The Axiomatic Basis for Transfinite Recursive Constructions and Applications to Topology

Nour Al Deen El Hatto

Spring 2026

Advisor: Dr. Alan Dow
Reader: Dr. Will Brian
Department of Mathematics and Statistics Honors Program
University of North Carolina at Charlotte

Contents

Introduction

Transfinite recursion and induction are among the central methods of set theory. They allow one to define objects stage by stage along the ordinals, and to analyze infinite structures by following the same transfinite hierarchy. These methods extend the familiar recursion and induction on the natural numbers to arbitrary ordinals by adding the limit step, allowing one to pass beyond the first infinite ordinal ω. However, this also introduces a new place for a construction to fail: a naive transfinite construction may proceed smoothly through the successor steps, and even through some limit stages, and still break down before reaching its intended end. To carry the construction through, one typically needs some additional idea or trick, and that idea is rarely obvious in advance.

This thesis is organized around that theme. The objects studied here come from several parts of set theory, some from combinatorial set theory and others closely connected with set-theoretic topology, but the underlying issue is the same: how can a transfinite construction be carried through to the end without breaking down at a limit stage? Throughout the thesis, we develop the tools needed to make such constructions succeed. These include carefully chosen inductive hypotheses, together with independent families and the related notions of independence modulo a filter or ultrafilter, as well as maximality of a filter with respect to a given independent family. Together, these provide a natural mechanism for controlling compatibility and ensuring that enough possibilities remain available at later stages of the construction.

We begin with the study of trees, since they provide natural examples in which the central theme of the thesis appears. In the construction of an Aronszajn tree, the successor steps are relatively straightforward, but the limit steps require a further trick in order for the induction to continue. The same theme appears again in the study of Souslin trees and Souslin lines, where transfinite constructions are used in proving the equivalence of the two notions.

A more central part of the thesis concerns reduced powers and ultrapowers of linear orders. Since set theory is deeply concerned with the study of order, ultrapowers provide a natural mechanism for producing new linear orders from old ones, and the resulting structures turn out to be remarkably rich. Łoś's theorem shows that an ultrapower is elementarily equivalent to its underlying structure, so first-order properties alone cannot distinguish the resulting orders. One major goal of the thesis is therefore to study these ultrapowers more closely. A key point is that these ultrapowers fail to be complete: in general, an increasing ω-sequence need not have a least upper bound. This leads naturally to the notion of a gap. To analyze such structures, we introduce invariants that go beyond first-order logic, such as the bounding number, the dominating number, and lower cofinality. The later sections return to these invariants in order to show that different ultrafilters can give rise to genuinely different ultrapowers.

Alongside the study of ultrapowers, we also study the ultrafilters themselves. Free ultrafilters cannot in general be written down explicitly, but their existence is guaranteed by the Ultrafilter Lemma. This leads naturally to further questions: how many free ultrafilters are there, and when should two ultrafilters be regarded as the same? To address the second question, we introduce Rudin–Keisler equivalence as a way of comparing ultrafilters. We then prove a theorem of Kunen showing that there exist two Rudin–Keisler incomparable ultrafilters. From this we prove a topological consequence about $\mathcal{P}(\omega)/\mathrm{fin}$, namely, that its homogeneity as a Boolean algebra does not extend to the space of ultrafilters on it: there exist two free ultrafilters that cannot be sent to one another by any automorphism of $\mathcal{P}(\omega)/\mathrm{fin}$.

Thus, although the topics of this thesis range from trees to ultrapowers to gap constructions, the underlying theme is consistent throughout. Transfinite recursion is not simply a formal procedure for building large objects. Its success depends on finding the right combinatorial principle to carry the construction through limit stages. The aim of this thesis is to illustrate that theme across several settings, while developing the set-theoretic tools needed to make such constructions work.

Preliminaries.

We begin by recalling the set-theoretic definitions and notation used throughout the thesis. Readers already comfortable with this material may omit this section on a first reading.

Orders and ordinals.

Definition. A *partial order* on a set P is a binary relation $\leq$ on P satisfying:

(1) (reflexivity) $p \leq p$ for every $p \in P$;
(2) (antisymmetry) if $p \leq q$ and $q \leq p$, then $p = q$;
(3) (transitivity) if $p \leq q$ and $q \leq r$, then $p \leq r$.

The pair $(P, \leq)$ is a *partially ordered set*, or *poset*. We write $p < q$ to mean $p \leq q$ and $p \neq q$.

An immediate example of a partial order is $(\mathbb{N}, \leq)$. At first glance, a partial order may not seem very different from the usual order on numbers. The key difference is that, in a partial order, not every pair of elements must be comparable.

For example, $(\mathcal{P}(\mathbb{N}), \subseteq)$ is a partially ordered set under inclusion. However, there exist subsets of $\mathbb{N}$ that cannot be compared: for instance, $\{1, 2, 3\}$ and $\{4, 5, 6\}$ satisfy neither $\{1, 2, 3\} \subseteq \{4, 5, 6\}$ nor $\{4, 5, 6\} \subseteq \{1, 2, 3\}$. Such elements are called *incomparable*. This brings us to our next definition.

Definition. A poset $(P, \leq)$ is a *total order*, or *linear order* if it has the following property:

(4) (trichotomy) for all $x, y \in P$, either $x < y$, $y < x$, or $x = y$.

While a total order guarantees that any two elements can be compared, there is another sense in which an order can be stronger. Namely, we may ask whether every nonempty subset has a least element. The usual order on $\mathbb{N}$ has this property, which is one of the reasons it behaves so nicely in induction proofs. This motivates the notions of *well-founded* and *well-ordered* sets.

Definition. A poset $(P, \leq)$ is *well-founded* if every nonempty subset of P has a minimal element.

Equivalently, there is no infinite strictly descending chain. For example, $(\mathcal{P}(\mathbb{N}), \subseteq)$ is not well-founded, since

$$\{0, 1, 2, \dots\} \supset \{1, 2, 3, \dots\} \supset \{2, 3, 4, \dots\} \supset \cdots$$

is an infinite strictly descending chain.

Definition. A *well-order* is a well-founded linear order.

Well-orders play a fundamental role in set theory. Every well-ordered set has a unique order type, and the ordinal numbers are introduced as canonical representatives of these order types. They are defined so that

$$\alpha < \beta \quad \text{if and only if} \quad \alpha \in \beta, \qquad \text{and} \qquad \alpha = \{\beta : \beta < \alpha\}.$$

Definition. An *ordinal* is a set α that is transitive (every element of α is also a subset of α) and well-ordered by $\in$.

We denote ordinals by lowercase Greek letters $\alpha, \beta, \gamma, \dots$, and write Ord for the class of all ordinals. Every ordinal is either a *successor*, of the form $\alpha + 1 = \alpha \cup \{\alpha\}$, or a *limit ordinal*, satisfying $\alpha = \sup\{\beta : \beta < \alpha\} = \bigcup \alpha$. We treat 0 as a limit ordinal, with $\sup \emptyset = 0$.

Theorem (Transfinite Induction)**.** *Let C be a class of ordinals such that:*

(i) *$0 \in C$;*
(ii) *if $\alpha \in C$, then $\alpha + 1 \in C$;*
(iii) *if λ is a nonzero limit ordinal and $\beta \in C$ for all $\beta < \lambda$, then $\lambda \in C$.*

Then C contains every ordinal.

As a theorem, transfinite induction shows that every ordinal is exactly one of the following: zero, a successor ordinal, or a limit ordinal. As a proof technique, it says that a property holding at 0 (base case), preserved by successors (successor case), and preserved at limits (limit case) must hold for all ordinals. This extends induction on the natural numbers, with the limit case added to handle ordinals not reached by repeatedly adding 1.

Theorem (Transfinite Recursion)**.** *Let G be a class function defined on all set-functions whose domain is an ordinal. Then there exists a unique function F on* Ord *such that*

$$F(\alpha) = G(F \restriction \alpha)$$

for every ordinal α.

Just as transfinite induction generalizes ordinary induction on ω to all ordinals, transfinite recursion generalizes recursion on ω to all ordinals. The two are companion principles: transfinite induction is a proof technique, used to verify that a property holds at every ordinal, while transfinite recursion is a construction technique, used to define an object at every ordinal.

Cardinals and cofinality.

Ordinals describe the order structure of a set, while cardinals measure its size. In the finite case, these two notions coincide: a well-ordered set with n elements has both order type n and cardinality n. But what happens in the infinite case?

The following definition generalizes the notion of cardinality to the infinite case.

Definition. Two sets X and Y have the same *cardinality*, written $|X| = |Y|$, if there exists a one-to-one mapping of X onto Y.

This definition leaves the finite case unchanged. In the infinite case, however, order type and cardinality need not coincide. For example, ω and $\omega + 5$ are not order-isomorphic. Nevertheless, they have the same cardinality. Indeed, define $f : \omega + 5 \to \omega$ by

$$f(\alpha) = \begin{cases} \alpha + 5, & \alpha < \omega, \\ n, & \alpha = \omega + n \text{ for } n < 5. \end{cases}$$

Then f is a bijection. Thus $\omega+5$ has order type $\omega+5$, but cardinality $|\omega+5| = |\omega|$.

Definition. Given two sets X and Y, we write $|X| \leq |Y|$ if there exists an injective map from X into Y. We write $|X| < |Y|$ if $|X| \leq |Y|$ and $|X| \neq |Y|$. In this case, we say that X has strictly smaller cardinality than Y.

This defines a relation on cardinalities. It is immediate that $\leq$ is reflexive and transitive. The next theorem shows that it is also antisymmetric, and hence a partial order.

Theorem (Cantor–Bernstein)**.** *If $|A| \leq |B|$ and $|B| \leq |A|$, then $|A| = |B|$.*

It follows from the Axiom of Choice, introduced later, that this order is linear. We state two additional theorems about cardinality before formally defining cardinals.

Theorem (Cantor)**.** *For every set X, $|X| < |\mathcal{P}(X)|$.*

Recall that $\mathcal{P}(X)$ denotes the power set of X. Cantor's theorem shows that there is no largest cardinality: from any set one obtains a strictly larger one by passing to its power set.

Theorem. *If $|A| = \kappa$, then $|\mathcal{P}(A)| = 2^\kappa$.*

Thus Cantor's theorem may be restated in cardinal notation as $\kappa < 2^\kappa$ for every cardinal κ.

For the finite case, clearly $|n| = |m|$ if and only if $n = m$, and so we define *finite cardinals* as the natural numbers, i.e., $|n| = n$ for all $n \in \mathbb{N}$. The next definition identifies cardinals with specific ordinals.

Definition. An ordinal α is called a *cardinal number* (or simply *cardinal*) if $|\alpha| \neq |\beta|$ for all $\beta < \alpha$.

Equivalently, If W is a well-ordered set, then there exists an ordinal α such that $|W| = |\alpha|$. Thus we let

$$|W| = \text{the least ordinal such that } |W| = |\alpha|.$$

As noted earlier, the finite cardinals are precisely the natural numbers. The first infinite cardinal is ω, also denoted by $\aleph_0$. Every infinite cardinal is a limit ordinal. More generally, the infinite cardinals are called the *alephs*, and are written

$$\aleph_0, \aleph_1, \aleph_2, \ldots, \aleph_\omega, \ldots,$$

where $\aleph_\alpha = \omega_\alpha$.

The cardinality of the continuum is $\mathfrak{c} = |\mathbb{R}| = |\mathcal{P}(\omega)| = 2^{\aleph_0}$. Thus $\mathfrak{c}$ is an uncountable cardinal, and hence $\mathfrak{c} \geq \aleph_1$. A natural question is whether $\mathfrak{c}$ is in fact the next cardinal after $\aleph_0$. This is the Continuum Hypothesis (CH), which asserts that $\mathfrak{c} = \aleph_1$. It is independent of ZFC: assuming ZFC is consistent, both CH and its negation are consistent with the axioms of set theory.

In many transfinite constructions, it is not enough to know the size of an ordinal; one must also know how it can be built up from below. This is measured by cofinality. Let $\alpha > 0$ be a limit ordinal. We say that an increasing β-sequence $\langle \alpha_\xi : \xi < \beta \rangle$, β a limit ordinal, is *cofinal* in α if $\lim_{\xi \to \beta} \alpha_\xi = \alpha$.

Definition. If α is an infinite limit ordinal, the *cofinality* of α is

$$\begin{aligned} \mathrm{cf}(\alpha) = {} & \text{the least limit ordinal } \beta \text{ such that there is an increasing} \\ & \beta\text{-sequence } \langle \alpha_\xi : \xi < \beta \rangle \text{ with } \lim_{\xi \to \beta} \alpha_\xi = \alpha. \end{aligned}$$

For example, $\mathrm{cf}(\omega) = \omega$, and more generally $\mathrm{cf}(\omega + \omega) = \omega$. A more interesting example is $\mathrm{cf}(\aleph_\omega) = \omega$, since $\aleph_\omega = \sup\{\aleph_n : n < \omega\}$. Thus even an uncountable cardinal may have countable cofinality. This distinction is important enough to deserve a name.

Definition. An infinite cardinal κ is *regular* if $\mathrm{cf}(\kappa) = \kappa$, and *singular* otherwise.

For instance, ω and ω_1 are regular, whereas $\aleph_\omega$ is singular. More generally, every successor cardinal is regular. In the transfinite arguments that follow, regularity will ensure that a cardinal cannot be reached from below in fewer than κ steps.

The Axiom of Choice.

Throughout this thesis we work in ZFC, the system consisting of the Zermelo –Fraenkel axioms together with the Axiom of Choice.

Definition (Axiom of Choice)**.** Every family of nonempty sets has a choice function.

If S is a family of sets and $\emptyset \notin S$, then a choice function for S is a function f such that $f(X) \in X$ for every $X \in S$.

The Axiom of Choice differs from other axioms of ZF by postulating the existence of a set without defining it. It is interesting to know whether mathematical statements can be proved without the Axiom of Choice. It turns out that the Axiom of Choice is independent of other axioms of set theory and that many mathematical theorems are unprovable in ZF without AC.

The Axiom of Choice admits several equivalent formulations. Different forms of it will be used throughout this thesis, often implicitly. We begin with the Well-Ordering Theorem.

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

Although a set may not come equipped with a natural well-ordering, the Well-Ordering Theorem allows us to impose one whenever needed. In this way, arbitrary sets may be treated as well-ordered for the purposes of transfinite arguments. For example, the real numbers are not naturally well-ordered, but the theorem guarantees a well-ordering of $\mathbb{R}$, and a transfinite induction can then be carried out over that ordering. Another equivalent statement is Zorn's Lemma.

Theorem (Zorn's Lemma). *If $(P, \leq)$ is a nonempty partially ordered set such that every chain in P has an upper bound, then P has a maximal element.*

While the Axiom of Choice is usually formulated in the language of set theory, Zorn's Lemma is often the more convenient form in other areas of mathematics. For instance, it is commonly used to prove that every vector space has a basis.

Filters and ultrafilters.

Filters and ultrafilters provide a way to formalize which subsets of a set are considered "large" and will play an important role in the constructions used later in this thesis.

Definition. A *filter* on nonempty set S is a collection $\mathcal{F}$ of subsets of S such that

(i) $S \in \mathcal{F}$ and $\emptyset \notin \mathcal{F}$,
(ii) if $X \in \mathcal{F}$ and $Y \in \mathcal{F}$, then $X \cap Y \in \mathcal{F}$,
(iii) if $X, Y \subset S$, $X \in \mathcal{F}$, and $X \subset Y$, then $Y \in \mathcal{F}$

Immediately from the definition we obtain some basic examples. One is the trivial filter, namely $\mathcal{F} = \{S\}$. Another is a *principal filter*: if X_0 is a nonempty subset of S, then $\mathcal{F} = \{X \subseteq S : X_0 \subseteq X\}$ is a filter on S. In the special case where $X_0 = \{s\}$ is a singleton, this is called the *principal ultrafilter* generated by s. Another important example, and one that will be used throughout the thesis, is the Fréchet filter on ω, namely the collection of all cofinite subsets of ω.

Definition. A filter $\mathcal{U}$ on a set S is an *ultrafilter* if

(iv) for every $X \subset S$, either $X \in \mathcal{U}$ or $S \setminus X \in \mathcal{U}$.

An immediate observation is that the principal filter generated by a point $s \in S$ is clearly an ultrafilter. In contrast, the Fréchet filter on ω is not an ultrafilter. To see this, let E be the set of even natural numbers and let $O = \omega \setminus E$ be the set of odd natural numbers. Then E and O partition ω, but neither is cofinite, so neither belongs to the Fréchet filter.

An ultrafilter that is not principal is called a *free ultrafilter*, or *nonprincipal ultrafilter*. Such objects cannot in general be written down explicitly; their existence is usually obtained indirectly, using Zorn's Lemma. In particular, the existence of free ultrafilters depends on some form of choice.

Theorem. *A filter $\mathcal{F}$ on S is an ultrafilter if and only if it is maximal.*

That is, $\mathcal{U}$ is an ultrafilter precisely when, for every filter $\mathcal{F}$ on S, if $\mathcal{U} \subseteq \mathcal{F}$, then $\mathcal{F} = \mathcal{U}$. Given any filter, a natural question is whether it can be extended to an ultrafilter; the answer is yes.

Theorem (Ultrafilter Lemma). *Every filter on a set S can be extended to an ultrafilter.*

The usual proof uses Zorn's Lemma, by considering the collection of all filters extending a given filter and ordering them by inclusion. Thus the Ultrafilter Lemma follows from the Axiom of Choice. However, it is strictly weaker: in general, the Ultrafilter Lemma does not imply the full Axiom of Choice.

1. Trees

We begin with trees. Before turning to transfinite arguments, we start with a simpler recursive construction that gives some initial exposure to ultrafilters and motivates the questions that arise in the rest of the section.

Definition 1.1. A *tree* T is a partially ordered set with a minimum element (called the *root*) such that the set of predecessors of every $t \in T$ is well-ordered. For an ordinal α, the α-th *level* of T, denoted T_α, is the set of all $t \in T$ whose predecessor set has order-type α. We say that T is *finitely branching* if every node has only finitely many immediate successors.

Lemma 1.2 (König's Lemma)**.** *Every countably infinite, finitely branching tree has an infinite chain.*

Proof. Let T be a countably infinite, finitely branching tree. Fix a bijection between T and ω, and let $\mathcal{U}$ be a free ultrafilter on ω. For each node $t \in T$, let $f^{\uparrow}(t) = \{s \in T : t \leq s\}$ denote the set of successors of t, including t itself, and identify this set with its image under the chosen bijection.

We recursively construct a chain $t_0 < t_1 < t_2 < \cdots$ such that $f^{\uparrow}(t_n) \in \mathcal{U}$ for every n. Let t_0 be the root of T. Then $f^{\uparrow}(t_0) = T$, so under the bijection it corresponds to all of ω, and hence $f^{\uparrow}(t_0) \in \mathcal{U}$.

Now suppose t_n has been chosen with $f^{\uparrow}(t_n) \in \mathcal{U}$. Let $s_1, \ldots, s_m$ be the immediate successors of t_n. Since $\mathcal{U}$ is free, no finite set belongs to $\mathcal{U}$, so $\{t_n\} \notin \mathcal{U}$, and therefore $f^{\uparrow}(t_n)\setminus\{t_n\} \in \mathcal{U}$. On the other hand, $f^{\uparrow}(t_n)\setminus\{t_n\} = f^{\uparrow}(s_1)\cup\cdots\cup f^{\uparrow}(s_m)$, and these sets are pairwise disjoint. Since an ultrafilter contains one piece of any finite partition of a set it contains, it follows that $f^{\uparrow}(s_i) \in \mathcal{U}$ for some $i \in \{1, \ldots, m\}$. Set $t_{n+1} = s_i$.

This recursion never terminates. If t_{n+1} were terminal, then $f^{\uparrow}(t_{n+1}) = \{t_{n+1}\}$ would be finite, contradicting the fact that every member of a free ultrafilter is infinite. Thus we obtain an infinite chain $t_0 < t_1 < t_2 < \cdots$, as required. □

As stated, König's Lemma may appear trivial. However, if we weaken the finite branching assumption to countable branching, the conclusion fails. Consider the tree T whose root has infinitely many immediate successors $s_1, s_2, s_3, \ldots$, where each s_n is the root of a finite chain of length n. Then T is infinite and countably branching, but every branch is finite, so T has no infinite chain.

Having established the countable case, we now examine the uncountable case, beginning with the following definition.

Definition 1.3. An ω_1-tree is an uncountable tree in which every level is countable.

We may wonder whether König's Lemma extends to the uncountable setting. In particular, can ultrafilters be used to prove that every ω_1-tree has an uncountable branch?

We attempt to mimic the ultrafilter proof of König's Lemma in the uncountable setting. First, note that every ω_1-tree has cardinality exactly ω_1: it has ω_1-many levels, each countable, so $|T| = \omega_1 \cdot \omega = \omega_1$. Fix a bijection $f\colon \omega_1 \to T$ and a free ultrafilter $\mathcal{U}$ on ω_1.

For a node $t \in T$, write $T_t = \{s \in T : s \leq t \text{ or } t \leq s\}$ for the set of all nodes comparable to t. (This should not be confused with the level notation T_α, which is indexed by an ordinal) Following the countable proof, we would like to recursively choose an increasing chain $t_0 < t_1 < t_2 < \cdots$ such that $f^{-1}(T_{t_n}) \in \mathcal{U}$ at each stage, ensuring that uncountably many nodes lie above t_n.

Suppose we have some method for choosing successors at each step, so that we build an increasing chain $t_0 < t_1 < t_2 < \cdots$ with $f^{-1}(T_{t_n}) \in \mathcal{U}$ for every n. For this to be meaningful, we need $\mathcal{U}$ to be *uniform*, meaning every member of $\mathcal{U}$ has cardinality ω_1; equivalently, no countable set belongs to $\mathcal{U}$. Without uniformity, $f^{-1}(T_{t_n}) \in \mathcal{U}$ could merely guarantee countably many nodes comparable to t_n, giving us no control over the length of the chain. With a uniform ultrafilter, this condition ensures that uncountably many nodes remain comparable to t_n at every stage.

Even so, to extend the chain past ω, we would need a node t_ω above all of the t_n's. In particular, t_ω must belong to T_{t_n} for every n, so any candidate for t_ω must lie in the intersection $\bigcap_{n\in\omega} T_{t_n}$. Thus, for the recursion to continue, we would need

$$\bigcap_{n\in\omega} f^{-1}(T_{t_n}) \in \mathcal{U}.$$

But is there any reason that this countable intersection should belong to the ultrafilter, or even be uncountable?

The answer is no. An ultrafilter closed under countable intersections is called *countably complete*, and a cardinal carrying such an ultrafilter is called a *measurable cardinal*. Measurable cardinals are examples of large cardinals, whose existence cannot be proved in ZFC. Moreover, if $\lambda \leq \kappa$ and if λ has a countably complete ultrafilter, then so does κ.

We claim that ω_1 is not measurable. Since ω_1 is the smallest uncountable cardinal and $\mathbb{R}$ is uncountable, $\omega_1 \leq \mathfrak{c} = |\mathbb{R}|$. By the contrapositive of the above, it suffices to show that $\mathfrak{c}$ is not measurable.

Let $\mathcal{U}$ be an ultrafilter on $\mathbb{R}$. Since $\mathbb{R}$ is in bijection with $(0,1)$, we may work on $(0,1)$ instead. For each $n \geq 1$, partition $(0,1)$ into the n intervals $\left[\frac{i}{n}, \frac{i+1}{n}\right)$ for $i < n$. Because $\mathcal{U}$ is an ultrafilter, one piece of each such finite partition must belong to $\mathcal{U}$. So for every n there is some $i_n < n$ such that $\left[\frac{i_n}{n}, \frac{i_n+1}{n}\right) \in \mathcal{U}$. The intersection $\bigcap_{n\in\omega}\left[\frac{i_n}{n}, \frac{i_n+1}{n}\right)$ is then at most a singleton, so $\mathcal{U}$ cannot be countably complete. Hence there is no countably complete ultrafilter on $\mathbb{R}$, and therefore $\mathfrak{c}$ is not measurable.

So the ultrafilter approach fails. But perhaps a different proof strategy could show that every ω_1-tree has an uncountable branch? The answer is no: the statement itself is false.

Definition 1.4. An *Aronszajn tree* is an ω_1-tree with no uncountable chains.

Theorem 1.5. *There exists an Aronszajn tree.*

Proof. The elements of our tree will be bounded, strictly increasing sequences of rationals $\langle q_\xi : \xi < \alpha\rangle$ for countable ordinals α, ordered by end extension: $\langle q_\xi : \xi < \alpha\rangle$ extends $\langle r_\xi : \xi < \beta\rangle$ if $\beta \leq \alpha$ and $q_\xi = r_\xi$ for all $\xi < \beta$. Any chain in such a tree is itself a strictly increasing sequence of rationals, and since $\mathbb{Q}$ is countable, so is any such chain. Thus our tree has no uncountable branches.

Naively choosing bounded rational sequences might lead us to a dead end. Suppose we have defined T_α for all $\alpha < \delta$. To populate T_δ, we need a δ-length sequence whose initial α-segment lies in T_α for every $\alpha < \delta$. But what if all such δ-sequences are cofinal in $\mathbb{R}$? Then none can be placed in our tree, T_δ is empty, and the construction halts. To ensure this does not happen, we impose an inductive hypothesis on the construction.

We construct the levels T_α by transfinite recursion on $\alpha < \omega_1$, maintaining the following inductive hypothesis at stage δ: for every $\alpha < \delta$, every $t \in T_\alpha$, and every real $r > \sup(t)$, there is, for every $\alpha < \beta < \delta$, an element $t < s \in T_\beta$ with $\sup(s) < r$. That is, any node can be extended to any later level constructed so far while staying below any prescribed real bound. As we shall see, this is precisely what allows the construction to survive limit stages. Set $T_0 = \{\langle\rangle\}$; the hypothesis holds vacuously.

For the successor step, suppose T_α has been constructed. For every $t \in T_\alpha$ let $T_{\alpha+1}$ be the set of all extensions $t^\frown\langle q\rangle$ for every rational $q > \sup(t)$. The level $T_{\alpha+1}$ is countable, since T_α and $\mathbb{Q}$ are both countable.

It remains to verify the inductive hypothesis at stage $\alpha + 1$. Fix $\beta \leq \alpha$, $t \in T_\beta$, and $r > \sup(t)$. We want an element of $T_{\alpha+1}$ extending t and still bounded by r. If $\beta < \alpha$, the inductive hypothesis at stage α gives some $s \in T_\alpha$ extending t with $\sup(s) < r$. If $\beta = \alpha$, take $s = t$. Now choose a rational q such that $\sup(s) < q < r$. Then $s^\frown\langle q\rangle \in T_{\alpha+1}$, it extends t, and $\sup(s^\frown\langle q\rangle) = q < r$. Hence the inductive hypothesis holds at stage $\alpha + 1$.

Now suppose δ is a limit ordinal and $\langle T_\alpha : \alpha < \delta\rangle$ has been constructed satisfying the inductive hypothesis. Since δ is countable, fix a strictly increasing sequence $\langle\gamma_k : k \in \omega\rangle$ cofinal in δ. Given $t \in T_{\gamma_0}$ and a real $r > \sup(t)$, the inductive hypothesis lets us recursively pick $t_{k+1} \in T_{\gamma_{k+1}}$ extending t_k with $\sup(t_{k+1}) < r$. The union $\bigcup_{k\in\omega} t_k$ is then a bounded strictly increasing δ-sequence of rationals extending t, with $\sup \leq r$.

At each of the ω-many steps in building such a chain, there are ω-many choices of which node to extend to, so the total number of chains is $|\omega^\omega| = \mathfrak{c}$. Including all of them would make T_δ uncountable. Instead, For each node $t \in \bigcup_{\alpha<\delta} T_\alpha$ and each rational $q > \sup(t)$, choose one chain extending t and staying below q, and put its union into T_δ. Since the set of such pairs (t, q) is countable, the level T_δ is countable.

It remains to verify the inductive hypothesis at stage δ. So fix $\beta < \gamma \leq \delta$, $t \in T_\beta$, and $r > \sup(t)$. If $\gamma < \delta$, this is exactly the inductive hypothesis already assumed below δ. If $\gamma = \delta$, choose a rational q with $\sup(t) < q < r$. By construction, T_δ contains a node extending t whose supremum is below q, and hence below r. Thus the inductive hypothesis continues to hold at stage δ.

After ω_1-many steps, $T = \bigcup_{\alpha<\omega_1} T_\alpha$ is a tree with ω_1-many levels, each countable, and no uncountable branch. $\square$

We now shift our attention to a problem closely related to Aronszajn trees. The construction above showed that an ω_1-tree can exist without any uncountable branches. But one can ask for more: can such a tree also have no uncountable antichains? Recall that a subset $A \subseteq T$ is called an *antichain* if no two elements of A have a common upper bound in T.

Definition 1.6. A *Souslin tree* is an Aronszajn tree in which every antichain is countable.

The question of whether a Souslin tree exists turns out to be equivalent to a classical problem about linear orders. To state it, we need the following notion.

Definition 1.7. A dense linearly ordered set P satisfies the *countable chain condition* (ccc) if every collection of pairwise disjoint open intervals in P is at most countable.

It is a classical result that $\mathbb{R}$ is, up to isomorphism, the unique linearly ordered set that is dense, unbounded, complete, and contains a countable dense subset — that is, it is *separable*. Note that separability implies ccc: if D is a countable dense subset of P and $\{I_\gamma\}$ is a collection of pairwise disjoint open intervals, then each I_γ contains a distinct element of D, so the collection is at most countable. One might ask: if P is a complete, dense, unbounded linearly ordered set satisfying the countable chain condition, must it be isomorphic to $\mathbb{R}$? In other words, can such a P be ccc without being separable? This is known as Souslin's problem.

Definition 1.8. A *Souslin line* is a dense linear order that satisfies the countable chain condition yet has no countable dense subset.

Thus Souslin's problem asks whether a Souslin line exists. We shall show that the existence of a Souslin line is equivalent to the existence of a Souslin tree.

Lemma 1.9. *If there is a Souslin tree, then there is one that satisfies the following*

(1) *the set above every node is uncountable;*
(2) *every node has infinitely many immediate successors.*

Proof. We begin with the first condition. Note that any uncountable subtree of a Souslin tree is again a Souslin tree. Let Let A be the set of minimal elements $t \in T$ such that $T_t = \{s \in T : t \leq s\}$ is countable. We claim that A is an antichain. Observe, if $a, b \in A$ and $a < b$, then $T_b \subseteq T_a$. Since T_a is countable, it follows that T_b is countable as well, contradicting the minimality of b.

Because T is Souslin, every antichain is countable, so A is countable. Now remove $A^{\uparrow} = \{s \in T : (\exists a \in A)\, a \leq s\}$. The remaining tree T_1 is still uncountable, hence again a Souslin tree. Moreover, T_t is uncountable for every node $t \in T_1$. Otherwise there would be some minimal $a \leq t$ such that T_a is countable. Then $a \in A$, so $t \in A^{\uparrow}$, contrary to the construction. In particular, no branch dies off in the pruned tree, since otherwise T_t would be countable for some node t.

We now ensure that every node branches. We select a subsequence of levels from T by recursion. Set $\alpha_0 = 0$. At stage ξ, choose $\alpha_{\xi+1}$ so that every node on level α_ξ has at least two successors on level $\alpha_{\xi+1}$. Such a level exists: T_t is uncountable for every t on level α_ξ, so if T_t never branches, that would produce an uncountable chain, contradicting the fact that T is Aronszajn. For each t on level α_ξ, let β_t be the first level above α_ξ at which t has at least two successors. Set

$\alpha_{\xi+1} = \sup\{\beta_t : t \in T_{\alpha_\xi}\}$; this is a countable supremum of countable ordinals, so $\alpha_{\xi+1} < \omega_1$.

At a limit stage δ, set $\alpha_\delta = \sup\{\alpha_\xi : \xi < \delta\}$. We take all nodes on level α_δ of the original tree T as our new level δ. This is countable because level α_δ of T is countable. Every node on a previously chosen level has an extension there, since T_s is uncountable for every s and hence s has successors on every level above it.

The subtree of T restricted to levels $\{\alpha_\xi : \xi < \omega_1\}$ is an ω_1-tree with countable levels in which every node has at least two immediate successors. Any antichain in this subtree is an antichain in T, hence countable, so it is still a Souslin tree. Finally, since every node branches and no branch dies, restricting further to nodes on limit levels (together with the root) yields a tree in which every node has infinitely many immediate successors. Thus the resulting tree is a Souslin tree satisfying both (1) and (2). □

Theorem 1.10. *There is a Souslin tree if and only if there is a Souslin line.*

Proof. We prove the forward implication. Let T be a Souslin tree as obtained by the previous lemma. For each $t \in T$, index the immediate successors of t by the nonzero integers: $\{s(t,n) : 0 \neq n \in \mathbb{Z}\}$ is the set of immediate successors of t. We think of t itself as representing $s(t,0)$.

Now define a new ordering $\leq$ on T, separate from the tree ordering $\leq_T$. Given $s, t \in T$, let r be the maximal extension of the root such that both s and t are above r. It may happen that $r = s$ or $r = t$. If $r = s$, set $n_s = 0$; otherwise let $0 \neq n_s \in \mathbb{Z}$ be the unique integer such that $s(r, n_s) \leq_T s$. Define n_t similarly. By maximality of r, we have $n_s \neq n_t$, since otherwise $s(r, n_s)$ would be a common extension of r below both s and t. We now declare that $s \leq t$ if and only if $n_s \leq n_t$ in the usual order on $\mathbb{Z}$.

The ordering is reflexive, antisymmetric, and transitive by construction. Moreover, any two nodes are comparable, since they share the root as a common predecessor. Thus it is a linear order. It is also easy to see that this order is dense.

For each $t \in T$, define $[t] = \{s \in T : t \leq_T s\}$. We now show that every nonempty open interval (a, b) in the linear order contains $[t]$ for some $t \in T$. Pick any c with $a < c < b$. Let r_1 be the last common ancestor of a and c, and r_2 the last common ancestor of c and b. Without loss of generality, suppose $r_1 \leq_T r_2$. At r_2, the nodes c and b diverge, with c passing through $s(r_2, n_c)$ and b through $s(r_2, n_b)$, where $n_c < n_b$. Set $t = s(r_2, n_c)$. Then every extension of t lies below b. Since $a < c$ and $r_1 \leq_T r_2$, every extension of t also lies above a. Hence $[t] \subseteq (a, b)$.

It follows that $[t]$ is open. For each $u \in [t]$, let $I_u = (s(u, -1), s(u, 1))$. By the previous argument, $I_u \subseteq [u] \subseteq [t]$. Conversely, if $x \in [t]$, then $x \in I_x$. Hence $[t] = \bigcup_{u \in [t]} I_u$, so $[t]$ is a union of open intervals, and therefore open.

We now verify the countable chain condition. Let $\{I_\gamma : \gamma \in \Gamma\}$ be a family of pairwise disjoint nonempty open intervals in $(T, \leq)$. By the previous paragraph, for each γ there is some $t_\gamma \in T$ such that $[t_\gamma] \subseteq I_\gamma$. Since the intervals are pairwise disjoint, the sets $[t_\gamma]$ are pairwise disjoint. Thus no two of the t_γ can have a common upper bound in the tree: if $u \geq_T t_\gamma$ and $u \geq_T t_\delta$ for $\gamma \neq \delta$, then $u \in [t_\gamma] \cap [t_\delta]$, contradicting disjointness. So $\{t_\gamma : \gamma \in \Gamma\}$ is an antichain in T. Since T is Souslin, every antichain is countable. Hence Γ is countable, and $(T, \leq)$ satisfies ccc.

Finally, we show that $(T, \leq)$ is not separable. Let D be any countable subset of T. Each element of D sits on some level of the tree, so D touches only countably many levels. Since T has ω_1-many levels, there is some level α above all elements of D. Pick any t on level α. Then $[t]$ is a nonempty open interval, and every element of $[t]$ is on level α or higher, so $[t] \cap D = \emptyset$. Thus D is not dense, and since D was arbitrary, $(T, \leq)$ is not separable.

Thus $(T, \leq)$ is a dense linear order satisfying ccc and having no countable dense subset. Therefore it is a Souslin line.

For the reverse implication, let S be a Souslin line. We construct a Souslin tree whose elements are closed nondegenerate intervals in S, ordered by reverse inclusion: for $I, J \in T$, declare $I \leq J$ if and only if $I \supseteq J$.

The nodes of T are constructed by recursion on $\alpha < \omega_1$. Set $I_0 = [a_0, b_0]$ for any $a_0 < b_0$ in S. At stage α, suppose the intervals $I_\beta = [a_\beta, b_\beta]$ have been chosen for all $\beta < \alpha$. The set $C = \{a_\beta : \beta < \alpha\} \cup \{b_\beta : \beta < \alpha\}$ is countable. If C were dense in S, then S would be separable, contradicting the assumption that S is a Souslin line. So there exist $a_\alpha < b_\alpha$ in S such that $[a_\alpha, b_\alpha]$ is disjoint from C; set $I_\alpha = [a_\alpha, b_\alpha]$. Since the construction at each stage depends only on the set of previously chosen endpoints, the successor and limit cases are identical.

The set $T = \{I_\alpha : \alpha < \omega_1\}$ is uncountable and partially ordered by $\supset$. If $\alpha < \beta$, then either $I_\alpha \supset I_\beta$ or I_α is disjoint from I_β. It follows that each α, $\{I \in T : I \supset I_\alpha\}$ is well-ordered by $\supset$ and thus T is a tree.

We shall show that T has no uncountable branches and no uncountable antichains. Then it is immediate that the height of T is at most ω_1; and since every level is an antichain and T is uncountable, we have that T has uncountable height.

For antichains, note that if $I, J \in T$ are incomparable, then neither contains the other, so $I \cap J = \emptyset$. Thus any antichain in T is a family of pairwise disjoint intervals in S, and hence is countable since S satisfies ccc.

For branches, let $\mathcal{B} \subseteq T$ be a branch. Then $\mathcal{B}$ is linearly ordered by reverse inclusion, so its intervals are nested. If $\mathcal{B}$ were uncountable, then the left endpoints of its intervals would form a strictly increasing uncountable sequence $\langle x_\alpha : \alpha < \omega_1 \rangle$. But then the open intervals $(x_\alpha, x_{\alpha+1})$ would be pairwise disjoint, contradicting ccc. Hence every branch in T is countable. Thus our tree is a Souslin tree. $\square$

It turns out that, unlike the Aronszajn tree whose existence we proved in ZFC, the existence of a Souslin tree is independent of ZFC: using the method of forcing, which we do not cover in this thesis, one can show that both the existence and the non-existence of a Souslin tree are each consistent with ZFC. Thus Souslin's problem cannot be settled from the standard axioms alone.

2. Reduced Powers and Ultrafilters

2.1. **Reduced Powers.** We now turn to the study of linear orders. Filters and ultrafilters provide a natural mechanism for producing new orders out of existing ones, and the resulting structures turn out to be remarkably rich. Our goal in this section is to develop the tools needed to describe these orders and to understand the gap structure they carry.

Given a linearly ordered set $(L, <)$ and a filter $\mathcal{F}$ on ω, the basic construction is the following. We write $\mathcal{L}$ for the structure $(L, <)$.

Definition 2.1. The *reduced power* of $\mathcal{L}$ by $\mathcal{F}$ is a relational structure on the base set L^ω. We let $<_\mathcal{F}$ denote the relation $f <_\mathcal{F} g$ meaning $\{n : f(n) < g(n)\} \in \mathcal{F}$. Similarly, define $=_\mathcal{F}$ by $f =_\mathcal{F} g$ if and only if $\{n : f(n) = g(n)\} \in \mathcal{F}$.

The two most common choices of filter, and the ones we focus on in this thesis, are the Fréchet filter and a free ultrafilter. The Fréchet filter consists of the cofinite subsets of ω; in this case, $<_\mathcal{F}$ and $=_\mathcal{F}$ are usually written as $<^*$ and $=^*$, and are called the *mod finite* ordering. When $\mathcal{F}$ is a free ultrafilter, the reduced power is called an *ultrapower.*

A natural first question is whether the reduced power inherits the most basic property of $\mathcal{L}$. In particular, if $\mathcal{L}$ is linearly ordered, when does the induced relation $<_\mathcal{F}$ remain a linear order? The next proposition shows that this occurs exactly when $\mathcal{F}$ is an ultrafilter.

Proposition 2.2. *If L has at least two elements, then $<_\mathcal{F}$ is a linear order if and only if $\mathcal{F}$ is an ultrafilter.*

Proof. Suppose $<_\mathcal{F}$ is a linear order. Fix $a < b$ in L and let $A \subseteq \omega$. Define $f, g \in L^\omega$ by

$$f(n) = \begin{cases} a & n \in A \\ b & n \notin A \end{cases} \qquad g(n) = \begin{cases} b & n \in A \\ a & n \notin A \end{cases}$$

Then $\{n : f(n) < g(n)\} = A$ and $\{n : g(n) < f(n)\} = \omega \setminus A$. Since $<_\mathcal{F}$ is a linear order, either $f <_\mathcal{F} g$ or $g <_\mathcal{F} f$, so either $A \in \mathcal{F}$ or $\omega \setminus A \in \mathcal{F}$. Since A was arbitrary, $\mathcal{F}$ is an ultrafilter.

Conversely, let $\mathcal{F}$ be an ultrafilter on ω and let $f, g \in L^\omega$. The sets $\{n : f(n) < g(n)\}$, $\{n : f(n) > g(n)\}$, and $\{n : f(n) = g(n)\}$ partition ω, so exactly one belongs to $\mathcal{F}$. Hence exactly one of $f <_\mathcal{F} g$, $g <_\mathcal{F} f$, or $f =_\mathcal{F} g$ holds. $\square$

A natural next question is whether the reduced power inherits stronger properties of $\mathcal{L}$. In particular, if $\mathcal{L}$ is an infinite well-ordered set, that is, every nonempty subset of L has a minimum element, does $<_\mathcal{F}$ still have this property? The next proposition shows that the answer is no. We first require a definition.

Definition 2.3. If S is any unary or binary operation on L, then there is a natural operation, let's call it S^ω, on L^ω; where S^ω is simply defined coordinatewise.

Proposition 2.4. *If $\mathcal{L}$ has any infinite chain, then there is a subset of L^ω that does not have a $<_\mathcal{F}$-minimum element.*

Proof. If $\{\ell_n : n \in \omega\}$ is a chain in $(L,<)$, then let $\vec{x}_0 \in L^\omega$ be the function sending n to ℓ_n.

Suppose that S is a unary operation defined on L that satisfies that $S(\ell_{n+1}) = \ell_n$ for all n. For completeness, assume that $S(\ell_0) = \ell_0$. For each $k \in \omega$, make sense of $\vec{x}_{k+1} = S^\omega(x_k)$. So we are defining a sequence of sequences:

$$\vec{x}_1 = S^\omega(\vec{x}_0), \quad \vec{x}_2 = S^\omega(\vec{x}_1), \quad \vec{x}_3 = S^\omega(\vec{x}_2), \quad \cdots$$

n	0	1	2	3	$\cdots$
$\vec{x}_0$	l_0	l_1	l_2	l_3	$\cdots$
$\vec{x}_1$	l_0	l_0	l_1	l_2	$\cdots$
$\vec{x}_2$	l_0	l_0	l_0	l_1	$\cdots$
$\vec{x}_3$	l_0	l_0	l_0	l_0	$\cdots$

The set $\{n : \vec{x}_{k+1}(n) < \vec{x}_k(n)\}$ is cofinite, hence belongs to the Fréchet filter, and also belongs to any free ultrafilter. $\square$

So far we have asked whether reduced powers preserve basic order-theoretic properties such as linearity and well-ordering. But these are only part of the picture. In the case of ultrapowers, a classical result known as Łoś's theorem tells us that every property expressible in first-order logic transfers from $\mathcal{L}$ to the ultrapower. Roughly, any statement built from elements, the order relation, logical connectives, and quantifiers that holds in $\mathcal{L}$ continues to hold in the ultrapower. In this precise sense, the ultrapower is indistinguishable from the original structure by first-order means. To say more, we need invariants that go beyond first-order expressibility. The following two do exactly this, and will play a central role in the study of gaps later in this thesis.

(1) The *unbounded number*, $\mathfrak{b}_{\mathcal{L}}$, is the minimal cardinality of a subset of L that has no upper bound. This is only of interest when the ordering is directed upward.
(2) The *dominating number*, $\mathfrak{d}_{\mathcal{L}}$, is the minimal cardinality of a dominating subset $D \subseteq L$, that is, a set D such that every $\ell \in L$ is below some member of D.

When $\mathcal{L} = (\omega^\omega, <^*)$, these are written simply as $\mathfrak{b}$ and $\mathfrak{d}$. For a structure such as $(\omega, <)$, both invariants are trivially ω. In reduced powers, however, they become nontrivial and will govern much of the gap structure studied later in this thesis.

Proposition 2.5. *Let $\mathcal{U}$ be a free ultrafilter on ω.*

(1) *For $<^*$ and for $<_{\mathcal{U}}$ alike, ω^ω is cofinal in each of $\mathbb{Z}^\omega$, $\mathbb{Q}^\omega$, and $\mathbb{R}^\omega$. In particular, all four structures share the same unboundedness and dominating numbers.*
(2) *If $(L,<)$ is a linear order with no maximum element, then $\mathfrak{b}_{\mathcal{U}} = \mathfrak{d}_{\mathcal{U}}$ in the ultrapower $(L^\omega, <_{\mathcal{U}})$.*
(3) *No strictly increasing ω-sequence in $(\mathbb{R}^\omega, <_{\mathcal{U}})$ has a least upper bound.*

Proof. For (1), let $f \in \mathbb{R}^\omega$. Define $g(n) = \max(\lceil f(n) \rceil + 1, 0)$. Then $g \in \omega^\omega$ and $f(n) < g(n)$ for every n, so $\{n : f(n) < g(n)\} = \omega \in \mathcal{F}$ for any filter $\mathcal{F}$. In particular $f <^* g$ and $f <_{\mathcal{U}} g$. Thus ω^ω is cofinal in $\mathbb{R}^\omega$ with respect to both $<^*$ and $<_{\mathcal{U}}$, and the same argument applies to $\mathbb{Q}^\omega$ and $\mathbb{Z}^\omega$.

For (2), by Proposition 2.2 the ultrapower $(L^\omega, <_{\mathcal{U}})$ is a linear order. If L has no maximum element, then neither does the ultrapower. Now every unbounded set is dominating, since in a linear order any element is below some member of an unbounded set. Conversely, every dominating set is unbounded, since an upper bound on a dominating set would bound every element of L^ω, producing a maximum. Hence $\mathfrak{b}_{\mathcal{U}} = \mathfrak{d}_{\mathcal{U}}$.

For (3), suppose $\langle f_n : n \in \omega\rangle$ is strictly $<_{\mathcal{U}}$-increasing in $(\mathbb{R}^\omega, <_{\mathcal{U}})$ and let g be a least upper bound. By modifying each f_n on a set outside $\mathcal{U}$, we may assume the sequence is pointwise monotone. For each n, define $U_n = \{m : f_n(m) < g(m)\} \in \mathcal{U}$. Since the sequence is pointwise monotone, $U_0 \supseteq U_1 \supseteq U_2 \supseteq \cdots$. The sets $U_n \setminus U_{n+1}$ are pairwise disjoint, and together with $\bigcap_{n\in\omega} U_n$ they partition U_0. Define $h \in \mathbb{R}^\omega$ by

$$h(m) = \begin{cases} \dfrac{f_n(m) + g(m)}{2} & \text{if } m \in U_n \setminus U_{n+1} \text{ for some } n \\ \dfrac{f_m(m) + g(m)}{2} & \text{if } m \in \displaystyle\bigcap_{n\in\omega} U_n \\ g(m) & \text{otherwise.} \end{cases}$$

Then $h(m) < g(m)$ on all of $U_0 \in \mathcal{U}$, so $h <_{\mathcal{U}} g$. For each fixed k, on $U_n \setminus U_{n+1}$ with $n \geq k$ we have $f_k(m) \leq f_n(m) < h(m)$, and on $\bigcap_n U_n$ with $m \geq k$ we have $f_k(m) \leq f_m(m) < h(m)$. Together these cover U_k up to a finite set, so $f_k <_{\mathcal{U}} h$. Thus h is an upper bound strictly below g, a contradiction. $\square$

Proposition 2.5 describes several features of the ultrapower. Part (1) shows that, just as ω is cofinal in $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$, the same cofinality persists in the reduced power: ω^ω is cofinal in $\mathbb{Z}^\omega$, $\mathbb{Q}^\omega$, and $\mathbb{R}^\omega$ under both the mod finite and ultrafilter orderings. Part (2) shows that in any ultrapower of a linear order with no maximum element, the unboundedness and dominating numbers coincide. Most striking, however, is part (3): the defining feature of the real line is completeness, yet the ultrapower $(\mathbb{R}^\omega, <_{\mathcal{U}})$ fails to carry this property. In fact, it is nowhere complete in the sense that no strictly increasing ω-sequence has a least upper bound. It is precisely this failure of completeness that gives rise to gaps in the reduced power, which we now turn to study.

Definition 2.6. Let $(L, <)$ be an ordered set. A pair (A, B) of subsets of L forms a *gap* if

(1) $a < b$ for every $a \in A$ and $b \in B$ (we write $A < B$), and
(2) the set $A \cup B$ is a $<$-chain,
(3) for every $c \in L$, either $c \leq a$ for some $a \in A$, or $b \leq c$ for some $b \in B$.

That is, no element of L lies strictly between all of A and all of B. We say that (A, B) is a (κ, λ)-*gap* if $\kappa = \mathfrak{b}_{(A,<)}$ and $\lambda = \mathfrak{b}_{(B,>)}$.

Having seen that bounded increasing ω-sequences in the ultrapower need not have least upper bounds, we now aim to quantify this failure: given such a gap, what is the cofinality coming down from above? We begin with the mod finite ordering.

Theorem 2.7. *If (A, B) is an (ω, λ)-gap or a (κ, ω)-gap with respect to $<^*$, then $\mathfrak{b} \leq \kappa, \lambda \leq \mathfrak{d}$.*

Proof. Let (A, B) be an (ω, λ)-gap. Let $\langle a_n : n \in \omega\rangle$ be a $<^*$-increasing sequence in A and cofinal in $(A, <^*)$. Replacing each a_n by an $=^*$-equivalent function if necessary, we may assume that for all $n, k \in \omega$, $a_n(k) < a_{n+1}(k)$. Similarly let $\langle b_\alpha : \alpha < \lambda\rangle$ be a $<^*$-decreasing sequence in B and cofinal in $(B, >^*)$.

For each $\alpha < \lambda$, define $f_\alpha \in \omega^\omega$ by letting $f_\alpha(n)$ be the least $k \in \omega$ such that $a_n(m) < b_\alpha(m)$ for all $m \geq k$. That is, $f_\alpha(n)$ measures how long it takes for b_α to dominate a_n. This is well-defined because $a_n <^* b_\alpha$. For $n \in \omega$, let $\bar{n} \in \omega^\omega$ denote the constant function with value n.

We first show that for all $n \in \omega$ and $\alpha < \beta < \lambda$, we have $\bar{n} <^* f_\alpha \leq^* f_\beta$. Fix $n \in \omega$ and $\alpha < \lambda$. If $f_\alpha(j) \leq n$, then by definition of $f_\alpha(j)$ we have $a_j(m) < b_\alpha(m)$ for all $m \geq f_\alpha(j)$, and in particular $a_j(n) < b_\alpha(n)$. Since for fixed n the sequence $j \mapsto a_j(n)$ is strictly increasing, this can happen for only finitely many j. Hence $f_\alpha(j) > n$ for all sufficiently large j, so $\bar{n} <^* f_\alpha$.

Now fix $\alpha < \beta < \lambda$. Since $b_\beta <^* b_\alpha$, choose $m_0 \in \omega$ such that $b_\beta(m) < b_\alpha(m)$ for all $m \geq m_0$. Also, by the first part, $\bar{m_0} <^* f_\beta$, so for all sufficiently large j we have $f_\beta(j) > m_0$. For such j, if $m \geq f_\beta(j)$, then $m \geq m_0$ and hence $a_j(m) < b_\beta(m) < b_\alpha(m)$. Therefore $f_\alpha(j) \leq f_\beta(j)$ for all sufficiently large j, and so $f_\alpha \leq^* f_\beta$.

We claim that the family $\{f_\alpha : \alpha < \lambda\}$ is $<^*$-unbounded. Suppose, for a contradiction, that $g \in \omega^\omega$ is a strictly increasing function and that $f_\alpha <^* g$ for all $\alpha < \lambda$. Modifying finitely many values of g if necessary, we may assume that $g(0) = 0$. Define $b_g \in \omega^\omega$ by

$$b_g(m) = a_n(m) \qquad \text{whenever } g(n) \leq m < g(n+1).$$

We show that $A <^* b_g <^* B$, contradicting that (A, B) is a gap. Fix $i \in \omega$. If $m \geq g(i+1)$ and $g(n) \leq m < g(n+1)$, then $n \geq i+1$, so by pointwise monotonicity $a_i(m) < a_n(m) = b_g(m)$. Hence $a_i <^* b_g$. Since $\langle a_n : n \in \omega\rangle$ is cofinal in $(A, <^*)$, it follows that $a <^* b_g$ for all $a \in A$.

Now fix $\alpha < \lambda$. Choose n_0 such that $f_\alpha(n) < g(n)$ for all $n \geq n_0$. If $m \geq g(n_0)$ and $g(n) \leq m < g(n+1)$, then $n \geq n_0$, so $m \geq g(n) > f_\alpha(n)$. Hence, by definition of $f_\alpha(n)$, $a_n(m) < b_\alpha(m)$. Therefore $b_g(m) = a_n(m) < b_\alpha(m)$ eventually, so $b_g <^* b_\alpha$. Since $\langle b_\alpha : \alpha < \lambda\rangle$ is cofinal in $(B, >^*)$, it follows that $b_g <^* b$ for all $b \in B$. This contradicts that (A, B) is a gap. Hence $\{f_\alpha : \alpha < \lambda\}$ is $<^*$-unbounded, and therefore $\mathfrak{b} \leq \lambda$.

Now let $\{g_\xi : \xi < \mathfrak{d}\} \subseteq \omega^\omega$ be a dominating family. For each $\alpha < \lambda$, choose $\xi_\alpha < \mathfrak{d}$ such that $f_\alpha \leq^* g_{\xi_\alpha}$. If $\mathfrak{d} < \lambda$, then since λ is regular, there is some $\xi < \mathfrak{d}$ such that $\Gamma = \{\alpha < \lambda : \xi_\alpha = \xi\}$ is cofinal in λ. But then for every $\beta < \lambda$, choose $\alpha \in \Gamma$ with $\beta \leq \alpha$. Since $\langle f_\alpha : \alpha < \lambda\rangle$ is $\leq^*$-increasing, we get $f_\beta \leq^* f_\alpha \leq^* g_\xi$, so g_ξ bounds the whole family $\{f_\alpha : \alpha < \lambda\}$, contradicting the previous paragraph. Thus $\lambda \leq \mathfrak{d}$. Therefore $\mathfrak{b} \leq \lambda \leq \mathfrak{d}$. The proof for a (κ, ω)-gap follows similarly. $\square$

So in the mod finite ordering on ω^ω, the cofinality of the gap is pinned between $\mathfrak{b}$ and $\mathfrak{d}$. It is therefore natural to ask what happens in the ultrapower of ω^ω.

Theorem 2.8. *Let $\mathcal{U}$ be an ultrafilter. There is a unique κ (and it is greater than ω) such that the $\mathcal{U}$-ultrapower of $(\omega^\omega, <_\mathcal{U})$ has an (ω,κ)-gap. Similarly any (λ,ω)-gap will also have $\lambda = \kappa$.*

Proof. The value of κ comes from the canonical gap (A_0, B_0) where A_0 is the countable set of constant functions and $b \in B_0$ iff for every $k \in \omega$, $\{n : b(n) > k\} \in \mathcal{U}$.

Suppose (A, B) is an (ω,λ)-gap, we may choose a sequence $\{a_k : k \in \omega\}$ that is strictly $<_\mathcal{U}$-increasing and $<_\mathcal{U}$-cofinal in A. However, our proof requires a stronger sequence, so we construct a new sequence $\{\dot{a}_k : k \in \omega\}$ such that $\dot{a}_k =_\mathcal{U} a_k$ for all k, and the sequence is pointwise strictly increasing in the following way.

Define $\dot{a}_k : \omega \to \omega$ inductively. Let $\dot{a}_0 = a_0$. Given $\dot{a}_0, \ldots, \dot{a}_k$, define

$$u_{k+1} = \bigcap_{i \le k} \{n : \dot{a}_i(n) < a_{k+1}(n)\}$$

Each set $\{n : \dot{a}(n)_i < a_{k+1}(n)\} \in \mathcal{U}$ because $\dot{a}_i =_\mathcal{U} a_i <_\mathcal{U} a_{k+1}$. Since $\mathcal{U}$ is closed under finite intersections, $u_{k+1} \in \mathcal{U}$. Now set

$$\dot{a}_{k+1}(n) = \begin{cases} a_{k+1}(n) & n \in u_{k+1} \\ \max\{\dot{a}_0(n), \ldots, \dot{a}_k(n)\} + 1 & n \notin u_{k+1} \end{cases}$$

Then $\dot{a}_{k+1} =_\mathcal{U} a_{k+1}$ and $\dot{a}_k(n) < \dot{a}_{k+1}(n)$ for all n. Using this new sequence, choose a $<_\mathcal{U}$-descending sequence $(b_\xi)_{\xi<\kappa}$ cofinal in B_0. For each $\xi < \kappa$, define $x_{b_\xi}(n) = \dot{a}_{b_\xi(n)}(n)$. We show that $\{x_{b_\xi} : \xi < \kappa\}$ lies above A. Fix $\dot{a}_k \in A$. Since $b_\xi \in B_0$, the set $\{n : b_\xi(n) > k\} \in \mathcal{U}$. On this set,

$$x_{b_\xi}(n) = \dot{a}_{b_\xi(n)}(n) > \dot{a}_k(n),$$

since the sequence $\dot{a}_k$ is pointwise strictly increasing. Hence $x_{b_\xi} >_\mathcal{U} \dot{a}_k$. Thus each x_{b_ξ} lies above all $\dot{a}_k$, and since this sequence is cofinal in A, it follows that $A <_\mathcal{U} \{x_{b_\xi} : \xi < \kappa\}$.

The mapping $b_\xi \mapsto x_{b_\xi}$ is order preserving. Since $(b_\xi)_{\xi<\kappa}$ is $<_\mathcal{U}$-decreasing, for $\gamma < \xi$, we have $b_\xi <_\mathcal{U} b_\gamma$, i.e. $S = \{n : b_\xi(n) < b_\gamma(n)\}$ lies in $\mathcal{U}$. For $n \in S$,

$$x_{b_\xi}(n) = \dot{a}_{b_\xi(n)}(n) < \dot{a}_{b_\gamma(n)}(n) = x_{b_\gamma}(n)$$

since the sequence $(\dot{a}_k(n))_{k\in\omega}$ is strictly increasing. Therefore $x_{b_\xi} <_\mathcal{U} x_{b_\gamma}$ and $\{x_{b_\xi} : \xi < \kappa\}$ is decreasing, it is left to show that the sequence is cofinal in B.

Assume for contradiction that the decreasing sequence $\langle x_{b_\xi} : \xi < \kappa\rangle$ is not cofinal in B. Then there exists $b \in B$ such that $b \le_\mathcal{U} x_{b_\xi}$ for all $\xi < \kappa$. We define a $\dot{b} : \omega \to \omega$ such that $x_{\dot{b}} \le_\mathcal{U} b$ and $\dot{b} \in B_0$. For each n, since the sequence $k \mapsto \dot{a}_k(n)$ is strictly increasing, it is unbounded in ω. Hence there exists a least $m(n) \in \omega$ such that $\dot{a}_{m(n)}(n) > b(n)$. Now set

$$\dot{b}(n) = \begin{cases} m(n) - 1 & \text{if } m(n) > 0 \\ 0 & \text{if } m(n) = 0 \end{cases}$$

By construction, for every n we have $\dot{a}_{\dot{b}(n)}(n) \le b(n)$, so if we define $x_{\dot{b}}(n) = \dot{a}_{\dot{b}(n)}(n)$, then $x_{\dot{b}} \le b$ pointwise, hence $x_{\dot{b}} \le_\mathcal{U} b$. To show $\dot{b} \in B_0$, fix any $k \in \omega$. On the set $\{n : \dot{a}_k(n) < b(n)\} \in \mathcal{U}$, the first index where $\dot{a}_{m(n)}(n) > b(n)$ must satisfy $m(n) \geq k+1$, hence $\dot{b}(n) = m(n) - 1 \geq k$. Therefore $\{n : \dot{b}(n) > k\} \in \mathcal{U}$, thus $\dot{b} \in B_0$. Now we get $x_{\dot{b}} \leq_{\mathcal{U}} b \leq_{\mathcal{U}} x_{b_\xi}$, but our mapping $b \mapsto x_b$ is order preserving so we conclude $\dot{b} \leq_{\mathcal{U}} b_\xi$ for all $\xi < \kappa$ contradicting that $(b_\xi)_{\xi<\kappa}$ is $<_{\mathcal{U}}$-decreasing cofinal in B_0.

Thus $\langle x_{b_\xi} : \xi < \kappa \rangle$ is cofinal in B, so $\lambda \leq \kappa$. Since we are working with a free ultrafilter, the same argument applies symmetrically, yielding $\kappa \leq \lambda$. Hence $\lambda = \kappa$. The statement for (λ, ω)-gaps follows a similar argument, completing the proof. □

Thus in the ultrapower, every (ω, λ)-gap, and likewise every (λ, ω)-gap, has the same cofinality coming down, namely κ. Unlike the mod finite ordering, here this cofinality is a single cardinal determined solely by the ultrafilter. It is therefore natural to regard κ as an invariant of $\mathcal{U}$, and to give it a name.

Definition 2.9. The *lower cofinality* of a free ultrafilter $\mathcal{U}$ on ω is the unique cardinal $\kappa_{\mathcal{U}}$ such that $(\omega^\omega, <_{\mathcal{U}})$ has an $(\omega, \kappa_{\mathcal{U}})$-gap.

2.2. **Ultrafilters.** A natural question is whether $\kappa_{\mathcal{U}}$ actually varies with the choice of ultrafilter. Since $\kappa_{\mathcal{U}}$ is an invariant of the ultrapower, this is one way to ask whether different ultrafilters can produce non-isomorphic ultrapowers. If every two ultrapowers were order-isomorphic, then $\kappa_{\mathcal{U}}$ would collapse to a single value and carry no distinguishing information. If instead the ultrapowers can differ, $\kappa_{\mathcal{U}}$ has a chance to witness that difference. The next theorem shows that under the Continuum Hypothesis, the ultrapowers of any two free ultrafilters on ω are order-isomorphic, and hence indistinguishable as orders.

Theorem 2.10 (CH)**.** *If $\mathcal{U}$ and $\mathcal{V}$ are ultrafilters on ω, then there is a bijection f from ω^ω to ω^ω (or from $\mathbb{R}^\omega$ to $\mathbb{R}^\omega$) that induces an isomorphism with respect to $<_{\mathcal{U}}$ to $<_{\mathcal{V}}$.*

Proof. We construct the desired isomorphism by building a sequence of partial isomorphisms. Assuming CH, we have $|\omega^\omega| = \omega_1$, so both ordered sets $(\omega^\omega, <_{\mathcal{U}})$ and $(\omega^\omega, <_{\mathcal{V}})$ can be enumerated as $\{a_\xi : \xi < \omega_1\}$ and $\{b_\xi : \xi < \omega_1\}$ respectively. We then construct the bijection by transfinite recursion.

Assume inductively that for each $\alpha < \omega_1$ we have constructed a countable partial isomorphism $\phi_\alpha : D_\alpha \to R_\alpha$ such that for all $x, y \in D_\alpha$,

$$x <_{\mathcal{U}} y \quad \text{iff} \quad \phi_\alpha(x) <_{\mathcal{V}} \phi_\alpha(y),$$

and whenever $x \in D_\alpha$, all iterated successor and predecessor shifts of x also lie in D_α, with the map preserving these shifts.

For $x \in \omega^\omega$, let $T(x) = 1 + x$ denote coordinatewise addition of 1, and let $S(x) = x - 1$ denote coordinatewise subtraction of 1, leaving coordinates equal to 0 unchanged. Mapping all such shifts ensures that immediate successor and predecessor relations are preserved during the construction, preventing dead-ends later in the recursion.

For the base stage, choose $\phi_0(a_0) = b_0$ and extend this by mapping all shifts:

$$\phi_0(T^n(a_0)) = T^n(b_0), \qquad \phi_0(S^n(a_0)) = S^n(b_0), \qquad n \in \omega.$$

This yields the initial partial isomorphism with countable domain and range.

For the successor step (forth direction), suppose ϕ_α has been constructed. Let a_β be the first element in the enumeration not yet contained in the domain. Consider the sets

$$L = \{\phi_\alpha(x) : x \in D_\alpha \text{ and } x <_{\mathcal{U}} a_\beta\}, \qquad R = \{\phi_\alpha(x) : x \in D_\alpha \text{ and } a_\beta <_{\mathcal{U}} x\}.$$

Both sets are countable. Since there are no (ω, ω)-gaps in the ultrafilter ordering, there exists b lying above every element of L and below every element of R. We then set $\phi_{\alpha+1}(a_\beta) = b$ and extend the map by including all successor and predecessor shifts. A symmetric construction yields the back direction of the successor step.

At a limit stage $\lambda < \omega_1$, define

$$\phi_\lambda = \bigcup_{\alpha<\lambda} \phi_\alpha.$$

Since the sequence $\langle \phi_\alpha : \alpha < \lambda \rangle$ is increasing, this union is well-defined. The domain and range remain countable unions of countable sets, hence countable, and order preservation is maintained because each ϕ_α is already order-preserving. Thus ϕ_λ is again a partial isomorphism.

Finally, define

$$\Phi = \bigcup_{\alpha<\omega_1} \phi_\alpha.$$

This is well-defined since the construction is increasing. By the forth and back steps, every element of each enumeration eventually appears in the domain and range, so Φ is a bijection from ω^ω onto ω^ω. Since each ϕ_α preserves the order, the union Φ also preserves the order, and hence induces an isomorphism between $(\omega^\omega, <_{\mathcal{U}})$ and $(\omega^\omega, <_{\mathcal{V}})$. This completes the proof. $\square$

But what if we do not assume the Continuum Hypothesis? What, then, can be said about ultrafilters and the ultrapowers they produce? To approach this question, we need a way to compare ultrafilters that is meaningful in our setting. This leads naturally to the following notion of equivalence.

Definition 2.11. Ultrafilters $\mathcal{U}$ and $\mathcal{V}$ are said to be RK-equivalent if there is a bijection (permutation) f on ω satisfying that $f[U] \in \mathcal{V}$ for every $U \in \mathcal{U}$.

This notion is useful in our context because it is reflected directly in the associated ultrapowers: RK-equivalent ultrafilters, as we shall show, give rise to order-isomorphic ultrapowers.

Proposition 2.12. *If f is a permutation on ω and if $\mathcal{U}$ is an ultrafilter on ω, then setting $\mathcal{V} = \{f[U] : U \in \mathcal{U}\}$ is an ultrafilter on ω and there is an isomorphism from $(\omega^\omega, <_{\mathcal{U}})$ to $(\omega^\omega, <_{\mathcal{V}})$. In particular $\mathfrak{d}_{\mathcal{U}} = \mathfrak{d}_{\mathcal{V}}$ and $\kappa_{\mathcal{U}} = \kappa_{\mathcal{V}}$.*

Proof. Since f is bijective, images of sets preserve finite intersections, inclusions, and complements, so $\mathcal{V} = \{f[U] : U \in \mathcal{U}\}$ is an ultrafilter. It remains to show that $(\omega^\omega, <_{\mathcal{U}})$ and $(\omega^\omega, <_{\mathcal{V}})$ are order isomorphic.

Define $T : \omega^\omega \to \omega^\omega$ by $T(g) = g \circ f^{-1}$; then T is bijective. We show that T is order preserving, i.e. if $h <_{\mathcal{U}} g$ then $T(h) <_{\mathcal{V}} T(g)$. Since

$$T(g)(m) = g(f^{-1}(m)) \quad \text{and} \quad T(h)(m) = h(f^{-1}(m)),$$

the inequality $h(f^{-1}(m)) < g(f^{-1}(m))$ holds whenever $m = f(n)$ for some n with $h(n) < g(n)$, because then

$$h(f^{-1}(f(n))) = h(n) < g(n) = g(f^{-1}(f(n))).$$

Thus the inequality holds for all $m \in f(A)$, where $A = \{n : h(n) < g(n)\}$. Whenever $A \in \mathcal{U}$, we have $h <_{\mathcal{U}} g$, and by definition $f[A] \in \mathcal{V}$, so $T(h) <_{\mathcal{V}} T(g)$. Therefore T is an isomorphism between $(\omega^\omega, <_{\mathcal{U}})$ and $(\omega^\omega, <_{\mathcal{V}})$. $\square$

We close this section with two themes that will guide what follows. On the one hand, we have introduced RK-equivalence as a way of comparing ultrafilters, and Proposition 2.12 shows that RK-equivalent ultrafilters give rise to isomorphic ultrapowers. Thus, if one hopes to find inherently different ultrapowers, one must first know that ultrafilters which are not RK-equivalent actually exist. This will be taken up later, when we develop further combinatorial machinery for constructing and distinguishing ultrafilters.

On the other hand, we have introduced the lower cofinality $\kappa_{\mathcal{U}}$ as an invariant of the ultrapower. If all ultrapowers were isomorphic, then $\kappa_{\mathcal{U}}$ would necessarily take a single value. If instead $\kappa_{\mathcal{U}}$ can vary with $\mathcal{U}$, then it provides a concrete way to witness non-isomorphic ultrapowers. We return to this question in Section 6, where different values of $\kappa_{\mathcal{U}}$ are realized.

3. Hausdorff gaps

In the earlier sections we carefully investigated $(\omega, \mathfrak{b})$-gaps. It is known that $\omega_1 \leq \mathfrak{b} \leq \mathfrak{d} \leq \mathfrak{c}$ but remarkably, no more about these inequalities can be said in ZFC: every combination of strict or non-strict inequalities is consistent.

The next theorem, due to Hausdorff, gives a gap whose existence is provable in ZFC alone via a transfinite construction of length ω_1.

Theorem 3.1. *There is an (ω_1, ω_1)-gap in $<^*$.*

Hausdorff's original construction is carried out in $(\mathcal{P}(\omega), \subseteq^*)$, the power set of ω ordered by almost inclusion, where $a \subseteq^* b$ means that $a \setminus b$ is finite. This is the natural setting for the argument, and we adopt it for the rest of the section. The connection with $<^*$ is given by identifying each set $a \subseteq \omega$ with its characteristic function $\chi_a \in \{0,1\}^\omega \subseteq \omega^\omega$. Under this identification,

$$a \subseteq^* b \quad \Longleftrightarrow \quad \chi_a \leq^* \chi_b.$$

Thus $(\mathcal{P}(\omega), \subseteq^*)$ embeds naturally into $(\omega^\omega, \leq^*)$. Moreover, if $f \in \omega^\omega$ satisfies $\chi_a \leq^* f \leq^* \chi_b$ for every $a \in \mathcal{A}$ and $b \in \mathcal{B}$, then $\{n \in \omega : f(n) \geq 1\}$ satisfies $a \subseteq^* \{n \in \omega : f(n) \geq 1\} \subseteq^* b$ for all such a and b. It follows that a gap in $(\mathcal{P}(\omega), \subseteq^*)$ yields a gap of the same type in $(\omega^\omega, \leq^*)$. It therefore suffices to carry out the construction here. With the setting fixed, we now define the structure we are after.

Definition 3.2. A Hausdorff gap is a pair, $\mathcal{A}, \mathcal{B}$, of indexed families of subsets of ω satisfying:

(1) $\mathcal{A} = \{a_\alpha : \alpha < \omega_1\}$,
(2) $\mathcal{B} = \{b_\alpha : \alpha < \omega_1\}$,

(3) for all $\alpha < \beta < \omega_1$, $a_\alpha \subset^* a_\beta \subset^* b_\beta \subset^* b_\alpha$,
(4) for all $\beta < \omega_1$ and $n \in \omega$, the set $\{\alpha < \beta : |a_\alpha \setminus b_\beta| < n\}$ is finite.

The definition follows the basics of a gap: $\mathcal{A}$ is a mod finite increasing chain, and $\mathcal{B}$ is a mod finite decreasing chain whose members mod finite contain every member of $\mathcal{A}$. The key condition is the last one, which is what guarantees that the pair $(\mathcal{A}, \mathcal{B})$ forms a mod finite gap, as we prove next.

Proposition 3.3. *If $(\mathcal{A}, \mathcal{B})$ is a Hausdorff gap, then there is no set Y that* ***splits the gap*** *in the sense that every member of $\mathcal{A}$ is mod finite contained in Y and every member of $\mathcal{B}$ mod finite contains Y.*

Proof. Let $Y \subseteq \omega$ and assume toward a contradiction that Y splits the gap, meaning $a_\alpha \subset^* Y$ and $Y \subset^* b_\alpha$ for every $\alpha < \omega_1$. For each α, choose $n_\alpha \in \omega$ such that $a_\alpha \setminus Y \subset n_\alpha$ and $Y \setminus b_\alpha \subset n_\alpha$.

Since ω_1 is uncountable and we only have countably many n_α there is some $n \in \omega$ for which $\Gamma = \{\alpha \in \omega_1 : n_\alpha = n\}$ is uncountable. Since Γ is uncountable, then there exist some $\beta \in \Gamma$ such that $\Gamma \cap \beta$ is infinite, i.e. β has infinitely many predecessors in Γ. Fix such a β. Since $\Gamma \cap \beta$ is infinite and condition (4) guarantees that only finitely many $\alpha < \beta$ satisfy $|a_\alpha \setminus b_\beta| < n$, we may pick $\alpha \in \Gamma \cap \beta$ with $|a_\alpha \setminus b_\beta| \geq n$.

However, since $\alpha, \beta \in \Gamma$, both $a_\alpha \setminus Y$ and $Y \setminus b_\beta$ are contained in n. Therefore

$$a_\alpha \setminus b_\beta \subseteq (a_\alpha \setminus Y) \cup (Y \setminus b_\beta) \subset n,$$

giving $|a_\alpha \setminus b_\beta| < n$, contradicting our choice of α. □

This establishes the no-splitting property for any Hausdorff gap; it remains to prove that such a family exists.

Proposition 3.4. *There is a Hausdorff gap.*

Proof. We construct the families $\{a_\alpha : \alpha < \omega_1\}$ and $\{b_\alpha : \alpha < \omega_1\}$ by transfinite recursion of length ω_1. Set $a_0 = \emptyset$ and $b_0 = \omega$.

Our inductive hypothesis at stage δ is that the partial families $\{a_\alpha : \alpha < \delta\}$ and $\{b_\alpha : \alpha < \delta\}$ satisfy conditions (3) and (4) of Definition 3.2, with ω_1 replaced by δ. In addition, for all $\beta < \alpha < \delta$, we have $a_\alpha \subset b_\alpha$, and both $b_\alpha \setminus a_\alpha$ and $a_\alpha \setminus a_\beta$ are infinite. This is trivially true at the base stage.

For the successor step, suppose $\delta = \beta + 1$. By the inductive hypothesis, $b_\beta \setminus a_\beta$ is infinite. Partition it into three infinite pieces C, D, and E, and define $a_{\beta+1} = a_\beta \cup C$ and $b_{\beta+1} = b_\beta \setminus D$. Then $a_\beta \subset a_{\beta+1} \subset b_{\beta+1} \subset b_\beta$. Since $E \subseteq b_{\beta+1} \setminus a_{\beta+1}$, the set $b_{\beta+1} \setminus a_{\beta+1}$ is infinite. Moreover, if $\alpha < \beta$, then $a_\alpha \subset^* a_\beta \subset a_{\beta+1}$ and $b_{\beta+1} \subset b_\beta \subset^* b_\alpha$, so condition (3) is preserved. For condition (4), note that $b_{\beta+1} \subset b_\beta$ implies $a_\alpha \setminus b_\beta \subset a_\alpha \setminus b_{\beta+1}$ for every $\alpha \leq \beta$. Hence, whenever $|a_\alpha \setminus b_{\beta+1}| < n$, we also have $|a_\alpha \setminus b_\beta| < n$. Therefore

$$\{\alpha < \beta + 1 : |a_\alpha \setminus b_{\beta+1}| < n\} \subseteq \{\alpha < \beta + 1 : |a_\alpha \setminus b_\beta| < n\},$$

and the latter set is finite by the inductive hypothesis.

For the limit step, fix a limit ordinal $\delta < \omega_1$ and assume by induction that $\{a_\alpha : \alpha < \delta\}$ and $\{b_\alpha : \alpha < \delta\}$ satisfy conditions (3) and (4) of Definition 3.2, together with the additional hypotheses that $a_\alpha \subset b_\alpha$, that $b_\alpha \setminus a_\alpha$ is infinite, and that $a_\alpha \setminus a_\beta$ is infinite whenever $\beta < \alpha < \delta$. Since δ is a countable limit ordinal, fix a strictly increasing sequence $\{\beta_k : k \in \omega\}$ cofinal in δ.

We now recursively choose a strictly increasing sequence of integers $\langle \ell_k : k \in \omega \rangle$ and define

$$b_\delta = \bigcup_{k \in \omega} (a_{\beta_k} \setminus \ell_k).$$

There are two requirements on ℓ_{k+1}. First, we require $a_{\beta_k} \setminus \ell_{k+1} \subset a_{\beta_{k+1}}$. This is achievable since $a_{\beta_k} \subset^* a_{\beta_{k+1}}$ by induction, so the difference $a_{\beta_k} \setminus a_{\beta_{k+1}}$ is finite; choose ℓ_{k+1} larger than every element of this difference (Recall that we identify each integer n with the set $\{0, \dots, n-1\}$, so $a_{\beta_k} \setminus \ell_{k+1}$ denotes the tail of a_{β_k} above ℓ_{k+1}). Second, we require that $a_{\beta_{k+1}} \setminus \ell_{k+1} \subset b_{\beta_m}$ for all $m \le k+1$. This is achievable because for each such m, $a_{\beta_{k+1}} \subset b_{\beta_{k+1}} \subset^* b_{\beta_m}$, so $a_{\beta_{k+1}} \setminus b_{\beta_m}$ is finite, and there are only finitely many $m \le k+1$; choose ℓ_{k+1} past the maximum of all these finite differences as well.

These two requirements together give $a_{\beta_k} \setminus \ell_{k+1} \subset a_{\beta_{k+1}} \subset b_{\beta_{k+1}}$, and more generally $b_\delta \setminus \ell_{k+1} \subset b_{\beta_{k+1}}$. To see the latter, note that a point in $b_\delta \setminus \ell_{k+1}$ belongs to some piece $a_{\beta_j} \setminus \ell_j$ and is $\ge \ell_{k+1}$. If $j \le k$, apply the first requirement repeatedly: $a_{\beta_j} \setminus \ell_{j+1} \subset a_{\beta_{j+1}}$, then $a_{\beta_{j+1}} \setminus \ell_{j+2} \subset a_{\beta_{j+2}}$, and so on, so any point of a_{β_j} that is $\ge \ell_{k+1}$ lies in a_{β_k}, and hence in $a_{\beta_k} \setminus \ell_{k+1} \subset a_{\beta_{k+1}} \subset b_{\beta_{k+1}}$. If $j \ge k+1$, the second requirement gives $a_{\beta_j} \setminus \ell_j \subset b_{\beta_m}$ for all $m \le j$; since $k+1 \le j$, we get $a_{\beta_j} \setminus \ell_j \subset b_{\beta_{k+1}}$.

Now we verify condition (4). Fix k and consider α with $\beta_k < \alpha \le \beta_{k+1}$; our goal is $|a_\alpha \setminus b_\delta| > k$. Since $b_\delta \setminus \ell_k \subset b_{\beta_{k+1}}$, any point of $a_\alpha \setminus b_{\beta_{k+1}}$ that is $\ge \ell_k$ is automatically outside b_δ. So if $|a_\alpha \setminus b_{\beta_{k+1}}| > \ell_k + k$, then at most ℓ_k of these points lie below ℓ_k, leaving more than k points of $a_\alpha \setminus b_{\beta_{k+1}}$ above ℓ_k, each of which is in $a_\alpha \setminus b_\delta$. Thus $|a_\alpha \setminus b_\delta| > k$.

It remains to handle the α's in $(\beta_k, \beta_{k+1}]$ for which $|a_\alpha \setminus b_{\beta_{k+1}}| \le \ell_k + k$. By condition (4) applied to $b_{\beta_{k+1}}$ (which holds by induction), taking $n = \ell_k + k + 1$, there are only finitely many $\alpha < \beta_{k+1}$ with $|a_\alpha \setminus b_{\beta_{k+1}}| < n$. So there are only finitely many exceptional α in the block $(\beta_k, \beta_{k+1}]$.

For each such α, we choose ℓ_{k+1} large enough that $(a_\alpha \setminus a_{\beta_k}) \cap [\ell_k, \ell_{k+1})$ has at least k elements. This is possible because $\alpha > \beta_k$, so $a_\alpha \setminus a_{\beta_k}$ is infinite by the extra inductive hypothesis, and since there are only finitely many exceptional α, a single choice of ℓ_{k+1} works for all of them. Every point of $(a_\alpha \setminus a_{\beta_k}) \cap [\ell_k, \ell_{k+1})$ lies outside b_δ: it is not in any earlier tail $a_{\beta_j} \setminus \ell_j$ with $j \le k$, because those are contained in a_{β_k} above ℓ_k by telescoping, and it is not in any later tail with $j \ge k+1$, because those begin only above ℓ_{k+1}. Hence $|a_\alpha \setminus b_\delta| > k$ for these α as well.

This completes the verification that $|a_\alpha \setminus b_\delta| > k$ for all $\alpha \in (\beta_k, \beta_{k+1}]$. Therefore $\{\alpha < \delta : |a_\alpha \setminus b_\delta| < n\} \subseteq \{\alpha \le \beta_n\}$, which is finite, and condition (4) holds for b_δ.

At this point we have built b_δ so that $a_\alpha \subset^* b_\delta \subset^* b_\alpha$ for all $\alpha < \delta$ and condition (4) is preserved. It remains to choose a_δ. Since δ is countable, this is a routine diagonalization. On each window $[\ell_k, \ell_{k+1})$, include into a_δ all points of $a_{\beta_k} \cap [\ell_k, \ell_{k+1})$, but omit at least one point of $b_\delta \setminus a_{\beta_k}$ from that window. Let a_δ be the resulting set. Then $a_\delta \subseteq b_\delta$ by construction, and $a_{\beta_k} \setminus a_\delta \subset \ell_k$ is finite for each k, so $a_{\beta_k} \subset^* a_\delta$ and hence $a_\alpha \subset^* a_\delta$ for all $\alpha < \delta$. Since we omit a point on each window, $b_\delta \setminus a_\delta$ is infinite. Finally, for any $\alpha < \delta$, choose k with $\alpha \le \beta_k$; then $a_\delta \setminus a_\alpha \supseteq a_{\beta_{k+1}} \setminus a_\alpha$ mod finitely many points, and the latter is infinite by induction, so $a_\delta \setminus a_\alpha$ is infinite. Thus the extra inductive hypotheses are preserved, completing the limit step.

After carrying out the construction for all $\delta < \omega_1$, the families $\{a_\alpha : \alpha < \omega_1\}$ and $\{b_\alpha : \alpha < \omega_1\}$ satisfy conditions (1)–(4) of Definition 3.2, so (A, B) is a Hausdorff gap. □

4. Independent families

In Section 2 we raised the question of how ultrafilters can differ, and how such differences may be reflected in the ultrapowers they produce. To pursue these questions, we need a way to construct ultrafilters with some control over their behavior. Independent families are the combinatorial tool we use to do this, and as we shall see later, they provide a natural way to build ultrafilters by transfinite induction. We begin by defining them.

Definition 4.1. A family $\{a_\xi : \xi < \kappa\} \subset \mathcal{P}(\omega)$ is called an independent family if for every disjoint pair of finite subsets F, H of κ, the set $\bigcap\{a_\xi : \xi \in F\} \setminus \bigcup\{a_\eta : \eta \in H\}$ is infinite.

We adopt the convention that for $a \subseteq \omega$, $a^1 = a$ and $a^0 = \omega \setminus a$. With this notation, the next proposition gives another way to think about independent families: every Boolean combination of distinct members of $\{a_\xi : \xi \in \kappa\}$ is infinite.

Proposition 4.2. *The family $\{a_\xi : \xi < \kappa\} \subset \mathcal{P}(\omega)$ is an independent family if and only if, for every finite $F \subset \kappa$ and function σ from F into 2, $a_\sigma = \bigcap\{a_\xi^{\sigma(\xi)} : \xi \in F\}$ is infinite.*

Proof. For the forward implication, fix finite $E \subset \kappa$ and $\sigma : E \to 2$. Let $F = \{\xi \in E : \sigma(\xi) = 1\}$ and $H = \{\xi \in E : \sigma(\xi) = 0\}$. Then

$$\bigcap_{\xi \in E} a_\xi^{\sigma(\xi)} = \Big(\bigcap_{\xi \in F} a_\xi\Big) \cap \Big(\bigcap_{\eta \in H} (\omega \setminus a_\eta)\Big) = \bigcap_{\xi \in F} a_\xi \setminus \bigcup_{\eta \in H} a_\eta,$$

so by independence the set is infinite.

For the reverse implication, fix disjoint finite $F, H \subset \kappa$. Let $E = F \cup H$ and define $\sigma : E \to 2$ by $\sigma \restriction F = 1$ and $\sigma \restriction H = 0$. Then again

$$\bigcap_{\xi \in E} a_\xi^{\sigma(\xi)} = \Big(\bigcap_{\xi \in F} a_\xi\Big) \cap \Big(\bigcap_{\eta \in H} (\omega \setminus a_\eta)\Big) = \bigcap_{\xi \in F} a_\xi \setminus \bigcup_{\eta \in H} a_\eta.$$

The left-hand side is infinite, hence so is the right-hand side. □

The condition for an independent family is a strong one, and it is not obvious that such a family even exists, let alone how large it can be. Note that if we can produce an independent family of some cardinality κ, then restricting to a subfamily yields an independent family of every smaller cardinality; so the real question is how large an independent family on ω can be. The next theorem shows that such a family exists and can be as large as $\mathfrak{c}$.

Theorem 4.3. *There is an independent family of cardinality $\mathfrak{c}$.*

Proof. Start with the set $2^{\mathbb{R}}$, the set of all 2-valued functions with domain $\mathbb{R}$. Let $\mathcal{I}$ be the family of all non-empty open intervals in $\mathbb{R}$ that have rational end-points. Fix an enumeration $\{I_k : k \in \omega\}$ of $\mathcal{I}$. Next, fix an enumeration $\{L_n : n \in \omega\}$ of $[\omega]^{<\omega}$, the set of all finite subsets of ω.

For each $n \in \omega$ (i.e. L_n), define the function $f_n \in 2^{\mathbb{R}}$ by

$$f_n(r) = \begin{cases} 1 & \text{if } r \in \bigcup\{I_k : k \in L_n\} \\ 0 & \text{if } r \notin \bigcup\{I_k : k \in L_n\} \end{cases}$$

It may happen that $f_n = f_m$ for some distinct values of n and m; if we worked harder we could prevent this, but we do not need to.

We index our independent family by the reals, i.e. $\{a_r : r \in \mathbb{R}\}$. So we need the definition of each a_r:

$$a_r = \{n : f_n(r) = 1\} \ .$$

Consider two disjoint finite sets of reals, F and H. We prove that $\bigcap\{a_r : r \in F\} \setminus \bigcup\{a_s : s \in H\}$ is not empty. For each $r \in F$ choose a $k_r \in \omega$ so that the interval I_{k_r} has r in it and is disjoint from H. Choose n so that $L_n = \{k_r : r \in F\}$. Then for each $r \in F$ we have $r \in \bigcup I_k : k \in L_n$, so $f_n(r) = 1$ and hence $n \in a_r$. For each $s \in H$, since every I_{k_r} is disjoint from H, we have $s \notin \bigcup I_k : k \in L_n$, so $f_n(s) = 0$ and hence $n \notin a_s$. Therefore

$$n \in \bigcap\{a_r : r \in F\} \setminus \bigcup\{a_s : s \in H\}$$

so the intersection is nonempty.

To see it is infinite, Fix $m \in \omega$ and pick $r_0 \in F$. Choose distinct $j_1, \ldots, j_m$ with $r_0 \in I_{j_i}$ and $I_{j_i} \cap H = \varnothing$. For each $i \leq m$ let

$$L^{(i)} = (L \setminus \{k_{r_0}\}) \cup \{j_i\},$$

and take n_i with $L_{n_i} = L^{(i)}$. Then $n_i \in \bigcap_{r\in F} a_r \setminus \bigcup_{s\in H} a_s$ for all i, and $n_1, \ldots, n_m$ are distinct; hence the intersection has size $\geq m$, so it is infinite. □

A family $\mathcal{F}$ of subsets of ω has the *strong finite intersection property* if every finite intersection of its members is infinite. Equivalently, $\mathcal{F}$ has this property exactly when, together with the Fréchet filter, it generates a free filter on ω. Since every free filter extends to a free ultrafilter, this gives a natural way to produce ultrafilters from suitable families of sets. In particular, every independent family has the strong finite intersection property, and we use this observation in the next corollary.

Corollary 4.4. *The cardinality of the set of ultrafilters on ω is equal to the cardinality of $2^{\mathbb{R}}$ (called $2^{\mathfrak{c}}$).*

Proof. For each $f \in 2^{\mathbb{R}}$, the family $\{a_r^{f(r)} : r \in \mathbb{R}\}$ has the strong finite intersection property, and so extends to a free ultrafilter, which we denote by $\mathcal{U}_f$. Moreover, if $f \neq g$, then $\mathcal{U}_f \neq \mathcal{U}_g$. Observe, if $f \neq g$, then there is some $r \in \mathbb{R}$ with $f(r) \neq g(r)$. Without loss of generality, $f(r) = 1$ and $g(r) = 0$. Then $a_r \in \mathcal{U}_f$ and $\omega \setminus a_r \in \mathcal{U}_g$, so $\mathcal{U}_f \neq \mathcal{U}_g$. This gives at least $2^{\mathfrak{c}}$ distinct ultrafilters.

On the other hand, every ultrafilter on ω is a subset of $\mathcal{P}(\omega)$, so the total number of ultrafilters is at most $|\mathcal{P}(\mathcal{P}(\omega))| = 2^{|\mathcal{P}(\omega)|} = 2^{\mathfrak{c}}$. Therefore the set of ultrafilters on ω has cardinality exactly $2^{\mathfrak{c}}$. □

So far we have shown that independent families exist in abundance. But the goal of this section is not merely to produce them; it is to use them to construct ultrafilters with prescribed properties. The machinery we develop from here on serves this purpose: the next few definitions and lemmas are not ends in themselves, but preparation for Theorem 4.9 (Kunen) and Corollary 4.14 (Frolík). The first step is to relativize the notion of independence to a background filter, which will let us track, as a filter grows, how much of an independent family remains compatible with it.

Definition 4.5. Let $\mathcal{F}$ be a filter on ω. A family $\{a_\alpha : \alpha \in I\}$ is *independent modulo* $\mathcal{F}$ if for every finite partial function $\sigma : I \to 2$ and every $F \in \mathcal{F}$, one has $F \cap a_\sigma \neq \emptyset$, where $a_\sigma = \bigcap_{\alpha \in \mathrm{dom}(\sigma)} a_\alpha^{\sigma(\alpha)}$.

Every independent family is, in particular, independent modulo the Fréchet filter. So this is a generalization of the earlier notion. We now record the first basic property of independence modulo a filter.

Lemma 4.6. *If $\mathcal{A}_I = \{a_\alpha : \alpha \in I\}$ is independent modulo a filter $\mathcal{F}$, then for every $Y \subset \omega$, either*

(1) *$\mathcal{A}_I$ is independent modulo the filter generated by $\mathcal{F} \cup \{Y\}$, or*
(2) *there is a finite $H \subset I$ so that $\mathcal{A}_{I \setminus H}$ is independent modulo the filter generated by the filter $\mathcal{F} \cup \{\omega \setminus Y\}$.*

Proof. Assume that (1) fails. Then $\mathcal{A}_I$ is not independent modulo the filter generated by $\mathcal{F} \cup \{Y\}$. Therefore there exists: a finite set $H \subseteq I$, a function $\sigma : H \to 2$, and a set $F \in \mathcal{F}$, such that $a_\sigma \cap F \cap Y = \emptyset$. Equivalently, $a_\sigma \cap F \subseteq \omega \setminus Y$. We will show that $\mathcal{A}_{I \setminus H}$ is independent modulo the filter generated by $\mathcal{F} \cup \{\omega \setminus Y\}$.

Let τ be any finite partial function from $I \setminus H$ into 2. Then $\sigma \cup \tau$ is a finite partial function from I into 2, and $a_{\sigma \cup \tau} = a_\sigma \cap a_\tau$. Since $\mathcal{A}_I$ is independent modulo $\mathcal{F}$, for every $F' \in \mathcal{F}$,

$$F \cap F' \cap a_{\sigma \cup \tau} = F \cap F' \cap a_\sigma \cap a_\tau \neq \emptyset$$

But $F \cap a_\sigma \subseteq \omega \setminus Y$, so

$$F' \cap (\omega \setminus Y) \cap a_\tau \supseteq F \cap F' \cap a_\sigma \cap a_\tau$$

Hence $F' \cap (\omega \setminus Y) \cap a_\tau \neq \emptyset$. Since $F' \in \mathcal{F}$ and τ was arbitrary, this shows that every set in the filter generated by $\mathcal{F} \cup \{\omega \setminus Y\}$ meets every a_τ. Therefore $\mathcal{A}_{I \setminus H}$ is independent modulo the filter generated by $\mathcal{F} \cup \{\omega \setminus Y\}$. $\square$

We now turn to the main construction of the section. Our aim is to build two ultrafilters that are RK-incomparable. We begin by isolating the target notion.

Definition 4.7. Ultrafilters $\mathcal{U}$ and $\mathcal{V}$ are said to be RK-incomparable if for any $f \in \omega^\omega$

$$f(\mathcal{U}) \neq \mathcal{V} \quad \text{and} \quad f(\mathcal{V}) \neq \mathcal{U}$$

Equivalently, $f(\mathcal{U}) \neq \mathcal{V}$ means that there is some $V \in \mathcal{V}$ such that $f^{-1}(V) \notin \mathcal{U}$.

To construct RK-incomparable ultrafilters, we will proceed by transfinite induction. At stage $\lambda < \mathfrak{c}$, we will have two filter bases $\mathcal{F}_\lambda$ and $\mathcal{G}_\lambda$, together with a set $I_\lambda \subseteq \mathfrak{c}$ of remaining indices; we write $\mathcal{A}_{I_\lambda}$ for the subfamily $\{a_\alpha : \alpha \in I_\lambda\}$, and maintain the invariant that $\mathcal{A}_{I_\lambda}$ is independent modulo each of $\mathcal{F}_\lambda$ and $\mathcal{G}_\lambda$. The key step is to extend the two filter bases while preserving this invariant, and at the same time to rule out one prescribed function $f \in \omega^\omega$ as a possible witness to RK-comparability. The next lemma provides exactly this step.

Lemma 4.8. *If f is any function from ω to ω, we can extend $\mathcal{F}_\lambda$ to $\mathcal{F}_{\lambda+1}$, $\mathcal{G}_\lambda$ to $\mathcal{G}_{\lambda+1}$ and remove a finite set from I_λ to obtain $I_{\lambda+1}$ so that*

(1) *$\mathcal{A}_{I_{\lambda+1}}$ is independent modulo each of $\mathcal{F}_{\lambda+1}$ and $\mathcal{G}_{\lambda+1}$,*
(2) *there is a set $V \in \mathcal{G}_{\lambda+1}$ such that $f^{-1}(V)$ is disjoint from some $U \in \mathcal{F}_{\lambda+1}$.*

Note that, by symmetry, we can similarly ensure that we can reverse the roles of $\mathcal{F}_{\lambda+1}$ and $\mathcal{G}_{\lambda+1}$ in the second item.

Proof. Fix $f : \omega \to \omega$. At stage λ, we have filter bases $\mathcal{F}_\lambda$, $\mathcal{G}_\lambda$ and a set $I_\lambda \subseteq \mathfrak{c}$ such that $\mathcal{A}_{I_\lambda}$ is independent modulo each of $\mathcal{F}_\lambda$ and $\mathcal{G}_\lambda$.

Choose $\alpha_\lambda \in I_\lambda$ and consider the two possible choices of either $V = a_{\alpha_\lambda}$ being in $\mathcal{G}_{\lambda+1}$ or its complement $V = \omega \setminus a_{\alpha_\lambda} = b_{\alpha_\lambda}$ being in $\mathcal{G}_{\lambda+1}$. By independence, either choice meets every member of $\mathcal{G}_\lambda$, and thus yields a filter-base extension of $\mathcal{G}_\lambda$. Let $I_{\lambda+1} = I_\lambda \setminus \{\alpha_\lambda\}$ (and we may delete finitely more indices later if needed), removing α_λ from I_λ preserves independence modulo $\mathcal{G}_{\lambda+1}$, since any Boolean combination involving V corresponds to a Boolean combination from $\mathcal{A}_{I_\lambda}$, and independence modulo $\mathcal{G}_\lambda$ already guarantees the required intersections.

Now we check that adding the corresponding set $U = \omega \setminus f^{-1}(V)$ to $\mathcal{F}_\lambda$ to obtain $\mathcal{F}_{\lambda+1}$ will preserve that $\mathcal{A}_{I_{\lambda+1}}$ is independent modulo $\mathcal{F}_{\lambda+1}$ as well.

First try $V = a_{\alpha_\lambda}$ and set $U = Y = \omega \setminus f^{-1}(V)$. Lemma 4.6 gives us two possibilities: either $\mathcal{A}_{I_{\lambda+1}}$ remains independent modulo the filter generated by $\mathcal{F}_\lambda \cup \{U\}$; in this case let $\mathcal{F}_{\lambda+1} = \mathcal{F}_\lambda \cup \{U\}$ and we are done. Otherwise Lemma 4.6 yields a finite $H_\lambda \subset I_{\lambda+1}$ such that $\mathcal{A}_{I_{\lambda+1} \setminus H_\lambda}$ is independent modulo the filter generated by $\mathcal{F}_\lambda \cup \{\omega \setminus Y\}$. In this case switch to $V = b_{\alpha_\lambda}$ and put $U = \omega \setminus Y = \omega \setminus f^{-1}(b_{\alpha_\lambda})$ into $\mathcal{F}_{\lambda+1}$. In either case the lemma follows. □

With this lemma in hand, we can now carry out the full construction.

Theorem 4.9 (Kunen)**.** *There are ultrafilters $\mathcal{U}$ and $\mathcal{V}$ that are RK-incomparable.*

Proof. Let $\mathcal{A} = \{a_\alpha : \alpha < \mathfrak{c}\}$ be an independent family. Since $|\omega^\omega| = |\mathcal{P}(\omega)| = \mathfrak{c}$, fix enumerations $\{f_\lambda : \lambda < \mathfrak{c}\}$ of ω^ω and $\{Y_\lambda : \lambda < \mathfrak{c}\}$ of $\mathcal{P}(\omega)$.

We recursively construct filter bases $\mathcal{F}_\lambda, \mathcal{G}_\lambda$ and sets $I_\lambda \subseteq \mathfrak{c}$ such that

(1) $\mathcal{A}_{I_\lambda}$ is independent modulo both $\mathcal{F}_\lambda$ and $\mathcal{G}_\lambda$,
(2) $|\mathfrak{c} \setminus I_\lambda| \le |\lambda|$.

Let $\mathcal{F}_0 = \mathcal{G}_0$ be the Fréchet filter and let $I_0 = \mathfrak{c}$. Since every Boolean combination determined by $\mathcal{A}$ is infinite, $\mathcal{A}$ is independent modulo the Fréchet filter.

At a successor stage λ, we first apply Lemma 4.8 to f_λ, and then apply the symmetric version with the roles of the two filters reversed. This yields extensions of the current filter bases and a smaller index set, still denoted $\mathcal{F}_{\lambda+1}, \mathcal{G}_{\lambda+1}, I_{\lambda+1}$, such that $\mathcal{A}_{I_{\lambda+1}}$ remains independent modulo both filters and there exist sets $U_\lambda \in \mathcal{F}_{\lambda+1}$

and $V_\lambda \in \mathcal{G}_{\lambda+1}$ with $U_\lambda \cap f_\lambda^{-1}(V_\lambda) = \emptyset$, and similarly $V'_\lambda \cap f_\lambda^{-1}(U'_\lambda) = \emptyset$ for some $U'_\lambda \in \mathcal{F}_{\lambda+1}$ and $V'_\lambda \in \mathcal{G}_{\lambda+1}$.

Next we decide whether Y_λ or its complement will be added to $\mathcal{F}_{\lambda+1}$. If either Y_λ or $\omega \setminus Y_\lambda$ already belongs to $\mathcal{F}_{\lambda+1}$, do nothing. Otherwise we apply Lemma 4.6 to $\mathcal{A}_{I_{\lambda+1}}$ and $\mathcal{F}_{\lambda+1}$; this allows us, after removing finitely many indices if necessary, to add either Y_λ or $\omega \setminus Y_\lambda$ while preserving independence modulo the resulting filter base. Do the same for $\mathcal{G}_{\lambda+1}$.

Since each application of Lemma 4.8 or 4.6 removes only finitely many indices, only finitely many indices are removed at stage λ. Hence $|\mathfrak{c} \setminus I_{\lambda+1}| \leq |\lambda + 1|$. By construction, $\mathcal{A}_{I_{\lambda+1}}$ is independent modulo both $\mathcal{F}_{\lambda+1}$ and $\mathcal{G}_{\lambda+1}$, so the inductive hypothesis is preserved.

At a limit stage δ, let $\mathcal{F}_\delta = \bigcup_{\lambda<\delta} \mathcal{F}_\lambda$ and $\mathcal{G}_\delta = \bigcup_{\lambda<\delta} \mathcal{G}_\lambda$, and define $I_\delta = \bigcap_{\lambda<\delta} I_\lambda$. Then $\mathfrak{c} \setminus I_\delta = \bigcup_{\lambda<\delta}(\mathfrak{c} \setminus I_\lambda)$, so $|\mathfrak{c} \setminus I_\delta| \leq |\delta|$. Since $\langle \mathcal{F}_\lambda : \lambda < \delta \rangle$ and $\langle \mathcal{G}_\lambda : \lambda < \delta \rangle$ are increasing chains, their unions $\mathcal{F}_\delta$ and $\mathcal{G}_\delta$ are again filters. Moreover, $\mathcal{A}_{I_\delta}$ is independent modulo both of them: if $F \in \mathcal{F}_\delta$, then $F \in \mathcal{F}_\lambda$ for some $\lambda < \delta$, and since $I_\delta \subseteq I_\lambda$, every Boolean combination from $\mathcal{A}_{I_\delta}$ meets F. The same argument applies to $\mathcal{G}_\delta$.

Finally, let $\mathcal{F} = \bigcup_{\lambda<\mathfrak{c}} \mathcal{F}_\lambda$ and $\mathcal{G} = \bigcup_{\lambda<\mathfrak{c}} \mathcal{G}_\lambda$. Since every subset of ω appears as some Y_λ, and at stage λ we ensured that either Y_λ or its complement belongs to each filter, both $\mathcal{F}$ and $\mathcal{G}$ are ultrafilters.

For each $\lambda < \mathfrak{c}$, by construction there are $U_\lambda \in \mathcal{F}_{\lambda+1} \subseteq \mathcal{F}$ and $V_\lambda \in \mathcal{G}_{\lambda+1} \subseteq \mathcal{G}$ such that $U_\lambda \cap f_\lambda^{-1}(V_\lambda) = \emptyset$, so $f_\lambda^{-1}(V_\lambda) \notin \mathcal{F}$. Hence $f_\lambda(\mathcal{F}) \neq \mathcal{G}$. Similarly, $f_\lambda(\mathcal{G}) \neq \mathcal{F}$. Since every $f \in \omega^\omega$ appears in the enumeration, $\mathcal{F}$ and $\mathcal{G}$ are RK-incomparable. $\square$

Another major result that follows from Kunen's theorem is Frolík's theorem. To state it, recall that the Boolean algebra $\mathcal{P}(\omega)/\text{fin}$ is homogeneous in the sense that for any pair of infinite, co-infinite subsets a, b of ω, there is an automorphism of $\mathcal{P}(\omega)/\text{fin}$ that sends $[a]$ to $[b]$, where $[a]$ denotes the equivalence class of a under agreement mod finite. A natural question is whether this homogeneity extends to the space of ultrafilters on $\mathcal{P}(\omega)/\text{fin}$: given any two free ultrafilters $\mathcal{U}, \mathcal{V}$ on ω, is there an automorphism φ of $\mathcal{P}(\omega)/\text{fin}$ that sends $\mathcal{U}$ to $\mathcal{V}$? Here $\varphi(\mathcal{U})$ means the set $\{\varphi([U]) : U \in \mathcal{U}\}$, and we are asking whether this equals $\{[V] : V \in \mathcal{V}\}$.

Frolík proved the answer is no. To prove this using Kunen's theorem, we first introduce a new way of constructing ultrafilters, the $\mathcal{U}$-limit, which will play a central role in the argument.

Definition 4.10. Let $\{a_n : n \in \omega\}$ be pairwise disjoint infinite subsets of ω. For each n, let $\mathcal{U}_n$ be an ultrafilter with $a_n \in \mathcal{U}_n$. Let $\mathcal{U}$ be any other ultrafilter on ω. We define the notion:

$$\mathcal{U}-\lim\{\mathcal{U}_n : n \in \omega\} = \{Y \subset \omega : \{n \in \omega : Y \in \mathcal{U}_n\} \in \mathcal{U}\}$$

and we next prove it is an ultrafilter.

Lemma 4.11. *The object $\tilde{\mathcal{U}} = \mathcal{U}-\lim\{\mathcal{U}_n : n \in \omega\}$ is an ultrafilter.*

Proof. Clearly $\omega \in \tilde{\mathcal{U}}$, since $\{n : \omega \in \mathcal{U}_n\} = \omega \in \mathcal{U}$. Similarly $\emptyset \notin \tilde{\mathcal{U}}$, since $\{n : \emptyset \in \mathcal{U}_n\} = \emptyset \notin \mathcal{U}$.

To prove closure under intersection, suppose that $Y_1, Y_2 \in \tilde{\mathcal{U}}$. Let $U_1 = \{n : Y_1 \in \mathcal{U}_n\} \in \mathcal{U}$ and $U_2 = \{n : Y_2 \in \mathcal{U}_n\} \in \mathcal{U}$. Then clearly $U_1 \cap U_2$ equals $\{n : Y_1 \cap Y_2 \in \mathcal{U}_n\} \in \mathcal{U}$. Thus $Y_1 \cap Y_2 \in \tilde{\mathcal{U}}$.

To prove upward closure, suppose $Y_1 \in \tilde{\mathcal{U}}$ and $Y_1 \subseteq Y_2$. Since each $\mathcal{U}_n$ is a filter, if $Y_1 \in \mathcal{U}_n$ then $Y_2 \in \mathcal{U}_n$, additionally $\mathcal{U}$ is a filter and hence $\{n : Y_1 \in \mathcal{U}_n\} \subseteq \{n : Y_2 \in \mathcal{U}_n\} \in \mathcal{U}$, therefore $Y_2 \in \tilde{\mathcal{U}}$.

Finally, let $Y \subseteq \omega$. Since each $\mathcal{U}_n$ is an ultrafilter, for every n exactly one of Y and $\omega \setminus Y$ belongs to $\mathcal{U}_n$. Thus $\{n : Y \in \mathcal{U}_n\} \cup \{n : \omega \setminus Y \in \mathcal{U}_n\} = \omega$ and these two sets are disjoint. Since $\mathcal{U}$ is an ultrafilter, exactly one of them is in $\mathcal{U}$. Hence exactly one of Y and $\omega \setminus Y$ belongs to $\tilde{\mathcal{U}}$. Therefore $\tilde{\mathcal{U}}$ is an ultrafilter. $\square$

Theorem 4.12. *Let $\{a_n : n \in \omega\}$ and $\{b_m : m \in \omega\}$ be two families of pairwise disjoint infinite subsets of ω; no assumption is made on how the a_n and b_m interact. For each n, let $\mathcal{U}_n$ and $\mathcal{V}_n$ be ultrafilters with $a_n \in \mathcal{U}_n$ and $b_n \in \mathcal{V}_n$. Suppose further that there are ultrafilters $\mathcal{U}$ and $\mathcal{V}$ such that*

$$\mathcal{W} = \mathcal{U}\text{--}\lim\{\mathcal{U}_n : n \in \omega\} = \mathcal{V}\text{--}\lim\{\mathcal{V}_n : n \in \omega\}.$$

Then either

(1) $\{n : (\exists m)\, b_m \in \mathcal{U}_n\} \in \mathcal{U}$, *or*
(2) $\{m : (\exists n)\, a_n \in \mathcal{V}_m\} \in \mathcal{V}$.

Intuitively, either the family $\{b_m : m \in \omega\}$ covers $\mathcal{U}$-many of the $\mathcal{U}_n$'s, or, symmetrically, the family $\{a_n : n \in \omega\}$ covers $\mathcal{V}$-many of the $\mathcal{V}_m$'s.

Proof. Suppose toward a contradiction that neither (1) nor (2) holds. Let $A = \{n : (\exists m)\, b_m \in U_n\}$ and $B = \{m : (\exists n)\, a_n \in V_m\}$. Then $A \notin \mathcal{U}$ and $B \notin \mathcal{V}$. Since $\mathcal{U}$ and $\mathcal{V}$ are ultrafilters, it follows that $U = \omega \setminus A \in \mathcal{U}$ and $V = \omega \setminus B \in \mathcal{V}$. Thus for every $n \in U$ and every m, $b_m \notin U_n$, and for every $m \in V$ and every n, $a_n \notin V_m$.

For each $n \in U$, let $\tilde{a}_n = a_n \setminus \bigcup\{b_k : k \leq n\}$. For each $m \in V$, let $\tilde{b}_m = b_m \setminus \bigcup\{a_k : k \leq m\}$. Since $b_k \notin U_n$ for every $k \leq n$, we have $\omega \setminus b_k \in U_n$, and hence $\tilde{a}_n \in U_n$ for every $n \in U$. Similarly, $\tilde{b}_m \in V_m$ for every $m \in V$.

The family $\{\tilde{a}_n : n \in U\} \cup \{\tilde{b}_m : m \in V\}$ is pairwise disjoint. Indeed, the $\tilde{a}_n$'s are pairwise disjoint since the a_n's are, and the $\tilde{b}_m$'s are pairwise disjoint since the b_m's are. If $n \in U$ and $m \in V$, then either $m \leq n$, in which case b_m was removed from $\tilde{a}_n$, or $n \leq m$, in which case a_n was removed from $\tilde{b}_m$. Hence $\tilde{a}_n \cap \tilde{b}_m = \emptyset$.

Let $A^* = \bigcup\{\tilde{a}_n : n \in U\}$ and $B^* = \bigcup\{\tilde{b}_m : m \in V\}$. For every $n \in U$ we have $\tilde{a}_n \in U_n$ and $\tilde{a}_n \subseteq A^*$, so $A^* \in U_n$. Thus $\{n : A^* \in U_n\} \supseteq U \in \mathcal{U}$, and hence $A^* \in W$. Similarly $B^* \in W$.

But A^* and B^* are disjoint, contradicting that W is an ultrafilter. Therefore one of (1) or (2) must hold. $\square$

We now connect the $\mathcal{U}$-limit construction to automorphisms of $\mathcal{P}(\omega)/\text{fin}$. The next theorem shows that if the $\mathcal{U}$-limit and $\mathcal{V}$-limit of a common family are related by an automorphism, then $\mathcal{U}$ and $\mathcal{V}$ are themselves RK-comparable. This is the bridge that will let Kunen's theorem defeat Frolík's question.

Suppose we have an automorphism φ of $\mathcal{P}(\omega)/\text{fin}$, free ultrafilters $\mathcal{U}$ and $\mathcal{V}$, a family of ultrafilters $\{\mathcal{U}_n : n \in \omega\}$, and a pairwise disjoint family of sets $\{a_n : n \in \omega\}$ satisfying:

(1) for each n, $a_n \in \mathcal{U}_n$,
(2) $\tilde{\mathcal{U}}$ is the $\mathcal{U}$-limit of $\{\mathcal{U}_n : n \in \omega\}$,
(3) $\tilde{\mathcal{V}}$ is the $\mathcal{V}$-limit of $\{\mathcal{U}_n : n \in \omega\}$,
(4) $\varphi(\tilde{\mathcal{U}}) = \tilde{\mathcal{V}}$.

Theorem 4.13. *Given the assumptions as listed above, there is a function $f \in \omega^\omega$ such that either*

$$f(\mathcal{U}) = \mathcal{V} \quad \text{or} \quad f(\mathcal{V}) = \mathcal{U}$$

i.e. $\mathcal{U}$ and $\mathcal{V}$ are RK-comparable.

Proof. For each n, let b_n be a set that is mod finite equal to $\varphi(a_n)$. Since a_n are disjoint and φ is an automorphism we can remove a finite set from each b_n, and assume $\{b_n : n \in \omega\}$ is pairwise disjoint. Also for each n, let $\mathcal{V}_n = \varphi(\mathcal{U}_n)$.

By assumption $\mathcal{W} = \varphi(\tilde{\mathcal{U}})$ is equal to $\tilde{\mathcal{V}}$. Notice that the ultrafilter $\mathcal{W}$ is therefore equal to the $\mathcal{V}$-limit of the family $\{\mathcal{U}_n : n \in \omega\}$, but moreover, because

$$\begin{aligned} X \in \varphi(\tilde{\mathcal{U}}) &\iff \varphi^{-1}(X) \in \tilde{\mathcal{U}} \\ &\iff \{n \in \omega : \varphi^{-1}(X) \in \mathcal{U}_n\} \in \mathcal{U} \\ &\iff \{n \in \omega : X \in \varphi(\mathcal{U}_n)\} \in \mathcal{U} \\ &\iff \{n \in \omega : X \in \mathcal{V}_n\} \in \mathcal{U}. \end{aligned}$$

$\mathcal{W}$ is also $\mathcal{U}$-limit of the family $\{\mathcal{V}_n : n \in \omega\}$.

Now we apply Theorem 4.12 to conclude that one of the following holds:

(1) $U = \{n : (\exists m)\ a_m \in \mathcal{V}_n\} \in \mathcal{U}$, or
(2) $V = \{m : (\exists n)\ b_n \in \mathcal{U}_m\} \in \mathcal{V}$.

Since the argument is the same, suppose $U \in \mathcal{U}$. For each $n \in U$, define $f(n) = m$ where $a_m \in \mathcal{V}_n$. For $n \notin U$, let $f(n) = 0$. We claim that $f(\mathcal{U}) = \mathcal{V}$.

It suffices to show that for every $U' \in \mathcal{U}$ with $U' \subseteq U$, we have $f(U') \in \mathcal{V}$. Since for any $X \in \mathcal{U}$, we can set $U' = X \cap U \in \mathcal{U}$. If $f(U') \in \mathcal{V}$, then $f(U') \subseteq f(X)$, and since $\mathcal{V}$ is upward closed, it follows that $f(X) \in \mathcal{V}$.

Let $U' \subset U$ be another element of $\mathcal{U}$. We first define $W = \bigcup\{a_m : m \in f(U')\}$ and show that it is an element of $\mathcal{W}$. Since for all $n \in U'$ we have $a_{f(n)} \in \mathcal{V}_n$ and W contains all such $a_{f(n)}$, clearly $\{n \in \omega : W \in \mathcal{V}_n\}$ contains U' and therefore is an element of $\mathcal{U}$. Meaning that W is in $\mathcal{U}$-limit of the set $\{\mathcal{V}_n : n \in \omega\}$ which is equal to $\mathcal{W}$, thus W is in $\mathcal{W}$. But now, $\mathcal{W}$ is also the $\mathcal{V}$-limit of the set $\{\mathcal{U}_m : m \in \omega\}$ and $W = \bigcup\{a_m : m \in f(U')\}$ hence it's in $\mathcal{V}$. we claim that $\{m : W \in \mathcal{U}_m\} = f(U')$ and thus it must be in $\mathcal{V}$.

If $m \in f(U')$, then $a_m \subseteq W$, and since $a_m \in \mathcal{U}_m$, by upward closure of $\mathcal{U}_m$ we have $W \in \mathcal{U}_m$. On the other hand, if $m \notin f(U')$, then $a_m \cap W = \varnothing$, since the family $\{a_m : m \in \omega\}$ is pairwise disjoint. Hence $a_m \subseteq \omega \setminus W$, so $\omega \setminus W \in \mathcal{U}_m$, and therefore $W \notin \mathcal{U}_m$. Thus $W \in \mathcal{U}_m$ if and only if $m \in f(U')$. This proves that $f(\mathcal{U}) = \mathcal{V}$. $\square$

With Theorem 4.13 in hand, Frolík's theorem now follows as a corollary, using Kunen's theorem.

Corollary 4.14 (Frolík–Kunen)**.** *There exist free ultrafilters $\mathcal{U}$ and $\mathcal{V}$ such that no automorphism of $\mathcal{P}(\omega)/\mathrm{fin}$ sends $\mathcal{U}$ to $\mathcal{V}$.*

Proof. By Theorem 4.9, there exist ultrafilters $\mathcal{U}$ and $\mathcal{V}$ that are RK-incomparable. Form the ultrafilters $\tilde{\mathcal{U}}$ and $\tilde{\mathcal{V}}$ as in Theorem 4.13, using the same family $\{\mathcal{U}_n : n \in \omega\}$. By the contrapositive of Theorem 4.13, since $\mathcal{U}$ and $\mathcal{V}$ are not RK-comparable, no automorphism φ of $\mathcal{P}(\omega)/\mathrm{fin}$ satisfies $\varphi(\tilde{\mathcal{U}}) = \tilde{\mathcal{V}}$. □

5. More about independence modulo a filter

Section 4 introduced independence modulo a filter and used it to construct ultrafilters with prescribed properties. In this short section we turn to a refinement: filters that are maximal with respect to this independence. The lemmas below develop the basic properties of such filters and will be used in Section 6.

Let $\mathcal{A}_I = \{a_\alpha : \alpha \in I\}$ be an independent family of infinite subsets of ω, and let $\mathcal{F}$ be a maximal filter such that $\mathcal{A}_I$ is independent modulo $\mathcal{F}$.

Lemma 5.1. *Let H be any finite subset of I and σ be any function from H into 2. Let a_σ denote the usual Boolean combination of the members of $\{a_\alpha : \alpha \in H\}$. Then the filter $\mathcal{F}_\sigma$ generated by $\mathcal{F} \cup \{a_\sigma\}$ is maximal with respect to $\mathcal{A}_{I\setminus H}$.*

Proof. We prove maximality by showing that any set satisfying the independence condition lies in $\mathcal{F}_\sigma$. Let $Y \subseteq \omega$ be such that for every $F \in \mathcal{F}$ and every usual set a_τ from the family $\mathcal{A}_{I\setminus H}$, $Y \cap F \cap a_\sigma \cap a_\tau \neq \emptyset$, i.e., Y meets every $F \cap a_\sigma \in \mathcal{F}_\sigma$ together with every Boolean combination from $A_{I\setminus H}$. We show that $Y \in \mathcal{F}_\sigma$.

Fix $F \in \mathcal{F}$ and let $a_{\tau'}$ be any usual set from $\mathcal{A}_I$. If $a_{\tau'}$ is compatible with σ, then $a_{\tau'} \cap a_\sigma = a_\sigma \cap a_\tau$ for a_τ from the previous paragraph, and hence Y meets $F \cap a_{\tau'}$. If $a_{\tau'}$ is incompatible with σ then $a_\sigma \cap a_{\tau'} = \emptyset$, so $a_{\tau'} \subseteq \omega \setminus a_\sigma$, and therefore $F \cap a_{\tau'} \subseteq Y \cup (\omega \setminus a_\sigma)$; It follows that $Y \cup (\omega \setminus a_\sigma)$ meets $F \cap a_{\tau'}$.

Thus $Y \cup (\omega \setminus a_\sigma)$ meets every set of the form $F \cap a_{\tau'}$, where $F \in \mathcal{F}$ and $a_{\tau'}$ is any usual set from the family $\mathcal{A}_I$. By maximality of $\mathcal{F}$, it follows that $Y \cup (\omega \setminus a_\sigma) \in \mathcal{F}$, and hence $[Y \cup (\omega \setminus a_\sigma)] \cap a_\sigma = Y \cap a_\sigma$ is a member of $\mathcal{F}_\sigma$ that is contained in Y. □

A set $Y \subseteq \omega$ is called *$\mathcal{F}$-positive* if $Y \cap F \neq \emptyset$ for every $F \in \mathcal{F}$. Equivalently, Y is $\mathcal{F}$-positive if and only if $\mathcal{F} \cup \{Y\}$ extends to a filter. We write $\mathcal{F}^+$ for the collection of $\mathcal{F}$-positive sets.

Lemma 5.2. *Let $\mathcal{F}$ be a filter that is maximal with respect to the family $\mathcal{A}_I$ is independent modulo $\mathcal{F}$. Then for every $\mathcal{F}$-positive set Y, there is a finite $H \subset I$ and a $\sigma : H \to 2$ such that $F \cap a_\sigma \subset^* Y$ for some $F \in \mathcal{F}$.*

Proof. Let Y be an $\mathcal{F}$-positive set. If $Y \in \mathcal{F}$, the lemma is trivial: take $F = Y$, and then $F \cap a_\sigma \subseteq Y$ for any boolean combination a_σ of the family.

Now suppose $Y \in \mathcal{F}^+$ but $Y \notin \mathcal{F}$. We claim that $\omega \setminus Y \notin \mathcal{F}$ and $\omega \setminus Y \in \mathcal{F}^+$. The first follows immediately since $Y \in \mathcal{F}^+$ means Y meets every member of $\mathcal{F}$, yet $Y \cap (\omega \setminus Y) = \emptyset$, so $\omega \setminus Y$ cannot belong to $\mathcal{F}$. For the second, if $\omega \setminus Y \notin \mathcal{F}^+$

then there exists $F \in \mathcal{F}$ with $F \cap (\omega \setminus Y) = \emptyset$, giving $F \subseteq Y$, and by upward closure of $\mathcal{F}$ this forces $Y \in \mathcal{F}$, a contradiction.

Since $\mathcal{F}$ is maximal and $\omega \setminus Y \notin \mathcal{F}$, adding $\omega \setminus Y$ to $\mathcal{F}$ violates the independence of $\mathcal{A}_I$ modulo $\mathcal{F}$. Hence there exist some $F \in \mathcal{F}$ and a suitable σ such that $\omega \setminus Y$ is disjoint from $F \cap a_\sigma$, i.e. there must be an $F \cap a_\sigma$ contained mod finite in Y. $\square$

Corollary 5.3. *If $\mathcal{F}$ is a maximal filter such that $\mathcal{A}_\mathfrak{c}$ is independent modulo $\mathcal{F}$, then every ultrafilter $\mathcal{U}$ that extends $\mathcal{F}$ has character $\mathfrak{c}$.*

Proof. Suppose for contradiction that $\mathcal{U}$ has a base $\{U_\alpha : \alpha < \kappa\}$ with $\kappa < \mathfrak{c}$. Each $U_\alpha \in \mathcal{F}^+$ since $\mathcal{U}$ extends $\mathcal{F}$, so by Lemma 5.2, for each $\alpha < \kappa$ there exist a finite $H_\alpha \subset \mathfrak{c}$, a function $\sigma_\alpha : H_\alpha \to 2$, and some $F_\alpha \in \mathcal{F}$ such that $F_\alpha \cap a_{\sigma_\alpha} \subseteq U_\alpha$.

Since each H_α is finite and $\kappa < \mathfrak{c}$, we have $\left|\bigcup_{\alpha<\kappa} H_\alpha\right| \leq \kappa < c$, so we may pick $\beta \in c \setminus \bigcup_{\alpha<\kappa} H_\alpha$, so that $\beta \notin H_\alpha$ for every $\alpha < \kappa$.

Since $\mathcal{U}$ is an ultrafilter, exactly one of a_β or $\omega \setminus a_\beta$ belongs to $\mathcal{U}$. Assume without loss of generality that $\omega \setminus a_\beta \in \mathcal{U}$. Since $\{U_\alpha\}$ is a base, there exists $\alpha < \kappa$ with $U_\alpha \subseteq \omega \setminus a_\beta$, meaning $U_\alpha \cap a_\beta = \emptyset$. But then $F_\alpha \cap a_{\sigma_\alpha} \cap a_\beta \subseteq U_\alpha \cap a_\beta = \emptyset$, yet $\beta \notin H_\alpha$ means a_β does not appear in the Boolean combination a_{σ_α}, so by independence of $\mathcal{A}_I$ modulo $\mathcal{F}$, the set $F_\alpha \cap a_{\sigma_\alpha} \cap a_\beta$ is nonempty. Contradiction. $\square$

6. Witnessing values for the lower cofinality of ultrapowers

In Section 2 we introduced the lower cofinality $\kappa_\mathcal{U}$ as an invariant of the ultrapower of ω^ω by $\mathcal{U}$, and asked whether it can distinguish different ultrapowers. Under CH the question collapses: all such ultrapowers are order-isomorphic, so $\kappa_\mathcal{U}$ takes a single value. Without CH the picture is less clear, and we introduced RK-equivalence as a first way of comparing ultrafilters: if two ultrafilters are RK-equivalent, then their ultrapowers are order-isomorphic, and hence have the same lower cofinality. Section 4 then showed, by Kunen's theorem, that not all ultrafilters are RK-equivalent.

What remains is the original question: does $\kappa_\mathcal{U}$ actually take different values as $\mathcal{U}$ varies? In this section we show that it does. For every regular cardinal κ with $\omega_1 \leq \kappa \leq \mathfrak{c}$, we construct an ultrafilter $\mathcal{U}_\kappa$ such that $\kappa_{\mathcal{U}_\kappa} = \kappa$. Since different values of $\kappa_\mathcal{U}$ force non-isomorphic ultrapowers, this yields a concrete family of ultrafilters whose ultrapowers are inherently different. The construction requires new definitions and a refinement of the independence machinery developed in Section 4, which we introduce next.

Let κ be an uncountable cardinal. A family $\mathcal{A} = \{a(\alpha, n) : \alpha \in \kappa, n \in \omega\}$ is a $\kappa \times \omega$*-independent mod finite family* if

(1) for each $\alpha \in \kappa$, the family $\{a(\alpha, n) : n \in \omega\}$ is a partition of ω into infinite sets,
(2) for every finite partial function $\sigma : \kappa \to \omega$, the set $a_\sigma = \bigcap\{a(\alpha, \sigma(\alpha)) : \alpha \in \mathrm{dom}(\sigma)\}$ is infinite.

The existence of such families for $\kappa \leq \mathfrak{c}$ follows by the same argument as Theorem 4.3, with ω in place of 2.

Definition 6.1. Let $\mathrm{Fn}(\kappa,\omega)$ denote the family of finite partial functions from κ into ω. A subset $\Sigma \subseteq \mathrm{Fn}(\kappa,\omega)$ is an *antichain* if its members are pairwise incompatible: for any distinct $\sigma,\sigma' \in \Sigma$, the union $\sigma \cup \sigma'$ is not a function.

An antichain Σ is *maximal* if adding any new finite partial function to Σ breaks the antichain property; equivalently, for every $\tau \in \mathrm{Fn}(\kappa,\omega)$, there is a $\sigma \in \Sigma$ such that $\sigma \cup \tau$ is a function.

Definition 6.2. Say that a subset D of $\mathrm{Fn}(\kappa,\omega)$ is *dense* if for every $\tau \in \mathrm{Fn}(\kappa,\omega)$, there is a $\sigma \in D$ that extends τ.

The rest of this section develops the machinery needed to prove Theorem 6.8. The next lemma records the basic connection between the two notions just introduced: every maximal antichain generates a dense set, and every dense set contains a maximal antichain.

Lemma 6.3. *For any maximal antichain $\Sigma \subset \mathrm{Fn}(\kappa,\omega)$, the set $D_\Sigma = \{\tau \in \mathrm{Fn}(\kappa,\omega) : (\exists\sigma \in \Sigma)\sigma \subset \tau\}$ is dense, and for any dense $D \subset \mathrm{Fn}(\kappa,\omega)$, there is a maximal antichain $\Sigma \subset D$ (and so $D \supset D_\Sigma$).*

Proof. For the first claim, let $\tau \in \mathrm{Fn}(\kappa,\omega)$. Since Σ is a maximal antichain, there exists $\sigma \in \Sigma$ such that $\sigma \cup \tau$ is a function. Pick any $\alpha \in \kappa \setminus \mathrm{dom}(\sigma \cup \tau)$ and any $n \in \omega$. Then $\sigma\cup\tau\cup\{(\alpha,n)\}$ properly extends σ, so it belongs to D_Σ, and it extends τ. Thus D_Σ is dense.

For the second claim, consider the poset of all antichains contained in D, ordered by inclusion. Every chain of antichains has an upper bound (its union), so by Zorn's lemma there is a maximal antichain $\Sigma \subseteq D$. To see that Σ is maximal in $\mathrm{Fn}(\kappa,\omega)$, let $\tau \in \mathrm{Fn}(\kappa,\omega)$. By density of D, there exists $\sigma' \in D$ extending τ. If σ' were incompatible with every member of Σ, we could add σ' to Σ and obtain a strictly larger antichain in D, contradicting maximality of Σ within D. Hence σ' is compatible with some $\sigma \in \Sigma$, and since $\tau \subseteq \sigma'$, it follows that τ is compatible with σ. Thus Σ is a maximal antichain in $\mathrm{Fn}(\kappa,\omega)$ contained in D. □

Lemma 6.4. *For every $\kappa \leq \mathfrak{c}$, every maximal antichain $\Sigma \subset \mathrm{Fn}(\kappa,\omega)$ is countable.*

Proof. By the construction in Theorem 4.3, with ω in place of 2, there is a countable set $\{f_k : k \in \omega\} \subseteq \omega^{\mathfrak{c}}$ such that every $\sigma \in \mathrm{Fn}(\mathfrak{c},\omega)$ is extended by some f_k. Fix such a family. Since $\kappa \leq \mathfrak{c}$, we have $\mathrm{Fn}(\kappa,\omega) \subseteq \mathrm{Fn}(\mathfrak{c},\omega)$. For each $\sigma \in \Sigma$, choose $k_\sigma \in \omega$ such that $\sigma \subseteq f_{k_\sigma}$.

We claim that $\sigma \mapsto k_\sigma$ is injective. Suppose $\sigma,\sigma' \in \Sigma$ and $k_\sigma = k_{\sigma'} = k$. Then $\sigma \subseteq f_k$ and $\sigma' \subseteq f_k$, so $\sigma \cup \sigma' \subseteq f_k$. In particular $\sigma \cup \sigma'$ is a function, which means σ and σ' are compatible. Since Σ is an antichain, $\sigma = \sigma'$. Hence Σ injects into ω, so Σ is countable. □

Lemma 6.5. *Suppose that $\mathcal{A}$ is a $\kappa\times\omega$-independent family and that $\mathcal{F}_0$ is a maximal filter such that $\mathcal{A}$ is independent modulo $\mathcal{F}_0$. For every $Y \subset \omega$, there is a maximal antichain Σ_Y of $\mathrm{Fn}(\kappa,\omega)$ satisfying that, for all $\sigma \in \Sigma$, there is an $F_\sigma \in \mathcal{F}_0$ such that either $F_\sigma \cap a_\sigma \subset^* Y$ or $F_\sigma \cap a_\sigma \cap Y =^* \emptyset$.*

Proof. Let

$$D = \{\sigma \in \mathrm{Fn}(\kappa,\omega) : \exists\, F \in F_0 \text{ s.t. } F \cap a_\sigma \subset^* Y \text{ or } F \cap a_\sigma \cap Y =^* \emptyset\}$$

It suffices to show that D is dense; the conclusion then follows by applying Lemma 6.3 to D, which yields a maximal antichain $\Sigma_Y \subset D$.

Fix $Y \subseteq \omega$ and $\tau \in \mathrm{Fn}(\kappa, \omega)$. Our goal is to extend τ to some $\sigma \in D$. By Lemma 5.1, the filter $\mathcal{F}_\tau$ generated by $\mathcal{F}_0 \cup \{a_\tau\}$ is maximal with respect to $\mathcal{A}_{\kappa \setminus \mathrm{dom}(\tau)}$. Since $\mathcal{F}_\tau$ is a filter, at least one of Y or $\omega \setminus Y$ is $\mathcal{F}_\tau$-positive; without loss of generality, say Y is. Apply Lemma 5.2 to obtain $\sigma' \in \mathrm{Fn}(\kappa \setminus \mathrm{dom}(\tau), \omega)$ and $F' \in \mathcal{F}_\tau$ satisfying $F' \cap a_{\sigma'} \subset^* Y$. Write $F' = F \cap a_\tau$ for some $F \in \mathcal{F}_0$, and set $\sigma = \tau \cup \sigma'$. Then σ extends τ and

$$F \cap a_\sigma = F \cap a_\tau \cap a_{\sigma'} = F' \cap a_{\sigma'} \subset^* Y,$$

so $\sigma \in D$. If instead $\omega \setminus Y$ is $\mathcal{F}_\tau$-positive, the same argument produces $\sigma \supseteq \tau$ with $F \cap a_\sigma \cap Y =^* \emptyset$, so again $\sigma \in D$. □

Proposition 6.6. *For every maximal antichain $\Sigma \subset \mathrm{Fn}(\kappa, \omega)$, the members of the family $\{a_\sigma : \sigma \in \Sigma\}$ are pairwise disjoint and the union is an element of $\mathcal{F}_0$.*

Proof. If $\sigma \neq \sigma'$ are in Σ, there is a coordinate ξ in their common domains such that $\sigma(\xi) \neq \sigma'(\xi)$. Therefore $a_\sigma, a_{\sigma'}$ are contained in $a(\xi, \sigma(\xi))$, respectively $a(\xi, \sigma'(\xi))$ which are disjoint.

Next let $Y = \bigcup\{a_\sigma : \sigma \in \Sigma\}$ and suppose that $Y \notin \mathcal{F}_0$. It follows from the maximality of $\mathcal{F}_0$ that there is some $\tau \in \mathrm{Fn}(\kappa, \omega)$ and some $F \in \mathcal{F}_0$ such that $F \cap a_\tau$ is disjoint from Y. Since Σ is maximal, there is a $\sigma \in \Sigma$ such that $\sigma \cup \tau$ is a function. In other words $a_\sigma \cap a_\tau = a_{\sigma \cup \tau}$ must meet F in an infinite set. Since $a_\sigma \subset Y$, this contradicts that $F \cap a_\tau$ is disjoint from Y. □

Definition 6.7. Let $\mathcal{Y}_\alpha$ be the family of all sets of the form $\bigcup\{a_\sigma : \sigma \in \Sigma\}$ where Σ is an antichain of $\mathrm{Fn}(\alpha, \omega)$.

With this machinery in place, we return to the main question of the section: can $\kappa_{\mathcal{U}}$ take different values as $\mathcal{U}$ varies? The following theorem answers yes, for every regular κ in the range $\omega_1 \leq \kappa \leq \mathfrak{c}$.

Theorem 6.8. *For any regular cardinal κ with $\omega_1 \leq \kappa \leq \mathfrak{c}$, there is an ultrafilter $\mathcal{U}_\kappa$ satisfying that $<_{\mathcal{U}_\kappa}$ has an (ω, κ)-gap.*

Proof. Let $\mathcal{A}$ be a $\kappa \times \omega$-independent family, and let $\mathcal{F}_0$ be a filter on ω that is maximal with respect to $\mathcal{A}$ being independent modulo $\mathcal{F}_0$. We will construct an ultrafilter $\mathcal{U}$ on ω such that $(\omega^\omega, <_{\mathcal{U}})$ has an (ω, κ)-gap. The lower side of the gap will consist of the constant functions $\bar{n}$ for $n \in \omega$. For the upper side, each coordinate $\alpha \in \kappa$ gives a partition $\{a(\alpha, n) : n \in \omega\}$ of ω; define $g_\alpha \in \omega^\omega$ by $g_\alpha(k) = n$ iff $k \in a(\alpha, n)$. For each $\alpha < \kappa$, let $\mathcal{A}_\alpha = \{a(\beta, m) : \alpha < \beta < \kappa,\, m \in \omega\}$ denote the subfamily omitting coordinates $\leq \alpha$. The goal is to build $\mathcal{U}$ so that $\bar{n} <_{\mathcal{U}} g_\alpha$ for all n and α, with nothing in between.

Define $\mathcal{H}_0 = \{h \in \omega^\omega : \bar{n} <_{\mathcal{F}_0} h \text{ for all } n \in \omega\}$. We want to split the gap formed by the constant functions and $\mathcal{H}_0$. Under $<_{\mathcal{F}_0}$, no function splits this gap; every function that is $\mathcal{F}_0$-above all constants is already in $\mathcal{H}_0$, much like the canonical gap of Theorem 2.8. But if we enlarge $\mathcal{F}_0$ to some $\mathcal{F}_1$, we may be able to split it, and g_0 is our chosen splitter. More generally, at stage α the function g_α will split the gap relative to $\mathcal{F}_\alpha$, and we continue until the filter becomes an ultrafilter.

For each $\alpha < \kappa$, $n \in \omega$, and $h \in \mathcal{H}_\alpha$, define $Y(\alpha, n, h) = \{k \in \omega : n < g_\alpha(k) < h(k)\}$. The point of $Y(\alpha, n, h)$ is that adding it to a filter forces $\bar{n} <_{\mathcal{F}} g_\alpha <_{\mathcal{F}} h$. Consider the family $\widetilde{\mathcal{F}}_0 = \{F \cap Y(0, n, h) : F \in \mathcal{F}_0,\, n \in \omega,\, h \in \mathcal{H}_0\}$. We claim that $\widetilde{\mathcal{F}}_0$ generates a filter. Indeed, a finite intersection of members of $\widetilde{\mathcal{F}}_0$ satisfies

$$(F_1 \cap Y(0, n_1, h_1)) \cap (F_2 \cap Y(0, n_2, h_2)) = (F_1 \cap F_2) \cap Y(0, n^*, h^*)$$

where $n^* = \max(n_1, n_2)$ and $h^* = \min(h_1, h_2)$ pointwise. Since $\mathcal{F}_0$ is closed under finite intersections, $F_1 \cap F_2 \in \mathcal{F}_0$, and since $\mathcal{H}_0$ is closed under finite minimum, $h^* \in \mathcal{H}_0$. Hence this is again a member of $\widetilde{\mathcal{F}}_0$.

Moreover, for every $\alpha < \kappa$, define $\mathcal{A}_\alpha = \{a(\beta, m) : \alpha < \beta < \kappa,\, m \in \omega\}$. The filter generated by $\widetilde{\mathcal{F}}_0$ is independent modulo $\mathcal{A}_0$. That is, we must show that $F \cap Y(0, n, h) \cap a_\tau$ is infinite for any $F \in \mathcal{F}_0$, $n \in \omega$, $h \in \mathcal{H}_0$, and $\tau \in \mathrm{Fn}(\kappa \setminus \{0\}, \omega)$. Fix such F, n, h, and τ. The set $Y(0, n, h)$ imposes two conditions on k: that $g_0(k) > n$ and that $g_0(k) < h(k)$. For the first, it suffices to restrict to a single cell $a(0, m)$ for any $m > n$, since every $k \in a(0, m)$ satisfies $g_0(k) = m > n$. The second condition then reduces to $h(k) > m$ on this cell, giving us

$$a(0, m) \cap \{k : h(k) > m\} \subseteq \{k : n < g_0(k) < h(k)\} = Y(0, n, h),$$

and therefore

$$F \cap Y(0, n, h) \cap a_\tau \supseteq F \cap \{k : h(k) > m\} \cap a_\tau \cap a(0, m).$$

Since $h \in \mathcal{H}_0$, we have $\{k : h(k) > m\} \in \mathcal{F}_0$, so $F \cap \{k : h(k) > m\} \in \mathcal{F}_0$. Since τ does not use 0 in its domain, $\tau \cup \{(0, m)\} \in \mathrm{Fn}(\kappa, \omega)$, and by independence of $\mathcal{A}$ modulo $\mathcal{F}_0$ the right-hand side is infinite.

Extend $\widetilde{\mathcal{F}}_0$ to a filter $\mathcal{F}_1$ such that $\mathcal{F}_1 \cap \mathcal{Y}_1$ is maximal. To do this, consider the poset of all filters extending $\widetilde{\mathcal{F}}_0$ that are generated by $\widetilde{\mathcal{F}}_0$ together with members of $\mathcal{Y}_1$, ordered by inclusion. Every chain has an upper bound (its union), so by Zorn's lemma there is a maximal element $\mathcal{F}_1$.

We claim that $\mathcal{A}_0$ is independent modulo $\mathcal{F}_1$. A typical member of $\mathcal{F}_1$ contains a set of the form $F \cap Y(0, n, h) \cap Y'$ where $F \in \mathcal{F}_0$, $h \in \mathcal{H}_0$, and $Y' = \bigcup_{j \in S} a(0, j) \in \mathcal{Y}_1$. Given $\tau \in \mathrm{Fn}(\kappa \setminus \{0\}, \omega)$, the argument proceeds as before: pick $m > n$ with $m \in S$, so that $a(0, m) \subseteq Y'$. Then $F \cap \{k : h(k) > m\} \cap a_\tau \cap a(0, m)$ is infinite by independence of $\mathcal{A}$ modulo $\mathcal{F}_0$, and every element of this set lies in $F \cap Y(0, n, h) \cap Y' \cap a_\tau$.

At a successor stage $\alpha + 1$, assume we have constructed a filter $\mathcal{F}_\alpha$ such that $\mathcal{A}_\alpha$ is independent modulo $\mathcal{F}_\alpha$ and $\mathcal{F}_\alpha \cap \mathcal{Y}_\alpha$ is maximal. Define $\mathcal{H}_\alpha = \{h \in \omega^\omega : \bar{n} <_{\mathcal{F}_\alpha} h \text{ for all } n \in \omega\}$ and set $Y(\alpha, n, h) = \{k \in \omega : n < g_\alpha(k) < h(k)\}$ for each $n \in \omega$ and $h \in \mathcal{H}_\alpha$. As in the base case, form the family

$$\widetilde{\mathcal{F}}_\alpha = \{F \cap Y(\alpha, n, h) : F \in \mathcal{F}_\alpha,\, n \in \omega,\, h \in \mathcal{H}_\alpha\}.$$

The same arguments show that $\widetilde{\mathcal{F}}_\alpha$ generates a filter and that $\mathcal{A}_{\alpha+1}$ is independent modulo it. We then extend $\widetilde{\mathcal{F}}_\alpha$ by Zorn's lemma to a filter $\mathcal{F}_{\alpha+1}$ such that $\mathcal{F}_{\alpha+1} \cap \mathcal{Y}_{\alpha+1}$ is maximal, and verify as before that $\mathcal{A}_{\alpha+1}$ is independent modulo $\mathcal{F}_{\alpha+1}$.

At a limit stage $\delta < \kappa$, let $\mathcal{F}_\delta = \bigcup_{\alpha < \delta} \mathcal{F}_\alpha$. This is a filter: since the sequence $(\mathcal{F}_\alpha)_{\alpha < \delta}$ is increasing, any $F, G \in \mathcal{F}_\delta$ lie in some common $\mathcal{F}_\gamma$, and hence $F \cap G \in \mathcal{F}_\gamma \subseteq \mathcal{F}_\delta$. Upward closure is immediate. We verify that $\mathcal{A}_\delta$ is independent modulo $\mathcal{F}_\delta$: let $\tau \in \mathrm{Fn}(\kappa \setminus \delta, \omega)$ and $F \in \mathcal{F}_\delta$. Then $F \in \mathcal{F}_\alpha$ for some $\alpha < \delta$. Since τ uses only coordinates $> \delta > \alpha$, independence of $\mathcal{A}_\alpha$ modulo $\mathcal{F}_\alpha$ implies that $F \cap a_\tau$ is infinite. As before, extend $\mathcal{F}_\delta$ by Zorn's Lemma so that $\mathcal{F}_\delta \cap \mathcal{Y}_\delta$ is maximal. Independence of $\mathcal{A}_\delta$ is preserved by the same argument as in the base case.

Finally, set $\mathcal{U} = \bigcup\{\mathcal{F}_\alpha : \alpha < \kappa\}$. We verify that $\mathcal{U}$ is an ultrafilter and that $(\{\bar{n} : n \in \omega\}, \{g_\alpha : \alpha \in \kappa\})$ is an (ω, κ)-gap in $<_\mathcal{U}$.

To show $\mathcal{U}$ is an ultrafilter, let $Y \subseteq \omega$. By Lemma 6.5, choose a maximal antichain $\Sigma_Y \subseteq \mathrm{Fn}(\kappa, \omega)$ that decides Y piece by piece: for each $\sigma \in \Sigma_Y$, there is $F_\sigma \in \mathcal{F}_0$ such that either $F_\sigma \cap a_\sigma \subset^* Y$ or $F_\sigma \cap a_\sigma \cap Y =^* \emptyset$. By Lemma 6.4, Σ_Y is countable, so the coordinates appearing in it are bounded below some $\alpha < \kappa$. By Proposition 6.6, the sets a_σ are pairwise disjoint and their union is in $\mathcal{F}_0$; therefore $\bigcup\{a_\sigma : \sigma \in \Sigma_Y\}$ is an element of $\mathcal{Y}_\alpha$. The point of the stage-α construction is that $\mathcal{F}_\alpha \cap \mathcal{Y}_\alpha$ is maximal, so this antichain decomposition forces a decision at stage α: one of Y or $\omega \setminus Y$ lies in $\mathcal{F}_\alpha \subseteq \mathcal{U}$. Hence $\mathcal{U}$ is an ultrafilter.

To prove the gap, for every $n \in \omega$ and $\alpha \in \kappa$, the set $Y(\alpha, n, h) \in \widetilde{\mathcal{F}}_\alpha \subseteq \mathcal{F}_{\alpha+1} \subseteq \mathcal{U}$ for any $h \in \mathcal{H}_\alpha$. Since $Y(\alpha, n, h) \subseteq \{k : g_\alpha(k) > n\}$, we have $\bar{n} <_\mathcal{U} g_\alpha$. Similarly, $Y(\alpha, n, h) \subseteq \{k : g_\alpha(k) < h(k)\}$, so $g_\alpha <_\mathcal{U} h$.

It remains to show that no function splits the gap. Suppose $f \in \omega^\omega$ satisfies $\bar{n} <_\mathcal{U} f$ for all $n \in \omega$. Then for each n, the set $\{k : f(k) > n\}$ belongs to $\mathcal{U}$, so it lies in some $\mathcal{F}_{\alpha_n}$. Since $\{\alpha_n : n \in \omega\}$ is countable and κ is regular and uncountable, there exists $\beta < \kappa$ such that $\alpha_n < \beta$ for all n. As the filters are increasing, it follows that $\{k : f(k) > n\} \in \mathcal{F}_\beta$ for every n, and hence $f \in \mathcal{H}_\beta$. By the construction at stage β, we have $g_\beta <_\mathcal{U} h$ for every $h \in \mathcal{H}_\beta$, and in particular $g_\beta <_\mathcal{U} f$. Therefore f cannot lie below every g_α, and so it cannot split the gap. $\square$