The Hyperuniverse

1.Vorlesung

The Hyperuniverse $\mathbb{H}$ is the collection of all countable transitive models of ZFC

The Hyperuniverse is interesting for 3 reasons:

• Much of set theory is about building transitive models of ZFC.
• By Löwenheim-Skolem, the first-order properties of these models all appear in models of the Hyperuniverse.
• The Hyperuniverse is closed under all techniques for building new countable transitive models from old ones and therefore provides the broadest range of possibilities for natural interpretations of set theory.

The Hyperuniverse is also a tool for understanding set-theoretic truth through the *The Hyperuniverse Programme*. The general idea of this programme is the following:

• Elements of the Hyperuniverse provide *possible pictures of* V which mirror all possible first-order properties of V.
• We can formulate natural criteria for *preferred* elements of the Hyperuniverse based on their status within the Hyperuniverse as a whole.
• Under the assumption that first-order properties of the real universe are mirrored by *preferred* elements of the Hyperuniverse, we can regard the first-order properties shared by these preferred universes as being "true" in V.

In this course we will however deal only with the the mathematical aspects of the Hyperuniverse. This work raises numerous issues in forcing, definability, large cardinals, determinacy and infinitary logic.

But first we have to clear up one point: Consistently with ZFC, the Hyperuniverse is empty! For a rich theory we therefore impose the assumption that every real belongs to a transitive model of ZFC. This assumption is rather modest from a mathematical point of view (it is much weaker than an inaccessible in consistency strength) yet it suffices to yield a robust structure.

I'll sometimes use the word "universe" to mean element of $\mathbb{H}$, i.e. a countable transitive model of ZFC.

Let's start with some simple observations about extensions of universes. First we need a pair of definitions.

Definition 1 *Suppose that M is a universe. A* width-extension *of M is a universe containing M with the same ordinals as M. A* height extension *of M is a universe N containing M such the same V_α's as M for α an ordinal of M.*

Thus if N is a height extension of M, either M equals N or M equals V^N_β where $\beta = \mathrm{Ord}(M)$.

Proposition 2 *(a) There is an element of $\mathbb{H}$ which is smallest under inclusion.*
(b) Every universe has continuum-many width-extensions.
(c) There are universes with no proper height-extensions.
(d) If there is a universe with an inaccessible cardinal then there is one with a proper height-extension.

Proof. (a) Let M be any universe. By Gödel the L of M is also a universe and equals L_β where $\beta = \mathrm{Ord}(M)$. Thus L_α where α is least so that L_α is a universe is contained in all universes.
(b) Let M be a universe and consider $\mathbb{P}$ = Cohen forcing. As M is countable there are reals which are $\mathbb{P}$-generio over M, and in fact we can build a perfect set of such $\mathbb{P}$-generics: List the dense subsets of $\mathbb{P}$ which belong to M as $(D_n \mid n < \omega)$ and define Cohen conditions $(p_s \mid s \in 2^{<\omega})$ so that p_{s*0}, p_{s*1} are incompatible extensions of p_s hitting D_n where n is the length of s. Any infinite branch b through $2^{<\omega}$ yields a Cohen-generic by taking the union of the corresponding conditions p_s for s a finite initial segment of b. If b_0, b_1 are distinct then we get distinct Cohen generics. It follows that we get continuum-many distinct Cohen width-extensions $M[b]$ as there are continuum-many b's and each width-extension contains only countably many b's. [Remark: With more care one can arrange that distinct b's are *mutually-generic* and therefore any two $M[b]$'s are distinct.]
(c) Suppose that the universe M has a proper height-extension N. Then in

N we can take a countable elementary submodel of M and form its transitive collapse. As M is uncountable in N (indeed it has size a strong limit cardinal of N) this transitive collapse witnesses that M is not the smallest universe. So the smallest universe has no proper height-extension.
(d) If M has an inaccessible κ then the universe V_κ^M has a proper height-extension, namely M. $\square$

Remark. Our background assumption that every real belongs to a universe is not sufficient to obtain the conclusion of (d) above. Indeed, this assumption holds if there is an uncountable transitive model of ZFC containing all of the reals but to obtain a proper height-extension one needs a universe with a V_α satisfying ZFC, a stronger property.

2.,3.Vorlesungen

Other notions of width-extension.

Let M, N be universes of the same ordinal height.

M is an *inner model* of N iff N is a width-extension of M
M is a *strong inner model* of N iff in addition N satisfies replacement for formulas with M as an additional unary predicate
M is a *definable inner model* of N iff in addition M is N-definable

Clearly "inner model" and "definable inner model" are transitive notions.

Proposition 3 *The notion of "strong inner model" is not transitive, and therefore the three notions of inner model are distinct.*

Proof. Start with $V_0 \vDash V = L$. Let C_0, C_1 be generic over V_0 for ∞-Cohen, the forcing that adds a Cohen class of ordinals. Also arrange that C_0, C_1 agree except on a cofinal subset of $\mathrm{Ord}(V_0)$ of ordertype ω. Force over (V_0, C_0) to add an $\aleph_{2\times\alpha+1}$-Cohen generic for α in C_0, using an Easton product. This is the model V_1. Then force over (V_1, C_0) to add an $\aleph_{2\times\alpha+1}$-Cohen generic for all α not in C_0 *using an Easton product defined in* V_0; this is the model V_2. (Note that V_2 is generic over V_0 for the Easton product of the $\aleph_{2\times\alpha+1}$-Cohen forcings for all α.) Finally, force over (V_2, C_1) to add an $\aleph_{2\times\alpha+2}$-Cohen generic for α in C_1; this is the model V_3.

Then (V_2, V_1) and (V_3, V_2) are models of ZFC but both C_0 and C_1 are definable in (V_3, V_1) so the latter is not a model of ZFC. $\square$

Universes M, N are *compatible* if they have a common width-extension.

Proposition 4 *There are incompatible universes of the same ordinal height.*

Proof. Let C be a real coding α. Build reals A, B which are Cohen generic over L_α and have the following property: Let $(k_n \mid n \in \omega)$ enumerate the places where A, B differ in increasing order; then $A(k_n) = 0$ iff n belongs to C. Then $L_\alpha[A], L_\alpha[B]$ are incompatible universes, as $C \leq_T (A, B)$. $\square$

A universe M of height α is a *node for comparability* iff every universe of height α is comparable with M, i.e., either contains M or is contained in M. M is a *node for compatibility* iff every universe of height α is compatible with M.

Obviously L_α is a node for comparability.

Proposition 5 *Suppose that M is a universe of height α which is a node for comparability. Then M equals L_α.*

Proof. There is an uncountable set of reals X such that any two distinct elements of X are mutually Cohen over L_α. If M is contained in $L_\alpha[R]$ for two distinct R in X then $M = L_\alpha$. Otherwise M must contain all but one element of X, contradicting its countability. $\square$

Open Question: Is L_α the only node for compatibility of height α? I.e., if M is a universe of height α which is compatible with all universes of height α, must M equal L_α?

Proposition 6 *Suppose that M has height α and contains an infinite subset A of ω which is sparse over L_α, i.e., such that for any $f : \omega \to \omega$ in L_α the interval $(n, f(n))$ is disjoint from A for infinitely many n in A. Then M is not a node for compatibility.*

Proof. Using A we can build a Cohen real R so that R codes any real (such as a code for α) on A. Then $L_\alpha[R]$ and M are incompatible universes.

Here are the details. Suppose that A is sparse over L_α. Let R be a real that codes α; we will code R into the pair (A, C) where C is Cohen over L_α. Suppose that D is dense for Cohen forcing and belongs to L_α. Consider the function f in L_α that given n chooses the least $f(n)$ so that any Cohen condition of length $n+1$ has an extension in D of length at most $f(n)$. As A is sparse we can choose n in A so that $f(n)$ is less than the least element of A greater than n. Define the Cohen condition p to be 0 up to and including n, extend it to a Cohen condition q in D of length at most $f(n)$ and then extend q to p_0 with 0's up to length n^*, where for some k, n^* is the k-th element of A and k codes a finite initial segment of the real R; finally p has length n^*+1 and assigns the value 1 at n^*. Then repeat this for all dense D in L_α, ensuring that if the resulting Cohen generic C assigns 1 on the k-th element of A then k codes a finite initial segment of R (and this happens for infinitely many k). Then using C and A we can recover infinitely many initial segments of R and therefore all of R. $\Box$

Corollary 7 *If M has height α and contains a function $f : \omega \to \omega$ that is unbounded over L_α (i.e. not dominated by a function in L_α), then M is not a node for compatibility.*

Proof (Lyubomyr). It suffices to show that the hypothesis implies that M contains a set which is sparse over L_α (and therefore the existence of an unbounded function is equivalent to the existence of a sparse set).

Aassume that f is strictly increasing and let A be the range of f. We claim that A is sparse over L_α. Let $g : \omega \to \omega$ in L_α be strictly increasing and such that $g(0) > 0$; we need to show that the set $C = \{n \mid g(f(n)) < f(n+1)\}$ is infinite. Set $h(0) = g(0)$ and $h(k+1) = g(h(k))$ for all k. (So $h(n) = g^{n+1}(0)$.) We show that the set $B = \{k \mid [h(k), h(k+1)) \cap A = \emptyset\}$ is infinite. Otherwise there exists $k_0 \in \omega$ such that $A \cap [h(k), h(k+1)) \neq \emptyset$ for all $k \geq k_0$, which implies $f(n) \leq h(n+k_0+1)$ for all n, contradicting the unboundedness of f over L_α.

Now pick any $k \in B$ and find $n_k \in \omega$ such that $f(n_k) < h(k) < h(k+1) \leq f(n_k+1)$. Then $g(f(n_k)) < g(h(k)) = h(k+1) \leq f(n_k+1)$, and hence $n_k \in C$. Since the map $k \mapsto n_k$ is injective, C is infinite. $\Box$

Jensen coding and minimality

Universes have width-extensions of a special form.

Theorem 8 *(Jensen) Suppose that M is a universe of height α. Then M has a width-extension of the form $L_\alpha[R]$ for some real R. Moreover, if M satisfies GCH then $H(\gamma)^M$ is definable over $L_\gamma[R]$ for each cardinal γ of M.*

A universe M is *minimal over a real* iff for some real R, M is the least universe (of any ordinal height) containing R.

Theorem 9 *Every universe has a width-extension which is minimal over a real.*

Proof. In light of Jensen's theorem we may assume that M is of the form $L_\alpha[R]$. Now force a club C of cardinals γ such that $L_\gamma[R]$ does not satisfy ZFC. Then collapse cardinals to ensure that all limit cardinals belong to C and apply Jensen's theorem again. The result is a model of the form $L_\alpha[R']$ in which ZFC fails in $L_\gamma[R']$ for all cardinals γ. $\square$

Now use:

Theorem 10 *(R.David-SDF) Suppose that $N = L_\alpha[R]$ is a model of ZFC, φ is a Σ_1 formula with parameter R and $N \vDash \varphi(\gamma)$ for every cardinal γ of N. Then for some real S, $L_\alpha[S]$ is a width-extension of N satisfying $\varphi(\delta)$ for every δ such that $L_\delta[S]$ models ZF$^-$.*

Apply this to the model $L_\alpha[R']$ and the formula $\varphi(\gamma) \equiv (L_\gamma[R] \nvDash \text{ZFC})$. This gives a real S such that ZFC fails in $L_\delta[R]$ for all δ and therefore $L_\alpha[S]$ is the least universe containing the real S. $\square$

4.,5.Vorlesungen

Truth in Universes

Until now we have compared universes under the relation of inclusion. We now extend this comparison to take into account what first-order (and certain second-order) properties hold in them.

A universe M of height α is *α-characterisable* iff for some sentence φ, M is the unique universe of height α satisfying φ.

Theorem 11 *Suppose that M is α-characterisable. Then: (a) M is an element of L_β where β is the least admissible greater than α. (b) If in addition α is a cardinal in L_β, then M must equal L_α.*

Proof. (a) Let φ witness that M is α-characterisable. Let $\mathcal{L}_\beta$ denote the admissible fragment of $L_{\omega_1\omega}$ determined by the admissible set L_β. Let ψ be the sentence in this fragment given by:

ZFC + φ
$\forall x(x$ is an ordinal iff $\bigvee_{\gamma<\alpha} x = \bar{\gamma})$ (where $\bar{\gamma}$ is a constant symbol denoting γ)

Then ψ is consistent and complete, and therefore has a model which is an element of L_β; this is the unique model of φ of height α.
(b) If α is a cardinal in L_β then all bounded subsets of α in L_β belong to L_α and therefore M is contained in L_α. □

A universe M is *characterisable* iff for some sentence φ, M is the unique universe satisfying φ (of any height). I mention the following without proof (as its proof is quite lengthy and technical).

Theorem 12 *There is a characterisable M of height the minimal model of ZFC which does not satisfy $V = L$.*

Absoluteness

In light of Theorem 11 we expect that a universe can be enlarged while preserving some of its first-order properties. Suppose that N is a width-extension of the universe M. A sentence φ with parameters in M is *absolute* beween M and N iff its truth value in M is the same as its truth value in N. An important case of such absoluteness is:

Theorem 13 *(Lévy-Shoenfield Absoluteness) Suppose that φ is Σ_1 with real parameters from M and N is a width-extension of M. Then φ is absolute between M and N.*

Proof. It is easy to see that Σ_1 formulas are upwards-persistent, i.e. if M satisfies φ then so must N. Conversely, if N satisfies φ then by Löwenheim-Skolem we can assume that the witness x to $\varphi = \exists x\psi(x)$ in N is hereditarily countable and is in fact a real. But then we can form a tree T on $\omega \times \omega_1^N$ such that T has an infinite branch and if (x, b) is an infinite branch through T then x codes a countable transitive model of $\psi(S)$ for some real S and b confirms that this model is well-founded. But T is definable in M and therefore has an infinite branch in M by the absoluteness of well-foundedness for binary relations. The result is a real R in M witnessing $\psi(R)$. □

Can we extend Lévy-Shoenfield absoluteness to include uncountable parameters? Clearly not, because if for example the parameter is ω_1^M then in some width-extension N we have that ω_1^M is countable and this is expressible as a Σ_1 sentence with parameter ω_1^M.

The obvious restriction to avoid this problem is to add the requirement that N preserves ω_1, i.e. $\omega_1^N = \omega_1^M$. However even with this restriction we cannot extend Lévy-Shoenfield in this way:

Theorem 14 *There is a Σ_1 formula with parameter ω_1^M and real parameters from M which is true in a width-extension of M but false in M.*

Proof. Suppose not. First consider the sentence: "There is a real R such that for $\alpha < \omega_1$, $L_\alpha[R]$ is not a model of ZFC". This is Σ_1 in the parameter ω_1 and holds in any width-extension of M which is minimal over some real. (By an earlier result, there exist such width-extensions.) So it holds in M. Choose a real R in M such that $L_\alpha[R]$ is not a model of ZFC for M-countable α. (Remark: By Löwenheim-Skolem there is in fact no ordinal α of M such that $L_\alpha[R]$ is a model of ZFC, but we will not need this here.) In particular ω_1^M is not inaccessible in M and therefore for some real S in M, every ordinal which is countable in M is also countable in $L[S]^M$. For α a countable ordinal of M let f_α be the $L[S]$-least surjection of ω onto α and for each n choose an ordinal α_n such that $f_\alpha(n) = \alpha_n$ for all α in a subset X_n of ω_1^M which is stationary in M.

Now for each n, M has a width-extension in which ω_1 is preserved and X_n contains a club. And this is expressible by a Σ_1 formula with parameters S and ω_1^M. Therefore in M each X_n contains a club and therefore there is a single club C contained in all of the X_n's. But then the surjection f_α is the same for all $\alpha \in C$, contradiction. □

The previous negative result still leaves the possibility of extending Lévy-Shoenfield absoluteness to Σ_1 formulas with parameter ω_1 (and no real parameters) for width-extensions which preserve ω_1.

M satisfies Lévy(ω_1) iff whenever a Σ_1 formula with parameter ω_1^M has a solution in an ω_1-preserving width-extension of M then it has a solution in M.

Theorem 15 *Assuming large cardinals, there exists an M satisfying Lévy(ω_1).*

Proof. We use both PD and Jensen coding. For any real R let $M(R)$ denote the minimal model of ZFC containing R. Using PD choose a real R such that if R is recursive in S then $M(R)$ and $M(S)$ have the same first-order theory.

We claim that $M(R)$ satisfies Lévy(ω_1). Indeed, suppose that φ is a Σ_1 formula with parameter ω_1^M and N is a width-extension of $M(R)$ which preserves ω_1 and in which φ is true. Let α be the ordinal height of $M(R) =$ the ordinal height of N. Apply Jensen coding to produce a real S such that N is contained in $L_\alpha[S]$ and N is a Σ_n-definable class in $L_\alpha[S]$ for some n. By further coding we can ensure that $L_\alpha[S]$ is the minimal model containing S. It is a fact about these codings that ω_1 is preserved when enlarging N to $L_\alpha[S]$. We may also assume that R is recursive in S.

Thus in $M(S)$ the following is true: "There is a Σ_n-definable inner model with the correct ω_1 in which φ is true". By the choice of R, this sentence is also true in $M(R)$. But then φ is true in $M(R)$ as it is true in an inner model of $M(R)$ with the correct ω_1. □

6.,7.Vorlesungen

More absoluteness

Consider the following natural strengthening of Lévy(ω_1).

M satisfies Lévy(ω_1, ω_2) iff whenever a Σ_1 formula with parameters ω_1^M, ω_2^M has a solution in an ω_1-preserving and ω_2-preserving width-extension of M then it has a solution in M.

Open question. Assuming large cardinals is there an M satisfying Lévy(ω_1, ω_2)?

The variants of Lévy absoluteness we've been considering are special cases of a much broader question.

Absoluteness question. Is there a universe M such that whenever N is a width-extension of M of a certain type and φ is a Σ_1 formula with parameters from M of a certain type, we have that M and N agree on the truth-value of φ?

Note the restriction to width-extensions "of a certain type". In the case of abasoluteness with parameter ω_1 we have seen that it is sufficient to restrict to

ω_1-preserving extensions, and clearly such a rsetriction is necessary. Another natural rsetriction is to cardinal-preserving width-extensions. What other restrictions on width-extensions arise in the study of absoluteness?

Definition 16 *Suppose that N is a width-extension of M. Then M* globally *covers N iff for some M-regular κ, if $f : \alpha \to M$ belongs to N then there is $g : \alpha \to M$ in M such that $f(i) \in g(i)$ and $g(i)$ has M-cardinality $< \kappa$ for all $i < \alpha$. In this case we also say that M* globally κ-covers *N.*

Remark. If M globally κ-covers N then it is easy to see that any regular cardinal of M of size at least κ is also regular in N. Also, in the definition of global covering, it suffices to consider functions from ordinals to ordinals, using the fact that every set can be wellordered.

Theorem 17 *(a) There is a universe M which satisfies Σ_1 absoluteness with subsets of ω_1^M in M as parameters (and therefore with reals in M and ω_1^M as parameters) for width-extensions which it globally ω_1-covers.*
(b) Assuming large cardinals, there is a universe M which satisfies Σ_1 absoluteness with subsets of ω_1^M in M as parameters (and therefore with reals in M and ω_1^M as parameters) for width-extensions which it globally covers and which preserve the stationarity of subsets of ω_1^M.

How does global covering facilitate absoluteness? The answer is revealed by the following.

Theorem 18 *(Bukovsky) (a) M globally covers N iff N is a set-generic extension of M.*
(b) M globally κ-covers N iff N is a κ-cc set-generic extension of M.

Given Bukovsky's Theorem we can easily explain why Theorem 17 is true. For (a), take M to be a model of MA_{ω_1}, Martin's axiom for ccc forcings and size ω_1 collections of dense sets. Then it is easy to see that if a Σ_1 formula φ with a subset of ω_1^M as parameter holds in a $\mathbb{P}$-generic extension where $\mathbb{P}$ is a ccc set-forcing, it must hold in M, as it suffices to meet ω_1-many dense subsets of $\mathbb{P}$ to ensure the truth of φ. By Bukovsky, M witnesses (a). For (b) the argument is the same if we can choose M to satisfy *MM*, Martin's Maximum, which asserts that one can meet ω_1-many dense sets for a set-forcing which prserves stationary subsets of ω_1. MM is known to be consistent relative to large cardinals, so again by Bukovsky we obtain (b).

Remark. In the above, *MM* can be replaced by the weaker *BMM*, Bounded Martin's Maximum, which only requires meeting ω_1-many maximal antichains of size ω_1. Whereas *MM* appears to require a supercompact, *BMM* can be forced from just a Woodin cardinal.

Proof of Bukovsky's Theorem. First the easy direction: Suppose that N is a κ-cc set-generic extension of M. We verify that M globally κ-covers N. Suppose that $f : \alpha \to M$ belongs to N, let $\dot{f}$ be a name for f and assume that the trivial condition forces that $\dot{f}$ is a total function from α into M. For each $i < \alpha$ let X_i be a maximal antichain of conditions deciding a value for $\dot{f}(i)$ and let $g(i)$ consist of the values of $\dot{f}(i)$ forced by the various conditions in X_i. Then $g : \alpha \to M$ is a function in M such that $g(i)$ has size $< \kappa$ in M and the trivial condition forces $\dot{f}(i) \in g(i)$ for each $i < \alpha$; so we have shown that the trivial condition forces the conclusion of global κ-covering for the function f.

It follows that if N is a $\mathbb{P}$-generic extension of M by some forcing $\mathbb{P}$ then M globally covers N: just take κ to be larger than any antichain in the forcing $\mathbb{P}$. So we have proved the directions right-to-left in (a) and (b).

Now we turn to the harder direction. Suppose that M globally κ-covers N. We will produce a κ-cc set-forcing $\mathbb{P}$ such that N is a $\mathbb{P}$-generic extension of M.

First note that it suffices to assume that N is of the form $M[A]$ where A is a set of ordinals: Assume that we have the result in this case and now let N be any width-extension of M which M globally κ-covers. Choose a set of ordinals A in N such that $M[A]$ contains all subsets of $(2^{<\kappa})^N$ in N. By assumption $M[A]$ is a κ-cc set-forcing extension of M. But we claim that $M[A]$ must equal all of N: Otherwise choose some set of ordinals B in N so that $M[A][B]$ is larger than $M[A]$. As M globally κ-covers N it follows that $M[A]$ globally κ-covers N (as in global κ-covering it suffices to consider functions from ordinals to ordinals) and therefore $M[A]$ globally κ-covers $M[A][B]$. By assumption $M[A][B]$ is a κ-cc set-generic extension of $M[A]$, which is larger than $M[A]$. But by choice of A, $M[A]$ contains all subsets of $(2^{<\kappa})^N$ of N and therefore all subsets of $(2^{<\kappa})^{M[A]}$ of $M[A]$. This contradicts the following general fact:

Fact. Suppose that $\mathbb{P}$ is a non-atomic κ-cc forcing in M (i.e. no $\mathbb{P}$-generic over M belongs to M). Then $\mathbb{P}$ adds a new subset of $(2^{<\kappa})^M$ over M.

To prove the *Fact*, assume that $\mathbb{P}$ is a non-atomic complete Boolean algebra and using the κ-cc form a non-atomic complete subalgebra $\mathbb{P}_0$ of size $(2^{<\kappa})^M$, by closing $\{0,1\}$ under size $< \kappa$ meets and joins as well as complements and a function which produces incompatible conditions below any nonzero condition. Then a $\mathbb{P}$-generic also adds a $\mathbb{P}_0$-generic, and the latter is a new subset of $(2^{<\kappa})^M$.

OK, so suppose now that M globally κ-covers $N = M[A]$ for some set of ordinals A. Choose an M-cardinal $\lambda = \lambda^{<\kappa}$ such that A is a subset of λ. We'll show that N is a generic extension of M by a κ-cc forcing of size λ.

The language $\mathcal{L}_\kappa^{QF,\lambda}(M)$

The formulas of $\mathcal{L}_\kappa^{QF,\lambda}(M)$ are defined inductively by:

1. *Basic formulas* $\alpha \in \dot{A}$, $\alpha \notin \dot{A}$ for $\alpha < \lambda$.

2. If $\Phi \in M$ is a size $< \kappa$ set of formulas then so are $\bigvee \Phi$ and $\bigwedge \Phi$.

As $\lambda = \lambda^{<\kappa}$ there are only λ-many formulas. We define an ordering of $\mathcal{L}_\kappa^{QF,\lambda}(M)$ as follows:

$B \subseteq \lambda$ *satisfies* φ iff φ is true when $\dot{A}$ is replaced by B.

For φ, ψ in $\mathcal{L}_\kappa^{QF,\lambda}(M)$:

$\varphi \leq \psi$ iff iff for all $B \subseteq \lambda$ (in a set-generic extension of M), if B satisfies φ then B also satisfies ψ.

The above is expressible in M and by Lévy absoluteness, $\varphi \leq \psi$ in M iff $\varphi \leq \psi$ in all width-extensions of M.

Now recall that M globally κ-covers N. Let f be a choice function in N on nonempty subsets Φ of $\mathcal{L}_\kappa^{QF,\lambda}(M)$ in M such that:

If A satisfies some $\psi \in \Phi$ then A satisfies $f(\Phi) \in \Phi$.

(If A satisfies no $\psi \in \Phi$ then $f(\Phi)$ can be any element of Φ.) Using a wellorder in M we can regard f as a function from some ordinal into M. Apply global κ-covering to get g in M so that for all nonempty subsets Φ of $\mathcal{L}_\kappa^{QF,\lambda}(M)$ in M, $g(\Phi) \subseteq \Phi$ has size $< \kappa$ and $f(\Phi) \in g(\Phi)$.

Consider the following set of formulas T in $\mathcal{L}^{QF,\lambda}_{\lambda^+}(M)$ (defined just like $\mathcal{L}^{QF,\lambda}_{\kappa}(M)$, but using size at most λ conjunctions and disjunctions):

$$T = \{(\bigvee \Phi \to \bigvee g(\Phi)) \mid \Phi \subseteq \mathcal{L}^{QF,\lambda}_{\kappa}(M),\ \Phi \in M\}$$

Let $\mathbb{P}$ be the forcing whose conditions are formulas φ of $\mathcal{L}^{QF,\lambda}_{\kappa}(M)$ such that some B satisfies all formulas of $T \cup \{\varphi\}$ (i.e. φ is "consistent with T").

Claim 1. $\mathbb{P}$ is κ-cc.

Proof. Suppose that Φ is a maximal antichain in $\mathbb{P}$. We show that $g(\Phi) = \Phi$, and therefore Φ has size $< \kappa$. It suffices to show that for any $\varphi \in \Phi$ there is some element ψ of $g(\Phi)$ such that $T \cup \{\varphi, \psi\}$ is consistent. Choose any $B \subseteq \lambda$ which satisfies $T \cup \{\varphi\}$ (this is possible because φ belongs to $\mathbb{P}$ and therefore $T \cup \{\varphi\}$ is consistent). As T includes the formula $\bigvee \Phi \to \bigvee g(\Phi)$ it follows that B also satisfies $\bigvee g(\Phi)$ and therefore ψ for some $\psi \in g(\Phi)$. So B satisfies $T \cup \{\varphi, \psi\}$ and therefore this set of formulas is consistent. $\Box$

Claim 2. Let $G(A)$ be the set of $\varphi \in \mathbb{P}$ such that A satisfies φ. Then $G(A)$ is $\mathbb{P}$-generic over M.

Proof. Suppose that Φ is a maximal antichain in $\mathbb{P}$. By Claim 1, Φ has size less than κ, so $\bigvee \Phi$ is a formula in $\mathcal{L}^{QF,\lambda}_{\kappa}(M)$. Now $T \cup \{\sim \bigvee \Phi\}$ is inconsistent, as otherwise $\sim \bigvee \Phi$ violates the maximality of Φ. As A satisfies the formulas in T it follows that A satisfies $\bigvee \Phi$ and therefore some some element of Φ. So $G(A)$ meets Φ. $\Box$

It now follows that $M[A]$ is a $\mathbb{P}$-generic extension of M, as $M[A] = M[G(A)]$. This completes the proof of Bukovsky's Theorem. $\Box$

8.,9.Vorlesungen

A refinemenent of Bukovsky's Theorem

Is there a similar characterisation with "κ-cc" replaced by "size at most κ"?

M κ-decomposes N iff every subset of M in N is the union of at most κ-many subsets, each of which belongs to M.

Theorem 19 *N is a size at most κ forcing extension of M iff M globally κ^+-covers and κ-decomposes N.*

Proof. For the easy direction, suppose that $N = M[G]$ where G is $\mathbb{P}$-generic and $\mathbb{P}$ has size at most κ. As $\mathbb{P}$ is κ^+-cc it follows that M globally κ^+-covers N. To show that M κ-decomposes N, suppose that $X \in N$ is a subset of M and choose $Y \in M$ that covers X. Let $\dot{X}$ be a name for X and for each $p \in G$ forcing that Y covers $\dot{X}$ let X_p consist of those $x \in Y$ such that p forces $x \in \dot{X}$. Then the X_p's give the desired κ-decomposition of X.

Conversely, suppose that M globally κ^+-covers and κ-decomposes N. By Bukovsky's Theorem, N is a $\mathbb{P}$-generic extension of M for some $\mathbb{P}$ which is κ^+-cc. We want to argue that $\mathbb{P}$ is equivalent to a forcing of size at most κ. We may assume that $\mathbb{P}$ is in fact a complete κ^+-cc Boolean algebra which we write as $\mathbb{B}$.

Write N as $M[G]$ where G is $\mathbb{B}$-generic over M. Take a $\mathbb{B}$-name for a κ-decomposition $\dot{G} = \bigcup_{i<\kappa} \dot{G}_i$ of $\dot{G}$, where each $\dot{G}_i$ is forced to belong to M. For each $i < \kappa$ let X_i be a maximal antichain of conditions in $\mathbb{B}$ which decide a specific value in M for $\dot{G}_i$. For each p in X_i let $p(\dot{G}_i)$ denote the value of $\dot{G}_i$ forced by p and $b(p)$ the meet of the conditions in $p(\dot{G}_i)$; $b(p)$ is a nonzero Boolean value because if G_p is generic below p then G_p must contain a condition below each element of $p(\dot{G}_i)$. Let D be the set of $b(p)$ for p in the union of the X_i's. The following *Claim* completes the proof.

Claim. D is dense in $\mathbb{B}$.

Proof of Claim. If q belongs to $\mathbb{P}$ then some r below q forces that q belongs o $\dot{G}_i$ for some i; we can assume that r extends some element p of X_i. But then q is extended by $b(p) \in D$. □

Question 1. In Theorem 19, can "globally κ^+-covers" be eliminated or replaced by "globally λ-covers for some λ"? It can be shown that the latter is possible if one adds the requirement that M just κ^+-covers N, i.e. that subsets of M of size at most κ in N are covered by sets of size at most κ in M.

Question 2. Is there a similar characterisation for κ-closed set-generic extensions?

Bukovsky for class forcing

The proof of Bukovsky's theorem suggests some interesting results concerning class forcing.

The main part of the proof was to show that if $N = M[A]$ where A is a set of ordinals in N and M globally κ-covers N then N is a κ-cc set-generic extension of M. Now let's explore what happens if instead A is a class of ordinals in N so that $N = M[A]$ and N with predicates for M and A satisfies ZFC.

Now we form the big language $\mathcal{L}_{\infty}^{QF,\infty}$ defined by:

1. *Basic formulas* $\alpha \in \dot{A}$, $\alpha \notin \dot{A}$ for all ordinals α.

2. If $\Phi \in M$ is any set of formulas then so are $\bigvee \Phi$ and $\bigwedge \Phi$.

In analogy to the Bukovsky proof we would like to form a class T of sentences in $\mathcal{L}_{\infty}^{QF,\infty}$ so that if $\mathbb{P}$ consists of those sentences in $\mathcal{L}_{\infty}^{QF,\infty}$ which are consistent with T (ordered in the natural way) then A is $\mathbb{P}$-generic over $(M[T], T)$ for the $(M[T], T)$-definable forcing $\mathbb{P}$. We want to choose T so that $\mathbb{P}$ will satisfy the ∞-cc, i.e. so that all $(M[T], T)$-definable antichains are sets in $M[T]$. For simplicity assume that M satisfies $V = L$, so that $N = M[A] = L[A]$ where the symbol L is being used here for L in the sense of N.

First I'll describe a way to achieve this if we allow T to depend not just on the model $N = L[A]$ but on the predicate A. Then I'll indicate how to do this in such a way that T depends only on N, giving rise to the *stability predicate* and the *stable core*.

Define the *A-stability predicate* as follows:

$S(A) = \{(n, \alpha, \beta) \mid n < \omega,\ \alpha < \beta$ are strong limit cardinals of $N = L[A]$ and $(H(\alpha)^N, A \cap \alpha)$ is Σ_n elementary in $(H(\beta)^N, A \cap \beta)\}$.

We remark that $S(A)$ is definable over (N, A) and in fact for any strong limit cardinal β of N and $k < n$, any Σ_k-definable subset of $H(\beta)^{L[S(A)]}$ with a predicate for $S(A) \cap \beta$ is Σ_n-definable over $H(\beta)^N$ with a predicate for $A \cap \beta$.

We form the theory T in the $\mathcal{L}_{\infty}^{QF,\infty}$ of $L[S(A)]$ as follows:

T consists of all axioms of the form

$\bigwedge(\Phi \cap H(\alpha)^N) \to \bigwedge(\Phi \cap H(\beta)^N)$,

where for some n, $\Phi \cap H(\beta)^{L[S(A)]}$ is Σ_k definable over $H(\beta)^{L[S(A)]}$ with a predicate for $S(A) \cap \beta$ and parameters from $H(\alpha)^{L[S(A)]}$ for some $k < n$ and (n, α, β) belongs to $S(A)$.

Note that by the above remark and the definition of $S(A)$, the axioms in T are all true when $\dot{A}$ is interpreted as A. Moreover, T is definable over $(L[S(A)], S(A))$. Let $\mathbb{P}$ be the class of sentences of the $\mathcal{L}_{\infty}^{QF,\infty}$ of $L[S(A)]$ which are consistent with T, ordered by $\varphi \le \psi$ iff every B satisfying $T \cup \{\varphi\}$ also satisfies ψ.

Claim. (a) Any $(L[S(A)], S(A))$-definable maximal antichain on $\mathbb{P}$ is an element of $L[S(A)]$.
(b) $G(A) = \{\varphi \mid \varphi$ is true when $\dot{A}$ is interpreted as $A\}$ is $\mathbb{P}$-generic over $(L[S(A)], S(A))$ for $(L[S(A)], S(A))$-definable maximal antichains.

Proof. (a) Suppose that X is an $(L[S(A)], S(A))$-definable maximal antichain on $\mathbb{P}$ and choose k so that this definition is Σ_k. Choose $n > k$; then X is Σ_n-definable over (N, A). Choose α strong limit so that $(H(\alpha)^N, A \cap \alpha)$ is Σ_n-elementary in (N, A) and contains the parameters in the $(L[S(A)], S(A))$ definition of X. Note that if $\beta > \alpha$ also has this property then $(n, \alpha, beta)$ belongs to the predicate $S(A)$. But then if Φ consists of the negations of the sentences in X, the theory T contains the axioms $\bigwedge(\Phi \cap H(\alpha)^N) \to \bigwedge(\Phi \cap H(\beta)^N)$ for unboundedly many β in $\mathrm{Ord}(N)$ and therefore any formula in which is T-incompatible with all formulas in $X \cap H(\alpha)^N$ is also T-incompatible with all formulas in X, showing that $X = X \cap H(\alpha)^N$ is a set in $L[S(A)]$.
(b) As in the Bukovsky proof, it is clear that $G(A)$ is generic for maximal antichains which belong to $L[S(A)]$; by (a) these are all of the $(L[S(A)], S(A))$-definable maximal antichains. □

To summarise: If $N = L[A]$ where A is a class of ordinals then N is an ∞-cc class-generic extension of $(L[S(A)], S(A))$ where $S(A)$ is the A-stability predicate. A similar argument shows that if $N = M[A]$ where A is a class of ordinals then N is an ∞-cc class-generic extension of $(L[S(M, A)], S(M, A))$ where $S(M, A)$ is the (M, A)-stability predicate, consisting of triples (n, α, β)

where $\alpha < \beta$ are strong limit cardinals in N and $(H(\alpha)^N, M \cap H(\alpha)^N, A \cap \alpha)$ is Σ_n-elementarry in $(H(\beta)^N, M \cap H(\beta)^N, A \cap \beta)$.

Note that the A-stability predicate depends on A and therefore the inner model $(L[S(A)], S(A))$ over which $N = L[A]$ is class-generic is not "canonical", as if $L[A] = L[B]$ it does not follow that $S(A)$ equals $S(B)$. Can we show that N is in fact class-generic over a "canonical" inner model with an N-definable wellorder?

To obtain a positive answer, define an improved A-stability predicate as follows. Again suppose that $N = L[A]$ where $A \subseteq \mathrm{Ord}(N)$ and (N, A) is a model of ZFC. For finite n, a strong limit cardinal α of N is *n-Admissible* if $H(\alpha)^N$ satisfies Σ_n-replacement. We define $S^+(A)$ to consist of all (n, α, β) where n is finite, $\alpha < \beta$ are n-Admissible strong limit cardinals of N and $(H(\alpha)^N, A \cap \alpha)$ is Σ_n-elementary in $(H(\beta)^N, A \cap \beta)$. The only difference between $S^+(A)$ and $S(A)$ is the further requirement of n-Admissibility. As with $S(A)$ we have that $G(A)$ is $\mathbb{P}^+$-generic over $(L[S^+(A)], S^+(A))$ where $\mathbb{P}^+$ is defined using $S^+(A)$ just as $\mathbb{P}$ was defined using $S(A)$.

Theorem 20 *Let S^+ denote $S^+(\emptyset)$. Then there exists an $A \subseteq Ord(N)$ such that $N = L[A]$ and $S^+(A) = S^+$. Therefore N is class-generic over $(L[S^+], S^+)$. The latter is called the "stable core" of N.*

Note that S^+ is N-definable and therefore $L[S^+]$ is contained in the HOD of N. As before we can similarly work with $S^+(M, A)$ for any inner model M and by working with $S^+(\mathrm{HOD}, \emptyset)$ we infer that N is a class-generic extension of (HOD, S^+), the inner model HOD with an additional N-definable predicate.

The class A of the theorem is forced over N with initial segments which are "sufficiently-generic" and as a result of this genericity preserve S^+ up to their length. The key lemma states that conditions can be extended aribitrarily, and this is where the condition of n-Admissibility is needed. The desired class A is obtained by forcing over N with these conditions.

10.,11.Vorlesungen

A converse to the Stable Core Theorem

Recall:

Theorem 21 *(Stable Core Theorem) Let S^+ denote the improved stability predicate, consisting of all (n, α, β) such that $\alpha < \beta$ are n-Admissible strong limit cardinals and $H(\alpha)$ is Σ_n elementary in $H(\beta)$. Then $V = L[S^+, G]$ where G is generic over $(L[S^+], S^+)$ for a definable forcing with the ∞-cc.*

There is a partial converse to this result. First note that the above conclusion still holds if we replace S^+ by any predicate S' contained in S^+ such that for each n there is an unbounded class X of α with $H(\alpha)$ Σ_n elementary in V and (n, α, β) in S' for all $\alpha < \beta$ in X. Say that such a predicate S' is *stability capturing.*

Theorem 22 *If $V = M[G]$ where G is generic for a forcing definable over an inner model (M, A) with the ∞-cc then there is an (M, A)-definable predicate S' such that for some (V, G)-definable club C, $S' \restriction C$ is stability capturing.*

Thus the definability in (M, A) of the Stability Predicate of V is "close" to being equivalent to the statement that V is a definable ∞-cc class-generic extension of (M, A).

Proof of Theorem 22. If $\mathbb{P}$ is the forcing of the hypothesis then we take C to be the club of α such that $G \cap H(\alpha)$ meets all maximal antichains on $\mathbb{P} \cap H(\alpha)$ which belong to $H(\alpha)$. For S' we take all (n, α, β) where all Σ_{n+1}-definable maximal antichains on $\mathbb{P} \cap H(\alpha)$ belong to $H(\alpha)$ (same for β). As $\mathbb{P}$ has the ∞-cc it follows that S' is stability capturing. $\square$

Another version of Bukovsky for class forcing arises if we consider models of Morse-Kelley Class Theory MK:

Theorem 23 *Suppose that $(M, \mathcal{C}^M) \subseteq (N, \mathcal{C}^N)$ are models of MK with global choice and $\mathcal{C}^M$ is definable in $(N, \mathcal{C}^N)$ (by a formula which quantifies over classes). Then each class in $\mathcal{C}^N$ belongs to a class-generic extension of $(M, \mathcal{C}^M)$ via an ∞-cc class forcing iff:*
$(*)$ For any $(N, \mathcal{C}^N)$-definable function f from $\mathcal{C}^M$ to M there is an $(M, \mathcal{C}^M)$-definable function g from $\mathcal{C}^M$ to M such that $f(x) \in g(x)$ for each $x \in \mathcal{C}^M$.

The proof is similar to that of the original Bukovsky theorem. If we go one type further to hyperclass theory, then the formulation of the result is even simpler (as instead of definable functions one can talk about functions that exist as hyperclasses).

Height-extensions and #-Generation

Recall that N is a *height-extension* of M if N contains M and $V_\alpha^N = V_\alpha^M$ for ordinals α in M. So either $M = N$ or $M = (V_\beta)^N$ for some β.

We'll use the notation $M \leq N$ for N is a height-extension of M and $M < N$ for N is a proper height-extension of M.

As we have pointed out, there are universes M with no proper height-extension, such as the minimal model. Of those with proper height-extensions there are two types:

Proposition 24 *(a) (Under mild assumptions) there are universes with height extensions of arbitrarily large countable height.*
(b) There are universes with proper height extensions but where there is a countable bound on the heights of such extensions.
(c) If M carries a definable wellorder and has a proper height-extension then M has a least one.

Proof. (a) Suppose that ω_1 is L-inaccessible. Then there is an elementary ω_1-chain of elementary submodels of L_{ω_1} and the elements of this chain are wellordered by height-extension.
(b) Let β be least so that L_β models ZFC and some $\alpha < \beta$ also models ZFC. Then β is countable in $L_{\beta+2}$ and therefore has no height-extension. So L_α has a proper height extension but none of height greater than β.
(c) If N is a proper height-extension of M then N contains $L_\alpha(M)$ where α is the height of N, and if M has a definable wellorder then this is a model of ZFC. So the least proper height-extension of M is of this form. □

We turn now to height-absoluteness. In analogy with the case of width-absoluteness we state:

Weak Height-absoluteness. M has height-extensions of arbitrarily large countable height and if a Σ_1 formula with parameters from M holds in some $N > M$ then it holds in some $\bar{N} < M$ containing those parameters.

This is indeed quite weak:

Proposition 25 *Suppose that M has height extensions of arbitrarily large countable height and for unboundedly many cardinals α in M, $H(\alpha)^M$ is a model of ZFC. Then M satisfies weak height-absoluteness.*

Proof. Suppose $M < N$ and let φ be a Σ_1 formula with parameters from M. Choose α in M so that the parameters in φ belongs to $H(\alpha^+)^M = H(\alpha^+)^N$. Then by Σ_1 reflection, φ holds in $H(\alpha^+)^M$ and therefore also in $H(\beta)^M$ where $\beta > \alpha$ and $H(\beta)^M$ is a model of ZFC. □

To obtain more height absoluteness we allow parameters that do not belong to M. Vaguely speaking:

Height-absoluteness. M has height-extensions of arbitrarily large countable height and if a Σ_1 formula with parameters from some $N > M$ holds in N then it holds in some $\bar{N} < M$ with "corresponding" parameters.

The meaning of "corresponding" parameters must be clarified.

First consider the case of parameters in $N > M$ which are subsets of M.

Proposition 26 *Suppose that $M < N$ and Ord(M) is regular (and therefore inaccessible) in N. Then if φ is a Σ_1 formula with a subset X of M as parameter which holds in N, then for some $\bar{M} < \bar{N} < M$, $\varphi(X \cap \bar{M})$ holds in $\bar{N}$.*

Thus in this case the parameter "corresponding" to X is simply the intersection of X with $\bar{M}$, an initial segment of X.

Proof. If $\varphi(X)$ holds in N, then using the regularity of Ord(M) in N, we can form a Σ_1 elementary submodel H of N containing X as an element whose intersection with M is some $\bar{M} < M$. Let $\bar{H}$ be the transitive collapse of H. Then $\bar{H}$ satisfies $\varphi(X \cap \bar{M})$. As Ord(M) is regular in N, there is some $\bar{N} < M$ which contains $\bar{H}$ and then $\varphi(X \cap \bar{M})$ also holds in $\bar{N}$. □

But we run into a new problem if we try to consider parameters in $N > M$ which are not subsets of M but instead sets of subsets of M. Suppose that $\mathcal{S}$ is such a parameter and N satisfies the Σ_1 formula φ with parameter $\mathcal{S}$. We want to assert that for some $\bar{M} < \bar{N} < M$, φ holds in $\bar{N}$ for a parameter $\bar{\mathcal{S}}$ "corresponding to" $\mathcal{S}$, which should be a set of subsets of $\bar{M}$. But it is not clear what $\bar{\mathcal{S}}$ should be; we cannot just take $\bar{\mathcal{S}} = \{X \cap \bar{M} \mid X \in \mathcal{S}\}$ for some $\bar{M} < M$ as for example $\varphi(\mathcal{S})$ could assert that $\mathcal{S}$ is the set of bounded subsets of Ord(M) in which case the latter parameter will contain subsets of Ord$(\bar{M})$ which are unbounded in Ord$(\bar{M})$.

An option that dates back to work of V. Marshall and Magidor is to instead take $\bar{\mathcal{S}}$ to be the image of $\mathcal{S}$ under the transitive collapse of an elementary submodel of (enough of) N whose intersection with M is transitive. Following this route leads us to supercompactness:

To clarify matters, think of subsets of M as elements of $H(\kappa^+)^N$. So now our parameter $\mathcal{S}$ is a subset of $H(\kappa^+)^N$ and we can form the structure $(H(\kappa^+)^N, \mathcal{S})$. A special case of our Σ_1 formula $\varphi(\mathcal{S})$ is one which asserts that this structure satisfies some first-order property.

Definition 27 *κ is* subcompact *if for any $\mathcal{S} \subseteq H(\kappa^+)$ there are $\bar{M} < M = H(\kappa)$, $\bar{\mathcal{S}} \subseteq H(\bar{\kappa}^+)$ and elementary $\pi : (H(\bar{\kappa}^+), \bar{\mathcal{S}}) \to (H(\kappa^+), \mathcal{S})$ with critical point $\bar{\kappa}$ sending $\bar{\kappa}$ to κ.*

So we can use subcompactness to provide a version of height-absoluteness with parameters contained in $H(\kappa^+)$. More generally, for α any cardinal greater than κ, define *α-subcompact* in a similar way, replacing $(H(\kappa^+), \mathcal{S})$ and $(H(\bar{\kappa}^+), \bar{\mathcal{S}})$ by $(H(\alpha), \mathcal{S})$ and $(H(\bar{\alpha}), \bar{\mathcal{S}})$ and requiring $\bar{\alpha} < \kappa$. Then this provides a version of height-absoluteness with parameters contained in $H(\alpha)$.

Theorem 28 *The following are equivalent:*
(a) κ is α-subcompact for all α.
(b) κ is supercompact.

Idea of Proof. For simplicity assume GCH and we show that for regular $\kappa < \alpha$, if κ is β-subcompact for a large enough β then it is α-supercompact, and if κ is α-supercompact then it is also α-subcompact.

Suppose that κ is β-subcompact for a large β. Then apply β-subcompactness to the structure $(H(\beta), \mathcal{S})$ where $\mathcal{S}$ is just $\{\alpha\}$. We then get an elementary $\pi : (H(\bar{\beta}), \{\bar{\alpha}\}) \to (H(\beta), \{\alpha\})$ with critical point $\bar{\kappa}$, sending $\bar{\kappa}$ to κ. But the range $\pi[\bar{\alpha}]$ of π on $\bar{\alpha}$ belongs to $H(\beta)$ so we get a supercompactness measure $\mathcal{U}$ on $P_{\bar{\kappa}}\bar{\alpha}$ defined by $X \in \mathcal{U}$ iff $\pi[\bar{\alpha}]$ belongs to $\pi(X)$; moreover this measure is in the domain of π. So $\bar{\kappa}$ is $\bar{\alpha}$-supercompact and by elementarity κ is α-supercompact.

Conversely, suppose that κ is α-supercompact and let $\pi : V \to M$ witness this. Then for any $\mathcal{S} \subseteq H(\alpha)$, π restricted to $(H(\alpha), \mathcal{S})$ belongs to M. So M sees that there is $\pi : (H(\alpha), \mathcal{S}) \to (H(\pi(\alpha)), \pi(\mathcal{S}))$ with critical point $\kappa < \pi(\kappa)$ sending κ to $\pi(\kappa)$ and so by elementarity, V sees that there is some

$\bar{\pi} : (H(\bar{\alpha}), \bar{\mathcal{S}}) \to (H(\alpha), \mathcal{S})$ with critical point some $\bar{\kappa}$ less than κ, sending $\bar{\kappa}$ to κ . So κ is α-subcompact. $\Box$

So if $M < N$ and $\mathrm{Ord}(M)$ is aupercompact in N we have a version of height-absoluteness between M and N that applies to any Σ_1 formula with any parameter from N.

Have we solved the problem of height-absoluteness? I don't think so, for several reasons.

One problem is that the replacement of the parameter $\mathcal{S}$ by the "corresponding parameter" $\bar{\mathcal{S}}$ is not "canonical", i.e. it depends on the choice of embedding π. There are apparently many witnesses π to subcompactness, yielding unrelated "corresponding parameters".

Second, height-absoluteness should be a property of *height* and not of *width*. So there should be initial segments of L which fulfill this property. Of course this will not be the case if we insist on supercompactness.

Third, and this is an issue for all strong forms of height-absoluteness, we would like the height-absoluteness of M to be independent of N. We cannot expect that $\mathrm{Ord}(M)$ is supercompact in all of its height-extensions, but is it well-motivated to only consider height-extensions in which this is the case?

So we take a different approach to the problem of height-absoluteness, extrapolating on the usual form of reflection provable in ZF. This leads to the theory of *#-generation*.

12.,13.Vorlesungen

#-generation

Before embarking on our analysis of height-absoluteness we should take note of the following: No *first-order* statement φ can be adequate to fully capture such absoluteness. This is simply because a first-order statement true in M will reflect to one of its rank initial segments and we are then naturally led from φ to the stronger first-order statement "φ holds both in M and in some rank initial-segment satisfying ZFC".

But how do we capture height-absoluteness with a non first-order axiom? We do this via a detailed analysis of the relationship between M and its height-extensions and height-restrictions (i.e. its rank initial-segments).

To save on notation, we'll use the symbol V not for the entire universe of sets but for the countable universe that we have been calling M. Thus we'll freely write V_α instead of V_α^V.

A special case of height-absoluteness is reflection, which says that properties of V "trasnfer" or "reflect" to rank initial segments V_κ. Standard reflection tells us that a single first-order property of V with parameters will hold in some V_κ which contains those parameters. It is natural to strengthen this to the simultaneous reflection of all first-order properties of V to some V_κ, allowing arbitrary parameters from V_κ. Thus we have reflected V to a V_κ which is an elementary submodel of V.

Repeating this process leads us to an increasing, continuous sequence of ordinals $(\kappa_i \mid i < \infty)$, whee ∞ denotes the ordinal height of V, such that the models $(V_{\kappa_i} \mid i < \infty)$ form a continuous chain $V_{\kappa_0} \prec V_{\kappa_1} \prec \cdots$ of elementary submodels of V whose union is all of V.

Let C be the set of the κ_i's, a proper class of V. We can apply reflection to V with C as an additional predicate to infer that properties of (V, C) also hold of some $(V_\kappa, C \cap \kappa)$. But the unboundedness of C is a property of (V, C) so we get some $(V_\kappa, C \cap \kappa)$ where $C \cap \kappa$ is unbounded in κ and therefore κ belongs to C. As a corollary, properties of V in fact hold in some V_κ where κ belongs to C. It is convenient to formulate this in its contrapositive form: If a property holds of V_κ for all κ in C then it also holds of V.

Now note that for all κ in C, V_κ can be *lengthened* (*height-extended*) to an elementary extension (namely V) of which it is a rank-initial segment. By the contrapositive form of reflection of the previous paragraph, V itself also has such a lengthening (height-extension) V^*.

But this is clearly not the end of the story. For the same reason we can also infer that there is a continuous increasing sequence of such lengthenings $V = V_{\kappa_\infty} \prec V^*_{\kappa_{\infty+1}} \prec V^*_{\kappa_{\infty+2}} \prec \cdots$ of length ω_1 For a further ease of notation, let us drop the *'s and write W_{κ_i} instead of $V^*_{\kappa_i}$ for $\infty < i$ and instead of V_{κ_i} for $i \leq \infty$. Thus V equals W_∞.

But which tower $V = W_{\kappa_\infty} \prec W_{\kappa_{\infty+1}} \prec W_{\kappa_{\infty+2}} \prec \cdots$ of lengthenings of V should we consider? Can we make the choice of this tower *canonical*?

Consider the entire sequence $W_{\kappa_0} \prec W_{\kappa_1} \prec \cdots \prec V = W_{\kappa_\infty} \prec W_{\kappa_{\infty+1}} \prec W_{\kappa_{\infty+2}} \prec \cdots$. The intuition is that all of these models resemble each other in the sense that they share the same first-order properties. Indeed by virtue of the fact that they form an elementary chain, these models all satisfy the same first-order sentences. But again in the spirit of "resemblance", the following should hold: For $i_0 < i_1$ regard $(W_{\kappa_{i_1}}, W_{\kappa_{i_0}})$ as the structure $(W_{\kappa_{i_1}}, \in)$ together with $W_{\kappa_{i_0}}$ as a unary predicate. Then it should be the case that any two such pairs $(W_{\kappa_{i_1}}, W_{\kappa_{i_0}})$, $(W_{\kappa_{j_1}}, W_{\kappa_{j_0}})$ (with $i_0 < i_1$ and $j_0 < j_1$) satisfy the same first-order sentences, even allowing parameters which belong to both $W_{\kappa_{i_0}}$ and $W_{\kappa_{j_0}}$. Generalising this to triples, quadruples and n-tuples in general we arrive at the following situation:

$(*)$ V occurs in a continuous elementary chain $W_{\kappa_0} \prec W_{\kappa_1} \prec \cdots \prec V = W_{\kappa_\infty} \prec W_{\kappa_{\infty+1}} \prec W_{\kappa_{\infty+2}} \prec \cdots$ of length $\infty + \infty$, where the models W_{κ_i} form a *strongly-indiscernible chain* in the sense that for any n and any two increasing n-tuples $\vec{i} = i_0 < i_1 < \cdots < i_{n-1}$, $\vec{j} = j_0 < j_1 < \cdots < j_{n-1}$, the structures $W_{\vec{i}} = (W_{\kappa_{i_{n-1}}}, W_{\kappa_{i_{n-2}}}, \cdots, W_{\kappa_{i_0}})$ and $W_{\vec{j}}$ (defined analagously) satisfy the same first-order sentences, allowing parameters from $W_{\kappa_{i_0}} \cap W_{\kappa_{j_0}}$.

We are getting closer to the desired axiom of #-generation. Surely we can impose higher-order indiscernibility on our chain of models. For example, consider the pair of models $W_{\kappa_0} = V_{\kappa_0}$, $W_{\kappa_1} = V_{\kappa_1}$. We can require that these models satisfy the same second-order sentences; equivalently, we require that $H(\kappa_0^+)^V$ and $H(\kappa_1^+)^V$ satisfy the same first-order sentences. But as with the pair $H(\kappa_0)^V$, $H(\kappa_1)^V$ we would want $H(\kappa_0^+)^V$, $H(\kappa_1^+)^V$ to satisfy the same first-order sentences *with parameters*. How can we formulate this? For example, consider κ_0, a parameter in $H(\kappa_0^+)^V$ that is second-order with respect to $H(\kappa_0)^V$; we cannot simply require $H(\kappa_0^+)^V \vDash \varphi(\kappa_0)$ iff $H(\kappa_1^+)^V \vDash \varphi(\kappa_0)$, as κ_0 is the largest cardinal in $H(\kappa_0^+)^V$ but not in $H(\kappa_1^+)^V$. Instead we need to replace the occurence of κ_0 on the left side with a "corresponding" parameter on the right side, namely κ_1, resulting in the natural requirement $H(\kappa_0^+)^V \vDash \varphi(\kappa_0)$ iff $H(\kappa_1^+)^V \vDash \varphi(\kappa_1)$. More generally, we should be able to replace each parameter in $H(\kappa_0^+)^V$ by a "corresponding" element of $H(\kappa_1^+)^V$. It is natural to solve this problem of "corresponding parameters" using embeddings.

Definition 29 *A structure $N = (N, U)$ is called a* $\#$ with critical point κ, *or just a $\#$, if the following hold:*
(a) N is a model of ZFC$^-$ (ZFC minus powerset) in which κ is both the largest cardinal and strongly inaccessible.
(b) (N, U) is amenable (i.e. $x \cap U \in N$ for any $x \in N$).
(c) U is a normal measure on κ in (N, U).
(d) N is iterable, i.e., all of the successive iterated ultrapowers starting with (N, U) are well-founded, yielding iterates (N_i, U_i) and Σ_1 elementary iteration maps $\pi_{ij} : N_i \to N_j$ where $(N, U) = (N_0, U_0)$.

We let κ_i denote the largest cardinal of the i-th iterate N_i. Also, LP(N_i denotes the V_{κ_i} of N_i (LP stands for *lower part.*) LP(N_i) is a model of ZFC.

Definition 30 *We say that a transitive model V of ZFC is* $\#$-generated *iff there is $N = (N, U)$, a $\#$ with iteration $N = N_0 \to N_1 \to \cdots$, such that V equals LP(N_∞) where ∞ denotes the ordinal height of V.*

$\#$-generation fulfills our requirements for height-absoluteness, with powerful consequences for reflection. L is $\#$-generated iff $0^\#$ exists, so this principle is compatible with $V = L$. If V is $\#$-generated via (N, U) then there are elementary embeddings from V to V which are canonically-definable through iteration of (N, U): In the above notation, any order-preserving map from the κ_i's to the κ_i's extends to such an elementary embedding. If $\pi : V \to V$ is any such embedding then we obtain not only the indiscernibility of the structures $H(\kappa_i^+)$, for all i but also of the structures $H(\kappa_i^{+\alpha})$ for any $\alpha < \kappa_0$ and more. Moreover, $\#$-generation evidently provides the maximum amount of height-absoluteness: If V is generated by (N, U) as LP(N_∞) where ∞ is the ordinal height of V, and x is any parameter in a further iterate $V^* = N_{\infty^*}$ of (N, U), then any first-order property $\varphi(V, x)$ that holds in V^* reflects to $\varphi(V_{\kappa_i}, \bar{x})$ in N_j for all sufficiently large $i < j < \infty$, where $\pi_{j,\infty^*}(\bar{x}) = x$. This implies any known form of height-absoluteness and summarizes the amount of reflection one has in L under the assumption that $0^\#$ exists, the maximum amount of reflection that L can have. This is reinforced by Jensen's *$\#$-generated coding theorem* which states that if V is $\#$-generated then V can be coded into a $\#$-generated model $L[x]$ for a real x where the given $\#$ which generates V extends to the natural generator $x^\#$ for the model $L[x]$.

From this we can conclude that $\#$-generated models have all of the large cardinal and reflection properties that L has when $0^\#$ exists.

#-generation also answers our question about which *canonical* tower of lengthenings of V to look at in height-absoluteness, namely the further lower parts of iterates of any # that generates V. And #-generation fully realizes the idea that V should look exactly like closed unboundedly many of its rank initial segments as well as its *canonical* lengthenings of arbitrary countable ordinal height.

In summary, #-generation stands out as a compelling formalization of the principle of *height-absoluteness*. It is not first-order (we have argued that no optimal height-absoluteness principle can be), however it is second-order in a very restricted way: For a countable V, the property of being a # that generates V is expressible by quantifying universally over the models $L_\alpha(V)$ as α ranges over the countable ordinals.

14.,15.Vorlesungen

We have argued that #-generation is the optimal formulation of absoluteness in height (height-maximality). But can we strengthen the claim that there are #-generated models? For example, is L #-generated, or equivalently, does $0^\#$ exist?

We'll see now that the existence of $0^\#$ in fact follows from the generalised form of Lévy absoluteness that we considered earlier. Recall:

M satisfies Lévy(ω_1) iff whenever a Σ_1 formula with parameter ω_1^M has a solution in an ω_1-preserving width-extension of M then it has a solution in M.

Theorem 31 *Assuming large cardinals, there exists an M satisfying Lévy(ω_1).*

We now show:

Theorem 32 *Assume Lévy(ω_1). Then $0^\#$ exists.*

A more sophisticated proof, due to Welch and myself, shows that this conclusion can be strengthened to "There are measurable cardinals in inner models of arbitrarily high Mitchell order".

Proof. Suppose that M is a universe satisfying that $0^\#$ does not exist. We show that there is a Σ^1_3 sentence (in 2nd order arithmetic) true in a class-forcing extension of M (satisfying ZFC) which does not hold in M. Now any

Σ^1_3 sentence can be translated into a Σ_1 sentence with parameter ω_1: If φ is Π^1_2 then:

$\exists x \varphi(x)$ iff
$\exists x(L_{\omega_1}[x] \vDash \varphi(x)$ iff
$\exists x \exists T(T \vDash \varphi(x)$ and $T = L_{\omega_1}[x]$,

and the last sentence above is Σ_1 with parameter ω_1.

The proof is based on the following result concerning L-definable partitions:

Theorem 33 *There exists an L-definable function n : L-Singulars $\to \omega$ such that if M satisfies $0^{\#}$ does not exist:*

1. *For some k, $M \models \{\alpha \mid n(\alpha) \leq k\}$ is Δ_2-stationary.*
2. *For each k there is a generic extension of M in which $\{\alpha \mid n(\alpha) \leq k\}$ is not Δ_2-stationary.*

Remark. "Δ_2-stationary in M" means: intersects every closed unbounded class of ordinals which is $\Delta_2(M)$-definable with parameters.

Proof. We define $n(\alpha)$. Let $\langle C_\alpha \mid \alpha$ L-singular$\rangle$ be an L-definable $\Box$-sequence: C_α is closed unbounded in α, ordertype $C_\alpha < \alpha$ and $\bar{\alpha} \in \lim C_\alpha \to C_{\bar{\alpha}} = C_\alpha \cap \bar{\alpha}$. Let ot C_α denote the ordertype of C_α. If ot C_α is L-regular then $n(\alpha) = 0$. Otherwise $n(\alpha) = n(\text{ot } C_\alpha) + 1$.

1 is clear, as otherwise there is a closed unbounded $C \subseteq L$-regulars definable in M, contradicting the Covering Theorem and the hypothesis that $0^{\#}$ does not exist in M.

Now we prove 2. Fix $n \in \omega$. In M let P consist of closed, bounded $p \subseteq$ ORD such that $\alpha \in p \to \alpha$ L-regular or $n(\alpha) \geq n+1$, ordered by $p \leq q$ iff p end extends q.

We claim that P is ∞-distributive in M. Suppose that $p \in P$ and $\langle D_\alpha \mid \alpha < \kappa\rangle$ is a definable sequence of open dense subclasses of P, κ regular. We wish to find $q \leq p$, $q \in D_\alpha$ for all $\alpha < \kappa$. Let C be the class of all strong limit cardinals β such that $D_\alpha \cap V_\beta$ is dense in $P \cap V_\beta$ for all $\alpha < \kappa$, a closed unbounded class of ordinals. It suffices to show that $C \cap \{\beta \mid n(\beta) \geq n+1\}$ has a closed subset of ordertype $\kappa+1$, for then p can be successively extended κ times meeting the D_α's, to conditions with maximum in $\{\beta \mid n(\beta) \geq n+1\}$; the final condition (at stage κ) extends p and meets each D_α.

Lemma 34 *Suppose $m \geq k$, α is regular and C is a closed set of ordertype $\alpha^{+m}+1$, consisting of ordinals greater than α^{+m} (where $\alpha^{+0} = \alpha$, $\alpha^{+(p+1)} = (\alpha^{+p})^+$). Then $C \cap \{\beta \mid n(\beta) \geq k\}$ has a closed subset of ordertype $\alpha^{+(m-k)}+1$.*

Proof. By induction on k. Suppose $k = 0$. Let $\beta = \max C$. Then β is singular and hence singular in L. So C_β is defined and $\lim(C_\beta \cap C)$ is a closed set of ordertype $\alpha^{+m}+1$ consisting of L-singulars. So $\lim(C_\beta \cap C) \subseteq C \cap \{\gamma \mid n(\gamma) \geq 0\}$ satisfies the lemma.

Suppose the lemma holds for k and let $m+1 \geq k+1$, C a closed set of ordertype $\alpha^{+(m+1)}+1$ consisting of ordinals greater than $\alpha^{+(m+1)}$. Let $\beta = \max C$. Then C_β is defined and $D = \lim(C_\beta \cap C)$ is a closed set of ordertype $\alpha^{+(m+1)}+1$. Let $\bar{\beta} = (\alpha^{+m} + \alpha^{+m} + 1)$st element of D. Then $\bar{D} = \{\text{ot } C_\gamma \mid \gamma \in D, (\alpha^{+m}+1)\text{st element of } D \leq \gamma \leq \bar{\beta}\}$ is a closed set of ordertype $\alpha^{+m}+1$ consisting of ordinals greater than α^{+m}. By induction there is a closed $\bar{D}_0 \subseteq \bar{D} \cap \{\gamma \mid n(\gamma) \geq k\}$ of ordertype $\alpha^{+(m-k)}+1$. But then $D_0 = \{\gamma \in D \mid \text{ot } C_\gamma \in \bar{D}_0\}$ is a closed subset of $C \cap \{\gamma \mid n(\gamma) \geq k+1\}$ of ordertype $\alpha^{+(m-k)}+1$. As $\alpha^{+(m-k)} = \alpha^{+((m+1)-(k+1))}$ we are done. $\Box$ (Lemma)

By the lemma, $C \cap \{\beta \mid n(\beta) \geq n\}$ has arbitrary long closed subsets for any n, for any closed unbounded $C \subseteq \text{ORD}$. It follows that P is ∞-distributive. Now to prove 2, we apply the forcing P to M, producing C witnessing the nonstationarity of $\{\alpha \mid n(\alpha) \leq n\}$, and then follow this with the forcing to code $\langle M, C\rangle$ by a real, making C Δ_2-definable. Of course this will not produce $0^\#$ as every successor to a strong limit cardinal is preserved in the coding. $\Box$

Proof of the Theorem. We use David's trick. Let φ_n be the sentence: $\exists R \forall \alpha$(If $L_\alpha[R] \models ZF^-$ then $L_\alpha[R] \models \beta$ a limit cardinal $\rightarrow \beta$ L-regular or $n(\beta) \geq n$). (This is equivalent to a Σ^1_3 sentence as it is of the form $\exists R\psi(R)$ where $\psi(R)$ is Π_1 in the sense of Lévy and hence equivalent to a Π^1_2 formula.) By Theorem (2) and cardinal collapsing (to guarantee that limit cardinals β are either L-regular or satisfy $n(\beta) \geq n$), M has a generic extension $L[R] \models \beta$ a limit cardinal $\rightarrow$

β L-regular or $n(\beta) \geq n$. Using David's trick we can in fact obtain φ_n in $L[R]$. $\square$

A variant of Lévy(ω_1)

We show that an interesting variant of Lévy(ω_1) is in fact equivalent to the existence of $0^\#$.

Let's say that $\mathbb{P}$ is an ω_1*-forcing* if it is a forcing with universe ω_1. We consider statements of the following form:

$(*)$ If $\mathbb{P}$ is a constructible ω_1-forcing and for each $p \in \mathbb{P}$ there is a $\mathbb{P}$-generic over L containing p in an ω_1-preserving width-extension of V then there is a $\mathbb{P}$-generic over L in V.

As with Lévy(ω_1) the above asserts that if a certain type of property holds in an ω_1-preserving width-extension of V then it already holds in V. If $(*)$ holds then we say that *V is L-saturated for ω_1-forcings.*

Theorem 35 *The following are equivalent:*
(a) V is L-saturated for ω_1-forcings.
(b) $0^\#$ exists.

Proof. (a) $\to$ (b) The existence of $0^\#$ is equivalent to the statement that every stationary constructible subset of ω_1 contains a CUB subset. Now use the following:

Fact. (Baumgartner) If X is a stationary constructible subset of ω_1 then there is a forcing $P \in L$ of L-cardinality ω_1 which preserves cardinals over V and adds a CUB subset to X. (P adds a CUB subset of X using "finite conditions".)

(b) $\to$ (a) Assume that $0^\#$ exists and suppose that P is a constructible forcing of L-cardinality ω_1 such that every condition in P belongs to a generic in an ω_1-preserving extension of V. We will show that there is a P-generic in V. Assume that the universe of P is exactly ω_1. Let P be of the form $t(\vec{i}, \omega_1, \vec{\infty})$ where $\vec{i} < \omega_1 < \vec{\infty}$ is a finite increasing sequence of indiscernibles and t is an L-term. We claim that if $\vec{i} < k_0 < k_1$ are countable indiscernibles and G_{k_0} is P_{k_0}-generic over L then there is G_{k_1} containing G_{k_0} which is P_{k_1}-generic over L, where $P_k = t(\vec{i}, k, \vec{\infty})$. If not, then player I wins the open game $\mathcal{G}(k_0, k_1, G_{k_0})$ where I chooses constructible dense subsets of P_{k_1} and II responds with increasingly strong conditions meeting these dense sets which are compatible with all conditions in G_{k_0}. The latter is a property of the model $L[G_{k_0}]$. Let $p \in P_{k_0}$ be a condition forcing that I wins $\mathcal{G}(k_0, k_1, G_{k_0})$. Then p forces that I wins $\mathcal{G}(k_2, k_3, G_{k_2})$, where $k_2 < k_3$ are any indiscernibles $\geq k_0$ and G_{k_2} denotes the P_{k_2}-generic. But now let G be a P-generic containing p in an ω_1-preserving extension of V. As G preserves ω_1 over V, there are indiscernibles $k_2 < k_3$ with $k_0 \leq k_2$ such that $G \cap k_2$ is P_{k_2}-generic and $G \cap P_{k_3}$ is P_{k_3}-generic, so clearly player II has a winning strategy in the game $\mathcal{G}(k_2, k_3, G \cap P_{k_2})$, in contradiction to the choice of p.

Now it is easy to build a P-generic: List the countable indiscernibles greater than $\vec{i}$ as $j_0 < j_1 < j_2 < \cdots$ and inductively choose P_{j_α}-generic G_α such that $\alpha < \beta$ implies $G_\alpha \subseteq G_\beta$. At the first step, G_{j_0} is an arbitrary P_{j_0}-generic. By the previous paragraph there is no difficulty at the successor steps, where one extends G_{j_α} to $G_{j_{\alpha+1}}$. At limit stages λ, the P_{j_λ}-genericity of the union G_{j_λ} of the G_{j_α}, $\alpha < \lambda$, follows by indiscernibility. The desired P-generic is the union of the G_{j_α}, $\alpha < \omega_1$. □

16.,17.Vorlesungen

The existence of (slightly more than) $0^\#$ also gives a strong form of L-saturation for class forcing. Work now in Gödel-Bernays class theory.

Suppose that $0^\#$ exists and say that a forcing $\mathbb{P}$ is an *L-forcing* if for some $A \subseteq L$, $\mathbb{P}$ is (L, A)-definable and (L, A) satisfies ZFC. The existence of $0^\#$ implies that all such A are definable in $L[0^\#]$ with ordinal parameters.

A cardinal κ is *α-Erdős* if whenever C is a club in κ and $f : [C]^{<\omega} \to \kappa$ is regressive (i.e. $f(a) < \min(a)$ for all a) then for some subset x of C of ordertype α, f is constant on $[x]^n$ for each n. We say that *Ord is α-Erdős* if this holds when C is a club in Ord and f is a class function.

Theorem 36 *Suppose that $0^\#$ exists and Ord is $\omega + \omega$-Erdős. If $\mathbb{P}$ is an L-forcing definable over $(L, A) \models$ ZFC which has a generic (over (L, A)), then there is such a generic G which is definable in a set-forcing extension of $L[0^\#]$. Moreover the model $L[G]$ is $\#$-generated.*

I.e., with slightly more than the existence of $0^\#$, we have that, modulo set-forcing, $L[0^\#]$ is "saturated" for L-forcings.

One would like to have a converse to this result, stating that if the universe is "saturated" for L-forcings modulo set-forcing, then $0^\#$ exists. For this it would suffice to have a version of Baumgartner's forcing to add a club to ω_1 with finite conditions that applies to large stationary classes (such as the "square-sequence dropping" classes $\{\alpha \mid n(\alpha) \geq k\}$ discussed earlier). Unfortunately, with the present state of knowledge, there are ways of adding clubs to ω_2 with finite conditions, but not to ω_3 and surely not to Ord.

One would also like to eliminate the assumption of an $\omega+\omega$-Erdős cardinal in Theorem 36, however something more than just the existence of $0^\#$ is needed for the last conclusion of the Theorem, regarding #-generation:

Theorem 37 *Suppose that $0^\#$ exists and M is a proper inner model of $L[0^\#]$. Then in M, for every ordinal α there is an α-Mahlo cardinal. But if there are no inaccessibles in $L[0^\#]$, there is a proper inner model of $L[0^\#]$ in which no cardinal α is α-Mahlo.*

If M is an inner model of $L[0^\#]$ in which no cardinal α is α-Mahlo then M cannot be #-generated, as #-generation implies the existence of such cardinals. So even if L is #-generated, there can be class-generic extensions of L which are inner models of $L[0^\#]$ but not #-generated.

Corollary 38 *Assume that $0^\#$ exists. Then:*
(a) There are L-forcings with no generic.
(b) There can be L-forcings with generics but no #-generated generic (i.e. no generic G such that $L[G]$ is #-generated).
(c) If there is an $\omega+\omega$-Erdős cardinal then every L-forcing with a generic has a #-generated generic.

For an example of (a) above, consider a forcing that adds a club through the L-singulars. This has no generic as I = the class of Silver indiscernibles is a club consisting of L-regulars.

However most "nice" L-forcings do have #-generated generics assuming just the existence of $0^\#$, provided that they are of "reverse Easton" type. To explain the distinction between "reverse Easton" and "forward Easton" forcings consider the following.

Proposition 39 *Suppose κ is L-regular and let $\mathbb{P}(\kappa)$ denote κ-Cohen forcing in L.*
(a) If κ has cofinality ω in $L[0^\#]$ then $\mathbb{P}(\kappa)$ has a generic over L.
(b) If κ has uncountable cofinality in $L[0^\#]$ then $\mathbb{P}(\kappa)$ has not generic over L.

Proof. Let j_n denote the first n Silver indiscernibles $\geq \kappa$.

(a) We use the fact that $\mathbb{P}(\kappa)$ is κ-distributive in L. Let $\kappa_0 < \kappa_1 < \ldots$ be an ω-sequence in $L[0^\#]$ cofinal in κ. Then any $D \subseteq P(\kappa)$ in L belongs to $\mathrm{Hull}(\kappa_n \cup j_n)$ for some n, where Hull denotes Skolem hull in L. As $\mathrm{Hull}(\kappa_n \cup j_n)$ is constructible of L-cardinality $< \kappa$ we can use the κ-distributivity of $P(\kappa)$ to choose $p_0 \geq p_1 \geq \ldots$ successively below any $p \in P(\kappa)$ to meet all dense $D \subseteq P(\kappa)$ in L.

(b) Note that in this case $\kappa \in \mathrm{Lim}\ I$, as otherwise $\kappa = \cup\{\kappa_n | n \in \omega\}$ where $\kappa_n = \cup(\kappa \cap \mathrm{Hull}(\bar\kappa + 1 \cup j_n)) < \kappa, \bar\kappa = \max(I \cap \kappa)$, and hence κ has $L[0^\#]$-cofinality ω. Suppose $G \subseteq \mathbb{P}(\kappa)$ were $\mathbb{P}(\kappa)$-generic over L. For any $p \in \mathbb{P}(\kappa)$ let $\alpha(p)$ denote the domain of p. Define $p_0 \geq p_1 \geq \ldots$ in G so that $\alpha(p_{n+1}) \in I$ and p_{n+1} meets all dense $D \subseteq \mathbb{P}(\kappa)$ in $\mathrm{Hull}(\alpha(p_n) \cup j_n)$. Then $p = \cup\{p_n | n \in \omega\}$ meets all dense $D \subseteq \mathbb{P}(\kappa)$ in $\mathrm{Hull}(\alpha \cup j)$ where $\alpha = \cup\{\alpha(p_n) | n \in \omega\} \in I$, $j = \cup\{j_n | n \in \omega\}$. But then p is $\mathbb{P}(\alpha)$-generic over L, as every constructible dense $\bar D \subseteq \mathbb{P}(\alpha)$ is of the form $D \cap \mathbb{P}(\alpha)$ for some D as above. So p is not constructible, contradicting $p \in G$. $\Box$

It follows that in the presence of $0^\#$ there can be no generic for the Easton product which adds an α-Cohen set to each L-regular α. However we can have generics for the reverse Easton *iteration* of α-Cohen. Recall that this is the iteration $(\mathbb{P}_\alpha \mid \alpha \in \mathrm{Ord})$ where for L-regular α, $\mathbb{P}_{\alpha+1} = \mathbb{P}_\alpha * \alpha$-Cohen, using Easton support (i.e. for L-inaccessible α, conditions are trivial on a final segment of α).

Theorem 40 *Assume that $0^\#$ exists. Then there is a generic over L for the reverse Easton iteration of α-Cohen.*

Proof. Recall that $\mathbb{P}(< \alpha)$ has a dense subset of L-cardinality $\leq (\alpha^+)^L$ for each α. By induction on $i \in I$ we define $G(\leq i) = G(< i) * G(i)$ to be $\mathbb{P}(\leq i)$-generic over L, where $\mathbb{P}(\leq i) = \mathbb{P}(< i) * \mathbb{P}(i)$, the first $i+1$ stages in the iteration defining $\mathbb{P}$. We will have: $i \leq j$ in $I \longrightarrow G(j)$ extends $G(i)$; this will enable us to get through limit stages. For $i = \min I$, take $G(\leq i)$ to be any $\mathbb{P}(\leq i)$-generic in $L[0^\#]$. If $G(\leq i)$ has been defined and i^* is the I-successor to i, then write $\mathbb{P}(< i^*)$ as $\mathbb{P}(\leq i) * \mathbb{P}[i+1, i^*)$ and as $\mathbb{P}(\leq i) \Vdash \mathbb{P}[i+1, i^*)$ is i^+-closed we can select $G[i+1, i^*)$ to be $\mathbb{P}[i+1, i^*)^{G(\leq i)}$-generic over $L[G(\leq i)]$ (the collection of dense sets that must be met is the countable union of subcollections of size i in $L[G(\leq i)]$, using the $\mathrm{Hull}(i \cup j_n)$'s as in the previous proof). Then $G(< i^*) = G(\leq i) * G[i+1, i^*)$ is $\mathbb{P}(< i^*)$-generic over L. We also choose $G(i^*)$ to be $\mathbb{P}(i^*)^{G(<i^*)}$-generic over $L[G(< i^*)]$, extending the condition $G(i)$ in this forcing.

For $i \in \mathrm{Lim}\ I$ take $G(< i)$ to be $\cup\{G(< j) | j \in I \cap i\}$, as in the previous proof $G(< i)$ is $\mathbb{P}(< i)$-generic over L. And we take $G(i) = \cup\{G(j) | j \in I \cap i\}$, which by our construction extends each $G(j)$, $j \in I \cap i$. Again we get genericity for $G(\leq i)$ from that of $G(\leq j)$, $j \in I \cap i$, as $G(< i)$, $G(i)$ extend $G(< j)$, $G(j)$ respectively for each $j \in I \cap i$. $\square$

But for $\#$-generation we want more.

Definition 41 *A class $A \subseteq L$* preserves indiscernibles *if I is a class of indiscernibles for the structure $\langle L[A], A \rangle$.*

Note that if G is $\mathbb{P}$-generic and preserves indiscernibles then $L[G]$ is $\#$-generated, with generating $\#$ equal to $L_{i_0}[G]$, where i_0 is the least indiscernible.

Theorem 42 *Assume that $0^\#$ exists. Then there is a generic over L for the reverse Easton iteration of α-Cohen which preserves indiscernibles.*

Proof. It suffices to build $H \subseteq L_{i_\omega}$ which is $\mathbb{P}(< i_\omega)$-generic over L_{i_ω} and such that $t(j_1 \dots j_n) \in H$ iff $t(j'_1 \dots j'_n) \in H$ whenever $j_1 < \dots < j_n$, $j'_1 < \dots < j'_n$ belong to $I \cap i_\omega$, $i_\omega = \omega^{\mathrm{th}}$ indiscernible. For then define $t(k_1 \dots k_n) \in G$ iff $t(i_1 \dots i_n) \in H$, $i_1 < \dots < i_n$ the first n indiscernibles. This is well-defined using the above property of H. And G is $\mathbb{P}$-generic over L: It suffices to consider predense $D \in L$ as $\mathbb{P}$ has the ∞-chain condition. Now write $D \in L$ as $s(l_1 \dots l_m)$, $l_1 < \dots < l_m$ in I, and then $\overline{D} = s(i_1 \dots i_m)$ is predense on $\mathbb{P}(< i_\omega)$. If $\bar{p} = t(i_1 \dots i_n) \in H$ meets $\overline{D}$ then $p = t(l_1 \dots l_m, l_{m+1} \dots l_n)$ meets D, where $l_m < l_{m+1} < \dots < l_n$ belong to I. Also $p \in G$ by definition of G. Finally, note that if $k_1 < \dots < k_m < l_1 < \dots < l_m$ and $l_1, \dots, l_m$ are in $\mathrm{Lim}\ I$, $k_1, \dots, k_m$ in I then for any φ, $\langle L[G], G \rangle \vDash \varphi(k_1 \dots k_m) \longleftrightarrow \varphi(l_1 \dots l_m)$ by the Truth Lemma and the fact that G obeys the same invariance property that characterized H. So I is a class of indiscernibles for $\langle L[G], G\rangle$.

Now we build H. Let $H_2 \subseteq \mathbb{P}(< i_2)$ be a $\mathbb{P}(< i_2)$-generic in $L[0^\#]$ and $H_1 = H_2 \cap \mathbb{P}(< i_1)$. We must now define $H_3 \subseteq \mathbb{P}(< i_3)$ to be $\mathbb{P}(< i_3)$-generic so that $t(i_1, \vec{j}) \in H_2$ iff $t(i_2, \vec{j}) \in H_3$, where $\vec{j}$ is an increasing sequence from $I - i_\omega$. Note that $H_2(i_1)$, a subset of i_1 generic over $L[H_1]$, is a condition in the i_2-Cohen forcing defined over $L[H_2]$; choose $H_3(i_2)$ to be a generic for this forcing extending $H_2(i_1)$. Now note that for each n there is $t_n(i_1, \vec{j}_n) = p_n \in H_2$ which reduces all predense $D \subseteq \mathbb{P}(< i_2)$ in $\mathrm{Hull}(i_1 \cup \{i_1, k_1 \dots k_n\})$ below i_1, where $i_\omega \le k_1 < \ldots < k_n$ belong to I, using the i_1^+-distributivity of $\mathbb{P}(> i_1)^{H_2(\le i_1)}$ in $L[H_2(\le i)]$. So if we define $H_3' = \{t_n(i_2, \vec{j}_n) | n \in \omega\}$ we have that H_3' reduces all predense $D \subseteq \mathbb{P}(< i_3)$, $D \in L$ below i_2. So the desired H_3 can be defined by $H_3 = \{p \in \mathbb{P}(< i_3) | p(\le i_2) \in H_3(\le i_2)$, p compatible with $H_3'\}$. By construction, $t(i_1, \vec{j}) \in H_2$ iff $t(i_2, \vec{j}) \in H_3$. Note that H_3 was uniquely determined by this last condition, once a choice of $H_3(i_2)$ was made.

H_4 is uniquely determined by $\mathbb{P}(< i_4)$-genericity and the condition $t(i_1, i_2, \vec{j}) \in H_3$ iff $t(i_2, i_3, \vec{j}) \in H_4$, as the forcing to add $H_3(i_2)$ is i_1^+-distributive (and the forcing to add $H_3(> i_2)$ is i_2^+-distributive). We must check that $t(i_1, i_3, \vec{j}) \in H_4$ iff $t(i_2, i_3, \vec{j}) \in H_4$. Now any condition in H_4 is extended by one of the form $p = (p_0, p_1)$ where $p_0 \in H_4(\le i_3)$ and $p_1 = t(i_3, \vec{j})$, as such p reduce all dense $D \subseteq \mathbb{P}(< i_4)$, $D \in L$ below i_3. So it suffices to show that $t(i_1, i_3, \vec{j}) \in H_4(\le i_3)$ iff $t(i_2, i_3, \vec{j}) \in H_4(\le i_3)$. By definition of H_4 we have $t(i_2, i_3, \vec{j}) \in H_4(\le i_3)$ iff $t(i_1, i_2, \vec{j}) \in H_3(\le i_2)$. But the latter implies that $t(i_1, i_2, \vec{j}) = t(i_1, i_3, \vec{j})$ and as $H_3(\le i_2)$ extends $H_2(\le i_1)$ we have that $H_4(\le i_3)$ extends $H_3(\le i_2)$. So $t(i_1, i_2, \vec{j}) \in H_3(\le i_2)$ iff $t(i_1, i_2, \vec{j}) \in H_4(\le i_3)$ iff $t(i_1, i_3, \vec{j}) \in H_4(\le i_3)$.

In general define H_{m+3} by the condition $t(i_m, i_{m+1}, \vec{j}) \in H_{m+2}$ iff $t(i_{m+1}, i_{m+2}, \vec{j}) \in H_{m+3}$. As above we get that H_{m+3} is $\mathbb{P}(< i_{m+3})$-generic and $t(i_1 \dots i_{m+1}, \vec{j}) \in H_{m+2}$ iff $t(i_1 \dots i_m, i_{m+2}, \vec{j}) \in H_{m+3}$. Finally let $H = \cup\{H_m | m \in \omega\}$. Then H is $\mathbb{P}(< i_\omega)$-generic over L and for any $k_1 < \ldots < k_{l+2} < \vec{j}$ in $I, k_{l+2} < i_\omega \le \vec{j}$ we have $t(k_1 \dots k_{l+1}, \vec{j}) \in H$ iff $t(k_1 \dots k_l, k_{l+2}, \vec{j}) \in H$. This is enough to imply that $t(\vec{k}_0) \in H$ iff $t(\vec{k}_1) \in H$ whenever $\vec{k}_0, \vec{k}_1$ are increasing sequences from $I \cap i_\omega$. □

18.,19.Vorlesungen

The MK Hyperuniverse

The Hyperuniverse we have been disucssing is the set of all countable transtiive models of ZFC. But one can also associate a Hyperuniverse to theories other than ZFC, such as the class theory MK. This is the theory with both sets and classes, where the sets obey ZFC plus replacement and comprehension for formulas which quantify over classes as well as global choice (the existence of a class which wellorders the sets). Then:

The MK-Hyperuniverse = the set of countable transitive models $(M, \mathcal{C})$ of MK

We've seen that a useful fact about the usual Hyperuniverse is the fact that any ZFC-universe M has a width-extension which is *minimal*, i.e. which is the smallest universe containing some real. Our aim now is to develop a similar result for MK-universes.

To obtain this result we need to develop a theory of forcing over MK-universes where the conditions are classes, and not sets, as so-called "hyperclass forcing". It turns out that the most effective way of doing this is to associate to an MK-universe an associated "companion" model of $\mathrm{ZFC}^- = \mathrm{ZFC} \setminus$ Powerset.

If $(M, \mathcal{C})$ is an MK-universe then we associate to it the transitive set M^+ consisting of the union of all transitive sets "coded" by a class in $\mathcal{C}$. Without going into the details of this "coding", a transitive set t is coded by a class $T \in \mathcal{C}$ if T is a wellfounded tree on a subclass of M which is isomorphic to the transitive closure of $\{t\}$ with the $\in$-relation, once nodes with isomorphic subtrees below them are identified with each other. For this coding it is useful to assume that $(M, \mathcal{C})$ is a *β-model*, which means that any relation in C which appears wellfounded in $(M, \mathcal{C})$ is in fact wellfounded in the real world.

It is our wish that the "companion" model M^+ be a model of ZFC^-. For this we need to assume that $(M, \mathcal{C})$ satisfies more than MK, namely the theory MK^*, which adds to MK the scheme of *class-bounding*:

If for each set x there is a class Y such that $\varphi(x, Y)$ then there is a single class Z such that for all x there is a y such that $\varphi(x, (Z)_y)$

where φ can be second-order with class parameters and $(Z)_y$, the "y-th slice of Z" is the set of z such that $(y, z) \in Z$. Using Global Choice it is not hard to show that class-bounding is equivalent over MK to *class-choice*, which says:

If for each set x there is a class Y such that $\varphi(x, Y)$ then there is a single class Z such that for all x, $\varphi(x, (Z)_x)$.

Now we have a nice way of translating between β-models of the second-order theory MK* and models of a first-order theory. The axioms of SetMK* are:

ZFC$^-$ (including the Bounding Principle)
There is a strongly inaccessible cardinal κ
κ is the largest cardinal (i.e. every set can be mapped injectively into κ)

Theorem 43 *(a) Suppose that $(M, \mathcal{C})$ is a β-model of MK*. Then M^+ as defined above is a transitive model of $SetMK^*$ such that if κ is the largest cardinal of M^+ then $M = V_\kappa^{M^+}$ and $\mathcal{C}$ consists of the subsets of M in M^+.*
(b) Suppose that M^+ is a transitive model of SetMK* with largest cardinal κ. Then $(M, \mathcal{C})$ is a β-model of MK* where $M = V_\kappa^{M^+}$ and $\mathcal{C}$ consists of all subsets of κ in M^+.
(c) The above transformations $(M, \mathcal{C}) \mapsto \mathcal{M}^+$ and $M^+ \mapsto (M, \mathcal{C})$ are inverses to each other.

Definable Hyperclass Forcing and MK**

To width-expand a β-model $(M, \mathcal{C})$ of MK to a minimal one (i.e. the least β-model of MK containing some real) requires use of an $(M, \mathcal{C})$-definable forcing whose conditions are classes. Our strategy is to replace $(M, \mathcal{C})$ by M^+ and then view this forcing as an M^+-definable class-forcing, applying techniques of class-forcing to the ZFC$^-$ model M^+. But there is one more remaining difficulty here, which can be illustrated by considering the following example:

Example. Suppose that M satisfies ZFC$^-$ and consider the class forcing $\mathbb{P}$ in M whose conditions are functions from an ordinal to 2, ordered by extension. Suppose that G is $\mathbb{P}$-generic over M. Then do M and $M[G]$ have the same sets, i.e. is $\mathbb{P}$ ∞-distributive for definable sequences of dense classes?

The answer would appear to be "yes", as the forcing $\mathbb{P}$ is clearly ∞-closed in M. But the difficulty in extending a given condition p to meet even ω-many dense classes $(D_n \mid n < \omega)$ is the need for a suitable form of *dependent choice* to choose $p = p_0 \geq p_1 \geq \cdots$ where p_{n+1} meets D_n.

Given a β-model $(M, \mathcal{C})$ of MK* we would like to apply the above forcing to the companion transitive model M^+ of ZFC$^-$; but to successfully do so we would like M^+ to satisfy (for first-order φ):

Definable κ-DC: If for all x there is a y such that $\varphi(x, y)$ then for any x there is a function f with domain κ such that $f(0) = x$ and for all $i > 0$, $\varphi(f \restriction i, f(i))$.

To obtain the κ-DC in M^+ in turn requires us to strengthen the axioms MK* by adding (for second-order φ):

Definable DC for Classes: If for each class X there is a class Y such that $\varphi(X, Y)$ then for any class X there exists a class Z such that $(Z)_0 = X$ and for all ordinals $i > 0$, $\varphi(Z \restriction i, (Z)_i)$,

where as before $(Z)_i = \{x \mid (i, x) \in Z\}$ and $Z \restriction i = \{(j, x) \mid j < i$ and $(j, x) \in Z\}$. The theory MK** is MK* together with *DC for Classes.* It is easy to verify that if $(M, \mathcal{C})$ is a β-model of MK** then the associated M^+ satisfies SetMK** and conversely, if starting with M^+ satisfying SetMK** we derive a β-model $(M, \mathcal{C})$ of MK* then in fact this latter model satisfies MK**.

OK, so now we are finally ready to start forcing over M^+. Let κ^* denote the ordinal height of M^+.

Step 1. We force over M^+ to get a SetMK**-model of the form $L_{\kappa^*}[A]$ where A is a subset of κ^*.

The forcing to produce A is simply the forcing mentioned earlier to add a Cohen class A of ordinals to M^+. The assumption of κ-DC is needed to show that this forcing is definably distributive and therefore does not add new sets. By genericity any set of ordinals in M^+ appears in $L_{\kappa^*}[A]$. We still have a model of SetMK^{++} with largest cardinal κ.

An important point however is the definability of the forcing relation, which follows from the special nature of the forcing: To determine if p forces $\sigma \in \tau$ for two names σ and τ, we simply see if $\sigma^q \in \tau^q$ for all extensions q of p of length greater than the ranks of σ and τ. In other words, forcing equals truth for atomic sentences and long enough conditions.

Step 2. By forcing we "reshape" A into another A', without adding sets, so that for no α between κ and κ^* does $(L_\alpha[A'], A' \cap \alpha)$ satisfy ZFC$^-$.

A condition is an initial segment of such an A' of length less than κ^*. One must show that this forcing is definably distributive, which is a special argument using sufficiently elementary submodels.

Step 3. We code A' into a subset X of κ, so that M^+ is now contained in $L_{\kappa^*}[X]$ and there is no α between κ and κ^* such that $L_\alpha[X]$ satisfies ZFC$^-$.

This is almost disjoint coding. The fact that A' is "reshaped" makes this possible. The forcing has small definable antichains (they are all sets) and the proof that the forcing relation is definable for "pretame" class forcings over models of ZFC can be adapted here, replacing use of the V-hierarchy by the $(L[A'], A')$-hierarchy.

Step 4. We add a club C of strong limit cardinals less than κ such that if $\bar{\kappa}$ belongs to C then there is no model of ZFC$^-$ of the form $L_\alpha[X \cap \bar{\kappa}]$ in which $\bar{\kappa}$ is strongly inaccessible.

This uses the fact that since there is no ZFC$^-$ model $L_\alpha[X]$ with $\alpha < \kappa^*$ in which κ is strongly inaccessible, the set of $\bar{\kappa}$ as above is "fat-stationary". So we are shooting a club through a fat-stationary set.

Step 5. We arrange that all limit cardinals less than κ belong to C using an Easton product of collapse forcings.

Step 6. We apply Jensen coding to get a model $L_{\kappa^*}[x]$ for some real x in which ZFC$^-$ holds, κ is strongly inaccessible and for no cardinal $\bar{\kappa} \leq \kappa$ is there an $\alpha < \kappa^*$ such that $L_\alpha[x]$ satisfies ZFC$^-$ and in which $\bar{\kappa}$ is strongly inaccessible. We still have a model of ZFC$^-$ thanks to the good behaviour of Jensen coding.

Step 7. Finally, we add a real y which ensures the above property of x not only at each *cardinal* less than or equal to κ but also at each *ordinal* less than or equal to κ, using a method due to David and myself.

That does it: Now we have a real y such that M^+ is contained in $L_{\kappa^*}[y]$ and the latter is the least transitive model of SetMK containing y. This is also a model of $SetMK^{**}$ and it follows that the MK**-model derived from $L_{\kappa^*}[y]$ is the least β-model of MK containing y.

20.Vorlesung

Minimality in the GB Hyperuniverse

Theorem 44 *Suppose that $(M,\mathcal{C})$ is a countable model of GB (Gödel-Bernays class theory). Then $(M,\mathcal{C})$ has an extension $(M^*,\mathcal{C}^*)$ with the same ordinals which for some real x is the smallest transitive model of GB containing x.*

Proof. The proof is unusual in that in the first step we force over a very "bad" ground model to code the elements of $\mathcal{C}$ into a single class, preserving GB. The proof then finishes in a standard way by applying the variant of Jensen coding needed to create minimal universes.

List the elements of $\mathcal{C}$ as $A_0, A_1, \ldots$, in an ω-sequence. We associate clubs $C_0, C_1, \ldots$ to this sequence as follows:

$C_0 =$ the club of $\alpha < \mathrm{Ord}(M)$ such that V_α^M is Σ_1 elementary in M.

$C_1 =$ the club of $\alpha < \mathrm{Ord}(M)$ such that V_α^M is Σ_2 elementary in M relative to the predicate (A_0, C_0).

In general, $C_{n+1} =$ the club of $\alpha < \mathrm{Ord}(M)$ such that V_α^M is Σ_{n+2} elementary in M relative to the predicate $(A_0, A_1, \ldots, A_n, C_0, \ldots, C_n)$.

And for each n let C_n^+ denote the successor elements of C_n.

Our goal is to force a class function $F : \mathrm{Ord}(M) \to 2$ which codes A_n at sufficiently large elements of C_n^+ and preserves ZFC over M. More precisely:

1. For each n and sufficiently large $i < \mathrm{Ord}(M)$, i belongs to A_n iff the value of F at the i-th element of C_n^+ is equal to 1.

2. (M, A) satisfies ZFC.

We take F to be generic over $(M, A_0, A_1, \ldots)$ for the forcing $\mathbb{P}$ consisting of pairs (p, n) where $|p|$ is an ordinal $< \mathrm{Ord}(M)$, $p : |p| \to 2$ and $n \in \omega$. When extending (p, n) to (q, k) we require that k is at least n and condition 1 above holds at all ordinals α which belong both to some C_m^+, $m \leq n$ and to the domain of $q \setminus p$. Note that the ground model $(M, A_0, A_1, \ldots)$ may not satisfy ZFC. However we force over this ground model with $\mathbb{P}$ anyway. The result is a function $F : \mathrm{Ord}(M) \to 2$ such that each A_n is definable over (M, F) with parameters. But we need a special argument for the preservation of ZFC when adding F to M.

Lemma 45 *For each n let $\mathbb{P}_n$ be the forcing consisting of conditions (p, n) in $\mathbb{P}$ with second coordinate n. Then any predense sublcass D of $\mathbb{P}_n$ which is Σ_n-definable over $(M, A_0, \ldots, A_n, C_0, \ldots, C_n)$ is also predense in $\mathbb{P}$.*

Proof. We want to extend a given $(p, k) \in \mathbb{P}$ below some (q, n) in D. We may sssume that k is at least n. Extend p so that $|p|$ is a limit point of C_k. Then extend (p, n) to a (q, n) extending an element of D of least possible length. As D is Σ_n-definable over $(M, A_0, \ldots, A_n, C_0, \ldots, C_n)$, the length of q is less than the least element of C_{n+1} greater than $|p|$. It follows that (q, k) extends both (p, k) and $(q, n) \in D$. $\Box$

By the Lemma, our $\mathbb{P}$-generic F is also $\mathbb{P}_n$-generic over the ground model $(M, A_0, \ldots, A_n, C_0, \ldots, C_n)$ for Σ_n-definable dense classes. As the forcing $\mathbb{P}_n$ preserves full replacement over this ground model (and $\mathbb{P}_n$ is Δ_1-definable over this ground model, adding no new sets) it follows that F preserves Σ_n replacement over $(M, A_0, \ldots, A_n, C_0, \ldots, C_n)$ and therefore Σ_n replacement over M, for each n. Thus F preserves ZFC over M.

So we have enlarged $(M, \mathcal{C})$ to a GB-model $(M, \mathcal{C}^*)$ where $\mathcal{C}^*$ consists just of the classes definable over (M, F). Now to finish the proof, apply Jensen coding to enlarge this further to a GB-model of the form $(L^M[x_0], \mathcal{C}^{**})$ where x_0 is a real and $\mathcal{C}^{**}$ are the classes definable over $L^M[x_0]$. Then apply the Beller-David result to enlarge $L^M[x_0]$ one last time to the least transitive GB-model conataining some real x. $\Box$

Remark. By further forcing we can enlarge the minimal model $L^M[x]$ above to a pointwise-definable model, by forcing $V = \mathrm{HOD}$ using a generic iteration and coding x into the GCH pattern on the $\aleph_n$'s. In this model every set is definable from ordinals, every ordinal is definable from x and x is definable. But it is not known if one can enlarge to a pointwise-definable model also satisfying $V = L[x]$ for some real x.