Self-Constructing Continua

Doug McLellan — 2020 Draft

Abstract

Recent interest in "potentialism," which conceives mathematics as growing over time, suggests that we try formalizing the idea of a mathematical universe that constructs itself in a gradual, organic way, with neither "work from outside," nor any "background model of set theory," nor any "inexplicable novelty" involved in the course of its growth. Taking *continuum* to mean the set of all reals in a standard inner model of set theory, we axiomatize the assertion that a family $\mathcal{F}$ of distinct continua constructs itself in such a way; we then present a forcing construction that outputs candidates $\mathcal{F}(B, \zeta, G)$ to satisfy our axioms, where B is a boolean algebra, ζ is a function determining which of B's subalgebras contribute continua to the family, and G is a generic ultrafilter on B. We derive a key necessary condition on B for $\mathcal{F}(B, \zeta, G)$ to satisfy the axioms. We show that Cohen and random-real algebras violate this condition, essentially because they have too many automorphisms. Rigid boolean algebras avoid this problem, but the example we give fails to yield a model of the axioms for other reasons. The most promising candidate B is a noncommutative analog of random-real forcing, which we discuss briefly.

1 Self-Construction Axioms

By the *constructive closure* $\mathbb{R}(X)$ of a set X of real numbers we mean, informally, all the real numbers that must exist given that X does; formally, we define $\mathbb{R}(X)$ for any set X as the set of all reals in $L(X)$, the constructible hierarchy seeded with X's transitive closure and then built up with definable operations using Gödel's well-known procedure. For background on $L(X)$—which underlies everything that follows—see the appendix, Section 12.

By a *continuum* we mean a set X of real numbers (henceforth construed as subsets of ω) that is constructively closed, meaning $\mathbb{R}(X) = X$. Note that $L(X)$ is the smallest standard inner ZF model having X as a member (see Lemma 3.2) and it follows that we could define a continuum equivalently as the set of all real numbers in some standard inner ZF model.

$L(\emptyset)$, usually just denoted L, is the constructible universe of sets. Its continuum $\mathbb{R}(\emptyset)$ is the least continuum. Gödel famously proved it consistent with ZF that this is the *only* continuum.

It is also consistent with ZF that we can arrange distinct continua so that each is the constructive closure of the smaller ones' union, as in Example 2.1 below. Such an arrangement suggests the real numbers being constructed out of each other—growing over time, as it were—or, if the continua are not *linearly* ordered by inclusion, then growing over spacetime, as *it* were.

The goal of this paper is to determine whether this growth metaphor can be pushed further, whether a set of continua might grow "organically" without the need for anyone to do any preliminary "arranging." From this intuitive idea we will develop two axioms whose satisfaction would entitle a set of continua to be deemed *self-constructing*. However, as this development is necessarily somewhat discursive, we prefer to remove it to an appendix (Section 11) and to take the two axioms as our starting point. To state them we must first define "self-collection."

Definitions. Recalling that a set $\mathcal{N}$ is *directed* (with respect to the inclusion ordering) if $X, Y \in \mathcal{N} \Rightarrow (\exists Z \in \mathcal{N})(X, Y \subseteq Z)$, we say that a set $\mathcal{N}$ of continua *self-collects into* X if the following hold:

(i) $X \notin \mathcal{N}$;

(ii) $\mathcal{N}$ is directed;

(iii) $\mathcal{N} \in L(X)$;

(iv) $(\neg\exists x \in X)(\bigcup\mathcal{N} \subseteq \mathbb{R}(x) \subset X)$;

(v) $X = \mathbb{R}(\bigcup\mathcal{N})$.

Axioms for a Self-Constructing Set $\mathcal{F}$ of Continua:

Self-Collection: $(\forall X \in L(\mathcal{F}))(X \in \mathcal{F} \iff \text{some } \mathcal{N} \subseteq \mathcal{F} \text{ self-collects into } X)$.

Foundation: $(\forall X \in \mathcal{F})(\forall x \in X)(\exists Y \in \mathcal{F})$

$$(x \in Y \subseteq X \text{ and } (\forall Z \in \mathcal{F})(Z \subset Y \Rightarrow x \notin Z)).$$

As motivation for calling these the "self-construction" axioms, let us content ourselves for now to note that Self-Collection captures an idea of continua naturally forming a collection that requires vanishingly little work to constructively close; and that Foundation keeps real numbers from slipping in inexplicably or unaccountably during $\mathcal{F}$'s growth. Again, the full motivation is developed in Section 11.

It is easy to check that $\mathcal{F} = \{\mathbb{R}(\emptyset)\}$ is a self-constructing family. We turn now to the question of whether any non-trivial $\mathcal{F}$ can satisfy our two axioms.

2 Forcing Singly-Generated Continua

The technique of forcing is uniquely suited to produce candidates for satisfying the self-construction axioms. To illustrate the challenges involved in using it successfully to this end, we introduce two relatively simple examples from a Cohen-forcing extension, which we call $\mathcal{C}_S$ and $\mathcal{C}_{\neg S}$.

Definition. A *singly-generated* continuum X is one that satisfies $X = \mathbb{R}(x)$ for some real number x (which is clearly never unique).

Example 2.1 *Let G be a generic filter for Cohen forcing (see Section 9.2); let $\mathcal{C}_S$ be the set of singly-generated continua in $L[G]$; let $\mathcal{C}_{\neg S}$ be the set whose members are $\mathbb{R}(\emptyset)$ and all non-singly-generated continua in $L[G]$.*

Remark. Let us pause to stress that the relative constructibility used in our definitions has been the "parentheses" version ([6], Definition 13.24). A forcing extension $L[G]$ is defined differently, as the class of interpretations that G gives to specified forcing names in L. This generalizes to a "square-brackets" version of relative constructibility ([6], Definition 13.22) which guarantees for arbitrary sets Z that $L[Z]$ satisfies the Axiom of Choice (AC), but not that Z is actually a member of $L[Z]$; for our purposes, this would not have been to the point. The square-brackets version is, however, the standard one to use when discussing forcing extensions, as we will throughout this paper, so it is well worth pointing out that in most of the cases we will consider, $L(Z)$ and $L[Z]$ come to the same thing:

Fact (Lemma 12.5): $L(Z) = L[Z]$ for any set $Z \subseteq L$; this holds notably when Z is a single real number, or any filter on a constructible poset. $\square$

Returning to our examples $\mathcal{C}_S$ and $\mathcal{C}_{\neg S}$, we cite the following facts, which are discussed in the appendix (Section 11):

- Both $\mathcal{C}_S$ and $\mathcal{C}_{\neg S}$ have the property that each of their members is the constructive closure of their smaller members' union.
- $\mathcal{C}_S$ satisfies our Foundation axiom but not Self-Collection.
- $\mathcal{C}_{\neg S}$ satisfies our Self-Collection axiom but not Foundation.

Given these two concrete examples, a natural strategy for forcing a self-constructing set of continua would be to choose one of them, and deduce what modifications would be neeeded to make it satisfy the axiom it does not already satisfy. We commit now to pursue this strategy with $\mathcal{C}_S$, first because we simply have no idea how to begin patching up $C_{\neg S}$'s failure to satisfy our Foundation axiom, and second, because the singly-generated continua in a forcing extension $L[G]$ have a useful simplifying feature. Namely, they are the continua of form $\mathbb{R}(G \cap D)$, where D is a countably-completely-generated subalgebra of the boolean algebra on which G is a filter. (This is proved in Lemma 3.8; it is closely related to Lemma 15.43 of [6], the well-known "intermediate model theorem.") We will ensure by fiat that this simplifying feature is always available to us:

Stipulation 2.2 *Henceforth every family $\mathcal{F}$ we consider as a candidate to satisfy the self-construction axioms will belong to some forcing extension $L[G]$, and each $\mathcal{F}$-member will be singly generated.*

Our central question thus becomes how a set of continua under Stipulation 2.2 can satisfy the Self-Collection axiom. In particular we ask: If $L[G]$ is a generic extension of L by a filter G on boolean algebra B, what are the conditions on B, and on a B-name $\dot{\mathcal{N}}$ for a set of singly-generated continua, under which G's interpretation of $\dot{\mathcal{N}}$ will self-collect into $\mathbb{R}(G)$?

3 Main Definitions, Facts, Questions

We now turn to the definitions needed to answer the question just stated, and also to pose tighter versions of it, incorporating some simplifying suppositions about B and $\dot{\mathcal{N}}$. We will register these versions as our Main Question at the end of this section.

3.1 Definitions and Notation

$L(y)$ is the union of all levels $L_\alpha(y)$ of the (relative) constructive hierarchy seeded with $\{y\}$'s transitive closure (that is, $L_0(y) \equiv \mathsf{TrCl}(\{y\})$); see Section 12 for details and background.

When $x \in L(y)$ we say that y *constructs* x. If also $y \in L(x)$, we say x and y are *interconstructible*.

We will not restate the basic terminology of forcing. We assume familiarity with the boolean-algebraic approach to forcing used in T. Jech's standard text [6], as well as in J. Bell's text [2]. We side with [2] in preferring the symbols $\wedge$ and $\vee$ for boolean meet and join, whereas [6] generally uses $\cdot$ and $+$.

In what follows, B is always an atomless complete boolean algebra in L, the constructible universe. We should say that the structure $\langle B, 0, 1, \vee, \wedge, \neg\rangle$ is our boolean algebra and that B is just its underlying set, but we will conflate the two, as is customary. We also assume B is countably completely generated (see the definition just below).

An *ACSA* of B is an atomless complete subalgebra of B. Switching "atomless" and "complete" would yield the more attractive acronym "CASA", but "complete subalgebra of B" is a syntactic unit: it denotes not just any subalgebra C of B that is complete, but only those C wherein the supremum of any subset of C is the same whether calculated in C or in B, and likewise for infima. See [6], chapter 7, under "Complete and Regular Subalgebras."

If $X \subseteq B$, X *completely generates* the subalgebra $\overline{X}$ that is the smallest complete subalgebra of B that includes X.

$\mathsf{ACSAs}(B)$ denotes $\{C \in L : C$ is a countably-completely-generated ACSA of $B\}$, where by *countably-completely-generated* we mean that $C = \overline{Q}$ for some countable $Q \subseteq C$. The restriction to such subalgebras is what will enforce Stipulation 2.2; see Lemma 3.8. As mentioned, we assume that B itself is countably-completely-generated.

C, D, E will always range over $\mathsf{ACSAs}(B)$.

B^+ means $B \setminus \{0\}$, that is, B without its least element.

G is an L-generic ultrafilter on B (equivalently, L-generic filter on B^+).

$C \leq_G D$ means $(G \cap D)$ constructs $(G \cap C)$, that is, $G \cap C \in L(G \cap D)$.

Recall that when Ψ is an assertion in B's forcing language (with L our ground model), Ψ's "boolean value" or "B-value", denoted $\|\Psi\|_B$ (or just $\|\Psi\|$ when B is understood from context) is the (unique) element of B that will be in G (for any generic G) if and only if Ψ is true in $L[G]$.

3.2 Main Facts About Relative Constructibility

Constructibility, relative and otherwise, is reviewed in more detail in the appendix, Section 12. The present section states some fundamental results, referring to that appendix for the proofs, and then derives from these results some basic facts about the interconstructibility of reals, continua, and generic filters.

Definition. A *standard* ZF model is one in which all elements are pure sets and "$\in$" is interpreted as the real membership relation among its elements. An *inner* ZF model is one that is transitive and has all the ordinals.

Definition. When y is any set, a *y-predicate* is a proposition $\Psi(x, c_1, \ldots, c_n)$ of the language of set theory, having one free variable x and finitely many constant symbols $c_1, \ldots, c_n$, considered together with fixed referents for the c_i chosen out of $L \cup \mathsf{TrCl}(\{y\})$. A y-predicate $\Psi(x, c_1, \ldots, c_n)$ is *y-absolute* if

$$\{x : \Psi(x, c_1, \ldots, c_n)\}$$

evaluates to the same set in any standard inner ZF model having y as a member.

Lemma 3.1 *There is a proposition Ψ such that $\Psi(x, \alpha, y)$ is, for any ordinal α and set y, a y-absolute predicate satisfying $\{x : \Psi(x, \alpha, y)\} = L_\alpha(y)$.*

The proof is discussed in the appendix (Section 12). $\Box$

Lemma 3.2 *$L(y)$ is the smallest standard inner ZF model having y as a member, in the sense that $L(y) \subseteq M$ for any other such model M.*

That $L(y)$ is a standard inner ZF model follows from a minor modification of the usual proof that L is such a model (see [6], Theorem 13.3, or [3], Theorem 1.2). The minimality claim for $L(y)$ follows from Lemma 3.1: If $L(y) \not\subseteq M$ then $L_\alpha(y) \notin M$ for some α, contradicting $L_\alpha(y)$'s definability by a y-absolute predicate. $\Box$.

Corollary 3.3 *The following are all easy consequences of Lemma 3.2:*

(a) If x constructs y (that is, $y \in L(x)$), then $L(y) \subseteq L(x)$.

(b) The relation "x constructs y" ($y \in L(x)$) is transitive.

(c) If x and y are interconstructible, then $L(x) = L(y)$ and $\mathbb{R}(x) = \mathbb{R}(y)$.

(d) The ordering $\leq_G$ on $\mathsf{ACSAs}(B)$ defined above (namely $C \leq_G D$ if and only if $G \cap C \in L(G \cap D)$) is guaranteed with boolean value 1 to be transitive; since it is also reflexive it is a preorder, so we may define $=_G$, $<_G$, etc. accordingly. □

It is a central fact about $L(y)$ that each of its members (not only each level $L_\alpha(y)$) is definable by a y-absolute predicate:

Lemma 3.4 *$z \in L(y)$ if and only if $z = \{x : \Psi(x, c_1, \ldots, c_n)\}$ for some y-absolute predicate $\Psi(x, c_1, \ldots, c_n)$.*

The proof is discussed in the appendix (Section 12). □

We now bring these central facts about relative constructibility to bear on our forcing construction (the generic filter G on the boolean algebra B).

Lemma 3.5 *For all (constructible) $X \subseteq B$, $G \cap \overline{X} \in L(G \cap X)$.*

By Lemma 3.4, $G \cap \overline{X} \in L(G \cap X)$ if and only if $G \cap \overline{X}$ is definable by a $(G \cap X)$-absolute predicate, which is what we will show.

$\overline{X}$, the boolean completion of X in B, was defined above as the smallest complete subalgebra of B containing X. The key to our proof is an alternative definition of $\overline{X}$ as the union of a transfinite hierarchy.

When $Y \subseteq B$, define $\phi(Y)$ to be Y's closure under arbitrary negations and joins in B, or more precisely:

$$\phi(Y) \equiv \{\neg z : z \in Y\} \cup \{\bigvee S : S \subseteq Y \text{ and } S \in L\}.$$

Now define a hierarchy by successive closures under ϕ:

$X_0 \equiv X$;

$X_{\alpha+1} \equiv \phi(X_\alpha)$;

$X_\alpha \equiv \bigcup_{\beta<\alpha} X_\beta$ when α is a limit ordinal.

It is easily checked that the union of all the X_α is the smallest complete subalgebra of B that includes X. In fact the growth of the X_α's must cease at some ordinal δ, lest $\overline{X}$'s cardinality exceed B's, so we may say that $\overline{X}$ is the union of the X_α's with α less than this δ. The function $\Xi : \alpha \mapsto X_\alpha$, with domain $\delta + 1$, is a member of L. (Lemmas 12.2 and 12.4 can be used to establish $\Xi \in L(\emptyset)$ rigorously.) Note that $\Xi(0) = X$ and $\Xi(\delta) = \overline{X}$.

Now for all $\alpha \le \delta$ define $\Gamma(\alpha) \equiv G \cap \Xi(\alpha)$. Plainly we have $\Gamma(0) = G \cap X$ and $\Gamma(\delta) = G \cap \overline{X}$. Because G is a generic ultrafilter on B, we will have for all $\alpha < \delta$:

$$\Gamma(\alpha+1) = \{\, b \in \Xi(\alpha+1) : \neg b \in \Xi(\alpha) \setminus \Gamma(\alpha), \text{ or } (\exists b' \in \Gamma(\alpha))(b' \le b) \,\}.$$

(We have used here the fact that if $b = \bigvee S$ for some $S \in L$, $S \subseteq \Xi(\alpha)$, then those elements of B^+ that are $\le$ some S-member form a subset of B^+ that is dense below b; hence by G's genericity, if $b \in G$, some S-member must be in $\Gamma(\alpha)$.)

Note now that $\Gamma(\alpha)$ is the union of previous $\Gamma(\beta)$'s when α is a limit ordinal; and that because the expresssion for $\Gamma(\alpha+1)$ above does not involve (either implicitly or explicitly) any unconstructible sets other than $\Gamma(\alpha)$, there is a proposition Ψ such that for all α, $\Psi(x, \Gamma(\alpha), \Xi, \alpha)$ is a $\Gamma(\alpha)$-absolute predicate that defines $\Gamma(\alpha+1)$. Lemmas 12.2 and 12.4 then ensure that $\Gamma(\delta)$ can be defined by a $\Gamma(0)$-absolute predicate, and since $\Gamma(\delta) = G \cap \overline{X}$ and $\Gamma(0) = G \cap X$, we are done. $\Box$.

Lemma 3.6 *Every real number $x \in L[G] \setminus L$ is interconstructible with $G \cap D$ for some $D \in \mathsf{ACSAs}(B)$, and vice-versa.*

Fix $D \in \mathsf{ACSAs}(B)$ and (since the definition of $\mathsf{ACSAs}(B)$ requires D to be countably completely generated) fix $\{d_n : n \in \omega\} \subseteq D$ that completely generates D. By Lemma 3.5, $G \cap D \in L(G \cap \{d_n : n \in \omega\})$. If we define the real number x by $\{n : d_n \in G\}$, then x is clearly interconstructible with $G \cap \{d_n : n \in \omega\}$, so by Corollary 3.3 we have $G \cap D \in L(x)$. (And clearly $x \in L(G \cap D)$ since $x = \{n : d_n \in G \cap D\}$.)

Conversely, fix an unconstructible real number $x \subseteq \omega$. Let $\dot{x}$ be a B-name for x such that $1 \Vdash$ "$\dot{x} \subset \omega$ is unconstructible", and let D be the subalgebra of B completely generated by $\{\|n \in \dot{x}\| : n \in \omega\}$. D will be atomless because an atom of D would force for all n either $n \in \dot{x}$ or $n \notin \dot{x}$, thus forcing x to be a member of the ground model L, contradicting the guarantee that $\dot{x}$ names an unconstructible real. As in the previous paragraph, x will be interconstructible with $D \cap G$. $\Box$

Lemma 3.7 *A real number x is interconstructible with its $\mathbb{R}(x)$.*

This is clear since $x \in \mathbb{R}(x) \subseteq L(\mathbb{R}(x))$, and $\mathbb{R}(x) \in L(x)$ by definition. $\Box$

Lemma 3.8 *The singly-generated continua in a forcing extension $L[G]$ are precisely the sets of form $\mathbb{R}(G \cap D)$, where $D \in \mathsf{ACSAs}(B)$ or $D = \emptyset$.*

$\mathbb{R}(\emptyset)$ is obviously singly-generated since we are construing real numbers as subsets of ω and $\emptyset$ is such a subset.

Fix $D \in \mathsf{ACSAs}(B)$. By Lemma 3.6, $D \cap G$ is interconstructible with some $x \subseteq \omega$; and by Corollary 3.3 (c), $\mathbb{R}(G \cap D)$ equals $\mathbb{R}(x)$, which is a singly-generated continuum. Conversely, fix an unconstructible singly-generated continuum $X \in L[G]$, and real number $x \subseteq \omega$ such that $X = \mathbb{R}(x)$. By the other direction of Lemma 3.6, x is interconstructible with some $G \cap D$, so again we will have $\mathbb{R}(x) = \mathbb{R}(G \cap D)$. $\Box$

Corollary 3.9 *G is interconstructible with $\mathbb{R}(G)$ (since we assume B is countably completely generated).*

Our assumption that B is itself countably-completely-generated entails $B \in \mathsf{ACSAs}(B)$, so by Lemma 3.6 G is interconstructible with some real number x. By Lemma 3.7, x is interconstructible with $\mathbb{R}(x)$. By transitivity, G is interconstructible with $\mathbb{R}(x)$. By Corollary 3.3, $\mathbb{R}(G) = \mathbb{R}(\mathbb{R}(x)) = \mathbb{R}(x)$. The claim follows. □

Corollary 3.10 (General form of candidate families) *Stipulation 2.2 is equivalent to the assertion that every family of continua we will consider as a candidate to self-construct will have form*

$$\mathcal{F}(B, \zeta, G) \equiv \mathbb{R}(\emptyset) \cup \{\mathbb{R}(G \cap C) : C \in \mathsf{ACSAs}(B) \text{ and } \zeta(C) \in G\},$$

where ζ is some mapping $\zeta : \mathsf{ACSAs}(B) \to B$. □

3.3 The Main Question

Plenty of forcing extensions $L[G]$ (including the Cohen-forcing extension in Example 2.1) have instances of $\mathcal{N}$ meeting requirements (i)-(iv) in the definition of self-collection into $\mathbb{R}(G)$. It is requirement (v), $\mathbb{R}(\bigcup \mathcal{N}) = \mathbb{R}(G)$, that is hard to satisfy; in order to focus on it we assume the following:

Stipulation 3.11 *$\mathcal{N}$ is a set of singly-generated continua meeting requirements (i)-(iv) in the definition of self-collecting into $\mathbb{R}(G)$, as well as Stipulation 2.2; furthermore, assume we have a name $\dot{\mathcal{N}}$ forced by $1 \in B$ to meet the foregoing.*

The question posed at the end of Section 2 now takes this form: Under what conditions will $\mathcal{N}$ meeting Stipulation 3.11 satisfy the final requirement for self-collection, $\mathbb{R}(\bigcup \mathcal{N}) = \mathbb{R}(G)$? This final requirement is, in light of the results of Section 3.2, equivalent to the following more perspicuous one:

Lemma 3.12 $\mathbb{R}(\bigcup \mathcal{N}) = \mathbb{R}(G) \iff G \in L(\bigcup \mathcal{N})$ *(under Stip. 3.11).*

Clause (iii) of Stipulation 3.11 is $\mathcal{N} \in L(\mathbb{R}(G))$. By Corollary 3.9 and Corollary 3.3 (c), $L(G) = L(\mathbb{R}(G))$. Thus $\mathcal{N} \in L(G)$, and since $L(G)$ satisfies the ZF axiom of union, $\bigcup \mathcal{N} \in L(G)$. If also $G \in L(\bigcup \mathcal{N})$, then G is interconstructible with $\bigcup \mathcal{N}$, and by Corollary 3.3 (c) we have $\mathbb{R}(\bigcup \mathcal{N}) = \mathbb{R}(G)$.

Conversely if $\mathbb{R}(\bigcup \mathcal{N}) = \mathbb{R}(G)$ then $\mathbb{R}(G) \in L(\bigcup \mathcal{N})$, whence by Corollary 3.9 and transitivity (Corollary 3.3 (b)) $G \in L(\bigcup \mathcal{N})$. □

This equivalence allows us to officially register our Main Question thus:

The Main Question ($\mathcal{N}$ Version, Forcing-Extension Perspective): *Assuming Stipulation 3.11, under what conditions does $G \in L(\bigcup \mathcal{N})$ hold (so that $\mathcal{N}$ fulfills the final requirement (v) for self-collection)?*

We will soon (Section 4) introduce a set Θ that is interconstructible with and more manageable than $\bigcup \mathcal{N}$, and we will find requirements on Θ equivalent to those of Stipulation 3.11 on $\mathcal{N}$. For the sake of keeping all versions of our Main Question in one place, we will register the second one now.

The Main Question (Θ Version, Forcing-Extension Perspective): *Assuming Stipulation 3.11, under what conditions does $G \in L(\Theta)$ hold?*

We will answer this version of the question in Theorem 7.8.

Note now that the question "Under what conditions does $\mathcal{N}$ self-collect into $\mathbb{R}(G)$?" is posed from the perspective of the forcing extension $L[G]$, and that we would ultimately like to answer a question posed from the perspective of the ground model L: "Under what conditions on B and on the B-name $\dot{\mathcal{N}}$ will $\dot{\mathcal{N}}$'s interpretation be *forced* to self-collect into $\mathbb{R}(G)$?"

The Main Question (Ground-Model Perspective): *Assuming Stipulation 3.11, how do we calculate $\|\dot{G} \in L(\bigcup \dot{\mathcal{N}})\|$ in terms of B and $\dot{\mathcal{N}}$?*

We will answer this version of the question in Theorem 8.7.

As to the question of whether these conditions can ever be met, the answer is yes; we will present in Section 10 an example of a set $\mathcal{N}$ of continua that self-collects. However, this example does not provide a model of the self-construction axioms since for each $X \in \mathcal{N}$, no subset of X's predecessors (that is, of X's proper sub-continua) self-collects into X. Prospects for obtaining models of the self-construction axioms are discussed in Section 10.

4 Equivalence Classes $\theta(X)$ Of Generic Filters

We now present the key to our analysis, which is a shift in perspective from reals and sets of reals, to sets of generic filters with which they are interconstructible. We know from Lemma 3.8 that a singly-generated continuum $Y \in L[G]$ is interconstructible with the filter $G \cap D$ for some subalgebra D of B. To apply this fact toward reconstructing G from $L(\bigcup \mathcal{N})$ (as the Main Question's first version demands), we might hope for a $\bigcup \mathcal{N}$-absolute way to define a set consisting of one such $G \cap D$ for each such $Y \in L(\bigcup \mathcal{N})$, and hope for the union of these $G \cap D$ to be all of G. However, multiple generic filters F on subalgebras of B might be interconstructible with Y, not all of them having form $G \cap D$. In general, the closest we can get to realizing the hope just expressed is to associate each singly-generated continuum $Y \in L(\bigcup \mathcal{N})$ with the *equivalence class* of filters interconstructible with Y. We call this equivalence class Y's "constructibility degree" of filters, and denote it $\theta(Y)$:

$$\theta(Y) \equiv \{F \in L(Y) : F \text{ is a } L\text{-generic filter on some } C \in \mathsf{ACSAs}(B) \text{ and } Y \in L(F)\}.$$

We now let Θ be an interconstructible stand-in for $\bigcup \mathcal{N}$ built with these constructibility degrees, defined in the following $\bigcup \mathcal{N}$-absolute way:

$$\Theta \equiv \{\theta(\mathbb{R}(x)) : x \in \bigcup \mathcal{N}\}.$$

Θ is indeed interconstructible with $\bigcup \mathcal{N}$, since in the other direction we have

$$\bigcup \mathcal{N} = \bigcup \{\mathbb{R}(M) : M \in \Theta\}.$$

It follows from the transitivity of constructibility (Corollary 3.3) that a given set is constructible from Θ if and only if it is constructible from $\bigcup \mathcal{N}$, and, in particular:

Lemma 4.1 *Under Stipulation 3.11, $\mathcal{N}$ satisfies $G \in L(\bigcup \mathcal{N})$ if and only if $G \in L(\Theta)$, where Θ is defined in terms of $\bigcup \mathcal{N}$ as above.* $\square$

This confirms the equivalence of the Main Question's first and second versions. Our hope for answering the latter — i.e. for finding a condition under which G is constructible from Θ — is a less naive version of the hope expressed at the outset of this section: it is to find something about Θ's structure that lets us deduce which filters in $\bigcup \Theta$ are genuinely of form $G \cap D$, as opposed to specious "copies," and lets us conclude that G is the union of all the former. Essentially all of our work from this point through Theorem 7.8 is dedicated to realizing this hope.

Since $\mathcal{N}$ meets Stipulation 3.11 with boolean (B) value 1, it is easy to verify that Θ as we have defined it from $\mathcal{N}$ will meet the following with boolean value 1:

Lemma 4.2 *Under Stipulation 3.11, the following assertions, corresponding to the first three requirements of the Self-Collection axiom, are forced by $1 \in B$:*

(i). $G \notin \bigcup \Theta$;

(ii). $(G \cap C, G \cap D \in \bigcup \Theta) \iff G \cap (\overline{C \cup D}) \in \bigcup \Theta$;

(iii). $\Theta \in L(G)$ [trivial since Θ is G's interpretation of a B-name]. $\square$

We now define a function $\chi : \mathsf{ACSAs}(B) \to B$ to serve as a "boolean characteristic function" for $\bigcup \Theta$, or more precisely, to encode for each D whether $G \cap D$ is a member of $\bigcup \Theta$:

$$\chi(D) \equiv \|\dot{G} \cap D \in \bigcup \dot{\Theta}\|.$$

Lemma 4.3 *Under Stipulation 3.11 the following hold:*

(i). $\chi(B) = 0$*;*

(ii). $(\forall C, D)(\chi(C) \wedge \chi(D) = \chi(\overline{C \cup D}))$.

This is easy to verify since each claim about χ is a straightforward translation of the corresponding one on Θ in Lemma 4.2. □

Clause (ii) of the preceding lemmas has the following useful consequence:

Lemma 4.4 *Clause (ii) on* Θ *(Lemma 4.2) entails*

$$(G \cap C \in \bigcup \Theta \text{ and } C \geq_G D) \Rightarrow G \cap D \in \bigcup \Theta;$$

similarly, clause (ii) on χ *(Lemma 4.3) entails*

$$(\forall C, D)(\chi(C) \wedge \|C \geq_G D\| \leq \chi(D)). \quad \square$$

Structure of the Equivalence Classes $\theta(X)$

We now show that generic filters on $\mathsf{ACSAs}(B)$-members in $L[G]$ are interconstructible if and only if they are "essentially isomorphic" copies of each other; consequently, the constructibility degrees $\theta(Y)$ will be equivalence classes of copies. (We should note that the basic idea here has long been understood; we introduce the term "copy" only as a convenience.)

Definitions:

$a, b \in B$ are *compatible* if $a \wedge b \neq 0$; similarly, a is compatible with $X \subseteq B$ if it is compatible with each X-member, and $X, Y \subseteq B$ are compatible if they are pairwise compatible.

The *upwards closure* of $X \subseteq B$ is $\{b \in B : (\exists x \in X)(x \leq b)\}$.

A *principal ideal* of $C \in \mathsf{ACSAs}(B)$ is $\{c \in C : c \leq e\}$ for some $e \in C^+$ (we put C^+ rather than C here because it will simplify matters to exclude the degenerate ideal $\{0\}$ by definition).

When $c \in C^+$, we use $C \restriction c$ to mean the principal ideal of C whose greatest member is c; and we write "$\max(I)$" for the greatest member of a principal ideal I.

Similarly, when ϕ is a function on C, we write $\phi \restriction c$ to mean ϕ's restriction to $C \restriction c$. (When context suggests that X is not a member of C, but rather a set likely to have a nonempty intersection with C, $\phi \restriction X$ will mean ϕ's restriction to $C \cap X$, as usual.)

Note that a principal ideal $C \restriction c$ can be considered a boolean algebra in its own right, having c as its greatest member, inheriting $\wedge$ and $\vee$ from C, and having a negation operation defined from C's negation $\neg$ as $e \mapsto c \wedge \neg e$.

When ϕ is a function, $\operatorname{dom}\phi$ and $\operatorname{ran}\phi$ denote its domain and range; and we adopt the common square-bracket notation $\phi[X]$ to mean $\operatorname{ran}(\phi \upharpoonright X)$.

If $b \in B$ and $A \subseteq B$, then $b \wedge [A]$ means $\{b \wedge a : a \in A\}$.

$\mathsf{Aut}(B)$ is the set of automorphisms of B (bijections from B onto itself that preserve boolean structure) in the ground model L.

A *partial automorphism* of B is a bijection between principal ideals of (possibly distinct) ACSAs of B that preserves boolean structure (when the principal ideals are considered as boolean algebras in their own right, as described above).

$\mathsf{ParAut}(B)$ is the set of B's partial automorphisms in L.

Definition of "copy": When $C, D \in \mathsf{ACSAs}(B)$, and F is a filter on C, and $\phi \in \mathsf{ParAut}(B)$ is an isomorphism from some principal ideal of C not disjoint from F to some principal ideal of D, we say that ϕ *copies* F to the upward closure in D of $\phi[F \cap \operatorname{dom}\phi]$, and that this upward closure is a *copy* of F, *via* ϕ.

The importance of copies is established in the following lemma, which shows that for any real number x, its "constructibility degree" $\theta(\mathbb{R}(x))$ of filters is just an equivalence class of copies.

Lemma 4.5 *For each singly-generated unconstructible continuum $X \in L[G]$,*

$(\exists C)(G \cap C \in \theta(X))$*, and*

$(\forall C)(G \cap C \in \theta(X) \;\Rightarrow\; \theta(X) = \{F : F \textit{ is a copy of } G \cap C\})$.

We will establish some preliminary lemmas in order to prove this.

Lemma 4.6 *If G_C, G_D are generic filters on (respectively) $C, D \in \mathsf{ACSAs}(B)$, then G_C constructs G_D (i.e. $G_D \in L(G_C)$) if and only if there exists $C' \in \mathsf{ACSAs}(B)$, $C' \subseteq C$, such that G_D is a copy of $G_C \cap C'$.*

The "if" direction is straightforward; we consider the "only if" direction. Suppose $G_D \in L(G_C)$. G_D is interconstructible with a real number r_D by Lemma 3.6. Therefore by transitivity of relative constructibility (Corollary 3.3 (b)), $r_D \in L(G_C)$. Again by Lemma 3.6, this time using our C for its B and our C_G for its G, we have r_D interconstructible with $L(G_C \cap C')$ for some $C' \in \mathsf{ACSAs}(C)$. Then by Corollary 3.3 (b) and (c), G_D is interconstructible with $G_C \cap C'$, and $L(G_D) = L(G_C \cap C')$.

Note that G_D and $(G_C \cap C')$ are subsets of L, so the parentheses and square-brackets versions of relative constructibility are equivalent for them (see Lemma 12.5), i.e. $L(G_D) = L[G_D]$ and $L(G_C \cap C') = L[G_C \cap C']$.

Since $G_D \in L[G_C \cap C']$, there exists a C'-name γ whose interpretation by $G_C \cap C'$ is G_D. Being a filter, G_D is upwards-closed in D, and therefore the following "nice" C'-name δ also yields G_D when interpreted by $G_C \cap C'$:

$$\delta \equiv \{(\check{d}, \bigvee\{\|\check{d}' \in \gamma\|_{C'} : d' \in D \text{ and } d' \leq d\}) : d \in D\}.$$

We will use δ rather than γ as a name for G_D because δ, considered as a function whose domain is the set of canonical forcing names for elements of D, satisfies $d \leq d' \Rightarrow \delta(\check{d}) \leq \delta(\check{d}')$. Similarly, let σ be a nice D-name whose interpretation by G_D is $G_C \cap C'$, and which satisfies $c \leq c' \Rightarrow \sigma(\check{c}) \leq \sigma(\check{c}')$.

For the remainder of this proof we will, for simplicity, permit ourselves to conflate elements of D with their canonical C'-names, and vice-versa, so that δ can be considered a function from D into C', and σ a function from C' into D. Therefore we may assert the following:

$$d \leq d' \Rightarrow \delta(d) \leq \delta(d'); \tag{1}$$
$$c \leq c' \Rightarrow \sigma(c) \leq \sigma(c'); \tag{2}$$
$$d \in G_D \iff \delta(d) \in G_C \cap C'; \tag{3}$$
$$c \in G_C \cap C' \iff \sigma(c) \in G_D. \tag{4}$$

From the last two properties we have

$$(\forall c \in C')(c \in G_C \iff \delta(\sigma(c)) \in G_C). \tag{5}$$

Now the genericity of G_C would guarantee a counterexample to (5) unless there existed $c \in G_C \cap C'$ such that:

$$\delta(\sigma(c')) = c' \text{ for all } c' \in C' \restriction c. \tag{6}$$

Similarly, because G_D is generic, we deduce the existence of $d \in G_D$ such that:

$$\sigma(\delta(d')) = d' \text{ for all } d' \in D \restriction d. \tag{7}$$

Since $d \in G_D$ we have $\delta(d) \in G_C$ so that $\delta(d) \wedge c \in G_C \cap C'$.

We claim that σ's restriction to $C' \restriction (\delta(d) \wedge c)$ is a $\mathsf{ParAut}(B)$-member that copies $(G_C \cap C')$ to G_D. It is an isomorphism, that is, $c_1 \leq c_2 \iff \sigma(c_1) \leq \sigma(c_2)$, for (2) above establishes one direction, and in the other, if $c_1 \not\leq c_2$ but $\sigma(c_1) \leq \sigma(c_2)$, then (6) and (1) entail a contradiction. Finally, the range of $\sigma \restriction \delta(d) \wedge c$ is all of $D \restriction \sigma(\delta(d) \wedge c)$; if $d' < d$ were a counterexample to this, $\delta(d')$ would be in the domain of $\sigma \restriction \delta(d) \wedge c$ by (1), but then $\sigma(\delta(d')) \neq d'$, violating (7). $\square$

Lemma 4.7 *If in addition to the supposition of Lemma 4.6 we have $L(G_D) = L(G_C)$, we can require there that C' be C itself.*

At the outset of our proof we chose C' to be any ACSA of C such that $(G_C \cap C')$ is interconstructible with G_D, so knowing $L(G_D) = L(G_C)$ allows us to choose C itself. $\Box$

Lemma 4.8 *A copy F' of a filter F on some $C \in \mathsf{ACSAs}(B)$ will be L-generic if and only if F is L-generic.*

Let ϕ copy F to F', and let $D \in \mathsf{ACSAs}(B)$ be such that $\operatorname{ran}\phi$ is a principal ideal of D, and F' is the upwards closure of $\phi[F \cap \operatorname{dom}\phi]$ in D.

We first show that F will be L-generic on C—that is, meet every constructible dense subset S of C^+—just if $F \cap \operatorname{dom}\phi$ is L-generic on $\operatorname{dom}\phi$. This is trivial if $\max(\operatorname{dom}\phi) = 1$, so assume otherwise. Note that $\max(\operatorname{dom}\phi) \in F$ by the "F not disjoint from $\operatorname{dom}\phi$" clause of the definition of "copy," and the fact that filters are upwards-closed.

Suppose F is L-generic on C and let S be any constructible dense subset of $\operatorname{dom}\phi$. Define $S' \equiv S \cup (C^+ \restriction \neg\max(\operatorname{dom}\phi))$. S' is dense in C^+, so there exists $s \in F \cap S'$; and this s must be in the S half of S' since otherwise, by the closure of F under $\wedge$, we would have $\max(\operatorname{dom}\phi) \wedge s = 0 \in F$, contradicting F's being a proper filter. Conversely, if $F \cap \operatorname{dom}\phi$ is L-generic on $\operatorname{dom}\phi$, and S is any constructible dense subset of C^+, define $S' \equiv \max(\operatorname{dom}\phi) \wedge [S]$; then S' is dense in $\operatorname{dom}\phi$, so there exists $s' \in F \cap S'$, which by definition of S' must be of form $s' = \max(\operatorname{dom}\phi) \wedge s$ for some $s \in S$, and this $s \in S$ must also be in F by the upward-closure of filters.

Likewise, F' is L-generic on D just if $F' \cap \operatorname{ran}\phi$ is L-generic on $\operatorname{ran}\phi$. The lemma then follows from the fact that ϕ is an isomorphism, which entails that a subset of $\operatorname{dom}\phi$ will be dense just if its image is dense in $\operatorname{ran}\phi$. $\Box$

Proof of Lemma 4.5:

"$\exists C$" clause: Just after defining $\theta(X)$ above, we noted that it always has a member of form $G \cap C$ (unless X is constructible, a case that Lemma 4.5's statement explicitly excludes).

"$\forall C$" clause: Letting C be arbitrary such that $G \cap C \in \theta(X)$, $G \cap C$ is of the form stated in the lemma because it is trivially a copy of itself; any other copy of $G \cap C$ is L-generic by Lemma 4.8 and is clearly interconstructible with $G \cap C$, so satisfies the definition of a $\theta(X)$-member. Finally, it follows from Lemma 4.6 and Lemma 4.7 that every $F \in \theta(X)$ is of this form. $\Box$

The following lemma gives more information about an important special case of Lemma 4.6, in which G_C and G_D are the restrictions to C and D of our generic G on B.

Lemma 4.9 $D \leq_G C \iff (\exists g \in G)(g \wedge [D] \subseteq g \wedge [C])$.

Recall that $D \leq_G C$ by definition means $G \cap D \in L(G \cap C)$, and that $L(G \cap C) = L[G \cap C]$ (Lemma 12.5, cited previously).

$\Rightarrow$: Because $G \cap D \in L[G \cap C]$, there is a C-name (a B-name involving only elements of C) whose interpretation by $G \cap C$ is $G \cap D$; call it δ. Because δ names the same set as the canonical B-name $\dot{G} \cap \check{D}$, we must have

$$\|\delta = \dot{G} \cap \check{D}\|_B \in G.$$

Let g denote $\|\delta = \dot{G} \cap \check{D}\|_B$; we verify that g meets the lemma's requirement, which is that for all $d \in D$ some $c \in C$ satisfies $c \wedge g = d \wedge g$. We claim in fact that for all d, $\|\check{d} \in \delta\|_C$ is a c that satisfies this. For if $\|\check{d} \in \delta\|_C \wedge g \neq d \wedge g$, then

$$g \wedge ((d \wedge \neg\|\check{d} \in \delta\|_C) \vee (\neg d \wedge \|\check{d} \in \delta\|_C))$$

is nonzero, and is an element of B below g that forces $\delta \neq \dot{G} \cap \check{D}$, impossible.

$\Leftarrow$: It is easy to verify that $\{d \in D : (\exists c \in G \cap C)(g \wedge c = g \wedge d)\}$ will evaluate to $G \cap D$ in $L(G \cap C)$. $\Box$

We can express the meaning of Lemma 4.9 by saying: there is some element g of information in the filter G, such that any *additional* information about G's extension that might be supplied by $G \cap D$, could just as well be supplied by $G \cap C$. This suggests the following terminology.

Definition: Call $b \in B^+$ an *information supplement for C's generic to construct D's generic* if $b \wedge [D] \subseteq b \wedge [C]$. Let $IS(C, D)$ denote the set of all such information supplements; it is clearly a downward-closed subset of B^+. With this definition we can restate Lemma 4.9 thus:

Corollary 4.10 $\|C \geq_G D\|_B = \bigvee IS(C, D)$. $\Box$

5 Self-Collection Condition In Terms Of Definability

We are now ready to begin answering the Main Question (in its "Θ version") stated at the end of Section 3.3, from the perspective of the generic extension $L[G]$. The question is under what conditions Θ obeying Stipulation 3.11 (and therefore all clauses of Lemma 4.2) constructs G, meaning $G \in L(\Theta)$.

Our starting point is an answer we have already found, in terms of definability: By Lemma 3.4, $G \in L(\Theta)$ if and only if G is definable by a Θ-absolute predicate. This answer is not very helpful by itself, since there is no obvious way to check whether this condition holds. We will now begin the work of translating it into a more easily verifiable condition, work which will culminate in Theorem 7.8. The first step is to obtain a somewhat tighter definability condition.

Lemma 5.1 *Under Stipulation 3.11, $G \in L(\Theta)$ if and only if*

$$G = \{x : \Psi(x, \Theta, G \cap C, c'_1, \ldots c'_n)\}$$

for some Θ-absolute predicate $\Psi(x, \Theta, G \cap C, c'_1, \ldots c'_n)$ such that each c'_i is a constant in L and $G \cap C \in \bigcup \Theta$.

The "only if" direction is the nontrivial one. Assume $G \in L(\Theta)$ and (invoking Lemma 3.4) that G is definable by a Θ-absolute predicate as

$$G = \{x : \Delta(x, c_1, \ldots c_n)\}.$$

We wish to show that there exists a Θ-absolute predicate of the more narrowly-defined form demanded in the lemma's statement, that defines the same set G.

The idea here is as follows. Δ's given constants $c_1, \ldots, c_n$ must refer, by the definition of Θ-absolute predicate, to sets that belong to L and/or to $\mathsf{TrCl}(\{\Theta\})$. We wish to find a *single* constant referring to some $(G \cap C) \in \mathsf{TrCl}(\{\Theta\})$, such that each c_i can be replaced in Δ by a subformula that defines c_i's referent using only Θ, this $G \cap C$, and finitely many constants *required to be members of* L, all of which will be parameters for our Ψ. What allows us to do this is the fact (established by Lemma 4.5) that each given c_i that is not already an L-member must be interconstructible with $(G \cap C_i)$ for some $C_i \in \mathsf{ACSAs}(B)$ such that $(G \cap C_i) \in \bigcup(\Theta)$. Clause (ii) of Lemma 4.2 on Θ (which translates the directedness requirement on $\mathcal{N}$) then guarantees the existence of a single $C \in \mathsf{ACSAs}(B)$ that includes each such C_i, and satisfies $(G \cap C) \in \bigcup \Theta$. Each of the $(G \cap C_i)$'s is then definable as the intersection with C_i of this one $(G \cap C)$. And each C_i is an L-member and so can be appended as another constant for our Ψ.

We formalize this argument in the appendix, Section 12. $\square$

A key consequence of Lemma 5.1 is that the predicate's Θ-absoluteness rules out the existence of certain filters G' whose restriction to a certain C is the same as G's. This is made precise in the following lemma.

Lemma 5.2 *If the situation stated in Lemma 5.1 holds, and is forced to hold (as it must be) by some $g \in G$, i.e.*

$$g \Vdash \dot{G} = \{\check{x} \in \check{B} : \Psi(\check{x}, \dot{\Theta}, \dot{G} \cap \check{C}, \check{c}'_1, \ldots \check{c}'_n)\}, \qquad (\star)$$

then there cannot exist in any ZF model an L-generic filter G' on B, $G' \neq G$, such that $g \in G'$, $G' \cap C = G \cap C$, and G' gives the same interpretation as G to the B-name $\dot{\Theta}$.

For if such a G' existed, and were used to interpret B-names to yield a generic extension $L[G']$, then the expression

$$\{x \in B : \Psi(x, \Theta, G \cap C, c'_1, \ldots, c'_n)\}$$

would evaluate there to $G' \neq G$, contradicting the Θ-absoluteness of $\Psi(\ldots)$. $\square$

6 Rigid inclusions of boolean algebras

We have just seen in Lemma 5.2 that in order for $G \in L(\Theta)$ to hold, there must not exist certain filters $G' \neq G$ having the same restriction to some subalgebra C as G. We will now see that, in the special case that the g in that lemma equals 1, this also rules out certain automorphisms of B that leave C fixed and would map G onto such a forbidden G'. This suggests the idea of a *rigid inclusion* of boolean algebras $C \subseteq B$, for which no such automorphisms exist. The one fact we will prove in this section makes this idea precise. Because this fact concerns only the special case just noted (i.e. $g = 1$) it will not be referenced later. We include it nonetheless because it shows in a particularly simple way the relevance to our work of rigid inclusions, whose importance will become clearer later: in Section 9.1 we will see that a lack of rigid boolean inclusions $C \subseteq B$, combined with a few other properties, suffices to prevent self-collection; and in Section 10 we will propose a way to obtain rigid boolean inclusions, as the most promising route to a model of the self-construction axioms.

Definitions. $C \subseteq B$ is a *rigid boolean inclusion* if no $\phi \in \mathsf{Aut}(B)$ other than the identity mapping leaves every C-member fixed. It is a *rigid boolean+χ inclusion* if no $\phi \in \mathsf{Aut}(B)$ other than the identity mapping leaves every C-member fixed and preserves χ, in the sense that $\phi(\chi(D)) = \chi(\phi[D])$ for all $D \in \mathsf{ACSAs}(B)$. (Note that if $C \subseteq B$ is a rigid boolean inclusion then it is also a rigid boolean+χ inclusion, for any χ.)

Fact 6.1 *Recalling that ($\star$) of Lemma 5.2 must hold for some C and g if we are to have $G \in L(\Theta)$, if ($\star$) holds for some C and $g = 1$, then $C \subseteq B$ is a rigid boolean+χ inclusion. If furthermore $\chi(D) = \|D <_G B\|$ for all $D \in \mathsf{ACSAs}(B)$, then $C \subseteq B$ is a rigid boolean inclusion.*

Suppose ($\star$) of Lemma 5.2 holds with $g = 1$ for some fixed C.

For the first claim, suppose $\phi \in \mathsf{Aut}(B)$ furnished a counterexample, i.e. ϕ is not the identity mapping but preserves χ and leaves each C member fixed. Let G be a generic ultrafilter on B such that when we define $G' \equiv \phi[G]$, we have $G' \neq G$. (This is possible because ϕ is not the identity mapping.) Consider G'. It too is L-generic, by Lemma 4.8. Since ϕ leaves every C member fixed, we have $G \cap C = G' \cap C$. And because ϕ exists in the ground model and satisfies $\phi(\chi(D)) = \chi(\phi[D])$ for all $D \in \mathsf{ACSAs}(B)$, $\dot{\Theta}$'s interpretation is the same under G' as under G, as is easily verified using χ's definition in terms of $\dot{\Theta}$. But then since ($\star$) holds with $g = 1$, G' violates Lemma 5.2; so the first claim can have no counterexample.

For the second claim, suppose $\chi(D) = \|D <_G B\|$ for all $D \in \mathsf{ACSAs}(B)$; it suffices to verify that this implies every $\phi \in \mathsf{Aut}(B)$ preserves χ, for then any counterexample ϕ to $C \subseteq B$ being a rigid boolean inclusion would also be a counterexample to its being a rigid boolean+χ inclusion.

Let $\phi \in \mathsf{Aut}(B)$ and $D \in \mathsf{ACSAs}(B)$ be arbitrary. From the definition of $IS(C, D)$ used in Corollary 4.10,

$$IS(D, B) = \{b \in B^+ : B \upharpoonright b = b \wedge [D]\}.$$

It follows from this and the fact that ϕ is a boolean isomorphism that for all b,

$$b \in IS(D, B) \iff \phi(b) \in IS(\phi[D], B),$$

which in turn implies

$$\phi(\bigvee IS(D, B)) = \bigvee IS(\phi[D], B).$$

Then by Corollary 4.10 and the fact that $\|D \geq_G B\| = \neg\|D <_G B\|$,

$$\phi(\neg\|D <_G B\|) = \neg\|\phi[D] <_G B\|.$$

Since ϕ preserves $\neg$,

$$\phi(\|D <_G B\|) = \|\phi[D] <_G B\|;$$

then our supposition that $\chi(D) = \|D <_G B\|$ and $\chi(\phi[D]) = \|\phi[D] <_G B\|$ implies

$$\phi(\chi(D)) = \chi(\phi[D]).$$

Since D was arbitrary, this means ϕ preserves χ, which, as we noted, is what it suffices to verify. $\square$

We do not know at present whether a *sufficient* condition for self-collection can be found in terms of the existence of rigid boolean+χ inclusions. Matters would be greatly simplified if it were. In our proof of Theorem 7.8, though, we will see that it may be possible to cobble together some of B's *partial* automorphisms—in a model larger than $L[G]$—to obtain a filter G' as in Lemma 5.2 that witnesses the failure of self-collection, yet is not isomorphic to G via any ground-model automorphism of B.

In any case, the existence of a rigid boolean+χ inclusion $C \subset B$ would not supply any obvious formula Ψ for defining G from Θ and $G \cap C$ (plus finitely many constants in L). In the next sections we will find a necessary and sufficient condition for self-collection that does provide such a Ψ.

7 Self-Collection Condition in Terms of Filters

The concept we use for turning Lemma 5.1's self-collection condition into a practically verifiable one is "generic refinability." We will define this concept in parallel for generic filters on ACSA's of B, and for $\mathsf{ParAut}(B)$-members. Having done so, we will derive a self-collection condition in terms of generically-refinable filters that can be checked from the generic-extension perspective; then in the next section we will translate it into a condition in terms of generically-refinable $\mathsf{ParAut}(B)$-members, which can be checked in the ground model perspective. A basic technical tool we will need is "one-element refinements."

7.1 One-element refinements

It will be important for us to be able to expand an ACSA by adding a single B-member to it and taking its completion. We use the following notation for this: When $C \in \mathsf{ACSAs}(B)$ and $b \in B$,

$$C^{+b} \equiv \{(c_1 \wedge b) \vee (c_2 \wedge \neg b) : c_1, c_2 \in C\}.$$

Lemma 7.1 *If $C \in \mathsf{ACSAs}(B)$ and $b \in B$, then $C^{+b} \in \mathsf{ACSAs}(B)$.*

This follows from routine boolean-algebraic manipulations. □

Refining. If F, F' are filters, we say F' *refines* F if $F \subseteq F'$. If $\phi, \phi' \in \mathsf{ParAut}(B)$, we say ϕ' *refines* ϕ if:

$\max(\operatorname{dom}\phi') \le \max(\operatorname{dom}\phi)$, and $\max(\operatorname{ran}\phi') \le \max(\operatorname{ran}\phi)$; and

$(\forall e \in \operatorname{dom}\phi)(\phi'(e \wedge \max(\operatorname{dom}\phi')) = \phi(e) \wedge \max(\operatorname{ran}\phi'))$.

Note the second clause implies $\max(\operatorname{dom}\phi') \wedge [\operatorname{dom}\phi] \subseteq \operatorname{dom}\phi'$.

Definition. When $b \in B$ and C is an ACSA of B — or, more generally, a subset of B that is closed under $\bigwedge$ (arbitrary meets) — we define

$$b \uparrow C \equiv \bigwedge\{c \in C : b \le c\}.$$

Given $\phi \in \mathsf{ParAut}(B)$ and some $b \in B^+$, $b \le \max(\operatorname{dom}\phi)$, a *one-element domain refinement* of ϕ by b is a refinement ϕ^{+b} of ϕ satisfying $\operatorname{dom}\phi^{+b} = b \wedge [\operatorname{dom}\phi]$ and $\max(\operatorname{ran}\phi^{+b}) = \phi(b \uparrow \operatorname{dom}\phi)$. We will now see that these requirements uniquely determine ϕ^{+b}.

Lemma 7.2 *If $C \in \mathsf{ACSAs(B)}$ and $b \in B^+$, the function $\psi : a \mapsto a \wedge b$ with $\operatorname{dom}\psi = C \restriction (b \uparrow C)$ is an isomorphism such that $\psi \in \mathsf{ParAut}(B)$, and $\psi(a) \uparrow C = a$ for all $a \in \operatorname{dom}\psi$.*

If we can show

$$(a \wedge b) \uparrow C = a \tag{*}$$

for all $a \in \operatorname{dom}\psi$, then the lemma's other claims will follow easily. Let $a \in \operatorname{dom}\psi$ be arbitrary, so that $a \in C$ and $a \le (b \uparrow C)$. Plainly (*) holds with $\le$ in place of $=$; we must show that this inequality cannot be strict. So suppose

$$(a \wedge b) \uparrow C = a' < a$$

holds for some a'. It would follow that $b \wedge a = b \wedge a'$. Then boolean laws entail

$$\begin{aligned} b &= (b \wedge a) \vee (b \wedge \neg a) \\ &= (b \wedge a') \vee (b \wedge \neg a) \\ &= b \wedge (a' \vee \neg a). \end{aligned}$$

Thus $b \le (a' \vee \neg a)$. But then also $b \le (a' \vee \neg a) \wedge (b \uparrow C)$; the right side of that inequality is strictly $< (b \uparrow C)$, contradicting the definition of $(b \uparrow C)$ as the least C-member that is $\ge b$. □

Lemma 7.3 *Given $\phi \in \mathsf{ParAut}(B)$ and some $b \in B^+$, $b \leq \max(\operatorname{dom}\phi)$, the one-element domain refinement ϕ^{+b} exists and is unique.*

Let C be any ACSA of B such that $\operatorname{dom}\phi \subseteq C$, and note that $b \wedge [\operatorname{dom}\phi]$ $(= b \wedge [C])$ is a principal ideal of C^{+b} defined above.

By Lemma 7.2, the map $\psi : a \mapsto a \wedge b$ is an isomorphism mapping $C \restriction (b \uparrow C)$ onto $b \wedge [C]$, the latter being a principal ideal of C^{+b}. Thus we may define ϕ^{+b} by composition:

$$\phi^{+b}(d) \equiv \phi(\psi^{-1}(d)),$$

with the domain of ϕ^{+b} set to $b \wedge [C]$. It is straightforward to check that ϕ^{+b} meets the definition of a refinement of ϕ, and that it is the unique such refinement satisfying the definition of a one-element refinement of ϕ by b. $\square$

Definition. In light of Lemma 7.3 we may use ϕ^{+b} to denote *the* one-element domain refinement of ϕ by b. Similarly, for $b \leq \max(\operatorname{ran}\phi)$, we define the *one-element range refinement* ϕ_{+b} of ϕ to be the refinement of ϕ satisfying $\operatorname{ran}\phi_{+b} = b \wedge [\operatorname{ran}\phi]$ and $\max(\operatorname{dom}\phi_{+b}) = \phi^{-1}(b \uparrow \operatorname{ran}\phi)$.

Lemma 7.4 *For all $C \in \mathsf{ACSAs}(B)$ and $b \in B^+$, $\|C =_G C^{+b}\| = 1$; so $\chi(C) = \chi(C^{+b})$.*

Suppose $b \in G$. By Lemma 7.2, the map $a \mapsto a \wedge b$ is a member of $\mathsf{ParAut}(B)$ mapping $C \restriction (b \uparrow C)$ onto $b \wedge [C]$, the latter being a principal ideal of C^{+b}. Since $b \in G$, neither the domain of this map nor its range is disjoint from G, so this map copies $G \cap C$ to $G \cap C^{+b}$, and its inverse copies $G \cap C^{+b}$ to $G \cap C$. Thus by Lemma 4.6, $G \cap C^{+b}$ and $G \cap C$ are interconstructible; whence by the definition of $=_G$ we have $C =_G C^{+b}$.

If on the other hand $\neg b \in G$, the same argument goes through with the map $a \mapsto a \wedge \neg b$. Thus both b and $\neg b$ force $C =_G C^{+b}$; thus 1 forces this too. $\chi(C) = \chi(C^{+b})$ then follows from Lemma 4.4. $\square$

7.2 Generic Refinability In Accordance With χ

Remark. The definitions of Γ_α, Γ, Φ_α, and Φ below are cumbersome, because they use recursion; it is a question we have not been able to settle whether there are simpler non-recursive definitions that would be equivalent. We will see that if $\Phi_1 = \Phi_0$, then $\Phi = \Phi_0$ (Lemma 9.4), but we have found no comparable simplification for other cases.

Accordance with χ. If F is an ultrafilter on $C \in \mathsf{ACSAs}(B)$, we say it *accords with* χ if $\chi(C)$ is in the upwards closure of F in B. For $\phi \in \mathsf{ParAut}(B)$, we say ϕ *accords with* χ if its domain and range are principal ideals of some C and D (respectively) such that $\max(\operatorname{dom}\phi) \leq \chi(C)$ and $\max(\operatorname{ran}\phi) \leq \chi(D)$. Note C and D may not be uniquely determined (since two distinct subalgebras of B can have a principal ideal in common), but the definition of χ ensures that if the requirement just stated holds for one choice of C, D, it holds for all choices. This is because whenever $\mathrm{dom}\,\phi$ is a principal ideal of both C and C', we have $\max(\mathrm{dom}\,\phi) \Vdash C =_G C'$ (by Lemma 4.9), so by Lemma 4.4 for χ,

$$\max(\mathrm{dom}\ \phi) \wedge \chi(C) = \max(\mathrm{dom}\ \phi) \wedge \chi(C'),$$

and similarly for the range of ϕ.

Generic refinability. $\phi \in \mathsf{ParAut}(B)$ is *generically refinable within a class* $W \subseteq \mathsf{ParAut}(B)$ if:

$(\forall C, D \in \mathsf{ACSAs}(B))(\forall S \subseteq B^+, S \in L$, with S dense in $B^+)$

$(\chi(C)$ compatible with $\max(\mathrm{dom}\,\phi) \Rightarrow (\exists \phi' \in W, s \in S)$

$(\phi'$ refines ϕ, and

$\max(\mathrm{ran}\,\phi') \leq s$, and

$\max(\mathrm{dom}\,\phi') \wedge [C] \subseteq \mathrm{dom}\,\phi'$, and

$(\max(\mathrm{ran}\,\phi') \leq \neg\chi(D)$, or $\max(\mathrm{ran}\,\phi') \wedge [D] \subseteq \mathrm{ran}\,\phi')))$.

Similarly, a filter $F \in \bigcup \Theta$ is *generically refinable in* (W, Θ), where $W \subseteq \bigcup \Theta$ if:

$(\forall D \in \mathsf{ACSAs}(B))(\forall M \in \Theta, S \subseteq B^+, S \in L$, with S dense in $B^+)$

$(\exists N \in \Theta, F' \in N \cap W, s \in S \cap F')$

$(M \in L(N)$ and F' refines F and

$(\neg\chi(D) \in F'$, or $F' \cap D$ is an L-generic filter on $D))$.

The set $\Gamma(\Theta)$ of fully refinable filters; the set Φ of fully refinable partial automorphisms

In the context of a particular χ, we use the above definitions to define a $\Gamma_\alpha(\Theta)$ hierarchy somewhat along the lines of $L_\alpha(\Theta)$, except here the sequences of levels shrink rather than grow:

$\Gamma_0(\Theta) \equiv \{F \in \bigcup \Theta : F$ accords with $\chi\}$;

$\Gamma_{\alpha+1}(\Theta) \equiv \{F \in \Gamma_\alpha : F$ is generically refinable in $(\Gamma_\alpha(\Theta), \Theta)\}$;

$\Gamma_\alpha(\Theta) \equiv \bigcap_{\beta<\alpha} \Gamma_\beta(\Theta)$ for limit α;

$\Gamma(\Theta) \equiv \bigcap_{\alpha \in Ord} \Gamma_\alpha(\Theta)$.

We call $\Gamma(\Theta)$ the set of *fully generically-refinable* filters (in the context of a given χ and Θ); similarly, we define the set Φ of fully generically-refinable partial automorphisms:

$\Phi_0 \equiv \{\phi \in \mathsf{ParAut}(B) : \phi$ accords with $\chi\}$;

$\Phi_{\alpha+1} \equiv \{\phi \in \Phi_\alpha : \phi$ is generically refinable within $\Phi_\alpha\}$;

$\Phi_\alpha \equiv \bigcap_{\beta<\alpha} \Phi_\beta$ for limit α;

$\Phi \equiv \bigcap_{\alpha \in Ord} \Phi_\alpha$.

Lemma 7.5 *Each $\Gamma(\Theta)$-member is generically refinable in $(\Gamma(\Theta), \Theta)$; and each Φ-member is generically refinable within Φ.*

The $\Gamma_\alpha(\Theta)$ hierarchy must stop shrinking at some ordinal α at which $\Gamma(\Theta) = \Gamma_\alpha(\Theta) = \Gamma_{\alpha+1}(\Theta)$. For this α we have by definition

$$\Gamma_{\alpha+1}(\Theta) = \{F \in \Gamma_\alpha : F \text{ is generically refinable in } (\Gamma_\alpha(\Theta), \Theta)\};$$

replacing $\Gamma_\alpha(\Theta)$ and $\Gamma_{\alpha+1}(\Theta)$ with $\Gamma(\Theta)$ here yields the claim about $\Gamma(\Theta)$. The claim about Φ is proved the same way. □

Lemma 7.6 *For all α, $\Phi_\alpha \in L$ and $\Gamma_\alpha(\Theta) \in L(\Theta)$; moreover, $\Phi \in L$ and $\Gamma(\Theta) \in L(\Theta)$.*

The proof is an application of our template for showing absoluteness of hierarchies, Lemmas 12.2 and 12.4. We show how to apply it to the $\Gamma(\Theta)$ hierarchy; the proof for Φ is similar. The template takes two predicates, $\Omega(\cdots)$ and $\Psi(\cdots)$, as its "inputs". For the first, let $\Omega(x', \Theta)$ be "$x' \in \bigcup\Theta$ and x' accords with χ"; and for the second, let $\Psi(x'', (Y, \Theta, \beta))$ be "$x'' \in Y \cap \bigcup\Theta$ and x'' is generically-refinable in $((Y \cap \bigcup\Theta), \Theta)$."

We affirm first that our $\Omega(x', \Theta)$ is a Θ-absolute predicate, since the union operation is absolute and χ is a member of L. We affirm secondly that our $\Psi(x'', (Y, \Theta, \beta))$ is (Y, Θ, β)-absolute: this is because the definition of generic refinability for filters quantifies only over Θ, Y, and constructible subsets of B^+.

With the claim about the $\Gamma_\alpha(\Theta)$ proved, the claim about $\Gamma(\Theta)$ follows from the observation that the levels $\Gamma_\alpha(\Theta)$ must stop shrinking when α reaches some ordinal δ, so that $\Gamma(\Theta) = \Gamma_\delta(\Theta)$. Similarly for Φ. □

The next lemma establishes that $\Gamma(\Theta)$ is not empty, and in particular has at least one member of form $G \cap C$ from each Θ-member.

Lemma 7.7 $(G \cap C) \in M \in \Theta \Rightarrow (G \cap C^{+\chi(C)}) \in M \cap \Gamma(\Theta)$.

Suppose $(G \cap C) \in M \in \Theta$. Note that since $G \cap C \in \bigcup\Theta$, $\chi(C) \in G$, by definition of χ.

If ϕ denotes the identity function on C, then the one-element range refinement $\phi_{+\chi(C)}$ (defined above) copies $G \cap C$ to $G \cap C^{+\chi(C)}$; therefore $G \cap C^{+\chi(C)} \in M$ by Lemma 4.5.

$\chi(C^{+\chi(C)}) = \chi(C)$ by Lemma 7.4, so $\chi(C^{+\chi(C)}) \in G \cap C^{+\chi(C)}$; thus $(G \cap C^{+\chi(C)}) \in \Gamma_0(\Theta)$.

Note this holds for all C, M as in the lemma's statement.

Now suppose there is a least ordinal $\delta > 0$ such that *some* C with $G \cap C \in \bigcup\Theta$ satisfies $(G \cap C^{+\chi(C)}) \notin \Gamma_\delta(\Theta)$, and fix one C witnessing this. Note δ must be a successor ordinal since otherwise $\Gamma_\delta(\Theta) = \bigcap_{\beta<\delta} \Gamma_\beta(\Theta)$, but $(G \cap C^{+\chi(C)}) \in \Gamma_\beta(\Theta)$ for all $\beta < \delta$.

Suppose that D, S, M witness $C^{+\chi(C)}$'s violation of the definition of generic refinability in $(\Gamma_{\delta-1}(\Theta), \Theta)$. Fix any $s \in S \cap G$ such that $s \leq \chi(D)$, if $\chi(D) \in G$, or otherwise such that $s \leq \neg\chi(D)$. (Such an s exists because G is generic.) We now need to find $N \in \Theta$ and $F' \in N$ that together with s will satisfy the condition of definition of generic refinability for D, S, M, contradicting the supposed violation of that definition.

Let E be such that $G \cap E \in M$ (some such E exists by Lemma 4.5). $G \cap E \in \bigcup \Theta$ implies by definition of χ that $\chi(E) \in G$. Let E' denote $E^{+\chi(E)+s}$. By Lemma 7.4, $\chi(E') = \chi(E)$, so $\chi(E') \in G$.

We consider two cases.

If $\neg\chi(D) \in G$, let E'' denote $\overline{C^{+\chi(C)} \cup E'}$; then clause (ii) of Lemma 4.2 on Θ, plus the fact that the meet of any pair of G-members is a G-member, entails $\chi(E'') \in G$. So by definition of χ, $G \cap E'' \in \bigcup \Theta$. Let $F'' = G \cap E''$ and let N be the unique Θ-member such that $F'' \in N$. By the leastness of δ, $F'' \in \Gamma_{\delta-1}(\Theta)$. So our F'', N satisfy the generic-refinability requirement for $C^{+\chi(C)} \in \Gamma_\delta(\Theta)$, contradiction.

If $\chi(D) \in G$, we argue similarly, but with E'' defined as $\overline{C^{+\chi(C)} \cup D \cup E'}$. We know that the χ-values of the three subalgebras used to define E'' are all in G; so as before we get $\chi(E'') \in G$, entailing $G \cap E'' \in \bigcup \Theta$. Letting $F'' = G \cap E''$ again, the same conclusions about F'' follow, and in addition we have that $F'' \cap D$ is a generic filter on D, fulfilling the last clause of generic refinability. So again F'', N satisfy the generic-refinability requirement for $C^{+\chi(C)} \in \Gamma_\delta(\Theta)$.

Thus no such ordinal δ can exist, and $C^{+\chi(C)} \in \Gamma(\Theta)$. $\square$

7.3 The self-collection condition in terms of filters

Theorem 7.8 (Self-collection condition in terms of filters) *Under Stipulation 3.11, $G \in L(\Theta)$ iff there exists $D \in \mathsf{ACSAs}(B)$ such that $(G \cap D) \in \Gamma(\Theta)$ and all $(G \cap D)$'s refinements in $\Gamma(\Theta)$ are mutually compatible, in which case the proposition "$x \in F \in \Gamma(\Theta)$ for some refinement F of $G \cap D$" serves as the predicate $\Psi(\ldots)$ for Lemma 5.1.*

For $D \in \mathsf{ACSAs}(B)$ satisfying $(G \cap D) \in \Gamma(\Theta)$, let Q_D be the set of all fully-refinable refinements of $G \cap D$, that is:

$$Q_D \equiv \{F \in \Gamma(\Theta) : G \cap D \subseteq F\}.$$

Since $\Gamma(\Theta) \in L(\Theta)$ (Lemma 7.6), $Q_D \in L(\Theta)$, and $\bigcup Q_D \in L(\Theta)$.

We begin with the "if" direction. If D exists as in the lemma's statement, then all Q_D-members are mutually compatible; we will show this entails $G = \bigcup Q_D$, which suffices. Note first that by Lemma 7.7, $G \cap D^{+\chi(D)}$ is a refinement of $G \cap D$ that is a member of $\Gamma(\Theta)$. Now for any $b \in B$, consider $(D^{+\chi(D)})^{+b}$. By Lemma 7.4, we have

$$\chi((D^{+\chi(D)})^{+b}) = \chi(D^{+\chi(D)}) = \chi(D);$$

thus $(D^{+\chi(D)})^{+b}$ already has its own χ-value as a member, and so again by Lemma 7.7, $(G \cap (D^{+\chi(D)})^{+b})$ is a refinement of $G \cap D$ in $\Gamma(\Theta)$, and so is a Q_D-member. So $b \in G$ iff $b \in \bigcup Q_D$, and since this holds for all b, $G = \bigcup Q_D$.

Next we prove the "only if" direction; supposing there is no D meeting the condition stated in the lemma, we will show that Θ does not construct G. Suppose towards a contradiction that Θ *did* construct G, and let Ψ, C, and $\vec{c}$ witness this in Lemma 5.1. Some $g \in G$ must force them to be witnesses; so let us choose $g \in G, g \leq \chi(C)$ such that

$$g \Vdash \{x : \Psi(x, \dot{\Theta}, \dot{G} \cap \check{C}, \vec{c})\} = \dot{G}.$$

We will now show that our suppositions permit the existence of a G' contradicting Lemma 5.2.

Let D be $(C^{+g})^{+\chi(C)}$. Note that $G \cap D \in \Gamma(\Theta)$ by Lemma 7.7. We will obtain our G' by forcing over the set Q_D ordered by reverse inclusion. Note that $L(G)$ is the ground model for this forcing. Let Y be a generic filter obtained by this forcing, and define G' to be $\bigcup Y$.

G' is clearly a filter on B^+. It is generic over L by the following argument. Let $S \in L$ be any dense subset of B^+. Because each $F \in Q_D$ satisfies $F \in \Gamma(\Theta)$, and each of its refinements in $\Gamma(\Theta)$ is in Q_D, any $F \in Y$ can be refined (by Lemma 7.5) to $F' \in Q_D$ with $F \cap S \neq \emptyset$. Thus, by genericity, some such F' is in Y, and $G' \cap S \neq \emptyset$.

It follows immediately from the definitions of D and of Q_D that $G \cap D = G' \cap D$; and since $g \in G \cap D$, also $g \in G' \cap D$.

To show that $G' \neq G$, note every $F \in Q_D$ can be refined to some $F' \in Q_D$ that is incompatible with G. For if some F were not so refinable, it would equal $G \cap D'$ for some $D' \supseteq D$, and all its refinements in $\Gamma(\Theta)$ would be mutually compatible (being compatible with G); D' would thus meet the requirements to be the D that we are supposing not to exist. By genericity/density, some F' inconsistent with G must therefore be in Y, whence $G' \neq G$.

To see that G' gives the same interpretation as G to $\dot{\Theta}$, consider first any unconstructible $M \in (\dot{\Theta})_G$. (For the remainder of this argument we use the notation $(\dot{\Theta})_G$ for G's interpretation of $\dot{\Theta}$, which we had been calling Θ, in order to avoid confusing it with $(\dot{\Theta})_{G'}$.) We will show that also $M \in (\dot{\Theta})_{G'}$.

Let $F \in Q_D$ be arbitrary. Since $F \in \Gamma(\Theta)$, there exists by Lemma 7.5 and the definition of generic refinability some $N \in (\dot{\Theta})_G$, and $F' \in N \cap \Gamma(\Theta)$, such that F' refines F and $M \in L(N)$. Since any $F \in Q_D$ is refinable to such an F', some such F' must, by Y's genericity, be a member of Y. Let $E \in \mathsf{ACSAs}(B)$ be the subalgebra on which F' is an ultrafilter. The definition of χ-accordance of filters along with $G' = \bigcup Y$ gives us $\chi(E) \in G'$; thus F' is a member of $\bigcup(\dot{\Theta})_{G'}$. Now since $M \in L(N)$, it follows from Lemma 4.9, the definition of Θ, and Lemma 4.5 that there must be an ACSA $E' \subseteq E$ such that $G' \cap E' \in M$. Lemma 4.4 applied to $(\dot{\Theta})_{G'}$ then entails that $G' \cap E'$ is a member of $\bigcup(\dot{\Theta})_{G'}$, so $M \in (\dot{\Theta})_{G'}$.

Conversely, fix an unconstructible $M \in (\dot{\Theta})_{G'}$. Since G', like G, is an L-generic ultrafilter on B, Lemma 7.7 applies to it, ensuring the existence of an E such that $G' \cap E \in M$. Now it would suffice to show that there exists an $F' \in Y$ such that F' refines $G' \cap E$, because then:

$F' \in Y \Rightarrow$

$F' \in Q_D \Rightarrow$

$F' \in \bigcup(\dot{\Theta})_G \Rightarrow$

$F' \in N \in (\dot{\Theta})_G$ for some $N \Rightarrow$

F' is a copy of some $G \cap A \in N$ via some $\phi \in \mathsf{ParAut}(B)$ (Lemma 4.5) $\Rightarrow$

$G' \cap E$ is a copy via (a restriction of) ϕ of $G \cap A'$ for some $A' \subseteq A \Rightarrow$

$G' \cap E \in \bigcup(\dot{\Theta})_G$ (Lemma 4.4) $\Rightarrow$

$M \in (\dot{\Theta})_G$.

To see that such an $F' \in Y$ does exist, consider that by the definition of generic refinability, any $F \in Q_D$ can be refined to some $F' \in Q_D$ such that either $\neg\chi(E) \in F'$, or $F' \cap E$ is a generic filter on E. By Y's genericity, then, some such F' is in Y, and since we already know $\chi(E) \in \bigcup Y$, the second alternative must hold, which entails that F' refines $G' \cap E$. □

8 Self-Collection Condition In Terms Of Isomorphisms

The preceding theorem gives a self-collection condition that can be checked in the generic extension, stated in terms of generic filters; we now prove lemmas that will let us translate this condition into a condition on the boolean algebra B (given in Theorem 8.7) that can be checked in the ground model, stated in terms of $\mathsf{ParAut}(B)$.

Definition. $\phi \in \mathsf{ParAut}(\mathsf{B})$ *fixes* $C \in \mathsf{ACSAs}(B)$ if it is a refinement of the identity function on C. This definition is equivalent to:

$$\max(\operatorname{dom}\phi) \wedge [C] \subseteq \operatorname{dom}\phi, \text{ and } (\forall b \in \operatorname{dom}\phi)(\phi(b) \uparrow C = b \uparrow C).$$

Lemma 8.1 *If ϕ copies $G \cap C$ to F and $b \in G$, ϕ^{+b} copies $G \cap C^{+b}$ to F.*

(We assume ϕ^{+b} is well defined, so $b \leq \max(\operatorname{dom}\phi)$.) Let D denote the ACSA on which F is an ultrafilter. We must show that for all $d \in D$,

$$(*) \quad (\exists c \in G \cap \operatorname{dom}\phi)(\phi(c) \leq d) \iff (\exists c' \in G \cap \operatorname{dom}\phi^{+b})(\phi^{+b}(c') \leq d).$$

We showed in Lemmas 7.2 and 7.3 that the map $\psi : c \mapsto b \wedge c$ defined on $C \restriction (b \uparrow C)$ is an isomorphism, and that $\phi^{+b}(c') = \phi(\psi^{-1}(c'))$ for all $c' \in \operatorname{dom}\phi^{+b}$. Since G is an ultrafilter on B and $b \in G$, we have for all $c \in \operatorname{dom}\phi$ that $c \in G \iff c \wedge b \in G$. Now every $c' \in \operatorname{dom}\phi^{+b}$ is of form $c \wedge b$ for some $c \in \operatorname{dom}\phi$, from which (*) follows. □

Lemma 8.2 *If ϕ copies $G \cap C$ to F, and F' refines F, and $F' \in \theta(Y)$ for some singly-generated continuum $Y \in L[G]$, there exist ϕ' and $D' \supseteq C$ such that ϕ' copies $G \cap D'$ to F' and ϕ' refines ϕ.*

By Lemma 4.5, *some* $\gamma \in \mathsf{ParAut}(B)$ copies *some* $G \cap C'$ to F'. Let $E \subseteq E'$ be the ACSAs on which F, F' are generic ultrafilters respectively. Let $D \subseteq C'$ be an ACSA of which $\gamma^{-1}[E \cap \text{ ran } \gamma]$ is a principal ideal, so that $D \cap G$ is interconstructible with $F' \cap E = F$. Thus $D =_G C$.

By Lemma 4.9 applied "in both directions," there exist $g_1, g_2 \in G$ such that $g_1 \wedge [D] \subseteq g_1 \wedge [C]$ and $g_2 \wedge [C] \subseteq g_2 \wedge [D]$. Since G is a filter, $g \equiv g_1 \wedge g_2 \in G$, and $g \wedge [C] = g \wedge [D]$.

Define $\gamma' \equiv \gamma^{+g}$, the one-element refinement of γ by g. If we define $D' \equiv \overline{C \cup C'}$, it is easy to check that $\operatorname{dom} \gamma'$ is a principal ideal of D', and that γ' refines ϕ, and copies $G \cap D'$ to F'. $\square$

Corollary 8.3 *If $F' \in \theta(Y)$ for some singly-generated continuum Y, and $G \cap C \subseteq F'$ for some $C \in \mathsf{ACSAs}(B)$, then there exists $D' \supseteq C$ and $\phi' \in \mathsf{ParAut}(B)$ such that ϕ' copies $G \cap D'$ to F', and ϕ' fixes C.*

In Lemma 8.2, let F be $G \cap C$, and let ϕ be the identity function on C. $\square$

Lemma 8.4 *In light of Stipulation 3.11's consequences for χ (Lemma 4.3), every $F \in \Gamma(\Theta)$ is the copy of some $G \cap E$ via some $\phi \in \Phi$.*

In $L[G]$, by Lemma 4.5, every generic filter F on any $D \in \mathsf{ACSAs}(B)$ is a copy via some $\phi \in \mathsf{ParAut}(B)$ of $G \cap E$ for some E. We must show that when $F \in \Gamma(\Theta)$, the ϕ and E can be chosen so that $\phi \in \Phi$. In fact we will show something stronger: given any choice of ϕ, E, there exists b such that $\phi^{+b} \in \Phi$ and ϕ^{+b} copies $G \cap E^{+b}$ to F.

Let $\Delta(\alpha)$ be the following assertion about an ordinal α:

$\Delta(\alpha) \iff$ *For every $F \in \Gamma(\Theta)$ and every ϕ copying some $G \cap E$ to F, there exists $b \in G$, $b \leq \max(\operatorname{dom} \phi)$, such that $\phi^{+b} \in \Phi_\alpha$.*

Any such ϕ^{+b} will copy $G \cap E^{+b}$ to F, by Lemma 8.1. Thus our "stronger" claim above will follow if $\Delta(\alpha)$ holds for all α (since $\Phi = \Phi_\alpha$ for some α). Let us suppose towards a contradiction that $\Delta(\alpha)$ fails for some least α. Fix F, ϕ, and E witnessing the failure, and let D be the ACSA on which F is an ultrafilter. There are then three cases.

Case I: $\alpha = 0$. Since $F \in \Gamma_0(\Theta)$, it accords with χ, so $\chi(D)$ is in F's upwards closure in B. Choose some $e \in \phi[G \cap \operatorname{dom} \phi]$ with $e \leq \chi(D)$. Let $b = \phi^{-1}(e)$. Then $\phi \restriction b = \phi^{+b}$, and $b \in G$. Now since $F \in \bigcup \Theta$, $G \cap E \in \bigcup \Theta$ and so $\chi(E) \in G$; thus $g \equiv \chi(E) \wedge b$ is in G, since G is a filter. It is then easy to check that ϕ^{+g} accords with χ, so $\phi^{+g} \in \Phi_0$. This contradicts the supposition that F, ϕ, E witness $\neg\Delta(0)$.

Case II: α is a successor ordinal. Since F, ϕ, and E supposedly witness $\neg\Delta(\alpha)$, we have for all $b \in G$ with $b \leq \max(\operatorname{dom}\phi)$ that $\phi^{+b} \notin \Phi_\alpha$. By the leastness of α, however, there does exist such a b satisfying $\phi^{+b} \in \Phi_{\alpha-1}$. Indeed, whenever $b \in G$ and $b \leq \max(\operatorname{dom}\phi)$, ϕ^{+b} will copy $G \cap E^{+b}$ to F (by Lemma 8.1), and $\Delta(\alpha-1)$ will hold for F, ϕ^{+b}, and E^{+b} too, so that there will exist $b' \leq b$ satisfying $\phi^{+b'} \in \Phi_{\alpha-1}$. For all such b', $\phi^{+b'}$ is in $\Phi_{\alpha-1}$ but fails to be generically refinable within $\Phi_{\alpha-1}$, and this failure must be witnessed by some C, D, S as in the definition of generic refinability (for isomorphisms); in particular $b' \wedge \chi(C) \neq 0$. Now by genericity we must have $b' \wedge \chi(C) \in G$ for some such b' having associated witnesses C, D, S (note this implies $\chi(C) \in G$). Fix such a b' and associated C, D, S.

Now define $C' \equiv \overline{C \cup E}$. Because $\chi(C), \chi(E) \in G$, stipulation (ii) on χ (from Lemma 4.3) entails $\chi(C') \in G$. Therefore there exists $M \in \Theta$ such that $G \cap C' \in M$. Since $F \in \Gamma(\Theta)$, Lemma 7.5 ensures that it refines generically in $(\Gamma(\Theta), \Theta)$. In particular (reading off the definition of generic refinability for filters) F refines to some $F' \in N \in \Theta$, for some N with $M \in L(N)$, such that $F' \cap S \neq \emptyset$, and $F' \in \Gamma(\Theta)$, and either $\neg\chi(D) \in F'$ or $F' \cap D$ is an L-generic filter on D.

Now by Lemma 8.2, this $F' \in \Gamma(\Theta)$ is a copy of $G \cap C''$ for some $C'' \supseteq C'$, via some ϕ' that refines ϕ. We know $\Delta(\alpha - 1)$ holds, so some ϕ'^{+e} with $e \in G$ and $e \leq \max(\operatorname{dom}\phi')$ satisfies $\phi'^{+e} \in \Phi_{\alpha-1}$. By Lemma 8.1, ϕ'^{+e} copies $C''^{+e} \cap G$ to F'. Consider the restriction ϕ'' of ϕ'^{+e} to some principal ideal of C''^{+e} such that $\max(\operatorname{ran}\phi'')$ is $\leq$ some member of $F' \cap S$. It is then straightforward to check that ϕ'' is a refinement of $\phi^{+b'}$ that satisfies the requirements of generic refinability within $\Phi_{\alpha-1}$ for our fixed C, D and S, contradicting our choices.

Case III: α is a limit ordinal. The argument here is similar. Recall that for limit α, Φ_α is just the intersection of the Φ_β's with $\beta < \alpha$. Recall our fixed F, ϕ, and E witnessing $\neg\Delta(\alpha)$: again, for all $b \in G$ with $b \leq \max(\operatorname{dom}\phi)$, $\phi^{+b} \notin \Phi_\alpha$. But now since α is a limit ordinal, this means for all such b there exists $\gamma < \alpha$, and thus a *least* $\gamma < \alpha$, such that $\phi^{+b} \notin \Phi_\gamma$; moreover this γ must be a successor ordinal. As before, for each such b, some C, D, S will witness $\phi^{+b} \in \Phi_{\gamma-1} \setminus \Phi_\gamma$ (for b's particular γ), so that in particular $b \wedge \chi(C) \neq 0$. As before, we have by genericity that $b \wedge \chi(C) \in G$ for some such b with its associated C, D, S. The contradiction is then obtained just as before, but with γ in place of α. □

Lemma 8.5 *If the F in Lemma 8.4 satisfies $G \cap D \subseteq F$ for some D, then the E and ϕ can be chosen such that ϕ fixes D.*

The proof of Lemma 8.4 showed that, given any choice of ϕ, E such that ϕ copies $G \cap E$ to F, there exists b such that $\phi^{+b} \in \Phi$ and ϕ^{+b} copies $G \cap E^{+b}$ to F. By Corollary 8.3, we may choose a ϕ that fixes D (is a refinement of the identity function on D), and the resulting ϕ^{+b} will also be a refinement of the identity on D. □

Lemma 8.6 *If $\phi \in \Phi$,* $\max(\operatorname{dom}\phi) \in G$, *and ϕ copies* $(G \cap \operatorname{dom}\phi)$ *to F, then* $F \in \Gamma(\Theta)$.

Note this is a converse to Lemma 8.4; in this direction, we will argue much the same way. Let $\Psi(\alpha)$ be the following assertion about an ordinal α:

$\Psi(\alpha) \iff$ *every filter F that is a copy of $G \cap E$ for some E, via some $\phi \in \Phi$, is a member of $\Gamma_\alpha(\Theta)$.*

Suppose toward a contradiction that $\Psi(\alpha)$ fails for some least α. Fix F, E, ϕ witnessing this failure, and let D denote the ACSA on which F is an ultrafilter.

Case I: $\alpha = 0$. Since ϕ accords with χ, max(ran ϕ) $\leq \chi(D)$. Since $\max(\operatorname{dom}\phi) \in G$, max(ran ϕ) $\in \phi[G \cap \operatorname{dom}\phi] \subseteq F$, so $\chi(D)$ is in the upward closure of F, which is just the definition of accordance with χ for F. Also $F \in \bigcup\Theta$ follows from Lemma 4.5 because $G \cap E \in \bigcup\Theta$ (since $\max(\operatorname{dom}\phi) \in G$ and $\max(\operatorname{dom}\phi) \leq \chi(E)$). So $F \in \Gamma_0(\Theta)$ after all.

Case II: α is a successor ordinal. By α's leastness, $F \in \Gamma_{\alpha-1}(\Theta)$; but we have supposed $F \notin \Gamma_\alpha(\Theta)$, so F fails to refine generically in $(\Gamma_{\alpha-1}(\Theta), \Theta)$. Let some D, S, M witness this failure, as in the definition of generic refinability of filters. In particular note that $M \in \Theta$; since also $G \cap E \in \bigcup\Theta$, there must exist by the directedness of $\bigcup\Theta$ (stipulation (ii) on Θ from Lemma 4.2) some $N \in \Theta$ and $C \in \mathsf{ACSAs}(B)$ such that $E \subseteq C$, and $G \cap C \in N$, and $M \in L(N)$. Fix such an N and C. Since $G \cap C \in \bigcup\Theta$, $\chi(C) \in G$, so $\chi(C)$ is compatible with $\max(\operatorname{dom}\phi)$ (which is also in G). Now apply the definition of ϕ's being generically refinable within Φ, using our present C, D and S, to obtain $s \in S$ and $\phi' \in \Phi$ refining ϕ, meeting the several requirements of that definition.

Let $A \in \mathsf{ACSAs}(B)$ have $\operatorname{ran}\phi'$ as a principal ideal. Then A^{+s} also has it as a principal ideal, because $\max(\operatorname{ran}\phi') \leq s$. Now consider the two cases in the final clause of the definition that was just invoked to obtain ϕ'. If $\max(\operatorname{ran}\phi) \leq \neg\chi(D)$, define

$$A' \equiv (A^{+s})^{+\neg\chi(D)};$$

otherwise we define

$$A' \equiv \overline{A^{+s} \cup D}.$$

In both cases we still have $\operatorname{ran}\phi$ a principal ideal of A''; in the latter case this is because $\max(\operatorname{ran}\phi') \wedge [D] \subseteq \operatorname{ran}\phi'$.

Now note that $F' \in \Gamma_{\alpha-1}(\Theta)$, by the way F' was obtained as a copy via $\phi' \in \Phi$ and because $\Psi(\alpha - 1)$ holds. It is then straightforward to verify, by $F' \in \Gamma_{\alpha-1}(\Theta)$, by the requirements that ϕ' meets, and by the definition of A'', that our N, F', and s meet all the requirements of the generic refinability definition for our D, S, M — the requirements that we supposed unsatisfiable.

Case III: There is no Case III here because the least α at which $\Psi(\alpha)$ fails cannot be a limit ordinal: if every filter F of the kind specified in the definition of $\Psi(\alpha)$ is a member of $\Gamma_\beta(\Theta)$ for all $\beta < \alpha$, then $F \in \Gamma_\alpha(\Theta)$, since the latter is just the intersection of all the $\Gamma_\beta(\Theta)$ with $\beta < \alpha$. □

Theorem 8.7 *Let X be the set of all $d \in B$ for which there exists $D \in \mathsf{ACSAs}(B)$, $d \leq \chi(D)$, such that every $\phi \in \Phi$ that fixes D and satisfies* $\max(\mathrm{dom}\,\phi) \wedge d > 0$ *is an identity mapping. Then under Stipulation 3.11, $\|G \in L(\Theta)\| = \bigvee X$.*

It suffices to show that there exists a $d \in X \cap G$ just if a D meeting the condition of Theorem 7.8 exists.

Suppose D witnesses $d \in X$, for some $d \in G$. Thus we have $d \leq \chi(D)$, so that $\chi(D) \in G$ and $D \cap G \in \bigcup \Theta$, and every $\phi \in \Phi$ that fixes D and satisfies $\max(\mathrm{dom}\,\phi) \wedge d > 0$ is an identity mapping. We claim that $D^{+\chi(D)}$ satisfies the condition of Theorem 7.8. We must show (reading off that condition) that $G \cap D^{+\chi(D)} \in \Gamma(\Theta)$ and all refinements $F \in \Gamma(\Theta)$ of $G \cap D^{+\chi(D)}$ are mutually compatible. The first requirement holds by Lemma 7.7. For the second, suppose that some $F, F' \in \Gamma(\Theta)$ violated it and (without loss of generality) that F is not compatible with G. By Lemma 8.4, F would have to be a copy, via some non-identity $\phi \in \Phi$, of $G \cap E$, for some E. By Lemma 8.5, we can also require that ϕ fix $D^{+\chi(D)}$, and hence D. Since ϕ copies $G \cap E$, we have $\max(\mathrm{dom}\,\phi) \in G$; thus $\max(\mathrm{dom}\,\phi) \wedge d$ belongs to G and so must be > 0. This ϕ contradicts our assumption that D witnessed $d \in X$.

Similarly, in the other direction, suppose D satisfying $(G \cap D) \in \Gamma(\Theta)$ meets the condition of Theorem 7.8: every $F \in \Gamma(\Theta)$ that refines $G \cap D$ is of form $G \cap E$ for some $E \supseteq D$. Choose some $d \in G$ that forces this assertion to hold. We claim that D witnesses $d \in X$. If not, some non-identity $\phi \in \Phi$ fixes D and satisfies $\max(\mathrm{dom}\,\phi) \wedge d > 0$. But by Lemma 8.6, if we had $\max(\mathrm{dom}\,\phi) \in G$, ϕ would copy $(G \cap \mathrm{dom}\,\phi)$ to a filter $F \in \Gamma(\Theta)$. Since ϕ is not an identity mapping, this F is not of form $G \cap E$. So $\max(\mathrm{dom}\,\phi) \wedge d$ would force this F to violate the assertion that d supposedly forced. □

9 Negative Results On Self-Construction

We will now establish that the self-collection criteria derived in the preceding sections are not trivial to satisfy. In particular, we will show that any $\mathcal{N}$ meeting Stipulation 3.11 fails to self-collect into $\mathbb{R}(G)$ if B has a "flexible homogeneity" property ensuring that it has certain partial automorphisms. Since this property makes no mention of $\mathcal{N}$ (nor of any structures defined from it, like Θ or χ), one upshot is that, even assuming B yields a rich structure of continua when used for forcing, it may just have too many partial automorphisms, so that no amount of cleverness in designing a B-name for $\mathcal{N}$ can suffice to make $\mathcal{N}$ self-collect into $\mathbb{R}(G)$. We will show that this is the case when B is a Cohen-real or random-real algebra.

9.1 A homogeneity property that precludes self-collection

We now begin arguments culminating with Theorem 9.6, to show that when our boolean algebra B has the "flexible homogeneity" property defined just below, B cannot yield singly-generated continua that self-collect into another such continuum.

Definitions:

The definitions and stipulations of Sections 2, 3, and 4 remain in force (for B, G, $\mathsf{ACSAs}(B)$, $\mathcal{N}$, Θ, χ, etc.).

B is *homogeneous* if it is isomorphic to each of its principal ideals.

B is *ACSA-homogeneous* if it is isomorphic to every principal ideal of every $C \in \mathsf{ACSAs}(B)$.

B is *flexibly homogeneous* if:

(1) B is ACSA-homogeneous;

(2) there exists $C \in \mathsf{ACSAs}(B)$ satisfying $1 \Vdash C \neq_G B$;

(3) for each C satisfying (2), $C \subseteq B$ is not a rigid boolean inclusion (see Section 6);

(4) each isomorphism between any C, C' satisfying (2) extends to an automorphism of B.

We will also apply the above terms to a principal ideal $B \upharpoonright b$ if it meets the relevant definitions when considered as a boolean algebra in its own right (with greatest element b and negation operation derived from B's negation as $x \mapsto b \wedge \neg x$).

Lemma 9.1 *If B is flexibly homogeneous, then each principal ideal of each $C \in \mathsf{ACSAs}(B)$ (including the "improper ideal," C itself) is also flexibly homogeneous.*

Immediate from clause (1) of flexible homogeneity. □

Lemma 9.2 *Given $C \in \mathsf{ACSAs}(B)$ and $b \in B^+$ such that $b \leq \chi(C)$, and assuming that $\|\mathcal{N}$ self-collects into $\mathbb{R}(G)\| = 1$, there exists $E \supset C$ such that*

$$b \wedge \chi(E) \wedge \|E >_G C\| > 0,$$

and each term in that conjunction is a member of E.

The main point here is that no continuum of form $\mathbb{R}(G \cap C)$ can be maximal among continua of that form contained in $\bigcup \mathcal{N}$; there must exist a strictly larger $\mathbb{R}(G \cap E) \subseteq \bigcup \mathcal{N}$.

The key is to show that some $D \in \mathsf{ACSAs}(B)$ must satisfy

$$b \wedge \chi(D) \wedge \|D \not\leq_G C\| > 0.$$

Suppose not. Then for all D,

$$(b \wedge \chi(D)) \Vdash D \leq_G C.$$

From this and the definitions of χ and Θ it follows that

$$b \Vdash (\mathbb{R}(G \cap D) \subseteq \bigcup \mathcal{N}) \Rightarrow (\mathbb{R}(G \cap D) \subseteq \mathbb{R}(G \cap C)),$$

and since that holds for all D, we have

$$b \Vdash \mathbb{R}(\bigcup \mathcal{N}) = \mathbb{R}(G \cap C).$$

From this and $\|\mathcal{N}$ self-collects into $\mathbb{R}(G)\| = 1$ it follows that

$$b \Vdash \mathbb{R}(G \cap C) = \mathbb{R}(G).$$

Now since $b \leq \chi(C)$, it follows from the definitions of Θ and χ in terms of $\mathcal{N}$ that

$$b \Vdash (\exists C')(\mathbb{R}(G \cap C) \subseteq \mathbb{R}(G \cap C') \text{ and } \mathbb{R}(G \cap C') \in \mathcal{N}),$$

which along with the previous step gives us

$$b \Vdash \mathbb{R}(G) \in \mathcal{N},$$

which would violate clause (i) of the definition of self-collection.

The foregoing permits us to fix an ACSA D such that when we define

$$e \equiv b \wedge \chi(D) \wedge \|D \not\leq_G C\|,$$

we have $e > 0$.

Now define $E = \overline{C \cup D}$. Plainly $C \subseteq E$, $C \leq_G E$, and $D \leq_G E$. Since $e \leq \chi(D)$, and $e \leq b \leq \chi(C)$, we have by (ii) of Lemma 4.3 that $e \leq \chi(E)$. We claim that $e \leq \|E >_G C\|$. This follows from the fact (Corollary 3.3) that $<_G$ is guaranteed with boolean value 1 to be a preorder: if some $e' \leq e$ forced $E =_G C$ and we had $e' \in G$, then in $L[G]$ we would have $E =_G C$ and $D \leq_G E$, entailing $D \leq_G C$, contradicting our choice of D.

Thus E meets all our requirements except possibly the requirement (needed for technical simplification later) that b, $\chi(E)$ and $\|E >_G C\|$ be members. To ensure this, simply replace E with $E^{+b+\chi(E)+\|E>_G C\|}$; Lemma 7.4 ensures that E's χ-value stays the same, and clearly the other requirements on E in the statement of the lemma will still hold when E is enlarged. □

Lemma 9.3 *Suppose that B is flexibly homogeneous and that $\|\dot{\mathcal{N}}$ self-collects into $\mathbb{R}(\dot{G})\| =$ 1; if $C \in \mathsf{ACSAs}(B)$ and $g \in G$ were witnesses to $\mathcal{N}$'s self-collection into $\mathbb{R}(G)$ as in Theorem 8.7, then there would exist $\phi \in \Phi_0$ with* $\max(\operatorname{dom}\phi) \leq g$ *that fixes C but is not an identity function.*

Note the fact that C, g witness Theorem 8.7 entails $g \leq \chi(C)$. Consider $B \restriction g$ and $C^{+g} \restriction g$. By Lemma 9.1 they are flexibly homogeneous. By Lemma 9.2, there exists $E \in \mathsf{ACSAs}(B \restriction g)$ such that $E \supset (C^{+g} \restriction g)$, and when we define

$$e \equiv \chi(E) \wedge \|E >_G (C^{+g} \restriction g)\|,$$

we have $0 < e \leq g$, and also $e \in E$.

By Lemma 9.1 $E \restriction e$ is flexibly homogeneous; in particular it satisfies clause (3), so there exists a nontrivial automorphism ϕ of $E \restriction e$ whose restriction to $C^{+e} \restriction e$ is the identity. And we have $\phi \in \Phi_0$ because $\max(\operatorname{ran}\phi) = \max(\operatorname{dom}\phi) = e \leq \chi(E')$, where E' is any ACSA of (all of) B such that $E' \restriction e = E \restriction e$. $\Box$.

Lemma 9.4 *If* $\Phi_1 = \Phi_0$ *then* $\Phi = \Phi_0$.

Suppose $\Phi_1 = \Phi_0$ but there is a least ordinal α (necessarily > 1) such that there exists $\phi \in \Phi_0 \backslash \Phi_\alpha$. Note α cannot be a limit ordinal since in that case Φ_α would be the intersection of the Φ_β for all $\beta < \alpha$. Thus there exists some $\phi \in \Phi_{\alpha-1}$ that fails to refine generically within $\Phi_{\alpha-1}$. However $\Phi_{\alpha-1} = \Phi_0$ and we already know ϕ refines generically within Φ_0, since $\phi \in \Phi_1$. $\Box$

Lemma 9.5 *Suppose* B *is flexibly homogeneous; if* $\|\dot{\mathcal{N}}$ *self-collects into* $\mathbb{R}(\dot{G})\|$ *were equal to* 1*, then* $\Phi = \Phi_0$.

By Lemma 9.4 it suffices to show that any $\phi \in \Phi_0$ satisfies $\phi \in \Phi_1$ — equivalently, refines generically within Φ_0. So fix $\phi \in \Phi_0$, and, for the parameters of the definition of generic refinability, fix an arbitrary dense subset $S \subseteq B^+$ and $C, D \in \mathsf{ACSAs}(B)$, such that when we define

$$a \equiv \chi(C) \wedge \max(\operatorname{dom}\phi),$$

we have $a > 0$. We will show that the ϕ' required in the definition does exist.

Let $C_0 \in \mathsf{ACSAs}(B)$ be such that $\operatorname{dom}\phi$ is a principal ideal of C_0, and likewise for D_0 with $\operatorname{ran}\phi$. (In general, variable names involving "C" will denote ACSA's involved with the domain of the ϕ' we are seeking, and likewise with "D" and the range of ϕ'.) Note that since $\phi \in \Phi_0$, we have $\max(\operatorname{dom}\phi) \leq \chi(C_0)$ and $\max(\operatorname{ran}\phi) \leq \chi(D_0)$.

Define $C_1 \equiv \overline{C_0 \cup C}$. By clause (ii) of Lemma 4.3, $\chi(C_1) = \chi(C_0) \wedge \chi(C)$. Therefore we have $a \leq \chi(C_1)$.

By Lemma 9.2, there exists $E_C \supseteq C_1$ such that, when we set

$$c \equiv a \wedge \chi(E_C) \wedge \|E_C >_G C_1\|,$$

we have $c > 0$, and each conjunct in c's definition is a member of E_C.

Now ϕ^{+c}, the one-element domain refinement of ϕ by c, is well-defined because $c \leq a \leq \max(\operatorname{dom}\phi)$. Note $\max(\operatorname{dom}\phi^{+c}) = c$, and

$$\max(\operatorname{ran}\phi^{+c}) \leq \max(\operatorname{ran}\phi) \leq \chi(D_0).$$

To fix the ACSA E_D of which $\operatorname{ran}\phi'$ will be a principal ideal, we consider two cases.

Case I: $(\chi(D) \wedge \max(\operatorname{ran} \phi^{+c})) > 0$.

In this case let s be any S-member $\leq (\chi(D) \wedge \max(\operatorname{ran} \phi^{+c}))$, and define

$$D_1 \equiv \overline{D_0 \cup D}.$$

By (ii) of Lemma 4.3, $\chi(D_1) = \chi(D_0) \wedge \chi(D)$. Also we have

$$s \leq \chi(D_1) = \chi(D_0) \wedge \chi(D),$$

because

$$s \leq \max(\operatorname{ran} \phi^{+c}) \leq \chi(D_0).$$

Case II: $\chi(D) \wedge \max(\operatorname{ran} \phi^{+c}) = 0$.

In this case let s be any S-member below max(ran ϕ^{+c}), and let $D_1 = D_0$ (so again we have $s \leq \chi(D_1)$).

Conclusion of argument for both cases:

Now use Lemma 9.2 to obtain $E_D \supseteq D_1$ such that when we define

$$d \equiv s \wedge \chi(E_D) \wedge \|E_D >_G D_1\|,$$

we have $d > 0$, and each conjunct in d's definition is a member of E_D.

Consider the isomorphism $(\phi^{+c})_{+d}$. Let C_0' denote its domain, and let D_0' denote its range. Let c' denote $\max(\operatorname{dom}((\phi^{+c})_{+d}))$; note $C_0' = c' \wedge [C_0]$.

Let $E_C' \equiv E_C \restriction c'$, and $E_D' \equiv E_D \restriction d$. We consider E_C' and E_D' as boolean algebras in their own right. It is easily verified that $C_0' \in \mathsf{ACSAs}(E_C')$ and $D_0' \in \mathsf{ACSAs}(E_D')$.

Since $C_1 \supseteq C_0$, $\|C_1 \geq_G C_0\| = 1$, so $c' \leq \|E_C >_G C_0\|$. Thus the boolean 1 of E_C' forces $E_C' >_G C_0'$.

Now E_C' is isomorphic to E_D' by ACSA-homogeneity of B; choose an isomorphism $\xi : E_C' \to E_D'$ witnessing this. Consider the map $\xi \circ ((\phi^{+c})_{+d})^{-1}$ with domain D_0'. Its range is some ACSA $E^\star$ of E_D'. Since C_0' is the range of $((\phi^{+c})_{+d})^{-1}$, and the boolean 1 of E_C' forces $E_C' >_G C_0'$, the boolean 1 of E_D' must (by the fact that ξ is an isomorphism) force $\|E^\star <_G E_D'\|$.

Consider the inverse of $\xi \circ ((\phi^{+c})_{+d})^{-1}$, which maps $E^\star$ onto D_0'. By clause (4) of flexible homogeneity, it extends to an automorphism ζ of E_D'. Now define $\phi' \equiv \zeta \circ \xi$ by composition. It is straightforward to check that ϕ' is the refinement of ϕ demanded by the definition of generic refinability for our chosen parameters S, C, D. $\square$

Theorem 9.6 *If $B \in L$ is a countably-completely-generated boolean algebra that is flexibly homogeneous, and G is a generic filter on B, and $x \in \mathbb{R}(G) \setminus \mathbb{R}(\emptyset)$, no set $\mathcal{N}$ of singly-generated continua in $L[G]$ self-collects into $\mathbb{R}(x)$, and thus no nontrivial set of such continua self-constructs.*

Suppose $\mathcal{N}$ did self-collect into $\mathbb{R}(x)$. We will first "zoom in" (if necessary) from all of B to a principal ideal of some ACSA, in order to reduce our situation to the one analyzed in our previous sections.

By Lemma 3.8 let B' be an ACSA of B such that $\mathbb{R}(G \cap B') = \mathbb{R}(x)$. By Lemma 7.4 we may assume that $\|\mathcal{N}$ self-collects into $\mathbb{R}(G \cap B')\| \in B'$. We wish to consider B' as our "outermost" algebra, and refer to $G \cap B'$ as G'.

Note that in passing to this (possibly) smaller algebra B' we do not bereave ourselves of all forcing names for $\mathcal{N}$, since the "$\mathcal{N} \in L(X)$" clause of the definition of self-collection ensures $\mathcal{N} \in L(\mathbb{R}(x)) = L(G \cap B')$, so that there is a B'-name for $\mathcal{N}$. Thus we assume Θ and χ to be defined relative to B' and to this B'-name.

Finally, in order to have the simplifying situation that "$\mathcal{N}$ self-collects into $\mathbb{R}(G \cap B')$" is forced by 1 (rather than just by some $g \in G$), let us zoom in further to the principal ideal

$$B'' \equiv B' \restriction \|\mathcal{N} \text{ self-collects into } \mathbb{R}(G \cap B')\|,$$

considered as a boolean algebra in its own right.

Now suppose that $D \in \mathsf{ACSAs}(B'')$ and $g \in G'$ witnessed the self-collection of $\mathcal{N}$ into $\mathbb{R}(x)$ in fulfillment of Theorem 8.7. Consider the ACSA D^{+g} and its principal ideal $D^{+g} \restriction g$. By Lemma 9.3, there exists $\phi \in \Phi_0$ that fixes D^{+g}, and is not an identity function. By Lemma 9.5, $\phi \in \Phi$. The existence of such a ϕ contradicts Theorem 8.7. □

9.2 Cohen forcing does not work in our construction

Theorem 9.7 *No nonempty set of singly-generated continua in a Cohen-forcing extension of L self-collects into a singly-generated continuum.*

The boolean completion of a Cohen-forcing poset satisfies the definition of flexible homogeneity, as we will now show by devoting one lemma to each of that definition's four clauses; the theorem just stated will then follow from Theorem 9.6.

Definitions. A *Cohen-forcing algebra* is a complete atomless boolean algebra with a countable dense subset. We will see that there is only one such algebra up to isomorphism. Let K denote it.

A *free boolean algebra on κ generators* is the boolean algebra F_κ having a subset $\{g_\alpha : \alpha < \kappa\}$ (called *free generators*) that (1) finitely generates F_κ in the sense that every F_κ-member is the output of some boolean operation on some finite set of g_α's, and (2) has "mutually independent" members in the sense that the meet of any finite number of g_α's and/or negations thereof will be nonzero (assuming of course that for no α are both g_α and $\neg g_\alpha$ among the conjuncts). Note that F_κ is a complete boolean algebra if and only if κ is finite.

These two defining properties easily entail that whenever A, B are free boolean algebras having sets A_0, B_0 of generators with the same cardinality, any bijection between A_0 and B_0 induces an isomorphism between A and B.

Lemma 9.8 *If K is a Cohen-forcing algebra and $X \subseteq K$ is a (possibly empty) set of free generators for a subalgebra X' that is dense in the complete subalgebra $\overline{X}$ that it completely generates (in K), and $\|\overline{X} \geq_G K\| = 0$, then there exists $Y \subseteq K \setminus X$, such that $X \cup Y$ is a set of free generators of a subalgebra of K that is dense therein.*

By definition K^+ has a countable dense subset $\{q_n : n \in \omega\}$. We will use this subset to iteratively fix K^+-members y_n so that the set $Y = \{y_n : n \in \omega\}$ will be as required by the lemma's statement. Let us use K_n to denote the subalgebra of K that will be completely generated by $X \cup \{y_i : i < n\}$. In the notation for one-element expansions from Section 7, we then have

$$K_0 = \overline{X};$$
$$K_1 = \overline{X}^{+y_0};$$
$$K_n = \overline{X}^{+y_0+y_1+\ldots+y_{n-1}}.$$

By repeated invocation of Lemma 7.4 we have $\|K_n =_G \overline{X}\| = 1$, and thus by transitivity (Corollary 3.3), $\|K_n \geq_G K\| = 0$, for all n.

Let our induction hypothesis be that the set $X \cup \{y_i : i < n\}$ is a set of free generators of a subalgebra that is dense in K_n. This clearly holds when $n = 0$.

We now begin the inductive defintions. At stage $n \geq 0$, let m be least such that no finite plus/minus meet from $X \cup \{y_i : i < n\}$ is $\leq q_m$. Some such q_m must exist, lest K_n^+ be dense in K^+, contradicting $\|K_n \geq_G K\| = 0$. Set $y_n^0 \equiv q_m$ (the superscript "0" here is a secondary index rather than an exponent). If $(y_n^0 \uparrow K_n) = 1$, then $y_n^0 \wedge k > 0$ for all $k \in K_n^+$; for if k were a counterexample, then $1 > \neg k \geq y_n^0$ would contradict $(y_n^0 \uparrow K_n)$'s definition as the *least* K_n-member that is $\geq y_n^0$. So in this case (namely $y_n^0 \uparrow K_n = 1$) we may simply set $y_n \equiv y_n^0$. Otherwise, consider that since $(y_n^0 \uparrow K_n) < 1$, we have $\neg(y_n^0 \uparrow K_n) > 0$. There must exist $k \in K^+$ such that $k \leq \neg(y_n^0 \uparrow K_n)$ but no K_n^+-member is $\leq k$, lest we have

$$\|K_n \geq_G K\| \geq \neg(y_n^0 \uparrow K_n),$$

violating $\|K_n \geq_G K\| = 0$. So let $y_n^1 \equiv y_n^0 \vee k$ for some such k; and, in general, iteratively define $y_n^{\alpha+1} \equiv y_n^\alpha \vee k$ for some k satisfying

$$k \leq \neg(y_n^\alpha \uparrow K_n),$$

and no K_n^+-member is $\leq k$, and define $y_n^\alpha \equiv \bigvee_{\beta<\alpha} y_n^\beta$ for limit ordinals α, until such time as $(y_n^\alpha \uparrow K_n) = 1$. Once this happens, set $y_n \equiv y_n^\alpha$. Note that $y_n \wedge (y_n^0 \uparrow K_n) = q_m$, so that there is a K_{n+1}^+-member that is $\leq q_m$.

It is then a matter of routine boolean algebra to verify that the induction hypothesis will hold at the next step, and that $Y \equiv \{y_n : n \geq 0\}$ will be as promised. □

Corollary 9.9 *A Cohen-forcing algebra K has a countable free subalgebra Y' that is dense in K.*

Invoke Lemma 9.8 for K with X empty. □

Corollary 9.10 *The Cohen-forcing algebra is unique up to isomorphism.*

Say A, B are Cohen-forcing algebras having (by Corollary 9.9) countably infinite free subalgebras A_0, B_0 that are dense in A and B respectively. Any isomorphism between sets of generators for A_0 and B_0 respectively will (as mentioned above) induce an isomorphism between A_0 and B_0; and because A_0 and B_0 are dense in A, B respectively, any isomorphism between A_0 and B_0 will extend uniquely to an isomorphism between A and B. □

Lemma 9.11 *The Cohen-forcing algebra satisfies ACSA-homogeneity (clause 1 of flexible homogeneity).*

This follows from Corollary 9.10 so long as each principal ideal $C \upharpoonright c$ of each ACSA C of a Cohen-forcing algebra K is itself a Cohen-forcing algebra. And this is the case because if X is any countable dense subset of K, then $\{x \uparrow C : x \in X \cap (K \upharpoonright c)\}$ is a countable dense subset of $C \upharpoonright c$. □

Lemma 9.12 *There exists an ACSA C of the Cohen-forcing algebra K satisfying $1 \Vdash C \neq_G K$ (clause 2 of flexible homogeneity).*

Obtain from Corollary 9.9 a countable free subalgebra Y' of K that is dense in K; let $\{y_n : n \in \omega\}$ be a set of free generators for Y'. Let C_{even} denote the ACSA of K completely generated by $\{y_n : n \text{ even}\}$; and likewise for C_{odd} and $\{y_n : n \text{ odd}\}$. It is then straightforward to show that $1 \Vdash C_{even} \not\geq_G C_{odd}$, and therefore $1 \Vdash C_{even} \neq_G K$. □

Lemma 9.13 *If $1 \Vdash C <_G K$ then there exists an automorphism of K that is not an identity mapping but whose restriction to C is an identity mapping (clause (3) of flexible homogeneity).*

Use Lemma 9.8 to obtain a set of free generators of a free subalgebra that is dense in C. Call this set X; and use Lemma 9.8 again, with this X and K, to obtain a set $Y \subseteq K \setminus X$ such that $X \cup Y$ is a set of free generators for a free subalgebra F that is dense in K. The automorphism of $X \cup Y$ obtained by interchanging y and $\neg y$, for some chosen $y \in Y$, and leaving the other generators invariant, induces an automorphism of F, which in turn induces the desired automorphism of K. □

Lemma 9.14 *Each isomorphism ϕ between any $C, C' \in \mathsf{ACSAs}(K)$ satisfying*

$$1 \Vdash C, C' \neq_G K$$

extends to an automorphism of K (clause (4) of flexible homogeneity).

Fix C, C', ϕ. Obtain X, Y for C as in Lemma 9.13. By isomorphy, X's image under ϕ will be a set X' of free generators for a free subalgebra dense in C'. Use Lemma 9.8 with this X' and K to obtain Y'. Let θ be any bijection from the countable set Y of generators onto the countable set Y' of generators. It is straightforward to show that $\phi \cup \theta$ induces an automorphism of K. $\Box$

Now Lemmas 9.11, 9.12, 9.13, 9.14 have established that all the clauses of the definition of flexible homogeneity hold for the Cohen-forcing algebra; so Theorem 9.7 is proved.

9.3 Random-real forcing does not work in our construction

A real number that is "random over L" in the sense established by R. Solovay is one that is interdefinable in a canonical way with a generic filter on the "random-real" algebra, namely the algebra R of Borel subsets of the interval $(0, 1)$ modulo Lebesgue-null sets; see Chapter 15 of [6]. Our goal is to show that R is, like the Cohen algebra, flexibly homogeneous, and so succumbs to Theorem 9.6.

Definitions:

A (strictly positive and normalized, or "probabilistic") *measure* on a boolean algebra B is a real-valued function μ on B that satisfies

(i) $\mu(0) = 0$;

(ii) $\mu(b) > 0$ for all $b \in B^+$;

(iii) for all pairwise incompatible b_n, $n = 0, 1, \ldots$,

$$\mu(\bigvee_{n \geq 0} b_n) = \sum_{n \geq 0} \mu(b_n);$$

(iv) $\mu(1) = 1$.

A *measure algebra* is a complete boolean algebra that carries a measure μ. The random-real algebra R is a countably-completely-generated measure algebra when μ is defined to be the Lebesgue measure (recall that R's members are equivalence classes of sets of reals modulo Lebesgue-null sets).

Two elements $x, y \in R$ are *μ-independent* if $\mu(x \wedge y) = \mu(x)\mu(y)$. Two subsets $X, Y \subseteq R$ are μ-independent if x, y are μ-independent for every pair $x \in X, y \in Y$.

When $X \subseteq R$, a *plus-minus meet from X* is an element of form

$$\pm\, x_1 \wedge \pm x_2 \wedge \cdots \wedge \pm x_n, \tag{8}$$

for some finite $n > 0$ and $x_i \in X$, where $\pm x$ means either x or $\neg x$, and for no i, j is $x_i = \neg x_j$.

We say $X \subseteq R$ has the "2^{-n} *property*" if, whenever x is a plus-minus meet of n different X-members and/or negations thereof, $\mu(x) = 2^{-n}$.

Lemma 9.15 *There exist subsets $\{e_i : i \in \omega\}, \{o_i : i \in \omega\} \subset R$, each having the 2^{-n} property, whose members are free generators of mutually μ-independent free subalgebras F_{even}, F_{odd} (respectively), such that $\overline{F_{even} \cup F_{odd}} = R$.*

It is well known that R has a subset $\{b_i : i \in \omega\}$ that completely generates R and has the 2^{-n} property. For a concrete example, consider that R is completely generated by the rational subintervals of $(0, 1)$; all these subintervals are completely generated just by those whose endpoints are terminating binary decimals, and these in turn are generated by the subsets of form

$$b_i \equiv \bigcup_{0 \leq n < 2^i} \left(\frac{2n}{2^{i+1}}, \frac{2n+1}{2^{i+1}}\right),$$

where $i \in \omega$. (Note that we have, as is customary in discussions of R, stopped writing "modulo null sets" explicitly.) It is easily checked that $\{b_i : i \in \omega\}$ has the 2^{-n} property.

Given these b_i, we first partition them into two subsets by defining, for all i, $e_i \equiv b_{2i}$ and $o_i \equiv b_{2i+1}$. We claim that these e_i and o_i are as required. Let F_{even}, F_{odd} be the subalgebras of R finitely generated by $\{e_i : i \in \omega\}$ and $\{o_i : i \in \omega\}$. Note that we have ensured $\overline{F_{even} \cup F_{odd}} = R$ because $\{b_i : i \in \omega\}$ completely generates R.

Since $\{e_i\}$ is a subset of $\{b_i\}$, it too has the 2^{-n} property; and from this it easily follows that F_{even} is a free subalgebra for which $\{e_i\}$ is a set of free generators. Similarly for F_{odd} and $\{o_i\}$.

It remains to show that F_{even} and F_{odd} are μ-independent, i.e. that $\mu(e \wedge o) = \mu(e)\mu(o)$ for all $e \in F_{even}, o \in F_{odd}$. Fix e and o. Consider that o is the output of a boolean expression using k-many distinct o_i's (and no other elements), namely $o_{i_1}, o_{i_2}, \ldots, o_{i_k}$. These elements are free generators of a finite boolean subalgebra $O_{fin} \subseteq F_{odd}$ having 2^k atoms, each of form

$$a = \pm o_{i_1} \wedge \pm o_{i_2} \wedge \cdots \wedge \pm o_{i_k},$$

(where $\pm x$ means either x or $\neg x$) and each satisfying $\mu(a) = 2^{-k}$. Any element of a finite boolean algebra is uniquely expressible as the join of finitely many atoms; in the case of our $o \in O_{fin}$, say m-many such atoms.

Likewise, e will be a member of some E_{fin} generated by l-many of the e_i's, and will be the join of finitely many (say n-many) E_{fin} atoms b, each satisfying $\mu(b) = 2^{-l}$. By the 2^{-n} property of the original $\{b_i\}$ $(= \{e_i\} \cup \{o_i\})$, $o \wedge e$ is the join of $m \times n$ many mutually incompatible elements of form $a \wedge b$, each satisfying $\mu(a \wedge b) = 2^{-k-l}$. By additivity of μ on incompatible elements, then,

$$\mu(o \wedge e) = (m \times n)2^{-k-l} = (m \times 2^{-k})(n \times 2^{-l}) = \mu(o)\, \mu(e). \quad \square$$

Definitions:

Fix subsets $\{o_i : i \in \omega\}$, $\{e_i : i \in \omega\}$, F_{odd}, F_{even} as in Lemma 9.15.

Let C_{even} denote $\overline{F_{even}}$ (F_{even}'s completion in R); and $C_{odd} = \overline{F_{odd}}$.

If $\{a_n : n < \omega\}$ is a sequence in R, let $\limsup a_n \equiv \bigwedge_{n=0}^{\infty} \bigvee_{k \geq n} a_k$ and $\liminf a_n \equiv \bigvee_{n=0}^{\infty} \bigwedge_{k \geq n} a_k$. If $\limsup_n a_n = \liminf_n a_n = a$, we say that $\{a_n : n < \omega\}$ *converges*, and let $\lim_n a_n = a$.

We restate the following two claims from [6] more or less verbatim, and without proof.

Fact 9.16 (Exercise 30.1 of [6]:) *If X is a subalgebra of a measure algebra R then the complete subalgebra of R σ-generated by X [which is what $\overline{X}$ denotes] consists of all limits of convergent sequences in X.* $\square$

Fact 9.17 (Exercise 30.2 of [6]:) *If μ is a measure on R and $a = \lim_n a_n$, then $\mu(a) = \lim_n \mu(a_n)$.* $\square$

Lemma 9.18 *If X, Y are μ-independent subalgebras of R then their completions $\overline{X}, \overline{Y}$ in R are μ-independent.*

Fix $x \in \overline{X}, y \in \overline{Y}$; we must verify $\mu(x \wedge y) = \mu(x)\mu(y)$. We show that this follows from Facts 9.16 and 9.17. Invoking 9.16, fix $\{x_n\}$ converging to x in the above sense, and $\{y_n\}$ converging to y.

Claim: $\lim_{n\to\infty} \mu(x \wedge \neg x_n) = 0$. Proof: let $\hat{x}_n$ denote $\bigwedge_{k \geq n} x_n$; then since $x = \liminf x_n$, x is the least upper bound of $\hat{x}_1 \leq \hat{x}_2 \leq \ldots$. By μ's countable additivity we must have $\lim_{n\to\infty} \mu(x - \hat{x}_n) = 0$. Now if the claim did not hold, there would exist $\epsilon > 0$ such that $\mu(x \wedge \neg x_n) > \epsilon$ for arbitrarily large n; but then $\mu(x \wedge \neg \hat{x}_n) > \epsilon$ for these same arbitrarily large n, contradiction.

Similarly $\lim_{n\to\infty} \mu(x_n \wedge \neg x) = 0$, lest $\limsup x_n = x$ fail; and likewise

$$\lim_{n\to\infty} \mu(y_n \wedge \neg y) = 0; \quad \lim_{n\to\infty} \mu(y \wedge \neg y_n) = 0.$$

From this it follows by simple boolean manipulation that

$$\lim_{n\to\infty} \mu(((x_n \wedge y_n) \wedge \neg(x \wedge y)) \vee ((x \wedge y) \wedge \neg(x_n \wedge y_n))) = 0.$$

Thus

$$\lim_{n\to\infty} |\mu(x_n \wedge y_n) - \mu(x \wedge y)| = 0$$

and

$$\mu(x \wedge y) = \lim_{n\to\infty} \mu(x_n \wedge y_n) = \lim_{n\to\infty} (\mu(x_n)\mu(y_n)) = \mu(x)\mu(y),$$

where the last equality holds by Fact 9.17. $\square$

Corollary 9.19 *C_{even} and C_{odd} are μ-independent.* $\Box$

Lemma 9.20 *R is ACSA-homogeneous (clause (1) of flexible homogeneity).*

Consider any $C \in \mathsf{ACSAs}(R)$. It is easy to verify conditions (i)-(iv) above for μ's restriction $\mu \restriction C$. The same is true for μ's restriction to any principal ideal $C \restriction c$ (we satisfy (iv) by defining $\mu'(b) \equiv \mu(b)/\mu(c)$, so that that $C \restriction c$, considered as a boolean algebra in its own right, carries the normalized measure μ'). Thus by Maharam's uniqueness-up-to-isomorphism theorem, $C \restriction c$ is isomorphic to R. $\Box$

Lemma 9.21 $1 \Vdash C_{even} \neq_G R$ *(clause (2) of flexible homogeneity).*

Suppose towards a contradiction that $\|C_{even} =_G R\| > 0$. By Lemma 4.9 there exists $b \in R^+$ such that $b \wedge [R] \subseteq b \wedge [C_{even}]$, that is,

$$(\forall o \in R)(\exists z \in C_{even})(b \wedge o = b \wedge z). \tag{9}$$

Recall the free generators o_i of F_{odd}, having the "2^{-n} property," whose completion in R is C_{odd}. Let $n > 0$ be arbitrary, and let M_n denote the set of all plus-minus meets of form

$$a = \pm o_0 \wedge \ldots \wedge \pm o_{n-1}$$

satisfying $b \wedge a > 0$. Enumerate M_n as $\{a_1, a_2, \ldots a_k\}$. Note that $i \neq j \Rightarrow a_i \wedge a_j = 0$.

Now for each $a_i \in M_n$, there exists by (9) some $z_i \in C_{even}$ such that $b \wedge z_i = b \wedge a_i$. For each $a_i \in M_n$ fix some such z_i. If for each i we define

$$z_i' \equiv z_i \wedge \bigwedge_{j \neq i} \neg z_j,$$

we will still have $b \wedge z_i' = b \wedge a_i$. For if $b \wedge z_i' < b \wedge z_i$, then for some $j \neq i$, we have $b \wedge z_i \wedge z_j > 0$. But this contradicts $b \wedge z_i \wedge z_j = b \wedge a_i \wedge a_j \leq a_i \wedge a_j = 0$.

Note that $\{z_i' : i \leq k\}$ is an antichain, and each z_i' belongs to C_{even}. We have:

$$b = (b \wedge a_1) \vee \ldots \vee (b \wedge a_k) = (b \wedge a_1 \wedge z_1') \vee \ldots \vee (b \wedge a_k \wedge z_k') \leq (a_1 \wedge z_1') \vee \ldots \vee (a_k \wedge z_k').$$

It follows that $\mu(b) \leq \mu((a_1 \wedge z_1') \vee \ldots \vee (a_k \wedge z_k')) = \mu(a_1 \wedge z_1') + \ldots + \mu(a_k \wedge z_k')$. Since $a_i \in C_{odd}$ and $z_i' \in C_{even}$, Corollary 9.19 implies

$$\mu(b) \leq \mu(a_1)\mu(z_1') + \ldots + \mu(a_k)\mu(z_k').$$

Now each a_i satisfies $\mu(a_i) = 2^{-n}$ (by the 2^{-n} property); and the z_i' are an antichain so that the mean value of the $\mu(z_i')$ terms in this expression is $\leq 1/k$. Thus we have

$$\mu(b) \leq k(2^{-n})(1/k) = 2^{-n}.$$

Since n was arbitrary, we have $\mu(b) \leq 2^{-n}$ for all n; thus $b = 0$, contradicting $b \in R^+$. $\Box$

To show that R satisfies clauses (3) and (4) of flexible homogeneity, we employ tools from the proof of Maharam's theorem in [6], Chapter 30, in particular the next two lemmas.

Lemma 9.22 (see [6], Lemma 30.5) *Let C be a complete subalgebra of R such that $1 \Vdash C \neq_G R$, and let ν be a measure on C such that $\nu(c) \leq \mu(c)$ for all $c \in C$. Then there exists some $d \in R$ such that*

$$\nu(c) = \mu(d \wedge c) \text{ for all } c \in C.$$

We refer to the proof of [6], Lemma 30.5. Note that $1 \Vdash C \neq_G R$ is equivalent to that lemma's assumption on C by our Lemma 4.9. □

Lemma 9.23 (see [6], Lemma 30.6) *Let C, D be complete subalgebras of R such that $1 \Vdash C, D \neq_G R$; and let f be a μ-preserving isomorphism of C onto D. Then for every $c \in R$ there exist $d \in R$ and a μ-preserving isomorphism $g \supseteq f$ of $\overline{C \cup \{c\}}$, the subalgebra of R generated by C and c, onto $\overline{D \cup \{d\}}$.*

This is a slight variant of [6], Lemma 30.6. Aside from variable symbol changes, our version is less general in that C and D are required to be subalgebras of the same algebra R, and are considered to carry measures that are restrictions of the same measure μ on the outer algebra; our version is more general in that we require that C, D satisfy $1 \Vdash C, D \neq_G R$ rather than the stronger condition that they be σ-generated by fewer than κ generators, where κ is the weight of R (in our case, $\aleph_0$). [6]'s proof uses the latter condition to justify invoking Lemma 9.22 (or rather 30.5 in [6]); but our weaker condition is manifestly sufficient for this purpose, as it is the condition explicitly required in the statement of Lemma 9.22. □

Lemma 9.24 *Let $C_0, D_0 \in \mathsf{ACSAs}(R)$ satisfy $1 \Vdash C_0, D_0 \neq_G R$; then each isometric isomorphism $f_0 : C_0 \to D_0$ extends to an isometric automorphism of R.*

We largely follow the proof of Theorem 30.1 of [6]. Recall our set $\{b_n : n \in \omega\} \subset R$ that completely generates R.

We build $C_0 \subseteq C_1 \subseteq C_2 \subseteq \ldots$ and $D_0 \subseteq D_1 \subseteq D_2 \subseteq \ldots$, and μ-preserving isomorphisms $f_0 \subseteq f_1 \subseteq f_2 \subseteq \ldots$ between the corresponding C_n's and D_n's, using an inductive back-and-forth construction. At stage $n+1$ where n is even, we apply Lemma 9.23 to $f_n : C_n \to D_n$ and $C_{n+1} \equiv \overline{C_n \cup \{b_{n/2}\}}$, obtaining $d_{n+1} \in R$ and the isometric isomorphism $f_{n+1} : C_{n+1} \to D_{n+1} \equiv \overline{D_n \cup \{d_{n+1}\}}$ such that $f_{n+1} \supseteq f_n$. At stage $n+1$ for odd n, we apply Lemma 9.23 to $f_n^{-1} : D_n \to C_n$ and $D_{n+1} \equiv \overline{D_n \cup \{b_{(n-1)/2}\}}$, obtaining $c_{n+1} \in R$ and the isometric isomorphism $f_{n+1}^{-1} : D_{n+1} \to C_{n+1} \equiv \overline{C_n \cup \{c_{n+1}\}}$ such that $f_{n+1}^{-1} \supseteq f_n^{-1}$.

After iterating n through all of $\mathbb{N}$ we set $C_\infty \equiv \bigcup_n C_n, D_\infty \equiv \bigcup_n D_n, f_\infty \equiv \bigcup_n f_n$, so that f_∞ is an isometric isomorphism $C_\infty \to D_\infty$. By Fact 9.16, every element of $\overline{C_\infty}$ is the limit of a convergent sequence in C_∞, and similarly for $\overline{D_\infty}$; so f_∞ extends to a unique isomorphism $f : \overline{C_\infty} \to \overline{D_\infty}$. Because our construction ensured that $b_n \in C_\infty$ and $b_n \in D_\infty$ for all n, and the b_n jointly generate R, we have $\overline{C_\infty} = \overline{D_\infty} = R$, so f is an automorphism of R that extends f_0. Fact 9.17 ensures that f is an isometry. □

Lemma 9.25 *Lemma 9.24 also holds with the word "isometric" stricken (clause (4) of flexible homogeneity).*

Let C_0, D_0, f_0 be as in Lemma 9.24's statement except that f_0 may not preserve μ. By Maharam's theorem, there exists an isometric isomorphism $\phi : C_{even} \to D_0$. By Lemma 9.24, ϕ extends to an isometric automorphism of R. Let F denote the image of C_{odd} under ϕ. Since C_{odd}, C_{even} are μ-independent, D_0 and F must be too, by isometry.

Now define another measure ν on $D_0 \cup F$, by $\nu(f_0(c)) \equiv \mu(c)$ for all $c \in C_0$, and $\nu(a) = \mu(a)$ for all $a \in F$. It follows from the fact that D_0 and F are μ-independent, and jointly generate R, that ν extends to a measure ν' on all of R. Clearly f_0 is an isometric isomorphism with respect to ν'; thus by Lemma 9.24 it extends to an isometric automorphism of R. □

Lemma 9.26 *For each C satisfying $1 \Vdash C \neq_G R$ there exists an automorphism on R that is not an identity mapping but whose restriction to C is an identity mapping (clause (3) of flexible homogeneity).*

Invoke Lemma 9.22 with the measure ν on C defined by $\nu(c) \equiv \mu(c)/2$, to obtain $d \in R$ such that $\mu(d \wedge c) = \nu(c) = \mu(c)/2$ for all $c \in C$. Let $C_0 = D_0 = \overline{C \cup \{d\}}$, and let f_0 be the unique μ-preserving isomorphism $C_0 \to D_0 (= C_0)$ that fixes C and exchanges d and $\neg d$. It is straightforward to verify that this f_0 exists and is unique, using the fact that

$$C_0 = D_0 = \{(c \wedge d) \vee (c' \wedge \neg d) : c, c' \in C\}$$

and the choice of d. Note that C_0 also satisfies $1 \Vdash C_0 \neq_G R$: if some g forced $R \leq_G C_0$ then by Lemma 4.9 at least one of $g \wedge d, g \wedge \neg d$ would force $R \leq_G C$.

Now it suffices to note that f_0 extends to an automorphism on R, by Lemma 9.25. □

Theorem 9.27 *No nonempty set of singly-generated continua in a random-real forcing extension of L self-collects into a singly-generated continuum.*

This now follows from Theorem 9.6 since we have established that R meets all clauses of the definition of flexible homogeneity (Lemmas 9.20, 9.21, 9.25, and 9.26). □

10 Prospects For Self-Collection

Theorem 8.7 showed that (under Stipulation 3.11) a forcing algebra B will allow some $\mathcal{N}$ in the generic extension to self-collect into a singly-generated continuum just if B avoids having partial automorphisms of a particular kind. Cohen forcing and random-real forcing algebras do have such partial automorphisms; we showed in Section 9 that this follows from their satisfaction of the flexible homogeneity property defined there. We now discuss ways that B might avoid this property and so avoid the kind of partial automorphisms that prevent self-collection. In particular we ask how B might violate clause (3) of flexible homogeneity by *rigidly including* subalgebras (a notion introduced in Section 6). The obvious heavy-handed approach would be to use boolean algebras that are themselves rigid, i.e. have *no* nontrivial automorphisms at all. We will see that this approach can yield sets of continua that self-collect, but that it has a difficulty that may prevent it from providing models of the self-construction axioms. We will then consider the question of whether some B might be nonrigid but still rigidly include subalgebras C; this approach is subtler, and we currently have no examples of B instantiating it, but we will discuss a potential path to obtaining examples.

A "rigid" approach.

Consider the following construction of P. Hájek, as described by R. Lubarsky in [8]:

> Hájek forces over L a sequence of reals $\langle a_n | n \in \omega\rangle$ [where a_{n+1} always constructs a_n but not a_{n+2}] which, when coded in some trivial manner into a real α, is a good candidate for the least upper bound of the c[onstructibility]-degrees of the a_n's. "Good candidate" means that any upper bound which does not collapse $\aleph_1$ also constructs the sequence. [...]
>
> Hájek's partial order is the ω-step iteration (with finite support) of the Jensen-Johnsbraten forcing [JJ]. The absoluteness comes about because any pair of JJ-generics (over the same ground model satisfying $V = L$) collapses $\aleph_1$. Also, a JJ-generic G over L is definable in $L[G]$. So if $\aleph_1^V = \aleph_1^L$ then the sequence of generics is definable as the ω-sequence $\langle a_0, a_1, \ldots\rangle$ such that a_{n+1} is the unique set satisfying the appropriate definition over $L[a_n]$.

From this construction we can define a set of distinct continua to serve as our $\mathcal{N}$:

$$\mathcal{N} = \{\mathbb{R}(a_n) : n \in \omega\}.$$

Now consider $\mathbb{R}(\alpha)$, where α is the real number cited above into which all the a_i have been "coded in some trivial manner" (like interleaving digits). We claim that all five requirements are met for $\mathcal{N}$ to self-collect into $\mathbb{R}(\alpha)$. Clearly $\mathbb{R}(\alpha) \notin \mathcal{N}$. $\mathcal{N}$ is linearly ordered (as $\mathbb{R}(a_0) \subset \mathbb{R}(a_1) \subset \ldots$) and so it is directed. By reversing the coding we can reconstruct the sequence of a_i's from α; thus $\mathcal{N} \in L(\mathbb{R}(\alpha))$ holds. For requirement (iv), note that any $x \in \mathbb{R}(\alpha)$ that constructs each a_n also constructs the sequence of all the a_n's (since x does not collapse $\aleph_1$), and so constructs α; thus $\mathbb{R}(x)$ cannot be a *proper* subset of $\mathbb{R}(\alpha)$. Finally, the sequence of a_i's is definable in $L(\bigcup\mathcal{N})$: "... the sequence of generics is definable as the ω-sequence $\langle a_0, a_1, \ldots\rangle$ such that a_{n+1} is the unique set satisfying the appropriate definition over $L[a_n]$;" and α is definable wherever the sequence $\langle a_0, a_1, \ldots\rangle$ is definable, so indeed $\mathbb{R}(\alpha) \in L(\bigcup\mathcal{N})$.

But this set $\mathcal{N} \cup \{\mathbb{R}(\alpha)\}$ of continua is not self-*constructing*: for any finite $n > 0$, $\mathbb{R}(a_n)$'s predecessors do not self-collect into $\mathbb{R}(a_n)$, since

$$\mathbb{R}(\bigcup_{i<n} \mathbb{R}(a_i)) = \mathbb{R}(\mathbb{R}(a_{n-1})) = \mathbb{R}(a_{n-1}) \neq \mathbb{R}(a_n).$$

One avenue for seekng a self-constructing set of continua would thus be to modify this construction to obtain a linearly *but also densely* ordered set of continua, from a forcing algebra that, like Hájek's, is free of "bad" partial automorphisms. Our work in this direction has been inconclusive, but we are prepared to say that the task is difficult at best.

The "rigid inclusion" approach.

A more subtle approach would be to seek a B that does have nontrivial partial automorphisms in Φ_0, but does not cleanly factor into complementary subalgebras in the way that ruled out self-collection in Theorem 9.6, so that $\Phi_1 \neq \Phi_0$. One might try to construct a nonrigid boolean algebra B that *rigidly includes* some ACSA's C, as this was defined in Section 6. We do not know whether such a B exists, but a *-algebraic analog was identified by R. Longo in the 1980s. His paper [7] exhibits inclusions $R \subset M$ of von Neumann algebras such that M's *-automorphisms are determined by their action on R; the same then holds of their respective sets of projection operators, considered as complete atomless orthomodular lattices. We hope to investigate whether the same will hold when we pass to the regular open algebras (i.e. boolean completions) of these lattices, considered as complete boolean algebras. We note in closing that forcing with these lattices can be considered a noncommutative version of random-real forcing.

11 Appendix: Development of Axioms From The Intuitive Description

Recall the definitions of *continuum* and of $\mathbb{R}(X)$ given at the outset. We wish to develop axioms for a set $\mathcal{F}$ of continua that answers to the following intuitive description:

A self-constructing set of continua grows gradually in such a way that each new continuum is just the closure under definable operations of the real numbers that had already arisen in the smaller continua, with neither "work from outside," nor any "background model of set theory," nor any "inexplicable novelty" involved in the course of its growth.

The requirement that each new continuum in $\mathcal{F}$ be "just the closure under definable operations of the real numbers that had already arisen in the smaller continua" has a natural formalization:

Axiom 0: $(\forall Y \in \mathcal{F})(Y = \mathbb{R}(\bigcup\{X \in \mathcal{F} : X \subset Y\}))$.

If in the context of $\mathcal{F}$ we define the *predecessors* of $Y \in \mathcal{F}$ to be those $\mathcal{F}$-members that are proper subsets of Y, we have the following noteworthy corollary to Axiom 0:

Corollary 11.1 *Any $Y \in \mathcal{F}, Y \neq \mathbb{R}(\emptyset)$, has infinitely many predecessors.*

For if $Y \neq \mathbb{R}(\emptyset)$ were a counterexample, there would exist a counterexample $X \neq \mathbb{R}(\emptyset)$ (either a predecessor of Y or Y itself) that either has no predecessors or has $\mathbb{R}(\emptyset)$ as its *only* predecessor. In the first case Axiom 0 implies $X = \mathbb{R}(\emptyset)$, a contradiction; in the second, we have $X = \mathbb{R}(\bigcup\{\mathbb{R}(\emptyset)\}) = \mathbb{R}(\mathbb{R}(\emptyset)) = \mathbb{R}(\emptyset)$, the same contradiction. □

We next try to understand what it might mean formally for $\mathcal{F}$ to grow "gradually" with no "work from outside" involved in the course of its growth. The clearest way to do this is to take a relatively simple model of Axiom 0 and pinpoint the ways it intuitively *fails* to grow this way.

Definitions. A *singly-generated* continuum X is one that satisfies $X = \mathbb{R}(x)$ for some real number x. As in Example 2.1, let G be a generic filter for Cohen forcing (see [6], Chapter 15); let $\mathcal{C}_S$ be the set of singly-generated continua in $L[G]$; let $\mathcal{C}_{\neg S}$ be the set whose members are $\mathbb{R}(\emptyset)$ and all non-singly-generated continua in $L[G]$.

(Note that by Lemma 12.5 we will always have $L(G) = L[G]$ in the cases under consideration in this section; the "parentheses" version is the relevant one for us in general, but the "brackets" version is the natural one in the context of forcing extensions.)

To see that $\mathcal{C}_S$ satisfies Axiom 0, and to weigh how much it deserves to be called "self-constructing," we need some facts about its structure.

Fact 11.2 *Each inner ZFC model of a Cohen forcing extension $L[G]$ has form $L(x)$ for some real number $x \in L[G]$, so (in light of Lemma 12.5) $\mathcal{C}_S$ can equivalently be defined as the set of continua of $L[G]$'s inner ZFC models.*

Thus $L[G]$'s inner ZFC models are all generated by individual real numbers, which can be partitioned into equivalence classes according to which model they generate. These classes are known as "degrees of constructibility" and it is mainly in this guise that this structure has been studied (see [1] for a thorough overview). $\mathcal{C}_S$'s most salient features are these:

Fact 11.3 *$\mathcal{C}_S$ is* densely ordered *([1], Corollary 1.2) in the sense that it has infinitely many intermediate members between any strictly ordered pair $X \subset Y$ of its members, and*

Fact 11.4 *$\mathcal{C}_S$ is* complemented *([1], Theorem 1.1) in the sense that for every pair $X, Y \in \mathcal{C}_S$ such that $X \subseteq Y$, there exists $X' \in \mathcal{C}_S$ such that $\mathbb{R}(X \cup X') = Y$ and $\mathbb{R}(X \cap X') = \mathbb{R}(\emptyset)$; furthermore, when x_1, x_2 are any real numbers witnessing $X = \mathbb{R}(x_1)$ and $X' = \mathbb{R}(x_2)$, we have $Y = \mathbb{R}(\{x_1, x_2\})$.*

Facts 11.3 and 11.4 entail that $\mathcal{C}_S$ does indeed satisfy Axiom 0: every $Y \in \mathcal{C}_S$ is constructible from the union of all its predecessors, because it is constructible (aside from the case $Y = \mathbb{R}(\emptyset)$) just from a *pair* x_1, x_2 of reals chosen from among its predecessors. But we argue that this construction from pairs of predecessors violates the spirit of "gradual growth" in our intuitive description of self-construction. In particular, when Y, X, X' are as in Fact 11.4 and none of them equals $\mathbb{R}(\emptyset)$, we make the following accusations:

—Y is *predetermined* by its predecessors X and X', since although $Y = \mathbb{R}(X \cup X')$ there exist by Fact 11.3 "intermediate" $Z, Z' \in \mathcal{C}_S$ satisfying $X \subset Z \subset Y$ and $X' \subset Z' \subset Y$; the growth of new real numbers from X to Z and from X' to Z' thus seems superfluous for the construction of Y.

—Y has a set of *mutually unrelated* predecessors that construct it: X and X' are unrelated to each other in the strong sense that they are mutually generic, and appealing to this pair's existence in a self-constructing family as justification for Y's existence there evokes "outside work" that is hand-picking X and X', and yoking them together to take their constructive closure.

—Y is *overdetermined* by its predecessors, being constructed independently as $Y = \mathbb{R}(\{x_1, x_2\})$ and as $Y = \mathbb{R}(\{x_3, x_4\})$ for some some real x_1, x_2, x_3, x_4 belonging respectively to four of Y's predecessor continua X_1, X_2, X_3, X_4, such that no X_i includes any of the others. A proof here requires more facts than just 11.3 and 11.4 but follows easily from this example: let $y \subseteq \omega$ be a Cohen real that witnesses Y's being singly generated, so $Y = \mathbb{R}(y)$; let x_1 be the set of y's even members, and x_2 be $y \setminus x_1$; let x_3 be the set of y's members that are divisible by three, and let x_4 be $y \setminus x_3$.

To satisfy our intuitive description, a member X of $\mathcal{F}$ ought instead to be constructed through the coming-together of smaller continua in $\mathcal{F}$ that naturally form a collection and require vanishingly little work to close.

Since we know (Corollary 11.1) that each non-least $X \in \mathcal{F}$ must have infinitely many predecessors, we suggest that continua "naturally form a collection" if they have already come together in $\mathcal{F}$ in all their finite combinations. Furthermore if the continuum X which such a collection generates is to be an $\mathcal{F}$-member, it ought to be possible to reconstruct the collection from X; otherwise some ambient model of set theory outside $L(X)$ would have been needed to specify the collection whose union constructed X. Finally, X ought to be the "limit" or "least upper bound" of the collection in the sense that no strictly smaller continuum $Y \subset X$ contains each member of the collection; otherwise $X \setminus Y$ could be regarded as a "gap" that would require externally-supplied "work" to surmount. If a set $\mathcal{N}$ of continua met these conditions, the emergence of $X = \mathbb{R}(\bigcup \mathcal{N})$ from $\mathcal{N}$ could be seen as the final step of a process each of whose infinitely many previous steps had already been completed, and thus as requiring vanishingly little work.

We propose the following formalization of this suggestion.

Definitions:

A set $\mathcal{N}$ is *directed* (with respect to the inclusion ordering) if $X, Y \in \mathcal{N} \Rightarrow (\exists Z \in \mathcal{N})(X, Y \subseteq Z)$. (This formalizes the idea that $\mathcal{N}$'s members have "already come together in all finite combinations.")

A set $\mathcal{N}$ of continua *self-collects into* X if all of the following hold:

(i) $X \notin \mathcal{N}$;

(ii) $\mathcal{N}$ is directed;

(iii) $\mathcal{N} \in L(X)$;

(iv) $(\neg\exists x \in X)(\bigcup\mathcal{N} \subseteq \mathbb{R}(x) \subset X)$;

(v) $X = \mathbb{R}(\bigcup\mathcal{N})$.

Draft Self-Collection Axiom: $(\forall X \in \mathcal{F})$(some $\mathcal{N} \subseteq \mathcal{F}$ self-collects into X).

Remark 1. Clause (iv) demands that no singly-generated continuum lie "between" $\mathcal{N}$ and X; arguably we should demand that no *non*-singly-generated continuum lie there either, but until this issue can be shown to cause problems, we are content to ignore it for the sake of simplicity.

Remark 2. Clause (iii) of self-collection prevents the Self-Collection Axiom from accessing any "ambient ZF model," beyond the model constructible from X, which would violate our intuitive description. One reason we cannot ignore such matters is discussed under "The Absoluteness Issue" below.

Remark 3. We could combine Axiom 0 with our Draft Self-Collection Axiom by requiring that the set of *all* X's predecessors in $\mathcal{F}$ self-collect into X. This would simplify matters. The reason we do not do this is that the set of predecessors of each $\mathcal{F}$-member would then have to be directed, and our intuitive description seems not to demand this: if a proper subset $\mathcal{N}$ of X's predecessors in $\mathcal{F}$ self-collects into X and constructs other continua $Y, Z \subseteq X$ as "by-products" (so to speak), such that Y and Z have no common superset in $\mathcal{N}$, this does not obviously violate our intuition about self-collection, even if Y and Z happen themselves to be members of $\mathcal{F}$.

Returning to our specific example $\mathcal{C}_S$, might the continuum $\mathbb{R}(x) \in \mathcal{C}_S$ of an unconstructible real x also equal $\mathbb{R}(\bigcup\mathcal{N})$ for some subset $\mathcal{N}$ of its predecessors that self-collects? The answer, given by our Theorem 9.7, is no:

Fact 11.5 (See Theorem 9.7) *Given any $\mathbb{R}(x) \in \mathcal{C}_S \setminus \mathbb{R}(\emptyset)$, no subset of $\mathbb{R}(x)$'s predecessors self-collects into $\mathbb{R}(x)$; therefore $\mathcal{C}_S$ does not satisfy Draft Self-Collection Axiom.*

Now given such an unconstructible $\mathbb{R}(x) \in \mathcal{C}_S$, $\mathbb{R}(x)$ *does* have increasing sequences (of length ω_1) of predecessors in $\mathcal{C}_S$ that meet requirements (i) through (iv) of self-collection, and are unbounded in the set of all its predecessors. So we may well ask, if $\mathcal{N}$ is such a sequence, what its $\mathbb{R}(\bigcup \mathcal{N})$ could be, if not $\mathbb{R}(x)$. The answer is that it is the continuum of an inner model that violates AC (the Axiom of Choice) and is thus not a member of $\mathcal{C}_S$.

This raises the question: in proposing $\mathcal{C}_S$ as a candidate to be a self-constructing set, did we have any justification for excluding this non-singly-generated continuum $\mathbb{R}(\bigcup \mathcal{N})$ from it in the first place? More generally: *which* subsets of a self-constructing set $\mathcal{F}$ get "considered as collections" so as to beget new continua in $\mathcal{F}$ via union-and-constructive-closure? If $\mathcal{F}$ had finite cardinality, we could have appealed to a maximality principle: *every* $\mathcal{F}$-subset undergoes the union-and-closure operation, resulting in an $\mathcal{F}$-member. But since $\mathcal{F}$ has infinite cardinality (Corollary 11.1) there is no clear fact of the matter about what subsets of $\mathcal{F}$ are available to undergo this. In general, it would depend on what "outer" (or "background" or "ambient") ZF model we took $\mathcal{F}$ to reside in—and any appeal to such a model would directly violate the intuition we began with.

Since naive maximality principles are excluded, arbitrariness may seem inescapable. For if $\mathcal{F}$ has subsets that "could" be considered as collections this way, but aren't—so that the continua they construct "would" belong to $\mathcal{F}$, but don't—there is evidently someone on the outside choosing which collections to consider *as* collections, and $\mathcal{F}$ cannot be regarded as *self*-constructing.

A partial solution we propose for this problem is to take the conditions of our Draft Self-Collection Axiom as necessary *and also sufficient* for an $\mathcal{F}$-subset to generate an $\mathcal{F}$-member. In other words, not only must each $\mathcal{F}$-member be constructed by a self-collecting subset of $\mathcal{F}$, but *every such subset* constructs an $\mathcal{F}$-member.

Self-Collection Ax.: $(\forall X \in L(\mathcal{F}))(X \in \mathcal{F} \iff \text{some } \mathcal{N} \subseteq \mathcal{F} \text{ self-collects into } X)$.

Remark. This version of the Self-Collection Axiom may not *fully* resolve the issue of non-absoluteness; see "The absoluteness issue" below.

Finally, our intuitive notion requires a principle of foundation to keep real numbers from emerging through "inexplicable novelty" or "work from outside" during the course of $\mathcal{F}$'s growth. To see how this problem can occur, we will consider a different subset $\mathcal{C}_{\neg S}$ (also defined above, in Example 2.1) of the continua in a Cohen forcing extension $L[G]$.

Fact 11.6 *$\mathcal{C}_{\neg S}$ can be defined equivalently as*

$$\{X \in L[G] : (\exists \mathcal{N} \subseteq \mathcal{C}_S)(\mathcal{N} \text{ self-collects into } X)\}.$$

We have already remarked that the empty set self-collects into the constructible continuum $\mathbb{R}(\emptyset)$. The $L(X)$ of any other continuum $X \in L[G]$ satisfies AC if and only if it has form $\mathbb{R}(x)$ for some real x (Fact 11.2): if it does, Fact 11.5 states that no such $\mathcal{N}$ self-collects into it; otherwise, it is straightforward to verify that $\mathcal{N}$ defined as $\{\mathbb{R}(x) : x \in X\}$ satisfies the definition of self-collection into X. □

$\mathcal{C}_{\neg S}$ satisfies Axiom 0 and the Self-Collection Axiom—but in an underhanded way, as we will now see, with the help of the following lemma:

Lemma 11.7 *For any distinct $Y, Y' \in \mathcal{C}_S$ such that $Y \subset Y'$, there exists in $L(Y')$ an unbounded subset $\mathcal{M}$ of Y''s predecessors (in $\mathcal{C}_S$) that has Y as a member and self-collects into $\bigcup \mathcal{M}$; thus by Fact 11.6, $\bigcup \mathcal{M} \in \mathcal{C}_{\neg S}$.*

By Fact 11.3, $L(Y')$ has no immediate predecessors in $\mathcal{C}_S$. Since $Y' \in \mathcal{C}_S$, the Axiom of Choice holds in $L(Y')$ (Fact 11.2). Thus Zorn's lemma entails there the existence of a maximal directed subset $\mathcal{M}$ of Y''s predecessors in $\mathcal{C}_S$, with $Y \in \mathcal{M}$. $\mathcal{M}$ clearly meets all but the last requirement for self-collection into Y'. But Fact 11.5 implies that $Y' \not\subseteq \mathbb{R}(\bigcup \mathcal{M})$; therefore

$$\mathcal{N} \equiv \{\mathbb{R}(x) : x \in \mathbb{R}(\bigcup \mathcal{M})\}$$

is a subset of Y''s predecessors such that $\mathcal{M} \subseteq \mathcal{N}$ and $\bigcup \mathcal{N} = \mathbb{R}(\bigcup \mathcal{M})$. It is clear that $\mathcal{N}$ is directed (since $x, y \in \mathbb{R}(\bigcup \mathcal{M})$ implies $z \in \mathbb{R}(\bigcup \mathcal{M})$ for some z interdefinable with $\{x, y\}$), so we must have $\mathcal{M} = \mathcal{N}$, lest $\mathcal{N}$ contradict $\mathcal{M}$'s maximality. Thus $\mathbb{R}(\bigcup \mathcal{M}) = \bigcup \mathcal{M}$, and it is easy to verify all the clauses required for $\mathcal{M}$ to self-collect into $\bigcup \mathcal{M}$. $\Box$

Lemma 11.8 *Within the Cohen forcing extension $L[G]$, given an arbitrary set X, some subset of $\mathcal{C}_S$ self-collects into X if and only if some subset of $\mathcal{C}_{\neg S}$ does.*

First suppose $\mathcal{N} \subseteq \mathcal{C}_{\neg S}$ self-collects into X. Define

$$\mathcal{N}' \equiv \{\mathbb{R}(x) : x \in \bigcup \mathcal{N}\} \subseteq \mathcal{C}_S;$$

we will verify that $\mathcal{N}'$ meets all the clauses of the definition of self-collection into X.

(i): $X \notin \mathcal{N}'$ because otherwise $X \in \mathcal{N}$: if $X = \mathbb{R}(x)$ for some $x \in \bigcup \mathcal{N}$, then this x is a member of some $Y \in \mathcal{N}$, so $\mathbb{R}(x) \subseteq Y$; but $Y \subseteq X = \mathbb{R}(x)$, so $Y = X$.

(ii): to verify directedness, take any $\mathbb{R}(x), \mathbb{R}(y) \in \mathcal{N}'$; we know $x, y \in \bigcup \mathcal{N}$, so by $\mathcal{N}$'s directedness, there must exist $Z \in \mathcal{N}$ such that $x, y \in Z$. Then letting $z \in Z$ be any real number interdefinable with $\{x, y\}$, we have $\mathbb{R}(z) \in \mathcal{N}'$, and $\mathbb{R}(x), \mathbb{R}(y) \subseteq \mathbb{R}(z)$.

(iii): follows from the transitivity of the relation "x constructs y." $\mathcal{N}'$ is constructible from $\mathcal{N}$ so the supposition that $\mathcal{N}$ self-collects into X, and in particular that it satisfies the $\mathcal{N} \in L(X)$ clause (iii) of self-collection into X, implies that $\mathcal{N}'$ satisfies this clause too.

(iv): plainly a continuum will be a counterexample to this clause for $\mathcal{N}$ if and only if it is a counterexample for $\mathcal{N}'$.

(v): It is clear that $\bigcup \mathcal{N}' = \bigcup \mathcal{N}$, so certainly $\mathbb{R}(\bigcup \mathcal{N}') = X$.

Conversely, suppose that $\mathcal{N}$ is now a subset of $\mathcal{C}_S$ that self-collects into X. No Y can be maximal in $\mathcal{N}$ since the latter is directed and has no bound in X's singly-generated sub-continua, so for each $Y \in \mathcal{N}$ choose a larger continuum $Y' \in \mathcal{N}$ such that $Y \subset Y'$. By Lemma 11.7 there exists a $\mathcal{C}_{\neg S}$-member M such that $Y \subseteq M \subseteq Y'$. Thus if we define

$$\mathcal{N}' \equiv \{M \in \mathcal{C}_{\neg S} : (\exists Y' \in \mathcal{N})(M \subseteq Y')\},$$

we will have $\bigcup \mathcal{N}' = \bigcup \mathcal{N}$. The other clauses can then be shown as in the proof of this lemma's other direction. $\square$

The above lemma along with Facts 11.5 and 11.6 entails:

Fact 11.9 *$\mathcal{C}_{\neg S}$ satisfies Axiom 0 and the Self-Collection Axiom.* $\square$

Nonetheless, we cannot regard $\mathcal{C}_{\neg S}$ as self-constructing, for it cannot account for *how* any of its real numbers "got constructed." Choose any non-least $X \in \mathcal{C}_{\neg S}$ and any unconstructible real $x \in X$. What accounts for x's presence in X? Could the process of closing X's predecessors' union under definable operations have newly generated x? No, x was already a member of one these predecessors. Consider: X is a non-least $\mathcal{C}_{\neg S}$-member and thus cannot be equal to $\mathbb{R}(x)$; so there exists $y \in X \setminus \mathbb{R}(x)$. $\mathbb{R}(\{x, y\})$ is the least continuum having x and y as members, so $\mathbb{R}(\{x, y\}) \subset X$. Now apply Lemma 11.7 with $\mathbb{R}(x)$ as Y and $\mathbb{R}(\{x, y\})$ as Y'; this yields a $\mathcal{C}_{\neg S}$-member that has x but is a proper sub-continuum of X. Thus any $X \in \mathcal{C}_{\neg S}$ that has x as a member can point to one of its predecessors and say "well, x was already possible there," and so on ad infinitum. The explanatory buck stops nowhere, and x seems like a case of inexplicable novelty, belying the idea that $\mathcal{C}_{\neg S}$ is self-constructing.

To rule such situations out, it would certainly suffice to use the following axiom: "Every real appearing in some member of $\mathcal{F}$ appears at a least such member." This may be unnecessarily strong, though, ruling out some inoffensive cases. Suppose $\mathcal{F}$-members X and Y are least for containing reals x and y respectively, and that neither $x \in Y$ nor $y \in X$; would anything mysterious necessarily be involved if there were two distinct $\mathcal{F}$-members Z_1 and Z_2 that were minimal for containing some z interdefinable with $\{x, y\}$? Our intuition is that there would not; a continuum that is minimal in $\mathcal{F}$ for having z as a member represents a way of building z out of "less complex" reals; the fact that two such ways exist is not at odds with the idea that $\mathcal{F}$ is self-constructing. The idea is violated not when there are *multiple* ways of constructing a real out of less complex reals, but when there are *no* ways. Hence the following axiom:

Foundation Axiom: $(\forall X \in \mathcal{F})(\forall x \in X)(\exists Y \in \mathcal{F})(x \in Y \subseteq X$ and $(\forall Z \in \mathcal{F})(Z \subset Y \Rightarrow x \notin Z))$.

Fact 11.10 *$\mathcal{C}_S$ satisfies our Foundation Axiom but $\mathcal{C}_{\neg S}$ does not.*

$\mathcal{C}_S$ satisfies our Foundation Axiom because for every $x \in \bigcup \mathcal{C}_S$, $\mathbb{R}(x)$ is the least $\mathcal{C}_S$-member having x as a member. $\mathcal{C}_{\neg S}$'s failure to satisfy the Axiom is clear from our discussion leading up to our statement of it. $\square$

Dropping Axiom 0; Defining Self-Construction

It can perhaps be shown that Axiom 0 is implied by our Self-Collection axiom. We have not been able to do so. However, even if there were a counterexample — a continuum $X \in \mathcal{F}$ some subset of whose $\mathcal{F}$-predecessors self-collect into it, but which is not constructed by the union of *all* its $\mathcal{F}$-predecessors — we could choose to regard it as a pathological case of self-collection, but not one that should dissuade us from calling $\mathcal{F}$ self-constructing (assuming it satisfies the Self-Collection and Foundation axioms). Thus, in the interest of simplicity, we prefer to omit Axiom 0 from our definition of self-construction.

Definition. A set $\mathcal{F}$ of continua is *self-constructing* if it satisfies the Self-Collection and Foundation Axioms.

The Absoluteness Issue

Although we will take the Self-Collection and Foundation axioms developed above to define a self-constructing set of continua, there is an issue that threatens to weaken their claim to do so.

To see the issue, suppose $\mathcal{F} \subset \mathcal{F}'$ are distinct sets of continua such that $\mathcal{F}$ satisfies the axioms in $L(\mathcal{F})$, and $\mathcal{F}'$ satisfies the axioms in $L(\mathcal{F}')$. There might exist in $L(\mathcal{F}')$ (or at least it is not obvious why there couldn't exist) a continuum Y, and a subset $\mathcal{N} \subset \mathcal{F}$, neither of which exist in $L(\mathcal{F})$, such that $\mathcal{N}$ self-collects into Y. Y would therefore be a member of $\mathcal{F}'$. But why was it not already a member of $\mathcal{F}$? We in $L(\mathcal{F}')$ "now realize" that the members of $\mathcal{N}$ were ready and waiting to self-collect into Y, but for some reason they did not. This recalls what we said earlier: If $\mathcal{F}$ has subsets that "could" be considered as collections this way, but aren't, so that the continua they construct "would" belong to $\mathcal{F}$, but don't, there is evidently someone on the outside choosing which collections to consider as such, and $\mathcal{F}$ cannot be regarded as *self*-constructing.

What we would like to have, then, is something about $\mathcal{F}$ guaranteeing that no new self-collecting subsets of it can ever be "discovered" later. Note that this requirement is not (at least in its most straightforward form) an axiom, because it is not just claiming that something holds in a particular ZF model. For present purposes, we will simply define a *strongly sealed* self-constructing set as one whose structure underwrites a guarantee of this type, and table the matter of *how* its structure might do this.

Definition. A self-constructing family $\mathcal{F}$ of continua is *strongly sealed* if its satisfaction of the Self-Collection Axiom is upwards-absolute in the sense that it will still hold when "$(\forall X \in L(\mathcal{F}))$" is replaced with "$(\forall X \in L(\mathcal{F}'))$" for any self-constructing family $\mathcal{F}' \supseteq \mathcal{F}$.

We choose to call this property "strongly sealed" in order to leave the term "sealed" available for a property that is weaker and perhaps better aligned with our intuitive notion. To see why the strongly-sealed property may be too strong, consider that cases of self-collection $X = \mathbb{R}(\mathcal{N})$ can fall into two categories. When $X = \bigcup \mathcal{N}$, no new real numbers are generated by self-collection, and we call this *sterile* self-collection. (Note all the self-collection in $\mathcal{C}_{\neg S}$ was sterile.) When X is strictly larger than $\bigcup \mathcal{N}$, we call it *generative* self-collection since new real numbers are generated. We submit that the "discovery" in $\mathcal{F}'$ of new sub-collections of $\mathcal{F}$ that self-collect in the *sterile* way would be innocuous.

Definition. A self-constructing family $\mathcal{F}$ of continua is *sealed* if for no larger self-constructing family $\mathcal{F}' \supseteq \mathcal{F}$ (in any ZF model containing $L(\mathcal{F})$) can there exist $\mathcal{N} \in L(\mathcal{F}')$, $\mathcal{N} \subseteq \mathcal{F}$, such that $\mathcal{N}$ *generatively* self-collects into some $X \notin \mathcal{F}$.

Since this definition of the property of being *sealed* quantifies over sets "in any ZF model containing $L(\mathcal{F})$," there is no straightforward way to axiomatize it. However, we restrict our attention in the body of this paper to self-constructing families that can by obtained by forcing over the constructible universe L; within the scope of this restriction, it may be possible to axiomatize the "sealed" property by quantifying over all the complete boolean algebras in L that could yield families larger than $\mathcal{F}$ if they were to be used for forcing. We will not pursue this here, though.

12 Appendix: Review of Relative Constructibility, Basic Facts

The best exposition of constructibility theory is probably the first pair of chapters in K. Devlin's book [3]. The theory is developed in a similar way but more tersely in Chapter 13 of T. Jech's standard set-theory text [6]. We will not reproduce this material here, for, as Devlin writes in his preface, "Constructibility theory is plagued with a large number of extremely detailed and potentially tedious arguments, involving such matters as investigating the exact logical complexity of various notions of set theory." What we will do here is recall enough of the theory to establish a few basic facts that will be fundamental to our work.

12.1 The Constructible Universe L

The axioms of ZF, Zermelo-Frankel set theory with no Choice axiom, capture faithfully but incompletely the intuitive notion of the hierarchy of all well-founded sets (assuming there is one unique such hierarchy). In particular ZF gives an incomplete answer to the question of what subsets an infinite set can have, and its answer is still incomplete if a Choice axiom (or its negation) is added. Of course ZF's Power Set axiom ensures that "all" subsets of X themselves form a set, and its Separation axioms ensure that any subset of X that can be defined (by a sentence of the formal language of set theory using finitely many set parameters) will be among "all" these subsets. But ZF leaves open the question of what other subsets, not definable in this way, might also exist.

The parsimonious position that no other subsets should exist is embodied in the *constructible universe* L. L is built by transfinite induction using the function $\mathsf{Def}(X)$, which yields the set of all definable subsets of a given set X. More precisely, $\mathsf{Def}(X)$ is the set of all sets of form

$$\{x \in X : \Psi^X(x, c_1, \ldots, c_n)\}$$

where Ψ is a proposition of the language of set theory having one free variable x, and $c_1, \ldots, c_n$ are some constants in X, and Ψ^X means Ψ's relativization to X (any unbounded quantifiers in Ψ are constrained to range over X). L is then defined as the union of all stages L_α of the *constructive hierarchy*:

$L_0 \equiv \emptyset$;

$L_{\alpha+1} \equiv \mathsf{Def}(L_\alpha)$;

$L_\alpha \equiv \bigcup_{\beta<\alpha} L_\beta$ when α is a limit ordinal;

$L \equiv \bigcup_{\alpha \in Ord} L_\alpha$.

L is a model of ZF ([3], Theorem 1.2; [6], Theorem 13.3) and also of the Axiom of Choice ([3], Theorem 3.8; [6], Theorem 13.18).

The construction of L relativizes quite simply to $L(X)$ for arbitrary sets X: specifically, we "seed" the hierarchy at stage 0 by letting $L_0(X)$ be the transitive closure of $\{X\}$. ($\mathsf{TrCl}(\{X\})$ is taken as the seed, rather than X, so that the resulting model will be transitive, which is a requirement to count as an "inner" model of ZF as that term is generally defined.) We then carry out the construction the same way:

$L_0(X) \equiv \mathsf{TrCl}(\{X\})$;

$L_{\alpha+1}(X) \equiv \mathsf{Def}(L_\alpha(X))$;

$L_\alpha(X) \equiv \bigcup_{\beta<\alpha} L_\beta(X)$ for limit α;

$L(X) \equiv \bigcup_{\alpha \in Ord} L_\alpha(X)$.

This definition can be found in [6] as Definition 13.24, where it is noted that $L(X)$ is the smallest inner model of ZF having X as a member. To *prove* that L is the smallest inner model ([6], Theorem 13.16) and that $L(X)$ is the smallest inner model having X as a member, we must show that they can be defined in an *absolute* way.

12.2 Absoluteness of L and $L(X)$

Definition. When $\Psi(x, Y)$ is a proposition of the language of set theory having exactly one free variable (x) and exactly one constant symbol (Y), we call Ψ a *Y-basic proposition*; when Y is understood to denote some specific set, we call $\Psi(x, Y)$ a *Y-basic predicate*; and if furthermore the expression

$$\{x : \Psi(x, Y)\}$$

defines the same set in any standard inner ZF model (having Y as a member!) in which it is evaluated, we call $\Psi(x, Y)$ a *Y-absolute basic predicate.*

Fact 12.1 *There is a Y-basic proposition Δ such that when Y is chosen to refer to any set, $\Delta(x, Y)$ is a Y-absolute basic predicate and $\{x : \Delta(x, Y)\}$ equals $\mathsf{Def}(Y)$ as defined above.*

To establish this fact we will simply reference [6], Lemma 13.14 and the initial segment of Chapter 13 leading up to it; or see [3], Lemma 2.4, and the pages leading up to it (where it is the formula $D(v, u)$ that formalizes Def). The terminology used in both references is slightly different from ours, but in both cases the formula's absoluteness is proved by showing that it has a low enough level of logical complexity, in terms of its quantifiers. We will not enter into the details here. □

The next step in showing the absoluteness of L and $L(Y)$ is to formalize the transfinite construction used in their definitions, and to show that each step preserves absoluteness. We will abstract the proof somewhat so that it can be reused for other hierarchies in our paper.

Lemma 12.2 (Absolute Definability Of Transfinite Hierarchies) *Suppose that Z is some fixed set, that $\Omega(x', Z)$ is a Z-absolute basic predicate for this Z, and that $\Psi(x'', (Y, Z, \beta))$ is a (Y, Z, β)-absolute basic predicate for any ordinal β and any set Y. Then for all ordinals α, the following is a (Z, α)-absolute basic predicate:*

$(\exists f)(f$ *is a function;*

$\mathrm{dom}(f) = \alpha + 1$;

$f(0) = \{x' : \Omega(x', Z)\}$;

$(\forall \beta < \alpha)(f(\beta + 1) = \{x'' : \Psi(x'', (f(\beta), Z, \beta))\})$;

$(\forall \beta \leq \alpha)(\beta$ *is a limit ordinal* $\Rightarrow f(\beta) = \bigcup_{\gamma<\beta} f(\gamma))$;

and $x \in f(\alpha)$).

This also holds when $\bigcup_{\gamma<\beta}$ is replaced with $\bigcap_{\gamma<\beta}$.

Note first that a function f defined by all the above clauses except "$x \in f(\alpha)$" will in fact always exist, by ZF's Axiom Schema of Replacement. Now suppose that for some α, the stated predicate fails to be (Z, α)-absolute. Then some *least* α witnesses this failure. But α cannot be 0 (since $\Omega(x', Z)$ is $(Z, 0)$-absolute), cannot be a successor ordinal (since then $f(\alpha - 1)$ would be $(Z, \alpha - 1)$-absolute by the induction hypothesis and $\Psi(x'', (f(\alpha - 1), Z, \alpha - 1)$ is $(f(\alpha - 1), Z, \alpha - 1)$-absolute), and cannot be a limit ordinal (since the union and intersection operations are absolute). □

Lemma 12.3 *For any ordinal α and any set Z, L_α is definable by an α-absolute basic predicate, and $L_\alpha(Z)$ is definable by a (Z, α)-absolute basic predicate.*

In both cases we plug propositions $\Omega(\cdots)$ and $\Psi(\cdots)$ into the template of Lemma 12.2. In the L_α case, let Z be $\emptyset$, and for $\Omega(x', Z)$ use "$x' \in \emptyset$" or any other always-false assertion, so that L_0 will be the empty set; and for $\Psi(x'', (Y, Z, \beta))$ use the proposition $\Delta(x'', Y)$ that (by Fact 12.1) defines $\mathsf{Def}(Y)$ and is a Y-absolute basic predicate for any Y.

Given any set Z, we obtain the Z-absolute predicate for $L_\alpha(Z)$ the same way, but using the given Z instead of $\emptyset$, and for $\Omega(x', Z)$ using the assertion "$x' \in \mathsf{TrCl}(\{Z\})$," which is clearly definable by a Z-absolute predicate. $\square$

From this lemma, the absoluteness of L and of $L(Z)$ follows: for any sets x and Z, the assertion $(\exists\alpha)(x \in L_\alpha(Z))$ will be either true in all standard inner ZF models, or false in all such models.

The absoluteness of L allows us to use L-members freely in any Y-absolute basic predicate without affecting its Y-absoluteness. More precisely, if $(c_1, \ldots c_n)$ are an ordered n-tuple of sets in L, and for some Y and Ψ, $\Psi(x, (Y, c_1, \ldots, c_n))$ is a $(Y, c_1, \ldots, c_n)$-absolute basic predicate, then

$$\{x : \Psi(x, (Y, c_1, \ldots, c_n))\}$$

defines the same set in *any standard inner model having Y as a member* — we need not specify "... having $(Y, c_1, \ldots, c_n)$ as a member" since every such model has $c_1, \ldots, c_n$ as members. Likewise, if any c_i is a member of $\mathsf{TrCl}(\{Y\})$, we need not specify "... having c_i as a member" once we have specified "having Y as a member." Therefore we may define an easier-to-use but equivalent absoluteness condition (which is the one given in Section 3.2):

Definition. When y is any set, a *y-predicate* is a proposition $\Psi(x, c_1, \ldots, c_n)$ of the language of set theory, having one free variable x and finitely many constant symbols $c_1, \ldots, c_n$, considered together with fixed referents for the c_i *chosen out of* $L \cup \mathsf{TrCl}(\{y\})$. A y-predicate $\Psi(x, c_1, \ldots, c_n)$ is *y-absolute* if

$$\{x : \Psi(x, c_1, \ldots, c_n)\}$$

evaluates to the same set in any standard inner ZF model having y as a member.

Lemma 12.4 *Lemma 12.2 still holds when $\Omega(x', Z)$ is allowed to be a Z-absolute (no longer necessarily basic) predicate $\Omega(x', c_1, \ldots, c_n)$, and $\Psi(x'', (Y, Z, \beta))$ is allowed to be a (Y, Z, β)-absolute predicate $\Psi(x'', c'_1, \ldots, c'_m)$; the predicate guaranteed by the lemma to be (Z, α)-absolute for all α will now be a (Z, α)-absolute predicate using the c_i's and c'_j's as constants.* $\square$

12.3 Proofs of constructibility lemmas from Section 3.2

Proof of Lemma 3.1: This follows from Lemma 12.3, which states the same thing in terms of absolute *basic* predicates. $\square$

Proof of Lemma 3.4:

The lemma's claim is: $z \in L(y)$ if and only if $z = \{x : \Psi(x, c_1, \ldots, c_n)\}$ for some y-absolute predicate $\Psi(x, c_1, \ldots, c_n)$. Note that we consider each c_i's referent as part of the predicate, and we will freely conflate the c_i's with their referents.

Clearly the "if" direction holds since $L(y)$ is a standard inner ZF model having y as a member.

For "only if," suppse $z \in L(y)$. By definition of the $L(y)$ hierarchy, z must appear as member of some lowest level $L_\alpha(y)$, and α cannot be a limit ordinal because limit levels of the hierarchy are just unions of previous levels. If $\alpha = 0$, then by the definition of $L_0(y)$, we have $z \in \mathsf{TrCl}(\{y\})$; thus z itself can legally be used as (the referent of) a constant in a y-absolute predicate. So we can simply let $\Psi(x, z)$ be the proposition "$x \in z$," and we have our y-absolute predicate that defines z.

The remaining case is that α is a successor ordinal. In this case, $L_\alpha(y) = \mathsf{Def}(L_{\alpha-1}(y))$, and therefore z must be defined by some formula Ψ relativized to $L_{\alpha-1}(y)$, using only a finite number of $L_{\alpha-1}(y)$-members $c_1, \ldots, c_n$ as constants.

Now we know (Lemma 3.1) that $L_{\alpha-1}(y)$ is itself definable by a y-absolute predicate using constants only for $\alpha - 1$ and $\mathsf{TrCl}(\{y\})$. Say $L_{\alpha-1}(y)$ is so defined by $\Delta(x, \alpha - 1, \mathsf{TrCl}(\{y\}))$. Then the predicate $\Omega(x, \alpha - 1, \mathsf{TrCl}(\{y\}), c_1, \ldots, c_n)$ defined by

$$\Omega(x) \iff (\exists M)((\forall v)(v \in M \iff \Delta(v, \alpha - 1, \mathsf{TrCl}(\{y\}))), \text{ and } \Psi^M(x, c_1, \ldots, c_n))$$

is a predicate (of free variable x) that defines z and is absolute for standard inner ZF models having $\mathsf{TrCl}(\{y\})$ and all the c_i as members. $\Omega(x)$ need not, however, be a y-predicate (as we have defined the term), since the c_i need not be members of $L \cup \mathsf{TrCl}(\{y\})$.

The key point, though, is that even if (the referent of) some c_i is not a member of $L \cup \mathsf{TrCl}(\{y\})$, it will still be a member of $L(y)$, and will have appeared first at some lower level $L_{\beta+1}(y)$. The task now is to replace any uses of such a constant c_i in Ω with a formula that was used to define its referent at step $\beta+1$ as a subset of $L_\beta(y)$. We will not formalize this replacement; we will simply note that it can be done, and that the result is a predicate $\Omega'(x)$ that will have a new constant c'_1 denoting β, and finitely many $c'_1, \ldots, c'_m$ denoting members of L_β, as new constant parameters. But the predicate will retain its absoluteness, and will still define z, and will no longer use c_i.

We now wish to iterate this replacement to remove any other c_i's — and now also any c'_i's — that are not members of of $L \cup \mathsf{TrCl}(\{y\})$. Doing so yields a tree of constants whose referents first appear at stages $L_\beta(y)$, and whose "replacement constants" appear at *lower levels* of the hierarchy. Because constants are replaced by formulas using constants from *lower* ordinal stages, the replacement process must terminate, with a final finite set of constants all in $L \cup \mathsf{TrCl}(\{y\})$, lest there be an infinite descending chain of ordinals. The final replacement predicate that uses all these constants is y-absolute and defines z. $\square$

12.4 The "square-brackets" version of relative constructibility

There is a second way to relativize the L construction to a set, namely the $L[X]$ construction (note the square brackets). This is in fact more commonly seen than the "parentheses" version, possibly because it is a more natural generalization of the generic extension construction $L[G]$ where G is a generic filter. (When G is a filter, $L[G]$ will indeed be the same whether calculated by the rules of generic extensions or as the hierarchy we will define just below, but we will not prove this equivalence here.)

$L[X]$ is defined as the union of all stages $L_\alpha[X]$, which are themselves defined exactly as the L_α's, but using an "augmented" definable-subsets operation, $\mathsf{Def}_X(Y)$. This operation is in turn defined just like $\mathsf{Def}(Y)$, except that its propositions Ψ may use, in addition to the usual syntax of set theory, an additional one-place predicate "$\underline{v \in X}$" that is true just if v is a member of X. In other words, we do not throw X directly into this construction as a set; instead we allow the definable operations used in the construction to consult a sort of oracle, to ask whether a given set v would be a member of X, if X existed. (And indeed, in certain cases, X will not be a member of $L[X]$.) So defined, $L[X]$ is a standard inner model of ZF, and moreover of ZFC.

Happily, in most of the cases we will consider, we will have $L[X] = L(X)$:

Lemma 12.5 $L[X] = L(X)$ *for any set* $X \subseteq L$.

Fix X such that $X \subseteq L$. The key point is that this ensures $X \in L[X]$, as follows. There must exist α such that $X \subseteq L_\alpha$, lest X be a proper class. By L_α's absoluteness, it is a member of the standard inner ZF model $L[X]$, and therefore of some $L_\beta[X]$. Since $L_\beta[X]$ is transitive, $X \subseteq L_\beta[X]$. It is then a member of $\mathsf{Def}_X(L_\beta[X])$, since it is definable within $L_\beta[X]$ using the "oracle" predicate, namely as $\{v \in L_\beta[X] : \underline{v \in X}\}$. Thus $X \in L_{\beta+1}[X] \subseteq L[X]$.

Now, as remarked above, $L(X)$ is the smallest ZF model (that is well-founded, transitive, and has all the ordinals) that has X as a member; since we have established $L[X]$ is such a model, we must have $L(X) \subseteq L[X]$.

In the other direction we cite Theorem 13.22 (iii) of [6]: If M is an inner model of ZF such that $X \cap M \in M$, then $L[X] \subseteq M$. $L(X)$ is such an M; therefore $L[X] \subseteq L(X)$. (Alternatively we could proceed from the observation that when X is a member of our model, we can replicate the "oracular" predicate $\underline{v \in X}$ with the "normal" membership predicate $v \in X$, and thereby reconstruct the Def_X hierarchy in our model.) □

Full proof of Lemma 5.1:

The "only if" direction of the lemma is what needs to be proved. Assume $G \in L(\Theta)$ and (invoking Lemma 3.4) that G is definable by a Θ-absolute predicate as

$$G = \{x : \Delta(x, c_1, \ldots, c_n)\}.$$

We will now show that there exists a Θ-absolute predicate of the more narrowly-defined form demanded in the lemma's statement, that defines the same set (namely G).

We claim it suffices to show that each c_i used by the given Δ can itself be defined by a Θ-absolute predicate of the more narrowly-defined form. Suppose for all $i \in \{1, \ldots, n\}$ that c_i were defined by a Θ-absolute predicate as

$$c_i = \{x : \Delta_i(x, \Theta, G \cap C_i, c_i')\} \qquad (*)$$

for some C_i with $G \cap C_i \in \bigcup \Theta$, and some constant $c_i' \in L$. By clause (ii) of Lemma 4.2, the ACSA C' defined by

$$C' \equiv \overline{C_1 \cup \ldots \cup C_n}$$

satisfies $G \cap C' \in \bigcup \Theta$. Thus for all i, since $C_i \in L$, we can alter $\Delta_i(\ldots)$ so that it continues to be a Θ-absolute predicate defining c_i, but instead of taking $G \cap C_i$ as a constant parameter, takes $G \cap C'$ and C_i as constant parameters, and uses "$(G \cap C') \cap C_i$" wherever the proposition Δ_i used "$G \cap C_i$". This yields a Θ-absolute predicate

$$\Delta_i'(x, \Theta, G \cap C', C_i, c_i')$$

that defines c_i and has the more narrowly-defined form. Once this is done for all the c_i, we can similarly transform the given $\Delta(\ldots)$ that defines G, replacing each c_i with the Δ_i' expression that defines it, yielding

$$G = \{x : \Delta'(x, \Theta, G \cap C', C_1, \ldots, C_n, c_1', \ldots, c_n')\},$$

where $\Delta'(\ldots)$ has the required narrower form.

To show that each c_i can indeed be defined by some Δ_i as in (*) above, we consider four cases. By definition of Θ-absolute predicate, c_i must be a member of L and/or $\mathsf{TrCl}(\{\Theta\})$.

Case 1 : $c_i \in L$. In this case set $c_i' = c_i$, let C_i be arbitrary, and let $\Delta_i(x, \ldots)$ simply be the proposition $x \in c_i'$.

Case 2: $c_i = \Theta$. In this set $c_i' = \emptyset$, let C_i be arbitrary, and let $\Delta_i(x, \ldots)$ simply be the proposition $x \in \Theta$.

Case 3: $c_i = \theta(X) \in \Theta$ for some continuum X. By Lemma 4.5, c_i will have a member $G \cap C_i$ for some C_i, and c_i will be the set of all copies of this $G \cap C_i$ on subalgebras in $\mathsf{ACSAs}(B)$, via functions in $\mathsf{ParAut}(B)$. Since $\mathsf{ACSAs}(B)$ and $\mathsf{ParAut}(B)$ are constructible sets, c_i is moreover *absolutely* definable from $G \cap C_i$ this way, as

$$c_i = \{x : \Delta_i(x, \Theta, G \cap C_i, \{\mathsf{ACSAs}(B), \mathsf{ParAut}(B)\})\},$$

for an appropriate proposition Δ_i that formally defines being a copy. Thus we can choose this C_i and Δ_i at our step i, and set $c_i' = \{\mathsf{ACSAs}(B), \mathsf{ParAut}(B)\}$.

Case 4: $c_i = F \in \theta(X) \in \Theta$ for some continuum X and some ultrafilter F. This is handled in essentially the same way as Case 3; here c_i' can be $\{C, \phi\}$, where C is the ACSA on which F is an ultrafilter and ϕ is a particular $\mathsf{ParAut}(B)$-member that copies $G \cap C_i$ to F.

Note that all the members of filters F that can appear in Case 4 are B-members and hence constructible; thus all other $c_i \in \mathsf{TrCl}(\{\Theta\})$ would have fallen under Case 1, so there are no more cases. □

13 Appendix: Self-constructing ZF$^-$ models as a solution to Russell's paradox

Since we cannot yet produce a model of the two self-construction axioms, it may seem premature to speculate about whether their conjunction with further axioms is satisfiable. There is, however, a third axiom suggesting a novel foundational approach that would dispatch the Russell-type paradoxes in a robust way. Since interest in this novel approach may serve as extra inducement to find models of the first two axioms, we judge that is worth dedicating a brief appendix to this further axiom.

A member of a self-constructing family $\mathcal{F}$ can be understood as a snapshot taken at a certain point in the organic growth process of the real numbers. With foundational questions in mind, we might consider broadening the scope of this intuitive idea so that it describes a self-constructing mathematical *universe*, each of whose members is the snapshot of *the totality of all mathematical structures* — not just real numbers — that are possible at a given point. In fact, we have already defined "continuum" in terms of just such a totality, namely, the relative constructible hierarchy $L(X)$ that the definition of $\mathbb{R}(X)$ references. Thus the simplest strategy to implement this envisioned broadening would be to keep the same two axioms for the collection $\mathcal{F}$ of continua, but then turn our attention from $\mathcal{F}$ itself to the collection

$$\mathcal{U} \equiv \{L(X) : X \in \mathcal{F}\}$$

of ZF models generated by each continuum in $\mathcal{F}$.

There are two related problems with this strategy; the first one could be called merely technical, but it points directly to the second one, which is fundamental. First, the collection $\mathcal{U}$ just defined cannot be a *set*, because each of its members $L(X)$ contains all the ordinals and so is a proper class. Now, if $L(X)$'s failure to be a set did cause an annoyance in some deduction we would like to make with it, we could presumably use some standard workaround using the fact that "$y \in L(X)$" is first-order expressible using just X as a parameter. But this annoyance would raise the question: are we justified in granting each member $L(X) \in \mathcal{U}$ "all the ordinals" in the first place? No, we claim — this goes against the original intuition we are axiomatizing, namely a collection of mathematical objects that constructs itself without anything given "in advance" or "from outside." In taking $L(X)$ to embody everything that can be constructed from X we are in fact surreptitiously giving ourselves this structure called "all the ordinals," for $L(X)$ is defined as the union of levels

$L_\alpha(X)$ of a relative constructible hierarchy, where α ranges over all the ordinals. How, why, and whence are all these α's given?

To address this concern we must grant ourselves only those levels $L_\alpha(X)$ indexed by α's that are themselves "constructively accessible from X." There are several approaches we could take to make this precise; here we will stick to just one and leave others unmentioned. Our approach applies only to singly generated continua, but since we have already committed to focus on such continua (Stipulation 2.2) this will not be a limitation for us in practice.

Letting X be a singly generated continuum, as witnessed by $X = \mathbb{R}(x)$ for some real x, and letting $\lambda(X)$ denote the smallest uncountable ordinal in $L(X)$, we consider an ordinal α to be *constructively accessible from* X if $\alpha < \lambda(X)$, and we consider the universe of sets $L^-(X)$ corresponding to the continuum X — the "broadening" of X — to be the union of all the $L_\alpha(x)$ for these ordinals:

$$L^-(X) \equiv L_{\lambda(X)}(x).$$

Note that all sets in $L^-(X)$ are countable and that the continuum X is a subset, but not a member, of $L^-(X)$. Thus $L^-(X)$ violates the powerset axiom and is not a model of ZF. It does, however, satisfy all the other axioms of ZF, which are collectively called ZF$^-$.

Lemma 13.1 *For any singly-generated continuum X, $L^-(X)$ is well defined (i.e. does not depend on the choice of x that generates X) and is the smallest standard model of ZF$^-$ whose collection of real numbers is X.*

We omit the proof. □

$L^-(X)$'s failure to satisfy ZF in full does not detract from its suitability to be considered a snapshot of the growth of a self-constructing universe, for the powerset axiom is notoriously non-constructive — that is, it asserts the existence of a set without specifying its members. Nor is the viewpoint unprecedented that all sets are countable in the true mathematical universe. D. Scott, one of the founding fathers of the modern (i.e. forcing) era of set theory, writes in his preface to [2]: "Perhaps we would be pushed in the end to say that all sets are *countable* (and that the continuum is not even a set) when at last all cardinals are absolutely destroyed. But really pleasant axioms have not been produced by me or anyone else, and the suggestion remains speculation. A new idea (or point of view) is needed ..." (D. Scott, in the foreword to [2], p. xv). We submit that the present appendix may provide such a point of view.

Since we are now considering a self-collecting set universe rather than just a self-collecting continuum, we must generalize our Foundation axiom, which ensures that real numbers do not slip in unaccountably, to ensure now that no sets of any kind slip in unaccountably. Concretely, this means replacing X, Y, Z, which represent continua in $\mathcal{F}$ in the original version of the axiom, with $L^-(X), L^-(Y), L^-(Z)$, which represent "totalities of mathematical possibilities." Thus:

Foundation Ax. (general version): $(\forall X \in \mathcal{F})(\forall s \in L^{-}(X))(\exists Y \in \mathcal{F})$
$(Y \subseteq X$ and $s \in L^{-}(Y)$ and $(\forall Z \in \mathcal{F})(Z \subset Y \Rightarrow s \notin L^{-}(Z)))$.

We use the term *Bergson history* (recalling the philosophy to which the creation of genuinely new possibilities is central) to denote a family $\mathcal{F}$ of continua that satisfies the Self-Collection Axiom, the general version of the Foundation Axiom, and the following:

Anti-Paradox Ax.: $(\exists X, Y \in \mathcal{F})(L^{-}(X) \in L^{-}(Y))$.

Let us emphasize how this last axiom dispatches the set-of-all-sets paradoxes: $L^{-}(X)$ is the totality of all sets at one point in the growth process that $\mathcal{F}$ embodies, and this totality *becomes* a set, quite harmlessly, at some later point in this process.

We note that the Anti-Paradox Axiom requires that, for some $X \subseteq Y$ in $\mathcal{F}$, the first uncountable cardinal $\aleph_1$ of $L(X)$ be collapsed to $\aleph_0$ in $L(Y)$. In terms of requirements on our forcing construction, this entails that some $C \in \mathsf{ACSAs}(B)$ must be a cardinal-collapsing algebra, and thus that B itself must be such an algebra.

References

[1] U. Abraham and R. Shore, "The Degrees of Constructibility of Cohen Reals", Proc. London Math. Soc., Volume s3-53, Issue 2, 1 September 1986, Pages 193–208

[2] J. Bell, *Set Theory: Boolean-Valued Models and Independence Proofs*, 3rd ed., Oxford University Press, 2005

[3] K. Devlin, *Constructibility*, Springer-Verlag, 1984

[4] D. H. Fremlin, *Measure Theory*, in the Cloud

[5] P. Hájek, "Some results on degrees of constructibility", in: Müller G.H., Scott D.S. (eds) Higher Set Theory. Lecture Notes in Mathematics, vol 669. Springer, Berlin, Heidelberg, 1978.

[6] T. Jech, *Set Theory*, 3rd Millennium ed., Springer, 2003

[7] R. Longo, "Simple injective subfactors." Adv. in Math., 63(2):152–171, 1987.

[8] R. Lubarsky, "There's No Forcing A Least Upper Bound," *Proceedings of the American Mathematical Society*, Vol. 105, No. 4 (Apr., 1989), pp. 998-1002