q \Vdash \sigma \notin \tau$
(3) For some P-name $\sigma \in \operatorname{dom}(\pi) \cup \operatorname{dom}(\tau)$, $q \Vdash \sigma \in \tau$ and $q \Vdash \sigma \notin \pi$

We claim D is dense: pick some $r \in P$. If $r \Vdash \pi = \tau$, then $r \in D$. If not, then by the definition of the forcing relation, there exists some σ and some $s \leq r$ where either condition (2) or condition (3) holds for s. Therefore, $G \cap D$ is nonempty. Pick some $p \in G \cap D$: by definition one of (1), (2), or (3) holds. But if either (2) or (3) were true for p, then by part (a) we would have $\pi_g \neq \tau_G$. So we must have $p \Vdash \pi = \tau$.

Now, we extend our result by applying induction on sentence complexity. The case for conjunctions is trivial: if $p \Vdash \varphi$ and $p \Vdash \psi$, then we have $M[G] \vDash \varphi$, and

[9]Here, we take a P-name τ to refer to the corresponding element $\tau_G \in M[G]$.

$M[G] \vDash \psi$, so it follows immediately that $M[G] \vdash \varphi \wedge \psi$. The other direction is almost as easy: if $M[G] \vDash \varphi \wedge \psi$, then some two elements must force φ and ψ respectively, so consider some element which is less than both.

Negation: By contrapositive, suppose we have $\neg[M[G] \vDash \neg\varphi]$, or equivalently $M[G] \vDash \varphi$. Then, for any $p \in G$, there would exist some $r \in G$ where $r \Vdash \varphi$ by the inductive hypothesis. So we could pick some $q \in G$ where $q \leq p, q \leq r$. This contradicts the condition for $p \Vdash \neg\varphi$ to be true.

In the other direction, suppose $M[G] \vDash \neg\varphi$. Then it is not the case that $M[G] \vDash \varphi$, so by part (a), there is no $p \in G$ where $p \Vdash \varphi$. Define D to be the set of $p \in P$ such that either $p \Vdash \neg\varphi$, or $p \Vdash \varphi$. We can easily see that D is dense: the condition for $p \Vdash \neg\varphi$ is that no lesser element forces φ, so if the condition fails then some such lesser element must exist. Therefore we can pick some $p \in G \cap D$. But by our assumption, p doesn't force φ: the only option left is that for any $q \leq p$, $\neg(q \Vdash \varphi)$, which means that $q \Vdash \neg\varphi$.

Existential quantifiers: suppose that $p \Vdash \exists x\varphi(x)$. Then the set of elements q such that $q \Vdash \varphi(\tau)$ for some P-name τ is dense below p. By lemma 4.4, its intersection with G is nonempty. But if we pick some element in the intersection, we have a $p \in G$ where $p \Vdash \varphi(\tau)$. It follows that $M[G] \vDash \varphi(\tau)$, which directly implies that $M[G] \vDash \exists x\varphi(x)$.

Finally, assume that $M[G] \vDash \exists x\varphi(x)$. It follows that for some τ, $M[G] \vDash \varphi(\tau)$. Apply the inductive hypothesis, and our proof is complete. □

Corollary 4.6. *Given any sentence φ and any $p \in G$ for a generic set G, there exists an element $q \leq p$ that decides φ: either $q \Vdash \varphi$ or $q \Vdash \neg\varphi$.*

Proof. Without loss of generality, suppose $M[G] \vDash \varphi$. Then some element $r \in G$ has $r \Vdash \varphi$, by the truth lemma. Let q be any element in G where $q \leq p, q \leq r$. □

4.2. **M[G] is a Model of ZFC.** We have four axioms left to verify: comprehension, replacement, powerset and choice. In fact, however, comprehension follows from the others:

Lemma 4.7. *Suppose that some model M satisfies pairing, union and replacement. Then M satisfies comprehension.*

Proof. Consider any set X and any formula $\varphi(x)$. For each $x \in X$, pairing implies that $\{x, x\} = \{x\}$ exists. Therefore, we can define a mapping F where $F(x) = \{x\}$ if $\varphi(x)$ holds, and otherwise $F(x) = \emptyset$. We observe that $\cup F(x) = \{x \in X \mid \varphi(x)\}$. □

There are multiple related but distinct ways of stating the axiom schema of replacement: the most general is that for a specific formula, the image of *any* set on which that formula acts as a function is itself a set. However, for our purposes, we can work with a specific set as the function's domain: if the universal statement fails, then $M[G]$ would prove some specific counterexample, and we will prove that it doesn't.

Theorem 4.8. *$M[G]$ models the axiom schema of replacement.*

Proof. Consider a set $A \in M[G]$, and a mapping F on A (that is, a mapping whose domain is all of A and which is not one-to-many). Since we assume F is definable, there exists a formula φ where $F(x) = y \Leftrightarrow \varphi(x, y)$.[10] Suppose $A = \tau_G$. Recall the

[10]The formula might include parameters other than x and y, but the proof is unchanged by their inclusion.

reflection theorem, which, given any formula, allows us to generate a set for which that formula is absolute. Consider the following formula:

$$\psi(\sigma, p) := \sigma \in \text{dom}(\tau) \land \exists\pi\,[p \Vdash \varphi(\sigma, \pi)]$$

This formula need not be satisfied for a given σ: if $\sigma_G \notin \tau_G$, then we have no guarantee that the function is defined on σ_G. However, we can use reflection to generate a set Q for which we have $\psi(\sigma, p) \Leftrightarrow \psi(\sigma, p)^Q$. In particular, we assume Q contains every name in the domain of τ and also the entirety of P. Now, we apply the comprehension axiom in M to construct the set we want:

$$\pi := \{\langle\theta, r\rangle \in Q \times P \mid \exists\langle\sigma, p\rangle \in \tau[r \Vdash \varphi(\sigma, \theta) \land r \leq p]\}$$

We claim that for any $x \in A$, there exists some $y \in \pi_G$ such that $F(x) = y$, or equivalently $\varphi(x, y)$. Suppose $x = \xi_G$, and that $\langle\xi, p\rangle \in \tau$ for some $p \in G$. Then since F is a mapping, by the truth lemma there exists some $q \in G$ and some name γ where $q \Vdash \varphi(\xi, \gamma)$. Therefore, $\psi(\xi, q)$ is satisfied, and in particular there exists some $\gamma \in Q$, by our reflection argument, that satisfies $q \Vdash \varphi(\xi, \gamma)$. Applying the definition of a filter, we can pick some $r \in G$ where $r \leq p$, $r \leq q$. Then by the definition of π, we have $\langle\gamma, r\rangle \in \pi$. Finally, $r \in G$, so $\gamma_G \in \pi_G$, and $q \Vdash \varphi(\xi, \gamma)$, so we're done.

In the other direction, consider an arbitrary $y \in \pi_G$. Then we have $y = \theta_G$, where $\langle\theta, r\rangle \in \pi$ and $r \in G$. By the definition of π, there exists some $p \geq r$ and some name σ where $\langle\sigma, p\rangle \in \tau$ and $r \Vdash \varphi(\sigma, \pi)$. But since $p \geq r$, $p \in G$ and $\sigma_G \in \tau_G$. So applying the truth lemma, we have $M[G] \vDash \varphi(\sigma, \theta)$, and hence $F(\sigma_G) = \theta_G$. Therefore, π_G is exactly the image $F(A)$. □

Corollary 4.9. *$M[G]$ satisfies the axiom schema of comprehension.* □

Note that so far, we haven't used the axiom of choice. If our ground model M merely satisfies ZF, then the extension will satisfy ZF as well. If it additionally satisfies choice, then so will the extension.

Definition 4.10. Recall the standard definition of an ordered pair: we interpret $\langle a, b\rangle = \{\{a\}, \{a, b\}\}$. For any P-names σ and τ, define:

$$\text{op}(\sigma, \tau) = \{\langle\{\langle, \sigma, 1\rangle\}, 1\rangle, \langle\{\langle\sigma, 1\rangle, \langle\tau, 1\rangle\}, 1\rangle\}$$

Then it's not hard to see that $\text{op}(\sigma, \tau)_G = \langle\sigma_G, \tau_G\rangle$.

Theorem 4.11. *$M[G] \vDash ZF$. Additionally, if M satisfies the axiom of choice, then $M[G] \vDash ZFC$.*

Proof. The only remaining axiom for ZF is powerset. Now that we've proven comprehension, it's enough to generate some superset of the relevant powerset: then we can use a suitable formula to restrict to the powerset itself. Let τ be some P-name, and let $Q = \mathcal{P}(\text{dom}(\tau) \times P)$, and let $\pi = Q \times \{1\}$. Now consider, any $c \in M[G]$ such that $c \subset \tau_G$. As usual, we have $c = \alpha_G$ for some P-name α.

Define $\xi = \{\langle\sigma, p\rangle : \sigma \in \text{dom}(\tau) \land p \Vdash \sigma \in \alpha\}$. This set exists in M by an application of replacement. We claim $\xi_G = c$. First, suppose that $\beta_G \in \xi_G$ for some name β, or equivalently $\langle\beta, p\rangle \in \xi$ for some $p \in G$. By the definition of ξ, we have that $p \Vdash \beta \in \alpha$, and $p \in G$, so this yields $\beta_G \in \alpha_G = c$.

In the other direction, pick some $y \in c$, and we have $y = \gamma_G$ for some name γ. By the truth lemma, some $p \in G$ must exist where $p \Vdash \gamma \in \alpha$. Since $c \subset \tau_G$, we can safely assume $\gamma \in \mathrm{dom}(\tau)$ — otherwise, pick whichever P-name for this set is included in $\mathrm{dom}(\tau)$. Therefore, $\langle \gamma, p\rangle \in \xi$, and $\gamma_G = y \in \xi_G$.

Finally, we prove that $M[G]$ satisfies the axiom of choice. Recall that the axiom of choice is equivalent to the statement that every set can be well-ordered, as proven for example in chapter 5 of [5]. Applying choice in M, we can enumerate $\mathrm{dom}(\tau) = \{\sigma^\xi : \xi < \alpha\}$ for some ordinal α. Let $R = \{\langle \mathrm{op}(\check{\xi}, \sigma^\xi), 1\rangle \mid \xi < \alpha\}$. Then R_G is a function from α onto $\mathrm{dom}(\tau)$. We can evidently use this to well-order τ_G, by associating any element with the least ordinal mapping to a name for that element. □

4.3. **Independence at Last.** At this point we know that $M[G]$ models ZFC for any generic set G, so we can begin exploring which models we get from different notions of forcing. The forcing poset to break the continuum hypothesis will ultimately be relatively straightforward. However, we need a guarantee that certain properties of our model will be preserved. Specifically, if we add $\aleph_2$ additional subsets of ω, we want $\aleph_2$ to designate the same cardinal in both M and $M[G]$. Remember in the following discussion that M and $M[G]$ are both transitive and have the same ordinals. These are basic results about the properties of forcing, and our discussion follows [7].

Definition 4.12. Let γ be a limit ordinal. The cofinality of γ, written $\mathrm{cf}(\gamma)$, is the minimal cardinality of a sequence (equivalently, a set) of ordinals less than γ whose least upper bound is γ.

We say γ is *regular* if $\mathrm{cf}(\gamma) = \gamma$. In particular, since γ must be a cardinal for this to hold, we say γ is a regular cardinal. If γ is not regular, we say γ is singular.

It will sometimes be convenient (as in our proof that $\mathrm{cf}(\mathrm{cf}(\gamma)) = \mathrm{cf}(\gamma)$ below) to assume that the relevant sequence is increasing. This can be true for any sequence, since a set of ordinals can be well-ordered.

Example 4.13. Consider any successor cardinal $\aleph_{\gamma+1}$. Then assuming the axiom of choice, $\aleph_{\gamma+1}$ is regular. We suppress the details, but note that the supremum of a sequence of ordinals is also its union. Then a sequence of length less than $\aleph_{\gamma+1}$ can have cardinality at most $\aleph_\gamma$, and each of the sets in the sequence will also be no larger than $\aleph_\gamma$. An elementary fact in cardinal arithmetic shows that $\aleph_\gamma \cdot \aleph_\gamma = \aleph_\gamma < \aleph_{\gamma+1}$.

Lemma 4.14. *Let γ be any limit ordinal. Then $cf(cf(\gamma)) = cf(\gamma)$, and hence $cf(\gamma)$ is regular.*

Proof. Trivially, $\mathrm{cf}(\mathrm{cf}(\gamma)) \leq \mathrm{cf}(\gamma)$, since one sequence converging to $\mathrm{cf}(\gamma)$ is the set of ordinals less than $\mathrm{cf}(\gamma)$. On the other hand: let $\{a_\xi \mid \xi < \mathrm{cf}(\gamma)\}$ be some sequence whose supremum is γ, and suppose $\mathrm{cf}(\mathrm{cf}(\gamma)) < cf(\gamma)$. Then for some $\beta < \mathrm{cf}(\gamma)$, there exists a sequence $\{b_\xi \mid \xi < \beta\}$ whose supremum is $\mathrm{cf}(\gamma)$. We can assume both sequences are increasing. But this means the sequence $\{a_{b_\xi} \mid \xi < \beta\}$ converges to γ, and $\beta < \mathrm{cf}(\gamma)$. Contradiction. □

Definition 4.15. Let P be a forcing poset.

(1) P preserves cardinals if, for any generic $G \subset P$ and any ordinal β, β is a cardinal in M if and only if β is a cardinal in $M[G]$.
(2) P preserves cofinalities if and only if for any generic $G \subset P$ and any limit ordinal γ, $\mathrm{cf}^M(\gamma) = \mathrm{cf}^{M[G]}(\gamma)$.

Lemma 4.16. *Let P be a forcing poset. Then P preserves cofinalities if, and only if, M and $M[G]$ have the same regular cardinals for any generic G. Furthermore, if P preserves cofinalities, then P preserves cardinals.*

Proof. Trivially, if P preserves cofinalities, then forcing over P will preserve regular cardinals. In the other direction, assume regular cardinals are preserved, pick some limit ordinal γ, and assume $\beta = \mathrm{cf}^M(\gamma)$. As we proved above, β is regular, and so by our assumption, β will also be regular in $M[G]$. Also, β is the order type of some sequence in $M[G]$ converging to γ, though we don't yet know whether it's the shortest one. Suppose we had a shorter sequence: Then we could define a new sequence by, for each term, picking the next largest element of the β-sequence. This would be a sequence of length less than β converging to β, contradicting the fact that β is regular.

Next, suppose P preserves all cofinalities, and consider some cardinal κ. If κ is finite or equal to ω, then it must be preserved between transitive models. Otherwise, κ is either a successor cardinal or a limit cardinal. If κ is a successor cardinal then it's regular, and will be preserved. Otherwise, κ is a limit cardinal: Then it's the supremum of a sequence of regular cardinals, and it must be a cardinal number since it's the first whose size is strictly greater than any of them.[11] □

The key condition, it turns out, is the following:

Definition 4.17. Let P be a forcing poset. P has the countable chain condition (ccc), if any collection of pairwise-incompatible elements (an antichain) is at most countable. In other words, if $X \subset P$ is such that for $p, q \in X$, $p \perp q$, then $|X| \leq \aleph_0$.

Lemma 4.18. *Suppose P has the countable chain condition in M, and that $A, B \in M$. If $f \in M[G]$ is a function from A to B, then there exists some $F : A \to \mathcal{P}(B)$ satisfying the following:*

(1) For any $a \in A$, $f(a) \in F(a)$.
(2) For any $a \in A$, $|F(a)| \leq \aleph_0$.

Note that M must *believe* that P is ccc: since M is countable, from outside M we will think that any of its posets have the condition, but that doesn't mean that every poset in M internally has the condition.

Proof. Suppose φ is a P-name for the function f. By the truth lemma, there exists a $p \in G$ such that $p \Vdash \varphi : \check{A} \to \check{B}$ — in other words, p forces the statement that f is a function from A to B. Note that since $A, B \in M$, we use their canonical names.

For any $a \in A$, define $f(A) = \{b \in B \mid \exists q \leq p\, [p \Vdash \varphi(\check{a}) = \check{b}]\}$. Evidently, $f(a) = c \in f(A)$: by the truth lemma some $q \in G$ exists where $q \Vdash f(a) = c$, and using the filter condition we can pick a q' where $q' \leq q, q' \leq p$. On the other hand, we claim every $f(A)$ is countable. Let $x, y \in f(A)$: then there exists some $r_x \leq p$ such that $r_x \Vdash f(a) = x$, and some $r_y \leq p$ such that $r_y \Vdash f(a) = y$. We already proved that no element forces a contradiction, so we must have $r_x \perp r_y$. But this means that the elements r_x for $x \in f(A)$ form an antichain: we conclude that there are only countably many of them, and hence that $f(A)$ is countable in M. □

[11] For example, $\aleph_\omega$ is the limit of a sequence of successor cardinals, all of which are regular. Any ordinal less than $\aleph_\omega$ must be strictly less than some $\aleph_n$, which is bounded by $\aleph_{n+1}$. So any ordinal less than $\aleph_\omega$ has strictly smaller cardinality, and $\aleph_\omega$ is a cardinal. The same argument applies, in turn, for the limit of a sequence of limit cardinals.

Theorem 4.19. *Suppose P has the countable chain condition in M. Then P preserves cofinalities, and hence also cardinals.*

Proof. By 4.16, it's enough to show that for any β, $(\beta \text{ is regular})^M \Leftrightarrow (\beta \text{ is regular})^{M[G]}$. In particular, any regular cardinal in $M[G]$ must be a regular cardinal in M, since a sequence in M which shows that β is singular would also be contained in $M[G]$.[12] So all we have to prove is that if β is regular in M, it's also regular in $M[G]$.

Suppose β is regular in M, but not in $M[G]$. Then there exists some sequence of ordinals less than β whose supremum β, with cardinality α strictly less than β. By the structure of the ordinals, we can interpret this sequence as being some subset $X \subset \beta$, where $|X| = \alpha$. By the definition of cardinality, this yields a function $f : \alpha \to X$, contained in $M[G]$, and we can interpret f as a function from α to β.

Now, apply the result we just proved: there must exist a function $F : \alpha \to \mathcal{P}(B)$, where $F \in M$, $f(a) \in F(a)$ and $|F(a)| \leq \aleph_0$. Let $U = \cup_{a\in\alpha} F(a)$: then $|U| \leq \aleph_0 \cdot \alpha$, which is less than β simply because $\alpha < \beta$ and β is infinite. On the other hand, $\sup(U) = \sup(f(a) : a \in \alpha) = \beta$ (note that $\sup(U) \leq \beta$ by the definition of its codomain). So U is a set of ordinals less than β whose cardinality is less than β, and its supremum is β. This contradicts the assumption that $(\beta \text{ is regular})^M$. □

The following construction, which helps prove a poset has the countable chain condition, can be found in [7].

Definition 4.20. Let A be a family of sets. A is a delta system with root R if for any $X, Y \in A$, $X \cap Y = R$.

Lemma 4.21 (Delta System Lemma)**.** *Let A be a family of finite sets such that $|A| = \kappa$ and κ is regular and uncountable. Then there exists some $B \subset A$ such that B forms a delta system and $|B| = \kappa$.*

Proof. Given that any $X \in A$ is finite, it can only have one of $\aleph_0$ possible cardinalities. If the number of sets having cardinality n is less than κ for every n, then since κ is regular their sum could not be equal to A. So pick some $D \subset A$ such that $|D| = \kappa$ and $|X| = n$ for all $X \in D$.

We prove by induction that D must contain a delta system of the right size. If $n = 1$, then D is a set of κ singleton sets, each containing a different element. They form a delta system whose root is the empty set.

For $n > 1$, and for any p, define $D_p = \{X \in D : \ p \in X\}$. Suppose that for some p, D_p has cardinality κ. Then consider the collection $\{X \setminus \{p\} : X \in D_p\}$: all the sets in this collection have size $n - 1$, so by induction they have a delta system. It follows that D_p as a whole has a delta system, whose root is the root of the smaller system with p added.

Otherwise, assume that $|D_p| < \kappa$ for every p. Observe that for any S, $\{X \in D \mid X \cap S \neq \emptyset\} = \cup_{p\in S} D_s$. Pick any $X_0 \in D$: we will use recursion to generate a delta system whose root is the empty set. Since X_0 is finite, $\{X \in D \mid X \cap X_0 \neq \emptyset\} = \cup_{p\in X_0} D_p$ is the union of less than κ sets of size less than κ, and since κ is regular must have cardinality strictly less than κ. So we can pick some X_1 where $X_0 \cap X_1 = \emptyset$. The same holds for any X_β where $\beta < \kappa$, so we generate a sequence of length κ all of whose members are pairwise disjoint. □

[12] To be slightly more formal: the relevant property is that there *does not* exist a function on some initial segment of β the supremum of whose range is β. Any such function, if it existed in M, would still possess the same relevant properties when interpreted as an object in $M[G]$.

Theorem 4.22. *If ZFC is consistent, then $ZFC + \neg CH$ is consistent.*

Proof. Let $P = Fn(\aleph_2 \times \omega, \{0,1\})$, and let G be a generic set over P. As we proved in example 3.9, this means that $\cup G := f$ is a function from $\aleph_2 \times \omega$ onto $\{0,1\}$. Furthermore, consider any distinct $\alpha, \beta \in \aleph_2$, and define $D(\alpha,\beta) = \{p \in P \mid \exists n \in \omega\, [p(\alpha, n) \neq p(\beta, n)]\}$. Clearly, $D(\alpha,\beta)$ is dense: for any $p \in P$, there exists some $m \in \omega$ which is not in the domain of p. Define $q = p \cup \{\langle\langle\alpha, m\rangle, 0\rangle, \langle\langle\beta, m\rangle, 1\rangle\}$. Then $q \in D(\alpha, \beta)$.

By the generic set condition, G must intersect every such $D(\alpha,\beta)$, and hence there must be some n where $f(\alpha, n) \neq f(\beta, n)$. For fixed $\alpha \in \aleph_2$, f defines a subset of ω, and this says that all such subsets are distinct. So there are at least $\aleph_2^M$ subsets of ω in $M[G]$.

All that remains to check is that $\aleph_2^M = \aleph_2^{M[G]}$, and by our argument above we only need verify P has the countable chain condition. Let $Q \subset P$ be any uncountable subset of P, and suppose Q is an antichain. Observe that for a given finite domain, there are only finitely many distinct functions. So given that Q is uncountable, we must have that $\{\text{dom}(q) \mid q \in Q\}$ is also uncountable. Let B be any subset of $\{\text{dom}(q) \mid q \in Q\}$ whose cardinality is $\aleph_1$.

By the delta system lemma, B has some subset, also of size $\aleph_1$, with root R, and since every domain is finite, R must be finite. In other words, for any $x, y \in B$, the domain on which both x and y are defined is exactly R. We've assumed that $B \subset Q$ and Q is an antichain, so for any such x, y, there must be some $r \in R$ on which they disagree (only an element in the domain of both functions can prevent them from being compatible). But there are only finitely many ways for two functions to disagree on the finite set R. So two of the functions must be compatible, and neither B nor Q is an antichain. This completes the proof. □

In order to prove that ZFC + CH is consistent, we employ a different method. Suppose $2^{\aleph_0} = \aleph_\alpha$: then we will "collapse" $\aleph_\alpha$ onto $\aleph_1$, yielding a model where the continuum hypothesis holds. By design, the relevant forcing notion will *not* preserve cardinality, and we don't want it to. But we want to make sure that our forcing doesn't add any additional subsets of ω, or else the set that we collapse onto $\aleph_1$ might not be the power set of ω in $M[G]$.

Definition 4.23. Let P be a forcing poset. P is countably closed if, for any descending sequence $p_0 > p_1 > p_2 \cdots$ contained in P, there exists some $q \in P$ which is a lower bound for the sequence.

Lemma 4.24. *If P is countably closed and $G \subset P$ is generic, then $M[G]$ contains no additional subsets of ω not already in M.*

Proof. Pick any $A \subset \omega$ in $M[G]$, and let α be a name for A. By the truth lemma, there exists some $p \in G$ where $p \Vdash \alpha \subset \check{\omega}$. Then either $0 \in A$, or $0 \notin A$, so for simplicity suppose the former. Applying the truth lemma and the condition that G is a filter, there exists some $p_0 \leq p$ where $p_0 \in G$, $p_0 \Vdash \check{0} \in \alpha$. Similarly, some $p_1 < p_0$ exists which decides whether 1 is contained in A, and so on. Let q be the lower bound of this sequence (since P is countably closed): then for any n, $q \leq p_n$ and hence q decides whether n is contained in A. Within M, we can define $A' = \{n \in \omega \mid q \Vdash \check{n} \in \alpha\}$, and clearly $A' = A$. So every subset of ω contained in $M[G]$ must already be in M. □

Theorem 4.25. *If ZFC is consistent, then $ZFC + CH$ is consistent.*

Proof. Let P be the set of partial functions from $\aleph_1$ to $2^{\aleph_0}$ defined on some countable domain, and suppose $G \subset P$ is generic. By the exact same argument we used above, $\cup G := f$ is a function from $\aleph_1$ onto $2^{\aleph_0}$. Therefore, within $M[G]$, $(2^{\aleph_0})^M$ has cardinality no greater than $\aleph_1$.

Furthermore, we note that P is countably closed. This is because for any decreasing ω-sequence of functions defined on countable domains, their union is the union of countably many countable sets and hence must also be defined on a countable domain. So $(2^{\aleph_0})^{M[G]} = (2^{\aleph_0})^M$ must have cardinality no greater than $\aleph_1$. And by Cantor's Theorem, its cardinality must be strictly greater than $\aleph_0$. □

5. Coda: What Forcing Can't Do

We already know that the forcing extension $M[G]$ is a model of ZFC, so given any statement provable in ZFC, there is no generic set that forces its opposite. On the other hand, forcing seems powerful and general enough that it's not immediately clear what would get in our way if we tried to force a sentence that goes against the axioms. Abstractly, we know that given *any* poset contained in M, we can successfully create a forcing extension. So what would happen, for instance, if we tried to force the natural numbers to have a countable powerset? Beyond the fact that it wouldn't work, can we say specifically what gets in our way? Along these lines, we conclude with two examples.

Example 5.1. Recall that since our model M is countable, there exists a countable ordinal A which we can identify as the supremum of its ordinals. Then A is clearly not in M. Furthermore, A corresponds to a subset A' of ω, and we proved that no notion of forcing can add the subset A' to our model.

The barrier, in Cohen's terminology, is that forcing "broadens" the set universe rather than lengthening it. As we proved above, the name of a set in M has rank at least as great as the set itself does in $M[G]$. Because our new model $M[G]$ is recursively constructed out of objects already present in M, it avoids disrupting the ordinal structure of M.

Example 5.2. Suppose we naively try to construct a model where Cantor's Theorem fails, and $2^{\aleph_0}$ is countable. The natural poset to use would be $P := Fn(\omega, 2^{\aleph_0})$. If G is generic over P, then $M[G]$ is a model in which $(2^{\aleph_0})^M$ is countable: as previously, $f := \cup G$ is a function from ω onto $(2^{\aleph_0})^M$, and hence f enumerates what M takes to be the power set of ω.

However, note that P is *not* countably closed, and hence may add new subsets of ω. In fact, using a classic diagonalization argument, we can explicitly name at least one subset added by G. For any $n \in \omega$, let $A_n = f(n)$ be the subset of ω corresponding through f to n. Define:

$$Q = \{n \in \omega \mid n \notin A_n\}$$

Then clearly $Q \neq A_m$ for any m, or else we would immediately arrive at a contradiction. But given G, we can easily define Q. So Q is a subset of ω added by the generic set G.

This proof differs in one important way from the standard proof of Cantor's Theorem: instead of arguing from contradiction, we use an actual, complete enumeration of subsets in one model to generate a subset actually contained in another.

Acknowledgments

Thanks to my mentor, Ronno Das, for perceptive feedback and help refining the ideas in my paper. Thanks to Professor Maryanthe Malliaris and Diego Bejarano-Rayo for invaluable insight and intuition that helped me attain a grip on the subject. Thanks to Noah Schoem for patient and detailed resolutions to a variety of set-theoretic confusions. And thanks, of course, to Peter May, whose tireless work makes the REU possible.

References

[1] Chang, C.C. and H. Jerome Keisler. Model Theory. Dover Publications. 1990.
[2] Chow, Timothy Y. A beginner's guide to forcing. `https://arxiv.org/abs/0712.1320`
[3] Cohen, Paul. Set Theory and the Continuum Hypothesis. W.A. Benjamin. 1966.
[4] Hamkins, Joel David. The set-theoretic multiverse: a model-theoretic philosophy of set theory. `http://lumiere.ens.fr/~dbonnay/files/talks/hamkins.pdf`
[5] Jech, Thomas. Set Theory: The Third Millennium Edition, revised and expanded. Springer. 2002.
[6] Kanamori, Akihiro: Levy and Set Theory, Annals of Pure and Applied Logic 140 (2006) 233 - 252.
[7] Kunen, Kenneth. Set Theory. College Publications. 2011.