FORCING: AN INTRODUCTION TO INDEPENDENCE PROOFS IN SET THEORY

JOSHUA YAGUPSKY

Abstract. In 1963, Paul Cohen invented the technique of forcing as a way to prove indpendence results in set theory. He used this method to demonstrate that the continuum hypothesis, which states that the cardinality of the real numbers $2^{\aleph_0}$ is the smallest uncountable cardinal, is independent from the standard ZFC axioms of set theory. That is, assuming ZFC is a consistent theory, it cannot tell us whether or not the continuum hypothesis is true. Cohen's method allows one to easily prove independence results and construct models of set theory with desired properties. In this paper, we will examine the forcing method and use it to prove the independence of the continuum hypothesis.

Contents

1. Introduction

The notion that infinite sets could come in different sizes or cardinalities was one of the main achievements of the early set theorists. Cantor famously showed using the diagonal argument that, not only are there different sizes of infinity, but that there are infinitely many such sizes. In particular, he showed that there are more real numbers, or equivalently more subsets of natural numbers, than there are natural numbers. Since the cardinality of the natural numbers is the smallest infinite cardinality, this raises the obvious question of whether the cardinality of the real numbers is the second smallest. Cantor hypothesized that there was no intermediate cardinality, a statement now called the *continuum hypothesis*, but was unable to prove it. In 1940, Gödel showed that the continuum hypothesis was consistent with the standard ZFC axioms of set theory, assuming the axioms themselves were consistent. Finally, in 1963, Paul Cohen proved that it was also consistent for the continuum hypothesis not to be true. His method for proving the independence of the continuum hypothesis, known as *forcing*, is extremely flexible and can be used for a wide variety of independence results. We'll use the forcing method in this paper to show that the continuum hypothesis is independent of the ZFC axioms.

2. ZFC Set Theory

2.1. **ZFC Axioms.** Before we can prove any independence results, we'll review the basics of set theory from the perspective of first-order logic. We will assume the reader is familiar with the basics of first-order logic: if not, [7] is a good reference. Zermelo-Frankel set theory with Choice, or ZFC, is a first-order theory over a simple language called the *language of set theory*. This language features no function symbols and only one binary relation symbol, namely $\in$. The elements of a structure for the language of set theory are called *sets*. In this language, we often use *bounded quantifiers*: formulas of the form $\forall x \in y.\varphi(x)$ or $\exists x \in y.\varphi(x)$. These are to be interpreted as $\forall x\,(x \in y \implies \varphi(x))$ and $\exists x\,(x \in y \wedge \varphi(x))$, respectively. We also abbreviate $\neg(x \in y)$ as $x \notin y$. The theory of ZFC can be generated by a computable set of formulas known as *axioms*, most of which are quite simple. Despite its simplicity, set theory has enormous expressive power, to the point that we can talk about first order logic and model theory from within set theory. The fact that we can do this means that ZFC can talk about its own models and their properties, such as whether or not ZFC is consistent. As a result of the Gödel incompleteness theorems, which we will not discuss in much detail here, if ZFC is consistent then ZFC cannot prove that it is consistent. However, ZFC can prove *relative consistency results*, which state that if ZFC is consistent so is some other theory, typically ZFC augmented by some additional sentences.

Definition 2.1 (ZFC)**.** ZFC set theory (Zermelo-Fraenkel Set Theory with Choice) is the first-order theory in the language of set theory generated by the following axioms:

(1) Extensionality: Two sets are equal if they have the same elements: $\forall x.\forall y.\,(\forall z.\,(z \in x \iff z \in y)) \implies x = y$.
(2) Union: For any set, the union of its elements is also a set: $\forall x.\exists y.\forall z.(z \in y \iff \exists a \in x.(z \in a))$.
(3) Foundation: Every nonempty set contains an element it is disjoint from: $\forall x.(\exists y.(y \in x) \implies \exists y \in x.\forall z \in y.(z \notin x))$.
(4) Power Set: For any set, there exists a set whose elements are exactly the subsets of the original set: $\forall x.\exists y.\forall z(z \in y \iff \forall a \in y(a \in x))$.

(5) Infinity: There is a set which contains a set with no elements and is closed under the map $x \mapsto x \cup \{x\}$:
$\exists x. ((\exists y \in x. \forall a(a \notin y)) \wedge \forall y \in x. \exists z \in x. \forall a(a \in z \iff (a \in y \vee a = y)))$.
(6) Choice: For any family of disjoint, nonempty sets, there exists another set which contains exactly one element from each member of the family:
$\forall x. (\forall y \in x.((\exists z. z \in y) \wedge \exists! w \in x. \exists a \in w.(a \in y)) \implies \exists w. \forall y \in x. \exists! a \in w(a \in y))$.
(7) Replacement: Let φ be a formula in the language of set theory with free variables x, y, $\overline{w}$, and z, where $\overline{w}$ is a finite list of variables $w_1, \dots w_n$. Then the following is an axiom: for all $w_1, \dots w_n$ and for any A, if for each element of A there exists a unique y such that $\varphi(x, y, \overline{w}, A)$ holds, then there is a set B such that for any y, $y \in B$ iff there is some $x \in A$ satisfying $\varphi(x, y, \overline{w}, A)$. B is said to be the *image of A under φ with parameters $\overline{w}$*:
$\forall \overline{w}. \forall A. ((\forall x \in A. \exists! y. \varphi(x, y, \overline{w}, A)) \implies \exists B. \forall y.(y \in B \iff \exists x \in A. \varphi(x, y, \overline{w}, A)))$.

Note that Replacement is an *axiom schema*: it is not a single axiom, but an infinite list of axioms. If φ satisfies the first part of the implication in Replacement, we call it a *functional*, and the set B we create is called the *image* of A under φ.

Lemma 2.2 (Pairing and Separation)**.** *The following are theorems of ZFC set theory:*

(1) Pairing: For any sets x and y, there is a set whose only elements are x and y:
$\forall x. \forall y. \exists z. \forall a.(a \in z \iff (a = x \vee a = y))$.
(2) Separation: Let φ be a formula in the language of set theory with free variables $x, \overline{w}, y$, where $\overline{w}$ is a finite list of variables $w_1, \dots w_n$. Then the following is a theorem of ZFC: for all $w_1, \dots w_n$ and for any A, there exists a set B whose elements consist of all and only those sets $x \in A$ such that $\varphi(x, \overline{w}, A)$ holds:
$\forall w_1, \dots w_n. \forall A. \exists B. \forall x. (x \in B \iff (x \in A \wedge \varphi(x, \overline{w}, A)))$.

Proof. We'll first show Pairing holds. The axiom of Infinity tells us there is a set containing a set that has no elements and is closed under the operation $x \mapsto x \cup \{x\}$. Applying this twice gives us that there exists a whose only elements are a set with no elements and a set whose only element is a set with no elements. We'll denote this set 2, for reasons that we'll discuss later. Let $\varphi(x, y, w_1, w_2)$ be the formula $(y = w_1 \wedge \neg \exists z. z \in x) \vee (y = w_2 \wedge \exists z. z \in x)$. This formula is a functional on 2 for any assignment of parameters w_1 and w_2, and the image of 2 under this functional is a set whose only elements are w_1 and w_2. This gives us Pairing. We can use a similar trick to prove Separation. Given a set, we want to obtain the set of all and only those elements which satisfy a formula φ. Infinity tells us that there is a set with no elements, and Pairing tells us that for any set a there is a set whose only element is a. Suppose we are peforming Separation on a set A. Given φ, we can write down a formula that will replace each element where φ holds with its singleton, and replace all other elements with the empty set. This formula is clearly a functional on A, and taking the union of its image will give us the desired set B. □

The ZFC axioms are quite unwieldy to deal with directly, and we rarely will except when proving certain structures are models of ZFC. The language of set theory is not suitable for expressing interesting set theoretic notions because it only contains one relation symbol and no constant symbols. Fortunately, there's nothing stopping us from augmenting the language of set theory with newly defined relation and function symbols. There is a fairly obvious way to do this: we can define a new n-ary relation symbol as a formula with n free variables, and an n-ary function symbol as a formula $\varphi(\overline{w}, y)$ with $n+1$ free variables such that ZFC proves that for all variable assignments for the first n variables there exists a unique y such that $\varphi(\overline{w}, y)$ holds. It's also not difficult to show that every sentence in this new language can be translated into an equivalent sentence in the original language of set theory. We can even interatively define new notions in terms of previously defined notions. Here are some of the notions that will be most helpful when proving results in ZFC:

Definition 2.3 (Basic Set-Theoretic Definitions)**.** Here we will define the most common function and relation symbols we will use in set theory. When defining the function symbols, we will briefly note how to prove that the definition sentence succeeds.

$x \subseteq y$: The binary relation $x \subseteq y$ is defined as $\forall z \in x.(z \in y)$. This allows Extensionality to be rewritten as $\forall x.\forall y.(x \subseteq y \wedge y \subseteq x) \implies x = y$.

$\emptyset$: The constant symbol $\emptyset$ is defined by the formula $\neg\exists x.(x \in \emptyset)$. Existence is guaranteed by Infinity, and uniqueness is guaranteed by Extensionality.

$\mathcal{P}(x)$: The unary function symbol $\mathcal{P}(x)$ is defined by the formula $\forall x.\forall y.(y \in \mathcal{P}(x) \iff y \subseteq x)$. Existence is guaranteed by Power Set, and uniqueness is guaranteed by Extensionality.

$\{x_1, \ldots x_n\}$: For each n, the n-ary function symbol $\{x_1, \ldots x_n\}$ is defined by the formula $\forall x_1, \ldots x_n.\forall x.(x \in \{x_1, \ldots x_n\} \iff (x = x_1 \vee \ldots x = x_n))$. Existence is guaranteed by repeatedly applying Infinity to obtain sets with exactly n distinct elements, and then using Replacement to obtain a set with the desired elements. Foundation implies no set can be an element of itself, so the set $x \cup \{x\}$ contains all of the elements of x together with a new element which is not in x. Uniqueness is guaranteed by Extensionality.

$\bigcup A$: The unary function symbol $\bigcup A$ is defined by the formula $\forall A.\forall x.(x \in \bigcup A \iff \exists a \in A.(x \in a))$. Existence is guaranteed by Union, and uniqueness is guaranteed by Extensionality. The binary function symbol $x \cup y$ is defined as $\bigcup\{x, y\}$.

$\bigcap A$: The unary function symbol $\bigcap A$ is defined by the formula $\forall A.\forall x.(x \in \bigcap A \iff \forall a \in A.(x \in a))$. Existence is guaranteed by Separation, and uniqueness is guaranteed by Extensionality. Note that this fails if A is empty, so as a convention we define $\bigcap \emptyset = \emptyset$. The binary function symbol $x \cap y$ is defined as $\bigcap\{x, y\}$.

$\{\tau \mid \varphi\}$: Let $\varphi(x, w_1, \ldots w_n)$ be a formula in the (possibly extended) language of set theory, and let $\tau(x, w_1, \ldots w_n)$ be a term in the extended language. Then the n-ary function symbol $\{\tau(x, w_1, \ldots w_n) \mid \varphi(x, w_1, \ldots w_n)\}$ is defined to be $\emptyset$ if there is no set containing all x such that $\varphi(x, w_1, \ldots w_n)$ holds, and is otherwise defined as the set whose elements are exactly each of the $\tau(x, w_1, \ldots w_n)$ for all and only those x satisfying $\varphi(x, w_1, \ldots w_n)$ holds. Existence is guaranteed by Replacement and Separation, and uniqueness by Extensionality. One case where this definition will always succeed is if our formula on the right hand side is of the form $(x \in A) \wedge \varphi(x, w_1, \ldots w_n, A)$ because of Separation. We often write

$$\{\tau(x, w_1, \ldots w_n, A), x \in A \mid \varphi(x, w_1, \ldots w_n, A)\}$$

as a shorthand for $\{\tau(x, w_1, \ldots w_n, A) \mid (x \in A) \wedge \varphi(x, w_1, \ldots w_n, A)\}$, in order to indicate that this definition is always guaranteed to succeed.

$A \setminus B$: The binary function symbol $A \setminus B$ is defined as $\{x \in A \mid x \notin B\}$.

(x, y) : The binary function symbol (x, y) is defined as $\{x, \{x, y\}\}$. It is a theorem of ZFC that $(x, y) = (x', y')$ iff $x = x'$ and $(y = y')$, justifying the use of this particular set as the definition of an ordered pair. We can also define the two projection operations $\pi_0(A) = \bigcap A$ and $\pi_1(A) = \bigcup A \setminus \pi_0(A)$: it is not hard to show that $\pi_0(x, y) = x$ and $\pi_1(x, y) = y$.

$A \times B$: The Cartesian product $A \times B$ is be defined as $\{(x, y) \mid x \in A \wedge y \in B\}$. The definition of the Cartesian product always succeeds, since we know that for any $y \in B$ the definition $\{(x, y) \mid x \in A\}$ succeeds, and $A \times B$ is clearly the same as $\bigcup \{\{(x, y) \mid x \in A\} \mid y \in B\}$. Subsets of $A \times B$ are called *binary relations* from A to B. The domain of a binary relation is the set of all zeroeth projections of each of its ordered pairs, and the range is the set of all first projections of each of its ordered pairs.

$\bigsqcup A$: The disjoint union $\bigsqcup A$ is defined as $\{(X, x) \mid X \in A, x \in X\}$. This definition always succeeds, since for any $X \in A$ the definition $\{(X, x) \mid x \in X\}$ always succeeds, and $\bigsqcup A$ is clearly the same as $\bigcup \{\{(X, x) \mid x \in X\} \mid X \in A\}$. Note that the family of sets consisting of each set $\{(X, x) \mid x \in X\}$ for all $X \in A$ is a pairwise disjoint family.

$f\colon A \longrightarrow B$: A set f is a *function* from a set A to a set B if it is a binary relation from A to B satisfying $\forall x \in A. \exists! y \in B. (x, y) \in f$. For each function f, we define a unary function symbol $f(x)$ to equal the empty set if x is not in the domain of f and the unique y such that $(x, y) \in f$ otherwise. We also define the image and preimage function symbols $f[X] = \{f(x) \mid x \in X\}$ and $f^{-1}[X] = \{x \in \operatorname{dom}(f) \mid f(x) \in X\}$. The set of all functions from A to B, denoted ${}^{A}B$, exists because we can write it as $\{f \in \mathcal{P}(A \times B) \mid f \text{ is a function}\}$. We often range over this set by writing $f\colon A \longrightarrow B$ rather than $f \in {}^{A}B$.

$\bigcap_{\varphi(x, \overline{w})} \tau(x, \overline{w})$: If $\varphi(x, \overline{w})$ is a formula in the language of set theory with $n+1$ free variables and $\tau(x, \overline{w})$ is a term with $n+1$ free variables then $\bigcap_{\varphi(x, \overline{w})} \tau(x, \overline{w})$ is a n-ary function symbol defined as the empty set if no x satisfies $\varphi(x, \overline{w})$, and otherwise as the intersection of all $\tau(x, \overline{w})$ satisfying $\varphi(x, \overline{w})$. This is well defined since for any x satisfying $\varphi(x, \overline{w})$ the intersection of $\tau(x, \overline{w})$ with all $\tau(y, \overline{w})$ such that y satisfies $\varphi(y, \overline{w})$ exists, and the result is the same no matter which x we pick.

$\mathbb{N}$: $\mathbb{N}$ is a constant symbol defined as $\bigcap_{\emptyset \in x \wedge \forall y \in x. (y \cup \{y\} \in x)} x$. Such an x does exist by the axiom of Infinity. We define sets for each of the natural numbers in the expected way: zero is assigned to the empty set, and if a natural number n is assigned a set x, then its successor is assigned $x \cup \{x\}$. Then it's easy to see by induction that $\mathbb{N}$ is the set of all natural numbers.

Indexed Family : If I is a set, a *family of sets indexed by* I is a function f with domain I. For each $i \in I$, we often write $f(i)$ as simply A_i. We write $\bigcup_{i \in I} A_i$ and $\bigcap_{i \in I} A_i$ to mean $\bigcup f[I]$ and $\bigcap f[I]$ respectively. We define the indexed disjoint union as $\bigsqcup_{i \in I} A_i = \{(i, x) \mid i \in I, x \in A_i\}$: this agrees with our previous disjoint union definition if we consider a set as a family indexed by itself. We define the indexed Cartesian product $\prod_{i \in I} A_i$ as $\left\{g\colon I \longrightarrow \bigcup_{i \in I} A_i \mid \forall i \in I. g(i) \in A_i\right\}$.

2.2. **Ordinals and Recursion.** We will now turn our attention to the notion of *well orderings* and *ordinals*. Well orders generalize the ordering relation on the natural numbers to all kinds of sets, and allow us to perform the same kind of inductive arguments and construct objects by recursion. Ordinals are special sets which serve as natural representatives of well-orders, and have useful set-theoretic properties. We will see that the whole universe of sets has a clean description in terms of recursion on ordinals. Ordinals also give us a natural way to talk about cardinalities of sets: the *well-ordering theorem*, which is a consequence of the axiom of choice, tells us that any set can be well-ordered, and therefore every set has a bijection with some ordinal. We will begin by examining the notion of a *transitive set*.

Definition 2.4 (Transitive Set)**.** A set A is said to be *transitive* if for all $a \in A$ and all $x \in a$ we have $x \in A$. In other words, A is transitive iff $\forall a \in A.(a \subseteq A)$.

Lemma 2.5 (Properties of Transitive Sets)**.** *A set A is transitive iff $\bigcup A \subseteq A$ iff $A \subseteq \mathcal{P}(A)$. If A is a transitive set, so are $\bigcup A$ and $\mathcal{P}(A)$. If B is a set of transitive sets, then $\bigcup B$, $B \cup \bigcup B$, and $\bigcap B$ are transitive.*

Proof. Just apply the definition of transitivity. □

Definition 2.6 (Well-Founded Relation)**.** Let A be a set, and let R be a binary relation on A. We say that (A, R) is *well-founded* if every nonempty subset $B \subseteq A$ has an *R-minimal element*: an element $b \in B$ such that no $c \in B$ satisfies cRb.

Well-founded relations are not quite as nice as well-orders, but they still have a useful notion of induction. Additionally, the axiom of Foundation tells us that the element relation $\in$ is well-founded on every set, as we will see in a moment. This gives us a useful induction on sets, called $\in$-induction.

Lemma 2.7 (Induction on Well-Founded Sets)**.** *Let A be a set and let R be a well-founded binary relation on R. If $\varphi(x, \overline{w})$ is a formula in the language of set theory and $\forall x \in A.\,((\forall y \in A.(yRx \implies \varphi(y, \overline{w}))) \implies \varphi(x, \overline{w}))$ then $\varphi(x, \overline{w})$ holds for all $x \in A$.*

Proof. By Separation, consider the set of all $x \in A$ such that $\varphi(x, \overline{w})$ does not hold. If this set is nonempty, it has an R-minimal element, which we'll call x. Since x is R-minimal, every $y \in A$ such that yRx must not be in this set, and therefore $\varphi(y, \overline{w})$ must hold. But by the inductive hypothesis, this means $\varphi(x, \overline{w})$ holds, which is a contradiction. Therefore, $\varphi(x, \overline{w})$ holds for all x. □

Theorem 2.8 ($\in$-Induction)**.** *For each set X, the relation $\in$ restricted to the elements of X is well-founded. Additionally, let $\varphi(x, \overline{w})$ be a formula in the language of set theory. Then the following sentence is a theorem of ZFC:*

$$\forall \overline{w}.\,((\forall x.((\forall y \in x.\varphi(y, \overline{w})) \implies \varphi(y, \overline{w}))) \implies \forall x.\varphi(x, \overline{w}))$$

Proof. First, we'll prove $\in$ is well-founded. Let A be any set, and let B be some nonempty subset of A. Foundation tells us there is an element of B which is disjoint from B: that is, there is some $b \in B$ such that any $c \in b$ is not in B. This is an $\in$-minimal element of B. Now suppose that the hypothesis of the $\in$-induction holds. Pick any set A: this will be well-founded with respect to $\in$, and the hypothesis of well-founded induction holds for A. Therefore, for each $x \in A$, φ holds, and by the hypothesis of $\in$-induction, A must satisfy φ as well. Therefore, every set satisfies φ. □

Definition 2.9 (Linear Order)**.** Let A be a set, and let $\prec$ be a binary relation on A. We say that $(A, \prec)$ is a (strict) *linear order* if the following requirements hold:

(1) Trichotomy: For all $a, b \in A$ exactly one of the statements $a \prec b$, $b \prec a$, or $a = b$ hold.
(2) Transitivity: For all $a, b, c \in A$, if $a \prec b$ and $b \prec c$ then $a \prec c$.

We define the non-strict version of the linear order to be $a \preceq b \iff (a \prec b) \vee (a = b)$.

Definition 2.10 (Well-Order)**.** Let A be a set, and let $\prec$ be a binary relation on A. We say that $(A, \prec)$ is a (strict) *well order* if it is a linear order and $\prec$ is well-founded on A. By Trichotomy, a minimal element of a subset $B \subseteq A$ is actually a *least element*: an element $b \in B$ such that every $c \in B$ satisfies $b \preceq c$.

Definition 2.11 (Ordinal)**.** A set A is called an *ordinal* if A is transitive and $(A, \in)$ is a well-ordered set.

Lemma 2.12 (Equivalent Definitions of Ordinals)**.** *Let A be any set. Then the following are equivalent:*

(1) A is an ordinal.
(2) A is transitive and $\in$ linearly orders A.
(3) A is transitive and $\in$ is trichotomous on A.
(4) A is transitive and all of its elements are ordinals.
(5) A is transitive and all of its elements are transitive.

Proof. If $\in$ linearly orders a set, we know that it well-orders the set since $\in$ is always wellfounded. Any transitive set linearly ordered by $\in$ is therefore an ordinal. If $\in$ is merely trichotomous, we can still prove transitivity. Suppose A is a transitive set and $\in$ is trichotomous on A, and suppose $a, b, c \in A$ satisfy $a \in b$ and $b \in c$. Applying Foundation to the set $\{a, b, c\}$ tells us that $c \neq a$ and $c \notin a$. But since $\in$ is trichotomous on A, $a \in c$, so $\in$ is transitive on A. Now if A is an ordinal, then each element $a \in A$ is also a subset of A, and is therefore well-ordered by $\in$. The set a is also transitive: if we have an element $b \in a$, then we have $b \in A$ and $b \subseteq A$. Therefore, any $c \in b$ is also in A, and since $\in$ is transitive on A, we also have $c \in a$. Thus, any ordinal is a transitive set of ordinals, and is also a transitive set of transitive sets. Now let A be transitive, and let each member of A also be transitive. We will say that elements $x, y \in A$ are incomparable if trichotomy fails. Suppose A is not an ordinal. Then A has incomparable elements: use Foundation to pick an x which is $\in$-minimal among the elements of A incomparable with some element of A, and apply Foundation again to pick a y which is $\in$-minimal among the elements of A incomparable with x. Any $z \in y$ is a subset of y because y is transitive, and by minimality z must be comparable with x. But since y is incomparable with x, we can rule out $z = x$ or $x \in z$, so $z \in x$. This works no matter which z we choose, so $x \subseteq y$. Applying the same logic in the other direction gives $y \subseteq x$ and therefore $x = y$, contradicting incomparability. $\in$ is therefore trichotomous on A, and since A is transitive, it is an ordinal. □

The ordinals are a good example of a *proper class*, which is a collection of sets too large to form a set. The class of all sets is also a proper class, since if there was a set of all sets, it would contain itself, contradicting Foundation. Similarly, if there was a set of all ordinals, then it would necessarily be transitive, since each element of an ordinal is an ordinal. Therefore, the set would be a transitive set of ordinals, and therefore an ordinal, which again contradicts Foundation. We will often use bounded quantifiers for proper classes: if we want to quantify over all ordinals, we can write $\forall \alpha \in \text{Ord}.\varphi(\alpha)$ to mean $\forall \alpha.(\alpha \in \text{Ord} \implies \varphi(\alpha))$, where $\alpha \in \text{Ord}$ is the formula saying that α is an ordinal. Our convention will be to implicitly quantify over ordinals when using a Greek letter variable.

Lemma 2.13 (Properties of Ordinals)**.** *For any ordinal α, its* successor $\alpha + 1 = \alpha \cup \{\alpha\}$ *is an ordinal, and is the smallest ordinal strictly greater (in the sense of $\in$) than α. The union of any set of ordinals is an ordinal. Every natural number is an ordinal, and aside from* 0*, every natural number is a successor of some smaller ordinal. The set of natural numbers $\mathbb{N}$ is an ordinal, denoted ω, and is the smallest* limit ordinal *(an ordinal which is not the empty set or the successor of some smaller ordinal). The class of ordinals* Ord *is linearly ordered by $\in$: any two ordinals are either equal, or one is an element of the other. It is even well-ordered: for any formula $\varphi(x, \overline{w})$, if $\exists \alpha.\varphi(\alpha, \overline{w})$ then there exists a least such α with respect to $\in$.*

Proof. Let α be an ordinal: it is a transitive set of transitive sets. We therefore know that $A = \alpha$ will be a set of transitive sets, and therefore $A \cup \bigcup A = \alpha \cup \{\alpha\} = \alpha + 1$ is transitive. Each element of this set is either α, which is transitive, or an element of α, which is transitive, meaning this must be an ordinal as well. Any ordinal β such that $\alpha \in \beta$ must satisfy $\{\alpha\} \subseteq \beta$ and $\alpha \subseteq \beta$ because β is transitive. Thus $\alpha + 1 \subseteq \beta$, and $\alpha + 1$ is the smallest ordinal greater than α. The empty set, or 0, is clearly an ordinal. Since ω is the smallest set that contains 0 and is closed under ordinal successor, it can only contain ordinals which are either zero or successors. ω is a limit ordinal since it could only be the successor of one of its elements, but ω is closed under successor. The union of a set of ordinals is transitive since it is a union of a set of transitive sets, and each element is a transitive set since each ordinal only has elements which are transitive. Thus, the union of a set of ordinals is an ordinal.

We'll now show that ordinals are linearly ordered. First, if α and β are ordinals, then we'll show $\alpha \in \beta$ iff $\alpha \subset \beta$. Clearly $\alpha \in \beta$ implies $\alpha \subseteq \beta$ by transitivity, and no set is a member of itself. If $\alpha \subset \beta$, then there is a smallest element $x \in \beta$ which is not in α. Every element of α must be strictly less than x by trichotomy: x cannot equal an element of α, and if x was less than an element of α then transitivty would imply $x \in \alpha$. This gives $\alpha \subseteq x$. Every element of x must be in α by minimality, so we have $x = \alpha$. Now suppose we have any two ordinals α and β. Let γ be their intersection: this is clearly an ordinal, as it is a transitive set of transitive sets. If $\alpha \setminus \gamma$ and $\beta \setminus \gamma$ are both nonempty, then we have $\gamma \in \alpha$ and $\gamma \in \beta$ by our previous result. Taking intersections gives $\gamma \in \gamma$, which is a contradiction, so one of the two sets must be empty. This means either $\alpha \subseteq \beta$ or $\beta \subseteq \alpha$, and therefore $\alpha \in \beta$ or $\alpha = \beta$ or $\beta \in \alpha$. This also tells us that ω is the smallest limit ordinal, since any smaller limit ordinal would be a natural number, which is always either 0 or a successor ordinal. The fact that the ordinals are well-ordered follows from the fact that each ordinal is a well-ordered set, so if some ordinal has a property, either none of its elements have that property, meaning it is the smallest such ordinal, or one of its elements does, in which case we can find the smallest such ordinal. □

We'll now examine how ordinals can encode all well-ordered sets.

Lemma 2.14 (Well-Order Isomorphism)**.** *Let $(A, \prec_A)$ be a well-ordered set. If $(B, \prec_B)$ is a linearly ordered set which is isomorphic to $(A, \prec_A)$, then $\prec_B$ is a well-order on B and the isomorphism is unique.*

Proof. We can find the minimum of any subset of B by taking the minimum of its image in A and mapping it back to B under the isomorphism. Any two isomorphisms are equal because there can be no smallest element where the two isomorphisms disagree. See [7] 63.7 for details. □

Definition 2.15 (Initial Segment)**.** Let $(A, \prec)$ be a well-order. If $a \in A$ is any element, then the *initial segment* of A below a is the set $A_a = \{x \in A \mid x < a\}$ together with the order relation $\prec$ restricted to A_a. Clearly it is also a well-order.

Lemma 2.16 (Isomorphisms with Initial Segments)**.** *No well-ordered set is isomorphic to an initial segment of itself.*

Proof. By induction we can show that any such isomorphism would satisfy $f(x) = x$ for all x, and therefore this would be the identity function. The image is therefore the whole set, not any initial segment. See lemma 63.9 in [7] for a full proof. □

Theorem 2.17 (Fundamental Theorem of Well-Orders)**.** *Every well-ordered set is isomorphic to a unique ordinal under the $\in$ relation, called the set's* order type*. This isomorphism is unique.*

Proof. Uniqueness of the isomorphism is guaranteed by the existence of an isomorphism, and the uniqueness of the ordinal comes from the fact that for any distinct ordinals α and β, one is an initial segment of the other, and there cannot be an isomorphism between a well-ordered set and one of its initial segments. For existence, we'll proceed by induction. If $(A, \prec)$ is our set, consider the set of all $a \in A$ such that the initial segment A_a is not isomorphic with an ordinal. Let x be the least element of this set. For all $y \prec x$, there is a unique isomorphism between A_y and some ordinal γ_y. Using Replacement, take the set off all such isomorphisms for each $y \prec x$ and take the union. If $z \prec y$, then the isomorphism from A_y to γ_y restricts to an isomorphism, which by uniqueness is the isomorphism from A_z to γ_z. Therefore, this union of functions will be a function. Its range will be a union of ordinals, and is therefore and ordinal. It is also clearly an isomorphism, since it is an isomorphism below each $y \prec x$. Therefore, we have an isomorphism between A_x and some ordinal, a contradiction. Therefore, each initial segment of A is uniquely isomorphic to an ordinal. Taking the union of all of these isomorphisms gives us an isomorphism from A to an ordinal. □

Ordinals, and well-founded relations more generally, can be used to define functions and terms via *recursion.* This generalizes the notion of defining functions on the natural numbers via recursion, and depends both on induction with well-founded relations and the Replacement schema. We'll first examine how to provide recursive definitions for functions on well-founded sets, before turning our attention to recursively defining terms on Ord.

Theorem 2.18 (Well-Founded Recursion)**.** *Let A be a set, and let $R \subseteq A^2$ be a well-founded relation. Suppose we have some term $\tau(x, y, \overline{w})$, with $n + 2$ free variables. Then for every choice of parameters $\overline{w}$ there is a unique function f defined on A such that $f(x) = \tau\left(x, f(\cdot, \overline{w}) \restriction_{\{y \in A \mid yRx\}}, \overline{w}\right)$ for all $x \in A$.*

Proof. Uniqueness is easy: if there were two distinct functions, there would be a minimal element where they disagreed, but they would agree below that element, contradicting the recursive definition. For existence, we'll say a subset of A is called an initial segment if it is closed under R: if an element x is in the initial segment and yRx then y is also in the initial segment. Note that this agrees with our definition of initial segments if R is a well-order. Call a function $f\colon I \longrightarrow A$ defined on an initial segment of A an *approximation* if it satisfies $f(x) = \tau\left(x, f\restriction_{\{y\in I \mid yRx\}}, \overline{w}\right)$ for all $x \in I$. Since I is closed under R, we can replace the $y \in I$ in the set comprehension with $y \in A$. Note that since R is also well-founded on I, if there is such an approximation for an initial segment it is unique. Define f as the union of all such approximations over all initial segments of A. We will show f is a function that satisfies the recursive property. If, for some $x \in A$, we have $(x, y) \in f$ and $(x, y') \in f$, then there must be approximations $g\colon I_g \longrightarrow A$ and $h\colon I_h \longrightarrow A$ defined on initial segments containing x with $g(x) = y$ and $h(x) = y'$. If we take the intersection of their domains, we get another initial segment containing x, and restricting both g and h to this intersection gives us two apprxomations on this initial segment. They therefore must agree on this initial segment, so $y = y'$.

f is therefore a function on some subset of A. This subset must be an initial segment, since it is the union of initial segments. We also know that f is an approximation, since its value at any $x \in R$ is the same as that of an approximation which equals f on some initial segment containing x. If f is not defined on all of A, we can find a minimal $x \in A$ where f is not defined. For any $y \in A$ such that yRx, f will be defined on y by minimality. Adding x to the domain of f will therefore preserve the fact that the domain is an initial segment. Define g as a function on the domain of f together with x such that $g(y) = f(y)$ when y is in the domain of f and set $g(x) = \tau\left(x, f\restriction_{\{y\in A \mid yRx\}}, \overline{w}\right)$. Clearly this is an approximation, so f must contain x in its domain, which is a contradiction. Therefore, f is defined on all of A. □

Corollary 2.19 (Mostowski Collapse Lemma)**.** *Let A be a set, and let R be a well-founded relation. If R is* extensional *on A, meaning $\forall a \in A. \forall b \in A.(\forall x \in A.(xRa \iff xRb) \implies a = b)$, then there exists a unique transitive set B and a unique function $f\colon A \longrightarrow B$ such that f is an isomorphism between (A, R) and $(B, \in)$. That is, f is a bijection and for all $a, b \in A$ we have aRb iff $f(a) \in f(b)$.*

Proof. Well-founded recursion tells us there is a unique function f defined on A such that for all $x \in A$ we have $f(x) = \{f(y) \mid y \in A \land yRx\}$. This condition is equivalent to $aRb \iff f(a) \in f(b)$ for all $a, b \in A$, so if we can show f is a bijection we will have uniqueness. Let B be the image of f. f is clearly surjective, so all we need to show is that f is injective and it is transitive. Injectivity can be proved by induction. Suppose that for all $y \in A$ such that yRx, f is injective on y, meaning that for any $z \in A$ if $f(z) = f(y)$ then $z = y$. We'll show that f is injective on x. Pick some z such that $f(x) = f(z)$. This means $\{f(y) \mid y \in A \land yRx\} = \{f(y) \mid y \in A \land yRz\}$. By the inductive hypothesis, f is injective on each y when yRx. Since these two sets are equal, we know that the set of $y \in A$ such that yRx equals the set of all $y \in A$ such that yRz. Since R is extensional, $x = z$, and f is injective on x. Therefore, f is injective on each $x \in A$. B is transitive since it is the image of f, and each $f(x)$ is a set consisting entirely of $f(y)$ for some $y \in A$, which is also in B. □

Because well-orders are always well-founded, we can perform recursion on well-ordered sets. This recursion method can also be extended to proper classes. However, we have to be somewhat careful, since this will be a term that we are defining, not a function, and we cannot quantify over terms. However, we can prove existence and uniqueness of the term we are defining as a *metatheorem* or *theorem schema*.

Theorem 2.20 (Transfinite Recursion for Classes)**.** *Suppose X is a class, or formula with one free variable, and R is a class relation, or formula with two free variables. Suppose R is well-founded on X, in the sense that every set whose elements are all in X has an R-minimal element, and set-like, which means for each $x \in X$ there exists a set whose elements are exactly all of the $y \in X$ such that yRx. Let $\tau(x, y, \overline{w})$ be a term with $n+2$ free variables. Then there is a term $\sigma(x, \overline{w})$ such that for every $x \in X$ we have $\sigma(x, \overline{w}) = \tau(x, \sigma(\cdot, \overline{w}) \restriction_{\{y \in X | yRx\}}, \overline{w})$, where σ restricted to a set is just the corresponding function. Our term σ is unique in the sense that if any other term satisfies these properties then it agrees with σ for all $x \in X$.*

Proof. Uniqueness follows from well-founded induction: if the two terms disagreed somewhere they would disagree on an R-minimal element of X, but they would have to agree on everything below this element, contradicting the recursive property. We can define the term explicitly as follows: we'll once again call a set of elements of X an *initial segment* if it is downward closed under R. There is a formula expressing that a set is an initial segment. Any initial segment is well-ordered by R, and importantly for each x in the initial segment the set of all $y \in X$ such that yRx is a subset of the initial segment. We know by well-founded recursion that there will exist a unique function on any initial segment which satisfies the recursive relation: call such a function an approximation. Every element of X is contained in an initial segment by induction. Therefore, we can define a term $\sigma(x, \overline{w})$ as the unique y such that every approximation on an initial segment which contains x has the value y on x. The argument that this y is unique and that our term satisfies the recursive property is exactly the same as for well-founded recursion on a set. □

We'll now provide the most important recursive definitions we'll use in set theory:

Definition 2.21 (Von Neumann Universe and Rank)**.** For any ordinal α, we define the α-th *level* by recursion:

$$V_\alpha = \bigcup_{\beta < \alpha} \mathcal{P}(V_\beta)$$

That is, the α-th level is the union of the powersets of each of the β-th levels for all ordinals $\beta < \alpha$. This definition by transfinite recursion works because the ordinals together with the $\in$ relation form a well-founded class, and $\in$ is clearly set-like. The *Von Neumann Universe* V is the proper class consisting of all sets which are elements of some V_α: in other words, $V = \bigcup_{\alpha \in \mathrm{Ord}} V_\alpha$. The *rank* of a set is defined by transfinite recursion: it is equal to the union of the successors of the ranks of each of its elements. This definition succeeds since the proper class of all sets together with the relation $\in$ is well-founded and set-like. By induction, it's easy to see that the rank of a set is always an ordinal. We denote the rank of a set $\mathrm{rank}\,(x)$.

Theorem 2.22 (Von Neumann Hierarchy Theorem)**.** *Every set is in the Von Neumann Universe, and the rank of a set x is the smalles ordinal α such that $x \subseteq V_\alpha$. Each V_α is a transitive set, and the rank of each ordinal is itself.*

Proof. Apply $\in$-induction, induction on α, and the recursive definition of the V_α. See proposition 64.22 in [7] for more details. □

2.3. **Cardinals.** We'll now turn to the main reason we are interested in studying set theory in the first place: infinite cardinalities. Often we think of cardinals as equivalence classes of sets under the *equinumerosity* relation: two sets are said to be equinumerous if there is a bijection between them. The axiom of Choice will be crucial here, because it will allow us to represent these equivalence classes as ordinals. Without the axiom of Choice, we can still define cardinals, but we run into several difficulties: the cardinals being linearly ordered, for instance, is equivalent to Choice.

Theorem 2.23 (Well-Ordering Theorem)**.** *Every set can be well-ordered.*

Proof. This is a well-known equivalent of the Axiom of Choice. For a proof, see theorem 69.6 in [7] for details. □

Definition 2.24 (Cardinality and Cardinals)**.** The *cardinality* of a set, denoted $|X|$, is the smallest ordinal α such that there exists a bijection from X to α. A *cardinal* is an ordinal α which equals its own cardinality.

Lemma 2.25 (Properties of Cardinals)**.** *The cardinality of any set is a cardinal. If X, Y are any sets, then $|X| \leq |Y|$ iff there is an injection from X to Y. An ordinal α is a cardinal iff there is no surjection from a smaller ordinal $\beta \in \alpha$ onto α. Every finite ordinal is a cardinal. The union of any set of cardinals is a cardinal: in particular, ω is a cardinal. Any infinite cardinal is a limit ordinal.*

Proof. If $|X| = \alpha$, and $|\alpha| = \beta$, then there is a bijection between β and α and a bijection between α and X. Thus, there is a bijection between X and β. Since $|X| = \alpha$ and there is a bijection from β to X, $\beta \geq \alpha$. Clearly, the cardinality of any ordinal is at most itself, so $\beta \leq \alpha$, and therefore $\beta = \alpha$. Thus, the cardinality of X, which is α, satisfies $|\alpha| = \alpha$, and is therefore a cardinal.

If $|X| < |Y|$, then $|X| \subseteq |Y|$. The inclusion map from $|X|$ to $|Y|$ will be an injection, and composing this with the bijections between $|X|$ and $|Y|$ gives us an injection from X to Y. Now if there is an injection from X to Y, this also gives us an injection from $|X|$ into $|Y|$. This gives us a bijection from $|X|$ to a subset of $|Y|$. This subset is well-ordered, and will therefore be isomorphic to some ordinal γ. We can see by induction that each element of γ will be at most the corresponding element in $|Y|$, so since $|Y|$ is an ordinal we have $\gamma \leq |Y|$. But composing our bijection between $|X|$ and the subset of $|Y|$ with the order isomorphism gives us a bijection between $|X|$ and γ, so $||X|| = |X| \leq \gamma$. Therefore, $|X| \leq |Y|$.

If α is an ordinal and no smaller ordinal $\beta < \alpha$ surjects onto α, then α is the smallest ordinal equinumerous with α, so α is a cardinal. If α is an ordinal and $\beta \in \alpha$ is an ordinal, then $\beta \subseteq \alpha$, and $|\beta| \leq |\alpha|$ by the injectivity of the inclusion map. If β surjects onto α, the axiom of choice tells us that this is equivalent to α injecting into β, or $|\beta| \leq |\alpha|$, so $|\beta| = |\alpha|$. Since $|\beta| \leq \beta < \alpha$, this tells us $|\alpha| < \alpha$, so α is not a cardinal.

We can see any finite ordinal is a cardinal by induction: 0 is clearly a cardinal, and if $n \in \mathbb{N}$ is a cardinal, then any surjection from some $m \leq n$ to $n+1$ is also a surjection from a subset of m missing one element to n. This subset has cardinality $< m$ since m is a cardinal and the subset injects into m, but then the inductive hypothesis tells us that we cannot have a surjection from any ordinal $< n$ onto n.

Let $(\kappa_i)_{i\in I}$ be a sequence of cardinals indexed by some set I, and let $\lambda = \bigcup_{i\in I}\kappa_i$. We will show λ is a cardinal as well. λ is clearly an ordinal since it is the union of ordinals, so suppose there is a surjection from some ordinal $\alpha < \lambda$ onto λ. Since α is in the union of all of the κ_i, it must be in at least one κ_i. The surjection from α to λ gives a surjection from α to κ_i: if f is the surjection, define g to be equal to f when $f(\beta) \in \kappa_i$ and 0 otherwise. Then κ_i cannot be a cardinal, since there is a surjection from an ordinal $\alpha \in \kappa_i$ onto κ_i. Since ω is the union of all ordinals in $\mathbb{N}$, each of which is a cardinal, ω is a cardinal. Note that if we have any successor ordinal $\alpha + 1$, where α is infinite, we can define a bijection between α and $\alpha + 1$: let $f(n) = n + 1$ when $n \in \omega$, let $f(\alpha) = 0$, and let $f(\beta) = \beta$ for any β which is infinite but not equal to α. Since $\alpha \geq \omega$, this will be a bijection, so $\alpha + 1$ cannot be a cardinal. Any infinite cardinal must therefore not be a successor and cannot be 0, so it is a limit ordinal. □

Theorem 2.26 (Cantor's Theorem)**.** *For any set X, $|X| < |\mathcal{P}(X)|$.*

Proof. Consider any function $f \colon X \longrightarrow \mathcal{P}(X)$: we will show it is not surjective. Define $A \subseteq X$ to be $\{x \in X \mid x \notin f(x)\}$. If $f(a) = A$ for some $a \in X$, then we have $a \in f(a)$ iff $a \in A$ iff $a \notin f(a)$, which is a contradiction. Therefore, f is not surjective. □

Corollary 2.27 (Arbitrarily Large Cardinals)**.** *There is no set of all cardinals.*

Proof. If A is a set of cardinals, then $\bigcup A$ is a cardinal which is at least as large as each element of A. Taking the cardinality of $\mathcal{P}(\bigcup A)$ gives a cardinal strictly larger than the cardinality of $\bigcup A$, which is strictly larger than each element of A. Therefore, A cannot contain this cardinal. □

Definition 2.28 (Aleph Numbers)**.** We define the cardinals $\aleph_\alpha$ for all ordinals α by transfinite recursion: $\aleph_\alpha$ is the smallest infinite cardinal strictly larger than each element of $\{\aleph_\beta \mid \beta < \alpha\}$. $\aleph_0$ is clearly ω, as ω is the smallest infinite cardinal. It is also clear that every infinite cardinal is an $\aleph_\alpha$ for some unique ordinal α: this can be proven by induction. If $\kappa = \aleph_\alpha$ where α is a a successor ordinal κ is called a *successor cardinal*; similarly, if α is a limit ordinal then κ is called a *limit cardinal.* The *successor* of a cardinal κ, denoted κ^+, is the smallest cardinal larger than κ: if κ is infinite, $\kappa = \aleph_\alpha$ and $\kappa^+ = \aleph_{\alpha+1}$.

Now that we've identified the infinite cardinals, we will define some arithmetic operations on them. These will be an extension of the notions of addition, multiplication, and exponentiation for natural numbers.

Definition 2.29 (Cardinal Arithmetic)**.** Let $(\kappa_i)_{i\in I}$ be an indexed family of cardinals. The *cardinal sum* $\sum_{i\in I}\kappa_i$ is defined as $|\bigsqcup_{i\in I}\kappa_i|$. Similarly, the *cardinal product* $\prod_{i\in I}\kappa_i$ is defined as the cardinality of the cartesian product $|\prod_{i\in I}\kappa_i|$. If κ and λ are cardinals, then κ^λ is the cardinality of the set of all functions from λ to κ. In particular, since a subset of κ can be encoded as a function from κ to 2 and vice versa, $2^\kappa = |\mathcal{P}(\kappa)|$. The *beth numbers* $\beth_\alpha$ are defined by transfinite recursion as the smallest infinite cardinal greater than or equal to each element of $\left\{2^{\beth_\beta} \mid \beta < \alpha\right\}$. It is easy to see that $\beth_0 = \aleph_0$ and $\beth_\alpha \geq \aleph_\alpha$ for every $\alpha \in \mathrm{Ord}$.

Lemma 2.30 (Properties of Cardinal Arithmetic)**.** *Sums, products, and exponents of cardinals are weakly monotonic: if we replace each cardinal with a cardinal greater than or equal to the original, we obtain a result greater than or equal to the original* *result. Cardinal arithmetic agrees with the arithmetic on natural numbers. The distributivity laws hold: for any cardinal λ and family of cardinals $(\kappa_i)_{i\in I}$, we have the following:*

$$\lambda \cdot \sum_{i\in I} \kappa_i = \sum_{i\in I} (\lambda \cdot \kappa_i)$$

$$\left(\prod_{i\in I} \kappa_i\right)^{\lambda} = \prod_{i\in I} \kappa_i^{\lambda}$$

$$\lambda^{\sum_{i\in I} \kappa_i} = \prod_{i\in I} \lambda_i^{\kappa}$$

$$\lambda^{\kappa_i \cdot \kappa_j} = (\lambda^{\kappa_i})^{\kappa_j}$$

Proof. These properties follow immediately by constructing a bijection between the relevant sets. For instance, a function from a disjoint union of sets to another set can be encoded as a sequence of functions from each set in the union to the other set, and vice versa. Once these properties are verified, the fact that they agree with regular arithmetic on natural numbers is proved by induction. □

Lemma 2.31 (Squares of Infinite Cardinals)**.** *For any infinite cardinal κ, we have $\kappa \cdot \kappa = \kappa$*

Proof. See lemma 68.10 in [7] for a proof. □

Corollary 2.32 (Sums and Products of Infinite Cardinals)**.** *If κ and λ are cardinals, at least one of which is infinite, then $\kappa + \lambda = \max\{\kappa, \lambda\}$. If we also have that neither κ nor λ are zero, then $\kappa \cdot \lambda = \max\{\kappa, \lambda\}$*

Proof. If at least one of κ and λ are infinite, then so is $\max\{\kappa, \lambda\}$. If neither are zero, then by monotonicity we have $\kappa = \kappa \cdot 1 \leq \kappa \cdot \lambda$, and $\lambda = 1 \cdot \lambda \leq \kappa \cdot \lambda$, so $\max\{\kappa, \lambda\} \leq \kappa \cdot \lambda$. Applying monotonicity again gives $\kappa \cdot \lambda \leq \max\{\kappa, \lambda\} \cdot \max\{\kappa, \lambda\} = \max\{\kappa, \lambda\}$, so $\kappa \cdot \lambda = \max\{\kappa, \lambda\}$. For the sum, we have that $\kappa + \lambda \geq \max\{\kappa, \gamma\}$ and $\kappa + \lambda \leq \max\{\kappa, \lambda\} \cdot 2 = \max\{\kappa, \lambda\}$. □

We'll now define the main result which we will prove is independent of ZFC:

Definition 2.33 (Continuum Hypothesis)**.** The *continuum hypothesis* is the sentence $\aleph_1 = \beth_1$, or equivalently $\aleph_1 = 2^{\aleph_0}$. $\beth_1$ is often called the *cardinality of the continuum*, since it is the cardinality of the real numbers.

2.4. **Relative Consistency, Absoluteness, and Transitive Models.** We now have a solid understanding of the ZFC axioms from the point of view of mathematical logic. We've seen that the axioms have a great deal of expressive power, and can describe essentially all of the everyday mathematical concepts we are familiar with, which are often stated in an intuitive set theory. Importantly, we can discuss first-order logic and model theory within ZFC. First order structures are sets with functions and relations, formulas can be encoded as natural numbers, and variable assignments are just functions. If $\mathcal{M}$ is some first-order structure with a variable assignment s and φ is a formula in the corresponding language, we write $\mathcal{M}, s \models \varphi$ to say that φ is true in $\mathcal{M}$ under variable assignment s. A first-order language can be encoded into the natural numbers, since sentences are finitary objects. We can even encode the notion of a mathematical proof in ZFC, since a proof is a finitary object consisting of a finite number of formulas put together in a particular way. This allows us to talk about ZFC within ZFC. For instance, there is a sentence Con(ZFC), which states that no natural number is a proof of a contradiction from the axioms of ZFC. Gödel famously showed that given a sufficiently expressive consistent theory T, there will be no proof of $\mathrm{Con}(T)$ from T. This means that, assuming ZFC is consistent, it cannot prove its own consistency. See part VII of [7] for a full proof of the incompleteness theorems. The best we can hope for is a *relative consistency proof*, which is a proof for a sentence of the form "if T is consistent, so is T together with some additional axiom A." ZFC is particularly well-suited for these types of proofs because it proves that first-order logic is sound and complete: a set of sentences, interpreted as a set of natural numbers, is consistent iff the set has a model, in the sense of a set with the necessary interpretations of the relation and function symbols. Since proofs in first-order logic can only mention a finite number of axioms, we obtain the *compactness theorem*: a first order theory has a model iff every finite subset has a model. See part III of [7] for a proof of the soundness and completeness theorems, and chapter 5 of [2] for a proof of the compactness theorem. To prove relative consistency results in ZFC, we therefore just have to show that a model of the theory we are interested in can be transformed into a model where an additional hypothesis holds true. In our case, forcing will take a model of ZFC and extend in a particular way. The sentences that are true in this new model will depend both on what is true in the original model and the method we use to extend it. We therefore will want to capture the relationships between the sentences which are true in the two models.

Definition 2.34 (Absoluteness)**.** Let $\mathcal{L}$ be a first order language, and let $\mathcal{M}$ and $\mathcal{N}$ be $\mathcal{L}$-structures, with $\mathcal{N}$ being a $\mathcal{L}$-substructure of $\mathcal{M}$. If φ is a formula, then φ is said to be *absolute* between $\mathcal{M}$ and $\mathcal{N}$ if for every variable assignment s in $\mathcal{N}$ we have $\mathcal{N}, s \models \varphi$ iff $\mathcal{M}, s \models \varphi$. Often we will work with defined relation and function symbols: we will say such a relation or function symbol is absolute between $\mathcal{M}$ and $\mathcal{N}$ if the definition sentence is absolute between the structures.

When we work with models of set theory, we will often work in models in which the interpretation of the element relation is the actual element relation $\in$. In other words, we will consider the structure where the symbol $\in$ on a set M is interpreted as the binary relation $M_{\in} = \{(x, y) \mid x, y \in M, x \in y\}$. This gives us a way we can naturally translate a formula φ into a formula about the structure $(M, \in)$, where M is some set:

Definition 2.35 (Relativization)**.** Let φ be a formula in the language of set theory, not extended by any definitions. The formula φ^x, where x is a variable not occuring in φ, is the formula obtained by replacing all the quantifiers in φ with quantifiers bound to the new free variable x. If M is some set, the formula φ^M is called the *relativization* of φ to M: it represents the sentence $(M, \in) \models \varphi$. We can also relativize defined terms by relativizing their definition sentences, and this will always produce a definition sentence in $(M, \in)$. This also gives us a way to talk about absoluteness between models: if $N \subseteq M$ a term or formula is absolute between $(M, \in)$ and $(N, \in)$ iff its relativizations are equivalent for formulas and equal for terms for all possible variable assignments in N. We can also say that a term τ or formula φ is absolute between some set M and the entire universe if $\tau^M = \tau$ or $\varphi^M \iff \varphi$ for all variable assignments in M.

Definition 2.36 (Transitive Model)**.** A *transitive model* of set theory is a transitive set $\mathbf{M}$ such that $(\mathbf{M}, \in)$ is a model of ZFC.

Transitive models of set theory are particularly nice to work with, because a wide class of sentences are absolute between transitive models and their transitive submodels. These are the Δ_0 formulas:

Definition 2.37 (Δ_0 Formula)**.** Let φ be a formula in the language of set theory. φ is said to be a Δ_0 *formula* if there is some formula ψ where all quantifiers are bounded and ZFC proves $\varphi \iff \psi$.

Theorem 2.38 (Δ_0 Absoluteness)**.** *If* $\mathbf{M}$ *is a transitive model of set theory and* φ *is a* Δ_0 *formula then* φ *is absolute between* $\mathbf{M}$ *and* V*. Similarly, if* $\mathbf{N}$ *is a transitive submodel of* $\mathbf{M}$*,* φ *is absolute between* $\mathbf{N}$ *and* $\mathbf{M}$*.*

Proof. Since φ is Δ_0 and equivalent over ZFC to a formula with only bounded quantifiers, we can assume φ only has bounded quantifiers. We'll first show $\varphi \iff \varphi^{\mathbf{M}}$ whenever we assign the free variables of φ to elements of $\mathbf{M}$. We proceed by induction on the complexity of φ. By the definition of relativization, the atomic formulas $x = y$ and $x \in y$ will be absolute, as will the conjunction, disjunction, negation, and conditional formed from formulas which are absolute. Now for the case of quantifiers, we know that since φ is Δ_0 they will be bounded quantifiers. This means we will have a formula of the form $\forall x \in y.\psi$ or $\exists x \in y.\psi$. We only will have to show the absoluteness for one of the quantifiers since we can just use negation to obtain the other, so we'll show it for $\exists x \in y.\psi$:

$$\begin{aligned}
(\exists x \in y.\psi)^{\mathbf{M}} &\equiv (\exists x.(x \in y \wedge \psi))^{\mathbf{M}} \\
&\equiv \exists x \in \mathbf{M}.(x \in y \wedge \psi^{\mathbf{M}}) \\
&\iff \exists x \in \mathbf{M}.(x \in y \wedge \psi) \\
&\iff \exists x \in y.(x \in \mathbf{M} \wedge \psi)
\end{aligned}$$

We can then drop the $x \in \mathbf{M}$ part due to transitivity: if $x \in y$ then since $y \in \mathbf{M}$ we have $x \in \mathbf{M}$ because $\mathbf{M}$ is transitive. The same reasoning holds when we want to show that φ is absolute between $\mathbf{N}$ and $\mathbf{M}$. $\square$

Corollary 2.39 (Absolute Notions for Transitive Models)**.** *In any transitive set, Extensionality holds. If* $\mathbf{M}$ *is a transitive model of ZFC, then the following notions are absolute between it and the entire universe or any transitive supermodel:*

- $x \subseteq y$.
- $\emptyset$, $\{x_1, \dots x_n\}$.
- $\bigcup A$, $\bigcap A$
- $\{\tau(x, w_1, \dots w_n) \mid x \in y \wedge \varphi(x)\}$, *where* τ *and* φ *are absolute notions between* $\mathbf{M}$ *and the entire universe or any transitive supermodel.*
- $A \setminus B$.
- (x, y), $A \times B$.
- x *is transitive.*
- x *is an ordinal,* x *is a successor/limit ordinal.*
- ω.
- x *is a finite/infinite ordinal.*
- f *is a function,* $\mathrm{dom}(f)$, $\mathrm{ran}(f)$, f *is injective,* f *is surjective.*

- *Any definition by transfinite recursion, so long as the term used in the recursion is absolute.*
- *x is a finite/infinite set.*

Proof. Extensionality always holds in a transitive set, because all of the elements of an element of a transitive set $\mathbf{M}$ will also be elements of $\mathbf{M}$. Therefore, the interpretation of $\subseteq$ in $\mathbf{M}$ will be the real $\subseteq$ relation, so Extensionality holds. All of the relations and terms are absolute because they are Δ_0 except for the last one, although some are a bit difficult to verify. See [4] for the details. The one statement which requires a bit more work is being a finite set or an infinite set. Suppose x is finite in $\mathbf{M}$. Then there is a bijection from x to some finite ordinal in $\mathbf{M}$. This will remain true in any transitive supermodel or even in the full universe, so x will be finite there as well. Now if x is infinite in $\mathbf{M}$, then there is an injection $f\colon \omega \longrightarrow x$ in $\mathbf{M}$, which will remain an injection in the larger universe. Thus, x is finite/infinite in $\mathbf{M}$ iff it is finite/infinite in the larger universe. □

Note that several items are missing from our list. Most importantly, it is not true that the set of all functions from a set A to a set B, the cardinality of a set, and the property of being a cardinal are absolute. All of these can be defined in terms of the power set operation, which is also not absolute between transitive models. The reason is that, although all of the elements of a set in a transitive model are required to be in the model, not all of its subsets are. In fact, as we will show in the next section, if there are transitive models of ZFC then there are countable transitive models. These models will contain the natural numbers, but not all of the subsets of natural numbers. Although this is a bit strange, it is quite helpful for our goal of proving that the continuum hypothesis is independent from ZFC, since it means that the size of the power set is not necessarily fixed by ZFC. We'll now turn to the question of how we can build new models of ZFC.

3. Forcing

In this section, we will develop the notion of *forcing*, which is the primary way we will prove independence results. The general idea of the method is to start with a model of ZFC, which we'll call $\mathbf{M}$. We will assume $\mathbf{M}$ is transitive so that we can leverage the facts of the $\in$ relation. We can modify $\mathbf{M}$ slightly to obtain a countable transitive set which is still a model of ZFC. This is quite surprising, but it has the interesting side effect that all infinite sets in our model are externally countable (although internally the model will "think" that some of these sets are uncountable). Therefore, there will exist a bijection in the full set theoretic universe between any two infinite sets of our model. The heart of forcing is to find a way to extend this model that adds in a new desired object, such as a bijection or additional subsets of some set, without disturbing the structure of the model. We'll first show how to construct our countable transitive model.

3.1. **The Countable Transitive Model M.** Before we begin with the forcing argument, we need to briefly discuss the metamathematics of forcing. All of our arguments will hinge on the existence of a countable transitive model of ZFC, which we will show can be constructed from a transitive model of ZFC. However, the consistency of ZFC only implies the existence of a model of ZFC, and this is not sufficient to construct a transitive model. In order for our forcing argument to work, we'll use countable transitive models of *large finite fragments of ZFC* rather than all of ZFC itself. As a consequence of the *reflection principle*, which we will prove in this section, every finite fragment of ZFC has a transitive model, namely some V_α. We will show how to reduce this to a countable transitive model. Now suppose we want to show that some statement φ is relatively consistent with ZFC. What forcing will allow us to do is, for any finite fragment T of $\mathrm{ZFC}+\varphi$, produce another finite fragment $T' \subseteq \mathrm{ZFC}$ and a procedure for turning a countable transitive model of T' into a countable transitive model of T. Now suppose ZFC entailed $\neg\varphi$. The theory $\mathrm{ZFC}+\varphi$ would therefore be inconsistent, and the compactness theorem of first order logic would tell us that some finite fragment $T \subseteq \mathrm{ZFC}+\varphi$. But then $T' \subseteq \mathrm{ZFC}$ would be inconsistent, since we know how to turn models of T' into models of T, and there are no models of T. Therefore, ZFC would prove its own inconsistency. Therefore, ZFC entails the relative consistency result if we can find such a translation of finite fragments. This leads us to adopt the following convention:

Convention 3.1 (Transitive Models of Set Theory)**.** We will use the terms "transitive model of set theory" and "countable transitive model of set theory" to refer to a (countable) transitive model of an arbitrarily large finite fragment of ZFC.

Now that we have the metamathematical argument out of the way, we can begin examining how to obtain a countable transitive model from a transitive model. We'll begin with the following definition and associated results from logic:

Definition 3.1 (Elementary Substructure)**.** Let $\mathcal{L}$ be a first order language, let $\mathcal{M}$ be an $\mathcal{L}$-structure, and let $\mathcal{N}$ be an $\mathcal{L}$-substructure of $\mathcal{M}$. $\mathcal{N}$ is called an *elementary substructure* of $\mathcal{M}$, denoted $\mathcal{N} \preceq \mathcal{M}$, if every formula φ in the language $\mathcal{L}$ is absolute between $\mathcal{N}$ and $\mathcal{M}$ for any variable assignment in $\mathcal{N}$. That is, for any variable assignment s for $\mathcal{N}$, we have $\mathcal{N}, s \models \varphi$ iff $\mathcal{M}, s \models \varphi$.

Lemma 3.2 (Tarski-Vaught Test)**.** *Let $\mathcal{L}$ be a first order language, let $\mathcal{M}$ be an $\mathcal{L}$-structure, and let $\mathcal{N}$ be an $\mathcal{L}$-substructure of $\mathcal{M}$. Suppose $S \subseteq \mathcal{L}$ is a set of formulas that is closed under taking subformulas: such a set is called* subformula closed*. Then every formula in S is absolute between $\mathcal{M}$ and $\mathcal{N}$ iff S satsifies the* Tarski-Vaught criterion*: for every formula in S of the form $\exists x.\varphi(x, y_1, \dots y_n)$ and every n-tuple of elements $b_1, \dots b_n \in \mathrm{dom}(\mathcal{N})$, if $\mathcal{M} \models \exists x.\varphi(x, b_1, \dots b_n)$ then there exists some $a \in \mathrm{dom}(\mathcal{N})$ such that $\mathcal{M} \models \varphi(a, b_1, \dots b_n)$. In particular, if $S = \mathcal{L}$, then $\mathcal{N} \preceq \mathcal{M}$ iff the Tarski-Vaught criterion holds for all formulas in $\mathcal{L}$.*

Proof. Induction on the structure of formulas in S. Most of the cases are straightforward, but the existential case requires the Tarski-Vaught criterion. See [2] for a full proof. □

Theorem 3.3 (Löwenheim-Skolem Theorem)**.** *Let $\mathcal{L}$ be a first order language with a countable signature, and let $\mathcal{M}$ be an $\mathcal{L}$-structure. Then for any countable set $A \subseteq \mathcal{M}$ there exists an elementary $\mathcal{L}$-substructure $\mathcal{N} \preceq \mathcal{M}$ such that $A \subseteq \mathrm{dom}(\mathcal{N})$ and $\mathrm{dom}(\mathcal{N})$ is countable.*

Proof. Fix a well-ordering of the domain of $\mathcal{M}$. We will construct a sequence of expanded languages $\mathcal{L}_0 \subseteq \mathcal{L}_1 \subseteq \dots$ such that $\mathcal{M}$ is an $\mathcal{L}_k$ structure for every $k \in \mathbb{N}$ using this well ordering. Set $\mathcal{L}_0 = \mathcal{L}$. Now suppose we've defined $\mathcal{L}_k$. To define $\mathcal{L}_{k+1}$, consider any formula $\varphi(x, y_1, \dots y_n) \in \mathcal{L}_k$ with $n+1$ free variables. Define a new n-ary function symbol f_φ and interpret it in $\mathcal{M}$ to be the least element $a \in \mathrm{dom}(\mathcal{M})$ in the fixed well-ordering such that if $\mathcal{M} \models \exists\varphi(x, b_1, \dots b_n)$ for $b_1, \dots b_n \in \mathrm{dom}(\mathcal{M})$ then $\mathcal{M} \models \varphi(f_\varphi(b_1, \dots b_n), b_1, \dots b_n)$, and the least element of $\mathrm{dom}(\mathcal{M})$ otherwise. Clearly, we have the following:

$$\mathcal{M} \models \forall y_1, \dots y_n \, (\exists x.\varphi(x, y_1, \dots y_n) \implies \varphi(f_\varphi(y_1, \dots y_n), y_1, \dots y_n))$$

This interpretation of f_φ is called an *Skolem function* for φ. Define the language $\mathcal{L}_{k+1}$ to be the first order language with a signature consisting of all relation and function symbols in $\mathcal{L}_k$ together with all function symbols f_φ for $\varphi \in \mathcal{L}_k$. The interpretation of $\mathcal{L}_{k+1}$ for $\mathcal{M}$ is defined to be the interpretation of $\mathcal{L}_k$ for $\mathcal{M}$ whenever we have a function or relation symbol in $\mathcal{L}_k$ and the Skolem function for φ whenever we have a function symbol f_φ for $\varphi \in \mathcal{L}_k$. This gives us a sequence of languages $\mathcal{L}_k$ which are increasing and their corresponding interpretations in $\mathcal{M}$. The interpretations clearly satisfy the property that the interpretation of a function or relation symbol in the language $\mathcal{L}_k$ is not modified for $\mathcal{L}_n$ for $n \geq k$. Therefore, we can define $\mathcal{L}_\omega$ to be a language whose signature consists of the union of the signatures of all of the $\mathcal{L}_k$, and define a corresponding interpretation of $\mathcal{L}_\omega$ for $\mathcal{M}$ by taking the union of the interpretations. Now our original language $\mathcal{L} = \mathcal{L}_0$ had a countable signature, and therefore had countably many formulas. Since every formula is assigned a Skolem function in the next language, if $\mathcal{L}_k$ has a countable signature so does $\mathcal{L}_{k+1}$. By induction, every language has a countable signature, and since the countable union of countable sets is countable, $\mathcal{L}_\omega$ has a countable signature as well. Consider any formula $\varphi \in \mathcal{L}_\omega$: since it is finite in length, it only mentions symbols from finitely many of the $\mathcal{L}_n$. Taking the maximum of these n shows us that it is in $\mathcal{L}_k$ for some k. Therefore, the Skolem function for the formula has a symbol in $\mathcal{L}_{k+1}$ and therefore in $\mathcal{L}_\omega$. Thus, our language $\mathcal{L}_\omega$ has a Skolem function for each of its formulas. We can enumerate the interpretations of the function symbols $(f_n)_{n=0}^\infty \colon \mathrm{dom}(\mathcal{M}) \longrightarrow \mathrm{dom}(\mathcal{M})$. Now take our countable set A and consider its closure under the functions $(f_n)_{n=0}^\infty$. We can see that this set is countable by setting $A_0 = A$, repeatedly taking the union of A_n with the images of A_n under every f_k to give A_{n+1}, and letting A_ω be the union of all of the A_n. This is a countable union of countable sets and images of countable sets, and is therefore countable. This set will be our $\mathrm{dom}(\mathcal{N})$. Give it an interpretation in $\mathcal{L}_\omega$ by restricting every relation symbol and function symbol to it. This will be an $\mathcal{L}_\omega$ substructure of $\mathcal{M}$, since $\mathrm{dom}(\mathcal{N})$ is closed under each f_n. Because our language $\mathcal{L}_\omega$ has a Skolem function for every formula, the Tarski-Vaught criterion holds, and $\mathrm{dom}(\mathcal{N})$ is an $\mathcal{L}_\omega$ elementary substructure of $\mathcal{M}$. Taking the $\mathcal{L}$-reduct of $\mathcal{N}$ gives $\mathcal{N} \preceq \mathcal{M}$. □

Corollary 3.4 (Extracting a Countable Transitive Model)**.** *Suppose* **M** *is a transitive model of set theory. Then it contains a countable transitive elementary submodel* **N***.*

Proof. Apply Löwenheim-Skolem to **M** and take $\emptyset$ to be the subset we are extending. This will give us a countable model whose interpretation of the element relation is $\in$, and is therefore a well-founded model. Applying Mostowski collapse gives us an isomorphic countable transitive model **N**. Since **N** is isomorphic to an elementary substructure of **M**, it will satisfy whatever finite fragment of ZFC is satisfied in **M**. □

The fact that there are countable models of ZFC at all is a bit strange, so we should take a moment to consider what this means. In a countable transitive model $\mathbf{M}$, there is a set $\mathcal{P}(\mathbb{N})^{\mathbf{M}} \in \mathbf{M}$ such that $\mathbf{M}$ "believes" that it is the power set of the natural numbers. However, since $\mathbf{M}$ is countable and transitive, this set can only contain countably many elements. Despite this, $\mathbf{M}$ is still a model of ZFC, so it also believes that there are strictly more sets of natural numbers than natural numbers. This is possible because what it means for $\mathcal{P}(\mathbb{N})^{\mathbf{M}}$ to have a greater cardinality than $\mathbb{N}$ in the model $\mathbf{M}$ is that there is no surjection from $\mathbb{N}$ to $\mathcal{P}(\mathbb{N})^{\mathbf{M}}$ *in the model* $\mathbf{M}$. There is certainly a bijection in the external universe V, since both sets have cardinality $\aleph_0$, but no such bijection can exist in the model $\mathbf{M}$. Thus, cardinality is not absolute between transitive models.

This suggests a method by which we could prove the independence of the continuum hypothesis. Start with a countable transitive model $\mathbf{M}$. If we can somehow add a bijection between $\mathcal{P}(\mathbb{N})^{\mathbf{M}}$ and $\aleph_1^{\mathbf{M}}$ to our model while preserving the fact that we have a model of set theory and that the power set of $\mathbb{N}$ and $\aleph_1$ do not change in our new model, then our new model will satisfy the continuum hypothesis. If we instead add in $\aleph_2^{\mathbf{M}}$ new subsets of natural numbers, which we can do since this is a countable number and there are only countably many subsets in our ground model, and we don't change the cardinality of $\aleph_2$, we will obtain a model of ZFC where there are at least $\aleph_2$ sets of natural numbers, so the continuum hypothesis will fail. Forcing will provide us a method for performing these extensions.

Before we proceed, we need to show that there are transitive models of ZFC, in the sense that any finite fragment of ZFC has a transitive model. We will even show that we can make such models contain an arbitrary set. Although this will not be useful for the purposes of extracting a countable transitive model, it will be useful later when proving results about the Replacement axiom holding in our forcing extension.

Theorem 3.5 (Reflection Principle)**.** *Let $\varphi(x_1, \ldots x_n)$ be a formula in the language of set theory. Then ZFC proves the following statement: for any ordinal α, there exists an ordinal $\beta \geq \alpha$ such that $\forall a_1, \ldots a_n \in V_\beta \left(\varphi^{V_\beta}(a_1, \ldots a_n) \iff \varphi(a_1, \ldots a_n)\right)$*

Proof. The proof idea is similar to the proof of Löwenheim-Skolem: close V_α under Skolem-like functions and use Tarski-Vaught to obtain the absoluteness. First, let S be the smallest set containing all of the subformulas of φ: such a set will be subformula closed. Since φ contains finitely many subformulas, S will be finite. Let $\varphi_1, \ldots \varphi_k$ be the formulas, with m free variables $x_1, \ldots x_m$. If we can find an ordinal $\beta \geq \alpha$ such that for any k and any free variable x_j we have $\forall a_1, \ldots a_{j-1}, a_{j+1}, \ldots a_m \in V_\beta. (\exists a_j.\varphi_k(\bar{a}) \implies \exists a_j \in V_\beta.\varphi_k(\bar{a}))$ then the Tarski-Vaught criterion will hold for V_β. Since the language of set theory has no function symbols, $(V_\beta, \in)$ is naturally a substructure, and therefore every formula in S, including φ, is absolute between V_β and the entire universe by the Tarski-Vaught test. To find such an ordinal, note that we can easily define a term $\tau(a_1, \ldots a_n)$ as the smallest ordinal such that that for every k, we have $\exists a_j.\varphi_k(\bar{a}) \implies \exists a_j \in V_\beta.\varphi_k(\bar{a})$. Because there are only finitely many formulas, this last condition is really a finite conjunction: otherwise we couldn't talk about each of the implications holding. Now set $\beta_0 = \alpha$, and set β_{n+1} to be the maximum of β_n and the image of V_{β_n} under τ. This image will exist by Replacement. Now, let β be the supremum of all of the β_n. This will be a limit ordinal, and since it is the supremum of all β_n, V_β will be the union of all of the V_{β_n}. Any assignment of free variables in V_β will really be an assignment in some V_{β_n}, and therefore if some existential statement holds under such an assignment there will be a witness at $V_{\beta_{n+1}} \subseteq V_\beta$. Thus, V_β will satisfy the Tarski-Vaught criterion. □

3.2. **Forcing Posets and Generic Filters.** Now that we have a countable transitive model $\mathbf{M}$ of ZFC, we need some method to extend $\mathbf{M}$ in a controlled manner so that it satisfies our desired logical properties. However, it is clear that we cannot simply add in any set we like: we must ensure that the resulting structure remains a model of ZFC. For instance, if we are to add a set G into $\mathbf{M}$, we must ensure that $\mathcal{P}(G)$ is also added, that all definable subseets of G are added, and so on. Furthermore, it is not clear that even if we succeed in producing an extended model that we can control any of its desired properties. For instance, suppose we want to construct a model where the continuum hypothesis holds. If we have a countable transitive model $\mathbf{M}$, we know that $\mathcal{P}(\omega)^{\mathbf{M}}$ and $\aleph_1^{\mathbf{M}}$ are both countable infinite sets. Even if we somehow add in a bijection between $\mathcal{P}(\omega)^{\mathbf{M}}$ and $\aleph_1^{\mathbf{M}}$ while maintaining the fact that our extended structure is still a model of ZFC, it might be the case that $\aleph_1$ in our new model is different from $\aleph_1$ in our original model, or that the power set of ω has changed. In either case, the continuum hypothesis may not hold.

Forcing takes a clever approach to avoid these problems. Instead of adding in an arbitrary new set G directly, we only add in a set which can in some sense be *approximated* from within $\mathbf{M}$. In other words, $\mathbf{M}$ contains a sketch of how to construct G consisting of all finite pieces of information we can have about G. It will also contain certain *requirements* that G must satisfy. For instance, if we wish to adjoin a bijection between two infinite sets $A, B \in \mathbf{M}$ in our model, we can approximate such a bijection by examining the set of all bijections between finite subsets of A and B in our model. In this case, one requirement that must be satisfied is that the resulting function is defined on all of A, and not just some subset of A. By ensuring that our set G obeys these requirements, we will be able to determine whether sentences in $\mathbf{M}[G]$ are true based on properties of the finite approximations within $\mathbf{M}$. This will give us both a great deal of generality in what we can adjoin to our model $\mathbf{M}$ while maintaining control over what is true in the extension.

We can make this notion more precise. Suppose we have such a set of finite approximations. Then we can order these approximations based whether one approximation gives us more information than another approximation. This gives rise to the following definition:

Definition 3.6 (Partially Ordered Set)**.** A *partially ordered set* or simply a *poset* is a pair $\mathbb{P} = (P, \preceq)$ where P is a set and $\preceq$ is a binary relation on P satisfying the following requirements:

(1) Reflexivity: For all $p \in P$, we have $p \preceq p$.
(2) Antisymmetry: For all $p, q \in P$, if $q \preceq p$ and $p \preceq q$ then $p = q$.
(3) Transitivity: For all $p, q, r \in P$, if $q \preceq p$ and $r \preceq q$ then $r \preceq p$.

We will often abuse notation and use $\mathbb{P}$ to refer to the underlying set P.

In forcing, the convention is that when we write $q \preceq p$ we mean that q gives more information than p. One can think of this as saying that the collection of sets G which satisfy q is a subset of those that satisfy p. In addition to our finite approximations forming a poset, we will want two additional requirements. First, there should be an *empty approximation* which gives us no information. Furthermore, we want our notion of forcing to add something new to our model: no single finite approximation should completely determine G. This gives rise to the following definition:

Definition 3.7 (Incompatible Elements)**.** Let $\mathbb{P}$ be a poset, and let $p, q \in \mathbb{P}$ be elements. p and q are said to be *incompatible*, denoted $p \perp q$, if there does not exist an $r \in \mathbb{P}$ such that $r \preceq p$ and $r \preceq q$.

Definition 3.8 (Notion of Forcing)**.** Let $\mathbb{P}$ be a poset. $\mathbb{P}$ is said to be a *notion of forcing* or a *forcing poset* if $\mathbb{P}$ satisfies the following requirements:

(1) $\mathbb{P}$ is *rooted*: There is a *least element.* That is, there is some $1_{\mathbb{P}} \in \mathbb{P}$ such that every $p \in \mathbb{P}$ satisfies $p \preceq 1_{\mathbb{P}}$
(2) $\mathbb{P}$ is *atomless*: For all $p \in \mathbb{P}$, there exist $q, r \in \mathbb{P}$ such that $q, r \preceq p$ and $q \perp r$.

The elements of a forcing poset are called *forcing conditions* or simply *conditions*, and the relation $q \preceq p$ is read "*q extends p*."

Now that we know the structure of our forcing conditions, we need to specify how they approximate our new set G. We will consider G to be the set of all conditions that approximate it. These conditions will need to satsify two requirements. First, if a condition is in G, then any condition weaker will also be in G. Second, any two conditions need to *cohere* together in some sense. This will ensure that the set we are approximating is well defined. The way we'll define this is by saying any two conditions in G will need to have a common extension in G. This gives rise to the following definition:

Definition 3.9 (Filter)**.** Let $\mathbb{P}$ be a poset. A *filter* is a subset $G \subseteq \mathbb{P}$ satisfying the following two conditions:

(1) Upper set: For any $p \in G$ and any $q \succeq p$, we have $q \in G$.
(2) Downward directedness: For any $p, q \in G$, there is an $r \in G$ such that $r \preceq p, q$.

Now it may be the case that the G we wish to add needs to satisfy certain requirements. For instance, suppose we want to add a bijection from $\aleph_1$ and $\mathcal{P}(\mathbb{N})$ into our model. Our forcing poset would consist of bijections between proper subsets of $\aleph_1$ and proper subsets of $\mathcal{P}(\mathbb{N})$, with one bijection f extending a bijection g when the domain of g is contained in that of f and f restricted to the domain of g agrees with g. It's clear that taking the union of a filter of such bijections will be another bijection from a subset of $\aleph_1$ to a subset of $\mathcal{P}(\mathbb{N})$. However, we want a bijection between $\aleph_1$ and $\mathcal{P}(\mathbb{N})$, not just a subset of each. We therefore need the resulting function to be surjective, which will only be true if each $A \in \mathcal{P}(\mathbb{N})$ is in the range of some condition of our filter. This kind of requirement has a special property: any condition has an extension that satisfies it, since we can just pick an element not in the domain of the condition and map it to A. From this, we have the following definitions:

Definition 3.10 (Dense Set)**.** Let $\mathbb{P}$ be a poset. A subset $D \subseteq \mathbb{P}$ is called *dense* if for any $p \in \mathbb{P}$, there exists a $q \in D$ such that $q \preceq p$.

Definition 3.11 (Generic Filter)**.** Let $\mathbb{P}$ be a poset, and let $\mathcal{D}$ be a family of dense subsets of $\mathbb{P}$. A filter G on $\mathbb{P}$ is called $\mathcal{D}$*-generic* if every $D \in \mathcal{D}$ contains a condition in G. We will shorten this by saying G *meets* every $D \in \mathcal{D}$. Often the set $\mathcal{D}$ will be clear from context, in which case we will just say that G is *generic*.

The fundamental result about generic filters is the following:

Theorem 3.12 (Rasiowa-Sikorski)**.** *Let $\mathbb{P}$ be a poset, let $p \in \mathbb{P}$ be any element, and let $\mathcal{D}$ be a countable family of dense subsets of $\mathbb{P}$. Then there exists a $\mathcal{D}$-generic filter G such that $p \in G$.*

Proof. We will use the axiom of choice to define a sequence of elements in $\mathbb{P}$. Let $p_0 = p$, and fix an enumeration $(D_n)_{n=0}^{\infty} \in \mathcal{D}$ of the dense sets. Suppose we know p_n. Since D_n is dense, there exists some $q \preceq p_n$ which is in D_n. We'll let this be our p_{n+1}. This gives us a sequence $(p_n)_{n=0}^{\infty} \in \mathbb{P}$. Define G as $G = \{q \in \mathbb{P} \mid \exists n \in \mathbb{N}(p_n \preceq q)\}$. Clearly G contains each p_n, so it contains p. G meets every $D \in \mathcal{D}$ since it contains every p_n, and $p_n \in D_{n-1}$ for $n > 0$. We can now show that G is a filter. If $q \in G$, then $p_n \preceq q$ for some n. If $q \preceq r$, then by transitivity $p_n \preceq r$, so $r \in G$. Now suppose $q, r \in G$. Then for some $m, n \in \mathbb{N}$, we have $p_n \preceq q$ and $p_m \preceq r$. Without loss of generality, assume $n \geq m$. By the definition of our sequence, $p_n \preceq p_m$, and $p_n \in G$. Therefore, $p_n \preceq q, r$, so q and r have a common extension in G. Thus, G is a $\mathcal{D}$-generic filter containing p. □

Convention 3.2. When proving facts about generic filters, we will occasionally say that all *sufficiently generic filters* of a poset satisfy a given property. This means that there exists a countable family of dense sets such that any filter generic with respect to that family satisfies the given property. This implication cannot be vacuous since there is always at least one generic filter by 3.12.

Although we can only use Rasiowa-Sikorski to prove the existence of generic filters for countable families of dense sets, this is no obstacle for our forcing proofs. Since we are working with a countable model $\mathbf{M}$, and a forcing poset $\mathbb{P} \in \mathbf{M}$, only countably many subsets of $\mathbb{P}$ will be sets in $\mathbf{M}$ (although internally there might be uncountably many). The filters which we will consider will be generic with respect to all dense subsets of $\mathbb{P}$ which are in the model $\mathbf{M}$, and this will be sufficient. Here is one more fact about generic filters which will be useful:

Definition 3.13 (Density Below an Element)**.** Let $\mathbb{P}$ be a poset, and let $p \in \mathbb{P}$ be an element. A set $D \subseteq \mathbb{P}$ is said to me *dense below* p if for any $q \preceq p$ there exists some $r \in D$ such that $r \preceq q$.

Lemma 3.14. *Suppose $\mathbb{P}$ is a poset, $p \in P$ is an element, and $D \subseteq \mathbb{P}$ is dense below p. Then any sufficiently generic filter G on $\mathbb{P}$ meets D if it contains p.*

Proof. Given our set D, define a set D^* to be the union of D and all elements in $\mathbb{P}$ which are incompatible with p. Consider any G which is generic relative to this set, and pick a $q \in D^* \cap G$. Suppose $p \in G$. Then p and q cannot be incompatible, since they are both in G, and therefore have a common extension in G. Therefore, $q \in D \cap G$, so G meets D. □

3.3. **$\mathbb{P}$-Names and the Forcing Extension $\mathbf{M}[G]$.** Now that we've proved some facts about generic filters, we can use the filter to construct our expanded model $\mathbf{M}[G]$, which we'll call *the forcing extension of* $\mathbf{M}$ *by* G. The notation we use is reminiscint of the notation for adjoining an element to a ring in algebra, and this is no coincidence. We will eventually show that $\mathbf{M}[G]$ is the smallest model of ZFC which extends $\mathbf{M}$ and contains G. Additionally, the method we use to construct $\mathbf{M}[G]$ is quite similar to the method used to adjoin an element to a ring. To adjoin an element to a ring, we consider all possible polynomials over that ring, and quotient by a particular polynomial with special properties. The polynomials represent "recipes" for elements in our new ring given our new element. To create our forcing extension, we will take our countable transitive model $\mathbf{M}$ and consider sets in $\mathbf{M}$ whose elements are "tagged" with elements in our forcing poset $\mathbb{P}$: these will be analogues of our polynomials. However, these elements will themselves be sets with tagged elements, which themselves will be sets with tagged elements, and so on. We call these tagged sets $\mathbb{P}$-*names*. Our generic filter G will then "quotient" the $\mathbb{P}$-names in such a way that we end up recovering both the original model $\mathbf{M}$ and the generic filter G. First, we'll construct the $\mathbb{P}$-names:

Definition 3.15 ($\mathbb{P}$-names)**.** Let $\mathbb{P}$ be a forcing poset. For any ordinal α, we define the α-th $\mathbb{P}$-*name level* by recursion:

$$V_\alpha^{\mathbb{P}} = \bigcup_{\beta<\alpha} \mathcal{P}(V_\beta^{\mathbb{P}} \times \mathbb{P})$$

Elements of the class $V^{\mathbb{P}} = \bigcup_{\alpha\in\mathrm{Ord}} V_\alpha^{\mathbb{P}}$ are called $\mathbb{P}$-*names*. Every $\mathbb{P}$-name consists of ordered pairs (a, p), where a is a $\mathbb{P}$-name and $p \in \mathbb{P}$ is a condition. We call the domain of a $\mathbb{P}$-name its set of *potential elements*. The *name rank* of a $\mathbb{P}$-name is the smallest ordinal strictly larger than the name ranks of all of its potential elements. If we have some model $\mathbf{M}$ and a forcing poset $\mathbb{P} \in \mathbf{M}$ we define $\mathbf{M}_\alpha^{\mathbb{P}}$, $\mathbf{M}^{\mathbb{P}}$, and $\mathrm{rank}_{\mathbb{P}}^{\mathbf{M}}(A)$ analogously via relativization. It should be noted that the $\mathbb{P}$-name rank is absolute between transitive models: this is easy to see by induction on the name rank.

Convention 3.3. Suppose $\mathbf{M}$ is a countable transitive model, and $\mathbb{P} \in \mathbf{M}$ is a forcing poset. When we refer to generic filters over $\mathbb{P}$, we are always referring to generic filters with respect to $\mathcal{D} = \{D \in \mathbf{M} \mid D \subseteq \mathbb{P}, D \text{ dense}\}$. This is a countable collection of sets since $\mathbf{M}$ is countable, so by 3.12 there will always be a generic filter G.

In our analogy, the $\mathbb{P}$-names act like polynomials: they denote potential sets that will exist in our model $\mathbf{M}[G]$ depending on what conditions are in our generic filter G. Here is the definition of the forcing extension:

Definition 3.16 (Forcing Extension)**.** Let $\mathbf{M}$ be a countable transitive model of ZFC, let $\mathbb{P} \in \mathbf{M}$ be a forcing poset, and let G be a generic filter over $\mathbb{P}$. We define a map eval_G on the $\mathbb{P}$-names $\mathbf{M}^{\mathbb{P}}$ recursively as follows:

$$\mathrm{eval}_G(A) = \{\mathrm{eval}_G(a) \mid \exists p \in G.(a, p) \in A\}$$

The *forcing extension* $\mathbf{M}[G]$ is the image of the $\mathbb{P}$-names $\mathbf{M}^{\mathbb{P}}$ under the map eval_G:

$$\mathbf{M}[G] = \left\{\mathrm{eval}_G(A) \mid A \in \mathbf{M}^{\mathbb{P}}\right\}$$

We are now in a position to prove some basic facts about the forcing extension $\mathbf{M}[G]$:

Definition 3.17 (Canonical Name)**.** Let $\mathbb{P}$ be a forcing poset, and let A be any set. The *canonical name* of A, denoted $\hat{A}_{\mathbb{P}}$, is a $\mathbb{P}$-name recursively defined as follows:

$$\hat{A}_{\mathbb{P}} = \{(\hat{a}_{\mathbb{P}}, 1_{\mathbb{P}}) \mid a \in A\}$$

Theorem 3.18 (Properties of the Forcing Extension)**.** *Let* $\mathbf{M}$ *be a countable transitive model, let* $\mathbb{P} \in \mathbf{M}$ *be a forcing poset, and let* G *be a generic filter over* $\mathbb{P}$*. Then* $G \notin \mathbf{M}$*. We have* $\mathbf{M} \subseteq \mathbf{M}[G]$ *and* $G \in \mathbf{M}[G]$*.* $\mathbf{M}[G]$ *is countable, transitive, and contains the same ordinals as* $\mathbf{M}$

Proof. First, we'll show $G \notin \mathbf{M}$. This will follow directly from the fact that $\mathbb{P}$ is an atomless poset. Suppose $G \in \mathbf{M}$. Then the set $H = \mathbb{P} \setminus G$ is also in $\mathbf{M}$ since $\mathbf{M}$ is a model. We will show that H is dense. Pick some $p \in \mathbb{P}$. Since $\mathbb{P}$ is atomless, there must exist $q, r \preceq p$ such that $q \perp r$. Since G is a filter, if q and r are both in G, then they must have a common extension in G, but this is impossible since q and r are incompatible. Therefore, at least one of q or r is in H. Since every $p \in \mathbb{P}$ has an extension in H, H is dense. Since H is dense and in our model $\mathbf{M}$, G must meet H, but this is impossible since $H = \mathbb{P} \setminus G$. Therefore, $G \notin \mathbf{M}$.

Next, we'll show $\mathbf{M} \subseteq \mathbf{M}[G]$. To do this, we'll show that for any $A \in \mathbf{M}$, the canonical name $\hat{A}_{\mathbb{P}}$ evaluates to A under eval_G. We'll do so by induction on rank of A. Suppose that for all a with $\text{rank}\,(a) < \alpha$, we have $\text{eval}_G\,(\hat{a}_{\mathbb{P}}) = a$, and $\text{rank}\,(A) = \alpha$. Then all of the elements of A have rank less than α. Therefore, we have the following:

$$\begin{aligned}
\text{eval}_G\left(\hat{A}_{\mathbb{P}}\right) &= \left\{\text{eval}_G(a) \mid \exists p \in G.(a,p) \in \hat{A}_{\mathbb{P}}\right\} \\
&= \left\{\text{eval}_G\left(\hat{b}_{\mathbb{P}}\right) \mid \exists p \in G.\,(p = 1_{\mathbb{P}} \wedge b \in A)\right\} \\
&= \left\{\text{eval}_G\left(\hat{b}_{\mathbb{P}}\right) \mid b \in A\right\} \\
&= \{b \mid b \in A\} \\
&= A
\end{aligned}$$

Since each $A \in \mathbf{M}$ is equal to the evaluation of a $\mathbb{P}$-name of M, namely $\hat{A}_{\mathbb{P}}$, $\mathbf{M} \subseteq \mathbf{M}[G]$. We can also find a $\mathbb{P}$-name that evaluates to G. Define $\Gamma = \{(\hat{p}_{\mathbb{P}}, p) \mid p \in \mathbb{P}\}$. We'll show this evaluates to G:

$$\begin{aligned}
\text{eval}_G(\Gamma) &= \{\text{eval}_G\, a \mid \exists p \in G.(a,p) \in \Gamma\} \\
&= \{\text{eval}_G\, \hat{p}_{\mathbb{P}} \mid \exists p \in G.(\hat{p}_{\mathbb{P}}, p) \in \Gamma\} \\
&= \{p \mid \exists p \in G.(\hat{p}_{\mathbb{P}}, p) \in \Gamma\} \\
&= \{p \mid p \in G\} \\
&= G
\end{aligned}$$

Finally, we'll show that $\mathbf{M}[G]$ is countable, transitive, and has the same ordinals as $\mathbf{M}$. For countability, note that $\mathbf{M}[G]$ is the image of the set $\mathbf{M}^{\mathbb{P}}$ under the map eval_G. Since $\mathbf{M}^{\mathbb{P}} \subseteq \mathbf{M}$ and $\mathbf{M}$ is countable, so is $\mathbf{M}^{\mathbb{P}}$, and therefore so is $\mathbf{M}[G]$. For transitivity, pick some $a \in \mathbf{M}[G]$, and pick some $b \in a$. By the definition of $\mathbf{M}[G]$, there exists some name $A \in \mathbf{M}^{\mathbb{P}}$ such that $a = \text{eval}_G(A)$. We know $\text{eval}_G(A) = \{\text{eval}_G(B) \mid \exists p \in G.(B,p) \in A\}$. Since $b \in a = \text{eval}_G(A)$, there exists some $B \in \mathbf{M}^{\mathbb{P}}$ such that $b = \text{eval}_G(B)$. Thus, b is the evaluation of a $\mathbb{P}$-name in $\mathbf{M}$, and therefore is in $\mathbf{M}[G]$. Since $\mathbf{M} \subseteq \mathbf{M}[G]$, the ordinals in $\mathbf{M}$ are also in $\mathbf{M}[G]$. However, we can see easily by induction that the name rank of a $\mathbb{P}$-name $A \in \mathbf{M}^{\mathbb{P}}$ is at least the rank of its evaluation $\text{eval}_G(A)$: the name rank of A is the smallest ordinal strictly larger than all of the potential elements of A, some of which have evaluations which might not appear in the evaluation $\text{eval}_G(A)$, meaning the rank of $\text{eval}_G(A)$ cannot be larger than the name rank of A. Any ordinal $\alpha \in \mathbf{M}[G]$ is the evaluation of a name in $\mathbf{M}^{\mathbb{P}}$ with some ordinal name rank $\beta \in \mathbf{M}$. But since $\text{rank}\,(\alpha) = \alpha$ and $\text{rank}\,(\alpha) \leq \beta$, we have either $\alpha = \beta$ or $\alpha \in \beta$. The transitivity of $\mathbf{M}$ means we have $\alpha \in \mathbf{M}$, meaning the forcing extension's ordinals are all in $\mathbf{M}$. □

3.4. **The Forcing Relations $\Vdash$ and $\Vdash^*$.** We are quite close to establishing that our forcing extension $\mathbf{M}[G]$ is a model of ZFC. Given that $\mathbf{M}[G]$ is a countable transitive set with the same ordinals as $\mathbf{M}$, we can prove several of the ZFC axioms hold in the forcing extension. Proving the existence of a particular set in the forcing extension simply requires us to find the right $\mathbb{P}$-name that will evaluate to the desired set under the generic filter G. However, we still have a few problems to deal with. Some of the axioms, such as Power Set and Replacement, require us to determine whether or not a property holds in the model $\mathbf{M}[G]$ holds based only on the fact that G contains a particular forcing condition. Additionally, determining whether the desired properties hold in the forcing extension, such as the continuum hypothesis, requires us to reason about the forcing extension $\mathbf{M}[G]$ *from within the ground model* $\mathbf{M}$. Fortunately, because our generic filter G can be approximated from within $\mathbf{M}$, truth in $\mathbf{M}[G]$ is reducible to truth in $\mathbf{M}$. This will allow us to verify that the forcing extension is a model of ZFC and that the desired properties hold. We will do this by introducing a logical notion which we'll call the *forcing relation*. Despite its semantic definition, we will show that it is definable from within $\mathbf{M}$ and that it admits a syntactic definition.

In order to talk about properties of the forcing extension from within the ground model, we'll use the fact that all of the elements of the forcing extension already exist in the ground model as $\mathbb{P}$-names. Depending on which conditions are in the generic filter G, the sets that the $\mathbb{P}$-names evaluate to might have different properties. In order to discuss what properties of the forcing extension hold given a condition in the generic filter, we'll use formulas in the language of set theory but restrict the variables to quantify over the $\mathbb{P}$-names. We can then ask questions about the evaluations of the $\mathbb{P}$-names in this language.

Definition 3.19 (Name Assignment)**.** Let $\mathbf{M}$ be a countable transitive model, and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. A variable assignment s for $\mathbf{M}$ is called a $\mathbb{P}$-name assignment if its range is in $\mathbf{M}^{\mathbb{P}}$. Given a generic filter $G \subseteq \mathbb{P}$, we define the variable assignment $\text{eval}_G\, s$ for $\mathbf{M}[G]$ to be the composition $(\text{eval}_G) \circ s$.

Definition 3.20 (Forcing)**.** Let $\mathbf{M}$ be a countable transitive model, and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. Suppose we have a condition $p \in \mathbb{P}$, a $\mathbb{P}$-name assignment s for $\mathbf{M}$, and a formula φ in the language of set theory. We say that p *forces* φ under the name assignment s, written $p, s \Vdash \varphi$, if for any generic filter G containing p we have $\mathbf{M}[G], \text{eval}_G\, s \models \varphi$.

Convention 3.4. Because the forcing relation is equivalent to a statement about satisfaction in a model, if we have two name assingments s and $\tilde{s}$ which agree on the free variables of φ, then for any $p \in \mathbb{P}$ we have $p, s \Vdash \varphi$ iff $p, \tilde{s} \Vdash \varphi$. Therefore, we often abuse notation in order to avoid talking about the name assignment s. Suppose φ has free variables $x_1, \ldots x_n$, and we have $\mathbb{P}$-names $A_1, \ldots A_n$. Then we write $p \Vdash \varphi(A_1, \ldots A_n)$ if for all name assignments s sending each x_k to A_k we have $p, s \Vdash \varphi$. This is equivalent to finding a single such variable assignment since the truth value of the forcing statement only depends on what values s takes on the free variables of φ.

We can now prove several properties of the forcing relation directly from the definition:

Theorem 3.21 (Properties of Forcing)**.** *The forcing relation satisfies the following properties:*

1. *If $p, s \Vdash \varphi$ then for any $q \preceq p$ we have $q, s \Vdash \varphi$.*
2. *If Γ is a set of formulas such that for any $\varphi \in \Gamma$ we have $p, s \Vdash \varphi$ and $\Gamma \models \psi$ then $p, s \Vdash \psi$.*
3. *For any formulas φ and ψ, $p, s \Vdash \varphi \wedge \psi$ iff $p, s \Vdash \varphi$ and $p, s \Vdash \psi$.*
4. *For any formula φ and any variable x, $p, s \Vdash \forall x.\varphi$ iff for any $\mathbb{P}$-name $A \in \mathbf{M}^{\mathbb{P}}$ we have $p, s\,[A/x] \Vdash \varphi$.*

Proof. If $p, s \Vdash \varphi$ then for any generic G containing p we have $\mathbf{M}[G], \mathrm{eval}_G\, s \models \varphi$. For any $q \preceq p$, any generic G containing q will also contain p because filters are upper sets. Therefore, every generic G containing q will also satisfy $\mathbf{M}[G], \mathrm{eval}_G\, s \models \varphi$, and therefore $q, s \Vdash \varphi$. If $p, s \Vdash \varphi$ for each $\varphi \in \Gamma$, then for any generic G containing p we have $\mathbf{M}[G], \mathrm{eval}_G\, s \models \Gamma$. Since $\Gamma \models \psi$, for every generic G containing p we have $\mathbf{M}[G], \mathrm{eval}_G\, s \models \psi$, and thus $p, s \Vdash \psi$.

Applying the previous result given the entailments $\varphi \wedge \psi \models \varphi$, $\varphi \wedge \psi \models \psi$, and $\{\varphi, \psi\} \models \varphi \wedge \psi$ gives the conjunction result. $p, s \Vdash \forall x.\varphi$ iff for each generic G containing p we have $\mathbf{M}[G], \mathrm{eval}_G\, s \models \forall x.\varphi$. By definition of $\models$, this is equivalent to $\mathbf{M}[G], (\mathrm{eval}_G\, s)[A/x] \models \varphi$ for every $A \in \mathbf{M}[G]$. Since every element of $\mathbf{M}[G]$ is the evaluation of some $\mathbb{P}$-name, this is equivalent to $\mathbf{M}[G], \mathrm{eval}_G\,(s[A/x]) \models \varphi$ for every $A \in \mathbf{M}^{\mathbb{P}}$ and every generic G containing p. This is equivalent to $p, s\,[A/x] \Vdash \varphi$ for every $\mathbb{P}$-name $A \in \mathbf{M}^{\mathbb{P}}$. $\square$

Unfortunately, we can't prove much more with our definition of forcing. We'd like to be able to find a nice requirement for when a condition forces a negation, for instance. The correct statement will turn out to be that $p, s \Vdash \neg\varphi$ iff for any $q \preceq p$ we have $q, s \nVdash \varphi$. However, we can only prove the forward direction with our current knowledge; the reverse direction requires us to show that for any formula φ and any $p \in G$ there is some $q \in G$ such that $q \preceq p$ and q "decides" φ under the variable assignment s. That is, we need to show that either $q, s \Vdash \varphi$ or $q, s \Vdash \neg\varphi$. We could do this if the set of such conditions was dense and in $\mathbf{M}$, since then we could just take a common extension of the condition with p. Unfortunately, we don't yet know that this set is definable from within $\mathbf{M}$, since it refers to facts about the model $\mathbf{M}[G]$, not $\mathbf{M}$. Additionally, we'd like to say that *any* formula which holds in the model $\mathbf{M}[G]$ is forced by some condition in G, since this will allow us to prove facts about $\mathbf{M}[G]$ from within $\mathbf{M}$. Accomplishing this is the heart of forcing, and will take quite a bit of work. Here is the general strategy: first, we'll examine more carefully what it means to force two $\mathbb{P}$-names to be equal. We will be able to define this relation via a formula in $\mathbf{M}$, and show that it satisfies all of the desireable properties. Then we will define a *purely syntactical* counterpart to our forcing relation, which we'll denote $\Vdash^*$. This will be defined recursively on the structure of formulas in the language of set theory, and can be thought of as recursively transforming a formula $\varphi(A_1, \dots A_n)$ in the language of set theory to another formula $\psi(p, A_1, \dots A_n)$ which will be satisfied in $\mathbf{M}$ iff $p \Vdash \varphi(A_1, \dots A_n)$.

Although we will not prove it here, it is important for metamathematical reasons that this transformation procedure is computable. We will then show that the semantic and syntactic definitions of forcing are equivalent, meaning that the forcing relation is definable in $\mathbf{M}$, and that any formula is satisfied in $\mathbf{M}[G]$ iff it is forced by some $p \in G$.

In order to define our syntactic forcing, we'll need to determine what it means for a condition in $\mathbb{P}$ to force the atomic sentences $x \in y$ and $x = y$. We'll first examine the case of $x = y$, since this will be the most difficult. Fortunately, because our forcing extension $\mathbf{M}[G]$ is transitive, we can use the fact that it satisfies extensionality to determine when the evaluation of two $\mathbb{P}$-names are equal. This will involve an induction on the name rank of two $\mathbb{P}$-names, so we'll define $\operatorname{rank}_{\mathbb{P}}(A, B)$ to be the maximum of the name ranks of A and B.

Definition 3.22 (Forcing Equality)**.** Let $\mathbf{M}$ be a countable transitive model, and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. We'll define a sequence of sets $S_\alpha^{\mathbf{M}} \in \mathbf{M}$ by transfinite recursion, where we have $S_\alpha^{\mathbf{M}} \subseteq \{(p, A, B) \mid p \in \mathbb{P}, A, B \in \mathbf{M}^{\mathbb{P}}, \operatorname{rank}_{\mathbb{P}}(A, B) \leq \alpha\}$. A triple (p, A, B) where $p \in \mathbb{P}$ and $A, B \in \mathbf{M}^{\mathbb{P}}$ of name rank at most α will be in $S_\alpha^{\mathbf{M}}$ iff the following two requirements hold:

(1) For any $q \preceq p$, if there is some $(a, p_a) \in A$ such that $q \preceq p_a$, then there exists some $r \preceq q$ and some $(b, p_b) \in B$ such that $r \preceq p_b$ and $(r, a, b) \in S^{\mathbf{M}}_{\operatorname{rank}_{\mathbb{P}}(a,b)}$.

(2) For any $q \preceq p$, if there is some $(b, p_b) \in B$ such that $q \preceq p_b$, then there exists some $r \preceq q$ and some $(a, p_a) \in A$ such that $r \preceq p_a$ and $(r, b, a) \in S^{\mathbf{M}}_{\operatorname{rank}_{\mathbb{P}}(a,b)}$.

This transfinite recursion is well defined since the name ranks of potential elements of A and B are strictly less than the name ranks of A and B.

Theorem 3.23 (Truth and Definability Theorem for Equality)**.** *Let $\mathbf{M}$ be a countable transitive model, and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. Then for any $\mathbb{P}$-names $A, B \in \mathbf{M}^{\mathbb{P}}$, we have the following:*

(1) Definability: For any $p \in \mathbb{P}$, $p \Vdash A = B$ iff $(p, A, B) \in S^{\mathbf{M}}_{\operatorname{rank}_{\mathbb{P}}(A,B)}$.

(2) Truth: For any generic filter $G \subseteq \mathbb{P}$, $\mathbf{M}[G] \models \operatorname{eval}_G A = \operatorname{eval}_G B$ iff some $p \in G$ forces $A = B$.

Proof. We will prove this by induction on $\operatorname{rank}_{\mathbb{P}}(A, B)$. Suppose that the statements hold for all $\mathbb{P}$-names of name rank less than $\operatorname{rank}_{\mathbb{P}}(A, B)$. Then we can show it will be true for $A = B$. First, assume some $p \in \mathbb{P}$ forces $A = B$. Then by 3.21 we know $p \Vdash A \subseteq B$ and $p \Vdash B \subseteq A$. We'll examine the $A \subseteq B$ part. Pick some $q \preceq p$, and pick any generic $G \subseteq \mathbb{P}$ containing q. We know $p \Vdash A \subseteq B$, and $p \in G$, so $\operatorname{eval}_G A \subseteq \operatorname{eval}_G B$. Now suppose there is some $(a, p_a) \in A$ such that $q \preceq p_a$. Then $p_a \in G$, so $\operatorname{eval}_G a \in \operatorname{eval}_G A$. Therefore, $\operatorname{eval}_G a \in \operatorname{eval}_G B$. This means there must exist some $(b, p_b) \in B$ such that $p_b \in G$ and $\operatorname{eval}_G a = \operatorname{eval}_G b$. Since a and b are of smaller name rank than $\operatorname{rank}_{\mathbb{P}}(A, B)$, there must exist an $r_0 \in G$ such that $r_0 \Vdash a = b$. Taking the common extension of q, r_0, p_b gives a condition $r \in G$ such that $r \preceq p_b$ and $r \Vdash a = b$, which by the induction hypothesis is equivalent to $(r, a, b) \in S^{\mathbf{M}}_{\operatorname{rank}_{\mathbb{P}}(a,b)}$. This proves the first half of the definition of $S^{\mathbf{M}}_{\operatorname{rank}_{\mathbb{P}}(A,B)}$, and the same logic applied to $p \Vdash B \subseteq A$ gives the second half. Therefore, $(p, A, B) \in S^{\mathbf{M}}_{\operatorname{rank}_{\mathbb{P}}(A,B)}$. Now suppose $(p, A, B) \in S^{\mathbf{M}}_{\operatorname{rank}_{\mathbb{P}}(A,B)}$. If we can show that $p \Vdash A \subseteq B$ and $p \Vdash B \subseteq A$, then we will have $p \Vdash A = B$ since $\mathbf{M}[G]$ is transitive for any generic filter G and extensionality holds for transitive sets. Pick some generic filter $G \subseteq \mathbb{P}$ such that $p \in G$. We'll show that the first clause of the definition of $S^{\mathbf{M}}_{\mathrm{rank}_{\mathbb{P}}(A,B)}$ leads to $\mathrm{eval}_G A \subseteq \mathrm{eval}_G B$. Pick some element of $\mathrm{eval}_G A$: by the definition of the forcing extension $\mathbf{M}[G]$ it is the evaluation of some $a \in \mathbf{M}^{\mathbb{P}}$. Therefore, there is some $p_a \in G$ such that $(a, p_a) \in A$. The first clause of the definition together with the inductive hypothesis imply that the set of $r \in \mathbb{P}$ such that there is some $(b, p_b) \in B$ with $r \preceq p_b$ and $r \Vdash a = b$ is dense below p_a, so because $p_a \in G$ we know G must meet this set. There is therefore some $r \in G$ and some $(b, p_b) \in B$ such that $r \preceq p_b$ and $r \Vdash a = b$. Since $r \preceq b$, we know $\mathrm{eval}_G b \in \mathrm{eval}_G B$, and since $r \Vdash a = b$, we know $\mathrm{eval}_G a = \mathrm{eval}_G b$. Therefore, every element in the evaluation of A is equal to some element in the evaluation of b, so $p \Vdash A \subseteq B$. Applying the same logic to the other clause gives $p \Vdash B \subseteq A$, and by the transitivity of $\mathbf{M}[G]$ we know $p \Vdash A = B$.

Now we'll deal with the truth requirement. Clearly if $p \Vdash A = B$, then the evaluations of A and B will be equal in any forcing extension. Now suppose that for some generic $G \subseteq \mathbb{P}$ we have $\mathbf{M}[G] \models \mathrm{eval}_G A = \mathrm{eval}_G B$, or equivalently $\mathrm{eval}_G A = \mathrm{eval}_G B$. If we can show that there exist conditions in G which force $A \subseteq B$ and $B \subseteq A$, their common extension will force $A = B$. We'll show the existence of a $p \in G$ which forces $A \subseteq B$. We know that $p \Vdash A \subseteq B$ is equivalent to the first clause, of the definition of $S^{\mathbf{M}}_{\mathrm{rank}_{\mathbb{P}}(A,B)}$. Consider the set of all $p \in \mathbb{P}$ which either force $A \subseteq B$ or there exists some $(a, p_a) \in A$ with $p_a \preceq p$ but for any $r \preceq p$ and every $(b, p_b) \in B$ with $r \preceq p_b$ we have $r \nVdash a = b$. This is clearly a dense set of conditions since either $p \Vdash A \subseteq B$ or some $q \preceq p$ satisfies the negation of the first clause of the definition of $S^{\mathbf{M}}_{\mathrm{rank}_{\mathbb{P}}(A,B)}$. It is also a set in $\mathbf{M}$ since forcing equality between names of rank less than $\mathrm{rank}_{\mathbb{P}}(A, B)$ is definable in $\mathbf{M}$ by the inductive hypothesis. Therefore, there is some $p \in G$ satisfying this requirement by genericity. Suppose the second part of the requirement holds. Then $\mathrm{eval}_G a \in \mathrm{eval}_G A$, and therefore $\mathrm{eval}_G a \in \mathrm{eval}_G B$. Therefore, there must be some $(b, p_b) \in B$ such that $p_b \in G$ and $\mathrm{eval}_G a = \mathrm{eval}_G b$. By the inductive hypothesis, this can only be true if some condition r_0 forces it to be true, but taking the common refinement of p, p_b, and r_0 contradicts the second part of the requirement. Therefore, p cannot satisfy the second part of the requirement, meaning $p \Vdash A \subseteq B$. Applying the second clause gives a condition forcing $B \subseteq A$, and therefore a condition forcing $A = B$. □

Although the argument was somewhat tedious, the preceding proof shows that $\mathbf{M}$ knows what conditions will force the evaluations of names to be equal, and that if two names evaluate to the same set there must be a single condition forcing this fact. From this, we can fully define our syntactic forcing relation $\Vdash^*$:

Definition 3.24 (Syntactic Forcing)**.** Let $\mathbf{M}$ be a countable transitive model, and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. If φ is a formula in the language of set theory, $p \in \mathbb{P}$ is a condition, and s is a name assignment, then we define the *syntactic forcing relation* $p, s \Vdash^* \varphi$ by recursion on the structure of φ:

- (=) If $\varphi \equiv (x = y)$ where x and y are variables, then $p, s \Vdash^* \varphi$ is defined as $(p, s(x), s(y)) \in S^{\mathbf{M}}_{\mathrm{rank}_{\mathbb{P}}(s(x),s(y))}$.
- (∈) If $\varphi \equiv (x \in y)$ where x and y are variables, then $p, s \Vdash^* \varphi$ is defined as $\forall q \preceq p. \exists r \preceq q. \exists (b, p_b) \in s(y). (r \preceq p_b \wedge (r, s[b/z] \Vdash^* x = z))$.
- (¬) If $\varphi \equiv (\neg\psi)$ where ψ is a formula in the language of set theory, then $p, s \Vdash^* \varphi$ is defined as $\forall q \preceq p. (q, s \nVdash^* \psi)$.

(∧) If $\varphi \equiv (\psi \wedge \theta)$ where ψ and θ are formulas in the language of set theory, then $p, s \Vdash^* \varphi$ is defined as $(p, s \Vdash^* \psi) \wedge (p, s \Vdash^* \theta)$.

(∀) If $\varphi \equiv (\forall x.\psi)$ where x is a variable and ψ is a formula in the language of set theory, then $p, s \Vdash^* \varphi$ is defined as $\forall A \in \mathbf{M}^{\mathbb{P}}. (p, s\,[A/x] \Vdash^* \psi)$.

The definitions for $\exists$, $\implies$, and $\vee$ are defined in terms of $\neg$, $\wedge$, and $\forall$ in the natural way.

Theorem 3.25 (Truth and Definability Theorem)**.** *Let $\mathbf{M}$ be a countable transitive model, and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. For any formula φ in the language of set theory and any name assignment s, we have the following:*

(1) Definability: For any $p \in \mathbb{P}$, $p, s \Vdash \varphi$ iff $p, s \Vdash^* \varphi$.

(2) Truth: For any generic filter $G \subseteq \mathbb{P}$, $\mathbf{M}[G], \mathrm{eval}_G\, s \models \varphi$ iff there exists some $p \in G$ such that $p, s \Vdash \varphi$.

Proof. Note that one direction of truth always holds: if some condition in the generic filter forces some formula, then the formula will hold in the forcing extension. We proceed by induction on the structure of the formula φ. First, assume φ is atomic. We've already dealt with the equality case, so we just need to check it for the $\in$ case. If $\varphi \equiv A \in B$ where A and B are $\mathbb{P}$-names, then the definition of $p \Vdash^* A \in B$ is $\forall q \preceq p.\exists r \preceq q.\exists (b, p_b) \in B. (r \preceq p_b \wedge (r \Vdash^* A = b))$. Since truth and definability holds for equality, this is equivalent to $\forall q \preceq p.\exists r \preceq q.\exists (b, p_b) \in B. (r \preceq p_b \wedge (r \Vdash A = b))$. Now suppose $p \Vdash A \in B$. Pick some $q \preceq p$ and some generic filter $G \subseteq \mathbb{P}$ containing q: we will have $\mathrm{eval}_G\, A \in \mathrm{eval}_G\, B$ since $p \in G$. This can only be true if there is some $(b, p_b) \in B$ such that $p_b \in G$ and $\mathrm{eval}_G\, A = \mathrm{eval}_G\, b$. Since truth and definability hold for equality, there must exist some $r_0 \in G$ that forces $A = b$. By taking the common extension of p_b, r_0, q, we get an $r \in G$ which is an extension of p_b and q and forces $A = b$. This is the definition of $p \Vdash^* A \in B$. For the other direction, suppose $p \Vdash^* A \in B$, and consider any generic $G \subseteq \mathbb{P}$ which contains p. Then the definition of $p \Vdash^* A \in B$ tells us that the set of conditions r such that there exists some $(b, p_b) \in B$ with $r \preceq p_b$ and $r \Vdash A = b$ is dense below p. This set is in $\mathbf{M}$ since forcing equality of names is definable in $\mathbf{M}$. Since $p \in G$, G must meet this set, so there is some $r \in G$ and $(b, p_b) \in B$ such that $r \preceq p_b$ and $r \Vdash A = b$. The first condition tells us $p_b \in G$, and therefore $\mathrm{eval}_G\, b \in \mathrm{eval}_G\, B$, and the second condition tells us that $\mathrm{eval}_G\, A = \mathrm{eval}_G\, b$. Therefore, $\mathrm{eval}_G\, A \in \mathrm{eval}_G\, B$.

Now for truth. Suppose $\mathrm{eval}_G\, A \in \mathrm{eval}_G\, B$. Then there must exist $(b, p_b) \in B$ such that $p_b \in G$ and $\mathrm{eval}_G\, A = \mathrm{eval}_G\, b$. By truth for $=$, some $p \in G$ forces $A = b$. If we take the common refinement of p and p_b in G, we obtain a condition such that any generic filter containing it forces $b \in B$ and $A = b$, so it forces $A \in B$.

Now for the inductive step. We'll show that the set of formulas in the language of set theory satisfying truth and definability is closed under $\neg$, $\wedge$, and $\forall$:

(¬) Suppose that $\varphi \equiv \neg\psi$, where ψ satisfies truth and definability. Then $p, s \Vdash^* \varphi$ iff every $q \preceq p$ we have $q, s \nVdash^* \psi$. Since forcing ψ is definable, this is equivalent to $q, s \nVdash \psi$. Suppose $p, s \Vdash \varphi$. Then if any extension $q \preceq p$ satisfied $q, s \Vdash \psi$, then we could construct a generic filter containing q and p which would satisfy ψ, contradicting the fact that $p, s \Vdash \neg\psi$. Therefore, if $p, s \Vdash \neg\psi$ then $p, s \Vdash^* \neg\psi$. Now suppose $p, s \Vdash^* \neg\psi$. Let G be a generic filter extending p. If $\mathbf{M}[G], \mathrm{eval}_G\, s \models \psi$, then some $q \in G$ forces ψ under s since truth holds for ψ. Taking the common extension of p and q gives a condition stronger than p which forces ψ, which contradicts the fact that

$p, s \Vdash^* \neg\psi$. Therefore, $\mathbf{M}[G], \mathrm{eval}_G\, s \models \neg\psi$ for every generic G containing p, or in other words $p, s \Vdash \neg\psi$.

Now for truth. Suppose for some generic filter $G \subseteq \mathbb{P}$ we have $\mathbf{M}[G], \mathrm{eval}_G\, s \models \neg\psi$. Consider the set of conditions in $\mathbb{P}$ which either force ψ or have no extension forcing ψ. This set is clearly dense: for any condition in $\mathbb{P}$, either some extension forces ψ, giving us an extension in this set, or no extension does, meaning it is already in the set. This set is definable in $\mathbf{M}$ since forcing ψ is, so G meets this set of conditions. No $p \in G$ can force ψ, since otherwise $\mathbf{M}[G], \mathrm{eval}_G\, s \models \psi$. Therefore, the other requirement must be satisfied, meaning there is some $p \in G$ having no extension forcing ψ. This means $p, s \Vdash^* \neg\psi$ or equivalently $p, s \Vdash \neg\psi$.

($\wedge$) Suppose $\varphi \equiv \psi \wedge \theta$, where ψ and θ both satisfy truth and definability. By definition, $p, s \Vdash^* \varphi$ iff $p, s \Vdash^* \psi$ and $p, s \Vdash^* \theta$. By the inductive hypothesis, this is equivalent to $p, s \Vdash \psi$ and $p, s \Vdash \theta$, which theorem 3.21 tells us is equivalent to $p, s \Vdash \varphi$. This proves definability. For truth, if for some generic filter $G \subseteq \mathbb{P}$ we have $\mathbf{M}[G], \mathrm{eval}_G\, s \models \varphi$ then $\mathbf{M}[G], \mathrm{eval}_G\, s \models \psi$ and $\mathbf{M}[G], \mathrm{eval}_G\, s \models \theta$. By truth for ψ and θ, there exist $q, r \in G$ such that $q, s \Vdash \psi$ and $r, s \Vdash \theta$. Taking the common extension of these two conditions gives us $p \in G$ such that $p, s \Vdash \psi$ and $p, s \Vdash \theta$, which implies $p, s \Vdash \varphi$.

($\forall$) Suppose $\varphi \equiv \forall x.\psi$, where ψ satisfies truth and definability. Then the definition of $\Vdash^*$ tells us that $p, s \Vdash^* \varphi$ iff for every $A \in \mathbf{M}^{\mathbb{P}}$ we have $p, s\,[A/x] \Vdash^* \psi$. By definability for ψ, this is equivalent to $p, s\,[A/x] \Vdash \psi$ for all $A \in \mathbf{M}^{\mathbb{P}}$, which by theorem 3.21 is equivalent to $p, s \Vdash \varphi$. This proves definability. For truth, suppose we have some generic filter $G \subseteq \mathbb{P}$ such that $\mathbf{M}[G], \mathrm{eval}_G\, s \models \varphi$. Then for any $A \in \mathbf{M}^{\mathbb{P}}$ we have $\mathbf{M}[G], \mathrm{eval}_G\, s\,[A/x] \models \psi$. Consider the set of conditions $p \in \mathbb{P}$ such that either $p, s\,[A/x] \Vdash \neg\psi$ for some $A \in \mathbf{M}^{\mathbb{P}}$ or for every $q \preceq p$ and $A \in \mathbf{M}^{\mathbb{P}}$ we have $q, s\,[A/x] \nVdash \neg\psi$. Clearly this set is dense, and since ψ satisfies definability so does $\neg\psi$, so this set is in $\mathbf{M}$. Therefore, G meets this set, say at p. There cannot exist $A \in \mathbf{M}^{\mathbb{P}}$ such that $p, s\,[A/x] \Vdash \neg\psi$, since otherwise this would be true in G, so for every $q \preceq p$ and $A \in \mathbf{M}^{\mathbb{P}}$ we have $q, s\,[A/x] \nVdash \neg\psi$. This means for every $A \in \mathbf{M}^{\mathbb{P}}$ we have $p, s\,[A/x] \Vdash \neg\neg\psi$, which implies $p, s\,[A/x] \Vdash \psi$. Since this holds for every $\mathbb{P}$-name A, we have $p, s \Vdash \varphi$.

□

Corollary 3.26 (Additional Properties of Forcing)**.** *Let $\mathbf{M}$ be a countable transitive model, and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. Then for any condition $p \in \mathbb{P}$ and any name assignment s, the following hold:*

($\vee$) *If φ and ψ are formulas in the language of set theory, then $p, s \Vdash \varphi \vee \psi$ iff for every $q \preceq p$ there exists an $r \preceq q$ such that $r, s \Vdash \varphi$ or $r, s \Vdash \psi$.*

($\implies$) *If φ and ψ are formulas in the language of set theory, then $p, s \Vdash \varphi \implies \psi$ iff for every $q \preceq p$ either $q \nVdash \varphi$ or there is some $r \preceq q$ such that $r \Vdash \psi$.*

($\exists$) *If φ is a formula in the language of set theory and x is a variable, then $p, s \Vdash \exists x.\varphi$ iff for every $q \preceq p$ there exists an $r \preceq q$ and a $\mathbb{P}$-name $A \in \mathbf{M}^{\mathbb{P}}$ such that $r, s\,[A/x] \Vdash \varphi$.*

Furthermore, for every formula φ in the language of set theory, the set of all conditions $p \in \mathbb{P}$ such that either $p, s \Vdash \varphi$ or $p, s \Vdash \neg\varphi$ is dense and in $\mathbf{M}$. Finally, $p, s \Vdash \varphi$ iff the set of conditions which force φ under s is dense below p and in $\mathbf{M}$.

Proof. The first three statements follow from the fact that $(\varphi \vee \psi)$ is equivalent to $\neg(\neg\varphi \wedge \neg\psi)$, $(\varphi \implies \psi)$ is equivalent to $\neg(\varphi \wedge \neg\psi)$, and $\exists x.\varphi$ is equivalent to $\neg\forall x.\neg\varphi$, and applying the relevant definition for syntactic forcing. The set of conditions that force φ or $\neg\varphi$ is the same as the set of conditions which either force φ or have no extension forcing φ, and this is dense. Finally, if $p, s \Vdash \varphi$ then the set of $q \in \mathbb{P}$ such that $q, s \Vdash \varphi$ is definable in $\mathbf{M}$ and contains every $q \preceq p$, so it is dense below p. If the set of q such that $q, s \Vdash \varphi$ is dense below p, then any generic filter $G \subseteq \mathbb{P}$ which contains p must meet this set. But then G would contain a condition forcing φ, so φ would be satisfied in the forcing extension. Since this holds for every generic G, $p, s \Vdash \varphi$. □

3.5. **The Model $\mathbf{M}[G]$.** Now that we have the forcing relation and the truth and definability theorem, we are in a position to show that the forcing extension $\mathbf{M}[G]$ is a countable transitive model. The one thing we need to worry about for this proof is the fact that $\mathbf{M}$ is not really a countable transitive model of all of ZFC, but only a countable transitive model for an arbitrarily large finite fragment of ZFC. Since we want $\mathbf{M}[G]$ to also satisfy an arbitrarily large finite fragment of ZFC, we'll show that for any finite fragment $T \subseteq$ ZFC we can find a corresponding finite fragment $T' \subseteq$ ZFC such that if $\mathbf{M} \models T'$ then $\mathbf{M}[G] \models T$. This is the content of the Fundamental Theorem of Forcing. However, we can use a simple trick to avoid worrying about finite fragments. If we show that $\mathbf{M}$ being a countable transitive model of all of ZFC implies that $\mathbf{M}[G]$ is as well, then this will be enough. By the truth and definability theorem, our forcing extension $\mathbf{M}[G]$ will satisfy some axiom of ZFC iff $\mathbf{M}$ satisfies the appropriate translation of the formula. Our proof that $\mathbf{M}[G]$ satisfies this ZFC axiom whenever $\mathbf{M} \models$ ZFC can then be understood as a proof that ZFC entails the translation, or a proof of the translation from ZFC by completeness of first-order logic. This proof will only use a finite number of axioms. Therefore, any finite fragment of ZFC that holds in every forcing extension can be proven to hold from some other finite fragment of ZFC which will hold in the ground model.

Theorem 3.27 (Fundamental Theorem of Forcing)**.** *Let $\mathbf{M}$ be a countable transitive model and let $\mathbb{P} \in \mathbf{M}$ be a forcing poset. For any generic filter $G \subseteq \mathbb{P}$, the forcing extension $\mathbf{M}[G]$ is a countable transitive model of ZFC.*

Proof. By our theorem 3.18, we know that $\mathbf{M}[G]$ is transitive and countable. Transitivity implies it satisfies Extensionality and Foundation. Since $\mathbf{M} \subseteq \mathbf{M}[G]$, and $\omega \in \mathbf{M}$, we know $\omega \in \mathbf{M}[G]$, giving us Infinity/Empty Set. We'll show Union holds. Suppose A is a $\mathbb{P}$-name. Define the $\mathbb{P}$-name $\tilde{A}$ as follows:

$$\tilde{A} = \{(a,p) \mid \exists (x,q) \in A.\, ((a,r) \in x, p \preceq q, r)\}$$

For any $A \in \mathbf{M}^{\mathbb{P}}$ this $\mathbb{P}$-name will be an element of $\mathbf{M}$ since $\mathbf{M}$ satisfies Replacement. We will show that $\operatorname{eval}_G \tilde{A} = \bigcup \operatorname{eval}_G A$. Suppose $b \in \mathbf{M}^{\mathbb{P}}$ is a $\mathbb{P}$-name such that $\operatorname{eval}_G b \in \operatorname{eval}_G \tilde{A}$. Then by definition, G contains a condition p stronger than conditions q and r such that $(b,r) \in x$ for some $x \in \mathbf{M}^{\mathbb{P}}$ and $(x,q) \in A$. Since p is stronger than both r and q, $\operatorname{eval}_G b \in \operatorname{eval}_G x \in \operatorname{eval}_G A$, so $\operatorname{eval}_G b \in \bigcup \operatorname{eval}_G A$. Similarly, if b is some $\mathbb{P}$-name whose evaluation is in the union of the evaluation of A, then there must be a potential element of A with condition q such that b is in this potential element with condition r, and $q, r \in G$. Letting p be their common refinement in G shows $b \in \tilde{A}$. Therefore, $\mathbf{M}[G]$ satisfies Union.

Now we'll show that any instance of Replacement will hold in $\mathbf{M}[G]$. Let $\varphi(x, y, \overline{w}, z)$ be a formula in the language of set theory which, for some assignment of parameters $\operatorname{eval}_G \overline{w}$, is a functional over some set $\operatorname{eval}_G A$. Consider the formula $\psi(x, p, \overline{w}, A)$ which states $A, x, \overline{w} \in \mathbf{M}^{\mathbb{P}}$, $p \in \mathbb{P}$, x is a potential element of A, and there is some $y \in \mathbf{M}^{\mathbb{P}}$ such that $p \Vdash \varphi(x, y, \overline{w}, A)$. Reflection tells us that for any set A and parameters $\overline{w}$, there is a transitive set $Q \in \mathbf{M}$ containing every condition in $\mathbb{P}$ and every potential element of A, and the formula $\psi(x, p, \overline{w}, A)$ is absolute between $\mathbf{M}$ and Q for $x \in Q$. This will only require finitely many instances of Replacement in $\mathbf{M}$. Define the name B to consist of all $(b, p) \in Q \times \mathbb{P}$ such that there exists some $(a, r) \in A$ with $p \preceq r$ and $p \Vdash \varphi(b, a, \overline{w}, A)$. We'll show that the evaluation of this B is the image of $\operatorname{eval}_G A$ under φ. Suppose we have $\operatorname{eval}_G a \in \operatorname{eval}_G A$, so $(a, r) \in A$ for some $r \in G$. Then the fact that φ is a functional means there is some $b \in \mathbf{M}^{\mathbb{P}}$ such that $\varphi\left(\operatorname{eval}_G b, \operatorname{eval}_G a, \overline{\operatorname{eval}_G w}, \operatorname{eval}_G A\right)$ holds in $\mathbf{M}[G]$. Some condition in G must force this to be true, and we can assume without loss of generality that this p is stronger than r. Because φ is absolute between $\mathbf{M}[G]$ and Q and $a \in Q$, we can assume $b \in Q$, and therefore $(b, p) \in B$. Thus, $\operatorname{eval}_G b \in \operatorname{eval}_G B$. The other direction is easy: every name whose evaluation is in the evaluation of B is forced to be the image of some element in the evaluation of A under φ. Thus, $\operatorname{eval}_G B$ is the image of $\operatorname{eval}_G A$ for any parameters $\overline{w}$.

Because we have Replacement and the ordinal 2 is in $\mathbf{M}[G]$ (since $\omega \in \mathbf{M}[G]$ and $\mathbf{M}[G]$ is transitive), we have Pairing, and then applying Replacement together with Pairing gives us Separation. Therefore, to prove Power Set, all we need to show is that for every $A \in \mathbf{M}^{\mathbb{P}}$ there exists a $B \in \mathbf{M}^{\mathbb{P}}$ such that every subset of $\operatorname{eval}_G A$ that is an element of $\mathbf{M}[G]$ is an element of $\operatorname{eval}_G B$, since we can just filter out all of the elements of $\operatorname{eval}_G B$ which are not subsets of $\operatorname{eval}_G A$ in the model. It takes finitely many instances of Separation/Replacement to show that for any $A \in \mathbf{M}^{\mathbb{P}}$ there exists $B \in \mathbf{M}^{\mathbb{P}}$ such that $B = \mathcal{P}(\operatorname{dom}(A) \times \mathbb{P}) \times \{1_{\mathbb{P}}\}$. In other words, B consists of all possible sets consisting of potential elements of A tagged by elements of $\mathbb{P}$, each tagged with $1_{\mathbb{P}}$. Consider some $C \in \mathbf{M}^{\mathbb{P}}$ such that $\operatorname{eval}_G C \subseteq \operatorname{eval}_G A$. Once again using a finite number of instances of Replacement in $\mathbf{M}$, we can show that the set $\hat{C} = \{(c, p) \mid c \in \operatorname{dom}(A) \wedge p \Vdash c \in C\}$ exists in $\mathbf{M}$. This is clearly a potential element of B, so if we can show $\operatorname{eval}_G \hat{C} = \operatorname{eval}_G C$, we're done. Pick any $c \in \mathbf{M}^{\mathbb{P}}$ such that $\operatorname{eval}_G c \in \operatorname{eval}_G C$. Then some $p \in G$ forces c to be in C, and since $(c, p) \in C$, $\operatorname{eval}_G c \in \operatorname{eval}_G \hat{C}$. Assume the converse: then some $p \in G$ forces $c \in C$, and therefore $\operatorname{eval}_G c \in \operatorname{eval}_G C$. Thus, $\operatorname{eval}_G \hat{C} = \operatorname{eval}_G C$, so $\mathbf{M}[G]$ satisfies Power Set.

Finally, we'll show that Choice holds. We'll equivalently prove the well-ordering theorem. Let $A \in \mathbf{M}^{\mathbb{P}}$ be a name: since Choice holds in $\mathbf{M}$, we can well-order the potential elements of A. We can then construct a name for a well-ordering of $\operatorname{eval}_G A$ by tagging each ordered pair of the canonical name for α and the α-th element with the condition $1_{\mathbb{P}}$. In $\mathbf{M}[G]$ this will evaluate to a set of pairs consisting of all elements of $\operatorname{eval}_G A$ and ordinals, giving us a well-ordering of $\operatorname{eval}_G A$. Thus, $\mathbf{M}[G]$ is a countable transitive model of ZFC. $\square$

Corollary 3.28 (Minimality of the Forcing Extension)**.** *For any countable transitive model $\mathbf{M}$ and generic filter G over a forcing poset $\mathbb{P} \in \mathbf{M}$, the forcing extension $\mathbf{M}[G]$ is the minimal transitive model extending $\mathbf{M}$ by G. That is, for any transitive model N of ZFC such that $\mathbf{M} \subseteq N$ and $G \in N$ will contain every element of $\mathbf{M}[G]$.*

Proof. Our model N will contain every $\mathbb{P}$-name in $\mathbf{M}$, and since it contains G, the evaluation map eval_G will be definable. Thus, for each $A \in \mathbf{M}^{\mathbb{P}}$, $A \in N$ and $\operatorname{eval}_G A \in N$. Thus, $\mathbf{M}[G] \subseteq N$. □

4. Independence Proofs

Forcing gives us a general way to construct new models of set theory: just pick the appropriate notion of forcing, and use the forcing relation to determine what is true in the new model. Now that we can construct new models of set theory with desired properties, we can examine how to use this method to prove the independence of the continuum hypothesis and related results. We'll first introduce some common forcing posets and examine their combinatorial properties, which will determine how the cardinals change from $\mathbf{M}$ to $\mathbf{M}[G]$. This will allow us to tightly control the cardinality of power sets, so we can force CH to be true or not in the extension.

4.1. **The Forcing Posets $\mathrm{F}_\kappa(A, B)$ and $\mathrm{B}_\kappa(A, B)$.** The forcing method requires us to first specify forcing posets. Recall that the forcing conditions should in some sense represent approximations of the object we are trying to construct. In most cases, we will want to construct a new map between two sets, possibly a bijection between them. The obvious posets are therefore functions or bijections defined on a subset of the first set. If p and q are such partial functions, we'll say $p \preceq q$ if the domain of q is a subset of the domain of p and p restricted to the domain of q equals q. We can also put restrictions on how large the domain of these functions are allowed to be. Here are the definitions of these posets:

Definition 4.1 (Partial Function Poset)**.** Let A be an infinite set, and let κ be an infinite cardinal such that $\kappa \leq |A|$. If B is any set with at least two elements, then we can define the *partial function poset* $\mathrm{F}_\kappa(A, B)$ to consist of all functions from a subset of A of size $< \kappa$ to B. The ordering relation for $\mathrm{F}_\kappa(A, B)$ is the extension relation on partial functions.

Definition 4.2 (Partial Bijection Poset)**.** Let A and B be infinite sets, and let κ be an infinite cardinal such that $\kappa \leq |A|, |B|$. Then we can define the *partial function poset* $\mathrm{B}_\kappa(A, B)$ to consist of all bijections from a subset of A of size $< \kappa$ to a subset of B. The ordering relation for $\mathrm{B}_\kappa(A, B)$ is the extension relation on partial functions.

It is trivial to see these are both forcing posets: the extension relation is a partial order, the empty function is a least element, and we can always extend a function's domain by one element and map it to different elements of B to obtain incompatible extensions of any condition. Now that we have a notion of forcing, we can examine what objects are added in the forcing extension:

Lemma 4.3 (Forcing with Partial Functions and Partial Bijections)**.** *Let* $\mathbf{M}$ *be a countable transitive model. Suppose* $A \in \mathbf{M}$ *is an infinite set,* $\kappa \in \mathbf{M}$ *is an infinite cardinal with* $\kappa \leq |A|$ *in* $\mathbf{M}$*, and* $B \in \mathbf{M}$ *has at least two elements. If* G *is a generic filter for* $\mathrm{F}_\kappa(A, B)$*, then the forcing extension* $\mathbf{M}[G]$ *contains a surjection from* A *to* B*. Similarly, if* B *is infinite and* $\kappa \leq |A|, |B|$ *in* $\mathbf{M}$*, then the forcing extension* $\mathbf{M}[G]$ *contains a bijection between* A *and* B *for a generic filter* G *on* $\mathrm{B}_\kappa(A, B)$*.*

Proof. We know the forcing extension will contain the generic filter G. In either case, for any $a \in A \in \mathbf{M}$, the set of conditions in either $\mathrm{F}_\kappa(A, B)$ or $\mathrm{B}_\kappa(A, B)$

containing a in their domain will be a dense set, since we can always extend a partial function or bijection to be defined on a without increasing the cardinality of the domain beyond κ or not having an element in B to map a to. Similarly, for any $b \in B$ the set of conditions that contain b in their image will also be dense. Now we can take the union of our generic filter G, which will be a set in $\mathbf{M}[G]$. This union must be a function, since if the union contained ordered pairs of the form (a, b) and (a, c) for $b \neq c$, these would have to come from incompatible conditions, which cannot exist in a filter. For both posets, this function will be a function on all of A, since some condition in G is defined on a for any $a \in A$. Similarly, it will be a surjection in both cases, since some condition in G has b in its image for any $b \in B$. Finally, in the case of partial bijections, the function will also be injective, since if it contained (a, b) and $(\tilde{a}, b)$ for $a \neq \tilde{a}$, these would come from partial bijections in G which are incompatible, since their union is not a bijection. Therefore, we have an injection and surjection from A to B, or a bijection between A and B. □

4.2. **Forcing CH: The Countable Closure Property.** While the forcing posets $\mathrm{F}_\kappa(A, B)$ and $\mathrm{B}_\kappa(A, B)$ add a surjection or bijection regardless of which κ we pick, not all κ will work for controlling the cardinality of the continuum. For instance, consider forcing with $\mathrm{F}_{\aleph_0}(\aleph_0, \aleph_1)$ in some countable transitive model $\mathbf{M}$. This gives a surjection from $\aleph_0$ and $\aleph_1$, but of course $\aleph_0$ is smaller than $\aleph_1$. This is because this only gives a surjection between $\aleph_0$ and $\aleph_1^{\mathbf{M}}$, not $\aleph_1^{\mathbf{M}[G]}$. We've *collapsed* the cardinality of $\aleph_1$. This leads to a problem when we want to force CH or ¬CH: if we're not careful, the cardinals in our new model may be different than what they were in the old model. Every cardinal in $\mathbf{M}[G]$ is clearly a cardinal in $\mathbf{M}$: they share the same ordinals, and if there is no map in $\mathbf{M}[G]$ from some ordinal to a smaller ordinal then there certainly can't be one in $\mathbf{M}$. However, this does not go the other way: we can force cardinals in $\mathbf{M}$ to no longer be cardinals in $\mathbf{M}[G]$ by adding a bijection to a smaller ordinal. However, depending on our choice of κ, we can exploit the combinatorial properties of the forcing posets to control which cardinalities are changed. We'll first examine forcing CH.

Definition 4.4 (Countable Closure Property)**.** Let $\mathbb{P}$ be a forcing poset. $\mathbb{P}$ is said to be *countably closed* if for any decreasing sequence of conditions $(p_n)_{n=0}^{\infty} \in \mathbb{P}$ there exists a condition $p^* \in \mathbb{P}$ such that $p^* \preceq p_n$ for all $n \in \mathbb{N}$. In other words, any decreasing sequence of conditions of length ω has a lower bound.

Theorem 4.5 ($\mathrm{B}_{\aleph_1}(A, B)$ is Countably Closed)**.** *The poset* $\mathrm{B}_{\aleph_1}(A, B)$ *is countably closed*

Proof. Consider a decreasing sequence of conditions $(f_n)_{n=0}^{\infty}$ in $\mathrm{B}_{\aleph_1}(A, B)$. This is a sequence of functions (or possibly bijections) between countable subsets of A and subsets of B, and the functions are increasing in the sense that each function extends all previous functions. Therefore, the union of the functions is a function, and if each function is a bijection, the union will be as well. The domain of this function will be the union of the domains of each of the functions in the sequence. Since each domain is countable and there are countably many functions, this set will also be countable, and this function will be in $\mathrm{B}_{\aleph_1}(A, B)$ □

The main fact about countably closed forcing posets is the following:

Lemma 4.6 (Preservation Lemma for Countably Closed)**.** *Let* $\mathbf{M}$ *be a countable transitive model, and let* $\mathbb{P} \in \mathbf{M}$ *be a poset with the countable closure property. If*

$G \subseteq \mathbb{P}$ is a generic filter and $(f\colon \omega \longrightarrow X) \in \mathbf{M}[G]$ is a function from ω to a set $X \in \mathbf{M}$ which is in the forcing extension, then $f \in \mathbf{M}$. In other words, forcing with a countably forcing poset adds no new functions from $\mathbb{N}$ to sets in $\mathbf{M}$.

Proof. We'll use the truth and definability theorem here. Let $\tilde{f}$ be a name for f. Since $f \in \mathbf{M}[G]$ is a function between α and X, both of which are sets in X, then the truth property tells us some condition $p \in G$ forces the statement "$\tilde{f}$ is a function from $\hat{\omega}_{\mathbb{P}}$ to $\hat{X}_{\mathbb{P}}$". As a formula, this says $p \Vdash \forall n \in \hat{\omega}_{\mathbb{P}}.\exists! x \in \hat{X}_{\mathbb{P}}.\, \mathrm{op}_{\mathbb{P}}(n, x) \in \tilde{f}$, where $\mathrm{op}_{\mathbb{P}}(n, x)$ is the canonical name for the ordered pair of names (n, x). In particular, we have $p \Vdash \forall n \in \hat{\omega}_{\mathbb{P}}.\exists x \in \hat{X}_{\mathbb{P}}.\, \mathrm{op}_{\mathbb{P}}(n, x) \in \tilde{f}$. Unwrapping the definition of this forcing statement tells us that for any $q \preceq p$ and $n \in \mathbb{N}$ there exists an $x \in X$ and some condition $r \preceq q$ such that r forces $\mathrm{op}_{\mathbb{P}}(\hat{n}_{\mathbb{P}}, \hat{x}_{\mathbb{P}}) \in \tilde{f}$. Now consider the set S of functions of functions in $\mathbf{M}$ from $\mathbb{N}$ to X. If we can show that the set of conditions which force $\tilde{f} \in \hat{S}_{\mathbb{P}}$ is dense below p, then we're done, since $p \in G$ and therefore our forcing extension will satisfy $\mathrm{eval}_G\, \tilde{f} = f \in S \in \mathbf{M}$, so by transitivity $f \in \mathbf{M}$. Pick some $q \preceq p$. We can recursively construct a decreasing sequence of conditions $(q_n)_{n=0}^{\infty}$ where $q_n \preceq q$ and a sequence $(x_n)_{n=0}^{\infty} \in X$ such that $q_n \Vdash \mathrm{op}_{\mathbb{P}}(\hat{n}_{\mathbb{P}}, \hat{x_n}_{\mathbb{P}}) \in \tilde{f}$. To construct a given q_n and x_n, we take a lower bound of all of the q_m for $m < n$, and use the fact that this is stronger than q to pick a q_n stronger than the lower bound and an $x_n \in X$ such q_n forces the function $\tilde{f}$ to equal x_n at n. Now using the countable closure condition, we can find a $q^* \preceq q_n$ for all $n \in \mathbb{N}$. For each $n < \mathbb{N}$, $q^* \Vdash \mathrm{op}_{\mathbb{P}}(\hat{n}_{\mathbb{P}}, \hat{x_n}_{\mathbb{P}}) \in \tilde{f}$. Since $q^* \preceq q$, q^* also forces $\tilde{f}$ to be a function, so q^* forces $\tilde{f}$ to be the same function as $n \mapsto x_n$. But this last function is in S, and therefore q^* forces $\tilde{f}$ to be in $\hat{S}_{\mathbb{P}}$. Since this construction worked for any $q \preceq p$, we have that the set of conditions forcing $\tilde{f}$ to be in $\hat{S}_{\mathbb{P}}$ is dense below p, so we're done. □

Theorem 4.7 (CH is Relatively Consistent with ZFC)**.** *Let $\mathbf{M}$ be a countable transitive model. Then there is a notion of forcing $\mathbb{P} \in \mathbf{M}$ such that for every generic filter $G \subseteq \mathbb{P}$ the forcing extension $\mathbf{M}[G]$ is a model for* ZFC + CH.

Proof. We want $2^{\aleph_0} = \aleph_1$ in our model, so we need to introduce a bijection between $\mathcal{P}(\mathbb{N})$ and $\aleph_1$. We know forcing with $\mathrm{B}_{\aleph_1}(\mathcal{P}(\mathbb{N}), \aleph_1)$ will introduce a bijection between $\mathcal{P}^{\mathbf{M}}(\mathbb{N})$ and $\aleph_1^{\mathbf{M}}$ in our forcing extension $\mathbf{M}[G]$. Since the poset is countably closed, we will add no new functions defined on ω, and therefore no new functions defined on countable sets in $\mathbf{M}$. Therefore, we add no new subsets of $\mathbb{N}$, since a subset can be encoded as a function from $\mathbb{N}$ to 2. Similarly, $\aleph_1$ in $\mathbf{M}$ will still be the smallest uncountable cardinal in the forcing extension, since we cannot add a surjection from a countable ordinal onto it. Therefore, our bijection between $\mathcal{P}^{\mathbf{M}}(\mathbb{N})$ and $\aleph_1^{\mathbf{M}}$ in $\mathbf{M}[G]$ will be a bijection between $\mathcal{P}(\mathbb{N})$ and $\aleph_1$ in $\mathbf{M}[G]$, meaning $\mathbf{M}[G]$ is a model of ZFC where the continuum hypothesis holds. □

4.3. **Forcing ¬CH: The Countable Antichain Condition.** We will now examine how to force ¬CH. This will require a different combinatorial property. Whereas before our concern was that our new model had more subsets of natural numbers than our ground model, here we want to make sure that $\aleph_2$ doesn't collapse down to a smaller cardinal. Our approach will be to add $\aleph_2$ new subsets of $\mathbb{N}$, and then show that this $\aleph_2$ is the same as the $\aleph_2$ in the new model. In other words, we will want to preserve cardinals *above* $\aleph_1$, rather than below $\aleph_1$.

Definition 4.8 (Antichain of a Poset)**.** Let $\mathbb{P}$ be any poset. An *antichain* of $\mathbb{P}$ is a subset of $\mathbb{P}$ where every pair of distinct elements are incompatible.

Definition 4.9 (Countable Antichain Condition)**.** Let $\mathbb{P}$ be a forcing poset, and let κ be an infinite cardinal. $\mathbb{P}$ is said to satisfy the *countable antichain condition* or c.a.c. if every antichain in $\mathbb{P}$ is countable.

Theorem 4.10 (Approximation Lemma for c.a.c. Forcing)**.** *Let* $\mathbf{M}$ *be a countable transitive model, let* $X, Y \in \mathbf{M}$ *be any sets, and let* $\mathbb{P} \in \mathbf{M}$ *be a forcing poset satisfying the c.a.c. If* $G \subseteq \mathbb{P}$ *is a generic filter, then for any function* $(f : X \longrightarrow Y) \in \mathbf{M}[G]$ *there exists a function* $(g : X \longrightarrow \mathcal{P}(Y)) \in \mathbf{M}$ *such that* $f(x) \in g(x)$ *for all* $x \in X$ *and each* $g(x)$ *is countable in* $\mathbf{M}$*.*

Proof. We'll use the truth and definability theorem. Let $\tilde{f}$ be a name for f in $\mathbf{M}$, and let $q \in G$ force "$\tilde{f}$ is a function between $\hat{X}_{\mathbb{P}}$ and $\hat{Y}_{\mathbb{P}}$". Such a q must exist by the truth property. We'll define $g(x)$ as the set of all $y \in Y$ such that some extension of q forces $\tilde{f}(\hat{x}_{\mathbb{P}}) = \hat{y}_{\mathbb{P}}$. This set is definable within $\mathbf{M}$ since the forcing relation is definable. For every x we clearly have $f(x) \in g(x)$, since the value of $f(x)$ will be forced by some condition in G, which we can take to be an extension of q. Now we just need to show that the cardinality of each $g(x)$ is countable. Fix some $x \in X$. For each $y \in g(x)$, use Choice to pick a unique $q_y \preceq q$ such that $\tilde{f}(\hat{x}_{\mathbb{P}}) = \hat{y}_{\mathbb{P}}$. We'll show that this set is an antichain. If q_y and q_z are compatible, then they have some common extension, which will force $\tilde{f}(\hat{x}_{\mathbb{P}}) = \hat{y}_{\mathbb{P}}$ and $\tilde{f}(\hat{x}_{\mathbb{P}}) = \hat{z}_{\mathbb{P}}$, so it will in particular force $\hat{y}_{\mathbb{P}} = \hat{z}_{\mathbb{P}}$. Since these are the canonical names for y and z, we must have $y = z$ and therefore $q_y = q_z$. This set is therefore an antichain, and must therefore be countable. Since this set has the same cardinality as $g(x)$, $|g(x)| \leq \omega$ for each $x \in X$. □

Corollary 4.11 (Cardinal Preservation for c.a.c. Forcing)**.** *Let* $\mathbf{M}$ *be a countable transitive model, and let* $\mathbb{P} \in \mathbf{M}$ *be a forcing poset satisfying the c.a.c. Then for any generic filter* $G \subseteq \mathbb{P}$, $\mathbf{M}$ *and* $\mathbf{M}[G]$ *have the same cardinals.*

Proof. $\mathbf{M}$ and $\mathbf{M}[G]$ have the same cardinals less than or equal to ω. Suppose α is an ordinal in $\mathbf{M}[G]$ (and therefore in $\mathbf{M}$) such that $\alpha > \omega$ and α is not a cardinal. Then there is a surjection onto α from an ordinal $\beta < \alpha$, which is clearly at least ω. We will show α is not a cardinal in $\mathbf{M}$ either. Let $f : \beta \longrightarrow \alpha$ be the surjection in the forcing extension. Since our forcing satisfies the countable antichain condition, there will be an approximation $g : \beta \longrightarrow \mathcal{P}(\alpha)$ in the ground model. If we take the union of the image of g, we will obtain α. Since each set $g(x)$ is countable and there are $|\beta|$ many such sets, the image has cardinality at most $\beta \times \omega = \beta$, since $\beta \geq \omega$. Therefore, there is a surjection in $\mathbf{M}$ from β onto α, so α isn't a cardinal in $\mathbf{M}$. We always know that cardinals in the forcing extension are cardinals in the ground model, so this shows the forcing extension and ground model must have the same cardinals. □

Theorem 4.12 ($\mathrm{F}_{\aleph_0}(A, B)$ is c.a.c.)**.** *The poset* $\mathrm{F}_{\aleph_0}(A, B)$ *satisfies the c.a.c. so long as* B *is countable.*

Proof. Suppose I is an antichain of conditions: we will show that it is countable. To do so, we will define a sequence of subsets $(A_n)_{n=0}^{\infty} \subseteq A$ by induction. We'll set $A_0 = \emptyset$. Suppose we've defined A_n. For each partial function from a finite subset of A_n, pick a $q \in I$ extending it, if such a q exists. A_{n+1} will be the union of A_n and the domains of each of these $q \in I$. By induction, we can see that each A_n is countable: the empty set has cardinality 0, and at successor stages, we are only adding at most the number of functions from all finite subsets of a countable set to another countable set. Since each such function is a finite subset of the Cartesian product, there will be at most $(\aleph_0 \cdot \aleph_0)^{<\aleph_0} = \aleph_0$ of them. A_{n+1} is therefore countable. The set $A_\omega = \bigcup_{n\in\mathbb{N}} A_n$ is also countable, as is the set of all partial functions with finite domain from A_ω to B. We'll show that each element of I is in this set. Pick some condition $p \in I$. Since the domain of p is a finite set, its intersection with A_ω is also finite. Since A_ω is the union of all of the A_n, and there are $\aleph_0$ of these sets, the fact that the domain of p is finite implies there exsits some $n \in \mathbb{N}$ such that for all $m \geq n$ we have $\mathrm{dom}(p) \cap A_m = \mathrm{dom}(p) \cap A_{m+1}$. If we consider the restriction of p to A_n, this is a function from a finite subset of A_n to B. Since this function is the restriction of some condition, namely p, to A_n, there must exist $q \in I$ whose restriction to A_n is the restriction of p to A_n and whose domain was added in A_{n+1}. We will show p and q are compatible: any x in the domain of both p and q is in A_{n+1} and therefore in $\mathrm{dom}(p) \cap A_{n+1} = \mathrm{dom}(p) \cap A_n$, and we know p and q agree on A_n. Therefore, p and q have a common extension: just take their union. Since p and q are both in I and are compatible, $p = q$, so $\mathrm{dom}(p) \subseteq A_{n+1} \subseteq A_\omega$. Thus, every element of I has domain in A_ω. The cardinality of I is at most the number of functions from finite subsets of A_ω to a countable set. Since A_ω is countable, we know this set has cardinality at most $(\aleph_0 \cdot \aleph_0)^{<\aleph_0} = \aleph_0$. □

Now that we have the antichain condition for our partial function poset, we can use it to force a model where there are $\geq \aleph_2$ subsets of $\mathbb{N}$:

Theorem 4.13 (¬CH is Relatively Consistent with ZFC)**.** *Let* $\mathbf{M}$ *be a countable transitive model. Then there is a notion of forcing* $\mathbb{P} \in \mathbf{M}$ *such that for every generic filter* $G \subseteq \mathbb{P}$ *the foring extension* $\mathbf{M}[G]$ *is a model for* $\mathrm{ZFC} + \neg\mathrm{CH}$.

Proof. We want $2^{\aleph_0} \geq \aleph_2$ in our model, so we need to introduce at least $\aleph_2$ new subsets of $\mathbb{N}$ into our model. A collection of $\aleph_2$ subsets can be encoded as an injective function from $\aleph_2$ to $2^{\aleph_0}$, so we will force with the poset $\mathrm{F}_{\aleph_0}(\aleph_2, \mathcal{P}(\mathbb{N}))$. This will introduce an injection from $\aleph_2^{\mathbf{M}}$ into $\mathcal{P}(\mathbb{N})\mathbf{M}$ in $\mathbf{M}[G]$. Since $\mathcal{P}^{\mathbf{M}}(\mathbb{N}) \subseteq \mathcal{P}^{\mathbf{M}[G]}(\mathbb{N})$, we will have an injection from $\aleph_2^{\mathbf{M}}$ into the power set of $\mathbb{N}$. We know that our poset satisfies the c.a.c. Therefore, the forcing extension will have the same cardinals as the ground model. In particular, $\aleph_2^{\mathbf{M}} = \aleph_2^{\mathbf{M}[G]}$, so in $\mathbf{M}[G]$ we will have $2^{\aleph_0} \geq \aleph_2$. This means $\mathbf{M}[G]$ will be a model of ZFC satisfying the negation of the continuum hypothesis. □

Acknowledgments

I would like to thank my mentor Miles Kretschmer for his incredible patience and guidance throughout the program. I wouldn't have begun to understand forcing without his help. I would also like to thank all of the professors and lecturers who took their summers to point the REU participants to often overlooked areas of mathematics. Finally, I'd like to give a special acknowledgement to Professor Peter May, whose tireless work made this incredible program possible.

References

[1] Denis R Hirschfeldt, Chitat Chong, Qi Feng, Theodore A Slaman, W Hugh Woodin, and Yue Yang. *Slicing the Truth.* WORLD SCIENTIFIC, 2014.
[2] Wilfrid Hodges. *A Shorter Model Theory.* Cambridge University Press, USA, 1997.
[3] Thomas Jech. *Set Theory: The Third Millennium Edition.* Springer, 2003.

[4] Kenneth Kunen. *Set Theory: An Introduction to Independence Proofs*. North-Holland, 1980.
[5] J. R. Shoenfield. Unramified forcing. *Axiomatic Set Theory*, 13:357–381.
[6] Nik Weaver. *Forcing for Mathematicians*. WORLD SCIENTIFIC, 2014.
[7] Richard Zach. *The Open Logic Text*. The Open Logic Project, 2012.