

GUIDE TO THE BRISTOL MODEL: GAZING INTO THE ABYSS

ASAF KARAGILA

ABSTRACT. The Bristol model is an inner model of $L[c]$, where c is a Cohen real, which is not constructible from a set. The idea was developed in 2011 in a workshop taking place in Bristol, but was only written in detail by the author in [8]. This paper is meant as a guide for those who want to get a broader view of the construction. We try to provide more intuition that might serve as a jumping board for those interested in this construction and in odd models of ZF. We also correct a few minor issues in the original paper, as well as prove new results. For example, that the Boolean Prime Ideal theorem fails in the Bristol model, as some sets cannot be linearly ordered. In addition to this we include a discussion on Kinna–Wagner Principles, which we think may play an important role in understanding the generic multiverse in ZF.

CONTENTS

1. Introduction	2
2. Symmetric extensions and other technical tools	3
3. The decoding apparatus	11
4. Cerberus, the three-headed guard of the underworld	20
5. Errata to “the Abyss”	23
6. Deconstructing constructibility	25
7. Gaps in the Multiverse	26
8. Kinna–Wagner Principles	27
9. Choice principles in Bristol	30
10. Open questions related to the Bristol model	31
References	33

Date: June 9, 2020.

2010 Mathematics Subject Classification. Primary 03E25; Secondary 03E35.

Key words and phrases. axiom of choice, symmetric extensions, iterations of symmetric extensions, Bristol model, generic multiverse, symmetric multiverse, Kinna–Wagner Principles.

The author was supported by the Royal Society grant no. NF170989.

1. INTRODUCTION

Mathematicians love classifications. We enjoy classifying objects into different categories, and for a good reason. Classifications teach us about abstract properties, and help us deepen our understanding of various objects and theories.

Set theorists are generally interested in models of set theory. If V satisfies ZFC, we want to classify models of set theory which lie between V and some generic extension, $V[G]$. In the case where “set theory” is understood as ZFC, Vopěnka’s theorem tells us exactly what the intermediate models are: they are generic extensions given by subforcings of the forcing which is used to introduce G over V .

On the other hand, when we are interested in classifying arbitrary models of ZF, instead, even if we assume that V satisfied ZFC, the task becomes significantly harder, and dare we say, nigh impossible. For a start, a generic extension of a model of ZFC cannot be a model of $\text{ZF} + \neg\text{AC}$. One might be inclined to say that such intermediate extension would still be a symmetric extension, which is a type of inner model of a generic extension defined using automorphisms of the forcing. While this is true under some additional conditions on the intermediate model, it turns out that if M is an intermediate model between V and $V[G]$, even if M is a symmetric extension of V , it might not be given by any forcing even remotely related to the one for which G was generic.

The reality is that intermediate models of ZF are far wilder than their ZFC-counterparts. The Bristol model is the first explicit example of such a model. This is a model intermediate to $L[c]$, where c is an L -generic Cohen real, which is not constructible from a set, let alone a symmetric extension of L (by any means, not just the Cohen forcing). While there is a semi-canonical Bristol model, modulo a particular choice of c , it is immediate from the construction that, in a very good sense of the word, most models intermediate to $L[c]$ are not even definable. We will clarify on this in section 7.

The idea for this model came about in a small 2011 workshop in Bristol on topics related to the HOD Conjecture. In attendance were Andrew Brooke-Taylor, James Cummings, Moti Gitik, Menachem Magidor, Ralf Schindler, Matteo Viale, Philip Welch, and W. Hugh Woodin, henceforth “the Bristol group”. The details were not written down in full, and the model remained as a folklore rumour until the author’s effort to formalise it. The details of the construction are given in [8], which was part of the author’s Ph.D. dissertation. This paper aims to give a bird’s view of the construction, from three different perspectives (for people coming from different walks of set theory). We will also correct a few minor mistakes in the original paper, and prove a handful of new theorems about the Bristol model, and about models of ZF in general.

1.1. Structure of this Paper. The Bristol model is presented in [8] as an iteration of symmetric extensions, starting from a Cohen real. The idea is to have, at successor steps, a “decoding mechanism” which is a symmetric extension over an intermediate step such that two properties hold: (1) the decoding mechanism has a generic (relative to the intermediate step) in the Cohen extension, and (2) the decoding mechanism only adds subsets of sufficiently high rank.

We will cover the basics of the technical tools in section 2. We will define symmetric extensions, and briefly outline the main ideas related to iterating them (or rather, why it is hard to iterate symmetric extensions). We will also discuss the combinatorial ideas needed for the decoding mechanism, both at successors of limits, as well as double successors.

After covering the preliminary tools, we will present the decoding mechanism, and the generic argument needed for the proof to work. In section 4 we explain

the three different approaches to constructing the Bristol model. All three are equivalent, but for different people some of these might be seen as “more natural” and can help understand the model better. We will not dive into the intimate details, though. The goal of this paper is to serve as a companion, and help provide not only the big picture of the construction, but also serve as a first step towards reading and understanding the construction’s details presented in [8].

Having discussed the construction of the model, we will then point towards some minor gaps and typos in the original [8]. Then we will discuss Kinna–Wagner Principles which we expect to play a role in the study of choiceless models such as the Bristol model. We will make some new observations, and suggest conjectures for future research. Finally, in section 9 we will prove that some sets in the Bristol model cannot be linearly ordered, and therefore the Boolean Prime Ideal theorem is false there. We finish the paper with a long list of open questions related to the Bristol model.

Acknowledgements. The author would like to express his deepest gratitude to Daisuke Ikegami for providing the opportunity to give a long tutorial on the Bristol model during the RIMS Set Theory Workshop in November 2019, “Set Theory and Infinity”, as well as to the audience, who sat and listened, asked questions, and pointed out difficulties, all of which contributed to this paper. We also want to thank David Asperó and Andrés E. Caicedo for providing thorough comments on early versions of this manuscript.

2. SYMMETRIC EXTENSIONS AND OTHER TECHNICAL TOOLS

In this paper, the term *forcing* will denote a preordered set with a maximum, denoted by $\mathbb{1}$, unless explicitly mentioned otherwise.¹ Of course, we will invariably think about a forcing as a partially ordered set, a separative one, in fact, knowing full well that this will not limit our generality.² The elements of $\mathbb{P}$ are called *conditions*, and using them we define $\mathbb{P}$ -names. We refer the reader to any of [4, 6, 11] for the basic methodology of forcing.

Let $\mathbb{P}$ be a notion of forcing, we follow the convention that if $p, q \in \mathbb{P}$, then $q \leq p$ indicates that q is a *stronger* condition, and we will often say that q *extends* p . Two conditions are *compatible* if they have a common extension, and they are *incompatible* otherwise.

Given a collection of $\mathbb{P}$ -names, $\{\dot{x}_i \mid i \in I\}$, that we want to transform into a name, we will denote by $\{\dot{x}_i \mid i \in I\}^\bullet$ the name $\{\langle \mathbb{1}, \dot{x}_i \rangle \mid i \in I\}$, and we say that a name is a $\bullet$ -name when it has this form. This extends naturally to ordered pairs, sequences, functions, etc. With this notation we can easily define the canonical names for ground model sets: $\check{x} = \{\check{y} \mid y \in x\}^\bullet$.

Given two $\mathbb{P}$ -names, $\dot{x}$ and $\dot{y}$, we say that $\dot{x}$ *appears* in $\dot{y}$ if there is some $p \in \mathbb{P}$ such that $\langle p, \dot{x} \rangle \in \dot{y}$. We will use a similar terminology stating that p appears in $\dot{y}$.

2.1. Symmetric extensions. As we remarked, a generic extension of a model of ZFC is again a model of ZFC. Symmetric extensions are intermediate models to generic extensions where the axiom of choice may fail.

Let $\mathbb{P}$ be a forcing, and let π be an automorphism of $\mathbb{P}$. The action of π extends to the $\mathbb{P}$ -names by recursion:

$$\pi\dot{x} = \{\langle \pi p, \pi\dot{y} \rangle \mid \langle p, \dot{y} \rangle \in \dot{x}\}.$$

¹We will, eventually, tickle class forcing.

²We still insist on the preorder definition, as it does make the definition of an iteration significantly more manageable.

Lemma (The Symmetry Lemma). *Let π be an automorphism of a forcing $\mathbb{P}$, and let $\dot{x}$ be a $\mathbb{P}$ -name. For every condition p ,*

$$p \Vdash \varphi(\dot{x}) \iff \pi p \Vdash \varphi(\pi \dot{x}).$$

Fix a group $\mathcal{G} \subseteq \text{Aut}(\mathbb{P})$. We say that $\mathcal{F}$ is a *filter of subgroups* on $\mathcal{G}$ if it is a filter on the lattice of subgroups, namely, it is a non-empty collection of subgroups which is closed under finite intersections and supergroups. We will, unless stated otherwise, assume it is a proper filter, i.e. the trivial group is not in $\mathcal{F}$.³ Finally, $\mathcal{F}$ is *normal* if whenever $\pi \in \mathcal{G}$ and $H \in \mathcal{F}$, then $\pi H \pi^{-1} \in \mathcal{F}$ as well. In most cases we are interested not necessarily in a filter, but in a filter base, and we will ignore the distinction between the two.

Call $\langle \mathbb{P}, \mathcal{G}, \mathcal{F} \rangle$ a *symmetric system* if $\mathbb{P}$ is a notion of forcing, $\mathcal{G}$ is a group of automorphisms of $\mathbb{P}$, and $\mathcal{F}$ is a normal filter of subgroups on $\mathcal{G}$. We shall fix a symmetric system $\langle \mathbb{P}, \mathcal{G}, \mathcal{F} \rangle$ for the rest of this subsection.

For a $\mathbb{P}$ -name, $\dot{x}$, let $\text{sym}_{\mathcal{G}}(\dot{x})$ denote the group $\{\pi \in \mathcal{G} \mid \pi \dot{x} = \dot{x}\}$. If it is the case that $\text{sym}_{\mathcal{G}}(\dot{x}) \in \mathcal{F}$, then we say that $\dot{x}$ is *$\mathcal{F}$ -symmetric*. And similarly, we say that $\dot{x}$ is *hereditarily $\mathcal{F}$ -symmetric* if being $\mathcal{F}$ -symmetric is hereditarily true for all names appearing in $\{\dot{x}\}^\bullet$.

We denote by $\text{HS}_{\mathcal{F}}$ the class of hereditarily $\mathcal{F}$ -symmetric names. We denote by $\Vdash^{\text{HS}}$ the relativisation of the forcing relation to HS : we restrict the quantifiers and free variables to this class. It is not hard to check that the Symmetry Lemma applies for $\Vdash^{\text{HS}}$, provided that we use automorphisms from $\mathcal{G}$.

Theorem. *Let $G \subseteq \mathbb{P}$ be a V -generic filter, and let M denote the interpreted class $\text{HS}_{\mathcal{F}}^G = \{\dot{x}^G \mid \dot{x} \in \text{HS}_{\mathcal{F}}\}$. Then M is a transitive class model of ZF such that $V \subseteq M \subseteq V[G]$. Moreover, $M \models \varphi(\dot{x}^G)$ if and only if there is some $p \in G$ such that $p \Vdash^{\text{HS}} \varphi(\dot{x})$.*

The class M is also called a *symmetric extension* of V . It turns out that M is a symmetric extension of V if and only if $M = V(x)$ for some $x \in V[G]$. We will discuss this in more detail in section 7.

We will omit $\mathcal{G}$ and $\mathcal{F}$ from the notation and terminology when they are clear from context, which is usually what is going to happen.

2.2. Example. Let $\mathbb{P}$ be $\text{Add}(\omega, \omega_1)$. Namely, $p \in \mathbb{P}$ is a finite partial function from $\omega_1 \times \omega \rightarrow 2$. We let our $\mathcal{G}$ be the group of permutations of ω_1 acting on $\mathbb{P}$ in the natural way: $\pi p(\pi \alpha, n) = p(\alpha, n)$. Finally, for $E \subseteq \omega_1$, let $\text{fix}(E)$ denote $\{\pi \in \mathcal{G} \mid \pi \restriction E = \text{id}\}$, and set $\mathcal{F} = \{\text{fix}(E) \mid E \in [\omega_1]^{<\omega_1}\}$.

For every $\alpha < \omega_1$, let $\dot{a}_\alpha$ be the name of the α th Cohen real, $\{\langle p, \check{n} \rangle \mid p(\alpha, n) = 1\}$, and set $\dot{A} = \{\dot{a}_\alpha \mid \alpha < \omega_1\}^\bullet$.

Claim 2.1. *For every $\pi \in \mathcal{G}$ and $\alpha < \omega_1$, $\pi \dot{a}_\alpha = \dot{a}_{\pi \alpha}$. Consequently, $\pi \dot{A} = \dot{A}$. $\square$*

As an immediate corollary, $\dot{a}_\alpha \in \text{HS}$ for each $\alpha < \omega_1$, as witnessed by $\text{fix}(\{\alpha\})$, and so $\dot{A} \in \text{HS}$ as well.

Theorem 2.2. $\mathbb{1} \Vdash^{\text{HS}} \dot{A}$ *cannot be well-ordered. Consequently, the real numbers cannot be well-ordered, and therefore $\mathbb{1} \Vdash^{\text{HS}} \neg \text{AC}$.*

Proof. Let $\dot{f} \in \text{HS}$, and suppose that p is a condition such that $p \Vdash^{\text{HS}} \dot{f}: \dot{A} \rightarrow \check{\eta}$ for some ordinal η . Let E be a countable set such that $\text{fix}(E) \subseteq \text{sym}(\dot{f})$. We may also assume that $\pi \in \text{fix}(E)$ satisfies $\pi p = p$ by adding a finite set to E , and replacing it with $E \cup \{\alpha \mid \exists n \langle \alpha, n \rangle \in \text{dom } p\}$.

³There is no point in using the improper filter when taking a symmetric extension. However, for the sake of generality it should be noted that this can be useful when iterating. We promise to never bring this up in the course of this paper again.

Fix $\alpha \notin E$, as $\mathbb{P}$ is a c.c.c. forcing, the set $X = \{\xi < \eta \mid \exists q \leq p : q \Vdash^{\text{HS}} \dot{f}(\dot{a}_\alpha) = \check{\xi}\}$ is countable. Note that $p \Vdash^{\text{HS}} \dot{f}(\dot{a}_\alpha) \in \check{X}$.

For any $\beta < \omega_1$, let π_β denote the 2-cycle $(\alpha \ \beta)$. Therefore,

$$\pi_\beta p \Vdash^{\text{HS}} \pi_\beta \dot{f}(\pi_\beta \dot{a}_\alpha) \in \pi_\beta \check{X}.$$

Easily, $\pi_\beta \in \text{fix}(E)$ if and only if $\beta \notin E$. So for all $\beta \notin E$, $p \Vdash^{\text{HS}} \dot{f}(\dot{a}_\beta) \in \check{X}$. In particular, p must force that $\dot{f}$ has a countable range. As p and $\dot{f}$ were arbitrary, we in fact have shown that every ordinal-valued function in the symmetric extension defined on A must have a countable range.

To finish the proof we appeal to the c.c.c. of the Cohen forcing again, noting that ω_1 is not collapsed, and therefore $\mathbb{1} \Vdash^{\text{HS}} |\dot{A}| \neq \aleph_0$. Therefore there is no injection from A into the ordinals, as wanted. $\square$

The standard arguments are usually presented in a slightly different way. We usually extend p to a condition q , which decides the value of $\dot{f}(\dot{a}_\alpha)$, and then show that we find $\pi \in \text{fix}(E)$ such that πq is compatible with q , and $\pi \alpha \neq \alpha$. This then shows that no extension of p can force $\dot{f}$ to be injective. Chain condition based arguments are not very common in results of this type, and we hope that this paper will help to popularise the idea.⁴

2.3. Iterations of symmetric extensions. Iterating symmetric extensions is not an easy task. While some ad-hoc constructions can be found in the literature from the very early 1970s,⁵ the only systematic development of such framework was done by the author in [9], and so far only for finite support iterations. The goal of this section is not to fully develop and explain this technique, but instead provide an intuition as to how the technique works, and what are the difficulties that need to be overcome in the case of the Bristol model.

The intuition behind iterations of symmetric extensions is as naive and simple as it can get. We want to have an iteration of forcing notions, in this case with finite support, and we want to identify a class of names which correspond to the intermediate model that we would get if we were to iterate symmetric extensions one step at a time.

Looking at a two-step iteration, say $\mathbb{P} * \dot{\mathbb{Q}}$, we also need to associate $\mathcal{G}$ and $\mathcal{F}$ to $\mathbb{P}$, and names for a group of automorphisms, $\mathcal{H}$, and a filter of subgroups $\mathcal{K}$. Now, a $\mathbb{P} * \dot{\mathbb{Q}}$ -name is going to be in the iterated symmetric extension if its projection to $\mathbb{P}$ is in $\text{HS}_{\mathcal{F}}$, and it is guaranteed to be interpreted as a name in $\text{HS}_{\mathcal{K}}$, that is a hereditarily symmetric name in the second step.

But we can weaken this slightly, and much like we only require that the $\mathbb{P}$ -name is guaranteed to be a name in $\text{HS}_{\mathcal{K}}$, we can require that the $\mathbb{P}$ -name is guaranteed to be equal to some name in $\text{HS}_{\mathcal{F}}$. That we, we are allowed to “mix” names from $\text{HS}_{\mathcal{F}}$ over an antichain.⁶

Consider, for example, Cohen’s first model, this is a model similar to the example above, replacing ω_1 by ω . Namely, we add an ω -sequence of Cohen reals, permute them, and consider finite stabilisers. Let $\dot{a}_n$ be the name for the n th real, then we can define the name

$$\dot{a} = \{\langle p, \check{m} \rangle \mid \exists n(p(n, 0) = 1 \wedge \forall k < n, p(k, 0) = 0 \wedge p(n, m) = 1)\}.$$

⁴See [10] for more examples of this sort.

⁵Examples include [12],[16],[17], and to some extent also [14].

⁶In general a class of names X has the mixing property if when $1 \Vdash \dot{x} \in X$, that is we can find a (pre-)dense set of conditions p and $\dot{x}_p \in X$ such that $p \Vdash \dot{x} = \dot{x}_p$, then $\dot{x} \in X$.

In clearer terms, this is the name for the first $\dot{a}_n$ which contains 0. We are guaranteed that $\dot{a}$ will be interpreted as one of the $\dot{a}_n$. But it is not hard to see that $\dot{a}$, as defined above, is not stabilised by fixing any finite subset of ω .⁷

Considering names like $\dot{a}$ seems like an unnecessary complication. But upon a closer examination of the general construction of iterations, we see this idea is somehow necessary. Indeed, the conditions of the iteration are usually defined to be $\langle p, \dot{q} \rangle$ such that $p \in \mathbb{P}$ and $\mathbb{1}_{\mathbb{P}} \Vdash \dot{q} \in \dot{\mathbb{Q}}$. This definition generalises to non-finite supports. We will call the forcing notions defined this way “Jech(-style) iterations”. Kunen, in his book, introduces iterations with a perhaps more naive approach: the conditions are pairs $\langle p, \dot{q} \rangle$ such that $p \in \mathbb{P}$ and $p \Vdash \dot{q} \in \dot{\mathbb{Q}}$. This definition fails to generalise nicely to infinite supports, but it is useful for understanding finite support iterations. We will use “Kunen(-style) iterations” to refer to forcing iterations defined this way.

Remark 2.3. It is worth pointing out at this point that we assume ZFC holds in our ground model. While the theory of symmetric extensions, as well as that of iterated forcing, can be developed reasonably well in ZF, it is not clear to what extent choice is truly necessary for developing the theory of iterated symmetric extensions. We utilise the mixing property quite significantly in the general theory, and while it is conceivable, and indeed it is our conjecture, that we can remove choice from the assumptions, there is a certain comfort and simplicity in assuming it. Moreover, as we want to start with L as our underlying model, ZFC is already a given. While $V = L$ is far too strong of an assumption, we will see that at the very least we will want GCH to hold, which implies choice anyway.

When working with Jech iterations, we can quickly see that even for the conditions of $\mathbb{P} * \dot{\mathbb{Q}}$ to be in the intermediate model, we need to allow this so-called “mixing property”. And so, if we require it to hold for the conditions, we will need to require it to also hold for the automorphisms of $\dot{\mathbb{Q}}$, etc., and so the idea itself is important. We can now proceed towards finding a combinatorial definition that will allow us to directly define the class of names, much like we did with HS in the case of a single symmetric extension.

Definition 2.4. Let $\mathbb{P}$ be a notion of forcing and let π be an automorphism of $\mathbb{P}$. We say that π *respects* a name $\dot{A}$ if $\mathbb{1}_{\mathbb{P}} \Vdash \pi \dot{A} = \dot{A}$. If $\langle \mathbb{P}, \mathcal{G}, \mathcal{F} \rangle$ is a symmetric system we say that $\dot{A}$ is $\mathcal{F}$ -*respected* if there is some $H \in \mathcal{F}$ such that every $\pi \in H$ respects $\dot{A}$.

Much like the definition of $\mathcal{F}$ -symmetric before it, this definition lends itself to a hereditary version. We will simply say “respected” when $\mathcal{F}$ is clear from context. If $\dot{A}$ carries an implicit structure (e.g., a forcing notion) then this structure is also required to be respected.

The idea is that being respected is “almost” being symmetric. We will soon weaken this property a bit further to accommodate the “mixing property” into the respected names.

Respect is the foremost necessary condition for developing a combinatorial characterisation of iterated symmetric extensions. Given a symmetric system $\langle \mathbb{P}, \mathcal{G}, \mathcal{F} \rangle$ and a $\mathbb{P}$ -name $\dot{\mathbb{Q}}$, in order to define an automorphism of $\mathbb{P} * \dot{\mathbb{Q}}$ using some $\pi \in \mathcal{G}$, the first thing we need to ensure is that π respects $\dot{\mathbb{Q}}$. Otherwise, $\langle p, \dot{q} \rangle \mapsto \langle \pi p, \pi \dot{q} \rangle$ is not an automorphism of $\mathbb{P} * \dot{\mathbb{Q}}$.

⁷This shows that typically HS does not have the mixing property. This is not a bad thing, though, as we are often concerned with particulars when working with symmetric extensions, and having to specify witnesses is a good thing.

Definition 2.5. Let $\langle \mathbb{P}, \mathcal{G}, \mathcal{F} \rangle$ be a symmetric system, and let $\langle \dot{\mathbb{Q}}, \dot{\mathcal{H}}, \dot{\mathcal{K}} \rangle^\bullet$ be a hereditarily respected name for a symmetric system. Suppose that $\pi \in \mathcal{G}$ and $\dot{\sigma}$ is a name such that $\mathbb{1}_{\mathbb{P}} \Vdash \dot{\sigma} \in \dot{\mathcal{H}}$, then we denote by $\int_{\langle \pi, \dot{\sigma} \rangle}$ the automorphism defined by $\langle p, \dot{q} \rangle \mapsto \langle \pi p, \pi(\dot{\sigma} \dot{q}) \rangle$.

Here we utilise the mixing property quite significantly, by defining $\dot{\sigma} \dot{q}$ to be the name guaranteed to be interpreted as the action of $\dot{\sigma}$ on the condition $\dot{q}$. If we were using Kunen iterations, we would have to define $\int_{\langle \pi, \dot{\sigma} \rangle}$ as a partial automorphism which is only defined when $p \Vdash \dot{\sigma} \in \mathcal{G}$. This is not a formal problem, but it makes the actual legwork a lot harder.

We will denote by $\mathcal{G} * \dot{\mathcal{H}}$ the group of all such automorphisms. In [9] we refer to this group as the *generic semidirect product*. Turning our attention to the filters $\mathcal{F}$ and $\dot{\mathcal{H}}$, we define a *support* to be $\langle \dot{H}_0, \dot{H}_1 \rangle$ where both of them are $\mathbb{P} * \dot{\mathbb{Q}}$ -names such that $\Vdash \dot{H}_0 \in \dot{\mathcal{F}}$ and $\dot{H}_1 \in \dot{\mathcal{H}}$. Note that we can always extend whatever condition to decide the actual group $\dot{H}_0$, and likewise we can decide the actual $\mathbb{P}$ -name for $\dot{H}_1$. However, using this approach allows us to take advantage of the mixing property.

We can now define the notion of respect relative to the supports. Namely, there is a pair $\langle \dot{H}_0, \dot{H}_1 \rangle$ such that whenever $\langle p, \dot{q} \rangle \Vdash \langle \pi, \dot{\sigma} \rangle \in \langle \dot{H}_0, \dot{H}_1 \rangle$, which is to say $\langle p, \dot{q} \rangle \Vdash \dot{\pi} \in \dot{H}_0$ and $\dot{\sigma} \in \dot{H}_1$, we have that $\langle p, \dot{q} \rangle \Vdash \int_{\langle \pi, \dot{\sigma} \rangle} \dot{A} = \dot{A}$. Whereas in Definition 2.4 we required that $\dot{H}_0$ would actually be a concrete group, here we allow a bit more leeway. The result is that the respected names, in this sense, are truly the closure of $\text{HS}_{\mathcal{F}}$ under mixing.

A symmetric iteration, or an iteration of symmetric extensions, is defined by specifying a sequence of names for symmetric systems, $\langle \dot{\mathbb{Q}}_\beta, \dot{\mathcal{G}}_\beta, \dot{\mathcal{F}}_\beta \mid \beta < \alpha \rangle$, and defining $\mathbb{P}_\alpha$ as the finite support iteration of the forcings $\dot{\mathbb{Q}}_\beta$; $\mathcal{G}_\alpha$ as the finite support generic semidirect product of the groups $\dot{\mathcal{G}}_\beta$; and $\mathcal{F}_\alpha$ as the collection of all supports.

We define supports in the general case as $\mathbb{P}_\alpha$ -names for sequences $\langle \dot{H}_\beta \mid \beta < \alpha \rangle$ such that $\Vdash_\alpha \dot{H}_\beta \in \dot{\mathcal{F}}_\beta$, and that $\Vdash_\alpha \{ \dot{\beta} \mid \dot{H}_\beta \neq \dot{\mathcal{G}}_\beta \}^\bullet$ is finite. The last part is crucial, as it allows us the flexibility in “knowing something has a finite definition” while not committing to its specifics just yet.

Let IS_α be the class of $\mathbb{P}_\alpha$ -names which are hereditarily respected. This class is now the class of names which will be interpreted in the intermediate model of stage α . So to complete our definition of a symmetric iteration we need to require that $\langle \dot{\mathbb{Q}}_\alpha, \dot{\mathcal{G}}_\alpha, \dot{\mathcal{F}}_\alpha \rangle^\bullet$ is in IS_α , and indeed that it is respected by all automorphisms in $\mathcal{G}_\alpha$, as we pointed out before. We also have a forcing relation $\Vdash_\alpha^{\text{IS}}$ which is defined by relativising the names and quantifiers to IS_α .

Using this definition we can show that when G is V -generic, then IS_α^G is a model of ZF, and if $\alpha = \beta + 1$, then this model is a symmetric extension of $\text{IS}_\beta^{G \restriction \beta}$ using the β th symmetric system. On the other hand, if we want to continue a symmetric iteration, that is of course possible, but we need to make sure that the generic filter used is not only IS_α^G -generic, but rather $V[G]$ -generic, and we may need to shrink the groups $\dot{\mathcal{G}}_\beta$ in a few places. This may present an issue if we want to continue our iteration by infinitely many steps, as we may need to shrink our groups infinitely many times, which might not be possible.

But here we arrive to our first obstacle when applying this to the Bristol model. We just required that the generic filter is not “pointwise”-generic, but rather V -generic for the entire iteration. While in the case of iterated forcing this is not an issue for the successor case,⁸ here we run into the first problem we had defining

⁸Although it can be an issue for the limit case.

the whole apparatus using mixing. The use of mixing allows us to use arbitrary antichains and predense sets to define the names in IS_α , and so we end up with names in IS_α which encode some generic information in them.

The “easy” way to leave this mess is to require that we relativise the definition pointwise. That is, at each step we only take names which are in the intermediate model. But this adds a layer of complexity when defining the actual forcing $\mathbb{P}_\alpha$, or the automorphisms in $\mathcal{G}_\alpha$, or using these in the same manner that we are used to when working with symmetric extensions.⁹ Even worse, while for finite support iterations all of these different constructions are equivalent, this is not the case if we want to extend our definition to other types of iterations.

We take a different route instead. Looking at products as a type of degenerate iterations, we may want to mimic this definition here. But a copy-paste approach is bound to result in just a product of symmetric extensions. While this is fine, it is not what we are looking for. We want to force over the symmetric extensions, but with a “very canonically defined forcing”. The idea is that we want to iterate $\mathbb{P} * \dot{\mathbb{Q}}$, and in $V[G]$, where $G \subseteq \mathbb{P}$, the forcing $\dot{\mathbb{Q}}^G$ is isomorphic to a forcing in V , and to some extent, this isomorphism does not even depend on G . This means that we are really taking a product. But in the symmetric extension given by $\mathbb{P}$, the forcing $\dot{\mathbb{Q}}^G$ is not isomorphic to any partial order in V , maybe because it cannot be well-ordered, or maybe due to a similar consideration. Nevertheless, it is distinct from forcings in V as far as HS^G is concerned.

We say that a symmetric iteration is *productive* if each $\dot{\mathbb{Q}}_\alpha$ is a $\bullet$ -name which allows us to use the Kunen-iterations for $\mathbb{P}_\alpha$, which are simply the products of these names. We also require that the names of $\dot{\mathcal{G}}_\alpha$ and $\dot{\mathcal{F}}_\alpha$ have similar properties, in that they are $\bullet$ -names and that $\mathbb{1}_\alpha$ decides equality related propositions.

Finally, in the definition of a support, needed for the definition of IS_α , we require more. We require that $\langle \dot{H}_\beta \mid \beta < \alpha \rangle$ is an *excellent support*, meaning that $\dot{H}_\beta$ is a name appearing in $\dot{\mathcal{F}}_\beta$, and in particular a $\mathbb{P}_\beta$ -name, and the finite set of non-trivial coordinates is decided in advance. This is in line with the previous demands: everything is decided in advance, this is “almost a product”.

Now that we have a condition on the iteration, we need a condition on the generic filters as well.

Definition 2.6. Suppose that $\langle \mathbb{P}, \mathcal{G}, \mathcal{F} \rangle$ is an iteration of symmetric extensions. We say that $D \subseteq \mathbb{P}$ is *symmetric* if there is an excellent support $\vec{H}$ such that whenever $p \Vdash \vec{\pi} \in \vec{H}$, then $p \Vdash D = \{\int_{\vec{\pi}} q \mid q \in D\}$. In other words, D is stable, as a set, under a large group of automorphisms.

We say that $G \subseteq \mathbb{P}$ is *symmetrically V -generic* if it is a filter meeting every symmetrically dense set in V .

It turns out that symmetrically generic filters are exactly the filters needed to interpret symmetric names. And we have the following theorem for productive iterations:¹⁰

Theorem 2.7. *Let $\mathbb{P}$ be a productive iteration, and let $p \in \mathbb{P}$ and $\dot{x} \in \text{IS}$ be some symmetric name. The following conditions are equivalent:*

- (1) $p \Vdash^{\text{IS}} \varphi(\dot{x})$.
- (2) For every symmetrically V -generic filter G such that $p \in G$, $\text{IS}^G \models \varphi(\dot{x}^G)$.
- (3) For every V -generic filter G such that $p \in G$, $\text{IS}^G \models \varphi(\dot{x}^G)$.

⁹This is not to say that this is not doable, or maybe even a better way of iterating symmetric extensions. We hope that these remarks will inspire more people to work on these problems.

¹⁰These include, of course, actual products, as well as single-step symmetric extensions.

It is not hard to check now that at least for the successor steps, the iteration of “pointwise” symmetrically generic filters is indeed a symmetrically generic filter.

We finish this overview with a preservation theorem.

Theorem 2.8. *Suppose that $\langle \dot{Q}_\alpha, \dot{\mathcal{G}}_\alpha, \dot{\mathcal{F}}_\alpha \mid \alpha < \delta \rangle$ defines a symmetric iteration, and G is V -generic for the iteration. Moreover, assume that for every $\alpha < \delta$, $\mathbb{1} \Vdash_\alpha \dot{Q}_\alpha$ is weakly homogeneous and $\dot{\mathcal{G}}_\alpha$ is rich enough to witness this.¹¹ Let η be an ordinal such that there exists $\alpha_0 < \delta$ that for any $\alpha \in (\alpha_0, \delta)$ the following equality holds:*

$$V_\eta^{\text{IS}_\alpha^{G \restriction \alpha}} = V_\eta^{\text{IS}_{\alpha+1}^{G \restriction \alpha+1}}.$$

Then $V_\eta^{\text{IS}_\delta^G} = V_\eta^{\text{IS}_{\alpha_0+1}^{G \restriction \alpha_0+1}}$. In other words, if no sets of rank $< \eta$ were added at successor steps, none were added at limit steps either.

This theorem is in stark contrast to the familiar case in the usual context of iterated forcing: iterating, with finite support, forcings which are not c.c.c. will collapse cardinals; and iterating non-trivial forcings, even if they are c.c.c., will add Cohen reals at limit steps. But in the case of symmetric iterations, even if the forcings are non-trivial, as long as they are homogeneous and do not add reals, the limit steps will not add reals either.

As a consequence, we can extend our apparatus now to an Ord-length iteration while preserving ZF in the resulting model. Moreover, the result holds for productive iterations with symmetrically generic filters, as one can state it in the language of forcing, rather than talking about V_η of various models. See also §9.2 of [9].

2.4. Permutable families and scales. The key mechanism in the construction of the Bristol model is “decoding a long sequence from a short sequence”. This can mean a sequence of length ω_1 from a Cohen real, or a sequence of length ω_{43} from one of length ω_{42} . We use almost disjoint families in successor steps to repeatedly decode these sequences, and we use a particular type of scale to succeed at this task when we are at limit steps. This will be as good a place as any to remind the reader that we work in ZFC, especially when thinking about these combinatorial objects that are used here.

Definition 2.9. Let κ be a regular cardinal, and fix a family $\mathcal{A} = \{A_\alpha \mid \alpha < \kappa^+\}$ of unbounded subsets of κ such that for $\alpha < \beta$, $\sup(A_\alpha \cap A_\beta) < \kappa$. For permutations $\pi: \kappa \rightarrow \kappa$ and $\Pi: \kappa^+ \rightarrow \kappa^+$ we say that π *implements* Π if $\pi^* A_\alpha =^* A_{\Pi(\alpha)}$ for all $\alpha < \kappa^+$.

Here we use $=^*$ to mean equality up to a bounded subset of κ . Which κ , of course, will be clear from the context, so we will spare the reader from using the symbol $=_{\kappa}^*$, or worse.

We are looking for an abstract property of an almost disjoint family which will ensure that it implements any bounded permutation of κ^+ , that is any permutation of κ^+ which is the identity on a tail can be implemented.

Definition 2.10. Let κ be a regular cardinal, $\{A_\alpha \mid \alpha < \kappa^+\} \subseteq [\kappa]^\kappa$ is called a *permutable family* if it is almost disjoint, that is for $\alpha \neq \beta$, $\sup(A_\alpha \cap A_\beta) < \kappa$, and for every $I \in [\kappa^+]^{<\kappa^+}$ there is a pairwise disjoint family $\{B_\xi \mid \xi \in I\}$ such that $B_\xi =^* A_\xi$ and

$$\alpha \in I \iff A_\alpha \cap \bigcup_{\xi \in I} B_\xi \text{ is unbounded in } \kappa.$$

We call $\{B_\xi \mid \xi \in I\}$ as in this definition a *disjoint approximation*, and in the case where $B_\xi \subseteq A_\xi$, we say it is a *disjoint refinement*.

¹¹We will say in this case that $\mathcal{G}_\alpha$ witnesses the homogeneity of $\mathbb{Q}_\alpha$.

Proposition 2.11. *If $\mathcal{A}$ is a permutable family of subsets of a regular cardinal κ , then it implements every bounded permutation of κ^+ .*

Proof. Let Π be a bounded permutation of κ^+ , fix $\eta < \kappa^+$ such that Π does not move any ordinals above η . Next, set $I = \eta$ and let $\{B_\xi \mid \xi < \eta\}$ be a disjoint refinement. Now let π be the function which is the order isomorphism from B_α to B_β when $\Pi(\alpha) = \beta$, and the identity elsewhere. Easily, π implements Π . $\square$

Having fixed a permutable family, if $\pi: \kappa \rightarrow \kappa$ implements Π , we will denote this by $\iota(\pi) = \Pi$.

Proposition 2.12. *Let κ be a regular cardinal, then a permutable family exists.*

Proof. Let $\langle T_\alpha \mid \alpha < \kappa^+ \rangle$ be a $\subseteq^*$ -increasing family of subsets of κ . Define A_α as $T_{\alpha+1} \setminus T_\alpha$, then $\{A_\alpha \mid \alpha < \kappa^+\}$ is a permutable family. $\square$

Remark 2.13. It should be pointed out that one can construct an increasing family of subsets from a permutable family. Recursively, set $T_0 = A_0$, $T_{\alpha+1} = T_\alpha \cup A_\alpha$, and for limit steps recursively construct T_α as the union of a disjoint refinement of $\{A_\beta \mid \beta < \alpha\}$, which exists by definition of a permutable family. As long as $\alpha < \kappa^+$ we can ensure that these disjoint refinements are also increasing in inclusion, which guarantees that T_α contains previous limit steps.

Definition 2.14. Given a permutable family on a regular cardinal κ , the *derived group* is the group $\mathcal{G}$ of all permutations of κ which implement a bounded permutation of κ^+ . The *derived filter* is the normal filter of subgroups on $\mathcal{G}$ generated by $\text{fix}(\mathcal{B})$, for a disjoint approximation $\mathcal{B}$, where $\text{fix}(\mathcal{B}) = \{\pi \in \mathcal{G} \mid \pi \upharpoonright \bigcup \mathcal{B} = \text{id}\}$.

While these definitions are given for regular cardinals, we will only use them in the basis case and successor case. For the limit case, where κ is a limit cardinal, we need to use a slightly different machinery, as the goal is to coalesce the information from previous steps and use it as a kind of “short sequence”. In some way, inaccessible cardinals are the “simpler case” compared to singular cardinals. Nevertheless, there is no need of separating the two.

Definition 2.15. Let λ be a limit cardinal, and let $\text{SC}(\lambda)$ denote $\{\mu^+ \mid \mu < \lambda\}$. Let $\{f_\alpha \mid \alpha < \lambda^+\}$ be a scale¹² in $\prod \text{SC}(\lambda)$. Given a sequence of permutations $\vec{\pi} = \langle \pi_\theta \mid \theta \in \text{SC}(\lambda) \rangle$ such that π_θ is a permutation of θ , we say that $\vec{\pi}$ *implements* a function $\Pi: \lambda^+ \rightarrow \lambda^+$ if for every $\alpha < \lambda^+$ and for every large enough θ ,

$$f_{\Pi(\alpha)}(\theta) = \pi_\theta f_\alpha(\theta).$$

We call a scale *permutable* if it implements every bounded permutation of λ^+ .

As with the case of permutable families, we denote by $\iota(\vec{\pi})$ the permutation Π that is implemented by $\vec{\pi}$.

The fact that inaccessible cardinals are the “simpler case” can be trivially seen as a consequence of the following proposition, while remembering that working in $V = L$ we always have the wanted cardinal arithmetic.

Proposition 2.16. *Suppose that λ is an inaccessible cardinal and $2^\lambda = \lambda^+$, then there is a permutable scale on λ .¹³*

This can be shown by a simple transfinite recursion, in a very similar fashion to the permutable family case.

¹²We actually only need to have an increasing sequence, it is irrelevant that it is also bounding.

¹³The assumption on 2^λ can be completely removed by simply limiting ourselves to increasing sequences instead of scales. Nevertheless, as we are working under GCH anyway, this is just simpler.

Proposition 2.17. *Suppose that λ is a singular cardinal and $\square_\lambda^*$, then there is a permutable scale on λ .*

The proof of this proposition can be found as Theorem 3.27 in [8]. The idea of the proof goes back to the Bristol group, and utilises the work of Cummings, Foreman, and Magidor in [2] where it is shown that $\square_\lambda^*$ implies the existence of “better scales”. The aforementioned Theorem 3.27 show that a better scale is in fact permutable.

The key point in proving a better scale is a permutable scale is that given any $I \in [\lambda^+]^{<\lambda^+}$, we can find a function $d: I \rightarrow \text{SC}(\lambda)$ such that $\{f_\alpha[d(\alpha), \lambda] \mid \alpha \in I\}$ is a family of pairwise disjoint sets. The proof of Theorem 3.27 also shows that even if we are only allowing π_θ to be a bounded permutation of θ , this is still enough to ensure that we can implement every bounded permutation of λ^+ using a permutable scale. We say that a sequence of permutation groups of each $\theta \in \text{SC}(\lambda)$ is *rich enough* if we can require π_θ to be in the relevant group when we find an implementing sequence.

Definition 2.18. Let λ be a limit cardinal, and for every $\theta \in \text{SC}(\lambda)$, let $\mathcal{G}_\theta$ be a rich enough group of permutations of θ . The *derived group* is the subgroup $\mathcal{G}$ of the full support product $\prod_{\theta \in \text{SC}(\lambda)} \mathcal{G}_\theta$ consisting of all sequences $\vec{\pi} = \langle \pi_\theta \mid \theta \in \text{SC}(\lambda) \rangle$ which implement a bounded permutation of λ^+ .

For $\eta < \lambda^+$ and $f \in \prod \text{SC}(\lambda)$ we let $K_{\eta, f}$ be the group

$$\{\vec{\pi} \in \mathcal{G} \mid \iota(\vec{\pi}) \restriction \eta = \text{id} \text{ and for all } \theta \in \text{SC}(\lambda), \pi_\theta \restriction f(\theta) = \text{id}\}.$$

The *derived filter* is the filter generated by $\{K_{\eta, f} \mid \eta < \lambda^+, f \in \prod \text{SC}(\lambda)\}$.

We can weaken the definition of $K_{\eta, f}$ and replace η by a bounded subset of λ^+ , i.e., $I \in [\lambda^+]^{<\lambda^+}$, and replace f by a sequence of bounded sets of each θ . But as the definition is complicated enough as it is, it is easier to just use η and f as upper bounds.

3. THE DECODING APPARATUS

Assume $V = L$ throughout this section. The Bristol model is constructed as a symmetric iteration, indeed a productive iteration. We will outline the construction of the different intermediate steps in this iteration, and the arguments needed for utilising the iterations apparatus.

Definition 3.1. A *Bristol sequence* is a sequence indexed by the ordinals such that for $\alpha = 0$ or α successor we have $\mathcal{A}_\alpha = \{A_\xi^\alpha \mid \xi < \omega_{\alpha+1}\}$ which is a permutable family on ω_α , and if α is a limit ordinal then we have $F_\alpha = \{f_\xi^\alpha \mid \xi < \omega_{\alpha+1}\}$ which is a permutable scale in the product $\prod \text{SC}(\omega_\alpha)$.

Fix a Bristol sequence. By assuming $V = L$ we not only have a Bristol sequence, indeed we have a canonical one, where we choose the $<_L$ -minimum permutable family or scale at each point. Our goal is to use these permutable objects and replace at each stage, α , a sequence of length ω_α by one of length $\omega_{\alpha+1}$. In particular the first step is to replace the Cohen real, which is a sequence of length ω , by a sequence of length ω_1 . But we want to be able and guarantee that the original sequence is **not** going to be definable from our longer sequence.

In this sense, we want to decode from a Cohen real a sequence of length ω_1 , which captures “some crucial bits” of the Cohen real, but not really all of it. Then we want to decode from this ω_1 sequence a new sequence of length ω_2 , forget the one of length ω_1 , and proceed.

3.1. Example: first steps. Let $\mathbb{P}$ be the Cohen forcing, and in this case we mean $p \in \mathbb{P}$ is a function from a finite subset of ω into 2. Let us omit the index from $\mathcal{A}_0$, as we are only concerned with the first step at the moment, so A_α denotes the α th set in the first permutable family. We let $\mathcal{G}$ and $\mathcal{F}$ denote the derived group and filter from $\mathcal{A}$. The action of $\mathcal{G}$ on $\mathbb{P}$ is the natural one:

$$\pi p(\pi n) = p(n).$$

We denote by $\dot{c}$ the canonical name for the Cohen real. For $A \subseteq \omega$ let $\mathbb{P} \restriction A$ denote the subforcing $\{p \in \mathbb{P} \mid \text{dom } p \subseteq A\}$, and let $\dot{c}_A$ denote the name

$$\{\langle p, \check{n} \rangle \mid p(n) = 1 \wedge \text{dom } p \subseteq A\}.$$

Of course, $\dot{c}_A$ is the canonical name of the generic real added by $\mathbb{P} \restriction A$. We have now that $\pi \dot{c}_A = \dot{c}_{\pi \restriction A}$. In general we say that a name $\dot{x}$ for a set of ordinals is *decent* if every name appearing in it is of the form $\check{\xi}$ for some $\xi \in \text{Ord}$. We say that a name for a set of ordinals is an *A-name* if it is a $\mathbb{P} \restriction A$ -name.

Proposition 3.2. *Suppose that $\dot{x} \in \text{HS}$ and $\mathbb{1} \Vdash \dot{x} \subseteq \omega$, then there is some disjoint approximation $\mathcal{B}$ and a decent $\bigcup \mathcal{B}$ -name $\dot{x}_*$ such that $\mathbb{1} \Vdash \dot{x} = \dot{x}_*$.*

Proof. Let $\mathcal{B}$ be a disjoint approximation such that $\text{fix}(\mathcal{B}) \subseteq \text{sym}(\dot{x})$ and define $\dot{x}_*$ as

$$\dot{x}_* = \left\{ \langle p, \check{n} \rangle \mid p \Vdash \check{n} \in \dot{x} \wedge \text{dom } p \subseteq \bigcup \mathcal{B} \right\}.$$

It is clear that $\mathbb{1} \Vdash \dot{x}_* \subseteq \dot{x}$, to show equality it is enough to prove that if $p \Vdash \check{n} \in \dot{x}$, then $p \restriction \bigcup \mathcal{B} \Vdash \check{n} \in \dot{x}$. Let $q \leq p \restriction \bigcup \mathcal{B}$.

By the very definition of a disjoint approximation $\bigcup \mathcal{B}$ is co-infinite, so we may find a finite set E such that $E \cap (\bigcup \mathcal{B} \cup \text{dom } p) = \emptyset$ and $|E| = |\text{dom } q \setminus \bigcup \mathcal{B}|$. Then let π be a permutation which maps $\text{dom } q \setminus \bigcup \mathcal{B}$ to E , and it is the identity elsewhere. Being a finitary permutation it implements the identity function, so indeed $\pi \in \mathcal{G}$, and by its very definition $\pi \in \text{fix}(\mathcal{B})$, so $\pi \dot{x} = \dot{x}$. So if q had forced $\check{n} \notin \dot{x}$, we would have $\pi q \Vdash \pi \check{n} \notin \pi \dot{x}$, which is the same as $\pi q \Vdash \check{n} \notin \dot{x}$. Alas, πq and p are clearly compatible, and so this is impossible. $\square$

We let G be a V -generic filter¹⁴ and let M denote the symmetric extension HS^G , as is standard, we will “omit the dot” to indicate the interpretation of a name, so c is going to be c^G , etc. The following is a very easy corollary from the above proposition.

Corollary 3.3. $c \notin M$. $\square$

This is where we start seeing the importance of the permutable family, as opposed to any almost disjoint family. If $\{X_\alpha \mid \alpha < \omega_1\}$ is an almost disjoint family, then $\{c \cap X_\alpha \mid \alpha < \omega_1\}$ is a family of mutually generic Cohen reals, any finitely many of them are mutually generic over any other finite subfamily. But in our case, where the almost disjoint family is in fact a permutable one we get **countable** mutual genericity. Any countable subset of $\{c \cap A_\alpha \mid \alpha < \omega_1\}$ is simultaneously mutually generic over any other countable subset (provided they are pairwise disjoint).

We can actually prove more. Nothing in the proof of Proposition 3.2 will change if we assume that $\dot{x}$ is a name of an arbitrary set of ordinals. This shows that every set of ordinals lies in an intermediate model given by $c \cap \bigcup \mathcal{B}$ for some disjoint approximation. Another way to see this fact is to note that $L[c]$ is, after all, only a Cohen extension by a single real. So every set of ordinals is constructible from a single real.

Corollary 3.4. $M \models \neg \text{AC}$.

¹⁴Or L -generic filter, to be explicit.

Proof. Suppose that $M \models \text{AC}$, then there is a set of ordinals A which codes $\mathbb{R}^M$ (e.g. by stacking the real numbers one after another, or by the usual coding of a set into a set of ordinals). Therefore there is $r \in M$ such that $\mathbb{R}^M = \mathbb{R}^{L[r]}$. By Proposition 3.2 we see this is impossible, indeed, if $\mathcal{B}$ is a disjoint approximation for which r has a $\bigcup \mathcal{B}$ -name, and A_α is such that $A_\alpha \cap \bigcup \mathcal{B}$ is finite, then $c \cap A_\alpha \notin L[r]$. $\square$

Of course, we can prove directly that $\mathbb{R}^M$ cannot be well-ordered in M , and this argument can be found in the proof of Theorem 2.7 in [8].

We can see Proposition 3.2 as somehow indicating not only that the reals of M are generated by countable parts of $\mathcal{A}$, but in fact if $\langle T_\alpha \mid \alpha < \omega_1 \rangle$ is a tower generated by $\mathcal{A}$, then we actually have that $\mathbb{R}^M$ is the increasing union of $\mathbb{R}^{L[c \cap T_\alpha]}$. This was the original approach of the Bristol group.

At this point, one might expect that the decoded sequence is $\langle c \cap A_\alpha \mid \alpha < \omega_1 \rangle$ or somehow $\langle c \cap T_\alpha \mid \alpha < \omega_1 \rangle$. But of course, this is not the case. For starters, we want to somehow “fuzzy out” some of the information as to guarantee that c is not constructible from the sequence. So instead of $c \cap A_\alpha$, we will look at $\mathbb{R}^{L[c \cap A_\alpha]}$. But more importantly, the decoded sequence is not even in M . Indeed, if we want this sequence to play the role of the Cohen real in the next step, that means that it needs to be forced into M instead.

Let us begin by understanding $\mathbb{R}^{L[c \cap A_\alpha]}$. In what way does this set “fuzzy out” some information? Well, for one, $c \cap A_\alpha$ is not the obvious real from which we construct this model. Indeed, any finite modification would work, and many more. In fact, any $\pi \in \mathcal{G}$ for which $\iota(\pi)(\alpha) = \alpha$ will satisfy that $\pi \dot{c}_{A_\alpha}$ is a name generating the same set of reals, as it is $c \cap A_\alpha$ up to a permutation of A_α and a finite set.

We say that a name $\dot{x}$ is an *almost A -name* if there exists B such that $A =^* B$ and $\dot{x}$ is a B -name. We now define

$$\dot{R}_\alpha = \{ \dot{x} \mid \dot{x} \text{ is a decent almost } A_\alpha\text{-name} \}^\bullet.$$

Proposition 3.5. *For every $\pi \in \mathcal{G}$, $\pi \dot{R}_\alpha = \dot{R}_{\iota(\pi)(\alpha)}$. In particular, $\dot{R}_\alpha \in \text{HS}$ for all α , and $\{ \dot{R}_\alpha \mid \alpha < \omega_1 \}^\bullet \in \text{HS}$ as well.*

Proof. Observe that $\pi^{\text{``}} \mathbb{P} \restriction A = \mathbb{P} \restriction \pi^{\text{``}} A$. Therefore if $\iota(\pi)(\alpha) = \beta$ we have that $\pi^{\text{``}} A_\alpha =^* A_\beta$, and so an almost A_α -name is moved to an almost A_β -name, so $\pi \dot{R}_\alpha = \dot{R}_{\iota(\pi)(\alpha)}$ as wanted. We now have that $\{ A_\alpha \}$ is a disjoint approximation for which $\text{fix}(\{ A_\alpha \}) \subseteq \text{sym}(\dot{R}_\alpha)$, and thus witnessing that $\dot{R}_\alpha \in \text{HS}$, and indeed $\mathcal{G} = \text{sym}(\{ \dot{R}_\alpha \mid \alpha < \omega_1 \}^\bullet)$ as wanted. $\square$

Similar arguments as we have seen so far also prove the following statement.

Proposition 3.6. *$\langle \dot{R}_\alpha \mid \alpha < \omega_1 \rangle^\bullet \notin \text{HS}$, but for every countable $I \in [\omega_1]^{<\omega_1} \cap L$, $\langle \dot{R}_\alpha \mid \alpha \in I \rangle^\bullet \in \text{HS}$.* $\square$

And here we arrive to the key point. Let ϱ denote the sequence $\langle R_\alpha \mid \alpha < \omega_1 \rangle$ and let $\dot{\varrho}$ denote its name. We will also write ϱ_I and $\dot{\varrho}_I$ when $I \subseteq \omega_1$ for the restriction of the sequence to I , similar to $\dot{c}_A$. Indeed, ϱ is going to be the decoded sequence, and it is of course going to be M -generic, but we need to find a suitable partial order. Proposition 3.6 provides us with a good clue: every initial segment of this sequence is in fact in M , so we can “safely” approximate this sequence.

It will be somewhat more convenient to use subsets of ω_1 , as we did with the Cohen forcing, rather than proper initial segments. This makes it easier to talk about A -names and almost A -names. And again for convenience (and so we can claim productivity, of course), we are also going to limit ourselves to subsets of ω_1 which are already in L .

Proposition 3.7. *Let $\dot{\mathbb{Q}}$ denote $\{ \pi \dot{\varrho}_I \mid \pi \in \mathcal{G}, I \in [\omega_1]^{<\omega_1} \}^\bullet$. Then $\dot{\mathbb{Q}} \in \text{HS}$, and indeed $\Vdash^{\text{HS}} \dot{\varrho}$ is HS-generic.*

In other words, ϱ is M -generic for $\mathbb{Q}$. The idea behind $\mathbb{Q}$ is that we want the smallest “reasonable” set which contains our generic filter (i.e. partial approximations of ϱ), and the easiest way to do that is to simply apply all permutations and obtain a set. But the true intuition behind $\mathbb{Q}$, and really behind the whole decoding apparatus, comes from understanding $\pi \varrho_I$.

The model M knows of the set $R = \{R_\alpha \mid \alpha < \omega_1\}$, it just does not know a well-ordering of this set. And we are trying to remedy that. As we know already every $\pi \in \mathcal{G}$ implements a permutation of ω_1 , which induces a permutation of R moving countably many points. So π shuffles R and thus modifies the range of ϱ_I . But R is an extremely impoverished set as far as M is concerned. This is not particularly important for the construction, and can be skipped entirely, but it is an interesting fact.

Proposition 3.8. *$M \models R$ is a strongly $\aleph_1$ -amorphous set. That is, R cannot be written as a union of two uncountable sets, and every uncountable partition of R has countably many non-singleton cells.*

Proof. Suppose that $\dot{X}, \dot{Y} \in \text{HS}$ and $p \Vdash^{\text{HS}} \dot{X}, \dot{Y} \subseteq \dot{R}$ and are uncountable”. Let $\mathcal{B}$ be a disjoint approximation such that $\text{fix}(\mathcal{B}) \subseteq \text{sym}(\dot{X}) \cap \text{sym}(\dot{Y})$ and such that $\text{dom } p \subseteq \bigcup \mathcal{B}$. By uncountability, we can extend p to some q for which there are α, β such that:

- (1) $q \Vdash^{\text{HS}} \dot{R}_\alpha \in \dot{X}$ and $\dot{R}_\beta \in \dot{Y}$.
- (2) $A_\alpha \cap \bigcup \mathcal{B}$ and $A_\beta \cap \bigcup \mathcal{B}$ are finite.

By enlarging one of the sets in $\mathcal{B}$, if necessary, we can also assume that $\text{dom } q \subseteq \bigcup \mathcal{B}$ as well. We can now find a permutation $\pi \in \text{fix}(\mathcal{B})$ such that $\iota(\pi)$ is the 2-cycle switching α and β .

Applying π to the first property of q we have that $\pi q \Vdash^{\text{HS}} \pi \dot{R}_\alpha \in \pi \dot{X}, \pi \dot{R}_\beta \in \pi \dot{Y}$. But since $\pi \in \text{fix}(\mathcal{B})$ we have that $\pi q = q$ and $\pi \dot{X} = \dot{X}$ and $\pi \dot{Y} = \dot{Y}$. This means that $q \Vdash^{\text{HS}} \dot{R}_\beta \in \dot{X}$ and $\dot{R}_\alpha \in \dot{Y}$. In particular q forces that $\dot{X}$ and $\dot{Y}$ are not disjoint.

Next we want to prove that every partition of R is almost entirely singletons. We will only sketch the idea behind the argument. If $\dot{S} \in \text{HS}$ and $p \Vdash^{\text{HS}} \dot{S}$ is a partition of $\dot{R}$ into uncountably many cells”, let $\mathcal{B}$ be an approximation such that $\text{fix}(\mathcal{B}) \subseteq \text{sym}(\dot{S})$. Pick α, β as above, so that we may switch between them without interfering with $\mathcal{B}$, and we can implement the 2-cycle $(\alpha \beta)$ without changing any given condition. This means that any point that was in the same cell as α must have moved to the cell containing β . But we only moved two points, so α and β must have been isolated as singletons. And therefore the only non-singleton cells come from a partition of $\mathcal{B}$ itself. $\square$

Remark 3.9. The above implies that every permutation of R in M only moves countably many points, of course, as the orbits define a partition. Some of these are new, as they can be encoded by some generic real, but this is irrelevant. We can prove that every permutation of R in M only moves countably many points with a direct argument in the style of Theorem 2.2: given a name for a permutation $\dot{f}$ and $\mathcal{B}$ a disjoint approximation such that $\text{fix}(\mathcal{B}) \subseteq \text{sym}(\dot{f})$, the c.c.c. condition ensures that there is δ such that for any α for which $A_\alpha \cap \bigcup \mathcal{B}$ is infinite, $\dot{R}_\delta$ is not in the same orbit as $\dot{R}_\alpha$. But now we can utilise the same strategy as we did before and move α to some other ordinal with a similar property, and therefore showing that either $\dot{f}$ is the identity on a cocountable set, or it is constant there.

Getting back to the matter at hand, we want to prove that ϱ is M -generic. Namely, if $D \subseteq \mathbb{Q}$ is a dense subset and $D \in M$, we want to prove that there is some $\alpha < \omega_1$ such that $\varrho_\alpha \in D$. In [8] we prove this by proving a technical lemma

about names of dense open sets, Lemma 2.12.¹⁵ In the paper we use this lemma also as a means for proving that $\mathbb{Q}$ is σ -distributive, and therefore does not add any new reals to the model (and as a corollary, it does not force AC back into the universe somehow), which also finishes the proof that ϱ is indeed the sequence we are looking for.

We will prove the genericity of ϱ using a simplified version of Lemma 2.12, and provide a separate argument for the distributivity.

Proposition 3.10. *ϱ is M -generic. In other words, if $\dot{D} \in \text{HS}$ and $p \Vdash \text{“}\dot{D} \text{ is a dense open subset of } \dot{\mathbb{Q}}\text{”}$, then there is some η such that $p \Vdash \dot{\varrho}_\eta \in \dot{D}$.*

Proof. Let $\dot{D}$ be a name as above, and let $\mathcal{B}$ be a disjoint approximation such that $\text{fix}(\mathcal{B}) \subseteq \text{sym}(\dot{D})$. Let $p \Vdash \text{“}\dot{D} \text{ is a dense open subset of } \dot{\mathbb{Q}}\text{”}$. Let α be large enough such that if $A_\xi \cap \bigcup \mathcal{B}$ is infinite, then $\xi < \alpha$. Our strategy is to find “enough” extensions of $\dot{\varrho}_\alpha$ so that one of them will be both $\dot{\varrho}_\eta$, and in $\dot{D}$.

First we prove the following claim: suppose that $q \leq p$ and $q \Vdash \pi \dot{\varrho}_A \in \dot{D}$, where $\alpha \subseteq A$ and $\iota(\pi) \restriction \alpha = \text{id}$. In other words, $\pi \dot{\varrho}_A$ is an extension of $\dot{\varrho}_\alpha$ which lies in $\dot{D}$. Then $q \Vdash \dot{\varrho}_A \in \dot{D}$. Of course, this is true because there is some $\tau \in \text{fix}(\mathcal{B})$ such that $\tau q = q$ and $\iota(\tau) = \iota(\pi)^{-1}$.

The first property, of course, is trivial. The second one is also easy to obtain because π implemented the identity up to α , and therefore $\iota(\pi)^{-1}$ will also be the identity up to α , and thus we can implement it using an automorphism in $\text{fix}(\mathcal{B})$.

This completes the proof of the claim since

$$\tau q = q \Vdash \tau \pi \dot{\varrho}_A = \dot{\varrho}_{\iota(\tau) \circ \iota(\pi) \restriction A} = \dot{\varrho}_A \in \tau \dot{D} = \dot{D}.$$

In turn, this is enough to prove the genericity of ϱ : find a maximal antichain below p of conditions which decide some $\pi \dot{\varrho}_A$ as above, and by openness we can assume each such A is in fact an ordinal. Let η be the supremum of this countable set of ordinals, and we have that $p \Vdash \dot{\varrho}_\eta \in \dot{D}$. $\square$

The final claim is that $\mathbb{Q}$ does not add new reals to M . This, as we remarked, ensures that $M[\varrho] \neq L[\mathbb{Q}]$. It has an added effect that $\mathbb{Q}$ is not adding any new sets of ordinals, or any countable sequences of ground model objects. In [8] the proof utilised a stronger version of the above proof which lets us intersect a countable sequence of dense open sets. Here we take a slightly different approach.¹⁶

Proposition 3.11. *$M \models \text{“}\mathbb{Q} \text{ is } \sigma\text{-distributive”}$. Namely, given $\langle D_n \mid n \in \omega \rangle \in M$ such that each D_n is a dense open subset of $\mathbb{Q}$, then $\bigcap_{n \in \omega} D_n$ is dense.*

Proof. In $L[c]$, $\mathbb{Q}$ is naturally isomorphic to $\text{Add}(\omega_1, 1)^L$. In L this forcing is σ -closed, and so in $L[c]$ it is still distributive. If $\langle D_n \mid n \in \omega \rangle$ is a sequence of dense open sets, then its intersection is still dense in $L[c]$, and therefore in M . $\square$

3.2. Outline of successor steps. The first two steps are fairly indicative of the standard successor step. The main change, of course, is that we need to understand what replaces $\mathbb{R}$, L , and c . We have a hint as to what replaces c , namely, the sequence ϱ that we ended up with after forcing with $\mathbb{Q}$. We also have an idea on what to replace $\mathbb{R}$ with, that would be $\mathcal{P}(\mathbb{R})$, and L is to be replaced by M itself.

We can make this much clearer if we recast the example above by replacing $\mathbb{R}$ with $V_{\omega+1}$. We can also replace c with a sequence of elements of V_ω , of course, but this seems to needlessly complicate things. After all, the case of ω is separate anyway.

¹⁵There is a minor mistake in the statement of the original lemma, see subsection 5.1 for details and corrections.

¹⁶The approach we take here is mentioned in a remark at the end of §2 of [8].

For a more uniform approach, we denote by M_α the α th step in the construction, which is a model of ZF intermediate between L and $L[c]$. We will also write ϱ_α to denote the generic sequence for $\mathbb{Q}_\alpha$, the α th forcing. So $\mathbb{Q}_0$ is Cohen forcing and ϱ_0 is c itself.

At each successor step we have $M_{\alpha+1}$ defined from ϱ_α , which was the M_α -generic sequence for $\mathbb{Q}_\alpha$. We will assume that while $\varrho_\alpha \notin M_{\alpha+1}$, for every $\xi < \omega_{\alpha+1}$, $\varrho_\alpha \restriction A_\xi^\alpha \in M_{\alpha+1}$, where A_ξ^α is the ξ th member of the permutable family we fixed in advance. Moreover, for every $I \in [\omega_{\alpha+1}]^{<\omega_{\alpha+1}} \cap L$, $\langle \varrho_\alpha \restriction A_\xi^\alpha \mid \xi \in I \rangle$ is in $M_{\alpha+1}$.

We now want to define $\mathbb{Q}_{\alpha+1}$ and $\varrho_{\alpha+1}$. For this we replace $\mathbb{R}^{L[c \cap A_\xi]}$ that we had in the first step with $V_{\omega+\alpha+2}^{M_{\alpha+1}[\varrho_\alpha \restriction A_\xi^\alpha]}$, let us denote this as R_ξ for now. The rest is more or less the same as above, relying, of course, on the recursive fact that any previous M_β were defined much in the same way as we are defining $\mathbb{Q}_{\alpha+1}$, $\varrho_{\alpha+1}$, and $M_{\alpha+2}$.

Let $\mathbb{Q}_{\alpha+1}$ denote the set of approximations of $\varrho_{\alpha+1} = \langle R_\xi \mid \xi < \omega_{\alpha+1} \rangle$ whose domains are in L . We are being vague, of course, as to what counts as “approximation” in this context. The idea is that we may permute the different R_ξ amongst themselves using a permutation of $\omega_{\alpha+1}$ which is coming from L .

The lemmas in the general case are exactly the same as we had before. The genericity of $\varrho_{\alpha+1}$ is proved by the same argument as Proposition 3.10, and while the distributivity argument is also similar, it is worth writing down.

Proposition 3.12. $M_{\alpha+1} \models \mathbb{Q}_{\alpha+1}$ is $\leq |V_\alpha^{M_{\alpha+1}}|$ -distributive. In particular, no new sets of rank $\alpha + 1$ are added.

Proof. As in Proposition 3.11, $L[c] \models \mathbb{Q}_{\alpha+1} \cong \text{Add}(\omega_{\alpha+1}, 1)^L$, the latter of which is $\leq \aleph_\alpha$ -distributive in $L[c]$. Suppose now that $\{D_x \mid x \in V_\alpha\} \in M_{\alpha+1}$ is a family of dense open subsets of $\mathbb{Q}_{\alpha+1}$. By the c.c.c. of the Cohen forcing, and the fact we only add a single real, $V_\alpha^{M_{\alpha+1}}$ has the same cardinality as V_α^L , which by GCH is $\aleph_\alpha$. Therefore $\bigcap \{D_x \mid x \in V_\alpha^{M_{\alpha+1}}\}$ is dense in $L[c]$ and thus in $M_{\alpha+1}$.

Suppose $\dot{f}$ is a $\mathbb{Q}_{\alpha+1}$ -name in $M_{\alpha+1}$ and $q \Vdash_{\mathbb{Q}_{\alpha+1}} \dot{f}: V_\alpha \rightarrow \dot{M}_{\alpha+1}$. Set D_x as $\{q' \leq q \mid q' \text{ decides } \dot{f}(\check{x})\}$, then every D_x is dense and open below q . By the density of their intersection, q has an extension $q' \in \bigcap \{D_x \mid x \in V_\alpha\}$, which has to decide the entire function $\dot{f}$. In particular, no new subsets of V_α are added. $\square$

Finally, we define $M_{\alpha+2}$ as the symmetric extension obtained by applying the derived group and filter using the permutable family $\mathcal{A}_{\alpha+1}$. If we now consider $V_{\omega+\alpha+2}^{M_{\alpha+1}[\varrho_{\alpha+1} \restriction A_\xi^{\alpha+1}]}$, this set is $M_{\alpha+2}$. Indeed, every proper initial segment of $\varrho_{\alpha+2}$ is in $M_{\alpha+2}$, where $\varrho_{\alpha+2}$ is the sequence of these $V_{\omega+\alpha+2}$, and it has length $\omega_{\alpha+2}$. We can now show that it is $M_{\alpha+2}$ -generic, etc., and thus all shall prosper.

The successor case can be found in [8] as §4.4, as well as §4.7 and §4.10 for successor of limit iterands. We are not separating the successor of limit steps in this text as the idea is the same with very minor variations, so as far as outlines go, there is essentially no difference.

3.3. Outline of limit steps. The limit step is divided into two parts. We have to contend with the iteration at a limit step, i.e. the finite support limit of the previous steps, and we have to deal with the limit step itself. An observant reader will notice that we did not utilise the full power of the framework of iterating symmetric extensions until now. Indeed, as far as successor steps are concerned any automorphism coming from coordinates before α itself will implement the identity function on the $(\alpha + 1)$ th iterand, rendering it moot.

It is here, at the limit, where we need to utilise the machinery as a whole. In fact, this machinery will do most of the heavy lifting at this stage. By Proposition 3.12

we have that the rank initial segments of the universe are stabilising, indeed for any $\beta > \alpha$ we have $V_{\omega+\alpha+1}^{M_{\alpha+1}} = V_{\omega+\alpha+1}^{M_\beta}$. The work left at this stage is making sure that the generic sequences we collected thus far are symmetrically generic, and setting up the stage for the limit step iterand. So it is a good idea to understand the limit step as a whole before proceeding to the details.

The main idea is that limit steps coalesce the information we have up to that point. Arriving to the limit is easy, as we said, the machinery of productive iterations is working for us there. But how do we proceed now? We are limited by two factors that we need to ensure continue to hold when we deal with the α th step:

- (1) $V_{\omega+\alpha}$ is stable. That is, no new sets of low rank are added, and
- (2) whatever we do is coherent with the other limit steps.

One simple way of ensuring this is by taking products of previous successor steps. This way, if $\alpha < \beta$ are two limit ordinals, then $\mathbb{Q}_\alpha$ is going to be, in some sense, a rank initial segment of $\mathbb{Q}_\beta$. But we can think about this from a different angle.

At each step, we gathered $V_{\omega+\xi}$ of various intermediate models, for $\xi < \alpha$, our limit ordinal. But these are smoothed out, in a sense, as we progress up the hierarchy, as each $V_{\omega+\xi+1}^{M_{\xi+2}}$ contains each $V_{\omega+\xi}^{M_{\xi+1}}$. But what if we could pick just one sequence, and remember it? In that case we are not going to add bounded sets to $V_{\omega+\alpha}$, at least not if we are being careful, and instead we only add this sequence. This idea should seem somewhat familiar to readers of all walks of set theory. After all, if we want to add a new subset to $\aleph_\omega$, it is easy to add Cohen subsets to each $\aleph_n$ first, and then choose a point from each one, creating a new cofinal sequence.¹⁷ Similar ideas, in one way or another, show up through Prikry-style forcings as well.

The coherence of limit steps has another very important use for the limit iterand. One of the subtle, but important properties we used in the successor steps is upward homogeneity.

Definition 3.13. We say that a two-step iteration $\mathbb{P} * \dot{\mathbb{Q}}$ is *upwards homogeneous* if whenever $\langle p, \dot{q} \rangle$ and $\langle p, \dot{q}' \rangle$ are two conditions, there is an automorphism π of $\mathbb{P}$ that respects $\dot{\mathbb{Q}}$ and such that $\pi p = p$ and $p \Vdash \pi \dot{q} = \dot{q}'$. In other words, we can move conditions in $\dot{\mathbb{Q}}$ by automorphisms of $\mathbb{P}$. In the context of iterating symmetric extensions we require that π comes from the relevant automorphism group.

So we want for the limit step that the iteration $\mathbb{P}_\alpha$ can move about conditions in $\mathbb{Q}_\alpha$. If $\mathbb{P}_\alpha$ is a finite support iteration, then either $\mathbb{Q}_\alpha$ needs to have some finitary flavour to its conditions, or somewhere along the iteration we had to condense the conditions from finitary to infinitary. Indeed, this is the very meaning of “coalescing” the successor steps.

To sum up, at the limit step of the iteration we use the properties of the iteration so far to ensure that the rank initial segments of the models stabilise, and indeed that the sequence of generic sequences is symmetrically generic for the iteration. We then want to have a forcing such that the sequences which lie in the product of the successor-step generics combine to form a generic for it, here the permutable scales will come in naturally, as we are concerned with products of increasingly longer sequences modulo the bounded ideal.

We return to the context of the Bristol model’s construction. We denote by $\mathbb{P}_\alpha$ the iteration up to α and by $\mathbb{Q}_\alpha$ the α th iterand, as we did before, and for now we will assume that those iterands were defined also for limit steps. If this proves to be somewhat confusing, the section can be read twice, first assuming $\alpha = \omega$.

Fact 3.14. *For every α , $\mathbb{P}_\alpha * \dot{\mathbb{Q}}_\alpha$ is upward homogeneous.*

¹⁷When forcing like this in the context of ZFC these cofinal sequences are added automatically, of course, but if one does a symmetric iteration, the cofinal sequences are not added. Then one can consider such a forcing in a more material sense.

We have seen this for the case of α being 0 or a successor.¹⁸ We will see the rest of the cases in this section as we progress through it. But for now it is easier to take this as a working assumption.

Proposition 3.15. *Let α be a limit ordinal, then $\langle \varrho_\beta \mid \beta < \alpha \rangle$ is symmetrically L -generic for $\mathbb{P}_\alpha$.*

This is essentially Proposition 4.4 ($\alpha = \omega$) and Proposition 4.15 (α an arbitrary limit ordinal) in [8]. We will prove this statement in subsection 5.2, as the original proof had a minor gap that needs to be corrected anyway.

But this means that we can understand the limit iterand fairly well now. We know that for $\beta < \alpha$, $V_{\omega+\beta+1}^{M_\beta} = V_{\omega+\beta+1}^{M_\alpha}$, and that not only we have a model of ZF which lies within $L[c]$, but that it is in fact an iteration of symmetric extensions, which means that we understand exactly the objects which lie within it and the truth value of statements about these objects from a forcing-theoretic point of view. We are now free to examine the iterand $\mathbb{Q}_\alpha$.

As we reiterate time and time again, we want to ensure that no sets of rank $\omega + \alpha$ are added. In the successor steps we did that by making sure that $\mathbb{Q}_\alpha$ is sufficiently distributive. For the limit step this will pose a problem. If we are to continue with our successor steps, then the $(\alpha + 1)$ th step needs to have a sequence of length $\omega_{\alpha+1}$. But without adding any sequences of length ω_α , this would mean that the sequence must have “mostly existed” already. So the forcing cannot be $<\omega_\alpha$ -distributive, let alone $\leq |V_{\omega+\alpha}|$ -distributive. In fact, if our plan is to add sequences of length α of sets of the form $V_{\omega+\beta}$ of some inner model, then at stages where $\text{cf}(\alpha) = \omega$ we must have added an ω -sequence.

The solution, as it turns out, is to not be distributive at all, but ensure that after applying the symmetric part of the step (rather than just the generic extension) we managed to remove any new set in $V_{\omega+\alpha}$. So even if we do not have a distributive forcing, we at least preserve the rank initial segments of the universe.

We define ϱ_α , where α is a limit, as a “copy of $\prod \text{SC}(\omega_\alpha)^L$ ”. This means that for every $f \in \prod \text{SC}(\omega_\alpha)^L$ we define $\varrho_{\alpha,f}$ to be the sequence $\langle \varrho_{\beta+1}(f(\beta+1)) \mid \beta < \alpha \rangle$, and ϱ_α is defined as $\langle \varrho_{\alpha,f} \mid f \in \prod \text{SC}(\omega_\alpha)^L \rangle$.

How should we define $\mathbb{Q}_\alpha$, then? The idea is always “bounded approximations”, but having the virtue of being a “two-dimensional object” this means that boundedness has two sides to it. Let us first deal with the one-dimensional counterparts: $\varrho_{\alpha,f}$.

If A is a subset of α , we write $\varrho_{\alpha,f} \restriction A = \langle \varrho_{\beta+1}(f(\beta+1)) \mid \beta \in A \rangle$, and so our recursive hypothesis tell us that $\varrho_{\alpha,f} \in M_\alpha$ for every f , and this lets us define $\mathbb{Q}_{\alpha,f}$ as the approximations of $\varrho_{\alpha,f}$. More rigorously, recall that we have defined the symmetric iteration $\mathbb{P}_\alpha$, along with the direct limit of the generic semidirect products, $\mathcal{G}_\alpha$. Then $\dot{\mathbb{Q}}_{\alpha,f}$ is the forcing ordered by reverse inclusion on the interpretation of

$$\left\{ \int_{\vec{\pi}} \dot{\varrho}_{\alpha,f} \restriction A \mid \sup A < \alpha, \int_{\vec{\pi}} \in \mathcal{G}_\alpha \right\}^\bullet.$$

For $E \subseteq \prod \text{SC}(\omega_\alpha)$ we may now define $\varrho_\alpha \restriction E = \langle \varrho_{\alpha,f} \mid f \in E \rangle$, and so if E is bounded in the product, i.e. there is $f \in \prod \text{SC}(\omega_\alpha)$ such that for all $g \in E$ and for all $\beta < \alpha$, $g(\beta+1) < f(\beta+1)$, and A is a bounded subset of α , our conditions are going to be approximations of ϱ_α , up to permutations of course, of the form $\langle \varrho_{\alpha,f} \restriction A \mid f \in E \rangle$, which we denote by $\varrho_\alpha \restriction (E, A)$. And so $\mathbb{Q}_\alpha$ is given by the name

$$\left\{ \int_{\vec{\pi}} \dot{\varrho}_\alpha \restriction (E, A) \mid E, A \text{ are bounded and } \int_{\vec{\pi}} \in \mathcal{G}_\alpha \right\}^\bullet.$$

Proposition 3.16. $\mathbb{P}_\alpha * \dot{\mathbb{Q}}_\alpha$ is upwards homogeneous.

¹⁸To be absolutely correct, we only talked about successors of 0 or other successor ordinals, but the successor of a limit will be just the same as before.

We do not prove this statement here, but the idea is to simply utilise the upwards homogeneity of the previous steps and “correct” the coordinates one by one. One might ask how do we deal with the case where $\alpha > \omega$, as a condition has seemingly infinitely many non-trivial coordinates. And the answer, as we repeatedly mention here, is utilising the previous limit cases where we condense this infinite amount of information, also in the form of $\mathcal{G}_\alpha$ being a subgroup of the full support product $\prod_{\beta < \alpha} \mathcal{G}_{\beta+1}$. The complete proof can be found as Propositions 4.5 (for $\alpha = \omega$) and 4.15 in [8]. The action of $\mathcal{G}_\alpha$ is coordinatewise, and it is important to stress at this point that we have this action where the initial segments of ϱ_α are “actual objects in M_α ”, so this is not applying automorphisms of a forcing, but rather applying permutations of each $\omega_{\beta+1}$.

Proposition 3.17. ϱ_α is M_α -generic for $\mathbb{Q}_\alpha$.

This again follows the same pattern as the successor steps, although here there is a notable complication in the case where $\alpha > \omega$. We will outline the proof.

Proof. Let $\dot{D} \in \text{IS}_\alpha$ be a name for a dense open subset of $\dot{\mathbb{Q}}_\alpha$, and let $\vec{H}$ be an excellent support witnessing that $\dot{D} \in \text{IS}_\alpha$. For each $\beta < \alpha$, H_β is a group of either the form $\text{fix}(\mathcal{B}_\beta)$ when β is 0 or successor, or K_{η_β, f_β} when β is itself a limit. We can use these to define bounded sets E and A such that whenever $\varrho_\alpha \upharpoonright (E, A)$ is extended to a condition in $\dot{D}$, we may permute this extension using the upwards homogeneity without changing $\dot{D}$.

In the case where $\alpha = \omega$, the set E is simple. For each $n < \omega$ we let $\mathcal{B}_n$ be a disjoint approximation such that $\text{fix}(\mathcal{B}_n) = H_n$, then $E = \prod_{n < \omega} \text{dom } \mathcal{B}_n$, where $\text{dom } \mathcal{B}_n = I$ such that $\mathcal{B}_n = \{B_\xi \mid \xi \in I\}$.¹⁹ Let A be the set $\{n < \omega \mid H_n \neq \mathcal{G}_n\}$, which is also bounded by the virtue of $\vec{H}$ being excellent, and we have ourselves the condition $\varrho_\omega \upharpoonright (E, A)$.

In the case where $\alpha > \omega$ we need to take into consideration some limit point $\delta < \alpha$, such that either $\delta + \omega = \alpha$, or δ is the smallest limit ordinal such that $\vec{H}$ is trivial above δ . We call such δ the condensation point of $\vec{H}$. So above δ there can be at most finitely many non-trivial coordinates, and only in the case where $\delta = \alpha + \omega$. We then let E_β for $\beta < \delta$ be decided by $H_\delta = K_{\eta_\delta, f_\delta}$ by taking $E_\beta = f_\delta(\beta)$, and above δ we do as with the case of ω .

Now we may proceed as before, using the fact that any extension of $\varrho_\alpha \upharpoonright (E, A)$ would have the form $\int_{\pi} \varrho_\alpha \upharpoonright (E', A')$, where π_β will necessarily have $\pi_\beta \upharpoonright E_\beta = \text{id}$, so we may find the needed automorphisms in $\vec{H}$ to complete the proof as in the successor case. $\square$

Now that we only need to take care of the preservation of $V_{\omega+\alpha}^{M_\alpha}$. Indeed, $M_\alpha[\varrho_\alpha]$ contains the sequences $\varrho_{\beta+1}$ for $\beta < \alpha$, all of which have rank smaller than $\omega + \alpha$. Luckily, the symmetries of $\mathbb{Q}_\alpha$ will help us get rid of these unwelcome sets.²⁰

Definition 3.18. Working in M_α , if $\dot{x}$ is a $\mathbb{Q}_\alpha$ -name we say that it is *bounded by* $f \in \prod \text{SC}(\omega_\alpha)^L$ if whenever $\int_{\pi} \varrho_\alpha \upharpoonright (E, A) \Vdash \dot{y} \in \dot{x}$, then we may replace E by $E \cap f \downarrow$, where $f \downarrow = \{g \in \prod \text{SC}(\omega_\alpha)^L \mid \forall \beta, g(\beta) < f(\beta)\}$. Similarly, if $\beta < \alpha$ we say that $\dot{x}$ is *bounded by* β if we can replace A by $A \cap \beta$.

Theorem 3.19. Suppose that $\dot{x} \in M_\alpha$ is a $\mathbb{Q}_\alpha$ -name such that every name appearing in $\dot{x}$ is $\dot{y}$ for some $y \in M_\alpha$.

- (1) If $\dot{x} \in \text{HS}_{\mathcal{F}_\alpha}$ and $K_{\eta, f} \subseteq \text{sym}(\dot{x})$, then $\dot{x}$ is bounded by f .

¹⁹In the original paper the proof of $\alpha = \omega$ is Lemma 4.7, and there is a minor mistake in the proof: B_n should be $\text{dom } \mathcal{B}_{n-1}$. Lemma 4.20, which is the general claim, has a correct proof.

²⁰This also highlights the importance of the assumptions in Theorem 2.8 being only about the symmetric extensions having the same V_α .

(2) If $\Vdash \text{rank}(\dot{x}) < \check{\omega} + \check{\alpha}$, then $\dot{x}$ is bounded by some $\beta < \alpha$.

In combination we have that if $\dot{x}$ is a symmetric name for a set of small rank, then all of its elements are decided by conditions with a uniform bound, meaning $\dot{x}$ is equal to an object in M_α . The proof of (1) is the standard homogeneity argument, and we have used it before in Theorem 2.2, so we will only outline the proof of (2).

Proof. Suppose that $\dot{x} \in \mathbb{S}_{\alpha+1}$, we denote by $[\dot{x}]$ the projection of the name to $\mathbb{S}_\alpha$. That is, $[\dot{x}]$ is a name in $\mathbb{S}_\alpha$ which is interpreted in M_α as a name in $\text{HS}_{\mathcal{F}_\alpha}^{M_\alpha}$. By the assumption that each name appearing in $\dot{x}$ is of the form $\check{y}$ for some $y \in M_\alpha$, we may assume that $[\check{y}]$, which appears in $[\dot{x}]$ (in the broad sense of the term) is a name in $\mathbb{S}_\beta$ for β such that $\omega + \beta$ is an upper bound on the forced rank of $\dot{x}$.

Let $\beta < \alpha$ be large enough such that $\vec{H}$ is trivial above β , where $\vec{H}$ witnesses that $[\dot{x}] \in \mathbb{S}_\alpha$. Now we can use automorphisms which only move coordinates above β to move any names of conditions, $\int_{\vec{\pi}} \varrho_\alpha \restriction (E, A)$, by changing their “content above β ” to any value. Thus, we may conclude that we may reduce A to $A \cap \beta$. $\square$

The complete proof of the theorem appears as Lemma 4.10 and Lemma 4.27 in [8]. This almost completes our decoding apparatus. We only need to worry about the $\varrho_{\alpha+1}$ now.

We define R_ξ for $\xi < \omega_{\alpha+1}$ to be $V_{\omega+\alpha+1}^{M_\alpha[\varrho_\alpha, f_\xi^\alpha]}$. That is, we use $\varrho_\alpha, f_\xi^\alpha$ as the “guide” for a new sequence in $V_{\omega+\alpha+1}$. Those who kept track can guess now that $\varrho_{\alpha+1}$ is $\langle R_\xi \mid \xi < \omega_{\alpha+1} \rangle$. We utilise the fact that the scale is permutable to ensure that any bounded part will be in $M_{\alpha+1}$, as well as the rest of the permutability apparatus to ensure the upwards homogeneity.

This is also the point where we see why the definition of $\mathbb{Q}_\alpha$ works in general, despite $\mathbb{P}_\alpha$ being a finite support iteration. The $\varrho_{\alpha, f}$ are initial segments of those ϱ_{β, f^*} that come in the future, and even if $\text{cf}(\alpha) > \omega$, we still end up with what we wanted to have.

With this we finish the discussion on the decoding apparatus. This is the main technical part of the construction. It is our sword and shield in our journey downwards. Now that we have that, we may venture deeper into the Hadean adventure that is the Bristol model.

4. CERBERUS, THE THREE-HEADED GUARD OF THE UNDERWORLD

Much like Cerberus, the construction of the Bristol model can be seen as a three-headed dog, guarding the realm of the underworld of models of ZF. The three heads of Cerberus represent the three causes of strife: nature, cause, and accident. The three heads of the Bristol model can be seen as representing three typical approaches to its construction: nature, cause, and accident. All lead us to the same construction, and indeed when getting down to brass tacks, the details become suspiciously similar in each approach. But the presentation of one may be more appealing to some readers over another.

Nature is the way in which the original group in Bristol went about to define the model: defining the von Neumann hierarchy by hand, and defining a model $L(X)$ where X is a class of sets in the Cohen extension. Cause is the way in which [8] presents the construction: defining an iteration of symmetric extensions, and finding a symmetrically generic filter at each step of the way, thus constructing an iteration and the von Neumann hierarchy of the Bristol model in tandem. Finally, Accident is the way in which we define a productive iteration of symmetric extensions, we study this iteration in an abstract manner, and then we find that by “complete accident” we can find all the symmetrically generic filters inside the Cohen extension.

Fix a Bristol sequence, a permutable family $\mathcal{A}_\alpha = \{A_\xi^\alpha \mid \xi < \omega_{\alpha+1}\}$ for $\alpha = 0$ or a successor, and a permutable scale $F_\alpha = \{f_\xi^\alpha \mid \xi < \omega_{\alpha+1}\}$ for α limit. We will define M , the Bristol model, in three different, yet equivalent ways. We will argue that it is a model of $\mathbf{ZF} + \forall x(V \neq L(x))$.

4.1. Nature. In here we will define the Bristol model one step at a time by defining its von Neumann hierarchy in $L[c]$. For this purpose it would be easier at times to use an increasing sequence modulo bounded sets, rather than the permutable families. Since the two are equivalent, we let $\{T_\xi^\alpha \mid \xi < \omega_{\alpha+1}\}$ denote a sequence obtained from $\mathcal{A}_\alpha$.

We define the ϱ_α and $V_{\omega+\alpha}^M$ in tandem, and we will omit the M from the superscript where possible, as the definitions will be complicated enough. Let $\varrho_0 = c$ and, as V_ω^M is just $V_\omega^L = L_\omega$, it is defined. Suppose that for α , ϱ_α and $V_{\omega+\alpha}$ were defined.

If α is 0 or a successor ordinal, define

$$(1) \quad V_{\omega+\alpha+1} = \bigcup_{\xi < \omega_{\alpha+1}} V_{\omega+\alpha+1}^{L(V_{\omega+\alpha}, \varrho_\alpha \restriction T_\xi^\alpha)},$$

$$(2) \quad \varrho_{\alpha+1} = \left\langle V_{\omega+\alpha+1}^{L(V_{\omega+\alpha}, \varrho_\alpha \restriction A_\xi^\alpha)} \mid \xi < \omega_{\alpha+1} \right\rangle.$$

If α is a limit ordinal, we define $V_\alpha = \bigcup_{\beta < \alpha} V_{\omega+\beta+1}$, and we define $\varrho_\alpha = \prod_{\beta < \alpha} \varrho_{\beta+1}$. And as before we write $\varrho_{\alpha, f}$ to indicate the “thread” of the function f in this product. To define the next step we need an analogue of the T_ξ^α , we write $\overline{\varrho_{\alpha, f}}$ to denote the sequence of $V_{\omega+\beta+1}^{L(V_{\omega+\beta}, \varrho_\beta \restriction T_{f(\beta)}^\beta)}$ for $\beta < \alpha$.

And we now define

$$(3) \quad V_{\omega+\alpha+1} = \bigcup_{\xi < \omega_{\alpha+1}} V_{\omega+\alpha+1}^{L(V_{\omega+\alpha}, \overline{\varrho_{\alpha, f_\xi^\alpha}})},$$

$$(4) \quad \varrho_{\alpha+1} = \left\langle V_{\omega+\alpha+1}^{L(V_{\omega+\alpha+1}, \varrho_\alpha \restriction f_\xi^\alpha)} \mid \xi < \omega_{\alpha+1} \right\rangle.$$

Finally, $M = \bigcup_{\alpha \in \text{Ord}} V_{\omega+\alpha+1}^M$. The handwritten notes passed on to us by some of the members of the Bristol group indicate that the original line of thought was about $\mathbb{R}^M$ and $\mathcal{P}(\mathbb{R})^M$, rather than $V_{\omega+1}$ and $V_{\omega+2}$. The arguments, moreover, as to why $V_{\omega+1} = V_{\omega+1}^{L(V_{\omega+1}, \varrho_1 \restriction T_\xi^\alpha)}$, i.e. why no reals are added when defining $\mathcal{P}(\mathbb{R})^M$, was not written down in these notes. Instead it merely suggests “condensation”.

While restoring the original arguments is certainly beyond us, these would be equivalent, more or less, to the arguments we presented at the first step of section 3.

4.2. Cause. In here we will define the Bristol model in tandem: define a symmetric extension, find it a generic in $L[c]$, define a symmetric extension again, repeat ad ordinalum.

This definition was used in [8], and you can very clearly see that this is the definition that we have in mind, as it very much dominates the approach given in section 3. As such, we really have done all the work ahead of time.

We define M_α, ϱ_α as in the decoding apparatus, i.e. as the generic objects for the symmetric iteration. We now define M simply as $\bigcup_{\alpha \in \text{Ord}} M_\alpha = \bigcup_{\alpha \in \text{Ord}} V_{\omega+\alpha}^{M_{\alpha+1}}$.

4.3. Accident. In here we will first define a class-forcing, and then argue that we can just happen to find symmetrically generic filters in $L[c]$.

Despite being the guiding view on the construction of the Bristol model, the actual argument for $M \models \mathbf{ZF}$ in [8] is the one rising from this approach, as it is less “ad-hoc” and more structural and general. This is also the reason why the proof

of the distributivity of successor steps in the original paper was proved directly, rather than the approach used in this paper, which is more in line with “Cause”.

We first define the iteration. Let $\mathbb{P}_0 = \{1\}$ and $\mathbb{Q}_0 = \text{Add}(\omega, 1)$, $\mathcal{G}_0$ and $\mathcal{F}_0$ are the derived group and filter. Finally, $\dot{\varrho}_0$ is the canonical name for the Cohen real.

Suppose that $\mathbb{P}_\alpha$ is the finite support iteration of the symmetric extensions defined so far. In the case where $\alpha = \beta + 1$, we define $\dot{\mathbb{Q}}_\alpha$ as the name

$$\left\{ \int_{\dot{\pi}} \left\langle \dot{V}_{\omega+\alpha}^{V[\varrho_\beta \upharpoonright A_\xi^\beta]} \mid \xi \in I \right\rangle^\bullet \mid \int_{\dot{\pi}} \in \mathcal{G}_\alpha, I \in [\omega_\alpha]^{<\omega_\alpha} \right\}^\bullet.$$

Where $\dot{V}_{\omega+\alpha}^{V[x]}$ denotes the name of the rank initial segment of the extension of IS_β by the set x . We then define $\dot{\varrho}_\alpha$ as the name for the generic of $\dot{\mathbb{Q}}_\alpha$.

If α is a limit ordinal, we define $\mathbb{P}_\alpha$ as the finite support iteration. We next define $\dot{\mathbb{Q}}_\alpha$ in the same spirit. For $f \in \prod \text{SC}(\omega_{\alpha+1})$ we define

$$\dot{\mathbb{Q}}_{\alpha,f} = \{ \int_{\dot{\pi}} \langle \dot{\varrho}_{\beta+1}(f(\beta+1)) \mid \beta \in A \rangle \mid \int_{\dot{\pi}} \in \mathcal{G}_\alpha, \sup A < \alpha \}^\bullet,$$

and we then define $\dot{\mathbb{Q}}_\alpha$ by adding the additional dimension of a bounded subset of $\prod \text{SC}(\omega_\alpha)$. As before, $\mathcal{G}_\alpha$ and $\mathcal{F}_\alpha$ are the derived group and filter. Finally, $\dot{\varrho}_\alpha$ is the name of the generic filter. Moving on, the definition $\dot{\mathbb{Q}}_{\alpha+1}$ and $\dot{\varrho}_{\alpha+1}$ is in line with what we have done so far.

The properties of the decoding apparatus imply the distributivity of each iterand. We need to be a bit more careful here, as we are not allowed to argue in $L[c]$, instead we carry on a recursive hypothesis that for successor steps $\mathbb{P}_\alpha$ satisfies a chain condition that allows us to prove that $\dot{\mathbb{Q}}_\alpha$ will be isomorphic to $\text{Add}(\omega_\alpha, 1)^L$.

The conditions of Theorem 2.8 are therefore satisfied. This implies that if G is any symmetrically L -generic for $\mathbb{P}$, the class-length iteration, then $\text{IS}^G \models \text{ZF}$, at the very least.

By recursion we can now show that each $\dot{\varrho}_\alpha$ has a symmetric interpretation inside $L[c]$, and therefore we may find an interpretation of the model which is intermediate between L and $L[c]$.

4.4. The basic properties of the Bristol model.

Proposition 4.1. $M \models \text{ZF}$.

Proof. We have two slightly different proofs here, the first we mentioned in the “Accident” approach, utilising Theorem 2.8.²¹ The second approach works well for “Nature” and “Cause”.

Since M is very clearly a transitive subclass of $L[c]$ that contains all the ordinals, it is enough to verify that it is closed under Gödel operations and that it is almost universal to conclude that it is a model of ZF. The first part is easy: if $x, y \in M$, there is some M_α , or some large enough $V_{\omega+\alpha}$ in the “Nature” approach, such that $x, y \in M_\alpha$ and therefore $\{x, y\}$, $x \times y$, etc. are all in M_α and therefore in M .

The second part is the fact mentioned at the end of “Nature” about condensation; or in the case of “Cause” this follows from the fact that for each α , $V_{\omega+\alpha}^{M_\alpha} = V_{\omega+\alpha}^{M_\beta}$ for all $\alpha \leq \beta$, and therefore the model is almost universal. $\square$

Proposition 4.2. $M \models \forall x (V \neq L(x))$.

Proof. If $x \in M$, then there is some α such that $x \in M_\alpha$, and therefore $L(x) \subseteq M_\alpha$. But since $M_\alpha \subsetneq M_{\alpha+1} \subseteq M$, we have that $L(x) \neq M$. In the “Nature” approach we need to be slightly more careful, as we have to verify that if $x \in V_\alpha$, then $V_{\alpha+1} \not\subseteq L(x)$. One way of doing this would be to argue that $\varrho_\alpha \upharpoonright A_\xi^\alpha$ are all generic over $L(V_\alpha)$, and therefore cannot be elements of this model. $\square$

²¹This is mentioned after Theorem 2.8, and proved as Theorem 9.4 in [9].

We will see other ways of deducing Proposition 4.2 later on, in ways that will be significantly more informative and more general.

5. ERRATA TO “THE ABYSS”

Here we correct some minor gaps and mistakes in the original paper. These mistakes repeat, once with the cases of $\alpha = 0$ and $\alpha = \beta + 2$, and once with the cases of $\alpha = \omega$ and a general limit ordinal.

We repeat the remark we made in footnote 19: in [8] the proof of Lemma 4.7, dealing with the genericity of ϱ_ω , contains a typo, whereas B_n is defined as $\bigcup \mathcal{B}_n$, and it should have been defined as $\text{dom } \mathcal{B}_{n-1}$. This is somewhat inconsequential, as the more general proof where α is a limit, in Lemma 4.20, is written correctly.

We also point out that there is a minor mistake in the construction’s induction hypothesis, where the requirement that $\mathbb{P}_\alpha$ satisfies the $\aleph_\alpha$ -c.c. will not hold for limit ordinals, only for 0 and successor ordinals. We can modify this by defining a notion of “symmetric chain condition”,²² but this is an unnecessary complication. We simply do not use the chain condition assumption for limit iterands.

5.1. Corrections to Lemma 2.12/4.1. Lemma 2.12 is read in the context of the first steps. Namely, $\mathbb{P}$ is Cohen forcing, $\dot{\mathbb{Q}}$ is the name of the forcing adding ϱ_1 , denoted in that lemma by σ . As such HS is the class of hereditarily symmetric names defined in the very first step.

Lemma ([8], Lemma 2.12). *Suppose that $\dot{D} \in \text{HS}$ and $p \Vdash \dot{D} \subseteq \dot{\mathbb{Q}}$ is a dense open set. There is some $\eta < \omega_1$ such that for every π and A such that $p \Vdash \pi \dot{\sigma}_A \in \dot{D}$ and $\eta \subseteq A$, if $\tau \dot{\sigma}_A$ is a condition such that $p \Vdash \pi \dot{\sigma}_\eta = \tau \dot{\sigma}_\eta$, then $p \Vdash \tau \dot{\sigma}_A \in \dot{D}$ as well.*

This lemma was used twice. The first consequence is the genericity of ϱ_1 ,²³ and the second is to show that $\dot{\mathbb{Q}}$ is σ -distributive. The proof is very similar to the proof of Proposition 3.10. Nevertheless, when we have two permutations π and τ , we want to move τ so that it agrees with π .

For this we need not only that $\iota(\tau) \restriction A = \iota(\pi) \restriction A$, but also that their inverses agree on A . In the case of the genericity of ϱ_1 we take $\pi = \text{id}$, in which case this property holds trivially. Indeed, this is the very strategy of the proof of Proposition 3.10.

To prove that $\dot{\mathbb{Q}}$ is σ -distributive we interpret the lemma above as saying that for a condition $p \in \mathbb{P}$, there is an ordinal η , such that deciding whether a condition from $\dot{\mathbb{Q}}$, whose domain is at least η , lies in $\dot{D}$ depends only on the permutation π . We then utilise the fact $\mathbb{P}$ is c.c.c. to show that this η can be bound uniformly depending on $\dot{D}$. Now, if $\dot{D}_n$ are names for a sequence of dense open sets, we can uniformly bound all of them by some η . Now if we take any condition whose domain is at least this η , we may extend it into each $\dot{D}_n$, but this means that such extension lies, in fact, in all the $\dot{D}_n$ simultaneously, and therefore the intersection is also dense.

Once we add the condition that $\iota(\pi)^{-1} \restriction A = \iota(\tau)^{-1} \restriction A$, the proof as written in [8] follows through. Alternatively, and perhaps more wisely, for proving the distributivity one should rely on the absoluteness proof that we used in this manuscript, which is also mentioned in the “Abyss”.

Lemma 4.1 is precisely the same, in the context of successor iterations. There is no need to repeat that which has been said.

²²For example, $\langle \mathbb{P}, \mathcal{G}, \mathcal{F} \rangle$ has κ -s.c.c. if every symmetrically dense open set contains a predense set of size $< \kappa$.

²³Or σ , in that context.

5.2. Corrections to Proposition 4.4/4.15. These two propositions are dealing with limit iterations. They show that $\langle \varrho_\beta \mid \beta < \alpha \rangle$ is symmetrically generic for $\mathbb{P}_\alpha$, with Proposition 4.4 dealing with the case $\alpha = \omega$, which was treated separately in the paper.

Proposition ([8], Proposition 4.4). *Suppose that $D \subseteq \mathbb{P}_\omega$ is a symmetrically dense open set, then there is a sequence $\langle \beta_n \mid n < \omega \rangle$ such that $\langle \dot{q}_n \restriction \beta_n \mid n < \omega \rangle \in D$.*

In the paper the proof takes a symmetrically dense open set D , and an excellent support $\vec{H}$ witnessing that. We then generate some sequence of domains such that when we extend the condition $\langle \varrho_\beta \restriction X_\beta \mid \beta < \alpha \rangle$ into D , then we can “correct” it using automorphisms which lie in $\vec{H}$. Namely, each coordinate of $\vec{H}$ is of the form $\text{fix}(\mathcal{B}_\beta)$, and we take X_β to be $\text{sup dom } \mathcal{B}_\beta + 1$, and therefore our starting condition is $\langle \varrho_\beta \restriction X_\beta \mid \beta < \alpha \rangle$.

The idea is fine, and it is somehow intuitively clear what should happen. However the proof described above, taken from the “Abyss”, will not work. The first hint is obvious: ϱ_β has domain ω_β , and X_β is a bounded subset of $\omega_{\beta+1}$. We can resolve this issue by considering $\varrho_{\beta+1} \restriction X_\beta$, but this raises the obvious question, what should we do with limit steps and with ϱ_0 ?

Let us prove this proposition in its general form.

Proposition 5.1 ([8], Proposition 4.15). *$\langle \varrho_\beta \mid \beta < \alpha \rangle$ is symmetrically L -generic for $\mathbb{P}_\alpha$.*

Proof. Let D be a symmetrically dense open set. Our goal is to find a condition of the form $\langle \varrho_\beta \restriction X_\beta \mid \beta < \alpha \rangle$, for appropriate X_β , in D .

Let $\vec{H}$ denote the excellent support witnessing that D is a symmetrically dense open set. For $\beta < \alpha$ let $\xi_{\beta+1}$ the ordinal defined by either,

- (1) when β is 0 or a successor let $\mathcal{B}_\beta$ be such that $H_\beta = \text{fix}(\mathcal{B}_\beta)$, and define $\xi_{\beta+1} = \text{sup dom } \mathcal{B}_\beta$;
- (2) when β is a limit let $\xi_{\beta+1}$ be an ordinal ξ such that for some $f \in \prod \text{SC}(\omega_\beta)$, $H_\beta = K_{\xi, f_\beta}$;
- (3) when $H_\beta = \mathcal{G}_\beta$, set $\xi_{\beta+1} = 0$.

For a limit ordinal β , let $E_\beta \subseteq \prod \text{SC}(\omega_\beta)$ be defined by $\prod_{\beta < \alpha} (\xi_{\beta+1} + 1)$, and let A_β denote $\text{sup}\{\gamma < \beta \mid H_{\gamma+1} \neq \mathcal{G}_{\gamma+1}\}$; if β is a limit ordinal such that $H_\gamma = \mathcal{G}_\gamma$ for all $\gamma \geq \beta$, set $E_\beta = A_\beta = \emptyset$. Now define Y_β as 0 for $\beta = 0$, ξ_β for successors, and (E_β, A_β) for limit ordinals.

The goal is to find the “domain” over which $\vec{H}$ must be the identity, and use that to start our journey into D . Let p be the condition $\langle \varrho_\beta \restriction Y_\beta \mid \beta < \alpha \rangle$. By its very definition, p is not moved by any $\vec{\pi} \in \vec{H}$. Using the density of D we can now extend p to a condition $p_0 \in D$. Moreover, since D is in L , we may assume that the first coordinate of p_0 agrees with ϱ_0 , i.e. the Cohen real c , and by shrinking H_0 if necessary, we may assume that $p_0(0)$ is not moved by automorphisms in H_0 .

If p_0 is not compatible with $\langle \varrho_\beta \mid \beta < \alpha \rangle$, let β be the least coordinate witnessing that. By the choice of Y_β we can find $\vec{\pi} \in H \restriction \beta$, and if β is not a limit ordinal then we may assume $\vec{\pi}$ has a single non-trivial coordinate too, such that by taking $p_1 = \int_{\vec{\pi}} p_0$, the following hold:

- (1) $\text{supp}(p_0) = \text{supp}(p_1)$,
- (2) $p_0 \restriction \beta = p_1 \restriction \beta$,
- (3) $p_1(\beta)$ is indeed compatible with ϱ_β .

The reason we can do that is exactly upwards homogeneity combined with our choice of p , which p_0 extended. As we did not increase the non-trivial coordinates when moving from p_0 to p_1 , we may proceed by recursion and after finitely many steps the process must halt with some $p_n \in D$. $\square$

Indeed, the main point is the fact that we may assume that the first coordinate, for which true genericity is already assumed, is compatible with the Cohen real, and then we can modify any further coordinates recursively using upward homogeneity, combined with the fact that we only need to change things outside of the domains we found, which means that the needed automorphisms can be found in $\tilde{H}$.

6. DECONSTRUCTING CONSTRUCTIBILITY

Now that we have constructed the Bristol model, and we have a good idea about how the construction works, we can ask the obvious question: do we really need $V = L$ in the ground model? The answer, of course, is not really.

We have merely used three assumptions: GCH , $\square_\lambda^*$ for singular λ , and global choice for fixing the Bristol sequence. Of those three, the last one can be easily dispensed, and we will discuss this in section 7. So we are left with only two assumptions which are compatible with a large range of models, including all known inner models from large cardinals below a subcompact.

This raises an interesting question, of course. What kind of large cardinals can the Bristol model accommodate, assuming they existed in the ground model?

Proposition 6.1. *If A is a set of ordinals in the Bristol model, then there is a real number r such that $A \in V[r]$, and moreover $r \in V$ or r is Cohen over V .*

Proof. Note that $V[A]$ is a model of ZFC intermediate to V and $V[c]$, and therefore $V[A]$ is equivalent to some $V[r]$. $\square$

This leads to an immediate corollary: if κ is a large cardinal defined by the existence (or lack thereof) of sets of ordinals, and this largeness is preserved by adding a Cohen real, then κ remains large in the Bristol model.

For example, if κ is Mahlo, then the stationary set of regular cardinals remain stationary in M , and thus κ remains Mahlo. If we define a weakly compact cardinal by stating that every colouring of $[\kappa]^2$ in 2 colours has a homogeneous subset, this too is given by sets of ordinals, and so it continues to hold in M .

We will see in section 9 that this can be extended from sets of ordinals to sets of sets of ordinals, and so on, as long as we iterate power sets less than κ times, the largeness remains. So for example, any measurable on a ground model measurable cardinal will have a unique extension in the Bristol model. Strong cardinals, defined by extenders, are also preserved.

6.1. Limitations. Despite this handsome accommodation of large cardinal assumptions in the ground model, as well as in the Bristol model itself, we can put a stop to this. Indeed, if κ is a supercompact cardinal, then there is a notable shortage of $\square_\lambda^*$ sequences for singular λ such that $\text{cf}(\lambda) < \kappa < \lambda$. We may ask ourselves, perhaps we can salvage permutable scales without having $\square_\lambda^*$?

Theorem 6.2 (The Bristol group). *Suppose that κ is a supercompact cardinal and $\lambda > \kappa > \text{cf}(\lambda)$. Then there are no permutable scales on $\prod \text{SC}(\lambda)$.*

Recall that we are still working under the assumption of GCH . Removing it will require adding $2^\lambda < 2^{\lambda^+}$, which itself is a harmless assumption as it holds for all strong limit cardinals above κ .

Proof. Let $F = \{f_\alpha \mid \alpha < \lambda^+\}$ be a scale in $\prod \text{SC}(\lambda)$, and let π be a permutation of λ^+ , not necessarily bounded, which is not implemented by any sequence of permutations, $\tilde{\pi}$. Pick $j: V \rightarrow W$ an elementary embedding with critical point κ witnessing that κ is at least λ^+ -supercompact, so in particular $j(\kappa) > \lambda^+$. We denote by $G = \{g_\alpha \mid \alpha < j(\lambda^+)\}$ the scale $j(F)$, which is a scale in $j(\prod \text{SC}(\lambda))$ which is equal to $\prod \text{SC}(j(\lambda))$.

Let $\nu = \sup j^{\omega}\lambda^+$, then ν is closed under $j(\pi)$, so we can let τ denote $j(\pi) \upharpoonright \nu$, which is a bounded permutation of $j(\lambda^+)$.

Suppose that there was a sequence of permutations $\vec{\tau}$ that implemented τ , then for $\alpha < \lambda^+$ and $\mu < \lambda$, we have $g_{j(\alpha)}(j(\mu^+)) = j(f_\alpha(\mu^+))$ and $\tau(j(\alpha)) = j(\pi(\alpha))$. Combining this with the assumption that $\vec{\tau}$ implements τ , we get that

$$\tau_{j(\mu^+)}(j(f_\alpha(\mu^+))) = j(f_{\pi(\alpha)}(\mu^+)).$$

We use this to define $\pi_{\mu^+}: \mu^+ \rightarrow \mu^+$. If $\tau_{j(\mu^+)}(j(\xi)) = j(\zeta)$ for some ζ , define $\pi_{\mu^+}(\xi) = \zeta$; otherwise $\pi_{\mu^+}(\xi) = \xi$. This is well-defined, since $\tau_{j(\mu^+)}$ is itself a permutation of $j(\mu^+)$, and so $j(\zeta) < j(\mu^+)$. This is also a permutation by similar reasons.

Finally, we claim that $\vec{\pi}$ implements π . Fix $\alpha < \lambda^+$, then for any large enough $\mu^+ < \lambda$, $\tau_{j(\mu^+)}(j(f_\alpha(\mu^+))) = j(f_{\pi(\alpha)}(\mu^+))$, which means that we defined π_{μ^+} to be exactly $f_{\pi(\alpha)}(\mu^+)$.

We have shown that if F can be used to implement all bounded permutations, then it can be used to implement all permutations, but that is definitely impossible on grounds of cardinal arithmetic. $\square$

7. GAPS IN THE MULTIVERSE

The Bristol model, as we said at first, is a particularly striking counterexample to our understanding of intermediate models when the axiom of choice is not assumed. Working in V 's meta-theory, we can consider the collection of all intermediate models between V and $V[c]$. We remarked before, those models which satisfy ZFC are exactly V itself or Cohen extensions of the form $V[r]$ for some $r \in \mathbb{R}^{V[c]}$.

What about models of ZF? We have, of course, symmetric extensions. These were studied extensively by Grigorieff in [3], and later by Usuba in [19]. Grigorieff proved that if M is a model such that $V \subseteq M \subseteq V[G]$, then M is a symmetric extension given by the same forcing used to add G if and only if $M = (\text{HOD}_{V \cup x})^{V[G]}$ for some set $x \in V[G]$. Usuba extended this and showed that in general, M is a symmetric extension of V if and only if $M = V(x)$ for some set x .

The construction of the Bristol model shows that there is indeed a difference between the two results. Indeed, by counting the number of possible automorphism groups and filters of groups we can see that there cannot be more than $\aleph_3$ distinct symmetric extensions of $\text{Add}(\omega, 1)$ over L .²⁴ Yet, the construction of the Bristol model goes through a proper class of steps, and even if we did not formally prove that each separate one is of the form $L(x)$, we may take $L(V_\alpha^M)$ as our models. By Usuba's result, these are all symmetric extensions of L , but of course most of them are not symmetric extensions where the forcing used is the Cohen forcing.

So when we consider the multiverse of symmetric extensions of L , even those that are landlocked inside $L[c]$, we seem to have two different options. But we want to study all intermediate models, and these include the Bristol model, which is very much not a symmetric extension of L , by any means, as Usuba's result indicate. So what can we say about the multiverse of ZF models?

First, let us ask, how many Bristol models are there? Of course, there is the canonical one, given by the $<_L$ -minimal Bristol sequence. But there are certainly more. Simply by removing some of the sets or functions in any given point in the Bristol sequence we invariably create a new Bristol model.

Proposition 7.1. *Assuming GCH and that $\square_\lambda^*$ holds for all singular λ , there is a class forcing which does not add sets whose generic is a Bristol sequence.* $\square$

²⁴We can improve this counting argument to show no more than $\aleph_2$, actually.

This is done, of course, by approximating the Bristol sequence with set-length options. This is also the way we can remove global choice from the assumptions. Indeed, if V was a model of ZFC, add a generic Bristol sequence, use it to define a Bristol model between V and $V[c]$, where c is a Cohen real, and promptly forget about this generic sequence.

As very clearly the Bristol model is definable from its Bristol sequence in $V[c]$, this means that there may be undefinable Bristol models. And indeed, this can very much be the case.

Theorem 7.2. *Suppose that V is a countable transitive model of $\text{ZFC} + \text{GCH} + \square_\lambda^*$ for all singular cardinals λ , and let c be a Cohen real over V . Then there are uncountably many Bristol models intermediate between V and $V[c]$.*

Proof. Given two Bristol sequences such that $\mathcal{A}_0$ and $\mathcal{A}'_0$ are the two permutable families on ω and such that $\bigcup \mathcal{A}_0 \cap \bigcup \mathcal{A}'_0$ is finite, the two Bristol models are distinct. This can be easily arranged, for example taking $\mathcal{A}_0$ to only have subsets of even integers and $\mathcal{A}'_0$ only has subsets of odd integers. This can be extended to any other step in the Bristol sequence. We can now easily construct uncountably many distinct Bristol models by simply considering with each subset of Ord^V how to modify a given, or indeed a generic, Bristol sequence. $\square$

This is in stark contrast to the case of intermediate models of ZFC of which there are only not just countably many, but there is a set of all the necessary generators in $V[c]$, and the same can be said about symmetric extensions given by the Cohen forcing itself, as Grigorieff's theorem shows. And while there is a proper class of symmetric extensions of V , it is still enumerated by the sets in $V[c]$, making it countable.

Therefore between V and $V[c]$ most models are models of ZF, they are not symmetric extensions of V , and in fact they are not definable in $V[c]$. To make matters worse, inside each Bristol model, we can find a different real, and use that real to interpret the Bristol sequence. Just as well, we may also use one of the $\mathcal{Q}_{\alpha+1} \restriction A_\xi^\alpha$ to interpret the Bristol model construction above a certain stage.

Which truly indicates that the Bristol models are intertwined through this multiverse of models. But to truly appreciate the Bristol model(s), and to understand a bit more its internal structure, we need to have a refined sense of choicelessness.

8. KINNA–WAGNER PRINCIPLES

The axiom of choice can be simply stated as “every set can be injected into an ordinal”, or in other words, “every set is equipotent with a set of ordinals”. This, in conjunction with the following theorem, makes a nice way of understanding models of ZFC.

Theorem (Balcar–Vopěnka). *Suppose that M and N are two models of ZF with the same sets of ordinals. If $M \models \text{ZFC}$, then $M = N$.*

Commonly this is stated when M and N are both models of ZFC, but that is in fact unnecessary. The proof relies on the fact that a relation on a set of ordinals can be coded as a set of ordinals. But can we extend the idea of this characterisation? Yes, yes we can.

Definition 8.1. We say that A is an α -set of ordinals, or simply “ α -set”, if there is an ordinal η such that $A \subseteq \mathcal{P}^\alpha(\eta)$.

Definition 8.2. Kinna–Wagner Principle for α , denoted by KWP_α , is the statement that every set is equipotent with an α -set. We write KWP to mean $\exists \alpha \text{ KWP}_\alpha$.

For $\alpha = 0$ this is simply AC. The principle for $\alpha = 1$ was defined by Kinna and Wagner, although formulated differently using selection functions, and was studied extensively. One of the immediate results is that KWP_1 implies that every set can be linearly ordered, but as the work of Pincus shows in [15], it is stronger and in fact independent of the Boolean Prime Ideal theorem (which also implies every set can be linearly ordered).

These general principles were defined by Monro in [12] for $\alpha < \omega$, and later extended by the author in [9]. Monro extended the result of Balcar and Vopěnka, and this result can be further extended as well.

Theorem 8.3 (The Generalised Balcar–Vopěnka–Monro Theorem).

Suppose that M and N are models of ZF with the same α -sets. If $M \models \text{KWP}_\alpha$, then $M = N$.

Proof. Note that for every α , we can code relations on α -sets by using α -sets in a robust and definable way by extending Gödel’s pairing function. So we can simply repeat the proof of Balcar–Vopěnka. If $x \in M$, we can encode $\text{tcl}(\{x\})$ and its membership relation as an α -set, X . So $X \in N$, and by decoding the membership relation and applying Mostowski’s collapse lemma, $x \in N$. Therefore $M \subseteq N$.

In the other direction, suppose that $V_\eta^M = V_\eta^N$, then we can encode it as an α -set in $M \cap N$. Now given any $x \subseteq V_\eta$ in N , i.e. $x \in V_{\eta+1}^N$, by looking at the α -set encoding V_η , we can identify the subset corresponding to x , in N . But as M and N share the same α -sets, this implies that this subset is in M as well, and we can therefore find $x \in M$. Now by transfinite induction, $V_\eta^M = V_\eta^N$ for all η , and equality ensues. $\square$

Monro proved in [12] that $\text{KWP}_{n+1} \rightarrow \text{KWP}_n$ for all n . This result was extended to show that $\text{KWP}_\omega \rightarrow \text{KWP}_n$ for all n by the author in [9], and was later extended as well by Shani in [18] to show that for all $\alpha < \omega_1$, $\text{KWP}_{\alpha+1} \rightarrow \text{KWP}_\alpha$.

Definition 8.4. *Small Violation of Choice* holds if there exists a set A such that for any set x there is an ordinal η and a surjection $f: \eta \times A \rightarrow x$. We write $\text{SVC}(A)$ to specify this set A , or SVC to mean $\exists A \text{ SVC}(A)$.

This axiom is due to Blass in [1], and it turns out to play an important role in the study of symmetric extensions.

Proposition 8.5. *Suppose that $\text{SVC}(A)$ holds, if A is equipotent with an α -set, then $\text{KWP}_{\alpha+1}$ holds.*

Proof. We may assume that A is an α -set itself, and so if $f: \eta \times A \rightarrow x$ is a surjection, by coding we may replace $\eta \times A$ by an α -set as well, say A_α . Now the function mapping $y \in x$ to $y_f = \{a \in A_\alpha \mid f(a) = y\}$ is injective, and each y_f is an α -set. Therefore x is equipotent to the $(\alpha + 1)$ -set $\{y_f \mid y \in x\}$. $\square$

Blass proved in [1] that SVC is equivalent to the statement “The axiom of choice is forceable [by a set forcing]”, and that SVC holds in every symmetric extension. Usuba showed in [19] that the latter implication can be reversed. Namely, $V \models \text{SVC}$ if and only if V is a symmetric extension of a model of ZFC.

Before we return to study the Bristol model, let us study Kinna–Wagner Principles a bit more in depth.

Theorem 8.6. *Suppose that $V \models \text{KWP}_\alpha$, if $V[G]$ is a generic extension, then $V[G] \models \text{KWP}_{\alpha^*}$, where $\alpha^* = \sup\{\beta + 2 \mid \beta < \alpha\}$.²⁵*

²⁵In other words, α^* is the successor of α when α itself is a successor, or α itself otherwise.

Proof. For every $x \in V[G]$ there is a name $\dot{x}$ in V , so $\dot{x} \restriction G = \{\langle p, y \rangle \in \dot{x} \mid p \in G\}$ in $V[G]$, and the interpretation map is a surjection onto x , which we can extend to a surjection from $\dot{x}$ itself. Therefore every set in $V[G]$ is the surjective image of a set in V .

If $\alpha < \alpha^*$ the above completes the proof, as we may assume that $\dot{x}$ was an α -set, and conclude that x is an $(\alpha + 1)$ -set. For $\alpha = 0$ or a limit ordinal, we use Lemma 8.7. $\square$

Lemma 8.7. *Let A be an α -set for $\alpha = 0$ or a limit ordinal. If $f: A \rightarrow x$ is a surjection, then there exists an α -set B which is equipotent with x .*

Proof. For $\alpha = 0$, A is a set of ordinals, so we may simply choose the least ordinal from the pre-image of each $y \in x$. Suppose that α is a limit ordinal, and define for $\beta < \alpha$, $A_\beta = \{a \in A \mid a \text{ is a } \beta\text{-set}\}$. For every $y \in x$, let β_y be the least β such that for some $a \in A_\beta$, $f(a) = y$.

Now define $B_y = \{a \in A_{\beta_y} \mid f(a) = y\}$, this is a $(\beta_y + 1)$ -set, and $y \neq y'$ implies $B_y \neq B_{y'}$. Therefore $B = \{B_y \mid y \in x\}$ is an α -set equipotent with x . $\square$

On the other hand, ground models need not satisfy the same KWP_α as their generic extension: as we have seen in the construction of the Bristol model, $L[c]$ has a proper class of ground models, $L(V_\alpha^M)$, and as we will see in the next section, the various KWP_α fails as we go up the hierarchy.

The next obvious question is whether or not Theorem 8.6 can be improved. Unfortunately, it cannot. The Cohen model famously satisfies KWP_1 ,²⁶ but as Monro demonstrated in [13], there is a generic extension of the Cohen model in which there exists an amorphous set, which cannot be linearly ordered, in particular, KWP_1 must fail in that generic extension. Nevertheless, by the theorem above, KWP_2 must hold. This leads us to the following conjecture.²⁷

Conjecture 8.8 (The α^* Conjecture). Suppose that KWP_{α^*} holds in every generic extension of V , then KWP_α must hold in V .

It is also easy to see that if $M \models \neg\text{KWP}$, then also any generic extension of M must satisfy this. Which points out to a particularly poignant feature of a generic multiverse, and indeed a *symmetric multiverse*,²⁸ KWP hold or fails uniformly throughout the entire multiverse.

Conjecture 8.9 (The Kinna–Wagner Conjecture). Suppose that $V \models \text{KWP}$ and G is a V -generic filter. If M is an intermediate model between V and $V[G]$ and $M \models \text{KWP}$, then $M = V(x)$ for some set x .

The Bristol model was an exercise in finding an intermediate model which is not constructible from a set. And we conjecture that having any such model, intermediate to a generic extension, will fail KWP , and vice versa: any intermediate model of KWP is constructible from a set over the ground model.

It may also be the case that KWP implies ground model definability, which is a notoriously difficult problem in ZF . Usuba proved that under certain conditions ground models are definable in ZF , but currently the only known models to satisfy these conditions are models satisfying SVC . Nevertheless, as α -sets can be used to characterise a model of ZF in the presence of KWP_α , it stands to reason that it may play a role in ground model definability as well.

²⁶An implicit proof can be found as Lemma 5.25 in [7].

²⁷These are not directly related to the Bristol model, so we include them here instead of section 10

²⁸Allowing symmetric extensions and grounds. See [19] for more information.

We can also define SVC_α to mean that we replace the ordinal, η , by an α -set. And in that case we can easily see that SVC_α is equivalent to “ KWP_α is forceable”. And one is now left wondering if SVC_α is equivalent to being a symmetric extension of a model satisfying KWP_α .

One can also take a different approach and define SVC_M , for a class M , where we may replace the ordinal η by a set from M , so $\text{SVC} = \text{SVC}_{\text{Ord}}$ and SVC_α is a shorthand for $\text{SVC}_{\mathcal{P}^\alpha(\text{Ord})}$. For this to be truly useful, we need to modify KWP_α so that 0-sets are subsets of M . These ideas may play a role in the ultimate answers regarding ground model definability in ZF, and we hope this discussion will help to inspire some of the readers to think about that.

Question 8.10. Is ground model definability equivalent to KWP ?

Note that this question is meaningful since as we observed, KWP is absolute through the generic, and indeed the symmetric, multiverse.

9. CHOICE PRINCIPLES IN BRISTOL

We want to investigate the failure of the axiom of choice in the Bristol model, M , and provide alternative proofs to the key property of the Bristol model, namely $\forall x(V \neq L(x))$.

Proposition 9.1 ([8], Theorem 5.5). *Let M denote the Bristol model, and M_α be the α th model in the construction. Suppose that $A \in M$ is an α -set, then $A \in M_{\alpha+1}$.*

We will not prove this proposition here, but the idea extends the homogeneity argument used in Proposition 3.2. Indeed, $L(V_{\alpha+1}^M)$ contains all the α -sets of M .

Corollary 9.2. *If $\beta > \alpha$, then $M_\beta \models \neg\text{KWP}_\alpha$. In particular, $M \models \neg\text{KWP}$.*

Proof. M, M_β and $M_{\alpha+1}$ have the same α -sets. If one of them would satisfy KWP_α , by The Generalised Balcar–Vopěnka–Monro Theorem, $M = M_\beta = M_{\alpha+1}$. $\square$

Corollary 9.3. *$M \models \neg\text{SVC}$, and consequently $M \models \forall x(V \neq L(x))$, as well as “the axiom of choice is not forceable (by a set forcing)”.* $\square$

It follows from this that at least for a proper class $A \subseteq \text{Ord}$, if $\alpha < \beta$ are both in A , then $\text{KWP}_\beta \not\rightarrow \text{KWP}_\alpha$. But we want to understand the gradation, or rather the degradation, of KWP through the construction.

Proposition 9.4. $M_1 \models \neg\text{KWP}_1 \wedge \neg\text{BPI}$

Proof. We proved that R , the set of $R_\xi = V_{\omega+1}^{L[c \cap A_\xi^0]}$ for $\xi < \omega_1$, is $\aleph_1$ -amorphous in M_1 in Proposition 3.8, and the same argument shows that it cannot be linearly ordered. In particular, it cannot be equipotent to any 1-set and it also witnesses that BPI fails.

Briefly, the argument starts by taking a name $\dot{z} \in \text{IS}_1$ and some $p \Vdash^{\text{IS}} \langle \dot{R}, \dot{z} \rangle^\bullet$ is a linearly ordered set”. Let $\mathcal{B}$ be a disjoint approximation such that $\text{fix}(\mathcal{B}) \subseteq \text{sym}(\dot{R})$ and $\text{dom } p \subseteq \bigcup \mathcal{B}$. We pick $\alpha, \beta \notin \text{dom } \mathcal{B}$ and distinct, and we let $q \leq p$ decide, without loss of generality, $\dot{R}_\alpha \dot{z} \dot{R}_\beta$. But now we can find π such that $\iota(\pi) = (\alpha \ \beta)$ and $\pi q = q$. $\square$

This is a remarkable point, as the elements of R themselves can be well-ordered separately. So you may think we can replace them by 0-sets, but the truth is that we cannot do that uniformly, and this forces us to treat them as 1-sets instead.

Proposition 9.5. $M_2 \models \neg\text{BPI}$.

Proof. We show that R still cannot be linearly ordered when passing to M_2 . Of course the forcing that led us there, $\mathbb{Q}_1$, linearly orders (and in fact well-orders) R . But it is easy to see that this well-order is promptly discarded, and instead we only remember bits and pieces of it in the form of $\varrho_1 \restriction A_\varepsilon^1$. To show that there is no linear ordering in M_2 we need to use the full power of the symmetric iteration. We are going to start with a rather naive attempt, which may not work, but we can identify the problem and circumvent it.

Suppose that $\dot{\prec} \in \mathbf{IS}_2$ is such that $\langle p, \dot{q} \rangle \Vdash_2^{\mathbf{IS}} \text{“}\dot{\prec} \text{ linearly orders } \dot{R}\text{”}$. Let $\mathcal{B}_0$ and $\mathcal{B}_1$ be disjoint approximations such that $\langle \text{fix}(\mathcal{B}_0), \text{fix}(\mathcal{B}_1) \rangle$ is a support of $\dot{\prec}$. Let $\alpha, \beta < \omega_1$ be such that $\alpha, \beta \notin \text{dom } \mathcal{B}_0 \cup \bigcup \mathcal{B}_1$ and moreover $\dot{R}_\alpha, \dot{R}_\beta$ are not mentioned in $\dot{q}$.

Let $\langle p', \dot{q}' \rangle$ be a condition extending $\langle p, \dot{q} \rangle$ such that $\langle p', \dot{q}' \rangle \Vdash_2^{\mathbf{IS}} \dot{R}_\alpha \dot{\prec} \dot{R}_\beta$. We would like to apply upwards homogeneity and consider $\pi \in \text{fix}(\mathcal{B}_0)$ which implements $(\alpha \beta)$ while also not moving p' . But we have to contend with the fact that $\pi \dot{q}'$ may have moved. Luckily, we know exactly where it moved to: $\pi \dot{q}$ simply permutes the range of $\dot{q}$, so if $\dot{R}_\alpha$ and $\dot{R}_\beta$ appear in $\dot{q}'$, which is the likely case, then we simply need to switch them back using some $\sigma \in \text{fix}(\mathcal{B}_1)$, that is an automorphism of $\dot{\mathbb{Q}}_1$, and that will be enough, since automorphisms of $\dot{\mathbb{Q}}_1$ do not change $\dot{R}_\alpha$ and $\dot{R}_\beta$.

Alas, we have a problem. For σ to exist, we need to make sure that $\dot{R}_\alpha$ and $\dot{R}_\beta$ appear in $\dot{q}'$ in coordinates which are not in $\bigcup \mathcal{B}_1$, otherwise we cannot move these coordinates at all. So this naive approach cannot work.

Luckily, we assumed that $\dot{R}_\alpha$ and $\dot{R}_\beta$ are not mentioned in our original $\dot{q}$. So to find $\dot{q}'$, first add both of these in coordinates that are not in $\bigcup \mathcal{B}_1$, and if this was not enough to decide how $\dot{\prec}$ will order them we can extend further to find $\dot{q}'$. In other words, we may assume without loss of generality that $\dot{q}'$ mentions $\dot{R}_\alpha$ and $\dot{R}_\beta$ in coordinates which are eligible to be switched from within $\text{fix}(\mathcal{B}_1)$.

Therefore, if $\langle p', \dot{q}' \rangle \Vdash_2^{\mathbf{IS}} \dot{R}_\alpha \dot{\prec} \dot{R}_\beta$, then also $\langle p', \dot{q}' \rangle \Vdash_2^{\mathbf{IS}} \dot{R}_\beta \dot{\prec} \dot{R}_\alpha$. Therefore $\langle p', \dot{q}' \rangle$ cannot force $\dot{\prec}$ to be a linear ordering, which is a contradiction, since it extends a condition which did force just that. $\square$

One may think that this is enough to prove that there are sets which cannot be linearly ordered in the Bristol model, as we just exhibited that the second symmetric extension will not linearly order R either. Alas, we already concluded that R is a 2-set, so a linear ordering of R will also be a 2-set. But we know that 2-sets are only determined in M_3 . But luckily, we are not very far behind completing this part of the journey.

Theorem 9.6. $M \models \neg \text{BPI}$.

Proof. Suppose that $\langle p, \dot{q}, \dot{r} \rangle \in \mathbb{P}_3$ is a condition that for some $\dot{\prec} \in \mathbf{IS}_3$ forces that $\dot{\prec}$ is a linear order of $\dot{R}$. We can actually run the proof of Proposition 9.5 again. First of all, π will not affect the condition extending $\dot{r}$ at all, but more importantly, σ can be chosen as a permutation moving only two points which means it implements the identity. So it also will not modify the condition extending $\dot{r}$.

And so as long as we were careful to choose the extension $\dot{q}'$ in a way that allows σ to be taken from $\text{fix}(\mathcal{B}_1)$, the argument is not affected. Therefore we showed that $M_3 \models \text{“}R \text{ cannot be linearly ordered”}$, and therefore M does as well. $\square$

10. OPEN QUESTIONS RELATED TO THE BRISTOL MODEL

There is still so much to learn about the Bristol model, both in the specific context of L , as well as many natural questions that come up from generalisations and details in the proof. We cannot possibly include all of these, but we will give four families of questions which are interconnected, but also seem to have independent interest.

10.1. The Bristol models in the multiverse.

Question 10.1. Is there a condition characterising the equivalence classes on Bristol sequences (definable or otherwise) based on the Bristol model they generate using a fixed Cohen real?

Question 10.2. Is the theory of any two Bristol models the same? Does the theory depend on the sequence or its properties?

Question 10.3. Are there any non-trivial grounds of any Bristol model?

Question 10.4. Is there a Bristol model which is definable in its generic extensions, or maybe is there one that is not definable in some of its generic extensions?

Question 10.5. Is there a generic extension of a Bristol model which itself is a Bristol model?

Question 10.6. Is there a maximal Bristol model, namely, is there a Bristol model $M \subseteq L[c]$ such that for any $x \in L[c] \setminus M$, $M(x) = L[c]$?

Question 10.7. Is it true in general that for $x \in L[c]$, either $L(x) = L[c]$ or there is a Bristol model M such that $x \in M$?

10.2. Large cardinals in the Bristol model.

Question 10.8. We saw that measurable cardinals remain measurable. Do they remain critical cardinals in the sense of [5]? What about weakly critical cardinals?

Question 10.9. Principles like $\square_\lambda^*$ are considered to witness failure of compactness. Suppose that λ is singular and no permutable scale exists on $\prod \text{SC}(\lambda)$. Can this compactness be harnessed to restart the Bristol model construction? (Note that a positive answer would indicate that Woodin's Axiom of Choice Conjecture is possibly false, which may imply also the eventual failure of the HOD Conjecture. As such the answer to this question is most likely negative, and a positive answer would be extremely hard to prove.)

Question 10.10. Suppose that elementary embeddings can be lifted and Woodin cardinals are preserved. Starting from strong enough hypotheses, can we construct Bristol model-like objects that satisfy AD?

10.3. Other type of Bristol models.

Question 10.11. Can we start the construction of the Bristol model with a different type of real? Clearly not every real is useful, minimal reals do not have intermediate models, for example. But what about reals that admit sufficiently many automorphisms and intermediate models such as random reals? What about "Cohen + condition" type of reals (Hechler, Mathias, etc.)? Will this also impact the type of forcing we need to do in the following steps (namely, will that force us to use something which is not isomorphic to $\text{Add}(\omega_\alpha, 1)^L$ in successor steps)?

Question 10.12. Can we start the construction with a Prikry-like forcings instead of a Cohen forcing?

Question 10.13. While there is no good definition for iteration of symmetric extensions with countable support, it is imaginable that for productive iterations such as the one used in the Bristol model this is doable by hand. What would this be? Can we have an ω_1 -Bristol model starting with an L -generic sequence for $\text{Add}(\omega_1, 1)$ for example?

10.4. Weak choice principles.

Question 10.14. Does DC hold in the Bristol model?

Question 10.15. Does ω have free ultrafilters in the Bristol model?

Question 10.16. Does $M_\alpha \models \text{KWP}_{\alpha*}$?

Question 10.17. Can any choice principles be forced over the Bristol model?

Question 10.18. Is countable choice true in the Bristol model? If so, is AC_{WO} , the axiom of choice for families that can be well-ordered, true? (Note that this will provide a positive answer about DC, as well as the lifting of elementary embeddings for measurable cardinals.)

Question 10.19. Say that A is x -amorphous if it cannot be well-ordered, and it cannot be written as the union of two sets that are not well-orderable. That is, for some κ , A is κ -amorphous. Are there any x -amorphous sets in the Bristol model?

REFERENCES

1. Andreas Blass, *Injectivity, projectivity, and the axiom of choice*, Transactions of the American Mathematical Society **255** (1979), 31–31.
 2. James Cummings, Matthew Foreman, and Menachem Magidor, *Squares, scales and stationary reflection*, J. Math. Log. **1** (2001), no. 1, 35–98. MR 1838355
 3. Serge Grigoriëff, *Intermediate submodels and generic extensions in set theory*, Ann. of Math. (2) **101** (1975), 447–490. MR 373889
 4. Lorenz J. Halbeisen, *Combinatorial Set Theory*, Springer Monographs in Mathematics, Springer, London, 2012, With a gentle introduction to forcing. MR 3025440
 5. Yair Hayut and Asaf Karagila, *Critical cardinals*, Israel Journal of Mathematics **236** (2020), no. 1, 449–472.
 6. Thomas Jech, *Set Theory*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, The third millennium edition, revised and expanded. MR 1940513
 7. Thomas J. Jech, *The Axiom of Choice*, North-Holland Publishing Co., Amsterdam-London; American Elsevier Publishing Co., Inc., New York, 1973, Studies in Logic and the Foundations of Mathematics, Vol. 75. MR 0396271
 8. Asaf Karagila, *The Bristol model: an abyss called a Cohen real*, J. Math. Log. **18** (2018), no. 2, 1850008, 37. MR 3878470
 9. ———, *Iterating symmetric extensions*, J. Symb. Log. **84** (2019), no. 1, 123–159. MR 3922788
 10. ———, *Realizing realizability results with classical constructions*, Bull. Symb. Log. **25** (2019), no. 4, 429–445. MR 4064113
 11. Kenneth Kunen, *Set Theory*, Studies in Logic (London), vol. 34, College Publications, London, 2011. MR 2905394
 12. G. P. Monro, *Models of ZF with the same sets of sets of ordinals*, Fund. Math. **80** (1973), no. 2, 105–110. MR 0347602
 13. ———, *On generic extensions without the axiom of choice*, Journal of Symbolic Logic **48** (1983), no. 1, 39–52.
 14. Douglass Bert Morris, *Adding total indiscernibles to models of set theory*, ProQuest LLC, Ann Arbor, MI, 1970, Thesis (Ph.D.)—The University of Wisconsin - Madison. MR 2620293
 15. David Pincus, *On the independence of the Kinna-Wagner principle*, Z. Math. Logik Grundlagen Math. **20** (1974), 503–516. MR 369066
 16. Gershon Sageev, *An independence result concerning the axiom of choice*, Annals of Mathematical Logic **8** (1975), no. 1-2, 1–184.
 17. ———, *A model of ZF + there exists an inaccessible, in which the Dedekind cardinals constitute a natural non-standard model of arithmetic*, Annals of Mathematical Logic **21** (1981), no. 2-3, 221–281.
 18. Assaf Shani, *Borel reducibility and symmetric models*, arXiv **1810.06722** (2018).
 19. Toshimichi Usuba, *Geology of symmetric grounds*, arXiv **1912.10246** (2019).
- Email address, Asaf Karagila: karagila@math.huji.ac.il
 URL: <http://karagila.org>

SCHOOL OF MATHEMATICS, UNIVERSITY OF EAST ANGLIA. NORWICH, NR4 7TJ, UK