Notes on Smullyan & Fitting, and on Forcing

Michael Weiss

March 24, 2019

I participated in a meetup devoted (initially) to going through Smullyan and Fitting's *Set Theory and the Continuum Problem* [25] (S&F for short). Later on we switched to other texts for forcing [10, 23].

I wrote these notes to support the meetup. Several sections are tethered closely to S&F (e.g., §§1–3, 5), but many others stand on their own; still others are a mixed bag (e.g., §20).

I wrote a companion set of notes giving general background in logic ("Basics of First-order Logic", referred to here as the Logic notes). These are not tied to S&F or indeed to any textbook.

I wish to thank my fellow participants, especially Bob Levine, Andras Pap, and Rob Sides.

1 S&F Errata

These are errors not listed in any of the (numerous) online errata (including the list by Olkhovikov).

1. p.49, near the end of the proof of Prop. 1.2: $x'Rb'$ should be $b'Rx'$, and $x \leq b$ should be $b \leq x$.
2. p.80, in proof of Prop. 1.3: $L(x)$ should be $L_<(x)$.
3. p.91, in Remarks: R_{n+1} doesn't have 2^n elements, instead it has $2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}$, with $(n-1)$ 2's in the tower.
4. p.104, in proof of Theorem 4.6: $g(x)$ should be $g''x$.
5. p.183, in proof of Theorem 2.8: $F(f \restriction \beta)$ should be $F(f(\beta))$.
6. p.195, in paragraph about hereditary cardinality: with their definition of the transitive closure, x is not its own transitive closure for a transitive set x, since $x \notin x$ but x is an element of its transitive closure.

2 Summary of Definitions

Transitive (2.1, p.17): A is **transitive** if $x \in y \in A \Rightarrow x \in A$.

Swelled (2.2, p.17): A is **swelled** if $x \subseteq y \in A \Rightarrow x \in A$.

Supercomplete (p.17): both transitive and swelled.

Nest (4.5, p.36): A is a **nest** if any two elements are comparable under $\subseteq$, so A is linearly ordered under $\subseteq$.

Chain (p.52): A nest that is also a set.

Superinductive under g (p.52): A is **superinductive** if

1. $\varnothing \in A$.
2. $x \in A \Rightarrow g(x) \in A$.

3. A is closed under chain unions: $C \subseteq A \Rightarrow \bigcup C \in A$ for any chain C.

Minimally Superinductive (p.53): A is superinductive but no proper subclass of A is superinductive.

Progressing (4.4, p.36): g is **progressing** if $x \subseteq g(x)$ for all x in the domain of g.

Strictly progressing (p.55): $x \subset g(x)$ for all x in the domain of g.

Slowly progressing (4.2, p.58): progressing and $g(x)$ contains at most one more element than x.

g-tower (p.54): M is a **g-tower** if M is minimally superinductive and g is progressing.

Slow g-tower (4.2, p.58): g-tower where g is slowly progressing.

$L_<$, $L_\leq$ (p.50): $L_<(x) = \{y : y < x\}$, $L_\leq(x) = \{y : y \leq x\}$.

Slow well-ordering (4.1, p.57): N is **slowly** well-ordered under $\subseteq$ if

1. $\varnothing$ is the least element of N.
2. If y is the successor of x, then y contains exactly one more element than x.
3. If x is a limit element, than $x = \bigcup L_<(x)$, i.e., x is the union of the set of all elements less than x.

type *M* (p.61): $(S, \leq)$ is of type M if every element of S is $\leq$ a maximal element of S. (S&F actually define this only for the special case where $\leq$ is $\subseteq$.)

***A* is of finite character** (p.61): $x \in A \Leftrightarrow$ all finite subsets of x are in A.

B* is bounded by *x (4.11, p.38): $(\forall b \in B) b \subseteq x$.

g-set (p.66): x is a g-set if x belongs to every class that is superinductive under g.

Lower Section (p.80): A lower section of a linear ordering$(A,<)$ is a B such that all elements of B are less than all elements of $A \setminus B$. Equivalently, $x < y \in B \Rightarrow x \in B$. Many people call this an **initial segment**. A lower section is proper if it's a proper subset.

Proper well-ordering (p.80): Every proper lower section is a set. The canonical example is On, the class of all ordinals.

Basic universe, Zermelo-(Fraenkel) universe (p.27,32,82): This requires quite a bit of discussion: see §3.

Relational system (p.127): $\Gamma = (A, R)$, a class A with a relation R on it.

a^* (p.127): Given a relational system (A, R) with $a \in A$, $a^* = \{x : xRa\}$ (in general, a class).

Components (p.127): In a relational system (A, R), the elements of a^* are the *components* of a.

Proper relational system (p.127): for all $a \in A$, a^* is a set.

Extensional relational system (p.128): for all $x, y \in A$, $x \neq y \Rightarrow x^* \neq y^*$.

$p(x)$ (p.130): For $x \subseteq A$, $p(x) = \{a \in A : a^* \subseteq x\}$.

Mostowski-Shepherdson Map (p.134): For a proper extensional well-founded relational system (A, R), the unique function F on A such that $F(x) = F''(x^*)$ for all $x \in A$. Usually just called the Mostowski collapsing map.

A-formula, pure formula (p.142): An A-formula can have elements of A (treated syntactically like constants); a pure formula has no such constants.

Elementary equivalence (p.144): Structures are elementarily equivalent if every closed pure formula holds in one $\Leftrightarrow$ it holds in the other.

Elementary subsystem (p.144): Say **A** is a substructure of **B**. **A** is an elementary substructure of **B** if every closed A-formula holds in **A** $\Leftrightarrow$ it holds in **B**. (S&F say subsystem where everyone else says either substructure or submodel.) Since their structures are all relational systems, 'substructure' means $A \subseteq B$ and $R_A = R_B \restriction A$.

Reflects (p.145): Say **A** is a substructure of **B**. **A** reflects **B** with respect to a pure formula $\varphi(\vec{x})$ if $\mathbf{A} \models \varphi(\vec{a}) \Leftrightarrow \mathbf{B} \models \varphi(\vec{a})$ for all $\vec{a} \in A$. **A** completely reflects **B** w.r.t. φ if **A** reflects **B** w.r.t. all subformulas of φ, including φ.

Henkin closed (p.145): Say **A** is a substructure of **B**. **A** is Henkin closed if for every A-formula $\varphi(x)$ with x as the only free variable, if $\mathbf{B} \models (\exists x)\varphi(x)$ then for some $a \in A$, $\mathbf{B} \models \varphi(a)$. So constants from the substructure are allowed, and witnesses can always be found in the substructure, but satisfaction is computed in the superstructure.

Order of a constructible set (p.156): The least ordinal α such that $L_{\alpha+1}$ contains the set. (Like rank, except for constructibility.)

Absolute (p.158): Given a class K and a closed K-formula φ, let φ^K be the result of replacing every $(\forall x)$ with $(\forall x \in K)$ and every $(\exists x)$ with $(\exists x \in K)$. Then φ is absolute over K if $V \models \varphi \Leftrightarrow V \models \varphi^K$. A pure formula $\varphi(\vec{x})$ is absolute over K if $\varphi(\vec{a})$ is absolute over K for all $\vec{a} \in K$. A formula is absolute if it is absolute over every *transitive* class.

Note that $K \models \varphi \Leftrightarrow V \models \varphi^K$.

Absolute upwards, downwards (p.179): $\varphi(\vec{x})$ is absolute upwards over K if for all $\vec{a} \in K$, $V \models \varphi^K(\vec{a}) \Rightarrow V \models \varphi(\vec{a})$. Replacing $\Rightarrow$ with $\Leftarrow$ gives absolute downwards.

Δ_0-formula (p.159): All quantifiers must be bounded, i.e., of the form $(\exists x \in y)$ or $(\forall x \in y)$ for some set y. Key fact: Δ_0-formulas are absolute.

Σ-formula (p.180): We're allowed unbounded existential quantifiers, but only bounded universal quantifiers. Key fact: Σ-formulas are absolute upwards.

Constructible class (p.163): A is constructible if (a) every $a \in A$ is constructible, and (b) $A \cap b$ is constructible for all constructible sets b.

Distinguished subclass (p.163): Given a class K, $A \subseteq K$ is distinguished if $A \cap b \in K$ for all $b \in K$. (So condition (b) in the previous item just says that A is distinguished for L.)

First-order swelled (p.169): K is first order swelled if for all $x \in K$, $\mathcal{P}^{1\text{st}}(x) \subseteq K$; here $\mathcal{P}^{1\text{st}}(x)$ is the set of all first-order definable subsets of x.

First-order universe (p.169): In the definition of Zermelo-Fraenkel universe, replace "swelled" with "first-order swelled", replace $\mathcal{P}(x)$ with $\mathcal{P}(x) \cap K$, and in the axiom of substitution (aka replacement), replace the function F with a first-order definable function.

3 The Axioms of Set Theory

First some page references:

P_1	15	Extensionality
P_2	16	Separation
A_1	17	Transitive
A_2	18	Swelled
A_3	18	Empty Set
A_4	19	Pairing
A_5	21	Union
A_6	23	Power Set
A_7	32	Infinity
A_8	82	Substitution (aka Replacement)
D	99	Foundation (Well-foundedness)
—	8	Choice
Ax1–Ax9	170–171	Standard ZF axioms
A_1–A_6	27	Basic universe
A_1–A_7	32	Zermelo universe
A_1–A_8	82	Zermelo-Fraenkel universe

S&F do things a bit differently from the standard treatment. The differences are technical but worth going into. Axioms Ax1 through Ax9 *are* standard. A model of these axioms is a universe for ZF set theory, in the usual sense. Also, A_3–A_8 correspond pretty directly with Ax3–Ax8, and axiom D corresponds to Ax9. So far so good.

You might think, naturally enough, that a Zermelo-Fraenkel universe is just a model for ZF set theory (without Foundation). Not so! P_1 corresponds to Ax1 (Extensionality) and P_2 to Ax2 (Separation). The Zermelo-Fraenkel universe of S&F replaces P_1 and P_2 with A_1 (V is transitive) and A_2 (V is swelled).

I think S&F wall off P_1 and P_2 from the rest of the ZF axioms because they apply to classes and not just to sets. (Now is probably a good time to recall the S&F convention (p.15) that lower case variables range over sets, upper case over classes.)

A_1 and A_2 are axioms neither of ZF nor of NBG, which they claim is the main system they use. Let's see why not. Informally, in ZF, they say this:

$$\begin{aligned} A_1 &: \quad x \in y \in V \to x \in V \\ A_2 &: \quad x \subseteq y \in V \to x \in V \end{aligned}$$

which looks OK. However, in ZF there is no constant V; the variables range implicitly over the universe of all sets. So there's no right-hand side—nothing to substitute for "$x \in V$". Stated in words: *everything* in ZF is a set, so the assertion "elements and subsets of sets are sets" means nothing.

The S&F treatment of NBG isn't really consistent. On the top of p.16 they forbid quantifying over all classes: "...we allow $\forall x$, $\exists x$, for x a *set* variable, but we do not allow $\forall A$, $\exists A$, where A is a *class* variable." At the end of that *very paragraph*, they quantify over classes in their statement of P_2! Anyway, let's write A_1 and A_2 as formulas in NBG, ignoring the "no class quantification" rule.

$$\begin{aligned} A_1 &: \quad (\exists Z) X \in Y \in Z \to (\exists U) X \in U \\ A_2 &: \quad (\exists Z) X \subseteq Y \in Z \to (\exists U) X \in U \end{aligned}$$

The first is a consequence of basic logic, the second a consequence of Power Set. In NBG, elements of classes are sets, so elements of elements are *ipso facto* sets; a subset of Y is an element of the power set of Y, and so is a set.[1]

[1] Possibly A_1 and A_2 are borrowed from Ackermann set theory [11, 14]. Ackermann's theory is equivalent to NBG with regard to *sets* (and both are equivalent to ZF), provided Foundation is included. But classes behave differently: some proper classes are elements of other classes in Ackermann's theory, while in NBG, "being an element of something" is the *definition* of a set. Sethood is a primitive concept in Ackermann's theory, most conveniently formalized by adding a class constant V representing the class of all sets. Ackermann's theory dispenses with most of the set-building axioms of NBG, replacing them with a set comprehension axiom schema. S&F never mention Ackermann set theory, nor use its terminology ("hereditary" instead of "supercomplete"), nor any of its other axioms. The Ackermann-style interpretation of A_1 and A_2 will not rescue Theorem 3.6 (see below).

I believe S&F really had the "inner model" situation in mind when they formulated A_1 and A_2. Only in this way can I make sense of their claims about Zermelo-Fraenkel universes. Let M be a subclass of V. When is $(M, \in)$ a model of the ZF axioms? If M is transitive, then it satisfies Extensionality because V does (trivial exercise). Likewise, if M is swelled then it inherits Separation from V (exercise).

The converses fail: Extensionality does not imply transitivity and Separation does not imply swellitude. This issue is not too bad for Extensionality (because of Mostowski collapse—see §19), but it's severe for swellitude. Briefly, P_2 only allows us to carve off subsets using first-order formulas. Consequently, P_2 only implies "first-order swellitude", which is defined in Ch.13 (p.169).

Swelled vs. first-order swelled comes to a head in Theorem 3.6 of Ch.8 (p.102). This asserts that a well-founded Zermelo-Fraenkel universe is either R_α for some α, or is R_Ω. Initially I read "Zermelo-Fraenkel universe" as "standard model of ZF (with Foundation)", as did Pollard in his review [20, §§2.7–2.8]. Pollard even gives a counter-example to this interpretation. If we read "Zermelo-Fraenkel universe" as "transitive and swelled subclass of V satisfying relativized versions of the other axioms", then the theorem is OK.

Where to lay the blame, Pollard or S&F? I lay it at S&F's doorstep. They nowhere give a clear statement of their intended interpretation, and Pollard's reading seems perfectly in keeping with their informal treatment of basic logical issues to this point. They introduce the notion of *inner model* only much later.

So far as I can see S&F never give an explicit label (like A_{10} or Ax10) to the Axiom of Choice.

4 Zermelo's Proof

For my money, Zermelo's original 1904 proof of the well-ordering theorem is easier to understand than the S&F proof. Here is a sketch. Background assumed: the definition of a well-ordering, including transfinite induction on a well-ordered set.

The axiom of choice is used in the following form: if S is a nonempty set, then there is a choice function c defined for all proper subsets $X \subset S$ with $c(X)$ being an element *outside* X:

$$c : \mathcal{P}(S) \smallsetminus \{S\} \to S; \quad c(X) \in S \smallsetminus X$$

We say that an irreflexive well-ordering $<$ on a subset $X \subseteq S$ is **good** if the following holds:

> For all $x \in X$, $c(L_<(x)) = x$.

Observe that x is the smallest element of $S \smallsetminus L_<(x)$.

Now let $(X, <)$ and $(Y, \prec)$ be two good well-orderings on subsets of S, with $\leq$ and $\preceq$ the reflexive versions. We prove the following by transfinite induction on the well-ordering $<$: for all $x \in X$,

1. either $L_\leq(x) \subseteq Y$ or $Y \subseteq L_\leq(x)$ (or both);
2. $L_\leq(x) \cap Y$ is ordered identically under $\leq$ and $\preceq$;
3. for any $y \in L_\leq(x) \cap Y$, $L_\preceq(y) \subseteq L_\leq(x)$.

Informally:

- Either $L_\leq(x)$ is an initial segment of Y, or Y is an initial segment of $L_\leq(x)$, or they are equal.

- The orderings $\leq$ and $\preceq$ agree on the common initial segment.

The induction is straightforward. Assume (1)–(3) hold for all $x' < x$. One possibility is that Y "runs out of elements" before $L_{\leq}(x)$ does, that is, $Y \subseteq L_{\leq}(x')$ for some $x' < x$. Then (1)–(3) for x' automatically imply (1)–(3) for x also. Otherwise, we must have $\bigcup_{x'<x} L_{\leq}(x') = L_{<}(x)$ contained in Y, identically ordered under $\leq$ and $\preceq$, and (by (3)) "with no $\prec$-gaps": if $y' \prec y \in L_{<}(x)$, then $y' \in L_{<}(x)$. So $L_{<}(x) = L_{\prec}(x)$, identically ordered under $<$ and $\prec$.

It's possible that $L_{<}(x)$ is *all* of Y, in which case it is easy to see that (1)–(3) hold for x. Otherwise $x = c(L_{<}(x)) = c(L_{\prec}(x))$ is the $<$-smallest element of $X \setminus L_{<}(x)$ and the $\prec$-smallest element of $Y \setminus L_{\prec}(x)$. In other words, you tack on the same element x to the initial segments $L_{<}(x) \subset X$ and $L_{\prec}(x) \subset Y$ to obtain $L_{\leq}(x)$ and $L_{\preceq}(x)$, and from this it follows immediately that (1)–(3) hold for x.

Now we let U be the union of the domains of *all* the good well-orderings. All these well orderings must agree on their common parts (if x and y both belong to the domains of $<$ and $\prec$, then $x < y$ iff $x \prec y$) so we have a well defined well-ordering on U. U must equal S, otherwise we could append $c(U)$ as the largest element of $U \cup \{c(U)\}$ and have a larger good well-ordering. So S is well-ordered. ▌

5 The S&F Proof of the Well-Ordering Theorem

In 1908 Zermelo gave a second proof of the well-ordering theorem; the S&F proof resembles it. Zermelo's first proof "builds up" the well-ordering (key step: taking the *union* of all the good well-orderings); his second proof "cuts down" to the well-ordering (key step: taking an *intersection* of a collection of sets). S&F likewise take an intersection.

I found the S&F proof confusing partly because of the ever-changing hypotheses in §§4.2–4.4. I liked it better when I strung the lemmas and theorems into a single argument, with a fixed scenario: S is a set with a choice function c choosing, for each proper subset $X \subset S$, an element $c(X)$ in the complement $S \setminus X$.

A sketch:

(1) Theorem 4.7 (p.59): We let $g : \mathcal{P}(S) \to \mathcal{P}(S)$ be defined by $g(x) = x \cup \{c(x)\}$ for proper subsets $x \subset S$, and let $g(S) = S$. So g is slowly progressing with S as its only fixed point.

(2) Theorem 4.6 (p.59): Now we let M be the intersection of all the g-superinductive subsets of $\mathcal{P}(S)$. ($\mathcal{P}(S)$ is obviously superinductive.) So M is minimally g-superinductive, and we can apply *proof by superinduction* (p.53) to it; this is sort of like transfinite induction, though not quite. (Compare the top of p.53 with Theorem 1.9 on p.50.) We need to show four things: (a) M is linearly ordered under $\subseteq$, i.e., M is a chain; (b) $S \in M$; (c) M is well-ordered under $\subseteq$; (d) the well-ordering of M under $\subseteq$ can be transferred to a well-ordering of $\bigcup M$; note that $\bigcup M = S$ by (b).

(3) To prove (a), we use Theorem 2.1 (double superinduction, p.53) plus Lemma 2.2 and Theorem 2.3 (p.54). I won't go through the details. But one remark: it's pretty easy to prove comparability just with ordinary (non-double) superinduction. All we need to show is preservation under tacking on an element at the end, and under the union of chains.

(4) The proof of (b) now falls out immediately: since M is a chain, $\bigcup M \in M$, and is obviously the largest element of M. So $g\left(\bigcup M\right) = \bigcup M$, and since S is the only fixed point of g, $S \in M$. (This is the argument of Theorems 2.5 and 2.6, pp.54–55.)

(5) To prove (c), let A be a subset of M. Note that (c) is the essence of

transfinite induction; we expect to use superinduction somehow. Suppose A has no least element; we will show that A is empty. Obviously $\varnothing \notin A$, and so $g(\varnothing) \notin A$, otherwise it would be a minimal (hence least) element of A. Likewise $g(g(\varnothing)) \notin A$, etc., and the union of the chain $\{\varnothing, g(\varnothing), g(g(\varnothing)) \ldots\}$ is not in A, etc. etc. To formalize this, let L be the set of all elements of M that are strictly less than (i.e., proper subsets of) all elements of A. Using the assumption that A has no least element, we conclude that $\varnothing \in L$, and that L is closed under g and under the union of chains. So $L = M$ and A is empty. This is the argument of Theorem 3.1 (p.56).

(6) A simple idea lies behind (d): for any linearly ordered set $(S, \leq)$, the 1–1 correspondence $a \leftrightarrow L_{\leq}(a)$ is order preserving (between $\leq$ and $\subseteq$). The range of this correspondence is all the initial segments of S that contain their endpoints. Our current situation: we have the set of initial segments (namely M), well-ordered under $\subseteq$, but we have yet to define $\leq$. But we can see how that has to be done: when $(S, \leq)$ is a linear ordering, $L_{\leq}(a)$ is the smallest initial segment containing a, and a is the largest element of $L_{\leq}(a)$. So define (as in the proof of Theorem 4.4, p.58) $F(a)$ to be the least element of M that contains a. Note that it is crucial that $g(x)$ adds only *one* element to x, otherwise we get a many-one correspondence and we obtain only a partial ordering on S. (This ordering is however **well founded**, i.e., every nonempty subset has a minimal element.)

The treatment in §§4.2–4.4 comes from disassembling this chain of argument, generalizing the hell out of the individual links, and then scrambling the order. IMO, Pollard's adjective "exasperating" applies. The presentation is justified only to the extent that the generalized links find later applications. Other than two instances in this chapter (Theorems 5.3, p.61, and Theorem 6.2, p.64), I've found none.

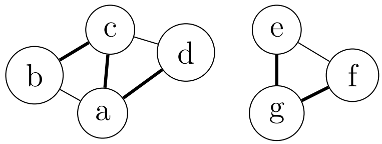

Figure 1: Spanning Forest (Thick Edges)

6 Zorn's Lemma

In pretty much every other math book I've seen, Zorn's Lemma is stated this way: if $(S, \leq)$ is a partial ordering in which every chain (linearly ordered subset) has an upper bound (that is not necessarily an element of the chain), then S has a maximal element. The apparently stronger statement, that for all $x \in S$ there is a maximal element $\geq x$, is a trivial consequence (just look at the partial order of all elements $\geq x$). Very often $\leq$ is $\subseteq$, but sometimes not (or at least not naturally).

Moore's book [19, §§3.4,4,4] details the history of maximal principles. He remarks, "The history of maximal principles, such as Zorn's Lemma, is strewn with multiple independent discoveries of fundamentally similar propositions." Hausdorff gets priority with papers in 1909 and 1914, but these were largely overlooked; his 1909 formulation closely resembles Zorn's lemma. Kuratowski in 1922 and Zorn in 1935 rediscovered maximal principles in slightly different forms.

The simplest standard algebraic example: every vector space has a Hamel basis, i.e., a linearly independent subset such that every element is a finite linear combination of elements from the basis. (Note that this is different from the typical Hilbert space basis, which allows convergent series instead of just finite linear combinations.)

A example from combinatorics: every undirected graph without hermits contains a spanning forest (see fig.1). Definitions: an undirected graph is a set of vertices V plus a subset E of the set of all unordered pairs of vertices, pictured as edges naturally. A hermit is a vertex untouched by any edge, i.e., $v \notin e$ for all $e \in E$. A forest is a graph containing no cycles. A spanning forest touches every vertex, i.e., the union of all the unordered pairs equals the set of all vertices: $\bigcup E = V$. When we say graph (V, E) contains (V', E'), we mean just that $V \supseteq V'$ and $E \supseteq E'$.

I'll use these examples to illustrate two remarks. First, proofs using the well-ordering theorem extend older inductive arguments in a natural way. Second, choices are still there, though they may be implicit.

Hamel Basis. The finite version: every finite-dimensional vector space V has a basis. Proof: choose a vector $e_1 \neq 0$. If $\{e_1\}$ doesn't generate V, then there is an e_2 such that $\{e_1, e_2\}$ is independent. Inductively, if $\{e_1, \ldots, e_k\}$ is independent, then either it generates V or there exists an e_{k+1} such that $\{e_1, \ldots, e_{k+1}\}$ is independent. The finite-dimensionality assumption means this can't go on forever, so eventually we get a basis.

Now drop the finite-dimensionality assumption. Modify the proof as follows: let $\{v_\alpha : \alpha < \beta\}$ be a well-ordering of V. We define a set B_κ of independent vectors for each $\kappa \leq \beta$ by transfinite induction, with B_κ chosen from the set $\{v_\alpha : \alpha < \kappa\}$. For a successor ordinal $\kappa + 1$, look at $B_\kappa \cup \{v_\kappa\}$. If this is independent, set $B_{\kappa+1} = B_\kappa \cup \{v_\kappa\}$. Otherwise set $B_{\kappa+1} = B_\kappa$. For a limit ordinal λ, set $B_\lambda = \bigcup_{\alpha<\lambda}\{B_\alpha\}$. Obviously B_κ is independent for all κ. For any vector $v_\kappa \in V$, either v_κ was added to B_κ to get $B_{\kappa+1}$, or else B_κ already generated v_κ, which is why v_κ wasn't added. So B_β generates V.

Here's the Zorn's lemma proof. Let $\mathcal{B}$ be the collection of all independent subsets of V. $(\mathcal{B}, \subseteq)$ is a partial order. If $\mathcal{C} \subseteq \mathcal{B}$ is a chain, then it's easy to show that $\bigcup \mathcal{C}$ is independent, so $(\mathcal{B}, \subseteq)$ satisfies the hypotheses of Zorn's lemma. Let B be a maximal element of $\mathcal{B}$. Then B generates V, for if $v \in V$ were not generated by B, then $B \cup \{v\}$ would be a larger independent set.

Note the modest part played by algebra in both proofs: essentially just the fact that for an independent B and $v \notin B$, $B \cup \{v\}$ is independent if and only if B does not generate v.

It's clear how the well-ordering proof grows out of the finite-dimensional proof. Next, consider that bases exist in abundance. Most trivial example: we can include e_1 or $2e_1$ in a basis, not both. In the well-ordering proof, we make the choice when we pick the well-ordering. In the Zorn's lemma proof, we choose a particular maximal element of $\mathcal{B}$.

Structurally, the spanning forest results look much the same despite the different mathematical fabric. Finite result: every finite graph without hermits has a spanning forest. Proof by induction: order the edges $\{e_1, \ldots, e_n\}$. Inductively define a set of edges F_k by setting $F_0 = \varnothing$, and at step k add in edge e_k if and only if this does not result in a cycle: $F_{k+1} = F_k \cup \{e_k\}$ if no cycle, otherwise $F_{k+1} = F_k$. So F_k remains a forest at each step. F_{n+1} must span the graph because for any vertex v and edge e_k touching v, if e_k isn't in F_{k+1}, it's because e_k was part of a cycle all of whose other edges belong to F_k. In other words, v is touched by some edge belonging to F_k.

Extend to an arbitrary graph without hermits: well-order the edges $\{e_\alpha : \alpha < \beta\}$. Define forests F_κ for each $\kappa \leq \beta$ as follows. Set $F_{\kappa+1} = F_\kappa \cup \{e_\kappa\}$ if this does not produce a cycle, otherwise $F_{\kappa+1} = F_\kappa$. For a limit ordinal λ, $F_\lambda = \bigcup_{\alpha<\lambda}\{F_\alpha\}$. F_β spans the graph by exactly the same argument as in the finite case.

For the Zorn's lemma argument, we let $\mathcal{F}$ be the collection of all forests contained in the graph. It is partially ordered under $\subseteq$, and a maximal element spans the graph, of course.

There are typically many spanning forests; for example, in fig.1, we could have added edge ab instead of ac or bc. The choice of well-ordering, or the choice of maximal element of $\mathcal{F}$, determines which spanning forest we end up with.

Contrast this with another result in graph theory: every vertex belongs to a maximum connected set of vertices, called its component. Proof: Let $C_0(v) = \{v\}$; let C_{n+1} be C_n plus all vertices joined to C_n by an edge; let $C(v) = \bigcup_{n<\omega}\{C_n\}$. No choices involved—the proof doesn't need AC. Characteristically, $C(v)$ is a *maximum* and not just *maximal*.

7 The Axiom of Determinacy

Just *denying* the axiom of choice doesn't buy you much. If you're going to throw away AC, you should add some powerful incompatible axiom in its place. The Axiom of Determinacy (AD) has been studied in this light.

Here's one formulation. Let S be $\mathbb{N}^{\mathbb{N}}$, i.e., the set of all infinite strings of natural numbers. Let $G \subseteq S$. Alice and Bob play a game where at step $2n$, Alice chooses a number s_{2n}, and at step $2n+1$, Bob chooses a number s_{2n+1}. If $s_0 s_1 s_2 \ldots \in G$, Alice wins, otherwise Bob wins. We say elements of G are *assigned to Alice*, and elements not in G are *assigned to Bob*. We'll call the infinite strings *results* (of the game). Rather than think of G as a set of results, think of it as a function $G : S \to \{\text{Alice}, \text{Bob}\}$.

A *strategy* for Alice tells her how to play each move. Formally, it's a function from the set of all number strings of finite even length to $\mathbb{N}$. Likewise, a strategy for Bob maps number strings of finite odd length to numbers. A game is *determined* if Alice or Bob has a winning strategy, i.e., if the player follows the strategy then that player will win. The Axiom of Determinacy says that each game is determined.

Interesting thing about the proof that AC $\Rightarrow \neg$AD: it's much easier using the well-ordering theorem instead of Zorn's lemma.

First note that there are $c = \aleph_0^{\aleph_0}$ strategies (lumping together both Alice and Bob strategies), likewise c results. Assuming AC, well-order the strategies $\{S_\alpha : \alpha < \omega_c\}$. Here ω_c is the least ordinal with cardinality c, so the set

$\{\alpha : \alpha < \kappa\}$ has cardinality less than c for each $\kappa < \omega_c$.

We construct a game G by inducting transfinitely through all the strategies, at step κ considering S_κ. Our goal is to assign some result to Alice or Bob that prevents S_κ from being a winning strategy. Say S_κ is an Alice strategy. Since we assign only one result at each step, fewer than c results have been assigned before step κ. However, there are c possible results if Alice follows S_κ, since Bob can play his numbers however he wants. So there exists a result where Alice follows S_κ but this result has not yet been assigned to either player. Assign it to Bob; this thwarts S_κ. If S_κ is a Bob strategy, just switch everything around. QED

The cardinality argument at the heart of this proof is harder to pull off with Zorn's lemma (though possible, of course). (The exact same argument works with bit strings instead of strings of natural numbers, but for some reason AD is generally stated using $\mathbb{N}^{\mathbb{N}}$ instead of $\mathcal{P}(\mathbb{N})$.)

8 Erdős's Theorem

In 1962, John Wetzel asked the following question. Say a family F of analytic functions on some common domain D is *pointwise countable* if for each $z \in D$, the set of values $\{f(z) : f \in F\}$ is countable. Does it follow that F is countable?

Erdős [9] answered the question very soon afterwards: pointwise countability implies countability if and only if the continuum hypothesis is *false*. It is so short and beautifully written that I won't repeat the proof, but I will say a couple of things about the key ideas. (It's also discussed in [1, Ch.19, pp.137–138].)

If $c > \aleph_1$, then we look at a family F of $\aleph_1$ functions. Suppose that for some z, all the $f(z)$ were distinct, as f ranges over F. That would prevent F from being pointwise countable. But since there are more $z \in D$ than $f \in F$, this is easy to achieve (details left as an exercise). Thus assuming ¬CH, pointwise countability fails for all uncountable families F. I.e., uncountability implies not pointwise countable. So in one direction, Erdős's theorem is a straightforward cardinality argument. (Incidentally, virtually the same argument shows that "pointwise finite" (obvious definition) implies that F is finite. No mention of CH or ¬CH needed.)

Next suppose $c = \aleph_1$. The goal is to construct an uncountable family F that is pointwise countable. We now have just as many $z \in D$ as $f \in F$. To exploit this, we well-order both D and F: $D = \{z_\alpha : \alpha < \omega_1\}$, $F = \{f_\beta : \beta < \omega_1\}$. Of course, F is constructed by transfinite induction up to ω_1. Here a crucial feature of ω_1 proves invaluable: all smaller ordinals are countable. When we are trying to define f_β, we have only countably many f_γ with $\gamma < \beta$ to worry about. Also, look at the set of values $\{f_\beta(z_\alpha) : \alpha, \beta < \omega_1\}$. If we fix α and restrict β to $\beta < \alpha$, we get a countable set; likewise if we fix β and restrict α to $\alpha < \beta$. With these ingredients, we can (a) insure that $\{f_\beta(z_\alpha) : \beta < \omega_1\}$ is countable for every z_α; (b) define f_β by a convergent series involving the z_α with $\alpha < \beta$ (countability helps a lot with convergent series); and (c) guarantee that all the f_β are distinct.

9 Other Remarks on Chapter 4

§4.7, Cowen's Theorem: I don't see the point of this result. Every other treatment I've seen has no difficulty defining *On*, the class of ordinals, with a formula of ZF, and proving everything you need to know about it. NBG is a conservative extension of ZF, so you get nothing extra by showing that *On* is a bona fide class. (Though I bet it wouldn't be hard to do anyway with the usual ZF formula.) Pollard hints as much in his review [20, §3.4]. Elsewhere [24] Smullyan makes the point that Cowen's theorem is proved without using Replacement (aka Substitution), but he never explains why

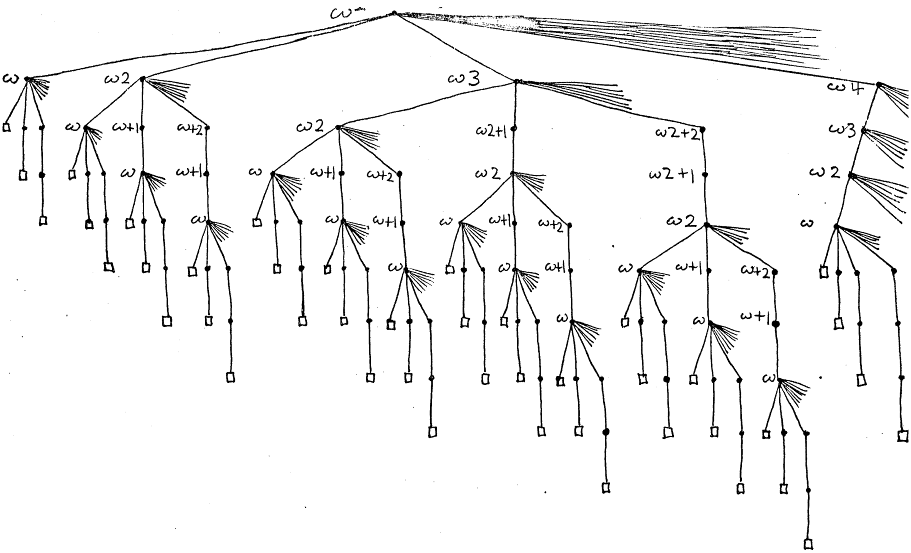

Figure 2: The Ordinal Tree for ω^2

one should care.

Fact: Robert Cowen, like Melvin Fitting, was Smullyan's Ph.D. student. Make of that what you will.

10 Chapter 5, Ordinals

This is a fairly short and easy chapter, given the results of Ch.4. S&F give a definition based on their notion of a superinductive set. von Neumann's original definition:

> α is an ordinal if α is transitive and well-ordered under $\in$.

I've checked several books (Cohen, Drake, Jech, and Kunen); they all use von Neumann's definition. (If you assume the axiom of foundation, aka regularity, then it's enough to assume α is linearly ordered. Drake does this.) This definition is equivalent to S&F's. All these treatments have the same flavor: throw transfinite induction at everything. Judged for clarity, motivation, and complexity, I find all the treatments (including that of S&F) pretty much the same (except for wading through the proof of Cowan's theorem).

One way to picture ordinals is via an "ordinal tree", as illustrated in fig.10. (Ordinal trees also have theoretical significance.)

The idea that the reals can be well-ordered is somewhat hard to stomach; it generated most of the initial skepticism about the axiom of choice. It seems impossible to picture a well-ordering of the reals. We can make this more precise.

Definition: say an ordinal α is **realizable in** $\mathbb{R}$ if there is a subset $A \subset \mathbb{R}$ whose order type, under the natural ordering of $\mathbb{R}$, is α. Examples:

1. ω is realized by $\mathbb{N} \subset \mathbb{R}$.
2. $\omega + 1$ is realized by $\{0, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}, \ldots\} \cup \{1\}$.
3. $\omega \cdot \omega$ is realized by

$$\begin{aligned} &\{0, \tfrac{1}{2}, \tfrac{3}{4}, \tfrac{7}{8}, \ldots\} \\ \cup &\{1, 1\tfrac{1}{2}, 1\tfrac{3}{4}, 1\tfrac{7}{8}, \ldots\} \\ \cup &\{2, 2\tfrac{1}{2}, 2\tfrac{3}{4}, 2\tfrac{7}{8}, \ldots\} \\ &\vdots \end{aligned}$$

4. $\omega \cdot \omega + 1$ is realized by

$$
\begin{aligned}
&\{0, \tfrac{1}{2}, \tfrac{3}{4}, \tfrac{7}{8}, \ldots\} \\
\cup\, &\{1, 1\tfrac{1}{4}, 1\tfrac{3}{8}, 1\tfrac{7}{16}, \ldots\} \\
\cup\, &\{1\tfrac{1}{2}, 1\tfrac{5}{8}, 1\tfrac{11}{16}, \ldots\} \\
\cup\, &\{1\tfrac{3}{4}, \ldots\} \\
\cup\, &\{1\tfrac{7}{8}, \ldots\} \\
&\;\vdots \\
\cup\, &\{2\}
\end{aligned}
$$

Etc. Of course, you can realize each of these ordinals in $\mathbb{R}$ in infinitely many ways. The basic trick: there is an order preserving bijection from $\mathbb{R}$ onto the open unit interval. If you "use up" all of the space from 0 to ∞ with an ordinal α, just squeeze it down so it fits into the unit interval and you then have room on the right to add more elements. I think of "realizable in $\mathbb{R}$" as a precise version of "can be pictured".

The following is a theorem: An ordinal is realizable in $\mathbb{R}$ iff it is countable. (I.e., α is realizable in $\mathbb{R}$ iff the set of ordinals from 0 to α is countable.) So $\aleph_1$, the first uncountable ordinal, is the first ordinal that "cannot be pictured". The proof of this theorem uses the basic trick just mentioned, plus the fact that countable union of countable sets is countable, plus the fact that if $\alpha < \aleph_1$, then (by definition) α is countable.

David Madore has a cool applet that can "draw" many ordinals [18]. John Baez wrote three breezy posts about countable ordinals [2].

Several countable ordinals have special properties. I will mention two: ε_0, and ω_1^{CK}.

Definition of ε_0: it's the limit of ω, ω^ω, ω^{ω^ω},... Gentzen showed that ε_0 is the so-called "proof strength" of Peano arithmetic (PA). Idea is that you can encode ordinals into $\mathbb{N}$; for example, if we define $\prec$ by (a) all even numbers are $\prec$ all odd numbers, and (b) otherwise $\prec$ is the same as $<$, then you can prove in PA that $\prec$ is a well-ordering. This encodes $\omega+\omega$. Gentzen showed that for any $\alpha < \varepsilon_0$, you can define a well-ordering in PA isomorphic to α, and prove it *is* a well-ordering. So you can "do transfinite induction in PA up to α". But Gentzen also proved that PA is consistent, using transfinite induction up to ε_0. By Gödel's second incompleteness theorem, it follows that transfinite induction up to ε_0 cannot be done in PA.

Church and Kleene looked at countable ordinals that can encoded in $\mathbb{N}$ using recursive functions. These are called *constructive* ordinals. Without going into details, the idea is that you can encode all the ordinals less than a constructive ordinal "in a computable way". Since there are only countably many partial recursive functions, it's not surprising that the set of constructive ordinals is countable. Hence the set of constructive ordinals is $\{\alpha < \omega_1^{\text{CK}}\}$ for some countable ordinal ω_1^{CK}. The CK stands for Church-Kleene. Mind you, the "ordinal encoding" must be computable for a constructive ordinal, but not necessarily *provably* computable; that's why ω_1^{CK} is much bigger than ε_0.

Turning to uncountable ordinals: Gödel gave an explicit definition for a well-ordering of the so-called *constructible* real numbers. (This definition, like the definition of *constructible*, is part of his relative consistency proof for AC.) Also Gödel showed that it is consistent with ZF to assume that all real numbers are constructible. On the other hand, using forcing it can be shown that it is consistent with ZFC to assume that there is no definable well-ordering of $\mathbb{R}$. (See §21.7.5.)

11 Chapters 6–8

For so-called "ordinary mathematics", you usually don't need to go any higher than rank $\omega\cdot 2$. Example: the complex Hilbert space $L^2(\mathbb{R}^3)$, beloved by analysts and physicists alike. We build up to this by contructing all the natural numbers, then the integers, rational numbers, and reals, ordered pairs and ordered triples of reals, then functions from the ordered triples to ordered pairs, and finally equivalence classes of these functions. (Recall that two L^2 functions are identified if they agree except for a set of measure zero; hence, the actual elements of $L^2(\mathbb{R}^3)$ are equivalence classes.) We have all the natural numbers by the time we get to R_ω. If $x, y \in R_\alpha$, then the ordered pair $(x, y) \in R_{\alpha+2}$. So if $A \in R_\alpha$, then any subset of $A \times A$ is in $R_{\alpha+3}$. Integers are often defined as sets of ordered pairs of natural numbers, rational numbers as sets of ordered pairs of integers, and real numbers as sets of rational numbers (Dedekind cuts). If I counted right, that puts all the elements of $\mathbb{R}$ in $R_{\omega+7}$ and any subset of $\mathbb{R}$ in $R_{\omega+8}$. The Hilbert space in question should belong to $R_{\omega+15}$, again if I haven't miscounted.

Although sets of high rank aren't "needed" by most mathematicians, it would be quite strange to impose a "rank ceiling" limitation on the power set axiom. Just maybe if $R_{\omega+1}$ were adequate—but it isn't.

Axiom E, §8.4 (p.103), is usually called "global choice" or "strong choice".

12 Cardinal Arithmetic

I don't have my usual complaints about the two proofs S&F give for $m^2 = m$ (for infinite m). But I think the proofs can be made even shorter and sweeter. I am inspired partly by Kaplansky's treatment [16].

First proof: let $|M| = m$. Assume M is infinite. Let $\mathcal{B}$ be the set of all bijections $f : U \to U \times U$ for $U \subseteq M$; we define $f \leq g$ if g extends f (with both f and g in $\mathcal{B}$). We know that $\mathcal{B}$ is not empty because M contains a countable set. Then $(\mathcal{B}, \leq)$ is a partial order satisfying all the requirements for Zorn's lemma, so let $f : U \to U \times U$ be a maximal element. Let $M = U \sqcup V$ (disjoint union) and let $u = |U|$, $v = |V|$. So $m = u + v$, and $u^2 = u$.

Case 1: $u \geq v$. Then $m = u + v = u$, and since $u^2 = u$, we have $m^2 = m$.

Case 2: $u \leq v$. (The cases overlap, but so what.) So V contains a set W of cardinality u. Now let's look at $U \sqcup W$, or more particularly its cross product with itself:

$$(U \sqcup W) \times (U \sqcup W) = (U \times U) \sqcup (U \times W) \sqcup (W \times U) \sqcup (W \times W)$$

All four terms on the right hand side have cardinality u^2, so the union of the last three terms has cardinality $3u^2 = u$ (because $u^2 = u$). So there is a bijection between W and that union. Combining this bijection with f gives a proper extension of f (to the domain $U \sqcup W$). Since f is maximal, this case can't happen.

Second proof: here S&F give a standard "rectangular" ordering of $On \times On$, giving rise to an order isomorphism they call P. Let m be a cardinal, considered as an ordinal. Since m is also the set of all previous ordinals, to show that $m^2 = m$ it's enough to show that the "inner square" (i.e., $M = \{(\alpha, \beta) : \alpha < m \text{ and } \beta < m\}$ maps onto the set of ordinals preceding m. Assume inductively that $d^2 = d$ for all cardinals $d < m$.

First we show that every ordinal $\gamma < m$ is the image of some $(\alpha, \beta) \in M$. This is easy: the inner square is the initial segment of all $(\alpha, \beta) < (0, m)$, and since order isomorphisms map initial segments to initial segments, if γ is omitted it follows that the image of the inner square is mapped into the set of predecessors of γ. In other words, the inner square has cardinality less than m. But this implies that $m^2 < m$, which is impossible.

Next we show that the image of the inner square is contained in the set of predecessors of m. (Note that this isn't true for *finite* m.) It's clearly enough to show that this holds for each *closed* square $\{(\alpha, \beta) : \alpha \leq \delta \text{ and } \beta \leq \delta\}$ for all $\delta < m$. Let $|\delta| = d$. Because m is a cardinal, $d < m$. The cardinality of the closed square is $d^2 + d + d$ (the inner square plus the two sides), which equals d by inductive hypothesis. So, by the same fact about the image of initial segments, the image of the closed square must be contained entirely

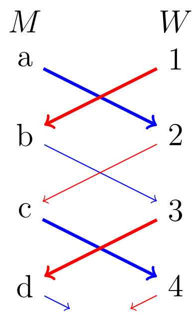

Figure 3: Schröder-Bernstein

in the set of predecessors of m (if not, we'd have $d \geq m$).

Minor variation on the last paragraph: the closed square for δ is the same as the inner square for $\delta + 1$, and if $\delta < m$, so is $\delta + 1$.

13 The Schröder-Bernstein Theorem

Here is some motivation for the first proof (p.108). In figure 3, sets M and W are indicated with the injections $f : M \to W$ (blue) and $g : W \to M$ (red). We want to construct a 1–1 pairing between all of M and W out of these two partial pairing. Each edge, red or blue, is a possible pairing. So a node like b has two possible partners: 1 or 3. However, a node like a has only one possible partner, namely 2. Likewise, 1's only possible partner is b.

That means we *have* to start by pairing off nodes without preimages: nodes $x \in M$ for which $g^{-1}(x)$ doesn't exist, and nodes $y \in W$ for which $f^{-1}(y)$ doesn't exist. Let's call this the *first round* of pairings.

Once we've paired off $a2$ and $1b$, it then turns out that c and 3 each have only one possible choice (second round). And so on. In fig.3, this is enough to complete the bijection. (Thickened edges indicate the pairings actually made.)

If you think about it, this procedure amounts to following the f and g edges backwards until you can't go any further. So we have two subsets of M and of W: if the backwards chain stops in M, the node is in M_1 or W_1; if it stops in W, the node is in M_2 or W_2. We use f to pair M_1 with W_1, and g to pair W_2 with M_2. These are rules (1) and (2) from p.109.

If you now ask who's left out, you end up with the sets M_3 and W_3: nodes belonging to even-length cycles or to doubly-infinite chains. Pair these off using f throughout—or using g throughout, but you have to use the same function for all the pairs.

Now for the second proof. First we deconstruct the Knaster-Tarski fixed point theorem. From the monotonicity of h, we see that

$$\varnothing \subseteq h(\varnothing) \subseteq h(h(\varnothing)) \subseteq \ldots$$

Let

$$B_\omega = \varnothing \cup h(\varnothing) \cup h(h(\varnothing)) \cup \ldots$$

It's not hard to see that $B_\omega \subseteq h(B_\omega)$, and so

$$B_\omega \subseteq h(B_\omega) \subseteq h(h(B_\omega)) \subseteq \ldots$$

Let $B_{\omega 2} = B_\omega \cup h(B_\omega) \cup h(h(B_\omega)) \cup \ldots$. And so on. Because ordinals go on strictly increasing forever but the $B's$ can't, being stuck inside a set, eventually we have a B_α with $h(B_\alpha) = B_\alpha$. This proof is longer and less elegant than the proof given in S&F, but it makes the application to the Schröder-Bernstein theorem less mysterious.

For some reason, S&F switch from M and N to A and B in their second proof. Anyway, let's look at $h(\varnothing) = A \smallsetminus g''(B)$. This is just the set of those nodes in A without a preimage in B: the "buck stops here" nodes in A

in the backwards chain! These nodes get paired in the first round. What about $h(h(\varnothing))$? Well,

$$h(h(\varnothing)) = A \smallsetminus g''(B \smallsetminus f''(h(\varnothing)))$$
$$f''(h(\varnothing)) = \text{nodes in } B \text{ paired in first round}$$

After removing all the first-round pairings and their edges, the nodes in $h(h(\varnothing))$ become the "buck stops here" nodes in A. (In fig.3, $h(\varnothing) = \{a\}$, and $h(h(\varnothing)) = \{a, c\}$.) The iterations of h track the successive rounds of pairings; after ω rounds, the union $h(\varnothing) \cup h(h(\varnothing)) \cup \ldots$ contains all the nodes in A that *have* to be paired using f, namely A_1. The second proof then uses) g for all the remaining pairings.

14 Sierpiński's Theorem: GCH implies AC

I had a look at the version of the proof in Cohen [4, §IV.12]. Sierpiński was a clever fellow, and he came up with a few tricks that would be hard to motivate. But S&F's habit of dicing proofs into little lemmas, presented in a seemingly random order, adds a layer of fog over the train of thought.

Here I will try to imagine how Sierpiński could have devised his proof. Cohen does offer one bit of intuition:

> The GCH is a rather strong assertion about the existence of various maps since if we are ever given that $A \leq B \leq P(A)$ then there must be a 1–1 map either from B onto A or from B onto $P(A)$. Essentially this means that there are so many maps available that we can well-order every set.

Let A be the set we wish to well-order. In our notation, GCH tells us that for any U,

$$A \preceq U \preceq P(A) \Rightarrow U \cong A \text{ or } U \cong \mathcal{P}(A)$$

If U is well-ordered, then $U \cong A$ and $U \cong \mathcal{P}(A)$ both imply that A can be well-ordered, the latter because A is naturally imbedded in $\mathcal{P}(A)$. But this is too simple an approach: $A \preceq U$ already makes A well-ordered for a well-ordered U, so if we could show the antecedent we'd be done—we wouldn't need GCH to finish the job.

Let's not assume U is well-ordered, but instead suppose it contains a well-ordered set. Say we could show that

$$A \preceq W + A \preceq \mathcal{P}(A)$$

for a well-ordered set W, where '+' stands for disjoint union. (That is, $W \times \{0\} \cup A \times \{1\}$, or some similar trick to insure disjointness.) Then we'd have

$$W + A \cong \mathcal{P}(A) \text{ or } W + A \cong A$$

Now, if $W + A \cong \mathcal{P}(A)$, then we *ought* to have $W \cong \mathcal{P}(A)$, just because A is "smaller" than $\mathcal{P}(A)$ (in some sense)—W should just absorb A, if W is "big enough". Also, if W is "big enough" then that should exclude the other arm of the choice, where $W + A \cong A$. And if $W \cong \mathcal{P}(A)$, then $\mathcal{P}(A)$ and so also A can be well-ordered, as we have seen.

At this point Hartog's theorem shows up at the door. This gives us a well-ordered set W with

$$W \preceq \mathcal{P}^4(A) \text{ and } W \not\preceq A$$

So we have

$$A \preceq W + A \preceq \mathcal{P}^4(A) + A \overset{?}{\cong} \mathcal{P}^4(A)$$

$W \not\preceq A$ excludes $W + A \cong A$, good. Let's postpone the issue of the '?'. Deal first with the problem that the bounds are not tight enough for GCH to apply. We fix that by looking at:

$$\mathcal{P}^3(A) \preceq W + \mathcal{P}^3(A) \preceq \mathcal{P}^4(A) + \mathcal{P}^3(A) \overset{?}{\cong} \mathcal{P}^4(A)$$

So if $W + \mathcal{P}^3(A) \cong \mathcal{P}^4(A)$, we *ought* to have $W \cong \mathcal{P}^4(A)$ and hence a well-ordering of A. What about the other case, $W + \mathcal{P}^3(A) \cong \mathcal{P}^3(A)$? Ah, then we have

$$\mathcal{P}^2(A) \preceq W + \mathcal{P}^2(A) \preceq W + \mathcal{P}^3(A) \cong \mathcal{P}^3(A)$$

and so we can repeat the argument: either $W+\mathcal{P}^2(A) \cong \mathcal{P}^3(A)$, which ought to make $\mathcal{P}^3(A)$ well-ordered and hence also A well-ordered; or $W+\mathcal{P}^2(A) \cong \mathcal{P}^2(A)$, in which case we repeat the argument yet again. Eventually we work our way down to

$$A \preceq W + A \preceq W + \mathcal{P}(A) \cong \mathcal{P}(A)$$

and $W + A \cong A$ is excluded since $W \not\preceq A$, and we are done.

All this relies on the intuition that if $W + M \cong \mathcal{P}(M)$, then we should have $W \cong \mathcal{P}(M)$: we used this with $M = \mathcal{P}^n(A)$ for $n = 0, \ldots, 3$. Well, we can prove something a little weaker.

Lemma: If $W + M \cong \mathcal{P}(M) \times \mathcal{P}(M)$, then $W \succeq \mathcal{P}(M)$.

Proof: Suppose $h : W + M \to \mathcal{P}(M) \times \mathcal{P}(M)$ is a bijection. Restrict h to M and compose with the projection to the second factor: $\pi_2 \circ (h{\restriction}M) : M \to \mathcal{P}(M)$. Cantor's diagonal argument shows that this map cannot be onto. (The fact that $\pi_2 \circ (h{\restriction}M)$ might not be 1–1 doesn't affect the argument.) So for some $s_0 \in \mathcal{P}(M)$, we know that $h(x)$ *never* takes the form $(-, s_0)$ for $x \in M$. In other words, the image of $h{\restriction}W$ must include all of $\mathcal{P}(M) \times \{s_0\}$. Therefore $\mathcal{P}(M) \times \{s_0\}$ can be mapped 1–1 to a subset of W. qed.

The missing pieces of the proof now all take the form of absorption equations. We know that $\mathcal{P}(M) \times \mathcal{P}(M) \cong \mathcal{P}(M + M)$—as an equation for cardinals, $2^{\mathfrak{m}}2^{\mathfrak{m}} = 2^{2\mathfrak{m}}$. If we had $2\mathfrak{m} = \mathfrak{m}$, that would take care of that problem. The $\overset{?}{\cong}$ above also takes the form $2^{\mathfrak{m}} + \mathfrak{m} \overset{?}{=} 2^{\mathfrak{m}}$, for $\mathfrak{m}$ the cardinality of $\mathcal{P}^3(A)$.

The general absorption laws for addition depend on AC. But we do have these suggestive equations even without AC:

$$\mathfrak{a} + \omega + 1 = \mathfrak{a} + \omega, \quad 2^{\mathfrak{m}+1} = 2 \cdot 2^{\mathfrak{m}}$$

and so if $\mathfrak{a} + \omega = \mathfrak{m}$ and $2^{\mathfrak{m}} = \mathfrak{b}$, then $2\mathfrak{b} = \mathfrak{b}$. So let's say we set $B = \mathcal{P}(A + \omega)$. Then we have $2B \cong \mathcal{P}(A + \omega + 1) \cong B$ (where $2B$, of course, is the disjoint union of B with itself). Also $B \preceq B + 1 \preceq 2B \cong B$, so $B \cong B + 1$ and so $\mathcal{P}(B) \cong 2\mathcal{P}(B)$. So if we replace A with B, then all gaps in the argument are filled and we conclude that B can be well-ordered. But obviously A can be imbedded in B, so A also can be well-ordered. QED.

Ernst Specker proved a "local" version of Sierpiński's Theorem: if $\mathfrak{m}$ and $2^{\mathfrak{m}}$ both satisfy CH, then $2^{\mathfrak{m}} = \aleph(\mathfrak{m})$.

15 Roadmap for Part 2

Part 2 (Consistency of the Continuum Hypothesis) rests on five foundation stones:

First-order logic: Concepts such as sentence, satisfaction, model, and definability all play key roles. S&F develop this material in Ch.11.

Constructible Sets: Gödel's universe of constructible sets is the centerpiece of the whole argument; it revolves around the notion of definability. Ch.12 introduces the constructible universe.

Absoluteness: Certain concepts of set theory are absolute, in that they do not depend on the surrounding model. Absoluteness is covered in Ch.14.

Reflection Principles: At two pivotal points, Gödel made use of a version of the Löwenheim-Skolem Theorem. The theorem says that we can always find a "small" submodel that "reflects" certain aspects of a larger model. S&F put this in Ch.11.

Mostowski Collapsing Lemma: At a key point in the argument, Gödel needed to remove "superfluous" sets from a model; he applied what is generally known as the Mostowski collapsing lemma. (In Cohen's book [4, p.96], he calls it "the trivial result... concerning $\in$-isomorphisms".)

The following sections give the argument in broad strokes.

16 The Constructible Universe

The constructible universe is traditionally denoted L. L is a subclass of V and is a proper class. Gödel proved three things about L:

1. All the axioms of ZF hold in L, i.e., L is a model of ZF.
2. $V = L$ holds in L, i.e., L is a model of the axiom "All sets are constructible". To quote Cohen: "This is a small but subtle point. It says that a constructible set is constructible when the whole construction is relativized to L."
3. $V = L \Rightarrow$ AC and $V = L \Rightarrow$ GCH are both provable in ZF.

So we can't prove not-GCH in ZF: if we could, it would have to hold in L, but GCH holds in L. Ditto for AC.

L is constructed according to the familiar transfinite scheme, using a function $\mathcal{F}$ (discussed below):

$$\begin{aligned} L_0 &= \varnothing \\ L_{\alpha+1} &= \mathcal{F}(L_\alpha) \\ L_\lambda &= \bigcup_{\alpha<\lambda} L_\alpha \\ L &= \bigcup_{\alpha\in On} L_\alpha \end{aligned}$$

Let A be a set. $\mathcal{F}(A)$ is a subset of $\mathcal{P}(A)$; it's the set of all sets that are *definable* using elements of A. A precise definition of *definable* requires the machinery of first-order logic. I'll just provide a few examples here.

The singleton $\{x\}$, the unordered pair $\{x, y\}$, the ordered pair $\langle x, y\rangle = \{\{x\}, \{x, y\}\}$, and the power set $\mathcal{P}(x)$ are all definable from x or from x and y:

$$\begin{aligned} z \in \{x\} &\Leftrightarrow z = x \\ z \in \{x, y\} &\Leftrightarrow z = x \vee z = y \\ z \in \langle x, y\rangle &\Leftrightarrow z = \{x\} \vee z = \{x, y\} \\ z \in \mathcal{P}(x) &\Leftrightarrow (\forall u)[u \in z \to u \in x] \end{aligned}$$

Each right-hand side is a first-order formula characterizing the elements of a set. For example, instead of writing $z = \{x\}$, we can write $(\forall t)[t \in z \leftrightarrow t = x]$. You should imagine the left-hand sides as abbreviations for the right-hand sides, so (for example) in the formal definition of $\langle x, y\rangle$, "$z = \{x\}$" and "$z = \{x, y\}$" have been expanded.

The prime examples of sets *not* obviously definable are choice functions. For example, it's easy to say what we desire of a choice function c for $\mathcal{P}(\mathbb{R})$, where $\mathbb{R} = \mathcal{P}(\omega)$:

$$(\forall s \subseteq \mathbb{R})[s \neq \varnothing \to c(s) \in s]$$

But this doesn't characterize c. (Here, "$c(s) \in s$" expanded out reads "$(\exists t)[\langle s,t\rangle \in c \wedge t \in s]$"; we'd also need to spell out the fact that c is a function with domain $\mathcal{P}(\mathbb{R}) \smallsetminus \{\varnothing\}$.)

For any set A, $x \subseteq A$ is *definable* over A if there is a first-order formula $\varphi(y, t_1, \ldots, t_n)$ and elements $a_1, \ldots, a_n \in A$ such that

$$z \in x \Leftrightarrow (z \in A) \wedge \varphi^A(z, a_1, \ldots, a_n)$$

where φ^A is φ *relativized* to A, i.e., all quantifiers in φ range only over A. As we said before, $\mathcal{F}(A)$ is the set of all subsets of A that are definable over A.

So $\mathcal{F}(A)$ is something like $\mathcal{P}(A)$, except we include only those sets where we can explicitly describe their criterion for membership. Cohen discusses how this notion arose from, but did not resolve, some philosophical concerns stemming from the paradoxes.

17 Absoluteness

Let's look again at the notion of definability. Rewrite it slightly:

$$x = \{z \in A : \varphi^A(z, a_1, \ldots, a_n)\}$$

Very clearly the right-hand side depends on A. In some cases, we can write a formula for x that is independent of A. For example:

$$x = \{y\} \Leftrightarrow y \in x \wedge (\forall z)[z \in x \rightarrow z = y]$$

Absoluteness is at the heart of Gödel's proof that $V = L$ holds in L. Failures of absoluteness present the main technical obstacles to showing that L satisfies ZF.

17.1 Three Non-absolute Notions

Let's look at three non-absolute notions:

$$x = \mathcal{P}(y)$$
$$z \text{ is uncountable}$$
$$x = \mathcal{P}_{\text{un}}(y) = \{z \subseteq y : z \text{ is uncountable}\}$$

Now think about relativizing these three notions to L and to the L_α's. Suppose x and y are both present in some L_α, and $L_\alpha \models x = \mathcal{P}(y)$. As we ascend to higher L_β's, the assertion "$x = \mathcal{P}(y)$" can become false[2], because new subsets of y can appear in L_β. The reverse switch from false to true can't happen: once we have a "witness" z to $x \neq \mathcal{P}(y)$ (i.e., $z \subseteq y$ but $z \notin x$ or $z \not\subseteq y$ but $z \in x$), it won't go away. (Note that the meaning of x and y can't change, because the L_α's are all transitive: by the time x shows up, all its elements have shown up. Ditto for y.)

Here's a slightly different way to look at it. Consider the sequence $x_\alpha = \mathcal{P}^{L_\alpha}(y)$. The sequence x_α is monotonically increasing (indeed, $\mathcal{P}^{L_\alpha}(y) = \mathcal{P}(y) \cap L_\alpha$). So the assertion $x = x_\alpha$ can flip from true to false as α increases[3]. It can't flip from false to true, because that would mean that x had some elements (subsets of y) that x_α was missing—but in that case x wouldn't have been an element of L_α in the first place.

It's a similar story for uncountability. We have:

$$L_\alpha \models z \text{ is uncountable} \Leftrightarrow L_\alpha \models \neg(\exists f) f{:}\omega \cong z$$

where $f{:}\omega \cong z$ is shorthand for saying that f is a bijection between ω and z. If a witness f to the countability of z appears in some higher L_β, then z "becomes countable", and remains countable after that.

---

[2] Just to be clear: by saying "$x = \mathcal{P}(y)$ becomes false", I mean that although $L_\alpha \models x = \mathcal{P}(y)$, $L_\beta \models x \neq \mathcal{P}(y)$. Here x and y are *fixed* elements of L, which both belong to L_α (and thus also to L_β).

[3] Again, to be clear, by "flip" I mean that $x = x_\alpha$ but $x \neq x_\beta$ for two ordinals $\alpha < \beta$.

Next we look at $x = \mathcal{P}_{\mathrm{un}}(y)$. Let $x_\alpha = \mathcal{P}_{\mathrm{un}}^{L_\alpha}(y)$. New uncountable subsets of y can appear at any time. But also, a subset of y that is uncountable in L_α can become countable later on. The upshot: $\mathcal{P}_{\mathrm{un}}^{L_\alpha}(y)$ can both gain and lose elements as α increases, and the truth-value of "$x = \mathcal{P}_{\mathrm{un}}(y)$" can switch back and forth over and over again.

We will look at the limiting behavior (i.e., $\mathcal{P}^L(y)$ and $\mathcal{P}_{\mathrm{un}}^L(y)$) in §20.

How about absolute formulas? We say that a formula $\varphi(\vec{u})$ is *absolute* if for all $\vec{a}$ in K

$$K \models \varphi(\vec{a}) \quad \Leftrightarrow \quad V \models \varphi(\vec{a})$$
$$(K \text{ a transitive class, } \vec{a} \subseteq K)$$

in other words, the truth-value of $K \models \varphi(\vec{a})$ doesn't depend on K, provided only that K is a transitive class. (To be precise, when we say $K \models \varphi(\vec{a})$, we mean $(K, \in) \models \varphi(\vec{a})$. Also, $\vec{a} \subseteq K$ means that all the a_i belong to K.)

S&F give a very clear treatment of absoluteness (see §12.2, p.158–163.) They give a long list of absolute formulas, but they also relate it to an important syntactic notion. A Δ_0 *formula* is one where all quantifiers are bounded, i.e., of the form $(\forall x \in y)$ or $(\exists x \in y)$. Δ_0 formulas are always absolute. They prove this in Theorem 2.2 (p.160).

Here's the intuition behind Δ_0 formulas. First we define the *transitive closure* of x to consist of x, plus the elements of x, plus the elements of the elements of x, etc. (See p.195.) Any transitive class containing x as an element contains its transitive closure. In fact, it's easy to see that the transitive closure of x is the smallest transitive class containing x as an element, and that the transitive closure is a set. OK: if $\varphi(\vec{a})$ is Δ_0, then to find out if K satisfies $\varphi(\vec{a})$, we only need to root around in the transitive closure of the a_i's. We never need to search through all of K.

Example 1: $x \subseteq y$ is Δ_0, since it can be written $(\forall z \in x) z \in y$. Example 2: $f : \omega \cong x$ is Δ_0. To verify this meticulously takes some work. Fortunately, S&F shoulder most of the burden, in their list of 35 Δ_0 formulas (p.159): item (17) is "x is an ordinal", and item (33) is $f{:}x \cong y$. We can describe $w = \omega$ with this formula:

$$w \text{ is an ordinal} \wedge (\forall y \in w)[y = \varnothing \vee (\exists x \in w)y = x^+]$$

(Item (13) is $y = x^+$.) But it's intuitively clear that $f{:}\omega \cong x$ is Δ_0. To check that f is 1–1 and onto, we just need to dig inside the guts of f and its domain and range. To check the claim that its domain (call it w) is really ω, we just need to crawl around inside w. In neither case do we need to climb outside the transitive closure of f and wander around the entire class K.

In contrast, we cannot check that $K \models x = \mathcal{P}(y)$ or that x is countable in K without surveying all of K, looking for (respectively) subsets of y and bijections f.

With this in mind, we turn to a syntactic analysis of our examples. Strictly speaking, we can follow S&F's treatment without doing this, but I think the effort will pay dividends in intuition.

As noted, $z \subseteq y$ and $f{:}\omega \cong z$ are Δ_0. So $x = \mathcal{P}(y)$ is of the form

$$(\forall z)\delta_1(x, y, z)$$

and "x is uncountable" is of the form

$$(\forall z)\delta_2(x, z)$$

where δ_1 and δ_2 are Δ_0.

Finally, $x = \mathcal{P}_{\text{un}}(y)$ is of the form $(\forall z)(\forall f)(\exists g)\delta_3(x, y, z, f, g)$, where δ_3 is Δ_0. Showing this takes a bit of work. First we break down $x = \mathcal{P}_{\text{un}}(y)$ into three parts:

$$\begin{aligned} & (\forall z)[z \in x \rightarrow z \subseteq y] \\ \wedge\ & (\forall z)[z \in x \rightarrow (\forall f)\neg\, f{:}\omega \cong z] \\ \wedge\ & (\forall z)[[z \subseteq y \ \wedge\ (\forall g)\neg\, g{:}\omega \cong z] \ \rightarrow\ z \in x] \end{aligned}$$

We ask, how could this conjunction fail to be true? This way:

$$\begin{aligned} &(\exists z)[z \in x \wedge z \not\subseteq y] \\ \vee\ &(\exists z)[z \in x \wedge (\exists f) f{:}\omega \cong z] \\ \vee\ &(\exists z)[z \subseteq y \wedge (\forall g)\neg\, g{:}\omega \cong z\ \wedge z \notin x] \end{aligned}$$

Now we use basic logical equivalences to move the quantifiers outwards. Say we have formulas φ, $\psi(u)$, and $\xi(u)$, where u does not appear in φ. Then

$$\begin{aligned} \varphi \wedge (\exists u)\psi(u) &\equiv (\exists u)[\varphi \wedge \psi(u)] \\ \varphi \vee (\exists u)\psi(u) &\equiv (\exists u)[\varphi \vee \psi(u)] \\ \varphi \wedge (\forall u)\psi(u) &\equiv (\forall u)[\varphi \wedge \psi(u)] \\ \varphi \vee (\forall u)\psi(u) &\equiv (\forall u)[\varphi \vee \psi(u)] \\ (\exists u)\psi(u) \vee (\exists u)\xi(u) &\equiv (\exists u)[\psi(u) \vee \xi(u)] \end{aligned}$$

So we can rewrite the failure of $x = \mathcal{P}_{\mathrm{un}}(y)$ as:

$$\begin{aligned} &(\exists z)(\exists f)(\forall g)\big[\\ &\quad (z \in x \wedge z \not\subseteq y) \\ &\vee\ (z \in x \wedge f{:}\omega \cong z) \\ &\vee\ (z \subseteq y \wedge \neg\, g{:}\omega \cong z\ \wedge z \notin x)\big] \end{aligned}$$

The stuff inside the brackets is Δ_0, so negating this gives us the form we claimed[4].

[4]Thinking in terms of witnesses makes this more picturesque. Suppose $x \neq \mathcal{P}_{\mathrm{un}}(y)$. In other words, x is accused of the crime of not being $\mathcal{P}_{\mathrm{un}}(y)$. The prosecution and defence must provide witness lists before the trial starts. The prosecution lists z and f; the defence, all the g's. Any one of the three disjuncts is sufficient to convict; let's imagine a trial lasting three days. The witness z is called each day. On the first day, if z testifies to being an element of x but not a subset of y, game over. But suppose z surprises Jack McCoy (the prosecutor) by being a subset of y. On the second day, f is also called, to testify to the countability of z; McCoy hopes to show that $z \in x$. Too bad for McCoy, f's testimony falls apart. On the third day, z is recalled and is shown to be both a subset of y, and not an element of x after all! The defence tries to argue that's ok, because z is countable. He calls up every single g to testify to being the required bijection, but each g fails. The jury convicts and McCoy repairs to the bar to have a drink with his ADA.

17.2 The Lévy Hierarchy

Summarizing the previous section:

$x = \mathcal{P}(y)$	is of the form	$(\forall u)\delta_1(x, y, u)$
x is uncountable	is of the form	$(\forall u)\delta_2(x, u)$
$x = \mathcal{P}_{\text{un}}(y)$	is of the form	$(\forall u)(\forall v)(\exists w)\delta_3(x, y, u, v, w)$

where δ_1, δ_2, and δ_3 are all Δ_0 formulas. This syntactic analysis fits into a scheme known as the Lévy Hierarchy. S&F cleverly avoid this machinery, but I think it clarifies matters.

Formulas of the form $(\forall u)\delta(\vec{x}, u)$ are called Π_1 or $\forall$ formulas; replace the $\forall$ with an $\exists$, and you've got a Σ_1 formula. (Of course, δ here stands for a Δ_0 formula.) More generally, any string of $\forall$'s is allowed at the front of a Π_1 formula, likewise any string of $\exists$'s at the front of a Σ_1 formula.

The negation of a Π_1 formula is Σ_1, and vice versa. Truth "propagates upwards" for Σ_1 formulas and "propagates downwards" for Π_1 formulas. You'll find a precise statement in §14.1, pp.179–181. The intuition is clear: if $K \models (\exists\vec{u})\delta(\vec{a}, \vec{u})$ for some $\vec{a} \subseteq K$ with K a transitive class, then we have *witnesses*—elements $\vec{c} \subseteq K$ such that $K \models \delta(\vec{a}, \vec{c})$. The witnesses cannot be impeached by enlarging K, because δ is Δ_0. So $(\exists\vec{u})\delta(\vec{a}, \vec{u})$ holds also in any transitive class containing K. Likewise for the downward propagation with Π_1 formulas.

$(\forall\vec{u})(\exists\vec{v})\delta(\vec{x}, \vec{u}, \vec{v})$ is a Π_2 formula, also called an $\forall\exists$ formula; its negation is a Σ_2 or $\exists\forall$ formula. A formula that looks like $(\forall\vec{u})(\exists\vec{v})(\forall\vec{w})\ldots\delta$, where there are n alternating quantifier blocks, is a Π_n formula; starting off with an existential block gives a Σ_n formula.

A formula equivalent to both a Σ_1 and a Π_1 formula will thus be absolute—but that word "equivalent" is the kicker. Equivalent in what sense? One answer: $\varphi(\vec{x})$ is Σ_n^{ZF} if there is a Σ_n formula $\psi(\vec{x})$ such that $\text{ZF} \vdash (\forall\vec{x})(\varphi(\vec{x}) \leftrightarrow \psi(\vec{x}))$; likewise for Π_n^{ZF}. If a formula is both Σ_n^{ZF} and Π_n^{ZF}, we say it's Δ_n^{ZF}.

So Δ_1^{ZF} formulas are absolute *between models of ZF.* (And of course, Σ_1^{ZF} formulas are absolute upwards, Π_1^{ZF} absolute downwards, but only between models of ZF.)

This tradeoff between admitting more formulas or more classes can take a variety of forms. I won't explore the full landscape, but a few aspects should be highlighted; I especially want to contrast the S&F approach with that found in many other books.

First let's look at the role of bounded quantifiers. Note that the S&F definition of a Σ formula (p.180) isn't quite the same as Σ_1. Σ formulas can have bounded quantifiers appearing anywhere, while in Σ_1 formulas they must appear on the inside. $(\forall x \in y)(\exists z)[z \in x]$ is Σ but not Σ_1, for example.

So long as we restrict attention to ZF-models, this distinction doesn't matter: all Σ formulas are Σ_1^{ZF}. Proof: the formula $(\forall x \in y)(\exists z)\varphi(x, y, z)$ is ZF-equivalent to $(\exists u)(\forall x \in y)(\exists z \in u)\varphi(x, y, z)$, by an argument involving ranks[5]. So we can migrate all bounded quantifiers to the inside.

S&F noted that Σ formulas are absolute upwards for *all* transitive classes, not just models of ZF[6]. Accordingly, they work with Σ formulas and transitive classes in Ch.14, even though many of the notions are Δ_1^{ZF}.

At the end of the day (i.e., in part C, pp.189–190), they want full absoluteness "over L", i.e., between L and V. They appeal to the "function-like" formula trick (Theorem 1.5, p.181): any formula that is absolute upwards between transitive classes $K \subseteq M$, and is function-like over both K and M, is in fact absolute between K and M. The key formulas of part C are Σ

[5]For each x in y, let α_x be least ordinal such that there is a z of rank α_x making $\varphi(x, y, z)$ true. Let $\xi = \sup_{x \in y} \alpha_x$, and set $u = V_\xi$. Later on I will call this sort of reasoning a *waiting argument.*

[6]See Prop.1.1(4), p.180. The idea is simple: in quantifhing $(\forall z \in y)\ldots$ with $y \in K$, the bounded quantifier never asks us to go outside K, for a transitive K. So if the assertion holds for K, it will hold for any transitive $M \supseteq K$.

and are function-like for models of ZF. Hence they are absolute between L and V.

Theorem 1.5 has a counterpart for the Lévy hierarchy. Suppose we have a Σ_1 formula

$$(\exists\vec{u})\varphi(\vec{x},y,\vec{u})$$

where $\varphi(\vec{x},y,\vec{u})$ is Δ_0. Suppose also that K is a transitive class, and $(\exists\vec{u})\varphi(\vec{x},y,\vec{u})$ is function-like for K. That is, for any $\vec{c}\subseteq K$, there is a unique d such that $K\models(\exists\vec{u})\varphi(\vec{c},d,\vec{u})$. Then our formula is equivalent to this Π_1 formula over K:

$$(\forall\vec{u})(\forall z)[\varphi(\vec{x},z,\vec{u})\rightarrow y=z]$$

I relegate the proof to a footnote[7]. It resembles the proof of Theorem 1.5.

Now suppose that $(\exists\vec{u})\varphi(\vec{x},y,\vec{u})$ is function-like for both K and M with $K\subseteq M$. Then it is equivalent in both classes to a Π_1 formula; we might say it is Δ_1 for K and M, and hence absolute between them.

The Π_n/Σ_n classification is not confined to set theory; in a more general context, quantifier-free formulas play the role of Δ_0 formulas. Historically, proofs in logic often began by reducing formulas to prenex normal form (i.e., all quantifiers in front). This isn't so widespread anymore. But induction on the "complexity" of formulas still pervades logic, and the Π_n/Σ_n classification is our deepest analysis of this complexity.

[7]Suppose the Σ_1 formula holds for $(\vec{c},d)$. If the antecedent holds in the Π_1 formula for some $\vec{u}$ with $\vec{x}=\vec{c}$ and $z=d'$, then the Σ_1 formula also holds for $(\vec{c},d')$. By the uniqueness hypothesis, $d=d'$ and the consequent holds for $(\vec{c},d)$ in Π_1 formula. That shows that the Σ_1 formula implies the Π_1 formula. For the other direction, suppose the Π_1 formula holds for $(\vec{c},d)$. By the existence part of "function-like", there must be a $\vec{u}$ and a d' making the Σ_1 formula true for $(\vec{c},d')$. The Π_1 formula tells us that $d=d'$, so the Σ_1 formula holds for $(\vec{c},d)$. This argument can be extended by induction to show that function-like Σ_n formulas are Π_n.

17.3 Absoluteness of Constructibility

Now we turn to the absolutness of the notion of constructibility. There is a formula Const(x) which says that x is constructible, and which holds in L iff it holds in V. Const(x) is not Δ_0, nor is it absolute over *all* transitive classes, so some subtleties come into play. (It is absolute between models of ZF.) S&F devote much of Ch.14 to these matters.

Herewith a brief sketch. I will lean more heavily on the corner bracket notation, make a clearer distinction between syntax and semantics, emphasize the role of names, give some related results for context, and handwave; I hope all this makes the discussion more readable.

Some notation, using my conventions. I will try to avoid conflict with S&F; often they just don't have a notation for something. Let a be any set.

Symbol	Meaning	Example
$\mathcal{L}$(ZF)	Language of ZF	
$\mathcal{L}_V$(ZF)	$\mathcal{L}$(ZF) augmented with names for all sets	
$\mathcal{L}_a$(ZF)	$\mathcal{L}$(ZF) augmented with names for all elements of a	
$\mathcal{E}$	codes of formulas of $\mathcal{L}_V$(ZF)	$\ulcorner\varphi(\vec{x},\vec{c})\urcorner$
$\mathcal{E}^a$	codes of formulas of $\mathcal{L}_a$(ZF)	$\ulcorner\varphi(\vec{x},\vec{c})\urcorner$ with $\vec{c}\subseteq a$
S	codes of sentences (closed formulas) of $\mathcal{L}_V$(ZF)	$\ulcorner\varphi(\vec{c})\urcorner$
$S(a)$	codes of sentences of $\mathcal{L}_a$(ZF)	$\ulcorner\varphi(\vec{c})\urcorner$ with $\vec{c}\subseteq a$
M	codes of monadic formulas of $\mathcal{L}_V$(ZF)	$\ulcorner\varphi(x,\vec{c})\urcorner$
$M(a)$	codes of monadic formulas of $\mathcal{L}_a$(ZF)	$\ulcorner\varphi(x,\vec{c})\urcorner$ with $\vec{c}\subseteq a$

$\mathcal{L}$(ZF) has only one predicate symbol '$\in$'; I also regard '$=$' as a basic logical symbol. For $\mathcal{L}_V(\mathrm{ZF})$, we could let sets name themselves, as S&F suggest. Since $\mathcal{L}(\mathrm{ZF})$ has no (built-in) constants, *name* and *constant* are synonymous here.

It is illuminating to consider three results together, all concerning the "definability of truth". There is a pure formula $\textsc{true}(x, y)$ (in $\mathcal{L}(\mathrm{ZF})$) defining truth for sentences in $\mathcal{L}_a(\mathrm{ZF})$. For any $n \in \mathbb{N}$, there is a pure formula $\textsc{true}_n(x)$ defining truth for sentences in $\mathcal{L}_V(\mathrm{ZF})$ of parsing depth at most n. But there is *no* pure formula $\textsc{true}(x)$ defining truth for sentences in $\mathcal{L}_V(\mathrm{ZF})$. The formula for $\mathcal{L}_a(\mathrm{ZF})$ is even Δ_1^{ZF}. Summarizing:

Yes (Δ_1^{ZF})	$a \models \varphi(\vec{c})$	iff	$V \models \textsc{true}(a, \ulcorner\varphi(\vec{c})\urcorner)$	$(\vec{c} \subseteq a)$
Yes	$V \models \varphi(\vec{c})$	iff	$V \models \textsc{true}_n(\ulcorner\varphi(\vec{c})\urcorner)$	(depth $\leq n$)
No	$V \models \varphi(\vec{c})$	iff	$V \models \textsc{true}(\ulcorner\varphi(\vec{c})\urcorner)$	

(I've written $a \models$ instead of $(a, \in) \models$ for brevity, likewise for V.) The last result is known as Tarski's theorem on the undefinability of truth. The first two are formalizations of Tarski's definition of truth.

Induction and Functions

Before plunging into details, some preliminaries. (S&F §14.2 has more.)

S&F define a function V_a instead of a predicate $\textsc{true}(a, -)$; $V_a(\ulcorner\varphi(\vec{c})\urcorner) = 1$ iff $a \models \varphi(\vec{c})$. Functions often have technical advantages over predicates.

Formally defining a function f means having a formula $\varphi(\vec{x}, y)$ for the relation $f(\vec{x}) = y$. If φ is Δ_1^{ZF}, then we say that f has a Δ_1^{ZF} definition (likewise for Σ_1 or Π_1 or pure or whatever). Most authors (unlike S&F) require functions to have domains that are sets, so the word *functional* is often used when this restriction is dropped. S&F say a formula $\varphi(x, y)$ is *function-like* over a transitive class K if it defines a functional defined on all of K.

As a reminder, Σ_1 definitions are absolute upwards, Π_1 absolute downwards. and Δ_1^{ZF} are absolute between models of ZF.

Example: the power set functional $\mathcal{P}(x)$ has a Π_1 definition because $y = \mathcal{P}(x)$ iff $(\forall z)[z \in y \leftrightarrow z \subseteq x]$, and $z \subseteq x$ is Δ_0. The ordered pair functional $\langle x, y\rangle$ is Δ_0 because

$$z = \langle x, y\rangle \text{ iff } (\exists u, v \in z)[z = \{u, v\} \wedge u = \{x\} \wedge v = \{x, y\}]$$

and $z = x \times y$ is Δ_0 because

$$\begin{aligned} z = x \times y \text{ iff } \quad & (\forall p \in x)(\forall q \in y)(\exists u \in z)[u = \langle p, q\rangle] \\ & \wedge (\forall u \in z)(\exists p \in x)(\exists q \in y)[u = \langle p, q\rangle] \end{aligned}$$

Special case: 0-ary functions, aka distinguished elements. The most important example is ω: $w = \omega$ iff w is transitive and the elements of w are linearly ordered under $\in$, because Foundation then implies that w is well-ordered under $\in$. So ω has a Δ_0 definition.

Definition by transfinite induction preserves Σ_1^{ZF}-ness, Π_1^{ZF}-ness, and hence Δ_1^{ZF}-ness. Proof: Suppose the functional $F(x)$ is Σ_1^{ZF}. Define F^* by

$$\begin{aligned} F^*(0) &= \varnothing \\ F^*(\alpha + 1) &= F(F^*(\alpha)) \\ F^*(\lambda) &= \bigcup_{\alpha<\lambda} F^*(\alpha) \end{aligned}$$

Then $y = F^*(\alpha)$ iff there is a function f with domain $\alpha + 1$ such that f satisfies the inductive demands for all $\beta \leq \alpha$ and $y = f(\alpha)$. Explicitly, but with some vernacular:

$$\begin{aligned} & \alpha \text{ is an ordinal and} \\ & (\exists f)\big[f \text{ is a function with domain } \alpha + 1 \text{ and} \\ & \quad f(0) = \varnothing \text{ and} \\ & \quad (\forall \beta \in \alpha + 1)[f(\beta + 1) = F(f(\beta))] \text{ and} \\ & \quad (\forall \text{ limit } \lambda \in \alpha + 1)[f(\lambda) = \bigcup_{\beta \in \lambda} f(\beta)] \text{ and} \\ & \quad y = f(\alpha)\big] \end{aligned}$$

Thus if F is Σ_1^{ZF}, so is F^*. On the other hand, if F is Π_1^{ZF}, we use a formula saying that *all* functions f with domain $\alpha+1$ and satisfying the inductive demands must have $f(\alpha)=y$. Explicitly:

$$
\begin{aligned}
&\alpha \text{ is an ordinal and} \\
&(\forall f)\big[f \text{ is a function with domain } \alpha+1 \text{ and} \\
&\quad f(0)=\varnothing \text{ and} \\
&\quad (\forall \beta\in\alpha+1)[f(\beta+1)=F(f(\beta))] \text{ and} \\
&\quad (\forall \text{ limit } \lambda\in\alpha+1)[f(\lambda)=\bigcup_{\beta\in\lambda} f(\beta)] \;\rightarrow \\
&\quad y=f(\alpha)\big]
\end{aligned}
$$

We can ring variations on this theme. Extra parameters can come along for the ride, i.e., replace $F(x)$ with $F(x,\vec{y})$ and $F^*(\alpha)$ with $F^*(\alpha,\vec{y})$. We don't have to start with $F^*(0)=\varnothing$, we can define a functional $F^*(\alpha,a)$ where a is the initial value (i.e., $F^*(0,a)=a$, $F^*(\alpha+1,a)=F(F^*(\alpha,a))$, etc.) We can include α as an additional argument to F, i.e., $F^*(\alpha+1)=F(\alpha,F^*(\alpha))$. "Ordinary" inductions that only go up to ω preserve $\Sigma_1^{\mathrm{ZF}}/\Pi_1^{\mathrm{ZF}}/\Delta_1^{\mathrm{ZF}}$-ness because "being ω" is Δ_0.

Example: the functional $y=\mathrm{tc}(x)$ (the transitive closure of x) is Δ_1^{ZF} because $z\in x$ is Δ_0 and tc has the inductive definition

$$
\begin{aligned}
\mathrm{tc}(0,x) &= \{x\} \\
\mathrm{tc}(n+1,x) &= \mathrm{tc}(n,x)\cup\bigcup_{z\in x}\mathrm{tc}(n,z) \\
\mathrm{tc}(\omega,x) &= \bigcup_{n\in\omega}\mathrm{tc}(n,x)
\end{aligned}
$$

and $\mathrm{tc}(x)=\mathrm{tc}(\omega,x)$. (Here I adhere to the S&F convention of including x as an element of $\mathrm{tc}(x)$.) Likewise rank is Δ_1^{ZF}, and more generally definitions by so-called $\in$-induction preserve $\Sigma_1^{\mathrm{ZF}}/\Pi_1^{\mathrm{ZF}}/\Delta_1^{\mathrm{ZF}}$-ness.

For industrial-strength use, one can develop a whole calculus to determine places in the complexity hierarchy. But we need very little of this.

I should mention that the S&F proof of Theorem 2.8 (p.183) contains an error, writing $F(f\restriction\beta)$ instead of $F(f(\beta))$.

I've included the ZF superscript throughout; this sweeps away any concerns about the placement of bounded quantifiers. S&F by contrast show only Σ-ness, which permits bounded quantifiers anywhere.

Formalizing Syntax in ZF

To formalize syntax in ZF, we code the base layer in a somewhat arbitrary but straightforward manner. Then we throw (ordinary) induction at it.

For the base layer, we need codes for all the *individuals*, i.e., the variables and the names of elements of V. S&F use $\langle 0, i\rangle$ to code v_i and $\langle 1, a\rangle$ to code the name for a. But we don't need to worry about those details, we'll just write $\ulcorner v_i \urcorner$ and $\ulcorner a \urcorner$.

Next, formulas. One can obviously code a formula as a finite sequence of codes of symbols and of individuals. S&F (and others) use a parse tree instead, a cleaner approach I think. They define the "tree-building" functionals

$$\begin{aligned}
(\ulcorner x \urcorner, \ulcorner y \urcorner) &\mapsto \ulcorner x \in y \urcorner \\
\ulcorner \varphi \urcorner &\mapsto \ulcorner \neg\varphi \urcorner \\
(\ulcorner \varphi \urcorner, \ulcorner \psi \urcorner) &\mapsto \ulcorner \varphi \wedge \psi \urcorner \\
\ulcorner \varphi \urcorner &\mapsto \ulcorner (\exists v_i)\varphi \urcorner
\end{aligned}$$

using ordered pairs and triples, with the last slot serving as "tag" to identify the type of each node (negation, conjunction, etc.) That makes it pretty obvious that these functionals are Δ_0.

The functional $a \mapsto \mathcal{E}^a$ takes a set a and finds all codes of formulas where the names all refer to elements of a. Like practically everything in the realm of syntax, we induct on the depth of parse trees. It's very similar to the definition of the transitive closure. We start with the base layer of atomic formulas, denoted $\mathcal{E}^a_0$. We want $\mathcal{E}^a$ to satisfy the inductive condition

$$\begin{aligned} &y \in \mathcal{E}^a \text{ iff} \\ &\qquad y \in \mathcal{E}^a_0 \text{ or} \\ &\qquad y = \ulcorner\neg\varphi\urcorner \text{ with } \ulcorner\varphi\urcorner \in \mathcal{E}^a \text{ or} \\ &\qquad y = \ulcorner\varphi \wedge \psi\urcorner \text{ with } \ulcorner\varphi\urcorner, \ulcorner\psi\urcorner \in \mathcal{E}^a \text{ or} \\ &\qquad y = \ulcorner(\exists v_i)\varphi(v_i)\urcorner \text{ with } \ulcorner\varphi(v_i)\urcorner \in \mathcal{E}^a \end{aligned} \tag{1}$$

To fit this into the F, F^* paradigm, we let F take a set x of codes of formulas, and throw in one application of any of the tree-building functionals. So $x \subseteq F(x)$, and if $\ulcorner\varphi\urcorner \in x$ then $\ulcorner\neg\varphi\urcorner \in F(x)$, etc. (S&F write $\Pi(x)$ instead of $F(x)$, for some reason.) Next, $F^*(0) = \mathcal{E}^a_0$, and $F^*(n+1)$ is defined inductively as $F(F^*(n))$, or in the S&F notation, $\mathcal{E}^a_{n+1} = \Pi(\mathcal{E}^a_n)$ (see Def 3.2 p.184). It follows from the generalities on induction that $\mathcal{E}^a$ is Δ^{ZF}_1.

The class $\mathcal{E}$ of codes of *all* formulas in $\mathcal{L}_V(\mathrm{ZF})$ is a proper class. It has a Σ^{ZF}_1 definition: $x \in \mathcal{E}$ iff $(\exists a)[x \in \mathcal{E}^a]$.

It's the same story for other aspects of syntax. For example, the substitution functional $(\ulcorner\varphi(x)\urcorner, \ulcorner c\urcorner) \mapsto \ulcorner\varphi(c)\urcorner$ has a Δ^{ZF}_1 definition.

I note in passing that the S&F treatment (pp.183–185) has their usual quota of slip-ups. For example, on p.185, the formula $t(m_1 \in m_2) = \{\ulcorner m_1\urcorner, \ulcorner m_2\urcorner\} \cap \omega$ is wrong, and so the proof of clause (2) in Lemma 3.4 is also wrong.

Formalizing Truth in ZF

We turn our attention to the two truth predicates we can formalize in ZF ($a \models \varphi(\vec{c})$, and $V \models \varphi(\vec{c})$ for depth$(\varphi) \leq n$) and the one we can't ($V \models \varphi(\vec{c})$).

For atomic sentences, the formalization is a breeze:

$$\begin{aligned}&\text{TRUE}_0(x) \text{ iff}\\&\qquad (x = \ulcorner c = d \urcorner \ \wedge\ c = d)\\&\qquad \vee\, (x = \ulcorner c \in d \urcorner \ \wedge\ c \in d)\end{aligned}$$

Let's start with the inductive definition we'd *like* for $V \models \varphi(\vec{c})$:

$$\begin{aligned}&\text{TRUE}(x) \text{ iff}\\&\qquad x \text{ is atomic and } \text{TRUE}_0(x)\\&\qquad \vee\ x = \ulcorner \neg\varphi \urcorner \ \wedge \neg\text{TRUE}(\ulcorner \varphi \urcorner)\\&\qquad \vee\ x = \ulcorner \varphi \wedge \psi \urcorner \ \wedge \text{TRUE}(\ulcorner \varphi \urcorner) \wedge \text{TRUE}(\ulcorner \psi \urcorner)\\&\qquad \vee\ x = \ulcorner (\exists x)\varphi(x) \urcorner \ \wedge (\exists d)\text{TRUE}(\ulcorner \varphi(d) \urcorner)\end{aligned} \tag{2}$$

Because of the circularity, this doesn't actually define a formula in $\mathcal{L}(\text{ZF})$; rather, it expresses the property we'd want the formula to have. Tarski's theorem tells us that no such formula exists.

First modification:

$$\begin{aligned}&\text{TRUE}_{n+1}(x) \text{ abbreviates}\\&\qquad x \text{ is atomic and } \text{TRUE}_0(x)\\&\qquad \vee\ x = \ulcorner \neg\varphi \urcorner \ \wedge \neg\text{TRUE}_n(\ulcorner \varphi \urcorner)\\&\qquad \vee\ x = \ulcorner \varphi \wedge \psi \urcorner \ \wedge \text{TRUE}_n(\ulcorner \varphi \urcorner) \wedge \text{TRUE}_n(\ulcorner \psi \urcorner)\\&\qquad \vee\ x = \ulcorner (\exists x)\varphi(x) \urcorner \ \wedge (\exists d)\text{TRUE}_n(\ulcorner \varphi(d) \urcorner)\end{aligned} \tag{3}$$

Imagine the right hand side expanded out repeatedly until we have a formula in $\mathcal{L}(\text{ZF})$. Put another way, the *induction* is outside ZF, although the longer and longer *formulas* belong to ZF.

Not just ever longer formulas: TRUE_n is Σ_n^{ZF}, because of the $(\exists d)$ embedded in it. (If we'd made $\forall$ fundamental and $\exists$ an abbreviation, then TRUE_n would be Π_n^{ZF}.)

As noted, $V \models \varphi$ iff $V \models \text{TRUE}_n(\ulcorner\varphi\urcorner)$, provided φ has depth $\leq n$. For the second modification of (2), we define a single formula $\text{TRUE}(x, y)$ such that $a \models \varphi(\vec{c})$ iff $V \models \text{TRUE}(a, \ulcorner\varphi(\vec{c})\urcorner)$. This time there is no restriction on the depth of $\varphi(\vec{c})$, but we do demand that $\vec{c} \subseteq a$.

We handle the circularity in (2) just as we did for $\mathcal{E}^a$ in (1). Recall that $S(a)$ is the set of codes of sentences of $\mathcal{L}_a(\text{ZF})$. Let $S_n(a)$ be the codes for sentences of depth $\leq n$. Let $T_n(a)$ be all the true sentences of $S_n(a)$, i.e., all that are satisfied by $(a, \in)$. We have an inductive definition of $T_n(a)$. $T_0(a)$ presents no issues: $x \in T_0(a)$ iff x is atomic and $\text{TRUE}_0(x)$.

$$\begin{aligned} x \in T_{n+1}(a) \text{ iff } x \in S_{n+1}(a) \text{ and } \Big[& \\ & x \text{ is atomic and } x \in T_0(a) \\ & \vee\ x = \ulcorner\neg\varphi\urcorner \ \wedge \ulcorner\varphi\urcorner \notin T_n(a) \\ & \vee\ x = \ulcorner\varphi \wedge \psi\urcorner \ \wedge \ulcorner\varphi\urcorner \in T_n(a) \ \wedge\ \ulcorner\psi\urcorner \in T_n(a) \\ & \vee\ x = \ulcorner(\exists x)\varphi(x)\urcorner \ \wedge (\exists d \in a)(\ulcorner\varphi(d)\urcorner \in T_n(a))\Big] \end{aligned} \tag{4}$$

It's no sweat to turn this into a function F such that $T_{n+1}(a) = F(T_n(a))$ for all $n \in \omega$. Moreover, F has a Δ_0 definition, because all the tree-building functionals are Δ_0. Note that the crucial existential quantifier, "$(\exists d \in a)$", is now *bounded*. So we have a Δ_1^{ZF} definition of the set $T(a)$ of true sentences in $S(a)$.

The S&F treatment is essentially the same, but they prefer to deal with the characteristic function of $T(a)$, which they call a *valuation function* and denote V_a.

I've emphasized the parallels between $\mathcal{E}^a$ and $T(a)$. Now let's highlight the differences. We defined the proper class $\mathcal{E}$ via the equivalence $x \in \mathcal{E} \equiv (\exists a) x \in \mathcal{E}^a$. Why doesn't this work for $T \subseteq S$, the proper class of true sentences about V? First hint of the problem: $V \models \varphi(\vec{c})$ and $(\exists a) a \models \varphi(\vec{c})$ are *not equivalent*, even when $\vec{c} \subseteq a$. We search for a culprit; the quantifier pleads guilty. Syntax doesn't care about the scope of a quantifier $(\exists x)$—it's just a node in the parse tree. But for semantics, the scope is central to the meaning. Put another way, when syntax examines the formula $\varphi(\vec{x}, \vec{c})$, it "sees" only the names explicitly present. Semantics considers *all* possible names when turning $(\exists x)\varphi(x)$ into $\varphi(d)$.

Absoluteness of *L*

Consider this list of relations. (S&F treat them in Prop.3.11 and Theorems 3.12 through 4.5, pp.188–190.)

1. $a \models \varphi(\vec{c})$, with $\vec{c} \subseteq a$.
2. $y = \{x \in a : a \models \varphi(x, \vec{c})\}$. (S&F: $\mathrm{Def}(\ulcorner\varphi\urcorner, y, a)$)
3. $y \in \mathcal{F}(a)$.
4. $y = \mathcal{F}(a)$.
5. $y \in L_\alpha$ and $y = L_\alpha$, where α is an ordinal. (S&F: $\mathcal{L}(\alpha, y) \equiv y \in L_\alpha$, $\mathcal{M}(\alpha, y) \equiv y = L_\alpha$)
6. $y \in L$. (S&F: $L(y)$ or $\mathrm{Const}(y)$)

These are all Δ_1^{ZF} except the last one. But $y \in L$ is absolute between V and L. Proof:

1. We've just seen that this is Δ_1^{ZF}, or rather, its equivalent $\textsc{true}(a, \ulcorner\varphi(x, \vec{c})\urcorner)$ is.
2. $y = \{x \in a : a \models \varphi(x, \vec{c})\}$ iff $(\forall z \in y)(a \models \varphi(z, \vec{c}))$ and $(\forall z \in a)[(a \models \varphi(z, \vec{c})) \to z \in y]$. So this is Δ_1^{ZF}.
3. $y \in \mathcal{F}(a)$ iff there is a monadic formula $\varphi(x, \vec{c})$ in $\mathcal{L}_a(\mathrm{ZF})$ such that $y = \{x \in a : a \models \varphi(x, \vec{c})\}$. Recall that $M(a)$ is the set of codes of such monadic formulas. So $y \in \mathcal{F}(a)$ iff

$$(\exists p \in M(a))[y = \{x \in a : \textsc{true}(a, p)\}]$$

The new feature: instead of a true bounded quantifier, we have something of the form $(\exists u \in f(z))\psi(y,z,u)$ where $f(z)$ and ψ are both Δ_1^{ZF}. But that's equivalent to $(\exists v)(\exists u \in v)[v = f(u) \wedge \psi(y,z,u)]$, which is Σ^{ZF}. It's also equivalent to $(\forall v)[v = f(u) \to (\exists u \in v)\psi(y,z,u)]$, which is Π_1^{ZF}.

4. $y = \mathcal{F}(a)$ iff $(\forall z \in y)[z \in \mathcal{F}(a)] \wedge (\forall z \in \mathcal{F}(a))[z \in y]$. We can handle the "kind-of" bounded quantifier, $(\forall z \in \mathcal{F}(a))$, much the same as we handled $(\exists p \in M(a))$ in the previous item.

5. Transfinite induction preserves Δ_1^{ZF}-ness.

6. $(\exists\alpha)[y \in L_\alpha]$. So $y \in L$ is Σ_1^{ZF}. So Const is upwards absolute from L to V. In the reverse direction, suppose $V \models \mathrm{Const}(s)$ for some s. Then for some ordinal α, $V \models s \in L_\alpha$. But L contains all ordinals, and being an ordinal is absolute, and L_α is absolute, so $L \models s \in L_\alpha$ and hence Const is downwards absolute from V to L.

This final item (6) is the goal of the whole argument. But as a matter of curiosity you might ask, is Const downwards absolute between models of ZF? How about upwards and downwards absoluteness for transitive classes in general?

S&F show that (1)–(6) are Σ, so all these notions are upwards absolute between transitive classes. We needed one feature of L, besides being a model of ZF, to establish downwards absoluteness from V: the fact that L contains all ordinals. Are there *any* standard models of ZF not containing all ordinals? The "yes" answer is one form of axiom SM. Using this, plus forcing, one can show Const is not downwards absolute between models of ZF.

It's much easier to show that Const is not downwards absolute between transitive classes. Consider $L_{\alpha+1}$ for some α. Suppose s has rank α, i.e., $s \in L_{\alpha+1} \setminus L_\alpha$. (For example, let $s = \alpha$ or $s = L_\alpha$.) So s is constructible, but not constructible "in $L_{\alpha+1}$". For if $L_{\alpha+1} \models \text{Const}(s)$, then $L_{\alpha+1} \models (\exists\beta) s \in L_\beta$, i.e., $s \in L_\beta$ for some $\beta \leq \alpha$. But we assumed s had rank α.

18 Reflection Principles

A reflection principle gives circumstances in which $K \models \varphi$ iff $M \models \varphi$, where $K \subseteq M$. "As above, so below."[8] The absoluteness of Δ_0 formulas *could* be called a reflection principle: if φ is Δ_0, then the transitivity of K and M is all we need. Usually though people reserve the term for the downward Löwenheim-Skolem theorem and its descendents.

The original downward Löwenheim-Skolem was the first really deep theorem of first-order logic. It says that if M is a structure for a countable language $\mathcal{L}$ and $K_0 \subseteq M$, K_0 countable, then there is a K with $K_0 \subseteq K \subseteq M$, K also countable, with K "reflecting" M in this sense: for any formula $\varphi(\vec{x})$ in $\mathcal{L}$ and any $\vec{a}$ in K,

$$K \models \varphi(\vec{a}) \Leftrightarrow M \models \varphi(\vec{a})$$

We say K is an *elementary substructure* of M. Standard notation: $K \preceq M$.

The absoluteness of Δ_0 formulas says that for a *highly restricted* class of formulas, truth is reflected, with very modest conditions on the pair $K \subseteq M$. The downward LS says that for *any* formula $\varphi(\vec{x})$ and any M, *we can find* a K that "reflects" M.

Ch.11 does a nice job going through all the details, including the variations needed for Gödel's results. The Logic notes discusses the downward Löwenheim-Skolem theorem in a general setting. In §20 I outline how reflection is used. Briefly, it compensates for failures of absoluteness. I have no criticisms of the S&F exposition in this chapter, so I'll just add some remarks around the margins.

[8] A quote from the Emerald Tablet of Hermes Trismegistus, a alchemical sacred text.

Surprisingly, S&F don't mention the famous Skolem paradox. By the class version of the theorem (Theorem 11.3.2, p.148), there is a countable model of any finite subset of ZF. Consider the theorem "$\mathcal{P}(\omega)$ is uncountable". The proof of this uses only a finite number of ZF axioms, so it has a countable model. What gives? Answer: in any countable model, $\mathcal{P}^M(\omega)$ possesses a 1–1 correspondence with ω, but the correspondence isn't *in the model.* (Skolem's paradox has many variations, all with similar resolutions.)

Theorem 11.2.1 (p.147) adds extensionality to the downward LS. Note that if we include '=' with our basic logical symbols, along with its logical axioms, then we don't need to make any special effort: the extensionality axiom is just another formula that gets reflected in Theorem 11.1.1 (p.145)

To conclude this section, let's prepare a segue to the Mostowski Collapsing Lemma. Suppose K_0 is one of the L_α's, and M is L. Now both L_α and L are transitive. But the iterative process of adding elements to $K_0 = L_\alpha$ might not preserve transitivity. It's not hard to insure that $K \supseteq L_\alpha$ is extensional, though. The Mostowski Lemma shows how to collapse K to a transitive K' with $L_\alpha \subseteq K' \subseteq L$. K' still reflects L because the collapsing map is an $\in$-isomorphism. (See Theorem 11.4.1, p.150, for details.)

19 Mostowski Collapsing Lemma

S&F define a relational system to be a pair (A, R) where A is a class and R is a relation on A. If R has certain properties, we get a mapping (the Mostowski collapsing map) $F : A \to M$ from A to a transitive class M. F is an isomorphism, in this sense: xRy if and only if $F(x) \in F(y)$. If xRy, we say x is a *component* of y, and we write y^* for the class of components of x.

The most important property demanded of R, as we will see, is *extensionality*: if $x \neq y$ then $x^* \neq y^*$. The two other requirements are properness (x^*

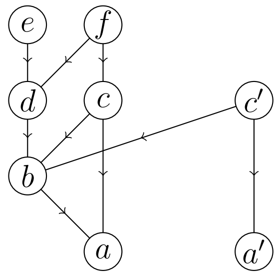

Figure 4: A Mostowski Map

$a \mapsto \varnothing$	$d \mapsto \{1\}$
$b \mapsto \{\varnothing\} = 1$	$e \mapsto \{\{1\}\}$
$c \mapsto \{\varnothing, 1\} = 2$	$f \mapsto \{\{1\}, 2\}$
$a' \mapsto \varnothing$	$c' \mapsto \{\varnothing, 1\} = 2$

is a set for all $x \in A$) and well-foundedness (R is a well-founded relation). If we drop extensionality, we can still define the Mostowski map, but we no longer get an isomorphism.

Figure 4 illustrates the Mostowski mapping F for a relational system (A, R). It helps to think of the elements of A as *names* or *labels* for sets. The mapping $x \mapsto F(x)$ sends the name x to the named set $F(x)$. In the figure, a is a name for the empty set, since it has no components. All the R-edges are explicit in the figure: although $bRdRe$, we do not have bRe.

The defining fact

$$F(x) = F''(x^*)$$

can also be expressed as

$$F(y) \in F(x) \leftrightarrow yRx$$

as we mentioned, or

> The names of the elements of $F(x)$ are precisely the components of x.

Fig.4 shows a non-extensional system: a' has the same components as a. This single failure of extensionality propagates upwards: while c and c' have different components, they still map to the same set because a' and a both name the same set. The figure shows a well-founded relation. Removing a' and c' makes it extensional. For the rest of this section, we deal with the modified (A, R), with $A = \{a, b, c, d, e, f\}$.

Now let's consider what $p(x)$ (p.130) means. The function $p(x)$ is defined on subsets of A. The definition says:

$$y \in p(x) \leftrightarrow y^* \subseteq x$$

In other words: y belongs to $p(x)$ iff the set of components of y is a subset of x.

Example 1: $x = \{f\}$. Then x has only two subsets, $\varnothing$ and $\{f\}$. Since $a^* = \varnothing$, and everything else has a component that is not f, we have $p(\{f\}) = \{a\}$.

Example 2: $x = f^* = \{c, d\}$. So $y \in p(f^*)$ iff all the components of y are c or d. This is true precisely for a, e, and f, so $p(f^*) = \{a, e, f\}$.

Example 3: $x = e^*$. Then $p(e^*) = \{a, e\}$.

We see that $p(x)$ is *sort of* the power set of x, but *not quite*. If $u \subseteq x$, we don't put u directly into $p(x)$: we look for a y with $y^* = u$. If we find one, we add y to $p(x)$, but if we don't, then u goes unrepresented.

If p truly represented the power set $\mathcal{P}$, we'd expect a commutative diagram like this:

$$\begin{array}{ccc} p(x) & \xrightarrow{F''} & \mathcal{P}(F''(x)) \\ \uparrow{\scriptstyle p} & & \uparrow{\scriptstyle \mathcal{P}} \\ x & \xrightarrow[F'']{} & F''(x) \end{array}$$

(We use F'' rather than F because x and $p(x)$ are both *sets of names* (subsets of A) rather than *names* (elements of A).) Do we have $F''(p(x)) = \mathcal{P}(F''(x))$? Not in general, because of unrepresented subsets. But we do have $F''(p(x)) \subseteq \mathcal{P}(F''(x))$.

Now let's take the special case where R is $\in$. $(A, \in)$ is automatically proper (Prop. 1.2, p.128); I always assume Foundation, so well-foundedness also comes for free. The key contrast occurs between extensionality and transitivity. Transitivity implies extensionality (easy to see), but not conversely. This is the whole point of the Mostowski map.

Here's how to think about it: when R is $\in$, then *all components are elements*, but *not all elements need be components*.

For example, here's a system $(A, \in)$ with the same graph as in fig.4 (with a' and c' removed):

$$\begin{array}{ll} a = \{5\} & d = \{b, 7\} \\ b = \{a, 6\} & e = \{d\} \\ c = \{a, b, 5\} & f = \{c, 5, 7\} \end{array}$$

I started with the image of F (i.e., $F(a) = \varnothing$, $F(b) = 1$, etc.) as the "main ingredients". Then I added 5, 6, and 7 as "spice". Since the spice does not belong to A, its addition doesn't change the graph. So b's only component is still a, c's only components are still a and b, etc. The map F "de-spices" the dish.

To conclude this section, let's highlight the critical results of the chapter. Theorem 4.2 (p.135) says that any extensional class A is $\in$-isomorphic to a transitive class M, and gives the defining property of the $\in$-isomorphism F:

$$F(x) = F''(x \cap A)$$

Here, $x \cap A$ are the components of x in the relational system $(A, \in)$. Theorem 4.1 generalizes Theorem 4.2 to any extensional well-founded proper relational system (A, R). This appears to be one of those useless but pleasant generalizations S&F are so fond of[9].

Note that if A is transitive and $x \in A$, then $x \cap A = x$. A simple transfinite induction produces Theorem 5.1 (p.136): if T is a transitive subclass of A, then F restricted to T is the identity.

The literature refers to Theorem 4.2 as the Mostowski Collapsing Lemma. "Collapsing" is justified by Theorem 6.7: if A is extensional and α is an ordinal in A, then $F(\alpha)$ is an ordinal and $F(\alpha) \leq \alpha$. In other words, F "collapses" ordinals. More generally, F collapses ranks: $\text{rank}(F(x)) \leq \text{rank}(x)$. This follows from Theorem 6.8 (p.139) plus Theorem 2.7 (p.132): the V-rank of $F(x)$ is the A-rank of x, which is less than or equal to the V-rank of x. The V-rank of a set is simply its rank (p.131). As a simple example, let $A = \{1, 2\}$. Then $1 \mapsto 0$ (i.e., $\varnothing$) and $2 = \{0, 1\} \mapsto \{F(1)\} = 1$.

20 Relative Consistency of AC and GCH

Now let's see how S&F put it all together. Recall that Gödel proved three main results about L:

1. L is a model of ZF.
2. $V = L$ holds in L.

[9]In principle, S&F could have appealed to Theorem 4.1 to construct the forcing model $M[G]$ in Ch.21 (p.292). But they find it easier to give a direct definition.

3. $V = L \Rightarrow$ AC and $V = L \Rightarrow$ GCH are both provable in ZF.

Let's start by assuming (1)–(3) and show, in detail, how this implies the relative consistency result: if ZF is consistent then ZF+AC+GCH is consistent. I want to highlight two aspects: (a) the role of absoluteness in (2); (b) syntax vs. semantics.

I'll use a common notation and write ZFL for ZF + $V = L$. From (3) we immediately conclude that the relative consistency of ZFL implies the relative consistency of ZF+AC+GCH (relative to ZF, naturally).

We want to show that if ZF is consistent, so is ZFL. First, the quick version, with 50% handwaving. If ZFL were inconsistent, that would mean that ZF $\vdash \neg V = L$. Since L is a model of ZF (by (1)), it follows that $L \models V \neq L$. But $L \models V = L$ (by (2)), so we've proven two contradictory things about L. If we assume ZF is "sound", i.e., all theorems of ZF are "true in V", then this is impossible, so ZFL must be consistent.

We can turn the screws a little tighter, just by being more careful about classes and sets. Suppose $\Lambda(x)$ is the formalization of $x \in L$. In other words, $\Lambda(x)$ is a mile-long formula (if written out in full) with one free variable (and no parameters) such that $V \models \Lambda(x)$ iff x is constructible. I gave a very cursory sketch of how to define Λ in §17, and S&F spend a lot more time on this in Ch.14.

The section "Relativization" in the Logic notes describe relativization: if ψ is a formula in ZF, then you obtain ψ^L by recursively applying these rules:

$$
\begin{aligned}
(\exists x)\varphi(x) &\Rightarrow (\exists x)[\Lambda(x) \wedge \varphi^L(x)] \\
(\forall x)\varphi(x) &\Rightarrow (\forall x)[\Lambda(x) \rightarrow \varphi^L(x)] \\
(\alpha \wedge \beta) &\Rightarrow \alpha^L \wedge \beta^L \\
(\alpha \vee \beta) &\Rightarrow \alpha^L \vee \beta^L \\
(\neg\varphi)^L &\Rightarrow \neg(\varphi^L) \\
x \in y &\Rightarrow x \in y \\
x = y &\Rightarrow x = y
\end{aligned}
$$

With a bit of syntactic sugar (vernacular), we could write the results in the first two rules as $(\exists x \in L)\varphi^L(x)$ and $(\forall x \in L)\varphi^L(x)$.

The point of relativization is this fact:

$$V \models \psi^L(\vec{a}) \;\Leftrightarrow\; L \models \psi(\vec{a}), \qquad \vec{a} \subseteq L$$

and so the formalization of $L \models \psi$ (or even $V \models (L \models \psi)$) is just ψ^L.

Items (1) and (2), formalized in ZF, are:

1. $\mathrm{ZF} \vdash \alpha^L$ for every axiom α of ZF.
2. $\mathrm{ZF} \vdash (\forall x)[\Lambda(x) \rightarrow \Lambda^L(x)]$.

Now suppose ZFL is inconsistent, so

$$\mathrm{ZF} \vdash (\exists x)\neg\Lambda(x)$$

Relativization preserves rules of inference and basic logical axioms. So (1) implies that any proof in ZF can be relativized to L: just replace every use of an axiom α with a proof of α^L. Apply this to the formula just displayed, to get

$$\mathrm{ZF} \vdash (\exists x)[\Lambda(x) \wedge \neg\Lambda^L(x)]$$

Now we make use of (2), and get:

$$\mathrm{ZF} \vdash (\exists x)[\Lambda^L(x) \wedge \neg\Lambda^L(x)]$$

a contradiction to the basic rules of logic.

So an inconsistency in ZFL can be transformed into one in ZF. The argument is purely syntactic—provability has replaced satisfaction. Only 15% handwaving! (For the S&F take on this, see §15.4, p.199.)

It's worth taking a moment to think about $\Lambda^L(x)$. This is *not* an infinite regress: in the mile-long formula $\Lambda(x)$, we replace the quantifiers as described. We use the *unrelativizied* Λ for this. For example, if we have $(\exists z)[\ldots]$ somewhere in the middle of $\Lambda(x)$, this first becomes $(\exists z)[\Lambda(z) \wedge \ldots]$, and then we apply relativization to "…", but we don't apply it to the just introduced $\Lambda(z)$.

Item (2) says that $\Lambda(x)$ is absolute between L and V. Normally absoluteness is an "if-and-only-if", but applied only to elements of the inner structure. So in one direction, we want "if $V \models \Lambda(a)$ and $a \in L$, then $V \models \Lambda^L(a)$"; in the other direction, "if $V \models \Lambda^L(a)$ and $a \in L$, then $V \models \Lambda(a)$". However, the formal version of "$a \in L$" is just "$V \models \Lambda(a)$", so this devolves to (2).

Just to hammer on the importance of absoluteness: suppose $\Lambda(x)$ was not absolute, like countability. For example, let's say that $\Lambda(x)$ had the form $(\exists y)\Phi(x, y)$, and for some $a \in L$, all the possible witnesses $y = b$ did not belong to L. So we would have a constructible set that L thought was not constructible—constructible, and thus belonging to L, but not "constructible in L". ZF could still prove $V \neq L$; neither V nor L would present a counterexample.

Now let's turn to the proofs of (1)–(3). We disposed of (2) in §17. As for ZFL $\Rightarrow$ AC, Gödel showed that L can be well-ordered. This shouldn't surprise us.

On the one hand, $L = \bigcup_\alpha L_\alpha$, so we can partially order constructible sets by their L-rank. New sets that appear during the induction step, $L_{\alpha+1} = \mathcal{F}(L_\alpha)$, all have definitions. Now, a definition can have parameters belonging to L_α:

$$x = \{y \in L_\alpha : L_\alpha \models \varphi(y, \vec{a})\},\ \vec{a} \subseteq L_\alpha$$

but that just means that each new set can be given a "card catalogue index" $(\varphi, \vec{a})$. If L_α has a well-ordering, so does the set of finite-length tuples $\vec{a}$, and φ is just a finite string of symbols selected from a countable alphabet. So no problems well-ordering $L_{\alpha+1}$. More careful analysis shows that all of L can be well-ordered by a formula $\psi(x, y)$. Moreover, $\psi(x, y)$ is absolute over L.

S&F give a detailed and lucid treatment of the proof of AC in §14.6 (pp.191–192). Note that this section does not depend on any of the preceding chapters in Part II—a good thing, since a key earlier result (Lemma 13.3.5, p.165) depends on it!

Now we come to the hard parts, where absoluteness fails. Let's start with the proof that $L \models$ ZF. The critical fact: any L-definable subset of L, say

$$x = \{z \in L : L \models \varphi(z, \vec{a})\} \text{ with } \vec{a} \subseteq L$$

is itself an element of L whenever x is a set. (It's not automatic that x is a set, but if it is, it belongs to L.)

In general, this is just a touch stronger than the notion of *first-order swelled* (p.170, see also §3 vis-a-vis swelled vs. first-ordered swelled). Let's call this *first-order really swelled.* It turns out that for L, the difference doesn't matter; we'll have more to say about that in a moment.

Assuming the first-order real swellitude of L, we can barrel through the ZF axioms like a snowplow through powder. Take Substitution (aka Replacement). If $F(x) = y$ is defined by $\varphi(x, y)$, then (relativizing to L)

$$(F^L)''(u) = \{y \in L : L \models (\exists x \in u)\varphi(x, y)\}$$

$(F^L)''(u)$ is a set because Substitution holds in V, so $(F^L)''(u)$ belongs to L and L satisfies Substitution. The other axioms likewise fall like wooden soldiers in a hurricane: see Theorem 13.2.1, pp.172–173.

First-order real swellitude must make you think of the definition of $\mathcal{F}$. Indeed, saying that L is first-order really swelled is just saying that $L = \mathcal{F}(L)$. Since $L_{\alpha+1} = \mathcal{F}(L_\alpha)$ and $L = \bigcup_\alpha L_\alpha$, it seems the first-order real swellitude of L should be a near triviality, a three line proof.

Not so! We highlight the problem with our old friend $\mathcal{P}_{\text{un}}(y)$, the set of uncountable subsets of y. We need to show that for any $y \in L$, $\mathcal{P}^L_{\text{un}}(y)$ is also in L. (We need to include only the *constructible* subsets of y that L thinks are uncountable.) That this *is* a set follows immediately from Power Set and Separation. Let

$$x_\alpha = \mathcal{P}^{L_\alpha}_{\text{un}}(y)$$

so $x_\alpha \in L_{\alpha+1}$. As we saw in §17, the x_α ordinal sequence can gain and lose elements as α increases. We need to show that for some α, $x_\alpha = x$.

As it happens, the x_α sequence eventually stabilizes: for any y, there is a β such that for all $\alpha \geq \beta$, $x_\alpha = x$. Let's see why. First look at the easier case of $x = \mathcal{P}^L(y)$, $x_\alpha = \mathcal{P}^{L_\alpha}(y)$. For every $z \in \mathcal{P}^L(y)$, there is an ordinal $\zeta(z)$ such that $z \in L_{\zeta(z)}$. Since $\mathcal{P}^L(y)$ is a set, there is an ordinal β greater than all the $\zeta(z)$'s. So $x = \mathcal{P}^L(y) \subseteq L_\beta$, and in fact is a subset of all L_α with $\alpha \geq \beta$. (I will call this the *waiting* argument. It occurs constantly in set theory, but I don't know of a standard name for it.) The relation $z \subseteq y$ is absolute, which implies that $x_\alpha = x \cap L_\alpha$: if L_α thinks z is a subset of y, so does L, and conversely, provided only that L_α knows about z in the first place. So $x_\alpha = x$ for all $\alpha \geq \beta$. In a phrase: "wait until all the constructible subsets of y have appeared".

Returning to $x_\alpha = \mathcal{P}^{L_\alpha}_{\text{un}}(y)$, $x = \mathcal{P}^L_{\text{un}}(y)$, we have a work a bit harder. First, we can find a β large enough so that $x \subseteq L_\beta$, again by the waiting argument. But as we saw in §17, we can still have $x_\beta \neq x$: x_β might have an element that is "really" countable, but looks uncountable to L_β. However, that's the only snag. We can wait until all countability witnesses have appeared! More precisely, if $z \subseteq y$ is countable, choose an $f_z : \omega \cong z$. Choose a γ so large that all the f_z's belong to L_γ. For any $\alpha \geq \max(\beta, \gamma)$, $x_\alpha = x$.

We have shown that L is closed under the $\mathcal{P}^L_{\mathrm{un}}$ operation. How about first-order real swellitude in general? You can see some kind of induction brewing, an induction on the complexity of the defining formula. S&F give two proofs, one using a reflection principle, the other an elaboration of the argument we just gave in a special case. Let's start with their first argument, the "non-reflection" proof.

A couple of preliminaries. The proof pits the complexity of formulas against the decomposition of L into levels: $L = \bigcup_{\alpha \in On} L_\alpha$, where each L_α is a constructible set. This fact enables the waiting argument: any subset of L is a subset of some L_α. S&F state this as property $\mathbf{C}_{11}$ of L (p.157), but with the explicit reference to levels removed:

> Every set of constructible sets is a *subset* of some constructible set.

Let's say that L has the *waiting property*. The waiting property of L pops up in Prop.12.3.3 (p.164) and Coro.13.2.2 (p.173), which both play important roles, also at the end of §13.3 and in Theorem 13.4.1 (both p.177), generalized to any transitive subclass of V. Coro.13.2.2 says that the waiting property plus first-order swellitude implies first-order real swellitude. (It's easy, a four line proof!) This is why the difference between real and "ordinary" first-order swellitude is inconsequential for L.

The proof of the first-order swellitude of L uses an induction on the complexity of formulas, which means looking at conjuctions (i.e., intersections), negations (i.e., complements), and existential quantifications. Now, the intersection of two constructible sets is a constructible set, but the complement $L \smallsetminus a$ of a constructible set is a proper class. One workaround: deal with $L_\alpha \setminus a$ as α ranges over all the ordinals. If a is a constructible set, then so is $L_\alpha \setminus a$, for all α. It's slicker to generalize this via Def.12.3.1 (p.163):

> $A \subseteq L$ is a constructible class iff $A \cap \ell$ is constructible for all $\ell \in L$ (and so in particular for all $\ell = L_\alpha$).

Though the levels L_α avoid explicit mention, they motivate the approach: think of A as the "limit" of the ordinal sequence $A \cap L_\alpha$. An easy exercise: $A \cap \ell$ constructible for all $\ell = L_\alpha$ implies $A \cap \ell$ constructible for all $\ell \in L$.

Constructible classes are closed under complementation (w.r.t. L) and intersection. First-order swellitude (real or ordinary) for L can now be rephrased: any $A \subseteq L$ that is definable over L is constructible. (We no longer have to stipulate that A is a set.)

With this throat-clearing out of the way, let's look at the key points of the "non-reflection" proof. The crucial point, of course, involves existential quantification.

Suppose A is a constructible class of $(n+1)$-tuples $\langle \vec{a}, b \rangle$. If A is defined by a formula $\psi(\vec{x}, z)$ over L (perhaps including parameters from L), then we need to show that this class of n-tuples is constructible:

$$\{\langle \vec{a} \rangle \in L^n : L \models (\exists z)\psi(\vec{a}, z)\}$$

This set is the projection of A onto its first n coordinates. Let's write $\mathrm{proj}_{\vec{x}} A$ for this, and $\varphi(\vec{x})$ for $(\exists z)\psi(\vec{x}, z)$.

We can grasp the essence of the argument by looking at our three special cases:

$$\begin{array}{ll}
P \subseteq L^2, & P = \{\langle x, y \rangle : L \models x = \mathcal{P}(y)\} \\
U \subseteq L, & U = \{x : L \models x \text{ is uncountable}\} \\
P_{\text{un}} \subseteq L^2, & P_{\text{un}} = \{\langle x, y \rangle : L \models x = \mathcal{P}_{\text{un}}(y)\}
\end{array}$$

All are definable over L, so first-order swellitude insists that they must be constructible classes. To check this, we look at their complements. Write C for $L \setminus U$, and $\neg P$ and $\neg P_{\text{un}}$ for the complements $L^2 \setminus P$ and $L^2 \setminus P_{\text{un}}$. They are defined by these formulas:

$$\begin{aligned}
\neg P: &\qquad L \models (\exists z)[(z \in x \wedge z \not\subseteq y) \vee (z \notin x \wedge z \subseteq y)] \\
C: &\qquad L \models (\exists y) y{:}\omega \cong x \\
\neg P_{\text{un}}: &\qquad L \models (\exists z)[(z \in x \wedge (z \not\subseteq y \vee z \in C)) \\
&\qquad\qquad \vee (z \notin x \wedge z \subseteq y \wedge z \in C)]
\end{aligned}$$

So these examples fit the $\varphi(\vec{x}) \equiv (\exists z)\psi(\vec{x}, z)$ template. We need to show all three classes are constructible, i.e., give constructible sets when intersected with any L_α.

Start with C, the class of sets that L thinks are countable. We noted earlier that

$$C_\alpha = \{x \in L_\alpha : L_\alpha \models x \text{ is countable}\}$$

isn't necessarily the same as

$$C \cap L_\alpha$$

because witnesses for c's countability aren't necessarily in L_α, even when c is. (We know that $C_\alpha \in L_{\alpha+1}$, but we need to show that $C \cap L_\alpha$ is in L.) But look at the class

$$F = \{\langle c, f\rangle \in L^2 : L \models f{:}\omega \cong c\}$$

Since $f{:}\omega \cong c$ is Δ_0, we *do* have

$$F \cap L_\beta = F_\beta = \{\langle c, f\rangle \in (L_\beta)^2 : L_\beta \models f{:}\omega \cong c\}$$

for every β. (This shows us that F is constructible.) By the waiting property of L, we can wait for an L_β ($\beta > \alpha$) containing a countability witness for every element of $C \cap L_\alpha$. So

$$(\text{proj}_x F_\beta) \cap L_\alpha = C \cap L_\alpha$$

But the left hand side of this equation is a subset definable over L_β, i.e., belongs to $L_{\beta+1}$. So it's constructible, although not necessarily in $L_{\alpha+1}$.

It's the same story for $\neg P$: the class

$$W = \{\langle x, y, z\rangle \in L^3 : z \text{ is a witness to } x \neq P^L(y)\}$$

is Δ_0, hence constructible. Using the waiting argument and projecting onto the first two coordinates, we get

$$(\text{proj}_{x,y} W_\beta) \cap L_\alpha = \neg P \cap L_\alpha$$

and so $\neg P$ is a constructible class.

Finally, $\neg P_{\text{un}}$. Let

$$W_{\text{un}} = \{\langle x, y, z\rangle \in L^3 : z \text{ is a witness to } x \neq P^L_{\text{un}}(y)\}$$

Although W_{un} is not Δ_0, it is constructible: it's the union of W with the class

$$\{\langle x, y, z\rangle \in L^3 : z \in x \wedge z \in C\}$$

With some fiddling, the constructibility of C yields the constructibility of W_{un}.

The rest of the argument gets a little messy, but the thrust should be clear. We want to find an L_γ, $\gamma > \alpha$, such that $\neg P \cap L_\alpha$ is definable over L_γ. We first find an L_β, $\beta > \alpha$, such that $W_{\text{un}} \cap L_\beta$ has all the witnesses we need for $\neg P \cap L_\alpha$. Since W_{un} is constructible, $W_{\text{un}} \cap L_\beta$ is definable over some L_γ with $\gamma > \beta$. That γ suffices to define $\neg P \cap L_\alpha$.

Now let's trace the thread of the general argument through S&F. The complexity induction takes place in Lemma 13.3.2 (p.174–175). (The S&F penchant for generalization causes them to replace L with a transitive class K, and to replace "constructible" with "distinguished".) The atomic base cases give no trouble (Prop.12.3.8, p.167). Negation and conjunction follow immediately from closure of constructible classes under complementation and intersection. Existential quantification is recast as projection to the first coordinate of a relation, which they call its domain. Theorem 13.3.3 (p.176) is the punchline: it says that L is a first-order universe. The proof simply notes that L satisfies all the hypotheses of Lemma 13.3.2. For closure under domains (i.e., projection), Prop.12.3.7 (p.166) is cited. So let's look at that.

Prop.12.3.7 is a special case of Lemma 12.3.6 (p.166), which in turn follows easily from Lemma 12.3.5 (p.165). This says that if A is a constructible class and Q is a function with certain properties, then $Q''(A)$ is constructible. For our application, Q is the projection $\text{proj}_{\vec{x}}$, and A is a constructible class of $(n+1)$-tuples $\langle \vec{x}, z \rangle$. The basic idea: given a subset b of $Q''(A)$, we wait until enough tuples $\langle \vec{x}, z \rangle$ show up in A to hit all of b. (I.e., we wait for an L_β such that if $\langle \vec{x} \rangle \in b$, then $\langle \vec{x}, z \rangle \in A \cap L_\beta$ for some z.) Some details are delegated to the proof of Prop.13.3.3 (p.164), which says that a class is constructible iff any subset of it sits inside a constructible subset of it. And the proof of *that* follows directly from the waiting property of L.

Now let's turn to the "reflection" proof, given on p.176. I won't have much to say about this, as their treatment is quite clear. Just note how the argument again revolves around the waiting argument, plus an induction on the complexity of formulas. (Most of this is off-loaded to the proof of the Montague-Lévy reflection principle, Theorem 5.3 (p.152).) I suspect that if you drilled down far enough, you'd find the two proofs are at heart the same.

Reflection principles come to rescue again when proving GCH for L. It turns out that the cardinality of L_α equals the cardinality of α for infinite α. Let's just look at CH; this gives the flavor. Now, $\omega \in L_{\omega+1}$, so the question is, how high up do we have to go in the constructible hierarchy to get all the constructible subsets of ω. Answer: no higher than L_{ω_1}, where ω_1 is the first uncountable cardinal. Given this, CH holds in L: $\mathcal{P}^L(\omega) \subseteq L_{\omega_1}$ and $|L_{\omega_1}| = |\omega_1| = \aleph_1$, so $|\mathcal{P}^L(\omega)| \leq \aleph_1$. The reverse inequality holds because

$|\mathcal{P}^L(\omega)| > \aleph_0$ and $\aleph_1$ is the next larger cardinal after $\aleph_0$. (Remember that we have AC in L, so all cardinals are comparable.)

Cohen is worth quoting (with a couple of tiny changes):

> Let us attempt to give some intuitive justification for why all sets of integers are constructible by countable ordinals. If $x \subseteq \omega$, $x \in L_\alpha$, then one can ask what are the essential properties of α which imply that $x \in L_\alpha$. Now x is determined by the truth values of the countably many statements "$n \in x$". For each n we can think of this as imposing one condition on α. Thus it is not unreasonable that these countably many conditions, if they can be satisfied by any α, can also be satisfied by a countable α. The mechanism for making this precise will be furnished by the Löwenheim-Skolem theorem which allows us to construct smaller sets having the same properties as larger sets.

There is an extra twist: applying the Löwenheim-Skolem theorem (or the Tarski-Vaught theorem, as S&F prefer to call it) doesn't quite finish the job. Suppose that L_β contains $\mathcal{P}^L(\omega)$. The Löwenheim-Skolem theorem spits out a superset S of $L_{\omega+1}$ "containing all constructible subsets of ω". The phrase in quotes means that for every constructible subset of ω, $(S, \in)$ contains (as an element) a set that *plays the role* of the subset. More precisely, if we have $a \in \mathcal{P}^L(\omega)$, then there is an $a' \in S$ that satisfies the same formulas over S that a satisfies over L.

As a consequence, S may not be transitive. It may contain ordinals of extravagantly high rank, higher than is really needed to construct "avatars" for all those subsets. The Mostowski collapsing lemma eliminates the chaff, leaving us with a lean mean $\alpha \leq \beta$, and an L_α of cardinality $\aleph_1$.

By the way, what S&F call the Gödel isomorphism theorem (Theorem 2.1, p.196) is called the Gödel condensation lemma by everyone else.

21 Forcing

Forcing in set theory comes encrusted with technicalities; the most "barnacle-free" setting for forcing is Peano arithmetic. The Logic notes treat this in §10, where you'll find any motivation or intuition I have to offer on forcing. I'll refer freely to §10 of the Logic notes (L/F for short) in this section.

Let's say then we're ready to tackle the classic Cohen results. I don't care for the treatment in S&F (see §22). So what should one read for forcing in set theory? At one time, I'd have said Shoenfield's paper [23], hands down. Now I also like Halbeisen's book [10]. Halbeisen provides a little more hand-holding, but the forcing stuff occurs in part III. You *can* plunge *in medias res* with some judicious skipping, but of course Shoenfield's treatment is more self-contained. Halbeisen and Shoenfield also differ on various technical details. So pros and cons. These notes will draw on both. Halbeisen's *Combinatorial Set Theory* doesn't have exercises *per se*, but he often leaves proofs to the reader; I will provide some of these.

One notational difference: Halbeisen adopts the so-called Jerusalem convention for conditions, where $p \geq q$ means that p extends q; Shoenfield uses the opposite American convention. I prefer the Jerusalem convention, and will use it throughout, silently transposing Shoenfield—even when quoting him directly.

A few generalities, and some historical background, before we get down to brass tacks. For his first application of forcing, Cohen constructed a model of ZF with a non-constructible subset of ω. He started with a countable model of ZFC (call it M; see footnote[10]), and threw in a new set $G \subseteq \omega$; this exists because M is countable and the true power set of ω isn't. As mentioned in L/F§10.7, the axioms of ZF then insist that you add of lots of

[10] Strictly speaking I should write $\mathbf{M}$ instead of M, and $\mathbf{M}[G]$ instead of $M[G]$. But these are standard models, i.e., $\in$ is the standard element relation, so I'm going with the less fussy notation.

other sets—sets like (G, G) and $G \times G$, to give two very simple examples. Write N for the new model of ZF generated this way (skipping over mounds of detail).

At first you'd think that's it—L relativized to M is contained in M and so doesn't contain G. So G is non-constructible in N, and we have our non-constructible subset of ω. But the problem of "hidden knowledge" rears its head. M doesn't contain G, but maybe that's because M doesn't contain all ordinals. As we know, $L = \bigcup_{\alpha\in\Omega} L_\alpha$, where Ω is the class of *all* ordinals. It's entirely possible that G belongs to some L_α with α an ordinal not in M. Let's say this ordinal α is somehow encoded inside G—by now we've seen a lot of clever ways to encode things. If that happens, then G may end up being "constructible in N".

To circumvent this problem, Cohen devised the notion of a generic set. He gave a construction of a set $M[G]$ containing both M and G. When G is generic, he showed that $M[G]$ is a model of ZF containing exactly the same ordinals as M, thus exactly the same constructible sets (since L_α is absolute for all α).

To define $M[G]$, Cohen adapted Gödel's definition of L. As noted in L/F§10.7, this swirled all the complexities of L and forcing together. Later work simplified matters, essentially replacing the L_α hierarchy with the V_α's. We'll look further into this in §21.11.

21.1 Useful Notions and Notations

Let P be a notion of forcing, i.e., a set P with a preorder $\leq$. (Halbeisen uses $\mathbb{P}$ for $(P, \leq)$, but I won't bother.) Let $p \in P$. The **downward cone** of p is $\wedge p = \{r : r \leq p\}$, and the **upward cone** is $\vee p = \{r : p \leq r\}$. The **downward closure** of a set X (denoted $\wedge X$) is the union of the downward cones of elements of X, and likewise for **upward closure** (denoted $\vee X$).

We restate some of the key concepts with these notions.

- $X \subseteq P$ is **open** iff it is a union of upward cones, i.e., contains the upward cone of each of its elements. So X is open iff it equals its upward closure.
- $X \subseteq P$ is **dense** iff it intersects every upward cone, which is true iff it intersects every nonempty open set.
- Recall that a filter is an $F \subseteq P$ that is downward closed and upward directed. F is downward closed iff it equals its downward closure, i.e., contains the downward cone of each of its elements, i.e., is a union of downward cones. F is upward directed iff the upward cones of any two of its elements intersect. Recall that p and q are compatible if their upward cones intersect (so they have a common extension), incompatible otherwise: $p \perp q$ iff $\vee p \cap \vee q = \varnothing$.
- Finally, recall that a filter is generic over $\mathbf{V}$ iff it intersects every dense open set belonging to $\mathbf{V}$. (Here $\mathbf{V}$ is a standard model of ZF.)

21.2 The Model $M[G]$

(Shoenfield §4; Halbeisen Ch.15, pp.341–343)

Forcing in set theory starts with a countable standard model M of ZF and a partial order P in M. (That is, both the set P and the order relation belong to M.) P is called a **notion of forcing**. Let $G \subseteq P$ be a generic filter. Next, a model $M[G]$ of ZF is constructed; $M[G]$ contains both M and G. With various choices for P, one can achieve all sorts of effects for $M[G]$.

To construct $M[G]$, we first define so-called **names**. The names belong to M, but they name elements of $M[G]$. This allows us to “talk about $M[G]$

inside M". A subtle but crucial aspect of forcing: we can *express facts* about $M[G]$ inside M, i.e., without knowing what G is; but we can *decide their truth values* only when we know[11] what G is.

Various approaches exist for the construction of names, but the simplest is surely Shoenfield's [23]. A name is a set of pairs $\langle a, p\rangle$, where p is a condition and a is a name. In other words, the following transfinite induction:

$$\begin{aligned} V_0^P &= \varnothing \\ V_{\alpha+1}^P &= \mathcal{P}(V_\alpha^P \times P) \\ V_\lambda^P &= \bigcup_{\alpha<\lambda} V_\alpha^P \\ V^P &= \bigcup_{\alpha\in\Omega^M} V_\alpha^P \end{aligned}$$

where Ω^M is the set of all ordinals belonging to M, i.e., $\Omega \cap M$. (Ω^M is a set in V, but a proper class according to M. Likewise, V^P is a subset of M, but a proper class according to M.)

Once G has been chosen, we define a mapping $K_G : V^P \to M[G]$ via induction on rank:

$$K_G(b) = \{K_G(a) : (\exists p \in G)\langle a, p\rangle \in b\}$$

In other words, a pair $\langle a, p\rangle$ says "you should put $K_G(a)$ into $K_G(b)$ if condition p is satisfied by G".

Shoenfield also writes $\bar{a}$ for $K_G(a)$ (following Cohen). Halbeisen adopts the distracting notion $\underset{\sim}{c}$ for names (i.e., elements of V^P), and writes $\underset{\sim}{c}[G]$ for Shoenfield's $K_G[a]$. I will stick with Shoenfield's $K_G(a)$ and $\bar{a}$.

Shoenfield actually doesn't explicitly reference the V_α^P hierarchy, instead applying the recursive definition of K_G to all of M. As he points out, this

[11] Here, "know" and "decide" are used in the "omniscient being" sense—nothing is implied about computability.

is nearly an example of a Mostowski collapsing map. Define $\in_G$ by

$$a \in_G b \Leftrightarrow (\exists p \in G)\langle a, p\rangle \in b$$

Applying the Mostowski collapsing technique to the structure $(M, \in_G)$ produces $(M[G], \in)$. That is, the nonstandard $\in_G$ relation on M becomes the standard $\in$ relation on $M[G]$. I said "nearly" above because the $\in_G$ relation isn't as a rule extensional, so the map $a \mapsto \bar{a}$ isn't an isomorphism. Note that in the definition of $K_G(b)$, if b contains elements not of the form $\langle a, p\rangle$, they are simply ignored by K_G. Such elements are "chaff", "non-coding DNA".

$M[G]$ is defined as:

$$M[G] = \{K_G(a) : a \in M\} = \{K_G(a) : a \in V^P\}$$

For each element $a \in M$, we construct a **canonical name**; I'll denote it by $\dot{a}$, instead of Halbeisen's $\underset{\cdot}{c}$ or Shoenfield's $\hat{a}$. (Halbeisen also uses his dot notation for other purposes, but I won't.) Assume P has a (unique) smallest condition $\emptyset$. (This was the empty bitstring in L/F.) Inductively define

$$\dot{b} = \{\langle \dot{a}, \emptyset\rangle : a \in b\}$$

("Unconditionally put $\bar{a}$ into $\bar{b}$ for all $a \in b$.") An easy induction shows that $K_G(\dot{a}) = a$ for all $a \in M$.

For G, we have the name:

$$\widehat{G} = \{\langle \dot{p}, p\rangle : p \in P\}$$

It is easy to see that $K_G(\widehat{G}) = G$. $\widehat{G}$ is a kind of "diagonal". Note that $\widehat{G}$ belongs to M, even if G doesn't: the definition of $\widehat{G}$ makes no explicit mention of G. It's the K_G mapping (which does know about G) that turns $\widehat{G}$ into G.

So $M[G]$ contains all of M and also contains G. In fact, any model containing M and G must contain $M[G]$ (exercise, or see Shoenfield p.365); in this sense, $M[G]$ is a little like an algebraic closure.

21.3 Forcing for Atomics: Shoenfield

(Shoenfield §5)

We now dive deeper into the pivotal task: defining $p \Vdash^* \varphi$, where φ is a closed formula in the so-called **forcing language**. This is just the language of ZF augmented with all the names we introduced in the previous section. In other words, we allow any element of V^P to occur in φ, treating a as if it were a constant denoting $\bar{a}$. (As noted above, we can even allow any $a \in M$, since Shoenfield defines $\bar{a}$ for all $a \in M$.)

Notation:

$$G \models \varphi(a_1, \ldots, a_k) \Leftrightarrow M[G] \models \varphi(\bar{a}_1, \ldots, \bar{a}_k)$$

(Shoenfield writes $\vdash_G \varphi$ for our $G \models \varphi$.) On the left hand side, we have a formula of the forcing language; on the right hand side, the standard notion of satisfaction for the language $\mathcal{L}_{M[G]}$, where elements of $M[G]$ are allowed as names representing themselves.

L/F§10.1 gave the inductive clauses for $p \Vdash^* \varphi \vee \psi$, $p \Vdash^* \neg\varphi$, and $p \Vdash^* (\exists x)\varphi(x)$; L/F§10.2 employed these in the proofs of the forcing lemmas. All this carries over almost without change. The differences crop up at the ground level. Instead of atomic statements $s = t$ and $G(n)$ (s and t closed terms, n a numeral), we have $a \in b$ and $a = b$, with a and b elements of V^P. Dealing with forcing for atomic statements demands more effort than any other part.

First, though, a few minor points. We define $p \Vdash^* (\exists x)\varphi(x)$ iff $p \Vdash^* \varphi(a)$ for some name a, since names in ZF play the role of numerals in PA. The Definability and Empty Condition Lemmas of L/F§10.2 take different forms.

L/F presented a picture (or better, a video) of forcing: a sequence of ever-longer finite bitstrings $p_0, p_1, \ldots$, whose union becomes the characteristic function of the generic set G. We can still reap valuable intuition from this imagery in the ZF setting. The simplest kind of ZF forcing essentially uses the same conditions[12], resulting in a non-constructible subset G of ω. In PA forcing, G was all we needed; here, we also get $K_G(a)$ for any name a. Each element $\langle a, p\rangle$ of b says, "If p occurs in the sequence $p_0, p_1, \ldots$, put $K_G(a)$ into $K_G(b)$".

Forcing for $a \in b$ asks whether $\bar{a} \in \bar{b}$, likewise for $a = b$. Let's take a simple example. Let p be a condition. Let a and b be names for two finite subsets of ω, say

$$a = \dot{3}$$
$$b = \{\langle \dot{3}, p\rangle\}$$

Do we have $K_G(a) \in K_G(b)$? Yes if and only if $p \in G$, because in that case $3 \in \bar{b}$. Now let's modify b a little, letting c be another name:

$$a = \dot{3}$$
$$c = \{\langle \dot{0}, p\rangle, \langle \dot{1}, p\rangle, \langle \dot{2}, p\rangle\}$$
$$b = \{\langle c, \emptyset\rangle\}$$

We are guaranteed that $\bar{a} = 3$ and that $\bar{c} \in \bar{b}$, but $\bar{c} = 3$ iff $p \in G$ (recalling that $3 = \{0, 1, 2\}$).

Reflecting on these two examples, we see that

$$K_G(a) \in K_G(b) \Leftrightarrow \text{ for some } c \text{ with } c \in_G b, K_G(a) = K_G(c).$$

So elementhood in $M[G]$ depends on equality in $M[G]$, but with names of lower rank (i.e., $\operatorname{rk} c < \operatorname{rk} b$).

How about $K_G(a) = K_G(b)$, or to wring more out of the previous example, $K_G(a) = K_G(c)$? We know that 0, 1, and 2 are elements of 3, but they belong to $K_G(c)$ iff $p \in G$. In general, $K_G(a) = K_G(b)$ iff $K_G(a) \subseteq K_G(b)$ and $K_G(a) \supseteq K_G(b)$, and:

$$K_G(a) \subseteq K_G(b) \Leftrightarrow K_G(c) \in K_G(b) \text{ for all } c \in_G a.$$

[12]Cohen's conditions were actually functions $f : S \to \{0, 1\}$, with S any finite subset of ω.

So equality in $M[G]$ depends on elementhood in $M[G]$, but again with names of lower rank (since $\operatorname{rk} c < \operatorname{rk} a$ here).

We are almost ready to give definitions for $p \Vdash^* a \in b$ and $p \Vdash^* a = b$, but we have one more matter to address: the way the p_n sequence "grows" into G. Start by supposing $G \models a \in b$ (which, recall, means $M[G] \models \bar{a} \in \bar{b}$). To preserve the truth of the Truth Lemma, we need $p \Vdash^* a \in b$ for some $p \in G$. Say $G \models a = c$ with $c \in_G b$. Induction on rank will, hopefully, give us $p \Vdash^* a = c$ for some $p \in G$, and $c \in_G b$ means that $\langle c, q\rangle \in b$ for some $q \in G$. Let r be a common extension of p and q. With the Extension Lemma, we should have $r \Vdash^* a = c$, but of course we won't necessarily have $\langle c, r\rangle \in b$. This, plus the notation $a \in_G b$, motivates the following two notations:

$$a \in_p b \Leftrightarrow \langle a, p\rangle \in b$$
$$a \in_{\wedge p} b \Leftrightarrow (\exists q \leq p)\langle a, q\rangle \in b$$

In other words: $a \in_p b$ means that there is an "instruction" in b saying, "if p is satisfied by G, then put $K_G(a)$ into $K_G(b)$". While $a \in_{\wedge p} b$ says, "if any condition immediately implied by p is satisfied by G, then put $K_G(a)$ into $K_G(b)$".[13]

We see that a kind of Extension Lemma holds, trivially, for $\in_{\wedge p}$: if $a \in_{\wedge p} b$ and $p' \geq p$, then $a \in_{\wedge p'} b$. Observe that the quantifier $(\exists c \in_{\wedge p} b)$ implicitly contains another quantifier: $(\exists q \leq p)(\exists c \in_q b)$. Fully written out, $(\exists c \in_{\wedge p} b) p \Vdash^* a = c$ would read

$$(\exists q)(\exists c)[q \leq p \ \wedge \ \langle c, q\rangle \in b \ \wedge \ p \Vdash^* a = c];$$

reason enough for the new notation. (If you take a moment, you'll also see that negations pass through as usual: $\neg(\exists c \in_{\wedge p} b)$ iff $(\forall c \in_{\wedge p} b)\neg$.)

[13]Another way to think about $a \in_{\wedge p} b$: suppose we "complete" b to b^* by throwing in any pair $\langle a, p\rangle$ whenever there is a $q \leq p$ with $\langle a, q\rangle \in b$. Because G (in the general case) is a filter and hence closed downwards, we see that $K_G(b^*) = K_G(b)$. Also $a \in_p b^*$ iff $a \in_{\wedge p} b$.

We can subsume all three notations, $a \in_G b$, $a \in_p b$, and $a \in_{\wedge p} b$, in the single notation $a \in_S b$ where $S \subseteq P$ is any set of conditions: $a \in_S b$ means $(\exists p \in S) a \in_p b$. As one more instance of this, $a \in_P b$ means that a is a potential element of b, in the sense that there is *some* condition that says a should be in b. Following Shoenfield, I'll write $\mathrm{Ra}(b) = \{a : a \in_P b\}$ for the set of potential elements of b. Example: $\mathrm{Ra}(\dot\omega) = \{\dot n : n \in \omega\}$, or more generally, $\mathrm{Ra}(\dot b) = \{\dot a : a \in b\}$.

Before plunging into details, I want to step back and repeat the key points with more hand waving. The forcing language allows denizens of M to *discuss* $M[G]$, but only "generically"—lacking specific knowledge of G, they can't say which statement hold. Even so, they know that G is generic, enabling them to conclude some facts (those forced by the empty condition).

Names constitute the foundation layer of the forcing language, and K_G assigns meaning to names. By the definition of K_G, something gets placed in $K_G(b)$ *only* by triggering an instruction $c \in_p b$; consequently, *facts about a name are determined by facts about names of lower rank.* This enables the inductive definitions of (strong) forcing for atomic sentences. (See Halbeisen pp.339 and 344 for similar comments.)

Shoenfield does something clever: he makes $a \neq b$ fundamental and $a = b$ derived, i.e., an abbreviation for $\neg a \neq b$. Why? Well, $G \models a = b$ iff $(\forall c)[G \models c \in a \leftrightarrow c \in b]$; universal quantifiers do not play as nicely with strong forcing as existential ones. (I'll say more when we reach the Truth Lemma.)

Here are the definitions (using $c \notin b$ as an abbreviation for $\neg c \in b$):

$$\begin{aligned}
&\text{(a) } p \Vdash^* a \in b : (\exists c \in_{\wedge p} b) p \Vdash^* a = c \\
&\text{(b) } p \Vdash^* a \neq b : (\exists c \in_{\wedge p} a) p \Vdash^* c \notin b \text{ or} \\
&\qquad\qquad\qquad\qquad (\exists c \in_{\wedge p} b) p \Vdash^* c \notin a
\end{aligned}$$

This is a definition by double transfinite induction, and the details are a bit tricky. The case $a \in b$ reduces to cases $a = c$ with $\mathrm{rk}\, c < \mathrm{rk}\, b$; the case $a \neq b$ reduces to cases $c \notin b$ with $\operatorname{rk} c < \operatorname{rk} a$, plus cases $c \notin a$ with $\operatorname{rk} c < \operatorname{rk} b$. Suppose we label the cases $a \in b$, $a \notin b$, $a \neq b$, and $a = b$ with the pair of ordinals $(\operatorname{rk} a, \operatorname{rk} b)$. Using $\succ$ to mean "reduces to", we can write symbolically:

$$\begin{aligned} \text{For } \in\text{:}\ (\alpha,\beta) &\succ (\alpha,\gamma) \text{ with } \gamma < \beta \\ \text{For } \neq\text{:}\ (\alpha,\beta) &\succ (\gamma,\beta) \text{ with } \gamma < \alpha \text{ plus} \\ &\succ (\gamma,\alpha) \text{ with } \gamma < \beta \end{aligned}$$

The cases $a = b$ and $a \notin b$ reduce to the unnegated forms via the definition of forcing for negations. Now if $\succ$ (or rather its reverse, $\prec$) were a well-ordering, we'd be golden. There *is* a standard well-ordering on $\Omega \times \Omega$, namely lexicographic:

$$(\alpha,\beta) \prec (\gamma,\delta) \Leftrightarrow \begin{cases} \beta < \delta \text{ or} \\ \beta = \delta \text{ and } \alpha < \gamma \end{cases}$$

This works for the first two cases, but not the third: $(\gamma,\alpha) \prec (\alpha,\beta)$ does *not* follow from $\gamma < \beta$.

We resolve this by well-ordering *unordered pairs* of ordinals. We label $a \in b$ and $a = b$ with $\{\operatorname{rk} a, \operatorname{rk} b\}$. For an unordered pair $\{\alpha,\beta\}$, let $s\{\alpha,\beta\}$ be the corresponding ordered pair (α',β') arranged so that $\alpha' \leq \beta'$. In other words,

$$s\{\alpha,\beta\} = \begin{cases} (\alpha,\beta) \text{ if } \alpha \leq \beta \\ (\beta,\alpha) \text{ if } \beta < \alpha \end{cases}$$

Define $\{\alpha,\beta\} \prec \{\gamma,\delta\}$ iff $s\{\alpha,\beta\} \prec s\{\gamma,\delta\}$ according to lexicographic order. Note that s maps the collection of unordered pairs of ordinals bijectively onto a subset of $\Omega \times \Omega$, namely all pairs (α,β) with $\alpha \leq \beta$. Any subclass of a well-ordered class is well-ordered under the same ordering, so we have a well-ordering of the class of unordered pairs of ordinals.

It remains to show that $\{\alpha,\gamma\} \prec \{\alpha,\beta\}$ if $\gamma < \beta$, and $\{\gamma,\beta\} \prec \{\alpha,\beta\}$

if $\gamma < \alpha$. This is an easy exercise. (The pesky "third case" becomes $\{\gamma, \alpha\} = \{\alpha, \gamma\} \prec \{\alpha, \beta\}$.)[14]

The Extension Lemma. As noted in L/F§10.2, to prove this, we mostly just turn the induction crank on formula complexity. Recall how this goes for negation: if $q \geq p \Vdash^* \neg\varphi$, then no extension of p strongly forces φ, so *a fortiori* no extension of q strongly forces φ, so $q \Vdash^* \neg\varphi$. Thus we have the lemma for $a = b$ and $a \notin b$ without further ado. Atomic sentences fall out easily from the "Extension Lemma" for $\in_{\wedge p}$. If $q \geq p \Vdash^* a \in b$, then $(\exists c \in_{\wedge p} b) p \Vdash^* a = c$, so $(\exists c \in_{\wedge q} b) q \Vdash^* a = c$, so $q \Vdash^* a \in b$. The proof for $a \neq b$ is similar. (No transfinite induction on rank!)

Next, the Definability Lemma. This says: if $\varphi(\vec{x})$ is a formula in the forcing language, then there is a formula $\text{FORCES}_\varphi(y, \vec{x})$ such that

$$p \Vdash^* \varphi(\vec{a}) \Leftrightarrow M \models \text{FORCES}_\varphi(p, \vec{a})$$

for all conditions p and $\vec{a}$ belonging to M. As a consequence, for any formula $\varphi(\vec{x})$ and any $\vec{a}$ belonging to M, the set

$$\{p \in P : p \Vdash^* \varphi(\vec{a})\}$$

is a set in M—it's a subset of the set P of conditions defined by a formula, so the Separation Axiom applies.

Shoenfield disposes of the Definability Lemma in two sentences (p.363). I'll add a remark, relating to the discussion in the Logic notes on definability of truth in ZF (§9.3), and of forcing in PA (L/F§10.2). We saw that logic walks a tightrope, with Tarski's undefinability theorem on one side and Tarski's definition of truth on the other. When quantifiers range over the entire domain (as is the case here for $\vec{a}$), we can't write a single formula

[14] To compare with Shoenfield, he gets by with a simpler transfinite induction on $\max(\text{rk}\, a, \text{rk}\, b)$. He pays for this by having to unwind the mutual recursion between $\neq$ and $\in$, complicating matters. (I am clear on how this goes on for the definition, but not for the proof of the Truth Lemma.)

TRUE giving truth for all formulas. The same holds for forcing. But various compromises work, one being a sequence of formulas TRUE_n for truth up to complexity n. It's pretty much the same story for forcing. That's the approach taken here.

The Definability Lemma implies that for any closed formula φ, this subset D_φ of P is a set in M:

$$D_\varphi = \{p \in P : p \Vdash^* \varphi \text{ or } p \Vdash^* \neg\varphi\}$$

D_φ is dense by the definition of strong forcing for negations. Let $\mathcal{D}$ be the set of all dense sets of condition belonging to M, i.e., $\mathcal{D} = \{D \subseteq P : D \in M \text{ and } D \text{ is dense}\}$. Shoenfield calls a filter ""generic over M" if it is generic for $\mathcal{D}$ (p.360).

We now have in hand all the key ingredients for this fact from L/F§10.4:

> If G is a generic filter for $\mathcal{D}$, then $G \Vdash^* \neg\varphi$ iff not $G \Vdash^* \varphi$.

(Here as in L/F, I write $G \Vdash^* \varphi$ to mean $(\exists p \in G) p \Vdash^* \varphi$.) This is all we need for one inductive step of the Truth Lemma:

$$G \models \neg\varphi \Leftrightarrow \text{ not } G \models \varphi \Leftrightarrow \text{ not } G \Vdash^* \varphi \Leftrightarrow G \Vdash^* \neg\varphi$$

This also tells us that if the Truth Lemma holds for $a \in b$ (for a specific a and b), then it holds for $a \notin b$; likewise for $a \neq b$, $a = b$.

The Truth Lemma for $a \in b$ and $a \neq b$ depends on a double transfinite induction, just like the definitions of $p \Vdash^* a \in b$ and $p \Vdash^* a \neq b$. We make use of the observation in the last paragraph. First, $a \in b$:

$$\begin{aligned} G \models a \in b &\Leftrightarrow (\exists c \in_G b) G \models a = c \\ &\Leftrightarrow (\exists c \in_G b) G \Vdash^* a = c \text{ (induction)} \\ &\Leftrightarrow (\exists q \in G)(\exists c \in_q b)(\exists p \in G) p \Vdash^* a = c \\ &\Leftrightarrow (\exists r \in G)(\exists c \in_{\wedge r} b) r \Vdash^* a = c \text{ (see below)} \\ &\Leftrightarrow (\exists r \in G) r \Vdash^* a \in b \text{ (definition)} \\ &\Leftrightarrow G \Vdash^* a \in b \end{aligned}$$

In the middle step, we let r be a common extension of p and q, and appeal to the Extension Lemmas.

Next, $a \neq b$:

$$\begin{aligned} G \models a \neq b &\Leftrightarrow (\exists c \in_G a) G \models c \notin b \text{ or vice versa} \\ &\Leftrightarrow (\exists c \in_G a) G \Vdash^* c \notin a \text{ or v.v.} \\ &\Leftrightarrow (\exists q \in G)(\exists c \in_q a)(\exists p \in G) p \Vdash^* c \notin b \text{ or v.v.} \\ &\Leftrightarrow (\exists r \in G)(\exists c \in_{\wedge r} b) r \Vdash^* c \notin b \text{ or v.v.} \\ &\Leftrightarrow (\exists r \in G) r \Vdash^* a \neq b \\ &\Leftrightarrow G \Vdash^* a \neq b \end{aligned}$$

If we had made $a = b$ fundamental, then we'd be dealing with "$(\forall c \in_q a)$" instead of the existential quantifier; this messes up the argument. In fact, a simple negation of the definition of $p \Vdash^* a \neq b$ is not what we want—that says "not $p \Vdash^* a \neq b$", not the same as "$p \Vdash^* \neg a \neq b$".

The rest of the section shows that the forcing lemmas hold also for weak forcing ($\Vdash$), and demonstrates what I've dubbed the Fundamental Property of Weak Forcing: $p \Vdash \varphi$ iff $G \models \varphi$ for all $G \ni p$. The proofs are essentially the same as in L/F§10.5.

Finally, the Empty Condition Lemma. (Shoenfield omits this, but it's occasionally nice to have.) Suppose $\varphi(x_1, \ldots, x_k)$ is formula in $L(\mathrm{ZF})$

where $x_1, \ldots, x_k$ are all the free variables. Also assume all quantifiers are "dotted", i.e., of the form $(\exists \dot{x})$ or $(\forall \dot{x})$.[15] Then $\emptyset \Vdash^* \varphi(\dot{a}_1, \ldots, \dot{a}_k)$ iff $M \models \varphi(a_1, \ldots, a_k)$, for any $a_1, \ldots, a_k \in M$. Sketch of proof: the pivotal fact is that any such $\varphi(\dot{a}_1, \ldots, \dot{a}_k)$ enjoys the "all or nothing" property: either all conditions strongly force it, or no conditions do. The proof is pretty much the same as in L/F§10.2, except for the atomic formulas $\dot{a} \in \dot{b}$ and $\dot{a} \neq \dot{b}$ and their negations. For that, we transfinitely induct on the unordered pairs of ranks $\{\operatorname{rk} a, \operatorname{rk} b\}$, as above. Key (but trivial) observation: the only pairs in $\dot{a}$ are of the form $\langle \dot{c}, \emptyset \rangle$ with $c \in a$, and likewise for $\dot{b}$. For example, in the inductive clause

$$p \Vdash^* \dot{a} \in \dot{b} \Leftrightarrow (\exists c \in_{\wedge p} \dot{b}) p \Vdash^* c = \dot{a}$$

we can replace the right hand side with

$$(\exists c \in b) p \Vdash^* \dot{c} = \dot{a}$$

and similarly for the defining clauses of $\dot{a} \neq \dot{b}$.

We need dotted quantifiers because the inductive clause

$$p \Vdash^* (\exists x)\varphi(x) \Leftrightarrow (\exists c \in M) p \Vdash^* \varphi(c)$$

results in a formula $\varphi(c)$ not necessarily of the right form.

The Fundamental Property of Weak Forcing furnishes a very simple proof of the Empty Condition Lemma, with $\Vdash^*$ replaced by $\Vdash$. Namely, $M[G] \models \varphi(\dot{a}_1, \ldots, \dot{a}_k)$ iff $M \models \varphi(a_1, \ldots, a_k \in M)$; this is because $K_G(\dot{a}) = a$, and the dotted quantifiers make the quantified variables range over M.

21.4 Halbeisen/Kunen Forcing

(Halbeisen Definition 15.8, p.348)

[15] These are vernacular, to be rewritten as $(\exists y)(\exists x)[y = \dot{x} \wedge \ldots]$ and $(\forall y)(\forall x)[y = \dot{x} \rightarrow \ldots]$, replacing all occurrences of $\dot{x}$ with y in $\ldots$.

Halbeisen (following Kunen [17]) defines $p \Vdash \varphi$ directly, without using strong forcing as an intermediate. (L/F§10.5 listed some reasons people prefer weak to strong forcing.) This demands hairier definitions (and proofs) for $p \Vdash a = b$ and $p \Vdash a \in b$. Here I will show the equivalence of the Shoenfield and Kunen definitions. Let's write (in this section only!) $\Vdash_{\mathrm{K}}$ for the Kunen definition, and $\Vdash_{\mathrm{S}}$ for Shoenfield's (weak) forcing, i.e., $\Vdash^* \neg\neg$. (Strong forcing will always be the Shoenfield definition.) Our goal: $p \Vdash_{\mathrm{K}} \varphi$ iff $p \Vdash_{\mathrm{S}} \varphi$.

Let's start with $p \Vdash_{\mathrm{S}} a = b$. Bear in mind that weak and strong forcing coincide for negated statements (see L/F§10.5). So $p \Vdash_{\mathrm{S}} a = b$ iff $p \Vdash^* a = b$ (since $a = b$ is really $\neg a \neq b$). We unwind $p \Vdash^* a = b$ for Shoenfield:

$$\begin{aligned}
&p \Vdash^* a = b \\
\Leftrightarrow &p \Vdash^* \neg a \neq b \\
\Leftrightarrow &(\forall r \geq p) \neg r \Vdash^* a \neq b \\
\Leftrightarrow &(\forall r \geq p) \neg \{[(\exists c \in_{\wedge r} a) r \Vdash^* c \notin b] \text{ or } [\text{vice versa}]\} \\
\Leftrightarrow &(\forall r \geq p) \{[(\forall c \in_{\wedge r} a) \neg r \Vdash^* \neg c \in b] \text{ and } [\text{v.v.}]\} \\
\Leftrightarrow &(\forall r \geq p) \{[(\forall c \in_{\wedge r} a) \neg(\forall q \geq r) \neg q \Vdash^* c \in b] \text{ and } [\text{v.v.}]\} \\
\Leftrightarrow &(\forall r \geq p) \{[(\forall c \in_{\wedge r} a)(\exists q \geq r) q \Vdash^* c \in b] \text{ and } [\text{v.v.}]\} \\
\Leftrightarrow &(\forall r \geq p) \{[(\forall c \in_{\wedge r} a)(\exists q \geq r)(\exists d \in_{\wedge q} b) q \Vdash^* c = d] \text{ and } [\text{v.v.}]\}
\end{aligned}$$

This defines $p \Vdash_{\mathrm{S}} a = b$ in terms of $q \Vdash_{\mathrm{S}} c = d$ with $\operatorname{rk} c < \operatorname{rk} a$ and $\operatorname{rk} d < \operatorname{rk} b$. So we have a definition by transfinite induction on $\max(\operatorname{rk} a, \operatorname{rk} b)$:

$$\begin{aligned}
&p \Vdash_{\mathrm{S}} a = b \\
\Leftrightarrow &(\forall r \geq p)(\forall c \in_{\wedge r} a)(\exists q \geq r)(\exists d \in_{\wedge q} b) q \Vdash_{\mathrm{S}} c = d \\
&\text{and vice versa}
\end{aligned}$$

(I've distributed the $(\forall r \geq p)$ over the two conjuncts, permissible by the rules of logic.)

Compare with Halbeisen Def.15.8 (p.348). Written out in our notation, this is:

$$\begin{aligned} &p \Vdash_{\mathrm{K}} a = b \\ \Leftrightarrow &(\forall c \in_s a)(\forall r \geq p)(\exists q \geq r)[q \geq s \rightarrow (\exists d \in_{\wedge q} b) q \Vdash_{\mathrm{K}} c = d] \\ &\text{and vice versa} \end{aligned}$$

The two right hand sides share a good deal, but are not identical. The key point: the quantifier "$(\forall c \in_{\wedge r} a)$" disavows any interest in $c \in_s a$ for which $s \not\leq r$; the phrase "$q \geq s \rightarrow$" accomplishes pretty much the same.

We'll show the equivalence by induction on $\max(\operatorname{rk} a, \operatorname{rk} b)$. Suppose $p \Vdash_{\mathrm{S}} a = b$, and let $c \in_s a$ and $r \geq p$ be arbitrary. Two cases: $s \leq r$ and $s \not\leq r$. In the first case, we have $c \in_{\wedge r} a$, so there is a $q \geq r$ such that $(\exists d \in_{\wedge q} b) q \Vdash_{\mathrm{S}} c = d$. By inductive hypothesis, $q \Vdash_{\mathrm{S}} c = d$ iff $q \Vdash_{\mathrm{K}} c = d$. In the second case, let $q = r$; then the antecedent $q \geq s$ is false. Either way, the Kunen version holds.

Suppose $p \Vdash_{\mathrm{K}} a = b$, and let $c \in_{\wedge r} a$ and $r \geq p$ be arbitrary. So there is an $s \leq r$ such that $c \in_s a$, and so if $q \geq r \geq s$, then the antecedent $q \geq s$ holds and we have $(\exists d \in_{\wedge q} b) q \Vdash_{\mathrm{K}} c = d$. Again induction tells us that $q \Vdash_{\mathrm{S}} c = d$ iff $q \Vdash_{\mathrm{K}} c = d$, so the Shoenfield version holds.

Next, $a \in b$. In L/F§10.5, we noted that $p \Vdash \varphi$ iff the set $\{q \geq p : q \Vdash^* \varphi\}$ is dense above p. So $p \Vdash_{\mathrm{S}} a \in b$ iff $\{q \geq p : (\exists c \in_{\wedge q} b) q \Vdash^* a = c\}$ is dense above p. But as we've just seen, $q \Vdash^* a = c$ iff $q \Vdash_{\mathrm{S}} a = c$ iff $q \Vdash_{\mathrm{K}} a = c$, so $p \Vdash_{\mathrm{S}} a \in b$ iff $p \Vdash_{\mathrm{K}} a \in b$.

The rest of the argument inducts on formula complexity. The facts we need are listed in L/F§10.5 (items (1)–(5)).

So with a little work, we've shown the equivalence of weak "Shoenfield" forcing and "Kunen" forcing.

21.5 $M[G]$ Revisited

(Shoenfield §2 and §6; Halbeisen Theorem 15.12 (p.353))

Now that $M[G]$ has been defined, we have to show it satisfies ZF, and ZFC if M does. You won't encounter much difficulty following the proofs in Shoenfield and Halbeisen step-by-step; here I'll outline a perspective that (hopefuly) makes the arguments natural, even inevitable.

Shoenfield gives background on ZF in §2. I'll add two remarks. First, Shoenfield's definition of $V(\alpha)$ is equivalent to this one, preferred by some authors:

$$V(\alpha+1) = S(V(\alpha))$$
$$V(\lambda) = \bigcup_{\alpha<\lambda} V(\alpha),\ \lambda \text{ a limit ordinal}$$

(Also, P for power set has become nearly standard, as opposed to Shoenfield's S.)

Second, note that (b) of Lemma 2.1 is just the Separation Axiom (aka Comprehension Axiom), and (c) is a consequence of the Union and Replacement (aka Substitution) Axioms. It's not hard to recover these two axioms from (b) and (c).

Turn to Shoenfield §6. As always, the contrast between denizens of M and of $M[G]$ looms large. We know ZF holds for M. Denizens of M can talk about $M[G]$ via names and the forcing language, but don't know which conditions actually hold in $M[G]$. This plays out in a strategy I'll dub "select from the warehouse". To show a set $X \subseteq M[G]$ exists as a set in $M[G]$, the M people first assemble an enclosing "warehouse of names". They should be 100% sure that any element of X has a name in the warehouse. Then they put together a "selector", i.e., a collection of pairs $\langle x, p\rangle$ such that G, perusing this collection, will select exactly the names for elements of X. A division of labor: the M denizens don't know the actual contents of X, but they don't have to, because G, who has the inside scoop, does the selecting.

The warehouse addresses the "size" aspects of ZF, insuring that we're not dealing with a proper class. Write w for the warehouse and s for the selector. Both w and s will be sets in M. Proving this for w will employ various arguments, but for s, we will always have $s \subseteq w \times P$, so s at least won't be "too big".

If the selector s is properly designed, we'll have $K_G(s) = X$. Here is a suggestive, if not quite meaningful, string of equivalences:

$$(\forall x \in w)[G \models x \in X \Leftrightarrow G \Vdash x \in X \Leftrightarrow (\exists p \in G)p \Vdash x \in X]$$

We are sure that every element of X has a name in w, so this (schematic) definition looks like the right tool for the job:

$$s = \{\langle x, p\rangle \in w \times P : p \Vdash x \in X\}$$

It remains only to translate "$p \Vdash x \in X$" into something mathematically precise, i.e., a formula in the forcing language. Doing this depends on what X is supposed to be. The Definability Lemma, coupled with the Separation Axiom for M, then shows that s is actually a set in M.

Let's see how this unfolds in Shoenfield §6. In Lemma 6.1, A corresponds to X, and w to $\mathrm{Ra}(a)$. Since $A \subseteq \bar{a}$ by assumption, each element of A belongs to $\bar{a}$ and was put there via a pair $\langle x, p\rangle \in a$, thus $x \in \mathrm{Ra}(a)$.

A being a class in $M[G]$ means that $\bar{x} \in A \Leftrightarrow G \models \Phi(x)$ for some formula Φ, so our schematic $p \Vdash x \in X$ becomes $p \Vdash \Phi(x)$. So we have our selector (denoted c by Shoenfield). Lemma 6.1 now follows cleanly from the "select from the warehouse" strategy, indeed looks like its embodiment.

Lemma 6.2 says that if $a \in M$ is a warehouse for $x \subseteq M[G]$, then $K_G(a \times \{1\})$ contains x. (Shoenfield's 1 is our $\emptyset$, the empty condition.) This is virtually the definition of a warehouse.

Next, Replacement. Shoenfield's clause (c) differs somewhat from the Replacement Axiom. On the one hand, (c) considers $\bigcup_{b\in F''(a)} b$ rather than $F''(a)$; on the other, (c) says only that this is contained in a set. (Clause (b), the Separation Axiom, then administers the coup de grâce.) Let's see how the argument goes for the usual Replacement Axiom. F is a functional in $M[G]$, which means that for some formula Φ, $F(\bar{x}) = \bar{y}$ iff $G \models \Phi(x, y)$. First we build the warehouse for $F''(\bar{a}) = \{F(\bar{b}) : \bar{b} \in \bar{a}\}$. The M inhabitants must approximate, not knowing when $G \models \Phi(b, c)$. They can be sure that every element of $\bar{a}$ has a name in $\mathrm{Ra}(a)$ (but not conversely). Suppose $b \in \mathrm{Ra}(a)$. They don't know the value of $F(\bar{b})$ in $M[G]$, but at least they know that if $G \models \Phi(b, c)$, then some p forces $\Phi(b, c)$. First idea: take the union of all the sets $\{c : p \Vdash \Phi(b, c)\}$ as p ranges over all conditions and b ranges over all of $\mathrm{Ra}(a)$. Problem: how do we know, for any particular b and p, that $\{c : p \Vdash \Phi(b, c)\}$ isn't "too big", i.e., a proper class? We don't, so we use Scott's trick: set

$$\rho(p, b) = \{c : p \Vdash \Phi(b, c)\ \&\ c \text{ has minimal rank for such a name}\}$$

In other words: If $p \Vdash \Phi(b, c)$, then we demand, for c's admission to $\rho(p, b)$, that $\mathrm{rk}\, c \leq \mathrm{rk}\, c'$ for all c' such that $p \Vdash \Phi(b, c')$. This is contained in a level V_α^M for some α, so it isn't "too big", and the Definability lemma shows that it's a set. Using Replacement for M twice shows that

$$w = \bigcup_{p\in P,\ b\in\mathrm{Ra}(a)} \rho(p, b)$$

is a set in M. So w is a warehouse (denoted d by Shoenfield).

That's all we need: Lemma 6.2 says that if we have a warehouse for $F''(\bar{a})$, then $F''(\bar{a})$ is contained in a set. $F''(\bar{a})$ has a defining formula, so the Separation Axiom (or Lemma 6.1) concludes the matter.

Since clause (c) differs from the Replacement Axiom in a couple of ways, Shoenfield's proof has an extra grace note (the transitivity of d), but this shouldn't occasion any hiccups.

Power Set: we need a warehouse for $\mathcal{P}^{M[G]}(\bar{a}) = \mathcal{P}(\bar{a}) \cap M[G]$. Say $\bar{b} \in \mathcal{P}(\bar{a}) \cap M[G]$. Rewriting $\bar{b} \in \mathcal{P}(\bar{a})$ as $\bar{b} \subseteq \bar{a}$, we see that $\bar{b}$ is a class contained in $\bar{a}$: membership in $\bar{b}$ is characterized by the formula "$x \in b$" in the forcing language. So by Lemma 6.1, $\bar{b} = \bar{c}$ with $c \subseteq \mathrm{Ra}(a) \times P$. (Recall that our P is Shoenfield's C.)

This tells us that $\mathcal{P}^M(\mathrm{Ra}(a) \times P)$ is a warehouse for $\mathcal{P}^{M[G]}(\bar{a})$: every element of the latter has a name in the former. Since Power Set holds in M, our warehouse is indeed a set in M. The selector is built as usual, since the formula "$x \subseteq a$" shows that the $M[G]$ power set is a class in $M[G]$. (Or just appeal to Lemma 6.1 once more.)

So far we haven't used the Axiom of Choice. However, if M satisfies AC, so does $M[G]$. I have nothing to add to Shoenfield's proof of this.

The treatment in Halbeisen is pretty much the same story, although he provides more details, and explicitly addresses all the ZF axioms.

21.6 Not CH

(Shoenfield §§2,7,10–11;
Halbeisen Lemma 14.3 & Coro.14.4 (pp.326–327), Ch.15 pp.360–366)

Herewith an outline of the flow of the argument, highlighting the key points.

No new ordinals or constructible sets

$M[G]$ has exactly the same ordinals as M. This falls out almost trivially from the construction, although it was a prime motivation for Cohen's concept of a generic set—the "hidden knowledge" problem, see p.70 of these notes. The burden has been shifted to showing that $M[G]$ satisfies the axioms.

As an immediate corollary to this plus the absoluteness of the constructible hierarchy, M and $M[G]$ have the same constructible sets.

Adding new subsets of ω

Now the fun begins, using various notions of forcing, i.e., partial orders $(P, \leq)$ of conditions. A so-called Cohen condition is a functions from a finite subset of ω to $\{0,1\} = 2$—really the same as a "formula" condition of L/F§10.1, except in the ZF context instead of PA. The generic set G amounts to a new subset of ω, "new" (i.e., not in M) by a simple density argument: given $X \subseteq \omega$ and a Cohen condition p, we can always extend p to disagree with the characteristic function of X. If X is in M, the set of "disagreeable extensions" is a dense set of conditions belonging to M, so G must contain one of them, and the set determined by G differs from X. By the previous point, the new subset of ω is not constructible in $M[G]$. So $M[G]$ is a model of $V \neq L$, Cohen's first independence result.

Generalizing, for any A, B in M, let $\mathrm{Fn}(A, B)$ the collection of all functions whose domains are finite subsets of A and whose ranges are contained in B. Such a condition gives a finite amount of information about a function $f : A \to B$. Again, simple density arguments show that a generic set determines a function in $M[G]$ from all of A onto B, and this function is not in M. Now let $\mathfrak{n}$ be an infinite cardinal in M, and let our notion of forcing be $\mathrm{Fn}(\mathfrak{n} \times \omega, 2)$. So we've added a function $\mathfrak{n} \times \omega \to 2$ to $M[G]$. That amounts to $\mathfrak{n}$ new subsets of ω—a simple density argument shows these new subsets are all different. That's *almost* a proof that $2^{\aleph_0} \geq \mathfrak{n}$ in $M[G]$, but one obstacle still blocks the conclusion.

Collapsing Cardinals

Let $\mathfrak{n}$ be an infinite cardinal in M. As our notion of forcing, use $\mathrm{Fn}(\omega, \mathfrak{n})$. So in $M[G]$, we have a function from ω *onto* $\mathfrak{n}$. This is called *collapsing cardinals.*

What if $\mathfrak{n}$ is actually uncountable? In that case, M is an uncountable model, and the there is no generic set over $\mathrm{Fn}(\omega, \mathfrak{n})$. The proof of the

Existence Lemma depends critically on the countability hypothesis. (See also Halbeisen Fact 14.2, p.325.)

Assume that M is countable, so $\mathfrak{n}$ is actually countable but just uncountable in M. In $M[G]$, $\mathfrak{n}$ is a countable ordinal (being an ordinal is absolute), i.e., $|\mathfrak{n}|^{M[G]} = \omega$, while $|\mathfrak{n}|^M = \mathfrak{n}$. So "being a cardinal" is not absolute.

Taking $\mathfrak{n} = 2^{\aleph_0}$, we learn that in $M[G]$, there are only countably many M-subsets of ω. Of course, $M[G]$ had a building boom and boasts *lots* of new ω subsets, but it can't have any new *constructible* ones, since M and $M[G]$ have the same constructible sets. So $M[G]$ satisfies "there are only countably many constructible subsets of ω".

Collapsing cardinals thus gives us positive results, but it's a hurdle to negating the Continuum Hypothesis. We saw above how to get an $M[G]$ with $\mathfrak{c}^{M[G]} \geq \mathfrak{n}^M$ (using the traditional notation $\mathfrak{c} = 2^{\aleph_0}$). But we want $\mathfrak{c}^{M[G]} \geq \mathfrak{n}^{M[G]}$. So we need to preserve the cardinal $\mathfrak{n}$.

Preserving Cardinals

We preserve cardinals by preserving cofinalities. The **cofinality** of an ordinal α, $\mathrm{cf}(\alpha)$, is the smallest cardinality of a cofinal subset[16] of α. (There are other equivalent definitions of $\mathrm{cf}(\alpha)$.)

The cofinality function is not absolute, but we do obviously have $\mathrm{cf}^{M[G]}(\alpha) \leq \mathrm{cf}^M(\alpha)$ for all α. It is a theorem that if the cofinality function is preserved ($\mathrm{cf}^{M[G]}(\alpha) = \mathrm{cf}^M(\alpha)$ for all α), then all cardinalities are preserved.[17]

The preservation of cofinalities follows from a certain combinatorial property of the notion of forcing P: the **countable chain condition** (ccc). This says that any antichain in P (set of mutually incompatible conditions) is at

[16]i.e., a $B \subseteq \alpha$ such that for any $\gamma < \alpha$, there is a $\beta \in B$ with $\beta \geq \gamma$.

[17]Halbeisen says this is an equivalence, but that's incorrect. The preservation of cardinalities tells you only that the range of the cofinality function is preserved. As Shoenfield remarks, Prikry forcing provides a counterexample.

most countable. Fortunately, $\text{Fn}(\mathfrak{n} \times \omega, 2)$ (see above) satisfies the ccc, so with this P, we have $\mathfrak{c} \geq \mathfrak{n}$.

We have to prove two things, then: (a) the ccc implies that cofinalities are preserved; (b) $\text{Fn}(\mathfrak{n} \times \omega, 2)$ satisfies the ccc.

(a) The key idea here is *bounded fan-out*, plus the usual trick of playing life in $M[G]$ off against the perspective of folks living in M. Let α be an ordinal in M. In $M[G]$, say $\bar{f} : \text{cf}^{M[G]}(\alpha) \to \alpha$ is bijection mapping $\text{cf}^{M[G]}(\alpha)$ to a cofinal subset of α. Say $\beta \in \text{cf}^{M[G]}(\alpha)$ and say $\bar{f}(\beta) = \gamma$. That is, $G \models f(\dot{\beta}) = \dot{\gamma}$. Thus for some $p \in G$, $p \Vdash f(\dot{\beta}) = \dot{\gamma}$. So while the M folk can't identify the true value of $\bar{f}(\beta)$, they *can* identify a set of *possible values*. The ccc tells us that this set of possible values is *countable*: if $p \Vdash f(\dot{\beta}) = \dot{\gamma}$ and $q \Vdash f(\dot{\beta}) = \dot{\delta}$ with $\gamma \neq \delta$, then p and q force inconsistent statements, and so are incompatible.

That's our bounded fan-out: from one $\beta \in \text{cf}^{M[G]}(\alpha)$, we have at most $\aleph_0$ possible values in α. Now, the set of possible values includes the set of true values, i.e., the actual range of $\bar{f}$, a cofinal subset of α. So in M, this set of possible values is a cofinal subset of α, and therefore has cardinality greater or equal to $\text{cf}^M(\alpha)$. So we have this inequality, holding in $M[G]$ (and for all infinite α):

$$\text{cf}^M(\alpha) \leq \text{cf}^{M[G]}(\alpha) \cdot \aleph_0 = \text{cf}^{M[G]}(\alpha)$$

We already know that $\text{cf}^{M[G]}(\alpha) \leq \text{cf}^M(\alpha)$, so ccc implies that cofinalities are preserved.

Shoenfield proves a more general result, demanding a more delicate argument, but the basic ideas remain the same.

(b) is our next item.

The Countable Chain Condition

Shoenfield and Halbeisen prove the ccc for $\text{Fn}(A, B)$, B any countable set, in (at least superficially) different ways. Shoenfield employs the following strategy: let I be an antichain in $\mathrm{Fn}(A, B)$. Define a sequence $A_0 \subseteq A_1 \subseteq \dots$ of subsets of A via induction so that $I \subseteq \mathrm{Fn}\left(\bigcup A_n, B\right)$ and $\mathrm{Fn}\left(\bigcup A_n, B\right)$ is countable. Observe that if $\bigcup A_n$ is countable, then so is $\mathrm{Fn}\left(\bigcup A_n, B\right)$. Let $A_0 = \varnothing$, so $\mathrm{Fn}(A_0, B)$ consists of just the empty function. For every $p \in \mathrm{Fn}(A_n, B)$, choose an extension $q \in I$, if there is one, and let A_{n+1} be the union of A_n plus the domains of all the chosen q's. It is obvious by the countability of B (and by induction) that all the A_n are countable, so $\bigcup A_n$ is countable.

Now pick a particular $f \in I$ and look at the growing sequence $\mathrm{dom}(f) \cap A_n$. Since $\mathrm{dom}(f)$ is finite, there is an n with $\mathrm{dom}(f) \cap A_n = \mathrm{dom}(f) \cap A_{n+1} = D$ (say). Perhaps $\mathrm{dom}(f) \subseteq A_n$; good, then $f \in \mathrm{Fn}(A_n, B)$. If not, then some other $g \in I$ must have been chosen at the way to extend $f \restriction D$—that is, we must have $f \restriction D = g \restriction D$ (we try to extend the domain of *every* $p \in \mathrm{Fn}(A_n, B)$), and also $\mathrm{dom}(f) \cap \mathrm{dom}(g) = D$ (no growth for $\mathrm{dom}(f) \cap A_k$ going from $k = n$ to $k = n + 1$). So f and g are compatible, since they agree on the overlap of their domains. But this is impossible since I is an antichain. We are left to conclude that $\mathrm{dom}(f) \subseteq A_n$, so $I \subseteq \mathrm{Fn}\left(\bigcup A_n, B\right)$.

(Shoenfield actually proves a more general result, allowing all functions whose domains have cardinality less than a fixed cardinal $\mathfrak{m}$, and requiring $|B| \leq \mathfrak{m}$. He also has a restriction on the cardinal $\mathfrak{m}$. He obtains the so-called $\mathfrak{m}^+$-chain condition instead of ccc. The same strategy is used, but the cardinal computations become more complicated.)

Halbeisen's proof of the ccc depends on the so-called Δ-system lemma[18] (Lemma 14.3, p.326; Coro.14.4, p.327, is the ccc). I have little to add. I found it helpful to regard the uncountable family $\mathcal{E}$ of finite sets as a bipartite graph, with a node $x \in \mathcal{E}$ linked to all its elements in $\bigcup \mathcal{E}$. We have finite fan-out (by hypothesis). The proof splits into two cases. Case 1: by trimming $\mathcal{E}$ to $\mathcal{E}'$, still uncountable, we achieve countable fan-in everywhere. Case 2: no matter how we trim $\mathcal{E}$ to an uncountable $\mathcal{E}'$, we always have

[18] Incidentally, people have started calling this the Sunflower Lemma.

uncountable fan-in at at least one $a \in \bigcup \mathcal{E}'$. Incidentally, the Δ-system lemma can be generalized to larger cardinalities.

Computing $2^{\aleph_0}$

From the above, we now know how to construct an $M[G]$ in which $2^{\aleph_0} \geq \mathfrak{n}$, for any infinite cardinal $\mathfrak{n}$ belonging to M. Can we make this an equality: $M[G] \models 2^{\aleph_0} = \mathfrak{n}$? Since $\left(2^{\aleph_0}\right)^{\aleph_0} = 2^{\aleph_0 \cdot \aleph_0} = 2^{\aleph_0}$, we immediately have the necessary condition $M[G] \models \mathfrak{n}^{\aleph_0} = \mathfrak{n}$. In fact we need $M \models \mathfrak{n}^{\aleph_0} = \mathfrak{n}$, since all functions $\omega \to \mathfrak{n}$ that belong to M also belong to $M[G]$, so $\left(\mathfrak{n}^{\aleph_0}\right)^M \leq \left(\mathfrak{n}^{\aleph_0}\right)^{M[G]}$. We also must have $\mathfrak{n} > \aleph_0$, but this follows from $\mathfrak{n}^{\aleph_0} = \mathfrak{n}$: $\mathfrak{n} = \mathfrak{n}^{\aleph_0} \geq 2^{\aleph_0} > \aleph_0$.

The notion of forcing $P = \mathrm{Fn}(\mathfrak{n} \times \omega, 2)$ gives $M[G] \models 2^{\aleph_0} = \mathfrak{n}$ whenever $M \models \mathfrak{n}^{\aleph_0} = \mathfrak{n}$. Shoenfield and Halbeisen give essentially the same proof, just dressed up differently. We start with Shoenfield's version.

The argument resembles the proof that $\mathrm{cf}^M(\alpha) \leq \mathrm{cf}^{M[G]}(\alpha)$. Once again we contrast what *is* true in $M[G]$ with what the M people *know* to be true. The key inequality to be demonstrated:

$$\underbrace{2^{\aleph_0}}_{\text{in } M[G]} \leq \underbrace{\mathfrak{n}^{\aleph_0} = \mathfrak{n}}_{\text{in } M}$$

Suppose $M[G] \models \bar{a} \subseteq \omega$. For any $n \in \omega$, the M folk approximate the answer to the question, "does n belong to $\bar{a}$?" with the set of conditions

$$\varphi_a(n) \stackrel{\text{def}}{=\!=} \{p : p \Vdash \dot{n} \in a\}$$

Basic idea: cook up an injective mapping from $\mathcal{P}^{M[G]}(\omega)$ to the set $\{\varphi_a : a \in M \,\&\, \bar{a} \subseteq \omega\}$, and then bound the cardinality of that set. The injectivity is is easy: suppose $\varphi_a = \varphi_b$; then $p \Vdash \dot{n} \in a$ iff $p \Vdash \dot{n} \in b$. Now, $G \models \dot{n} \in a$ iff there is a $p \in G$ with $p \Vdash \dot{n} \in a$, likewise for b, thus $G \models \dot{n} \in a$ iff $G \models \dot{n} \in b$.

One small snag: the composition $\bar{a} \mapsto a \mapsto \varphi_a$ requires us to choose a unique name a for each element $\bar{a}$ of $\mathcal{P}^{M[G]}(\omega)$. Now, in proving that Power Set holds in $M[G]$, we showed that power sets have "warehouses". Applied to ω, this says there is a set $W \in M$ that has a name (guaranteed!) for each element of $\mathcal{P}^{M[G]}(\omega)$. Since AC holds in M (by assumption), we can well-order W. Then we can specify $\bar{a} \mapsto a$ uniquely by choosing the first possible name in W.

So we have an injective map from $\mathcal{P}^{M[G]}(\omega)$ to the set $\{\varphi_a : a \in W\}$. We need this to be a function in $M[G]$. The Definability Lemma gives the definability of $a \mapsto \varphi_a$.[19] For $\bar{a} \mapsto a$, we use the fact that $a \mapsto \bar{a}$ is definable in $M[G]$ by an obvious transfinite induction, so $\bar{a} \mapsto a$ is definable (using the well-ordering of W as a parameter). It's easy to see that the domain and range are sets in $M[G]$, so we have a function in $M[G]$.

Up to now the argument has straddled $M[G]$ and M; the rest takes place entirely in M. We want to bound the size of the set $\{\varphi_a : a \in W\}$. Two obvious ideas don't work. A short computation gives $2^{\mathfrak{n}}$ for $|W|$; alternately, each φ_a is a function $\omega \to \mathcal{P}^M(P)$, but this also gives the bound $(2^{\mathfrak{n}})^{\aleph_0} = 2^{\mathfrak{n}}$.

To get the sharper bound $\mathfrak{n}$, we look more closely at $\varphi_a(n)$. As the set of conditions forcing $\dot{n} \in a$, it enjoys quite a bit of structure. The next few facts depend only on $\varphi_a(n)$ being the set of conditions forcing a statement, not what that statement is. To highlight this, let $A_+ = \varphi_a(n)$ and let $A_- = \{p : p \Vdash \neg\dot{n} \in a\}$.

The Extension Lemma tells us that A_+ and A_- contain the upward cones of any of their conditions—i.e., they are both open sets. Obviously A_+ and

[19] A little more detail, since we will soon need to know that $a \mapsto \varphi_a$ is a function in M. The Definability Lemma tells us that $\{\langle p, n, a\rangle \in P \times \omega \times W : p \Vdash \dot{n} \in a\}$ is a set in M. We now do some reshuffling of that set of triples. First we make it into a set of pairs $\{\langle a, f(a)\rangle : a \in W\}$ where $f(a)$ is the set of pairs $\{\langle n, p\rangle : p \Vdash \dot{n} \in a\}$. Then we replace each $f(a)$ with the set $\{\langle n, \varphi_a(n)\rangle : n \in \omega\}$. The ZF axioms authorize this reorganization. At each step we create "functionhood" by this sort of transformation: $\{\langle x, y_1\rangle, \langle x, y_2\rangle\} \to \{\langle x, \{y_1, y_2\}\rangle\}$.

A_- are disjoint.

We are using weak forcing; this gives us another fact, namely that for any p, exactly one of the following holds:

- $\vee p \subseteq A_+$, i.e., $p \Vdash \dot{n} \in a$;
- $\vee p \subseteq A_-$, i.e., $p \Vdash \neg\dot{n} \in a$;
- $\vee p$ intersects both A_+ and A_-.

Rephrased, either $\vee p$ is contained in A_+, or every element of $\vee p$ is incompatible with every element of A_+, or $\vee p$ has extensions q_+ and q_- such that $\vee q_+$ is contained in A_+ and every element of $\vee q_-$ is incompatible with every element of A_+.[20]

One more observation: if $B \subseteq A_+$ is an antichain, maximal among all antichains contained in A_+, then every element of A_+ is compatible with some element of B. (If not, B wouldn't be maximal.) Any element of A_- is of course incompatible with all elements of A_+, *a fortiori* is incompatible with all elements of B. So we conclude that

$$p \in A_+ \Leftrightarrow (\forall q \geq p)(\exists r \in B)[q \text{ is compatible with } r]$$

It follows that B completely determines A_+.

That's all we need! We can now bound the cardinality of $\{\varphi_a : a \in W\}$, calculating inside M. Since $\varphi_a(n)$ is determined by an antichain contained

[20] A fanciful analogy: think of conditions as voting on a ballot question "ξ or $\neg\xi$?" If $p' \geq p$, think of p' as a future version of p. If $p \Vdash \xi$, then p has decided for ξ; likewise for $p \Vdash \neg\xi$. Once a voter decides, it doesn't change its mind—that's the Extension Lemma. An undecided voter has extensions ("possible futures"), voting "yes" in one future and "no" in another—both "yes" and "no" are compatible with the voter at present. That's the *only* way to be undecided, so a voter who will *never* decide "no" in any possible future must have already decided "yes".

in it, *a fortiori* contained in P, the number of possible values of $\varphi_a(n)$ is at most the number of such antichains. By ccc, the number of such antichains is less than or equal to the number of countable subsets of P. Since $P = \mathrm{Fn}(\mathfrak{n}\times\omega, 2)$, $|P| \leq \aleph_0 \cdot \mathfrak{n} = \mathfrak{n}$ (each element of P being a function with a finite domain, and there being only a finite number of functions per domain). The number of countable subsets of P is therefore at most $\mathfrak{n}^{\aleph_0} = \mathfrak{n}$ (because each countable subset is the range of a function with domain ω). Finally, the function φ_a is specified by choosing $\aleph_0$ values, with at most $\mathfrak{n}$ choices for each value. Thus, at most $\mathfrak{n}^{\aleph_0} = \mathfrak{n}$ possible φ_a's.

Halbeisen slices the same loaf of ideas at a different angle. For any $x \in M$ with $\bar{x} \subseteq \omega$ he constructs a *nice name* for it, which I'll denote by $\ddot{x}$. For any G we have $G \models x = \ddot{x}$, and the cardinality of the set of nice names is at most $\mathfrak{n}$. Effectively Halbeisen provides a smaller warehouse than our W.

The nice name for x is defined as follows:

$$\begin{aligned} \Delta_{\dot{n}\in x} &= \{p : p \Vdash \dot{n} \in x \text{ or } p \Vdash \dot{n} \notin x\} \\ A_n &= \text{ a maximal antichain contained in } \Delta_{\dot{n}\in x} \\ \ddot{x} &= \{\langle \dot{n}, p\rangle : p \in A_n \text{ and } p \Vdash \dot{n} \in x\} \end{aligned}$$

$\Delta_{\dot{n}\in x}$ is dense (trivial exercise). It's easy to see that $G \models \ddot{x} \subseteq x$ for any G, just by looking at the transfinite induction defining $K_G(\ddot{x})$. (We appeal to the Truth Lemma: if $G \ni p \Vdash \dot{n} \in x$ then $G \models \dot{n} \in x$.) In the other direction, suppose $G \models \dot{n} \in x$. So for some $p \in G$, $p \Vdash \dot{n} \in x$. By Fact 15.6(b) (p.346) of Halbeisen, G intersects A_n.[21] Say $p_1 \in A_n \cap G$. We must have $p_1 \Vdash \dot{n} \in x$ or $p_1 \Vdash \dot{n} \notin x$ because $p_1 \in \Delta_{\dot{n}}$, and p_1 and p must be compatible because they both belong to G. That rules out $p_1 \Vdash \dot{n} \notin x$, so $p_1 \Vdash \dot{n}$ and $\langle \dot{n}, p_1\rangle \in \ddot{x}$. Hence $G \models \dot{n} \in \ddot{x}$. We conclude that $G \models x \subseteq \ddot{x}$.

(Side note: since $G \models x = \ddot{x}$ for all generic G, we have $\emptyset \Vdash x = \ddot{x}$, by

[21]Note first that $\vee A_n$ is open dense in P: for any $r \in P$ we have an $s \geq r$ with $s \in \Delta_{\dot{n}\in x}$, and s must be compatible with some element of A_n (else A_n would not be maximal in $\Delta_{\dot{n}\in x}$), thus there is a $t \geq s \geq r$ with $t \in \vee A_n$. Therefore G intersects $\vee A_n$, and since G is downward closed, G contains an element of A_n.

the Fundamental Property of Weak Forcing. Although the M people can't "see" G directly, they do know what forces what. Halbeisen gives a direct argument that $\emptyset \Vdash x = \ddot{x}$, without going through G; I find the version above more intuitive, but your mileage may vary.)

It remains to calculate an upper bound for the number of nice names. A nice name is a countable set of pairs $\langle \dot{n}, p \rangle$ with all the p's for a given n belonging to A_n. By the ccc, there are only countably many such p's, hence a nice name is a countable subset of the set $\mathrm{Ra}(\dot{\omega}) \times P$, which has cardinality $\aleph_0 \times \mathfrak{n}$. Hence, as before, we get the upper bound $\mathfrak{n}^{\aleph_0} = \mathfrak{n}$.

Coda

All the above generalizes readily, with an arbitrary infinite cardinal $\mathfrak{m}$ in M taking the place of $\aleph_0$. The notion of forcing $\mathrm{Fn}(\mathfrak{n} \times \mathfrak{m}, 2)$ satisfies the ccc, so $M[G]$ preserves cardinals and $M[G] \models 2^{\mathfrak{m}} = \mathfrak{n}$ provided $M \models \mathfrak{n}^{\mathfrak{m}} = \mathfrak{n}$. Shoenfield does the general version.

We can always assume our ground model M satisfies GCH; if not, just look at the submodel of constructible sets in M. If we construct an $M[G]$ satisfying $2^{\mathfrak{m}} = \mathfrak{n}$, what does this do to the other cardinal powers? For example, if $\mathfrak{p} < \mathfrak{m}$, then $\mathfrak{n} \times \mathfrak{p} \subseteq \mathfrak{n} \times \mathfrak{m}$, so $\mathrm{Fn}(\mathfrak{n} \times \mathfrak{p}, 2) \subseteq \mathrm{Fn}(\mathfrak{n} \times \mathfrak{m}, 2)$. From this it readily follows that $M[G] \models 2^{\mathfrak{p}} = \mathfrak{n}$ for all infinite $\mathfrak{p} < \mathfrak{m}$ when forcing with $\mathrm{Fn}(\mathfrak{n} \times \mathfrak{m}, 2)$.

Solovay constructed an $M[G]$ in which $\mathfrak{m}$ is the *first* cardinal for which GCH fails, for any regular $\mathfrak{m}$ (and $M \models \mathrm{GCH}$). Central to Solovay's result is a lemma, quite interesting in its own right. I state a special case. Suppose P has the property that for every sequence of conditions $p_0 \leq p_1 \leq \dots$ there is a q with $q \geq p_n$ for all n. (We say P is **σ-closed**.) Then for any generic G and any $X \in M$, $({}^{\omega}X)^{M[G]} = ({}^{\omega}X)^{M}$, i.e., all functions $f : \omega \to X$ belonging to $M[G]$ already belong to M.

As a corollary, $\mathcal{P}^{M[G]}(\omega) = \mathcal{P}^{M}(\omega)$, so a σ-closed notion of forcing adds no new reals. The general version of Solovay's lemma is just the ticket to add lots of new subsets of $\mathfrak{m}$ while adding no new subsets of $\mathfrak{p}$ for any $\mathfrak{p} < \mathfrak{m}$. Also, of course, you can't use $\mathrm{Fn}(A, B)$ as your notion of forcing. (Shoenfield has full details.)

Key idea for Solovay's lemma, in the special case: suppose $\bar{f} : \omega \to X$ is in $M[G]$, i.e., $G \models f : \dot{\omega} \to \dot{X}$. So to speak, "$G$ knows that $\bar{f}$ is a function from ω to X". By the Truth Lemma, there is a $q \in G$ with $q \Vdash f : \dot{\omega} \to \dot{X}$, but that's not good enough—knowing q doesn't necessarily tell you the value of $\bar{f}(n)$ for any particular n.

However, for every n there is a unique $x_n \in X$ such that $M[G] \models \bar{f}(n) = x_n$, and hence a $p_n \in G$ such that

$$p_n \Vdash \langle \dot{n}, \dot{x}_n \rangle \in f \text{ and this is the only pair in } f \text{ with first slot} = \dot{n}$$

or more formally

$$p_n \Vdash \langle \dot{n}, \dot{x}_n \rangle \in f \wedge (\forall y)[\langle \dot{n}, y \rangle \in f \to y = \dot{x}_n]$$

Now using the upward directedness of G we construct a sequence $p'_0 \leq p'_1 \leq \ldots$ with all $p'_n \in G$ and $p'_n \geq p_n$ for all n. Then we let q be an upper bound for the sequence. So q forces all those statements about $\bar{f}(n) = x_n$. In other words, q wraps up all the relevant info from G into one tidy package. Finally we consider the formula

$$\Phi(q, n, x, f) \equiv q \Vdash \langle \dot{n}, x \rangle \in f \wedge (\forall y)[\langle \dot{n}, y \rangle \in f \to y = x]$$

This defines a class in M by the Definability Lemma, and so

$$\{\langle n, x \rangle : n \in \omega \wedge x \in X \wedge \Phi(q, n, x, f)\}$$

is a set in M (using f and q as parameters). But this set is precisely $\bar{f}$ (as a set of ordered pairs).

Halbeisen offers another curious consequence of Solovay's lemma: you can use forcing to show the consistency of CH, bypassing the constructible universe entirely. Do not assume the ground model satisfies CH, let alone GCH.

Say $M \models 2^{\aleph_0} = \mathfrak{n}$. With the right notion of forcing, you can collapse the cardinal $\mathfrak{n}$ to $\aleph_1$, while not adding any new reals. So $2^{\aleph_0}$ is the same in M and $M[G]$, but since $|\mathfrak{n}|^{M[G]} = (\aleph_1)^{M[G]}$, $M[G] \models 2^{\aleph_0} = \aleph_1$. Forcing giveth, and forcing taketh away!

21.7 Not AC

(Shoenfield §§8,9; Halbeisen Chs.8,17)

If M satisfies AC, then so does every generic extension $M[G]$ (see §21.5). One circumvents this difficulty by finding an intermediate model N where AC fails, $M \subseteq N \subseteq M[G]$. Shoenfield uses the hereditarily ordinal definable (HOD) sets, Halbeisen the method of symmetric models.

21.7.1 Preliminary Definitions

Two sets are **equinumerous** if there is a bijection between them. Without Choice, we can't employ the usual definition of $|A|$ as the least ordinal equinumerous with A. The customary substitute: $|A|$ is set of all sets of least rank that are equinumerous with A. We define $|A| \leq |B|$ if there is an injection from A to B. Equivalently, $|A| \leq |B|$ if there is an injection from some $A' \in |A|$ to some $B' \in |B|$, or again equivalently, such injections exists for *all* $A' \in |A|$ and $B' \in |B|$. We let $\aleph_0 = |\omega|$.

A set is **finite** if it is equinumerous to a finite ordinal, **infinite** otherwise. A trivial induction shows that if A is infinite, then for all $n \in \omega$, there is an injection from n into A. A set A is **Dedekind-infinite** if it is equinumerous with a proper subset, i.e., if there's an injection $A \to A$ whose range is a proper subset of A.[22] Halbeisen also uses the lovely term **transfinite**, as

[22]Dedekind proposed this definition in his 1888 essay, "Was sind und was sollen die Zahlen?" ("What are numbers and what should they be?"). In a footnote he defended

will I. Naturally, **Dedekind-finite** means *not Dedekind-infinite*. An easy induction shows that all finite sets are Dedekind-finite; contrapositively, all transfinite sets are infinite.

Fact: A is transfinite iff there is an injection $\omega \to A$. We get one direction without breathing hard: an injection $h : \omega \to A$ enables us to define an injection from the range of h to a proper subset of the range; we extend to all of A by leaving all other elements of A fixed. This produces a bijection between A and a proper subset. For the other direction, suppose $f : A \to A$ is an injection onto a proper subset of A. Let $a \in A$ lie outside the range of f, and look at the iterates of f, $\{f^n(a) : n \in \omega\}$. The map $n \mapsto f^n(a)$ is an injection $\omega \to A$. Proof: if $f^n(a) = f^{n+k}(a)$ for some $k > 0$, then by induction on k and the injectivity of f, $a = f^0(a) = f^k(a)$. But this would mean that a is the f-image of $f^{k-1}(a)$, contrary to assumption.

A nice summary:

$$A \text{ is infinite} \Leftrightarrow (\forall n \in \omega)|n| \leq |A|$$
$$A \text{ is transfinite} \Leftrightarrow \aleph_0 \leq |A|$$

21.7.2 Consequences of Choice

Before delving into models where AC fails, let's look at some of its consequences.

Countable Choice. This says that there is a choice function for any countable family of nonempty sets. Formally, if A is a function with domain ω and $A(n) \neq \varnothing$ for all n, then there is a function $c : \omega \to \bigcup A(n)$ such that $c(n) \in A(n)$ for all n.

his definition: "All other attempts that have come to my knowledge to distinguish the infinite from the finite seem to me to have met with so little success that I think I may be permitted to forego any criticism of them."

Countable Unions. A countable union of countable sets is countable. Minor variations are equivalent, such as substituting "countably infinite" for "countable" throughout, or insisting that the sets be pairwise disjoint. Formally (using "countably infinite" and demanding disjointness) if there is a bijection $f_n : \omega \to A(n)$ for each n, then there is a bijection $f : \omega \to \bigsqcup A(n)$.

Countable Unions follows from Countable Choice. The bijection f amounts to the usual bijection $\omega \leftrightarrow \omega \times \omega$, once we've chosen a bijection f_n for each n. So we let $F(n)$ be the set of bijections $\omega \leftrightarrow A(n)$, and apply Countable Choice to the family of $F(n)$'s.

Dedekind-finite implies Finite. Contrapositively, infinite implies transfinite. This also follows from Countable Choice. Suppose A is infinite. Let $[A]^n$ be the set of all subsets of A of size n (i.e., equinumerous to n). We know each $[A]^n \neq \varnothing$ by our remarks above. Use Countable Choice to obtain a function F with $F(n) \in [A]^n$. So each $F(n)$ is a subset of size n. At first you might think another application of Countable Choice does the trick: for $n > 0$ let $c(n) \in F(n)$ be a choice function, and let $c(0)$ be arbitrary. While c maps ω into A, it might not be injective. To resolve this, let s_n be a sequence of positive integers such that $s_{n+1} > \sum_{k=0}^{n} s_k$ for all n; for example, we could set $s_n = 2^n$. Then $|F(s_{n+1})| > |\bigcup_{k=0}^{n} F(s_k)|$, so the sets $G(0) = F(s_0)$, $G(n+1) = F(s_{n+1}) \smallsetminus \bigcup_{k=0}^{n} F(s_k)$ are all nonempty and disjoint. If $c(n) \in G(n)$ is a choice function, then c is injective.

Halbeisen constructs models in which each of these consequences fails. Exploring the topography of implications became quite a cottage industry in the aftermath of forcing. (Jech [12] is a comprehensive reference; Moore [19, Appendix 2] has tables summarizing the facts.)

21.7.3 Intuition

Let's start with a strategy that doesn't quite work, and then see how to tweak it. Suppose we add a set $A = \{a_0, a_1, \ldots\}$ to M, obtaining a model we'll denote $M[A]$. (For the moment, don't worry about what the a_n's are.) Given any permutation π of A, it stands to reason that we can extend π to all of $M[A]$ by induction on rank. Let $\mathcal{G}$ be the group of all these π's. Each π is an $\in$-automorphism on $M[A]$: $x \in y$ iff $\pi x \in \pi y$. Furthermore, M remains pointwise fixed under all of $\mathcal{G}$, at least if we extend the permutations in a natural way.

Let N be the collection of elements of $M[A]$ remaining fixed under $\mathcal{G}$, i.e., $N = \{x \in M[A] : (\forall \pi \in \mathcal{G})\pi x = x\}$. N is a set in "the real world", but not a set in $M[A]$, or even necessarily a class in $M[A]$. Still, *suppose* that N is a model of ZF. (Sort of like finding fixed fields in Galois theory.) We'll have $M \subseteq N$ and also $A \in N$. But no well-ordering of A can belong to N because it obviously won't remain fixed when we scramble A.

Now for the chief obstacle: none of the a_n's belong to N either! N fails to be transitive, making hash of Extensionality. (A has no elements but isn't the empty set.) The tweak: don't insist that the elements of N be fixed under *all* of $\mathcal{G}$, just under *large enough subgroups* of $\mathcal{G}$. Put another way: the elements of N don't have to be *fully symmetric*, just *symmetric enough*.

Halbeisen lays out the central definitions clearly in Ch.8 (p.193) and Ch.17 (p.384). First, a family $\mathcal{F}$ of subgroups of $\mathcal{G}$ is a **normal filter** if for all subgroups H, K of $\mathcal{G}$ we have:

1. $\mathcal{G} \in \mathcal{F}$.
2. If $H \in \mathcal{F}$ and $H \subseteq K$, then $K \in \mathcal{F}$.
3. If $H \in \mathcal{F}$ and $K \in \mathcal{F}$ then $H \cap K \in \mathcal{F}$.
4. If $\pi \in \mathcal{G}$ and $H \in \mathcal{F}$ then $\pi H \pi^{-1} \in \mathcal{G}$.

$\mathcal{F}$ is the formalization of the notion "large enough subgroup": all subgroups belonging to $\mathcal{F}$ are "large enough".

Next, define $\mathrm{sym}(x) = \{\pi \in \mathcal{G} : \pi x = x\}$. We say x is **symmetric** if $\mathrm{sym}(x) \in \mathcal{F}$. Now, $\mathrm{sym}(a_n)$ is the group of permutations that don't move a_n. If we insist that $\mathrm{sym}(a_n) \in \mathcal{F}$ for all n, then all the a_n's will be symmetric.[23]

For any $S \subseteq A$, let $\mathrm{fix}(S) = \{\pi \in \mathcal{G} : (\forall s \in S)\pi s = s\}$. (Unlike $\mathrm{sym}(S)$, $\mathrm{fix}(S)$ demands that its automorphisms leave S *pointwise* fixed.) By taking finite intersections (clause 3 above), we see that $\mathcal{F}$ must include not only $\mathrm{sym}(a_n)$ for all a_n, but $\mathrm{fix}(S)$ for all finite subsets S of A—at least if we want all the a_n's to be symmetric. OK, set

$$\mathcal{F} = \{H \text{ subgroup of } \mathcal{G} : H \supseteq \mathrm{fix}(S) \text{ for some finite } S \subseteq A\}$$

This $\mathcal{F}$ is a normal filter; using it, both A and all the a_n's are symmetric. Some thought reveals that any finite or cofinite subset of A is symmetric for this $\mathcal{F}$.

More generally, we can replace the collection of finite subsets of A with any idea of a "small enough" set. A **normal ideal** I (p.194) is a family of subsets satisfying:

1. $\varnothing \in I$.
2. If $F \subseteq E \in I$ then $F \in I$.
3. If $E, F \in I$ then $E \cup F \in I$.
4. If $\pi \in \mathcal{G}$ and $E \in I$ then $\pi E \in I$.

Halbeisen also demands that $\{a\} \in I$ for all $a \in A$.

[23]Halbeisen makes this part of his definition of a normal filter (p.193). In a cut&paste error, this clause also appears on p.384, where it doesn't make sense.

Halbeisen says, "let $\mathcal{F}$ be the filter on $\mathcal{G}$ generated by the subgroups $\{\text{fix}(E) : E \in I\}$." We can be more precise: $K \in \mathcal{F}$ iff $K \supseteq \text{fix}(E)$ for some $E \in I$. On the one hand, all such K's must belong to $\mathcal{F}$ by clause (2) in the definition of a normal filter. On the other hand, the family of all such K's constitutes a normal filter because of two easily demonstrated facts:

$$\text{fix}(E \cup F) = \text{fix}(E) \cap \text{fix}(F)$$
$$\text{fix}(\pi E) = \pi \text{fix}(E) \pi^{-1}$$

The notion of *support* (p.194) crystallizes the notion of "symmetric enough": x is symmetric enough if there is a "small enough" set E (that is, an $E \in I$) such that any automorphism leaving E pointwise fixed doesn't move x (i.e., $\alpha x = x$). In other words: if we *lock down* the set E, that will insure that x remains fixed. In this case, we say that E is a **support** for x. Once you've gotten used to the framework, the phrase "finite support" shorthands the discussion.

Finally we let N be all the symmetric ("symmetric enough") elements of $M[A]$. With this set-up, it turns out that N satisfies ZF. On the other hand, if f is a bijection between A and an ordinal, then $\pi f = f$ only if π is the identity. (Proof: transfinite induction.) So f isn't symmetric and A has no well-ordering in N.

We can sharpen the conclusion: there is no injection $f : \omega \to A$ in N. No such f could have finite support: you can always scramble A, keeping any finite subset pointwise fixed, in a manner that changes f. So A is infinite but not transfinite, and a relatively weak consequence of AC fails.

Now let's worry about the nature of the a_n's. "ZF with atoms" (ZFA, also called "ZF with urelements") supplies the platform for permutation models. Briefly, atoms are primitive objects, not containing any elements. In ZFA, the cumulative hierarchy starts with the set of all atoms at the ground level instead of with $\varnothing$. Halbeisen gives more details in Ch.8. One crucial point: a model of ZFA is *not* a model of ZF (unless there are no atoms). Technically, ZFA has a modified form of Extensionality. More fundamentally, in ZF everything belongs to some V_α; in ZFA, we have to expand this to $V_\alpha[A]$ to get everything.

Fraenkel introduced these **permutation models** in 1922; later Lindenbaum and Mostowski perfected them, so people also call them Fraenkel-Mostowski or FM models. People call the $M[A]$ outlined above, with $A = \{a_0, \ldots\}$, the **Basic Fraenkel Model**. By varying the set A, the group $\mathcal{G}$, and the normal filter $\mathcal{F}$, you can achieve many effects; we'll see a few below.

Permutation models tell us nothing *directly* about ZF, since as noted they are models of ZFA, not of ZF. Is there a well-ordering of $\mathcal{P}(\omega)$? Permutation models shed no direct light on this. (The weasle word "direct" is a concession to the Jech-Sochor construction, which embeds large fragments of permutation models into symmetric models. Halbeisen describes this.)

Digression: symmetry burst upon the mathematical world stage with Galois theory. Let's compare and contrast. Typically in Galois theory we look at a tower of fields $F = F_0 \subseteq \ldots \subseteq F_n = E$, and study the automorphism groups $\mathrm{Aut}(E/F_k)$, $\mathrm{Aut}(F_k/F)$, and $\mathrm{Aut}(F_{k+1}/F_k)$. The group $\mathrm{Aut}(E/F)$ enjoys maximum symmetry; somehow this symmetry must be attained piecemeal as we go up the tower. That's how the famous unsolvability of the quintic is proved: no tower of radical extensions can supply enough symmetry. Two other well-known results adhere to the same pattern: the irreducible case of the cubic, and the impossibility of various geometrical constructions.

In a slogan: you can't chain a highly symmetrical object to a less symmetrical construction.

The same slogan works for $M[A]$ and the injection $f : \omega \to A$. The support of f is its infinite range, so it is insufficiently symmetrical to gain entrance to the submodel N. You can push the analogy a little ways, but not too far. The splitting field of a polynomial matches up nicely with $M[A]$, and the set of roots with A. Just as permutations of the roots determine the automorphisms, so too with permutations of A. The lattice of subgroups of

the automorphism group plays a big role in both contexts: the Galois correspondence pairs up the subgroups with the lattice of intermediate fields, and for permutation models, the normal filter lives inside the lattice. But the core of Galois theory is the Galois correspondence, associating a subgroup with the subfield of fixed elements. As we've seen for permutation models, the class of all fixed elements of the automorphism group isn't even transitive, let alone a submodel. Many other fundamental aspects of Galois theory (factor groups, degrees of extensions) don't seem to carry over. The analogy apparently leads nowhere.
End digression.

Cohen demonstrated the consistency of ZF + ¬AC with a generic model where the a_n's are all generic subsets of ω (so-called Cohen reals). Given everything we've seen about generic models, you can guess how this goes. First, we transfer the discussion to M, since M can talk about $M[G]$. Instead of a group of permutations of the atoms, we start with a group of automorphisms of the notion of forcing P. That is, $\alpha : P \to P$ and $p \leq q$ iff $\alpha p \leq \alpha q$ for all conditions p and q. We then extend each such α to an automorphism on the set of all names. (Instead of allowing all elements of M to function as names, we confine our attention to the class denoted V^P in §21.2.) We can now define **symmetric names** pretty much the same way, and let N be $\{K_G(x) : x \text{ is a symmetric name}\}$. (Actually we have to use *hereditarily symmetric names*, in order to get a transitive model.) N is a **symmetric submodel** of $M[G]$.

Let's explore the differences between $M[A]$ and $M[G]$ a little further. Given any permutation of A, we obtain a true $\in$-automorphism of $M[A]$. This is the core of the permutation-model independence proofs: we look at the submodel of "sufficiently symmetric" elements of $M[A]$, under these automorphisms. In contrast, an easy induction on rank shows that the only $\in$-automorphism of *any* standard model of ZF is the identity. Even stronger: the only $\in$-isomorphism from one standard ZF model onto another is the identity (same proof). (Remember that the elements of a standard model are all sets in the "real world" of V, so it makes sense to say that x and αx

are equal, even if they belong to different models.)

On the other hand, if α is an automorphism of P, extended to the class V^P of names, we have:

$$a \in_p b \Leftrightarrow \alpha a \in_{\alpha p} \alpha b$$
$$p \Vdash \varphi(c_1, \ldots, c_k) \Leftrightarrow \alpha p \Vdash \varphi(\alpha c_1, \ldots, \alpha c_k)$$
$$(c_1, \ldots, c_k \text{ all names})$$
$$K_G(a) \in K_G(b) \Leftrightarrow K_{\alpha G}(\alpha a) \in K_{\alpha G}(\alpha b)$$

You might think we can get an isomorphism $M[G] \to M[\alpha G]$ via $K_G(a) \mapsto K_{\alpha G}(\alpha a)$. Indeed we can—this map is well-defined—but it's just the identity. So $M[G] = M[\alpha G]$ for all G, and $M[G]$ possesses no non-trivial automorphisms.

Instead of automorphisms on $M[G]$, we work with automorphisms on V^P. We have a commutative diagram:

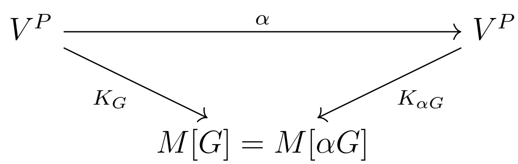

Lifting the discussion from symmetry-poor $M[G]$ to symmetry-rich V^P means relying on the forcing relation, but that's old hat by now.

One last bout of handwaving before we descend to specifics. Why do permutation and symmetric models satisfy ZF? What divides AC from all the other axioms, so that they continue to hold and AC fails?

The ZF axioms, except for Extensionality, Foundation, and Infinity, construct a new set from given sets. AC asserts the *existence* of a set partially specified by other sets, without giving a *construction*—precisely the feature that rubbed many mathematicians the wrong way when Zermelo first proposed it.

Halbeisen justifies the Separation Axiom on p.385. This creates a new set $\{\bar{v} \in \bar{u} : \varphi(\bar{v}, \bar{a}_1, \ldots, \bar{a}_n)\}$ from existing sets $\bar{u}$ and $\bar{a}_1, \ldots, \bar{a}_n$. The pivotal computation appears near the bottom of the page, and relies on the equivalence

$$p \Vdash \varphi(v, a_1, \ldots, a_n) \Leftrightarrow \alpha p \Vdash \varphi(\alpha v, \alpha a_1, \ldots, \alpha a_n)$$

along with the observation that as $\langle v, p\rangle$ ranges over all pairs in $V^P \times P$, so does $\langle \alpha v, \alpha p\rangle$.

The new set partakes of symmetries of its ingredients, as it were. We can see this even more clearly with the permutation model $N \subseteq M[A]$: applying an $\in$-automorphism α to $\{v \in u : \varphi(v, a_1, \ldots, a_n)\}$ yields

$$\{\alpha v \in \alpha u : \varphi(\alpha v, \alpha a_1, \ldots, \alpha a_n)\}$$

and if u and all the a_i's are left fixed by α, this equals

$$\{\alpha v \in u : \varphi(\alpha v, a_1, \ldots, a_n)\}$$

the same as $\{v \in u : \varphi(v, a_1, \ldots, a_n)\}$.

Likewise for the other constructions sanctioned by the ZF (or ZFA) axioms. Take Power Set: $v \subseteq u$ iff $\alpha v \subseteq \alpha u$, and if $\alpha u = u$, then we get $\alpha\mathcal{P}(u) = \mathcal{P}(u)$. Replacement and Union look much like Separation. For symmetric models, we work with names and forcing conditions instead of $\in$-automorphisms, but otherwise it's the same story.

When we turn to AC, a very different tale. If $w : \gamma \leftrightarrow u$ is a bijection between an ordinal γ and u, and if $\alpha u = u$, generally speaking we do *not* have $\alpha w = w$. The elements of γ remain fixed but the elements of u get scrambled. So w does not partake of the symmetries of u. Or say c is a choice function on u, i.e., $c(v) \in v$ for all nonempty $v \subseteq u$. It can certainly happen that $\alpha v = v$ without $\alpha c(v) = c(v)$. Arbitrary choices don't track symmetries.

That's the symmetric model approach. I postpone discussion of the HOD approach.

21.7.4 Permutation and Symmetric Models

Halbeisen treats permutation models in Ch.8 and symmetric models in Ch.17. I won't repeat all the technical details here, just offer some supplemental intuition (i.e., hand-waving).

The Basic Fraenkel and Cohen Models

Halbeisen describes the basic Fraenkel model (BFM for short) on pp.195–196, and its companion the basic Cohen model (BCM for short) on p.386–388. We delineated the BFM above: the finite-support permutation model with the set of atoms $A = \{a_0, a_1, \ldots\}$ and all permutations of A. The BCM replaces this with a set A of so-called Cohen reals: each $\bar{a}_n$ is a generic subset of ω. We saw in §21.6 that the forcing notion $\mathrm{Fn}(\mathfrak{n} \times \omega, 2)$ adds $\mathfrak{n}$ new reals, so we use $\mathrm{Fn}(\omega \times \omega, 2)$ to add the new Cohen reals to $M[G]$. We need a name $a_n \in M$ for $\bar{a}_n \in M[G]$.[24] A condition p maps a finite subset of $\omega \times \omega$ to $\{0, 1\}$, with $p(n, k) = 1$ iff $\bar{a}_n$ is supposed to contain k. Now,

$$\bar{a}_n = \{\bar{b} : b \in_p a_n \text{ with } p \in G\}$$

so we let $b = \dot{k}$, $\bar{b} = k$, and include the pair $\langle \dot{k}, p \rangle$ just when $p(n, k) = 1$. This yields the name a_n Halbeisen gives on p.386:

$$a_n = \{\langle \dot{k}, p \rangle : k \in \omega \;\&\; p(n, k) = 1\}$$

The name A uses the empty condition to include all the $\bar{a}_n$'s unconditionally into $\bar{A}$.

Halbeisen specifies the group $\mathcal{G}$ for the BCM on p.387. For each π permuting ω, we need an automorphism α_π of P, inducing an automorphism of V^P (also denoted α_π). We want $\alpha_\pi a_n = a_{\pi n}$. So any given p should make the same claims regarding "$k \in \alpha_\pi a_n$" and "$k \in a_{\pi n}$", i.e., $\alpha_\pi p(n, k) \equiv p(\pi n, k)$.

[24] Remember how my notation differs from Halbeisen's: I use a for the name and $\bar{a}$ for the corresponding element of $M[G]$. For $a \in M$, $\dot{a}$ is the canonical name for a. See §21.2.

This results in Halbeisen's formula for α_π:

$$\alpha_\pi p = \{\langle\langle \pi n, k\rangle, i\rangle : \langle\langle n, k\rangle, i\rangle \in p\}$$

We define the notion of **support** nearly as before: $E \subseteq \omega$ is a support for a name x if locking down the elements of E insures that x is left fixed. (Formally, if $\pi \in \text{fix}(E)$ implies $\alpha_\pi x = x$.) A name x is **symmetric** if it has finite support.

Finally, a name a is hereditarily symmetric when not just a but all its potential elements are symmetric, and so on all the way down. "All the way down" really means "induction on rank". As for "potential element", we mentioned this in §21.3: we call b a potential element of a when $b \in_P a$ (meaning $b \in_p a$ for some condition p). HS is the collection of all hereditarily symmetric names, and $N = \{\bar{a} : a \in \text{HS}\}$.

Symmetry in action: N violates AC, indeed violates the implication "infinite implies transfinite", for both the BFM and the BCM. The handwaving proof is very short: if $f : \omega \to A$ is 1–1 for some $f \in N$, then f cannot have finite support, for we can always move something in the range of f while locking down any given finite set.

For the BFM, the details look like this. If f were symmetric it would have finite support $E \subseteq A$. But there must be an $m \in \omega$ such that $f(m) = a_n \in A \smallsetminus E$ because f is 1–1, and thus a permutation $\pi : A \to A$ that leaves E pointwise fixed while moving a_n. Extending π to an automorphism of the BFM (still denoted π), we must have $(\pi f)(m) = \pi a_n$. Since f has support E and π leaves E pointwise fixed, $\pi f = f$. But then $f(m) = a_n$ and $f(m) = \pi a_n \neq a_n$. This shows that f is not a function.

For the BCM, we lift the argument from $\bar{A}$ to the realm of names. Let f be an HS name for $\bar{f} : \omega \to \bar{A}$. Say f has finite support $E \subseteq \omega$. So there is an $m \in \omega$ and an $n \in \omega \smallsetminus E$ such that $\bar{f}(m) = \bar{a}_n$; also a permutation π that fixes E pointwise while sending n to $\pi n \neq n$. From π we get the automorphism α_π on the set of conditions, extended to the set of all names.

Now α_π doesn't monkey around with $\bar{f}(m)$ directly; instead, we argue that $p \Vdash f(\dot{m}) = a_n$ for some $p \in G$. Applying α_π, we have

$$\alpha_\pi p \Vdash (\alpha_\pi f)(\alpha_\pi \dot{m}) = \alpha_\pi a_n$$

However, $\alpha_\pi f = f$ because f has support E, $\alpha_\pi \dot{m} = \dot{m}$ because $\dot{m}$ is fully symmetric, and it is easily shown that $\alpha_\pi a_n = a_{\pi n}$ (Halbeisen p.387). So we have

$$\alpha_\pi p \Vdash f(\dot{m}) = a_{\pi n}$$

To get a contradiction, we need a common extension r of p and $\alpha_\pi p$ forcing four things:

$$\begin{aligned} &r \Vdash f(\dot{m}) = a_n \\ &r \Vdash f(\dot{m}) = a_{\pi n} \\ &r \Vdash a_n \neq a_{\pi n} \\ &r \Vdash f \text{ is a function} \end{aligned}$$

In fact, $\emptyset \Vdash a_n \neq a_{n'}$ whenever $n \neq n'$ (because the set of conditions with $p(n,k) \neq p(n',k)$ for some k is dense). The assumption that $\bar{f}$ is a function means that $q \Vdash$ "f is a function" for some $q \in G$; q is compatible with p because both belong to G. Finally, how can we insure that $\alpha_\pi p$ is compatible with p and q? Answer: π has access to lots of "virgin territory", areas in the $\omega \times \omega$ bit matrix outside the combined footprint of p and q. In less florid language, π can exchange n with some n' outside the domains of p and q, i.e., an n' for which $p(n',k)$ is undefined for all k, ditto for q. (See fig.5). So $\emptyset$, q, p, and $\alpha_\pi p$ are all compatible and there is an r forcing four contradictory statements.

Halbeisen actually gives a different (very slick) argument for the BFM: see Prop.8.3, p.196. However, I wanted to highlight the parallels and differences between the BFM and the BCM.

I want to sound a few recurring notes. We'll see the "virgin territory" argument again. $\bar{A}$ is transfinite (indeed, countable) in $M[G]$; the symmetry

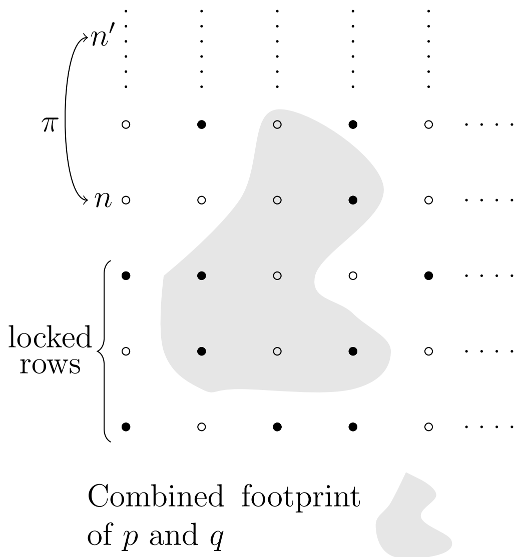

Figure 5: The "Virgin Territory" Argument

A small portion of the (n, k) bit-matrix is shown. The "footprint" (combined domains) of conditions p and q is shaded. Locked rows cannot be moved. The permutation π exchanges row n with a row n' outside the footprint; thus, $\alpha_\pi p$ is compatible with both p and q. See text for more details.

requirements act as a gatekeeper, allowing $\bar{A}$ into N but keeping out all the injections $\omega \to \bar{A}$ that are found in $M[G]$. Finally, I reiterate a fine point mentioned earlier. The α_π's permute the names a_n, not the sets $\bar{a}_n$. You can certainly use π to *define* a permutation of the $\bar{a}_n$'s, but this won't come from an $\in$-automorphism of N: all elements of ω are left fixed by any $\in$-automorphism, hence any subset of ω is also left fixed.

A Bunch of Pairs

In these models, Countable Choice fails even when we are choosing from a set of unordered pairs.

Consider a minor tweak to the BFM: instead of allowing all permutations, permit only those that swap adjacent even/odd items. Let the atoms be $\{a_0, b_0, a_1, b_1, \ldots\}$; the group of permutations is generated by the transpositions $(a_n b_n)$. The normal filter consists of subgroups locking down all a_i, b_i for i less than some given n. There is no choice function for the set of pairs $\{\{a_0, b_0\}, \{a_1, b_1\}, \ldots\}$. Halbeisen presents this in Ch.8 (p.197) where he calls it the Second Fraenkel Model.

For the BCM, it's instructive to start off with a modification that *doesn't* work, and then fix it. Use the same notion of forcing, but let $\mathcal{G}$ be the group of permutations of ω generated by the transpositions $(2n, 2n+1)$. Let a_n be the name associated with row $2n$, b_n with row $2n+1$. Using the empty condition, it's easy to concoct HS names for the unordered pairs $\{a_n, b_n\}$ and the set of all such pairs.

The symmetric model N thus contains the set of all unordered pairs $\{\bar{a}_n, \bar{b}_n\}$. Each $\bar{a}_n$ and $\bar{b}_n$ is a generic subset of ω. However, N also contains a choice function for this set of pairs! The correspondence $\mathcal{P}(\omega) \leftrightarrow \mathbb{R}$ makes this obvious: just choose the smaller real number from each pair. Without using this correspondence, find the first difference in the bitstrings for $\bar{a}_n$ and $\bar{b}_n$, and choose the element with a 0 in that spot.

What went wrong? Answer: the virgin territory argument. If $p \Vdash c(\dot{n}) = a_n$, for example, then the footprint of p will (very likely) include the first-difference bits of both $\bar{a}_n$ and $\bar{b}_n$. The permutation π exchanging $2n$ with $2n+1$ yields an automorphism α_π changing $p \Vdash c(\dot{n}) = a_n$ into $\alpha_\pi p \Vdash c(\dot{n}) = b_n$. But p and $\alpha_\pi p$ are not compatible.

We fix this by going up one level, creating a set of pairs $\{\{\bar{U}_0, \bar{V}_0\}, \{\bar{U}_1, \bar{V}_1\}, \ldots\}$, with each $\bar{U}_n$ and $\bar{V}_n$ a set of reals. (Or what amounts to the same thing, a set of subsets of ω.) That means we use $\mathrm{Fn}(\omega \times \omega \times \omega, 2)$ as our notion of forcing, with $p(2n, k, j)$ dictating whether j should belong to $\bar{a}_{nk}$ and $p(2n+1, k, j)$ doing the same for $\bar{b}_{nk}$. Then $\bar{U}_n$ will equal $\{\bar{a}_{n0}, \bar{a}_{n1}, \ldots\}$ and $\bar{V}_n$ will equal $\{\bar{b}_{n0}, \bar{b}_{n1}, \ldots\}$. The group of permutations is a bit more complicated. Besides switching U_n with V_n, we need to insure that the a_{nk}'s making up U_n can be permuted among themselves, likewise for b's; otherwise, N could choose between $\bar{U}_n$ and $\bar{V}_n$ by looking at the "first" sets $\bar{a}_{n0}$ and $\bar{b}_{n0}$. (In other words, we need to exclude choice functions for $\{\bar{U}_0, \bar{U}_1, \ldots\}$ and $\{\bar{V}_0, \bar{V}_1, \ldots\}$.) I leave the details as an exercise.

Violating the Countable Union Theorem

Halbeisen gives a symmetric model for this on pp.389–391. Just like the BCM, we add Cohen reals $\{a_0, a_1, \ldots\}$; our notion of forcing is $\mathrm{Fn}(\omega\times\omega, 2)$, as before, with $p(n,k)$ dictating whether k should be an element of a_n.

Our goal is not to render the $\bar{a}_n$ indistinguishable. Instead, we'll associate a countable set $\bar{U}_n$ with each $\bar{a}_n$ in a manner that makes the union $\bar{U} = \bigcup \bar{U}_n$ uncountable. Recall why the Countable Union Theorem needs AC: we have to choose a bijection $\omega \leftrightarrow \bar{U}_n$ for each n. In our symmetric model N, there will *be* a bijection for each $\bar{U}_n$, it just won't be possible to *choose* one simultaneously for all n.

We achieve this by letting $\bar{U}_n$ be the set of all subsets of ω that differ from $\bar{a}_n$ in finitely many places. (The usual density argument shows that any two $\bar{a}_n$'s differ in infinitely many places, so the $\bar{U}_n$ are pairwise disjoint, although this won't matter.) Using the customary notation $X\triangle Y$ for the symmetric difference $X\smallsetminus Y \cup Y\smallsetminus X$, $\bar{U}_n$ is $\{\bar{u} \subseteq \omega : |\bar{u}\triangle\bar{a}_n| < \aleph_0\}$.

We can effectively enumerate the finite subsets of ω, say $(F_j : j \in \omega)$—this enumeration belongs to the ground model M. Then we let $\bar{u}_{nj}$ be $\bar{a}_n$ with the "bits flipped" exactly at the places in F_j. If we regard subsets of ω as characteristic functions, then flipping bits is just modulo 2 addition. Our model will include $\bar{U}_n = \{\bar{u}_{nj} : j \in \omega\}$ for all $n \in \omega$.

Having $(F_j : j \in \omega)$ in the ground model means that once we pick *one* element of $\bar{U}_n$, we'll have an effective enumeration of *all* the elements of $\bar{U}_n$—that's what makes $\bar{U}_n$ countable. To prevent $\bar{U}$ from being countable, we have to make all the elements of a given $\bar{U}_n$ "look alike", so there's no choice function for the family $\{\bar{U}_n : n \in \omega\}$. This determines our group of automorphisms: we want the ability to "flip the bits" in a finite number of places, for each of the $\bar{a}_n$'s. Let f be a "flip function": $f : \omega \times \omega \to 2$ with finite support (i.e., $f(n,k) = 1$ for only finitely many pairs (n,k)). Define

$$(\alpha_f p)(n,k) = p(n,k) \triangle f(n,k)$$

where $\triangle$ is addition modulo 2. If $p(n,k)$ is undefined, then so is $p(n,k)\triangle f(n,k)$. (Although this won't matter, α_f has period 2: $\alpha_f(\alpha_f p) = p$.) Each α_f is an automorphism of P, and our group of automorphisms is the set of all the α_f's.

Now focus attention on a single row of the (n,k) matrix of bits (i.e., consider a fixed value of n). The name a_n goes with that row. If G is a generic filter, then restricting each $p \in G$ to that row gives all the info specifying $\bar{a}_n$. For motivation, we consider a map determined by α_f (not an $\in$-automorphism), sending each $\bar{a}_n$ to another subset of ω; call it $\alpha_f \bar{a}_n$. $K_G(a_n) = \bar{a}_n$ (notational convention), and $K_{\alpha_f G}(a_n) = \alpha_f \bar{a}_n$. So $\alpha_f \bar{a}_n$ differs from $\bar{a}_n$ in only finitely many places. Even better: as f ranges through all possible flip functions, $\alpha_f \bar{a}_n$ ranges over all the elements of $\bar{U}_n$. This thwarts any attempt to choose an element of $\bar{U}_n$ for each n, at least in the symmetric submodel N. On the other hand, the sequence of $\bar{U}_n$'s is immune to the α_f's, allowing the enumeration $(\bar{U}_0, \bar{U}_1, \ldots)$ to survive inside N.

Next we specify the normal ideal generating the normal filter. A set belongs to the normal ideal iff it is contained in a finite number of rows. (More formally, E_0 is in the normal ideal iff $E_0 \subseteq E \times \omega$ where E is a finite subset of ω.) So our normal filter permits us to "lock down" the bits in a finite number of rows: a name x will be symmetric if there is a finite $E \subseteq \omega$ such that $\alpha_f x = x$ whenever $f(n,k) = 0$ for all $n \in E$ and all k. This definition insures that the names

$$a_n = \{\langle \dot{k}, p \rangle : k \in \omega \ \&\ p(n,k) = 1\}$$

and

$$u_{nj} = \{\langle \dot{k}, p \rangle : k \in \omega \ \&\ p(n,k) \triangle F_j(k) = 1\}$$

are all symmetric (treating F_j as the characteristic function of the j-th finite subset of ω).

We show next that we have symmetric names for $\bar{U}_n$, for the sequence $(\bar{U}_0, \bar{U}_1, \dots)$, and for its union $\bar{U}$. That's why $\bar{U}$ is a countable union of countable sets in N! Now, even though we could lock down row n to make U_n symmetric, we can't lock down *all* the rows to make $(U_0, U_1, \dots)$ symmetric. However, the α_f's permute the names u_{nj}—that saves us. Here's one possible symmetric name U_n:

$$U_n = \{\langle u_{nj}, \emptyset\rangle : j \in \omega\}$$

The symmetry of the names has a counterpart in $M[G]$: any α_f just permutes the elements of any $\bar{U}_n$, leaving $\bar{U}_n$ fixed as a set.

Finally, setting $U = \bigcup U_n$ gives a symmetric name for $\bar{U}$. $\bar{U}$ is not countable in N. If we had a symmetric name h for a surjection $\bar{h} : \omega \twoheadrightarrow \bar{U}$, then we'd reach a contradiction using a "virgin territory" argument. First, let E be a finite support for h. Since $\bar{h}$ is surjective, the range of $\bar{h}$ includes some $\bar{a}_n$ outside the locked down rows in E. Thus $p \Vdash h(\dot{m}) = a_n$ for some $m, n \in \omega$ and some condition $p \in G$. Let $G \ni q \Vdash$ "h is a function". We need an α_f with $\alpha_f p$ compatible with p and q, and with $\emptyset \Vdash \alpha_f a_n \neq a_n$. Also α_f must respect the locked down rows, i.e., f must be 0 on all these rows. This will give us the contradiction

$$\begin{aligned} p &\Vdash h(\dot{m}) = a_n \\ \alpha_f p &\Vdash h(\dot{m}) = \alpha_f a_n \\ \emptyset &\Vdash \alpha_f a_n \neq a_n \\ q &\Vdash \text{“}h \text{ is a function”} \end{aligned}$$

Just let f flip a single bit, belonging to row n and outside the finite footprint (i.e., domain) of p. Not only are p and $\alpha_f p$ compatible—they're the same! Since p and q both belong to G, they're compatible. The set of conditions r such that $r \Vdash \alpha_f a_n \neq a_n$ is easily seen to be dense, so $\emptyset \Vdash \alpha_f a_n \neq a_n$. We are done.

Let's see how this all plays out in Halbeisen. He introduces the names for a_n, U_n, and U on p.389. His name for U_n

$$U_n = \{u \subseteq \mathrm{Ra}(\dot{\omega}) \times P : \emptyset \Vdash |u \triangle a_n| < \dot{\aleph}_0\}$$

is equivalent to ours, in the sense that they both map to the same $\bar{U}_n$. He lets F be the restriction of a flip function f to a given row n_0, and writes α_{π_F, n_0} for the restriction of α_f to that row. He makes a minor error in defining the group $\mathcal{G}$: with his definition, $\mathcal{G}$ is not closed under composition, but the group generated by his $\mathcal{G}$ is just the group of all the α_f's.

On p.390 Halbeisen observes that $\alpha U_n = U_n$ for all $\alpha \in \mathcal{G}$ and all U_n (ditto for U) because α just permutes the u_{nj} (without changing n). He introduces the $(F_j : j \in \omega)$ enumeration of finite sets to show that each $\bar{U}_n$ is countable. Lastly the "virgin territory" argument occurs on pp.390–391, with $F = \{k\}$; here (n_0, k) is the bit outside the locked rows and the footprint of p (p_0 for Halbeisen).

Sounding a recurring note, $\bar{U}$ is countable in $M[G]$, at least if M (and hence $M[G]$) satisfies ZFC.

The Feferman-Lévy Model

In this model (pp.392–395), $\mathcal{P}(\omega)$ is a countable union of countable sets. Halbeisen's treatment is (I feel) not as clear as it could be; here's my version.

This model collapses the cardinals $\aleph_n$ for all $n \in \omega$, while preserving $\aleph_\omega$ as a cardinal. Thus $\aleph_\omega$ becomes the new $\aleph_1$. In the ground model, $\aleph_\omega = \bigcup_{n\in\omega} \aleph_n$, and this remains true as a relation on ordinals (which are absolute). So $\aleph_1$ is a singular cardinal. Start with a ground model satisfying GCH. It should come as no surprise that $\mathcal{P}(\omega)$, originally having cardinality $\aleph_1$, becomes a countable union of countable sets. For this example, I will employ $\aleph$ only for cardinality computations; ω_n will denote $(\omega_n)^M$ throughout (M being the ground model), even in the context of extensions of M. In other words, I'm treating ω_n as an "absolute name" for a specific ordinal, regardless of its cardinality. (Ditto ω_ω.)

Assume the ground model M satisfies AC and GCH. Recall how you collapse a cardinal $\mathfrak{n}$: you let $P = \text{Fn}(\omega, \mathfrak{n})$, then an easy density argument shows that $M[G]$ contains a function from ω onto $\mathfrak{n}$. To collapse all the cardinals ω_n at once, let P be this notion of forcing:

$$P = \{p \in \text{Fn}(\omega \times \omega, \omega_\omega) : (\forall n, k) p(n, k) \in \omega_n\}$$

So P combines all the $\text{Fn}(\omega, \omega_n)$ into one. From the usual density arguments, $M[G]$ has surjections (and hence bijections) $\omega \to \omega_n$ for each n. Since AC also holds in $M[G]$, it follows that $M[G]$ has a bijection $\omega \to \omega_\omega$. We want to install an automorphism group $\mathcal{G}$ as a gatekeeper to keep out this last bijection.

For the group $\mathcal{G}$, we allow all permutations that permute each row within itself: for each n, $\pi(n, k) = (n, \pi_n(k))$ for some permutation π_n of ω. Let H_n be the subgroup of permutations that fix all rows before n: $\pi_m \equiv \text{id}$ for all $m < n$. (To be fussy, I should say "α_π's corresponding to permutations π" etc., but you know what I mean.) Our normal filter $\mathcal{F}$ is generated by the H_n's, i.e., $H \in \mathcal{F}$ iff $H \supseteq H_n$ for some n. We will say that a name x has *support* H_n if $\text{sym}(x) \supseteq H_n$. Let N be the corresponding symmetric model.

The key lemma falls out of a "virgin territory" argument. If p is a condition, then $p \restriction n$ stands for $p \restriction (n \times \omega)$, i.e., $p \restriction n$ is defined only for (m, k) with $m < n$, where it agrees with p (so $p \geq p \restriction n$).

Finite support lemma: If $\varphi(x_1, \ldots, x_r)$ is a formula with free variables as shown, and $c_1, \ldots, c_r$ are symmetric names all having support H_n, and if $p \Vdash \varphi(c_1, \ldots, c_r)$, then $p \restriction n \Vdash \varphi(c_1, \ldots, c_r)$.

Proof: If $p \restriction n$ does not force $\varphi(c_1, \ldots, c_r)$, then $q \Vdash \neg\varphi(c_1, \ldots, c_r)$ for some $q \geq p \restriction n$. Choose an $\alpha_\pi \in H_n$ such that $\alpha_\pi q$ is compatible with p. (Virgin territory: for any (n', k) such that $p(n', k) \neq q(n', k)$, just exchange (n', k) with an (n', k') at which p is undefined. That makes $\alpha_\pi q$ compatible with p. We must have $n' \geq n$ because p and q both extend $p{\restriction}n$, so none of the locked rows are affected.)

Applying α_π to $q \Vdash \neg\varphi(c_1, \ldots, c_r)$, we have

$$\alpha_\pi q \Vdash \neg\varphi(\alpha_\pi c_1, \ldots, \alpha_\pi c_r)$$

but since all the c_i's have support H_n, this is the same as

$$\alpha_\pi q \Vdash \neg\varphi(c_1, \ldots, c_r)$$

So we have two supposedly compatible conditions p and $\alpha_\pi q$ forcing contradictory statements, contradiction. QED

Let P_n be those conditions with *support* n, i.e., domains contained in $n \times \omega$. The lemma tells us that one kind of finite support implies another kind.

Suppose x is a symmetric name for a subset of ω, with support H_n. Applying first the Truth Lemma and then the Finite Support Lemma, we conclude that for any $k \in \omega$, this holds: if $M[G] \models k \in \bar{x}$, i.e., $G \models \dot{k} \in x$, then $p \Vdash \dot{k} \in x$ for some $p \in G$, and so $p{\restriction}n \Vdash \dot{k} \in x$. This suggests the following definition for any x with support H_n (and $\bar{x} \subseteq \omega$):

$$\ddot{x} = \{\langle \dot{k}, p\rangle : p \in P_n \;\&\; p \Vdash \dot{k} \in x\}$$

It follows easily that $K_G(x) = K_G(\ddot{x})$. You may recall the "nice names" of §21.6 (p.96, these notes); these $\ddot{x}$'s share a family resemblance.[25]

Recall from §21.5 that $W \stackrel{\text{def}}{=\!=} \mathcal{P}^M(\text{Ra}(\dot{\omega}) \times P)$ is a warehouse for $\mathcal{P}^{M[G]}(\omega)$. That is, if $X \subseteq \omega$ in $M[G]$, then $X = \bar{x}$ for some $x \in W$. We need more

[25]On p.393, Halbeisen gives a less nice version of $\ddot{x}$, called "canonical names". Why the (unnecessary) extra hair? I think two issues are in play. In proving $K_G(x) \subseteq K_G(\ddot{x})$, we use the fact that $p \in G$ implies $p{\restriction}n \in G$. For $K_G(\ddot{x}) \subseteq K_G(x)$, this fact is needed: $k \in K_G(\ddot{x})$ implies that $\langle \dot{k}, p\rangle \in \ddot{x}$ for some $p \in G$. This falls right out of the definition of K_G. For names in general, $K_G(a) \in K_G(b)$ is a shaggier tale—we saw this in §21.3. But the simple structure of $\ddot{x}$ wards off such complications.

from W: we need it to contain a symmetric name for X if one exists at all. Suppose $X = \bar{x}$ where x has support H_n. Define:

$$x^* = \{\langle \dot{k}, p\rangle : k \in \omega \ \&\ p \Vdash \dot{k} \in x\}$$

We immediately have $K_G(x^*) = K_G(x)$ and $x^* \in W$. Furthermore, if $\alpha \in H_n$, then $\alpha x^* = x^*$, because

$$\begin{aligned} \alpha x^* &= \{\langle \dot{k}, \alpha p\rangle : k \in \omega \ \&\ p \Vdash \dot{k} \in x\} \\ &= \{\langle \dot{k}, \alpha p\rangle : k \in \omega \ \&\ \alpha p \Vdash \dot{k} \in \alpha x\} \\ &= \{\langle \dot{k}, \alpha p\rangle : k \in \omega \ \&\ \alpha p \Vdash \dot{k} \in x\} \\ &= x^* \end{aligned}$$

using the fact that as p ranges over all conditions, so does αp.

Now define:

$$R_n = \{\langle \ddot{x}, \emptyset\rangle : x \in W \text{ is a name with support } H_n\}$$

It follows that $\bigcup \bar{R}_n = \mathcal{P}(\omega) \cap N$. We have a few things to check: the R_n are all sets in M, are all in HS, and the sequence $(R_n : n \in \omega)$ is in HS. The Definability Lemma tells us that $x \mapsto \ddot{x}$ is a functional in M. For any n, the set of all names in the warehouse with support H_n is itself a set, because $\mathcal{G}$ belongs to M (by our standing conventions for symmetric models). R_n results from applying a functional to a set in M, and so is itself a set in M.

OK, how about symmetry? We know that $p \Vdash \dot{k} \in x$ iff $\alpha p \Vdash \dot{k} \in \alpha x$. This implies that $\alpha\ddot{x} = (\alpha x)^{\cdot\cdot}$. Also, if x has support H_n so does αx. Therefore $\alpha R_n = R_n$, for all $\alpha \in \mathcal{G}$, not just all $\alpha \in H_n$. This makes the sequence $(R_n : n \in \omega)$ symmetric. As for hereditary symmetry, $\dot{k}$ has support $\mathcal{G}$, $\ddot{x}$ has support H_n whenever x does, and we've just seen that R_n has support $\mathcal{G}$. So everything we need is in HS, and $N \models \mathcal{P}(\omega) = \bigcup \bar{R}_n$.

Next we show that each ω_n is countable in N. We already observed that we have surjective functions $\bar{h}_n : \omega \to \omega_n$ in $M[G]$ for each n. What do the names look like? If $p(n,k) = \gamma \in \omega_n$, then p says that $\bar{h}_n(k)$ should equal γ. The name for $\bar{h}_n$ makes use of $\mathrm{op}(x,y)$ (Halbeisen p.343) which satisfies $K_G(\mathrm{op}(x,y)) = \langle K_G(x), K_G(y)\rangle$.

$$h_n = \{\langle \mathrm{op}(\dot{k},\dot{\gamma}), p\rangle : p(n,k) = \gamma\}$$

Locking row n is enough to make this immune to permutations, so h_n is symmetric (indeed, HS) and $\bar{h}_n$ belongs to N. Now in general, to go from a surjection $h : A \to B$ to a bijection $A \leftrightarrow B$ requires AC. But in the special case where $A = \omega$, no problem: just "cross out duplicates" in the enumeration of B.[26]

Next on the agenda: each $\bar{R}_n$ is countable in N. First we compute $|R_n|$, using a familiar technique: calculate inside M, where AC holds, and then apply what we learn to N. (The computation may remind you of the calculation of $2^{\aleph}_0$ in §21.6.) First, each condition in P_n is basically a finite sequence of elements of ω_n, so $|P_n| = \aleph_n$. Next, each $\ddot{x}$ in R_n is a subset of $\mathrm{Ra}(\dot{\omega}) \times P_n$, so we have $|R_n| \leq 2^{\aleph_n} = \aleph_{n+1}$, since GCH holds in M. In fact, $|R_n| = \aleph_{n+1}$; I leave it as an exercise to get the lower bound, but it doesn't really matter, as we'll see in a moment.

Suppose that $|R_n| = \aleph_m$ in M for some m, i.e., there's a bijection $\omega_m \leftrightarrow R_n$ in M. In N there is a bijection $\omega \leftrightarrow \omega_m$, as we've seen, and $M \subseteq N$, so in N we have a bijection $\omega \leftrightarrow R_n$.

So R_n is countable in N. How about $\bar{R}_n$? Since $\bar{R}_n = K_G(R_n)$, our plan is to compose $\omega \to R_n$ with K_G, getting a surjection, and then remove duplicates. One more niggle: the map K_G belongs to $M[G]$, but not to N.[27] We show next that $K_G{\restriction}R_n$ *does* belong to N.

[26]Formally, let $e(0) = 0$ and set inductively $e(n+1) =$the first $j > n$ such that $h(j)$ is not equal to any of $h(e(0)), \ldots, h(e(n))$. Then $f(n) = h(e(n))$ is a bijection between ω and B. This works for maps from any ordinal, not just from ω.

[27]Recall how Shoenfield proves that AC holds in $M[G]$ (p.365). If K_G belonged to N, the same argument would make AC hold in N as well.

Here's a name for $K_G{\restriction}A$, for any $A \in M$.

$$\hat{K}(A) = \{\langle \text{op}(\dot{a}, a), \emptyset \rangle : a \in A\}$$

Applying K_G to $\hat{K}(A)$, it unconditionally places $\langle K_G(\dot{a}), K_G(a)\rangle = \langle a, \bar{a}\rangle$ into the result for each $a \in A$. In other words, $K_G(\hat{K}(A))$ is just $K_G{\restriction}A$. Now if $A = R_n$, then locking all rows below n guarantees that each $a \in R_n$ will stay fixed. (I.e., $\alpha \in H_n \Rightarrow \alpha a = a$ for all $a \in R_n$.) This implies that $\hat{K}(R_n)$ is symmetric.

Conclusion: in N, $\mathcal{P}(\omega)$ is a countable union of countable sets. Feferman and Lévy also showed that ω_ω is a cardinal in N, and thus $|\omega_\omega| = \aleph_1$. The proof again is reminiscent of the computation of $2^{\aleph_0}$ in §21.6. Suppose $\bar{h} : \omega \twoheadrightarrow \omega_\omega$ is a surjection in $M[G]$. (We noted earlier that ω_ω is countable in $M[G]$, since AC holds there.) Hence there is a condition $G \ni p \Vdash h : \omega \twoheadrightarrow \omega_\omega$. Since $\bar{h}$ is surjective, for each $\gamma \in \omega_\omega$ there is a $k_\gamma \in \omega$ and a $q_\gamma \geq p$ such that $q_\gamma \Vdash h(\dot{k}_\gamma) = \dot{\gamma}$. The rest of the proof takes place entirely inside M. By a pigeonhole argument, there is a single $k \in \omega$ such that $q_\gamma \Vdash h(\dot{k}) = \dot{\gamma}$ for $\aleph_\omega$ different values of γ, let's say for all $\gamma \in \Gamma$. Clearly all the conditions in $\{q_\gamma : \gamma \in \Gamma\}$ are incompatible.

Suppose that h is symmetric, with support H_n. By the Finite Support Lemma, we have $q_\gamma{\restriction}n \Vdash h(\dot{k}) = \dot{\gamma}$ for all $\gamma \in \Gamma$. So all the $q_\gamma{\restriction}n$ are incompatible; in particular, they are all different. But they all belong to P_n, and as we saw, $|P_n| = \aleph_n$. So there cannot be $\aleph_\omega$ different $q_\gamma{\restriction}n$'s.

I like the "venue shopping" in this proof. Up in $M[G]$, $\bar{h}(k)$ can have only one value—no incompatibilities permitted. $M[G]$ avoids inconsistency, and allows $\bar{h}$ to be surjective, by making ω_ω countable. No pigeonhole argument here! Back down in M, differing cardinalities create a "collision" of k_γ's, resulting in massive incompatibilities. But since we're now dealing with forcing conditions and not function values, again no inconsistency. The hypothesis of a symmetric name for $\bar{h}$ yields an inconsistency involving cardinalities. Finally, since M enjoys both AC and GCH, cardinal computations are a breeze.

Halbeisen indicates another proof that $(\omega_1)^N$ is singular (Related Results 97, p.401).[28] Fact: Any ZF model N contains a surjection of $\mathcal{P}(\omega)$ onto $(\omega_1)^N$. Proof: $\mathcal{P}(\omega)$ is bijectively equivalent to $\mathcal{P}(\omega \times \omega)$. If a subset of $\omega \times \omega$ is a well-ordering of ω, map it to its order type; otherwise, map it to 0. This is clearly surjective. No part of this argument uses AC. Since $\mathcal{P}(\omega) = \bigcup \bar{R}_n$ in N, we can use the surjection $\mathcal{P}^N(\omega) \to (\omega_1)^N$ to express $(\omega_1)^N$ as a countable union of countable ordinals.

Halbeisen gives one more symmetric model, plus the Jech-Sochor construction for converting permutation models into symmetric ones. But I think we've seen enough to taste the flavor.

21.7.5 HOD Models

A set x is **ordinal definable** (OD) if for some pure $\varphi(x, y_1, \ldots, y_n)$, x is the unique set satisfying

$$\varphi(x, \alpha_1, \ldots, \alpha_n)$$

for some ordinals $\alpha_1, \ldots, \alpha_n$. Equivalently, x is OD if there is a pure formula $\psi(z, y_1, \ldots, y_n)$ such that

$$x = \{z : \psi(z, \alpha_1, \ldots, \alpha_n)\}$$

and again $\alpha_1, \ldots, \alpha_n$ are ordinals. The levels V_α of the cumulative hierarchy offer the easiest example, since

$$V_\alpha = \{z : \mathrm{rk}(y) < \alpha\}$$

If x is OD and all its elements are OD and so on all the way down, then we say x is **hereditarily ordinal definable** (HOD). Formally, x is HOD if x and all the elements of its transitive closure are OD. You'd probably like an example of something OD but not HOD; we'll return to this momentarily.

[28]Here $(\omega_1)^N$ means "ω_1 in N", the explicit superscript N overriding the "absolute name" convention introduced at the start of this example.

One easy observation: since all the V_α's are OD, the assumption that all OD sets are HOD implies that *all* sets are HOD.

Gödel first suggested the OD and HOD classes, as a quicker route to the relative consistency of AC. Myhill and Scott worked out the details. If V satisfies ZF, then HOD is a transitive inner model satisfying ZFC. HOD satisfies ZF because, roughly speaking, the "big" axioms (Union, Separation, Replacement, Power Set) give explicit constructions of new sets from old. If the old sets have definitions from ordinals, it's easy top it off with the construction of the new set. For example, Power Set. Suppose x is in HOD, so $x = \{y : \varphi(y, \vec{\alpha})\}$ for ordinals $\vec{\alpha}$. Then $\mathcal{P}(x) = \{u : (\forall y \in u)\varphi(y, \vec{\alpha})\}$. (Compare the justification of the ZF axioms for permutation and symmetric models.)

HOD has a definable well-ordering (definable without parameters); this is stronger than AC. Proof: just use lexicographic ordering on $\langle \ulcorner\varphi\urcorner, \vec{\alpha}\rangle$, where $\ulcorner\varphi\urcorner$ as usual is a coding of φ. (Even easier than the proof that the constructible universe has a definable well-ordering!) So every subclass of HOD has a definable well-ordering, but the converse also holds: if a transitive class K has a definable well-ordering, it is a subclass of HOD. Proof: if K is a proper class, we can use the well-ordering to define a mapping F from Ω onto K; then obviously every $F(\alpha) \in K$ is OD, and since K is transitive, all elements of K are HOD. If K is a set, then we have a mapping from some ordinal γ onto K, and the rest of the argument is unchanged. Corollary: since L is a transitive class with a definable well-ordering, all constructible sets are HOD. Since $V = L$ is consistent with ZF, it's consistent to assume that all sets are HOD. In our quest for an OD but non-HOD set, we see the best we can hope for is a model containing such a set.

The very first Cohen model, the one that added a single new subset of ω, fulfills our desire. The "slipperiness" of the new real prevents it from being OD, just as it couldn't be constructible. (We'll go into detail shortly.) So $\mathcal{P}(\omega)$ in this model is OD but not HOD.

If M satisfies AC, so does $M[G]$ for any generic G; we learned this in §21.5. Hence there is a bijection $f : \gamma \to \mathcal{P}(\omega)$ for some ordinal γ. If this f were OD, then every $f(\alpha)$ would also be OD. We see that the Cohen model with a single new real has a well-ordering of $\mathcal{P}(\omega)$, but no OD well-ordering of it, let alone one definable without parameters.

Now let's turn to Shoenfield's treatment in §9. He has automorphisms of the set P of conditions (which he calls C). He notes that $M[G] = M[\pi G]$ for any automorphism π, as we observed above: see the commutative diagram illustrating the equation

$$K_G(a) = K_{\pi G}(\pi a)$$

Suppose $\pi a = a$. Then $K_G(a) = K_{\pi G}(\pi a) = K_{\pi G}(a)$. If $K_G(a) = K_{\pi G}(a)$ for all π in some set $\mathfrak{A}$ of automorphisms, then Shoenfield says that a is an **$\mathfrak{A}$-invariant name**. With this definition, he avoids explicitly extending the automorphisms from P to the class of names. I don't see a easy way to reverse the implication—to show that an $\mathfrak{A}$-invariant name is left fixed by all $\pi \in \mathfrak{A}$—but in practice, this symmetry holds for all of Shoenfield's applications.

In a virgin territory argument, we would find an automorphism α making p and αq compatible for some given conditions p and q. Shoenfield says that P is **$\mathfrak{A}$-homogeneous** if this is always possible with $\pi^{-1} = \alpha$, $\pi \in \mathfrak{A}$. (Don't stress over the inverse—Shoenfield's $\mathfrak{A}$ is always a group in practice.) He notes that $\mathrm{Fn}(A, B)$ ($H(A, B)$ in his notation) is $\mathfrak{A}$-homogeneous with $\mathfrak{A}$ the group of all automorphisms induced by permutations of A. He encapsulates the virgin territory argument in Lemma 9.2:

> If $\mathfrak{A} \in M$ and P is $\mathfrak{A}$-homogeneous and all the names in the closed formula $\varphi(\vec{a})$ are $\mathfrak{A}$-invariant, then
>
> $$G \models \varphi(\vec{a}) \Leftrightarrow \emptyset \Vdash \varphi(\vec{a})$$

Replacing "$\mathfrak{A}$-invariant" with "symmetric under $\mathfrak{A}$", this should look quite familiar. If $G \models \varphi(\vec{a})$, then some $G \ni p \Vdash \varphi(\vec{a})$. If $\emptyset \nVdash \varphi(\vec{a})$, then some

$q \Vdash \neg\varphi(\vec{a})$. Find an α such that p and αq are compatible. Since $q \Vdash \neg\varphi(\vec{a})$, $\alpha q \Vdash \neg\varphi(\alpha\vec{a})$, and since the $\vec{a}$'s are symmetric under $\mathfrak{A}$, $\alpha q \Vdash \neg\varphi(\vec{a})$. So φ and $\neg\varphi$ are forced by supposedly compatible conditions. Shoenfield's argument is just a touch more intricate, but essentially the same.

Lemma 9.3 and its corollary Theorem 9.1 are straightforward. They become even clearer if we start with a special case: if $\bar{u} \subseteq \omega$ is OD in $M[G]$, then in fact $\bar{u} \in M$. For in that case, $G \models \dot{n} \in u$ iff $G \models \varphi(\dot{n}, \dot{\alpha}_1, \ldots, \dot{\alpha}_r)$ iff $\emptyset \Vdash \varphi(\dot{n}, \dot{\alpha}_1, \ldots, \dot{\alpha}_r)$, and so

$$\bar{u} = \{n \in \omega : \emptyset \Vdash \varphi(\dot{n}, \dot{\alpha}_1, \ldots, \dot{\alpha}_r)\}$$

a set in M by the Definability Lemma. (This already covers our comments about the Cohen model with one new real.) To get Lemma 9.3, make two changes to the argument. First, we don't know that all the elements of $\bar{u}$ already belong to M, so we reason about $\bar{u} \cap M$. Second, φ could contain other invariant names besides $\dot{\alpha}_1, \ldots, \dot{\alpha}_r$; these are the names of the v_i's.

Recall how the independence proofs went for the symmetric models, say the Basic Cohen Model (BCM). We have a group $\mathcal{G}$ of automorphisms, and a normal filter $\mathcal{F}$ of subgroups. $\mathcal{G}$ is essentially the same as the group of permutations of ω; I'll make no distinction between the ω-permutation and the corresponding automorphism of P or of V^P. Just note that $\pi \in \mathcal{G}$ (as automorphism of V^P) permutes the names $\{a_0, a_1, \ldots\}$ the same way it (as permutation of ω) permutes $\{0, 1, \ldots\}$.

For the BCM, we took P to be the set $\mathrm{Fn}(\omega \times \omega \to 2)$. Shoenfield's $H(\omega, 2)^\omega$ amounts to the same thing: $p_n(k)$ instead of $p(n, k)$. Note that $H(\omega, 2)^\omega$ is a *weak power* as defined at the top of p.368, where all but finitely many of the p_n's must be the empty condition.

The G_n of Shoenfield are the same as the $\bar{a}_n$ of the BCM, and Shoenfield's H is $\{\bar{a}_0, \bar{a}_1, \ldots\}$. In the BCM, the $\bar{a}_n$ don't have *fully* symmetric names, an obstacle hurdled via the normal filter of subgroups. We have a similar hitch in the HOD model: none of the G_n are OD. (Otherwise, by Lemma

9.3, the G_n would belong to M, but a standard density argument shows they don't.) We circumvent this problem by generalizing the notion of OD. A set x is **OD from** $v_1, \ldots, v_k$ if there are ordinals $\alpha_1, \ldots, \alpha_n$ and a pure formula $\varphi(x, y_1, \ldots, y_n, z_1, \ldots, z_k)$ such that x is the unique set satisfying

$$\varphi(x, \alpha_1, \ldots, \alpha_n, v_1, \ldots, v_k)$$

Equivalently,

$$x = \{z : \psi(z, \alpha_1, \ldots, \alpha_n, v_1, \ldots, v_k)\}$$

for some pure formula $\psi(x, y_1, \ldots, y_n, z_1, \ldots, z_k)$. (Shoenfield's definition on p.369 is equivalent.) Also, x is **OD over** w if x is OD from $w, v_1, \ldots, v_k$ for some $v_1, \ldots, v_k \in w$.

Shoenfield lets N be all the sets HOD over $H = \{G_0, G_1, \ldots\}$. This N corresponds to the BCM. Continuing the dictionary, an OD definition from $G_0, \ldots, G_{n-1}$ corresponds to a symmetric name with support n, i.e., one left fixed by any automorphism fixing rows 0 to $n-1$. HOD corresponds to HS (hereditarily symmetric). As we will see in a moment, Lemma 9.3 plays the gatekeeper role of the symmetry requirement, and the Product Theorem of §8 takes over the job of the normal filter of subgroups.

Let's recap the basic plot of the proof that $\mathcal{P}(\omega)$ has no well-ordering in the BCM. If f is a symmetric name for a function from an ordinal onto $\mathcal{P}(\omega)$, then f is left fixed by any automorphism fixing the rows below row n, for some n. Some $\bar{f}(\gamma)$ equals $\bar{a}_n$; we obtain a contradiction with a virgin territory argument, finding an automorphism that moves a_n while leaving the locked rows unmoved. Now translate to N. The symmetric name becomes a definition of $\bar{f}$ that is OD from $G_0, \ldots, G_{n-1}, H$. This implies that $\bar{f}(\gamma) = G_n$ is OD from $G_0, \ldots, G_{n-1}, H$. The split between the locked rows and the other rows becomes a decomposition of N via the Product Theorem. Specifically, C' represents rows below n, C'' rows n and above. In $M' = M[G']$, the rows below n are effectively locked because $G_0, \ldots, G_{n-1}$ belong to it. The Product Theorem tells us that $M[G] = M'[G'']$, with G'' generic over M'. Shoenfield's set $\mathfrak{A}$ of automorphisms is just the subgroup H_n of automorphisms fixing the locked rows. The locked status of rows below n translates into invariant names for $G_0, \ldots, G_{n-1}$. The $\mathfrak{A}$-homogeneity of C'' permits us to apply Lemma 9.3, concluding that G_n belongs to M'. On the other hand, G_n is generic over M', and we have our contradiction.

Personally, I find the symmetric model proofs more intuitive.

21.8 Shoenfield Miscellany

I don't find the treatment in Shoenfield especially hard, but here are some remarks to reduce friction to a minimum.

p.358

The union of $<\mathfrak{m}$ sets of cardinality $<\mathfrak{m}$ has cardinality $<\mathfrak{m}$, if $\mathfrak{m}$ is regular. Let $U = \bigcup_{\alpha<\mathfrak{n}} P_\alpha$ be such a union, with $\mathfrak{p}_\alpha = |P_\alpha|$, and $\mathfrak{n} < \mathfrak{m}$ and all $\mathfrak{p}_\alpha < \mathfrak{m}$. Let $W = \bigsqcup_{\alpha<\mathfrak{n}} \mathfrak{p}_\alpha$ be the disjoint union, i.e., $W = \{(\chi, \alpha) : \chi < \mathfrak{p}_\alpha, \alpha < \mathfrak{n}\}$. Obviously there is a map from W onto U, so it's enough to show $|W| < \mathfrak{m}$. Lexicographically order W, with all elements of $\mathfrak{p}_\beta$ preceding all elements of $\mathfrak{p}_\alpha$ when $\beta < \alpha$. There is a map $\mathfrak{n} \to W$ with a cofinal image: map $\alpha \mapsto (0, \alpha)$, and note that if (χ, β) is any element of W, then $(0, \alpha)$ comes after it for any $\alpha > \beta$. This in turn gives us a cofinal subset of the order type of W (call it γ) of cardinality $\mathfrak{n}$. Since $\mathfrak{m}$ is regular and $\mathfrak{n} < \mathfrak{m}$, it follows that $\gamma < \mathfrak{m}$, and so $|W| \leq \gamma < \mathfrak{m}$.

${}^{\mathfrak{n}}\mathfrak{m} = \bigcup_{\alpha<\mathfrak{m}} {}^{\mathfrak{n}}\alpha$. If $f : \mathfrak{n} \to \mathfrak{m}$, then the image of f has an upper bound less than $\mathfrak{m}$, say $f(\mathfrak{n}) \subseteq \alpha < \mathfrak{m}$. But then $f \in {}^{\mathfrak{n}}\alpha$.

$|\alpha|^{\mathfrak{n}} \leq 2^{|\alpha|\cdot\mathfrak{n}}$. Indeed, $|\alpha| < 2^{|\alpha|}$, so $|\alpha|^{\mathfrak{n}} \leq \left(2^{|\alpha|}\right)^{\mathfrak{n}} = 2^{|\alpha|\cdot\mathfrak{n}}$.

p.375

The point of the top paragraph bears repeating. If we say $\mathfrak{c} = \aleph_1$ in M but

$\mathfrak{c} = \aleph_2$ in $M[G]$, what do these equations really mean, formally speaking? What does it mean to say that $\mathfrak{c}$ changes going from M to $M[G]$, but $\aleph_1$ and $\aleph_2$ stay the same? Shoenfield's answer: imagine adding new constants $\mathfrak{c}$, $\aleph_1$, and $\aleph_2$ to the language of ZF, along with axioms describing these constants: "$\mathfrak{c}$ is the cardinality of $\mathcal{P}(\omega)$", "$\aleph_1$ is the first uncountable ordinal", "$\aleph_2$ is the second uncountable ordinal". (Any knowledgeable masochist could translate these prose descriptions into assertions in the language of ZF augmented with the new constants.) In ZFC plus the new axioms we can prove that the descriptions single out items uniquely, that the items are cardinals, and the other stuff needed for Theorem 11.1. The preservation of cardinals means that $(\aleph_1)^M = (\aleph_1)^{M[G]}$ and $(\aleph_2)^M = (\aleph_2)^{M[G]}$, but obviously the constant $\mathfrak{c}$ is not preserved. (Cohen [4, p.135] makes similar remarks.)

21.9 Halbeisen, Fact 14.3, p.278

Definition: $h : P \to Q$ is a **dense embedding** if

- $(\forall p_0, p_1 \in P)[p_0 \leq p_1 \leftrightarrow h(p_0) \leq h(p_1)]$, and
- $(\forall q \in Q)(\exists p \in P)[q \leq h(p)]$. That is, $h(P)$ is dense.

Note that despite the term "embedding", h might not be injective. But if we let $p_0 \equiv p_1$ be the equivalence relation $p_0 \leq p_1 \wedge p_1 \leq p_0$, then $p_0 \equiv p_1 \leftrightarrow h(p_0) \equiv h(p_1)$.

Fact 14.3: Let P and Q be any two forcing notions. If there exists a dense embedding $h : P \to Q$ belonging to $\mathbf{V}$, then P and Q are equivalent. In fact, if $G \subseteq P$ is P-generic over $\mathbf{V}$, then the set

$$H = \{q \in Q : (\exists p \in G) q \leq h(p)\} \tag{5}$$

is Q-generic over $\mathbf{V}$ and $\mathbf{V}[G] = \mathbf{V}[H]$. Conversely, if a set $H \subseteq Q$ is Q-generic over $\mathbf{V}$, then the set

$$G = \{p \in P : h(p) \in H\} \tag{6}$$

is P-generic over $\mathbf{V}$ and $\mathbf{V}[H] = \mathbf{V}[G]$.

Proof: First note that eq.5 says that H is the downward closure of $h(G)$, while eq.6 says that G is $h^{-1}(H)$.

(1): Assume G is P-generic and let $H = \wedge h(G)$. We show H is Q-generic.

(1a): H is a filter. Obviously H is downward closed. Suppose $q_1, q_2 \in H$, say $q_1 \leq h(p_1)$, $q_2 \leq h(p_2)$, with $p_1, p_2 \in G$. Since G is a filter $p_1, p_2 \leq p$ for some $p \in G$, so $q_1, q_2 \leq h(p)$ and $h(p) \in H$. So H is upward directed.

(1b): H is generic. Suppose $E \subseteq Q$ is dense open and belongs to $\mathbf{V}$. Let $D = h^{-1}(E)$. D belongs to $\mathbf{V}$ because E and h do. D is dense:

$$\begin{aligned} p' \in P &\Rightarrow (\exists q \in E)\, h(p') \leq q \quad (E \text{ is dense}) \\ &\Rightarrow (\exists p)\, h(p') \leq q \leq h(p) \quad (h(P) \text{ is dense}) \\ &\Rightarrow h(p') \leq h(p) \in E \quad (E \text{ is open}) \\ &\Rightarrow p' \leq p \in D \end{aligned}$$

D is open:

$$\begin{aligned} D \ni p' \leq p &\Rightarrow E \ni h(p') \leq h(p) \\ &\Rightarrow h(p) \in E \quad (E \text{ is open}) \\ &\Rightarrow p \in D \end{aligned}$$

Therefore $G \cap D \neq \varnothing$; if $p \in G \cap D$, then $h(p) \in H \cap E$. So H meets every dense open set in $\mathbf{V}$.

(1c): $\mathbf{V}[G] = \mathbf{V}[H]$. The argument will use this fact about G and H, and *only* this fact: $p \in G \Leftrightarrow h(p) \in H$. (I.e., $G = h^{-1}(H)$.) The $\Rightarrow$ direction is immediate. For $\Leftarrow$, say $h(p) \in H$, so $h(p) \leq h(p')$ for some $p' \in G$, so $p \leq p' \in G$ and so $p \in G$ because G is downward closed.

For any $x \in \mathbf{V}^P$ define $\hat{x}$ inductively by

$$\hat{x} = \{\langle \hat{y}, h(p) \rangle : \langle y, p \rangle \in x\}$$

We will show (by induction) that $\hat{x}[H] = x[G]$; this implies that $\mathbf{V}[G] \subseteq \mathbf{V}[H]$.

For starters, transcribing the definitions gives us:

$$x[G] = \{y[G] : (\exists p \in G)\langle y, p \rangle \in x\}$$
$$\hat{x}[H] = \{v[H] : (\exists q \in H)\langle v, q \rangle \in \hat{x}\}$$

By the definition of $\hat{x}$, $\langle v, q \rangle$ must be of the form $\langle \hat{y}, h(p) \rangle$ with $\langle y, p \rangle \in x$, so the second equation becomes

$$\hat{x}[H] = \{\hat{y}[H] : (\exists h(p) \in H)\langle y, p \rangle \in x\}$$

By inductive hypothesis, $\hat{y}[H] = y[G]$. Also we noted that $h(p)$ is in H iff p is in G. So:

$$\hat{x}[H] = \{y[G] : (\exists p \in G)\langle y, p \rangle \in x\} = x[G]$$

To show that $\mathbf{V}[H] \subseteq \mathbf{V}[G]$, we will define $\tilde{u}$ for any $u \in \mathbf{V}^Q$ and show that $\hat{\tilde{u}} = u$. It then follows that

$$u[H] = \hat{\tilde{u}}[H] = \tilde{u}[G]$$

so $\mathbf{V}[H] \subseteq \mathbf{V}[G]$.

Define inductively

$$\tilde{u} = \{\langle \tilde{v}, p \rangle : \langle v, h(p) \rangle \in u\}$$

Then

$$\begin{aligned}
\hat{\tilde{u}} &= \{\langle \hat{y}, h(p) \rangle : \langle y, p \rangle \in \tilde{u}\} \\
&= \{\langle \hat{y}, h(p) \rangle : \langle y, p \rangle = \langle \tilde{v}, p \rangle \text{ with } \langle v, h(p) \rangle \in u\} \\
&= \{\langle \hat{\tilde{v}}, h(p) \rangle : \langle v, h(p) \rangle \in u\} \\
&= \{\langle v, h(p) \rangle : \langle v, h(p) \rangle \in u\} = u
\end{aligned}$$

(2): Assume H is Q-generic and let $G = h^{-1}(H)$. We show G is P-generic.

(2a): G is a filter. G is downward closed because if $p' \leq p$ with $h(p) \in H$, then $h(p') \leq h(p) \in H$ and so $h(p') \in H$ because H is downward closed. So $p' \in G$.

G is upward directed:

$$\begin{aligned} p_1, p_2 \in G &\Rightarrow h(p_1), h(p_2) \in H \\ &\Rightarrow (\exists q \in H)\, q \geq h(p_1), h(p_2) \quad (H \text{ is upward directed}) \\ &\Rightarrow (\exists h(p) \in H)\, h(p) \geq q \quad (\text{see below}) \\ &\Rightarrow p \geq p_1, p_2 \end{aligned}$$

The crucial line above says that not only is there an $h(p) \geq q$, but we can also demand $h(p) \in H$. The first assertion holds because h is a dense embedding, so $h(P)$ is dense. For the second, we appeal to the corollary to Fact 14.7 (see §21.10), and the assumption that H is generic. (Although Fact 14.7 appears after Fact 14.3 in Halbeisen, its proof is independent.)

(2b): G is generic. Let D be open dense. Let $E = \vee h(D)$. Obviously E is open. E is dense:

$$\begin{aligned} q \in Q &\Rightarrow (\exists p)\, h(p) \geq q \quad (h(P) \text{ is dense}) \\ &\Rightarrow (\exists p' \in D)\, p' \geq p \quad (D \text{ is dense}) \\ &\Rightarrow E \ni h(p') \geq h(p) \geq q \end{aligned}$$

So there is a $q' \in H \cap E$ since H is generic. Since $E = \vee h(D)$, there is a $q'' \in h(D)$ with $q'' \leq q'$ and $q'' \in h(D)$. Since H is downward closed, $q'' \in H$. So $q'' \in h(D) \cap H$, i.e., there is a $p'' \in D$ with $h(p'') \in H$, which means that $p'' \in D \cap G$. So G meets every open dense set.

(2c): $\mathbf{V}[H] = \mathbf{V}[G]$. The proof of (1c) applies without change, since the only fact used about G and H is that $G = h^{-1}(H)$.

21.10 Halbeisen, Fact 14.7, p.280

Definition: D is **dense above** p iff for any $p' \geq p$, D meets the upward cone of p'.

Fact 14.7: Let G be a filter containing p. G is generic over $\mathbf{V}$ iff G meets every D (in $\mathbf{V}$) that is dense above p.

Proof: In one direction this is trivial (given Fact 14.6), since if D is dense then it is dense above p. So if G meets every set in $\mathbf{V}$ that is dense above p, then G meets every dense set (in $\mathbf{V}$).

In the other direction, suppose D (in $\mathbf{V}$) is dense above $p \in G$. There is a maximal antichain A (in $\mathbf{V}$) containing p. Consider $D' = D \cup \vee(A \smallsetminus \{p\})$. D' is dense: given $q \in P$, $\vee q$ must meet $\vee A$, otherwise we could add q to A to get a strictly larger antichain. If $\vee q$ meets $\vee(A \smallsetminus \{p\})$, then *a fortiori* $\vee q$ meets D'. Otherwise, $\vee q$ meets $\vee p$, so there is a p' with $q \leq p'$ and $p \leq p'$. But since D is dense above p, there is a $q' \in D$ with $q' \geq p' \geq p$. So again $\vee q$ meets D'.

Clearly D' is in $\mathbf{V}$, so G meets D'. Now the upward cones of any two elements of A are disjoint, and the upward cones of any two elements of G meet (since filters are upward directed). Since G contains p, it follows that G is disjoint from $\vee(A \smallsetminus \{p\})$, and thus G meets D.

Corollary: If G is generic with $p \in G$, and D is dense (in $\mathbf{V}$), then there is a $q \in G \cap D \cap \vee p$.

Proof: $D \cap \vee p$ is dense above p.

21.11 Other Approaches

Cohen

I refer mainly to *Set Theory and the Continuum Hypothesis* [4]. I offer a quick sketch, out of historical interest.

Cohen's treatment differs from Shoenfield's chiefly in the names (or labels) for elements of the model. Cohen defines a *label space* S by transfinite induction, patterned after Gödel's L_α hierarchy. Recall the crucial inductive definition for L: x belongs to $L_{\alpha+1}$ if there is a formula $A(y, c_1, \ldots, c_m)$ with one free variable and with constants $c_1, \ldots, c_m$ belonging to L_α, such that

$$x = \{y \in L_\alpha : L_\alpha \models A(y, c_1, \ldots, c_m)\}$$

Let's say we try to sprinkle new symbols into this construction, representing new generic sets. Unlike L, we won't induct through *all* the ordinals, just those belonging to the ground model M; let α_0 be the suprenum of these ordinals. The label space S will be $\bigcup_{\alpha<\alpha_0} S_\alpha$. The roster of S_0 varies depending on the application. (Simplest case: S_0 contains a single generic symbol a.) If we have a formula $A(y, c_1, \ldots, c_m)$, with all $c_i \in S_\alpha$, we're going to want a set corresponding to it in the final model N. So we let the formula itself be a new label in $S_{\alpha+1}$.

Note that the bound quantifiers in this $A(y, c_1, \ldots, c_m)$ are implicitly restricted to level α. Cohen makes this explicit by writing $\exists_\alpha$ and $\forall_\alpha$. He introduces *limited statements*, where all quantifiers are restricted this way, and *unlimited statements* with the quantifiers ranging over the whole model N.

In addition to the "formula" labels, we have labels (formal symbols) for generic sets; these can show up at any level. Conditions are finite sets of formulas $c_1 \in c_2$, with c_2 a generic symbol and c_1 belonging to an earlier level.

Now we're ready to let the machinery roll. Cohen defines forcing inductively, much as we've seen above. He shows, using the countability of M, that there is a *complete sequence* $\{P_n\}$ of conditions—basically a generic filter. $\{P_n\}$ determines a map $c \mapsto \bar{c}$ with domain S. The complete sequence determines $\bar{c}$ for generic symbols c in the obvious way. For a formula label $x = A(y, c_1, \ldots, c_m)$ belonging to $S_{\alpha+1}$,

$$\bar{x} = \{y \in \bar{S}_\alpha : \bar{S}_\alpha \models A(y, \bar{c}_1, \ldots, \bar{c}_m)\}$$

where $\bar{S}_\alpha = \{\bar{s} : s \in S_\alpha\}$. The final model N is the image of S under the map $c \mapsto \bar{c}$.

I've left out many details. Inductions first run over limited statements, then once more for the unlimited case. Forcing for atomic statements took some juggling in §21.3; this becomes even more acrobatic. The CH arguments look nearly the same, with the countable chain condition playing an essential role. Curiously, the independence of AC takes a simpler form, a blend of the symmetric and HOD approaches. The model has generic labels $\{a_0, \ldots\}$ and V; conditions are adjusted to insure that $\bar{V} = \{\bar{a}_0, \ldots\}$ in N. We consider permutations of $\{a_0, \ldots\}$ and extend these to automorphisms of S, automorphisms in the following sense:

$$p \Vdash \varphi \Leftrightarrow \pi p \Vdash \pi\varphi$$

All elements of S are "symmetric enough" automatically! For generic symbols this is obvious. For formula labels it also falls out easily: to keep $A(y, c_1, \ldots, c_m)$ fixed, it's enough to keep $c_1, \ldots, c_m$ fixed. We can always achieve this by locking a finite number of the a_i's in place. This is analogous to the invariance of HOD definitions.

Boolean-Valued Models

I give a cursory sketch; for more details, consult any of [3, 13, 17, 21, 26].

Let B be a complete boolean algebra, let D be some set, and let $f : D \to B$. If B is the simplest boolean algebra $\{0, 1\}$, then f amounts to a subset of D. If B is more complex, then f is a kind of "generalized subset". The collection of all these f's gives a "generalized power set" of D.

Let's say we build the cumulative hierarchy, but using the generalized power set instead of the usual one. So $V_0^B = \varnothing$; $V_{\alpha+1}^B$ is the generalized power set of V_α^B; $V_\lambda^B = \bigcup_{\alpha<\lambda} V_\alpha^B$ for a limit ordinal λ; $V^B = \bigcup_{\alpha<\Omega} V_\alpha^B$.

You can regard the elements of V^B as "names" or "labels". Let's say we allow such labels to appear in ZF formulas; this gives us a *forcing language.* We associate a boolean value in B with any closed formula of the forcing language, say $\|\varphi\|$. Assigning boolean values to atomic formulas $a = b$ and $a \in b$ requires a transfinite induction; it's similar to the definitions of $p \Vdash a = b$ and $p \Vdash a \in b$.

Suppose you want to obtain a classical model. First step: replace V with a standard model M of ZF. Construct M^B just like V^B, except iterating only through the ordinals in M. Assume you have a boolean algebra homomorphism $B \to \{0, 1\}$; this tells you which boolean values are "true" and which are "false". Such homomorphisms correspond canonically to so-called *ultrafilters* of B; these are maximal filters. With an ultrafilter in hand, define equivalence classes in M^B (put a and b in the same class if $\|a = b\|$ is a "true" boolean value). Define an "element-of" relation on these classes ($[a]E[b]$ iff $\|a \in b\|$ is "true"). You'll get a well-founded model, *provided* the ultrafilter satisfies a certain *genericity* property. If so, apply Mostoski collapse. The existence of a generic ultrafilter U is assured when M is a countable model (the usual proof). Final result: a model $M[U]$ just like the model $M[G]$ in the Shoenfield construction.

Much of the allure of boolean-valued models, however, is that you can obtain relative consistency results without assuming the existence of a standard model (or any countability assumptions). For example, with the right boolean algebra B, all the axioms of ZF will have boolean value 1 but $V = L$ will not. The rules of inference "preserve the value 1", so to speak, so a proof in ZF will always lead to formulas with boolean value 1. Likewise for AC or CH or whatever.

Smullyan & Fitting

The reader who has patiently gone through parts I and II of S&F may be reluctant to abandon them in the final stretch without a backwards glance. I do recommend the last chapter (Ch.22), a well-written survey of the many guises of forcing. This situates their modal logic approach in a wider context (see the "syntactic sugar" remarks on p.302). I offer here a brief sketch of the main points of contact with (my preferred) Shoenfield approach.

Start with the definition of a *frame* $\langle \mathcal{G}, \mathcal{R} \rangle$ (Def.2.1, p.209). This corresponds to a notion of forcing $\langle P, \leq \rangle$: a *possible world* is just another name for a condition, and $p\mathcal{R}q$ means that q extends p.

S&F's version of forcing (Def.16.2.1, p.209) is strong forcing (which they denote $\Vdash$ instead of our $\Vdash^*$). However, they introduce weak forcing in §16.4 (p.215). We have these equivalences:

$$p \Vdash^* \Box\varphi \Leftrightarrow (\forall q \geq p) q \Vdash^* \varphi$$
$$p \Vdash^* \Diamond\varphi \Leftrightarrow (\exists q \geq p) q \Vdash^* \varphi$$
$$p \Vdash^* \Box\Diamond\varphi \Leftrightarrow \{q : q \Vdash^* \varphi\} \text{ is dense above } p$$

Recall that the last line is also equivalent to p weakly forcing φ, which is equivalent to φ holding in all generic models satisfying p. This clarifies their definition of the "classical embedding" in §16.4.

Ch.17 gets down to the brass tacks of modal models of ZF. In other words, it defines the domain of the models and defines $p \Vdash^* \varphi$. You will recognize Def.17.1.1, as giving the same domain as in the Shoenfield universe of names V^P (§21.2), with the trivial variation of using $\langle p, x \rangle$ instead of $\langle x, p \rangle$. You may recall the tricky double transfinite induction required to define forcing for atomics, in §21.3. The S&F approach shares DNA, without looking exactly the same. First, their $p \Vdash f \varepsilon g$ is the same as our $f \in_p g$ (Def.17.1.3, p.224). Next, they carry out the inductive definition of $p \Vdash f = g$ via a stratified notation, writing $f \approx_\alpha g$ for "equality at level α" and $f \approx g$ if $f \approx_\alpha g$ for some ordinal α. Establishing the properties of $\approx$, and using it to define $p \Vdash f \in g$ (with attendant properties) occupies most of §§17.1–17.2.

The rest of Ch.17 is given over to showing that we have a modal model of ZF.

Chs.18–20 present the usual independence results ($V = L$, AC, GCH) in essentially the usual way. One key difference from the treatment in Shoenfield's paper (and Cohen's book): generic sets play no role, only forcing. The S&F modal approach belongs to a long tradition (beginning with Cohen) of independence proofs based purely on a syntactic analysis of forcing. For another example of this style, see Shoenfield's book [22, §§9.8–9.9].

In Ch.21, S&F show how to obtain classical models from modal models. Here we finally encounter generic filters, the Truth Lemma, and all that apparatus. Finally, Ch.22 paints, in broad strokes, Cohen's original approach, boolean-valued models, Shoenfield's unramified forcing, and how these all relate to each other and to modal models.

21.12 Historical Remarks: Syntax vs. Models

Toward the end of *Set Theory and the Continuum Hypothesis*, in a section titled "Avoiding SM", Cohen sketches a way to avoid assuming that ZF has a standard model. The idea is to study the relation "$p \Vdash \varphi$" purely syntactically. (Gödel's monograph on AC and CH also stressed the syntactical viewpoint.)

Cohen then adds:

> Although this point of view may seem like a rather tedious way of avoiding models, it should be mentioned that in our original approach to forcing this syntactical point of view was the dominant point of view, and models were later introduced as they appeared to simplify the exposition.

Based on this quote, S&F write, "Cohen's first, unpublished, approach was purely syntactic." (p.300) Not so. In fact, Cohen only began to make real progress once he rejected the syntactic approach in favor of standard models. Once he got far enough along, syntax reentered the picture via forcing. I base these conclusions on two talks Cohen gave later in life, providing an extensive account of how he discovered forcing. Here are some quotes, starting with "The Discovery of Forcing" [5].

> ...in my own work, one of the most difficult parts of proving independence results was to overcome the psychological fear of thinking about the existence of various models of set theory as being natural objects in mathematics about which one could use natural mathematical intuition. (p.1072)
>
> In my own work on CH, I never was able to successfully analyze proofs as a combinatorial "game" played with symbols on paper. (p.1080)
>
> At first [trying to prove the independence of AC] I tried various devices..., but soon I found myself enmeshed in thinking about the structure of proofs. At this point I was not thinking about models, but rather syntactically... The question I faced was this: How to perform any kind of "induction" on the length of a proof.... It was at this point that I realized the connection with the models, specifically standard models. Instead of thinking about proofs, I would think about the formulas that defined sets, these formulas might involve other sets previously defined, etc. So if one thinks about sets, one sees that the induction is on the rank (pp.1088–1089)
>
> Nevertheless the feeling of elation was that I had eliminated many wrong possibilities by totally deserting the proof-theoretic approach. I was back in mathematics, not in philosophy. (p.1089)

Cohen frequently uses "philosophy" in these talks to mean a syntactical viewpoint (cf. the formalist credo, that mathematics is just a game with symbols).

The second talk [6] tells a similar story.

> When I began my own attempts at solving the CH problem in 1962, I was strongly motivated by the idea of constructing a model... Of course, I eventually realized that Gödel's "syntactical approach" [for L]... was not very different. After my work was completed, I realized that a prejudice against models was widespread among logicians and even led to serious doubts about the correctness of my work. Nevertheless, as I was trying to construct interesting new models of Set Theory, I had to start with a given model of it.
>
> ... A third theme of my work, namely the analysis of "truth", is made precise in the definition of "forcing". To people who read my work for the first time, this may very well strike them as the essential ingredient... Although I emphasized the "model" approach in my earlier remarks, this point of view is more akin to the "syntactical" approach, which I associate with Gödel. At least as I first presented it, the approach depends heavily on an analysis and ordering of statements and, by implication, on an analysis of proof.
>
> Initially, I was not very comfortable with this approach; it seemed to be too close to philosophical discussions, which I felt would not ultimately be fruitful... Of course, in the final form, it is very difficult to separate what is [model] theoretic and what is syntactical. As I struggled to make these ideas precise, I vacillated between two approaches: the model theoretic, which I regarded as roughly more mathematical, and the syntactical-forcing, which I thought of as more philosophical.

Articles by Moore [19] and Kanamori [15] present a broader context. Briefly, syntax predominated in the first half of the 20th century, despite notable results in model theory (completeness and the downward Löwenheim-Skolem theorems). Gödel favored a very precise, even pedantic style; this made his monograph on AC and GCH difficult to read, and obscured its model theoretic underpinnings. You can see the philosophical impetus: pure syntax is the "safest" ground. Even a committed finitist will assent to Gödel's purely syntactic proof that Con(ZF) implies Con(ZF+AC+GCH).

This style was antithetical to Cohen's mode of thought. Kanamori gives an example relating to the minimal model of ZF. Shepherdson first obtained this in 1953; Cohen rediscovered it a decade later. Kanamori contrasts the two styles:

> The Shepherdson and Cohen papers are a study in contrasts that speaks to the historical distance and the coming breakthrough. While both papers are devoted to establishing essentially the same result, the former takes 20 pages and latter only 4. Shepherdson labors in the Gödelian formalism with its careful laying out of axioms and propositions in first-order logic, while Cohen proceeds informally and draws on mathematical experience. Shepherdson works out the relativization of formulas, worries about absoluteness and comes down to the minimal model, while Cohen takes an algebraic closure.

Once Cohen had the independence results in hand, he revisited the land of pure syntax, partly to placate philosophical qualms. Kanamori:

> Cohen did appreciate that starting with a standard model of ZF is formally more substantial than assuming merely the consistency of ZF, and he indicated, as a separate matter, a syntactic way to pare down his arguments into formal relative consistency statements of the type $\mathrm{Con(ZF)} \to \mathrm{Con(ZF} + \neg\mathrm{AC)}$.

Cohen's 1966 book [4] recast Gödel's monograph in a much more accessible form. Kanamori:

> The middle third of the monograph is an incisive account of axiomatic set theory through Gödel's work on L to the minimal model, an account shedding Gödel's formal style and pivoting to a modern presentation.

More generally, Cohen's *Set Theory and the Continuum Hypothesis* remains one of the best places to encounter logic and axiomatic set theory for the first time. As Martin Davis writes in the introduction to the new Dover edition,

> In writing this book, Cohen set himself the goal of developing his proof of the independence of CH without presupposing any knowledge of mathematical logic. He has accomplished this in less than 200 pages. En route he has covered the main methods and results of the subject. He did not aim for a polished work in which every last detail had been meticulously worked out. Rather he has presented the broad sweep of the subject underlining the intuitive ideas that more polished expositions sometimes hide.

Ironically, because of the simplifications discovered after the book's publication, the one exception is the chapter on forcing.

22 Review of Smullyan and Fitting

The first sentence of Pollard's review [20] sums up my feelings perfectly: "This rewarding, exasperating book. . . " On balance, I found it more exasperating than rewarding. But it does have its charms.

As I mentioned at the start of these notes, I participated in a meetup group that went through the first two parts of S&F. My fellow participants possessed considerable mathematical knowledge and sophistication, but had only slight prior acquaintance with mathematical logic and none with axiomatic set theory. The opinions here are my strictly my own, but they are, I think, informed by my experience in the meetup. If I had just skimmed the book, glancing at familiar material, I would probably have a more positive impression.

I will begin with the book's minuses, so as to end on a positive note.

Sloppiness

A major drawback for self-study. Some of the errors are clearly just typos, but many are not so easily fixed. We frequently found mistakes not listed in any of the online lists of errata (themselves not short for a book of this size), or mentioned in Pollard's review.

The sloppiness goes beyond specific errors. As Pollard observes, the authors don't want to take '=' as a primitive relation, but they never define it; each of the two standard definitions is inconsistent with the text at one point or another. Another example: their treatment of classes. Some authors opt for ZF, others for NBG, but S&F occupy an uneasy middle ground. It often seems that they couldn't make up their minds. I noted an especially egregious example in §3 (class quantifiers). Even minor matters (like the haphazard numbering of the axioms—see Pollard and §3) betray a pervasive inattention to details.

Logic Background

S&F fall prey to the seductive notion that you can understand one specific first-order theory, ZF (or NBG), without first knowing what the phrase "a first-order theory" means. They say in the preface:

> Our book is intended as a text for advanced undergraduates

> and graduates in mathematics and philosophy. It is close to self-contained, involving only a basic familiarity with the logical connectives and quantifiers. A semester course in symbolic logic is more than enough background.

At best, this may be technically true. In practice, to appreciate the issues you need familiarity with quite a few notions from first-order logic and model theory.[29]

S&F put off a serious discussion of first-order syntax and semantics until Ch.11 (p.141–144), in part II. Part I takes place in the land of naive set theory, with confusing forays into axiomatic territory.

This fuzziness regarding syntax and semantics spills over onto content; it's not just style. "Axioms" A_1 and A_2 (p.17) prove especially troublesome, as I detailed in §3. I gave my best guess at their intended interpretation, relying on uses occurring 85 pages later, and a concept that first appears (implicitly) over 130 pages later, in Ch.12. But a reader shouldn't have to engage in such detective work. All this fogginess wafts up from the authors' refusal to be explicit (until much later) about the syntax and semantics of first-order logic.

Unhelpful originality

Sean Carroll, in his preface to *Spacetime and Geometry: An Introduction to General Relativity*, writes:

> An intentional effort has been made to prefer the conventional over the idiosyncratic. . . Since the goal of the book is pedagogy rather than originality, I have often leaned heavily on other books (listed in the bibliography) when their expositions seemed perfectly sensible to me.

[29]This became quite clear during the meetup. I expanded a short talk into the companion to these notes, "Basics of First-order Logic", to provide the needed background.

Would that S&F had followed this philosophy! They sometimes seem allergic to convention.

Their choice of terms manifests this in a minor way. The downward Löwenheim-Skolem Theorem becomes the Tarski-Vaught Theorem. Mostowski collapse becomes the Mostowski-Shepherdson mapping. The Gödel condensation lemma becomes Gödel's isomorphism theorem. They may have historical justification for some of these attributions, but surely a parenthetical "(usually called the downward Löwenheim-Skolem Theorem)" wouldn't have hurt. (Cummings' review [7] also complains about this.)

In the preface, S&F call out some unusual features of their treatment of ordinals and the well-ordering theorem (Chs.4 and 5). Pollard's sardonic comment deserves to be quoted:

> The authors spare reviewers the task of crafting laudatory prose to describe this chapter. It offers, they say, "a particularly smooth and intuitive development of the ordinals". Indeed, we are guaranteed "a beautifully natural and elegant treatment". The authors cannot help but remark, "It is high time this neat approach should be known!" Any praise this reviewer might offer would be superfluous.

I found their treatment far *less* clear and intuitive than the typical one. Expositional quirks bear most of the blame (see below). The S&F approach is "top-down", the usual one "bottom-up": see my remarks in §5 about Zermelo's two proofs of the well-ordering theorem. Now, 90% of the development looks pretty much the same with either approach: throw transfinite induction (or superinduction) at everything in sight. For the remaining 10%, I see no intrinsic advantage to top-down over bottom-up.[30]

[30] S&F's implemention of the top-down approach demands Cowen's theorem. As Pollard points out (a) its use is easily avoidable (p.478), and (b) a key definition is "a disaster" (p.476), thanks to a (fixable) careless error (not just a typo). The proof of Cowen's theorem is neither short nor sweet.

The preface notes another special feature of their approach:

> Our main novelty lies in the use of Smullyan's double induction and double superinduction principles... It is high time that these should appear in a textbook.

I've found *exactly one* application of double superinduction in the book, and none elsewhere. This is in the proof of comparability for ordinals. Other treatments prove comparability with ordinary transfinite induction—it's a just a touch trickier than many such arguments. You could extract the proof of double superinduction from the comparability argument. Is the game worth the candle? S&F seem to think so.[31]

S&F's biggest claim to originality is Part III, the treatment of forcing via S4 modal logic. As Pollard [20] diplomatically says, "This approach will be welcomed by scholars who have struggled with forcing, but are comfortable with modal logic." Not a large cohort, I'll bet! Pollard continues, "Many students, however, may lack both the motivation and the background to make sense of it." I offer a brief sketch of their approach in §21.11. My own two cents: unless you have a special affinity for modal logic, don't bother. (Pollard has a more positive assessment. Cummings [7] on the other hand writes, "I do not feel that the pedagogical advantages of using modal logic to present forcing outweigh the drawbacks. In the modal approach the student is called on to master a new formal language and a new semantics before even embarking on the study of forcing; this seems likely to make a hard subject even more confusing.")

Slice and dice, and the pointless generalization

You can take almost any step in a train of mathematical reasoning and

[31] In [24] Smullyan says, "It was as a result of searching for a simplified proof of Cowen's Theorem that we uncovered Theorem B"; Theorem B is double superinduction as presented in S&F. While I haven't studied this paper in detail, it seems to cover the same ground as in the book.

extract some kind of generalization from it. Just look at the precise assumptions used in that step, invent new terms as needed, more new terms for the consequences, and voila! A new proposition! Nine times out of ten it isn't worth the trouble.

S&F couple this sort of generalization with "slice and dice": scrambling the natural order of the links in a chain of argument. What should be $A \Rightarrow B \Rightarrow C \Rightarrow D \Rightarrow E$ becomes $C' \Rightarrow D_0$; $A^* \Rightarrow B''$; $D' \Rightarrow \overline{E}$; $B_1 \Rightarrow C^*$. (This burdens the reader with recognizing that B_1 is a special case of B'', etc.)

Their proof of Zermelo's well-ordering theorem exhibits these quirks at full strength. In §5 I've given what I feel is a more natural ordering of the steps. Ch.13 on the constructible universe also lacks a clear narrative thread. Their addiction to generalization shows up in an amusing way. In §12.3, they introduce Gödel's notion of a *constructible class*. In §13.2, they generalize this to a *distinguished subclass*. Two pages later, in §13.3, they take pains to point out that the proof of Lemma 3.2 still works if *distinguished* is replaced with the yet more general *special*. No other application of these more general notions is ever made.

The slice and dice technique has an unfortunate corollary: a population explosion in terminology. I wrote §2 of these notes to help keep track of terms like *progressing*, *strictly progressing*, *slowly progressing*, *g-tower*, *slow g-tower*, *slow well-ordering*, etc.—all slight variations on the same cluster of ideas. If at each link we look for the weakest possible hypotheses to justify that implication, the thrust of the argument is lost in a profusion of terms and ever-shifting assumptions.

I do not find fault with all their generalizations. The definition of *ordinal hierarchy* is amply justified by its three uses: the R_α's in Ch.7, the L_α's in Ch.12, and the $R^{\mathcal{G}}_\alpha$'s in Ch.17. The *relational system* concept of Ch.10, although single-use in this book, is pleasingly simple and may furnish some insight into Mostowski collapsing maps. (S&F may have missed a bet. Shoenfield used Mostowski collapsing maps (or nearly so) to construct generic extensions; very likely S&F could have used relational systems in Ch.21 for the same purpose.)

Now for the plusses.

Detailed arguments

Perhaps you've never had this experience: you get stuck for a day or two in the middle of reading a proof, only to only to realize that a casual remark two pages back supplies the missing link. Or that a slightly different argument handles the case where $A = \varnothing$.

S&F provide detailed, step-by-step proofs; except for typos and such, "stuck in the middle" syndrome seems unlikely.

Notation/prose balance

One is naturally tempted, in a book on axiomatic set theory, to go heavy on notation. The reader must deal with formal logic; why not take advantage of its precision and concision?

Except for part III, S&F strike a near-perfect balance; if anything they sometimes push the needle a touch too far toward the "prose" end of the dial. (I wish they had made use of the satisfaction symbol.) As a small example, consider how Drake [8, p.137] states Gödel's result for GCH in L:

$$\mathrm{ZFC} \vdash \mathrm{Init}(\kappa) \to \forall x(x \in H(\kappa) \wedge x \in L \to x \in L_\kappa)$$

The S&F version (p.194):

> For any infinite cardinal c, every constructible subset of L_c is an element of L_{c^*}.

The notation-heavy style demands a steady level of "decoding" from the reader, which gets tiring over time. (And Drake is a very good textbook.)

Technical felicities

Some parts of S&F are very nicely done; their treatment of absoluteness and related concepts in §12.2 and §§14.1–14.2 is particularly clear and meticulous. They employ a trick in Ch.14 (favoring Σ over Δ_1^{ZF}; see §17.2) that I found technically sweet. The last part of Ch.14, proving that L has a definable well-ordering, is a gem of exposition.

Conversational style

In the bad old days, Edmund Landau's *Foundations of Analysis* was regarded as the epitome of mathematical style: Definition-Theorem-Proof, with none of that frilly motivation and hand-holding. The famous Bourbaki series solidified this sort of exposition.

The past few decades have seen a welcome trend away from those chilly tomes, but not all authors can wave their hands with aplomb. Fortunately our authors can. The very best sections of S&F are the chatty ones.

I've quoted Pollard's review [20] several times already. His view of S&F's first chapter seems a little jaundiced:

> The opening chapter provides a happy-go-lucky introduction to size comparisons between infinite sets. The use of the first person singular helps to warm an atmosphere already sweetened with cozy good humor (perplexing though the 'I' might be in a work with two authors). If this chapter is meant to entice and entertain readers rather than to instruct them, then it succeeds admirably.

In fact, the opening chapter covers pretty standard material: countability of $\mathbb{N} \times \mathbb{N}$, of $\mathbb{Q}$, of the set of all finite subsets of $\mathbb{N}$, and the uncountability of $\mathcal{P}(\mathbb{N})$. S&F present this by way of an engaging metaphor; this would not be out of place in a pop math book. The chapter also covers Russell's paradox, and economically sketches the transition from Frege's (inconsistent) set theory to Zermelo's axioms. Perhaps Pollard was put off by the elementary nature of this chapter; most texts on the continuum hypothesis get off to a faster start.

Nearly every chapter sets the stage with a paragraph or two of well-chosen words. Chapter 5 introduces ordinals with suitable fanfare. The opening section of part III categorizes the approaches to forcing as "non-classical logic" vs. "non-inner model" with some hand-waving of the first order. When S&F do turn to non-inner models, they don't just plunge into the machinery of dense and generic sets over a partial order, but motivate it via Cohen's concrete construction of a complete sequence. Finally, the closing chapter presents a brief sketch of the history of forcing, beautifully done. (Though I do disagree with one of their historical assertions: see §21.12.)

Beyond these specific passages, the *desire* to be reader-friendly pervades the book. Better to have that goal and sometimes fall short, than not to have it at all.

References

[1] Martin Aigner and Günter M. Ziegler. *Proofs from THE BOOK*. Springer-Verlag, 6th edition, 2018.

[2] John Baez. Large countable ordinals (part 1). https://johncarlosbaez.wordpress.com/2016/06/29/large-countable-ordinals-part-1/.

[3] John L. Bell. *Set Theory: Boolean-valued Models and Independence Proofs*, volume 47 of *Oxford Logic Guides*. Oxford University Press, 3rd edition, 2005.

[4] Paul J. Cohen. *Set Theory and the Continuum Hypothesis*. W. A. Benjamin, 1966.

[5] Paul J. Cohen. The Discovery of Forcing. *Rocky Mountain Journal of Mathematics*, 32(4):1071–1100, 2002.

[6] Paul J. Cohen. My Interaction with Kurt Gödel: The man and his work. In *Kurt Gödel and the Foundations of Mathematics: Horizons of Truth*. Cambridge University Press, 2009. Included in the Dover 2012 edition of *Set Theory and the Continuum Hypothesis*.

[7] James Cummings. Review of Set Theory and the Continuum Problem. *Journal of Symbolic Logic*, 64(1):401–403, 1999.

[8] Frank R. Drake. *Set Theory: An Introduction to Large Cardinals*. North-Holland, 1974.

[9] Paul Erdős. An interpolation problem associated with the continuum hypothesis. *Michigan Mathematics Journal*, 11(3):9–10, 1964. https://projecteuclid.org/euclid.mmj/1028999028.

[10] Lorenz J. Halbeisen. *Combinatorial Set Theory*. Springer, 2nd edition, 2017.

[11] M. Randall Holmes. Alternative Axiomatic Set Theories. *Stanford Encyclopedia of Philosophy*, 2017. https://plato.stanford.edu/entries/settheory-alternative/#AckeSetTheo.

[12] Thomas Jech. *The Axiom of Choice*. Dover, 1973.

[13] Thomas Jech. *Set Theory: The Third Millennium Edition*. Springer, 2002.

[14] Akihiro Kanamori. Levy and set theory. *Annals of Pure and Applied Logic*, 140:233–252, 2006.

[15] Akihiro Kanamori. Cohen and set theory. *The Bulletin of Symbolic Logic*, 14(3):351–378, 2008.

[16] Irving Kaplansky. *Set Theory and Metric Spaces*. Allyn and Bacon, 1972.

[17] Kenneth Kunen. *Set Theory: An Introduction to Independence Proofs*. Elsevier, 1980.

[18] David Madore. An applet to draw ordinals. http://www.madore.org/˜david/math/drawordinals.html.

[19] Gregory Moore. *Zermelo's Axiom of Choice: Its Origins, Development, and Influence*. Springer, 1982.

[20] Stephen Pollard. Review of Set Theory and the Continuum Problem. *Notre Dame Journal of Formal Logic*, 38(3):475–480, 1997. http://projecteuclid.org/euclid.ndjfl/1039700750.

[21] J. Barkley Rosser. *Simplified Independence Proofs*. Academic Press, 1969.

[22] J. R. Shoenfield. *Mathematical Logic*. Addison-Wesley, 1967.

[23] J. R. Shoenfield. Unramified Forcing. In Dana Scott, editor, *Axiomatic Set Theory*, volume 13 part 1, pages 357–381. American Mathematical Society, 1971.

[24] Raymond M. Smullyan. Some New Double Induction and Superinduction Principles. *Studia Logica*, 49(1):23–30, 1990.

[25] Raymond M. Smullyan and Melvin Fitting. *Set Theory and the Continuum Problem*. Dover, 2010.

[26] Gaisi Takeuti and Wilson M. Zaring. *Axiomatic Set Theory*. Springer, 1973.