Combinatorial Cardinal Characteristics of the Continuum

Andreas Blass[1]
Mathematics Department
University of Michigan
Ann Arbor, MI 48109–1109, U.S.A.
ablass@umich.edu

November 24, 2003

[1]Partially supported by NSF grant DMS–9505118

1. Introduction

The first theorem about cardinal characteristics of the continuum is Cantor's classical result [37] that the cardinality $\mathfrak{c} = 2^{\aleph_0}$ of the continuum is strictly larger than the cardinality $\aleph_0$ of a countably infinite set. The distinction between $\aleph_0$ and $\mathfrak{c}$ was soon put to good use, especially in real analysis, where countable sets were shown to have many useful properties that cannot be extended to sets of cardinality $\mathfrak{c}$. Here are a few familiar examples; more examples are implicit throughout this chapter.

- Countably many nowhere dense sets cannot cover the real line. (The Baire Category Theorem)
- If countably many sets each have Lebesgue measure zero then so does their union.
- Given countably many sequences of real numbers, there is a single sequence that eventually dominates each of the given ones.
- Let countably many bounded sequences of real numbers $S_k = \langle x_{k,n} \rangle_{n \in \omega}$ be given. There is an infinite subset A of ω such that all the corresponding subsequences $S_k \restriction A = \langle x_{k,n} \rangle_{n \in A}$ converge.

Each of these results becomes trivially false if the hypothesis of countability is weakened to allow cardinality $\mathfrak{c}$. It is natural to ask whether the hypothesis can be weakened at all and, if so, by how much. For which uncountable cardinals, if any, do these results remain correct?

If the continuum hypothesis (CH) is assumed, the answer is trivial. The results are false already for $\aleph_1$ because $\aleph_1 = \mathfrak{c}$. But the continuum hypothesis, though not refutable from the usual (ZFC) axioms of set theory, is also not provable from them, so one can reasonably ask what happens if CH is false. Then there are cardinals strictly between $\aleph_0$ and $\mathfrak{c}$, and it is not evident whether the results cited above remain valid when "countable" is replaced by one of these cardinals.

Not only is it not evident, but it is not decidable in ZFC. For example, it is consistent with ZFC that $\mathfrak{c} = \aleph_2$ and all the cited results remain correct for $\aleph_1$, but it is also consistent that $\mathfrak{c} = \aleph_2$ and all the cited results fail for $\aleph_1$. It may seem that this undecidability prevents us from saying anything useful about extending the results above to higher cardinals. Fortunately, though little can be said about extending any one of these results, there are surprising and deep connections between extensions of different results. For example, if the Lebesgue measure result quoted above remains true for a cardinal κ, then so do the results about Baire category and about eventual domination.

A major goal of the theory of cardinal characteristics of the continuum is to understand relationships of this sort, either by proving implications

like the one just cited or by showing that other implications are unprovable in ZFC. The cardinal characteristics are simply the smallest cardinals for which various results, true for $\aleph_0$, become false. (The characteristics corresponding to the four results cited above are called $\mathbf{cov}(\mathcal{B})$, $\mathbf{add}(\mathcal{L})$, $\mathfrak{b}$, and $\mathfrak{s}$, respectively, so the implication at the end of the preceding paragraph would be expressed by the inequalities $\mathbf{add}(\mathcal{L}) \leq \mathbf{cov}(\mathcal{B})$ and $\mathbf{add}(\mathcal{L}) \leq \mathfrak{b}$.) We shall be concerned here only with results about $\aleph_0$ that are false for $\mathfrak{c}$, so the characteristics we consider lie in the interval from $\aleph_1$ to $\mathfrak{c}$, inclusive.

A second goal of the theory, which we touch on only briefly here, is to find situations, in set theory or other branches of mathematics, where cardinal characteristics arise naturally. Wherever a result involves a countability hypothesis, one can ask whether it extends to some uncountable cardinals. Quite often, one can extend it to all cardinals below some previously studied characteristic. (Of course, if the result fails for $\mathfrak{c}$, one can simply use it to define a new characteristic, but this is of little value unless one can relate it to more familiar characteristics or at least give a simple, combinatorial description of it.) Such applications are fairly common in set-theoretic topology — notice that the two standard survey articles on cardinal characteristics, [41] and [110], appeared in topology books. They are becoming more common in other branches of mathematics as these branches come up against set-theoretic independence results.

We digress for a moment to comment on the meaning of "continuum" in the name of our subject. In principle, "continuum" refers to the real line $\mathbb{R}$ or to an interval like $[0,1]$ in $\mathbb{R}$, regarded as a topological space. It is, however, common practice in set theory to apply the word also to spaces like ${}^{\omega}2$, ${}^{\omega}\omega$ and $[\omega]^{\omega}$. Here ${}^{\omega}X$ means the space of ω-sequences of elements of X, topologized as a product of discrete spaces. Thus, ${}^{\omega}2$ consists of sequences of zeros and ones; it may be identified with the power set $\mathcal{P}\omega$ of ω. $[\omega]^{\omega}$ is the subspace of $\mathcal{P}\omega$ consisting of the infinite sets. All these spaces are equivalent for many purposes, since any two become homeomorphic after removal of suitable countable subsets. We remark in particular that there is a continuous bijection from ${}^{\omega}\omega$ to $[0,1)$, whose inverse is continuous except at dyadic rationals. This bijection, which takes the sequence $(a_0, a_1, \dots) \in {}^{\omega}\omega$ to the number whose binary expansion is a_0 ones, a zero, a_1 ones, a zero, ..., also behaves nicely with respect to measure. Lebesgue measure on $[0,1)$ corresponds to the product measure on ${}^{\omega}\omega$ obtained from the measure on ω giving each point n the measure 2^{-n-1}. Similarly, the obvious "binary notation" map from ${}^{\omega}2$ onto $[0,1]$, which fails to be one-to-one only over the dyadic rationals, makes Lebesgue measure correspond to the product measure on ${}^{\omega}2$ obtained from the uniform measure on 2. In view of correspondences like these, we shall, without further explanation, apply cardinal characteristics like $\mathbf{cov}(\mathcal{B})$ and $\mathbf{add}(\mathcal{L})$ to all these versions of the continuum (with the corresponding measures), although they were defined in terms of $\mathbb{R}$ (with Lebesgue measure).

Another aspect of our subject's name also deserves a brief digression. Are these cardinals really characteristics of the continuum, or do they depend on more of the set-theoretic universe? Of course they depend on the class of cardinals; a characterstic that ceases to be a cardinal in some forcing extension will obviously cease to be a characteristic there also. So a more reasonable question would be whether the characteristics are determined by the continuum and the cardinals. More specifically, can cardinal characteristics of the continuum be changed by a forcing that neither adds reals nor collapses cardinals? Mildenberger [77] has shown that, for certain characteristics, such changes are possible but only in the presence of inner models with large cardinals.

We adopt the following standard notations for dealing with "modulo finite" notions on the natural numbers. First, $\forall^\infty x$ means "for all but finitely many x"; here x will always range over natural numbers, so the quantifier is equivalent to "for all sufficiently large x." Similarly $\exists^\infty x$ means "for infinitely many x" or equivalently "there exist arbitrarily large x such that." Notice that these quantifiers stand in the same duality relation as simple $\forall$ and $\exists$, namely $\neg\forall^\infty x$ is equivalent to $\exists^\infty x\neg$. An asterisk is often used to indicate a weakening from "for all" to "for all but finitely many." In particular, for subsets X and Y of ω, we write $X \subseteq^* Y$ to mean that X is almost included (or included modulo finite) in Y, i.e., $\forall^\infty x\,(x \in X \implies x \in Y)$. Similarly, for functions $f, g \in {}^\omega\omega$, we write $f \leq^* g$ to mean $\forall^\infty x\,(f(x) \leq g(x))$. We often use "almost" to mean modulo finite sets. For example, an almost decreasing sequence of sets is one where $X_m \supseteq^* X_n$ whenever $m < n$.

We use the standard abbreviations (some already mentioned above): ZFC for Zermelo-Fraenkel set theory including the axiom of choice, CH for the continuum hypothsis ($\mathfrak{c} = \aleph_1$), and GCH for the generalized continuum hypothesis ($2^{\aleph_\alpha} = \aleph_{\alpha+1}$ for all cardinals $\aleph_\alpha$).

2. Growth of Functions

The ordering $\leq^*$ on ${}^\omega\omega$ provides two simple but frequently useful cardinal characteristics, the dominating and (un)bounding numbers.

2.1 Definition A family $\mathcal{D} \subseteq {}^\omega\omega$ is *dominating* if for each $f \in {}^\omega\omega$ there is $g \in \mathcal{D}$ with $f \leq^* g$. The *dominating number* $\mathfrak{d}$ is the smallest cardinality of any dominating family, $\mathfrak{d} = \min\{|\mathcal{D}| : \mathcal{D} \text{ dominating}\}$.

2.2 Definition A family $\mathcal{B} \subseteq {}^\omega\omega$ is *unbounded* if there is no single $f \in {}^\omega\omega$ such that $g \leq^* f$ for all $g \in \mathcal{B}$. The *bounding number* $\mathfrak{b}$ (sometimes called the *unbounding number*) is the smallest cardinality of any unbounded family.

2.3 Remark Had we used the "everywhere" ordering ($f \leq g$ if $\forall x\,(f(x) \leq g(x))$) instead of the "almost everywhere" ordering, $\mathfrak{d}$ would be unchanged,

as any dominating $\mathcal{D}$ could be made dominating in the new sense by adding all finite modifications of its members. But $\mathfrak{b}$ would drop down to $\aleph_0$, as the constant functions form an unbounded family in the new sense.

Both $\mathfrak{b}$ and $\mathfrak{d}$ would be unchanged if in their definitions we replaced ${}^{\omega}\omega$ with ${}^{\omega}\mathbb{R}$ or with the set of sequences from any linear ordering of cofinality ω.

The following theorem gives all the constraints on $\mathfrak{b}$ and $\mathfrak{d}$ that are provable in ZFC.

2.4 Theorem $\aleph_1 \leq \mathit{cof}(\mathfrak{b}) = \mathfrak{b} \leq \mathit{cof}(\mathfrak{d}) \leq \mathfrak{d} \leq \mathfrak{c}$.

Proof. That $\aleph_1 \leq \mathfrak{b}$ means that, for every countably many functions $g_n : \omega \to \omega$, there is a single $f \geq^* $ all of them. Such an f is given by $f(x) = \max_{n \leq x} g_n(x)$.

To prove that $\mathfrak{b} \leq \mathrm{cof}(\mathfrak{d})$, let $\mathcal{D}$ be a dominating family of size $\mathfrak{d}$, and let it be decomposed into the union of $\mathrm{cof}(\mathfrak{d})$ subfamilies $\mathcal{D}_\xi$ of cardinalities $< \mathfrak{d}$. So there is, for each ξ, some f_ξ not dominated by any $g \in \mathcal{D}_\xi$. There can be no f dominating all the f_ξ, for such an f would not be dominated by any $g \in \mathcal{D}$. So $\{f_\xi : \xi < \mathrm{cof}(\mathfrak{d})\}$ is unbounded.

The proof that $\mathrm{cof}(\mathfrak{b}) = \mathfrak{b}$ is similar, and the rest of the theorem is obvious. ⊣

Hechler [56] has shown that, if P is a partially ordered set in which every countable subset has an upper bound, then P can consistently be isomorphic to a cofinal subset of $({}^{\omega}\omega, \leq^*)$. More precisely, given any such P, Hechler constructs a ccc forcing extension of the universe where there is a strictly order-preserving, cofinal embedding of P into $({}^{\omega}\omega, \leq^*)$. (Hechler's proof, done soon after the invention of forcing, has been reworked, using a more modern formulation, by Talayco in [108, Chapter 4] and by Burke in [35].) Hechler's result implies that the preceding theorem is optimal in the following sense.

2.5 Theorem *Assume GCH, and let* $\mathfrak{b}'$, $\mathfrak{d}'$, *and* $\mathfrak{c}'$ *be any three cardinals satisfying*

$$\aleph_1 \leq \mathit{cof}(\mathfrak{b}') = \mathfrak{b}' \leq \mathit{cof}(\mathfrak{d}') \leq \mathfrak{d}' \leq \mathfrak{c}'$$

and $\mathit{cof}(\mathfrak{c}') > \aleph_0$. *Then there is a ccc forcing extension of the universe satisfying* $\mathfrak{b} = \mathfrak{b}'$, $\mathfrak{d} = \mathfrak{d}'$, *and* $\mathfrak{c} = \mathfrak{c}'$.

Proof. Apply Hechler's theorem to $P = [\mathfrak{d}']^{<\mathfrak{b}'}$ partially ordered by inclusion. The regularity of $\mathfrak{b}'$ implies that any $< \mathfrak{b}'$ elements in P have an upper bound, but some $\mathfrak{b}'$ elements (e.g., distinct singletons) do not. From $\mathrm{cof}(\mathfrak{d}') \geq \mathfrak{b}'$ and GCH we get that $|P| = \mathfrak{d}'$. Fewer than $\mathfrak{d}'$ elements of P cannot be cofinal, for their union (as sets) has cardinality smaller than $\mathfrak{d}'$.

These observations imply that $\mathfrak{b} = \mathfrak{b}'$ and $\mathfrak{d} = \mathfrak{d}'$ in the forcing extension given by Hechler's theorem. Finally, to get $\mathfrak{c} = \mathfrak{c}'$, adjoin $\mathfrak{c}'$ random reals; these will not damage $\mathfrak{b}$ or $\mathfrak{d}$, as the ground model's ${}^{\omega}\omega$ is cofinal in the ${}^{\omega}\omega$ of any random real extension. $\dashv$

To see that $\mathfrak{b} < \mathfrak{d}$ is consistent, it is not necessary to invoke Hechler's theorem. The original Cohen models [39] for the negation of CH have $\mathfrak{b} = \aleph_1$ and $\mathfrak{d} = \mathfrak{c}$. In fact, if one adjoins $\kappa \geq \aleph_1$ Cohen reals (by the usual product forcing) to any model of set theory, then the resulting model has $\mathfrak{b} = \aleph_1$ while $\mathfrak{d}$ becomes at least κ.

The contrary situation, that $\mathfrak{b} = \mathfrak{d}$, has the following useful characterization.

2.6 Theorem $\mathfrak{b} = \mathfrak{d}$ *if and only if there is a scale in* ${}^{\omega}\omega$*, i.e., a dominating family well-ordered by* $\leq^*$*.*

Proof. If $\mathcal{D} = \{f_\xi : \xi < \mathfrak{b}\}$ is a dominating family of size $\mathfrak{b}$, then we obtain a scale $\{g_\xi : \xi < \mathfrak{b}\}$ by choosing each g_ξ to dominate f_ξ and all previous g_η $(\eta < \xi)$; this can be done because we need to dominate fewer than $\mathfrak{b}$ functions at a time.

Conversely, if there is a scale, choose one and let $\mathcal{B}$ be an unbounded family of size $\mathfrak{b}$. By increasing each element of $\mathcal{B}$ if necessary, we can arrange for $\mathcal{B}$ to be a subset of our scale. But then, being unbounded, it must be cofinal in the well-ordering $\leq^*$ of the scale. Therefore it is a dominating family. $\dashv$

There are several alternative ways of looking at $\mathfrak{b}$ and $\mathfrak{d}$. We present two of them here and refer to [41], [55] and [57] for others.

The first of these involves the "standard" characteristics of an ideal, defined as follows.

2.7 Definition Let $\mathcal{I}$ be a proper ideal of subsets of a set X, containing all singletons from X.

- The *additivity* of $\mathcal{I}$, $\mathbf{add}(\mathcal{I})$, is the smallest number of sets in $\mathcal{I}$ with union not in $\mathcal{I}$.
- The *covering number* of $\mathcal{I}$, $\mathbf{cov}(\mathcal{I})$, is the smallest number of sets in $\mathcal{I}$ with union X.
- The *uniformity* of $\mathcal{I}$, $\mathbf{non}(\mathcal{I})$, is the smallest cardinality of any subset of X not in $\mathcal{I}$.
- The *cofinality* of $\mathcal{I}$, $\mathbf{cof}(\mathcal{I})$ is the smallest cardinality of any subset $\mathcal{B}$ of $\mathcal{I}$ such that every element of $\mathcal{I}$ is a subset of an element of $\mathcal{B}$. Such a $\mathcal{B}$ is called a *basis* for $\mathcal{I}$.

It is easy to check that both $\mathbf{cov}(\mathcal{I})$ and $\mathbf{non}(\mathcal{I})$ are $\geq \mathbf{add}(\mathcal{I})$ and $\leq \mathbf{cof}(\mathcal{I})$. In fact, $\mathbf{add}(\mathcal{I})$ is a lower bound for the cofinalities $\mathrm{cof}(\mathbf{non}(\mathcal{I}))$ and $\mathrm{cof}(\mathbf{cof}(\mathcal{I}))$ also. In this chapter, $\mathcal{I}$ will always be a σ-ideal, so its additivity (and therefore the other three characteristics) will be uncountable. Furthermore, $\mathcal{I}$ will have a basis consisting of Borel sets; since there are only $\mathfrak{c}$ Borel sets, the cofinality (and therefore the other three characteristics) will be $\leq \mathfrak{c}$. (That the other three characteristics are $\leq \mathfrak{c}$ follows already from the simpler fact that the underlying set X is the continuum.)

The ideal relevant to the present section is the σ-ideal $\mathcal{K}_\sigma$ generated by the compact subsets of ${}^\omega\omega$, i.e., the ideal of sets coverable by countably many compact sets. Its connection with $\leq^*$ was pointed out by Rothberger in [90].

2.8 Theorem $\mathbf{add}(\mathcal{K}_\sigma) = \mathbf{non}(\mathcal{K}_\sigma) = \mathfrak{b}$ *and* $\mathbf{cov}(\mathcal{K}_\sigma) = \mathbf{cof}(\mathcal{K}_\sigma) = \mathfrak{d}$.

Proof. Since a subset of the discrete space ω is compact if and only if it is finite, the Tychonoff theorem implies that a subset of ${}^\omega\omega$ is compact if and only if it is closed and included in a product of finite subsets of ω. There is no loss of generality in taking the finite subsets to be initial segments, so we find that all sets of the form

$$\{f \in {}^\omega\omega : f \leq g\} = \prod_{n\in\omega} [0, g(n)]$$

are compact and every compact set is included in one of this form. It follows that all sets of the form $\{f \in {}^\omega\omega : f \leq^* g\}$ (with $\leq^*$ instead of $\leq$) are in $\mathcal{K}_\sigma$ and every set in $\mathcal{K}_\sigma$ is a subset of one of these. (The last uses that $\mathfrak{b} \geq \aleph_1$ to show that countably many bounds g for countably many compact sets are all $\leq^*$ a single bound.)

This connection between $\mathcal{K}_\sigma$ and $\leq^*$ easily implies the theorem. ⊣

Recalling that ${}^\omega\omega$ is homeomorphic, via continued fraction expansions, to the space of irrational numbers $\mathbb{R} - \mathbb{Q}$ (topologized as a subspace of $\mathbb{R}$), we see that the theorem remains valid if we interpret $\mathcal{K}_\sigma$ as the σ-ideal generated by the compact subsets of $\mathbb{R}-\mathbb{Q}$. In particular, $\mathfrak{d}$ is characterized as the minimum number of compact sets whose union is $\mathbb{R} - \mathbb{Q}$. (Here the choice of "continuum" is important. The corresponding cardinals for the spaces ${}^\omega 2$, $[0,1]$, and $\mathbb{R}$ are clearly 1, 1, and $\aleph_0$, respectively.)

Yet another way of looking at the ordering $\leq^*$ and the associated cardinals $\mathfrak{b}$ and $\mathfrak{d}$ involves partitions of ω into finite intervals. (The earliest reference I know for this idea is Solomon's [102].)

2.9 Definition An *interval partition* is a partition of ω into (infinitely many) finite intervals I_n $(n \in \omega)$. We always assume that the intervals are numbered in the natural order, so that, if i_n is the left endpoint of I_n then

$i_0 = 0$ and $I_n = [i_n, i_{n+1})$. We say that the interval partition $\{I_n : n \in \omega\}$ *dominates* another interval partition $\{J_n : n \in \omega\}$ if $\forall^\infty n \, \exists k \, (J_k \subseteq I_n)$. We write IP for the set of all interval partitions.

2.10 Theorem $\mathfrak{d}$ *is the smallest cardinality of any family of interval partitions dominating all interval partitions.* $\mathfrak{b}$ *is the smallest cardinality of any family of interval partitions not all dominated by a single interval partition.*

Proof. We prove only the first statement, as the second can be proved similarly or deduced from the proof of the first using the duality machinery of Section 4.

Suppose first that we have a family $\mathcal{F}$ of interval partitions dominating all interval partitions. To each of the partitions $\{I_n = [i_n, i_{n+1}) : n \in \omega\}$ in $\mathcal{F}$, associate the function $f : \omega \to \omega$ defined by letting $f(x)$ be the right endpoint of the interval after the one containing x; thus if $x \in I_n$ then $f(x) = i_{n+2} - 1$. We shall show that these functions f form a dominating family, so $\mathfrak{d} \leq |\mathcal{F}|$. Given any $g \in {}^\omega\omega$, the required f dominating g is obtained as follows. Form an interval partition $\{J_n = [j_n, j_{n+1}) : n \in \omega\}$ such that whenever $x \leq j_n$ then $g(x) < j_{n+1}$; it is trivial to do this by choosing the j_n inductively. Let $\{I_n = [i_n, i_{n+1}) : n \in \omega\}$ in $\mathcal{F}$ dominate this $\{J_n : n \in \omega\}$, and let f be the function associated to $\{I_n : n \in \omega\}$. To see that $g(x) \leq f(x)$ for all sufficiently large x, we chase through the definitions as follows. Let n be the index such that $x \in I_n$ and let (since x is sufficiently large) k be an index such that $J_k \subseteq I_{n+1}$. Then as $x \leq j_k$, we have $g(x) \leq j_{k+1} - 1 \leq i_{n+2} - 1 = f(x)$. This completes the proof that $\mathfrak{d} \leq |\mathcal{F}|$.

To produce a dominating family of interval partitions of cardinality $\mathfrak{d}$, we begin with a dominating family $\mathcal{D}$ of cardinality $\mathfrak{d}$ in ${}^\omega\omega$, and we associate to each $g \in \mathcal{D}$ an interval partition $\{J_n = [j_n, j_{n+1}) : n \in \omega\}$ exactly as in the preceding paragraph. To show that the resulting family of $\mathfrak{d}$ interval partitions dominates all interval partitions, let an arbitrary interval partition $\{I_n = [i_n, i_{n+1}) : n \in \omega\}$ be given, associate to it an $f \in {}^\omega\omega$ as in the preceding paragraph, and let $g \in \mathcal{D}$ be $\geq^* f$. We shall show that the $\{J_n : n \in \omega\}$ associated to this g dominates $\{I_n : n \in \omega\}$. For any sufficiently large n, we have $f(j_n) \leq g(j_n) \leq j_{n+1} - 1$. By virtue of the definition of f, this means that the next I_k after the one containing j_n lies entirely in J_n. $\dashv$

3. Splitting and Homogeneity

In this section, we treat several characteristics related to the "competition" between partitions trying to split sets and sets trying to be homogeneous

for partitions. We begin with a combinatorial definition of a characteristic already mentioned, from an analytic point of view, in the introduction.

3.1 Definition A set $X \subseteq \omega$ *splits* an infinite set $Y \subseteq \omega$ if both $Y \cap X$ and $Y - X$ are infinite. A *splitting family* is a family $\mathcal{S}$ of subsets of ω such that each infinite $Y \subseteq \omega$ is split by at least one $X \in \mathcal{S}$. The *splitting number* $\mathfrak{s}$ is the smallest cardinality of any splitting family.

Having defined $\mathfrak{s}$ differently in the introduction, we hasten to point out that the definitions are equivalent.

3.2 Theorem $\mathfrak{s}$ *is the minimum cardinality of any family of bounded* ω*-sequences* $S_\xi = \langle x_{\xi,n} \rangle_{n \in \omega}$ *of real numbers such that for no infinite* $Y \subseteq \omega$ *do all the corresponding subsequences* $S_\xi \restriction Y = \langle x_{\xi,n} \rangle_{n \in Y}$ *converge. The same is true if we consider only sequences consisting of just zeros and ones.*

Proof. The second assertion, where all S_ξ are in ${}^\omega 2$, is a trivial rephrasing of the definition of $\mathfrak{s}$; just regard the sequences S_ξ as the characteristic functions of the sets in a splitting family. The key point is that, for the characteristic function of X, convergence means eventual constancy, and so convergence of its restriction to Y means that Y is not split by X.

Half of the first assertion follows immediately from the second. To prove the other half of the first assertion, use the fact that a bounded sequence of real numbers converges if (though not quite only if) for each k the sequence of k^{th} binary digits converges. ⊣

The last part of the preceding proof implicitly used the fact that $\mathfrak{s}$ is uncountable. We omit the easy, direct proof of this, because it will also follow from results to be proved later ($\aleph_1 \leq \mathfrak{t} \leq \mathfrak{h} \leq \mathfrak{s}$; see Section 6).

Theorem 2.10 makes it easy to relate $\mathfrak{s}$ to $\mathfrak{d}$.

3.3 Theorem $\mathfrak{s} \leq \mathfrak{d}$.

Proof. By Theorem 2.10, fix a family of $\mathfrak{d}$ interval partitions dominating all interval partitions. To each partition $\Pi = \{I_n : n \in \omega\}$ in this family, associate the union $\varphi(\Pi) = \bigcup_n I_{2n}$ of its even-numbered intervals. We shall show that these $\mathfrak{d}$ sets $\varphi(\Pi)$ constitute a splitting family. To this end, consider an arbitrary infinite subset X of ω. Associate to it an interval partition $\psi(X)$ in which every interval contains at least one member of X. Our dominating family of interval partitions contains a Π that dominates $\psi(X)$. But then each interval of Π, except for finitely many, includes an interval of $\psi(X)$ and therefore contains a point of X. It follows immediately that both $\varphi(\Pi)$ and its complement (the union of the odd-numbered intervals) contain infinitely many points of X. So $\varphi(\Pi)$ splits X. ⊣

We record for future reference the basic property of the constructions φ and ψ that makes the preceding proof work: For any interval partition Π

and any infinite $X \subseteq \omega$,

$$\Pi \text{ dominates } \psi(X) \quad \Longrightarrow \quad \varphi(\Pi) \text{ splits } X.$$

The inequality in the theorem can consistently be strict. For example, if one adds $\kappa > \aleph_0$ Cohen reals to a model of set theory, then in the resulting model $\mathfrak{d} \geq \kappa$ (as remarked earlier) while $\mathfrak{s} = \aleph_1$ because any $\aleph_1$ of the added Cohen reals constitute a splitting family.

The splitting number is the simplest of a family of characteristics defined in terms of structures that are not simultaneously homogeneous (modulo finite) on any one infinite set. For $\mathfrak{s}$, the "structures" are two-valued functions and "homogeneous" simply means constant. Other notions of structure and homogeneity are suggested by various partition theorems. We shall characterize the analog of $\mathfrak{s}$ arising from Ramsey's theorem and briefly mention a few other analogs afterward.

3.4 Definition A set $H \subseteq \omega$ is *homogeneous* for a function $f : [\omega]^n \to k$ (a partition of $[\omega]^n$ into k pieces) if f is constant on $[H]^n$. H is *almost homogeneous* for f if there is a finite set F such that $H - F$ is homogeneous for f. $\mathfrak{par}_n$ is the smallest cardinality of any family of partitions of $[\omega]^n$ into two pieces such that no single infinite set is almost homogeneous for all of them simultaneously.

We note that $\mathfrak{par}_1$ is simply $\mathfrak{s}$ and that the definition of $\mathfrak{par}_n$ would be unchanged if we allowed partitions into any finite number of pieces (for any such partition could be replaced with the finitely many coarser partitions into two pieces). We note also that the use of *almost* homogeneity in the definition is essential; it is easy to produce countably many partitions with no common infinite homogeneous set.

3.5 Theorem *For all integers $n \geq 2$, $\mathfrak{par}_n = \min\{\mathfrak{b}, \mathfrak{s}\}$.*

Proof. Notice first that $\mathfrak{par}_n \leq \mathfrak{par}_m$ if $n \geq m$, because any partition $[\omega]^m \to 2$ can be regarded as a partition of $[\omega]^n$ ignoring the last $n - m$ elements of its input. In particular, we have $\mathfrak{par}_n \leq \mathfrak{s}$, and if we show $\mathfrak{par}_2 \leq \mathfrak{b}$ then the $\leq$ direction of the theorem will be proved. For the $\geq$ direction, we must consider arbitrary n, but in fact we shall confine attention to $n = 2$ since the general case is longer but not harder.

To show $\mathfrak{par}_2 \leq \mathfrak{b}$, let $\mathcal{B} \subseteq {}^\omega\omega$ be an unbounded family of size $\mathfrak{b}$, assume without loss of generality that each $g \in \mathcal{B}$ is monotone increasing, and associate to each such g the partition of $[\omega]^2$ that puts a pair $\{x < y\}$ into class 0 if $g(x) < y$ and into class 1 otherwise. We shall show that no infinite $H \subseteq \omega$ is almost homogeneous for all these partitions simultaneously. Notice first that a homogeneous set of class 1 must be finite since, if x is its first element, then all the other elements are majorized by $g(x)$. So suppose,

toward a contradiction, that H is infinite and almost homogeneous of class 0 for all the partitions associated to the functions $g \in \mathcal{B}$. Consider the function h sending each natural number x to the second member of H above x. For each x, we have $x < y < h(x)$ with both y and $h(x)$ in H. By almost homogeneity of H, we have, for each $g \in \mathcal{B}$ and for all sufficiently large x, $g(y) < h(x)$ and thus, by monotonicity of g, $g(x) < h(x)$. Thus, $g \leq^* h$ for all $g \in \mathcal{B}$, contrary to our choice of $\mathcal{B}$.

To show $\mathfrak{par}_2 \geq \min\{\mathfrak{b}, \mathfrak{s}\}$, suppose we are given a family of $\kappa < \min\{\mathfrak{b}, \mathfrak{s}\}$ partitions $f_\xi : [\omega]^2 \to 2$; we must find an infinite set almost homogeneous for all of them. First, consider the functions

$$f_{\xi,n} : \omega \to 2 : x \mapsto f_\xi\{n, x\}.$$

(This is undefined for $x = n$; define it arbitrarily there.) Since the number of these functions is $\kappa \cdot \aleph_0 < \mathfrak{s}$, there is an infinite $A \subseteq \omega$ on which they are almost constant; say $f_{\xi,n}(x) = j_\xi(n)$ for all $x \geq g_\xi(n)$ in A. Furthermore, since $\kappa < \mathfrak{s}$ we can find an infinite $B \subseteq A$ on which each j_ξ is almost constant, say $j_\xi(n) = i_\xi$ for all $n \geq b_\xi$ in B. And since $\kappa < \mathfrak{b}$ we have a function h majorizing each g_ξ from some integer c_ξ on. Let $H = \{x_0 < x_1 < \dots\}$ be an infinite subset of B chosen so that $h(x_n) < x_{n+1}$ for all n. Then this H is almost homogeneous for each f_ξ. Indeed, if $x < y$ are elements of H larger than b_ξ and c_ξ, then $y > h(x) \geq g_\xi(x)$ and so $f_\xi(\{x, y\}) = f_{\xi,x}(y) = j_\xi(x) = i_\xi$. ⊣

One can define characteristics analogous to $\mathfrak{par}_n$ using stronger partition theorems in place of Ramsey's theorem, for example Hindman's finite sums theorem [58] or the Galvin-Prikry theorem [48] and its extension to analytic sets by Silver [101]. It is not difficult to see that these characteristics are bounded above by $\min\{\mathfrak{b}, \mathfrak{s}\}$. The Silver and (a fortiori) the Galvin-Prikry variants of $\mathfrak{par}$ are easily seen to be bounded below by the characteristic $\mathfrak{h}$ defined in Section 6. Eisworth has also obtained (private communication) a lower bound of the form $\min\{\mathfrak{b}, \mathfrak{s}'\}$, where $\mathfrak{s}'$ is the following variant of $\mathfrak{s}$. A cardinal κ is $< \mathfrak{s}'$ if, for any κ reals, there exist

1. a transitive model N of enough of ZFC containing the given reals,
2. $\mathcal{U} \in N$ such that N satisfies "$\mathcal{U}$ is a non-principal ultrafilter on ω," and
3. an infinite $a \subseteq \omega$ almost included in every member of $\mathcal{U}$.

Eisworth's proof uses forcing techniques from [60], but a direct combinatorial proof can be based on [17, Theorem 4]. Note that, if we weakened requirement (2) in the definition of $\mathfrak{s}'$ to say only that $\mathcal{U}$ is a non-principal ultrafilter in the Boolean algebra of subsets of ω in N (but $\mathcal{U}$ need not be in

N), then the cardinal defined would be simply $\mathfrak{s}$. It is not known whether $\mathfrak{s}' < \mathfrak{s}$ is consistent.

For the variant of $\mathfrak{par}$ based on Hindman's theorem, the best lower bound known to me is the characteristic $\mathfrak{p}$ defined in Section 6. The proof that this is a lower bound uses the construction from Martin's axiom mentioned in [16, page 93], the observation that Martin's axiom is applied here to a σ-centered poset, and Bell's theorem (Theorem 7.12 below).

One can also consider weaker sorts of homogeneity. For example, define $\mathfrak{par}_{1,c}$ to be the smallest cardinality of a family $\mathcal{F}$ of functions $f : \omega \to \omega$ such that there is no single infinite set $A \subseteq \omega$ on which all the functions from $\mathcal{F}$ are almost one-to-one or almost constant, where "almost" means, as usual, except at finitely many points in A. (The subscript $1, c$ refers to the canonical partition theorem for sets of size 1.) Each function f gives rise to a partition $f' : [\omega]^2 \to 2$, where $f'(\{x, y\}) = 0$ just when $f(x) = f(y)$. The sets where f is one-to-one or constant are the homogeneous sets of f', so $\mathfrak{par}_{1,c} \geq \mathfrak{par}_2$. In fact equality holds here, because $\mathfrak{par}_{1,c}$ is $\leq$ both $\mathfrak{s}$ and $\mathfrak{b}$. To see the former, associate to each set X from a splitting family its characteristic function. To see the latter, fix a family of $\mathfrak{b}$ interval partitions not dominated by any single interval partition (by Theorem 2.10) and associate to each of these partitions a function f constant on exactly the intervals of the partition. Since such an f is not constant on any infinite set, it suffices to show that there is no infinite A on which each f is almost one-to-one. But if there were such an A, then we could build an interval partition in which each interval contains at least three elements of A, and this partition would dominate all the partitions in our chosen, allegedly undominated family.

We now shift our focus from counting partitions to counting candidates for homogeneous sets.

3.6 Definition A family $\mathcal{R}$ of infinite subsets of ω is *unsplittable* if no single set splits all members of $\mathcal{R}$. It is *σ-unsplittable* if no countably many sets suffice to split all members of $\mathcal{R}$. The *unsplitting number* $\mathfrak{r}$, also called the *refining* or *reaping number*, is the smallest cardinality of any unsplittable family. The *σ-unsplitting number* $\mathfrak{r}_\sigma$ is the smallest cardinality of any σ-unsplittable family.

Obviously, $\mathfrak{r} \leq \mathfrak{r}_\sigma$. It is not known whether strict inequality here is consistent with ZFC.

We omit the proof of the following theorem since it involves nothing beyond what went into the proof of Theorem 3.2.

3.7 Theorem *$\mathfrak{r}_\sigma$ is the minimum cardinality of any family of infinite sets $Y \subseteq \omega$ such that, for each bounded sequence $\langle x_n \rangle_{n \in \omega}$ of real numbers, the restriction $\langle x_n \rangle_{n \in Y}$ to some Y in the family converges. If we consider only*

sequences of zeros and ones, then the corresponding minimum cardinality is $\mathfrak{r}$.

We emphasize that, although in Theorem 3.2 the cardinal was the same for real-valued sequences as for two-valued sequences, the analogous equality in the present theorem is an open problem.

3.8 Theorem $\mathfrak{b} \leq \mathfrak{r}$.

Proof. As in the proof of Theorem 3.3, let φ be the operation sending any interval partition to the union of its even-numbered intervals, and let ψ be an operation sending any infinite subset X of ω to an interval partition in which every interval contains at least one member of X. Let $\mathcal{R}$ be an unsplittable family of $\mathfrak{r}$ infinite subsets of ω; thanks to Theorem 2.10, we can complete the proof by showing that no interval partition Π dominates all the partitions $\psi(X)$ for $X \in \mathcal{R}$. But, as we showed in the proof of Theorem 3.3 and recorded for reference immediately thereafter, if Π dominated all these $\psi(X)$, then $\varphi(\Pi)$ would split every $X \in \mathcal{R}$, contrary to the choice of $\mathcal{R}$. $\dashv$

We next introduce the homogeneity cardinals associated to Ramsey's theorem and the "one-to-one or constant" theorem. As in the discussion of partition counting, we could define homogeneity cardinals from Hindman's theorem, the Galvin-Prikry theorem, etc., but (as there) not much could be said about them.

3.9 Definition $\mathfrak{hom}_n$ is the smallest size of any family $\mathcal{H}$ of infinite subsets of ω such that every partition of $[\omega]^n$ into two pieces has an almost homogeneous set in $\mathcal{H}$. $\mathfrak{hom}_{1,c}$ is the smallest size of any family $\mathcal{H}$ of infinite subsets of ω such that every function $f : \omega \to \omega$ is almost one-to-one or almost constant on some set in $\mathcal{H}$.

This definition would be unchanged if we deleted "almost," for we could put into $\mathcal{H}$ all finite modifications of its members. Notice that $\mathfrak{hom}_1 = \mathfrak{r}$ and that $\mathfrak{hom}_n \geq \mathfrak{hom}_m$ if $n \geq m$ (the reverse of the corresponding inequality for $\mathfrak{par}$).

3.10 Theorem *For all integers* $n \geq 2$, $\mathfrak{hom}_n = \max\{\mathfrak{d}, \mathfrak{r}_\sigma\}$. *In addition,* $\max\{\mathfrak{d}, \mathfrak{r}\} \leq \mathfrak{hom}_{1,c} \leq \max\{\mathfrak{d}, \mathfrak{r}_\sigma\}$.

Proof. Although this proof contains only one idea not already in the proof of Theorem 3.5 and the subsequent discussion of $\mathfrak{par}_{1,c}$, we repeat some of the earlier ideas to clarify why we now have $\mathfrak{r}$ in one assertion and $\mathfrak{r}_\sigma$ elsewhere.

To show that $\max\{\mathfrak{d}, \mathfrak{r}\} \leq \mathfrak{hom}_{1,c}$, we assume that $\mathcal{H}$ is as in the definition of $\mathfrak{hom}_{1,c}$, and we show that its cardinality is $\geq$ both $\mathfrak{r}$ and $\mathfrak{d}$. For the former, we find that $\mathcal{H}$ is unsplittable because if X splits H then the characteristic function of X is neither almost one-to-one nor almost constant

on H. For the comparison with $\mathfrak{d}$, associate to each $H \in \mathcal{H}$ an interval partition Π_H such that each of its intervals contains at least three members of H. By Theorem 2.10, we need only check that every interval partition Θ is dominated by such a Π_H. Given Θ, let f be constant on exactly its intervals, and find $H \in \mathcal{H}$ on which f is almost one-to-one (as f is not constant on any infinite set). But then any interval of Π_H (except for finitely many) contains three points from H, all from different intervals of Θ, so it must contain a whole interval of Θ. So we have the required domination.

Next, we show that $\mathfrak{hom}_2 \leq \max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$ by constructing an $\mathcal{H}$ of size $\max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$ with the homogeneity property required in the definition of $\mathfrak{hom}_2$. (Note the similarity of this construction with the argument proving $\mathfrak{par}_2 \geq \min\{\mathfrak{b}, \mathfrak{s}\}$.) Let $\mathcal{D} \subseteq {}^\omega\omega$ be a dominating family of size $\mathfrak{d}$. Let $\mathcal{R}$ be a σ-unsplittable family of size $\mathfrak{r}_\sigma$. For each $A \in \mathcal{R}$, let $\mathcal{R}_A$ be an unsplittable family of $\mathfrak{r}$ subsets of A. For each $h \in \mathcal{D}$, each $A \in \mathcal{R}$, and each $B \in \mathcal{R}_A$, let $H = H(h, A, B)$ be an infinite subset of B such that, for any $x < y$ in H, $h(x) < y$. The family $\mathcal{H}$ of all these sets $H(h, A, B)$ has size at most $\max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$, and we shall now show that it contains an almost homogeneous set for every partition $f : [\omega]^2 \to 2$. Given f, define (as in the proof of Theorem 3.5) $f_n : \omega \to 2 : x \mapsto f\{n, x\}$. As $\mathcal{R}$ is σ-unsplittable, it contains an A on which each f_n is almost constant, say $f_n(x) = j(n)$ for all $x \geq g(n)$ in A. The function $j : A \to 2$ is almost constant on some B in the unsplittable family $\mathcal{R}_A$, say $j(n) = i$ for all $n \geq b$ in B. And $\mathcal{D}$ contains an h dominating g, say $h(x) \geq g(x)$ for all $x \geq c$. It is now routine to check (as in the proof of Theorem 3.5) that f is constant with value i on all pairs of elements larger than b and c in $H(h, A, B)$.

The proof that $\mathfrak{hom}_n \leq \max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$ for $n > 2$ is similar to the preceding but uses n rather than two nestings of σ-unsplittable families (with no σ needed for the last one). We omit the details.

The preceding arguments, along with the observation that "one-to-one or constant" is a special case of homogeneity for partitions of pairs, establish that

$$\max\{\mathfrak{r}, \mathfrak{d}\} \leq \mathfrak{hom}_{1,c} \leq \mathfrak{hom}_2 \leq \mathfrak{hom}_3 \leq \cdots \leq \max\{\mathfrak{r}_\sigma, \mathfrak{d}\}.$$

All that remains to be proved is that $\mathfrak{r}_\sigma \leq \mathfrak{hom}_2$, and this requires a method not involved in Theorem 3.5. The following argument is due to Brendle [31]. (Shelah had previously established the corresponding result for $\mathfrak{hom}_3$.)

Let $\mathcal{H}$ be as in the definition of $\mathfrak{hom}_2$, and let countably many functions $f_n : \omega \to 2$ be given. We seek a set in $\mathcal{H}$ on which each f_n is almost constant. Define, for each $x \in \omega$, the sequence of zeros and ones $\widehat{x} = \langle f_n(x) \rangle_{n \in \omega}$, so $\widehat{x}_n = f_n(x)$. Then define a partition of $[\omega]^2$ by putting $\{x < y\}$ into class 0 if $\widehat{x}$ lexicographically precedes $\widehat{y}$ and into class 1 otherwise. Let $H \in \mathcal{H}$ be almost homogeneous for this partition, let H' be a homogeneous set obtained by removing finitely many elements from H, and from now on let x and y range only over elements of H'. Suppose H' is homogeneous for

class 0. (The case of class 1 is analogous.) Then as x increases, $\widehat{x}_0$ can only increase. That is, if the value of $f_0(x)$ ever changes, then it changes from 0 to 1 and remains constant forever after. Once $\widehat{x}_0$ has stabilized, $\widehat{x}_1$ can only increase and must therefore stabilize. Continuing in this way, we see that, as x increases through values in H', each $\widehat{x}_n$ eventually stabilizes. This means that each $f_n(x)$ is almost constant on H' and therefore on H, as required. ⊣

3.11 Remark The last paragraph of this proof is similar to the proof that cardinals κ satisfying the partition relation $\kappa \longrightarrow (\kappa)^2_2$ are strong limit cardinals. The nature of the stabilization, where each component moves at most once after all its predecessors have stabilized, is also reminiscent of the proof that all requirements are eventually satisfied in a finite-injury priority argument.

4. Galois-Tukey Connections and Duality

We interrupt the description and discussion of particular cardinal characteristics in order to set up some machinery that is useful for describing many (though not all) of the characteristics and the relationships between them. This machinery was isolated by Vojtáš [111] under the name of "generalized Galois-Tukey connections"; the basic ideas had been used, but neither isolated nor named, in earlier work of Fremlin [46] and Miller (unpublished). The definitions of many cardinal characteristics have the form "the smallest cardinality of any set Y (of objects of a specified sort) such that every object x (of a possibly different sort) is related to some $y \in Y$ in a specified way." And many proofs of inequalities between such cardinals involve the construction of maps between the various sorts of objects involved in the definitions. This is formalized as follows.

4.1 Definition A triple $\mathbf{A} = (A_-, A_+, A)$ consisting of two sets $A_\pm$ and a binary relation $A \subseteq A_- \times A_+$ will be called simply a *relation*. In connection with such a relation, we call A_- the set of *challenges* and A_+ the set of *responses*; we read xAy (meaning $(x, y) \in A$) as "response y *meets* challenge x."

4.2 Definition The *norm* $\|\mathbf{A}\|$ of a relation $\mathbf{A} = (A_-, A_+, A)$ is the smallest cardinality of any subset Y of A_+ such that every $x \in A_-$ is related by A to at least one $y \in Y$. That is, it is the minimum number of responses needed to meet all challenges.

The definitions of cardinal characteristics in the preceding sections (as well as many others) amount to norms of relations. Furthermore, characteristics tend to come in pairs whose relations are dual to each other in the following sense.

4.3 Definition If $\mathbf{A} = (A_-, A_+, A)$ then the *dual* of $\mathbf{A}$ is the relation $\mathbf{A}^\perp = (A_+, A_-, \neg\breve{A})$ where $\neg$ means complement and $\breve{A}$ is the converse of A; thus $(x,y) \in \neg\breve{A}$ if and only if $(y,x) \notin A$.

4.4 Example Let $\mathfrak{D}$ be the relation $({}^\omega\omega, {}^\omega\omega, <^*)$. Then $\|\mathfrak{D}\| = \mathfrak{d}$ and $\|\mathfrak{D}^\perp\| = \|({}^\omega\omega, {}^\omega\omega, \not>^*)\| = \mathfrak{b}$. By Theorem 2.10, the same equations hold if we replace $\mathfrak{D}$ with $\mathfrak{D}' = (IP, IP, \text{is dominated by})$. (Recall from Definition 2.9 that IP is the set of interval partitions.)

Let $\mathfrak{R}$ be the relation $(\mathcal{P}(\omega), [\omega]^\omega, \text{does not split})$. Then $\|\mathfrak{R}\| = \mathfrak{r}$ and $\|\mathfrak{R}^\perp\| = \mathfrak{s}$.

Let $\mathfrak{Hom}_n$ be the relation $(P, [\omega]^\omega, H)$ where P is the set of partitions $f : [\omega]^n \to 2$ and where fHX means that X is almost homogeneous for f. Then $\|\mathfrak{Hom}_n\| = \mathfrak{hom}_n$ and $\|\mathfrak{Hom}_n{}^\perp\| = \mathfrak{par}_n$.

Let $\mathcal{I}$ be an ideal of subsets of X. Let $\mathbf{Cov}(\mathcal{I})$ be the relation $(X, \mathcal{I}, \in)$ and let $\mathbf{Cof}(\mathcal{I})$ be the relation $(\mathcal{I}, \mathcal{I}, \subseteq)$. Then $\|\mathbf{Cov}(\mathcal{I})\| = \mathbf{cov}(\mathcal{I})$, $\|\mathbf{Cov}(\mathcal{I})^\perp\| = \mathbf{non}(\mathcal{I})$, $\|\mathbf{Cof}(\mathcal{I})\| = \mathbf{cof}(\mathcal{I})$, and $\|\mathbf{Cof}(\mathcal{I})^\perp\| = \mathbf{add}(\mathcal{I})$.

In general, we name the relation corresponding to a characteristic by capitalizing the name of the characteristic, except when another name is readily available, e.g., as the dual of a previously defined relation.

4.5 Remark We remarked earlier that the definition of $\mathfrak{d}$ would be unaffected if we replaced $\leq^*$ by $\leq$. That is, $\mathfrak{d}$ is the norm not only of the $\mathfrak{D}$ defined above but also of $({}^\omega\omega, {}^\omega\omega, \leq)$. The dual of this last relation, however, has norm $\aleph_0$, not $\mathfrak{b}$.

Similar remarks apply to $\mathfrak{R}$ and $\mathfrak{Hom}$. It was for the sake of duality that we used "modulo finite" even in definitions where it could have been left out.

The following example indicates another situation where a change in a relation does not affect its norm but might affect the norm of the dual.

4.6 Example Let $\mathfrak{R}_\sigma$ be the relation $({}^\omega\mathcal{P}(\omega), [\omega]^\omega, \text{does not split})$, where an ω-sequence of sets is said to split X if at least one term in the sequence splits X. Then $\|\mathfrak{R}_\sigma\| = \mathfrak{r}_\sigma$ and $\|\mathfrak{R}_\sigma{}^\perp\| = \mathfrak{s}$.

Thus, both $\mathfrak{r}$ and $\mathfrak{r}_\sigma$ can be regarded as duals of $\mathfrak{s}$. Duality is well-defined on relations but in general not on characteristics.

4.7 Remark For any relation $\mathbf{A}$, one can define a relation $\mathbf{A}_\sigma$ that is related to $\mathbf{A}$ as $\mathfrak{R}_\sigma$ in the preceding example is related to $\mathfrak{R}$. That is,

$$\mathbf{A}_\sigma = ({}^\omega A_-, A_+, A_\sigma)$$

where $fA_\sigma a$ means that $f(n)Aa$ for all $n \in \omega$. Thus, $\|\mathbf{A}_\sigma\|$, also written $\|\mathbf{A}\|_\sigma$, is the minimum number of answers needed so that every countably

many challenges can be met simultaneously by a single one of these answers. For some relations, the σ construction produces nothing new; for example, $\mathfrak{d}_\sigma = \mathfrak{d}$. But for other relations, interesting new characteristics arise in this way. We already mentioned $\mathfrak{r}_\sigma$ above; $\mathfrak{s}_\sigma$ is studied in, for example, [63] and [72].

Clearly, $\|\mathbf{A}_\sigma\| \geq \|\mathbf{A}\|$. Whether the reverse inequality is provable in ZFC or whether strict inequality is consistent is, as we mentioned above, an open problem for $\mathbf{A} = \mathfrak{R}$. It is also open for $\mathbf{A} = \mathfrak{R}^\perp$; that is, it is not known whether $\mathfrak{s}_\sigma > \mathfrak{s}$ is consistent. On the other hand, it is known that $\mathbf{cov}(\mathcal{L})_\sigma > \mathbf{cov}(\mathcal{L})$ is consistent. See Bartoszyński's chapter in this handbook for a proof that $\mathbf{cov}(\mathcal{L})$ can consistently have countable cofinality; it is easy to see that no $\|A_\sigma\|$ can have countable cofinality.

Notice that the transformation $\mathbf{A} \mapsto \mathbf{A}_\sigma$ does not commute with duality. Indeed, in all non-trivial cases, $(\mathbf{A}_\sigma)^\perp$ has the same norm as $\mathbf{A}^\perp$, whereas, as indicated above, $(\mathbf{A}^\perp)_\sigma$ may well have a different norm.

The next definition captures the construction used in the proofs of many cardinal characteristic inequalities.

4.8 Definition A *morphism* from one relation $\mathbf{A} = (A_-, A_+, A)$ to another $\mathbf{B} = (B_-, B_+, B)$ is a pair $\varphi = (\varphi_-, \varphi_+)$ of functions such that

- $\varphi_- : B_- \to A_-$
- $\varphi_+ : A_+ \to B_+$
- For all $b \in B_-$ and $a \in A_+$, if $\varphi_-(b) A a$ then $b B \varphi_+(a)$.

We use the terminology "morphism" instead of Vojtáš's "generalized Galois-Tukey connection" partly for brevity and partly because our convention differs from his as to direction. A morphism from $\mathbf{A}$ to $\mathbf{B}$ is a generalized Galois-Tukey connection from $\mathbf{B}$ to $\mathbf{A}$.

It is clear from the definitions that if $\varphi = (\varphi_-, \varphi_+)$ is a morphism from $\mathbf{A}$ to $\mathbf{B}$ then $\varphi^\perp = (\varphi_+, \varphi_-)$ is a morphism from $\mathbf{B}^\perp$ to $\mathbf{A}^\perp$.

Relations and morphisms form (as the name "morphism" suggests) a category in an obvious way, and we shall use the notation $\varphi : \mathbf{A} \to \mathbf{B}$ for morphisms. The category has products and coproducts, but these seem to be of little relevance to cardinal characteristics. Duality is a contravariant involution.

4.9 Theorem *If there is a morphism* $\varphi : \mathbf{A} \to \mathbf{B}$ *then* $\|\mathbf{A}\| \geq \|\mathbf{B}\|$ *and* $\|\mathbf{A}^\perp\| \leq \|\mathbf{B}^\perp\|$.

Proof. It suffices to prove the first inequality, as the second follows by applying the first to the dual morphism $\varphi^\perp$.

Let $X \subseteq A_+$ have cardinality $\|\mathbf{A}\|$ and contain responses meeting all challenges in A_-. Then $Y = \varphi_+(X) \subseteq B_+$ has cardinality $\leq \|\mathbf{A}\|$, so we need only check that it contains responses meeting all challenges from B_-. Given $b \in B_-$, find in X a response x meeting $\varphi_-(b)$. Then $\varphi_+(x)$ is in Y and meets b because, by definition of morphism, $\varphi_-(b)Ax$ implies $bB\varphi_+(x)$. $\dashv$

Morphisms and Theorem 4.9 were implicit in several proofs of inequalities in the preceding sections. For example, the proof of Theorem 2.10 exhibits morphisms in both directions between $\mathfrak{D}$ and $\mathfrak{D}' = (IP, IP, \text{is dominated by})$, where IP is the set of all interval partitions. Both morphisms consist of the same two maps (in opposite order). One map sends any interval partition to the function sending any natural number x to the right endpoint of the next interval of the partition after the interval containing x. The other sends any function $f \in {}^{\omega}\omega$ to an interval partition $\{[j_n, j_{n+1}) : n \in \omega\}$ such that $f(x) < j_{n+1}$ for all $x \leq j_n$. The existence of this pair of morphisms implies not only that $\mathfrak{d} = \|\mathfrak{D}'\|$, but also, by duality, $\mathfrak{b} = \|\mathfrak{D}'^{\perp}\|$. The latter is the second assertion of Theorem 2.10, whose proof we omitted earlier.

The preceding example is somewhat atypical in that the same maps give morphisms in both directions between the same relations. Usually, one has a morphism in only one direction, and therefore an inequality rather than equality between cardinal characteristics. For example, the essential point in the proof of $\mathfrak{s} \leq \mathfrak{d}$ (Theorem 3.3), can be expressed by saying that the functions φ and ψ defined in that proof constitute a morphism $(\psi, \varphi) : \mathfrak{D}' \to \mathfrak{R}^{\perp}$. It follows that they also constitute a morphism $(\varphi, \psi) : \mathfrak{R} \to \mathfrak{D}'^{\perp}$, so we have $\mathfrak{b} \leq \mathfrak{r}$ (Theorem 3.8). Morphisms, duality, and Theorem 4.9 codify the observation that Theorems 3.3 and 3.8 have "essentially the same proof."

If $\mathcal{I}$ is an ideal on X containing all singletons, then in view of Example 4.4, the inequalities $\mathbf{add}(\mathcal{I}) \leq \mathbf{cov}(\mathcal{I}) \leq \mathbf{cof}(\mathcal{I})$ and $\mathbf{add}(\mathcal{I}) \leq \mathbf{non}(\mathcal{I}) \leq \mathbf{cof}(\mathcal{I})$ follow from the existence of morphisms from $\mathbf{Cof}(\mathcal{I}) = (\mathcal{I}, \mathcal{I}, \subseteq)$ to both $\mathbf{Cov}(\mathcal{I}) = (X, \mathcal{I}, \in)$ and its dual $\mathbf{Cov}(\mathcal{I})^{\perp} = (\mathcal{I}, X, \not\ni)$. The first of these can be taken to be (S, id), where S is the singleton map $x \mapsto \{x\}$ and id is the identity map. The second can be taken to be (id, N), where N sends each $I \in \mathcal{I}$ to some element of $X - I$.

The inequalities $\|\mathbf{A}_{\sigma}\| \geq \|\mathbf{A}\|$, for all $\mathbf{A}$, also arise from morphisms $\mathbf{A}_{\sigma} \to \mathbf{A}$. The map on challenges sends each $a \in A_-$ to the constant function $\omega \to A_-$ with value a, and the map on responses is the identity function.

The inequalities $\mathfrak{par}_n \leq \mathfrak{b}$ and $\mathfrak{par}_n \leq \mathfrak{s}$ in Theorem 3.5 and their duals $\mathfrak{hom}_n \geq \mathfrak{d}$ and $\mathfrak{hom}_n \geq \mathfrak{r}$ in Theorem 3.10 are also given by morphisms, as an inspection of the proofs will show. The same goes for Brendle's improvement of the last of these inequalities, with $\mathfrak{r}_{\sigma}$ in place of $\mathfrak{r}$, and the same goes for the analogous inequalities for $\mathfrak{par}_{1,c}$ and $\mathfrak{hom}_{1,c}$.

But the same cannot be said (yet) for the reverse inequalities, $\mathfrak{par}_n \geq \min\{\mathfrak{b}, \mathfrak{s}\}$ and its dual $\mathfrak{hom}_n \leq \max\{\mathfrak{d}, \mathfrak{r}_{\sigma}\}$, simply because the minimum

and maximum here are not (yet) realized as the norms of natural relations. There are, fortunately, several ways to combine two relations into a third whose norm is the maximum (or the minimum) of the norms of the first two. Two of these provide what we need in order to present in terms of morphisms the proofs of the inequalities just cited; we present a third combination along with these two because of its category-theoretic naturality.

To avoid trivial exceptions, we assume in the following that, in the relations (A_-, A_+, A) under consideration, the sets $A_\pm$ are not empty. We also adopt the convention of using a boldface letter for the relation whose components are denoted by the corresponding lightface letter; thus $\mathbf{A} = (A_-, A_+, A)$.

4.10 Definition The *categorical product* $\mathbf{A} \times \mathbf{B}$ is $(A_- \sqcup B_-, A_+ \times B_+, C)$, where $\sqcup$ means disjoint union and where $xC(a,b)$ means xAa if $x \in A_-$ and xBb if $x \in B_-$.

The *conjunction* $\mathbf{A} \wedge \mathbf{B}$ is $(A_- \times B_-, A_+ \times B_+, K)$, where $(x,y)K(a,b)$ means xAa and yBb.

The *sequential composition* $\mathbf{A};\mathbf{B}$ is $(A_- \times {}^{A_+}B_-, A_+ \times B_+, S)$, where the superscript means a set of functions and where $(x,f)S(a,b)$ means xAa and $f(a)Bb$.

The dual operations are the *categorical coproduct* $\mathbf{A} + \mathbf{B} = (\mathbf{A}^\perp \times \mathbf{B}^\perp)^\perp$, the *disjunction* $\mathbf{A} \vee \mathbf{B} = (\mathbf{A}^\perp \wedge \mathbf{B}^\perp)^\perp$, and the *dual sequential composition* $\mathbf{A} \dot{;} \mathbf{B} = (\mathbf{A}^\perp ; \mathbf{B}^\perp)^\perp$.

The two categorical operations are, as their names suggest, the product and coproduct in the category of relations and morphisms.

The conjunction was called the product in a preprint version of [111] and has therefore sometimes been called the *old product*. It is a sort of parallel composition. A challenge consists of separate challenges in both components and a (correct) response consists of (correct) responses in both components separately.

Sequential composition describes a two-inning game between the challenger and the responder. The first inning consists of a challenge x in $\mathbf{A}$ followed by a response a there; the second inning consists of a challenge $f(a)$ in $\mathbf{B}$, which may depend on the previous response a, followed by a response b there. To model this in a single inning, we regard the whole function f as part of the challenge. As in the case of conjunction, a correct response in the sequential composition must be correct in both components.

Notice that one can obtain a description of disjunction by simply changing the last "and" to "or" in the definition of conjunction. The dualization of sequential composition is more complicated; not only does "and" become "or" but the functional dependence changes so that the response in $\mathbf{B}$ can depend on the challenge in $\mathbf{A}$.

The following theorem describes the effect of these operations on norms. Its proof is quite straightforward and therefore omitted.

4.11 Theorem *1.* $\|\mathbf{A} \times \mathbf{B}\| = \max\{\|\mathbf{A}\|, \|\mathbf{B}\|\}$.

2. $\max\{\|\mathbf{A}\|, \|\mathbf{B}\|\} \leq \|\mathbf{A} \wedge \mathbf{B}\| \leq \|\mathbf{A}\| \cdot \|\mathbf{B}\|$.

3. $\|\mathbf{A}; \mathbf{B}\| = \|\mathbf{A}\| \cdot \|\mathbf{B}\|$.

4. $\|\mathbf{A} + \mathbf{B}\| = \min\{\|\mathbf{A}\|, \|\mathbf{B}\|\}$.

5. $\|\mathbf{A} \vee \mathbf{B}\| = \min\{\|\mathbf{A}\|, \|\mathbf{B}\|\}$.

6. $\|\mathbf{A} \dot{;} \mathbf{B}\| = \min\{\|\mathbf{A}\|, \|\mathbf{B}\|\}$.

When the norms are infinite, maxima and products are the same, so the second and third items in the theorem simplify to $\|\mathbf{A} \wedge \mathbf{B}\| = \|\mathbf{A}; \mathbf{B}\| = \max\{\|\mathbf{A}\|, \|\mathbf{B}\|\}$. (In the finite case there is no such simplification. Both of the inequalities involving $\|\mathbf{A} \wedge \mathbf{B}\|$ can be strict; consider $\mathbf{A} = \mathbf{B} = (3, 3, \neq)$.)

4.12 Example In the proof of Theorem 3.10, the part showing that $\mathfrak{hom}_2 \leq \max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$ actually gives a morphism from $\mathfrak{R}_\sigma; (\mathfrak{R} \wedge \mathfrak{D})$ to $\mathfrak{Hom}_2$, as detailed below. By Theorems 4.9 and 4.11, the existence of such a morphism implies both $\mathfrak{hom}_2 \leq \max\{\mathfrak{r}_\sigma, \mathfrak{r}, \mathfrak{d}\} = \max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$ and $\mathfrak{par}_2 \geq \min\{\mathfrak{s}, \mathfrak{b}\}$ (the part of Theorem 3.5 that really involves all three cardinals simultaneously).

To exhibit the morphism implicit in the proof of Theorem 3.10, we first describe $\mathfrak{R}_\sigma; (\mathfrak{R} \wedge \mathfrak{D})$. Following the definitions, we find that a challenge here amounts to a triple (S, F, G) where S is an ω-sequence of subsets S_n of ω, F is a function assigning to each infinite $A \subseteq \omega$ a subset $F(A)$ of ω, and G is a function assigning to each such A a function $G(A) \in {}^\omega\omega$. A response is a triple (A, B, h) where A and B are infinite subsets of ω and $h \in {}^\omega\omega$. The response (A, B, h) meets the challenge (S, F, G) if (1) A is not split by any component S_n of S, (2) B is not split by $F(A)$, and (3) $G(A) <^* h$. Using the notation (f_n, j, g, H) of the proof of Theorem 3.10 and the notation e_A for the increasing enumeration of an infinite $A \subseteq \omega$, we can describe the morphism from $\mathfrak{R}_\sigma; (\mathfrak{R} \wedge \mathfrak{D})$ to $\mathfrak{Hom}_2$ as follows. The "challenge" part sends any partition $f : [\omega]^2 \to 2$ to (S, F, G), where S_n has characteristic function f_n, where $F(A)$ has characteristic function $j \circ e_A$, and where $G(A) = g$. (The j and g in the proof of Theorem 3.10 depend on A.) The "response" part of the morphism sends a triple (A, B, h) to $H(h, A, e_A(B))$. The verification that these two operations constitute a morphism is as in Theorems 3.5 and 3.10. (The need for e_A in the present discussion but not in the earlier proofs results from our tacit use, in the earlier proofs, of the equivalence between splitting phenomena in ω and the analogous phenomena in any infinite subset A. e_A serves to make the equivalence explicit.)

We remark that the formal structure, $\mathfrak{R}_\sigma; (\mathfrak{R} \wedge \mathfrak{D})$, reflects the intuitive structure of the proof of Theorem 3.5. That proof invoked the hypothesis $\kappa < \mathfrak{s}$ twice (corresponding to $\mathfrak{R}_\sigma$ and $\mathfrak{R}$) and $\kappa < \mathfrak{b}$ once (corresponding to $\mathfrak{D}$). The first use of $\kappa < \mathfrak{s}$ logically precedes the other two (corresponding

to sequential composition) because the unsplit set A obtained at the first step is used to produce the j and g for the other two steps. The second use of $\kappa < \mathfrak{s}$ and the use of $\kappa < \mathfrak{b}$ can proceed in parallel, as neither depends on the other (corresponding to conjunction).

4.13 Example Sequential composition also occurs naturally in much simpler situations. Consider, for example, the following variant of unsplitting: $\mathfrak{R}_3 = ({}^{\omega}3, [\omega]^{\omega}, \text{is almost constant on})$. Its norm $\mathfrak{r}_3$ is the minimum number of infinite subsets of ω not all split by a single partition of ω into three pieces. This cardinal is easily seen to be equal to $\mathfrak{r}$, but one direction of the proof involves a sequential composition. A "3-unsplittable" family is obtained by starting with an unsplittable family and then forming, within each of its sets, a further unsplittable family. The union of the latter families is then 3-unsplittable (and even 4-unsplittable). In terms of morphisms, one obtains $\mathfrak{R}; \mathfrak{R} \to \mathfrak{R}_3$ (as well as the trivial $\mathfrak{R}_3 \to \mathfrak{R}$).

Equipped with the concept of morphism, we can address an issue that was glossed over in the introduction. If one believes the continuum hypothesis (CH), then the theory of cardinal characteristics becomes trivial, for they are all equal to $\aleph_1$. Nevertheless, there is non-trivial combinatorial content in proofs like those of Theorems 2.10 and 3.3, even if CH holds and makes the theorems themselves trivial. That combinatorial content is used to construct the morphisms $\mathfrak{D} \leftrightarrow \mathfrak{D}' \to \mathfrak{R}^{\perp}$, so one might hope that the existence of such morphisms is what the argument "really" proves, a non-trivial result even in the presence of CH. Yiparaki [113] showed that this hope is not justified; CH implies not only the equality of all our cardinal characteristics but also the existence of morphisms in both directions between the corresponding relations. The last part of the following theorem embodies this result.

4.14 Theorem *Let* $\mathbf{A} = (A_-, A_+, A)$ *and* $\mathbf{B} = (B_-, B_+, B)$ *be two relations and let* κ *be an infinite cardinal.*

1. $\|\mathbf{A}\| \leq \kappa$ *if and only if there is a morphism from* $(\kappa, \kappa, =)$ *to* $\mathbf{A}$.
2. *If* $\|\mathbf{A}\| = |A_+| = \kappa$, *then there is a morphism from* $\mathbf{A}$ *to* $(\kappa, \kappa, <)$.
3. *If* $\|\mathbf{A}^{\perp}\| = |A_-| = \kappa$, *then there is a morphism from* $(\kappa, \kappa, <)$ *to* $\mathbf{A}$.
4. *If* $\|\mathbf{A}\| = |A_+| = \|\mathbf{B}^{\perp}\| = |B_-| \geq \aleph_0$, *then there is a morphism from* $\mathbf{A}$ *to* $\mathbf{B}$.

Proof. The "if" direction of (1) is immediate from Theorem 4.9 and the fact that $\|(\kappa, \kappa, =)\| = \kappa$. For the "only if" direction, let $\varphi_+ : \kappa \to A_+$ enumerate a set of at most κ responses meeting all challenges; then for each challenge $a \in A_-$ let $\varphi_-(a)$ be any $\alpha < \kappa$ such that $\varphi_+(\alpha)$ meets a.

For (2), let $\varphi_+ : A_+ \to \kappa$ be any one-to-one map. Then, for any $\alpha < \kappa$, the set $\{a \in A_+ : \varphi_+(a) \leq \alpha\}$ has cardinality smaller than $\kappa = \|\mathbf{A}\|$, so

some challenge in A_- has no correct response in this set. Let $\varphi_-(\alpha)$ be any such challenge.

Note that (2) remains true if we replace $(\kappa, \kappa, <)$ with $(\kappa, \kappa, \leq)$. Then dualization gives (3).

Finally, to prove (4), just compose the morphisms $\mathbf{A} \to (\kappa, \kappa, <) \to \mathbf{B}$ given by (2) and (3). $\dashv$

If CH holds, then part 4 of this theorem applies to most of the relations in Example 4.4 above, for the cardinals involved are $\aleph_1$. The only exceptions are $\mathbf{Cov}(\mathcal{I})$ and $\mathbf{Cof}(\mathcal{I})$, but even here we can (indirectly) apply part 4 when $\mathcal{I}$ is the ideal of measure zero sets or the ideal of meager sets in $\mathbb{R}$ or a similar ideal. More precisely, if $\mathcal{I}$ is an ideal on $\mathbb{R}$ and $\mathcal{I}$ has a cofinal subset $\mathcal{I}_0$ of size $\leq \mathfrak{c}$, then part 4 applies directly to variants of $\mathbf{Cov}(\mathcal{I})$ and $\mathbf{Cof}(\mathcal{I})$ with $\mathcal{I}$ replaced by $\mathcal{I}_0$. But it is trivial to check that there are morphisms in both directions between these variants and the original relations. In effect then, part 4 provides morphisms in both directions between any two of the relations we are considering; CH trivializes not only the inequalities between cardinal characteristics but also the morphisms between the corresponding relations.

Nevertheless, there is still some hope of using morphisms to describe the combinatorial content of the theory in a way that makes good sense even when CH holds. This hope is based on the observation that the morphisms given by Theorem 4.14 are highly non-constructive; they involve well-orderings of the continuum (and similar sets). By contrast, the morphisms given by the proofs of cardinal characteristic inequalities are much better behaved. They consist of Borel maps with respect to the usual topologies on the sets involved (like ${}^\omega\omega$ and $\mathcal{P}(\omega)$). Two clarifications are in order here. One is that, when the sets involved are bases $\mathcal{I}_0$ for some ideals, as in the preceding paragraph, then the sets in $\mathcal{I}_0$ should be coded by reals in some standard way. For example, if $\mathcal{I}$ is the ideal of meager (resp. measure zero) sets in $\mathbb{R}$, then $\mathcal{I}_0$ can be taken to consist of the F_σ (resp. G_δ) members of $\mathcal{I}$, and there are well-known ways of coding such sets (or arbitrary Borel sets) by reals. The second clarification is that Pawlikowski and Recław have shown [84] that, with suitable coding, the morphisms can be taken to consist of continuous maps; nevertheless, we shall continue to use "Borel" as our main criterion of simplicity.

The existence of Borel morphisms seems to serve well as a codification of the combinatorial content of proofs of cardinal characteristic inequalities. On the one hand, the usual proofs provide Borel morphisms. On the other hand, when an inequality is not provable then, although it may hold in specific models and even have morphisms attesting to it (e.g., in models of CH), there will never be Borel morphisms attesting to it. The following theorem establishes this last fact for the particular unprovable inequality $\mathfrak{d} \leq \mathfrak{s}$. Similar arguments can be given for other unprovable inequalities,

but they usually involve notions of forcing more complicated than the Cohen forcing used here. We remark that the theorem proves a bit more than was claimed above; a morphism φ attesting to $\mathfrak{d} \leq \mathfrak{s}$ cannot have even one of its two constituent functions $\varphi_\pm$ Borel. (The weaker result that $\varphi_\pm$ cannot both be Borel in this situation was established in [22].)

4.15 Theorem *If φ is a morphism $\mathfrak{R}^\perp \to \mathfrak{D}$, then neither φ_+ nor φ_- is a Borel function.*

Proof. Recalling the definitions of $\mathfrak{R}$ and $\mathfrak{D}$, we see that

$$\begin{aligned} \varphi_- : {}^\omega\omega &\to [\omega]^\omega, \\ \varphi_+ : \mathcal{P}(\omega) &\to {}^\omega\omega, \text{ and} \\ \varphi_-(a) \text{ is split by } b &\implies a <^* \varphi_+(b). \end{aligned}$$

Suppose first that φ_- were a Borel function, with code p (in a standard coding system for Borel sets and functions). Adjoin to the universe a Cohen-generic function $c : \omega \to \omega$, and define $d = \tilde{\varphi}_-(c)$, where $\tilde{\varphi}_-$ is the Borel function coded by p in $V[c]$. Thus $d \in [\omega]^\omega$ in $V[c]$. The ground model reals form a splitting family in the Cohen extension $V[c]$ (because they form a non-meager family there; see Subsection 11.3 and the proof of Theorem 5.19 below). So there is a real $r \in V \cap \mathcal{P}(\omega)$ that splits d. In the ground model V, let $g = \varphi_+(r)$ and notice that, because φ is a morphism,

$$\forall x \in {}^\omega\omega \, [\varphi_-(x) \text{ is split by } r \implies x <^* g].$$

This is a Π^1_1 statement about r, g, and the code p of φ_-. So it remains true in $V[c]$, where p codes $\tilde{\varphi}_-$ and where x can take c as a value. Thus we find, in $V[c]$, since $\tilde{\varphi}_-(c) = d$ is split by r, that $c <^* g$. But this is absurd; a Cohen-generic $c \in {}^\omega\omega$ cannot be dominated by a g from the ground model. This contradiction shows that φ_- cannot be a Borel map.

Now suppose instead that φ_+ were a Borel function, with Borel code p. Let $c \in \mathcal{P}(\omega)$ be Cohen-generic and let $e = \tilde{\varphi}_+(c)$, where $\tilde{\varphi}_+$ is the Borel function coded by p in $V[c]$. Thus $e \in {}^\omega\omega$ in $V[c]$. The ground model reals are unbounded in ${}^\omega\omega$ in a Cohen extension, so fix $r \in V \cap {}^\omega\omega$ with $r \not<^* e$. Let $q = \varphi_-(r)$, an infinite subset of ω in V. Because φ is a morphism,

$$\forall x \in \mathcal{P}(\omega) \, [q \text{ is split by } x \implies r <^* \varphi_+(x)].$$

As before, this is a Π^1_1 statement about q, r, and p, so it remains true in $V[c]$. There c is a possible value of x and p codes $\tilde{\varphi}_+$, so from $r \not<^* e = \tilde{\varphi}_+(c)$ we can infer that q is not split by c. This is absurd, as every infinite subset of ω in the ground model V is split by the Cohen subset c of ω. ⊣

The use of Borel morphisms can also clarify the need for sequential (and other) composition operations on relations. Specifically, a forcing argument

is used in [22] to show that some naturally occurring morphisms involving sequential compositions (e.g., the proof of Theorem 5.6 below) cannot be simplified to use conjunctions or products or even sequential composition in a different order. Mildenberger [76] and Spinas [106] have obtained similar results by combinatorial methods in some cases where the forcing method of [22] does not apply. A forcing argument in [27] shows that the sequential composition $\mathfrak{R}_\sigma ; (\mathfrak{R} \wedge \mathfrak{D})$ used in the proof of $\mathfrak{hom}_2 \leq \max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$ cannot be replaced by simply $\mathfrak{R}_\sigma \wedge \mathfrak{D}$. But other potential simplifications in this problem and similar simplifications in other problems, though they seem unlikely, have not been proved impossible.

4.16 Remark Let $\mathbf{A}$ and $\mathbf{B}$ be relations where $A_\pm$ and $B_\pm$ are sets of reals. Call a morphism $\varphi : \mathbf{A} \to \mathbf{B}$ *semi-Borel* (on the positive side) if φ_+ is a Borel function. Thus, Theorem 4.15 asserts that certain morphisms cannot be semi-Borel.

Call a set X of reals *small with respect to* $\mathbf{A}$ if there is no semi-Borel morphism from $(X, X, =)$ to $\mathbf{A}$. Without "semi-Borel," this definition would say simply that $|X| < \|\mathbf{A}\|$, by the first part of Theorem 4.14. With "semi-Borel" smallness is a weaker notion, related to the topological (or Borel) structure of X, not just to its cardinality. It can be expressed as "no image of X under a Borel function to A_+ contains responses meeting all challenges from A_-."

The smallness properties associated in this way to the relations involved in Cichoń's diagram were introduced and studied by Pawlikowski and Recław [84], who connected them with various classical smallness properties of sets of reals. Bartoszyński's chapter in this handbook contains extensive information about this topic.

5. Category and Measure

Despite their origins in real analysis, Baire category and Lebesgue measure are, to a large extent, combinatorial notions. As such, they have close ties with some of the objects discussed in the preceding sections. We give here a rather cursory presentation of some of these combinatorial aspects of category and measure. For a more complete treatment, see Bartoszyński's chapter in this handbook and the book [5] of Bartoszyński and Judah.

Recall Definition 2.7 of the four cardinal characteristics **add**, **cov**, **non**, **cof** associated to any proper ideal (containing all singletons) on any set. We shall be interested in these and in the corresponding relations ($\mathbf{Cof}^\perp$, **Cov**, $\mathbf{Cov}^\perp$, and **Cof**, respectively, from Example 4.4) when the ideal is either the σ-ideal of meager (also called first category) sets or the σ-ideal of sets of Lebesgue measure zero (also called null sets). We use $\mathcal{B}$ and $\mathcal{L}$ respectively to denote these two ideals. (The notation stands for "Baire" and

"Lebesgue"; other authors have used $\mathcal{C}$ for "Category," $\mathcal{K}$ for "Kategorie," $\mathcal{M}$ for "meager," $\mathcal{M}$ for "measure," and $\mathcal{N}$ for "null.") As indicated in the introduction, we do not distinguish notationally between the meager ideals on various versions of the continuum, $\mathbb{R}$, ${}^{\omega}2$, ${}^{\omega}\omega$, etc., and similarly for measure. The various versions of each cardinal characteristic are equal; the various versions of each relation admit morphisms in both directions. We tolerate an additional, equally innocuous ambiguity by not distinguishing between an ideal and a basis for it. Thus, we may pretend that $\mathcal{B}$ consists of meager F_σ sets and that $\mathcal{L}$ consists of G_δ null sets. If we discuss Borel morphisms, we further identify F_σ and G_δ sets with some standard encoding as reals.

We begin our treatment of Baire category by giving a convenient combinatorial description of meagerness in the space ${}^{\omega}2$. This idea was introduced in a more specialized context by Talagrand [107].

5.1 Definition A *chopped real* is a pair (x, Π), where $x \in {}^{\omega}2$ and Π is an interval partition of ω. Recall that we introduced the notation IP for the set of all interval partitions; we write CR for the set ${}^{\omega}2 \times IP$ of chopped reals. A real $y \in {}^{\omega}2$ *matches* a chopped real (x, Π) if $x \restriction I = y \restriction I$ for infinitely many intervals $I \in \Pi$.

5.2 Theorem *A subset M of ${}^{\omega}2$ is meager if and only if there is a chopped real that no member of M matches.*

Proof. The set of reals y that match a given chopped real $(x, \{I_n : n \in \omega\})$ is

$$\mathrm{Match}(x, \{I_n : n \in \omega\}) = \bigcap_k \bigcup_{n \geq k} \{y : x \restriction I_n = y \restriction I_n\},$$

the intersection of countably many dense open sets. So $\mathrm{Match}(x, \Pi)$ is comeager, and the "if" part of the theorem follows.

To prove "only if," suppose M is meager, and fix a countable sequence of nowhere dense sets F_n that cover M. Note that, for the standard (product) topology on ${}^{\omega}2$, to say that a set F is nowhere dense means that for every finite sequence $s \in {}^{<\omega}2$ there is an extension $t \in {}^{<\omega}2$ such that no $y \in F$ extends t. Note also that the union of finitely many nowhere dense sets is nowhere dense, so we can and do arrange that $F_n \subseteq F_{n+1}$ for all n. Then we can complete the proof by constructing a chopped real $(x, \{I_n : n \in \omega\})$ such that, for each n, no real in F_n agrees with x on I_n. This suffices because then any y that matches $(x, \{I_n : n \in \omega\})$ will be outside infinitely many F_n, hence outside them all by monotonicity, and hence outside M.

To define I_n and $x \restriction I_n$, suppose the earlier I_k $(k < n)$ are already defined and are contiguous intervals. So we know the point m where I_n should start. I_n will be the union of 2^m contiguous subintervals J_i $(i < 2^m)$ defined as follows. List all the functions $m \to 2$ as u_i $(i < 2^m)$. By

induction on i, choose J_i and $x \restriction J_i$ so that no element of F_n is an extension of $u_i \cup \bigcup_{j \leq i}(x \restriction J_j)$. These choices are possible because F_n is nowhere dense. Finally, let $I_n = \bigcup_{j<2^m} J_i$; having already defined each $x \restriction J_i$, we have determined $x \restriction I_n$.

If y agrees with x on I_n, then y extends $u_i \cup \bigcup_{j \leq i}(x \restriction J_j)$ for some i, namely the i such that $u_i = y \restriction m$. Therefore, $y \notin F_n$, as required. $\dashv$

The theorem shows that the sets $\mathrm{Match}(x, \Pi)$ form a base for the filter of comeager sets and so their complements form a base for the ideal $\mathcal{B}$. We may therefore confine attention to these complements when discussing the cardinal characteristics of $\mathcal{B}$ and the associated relations. In this connection, it is useful to have the following combinatorial formulation of the inclusion relation between these sets; we leave the straightforward proof to the reader.

5.3 Proposition *$\mathrm{Match}(x, \Pi) \subseteq \mathrm{Match}(x', \Pi')$ if and only if for all but finitely many intervals $I \in \Pi$ there exists an interval $J \in \Pi'$ such that $J \subseteq I$ and $x' \restriction J = x \restriction J$.*

We shall say that (x, Π) *engulfs* (x', Π') when the equivalent conditions in the proposition hold.

Thus, we have morphisms in both directions between $\mathbf{Cof}(\mathcal{B})$ and

$$\mathbf{Cof}'(\mathcal{B}) = (CR, CR, \text{is engulfed by}),$$

as well as morphisms in both directions between $\mathbf{Cov}(\mathcal{B})$ and

$$\mathbf{Cov}'(\mathcal{B}) = ({}^{\omega}2, CR, \text{does not match}).$$

Notice that if (x, Π) engulfs (x', Π') then Π dominates Π'. Combining this with the characterization of $\mathfrak{d}$ and $\mathfrak{b}$ in Theorem 2.10 and the characterization of $\mathbf{add}(\mathcal{B})$ and $\mathbf{cof}(\mathcal{B})$ in Example 4.4, we obtain the following inequalities.

5.4 Corollary $\mathbf{add}(\mathcal{B}) \leq \mathfrak{b}$ *and* $\mathfrak{d} \leq \mathbf{cof}(\mathcal{B})$.

Another relation between the characteristics from Section 2 and the characteristics of Baire category follows from Theorem 2.8.

5.5 Proposition $\mathfrak{b} \leq \mathbf{non}(\mathcal{B})$ *and* $\mathbf{cov}(\mathcal{B}) \leq \mathfrak{d}$.

Proof. In ${}^{\omega}\omega$, any set of the form $\{f : f \leq g\}$ is clearly nowhere dense (because every finite sequence in ${}^{<\omega}\omega$ has an extension in ${}^{<\omega}\omega$ with some values greater than the corresponding values of g). The proof of Theorem 2.8 shows, therefore, that all compact sets in ${}^{\omega}\omega$ are nowhere dense and therefore $\mathcal{K}_\sigma \subseteq \mathcal{B}$. That immediately implies $\mathbf{cov}(\mathcal{K}_\sigma) \geq \mathbf{cov}(\mathcal{B})$ and

$\mathbf{non}(\mathcal{K}_\sigma) \leq \mathbf{non}(\mathcal{B})$. (Indeed, whenever $\mathcal{I} \subseteq \mathcal{J}$ are ideals, we have a morphism $\mathbf{Cov}(\mathcal{I}) \to \mathbf{Cov}(\mathcal{J})$ given by the identity map on challenges and the inclusion map on responses.) Now Theorem 2.8 completes the proof. ⊣

All ZFC-provable inequalities among $\mathfrak{b}$, $\mathfrak{d}$, and the four characteristics of $\mathcal{B}$ are obtainable by transitivity from the preceding corollary and proposition and the general facts that $\mathbf{add} \leq \mathbf{cov} \leq \mathbf{cof}$ and $\mathbf{add} \leq \mathbf{non} \leq \mathbf{cof}$ for any nontrivial ideal. There are, however, two additional relations due to Miller [79] and Truss [109], each involving three of these cardinals.

5.6 Theorem *1. There is a morphism from* $(\mathbf{Cov}'(\mathcal{B}))^\perp; \mathfrak{D}'$ *to* $\mathbf{Cof}'(\mathcal{B})$.

2. $\mathbf{cof}(\mathcal{B}) = \max\{\mathbf{non}(\mathcal{B}), \mathfrak{d}\}$.

3. $\mathbf{add}(\mathcal{B}) = \min\{\mathbf{cov}(\mathcal{B}), \mathfrak{b}\}$.

Proof. Recall that $\mathfrak{D}' = (IP, IP, \text{is dominated by})$ where IP is the set of interval partitions, that $\|\mathfrak{D}'\| = \mathfrak{d}$, and that $\|\mathfrak{D}'^\perp\| = \mathfrak{b}$. Thus, if we prove part 1 of the theorem, then the $\leq$ half of part 2 and the $\geq$ half of part 3 will follow by Theorems 4.9 and 4.11. The other halves of parts 2 and 3 were already established, so we need only prove part 1.

A morphism φ as claimed in part 1 would consist of a function φ_- from the set CR of chopped reals to $CR \times {}^{(\omega 2)}IP$ and a function φ_+ from ${}^\omega 2 \times IP$ to CR, satisfying an implication to be exhibited after we simplify notation a bit. As a map into a product, φ_- consists of two maps, $\alpha : CR \to CR$ and $\beta : CR \to {}^{(\omega 2)}IP$. We shall take α and φ_+ to be identity maps. (Recall that $CR = {}^\omega 2 \times IP$, so this makes sense.) It remains to define β so as to satisfy the required implication, which now reads: For all $x \in CR$, all $y \in {}^\omega 2$, and all $\Pi \in IP$,

$$[y \text{ matches } x \text{ and } \Pi \text{ dominates } \beta(x)(y)] \implies [(y, \Pi) \text{ engulfs } x].$$

It does not matter how we define $\beta(x)(y)$ when y does not match x. If y does match x, i.e., if there are infinitely many intervals I in the partition component of the chopped real x on which x and y agree, then we define $\beta(x)(y)$ to be some interval partition each of whose intervals includes at least one such I. ⊣

5.7 Remark It is easy to specify the β in the last part of the proof more explicitly so that $\beta(x)(y)$ is a Borel function of x and y; since the other components of φ are trivial, we can say that part 1 of the theorem is witnessed by a Borel morphism. It is shown in [22] that one cannot get a Borel morphism in part 1 if one replaces the sequential product there with the categorical product, or the conjunction, or the sequential product in the other order.

Before turning from category to measure, we give an elegant, combinatorial description of $\mathbf{cov}(\mathcal{B})$, due to Bartoszyński [4].

5.8 Definition Call two functions $x, y \in {}^{\omega}\omega$ *infinitely equal* if $\exists^{\infty} n\,(x(n) = y(n))$ and *eventually different* otherwise, i.e., if $\forall^{\infty} n\,(x(n) \neq y(n))$.

5.9 Theorem *1.* $\mathbf{cov}(\mathcal{B}) = \|({}^{\omega}\omega, {}^{\omega}\omega, \textit{eventually different})\|$.

2. $\mathbf{non}(\mathcal{B}) = \|({}^{\omega}\omega, {}^{\omega}\omega, \textit{infinitely equal})\|$.

Proof. We prove only part 1 as part 2 is dual to it. The $\leq$ direction is clear once one observes that, for any $x \in {}^{\omega}\omega$, the set of $y \in {}^{\omega}\omega$ eventually different from x is meager. (In fact, sending x to this set defines half of a morphism from the relation on the right of part 1 to $\mathbf{Cov}(\mathcal{B})$ (when the reals are taken to be ${}^{\omega}\omega$); the other half of the morphism is the identity map.)

To prove the $\geq$ direction of part 1, we show how to match, with a single real y, all the chopped reals in a family $\{(x_\alpha, \Pi_\alpha) : \alpha < \kappa\}$, where

$$\kappa < \|({}^{\omega}\omega, {}^{\omega}\omega, \text{eventually different})\|.$$

Note that the norm here is trivially $\leq \mathfrak{d}$ (there's a morphism from $\mathfrak{D}$ consisting of the identity map in both directions). So by Theorem 2.10 there is an interval partition Θ not dominated by any Π_α.

Temporarily fix an arbitrary $\alpha < \kappa$. Non-domination means that Π_α has infinitely many intervals that include no interval of Θ and are therefore covered by two adjacent intervals of Θ. Call a pair of adjacent intervals of Θ good if they cover an interval of Π_α; so there are infinitely many good pairs.

Define a function f_α on ω as follows. $f_\alpha(n)$ is obtained by taking $2n+1$ disjoint good pairs, taking the union of the two intervals in each pair to obtain $2n+1$ intervals $J_0, \ldots, J_{2n}$, and then forming the set of restrictions of x_α to these intervals:

$$f_\alpha(n) = \{x_\alpha \restriction J_0, \ldots, x_\alpha \restriction J_{2n}\}.$$

Note that, although the values of f_α are not natural numbers, they can be coded as natural numbers.

Now un-fix α. By our hypothesis on κ, find a function g infinitely equal to each f_α. Without harming this property of g, we can arrange that, for each n, $g(n)$ is a set of $2n+1$ functions, each mapping an interval of ω to 2. Furthermore, we can arrange that these $2n+1$ intervals are disjoint and each of them is the union of two adjacent intervals of Θ. (Any n for which $g(n)$ is not of this form could not contribute to the agreement between g and any f_α, so we are free to modify $g(n)$ arbitrarily.)

We define a function $y : \omega \to 2$ by recursion, where at each stage we specify the restriction of y to a certain pair of adjacent intervals in Θ. After stages 0 through $n-1$ are completed, y is defined on only $2n$ intervals of Θ, so at least one of the $2n+1$ members of $g(n)$, say $z(n)$, has its domain J disjoint from where y is already defined. Extend y to agree with $z(n)$ on J. This completes the recursion; if there are places where y never gets defined, define it arbitrarily there.

To complete the proof, we show that y matches every (x_α, Π_α). Consider any α and any one of the infinitely many n for which $g(n) = f_\alpha(n)$. At stage n of the construction of y, we ensured that y extends some $z(n) \in g(n) = f_\alpha(n)$. But the construction of $f_\alpha(n)$ ensures that $z(n)$ is the restriction of x_α to an interval (the union of a good pair of intervals from Θ) that includes an interval of Π_α. Thus y agrees with x_α on that interval of Π_α. Since this happens for infinitely many n, y matches (x_α, Π_α). ⊣

5.10 Remark The preceding proof exhibits a morphism from $\mathbf{Cov}'(\mathcal{B})$ to $\mathfrak{D}'_{\breve{}}({}^\omega\omega, {}^\omega\omega, \text{eventually different})$. Ignoring the coding needed to make f_α and g functions into ω, we can say that the "challenge" half of the morphism is the construction of y from Θ and g and the "response" half of the morphism sends any (x, Π) (where we omit the α subscripts needed in the proof but not here) to the pair consisting of Π and the function that maps any Θ not dominating Π to the f as in the proof (and maps Θ's that dominate Π arbitrarily).

It is an open problem whether one can omit the "$\mathfrak{D}'_{\breve{}}$" part, i.e., whether there is a Borel morphism from $\mathbf{Cov}'(\mathcal{B})$ to $({}^\omega\omega, {}^\omega\omega, \text{eventually different})$. An essentially equivalent question is whether any forcing that adds a real (in ${}^\omega\omega$) infinitely equal to all ground model reals (called a "half-Cohen" real) must add a Cohen real. The proof above shows that if one first adds an unbounded real and then a half-Cohen real over the resulting model, the final model contains a Cohen real over the ground model.

We now turn to Lebesgue measure (and equivalent measures on ${}^\omega 2$, ${}^\omega\omega$, etc.) and its connections with Baire category. The first such connection was given by Rothberger [89].

5.11 Theorem $\mathbf{cov}(\mathcal{B}) \leq \mathbf{non}(\mathcal{L})$ *and* $\mathbf{cov}(\mathcal{L}) \leq \mathbf{non}(\mathcal{B})$.

Proof. Let Π be the interval partition whose n^{th} interval I_n has $n+1$ elements for all n. Define a binary relation R on ${}^\omega 2$ by letting xRy mean that $x \restriction I_n = y \restriction I_n$ for infinitely many n, i.e., that y matches the chopped real (x, Π). Notice that R is symmetric and, for every x, the set $R_x = \{y : xRy\}$ is a comeager set of measure zero. ("Comeager" was proved in Theorem 5.2. The calculation for "measure zero" consists of noticing that, once x is fixed, the y's that agree with it on I_n form a set of measure $2^{-(n+1)}$, so the y's

that agree with x on at least one I_n beyond I_k form a set of measure at most 2^{-k}, and so the y's that do this for all k form a set of measure zero.)

Thus, letting $\mathbf{R} = ({}^\omega 2, {}^\omega 2, R)$, we have morphisms $\varphi : \mathbf{R} \to \mathbf{Cov}(\mathcal{L})$ and $\psi : \mathbf{R}^\perp \to \mathbf{Cov}(\mathcal{B})$, where φ_+ and ψ_+ send x to R_x and ${}^\omega 2 - R_x$ respectively and where both φ_- and ψ_- are the identity on ${}^\omega 2$. Composing each of these morphisms with the dual of the other, we get morphisms $\mathbf{Cov}(\mathcal{B})^\perp \to \mathbf{Cov}(\mathcal{L})$ and $\mathbf{Cov}(\mathcal{L})^\perp \to \mathbf{Cov}(\mathcal{B})$. Since $\mathbf{cov} = \|\mathbf{Cov}\|$ and $\mathbf{non} = \|\mathbf{Cov}^\perp\|$ for both ideals, the theorem follows. $\dashv$

5.12 Remark The relation R in the preceding proof could be replaced by any relation of the form "$x \oplus y \in M$" where $\oplus$ is addition modulo 2 and M is any comeager set of measure zero. For example, M could be the set of 0-1 sequences in which the density of 1's in initial segments does not approach $1/2$.

In this form, the proof generalizes to any pair of translation-invariant (with respect to $\oplus$) ideals that concentrate on disjoint sets.

The rest of our discussion of measure characteristics is based on a combinatorial characterization, due to Bartoszyński [3], of $\mathbf{add}(\mathcal{L})$. To formulate it, we need the following terminology.

5.13 Definition A *slalom* is a function S assigning to each $n \in \omega$ a set $S(n) \subseteq \omega$ of cardinality n. We say that a real $x \in {}^\omega\omega$ *goes through* slalom S if $\forall^\infty n\, (x(n) \in S(n))$.

5.14 Theorem $\mathbf{add}(\mathcal{L})$ *is the smallest cardinality of any family $\mathcal{F} \subseteq {}^\omega\omega$ such that there is no single slalom through which all the members of $\mathcal{F}$ go.*

For the proof, we refer to Bartoszyński's original paper [3], his chapter in this handbook, his book with Judah [5, Theorem 2.3.9], or Fremlin's article [46].

5.15 Remark The theorem would remain true if we modified the definition of "slalom" by requiring $S(n)$ to have cardinality $f(n)$ instead of n; here f can be any function $\omega \to \omega$ that grows without bound. We refer to this modified notion of slalom as an f-slalom (or $f(n)$-slalom). Suppose, for example, that κ is a cardinal such that every κ functions in ${}^\omega\omega$ go through a single f-slalom. To show that any κ functions x_α go through a single slalom in the original sense, partition ω into intervals such that the nth interval starts at or after $f(n)$. Let $y_\alpha(n)$ be (or code) the restriction of x_α to the nth interval. From an f-slalom through which all the y_α go, one easily gets a slalom in the orginal sense through which all the x_α go.

Despite this observation, it is not true that one could simply omit the cardinality bound (n or $f(n)$) in the definition of slalom and merely require each $S(n)$ to be finite. Indeed, with this weakening, the cardinal described in the theorem would be simply $\mathfrak{b}$, which can be strictly larger than $\mathbf{add}(\mathcal{L})$.

As indicated for example in Oxtoby's book [82], there are a great many similarities between Baire category and Lebesgue measure. The following inequality, due to Bartoszyński [3] and independently but a bit later to Raisonnier and Stern [88], was an early indication that the symmetry is not so extensive as one might have thought. (The first indication of this was Shelah's proof [97] that ZF (without choice) plus "all sets of reals have the Baire property" is consistent if ZF is, whereas the consistency of ZF plus "all sets of reals are Lebesgue measurable" requires the consistency of an inaccessible cardinal.) The theme of measure-category asymmetry is developed much further in the book [5].

5.16 Theorem $\mathbf{add}(\mathcal{L}) \leq \mathbf{add}(\mathcal{B})$.

Proof. In view of Theorem 5.6, it suffices to prove that $\mathbf{add}(\mathcal{L}) \leq \mathfrak{b}$ and $\mathbf{add}(\mathcal{L}) \leq \mathbf{cov}(\mathcal{B})$. The former is immediate, in view of Theorem 5.14, for a family of reals going through a single slalom is obviously bounded. (It should be mentioned that the inequality $\mathbf{add}(\mathcal{L}) \leq \mathfrak{b}$ was originally proved by Miller [80] before Theorem 5.14 was known.) For the second inequality, we use Theorem 5.9.

If $\kappa < \mathbf{add}(\mathcal{L})$ and if we are given κ functions $x_\alpha \in {}^\omega\omega$, we must find a single function y infinitely equal to them all. Fix an interval partition Π whose n^{th} interval I_n has cardinality $\geq n$. To each x_α associate the function $x'_\alpha \in {}^\omega\omega$ where $x'_\alpha(n)$ codes (in some standard way) $x_\alpha \restriction I_n$. Let S be a slalom through which all the x'_α go. We may assume that all n elements of $S(n)$ code functions $I_n \to \omega$, for any other elements can be replaced with such codes without harming the fact that all x'_α go through S. For each n, choose a function $y_n : I_n \to \omega$ that agrees at least once with each of the n members of $S(n)$; this is trivial to arrange, since $|I_n| \geq n$. Then the union of all the y_n is the desired y. Indeed, every x_α agrees with y at least once in each I_n except for finitely many. ⊣

5.17 Remark The preceding proof, though short, has a defect from the point of view of morphisms between relations. Because it relies on theorems 5.6 and 5.9, it involves sequential compositions. In fact, it provides a Borel morphism from $\mathbf{Cof}(\mathcal{L}); \mathbf{Cof}(\mathcal{L}); \mathbf{Cof}(\mathcal{L})$ to $\mathbf{Cof}(\mathcal{B})$. The presence of these sequential compositions is an artifact of the proof. Bartoszyński's chapter in this handbook contains a different proof, giving a Borel morphism from $\mathbf{Cof}(\mathcal{L})$ to $\mathbf{Cof}(\mathcal{B})$.

Since the proof gave a morphism, we also have the dual result.

5.18 Corollary $\mathbf{cof}(\mathcal{B}) \leq \mathbf{cof}(\mathcal{L})$.

Our discussion of the four standard characteristics (**add**, **cov**, **non**, and **cof**) of measure and category, along with $\mathfrak{b}$ and $\mathfrak{d}$, is now complete, in the following strong sense. If one assigns to each of these ten characteristics one of the values $\aleph_1$ and $\aleph_2$, and if the assignment is consistent with the equations and inequalities proved above, then that assignment is realized in some model of ZFC. We shall comment on a few of these models in Section 11 below, but we refer to [6] or [5, Chapter 7] for all the details.

The inequalities between these ten cardinal characteristics are summarized in the following picture, known as *Cichoń's diagram*, in which one goes from larger to smaller cardinals by moving down or to the left along the arrows. (A 45° counterclockwise rotation would produce a Hasse diagram in the customary orientation. We've drawn the arrows in the direction of the morphisms between the corresponding relations, hence from larger to smaller characteristics.)

$$\begin{array}{ccccccc}
\mathbf{cov}(\mathcal{L}) & \longleftarrow & \mathbf{non}(\mathcal{B}) & \longleftarrow & \mathbf{cof}(\mathcal{B}) & \longleftarrow & \mathbf{cof}(\mathcal{L}) \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
 & & \mathfrak{b} & \longleftarrow & \mathfrak{d} & & \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\mathbf{add}(\mathcal{L}) & \longleftarrow & \mathbf{add}(\mathcal{B}) & \longleftarrow & \mathbf{cov}(\mathcal{B}) & \longleftarrow & \mathbf{non}(\mathcal{L})
\end{array}$$

To conclude this section, we point out an elementary connection between the covering and uniformity numbers studied here and the splitting and refining numbers from Section 3.

5.19 Theorem $\mathfrak{s} \leq \mathbf{non}(\mathcal{B}), \mathbf{non}(\mathcal{L})$ *and* $\mathfrak{r} \geq \mathbf{cov}(\mathcal{B}), \mathbf{cov}(\mathcal{L})$.

Proof. For any infinite $A \subseteq \omega$, the sets $X \subseteq \omega$ that fail to split A form a meager, measure-zero set U_A. Then the function $A \mapsto U_A$ and the identity function on $\mathcal{P}(\omega)$ constitute a morphism from $\mathfrak{R}$ to $\mathbf{Cov}(\mathcal{B})$ and also to $\mathbf{Cov}(\mathcal{L})$. ⊣

6. Sparse Sets of Integers

This section is primarily about two characteristics, $\mathfrak{t}$ and $\mathfrak{h}$, related to the idea of thinning out infinite subsets of ω, i.e., replacing them by subsets, usually so as to achieve some useful property like homogeneity for some partition. $\mathfrak{t}$ is concerned with the (transfinite) thinning process itself; $\mathfrak{h}$ focuses on what can be achieved by iterated thinning. We shall also briefly consider two characteristics, $\mathfrak{p}$ and $\mathfrak{g}$, whose definitions resemble those of $\mathfrak{t}$ and $\mathfrak{h}$, though their most significant properties are treated only in later sections.

We begin with the definition and simplest properties of $\mathfrak{t}$.

6.1 Definition A *pseudointersection* of a family $\mathcal{F}$ of sets is an infinite set that is $\subseteq^*$ every member of $\mathcal{F}$.

6.2 Definition A *tower* is an ordinal-indexed sequence $\langle T_\alpha : \alpha < \lambda \rangle$ such that:

1. Each T_α is an infinite subset of ω.
2. $T_\beta \subseteq^* T_\alpha$ whenever $\alpha < \beta < \lambda$.
3. $\{T_\alpha : \alpha < \lambda\}$ has no pseudointersection.

The *tower number* $\mathfrak{t}$ is the smallest λ that is the length of a tower.

6.3 Remark Hechler [54] has constructed a model where many regular cardinals occur as the lengths of towers.

Some authors define "tower" using only the first two clauses in the definition above, i.e., an almost decreasing sequence in $[\omega]^\omega$; what we call a tower, they would call an inextendible tower. Also, some authors take towers to be almost increasing sequences of co-infinite subsets of ω rather than almost decreasing sequences of infinite sets.

We shall not always be as careful as we were in clause 3 of the definition about the distinction between a sequence like $\langle T_\alpha : \alpha < \lambda \rangle$ and the set $\{T_\alpha : \alpha < \lambda\}$ of its terms.

6.4 Proposition $\mathfrak{t}$ *is a regular uncountable cardinal.*

Proof. Regularity is clear since any cofinal subsequence of a tower is a tower. To show that there can be no tower $\langle T_n : n \in \omega \rangle$ of length ω, note that we could form an infinite set X by taking any element of T_0, any different element of $T_0 \cap T_1$, any different element of $T_0 \cap T_1 \cap T_2$, etc., since all these sets are infinite. This X would be a pseudointersection, violating requirement 3 in the definition of tower. ⊣

Before continuing with further properties of $\mathfrak{t}$, we introduce $\mathfrak{h}$, its basic properties, and its connection with $\mathfrak{t}$.

6.5 Definition A family $\mathcal{D} \subseteq [\omega]^\omega$ is *open* if it is closed under almost subsets. It is *dense* if every $X \in [\omega]^\omega$ has a subset in $\mathcal{D}$. The *distributivity number* $\mathfrak{h}$ is the smallest number of dense open families with empty intersection.

6.6 Remark The open sets as defined here constitute a topology on $[\omega]^\omega$, which we call the *lower topology*. Density as defined here agrees with topological density in the lower topology as long as $\mathcal{D}$ is closed under finite modifications (for example if it is open). Analogous definitions can be made for any pre-ordered set in place of $([\omega]^\omega, \subseteq^*)$.

The name "distributivity number" comes from viewing $([\omega]^\omega, \subseteq)$ as a notion of forcing and asking how distributive the associated complete Boolean algebra is. Standard techniques from forcing theory show that the answer is given by $\mathfrak{h}$. More precisely, Boolean meets of fewer than $\mathfrak{h}$ terms distribute over arbitrary (finite or infinite) joins, but meets of $\mathfrak{h}$ terms need not distribute even over binary joins. Equivalently, in a forcing extension by $([\omega]^\omega, \subseteq)$, $\mathfrak{h}$ has new subsets but smaller ordinals do not (not even new functions into the ordinals). We shall see later (6.20) that this forcing extension collapses $\mathfrak{c}$ to $\mathfrak{h}$ if $\mathfrak{h} < \mathfrak{c}$.

6.7 Proposition *The intersection of any fewer than $\mathfrak{h}$ dense open families is dense open. $\mathfrak{h}$ is a regular cardinal.*

Proof. The second sentence follows immediately from the first. (It also follows from the remark about distributivity.) To show that the intersection of fewer than $\mathfrak{h}$ dense open families $\mathcal{D}_\alpha$ is dense open, note first that it is obviously open. As for density, let X be any infinite subset of ω and consider the families $\mathcal{D}'_\alpha = \{Y \in \mathcal{D}_\alpha : Y \subseteq X\}$. These are fewer than $\mathfrak{h}$ dense open families of subsets of X, so they have a common member Y. That is, $Y \subseteq X$ and $Y \in \bigcap_\alpha \mathcal{D}_\alpha$. ⊣

The definition of $\mathfrak{t}$ is essentially about the process of thinning out infinite subsets of ω by repeatedly passing to (almost) subsets. If one attempts to iterate such a thinning process transfinitely, the definition of $\mathfrak{t}$ ensures that one will not get stuck at limit stages of cofinality $< \mathfrak{t}$.

The definition of $\mathfrak{h}$ addresses the same idea from the point of view of what such thinning can achieve. A dense open family is one that one can get into, from an arbitrary infinite subset of ω, by passing to a subset (and subsequent passages to further (almost) subsets will not undo this achievement). The next proposition is just the result of comparing these intuitions that stand behind $\mathfrak{t}$ and $\mathfrak{h}$.

6.8 Proposition $\mathfrak{t} \leq \mathfrak{h}$.

Proof. Suppose $\kappa < \mathfrak{t}$, and let κ dense open families $\mathcal{D}_\alpha$ $(\alpha < \kappa)$ be given; we must find a set in their intersection. Define an almost decreasing sequence $\langle T_\alpha : \alpha \leq \kappa \rangle$ by the following recursion. $T_0 = \omega$. $T_{\alpha+1}$ is any subset of T_α that is in $\mathcal{D}_\alpha$; this exists because $\mathcal{D}_\alpha$ is dense. If $\lambda \leq \kappa$ is a limit ordinal, then T_λ is any pseudointersection of $\{T_\alpha : \alpha < \lambda\}$; this exists because $\kappa < \mathfrak{t}$ so $\{T_\alpha : \alpha < \lambda\}$ cannot be a tower, yet the previous steps ensured that it satisfies the first two requirements for a tower. Since $T_\kappa \subseteq^* T_{\alpha+1}$ for all $\alpha < \kappa$, we have, thanks to openness, that T_κ is in all the families $\mathcal{D}_\alpha$. ⊣

It is consistent with ZFC to have $\mathfrak{t} < \mathfrak{h}$. In fact, Dordal [40] built a model where $\mathfrak{h} = \aleph_2 = \mathfrak{c}$ but there are no towers of length ω_2.

Upper bounds for $\mathfrak{h}$, and therefore also for $\mathfrak{t}$, can be obtained by considering specific examples of dense open families. One family of examples consists of $\{X \in [\omega]^\omega : X \text{ is not split by } Y\}$ for arbitrary Y. Another consists of $\{X \in [\omega]^\omega : \forall^\infty x \in X\, \forall y \in X \text{ (if } x < y \text{ then } f(x) < y)\}$ for arbitrary $f : \omega \to \omega$. Using these, one easily obtains the following proposition, but we give another proof to suggest another class of examples.

6.9 Theorem $\mathfrak{h} \leq \mathfrak{b}, \mathfrak{s}$.

Proof. By Theorem 3.5 it suffices to show $\mathfrak{h} \leq \mathfrak{par}_2$. So let $\kappa < \mathfrak{h}$ partitions f_α of $[\omega]^2$ be given; we must find an infinite set almost homogeneous for them all. For each α, let $\mathcal{D}_\alpha$ be the family of all infinite subsets of ω that are almost homogeneous for f_α. Then $\mathcal{D}_\alpha$ is dense open, thanks to Ramsey's theorem. So there is a set H common to all the $\mathcal{D}_\alpha$. ⊣

6.10 Remark The same proof shows that one can get simultaneous almost homogeneity for fewer than $\mathfrak{h}$ partitions of more complicated sorts, provided one has the analog of Ramsey's theorem to ensure density. Thus, for example, Silver's partition theorem for analytic sets [101] implies that any $< \mathfrak{h}$ partitions of $[\omega]^\omega$ into an analytic and a coanalytic piece have a common infinite almost homogeneous set.

By Proposition 6.8, the upper bounds on $\mathfrak{h}$ apply also to $\mathfrak{t}$, but for $\mathfrak{t}$ we can improve $\mathfrak{b}$ to $\mathbf{add}(\mathcal{B})$. In order to prove this, we need the following lemma, in which $\mathbb{Q}$ denotes the set of rational numbers and "dense" has its usual topological (or order-theoretic) meaning for subsets of $\mathbb{Q}$. Both the lemma and the subsequent theorem are from [92] (stated for special cases but the proofs work in general); they were rediscovered in [85].

6.11 Lemma *Suppose* $\lambda < \mathfrak{t}$ *and* $\langle T_\alpha : \alpha < \lambda \rangle$ *is an almost decreasing sequence of dense subsets of* $\mathbb{Q}$. *Then there exists a dense* $X \subset \mathbb{Q}$ *that is almost included in every* T_α.

Proof. In each interval I with rational endpoints, consider the almost decreasing sequence $\langle T_\alpha \cap I : \alpha < \lambda \rangle$ of infinite subsets of I. As it is too short to be a tower, there is an infinite $Y_I \subset I$ almost included in all the T_α. (The union of all the Y_I is dense, but it need not be $\subseteq^* T_\alpha$, so we must work a bit harder to get X.) Enumerate each Y_I as an ω-sequence $\langle y_{I,n} \rangle$. For each α let $f_\alpha(I) \in \omega$ be an upper bound for the finitely many n such that $y_{I,n} \notin T_\alpha$. Since $\lambda < \mathfrak{t} \leq \mathfrak{b}$ (and the set of intervals I is countable), let g be a function to ω from the set of rational intervals such that g dominates all the f_α's. Then

$$X = \bigcup_I \{y_{I,n} : n > g(I)\}$$

is dense in $\mathbb{Q}$ (because it contains almost all of each Y_I) and is almost included in each T_α (for $X - T_\alpha$ consists of finitely many elements from each of the finitely many Y_I where $g(I) < f_\alpha(I)$). $\dashv$

6.12 Theorem $\mathfrak{t} \leq \mathbf{add}(\mathcal{B})$.

Proof. We must show that if $\kappa < \mathfrak{t}$ then the intersection of any κ dense open subsets G_α ($\alpha < \kappa$) of $\mathbb{R}$ is comeager. We begin by defining an almost decreasing sequence $\langle T_\alpha : \alpha \leq \kappa \rangle$ of dense subsets of $\mathbb{Q}$. Start with $T_0 = \mathbb{Q}$. At limit stages, apply the lemma. At successor stages, set $T_{\alpha+1} = T_\alpha \cap G_\alpha$; this is dense because it is the intersection of two dense sets one of which is open. Note that T_κ, being $\subseteq^*$ each $T_{\alpha+1}$ ($\alpha < \kappa$) is also $\subseteq^*$ each G_α.

For $t \in T_\kappa$ and $\alpha < \kappa$, define $f_\alpha(t) \in \omega$ to be some n such that $(t - \frac{1}{n}, t + \frac{1}{n}) \subseteq G_\alpha$ if $t \in G_\alpha$, and 0 otherwise. Since T_κ is countable and $\kappa < \mathfrak{t} \leq \mathfrak{b}$, there is a $g : T_\kappa \to \omega$ dominating all the f_α's.

For each finite $F \subseteq T_\kappa$, let

$$U_F = \bigcup_{t \in T_\kappa - F} (t - \frac{1}{g(n)}, t + \frac{1}{g(n)}).$$

Then U_F is dense, because it almost includes T_κ, and it is obviously open; since there are only countably many F's, $\bigcap_F U_F$ is comeager, and it remains only to prove that this intersection is included in the intersection of the G_α's. In fact, each G_α includes one of the U_F's; given α just take F to contain the finitely many $t \in T_\kappa - G_\alpha$ and the finitely many t where $g(t) < f_\alpha(t)$. $\dashv$

6.13 Remark By a countable support iteration of Mathias forcing over a model of CH, one obtains a model where $\mathfrak{h} = \aleph_2$ but $\mathbf{cov}(\mathcal{B})$ and therefore $\mathbf{add}(\mathcal{B})$ are only $\aleph_1$ (as no Cohen reals are produced). Thus, the preceding theorem cannot be improved by putting $\mathfrak{h}$ in place of $\mathfrak{t}$.

The next theorem can be viewed as another upper bound on $\mathfrak{t}$.

6.14 Theorem *If* $\aleph_0 \leq \kappa < \mathfrak{t}$ *then* $2^\kappa = \mathfrak{c}$.

Proof. We need only check that $2^\kappa \leq \mathfrak{c}$, and we do this by building a complete binary tree of $\kappa + 1$ levels, whose nodes are distinct subsets of ω. More precisely, we associate to every sequence $\eta \in {}^{\leq\kappa}2$ an infinite subset T_η of ω in such a way that:

1. If η is an initial segment of θ, then $T_\theta \subseteq^* T_\eta$.
2. If neither of η and θ is an initial segment of the other, then $T_\eta \cap T_\theta$ is finite.

The construction is by recursion on the length of η, starting with $T_\emptyset = \omega$. At successor stages, we define $T_{\eta^\frown\langle 0\rangle}$ and $T_{\eta^\frown\langle 1\rangle}$ to be any two disjoint, infinite subsets of T_η. Finally, for θ of limit length λ, we observe that $\langle T_{\theta\restriction\alpha} : \alpha < \lambda\rangle$ is an almost decreasing sequence but cannot be a tower because $\lambda \leq \kappa < \mathfrak{t}$. So there is an infinite X almost included in all these $T_{\theta\restriction\alpha}$; any such X can serve as T_θ.

It is immediate that the construction has the desired properties. In particular, the 2^κ sets T_η, for all η of length κ, are infinite and almost disjoint and therefore certainly distinct. ⊣

6.15 Corollary $\mathfrak{t} \leq cof(\mathfrak{c})$.

Proof. If $\kappa < \mathfrak{t}$ then, by Theorem 6.14 and König's theorem, $\mathrm{cof}(\mathfrak{c}) = \mathrm{cof}(2^\kappa) > \kappa$. ⊣

Returning to consider $\mathfrak{h}$ in more detail, we first give an alternative way to view dense open families of subsets of ω.

6.16 Definition An *almost disjoint* family is a family of infinite sets, every two of which have finite intersection. A *maximal almost disjoint* (MAD) family is an infinite almost disjoint family of subsets of ω, maximal with respect to inclusion.

6.17 Remark Note that MAD families are required to be infinite; in the absence of this requirement, any partition of ω into finitely many infinite sets would count as MAD. Note also that, if $\mathcal{A}$ is MAD and X is any infinite subset of ω, then $X \cap A$ is infinite for at least one $A \in \mathcal{A}$.

6.18 Proposition *If $\mathcal{A}$ is a MAD family, then $\mathcal{A}{\downarrow} = \{X \in [\omega]^\omega : \exists A \in \mathcal{A}\,(X \subseteq^* A)\}$ is dense open. Every dense open family includes one of this form.*

Proof. The first statement is proved by routine checking of definitions. For the second, let $\mathcal{D}$ be dense open, and let $\mathcal{A}_0$ be an infinite, almost disjoint subfamily of $\mathcal{D}$; for example, take some $X \in \mathcal{D}$ and partition it into infinitely many infinite pieces. By Zorn's Lemma, let $\mathcal{A} \supseteq \mathcal{A}_0$ be an almost disjoint family included in $\mathcal{D}$ and maximal among such families. We claim that $\mathcal{A}$ is maximal among all almost disjoint families, not just those included in $\mathcal{D}$. Once we establish this claim, we will have $\mathcal{A}$ MAD and $\mathcal{A}{\downarrow} \subseteq \mathcal{D}$ as required.

To establish maximality, consider any $X \in [\omega]^\omega$. As $\mathcal{D}$ is dense, it contains a subset Y of X. As $\mathcal{A}$ is maximal among almost disjoint subfamilies of $\mathcal{D}$, it contains a set A that has infinite intersection with Y and therefore also with X. ⊣

6.19 Corollary $\mathfrak{h}$ *is the minimum number of MAD families such that, for each* $X \in [\omega]^\omega$*, one of these families contains at least two sets whose intersections with* X *are infinite.*

Proof. X has infinite intersection with at least two sets from a MAD family $\mathcal{A}$ if and only if $X \notin \mathcal{A}{\downarrow}$. With this observation, the corollary follows immediately from the proposition and the definition of $\mathfrak{h}$. $\dashv$

The following theorem of Balcar, Pelant, and Simon [1] was the original motivation for the introduction of $\mathfrak{h}$. A tree of the sort described by this theorem is called a *base matrix tree* (for $[\omega]^\omega$). The theorem would become false if $\mathfrak{h}$ were replaced by any smaller cardinal. The symbol $\mathfrak{h}$ was chosen to refer to the "height" of the base matrix tree.

6.20 Theorem *There is a family* $\mathcal{T} \subseteq [\omega]^\omega$ *with the following properties.*

1. *Ordered by reverse almost inclusion,* $\mathcal{T}$ *is a tree of height* $\mathfrak{h}$ *with root* ω*.*
2. *Each level of* $\mathcal{T}$*, except for the root, is a MAD family.*
3. *Every* $X \in [\omega]^\omega$ *has a subset in* $\mathcal{T}$*.*

Proof. Let $\mathcal{D}_\alpha$ for $\alpha < \mathfrak{h}$ be dense open families with no common member. We define the levels $\mathcal{T}_\alpha$ of the desired tree inductively as follows. At level 0, put ω. At a limit level $\lambda < \mathfrak{h}$, use Proposition 6.7 to obtain a dense open family included in all $\mathcal{T}_\alpha{\downarrow}$ for $\alpha < \lambda$. By Proposition 6.18, shrink this to a dense open family of the form $\mathcal{A}{\downarrow}$, and let that $\mathcal{A}$ be $\mathcal{T}_\lambda$.

At an odd-numbered successor stage, say $2\alpha+1$, choose $\mathcal{T}_{2\alpha+1}$ as a MAD family included in both $\mathcal{T}_{2\alpha}{\downarrow}$ and $\mathcal{D}_\alpha$. This can be done by Propositions 6.7 and 6.18.

At an even-numbered successor stage, say $2\alpha + 2$, proceed as follows. Call a set $X \in [\omega]^\omega$ *active* at this stage if it has infinite intersection with $\mathfrak{c}$ members of $\mathcal{T}_{2\alpha+1}$. Assign to each active X some $\psi(X) \in \mathcal{T}_{2\alpha+1}$ that has infinite intersection with X, and do this in such a way that ψ is one-to-one. This is easily done by arranging all the active X's in a well-ordered sequence, of length $\leq \mathfrak{c}$, and defining ψ by recursion along this ordering. At each stage of the recursion, there are $\mathfrak{c}$ elements of $\mathcal{T}_{2\alpha+1}$ that have infinite intersection with the current X and fewer than $\mathfrak{c}$ of them have already been assigned as earlier values of ψ, so there are plenty of candidates left to serve as $\psi(X)$. Once ψ has been defined, partition each $Y \in \mathcal{T}_{2\alpha+1}$ into two infinite pieces Y' and Y'', subject to the requirement that if $Y = \psi(X)$ for some (unique) X then $Y' \subseteq X$. Then let $\mathcal{T}_{2\alpha+2}$ consist of these sets Y' and Y'' for all $Y \in \mathcal{T}_{2\alpha+1}$.

This completes the construction of $\mathcal{T}$. The first two parts of the theorem are clear, and the third will be clear once we show that every $X \in [\omega]^\omega$ is

active at some stage $2\alpha + 2$. To this end, we consider a fixed X and we build a binary subtree of $\mathcal{T}$, of height ω, as follows. Its root is the root ω of $\mathcal{T}$. After its n^{th} level has been constructed, consisting of 2^n nodes Z of $\mathcal{T}$, all at the same level of $\mathcal{T}$, say level α_n, and all having infinite intersection with X, we produce the next level as follows. For each node Z of level n in our subtree, $Z \cap X$ is an infinite set and cannot be in all the dense open families $\mathcal{D}_\xi$ as these were chosen to have no common member. Since $\mathcal{T}_{2\xi+1} \subseteq \mathcal{D}_\xi$, $Z \cap X$ must meet at least two sets in $\mathcal{T}_\beta$ for all sufficiently large β. Choose a β that is sufficiently large in this sense for all 2^n Z's; call it α_{n+1}, and let the successors of each Z at this level be two nodes that meet $Z \cap X$ infinitely. Note that these are necessarily $\subseteq^* Z$ (for otherwise they would be almost disjoint from Z), so we are getting a subtree of $\mathcal{T}$. After the subtree has been constructed, use the fact that $\mathfrak{h}$ is uncountable and regular to see that the supremum, say γ, of the α_n's is still $< \mathfrak{h}$, so there are a γ^{th} and a $(\gamma+1)^{\text{st}}$ level of $\mathcal{T}$.

For each path p through our subtree, the nodes along that path, intersected with X, form an almost-decreasing ω-sequence, so there is an infinite $X' \subseteq X$ almost included in all of them (as $\mathfrak{t} > \omega$). That X' has infinite intersection with some node Y in level $\gamma + 1$ of $\mathcal{T}$, because the level is a MAD family. This Y has infinite intersection with each of the nodes Z along the path p, so Y is almost included in each of these Z's (because $\mathcal{T}$ is a tree). Since distinct nodes at the same level are almost disjoint, distinct paths p must lead to distinct nodes Y. So we have $\mathfrak{c}$ nodes Y at level $\gamma+1$, all meeting X infinitely. Since γ is a limit ordinal, $\gamma = 2\gamma$ and X is active at stage $2\gamma + 2$. ⊣

6.21 Remark Clause 3 of the theorem implies that forcing with $([\omega]^\omega, \subseteq^*)$ is equivalent to forcing with $(\mathcal{T}, \subseteq^*)$. It is not difficult to modify the construction of the base matrix tree so that each node has $\mathfrak{c}$ immediate successors. Then this forcing clearly adjoins a function from $\mathfrak{h}$ onto $\mathfrak{c}$. Since $\mathfrak{h}$ is not collapsed and no reals are added (because of the distributivity), we find that $\mathfrak{h}$ is the cardinality of the continuum in the forcing extension by $([\omega]^\omega, \subseteq^*)$.

We introduce a cardinal $\mathfrak{p}$, a slight modification of $\mathfrak{t}$, that is often useful because of its connection with forcing; see Theorem 7.12 below. Notice that it makes sense to ask about pseudointersections of families more general than towers. An obvious necessary condition for a family to have a pseudointersection is the strong finite intersection property defined below; $\mathfrak{p}$ measures the extent to which this necessary condition is also sufficient.

6.22 Definition A family $\mathcal{F}$ of infinite sets has the *strong finite intersection property* (SFIP) if every finite subfamily has infinite intersection. The

pseudointersection number $\mathfrak{p}$ is the smallest cardinality of any $\mathcal{F} \subseteq [\omega]^\omega$ with SFIP but with no pseudointersection.

Since a tower is a family with SFIP and no pseudointersection, we immediately get half of the following proposition. The other half, that $\mathfrak{p}$ is uncountable, is proved exactly as for $\mathfrak{t}$ (and is improved in Proposition 6.24 below).

6.23 Proposition $\aleph_1 \leq \mathfrak{p} \leq \mathfrak{t}$.

It is not known whether $\mathfrak{p}$ can be strictly smaller than $\mathfrak{t}$, but the next theorem shows that, for this to happen, $\mathfrak{p}$ would have to be at least $\aleph_2$ and (therefore) $\mathfrak{t}$ at least $\aleph_3$. To prove the theorem, we need a proposition that will be useful again later (in the proof of Theorems 8.13 and 9.25) and is of some interest in its own right as it can serve as a characterization of $\mathfrak{d}$ (the $\mathfrak{d}$ in the hypothesis is easily seen to be optimal). The theorem and a version of the proposition are in [91]; a form of the proposition closer to the present one is in [65].

6.24 Proposition *Suppose $\langle C_n : n \in \omega \rangle$ is a decreasing (or almost decreasing) sequence of infinite subsets of ω, and suppose $\mathcal{A}$ is a family of fewer than $\mathfrak{d}$ subsets of ω such that each set in $\mathcal{A}$ has infinite intersection with each C_n. Then $\{C_n : n \in \omega\}$ has a pseudointersection B that has infinite intersection with every set in $\mathcal{A}$.*

Proof. We may assume $\langle C_n : n \in \omega \rangle$ is decreasing, for if it is only almost decreasing then we can replace each C_n with $\bigcap_{k \leq n} C_k$ without affecting the other hypotheses or the conclusion, as each new C_n differs only finitely from the old.

For any $h \in {}^\omega\omega$, let $B_h = \bigcup_{n \in \omega}(C_n \cap h(n))$. Each C_n includes all but the first n terms of this union, so B_h is a pseudointersection of the C_n's. It remains to choose h so that $A \cap B_h$ is infinite for all $A \in \mathcal{A}$.

For each such A, let $f_A(n)$ denote the nth element of the infinite set $A \cap C_n$. Observe that, if $h(n) > f_A(n)$ for some A and n, then $A \cap B_h$ has cardinality at least n, for it contains the first n elements of $A \cap C_n$. So B_h can serve as the B in the proposition provided $\forall A \in \mathcal{A}\, \exists^\infty n\, (h(n) > f_A(n))$. But there are fewer than $\mathfrak{d}$ functions f_A, so there is an h not dominated by any of them. $\dashv$

6.25 Theorem *If $\mathfrak{p} = \aleph_1$ then $\mathfrak{t} = \aleph_1$.*

Proof. Since $\mathfrak{t} \leq \mathfrak{h} \leq \mathfrak{b} \leq \mathfrak{d}$, the result is immediate if $\mathfrak{d} = \aleph_1$. So we assume for the rest of the proof that $\mathfrak{d} > \aleph_1$.

By hypothesis we have a family $\mathcal{A} = \{A_\alpha : \alpha < \aleph_1\}$ with SFIP but with no pseudointersection, and we may assume that it is closed under finite

intersections. We build a tower $\langle T_\alpha : \alpha < \aleph_1 \rangle$ of length $\aleph_1$ by recursion, ensuring at each stage that T_α has infinite intersection with each $A \in \mathcal{A}$ and that $T_{\alpha+1} \subseteq A_\alpha$. We start with $T_0 = \omega$, and at countable limit stages λ we continue the tower by applying the proposition (with $\langle C_n : n \in \omega \rangle$ being a cofinal subsequence of $\langle T_\alpha : \alpha < \lambda \rangle$). At successor stages we set $T_{\alpha+1} = T_\alpha \cap A_\alpha$. It is easy to verify that this defines an almost decreasing sequence with the claimed properties. It is a tower, because any pseudointersection of the T_α's would also be a pseudointersection of the A_α's. $\dashv$

We close this section by discussing the groupwise density number $\mathfrak{g}$, a close relative of $\mathfrak{h}$. More information about it, including the motivation for its definition, is in Section 8.

6.26 Definition A family $\mathcal{G} \subseteq [\omega]^\omega$ is *groupwise dense* if it is open in the lower topology (i.e., closed under almost subsets) and, for every interval partition Π, some union of (infinitely many) intervals of Π belongs to $\mathcal{G}$. The *groupwise density number* $\mathfrak{g}$ is the smallest number of groupwise dense families with empty intersection.

It is conventional, though perhaps unnatural, to include closure under almost subsets in the definition of "groupwise dense" even though it is not in the definition of "dense." The first part of the following proposition gives a convenient synonym for "groupwise dense," namely "nonmeager open" where non-meager refers to the usual topology of $[\omega]^\omega$ (as a subspace of $\mathcal{P}\omega \cong {}^\omega 2$) whereas "open" refers to the lower topology.

6.27 Proposition *1. A family $\mathcal{G} \subseteq [\omega]^\omega$ is groupwise dense if and only if it is closed under almost subsets and nonmeager in the standard topology.*

2. The intersection of any fewer than $\mathfrak{g}$ groupwise dense families is groupwise dense.

3. $\mathfrak{g}$ is regular.

4. $\mathfrak{h} \leq \mathfrak{g} \leq \mathfrak{d}$.

Proof. Identify $[\omega]^\omega$ with a cocountable subset of ${}^\omega 2$ via characteristic functions. Let $\mathcal{G} \subseteq [\omega]^\omega$ be closed under almost subsets. By Theorem 5.2, it is nonmeager if and only if it contains enough reals to match each chopped real (x, Π). Thanks to closure under subsets, it suffices to match those chopped reals whose first component x is the identically 1 function. But matching these chopped reals is precisely what the definition of groupwise density requires.

For part 2, suppose $\mathcal{G}_\alpha$ are fewer than $\mathfrak{g}$ groupwise dense families. Their intersection $\mathcal{G}$ is clearly closed under almost subsets, so consider an arbitrary

interval partition $\Pi = \{I_n : n \in \omega\}$. We must find an infinite set $X \subseteq \omega$ such that $\bigcup_{n\in X} I_n \in \mathcal{G}$. That is, we must find an X common to the families $\mathcal{H}_\alpha = \{X \in [\omega]^\omega : \bigcup_{n\in X} I_n \in \mathcal{G}_\alpha\}$. Since there are fewer than $\mathfrak{g}$ of these families, it suffices to prove that each of them is groupwise dense. This is a routine verification; to see that $\mathcal{H}_\alpha$ contains a union of intervals from a partition Θ, use that $\mathcal{G}_\alpha$ contains a union of intervals from the partition $\{\bigcup_{n\in J} I_n : J \in \Theta\}$.

The regularity of $\mathfrak{g}$ follows immediately from part 2.

Every groupwise dense family $\mathcal{G}$ is dense. Indeed, given any infinite $X \subset \omega$, we can form an interval partition in which each interval contains a member of X. Then $\mathcal{G}$ contains a union of such intervals and therefore, being closed under subsets, contains an infinite subset of X. This observation immediately gives $\mathfrak{h} \leq \mathfrak{g}$.

Finally, to show that $\mathfrak{g} \leq \mathfrak{d}$, let $\mathcal{D}$ be a dominating family of size $\mathfrak{d}$, and associate to each $f \in \mathcal{D}$ the set

$$\mathcal{G}_f = \{X \in [\omega]^\omega : \exists^\infty n\, (X \cap [n, f(n)) = \emptyset)\}.$$

Then $\mathcal{G}_f$ is groupwise dense; given any interval partition Π, we can form an element of $\mathcal{G}_f$ by taking infinitely many intervals from Π, separated by gaps so long that each gap includes $[n, f(n))$ for some n. But there can be no X common to all the $\mathcal{G}_f$, for if there were then the function sending each $n \in \omega$ to the next larger member of X would not be dominated by any $f \in \mathcal{D}$. So we have $\mathfrak{d}$ groupwise dense families with empty intersection. $\dashv$

6.28 Remark This proposition shows that, in the lattice of subsets of $[\omega]^\omega$ closed downward with respect to $\subseteq^*$, the non-meager sets form a filter, indeed a $< \mathfrak{g}$-complete filter. This may be somewhat surprising, since in the lattice (Boolean algebra) of all subsets of $[\omega]^\omega$, two non-meager sets can be disjoint; in fact there are $\mathfrak{c}$ pairwise disjoint non-meager sets.

6.29 Remark The characteristics studied in this section, as well as $\mathfrak{s}$ and $\mathfrak{r}$ from Section 3 above and $\mathfrak{a}$ from Section 8 below, have interesting analogs in structures other than $[\omega]^\omega$. One example is the system of dense subsets of $\mathbb{Q}$, ordered by $\subseteq^*$. Little is known about these characteristics, but Lemma 6.11 says that the tower number in this situation is no smaller than the ordinary tower number.

Another example is the set of partitions of ω into infinitely many pieces, ordered by "coarser than modulo finite." Several characteristics of this sort have been studied by Krawczyk (unpublished).

7. Forcing Axioms

Forcing axioms are combinatorial statements designed to express what is achieved by certain sorts of iterated forcing constructions. They serve to

hide such constructions in a "black box"; instead of showing that a statement of interest can be forced by such a construction, one derives it from the combinatorial principle. The oldest and still the most frequently used of these principles is Martin's axiom, introduced in [73]. To state it, we need some terminology from forcing theory.

7.1 Definition Let $(P, \leq)$ be a nonempty partially ordered set. Two elements $p, q \in P$ are *compatible* if they have a common lower bound and *incompatible* otherwise. An *antichain* is a set of pairwise incompatible elements. P satisfies the *countable chain condition* (ccc) or *countable antichain condition* if all its antichains are countable. More generally, P satisfies the $< \kappa$-chain condition if all its antichains have cardinalities $< \kappa$.

A subset $D \subseteq P$ is *dense* if every element of P is $\geq$ an element of D. If $\mathcal{D}$ is a family of dense subsets of P, then $G \subseteq P$ is *$\mathcal{D}$-generic* if it is closed upward and directed downward (every two members have a common lower bound) and intersects every $D \in \mathcal{D}$.

7.2 Definition *Martin's axiom* (MA) is the statement that, if $\mathcal{D}$ is a family of fewer than $\mathfrak{c}$ dense subsets of a partial order P with ccc, then there is a $\mathcal{D}$-generic filter $G \subseteq P$. More generally, if κ is a cardinal and $\mathcal{K}$ is a class of nonempty partial orders, then we write $\mathrm{MA}_\kappa(\mathcal{K})$ for the statement that every family $\mathcal{D}$ of κ dense subsets in a member P of $\mathcal{K}$ admits a $\mathcal{D}$-generic $G \subseteq P$. $\mathrm{MA}_{<\kappa}(\mathcal{K})$ is defined analogously. One omits the subscript when it is "$< \mathfrak{c}$" and one omits the class $\mathcal{K}$ when it is the class of ccc posets.

Thus, MA is $\mathrm{MA}_{<\mathfrak{c}}(\mathrm{ccc})$. Some authors write MA_κ to mean what we have called $\mathrm{MA}_{<\kappa}$.

MA describes the model obtained by a finite support forcing iteration, of length some uncountable $\kappa = \kappa^{<\kappa}$, in which all ccc posets (of the extension) of size $< \kappa$ are used as forcing conditions during the iteration. This iteration, which is itself a ccc forcing, produces a model of MA and $\mathfrak{c} = \kappa$. Thus, MA is consistent with the continuum being arbitrarily large. Although only small (smaller than κ) posets were used during the iteration, a reflection argument (essentially the Löwenheim-Skolem theorem; see the second preliminary simplification in the proof of Theorem 7.12 below) shows that all ccc posets, not only the small ones, acquire generic sets with respect to small families of dense subsets. For details about this, see [105] or [67, Section VIII.6].

For orientation, we mention that:

- MA_ω(all posets) is provable in ZFC, and therefore CH implies MA.
- $\mathrm{MA}_{\aleph_1}$(all posets) is refutable. Take P to be ${}^{<\omega}\aleph_1$ ordered by reverse inclusion, take $D_\alpha = \{p \in P : \alpha \in \mathrm{ran}(p)\}$, and observe that a generic G would give a map $\bigcup G$ of ω onto $\aleph_1$.

- $\mathrm{MA}_\mathfrak{c}$(Cohen) is refutable, where "Cohen" refers to the single poset ${}^{<\omega}2$ ordered by reverse inclusion. For each $f : \omega \to 2$, let $D_f = \{p : p \not\subseteq f\}$ and observe that a generic G would give a function $\bigcup G : \omega \to 2$ different from every f.

The last two of these observations indicate why MA refers only to ccc posets and only to $< \mathfrak{c}$ dense sets.

The effect of MA on cardinal characteristics of the continuum is to make them large, as the next two theorems and their corollaries show. These results are from [73].

7.3 Theorem *MA implies* $\mathbf{add}(\mathcal{L}) = \mathfrak{c}$.

Proof. Suppose $\kappa < \mathfrak{c}$ and we are given κ sets $N_\alpha \subseteq \mathbb{R}$ $(\alpha < \kappa)$ of measure zero. We must show, assuming MA, that their union has measure 0. It suffices to find, for each positive ε, a set of measure $\leq \varepsilon$ that includes all the N_α as subsets.

Given ε, let P be the set of open subsets of $\mathbb{R}$ having measure $< \varepsilon$, and order P by reverse inclusion. In order to apply MA to this P, we first verify the ccc. Let uncountably many elements p of P be given. Inside each of these open sets, find a finite union $q(p)$ of open intervals with rational endpoints, large enough so that $\mu(p - q(p)) < \varepsilon - \mu(p)$. Notice that this implies $\mu(p - q(p)) < \frac{1}{2}(\varepsilon - \mu(q(p)))$. There are only countably many possibilities for $q(p)$, so two (in fact uncountably many) of the $q(p)$ must be the same q. But then the union of the two corresponding p's has measure $< \varepsilon$ (because it consists of q plus the two remainders $p - q$, and each remainder has measure less than half of $\varepsilon - \mu(q)$), so it is in P and is a common lower bound for those two p's. Thus, an uncountable family of p's cannot be an antichain.

For each of the given N_α's, let $D_\alpha = \{p \in P : N_\alpha \subseteq p\}$, and notice that this is a dense subset of P (because a set of measure zero is included in open sets of arbitrarily small measure). Since $\kappa < \mathfrak{c}$, MA provides a generic G meeting all the D_α. Then $\bigcup G$ includes all the N_α. Furthermore, as a directed union of open sets of measure $< \varepsilon$, this $\bigcup G$ has measure $\leq \varepsilon$. $\dashv$

7.4 Corollary *MA implies that all the cardinals in Cichoń's diagram are equal to* $\mathfrak{c}$ *and that* $\mathfrak{r} = \mathfrak{c}$.

7.5 Remark The partial ordering used in the proof of the theorem is called the *amoeba* order. To understand the name, visualize the open sets in three dimensional space instead of $\mathbb{R}$, and visualize the proof of density of D_α as extruding a tentacle[1] from a given open set to engulf N_α.

The proof of ccc for the amoeba actually establishes the stronger property of being σ-linked in the sense of the following definition.

[1] It has been pointed out to me that an amoeba has pseudopodia, not tentacles. But it seems easier to visualize tentacles.

7.6 Definition In a partial order, a subset S is called *linked* if every two of its elements are compatible. It is *n-linked* if every n of its members have a common lower bound. It is *centered* if every finitely many of its members have a common lower bound. *σ-linked* means the union of countably many linked subsets. *σ-n-linked* and *σ-centered* are defined analogously.

Clearly, σ-centered implies σ-linked, which in turn implies ccc. As n increases, σ-n-linked becomes stronger but still remains weaker than σ-centered.

In the proof of Theorem 7.3, we essentially showed that the amoeba is σ-linked, as witnessed by the countably many sets $\{p : q(p) = q\}$, where q ranges over finite unions of rational intervals and where $q(p)$ is defined for all p as it was defined in the proof above for p in the supposed antichain. A similar argument shows that the amoeba is σ-n-linked for all n. But it is not σ-centered.

Instead of working directly with sets of measure zero, one can prove the preceding theorem by using Theorem 5.14, which described $\mathbf{add}(\mathcal{L})$ in terms of slaloms. Given fewer than $\mathfrak{c}$ functions $\omega \to \omega$, one needs a slalom through which all of them go. This is obtainable by applying MA to a poset P consisting of pieces of slaloms. Specifically, a member p of P is a function on ω assigning to each n a finite set of natural numbers, such that, for some k, $|p(n)|$ is n for $n < k$ and k for $n \geq k$. The ordering is componentwise reverse inclusion, and the relevant dense sets are $\{p : \forall n > k\, (f(n) \in p(n))\}$, where k witnesses that $p \in P$ and where f is one of the given functions that should go through our slalom. This forcing is called localization forcing in [5, Section 3.1].

7.7 Theorem *MA implies* $\mathfrak{p} = \mathfrak{c}$.

Proof. Suppose $\mathcal{F} \subseteq [\omega]^\omega$ has the strong finite intersection property and $|\mathcal{F}| < \mathfrak{c}$. To find a pseudointersection X for $\mathcal{F}$, we apply MA to the following poset P. A member of P is a pair (s, F) where s is a finite subset of ω and F is a finite subset of $\mathcal{F}$. (The "meaning" of (s, F) is that the desired X should include s and should, except for s, be included in each $A \in F$.) The ordering puts $(s', F') \leq (s, F)$ if

$$s \text{ is an initial segment of } s', \quad F' \supseteq F, \quad \text{and} \quad \forall A \in F\, (s' - s \subseteq A).$$

Any two pairs with the same first component are compatible, as one can just take the union of the second components. (In fact, any finitely many pairs with the same first component have a common lower bound. So this ordering is σ-centered.) For each $A \in \mathcal{F}$, the set $D_A = \{(s, F) \in P : A \in F\}$ is dense. So is $D_n = \{(s, F) \in P : |s| > n\}$ because of the SFIP of $\mathcal{F}$. MA provides a generic G meeting all these dense sets. Let $X = \bigcup_{(s,F)\in G} s$. This is infinite because G meets each D_n. To see that it is almost included in

each $A \in \mathcal{F}$, use that G and D_A have a common member (s_0, F_0). That means $A \in F_0$, and we shall show that $X - s_0 \subseteq A$. Any member k of $X - s_0$ is in $s - s_0$ for some (s, F) in G, and, as G is directed downward, it contains some $(s', F') \leq$ both (s, F) and (s_0, F_0). Then $k \in s - s_0 \subseteq s' - s_0 \subseteq A$, as required. $\dashv$

7.8 Remark The forcing used in the preceding proof is called *Mathias forcing with respect to* $\mathcal{F}$. One can equivalently view it as consisting of pairs (s, A) where A is the intersection of finitely many sets from $\mathcal{F}$; in this form, the ordering $(s', A') \leq (s, A)$ is defined by

$$s \text{ is an initial segment of } s', \quad A' \subseteq A, \quad \text{and} \quad s' - s \subseteq A.$$

Mathias forcing (without respect to any $\mathcal{F}$) means the similarly defined poset where the second components A can be arbitrary infinite subsets of ω. In contrast to Mathias forcing with respect to an $\mathcal{F}$ with SFIP, this Mathias forcing does not satisfy the ccc. It can be viewed as a two-step forcing iteration, where the first step is forcing with $([\omega]^\omega, \subseteq^*)$, which adjoins a generic ultrafilter $\mathcal{U}$ on ω, and the second step is Mathias forcing with respect to $\mathcal{U}$.

7.9 Corollary *MA implies* $\mathfrak{p} = \mathfrak{t} = \mathfrak{h} = \mathfrak{g} = \mathfrak{s} = \mathfrak{c}$.

Thus, all the characteristics we have discussed are equal to $\mathfrak{c}$ if MA holds. The proofs actually show a bit more, if we introduce new characteristics related directly to MA.

7.10 Definition For any class $\mathcal{K}$ of posets, $\mathfrak{m}(\mathcal{K})$ is the smallest κ for which $\mathrm{MA}_\kappa(\mathcal{K})$ is false. If $\mathcal{K}$ is the class of ccc posets, we omit mention of it and write simply $\mathfrak{m}$.

Thus MA is the statement $\mathfrak{m} = \mathfrak{c}$. Clearly,

$$\mathfrak{m} \leq \mathfrak{m}(\sigma\text{-linked}) \leq \mathfrak{m}(\sigma\text{-3-linked}) \leq \cdots \leq \mathfrak{m}(\sigma\text{-centered}) \leq \mathfrak{m}(\text{Cohen}).$$

See [68] for a model where $\mathfrak{m} < \mathfrak{m}(\sigma\text{-linked})$; similar techniques can be used to get strict inequalities between other such variants of $\mathfrak{m}$.

The proofs of the last two theorems and our remarks about the σ-linked and σ-centered properties of the posets in the proofs establish the following inequalities.

7.11 Corollary $\mathfrak{m}(\sigma\text{-}linked) \leq \mathbf{add}(\mathcal{L})$ *and* $\mathfrak{m}(\sigma\text{-}centered) \leq \mathfrak{p}$.

Of course, one could be even more specific about the posets used; for example the proof above shows that $\mathfrak{m}(\text{amoeba}) \leq \mathbf{add}(\mathcal{L})$. In fact, equality holds here; see [5, Theorem 3.4.17].

The second half of the last corollary can also be improved to an equality, Bell's theorem [12].

7.12 Theorem $\mathfrak{m}(\sigma\text{-centered}) = \mathfrak{p}$.

Proof. In view of Corollary 7.11, it suffices to consider an arbitrary σ-centered poset P, say the union of centered parts C_n, and to find a $\mathcal{D}$-generic G for a prescribed family $\mathcal{D}$ of fewer than $\mathfrak{p}$ dense subsets of P. It is convenient to begin with several simplifications of the problem.

First, we may assume that each $D \in \mathcal{D}$ is closed downward, because closing the dense sets will not affect $\mathcal{D}$-genericity.

Second, we may assume that $|P| < \mathfrak{p}$. Indeed, suppose the theorem were proved in this case, and suppose we are given the situation above with $|P| \geq \mathfrak{p}$. By the Löwenheim-Skolem theorem, the structure $(P, \leq, C_n, D)_{n\in\omega, D\in\mathcal{D}}$ has an elementary substructure P' of cardinality $< \mathfrak{p}$. Then P' is σ-centered and $\mathcal{D}' = \{P' \cap D : D \in \mathcal{D}\}$ is a family of $< \mathfrak{p}$ dense subsets, so there is a $\mathcal{D}'$-generic $G \subseteq P'$. The upward closure of G' in P is then $\mathcal{D}$-generic, as required.

Third, instead of producing a generic G, it suffices to produce a linked L meeting every $D \in \mathcal{D}$. Indeed, suppose we could always do this. Then, given P and $\mathcal{D}$ as above, we enlarge $\mathcal{D}$ by adjoining the sets

$$D_{p,q} = \{r \in P : r \text{ is incompatible with } p \text{ or with } q, \text{ or } r \leq p, q\},$$

which are easily seen to be dense. If L is linked and meets all the sets in $\mathcal{D}$ and all the $D_{p,q}$, then the upward closure G of L is $\mathcal{D}$-generic. The only thing to check is that it is directed downward. To find a common lower bound for any $p, q \in G$, we may, by lowering both, assume that p and q are in L. Let $r \in L \cap D_{p,q}$. Then r cannot be incompatible with p or with q as L is linked; so $r \leq p, q$, as required.

Fourth, we may assume that, for each $n \in \omega$, there is some $D_n \in \mathcal{D}$ disjoint from C_n. Otherwise, C_n could serve as the required L.

Fifth, we may assume that $\mathcal{D}$ is closed under finite intersections. Closing it in this way does no harm, because the cardinality $|\mathcal{D}|$ will not be increased (unless it was finite — a trivial case) and the intersection of any finitely many dense, downward-closed sets is again dense and downward closed.

After all these simplifications, we begin the real work of the proof. For each $p \in P$ and each $D \in \mathcal{D}$, let $A(p, D)$ be the set of those $n \in \omega$ such that some member of $C_n \cap D$ is $\leq p$.

I claim that, for each $k \in \omega$, the family $\mathcal{F}_k = \{A(p, D) : p \in C_k \text{ and } D \in \mathcal{D}\}$ has the strong finite intersection property. By our fourth simplification, it suffices to show that each finite subfamily $\mathcal{F}_k^0$ of $\mathcal{F}_k$ has nonempty intersection, for we could always include in $\mathcal{F}_k^0$ sets of the form $A(p, D_n)$ for any finitely many of the D_n and so keep any finitely many n's out of the intersection. By our fifth simplification, we may assume that the sets in $\mathcal{F}_k^0$ are $A(p_i, D)$ for various $p_i \in C_k$ but just one $D \in \mathcal{D}$, for different D's could be replaced with their intersection. As C_k is centered, the p_i's have a lower

bound p, and below that we can find a member q of the dense set D. If $q \in C_n$ then $n \in \bigcap \mathcal{F}_k^0$. This completes the verification of the claim.

Since $|\mathcal{F}_k| \leq |P| \cdot |\mathcal{D}| < \mathfrak{p}$ by our second simplification, $\mathcal{F}_k$ has a pseudointersection A_k.

Next, we define several labelings of the ω-branching tree ${}^{<\omega}\omega$ of height ω, a primary labeling by natural numbers and, for each $D \in \mathcal{D}$, a secondary labeling by members of P. In the primary labeling, the label of the root is (arbitrarily chosen as) 0, and if a node has label k then the labels of its immediate successors are the numbers in A_k (once each). The secondary labeling associated to a particular $D \in \mathcal{D}$ is defined as follows. The secondary label of the root is an arbitrary element of C_0. If a node has been given secondary label p and if an immediate successor of it has primary label n, then the secondary label of that successor is to be an element of $C_n \cap D$ that is $\leq p$ in P, provided such an element exists, i.e., provided $n \in A(p, D)$ — in this case we call that successor node "good" for D. If no such element exists, then the secondary label of that successor node is chosen arbitrarily from C_n and the node is called "bad" for D. Notice that, whether a node is good or bad, its secondary label is always in C_n where n is its primary label.

Because A_k is a pseudointersection for $\mathcal{F}_k$, all but finitely many of the immediate successors of any node are good for any particular $D \in \mathcal{D}$. For each node s and each $D \in \mathcal{D}$, let $f_D(s)$ be a number so large that all the nodes $s^\frown\langle m\rangle$ for $m \geq f_D(s)$ are good for D. Since $|\mathcal{D}| < \mathfrak{p} \leq \mathfrak{b}$ (and since the tree has only countably many nodes), there is $g : {}^{<\omega}\omega \to \omega$ that is $>^*$ all the f_D.

Using g, we define a path X through the tree ${}^{<\omega}\omega$ by starting at the root and, after reaching a node s, proceeding to $s^\frown\langle g(s)\rangle$. Our choice of g ensures that, for each $D \in \mathcal{D}$, all but finitely many nodes along the path X are good for D. Choose, for each D, a node s_D on X such that it and all later nodes on X are good for D, and let p_D be its secondary label associated to D. Thus $p_D \in D$. This guarantees that $L = \{p_D : D \in \mathcal{D}\}$ meets every $D \in \mathcal{D}$.

To complete the proof, we verify that L is linked. Consider any two elements $p_D, p_{D'} \in L$. If $s_D = s_{D'}$ then both of $p_D, p_{D'}$ are in the same C_n, where n is the primary label of s_D, so they are compatible because C_n is centered. Suppose therefore that s_D occurs before $s_{D'}$ along the path X. By choice of s_D, all the nodes along the path X from s_D to $s_{D'}$ are good for D, so the secondary labeling associated to D puts at the node $s_{D'}$ a label q that is $\leq p_D$. But, being in the same C_n, q and $p_{D'}$ are compatible. Therefore so are p_D and $p_{D'}$. ⊣

There is a similar (but easier) result about countable partial orders.

7.13 Theorem $\mathfrak{m}(\mathit{Cohen}) = \mathfrak{m}(\mathit{countable}) = \mathbf{cov}(\mathcal{B})$.

Proof. Since Cohen forcing is a countable poset, MA_κ(countable) implies MA_κ(Cohen). We shall complete the proof by showing that MA_κ(Cohen) implies $\kappa < \mathbf{cov}(\mathcal{B})$ and that this in turn implies MA_κ(countable).

Assume MA_κ(Cohen) and let κ nowhere dense subsets X_α of ${}^\omega 2$ be given. We must show that the X_α do not cover ${}^\omega 2$. For each α, let D_α be the set of those $s \in {}^{<\omega}2$ that have no extensions in X_α. Because X_α is nowhere dense (in the topological sense), D_α is dense (in the partial order sense). So MA_κ(Cohen) gives us a generic $G \subseteq {}^\omega 2$ meeting every D_α. Then $\bigcup G \in {}^\omega 2$ is in none of the X_α.

Finally, assume $\kappa < \mathbf{cov}(\mathcal{B})$, and let κ dense subsets D_α of a countable poset P be given. Let $T : {}^\omega P \to {}^\omega P$ be the transformation that turns any sequence $x \in {}^\omega P$ into a (weakly) decreasing sequence $T(x)$ in a greedy way; that is, $T(x)(0) = x(0)$, and $T(x)(n+1) = x(n+1)$ if this is $\leq T(x)(n)$ in P, and otherwise $T(x)(n+1) = T(x)(n)$. We similarly define T on finite sequences instead of infinite ones.

The sets $U_\alpha = \{x \in {}^\omega P : \exists n\, (T(x)(n) \in D_\alpha)\}$ are dense open subsets of ${}^\omega P$. To verify density, consider any nonempty $s \in {}^{<\omega}P$, let p be the last term of $T(s)$, and let $q \leq p$ be in D_α. Then every extension of $s^\frown\langle q\rangle$ is in U_α.

As $\kappa < \mathbf{cov}(\mathcal{B})$, there is an x in the intersection of all the U_α. Then the range of the decreasing sequence $T(x)$ meets every D_α, and the upward closure of this range is therefore the desired generic set. ⊣

As an application of Bell's theorem 7.12, we give an analog of Proposition 6.24, weakening the hypothesis of countability (of the list of C's) and strengthening the hypothesis of cardinality $< \mathfrak{d}$ (for $\mathcal{A}$) by replacing both with the hypothesis of cardinality $< \mathfrak{p}$.

7.14 Theorem *Suppose $\mathcal{C}$ and $\mathcal{A}$ are families of $< \mathfrak{p}$ subsets of ω, and suppose every intersection of finitely many sets from $\mathcal{C}$ and one set from $\mathcal{A}$ is infinite. Then $\mathcal{C}$ has a pseudointersection B that has infinite intersection with each set in $\mathcal{A}$.*

Proof. Let P be Mathias forcing with respect to $\mathcal{C}$, as defined in the proof of Theorem 7.7 and the remark following it. As shown there, this is σ-centered, and for each $C \in \mathcal{C}$ the set $D_C = \{(s,F) \in P : C \in F\}$ is dense. Furthermore, for each $A \in \mathcal{A}$ and each $n \in \omega$, the set $D_{A,n} = \{(s,F) \in P : |s \cap A| > n\}$ is dense because each intersection of finitely many sets from $\mathcal{C}$ and one set from $\mathcal{A}$ is infinite.

As both $|\mathcal{C}|$ and $|\mathcal{A}|$ are $< \mathfrak{p} = \mathfrak{m}(\sigma\text{-centered})$, P has a generic subset G meeting all these D_C and $D_{A,n}$. As in the proof of Theorem 7.7, we define $B = \bigcup_{(s,F)\in G} s$ and we find that this is a pseudointersection of $\mathcal{C}$. Furthermore, it has infinite intersection with each $A \in \mathcal{A}$ because G meets

each $D_{A,n}$. ⊣

As a consequence, we obtain that $\mathfrak{p}$, like its relatives $\mathfrak{t}$, $\mathfrak{h}$, and $\mathfrak{g}$, is regular, but the proof is trickier than for the relatives. This proof is taken from [47, Section 21], where it is attributed to Szymański.

7.15 Theorem $\mathfrak{p}$ *is regular.*

Proof. Suppose $\mathfrak{p}$ were singular with cofinality $\lambda < \mathfrak{p}$. Let $\mathcal{A}$ be a family of $\mathfrak{p}$ subsets of ω having the strong finite intersection property but having no pseudointersection. Express $\mathcal{A}$ as the union of an increasing λ-sequence of subfamilies $\mathcal{A}_\alpha$, each of cardinality $< \mathfrak{p}$. To simplify later considerations, we assume without loss of generality that $\mathcal{A}$ and all the $\mathcal{A}_\alpha$ are closed under finite intersections.

In this situation, we have the following improvement of Theorem 7.14. If $\mathcal{C}$ is any family of fewer than $\mathfrak{p}$ sets such that every intersection of finitely many sets from $\mathcal{C}$ and one set from $\mathcal{A}$ is infinite, then $\mathcal{C}$ has a pseudointersection B whose intersection with each set from $\mathcal{A}$ is infinite. (The improvement is that $\mathcal{A}$ has size $\mathfrak{p}$ rather than strictly smaller size.) To prove this, note first that each $\mathcal{C} \cup \mathcal{A}_\alpha$ has the SFIP and has size $< \mathfrak{p}$, so it has a pseudointersection Z_α. Then apply Theorem 7.14 with $\{Z_\alpha : \alpha < \lambda\}$ in the role of $\mathcal{A}$.

We intend to build an almost decreasing $\lambda + 1$-sequence $\langle C_\alpha : \alpha \leq \lambda \rangle$ such that each C_α for $\alpha < \lambda$ is a pseudointersection of $\mathcal{A}_\alpha$. If we can do this then, because the C_α are almost decreasing and the $\mathcal{A}_\alpha$ are increasing and cover $\mathcal{A}$, C_λ will be a pseudointersection of $\mathcal{A}$, a contradiction.

We define the C_α by recursion. To make the recursion work, we carry along the additional requirement that each C_α must have infinite intersection with every member of $\mathcal{A}$.

Suppose $\alpha \leq \lambda$ and C_β is already defined for all $\beta < \alpha$ in such a way that our requirements are satisfied. We need to define C_α so that it is $\subseteq^*$ each previous C_β, it is $\subseteq^*$ each member of $\mathcal{A}_\alpha$, and it has infinite intersection with every member of $\mathcal{A}$. Such a set is produced by applying the improvement above of Theorem 7.14 with $\mathcal{C} = \{C_\beta : \beta < \alpha\} \cup \mathcal{A}_\alpha$, provided the hypothesis of that improvement is satisfied. So we need only check that every intersection X of finitely many C_β's ($\beta < \alpha$) and finitely many members of $\mathcal{A}_\alpha$ and one member of $\mathcal{A}$ is infinite. Since the C_β's are almost decreasing, since $\mathcal{A}_\alpha \subseteq \mathcal{A}$, and since $\mathcal{A}$ is closed under finite intersection, such an X almost includes $C_\beta \cap A$ for some $\beta < \alpha$ and some $A \in \mathcal{A}$. The induction hypothesis guarantees that $C_\beta \cap A$ and therefore X are infinite. ⊣

7.16 Remark This section has dealt almost exclusively with forcing axioms for the class of ccc posets and subclasses, because these are the forcing

axioms most relevant to cardinal characteristics. To avoid giving a completely unbalanced picture, however, we should at least mention that numerous other forcing axioms have been considered. The most popular of these is the *proper forcing axiom* PFA, which is $\mathrm{MA}_{\aleph_1}(\text{proper})$. Proper forcing was defined by Shelah [96, Chapter III], who showed that it permits countable-support iterations without collapsing $\aleph_1$; see Abraham's chapter in this handbook. PFA summarizes the result of a countable-support iteration of all small proper posets. Unlike the construction of a model for MA, where the improvement from small ccc posets to all ccc posets was handled by a Löwenheim-Skolem argument, the construction of a model for PFA uses a supercompact cardinal in the ground model to get the necessary reflection property for the corresponding improvement from small to all.

8. Almost Disjoint and Independent Families

This section is devoted to families of subsets of ω with various special properties, and particularly to those families that are maximal with respect to these properties.

Recall from Section 6 that an almost disjoint family is a family of infinite sets whose pairwise intersections are finite, and that the phrase "maximal almost disjoint (MAD) family" refers to an *infinite* family of subsets of ω maximal with respect to almost disjointness.

Although a set of size κ clearly cannot support a family of more than κ disjoint sets, the situation for almost disjoint sets is quite different.

8.1 Proposition *On any countably infinite set, there is a family of* $\mathfrak{c}$ *almost disjoint subsets.*

Proof. It clearly does not matter which countably infinite set we consider. Choosing the binary tree ${}^{<\omega}2$ as the ambient set, we can use its $\mathfrak{c}$ branches as the almost disjoint family. ⊣

8.2 Remark There are at least two other similar and equally easy proofs of this proposition. One uses the set of rationals as the ambient set and assigns to every real r a sequence of rationals converging to r; sequences with different limits are clearly almost disjoint. Another uses $\omega \times \omega$ as the ambient set and assigns to each positive real r the set $\{(n, \lfloor rn \rfloor) : n \in \omega\}$.

The proposition and Zorn's lemma imply the existence of a MAD family of cardinality $\mathfrak{c}$, but there may also be smaller MAD families. For example, it is shown in [67, Theorem VIII.2.3] that if one adds any number of Cohen reals to a model of CH, then the resulting model has a MAD family of size $\aleph_1$; see also Section 11. Hechler [54] gives a model with MAD families of many different cardinalities.

8.3 Definition The *almost disjointness number* $\mathfrak{a}$ is the smallest cardinality of any MAD family.

8.4 Proposition $\mathfrak{b} \leq \mathfrak{a}$.

Proof. Let $\mathcal{A}$ be a MAD family of size $\mathfrak{a}$, let C_n $(n \in \omega)$ be any countably many members of it, and let $\mathcal{A}'$ be the rest of $\mathcal{A}$. By making finite changes to each C_n, we can arrange that these sets are really disjoint, not just almost disjoint, and that they partition ω. By a suitable bijection, identify ω with $\omega \times \omega$ in such a way that C_n is the column $\{n\} \times \omega$. Each $A \in \mathcal{A}'$ has only finitely many elements per column, so we can define $f_A : \omega \to \omega$ to be the function whose graph is the upper boundary of A. If there were a function $g : \omega \to \omega$ that is $>^*$ all the f_A, then its graph would be almost disjoint from all $A \in \mathcal{A}'$ and all C_n, contrary to maximality of $\mathcal{A}$. So the f_A's constitute an unbounded family of size $\mathfrak{a}$. $\dashv$

Shelah [98] showed that $\mathfrak{b} < \mathfrak{a}$ is consistent. He also showed there that, if we define $\mathfrak{a}_s$ like $\mathfrak{a}$ except that we use $\omega \times \omega$ as the ambient set and require the MAD family to consist of graphs of partial functions, then $\mathfrak{a} < \mathfrak{a}_s$ is consistent. Brendle has pointed out the following alternative proof of the consistency of $\mathfrak{a} < \mathfrak{a}_s$. By part 2 of Theorem 5.9, we have $\mathbf{non}(\mathcal{B}) \leq \mathfrak{a}_s$. We shall see in Section 11 that the random real model (obtained by forcing with a large measure algebra over a model of GCH) has $\mathbf{non}(\mathcal{B}) = \mathfrak{c} > \aleph_1$ and $\mathfrak{a} = \aleph_1$. Therefore it has $\mathfrak{a} < \mathfrak{a}_s$.

Little else is known about connections between $\mathfrak{a}$ and other cardinal characteristics, but Shelah has shown in [99] that $\mathfrak{a} > \mathfrak{d}$ is consistent.

8.5 Remark Proposition 8.1 can be used to evaluate the "dual" of $\mathfrak{h}$. Unlike the definitions of $\mathfrak{t}$ and $\mathfrak{p}$, that of $\mathfrak{h}$ fits the "norm of relation" format discussed in Section 4. Indeed, $\mathfrak{h} = \|([\omega]^\omega, DO, \notin)\|$ where DO is the family of dense open subsets of $[\omega]^\omega$. (There is an important difference between this relation and those associated to cardinal characteristics in Sections 4 and 5. The elements of DO cannot be coded by reals, nor does DO possess a nice base whose elements can be coded by reals.) It is natural to ask about the norm of the dual relation, i.e., the minimum size of a family $\mathcal{X} \subseteq [\omega]^\omega$ such that every dense open family $\mathcal{D}$ intersects $\mathcal{X}$. It follows from Proposition 8.1 that this cardinal is $\mathfrak{c}$. In fact, the same also holds for the dual of $\mathfrak{g}$, by nearly the same proof.

8.6 Theorem *The minimum number of sets in* $[\omega]^\omega$ *meeting every dense open family, or even every groupwise dense family, is* $\mathfrak{c}$.

Proof. Suppose $\mathcal{X} \subseteq [\omega]^\omega$ has cardinality $< \mathfrak{c}$; we shall find a groupwise dense (hence dense open) $\mathcal{D} \subseteq [\omega]^\omega$ disjoint from $\mathcal{X}$. Let $\mathcal{D} = \{Y \in [\omega]^\omega : \forall X \in \mathcal{X}\,(X \not\subseteq^* Y)\}$. This $\mathcal{D}$ is clearly disjoint from $\mathcal{X}$ and closed under almost

subsets, so we need only check that, for any interval partition $\{I_n : n \in \omega\}$, the union of some infinitely many of its intervals is in $\mathcal{D}$. Let $\mathcal{A}$ be a family of $\mathfrak{c}$ almost disjoint subsets of ω, and for each $A \in \mathcal{A}$ let $A' = \bigcup_{n\in A} I_n$. Then the A' are also almost disjoint, so no two of them can almost include the same $X \in \mathcal{X}$. Since there are more A''s than X's, some A' must not almost include any X, i.e., some A' must be in $\mathcal{D}$. ⊣

8.7 Corollary *cof*$(\mathfrak{c}) \geq \mathfrak{g}$.

Proof. Let $[\omega]^\omega = \bigcup_{\alpha<\mathrm{cof}(\mathfrak{c})} \mathcal{X}_\alpha$, where each $|\mathcal{X}_\alpha| < \mathfrak{c}$. By the theorem, there are groupwise dense families $\mathcal{D}_\alpha$ each disjoint from the corresponding $\mathcal{X}_\alpha$. No set can belong to all the $\mathcal{D}_\alpha$, for it would then belong to no $\mathcal{X}_\alpha$. So we have cof$(\mathfrak{c})$ groupwise dense families with empty intersection. ⊣

Notice that this corollary subsumes Corollary 6.15. The intermediate result that cof$(\mathfrak{c}) \geq \mathfrak{h}$ was proved in [1]. Among the familiar cardinal characteristics of the continuum, $\mathfrak{g}$ is the largest one known (to me) to be a lower bound for cof$(\mathfrak{c})$. In particular, it is consistent that $\mathfrak{b} > \mathrm{cof}(\mathfrak{c})$ and it is consistent that $\mathfrak{s} > \mathrm{cof}(\mathfrak{c})$. For the former, start with a model satisfying MA and $\mathfrak{c} = \aleph_2$ and GCH at all larger cardinals, and adjoin $\aleph_{\aleph_1}$ random reals. Then $\mathfrak{c} = \aleph_{\aleph_1}$ and $\mathfrak{b}$, unaffected by the random reals, is $\aleph_2 > \mathrm{cof}(\mathfrak{c})$. For the latter, start with a model of $\mathfrak{c} = \aleph_{\aleph_1}$, and do an $\aleph_2$-stage, finite-support iteration of Mathias forcing with respect to (arbitrarily chosen) ultrafilters. The finite-support iteration automatically adds Cohen reals at limit stages of cofinality ω and choosing one of them at each stage provides a splitting family of size $\aleph_2$. There is no smaller splitting family, because any $\aleph_1$ reals lie in an intermediate extension and fail to split the subsequently added Mathias reals.

8.8 Definition A family $\mathcal{I}$ of subsets of ω is *independent* if the intersection of any finitely many members of $\mathcal{I}$ and the complements of any finitely many other members of $\mathcal{I}$ is infinite.

The "infinite" at the end of the definition could be equivalently replaced with "nonempty" if we assumed that $\mathcal{I}$ is infinite.

8.9 Proposition *There is an independent family of cardinality* $\mathfrak{c}$.

Proof. Let C be the set of finite subsets of $\mathbb{Q}$. Since C is countably infinite, it suffices to find $\mathfrak{c}$ independent subsets of C. For each real r, let

$$E_r = \{F \in C : |F \cap (-\infty, r)| \text{ is even}\}.$$

To see that these sets E_r are independent, let any finitely many distinct reals $r_1, \ldots, r_k, s_1, \ldots, s_l$ be given. We must find an $F \in C$ that belongs

to all the E_{r_i} and none of the E_{s_j}. But this is easy; F consists of 0 or 1 rationals from each of the $(k+l+1)$ intervals into which the r's and s's partition $\mathbb{R}$, the choice of 0 or 1 being made so as to get the right parities. $\dashv$

8.10 Remark The preceding proposition is due to Fichtenholz and Kantorovich [45]. It was generalized by Hausdorff [53] who showed that any infinite cardinal κ has 2^κ independent subsets.

Hausdorff's construction (for $\kappa = \aleph_0$) uses the countable set C of pairs (a, B) where a ranges over finite subsets of ω and B ranges over subsets of $\mathcal{P}(a)$. To each $X \subseteq \omega$ associate the subset $\{(a, B) \in C : a \cap X \in B\}$ of C. It is easy to verify that all these subsets are independent.

The corresponding generalization of Proposition 8.1 fails. Baumgartner [8, Theorem 5.6] showed that $\aleph_1$ need not have $2^{\aleph_1}$ uncountable subsets with pairwise countable intersections.

The proposition and Zorn's lemma provide a maximal independent family of size $\mathfrak{c}$, but there may be smaller maximal independent families.

8.11 Definition The *independence number* $\mathfrak{i}$ is the smallest cardinality of any maximal independent family of subsets of ω.

No upper bounds (except for the trivial $\mathfrak{c}$) are known for $\mathfrak{i}$, but there are two lower bounds.

8.12 Proposition $\mathfrak{r} \leq \mathfrak{i}$

Proof. Let $\mathcal{I}$ be a maximal independent family, and let $\mathcal{R}$ consist of all the sets obtainable by intersecting finitely many sets from $\mathcal{I}$ and the complements of finitely many others. The definition of independence ensures that $\mathcal{R} \subseteq [\omega]^\omega$, and $\mathcal{R}$ must be unsplittable because if X were to split all its members then $\mathcal{I} \cup \{X\}$ would be independent, contrary to the maximality of $\mathcal{I}$. So $|\mathcal{R}| \geq \mathfrak{r}$, from which it follows that $|\mathcal{I}| \geq \mathfrak{r}$. $\dashv$

The following more difficult estimate of $\mathfrak{i}$ is due to Shelah [110, Appendix by Shelah]. The proof we give, a simplification of Shelah's, is from [20]; the simplification was found independently by Bill Weiss.

8.13 Theorem $\mathfrak{d} \leq \mathfrak{i}$.

Proof. Suppose $\mathcal{I}$ is an independent family of cardinality $< \mathfrak{d}$; we shall show that it is not maximal. Throughout the proof, we let $\mathcal{X}$ and $\mathcal{Y}$ stand for finite, disjoint subfamilies of $\mathcal{I}$; thus, the independence of $\mathcal{I}$ means that $\bigcap \mathcal{X} - \bigcup \mathcal{Y}$ is always infinite, and our goal is to find Z such that each $\bigcap \mathcal{X} - \bigcup \mathcal{Y}$ meets both Z and $\omega - Z$ in an infinite set.

Select any countably many sets $D_n \in \mathcal{I}$, and let $\mathcal{I}'$ be the rest of $\mathcal{I}$. Write D_n^0 for D_n and write D_n^1 for $\omega - D_n$. For each $x : \omega \to 2$, apply Proposition 6.24 with

$$C_n = \bigcap_{k<n} D_k^{x(k)}$$

and

$$\mathcal{A} = \{\bigcap \mathcal{X} - \bigcup \mathcal{Y} : \mathcal{X}, \mathcal{Y} \text{ finite disjoint subfamilies of } \mathcal{I}'\}.$$

Independence of $\mathcal{I}$ gives the hypotheses of the proposition. So we get $B_x \subseteq \omega$ with:

1. $B_x \subseteq^* \bigcap_{k<n} D_k^{x(k)}$ for all n.
2. B_x has infinite intersection with each $\bigcap \mathcal{X} - \bigcup \mathcal{Y} \in \mathcal{A}$.

It follows from (1) that the B_x's for distinct x are almost disjoint.

Fix two disjoint, countable, dense (in the usual topology) subsets Q and Q' of ${}^{\omega}2$. Removing finitely many elements from B_x for each $x \in Q \cup Q'$, we can arrange that these countably many B_x's are really disjoint, not just almost disjoint. Set

$$Z = \bigcup_{x \in Q} B_x \quad \text{and} \quad Z' = \bigcup_{x \in Q'} B_x.$$

So Z and Z' are disjoint. We shall show that Z has infinite intersection with every $\bigcap \mathcal{X} - \bigcup \mathcal{Y}$; the same argument applies to Z', so $\omega - Z$ will also have infinite intersection with every $\bigcap \mathcal{X} - \bigcup \mathcal{Y}$, and so the proof will be complete.

Let finite, disjoint $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{I}$ be given, and let $\mathcal{X}'$ and $\mathcal{Y}'$ be their intersections with $\mathcal{I}'$. Fix n so large that, if D_k is in $\mathcal{X}$ or $\mathcal{Y}$ then $k < n$. Using the density of Q, fix $x \in Q$ such that if D_k is in $\mathcal{X}$ or $\mathcal{Y}$ then $x(k)$ is 0 or 1, respectively. Thus,

$$\begin{aligned}
\bigcap \mathcal{X} - \bigcup \mathcal{Y} &= (\bigcap \mathcal{X}' - \bigcup \mathcal{Y}') \cap \bigcap_{k : D_k \in \mathcal{X} \cup \mathcal{Y}} D_k^{x(k)} \\
&\supseteq (\bigcap \mathcal{X}' - \bigcup \mathcal{Y}') \cap \bigcap_{k<n} D_k^{x(k)} \\
&\supseteq^* (\bigcap \mathcal{X}' - \bigcup \mathcal{Y}') \cap B_x.
\end{aligned}$$

The last intersection here is infinite, by property (2) of B_x. It is included in Z because $x \in Q$. So we have an infinite set almost included in $Z \cap (\bigcap \mathcal{X} - \bigcup \mathcal{Y})$, and the proof is complete. $\dashv$

9. Filters and Ultrafilters

This section is devoted to filters and ultrafilters on ω. We begin by summarizing the terminology we use. Note that we require all filters to contain the cofinite sets, so all our ultrafilters are non-trivial.

9.1 Definition A *filter* (on ω) is a family $\mathcal{F} \subseteq \mathcal{P}\omega$ that contains all cofinite sets but not the empty set, is closed under supersets, and is closed under finite intersections. An *ultrafilter* (on ω) is a filter with the additional property that, for any $X \subseteq \omega$, either X or its complement belongs to $\mathcal{F}$. A *base* for a filter $\mathcal{F}$ is a subfamily of $\mathcal{F}$ containing subsets of all the sets in $\mathcal{F}$.

We occasionally stretch the meaning of "base" of $\mathcal{F}$ to apply to a family $\mathcal{B}$ such that for every $F \in \mathcal{F}$ there is $B \in \mathcal{B}$ with $B \subseteq^* F$ (rather than $B \subseteq F$). This stretching will make no essential difference but will simplify a few statements.

We shall need the following well-known consequences of the definition. A subset $\mathcal{X}$ of $\mathcal{P}\omega$ is included in a filter if and only if it has the strong finite intersection property, and then the smallest filter including $\mathcal{X}$ consists of the almost supersets of intersections of finite subfamilies of $\mathcal{X}$. We say that $\mathcal{X}$ generates this filter.

An ultrafilter is the same thing as a maximal filter; thus by Zorn's lemma every family with SFIP is included in an ultrafilter. Since an ultrafilter contains a set $X \subseteq \omega$ if and only if it does not contain $\omega - X$, it follows that a family $\mathcal{Y} \subseteq \mathcal{P}\omega$ is disjoint from some ultrafilter if and only if no finitely many members of $\mathcal{Y}$ almost cover ω.

9.2 Definition Let $\mathcal{F}$ be a subset of $\mathcal{P}\omega$ (usually a filter, but the definition makes sense in general) and let $f : \omega \to \omega$. Then $f(\mathcal{F})$ is defined to be $\{X \subseteq \omega : f^{-1}(X) \in \mathcal{F}\}$.

If $\mathcal{F}$ is a filter or an ultrafilter, then so is $f(\mathcal{F})$ provided it contains all cofinite sets. This proviso is automatically satisfied if f is finite-to-one, which will usually be the case in what follows.

9.3 Definition A filter $\mathcal{F}$ is *feeble* if, for some finite-to-one $f : \omega \to \omega$, $f(\mathcal{F})$ consists of only the cofinite sets.

The cofinite sets constitute the smallest filter, so feeble filters should also be thought of as small. They are at the opposite extreme from ultrafilters.

9.4 Proposition *The following are equivalent for any filter $\mathcal{F}$ on ω.*

1. *$\mathcal{F}$ is feeble.*

2. *There is a partition of ω into finite sets such that every set in $\mathcal{F}$ intersects all but finitely many pieces of the partition.*

3. *There is an interval partition as in 2 above.*

4. *$\{\omega - X : X \in \mathcal{F}\}$ is not groupwise dense.*

5. *$\mathcal{F}$ is meager (in the usual topology on $\mathcal{P}\omega \cong {}^{\omega}2$).*

Proof. The equivalence of statements 1 and 2 is immediate if one views the pieces of a partition (as in 2) as the sets on which a finite-to-one function (as in 1) is constant. If a partition Π is as in 2, then we can find an interval partition Θ, each of whose intervals includes at least one piece of Π; then Θ works in 3. The equivalence of 3 and 4 is just the definition of groupwise density. Finally, the equivalence of 4 and 5 follows from Proposition 6.27 because complementation ($X \mapsto \omega - X$) is a homeomorphism from $\mathcal{P}\omega$ to itself and thus preserves meagerness. ⊣

We next consider how many sets are needed to generate a large filter, where "large" can have a strong interpretation — ultrafilter — or a weak one — non-feeble filter. The former gives a new cardinal characteristic, while the latter gives a new view of an old characteristic. Notice that any infinite generating set for a filter yields a base of the same cardinality just by closing under finite intersections. So we can equivalently ask about cardinalities of bases.

We begin with a result similar to Propositions 8.1 and 8.9, namely that ultrafilter bases can be large. Of course, any ultrafilter is a base for itself and has cardinality $\mathfrak{c}$; the following proposition, due to Pospíšil [87], shows that for some ultrafilters there are no smaller bases.

9.5 Proposition *There is an ultrafilter on ω every base of which has cardinality $\mathfrak{c}$.*

Proof. Let $\mathcal{I}$ be an independent family of size $\mathfrak{c}$, by Proposition 8.9, and let $\mathcal{X}$ consist of

- all sets in $\mathcal{I}$ and
- the complements of all sets of the form $\bigcap \mathcal{C}$ with $\mathcal{C}$ an infinite subset of $\mathcal{I}$.

The independence of $\mathcal{I}$ easily implies that $\mathcal{X}$ has the SFIP, so there is an ultrafilter $\mathcal{U} \supseteq \mathcal{X}$. Suppose, toward a contradiction, that $\mathcal{U}$ had a base $\mathcal{Y}$ of cardinality $< \mathfrak{c}$. As each set in $\mathcal{I}$ has a subset in $\mathcal{Y}$ and $|\mathcal{I}| > |\mathcal{Y}|$, there must be infinitely many sets in $\mathcal{I}$ all including the same $Y \in \mathcal{Y}$. Then the intersection of these infinitely many sets from $\mathcal{I}$ is in $\mathcal{U}$ (because it includes

Y), but its complement is in $\mathcal{X}$ and thus also in $\mathcal{U}$. This contradiction completes the proof. $\dashv$

Nevertheless, it is consistent that some ultrafilters have bases of cardinality smaller than $\mathfrak{c}$.

9.6 Definition $\mathfrak{u}$, sometimes called the *ultrafilter number*, is the minimum cardinality of any ultrafilter base.

Kunen [67, Chapter 8, Ex. A10] built a finite-support iterated forcing model where $\mathfrak{c} = \aleph_{\aleph_1}$ but $\mathfrak{u} = \aleph_1$. Baumgartner and Laver [11] showed that an $\aleph_2$-step, countable-support iteration of Sacks forcing (over a model of GCH) produces a model where certain ultrafilters in the ground model (the selective ones) generate ultrafilters in the extension. Thus, their model has $\mathfrak{u} = \aleph_1$ while $\mathfrak{c} = \aleph_2$.

An ultrafilter base is an unsplittable family, for if X were to split it then neither X nor $\omega - X$ could be in the ultrafilter it generates. Thus, we immediately have the following inequality.

9.7 Proposition $\mathfrak{r} \leq \mathfrak{u}$.

In most known models, $\mathfrak{r} = \mathfrak{u}$, but Goldstern and Shelah [51] showed that the inequality can be strict. A stronger connection between $\mathfrak{r}$ and ultrafilters is given by the following result of Balcar and Simon [2].

9.8 Definition A *pseudobase* or *π-base* for a filter $\mathcal{F}$ on ω is a family $\mathcal{X} \subseteq [\omega]^\omega$ such that every set in $\mathcal{F}$ has a subset in $\mathcal{X}$.

This differs from the notion of base only in that $\mathcal{X}$ need not be a subfamily of $\mathcal{F}$.

9.9 Proposition $\mathfrak{r}$ *is the minimum cardinality of any ultrafilter pseudobase.*

Proof. A family $\mathcal{X} \subseteq [\omega]^\omega$ is an ultrafilter pseudobase if and only if there is an ultrafilter disjoint from

$$\mathcal{Y} = \{Y \subseteq \omega : Y \text{ has no subset in } \mathcal{X}\}.$$

As mentioned above, this is equivalent to saying that ω is not almost covered by finitely many sets from $\mathcal{Y}$. Equivalently, whenever ω is partitioned into finitely many pieces, one of the pieces must have an almost subset in $\mathcal{X}$. This means that $\mathcal{X}$ must be unsplittable, 3-unsplittable (in the sense of Example 4.13), ..., n-unsplittable for all finite n. On the one hand, mere unsplittability requires $\mathcal{X}$ to have cardinality at least $\mathfrak{r}$. On the other hand we can, as in Example 4.13, produce an n-unsplittable family of size $\mathfrak{r}$ for

each n and then take the union of these families to obtain an $\mathcal{X}$ as above of size $\mathfrak{r}$. ⊣

We now consider what is needed to generate a non-feeble filter. The first part of the following theorem is essentially due to Solomon [102]; the second part is an unpublished result of Petr Simon.

9.10 Theorem *Every filter on ω generated by fewer than $\mathfrak{b}$ sets is feeble, but there is a non-feeble filter generated by $\mathfrak{b}$ sets.*

Proof. Consider first a filter generated by fewer than $\mathfrak{b}$ sets, and associate to each of these generators A an interval partition Π_A chosen so that each interval in the partition contains an element of A. By Theorem 2.10, there is a single interval partition dominating all these Π_A's. It clearly satisfies statement 3 in Proposition 9.4, so our filter is feeble.

To produce a non-feeble filter generated by $\mathfrak{b}$ sets, we distinguish two cases, according to whether $\mathfrak{b} = \mathfrak{d}$.

If $\mathfrak{b} = \mathfrak{d}$, invoke Theorem 2.10 to get a $\mathfrak{b}$-indexed family of interval partitions Π_α ($\alpha < \mathfrak{b}$) dominating all interval partitions. We build the desired filter and a generating family $\mathcal{X}$ for it by a recursion of length $\mathfrak{b}$, starting with the family of cofinite sets, and adding at most one new set to $\mathcal{X}$ at each stage. At stage α, see whether the filter $\mathcal{F}_\alpha$ generated by the sets already put into $\mathcal{X}$ contains a set disjoint from infinitely many intervals of Π_α. If so, do nothing at stage α. If not, put into $\mathcal{X}$ the union of the even-numbered intervals of Π_α, and note that the SFIP of $\mathcal{X}$ is preserved. In either case, our final filter will contain a set missing infinitely many intervals of Π_α. After all $\mathfrak{b}$ steps have been completed, we have a filter that is not feeble because any interval partition as in statement 3 of Proposition 9.4 could not be dominated by any Π_α.

There remains the case that $\mathfrak{b} < \mathfrak{d}$. Let $\mathcal{B}$ be an unbounded family of size $\mathfrak{b}$ in ${}^\omega\omega$; without loss of generality, assume it is closed under forming the pointwise maximum of two functions and assume each function $g \in \mathcal{B}$ is non-decreasing. Since $|\mathcal{B}| < \mathfrak{d}$, let $f \in {}^\omega\omega$ be non-decreasing and dominated by no member of $\mathcal{B}$. Thus, the sets

$$X_g = \{n \in \omega : g(n) < f(n)\} \qquad (g \in \mathcal{B})$$

are infinite. The family $\{X_g : g \in \mathcal{B}\}$ is closed under finite intersections (because $\mathcal{B}$ is closed under maxima), so it is a base for a filter $\mathcal{F}$. To complete the proof, we suppose that $\mathcal{F}$ is feeble and we deduce a contradiction.

Suppose therefore that $\{I_n : n \in \omega\}$ is an interval partition such that each set in $\mathcal{F}$ meets all but finitely many I_n's. Define $f' : \omega \to \omega$ by letting $f'(k)$ be the value of f at the right endpoint of the next I_n after the one containing k. Consider an arbitrary $g \in \mathcal{B}$ and a k so large that X_g, being in $\mathcal{F}$, meets the next interval I_n after the one containing k. Calling that

interval $[a,b]$ and letting c be in its intersection with X_g, we have, since f and g are non-decreasing,

$$g(k) \leq g(c) < f(c) \leq f(b) = f'(k).$$

Thus, $g <^* f'$; since g was an arbitrary element of $\mathcal{B}$, we have a contradiction to the fact that $\mathcal{B}$ is unbounded. ⊣

9.11 Remark The first part of the preceding proof actually shows that a filter with a pseudobase of size $< \mathfrak{b}$ must be feeble.

It is easy to see that every filter $\mathcal{F}$ is the intersection of some ultrafilters, in fact of at most $\mathfrak{c}$ ultrafilters. Indeed, for each $A \in \mathcal{P}\omega - \mathcal{F}$, the family $\mathcal{F} \cup \{\omega - A\}$ has the SFIP and is therefore included in an ultrafilter $\mathcal{U}_A$. The intersection of these $\mathcal{U}_A$'s is $\mathcal{F}$.

The next two propositions contain information about how many ultrafilters must be intersected in order to get filters that are small in one or another sense. The first one, due to Plewik [86], is another application of Proposition 8.1.

9.12 Proposition *The intersection of fewer than* $\mathfrak{c}$ *ultrafilters is not feeble.*

Proof. Suppose the feeble filter $\mathcal{F}$ is the intersection of ultrafilters $\mathcal{U}_\alpha$. Let f be a finite-to-one function such that $f(\mathcal{F})$ consists only of the cofinite sets. Let $\mathcal{A}$ be a family of $\mathfrak{c}$ almost disjoint subsets of ω. For each $A \in \mathcal{A}$, we have $\omega - A \notin f(\mathcal{F})$ (as A is infinite), so $\omega - f^{-1}(A) = f^{-1}(\omega - A) \notin \mathcal{F}$, so $\omega - f^{-1}(A) \notin \mathcal{U}_\alpha$ for at least one α, and so $f^{-1}(A) \in \mathcal{U}_\alpha$. But the sets $f^{-1}(A)$ are almost disjoint because f is finite-to-one. So no two can be in the same $\mathcal{U}_\alpha$. Therefore there must be at least as many $\mathcal{U}_\alpha$'s as there are A's, namely $\mathfrak{c}$. ⊣

9.13 Proposition *There are* $\mathfrak{d}$ *ultrafilters whose intersection is not sent to an ultrafilter by any finite-to-one function.*

Proof. By Theorem 2.10, choose a family of $\mathfrak{d}$ interval partitions dominating all interval partitions, and associate to each $\Pi = \{I_n : n \in \omega\}$ in this family two ultrafilters $\mathcal{U}_\Pi$ and $\mathcal{V}_\Pi$ such that one contains $A_\Pi = \bigcup_n I_{8n}$ and the other contains $B_\Pi = \bigcup_n I_{8n+4}$. We shall show that the $\mathfrak{d}$ ultrafilters $\mathcal{U}_\Pi$ and $\mathcal{V}_\Pi$ are as required.

Suppose, to the contrary, that their intersection $\mathcal{F}$ is mapped to an ultrafilter by a finite-to-one map f. Let Θ be an interval partition such that each of the finite fibers $f^{-1}(\{n\})$ is included in the union of two adjacent intervals of Θ. (Simply build Θ inductively so that the right end of each interval is greater than all elements of all fibers whose left ends were in the

previous interval.) Let Π be an interval partition in our originally chosen family that dominates Θ. Then each interval of Θ, except for finitely many, is included in the union of two consecutive intervals of Π. It follows that each fiber of f, except for finitely many, is covered by four consecutive intervals of Π and therefore cannot meet both A_Π and B_Π. So $f(A_\Pi)$ and $f(B_\Pi)$ are almost disjoint and $f(\mathcal{F})$, being an ultrafilter, must contain the complement of one of them, say $\omega - f(A_\Pi)$. But then this complement would be in $f(\mathcal{U}_\Pi)$, which is absurd as $A_\Pi \in \mathcal{U}_\Pi$. $\dashv$

We shall next present some consequences of the inequality $\mathfrak{u} < \mathfrak{g}$. This inequality was introduced in [23] (where $\mathfrak{g}$ was first defined) as a "black box" summary of the crucial properties of the models, due to Shelah [24, 25], in which every two ultrafilters have a common finite-to-one image. Since then, numerous additional consequences and reformulations of $\mathfrak{u} < \mathfrak{g}$ have been found, and we present some of them here.

9.14 Definition For any family $\mathcal{F} \subseteq [\omega]^\omega$, we write $\sim\mathcal{F}$ for its complement and $\mathcal{F}\sim$ for the family of complements of its members.

$$\sim\mathcal{F} = [\omega]^\omega - \mathcal{F} \quad \text{and} \quad \mathcal{F}\sim = \{\omega - X : X \in \mathcal{F}\}.$$

We write $\breve{\mathcal{F}}$ for the dual family $\sim\mathcal{F}\sim = \{X \in [\omega]^\omega : \omega - X \notin \mathcal{F}\}$.

If $\mathcal{F}$ is closed under supersets then $\breve{\mathcal{F}}$ consists of just those $X \in [\omega]^\omega$ that intersect every member of $\mathcal{F}$. If $\mathcal{F}$ is a filter then $\mathcal{F} \subseteq \breve{\mathcal{F}}$, with equality holding exactly when $\mathcal{F}$ is an ultrafilter.

9.15 Lemma *Suppose that $\mathcal{X}, \mathcal{Y} \subseteq [\omega]^\omega$, that $|\mathcal{X}| < \mathfrak{g}$, and that $\mathcal{Y}\sim$ is groupwise dense. Then there is a finite-to-one $f : \omega \to \omega$ such that*

$$\forall X \in \mathcal{X}\, \exists Y \in \mathcal{Y}\, (f(Y) \subseteq f(X)).$$

Proof. For each $X \in \mathcal{X}$ define

$$\mathcal{G}_X = \{Z \in [\omega]^\omega : \exists Y \in \mathcal{Y}\, \forall a, b \in Z\, (\text{if } [a,b) \cap Y \neq \emptyset \text{ then } [a,b) \cap X \neq \emptyset)\}.$$

We verify that $\mathcal{G}_X$ is groupwise dense. $\mathcal{G}_X$ is clearly closed under subsets, and it is closed under finite modifications because $\mathcal{Y}$ is. Now suppose Π is any interval partition. Coarsening it, we may assume that each of its intervals contains an element of X. As $\mathcal{Y}\sim$ is groupwise dense, it contains a union of infinitely many intervals of Π. Call that union Z, and call its complement, which is in $\mathcal{Y}$, Y. We show that $Z \in \mathcal{G}_X$, witnessed by Y. Suppose $a < b$ are in Z and there is an element of Y in $[a,b)$. That means that a whole interval of Π must lie between a and b, and that interval contains a member of X. This completes the proof that $\mathcal{G}_X$ is groupwise dense.

Since there are fewer than $\mathfrak{g}$ X's in $\mathcal{X}$, there is a Z common to all the $\mathcal{G}_X$'s. Fix such a Z and define $f : \omega \to \omega$ by letting $f(n)$ be the number of elements of Z that are $\leq n$. Thus f is finite-to-one, being constant on the intervals $[a, b)$ where $a < b$ are consecutive in Z. For each $X \in \mathcal{X}$, the fact that $Z \in \mathcal{G}_X$ implies that there is $Y \in \mathcal{Y}$ with $f(Y) \subseteq f(X)$, as required. $\dashv$

9.16 Theorem *Assume $\mathfrak{u} < \mathfrak{g}$. For any filter $\mathcal{F}$ on ω either $\mathcal{F}$ is feeble or there is a finite-to-one f such that $f(\mathcal{F})$ is an ultrafilter.*

Proof. Apply the lemma with $\mathcal{X}$ being an ultrafilter base of cardinality $< \mathfrak{g}$ and $\mathcal{Y}$ being $\mathcal{F}$. If $\mathcal{F}$ is not feeble, then $\mathcal{Y}{\sim}$ is groupwise dense by Proposition 9.4, so the lemma provides a finite-to-one f such that $f(X) \in f(\mathcal{F})$ for all $X \in \mathcal{X}$ and therefore for all X in the ultrafilter $\mathcal{U}$ generated by $\mathcal{X}$. Thus, the ultrafilter $f(\mathcal{U})$ is included in the filter $f(\mathcal{F})$. Since ultrafilters are maximal filters, the inclusion cannot be proper, and $f(\mathcal{F})$ is an ultrafilter. $\dashv$

9.17 Remark The conclusion of this theorem is called the principle of *filter dichotomy*. It is not known whether it implies $\mathfrak{u} < \mathfrak{g}$.

The hypothesis of the theorem can be replaced by the apparently weaker $\mathfrak{r} < \mathfrak{g}$. Indeed, if $\mathcal{X}$ is not an ultrafilter base but merely unsplittable, the proof above provides a finite-to-one f such that $f(\mathcal{F})$ is also unsplittable. But an unsplittable filter is an ultrafilter.

The improvement is, however, illusory, for Mildenberger has shown that the inequalities $\mathfrak{u} < \mathfrak{g}$ and $\mathfrak{r} < \mathfrak{g}$ are equivalent. In fact, she proved $\mathfrak{r} \geq \min\{\mathfrak{u}, \mathfrak{g}\}$.

9.18 Corollary *Assume $\mathfrak{u} < \mathfrak{g}$ (or just filter dichotomy). For every two ultrafilters $\mathcal{U}$ and $\mathcal{V}$ on ω, there is a finite-to-one function f with $f(\mathcal{U}) = f(\mathcal{V})$.*

Proof. Apply filter dichotomy to the filter $\mathcal{U} \cap \mathcal{V}$. It is not feeble, and any f that maps it to an ultrafilter must map both $\mathcal{U}$ and $\mathcal{V}$ to the same ultrafilter. $\dashv$

9.19 Remark The conclusion of this corollary is called the principle of *near coherence of filters* (NCF). The name refers to the easily equivalent formulation: For any two filters $\mathcal{F}$ and $\mathcal{G}$ on ω, there is a finite-to-one function f such that $f(\mathcal{F})$ and $f(\mathcal{G})$ are coherent in the sense that their union generates a filter.

NCF implies $\mathfrak{u} < \mathfrak{d}$, but it is not known whether it implies $\mathfrak{u} < \mathfrak{g}$ or even filter dichotomy.

Corollary 9.18 can be improved to handle not just two ultrafilters but any number $< \mathfrak{c}$, by essentially the same proof, using Proposition 9.12 to ensure that the intersection filter is not feeble. NCF alone implies the improvement to $< \mathfrak{d}$ ultrafilters. It is also equivalent to the statement that every ultrafilter has a finite-to-one image that is generated by $< \mathfrak{d}$ sets. See [15] for these results and more information on NCF.

9.20 Corollary *If $\mathfrak{u} < \mathfrak{g}$ (or just filter dichotomy) then $\mathfrak{b} = \mathfrak{u}$ and $\mathfrak{d} = \mathfrak{c}$.*

Proof. Without any hypothesis, we have $\mathfrak{b} \leq \mathfrak{r} \leq \mathfrak{u}$ and $\mathfrak{d} \leq \mathfrak{c}$. If we assume filter dichotomy then Proposition 9.13 provides a feeble filter that is the intersection of $\mathfrak{d}$ ultrafilters, and then Proposition 9.12 says that $\mathfrak{d} \geq \mathfrak{c}$.

Theorem 9.10 gives a non-feeble filter generated by $\mathfrak{b}$ sets. By filter dichotomy, some image of this filter, which is also generated by $\mathfrak{b}$ sets (the images of the previous generators), is an ultrafilter. So $\mathfrak{u} \leq \mathfrak{b}$. ⊣

9.21 Remark The conclusion $\mathfrak{d} = \mathfrak{c}$ can be strengthened to $\mathfrak{g} = \mathfrak{c}$ under the hypothesis $\mathfrak{u} < \mathfrak{g}$; see [19].

The following result of Laflamme [69] extends the preceding dichotomy to a trichotomy for all upward-closed families. Its conclusion is in fact equivalent to $\mathfrak{u} < \mathfrak{g}$ but we omit the proof of this; see [19].

9.22 Theorem *Assume $\mathfrak{u} < \mathfrak{g}$. For any family $\mathcal{Y} \subseteq [\omega]^\omega$ that is closed under almost supersets, there is a finite-to-one $f : \omega \to \omega$ such that one of the following holds:*

- *$f(\mathcal{Y})$ contains only cofinite sets.*
- *$f(\mathcal{Y}) = [\omega]^\omega$.*
- *$f(\mathcal{Y})$ is an ultrafilter.*

Proof. Let $\mathcal{Y}$ be as in the theorem and let $\mathcal{X}$ be an ultrafilter base of cardinality $< \mathfrak{g}$. If $\mathcal{Y}{\sim}$ is not groupwise dense, then we have (by definition of groupwise dense) the first alternative in the theorem, and if ${\sim}\mathcal{Y}$ is not groupwise dense, then we have the second alternative. So we assume that both $\mathcal{Y}{\sim}$ and ${\sim}\mathcal{Y} = \breve{\mathcal{Y}}{\sim}$ are groupwise dense. The former lets us apply Lemma 9.15 to obtain a finite-to-one g such that each $g(X)$ with $X \in \mathcal{X}$ includes some $g(Y)$ with $Y \in \mathcal{Y}$. If $\mathcal{U}$ is the ultrafilter with base $\mathcal{X}$, then we have that $g(\mathcal{U}) \subseteq g(\mathcal{Y})$. Since finite-to-one images preserve groupwise density and commute with complementation, we also have that $g(\breve{\mathcal{Y}}){\sim}$ is groupwise dense, so we can apply the lemma with the base $\{g(X) : X \in \mathcal{X}\}$ of $g(\mathcal{U})$ in the role of $\mathcal{X}$ and with $g(\breve{\mathcal{Y}})$ in the role of $\mathcal{Y}$. We obtain a finite-to-one h such that $hg(\mathcal{U}) \subseteq hg(\breve{\mathcal{Y}}) = (hg(\mathcal{Y}))^\breve{}$. Since dualization $(\breve{})$

reverses inclusions and fixes ultrafilters, we get $hg(\mathcal{U}) \supseteq hg(\mathcal{Y})$. The reverse inequality follows from $g(\mathcal{U}) \subseteq g(\mathcal{Y})$. So the finite-to-one map hg sends $\mathcal{Y}$ to an ultrafilter. $\dashv$

We conclude this section with a brief discussion of some special sorts of ultrafilters. The theory of these ultrafilters is quite extensive, but we shall consider only those aspects that directly involve some of the cardinal characteristics defined earlier.

9.23 Definition An ultrafilter $\mathcal{U}$ on ω is *selective* if every function $f : \omega \to \omega$ becomes either one-to-one or constant when restricted to some set in $\mathcal{U}$. It is a *P-point* if every function $f : \omega \to \omega$ becomes either finite-to-one or constant when restricted to some set in $\mathcal{U}$. It is a *Q-point* if every finite-to-one function $f : \omega \to \omega$ becomes one-to-one when restricted to some set in $\mathcal{U}$.

9.24 Remark Clearly, an ultrafilter is selective if and only if it is both a P-point and a Q-point.

The name "selective" refers to the fact that, when ω is partitioned into pieces that are not in $\mathcal{U}$ then some set in $\mathcal{U}$ selects one element per piece. Selective ultrafilters are also called Ramsey ultrafilters, because Kunen showed (see [28]) that, if $\mathcal{U}$ is selective and $f : [\omega]^k \to 2$, then some set in $\mathcal{U}$ is homogeneous for f. Thus, any pseudobase for a selective ultrafilter must have cardinality at least $\mathfrak{hom} = \max\{\mathfrak{r}_\sigma, \mathfrak{d}\}$. Selective ultrafilters are also called RK-minimal, for they are minimal in the Rudin-Keisler ordering defined by putting $f(\mathcal{U}) \leq \mathcal{U}$ for all ultrafilters $\mathcal{U}$ and all mappings f.

An ultrafilter $\mathcal{U}$ is a P-point if and only if every decreasing (or almost-decreasing) ω-sequence of sets from $\mathcal{U}$ has a pseudo-intersection in $\mathcal{U}$. To prove the equivalence of this with the definition above, just arrange that $f(n)$ is constant exactly on the differences of consecutive sets in the decreasing sequence. (One can assume without loss of generality that the sequence begins with ω and that its intersection is empty.) There is a general topological concept of P-point (see for example [94, 49]), namely a point (in a topological space) such that every countable intersection of open neighborhoods of it includes another open neighborhood of it. When applied to the topological space $\beta\omega - \omega$, the Stone-Čech remainder of the discrete space ω, whose points are naturally identified with (non-trivial) ultrafilters on ω, this topological notion becomes the concept defined above. The "P" in "P-point" refers to prime ideals (in rings of functions); see [49, Exercises 4J and 4L].

The "Q" in "Q-point" was chosen because it's next to "P" in the alphabet. Q-points are also called rare ultrafilters.

There are ultrafilters that are neither P-points nor Q-points. Indeed, if $\mathcal{U}$ is any ultrafilter on ω then

$$\mathcal{V} = \{X \subseteq \omega^2 : \{a : \{b : \langle a, b\rangle \in X\} \in \mathcal{U}\} \in \mathcal{U}\}$$

is an ultrafilter on ω^2. It is not a P-point because the first projection $\omega^2 \to \omega$ is neither finite-to-one nor constant on any set in $\mathcal{V}$. It is not a Q-point because the second projection is finite-to-one on a set in $\mathcal{V}$, namely $\{\langle a, b\rangle : a < b\}$, but not one-to-one on any set in $\mathcal{V}$.

The existence of P-points, Q-points, and selective ultrafilters is more problematic. W. Rudin [94] showed that CH implies the existence of P-points, and other existence results followed, with the hypothesis weakened to MA or even to $\mathfrak{p} = \mathfrak{c}$ once these axioms had been formulated; see for example [28], [13], [14], [74], and [93].

But some hypotheses beyond ZFC are needed for such existence results. Kunen [66] showed that adding $\aleph_2$ random reals to a model of GCH produces a model with no selective ultrafilters. Miller [78] showed that an $\aleph_2$-step, countable support iteration of Laver forcing over a model of GCH produces a model with no Q-points. And Shelah [96, Section VI.4], [112] produced a model with no P-points by iterating a product of Grigorieff forcings.

We shall be concerned here with conditions for the existence of these special ultrafilters. It turns out that cardinal characteristics can be used to give necessary and sufficient conditions for the extendibility, to special ultrafilters, of all filters with sufficiently small bases. Thus, they provide sufficient, though not necessary, conditions for the mere existence of special ultrafilters. The first result of this sort is due to Ketonen [65], who showed that $\mathfrak{c} = \mathfrak{d}$ implies the existence of P-points, by a proof that essentially gives the following result.

9.25 Theorem *1. If $\mathfrak{c} = \mathfrak{d}$ then every filter generated by fewer than $\mathfrak{c}$ sets is included in some P-point.*

2. There is a filter generated by $\mathfrak{d}$ sets that is not included in any P-point.

3. Every ultrafilter generated by fewer than $\mathfrak{d}$ sets is a P-point.

Proof. For part 1, assume $\mathfrak{c} = \mathfrak{d}$, let $\mathcal{F}$ be a filter generated by fewer than $\mathfrak{c}$ sets, and let $\langle S^\alpha : \alpha < \mathfrak{c}\rangle$ be an enumeration of all decreasing ω-sequences of infinite subsets of ω, $S^\alpha = \langle S^\alpha_0 \supseteq S^\alpha_1 \supseteq \ldots\rangle$. We shall define an increasing sequence $\langle \mathcal{F}^\alpha : \alpha \leq \mathfrak{c}\rangle$, starting with $\mathcal{F}^0 = \mathcal{F}$, taking unions at limit stages, and at successor stages adding one new generator to the filter in such a way that either the new generator is a pseudointersection of S^α or it is the complement of some S^α_n. Of course, we must make sure that the newly added generator a stage $\alpha + 1$ has infinite intersection with every set in $\mathcal{F}^\alpha$, so that $\mathcal{F}^{\alpha+1}$ will be a filter. But this is not difficult. If, for some n, $S^\alpha_n \notin \mathcal{F}^\alpha$,

then $\omega - S_n^\alpha$ can be added. If, on the other hand, $S_n^\alpha \in \mathcal{F}^\alpha$ for all n, then, because $\mathcal{F}^\alpha$ is generated by fewer than $\mathfrak{d}$ sets, Proposition 6.24 provides a pseudointersection of S^α that has infinite intersection with every generator of $\mathcal{F}^\alpha$ and hence with every set in $\mathcal{F}^\alpha$. That pseudointersection can serve as the new generator for $\mathcal{F}^{\alpha+1}$. Thus, the construction of the sequence of filters can be carried out, and it clearly ensures that any ultrafilter extending $\mathcal{F}^\mathfrak{c}$ is a P-point.

For part 2, consider the filter on ω^2 generated by the sets $\{\langle a,b\rangle : a \geq n\}$ for $n \in \omega$ and the sets $\{\langle a,b\rangle : b > f(a)\}$ for f in a dominating family $\mathcal{D} \subseteq {}^\omega\omega$ of cardinality $\mathfrak{d}$. An ultrafilter extending this filter cannot be a P-point, for any set on which the first projection $\omega^2 \to \omega$ is constant or finite-to-one is disjoint from a set in the filter and is therefore not in the ultrafilter.

For part 3, let $\mathcal{U}$ be an ultrafilter generated by fewer than $\mathfrak{d}$ sets and let $S = \langle S_n\rangle$ be a decreasing sequence of sets from $\mathcal{U}$. As in the proof of part 1, Proposition 6.24 provides a pseudointersection of S that meets every generator of $\mathcal{U}$. But as $\mathcal{U}$ is an ultrafilter, it follows that this pseudointersection is in $\mathcal{U}$. ⊣

Canjar [36] proved the following analogous result for selective ultrafilters. It was also found independently by Bartoszyński and Judah; see [5, Section 4.5.B].

9.26 Theorem *1. If $\mathfrak{c} = \mathbf{cov}(\mathcal{B})$ then every filter generated by fewer than $\mathfrak{c}$ sets is included in some selective ultrafilter.*

2. There is a filter generated by $\mathbf{cov}(\mathcal{B})$ sets that is not included in any selective ultrafilter.

Proof. For part 1, we proceed as in the corresponding proof for P-points, using an enumeration $\langle f^\alpha : \alpha < \mathfrak{c}\rangle$ of ${}^\omega\omega$ in place of the enumeration of decreasing sequences S^α. At stage α we have a filter $\mathcal{F}^\alpha$ with a basis $\mathcal{X}$ of fewer than $\mathbf{cov}(\mathcal{B})$ sets and we wish to form $\mathcal{F}^{\alpha+1}$ by adding one new generator, a set on which f^α is one-to-one or constant. If some set of the form $(f^\alpha)^{-1}(\{n\})$ (on which f is constant) has infinite intersection with every set in $\mathcal{F}^\alpha$, then it can serve as the new generator. So from now on we assume that this is not the case. We intend to find a "selector" $g \in \prod_{n\in R}(f^\alpha)^{-1}(\{n\})$, where $R = \mathrm{ran}(f^\alpha)$, such that for each generator $X \in \mathcal{X}$ of $\mathcal{F}^\alpha$ we have $\exists^\infty n\,(g(n) \in X)$. Once we have such a g, its range can clearly serve as the new generator for $\mathcal{F}^{\alpha+1}$. To obtain g, notice first that the space $\prod_{n\in R}(f^\alpha)^{-1}(\{n\})$ from which we want to choose it is a product of countable (possibly finite) discrete sets, so it is not covered by fewer than $\mathbf{cov}(\mathcal{B})$ meager sets. But for each $X \in \mathcal{X}$, those g that fail to have infinitely many values in X form a meager set. So, since $|\mathcal{X}| < \mathbf{cov}(\mathcal{B})$, the desired g exists.

Part 2 is immediate from part 2 of the preceding theorem if $\mathbf{cov}(\mathcal{B}) = \mathfrak{d}$, so we may assume for the rest of the proof that $\mathbf{cov}(\mathcal{B}) < \mathfrak{d}$ (recall Proposition 5.5). By Theorem 5.2, fix a family of $\mathbf{cov}(\mathcal{B})$ chopped reals (x_α, Π_α) such that no single real matches them all. Assume without loss of generality that every finitely many of these chopped reals are engulfed by another chopped real from the chosen family, i.e., the family is directed upward with respect to the engulfing relation. Since we are assuming $\mathbf{cov}(\mathcal{B}) < \mathfrak{d}$, there is an interval partition $\Theta = \{J_n : n \in \omega\}$ not dominated by any of the Π_α. This implies that each Π_α has infinitely many intervals I_k that do not include any interval of Θ; such an I_k is covered by $J_n \cup J_{n+1}$ for some n.

Let Z be the set of functions z whose domains are unions of two consecutive intervals of Θ and whose values are 0's and 1's. For $z \in Z$, let $p(z)$ be the n such that $\mathrm{dom}(z) = J_n \cup J_{n+1}$. Thus, $p : Z \to \omega$ is finite-to-one. Let $\mathcal{F}$ be the filter on Z generated by the sets $\{z \in Z : p(z) > n\}$ for all $n \in \omega$ and the sets

$$A_\alpha = \{z \in Z : \exists I \in \Pi_\alpha \, (I \subseteq \mathrm{dom}(z) \text{ and } z \restriction I = x_\alpha \restriction I)\}$$

for all α. We must check that these sets have the SFIP, so consider any finitely many of them. We may assume only one of them is of the form $\{z \in Z : p(z) > n\}$; if there are more, keep only the one with the largest n as it's a subset of the others. Thanks to our assumption that any finitely many (x_α, Π_α) are engulfed by another, we may also assume that only one A_α is involved, for if (x_β, Π_β) engulfs certain (x_α, Π_α)'s, then the corresponding A_β is almost included in the corresponding A_α's. So our task is simply to check that each A_α contains z's with arbitrarily large $p(z)$. But this follows immediately from the fact that infinitely many intervals of Π_α are included in sets of the form $J_n \cup J_{n+1}$.

So $\mathcal{F}$ is a filter generated by $\mathbf{cov}(\mathcal{B})$ sets. Let $\mathcal{U}$ be any ultrafilter extending $\mathcal{F}$. p is a finite-to-one function, so it is certainly not constant on any set in $\mathcal{U}$. Suppose it were one-to-one on some set $X \in \mathcal{U}$. One of $X_0 = \{x \in X : p(x) \text{ even}\}$ and $X_1 = \{x \in X : p(x) \text{ odd}\}$ is in $\mathcal{U}$; say it's X_i. Then the union g of all the members of X_i is a partial function from ω to 2 such that each Π_α contains infinitely many intervals on which g agrees with x_α (because X_i meets all sets in $\mathcal{F}$). Any extension of g to a total function $\omega \to 2$ therefore matches all the (x_α, Π_α), contrary to our choice of these chopped reals. So p is not one-to-one on any set in $\mathcal{U}$. ⊣

By analogy with part 3 of Theorem 9.25, one might expect Theorem 9.26 to assert that every ultrafilter generated by fewer then $\mathbf{cov}(\mathcal{B})$ sets is selective. Though true, that assertion is vacuous, since Theorem 5.19 and Proposition 9.7 give $\mathbf{cov}(\mathcal{B}) \leq \mathfrak{r} \leq \mathfrak{u}$.

Canjar [36] also obtained an analogous result for Q-points.

9.27 Theorem *1. If* $\mathbf{cov}(\mathcal{B}) = \mathfrak{d}$ *then every filter generated by fewer than* $\mathfrak{d}$ *sets can be extended to a Q-point.*

2. If $\mathbf{cov}(\mathcal{B}) < \mathfrak{d}$ *then there is a filter generated by* $\mathbf{cov}(\mathcal{B})$ *sets that is not included in any Q-point.*

The proof of Theorem 9.26 also establishes Part 2 of the present theorem, and Part 1 is established similarly to Parts 1 of Theorems 9.25 and 9.26.

10. Evasion and Prediction

The terminology of prediction and evasion and the evasion number $\mathfrak{e}$ were introduced in [21] on the basis of motivation from algebra. Since then, several variants have been studied, particularly in [30, 34], but we begin with the original version.

10.1 Definition A *predictor* is a pair $\pi = (D, \langle \pi_n : n \in D \rangle)$ where $D \in [\omega]^\omega$ and where each $\pi_n : {}^n\omega \to \omega$. This predictor π *predicts* a function $x \in {}^\omega\omega$ if, for all but finitely many $n \in D$, $\pi_n(x \restriction n) = x(n)$. Otherwise, x *evades* π. The *evasion number* $\mathfrak{e}$ is the smallest cardinality of any family $\mathcal{E} \subseteq {}^\omega\omega$ such that no single predictor predicts all members of $\mathcal{E}$.

We may identify a predictor $(D, \langle \pi_n : n \in D \rangle)$ with $\bigcup_{n \in D} \pi_n$, a partial function from ${}^{<\omega}\omega$ to ω.

The idea behind the definition is that the values $x(n)$ of an unknown $x \in {}^\omega\omega$ are being revealed one at a time (in order) and we are trying to guess some of these values just before they are revealed. A predictor $(D, \langle \pi_n : n \in D \rangle)$ is a strategy for predicting $x(n)$, for each $n \in D$, after we have seen $x \restriction n$, and it predicts x if it is successful in the sense that almost all of its predictions about x are correct.

Clearly, it would make no difference if we defined predictors with $\pi_n : {}^nC \to C$ and used them to predict functions in ${}^\omega C$ for any countably infinite set C.

What was directly relevant to the algebraic subject of [21] was not $\mathfrak{e}$ but a variant, the *linear evasion number* $\mathfrak{e}_l$, whose definition is similar except that the components of a predictor are *linear* functions $\pi_n : \mathbb{Z}^n \to \mathbb{Q}$ and the functions being predicted are in ${}^\omega\mathbb{Z}$. Thus a remnant of algebra (linearity) was mixed with the combinatorics. Fortunately, it is proved in [34] that $\mathfrak{e}_l = \min\{\mathfrak{e}, \mathfrak{b}\}$, so the algebra can be eliminated in favor of pure combinatorics.

Several additional variants were defined in [30] by restricting the possible values of the functions being predicted, as follows.

10.2 Definition Let $f : \omega \to \omega - \{0, 1\}$. Let $\mathfrak{e}_f$ be the smallest cardinality of any family $\mathcal{E} \subseteq \prod_{n \in \omega} f(n)$ such that no single predictor predicts all

members of $\mathcal{E}$. When f is the constant function with value $n \geq 2$, we write $\mathfrak{e}_n$ instead of $\mathfrak{e}_f$. The *unbounded evasion number* $\mathfrak{e}_{\rm ubd}$ is the minimum of $\mathfrak{e}_f$ over all functions f as above.

Clearly, $\mathfrak{e}_f \geq \mathfrak{e}_g$ whenever $f \leq g$, and $\mathfrak{e}_{\rm ubd} \geq \mathfrak{e}$. The following theorem from [30] summarizes relationships between these variants and the original $\mathfrak{e}$ (as well as $\mathfrak{b}$ and $\mathfrak{s}$).

10.3 Theorem *1.* $\mathfrak{e}_n = \mathfrak{e}_2$ *for all* $n \geq 2$.

2. $\mathfrak{e}_2 \geq \mathfrak{s}$.

3. $\mathfrak{e} \geq \min\{\mathfrak{e}_{ubd}, \mathfrak{b}\}$.

4. It is consistent that $\mathfrak{e} < \mathfrak{e}_{ubd}$.

5. It is consistent that $\mathfrak{e}_{ubd} < \mathfrak{e}_2$.

Proof. We only sketch the proofs, referring to [30] for details.

For part 1, the idea is to predict a function $x : \omega \to n$ (where $n \geq 3$) by predicting the two functions $k \mapsto x(k) \bmod 2$ and $k \mapsto \lfloor x(k)/2 \rfloor$, whose ranges are smaller than n. More precisely, after predicting the former on some D, one predicts (on some $D' \subseteq D$) the restriction of the latter to D.

For part 2, we show that a family $\mathcal{E} \subseteq {}^{\omega}2$ that is not splitting (when viewed in $\mathcal{P}\omega$) can be predicted. If X is an infinite set on which each $x \in \mathcal{E}$ is almost constant, then let π be the predictor, with domain $D = X - \{\min X\}$, predicting that x will take, at any point of D, the same value that it took at the last previous member of X. This guess is right almost always, for every $x \in \mathcal{E}$.

For part 3, the idea is that any fewer than $\min\{\mathfrak{e}_{\rm ubd}, \mathfrak{b}\}$ functions can be predicted by first dominating them with some f (as there are $< \mathfrak{b}$ of them) and then regarding them as functions in $\prod_{n\in\omega} f(n)$, where they can be predicted (as there are fewer than $\mathfrak{e}_{\rm ubd} \leq \mathfrak{e}_f$ of them). Some care is needed as each function is below f only almost everywhere.

Part 4 is proved by an iterated forcing argument, where each step is a σ-centered forcing adding a predictor that predicts all ground-model elements of $\prod_{n\in\omega} f(n)$ for some f. A condition consists of a finite part of the desired predictor plus a promise to predict correctly all later values of finitely many functions. A finite-support iteration of this clearly makes $\mathfrak{e}_{\rm ubd}$ large in the extension. We omit the hard part of the proof, namely showing that $\mathfrak{e}$ does not become large.

For part 5, iterate Mathias forcing with countable supports for $\aleph_2$ steps over a model of GCH. The resulting model has $\mathfrak{h} = \mathfrak{c} = \aleph_2$, so both $\mathfrak{b}$ and $\mathfrak{s}$ are $\aleph_2$. By part 2, we have $\mathfrak{e}_2 = \aleph_2$. On the other hand, the forcing adds no Cohen reals, so $\mathbf{cov}(\mathcal{B}) = \aleph_1$. We shall see below (Table 2 and its

explanation) that $\mathfrak{e} \leq \mathbf{cov}(\mathcal{B})$. So by part 3 we have $\min\{\mathfrak{e}_{\rm ubd}, \mathfrak{b}\} \leq \aleph_1$. Since $\mathfrak{b} = \aleph_2$, we must have $\mathfrak{e}_{\rm ubd} = \aleph_1$. $\dashv$

Returning from the discussion of these variants to the original $\mathfrak{e}$, we have the following results.

10.4 Theorem *1.* $\mathbf{add}(\mathcal{L}), \mathfrak{p} \leq \mathfrak{e} \leq \mathbf{non}(\mathcal{B}), \mathbf{cov}(\mathcal{B})$.

2. It is consistent that $\mathfrak{e} < \mathbf{add}(\mathcal{B})$.

3. It is consistent that $\mathfrak{b} < \mathfrak{e}$.

In part 1, the inequality involving $\mathbf{cov}(\mathcal{B})$ is an unpublished result of M. Kada. The rest of part 1 is from [21]. Parts 2 and 3 are from [30] and [34] respectively.

Proof. The upper bound of $\mathbf{cov}(\mathcal{B})$ will follow from Tables 2 and 3 and their justification below. The upper bound of $\mathbf{non}(\mathcal{B})$ follows from the observation that any predictor can predict only a meager set of functions. (The set of functions predicted by any π also has measure zero in the standard measure, described in the introduction, on ${}^{\omega}\omega$. So $\mathbf{non}(\mathcal{L})$ is also an upper bound, but this is a weaker bound than $\mathbf{cov}(\mathcal{B})$ by Theorem 5.11.)

To prove the lower bound of $\mathfrak{p}$, we use Theorem 7.12. We assume $\mathrm{MA}_\kappa(\sigma\text{-centered})$ and show that any family $\mathcal{H}$ of κ functions can be predicted by some predictor (D, π). Let P be the set of triples (d, p, F) where d is a finite subset of ω, p is a finite partial function into ω whose domain consists of sequences from ${}^{n}\omega$ for $n \in d$, and F is a finite subset of $\mathcal{H}$. (The "meaning" of (d, p, F) is that d is an initial segment of D, p is a finite part of π, and the functions in F will be predicted correctly at all points of $D - d$.) Partially order P by putting $(d', p', F') \leq (d, p, F)$ if d is an initial segment of d', $p \subseteq p'$, $F \subseteq F'$, and whenever $n \in d' - d$ and $x \in F$ then $p'(x \restriction n)$ is defined and equal to $x(n)$. Any finitely many elements with the same first and second components have a lower bound, obtained by taking the union of the third components. So $\mathrm{MA}_\kappa(\sigma\text{-centered})$ provides $G \subseteq P$ generic with respect to the dense sets $\{(d, p, F) \in P : x \in F\}$ for all $x \in \mathcal{H}$, $\{(d, p, F) \in P : s \in \mathrm{dom}(p) \text{ or } n \notin d, n < \max d\}$ for all $n \in \omega$, $s \in {}^{n}\omega$, and $\{(d, p, F) \in P : |d| \geq n\}$ for all $n \in \omega$. (For proving the density of the last of these, the idea is that, starting with any (d, p, F), we can enlarge d by choosing m so large that all the $x \restriction m$ for $x \in F$ are distinct and then adjoining m to d and enlarging p as required by the definition of $\leq$. The choice of m ensures that the required enlargements of p do not conflict.) Then by letting D and π be the unions of the first components and second components, respectively, of the triples in G, we obtain a predictor predicting all the functions in $\mathcal{H}$.

To prove the lower bound of $\mathbf{add}(\mathcal{L})$, suppose we are given a family $\mathcal{H}$ of fewer than $\mathbf{add}(\mathcal{L})$ functions $x : \omega \to \omega$. Let $\{I_n : n \in \omega\}$ be the interval

partition where $|I_n| = n+1$. To each $x \in \mathcal{H}$ associate the function defined by $x'(n) = x \restriction I_n$. By Theorem 5.14, we can assign to each n a set $S(n)$ consisting of n functions $I_n \to \omega$ in such a way that $\forall x \in \mathcal{H} \forall^\infty n\, (x'(n) \in S(n))$. Any n functions produce at most $n-1$ branching points, i.e., points k where two of the functions first differ. So there is some $i_n \in I_n$ that is not a branching point for any of the n functions in $S(n)$. So we can define a predictor with $D = \{i_n : n \in \omega\}$ by setting $\pi(s) = z(i_n)$ if s has length i_n and $z \in S(n)$ and s agrees with z on $I_n \cap i_n$. (Extend p arbitrarily to those s whose length is in D but which admit no such z.) This π predicts all $x \in \mathcal{H}$ because the associated x' have almost all their values in $S(n)$.

This completes (modulo Tables 2 and 3) the proof of part 1. For parts 2 and 3, we only indicate the forcings used, referring to [30, 34] for the hard parts of the proofs.

Part 2 is proved by a finite-support iteration of Hechler forcing. Since this adds Cohen reals and dominating reals, both $\mathbf{cov}(\mathcal{B})$ and $\mathfrak{b}$ and therefore also their minimum $\mathbf{add}(\mathcal{B})$ are large in the extension. The hard part of the proof is to show that $\mathfrak{e}$ remains small.

Part 3 is proved by a finite-support iteration where each step adds a predictor that predicts all ground model reals. As in the proof of $\mathfrak{p} \leq \mathfrak{e}$ above, a condition consists of a finite part of the desired predictor together with finitely many functions that are to be predicted correctly at all later points. This forcing clearly makes $\mathfrak{e}$ large; the hard part is to prove that $\mathfrak{b}$ remains small. ⊣

10.5 Remark Laflamme has improved the inequality $\mathfrak{p} \leq \mathfrak{e}$ in Theorem 10.4 to $\mathfrak{t} \leq \mathfrak{e}$. In [70, Prop. 2.3] he shows that $\mathfrak{t} \leq \mathfrak{e}_{ubd}$, and he mentions that $\mathfrak{t} \leq \mathfrak{e}$ follows via part 3 of Theorem 10.3.

We turn next to some additional variations on the theme of prediction and evasion. These variations turn out to be closely connected to cardinals studied in previous sections. We consider three sorts of variations, singly and in combination.

First, the predictor could guess less information than the exact value of the $x(n)$ being predicted. Thus, we consider predictors $(D, \langle \pi_n : n \in D \rangle)$ where each $\pi_n : {}^n\omega \to \mathcal{P}\omega$, and we consider that $x \in {}^\omega\omega$ is predicted by such a π if $\forall^\infty n\, (x(n) \in \pi_n(x \restriction n))$. To avoid trivialities, the sets that occur as values of π_n must be small in some sense. (The predictor whose values are all equal to ω predicts every x.) We shall consider the following six possibilities for the values of π_n.

- Singletons. (This is the case considered above.)
- Sets of cardinality k for some fixed $k \in \omega$.

- Sets of cardinality $f(n)$, where f is a function $\omega \to \omega$ that tends to infinity.
- Finite sets.
- Co-infinite sets.
- Proper subsets of ω.

Thus, we shall refer to "single-valued" predictors, "k-valued" predictors, etc. Each type of predictor gives rise to an evasion number, namely the minimum number of functions not all predicted by a single predictor of that type.

Clearly, as the predictor's guesses become less specific (as we go down the list above), prediction becomes easier, evasion harder, and the evasion number larger.

Notice also that we could replace "finite sets" as values for π with "initial segments of ω" without affecting the evasion number, for given any predictor π of one sort we can trivially produce a predictor π' of the other sort predicting all the functions predicted by π. For the same reason, we can replace "proper subsets of ω" with "co-singletons."

The next variation concerns which values of x a predictor must guess correctly in order to predict x; it was also considered by Kada [62]. The definitions above permit the predictor to specify an infinite set D and guess $x(n)$ only for $n \in D$; it predicts x if almost all of these guesses are correct. We can make the definition more restrictive by requiring $D = \omega$. This variation will be called *global* prediction, and the original version will, when we want to emphasize the difference, be called *local* prediction.

Alternatively, we can make the definition less restrictive by saying that π predicts x if infinitely many (rather than almost all) of the guesses are right. We refer to this as *infinite* prediction. Notice that in this situation one might as well take $D = \omega$, because extending a predictor to a larger D can only increase the collection of functions it predicts. Thus, for both global and infinite prediction, we usually regard a predictor as either a sequence $\langle \pi_n \rangle_{n \in \omega}$ or as the union of such a sequence, $\pi : {}^{<\omega}\omega \to \omega$.

Clearly, as we move from global to local to infinite prediction, prediction becomes easier, evasion harder, and the evasion number larger.

The final variation that we consider here is to make $\pi_n(s)$ independent of s. In other words, the predictor is not allowed to see $x \restriction n$ but only knows n when guessing $x(n)$. Thus, the predictor is essentially just a function π on ω or D, taking "small" values in one of the senses above. We refer to such predictors as *non-adaptive* while predictors of the original sort are *adaptive*. Clearly, adaptive prediction is easier than non-adaptive prediction, evasion harder, and the evasion number larger.

The six choices for "small," the three choices global or local or infinite, and the two choices non-adaptive or adaptive give 36 evasion numbers, one

of which (singleton, local, adaptive) is $\mathfrak{e}$. Many of the others coincide with cardinals discussed earlier, and for the rest there are bounds in terms of such cardinals. This information is summarized in the following tables. The first column of each table lists the six species of smallness, with G representing a typical guess for $x(n)$.

Our remarks above imply that the entries in each table increase (weakly) from top to bottom and from left to right; also, as we go from one table to the next (global to local to infinite), the entries in any single position increase (weakly). We shall usually refer to these facts as "monotonicity" without going into any more detail.

	Non-adaptive	Adaptive
$\|G\| = 1$	2	$\aleph_1$
$\|G\| = k$	$k+1$	$\mathfrak{m}(\sigma\text{-}k\text{-linked}) \leq ? \leq \mathbf{add}(\mathcal{L})$
$\|G\| = f(n)$	$\mathbf{add}(\mathcal{L})$	$\mathbf{add}(\mathcal{L})$
G finite	$\mathfrak{b}$	$\mathfrak{b}$
$\omega - G$ infinite	$\mathbf{non}(\mathcal{B})$	$\mathbf{non}(\mathcal{B})$
$G \subsetneq \omega$	$\mathbf{non}(\mathcal{B})$	$\mathbf{non}(\mathcal{B})$

Table 1: Evasion Numbers for Global Prediction

	Non-adaptive	Adaptive
$\|G\| = 1$	2	$\mathfrak{e}$
$\|G\| = k$	$k+1$	$\mathfrak{e} \leq ? \leq \mathbf{cov}(\mathcal{B}), \mathbf{non}(\mathcal{B})$
$\|G\| = f(n)$	$\min\{\mathfrak{e}, \mathfrak{b}\}$	$\mathfrak{e} \leq ? \leq \mathbf{cov}(\mathcal{B}), \mathbf{non}(\mathcal{B})$
G finite	$\mathfrak{b}$	$\mathfrak{e}, \mathfrak{b} \leq ? \leq \mathfrak{d}, \mathbf{non}(\mathcal{B})$
$\omega - G$ infinite	$\mathbf{non}(\mathcal{B})$	$\mathbf{non}(\mathcal{B})$
$G \subsetneq \omega$	$\mathbf{non}(\mathcal{B})$	$\mathbf{non}(\mathcal{B})$

Table 2: Evasion Numbers for Local Prediction

	Non-adaptive	Adaptive
$\|G\| = 1$	$\mathbf{cov}(\mathcal{B})$	$\mathbf{cov}(\mathcal{B})$
$\|G\| = k$	$\mathbf{cov}(\mathcal{B})$	$\mathbf{cov}(\mathcal{B})$
$\|G\| = f(n)$	$\mathbf{cov}(\mathcal{B})$	$\mathbf{cov}(\mathcal{B})$
G finite	$\mathfrak{d}$	$\mathfrak{d}$
$\omega - G$ infinite	$\mathfrak{c}$	$\mathfrak{c}$
$G \subsetneq \omega$	$\mathfrak{c}$	$\mathfrak{c}$

Table 3: Evasion Numbers for Infinite Prediction

The question marks in four of the entries indicate that I do not know the values of these evasion numbers but only the indicated bounds and the

result of Mildenberger (unpublished) that the following three cardinals are equal:

- $\mathfrak{e}$,
- the smaller of $\mathfrak{e}_2$ and the question mark in the "$|G| = k$" line of Table 2,
- the smaller of $\mathfrak{e}_{\mathrm{ubd}}$ and the question mark in the "$|G|$ finite" line of Table 2.

One could regard these entries with question marks as defining four more cardinal characteristics. On the other hand, one might regard the entry $\mathfrak{e}$ in Table 2 as a euphemism for a question mark with the bounds given in Theorem 10.4. The difference between $\mathfrak{e}$ and the question marks is that the former has been studied enough to indicate that it differs from the previously studied characteristics, while the question marks might well reduce to something simpler.

In the following paragraphs, we give reasons for the table entries, leaving some details to the reader.

2 and $k+1$ For both global and local prediction, $k+1$ distinct constant functions evade any non-adaptive predictor of k-element sets. And any k functions clearly can be predicted.

$\mathbf{add}(\mathcal{L})$ In the non-adaptive column of Table 1, the occurrence of $\mathbf{add}(\mathcal{L})$ expresses Theorem 5.14 and the remark following it. The occurrence in the adaptive column comes from the fact that an adaptive $f(n)$-valued predictor π gives rise to a non-adaptive $f'(n)$-valued predictor π' (for a larger f') such that all functions globally predicted by π are also globally predicted by π'. Given π, associate to each $s \in {}^{<\omega}\omega$ and each natural number n the set $\pi_s(n)$ of all possible values of $x(n)$ for functions $x \in {}^{\omega}\omega$ that start with s and are correctly predicted by π thereafter. (That is, $x(k)$ is $s(k)$ for $k < \mathrm{length}(s)$ and $\pi(x \restriction k)$ for all larger k.) Also fix an enumeration of ${}^{<\omega}\omega$ in an ω-sequence. Let $\pi'(n)$ be the union of the sets $\pi_s(n)$ as s ranges over the first n elements of ${}^{<\omega}\omega$. It is easy to find an appropriate f' depending only on f and to verify that π' globally predicts everything that π does.

$\mathfrak{b}$ We may take a finite-valued predictor's guesses to be initial segments of ω, i.e., natural numbers, a guess being correct if it is greater than the actual value of the function being guessed. In this light, the occurrence of $\mathfrak{b}$ in the non-adaptive column of Table 1 expresses just the definition of $\mathfrak{b}$. The occurrence in the adaptive column is justified by an argument analogous to that in the discussion of $\mathbf{add}(\mathcal{L})$ above.

As for the occurrence in Table 2, consider any unbounded family $\mathcal{E}$ of $\mathfrak{b}$ non-decreasing functions. We shall see that they evade local prediction by any non-adaptive, finite-valued predictor (π, D). As above, we assume the values of π are natural numbers. Define $\pi' : \omega \to \omega$ by letting $\pi'(n)$ be the

value of π at the next member of D after n. By our choice of $\mathcal{E}$, it contains a member x not dominated by π'. Since x is non-decreasing, if (π, D) locally predicted it then for all sufficiently large $n \in \omega$ we would have, letting k be the next element of D after n,

$$x(n) \leq x(k) < \pi(k) = \pi'(n).$$

This contradicts the choice of x, so x evades (π, D).

$\boxed{\mathbf{non}(\mathcal{B})}$ Let us consider first the bottom row in Tables 1 and 2, where the guesses are proper subsets of ω. Without loss of generality, we may assume that the guesses are complements of singletons. We'll write $\tilde{\pi}(n)$ for the number absent from $\pi(n)$ (and similarly for $\tilde{\pi}(s)$ in the adaptive case). Part 2 of Theorem 5.9 says that the bottom entry in the non-adaptive column of Table 1 is $\mathbf{non}(\mathcal{B})$. By monotonicity, the other entries in the bottom row of Tables 1 and 2 are no smaller. They are no larger because any predictor predicts globally or locally only a meager subset of ${}^{\omega}\omega$.

To justify the next-to-bottom row in Tables 1 and 2, where the guesses are co-infinite, it suffices, thanks to monotonicity, to show that a family $\mathcal{F}$ of fewer than $\mathbf{non}(\mathcal{B})$ functions cannot evade global prediction by non-adaptive co-infinite predictors. Fix a map $p : \omega \to \omega$ such that every $p^{-1}\{n\}$ is infinite. The fewer than $\mathbf{non}(\mathcal{B})$ functions $p \circ f$ for $f \in \mathcal{F}$ are globally predicted by a predictor π of proper subsets of ω (by Theorem 5.9); so the functions in $\mathcal{F}$ are predicted by $p^{-1} \circ \pi$, whose values are co-infinite.

$\boxed{\mathfrak{c}}$ Clearly, all evasion numbers are $\leq \mathfrak{c}$, since any predictor, even an adaptive predictor of co-singletons, can be completely evaded by some function. On the other hand, to evade infinite prediction even by non-adaptive predictors π of co-singletons requires $\mathfrak{c}$ functions, because one needs functions eventually equal to any prescribed $f : \omega \to \omega$ (giving the values omitted by the predictor). To obtain the same result with "co-infinite" in place of "co-singleton," use the same "compose with p" trick as in the discussion of $\mathbf{non}(\mathcal{B})$ above.

$\boxed{\mathfrak{d}}$ As in the discussion of $\mathfrak{b}$ above, we may assume that predictors give natural numbers, intended as upper bounds for the values to be guessed. Then the $\mathfrak{d}$ in the non-adaptive column of Table 3 is justified by the definition of $\mathfrak{d}$. To see that $\mathfrak{d}$ functions suffice to evade even adaptive prediction, take a family of $\mathfrak{d}$ adaptive predictors that dominate all the adaptive predictors, and choose for each of these predictors some function evading it.

$\boxed{\mathbf{cov}(\mathcal{B})}$ Of the six occurrences of $\mathbf{cov}(\mathcal{B})$ in Table 3, the top one in the non-adaptive column expresses Part 1 of Theorem 5.9. To justify the rest, it suffices by monotonicity to check that $\mathbf{cov}(\mathcal{B})$ functions suffice to evade infinite prediction by adaptive predictors whose guesses at n have cardinality $f(n)$. We do this first for non-adaptive predictors, by a modification of the argument for Theorem 5.9, and then we show how to extend the result to

the adaptive case. We may assume $\mathbf{cov}(\mathcal{B}) < \mathfrak{d}$, for otherwise the desired information follows by monotonicity from the $\mathfrak{d}$'s in the next row of Table 3.

Fix $\mathbf{cov}(\mathcal{B})$ chopped reals (x_α, Π_α) with no single $y \in {}^\omega 2$ matching them all (by Theorem 5.2). Since $\mathbf{cov}(\mathcal{B}) < \mathfrak{d}$, fix an interval partition Θ not dominated by any of the Π_α. As in the proof of Theorem 5.9, this means that every Π_α contains infinitely many intervals each covered by two consecutive intervals of Θ.

Define $g(n) = 2 \cdot \sum_{k \leq n} f(k) - 1$. For each α and each n, we define a set $q_\alpha(n)$ as follows. Find $g(n)$ disjoint pairs of consecutive Θ-intervals, each pair covering a Π_α-interval; let the unions of these pairs be $J_0, \ldots, J_{g(n)-1}$. Then let $q_\alpha(n) = \{x_\alpha \restriction J_0, \ldots, x_\alpha \restriction J_{g(n)-1}\}$. So each $q_\alpha(n)$ is a set of $g(n)$ functions into 2, each having as domain the union of two consecutive Θ-intervals, and such that the domains of different members of $q_\alpha(n)$ are disjoint.

By coding their values as natural numbers, we can regard the q_α as functions $\omega \to \omega$. We claim that these $\mathbf{cov}(\mathcal{B})$ functions evade infinite prediction by any non-adaptive, $f(n)$-valued predictor, i.e., any f-slalom.

Suppose this failed. So there is a function S assigning to each $n \in \omega$ a set $S(n)$ of $f(n)$ elements such that for each α we have $\exists^\infty n \, (q_\alpha(n) \in S(n))$. Without loss of generality, each element s of $S(n)$ is a set of $g(n)$ functions into 2, each having as domain the union of two consecutive Θ-intervals, and such that the domains of different members of s are disjoint. Now define $y \in {}^\omega 2$ by the following recursion, defining y on $2 \cdot f(n)$ Θ-intervals at step n.

Suppose steps 0 through $n-1$ have been completed, so y is already defined on $2 \cdot \sum_{k<n} f(k)$ Θ-intervals. From each $s \in S(n)$, remove those partial functions whose domains overlap the set where y is already defined. That removes at most $2 \cdot \sum_{k<n} f(k)$, so at least $2 \cdot f(n) - 1$ are left, since s had cardinality $g(n)$. Go through the $f(n)$ sets so obtained (one from each $s \in S(n)$) in some order, picking one function from each, making sure that the domain of each chosen function is disjoint from the domains of the previously chosen functions. Since each of the domains is the union of two consecutive Θ-intervals, each domain can overlap at most two others. Thus, there are at least $2 \cdot f(n) - 1$ options for the first choice, at least $2 \cdot f(n) - 3$ for the second, and so on down to at least 1 option for the $f(n)^{\text{th}}$ choice. So all the choices can be made. Then extend y to agree with each of the chosen functions on its domain. This completes step n of the recursion. After all steps are completed, if y is not defined on all of ω, extend it arbitrarily.

For each α, there are infinitely many n with $q_\alpha(n) \in S(n)$, so $q_\alpha(n)$ is one of the s's considered at step n in the definition of y. So some element z of $q_\alpha(n)$ becomes part of y. But that z is $x_\alpha \restriction J$ for some J that includes an interval of Π_α. So y matches each (x_α, Π_α), contrary to our choice of these chopped reals. This contradiction shows that the (coded) q_α are evasive as claimed.

It remains to extend the result to adaptive predictors whose guesses have size $f(n)$. Such a predictor is a function $\pi : {}^{<\omega}\omega \to [\omega]^{<\omega}$ (with $\pi(s) \in [\omega]^{f(n)}$ if $s \in {}^{n}\omega$). Identifying the domain ${}^{<\omega}\omega$ with ω via some bijective coding, we can view every such π as a non-adaptive predictor whose guesses have size $f'(n)$ for a certain f' (that depends on f and the coding). Applying the preceding argument to these non-adaptive predictors, and then reversing the coding process, we get a family $\mathcal{E}$ of $\mathbf{cov}(\mathcal{B})$ functions ${}^{<\omega}\omega \to \omega$ such that, for every adaptive predictor π as above,

$$\exists z \in \mathcal{E}\, \forall^{\infty} s \in {}^{<\omega}\omega\, (z(s) \notin \pi(s).)$$

Use each $z \in \mathcal{E}$ to recursively define a $z' : \omega \to \omega$ by $z'(n) = z(z' \restriction n)$. Then the family $\mathcal{E}' = \{z' : z \in \mathcal{E}\}$ evades infinite prediction by any π as above. Indeed, with π and z as above, we have, for all but finitely many n, that

$$z'(n) = z(z' \restriction n) \notin \pi(z' \restriction n).$$

$\boxed{\mathfrak{e}}$ The entry $\mathfrak{e}$ in Table 2 is just the definition of $\mathfrak{e}$. In view of what we just proved about $\mathbf{cov}(\mathcal{B})$, we get $\mathfrak{e} \leq \mathbf{cov}(\mathcal{B})$ by monotonicity. This completes the proof of Theorem 10.4 above.

$\boxed{\aleph_1}$ Any countably many functions $h_i : \omega \to \omega$ are globally predicted by the adaptive predictor defined by requiring $\pi(s) = h_i(n)$ if s has length n and i is the first index with $h_i \restriction n = s$. Such a π predicts each h_i accurately at all n beyond the points where h_i first differs from the earlier h_j's.

On the other hand, any adaptive predictor whose values are singletons can globally predict only countably many functions. Indeed, a function h globally predicted by such a π is completely determined by the finite part of h consisting of the values not correctly guessed by π.

$\boxed{\mathfrak{m}(\sigma\text{-}k\text{-linked}) \leq ? \leq \mathbf{add}(\mathcal{L})}$ Monotonicity gives the upper bound of $\mathbf{add}(\mathcal{L})$. To establish the lower bound, we assume $\mathrm{MA}_{\kappa}(\sigma\text{-}k\text{-linked})$ and we prove that any family $\mathcal{H}$ of κ functions $\omega \to \omega$ can be globally predicted by an adaptive predictor π with k-element guesses, i.e., $\pi : {}^{<\omega}\omega \to [\omega]^k$. Let P be the set of pairs (p, F) where p is a finite partial map from ${}^{<\omega}\omega$ to $[\omega]^k$ and F is a finite subset of $\mathcal{H}$ with the following "branching restriction": For any two functions in F, if s is their longest common initial segment, then $s \in \mathrm{dom}(p)$. (The "meaning" of (p, F) is that p is part of the desired predictor and that each function in F will be guessed correctly except possibly at those places where p is defined.) Partially order P by putting $(p', F') \leq (p, F)$ if $p \subseteq p'$, $F \subseteq F'$, and, whenever $f \in F$ and $f \restriction n \in \mathrm{dom}(p') - \mathrm{dom}(p)$ then $f(n) \in p'(f \restriction n)$.

This partial ordering is σ-k-linked because any k elements (p, F_i) with the same first component have a common lower bound, constructed as follows. First form $(p, \bigcup_i F_i)$. If this is not the desired lower bound, it is because the branching restriction is violated. So there are some $s \notin \mathrm{dom}(p)$ that are

the largest common initial segments of some $f \in F_i$ and $g \in F_j$. Then $i \neq j$ because each of the (p, F_i) satisfied the branching restriction. So any such s, say of length n, is an initial segment of at most k members of $\bigcup_i F_i$. But then we can extend p by defining $p(s)$ to be a k-set containing the values at n of those $\leq k$ members of $\bigcup_i F_i$. Doing this for each such s, we get the desired lower bound.

Applying $\mathrm{MA}_\kappa(\sigma\text{-}k\text{-linked})$, we get a set $G \subseteq P$ generic with respect to the dense sets $\{(p, F) \in P : s \in \mathrm{dom}(p)\}$ for all $s \in {}^{<\omega}\omega$ and $\{(p, F) \in P : f \in F\}$ for all $f \in \mathcal{H}$. (The former is dense thanks to the branching restriction. To verify the density of the latter, given any (p, F) and any $f \in \mathcal{H} - F$, first form $(p, F \cup \{f\})$. If the branching restriction is violated, extend p so as to be defined at the new branching locations. Here we need that $k \geq 2$.) Let π be the union of all the first components of the pairs $(p, F) \in G$. It is routine to check (as in the proof of Theorem 7.7) that this π is an adaptive predictor with k-set guesses, globally predicting every function from $\mathcal{H}$.

$\boxed{\mathfrak{e} \leq ? \leq \mathbf{cov}(\mathcal{B}), \mathbf{non}(\mathcal{B})}$ Monotonicity implies all three inequalities.

$\boxed{\mathfrak{e}, \mathfrak{b} \leq ? \leq \mathfrak{d}, \mathbf{non}(\mathcal{B})}$ Again, monotonicity implies all four inequalities.

$\boxed{\min\{\mathfrak{e}, \mathfrak{b}\}}$ This is [34, Lemma 2.5]. Monotonicity gives the upper bound $\mathfrak{b}$. The proof that $\mathfrak{e}$ is also an upper bound is essentially the same as the proof of $\mathbf{add}(\mathcal{L}) \leq \mathfrak{e}$ in Theorem 10.4. The only difference is that here we are dealing with "partial slaloms," i.e., functions S defined on some infinite $D \subseteq \omega$ and satisfying $|S(n)| = f(n)$ for all $n \in D$. Instead of predicting at all i_n, as in the earlier argument, we now predict at i_n for $n \in D$.

To prove that $\min\{\mathfrak{e}, \mathfrak{b}\}$ is also a lower bound, let $\mathcal{H}$ be a family of fewer than $\min\{\mathfrak{e}, \mathfrak{b}\}$ functions; we must find a partial slalom (in the sense defined above) such that each $h \in \mathcal{H}$ satisfies $\forall^\infty n \in D\,(h(n) \in S(n))$. Since there are fewer than $\mathfrak{b}$ functions in $\mathcal{H}$, we can find a single, strictly increasing $g : \omega \to \omega$ that dominates them all. Let $\{I_n : n \in \omega\}$ be an interval partition such that, if a is the left endpoint of any I_n, then there is some $i_n \in I_n$ with $f(i_n) \geq g(a)^a$. To each $h \in \mathcal{H}$ associate the function defined by $h'(n) = h \restriction I_n$. Since the number of such h' is $< \mathfrak{e}$, there is an adaptive predictor of singletons (D', π') that locally predicts all the h'; that is,

$$\forall h \in \mathcal{H}\, \forall^\infty n \in D'\ (h'(n) = \pi'(h' \restriction n))\,.$$

Do the following for each $n \in D'$. Let a be the left endpoint of I_n, and recall that our interval partition was chosen so that $g(a)^a \leq f(i_n)$ for some $i_n \in I_n$. Consider all functions s from a into $g(a)$; there are exactly $g(a)^a$, and thus no more than $f(i_n)$, of them. Each gives an s' by $s'(m) = s \restriction I_m$ for $m < n$. Then $\pi'(s')$ is some function $I_n \to \omega$; evaluate it at i_n. Doing this for each s gives no more than $f(i_n)$ numbers; let $S(i_n)$ be the set of these numbers.

Doing this for all $n \in D'$, we get a partial slalom defined on $D = \{i_n : n \in D'\}$. For each $h \in \mathcal{H}$, if π' predicted $h'(n)$ correctly (where $n \in D'$), then $h(i_n) \in S(i_n)$. So we have the desired partial slalom.

10.6 Remark The variants of evasion discussed at the beginning of this section ($\mathfrak{e}_g$ and $\mathfrak{e}_{\mathrm{ubd}}$) can be combined with some of the variants in Tables 1 to 3. Finite and co-infinite predictors no longer make sense. When, as in the case of $\mathfrak{e}_g$, the functions to be predicted are bounded by a fixed g, we need to pay attention to the function f in the $|G| = f(n)$ lines of the tables; it is no longer the case that any function tending to infinity is equivalent to any other. Also, in this situation, the co-singleton case becomes a special case of $|G| = f(n)$ with $f(n) = g(n) - 1$. Thus, we would have three-line tables for these variants. We omit any further discussion of these, since little is known about them beyond carrying over some of the arguments presented above.

Another variation, lying between global and local, was introduced by Kamo [64]. Say that a function $\pi : {}^{<\omega}\omega \to \omega$ *constantly predicts* $x : \omega \to \omega$ if there is $n \in \omega$ such that, with finitely many exceptions, any interval $[m, m+n)$ of length n contains some k such that $x(k) = \pi(x \restriction k)$. This concept has been studied further by Kamo, Kada, and Brendle; see for example [32] and the references there.

Finally, all the evasion cardinals considered in this section have duals of the form: the smallest number of predictors needed to predict all functions. These too have been little studied, but there is one remarkable result concerning the number of f-slaloms needed to globally predict all members of $\prod_n g(n)$. Goldstern and Shelah [52] showed that this cardinal can vary with f and g and in fact that in some models of set theory uncountably many cardinals are of this form (infinitely many with recursive f and g).

11. Forcing

In this final section, we describe the effect of various forcing constructions on cardinal characteristics. We shall discuss only the most commonly used forcing notions and their most natural iterations; for a far more extensive discussion, see [5, Chapters 3, 6, and 7].

Most of the forcing notions we consider are designed to add a real with some prescribed properties, and the properties are often closely connected with some Borel relation $\mathbf{A} = (A_-, A_+, A)$ (where we use the notation of Section 4). Specifically, we say that a real x in a forcing extension *solves* $\mathbf{A}$ (over the ground model) if $x \in \tilde{A}_+$ and $(a, x) \in \tilde{A}$ for all $a \in A_-$ in the ground model. Here $\tilde{A}$ denotes the relation in the extension having the same Borel code as A has in the ground model, and similarly for $\tilde{A}_+$ etc., but we shall often omit the tilde since no confusion will result.

If there is a morphism $\varphi : \mathbf{A} \to \mathbf{B}$ whose φ_+ component is Borel, so that $\tilde{\varphi}_+$ makes sense, and if x solves $\mathbf{A}$ then $\tilde{\varphi}_+(x)$ solves $\mathbf{B}$. Indeed, given any $b \in B_-$ in the ground model, let $a = \varphi_-(b) \in A_-$. The statement

$$\forall u \in A_+ \, (aAu \implies bB\varphi_+(u))$$

is true in V and absolute when expressed in terms of the Borel codes of A_+, A, B, and φ_+. Thus it is true in any forcing extension that

$$\forall u \in \tilde{A}_+ \, (a\tilde{A}u \implies b\tilde{B}\tilde{\varphi}_+(u)).$$

Since a is in the ground model, we have $a\tilde{A}x$ and therefore $b\tilde{B}\tilde{\varphi}_+(x)$ as claimed.

Notice that we do not need A_-, B_- or φ_- to be Borel in the preceding discussion.

It is easy to check that if x solves $\mathbf{A}$ and y solves $\mathbf{B}$ then (x, y) solves the conjunction $\mathbf{A} \wedge \mathbf{B}$ and the product $\mathbf{A} \times \mathbf{B}$. For sequential composition, the situation is more complicated, because even if $\mathbf{A}$ and $\mathbf{B}$ are Borel, the set of challenges in $\mathbf{A}; \mathbf{B}$ is of higher type, so this relation cannot be Borel. However, if we have a morphism $\varphi : \mathbf{A}; \mathbf{B} \to \mathbf{C}$ then under suitable Borelness hypotheses we can conclude, by a proof very similar to that above, that if $V \subseteq V' \subseteq V''$, if $x \in V'$ solves $\mathbf{A}$ over V, and if $y \in V''$ solves $\mathbf{B}$ over V' then $\tilde{\varphi}_+(x, y)$ solves $\mathbf{C}$ over V. Most of the "suitable Borelness hypotheses" are the ones obviously needed for the statement to make sense: A_+, A, B_+, B, B_-, and φ_+ must be Borel. (B_-, unlike A_-, must be Borel so that solving $\mathbf{B}$ over V', not over V, makes sense.) But one additional Borelness hypothesis is needed for the proof. If we regard $\varphi_- : C_- \to A_- \times {}^{A_+}B_-$ as a pair of functions $\alpha : C_- \to A_-$ and $\beta : C_- \to {}^{A_+}B_-$, and if we regard β as $\beta' : C_- \times A_+ \to B_-$ (where $\beta'(c, a) = \beta(c)(a)$), then we need that β' is Borel. We leave the details to the reader.

Most of the iterations we consider will be either finite-support iterations of ccc forcing notions or countable-support iterations of proper forcing notions. For general information about iterations, see Abraham's chapter in this handbook or [61, 9, 96]. All the proper forcing notions considered below satisfy Baumgartner's Axiom A [9], which is stronger and usually easier to check than properness. We usually write V for the ground model and V_α for the model obtained after α stages of an iteration.

11.1. Finite-Support Iteration and Martin's Axiom

A finite-support iteration of ccc forcing is equivalent to a single ccc forcing [105] and therefore preserves cardinals. Also, if the length λ of the iteration has uncountable cofinality, then every real in the final extension V_λ is already in an intermediate extension V_α, $\alpha < \lambda$. If, cofinally often in such an iteration, one adjoins a real solving $\mathbf{A}^\perp$ over the previous model, then in V_λ

the norm $\|\mathbf{A}\|$ will be at least $\mathrm{cof}(\lambda)$. Indeed, given any fewer than $\mathrm{cof}(\lambda)$ members of A_+ in V_λ, we can find an $\alpha < \lambda$ such that all these reals are in V_α; increasing α if necessary, we can, by hypothesis, arrange that $V_{\alpha+1}$ contains a real $x \in A_-$ solving $\mathbf{A}^\perp$ over V_α. But that means in particular that x is A-related to none of our given fewer than $\mathrm{cof}(\lambda)$ reals.

The preceding remarks indicate a way to make a characteristic $\|\mathbf{A}\|$ large, namely iterate a ccc forcing that solves $\mathbf{A}^\perp$, with finite support, for λ stages, where λ is regular and large.

Applying this method with all ccc forcings of size $< \lambda$ (in all the intermediate models) suitably interleaved, one obtains a model of MA and $\mathfrak{c} = \lambda$ provided GCH held in the ground model. If one uses only σ-centered posets in the iteration, then one obtains a model of MA(σ-centered), i.e., $\mathfrak{p} = \mathfrak{c}$ (see Theorem 7.12), but MA fails and in fact $\mathbf{cov}(\mathcal{L}) = \aleph_1$ (see [5, Section 6.5D]. Similar constructions give models satisfying various fragments of MA while violating others; see Appendix B1 of [47] and the references there.

To prove independence results in the theory of cardinal characteristics, one needs techniques for making one characteristic large while keeping another small. As indicated above, it is not difficult to make a chosen characteristic large, but it is usually difficult to prove that another characteristic remains small. In fact, some characteristics cannot be kept small in a non-trivial finite-support iteration. The reason is that such an iteration always introduces Cohen reals at all limit stages of cofinality ω. Cohen reals solve various Borel relations (see below), notably $\mathbf{Cov}(\mathcal{B})^\perp$, and therefore finite support iterations cannot help making certain characteristics, notably $\mathbf{cov}(\mathcal{B})$, large.

11.2. Countable-Support Proper Iteration

A countable-support iteration of proper forcing is equivalent to a single proper forcing [96, Theorem 3.2] and therefore preserves $\aleph_1$. For our purposes, it will be important to also preserve larger cardinals, and this is usually ensured by an appeal to [96, Theorem 4.1], which gives the $< \aleph_2$-chain condition provided (1) CH holds in the ground model, (2) the forcing notion used to produce $V_{\alpha+1}$ from V_α has cardinality at most $\mathfrak{c}$ in V_α, and (3) the length of the iteration is at most ω_2. The first two of these provisos will be satisfied automatically in the situations we are interested in, but the third is a real impediment. This limitation on the length of the iteration prevents us from making the continuum arbitrarily large with countable-support iterations; only $\mathfrak{c} = \aleph_2$ can be achieved. It is shown in [11] that iterating Sacks forcing (which is proper) with countable support for $\omega_2 + 1$ steps collapses $\aleph_2$. Also, it is pointed out in [50, Remark 0.3] that a countable-support ω_1-stage iteration of any non-trivial forcings will collapse $\mathfrak{c}$ to $\aleph_1$.

Our inability to produce larger values of $\mathfrak{c}$ with the kind of detailed control available for countable-support iterations has prevented the solution of

several problems. For example, although we have models with no P-points and models with no Q-points (both obtained by countable-support proper iterations), we do not know how to achieve both simultaneously. By Theorems 9.25 and 9.27, such a model would need to have $\mathbf{cov}(\mathcal{B}) < \mathfrak{d} < \mathfrak{c}$ and therefore $\mathfrak{c} \geq \aleph_3$. Similarly, we have no model for $\mathfrak{p} < \mathfrak{t}$; by Theorem 6.25, such a model would need to have $\aleph_2 \leq \mathfrak{p} < \mathfrak{t}$ and therefore $\mathfrak{c} \geq \aleph_3$. (Brendle has pointed out, however, that there is no a priori reason why a model of $\mathfrak{p} < \mathfrak{t}$ could not be produced by finite support iteration. This contrasts with the situation for producing a model with neither P-points nor Q-points; here finite support iteration has no chance because the Cohen reals it introduces make $\mathbf{cov}(\mathcal{B})$ large, and then Theorem 9.26 produces a selective ultrafilter.)

The "$\aleph_3$ barrier" is widely regarded as merely a technical problem. It has, however, resisted our efforts long enough to suggest that perhaps our inability to produce certain models is caused not by our technical deficiencies but by the non-existence of the models.

In the rest of this section, countable-support iterations will always be of the sort discussed above; that is, GCH will hold in the ground model, each step will be a proper forcing notion of cardinality at most $\mathfrak{c}$, and the length of the iteration will be ω_2. Thus, all cardinals are preserved. Furthermore, every real in the final model V_{ω_2} is already in some intermediate model V_α, $\alpha < \omega_2$. Thus, as with finite-support iterations, we can increase a characteristic $\|\mathbf{A}\|$ (but only up to $\aleph_2$) by cofinally often adding reals that solve $\mathbf{A}^\perp$. To prove independence results, we want to simultaneously keep some other characteristic small, and for this purpose there are a large number of powerful preservation theorems; see [96, 50, 44]. For example, in a countable-support proper iteration, if each $V_\alpha \cap {}^\omega\omega$ is a dominating family in $V_{\alpha+1}$ then $V \cap {}^\omega\omega$ is dominating in V_λ. In other words, if $V_{\alpha+1}$ never contains a real solving $\mathfrak{D}^\perp$ over V_α, then $\mathfrak{d}$ remains $\aleph_1$ in the final model.

11.1 Remark Zapletal [114] has shown that, under a strong large cardinal assumption (a proper class of measurable Woodin cardinals), many cardinal characteristics $\mathfrak{y}$ admit an optimal notion of forcing $P_\mathfrak{y}$ to make them large. Optimality means that, if $\mathfrak{x}$ is any tame characterstic and $\mathfrak{x} < \mathfrak{y}$ can be forced by some set forcing notion, then it is forced by $P_\mathfrak{y}$. The notion of tameness used here is somewhat more general than being the norm of a projective relation in that it permits some additional restrictions on the set $Y \subseteq A_+$ in Definition 4.1 of norms. All norms of Borel relations are tame, and so are, for example, $\mathfrak{p}$, $\mathfrak{t}$, and $\mathfrak{u}$, but not, for example, $\mathfrak{g}$.

Zapletal gives the following specific examples (among others) of optimal forcings for certain characteristics. See the following subsections for descriptions of these forcings. Cohen forcing is optimal for $\mathbf{cov}(\mathcal{B})$. Random forcing is optimal for $\mathbf{cov}(\mathcal{L})$. Sacks forcing is optimal for $\mathfrak{c}$. Laver forcing is optimal for $\mathfrak{b}$. Mathias forcing is optimal for $\mathfrak{h}$. Miller forcing is optimal for $\mathfrak{d}$.

11.3. Cohen Reals

The Cohen forcing poset, ${}^{<\omega}2$ ordered by reverse inclusion, adjoins a real $c : \omega \to 2$ (namely the union of the conditions in the generic set) that matches every chopped real (x, Π) from the ground model. Indeed, for each (x, Π) and each $n \in \omega$, the forcing conditions that agree with x on at least one interval of Π beyond n form a dense set in the ground model, so by genericity one of them must be included in c. Thus, a Cohen real solves $\mathbf{Cov}(\mathcal{B})^{\perp}$. (In fact, this characterizes Cohen reals.)

The usual way to iterate Cohen forcing is with finite support. Since the forcing poset is absolute, finite-support iteration and finite-support product are equivalent. The resulting model (when the ground model satisfies GCH) is usually called "the Cohen model" independently of the number λ of factors; for more precision, one says "the λ Cohen real model." This is the model used by Cohen [39] for his proof of the independence of GCH. Because of the ccc, every real in the Cohen model is already in the intermediate model generated by (the restriction of the generic filter to) some countable sub-product. Such a countable product (indeed, any countable atomless forcing notion) is equivalent to the single forcing ${}^{<\omega}2$. Thus any real in the Cohen model is in a submodel generated by a single Cohen real.

Since a Cohen real solves $\mathbf{Cov}(\mathcal{B})^{\perp}$, the λ Cohen real model (for any uncountable regular λ) has $\mathbf{cov}(\mathcal{B}) = \lambda = \mathfrak{c}$. It follows that all cardinals in the right half of Cichoń's diagram equal λ in this model. Furthermore, since

$$\mathbf{cov}(\mathcal{B}) \leq \mathfrak{r} \leq \mathfrak{u}, \mathfrak{i},$$

all these cardinals also equal $\mathfrak{c}$ in the Cohen model. (One can also see directly that a Cohen real splits all ground model reals, so $\mathfrak{r} = \lambda$.)

On the other hand, $\mathbf{non}(\mathcal{B}) = \aleph_1$ in the Cohen model, the set of ground model reals being non-meager. To prove this, we must show that every chopped real (x, Π) in the extension is matched by some ground model real. By our remarks above, we may assume that (x, Π) is in the forcing extension by a single Cohen real. In the ground model, we construct a real y such that for no condition $p \in {}^{<\omega}2$ and natural number n can p force "y does not agree with x on any interval of Π beyond n." Such a y is easily built by a recursion of length ω in which each step defines $y(k)$ for finitely many k and takes care of one pair (p, n). Taking care of (p, n) means to proceed as follows. Extend p to a condition q deciding a particular value for the restriction of x to the first interval $I \in \Pi$ whose left endpoint is greater than n and greater than all points already in the domain of y. Then extend y to agree with that restriction of x. Thus, q forces that y and x agree on an interval of Π beyond n, so p cannot force the contrary. (Note that this proof shows more than claimed. Not only the set of all ground model reals but any non-meager set in the ground model remains non-meager in a Cohen extension.)

In fact, $\mathbf{non}(\mathcal{B}) = \aleph_1$ holds in any model obtained by adjoining at least $\aleph_1$ Cohen reals to any ground model whatsoever. The reason is that $\aleph_1$ Cohen reals constitute a non-meager set.

From $\mathbf{non}(\mathcal{B}) = \aleph_1$, it immediately follows that all cardinals in the left half of Cichoń's diagram are $\aleph_1$. Furthermore, we have

$$\mathbf{non}(\mathcal{B}) \geq \mathfrak{b} \geq \mathfrak{h} \geq \mathfrak{t} \geq \mathfrak{p} \geq \mathfrak{m}, \quad \mathbf{non}(\mathcal{B}) \geq \mathfrak{s}, \quad \text{and} \quad \mathbf{non}(\mathcal{B}) \geq \mathfrak{e},$$

so all these cardinals are also $\aleph_1$ in the Cohen model.

Kunen showed [67, Theorem VIII.2.3] that $\mathfrak{a} = \aleph_1$ in the Cohen model. The idea is to construct, by transfinite induction in the ground model (where CH is available) a MAD family that remains MAD when one adds a Cohen real to the universe. It therefore remains MAD in any Cohen extension, since a failure to remain MAD would be witnessed by a single real. We omit the construction, since a similar one is given in the discussion of random reals below.

Finally, we cite from [18] the result that $\mathfrak{g} = \aleph_1$ in the Cohen model (or indeed in any model obtained by adjoining at least $\aleph_1$ Cohen reals to any model at all).

11.4. Random Reals

The notion of forcing to add one random real is the Boolean algebra of Borel sets modulo sets of Lebesgue measure zero (in any of $[0,1]$, $\mathbb{R}$, ${}^{\omega}2$, ${}^{\omega}\omega$; they are all equivalent). (Here and in general, when one refers to a Boolean algebra as a notion of forcing, one means the algebra minus its zero element.) Random forcing was introduced by Solovay [103, 104]. A generic G determines a real r, called "random," such that, if B is any Borel set in the ground model, then $r \in \tilde{B}$ if and only if $[B] \in G$. (For basic intervals B, this is the definition of r; for other B it is a theorem.) Thus, r solves $\mathbf{Cov}(\mathcal{L})^{\perp}$. This property characterizes random reals.

Although random forcing can be iterated with finite support or with countable support (being ccc and therefore proper), the most common way to add many random reals uses a large measure algebra, namely the algebra of Borel subsets modulo measure zero sets in ${}^{I}2$ for large I. The measure here is the product measure induced by the uniform measure on 2. This forcing adds a random function $f : I \to 2$ whose restrictions to countable subsets of I in V amount to random reals. One often starts with a ground model satisfying GCH, takes $I = \lambda \times \omega$, and regards the forcing as adding the λ random reals $r_\alpha : \omega \to 2 : n \mapsto f(\alpha, n)$. Any real in this λ random reals model is in the submodel generated by countably many of the r_α, and this submodel is equivalent to one obtained by adjoining a single random real to the ground model.

Because a random real solves $\mathbf{Cov}(\mathcal{L})^{\perp}$, the λ random reals model has (for uncountable regular λ) $\mathbf{cov}(\mathcal{L}) = \lambda = \mathfrak{c}$. Therefore, all the cardinals in

the top row of Cichoń's diagram equal $\mathfrak{c}$ in this model, and so do $\mathfrak{r}$, $\mathfrak{u}$, and $\mathfrak{i}$.

On the other hand, $\mathfrak{d}$ and $\mathbf{non}(\mathcal{L})$ are both $\aleph_1$, as in the ground model. More generally, if uncountably many random reals are added (with the usual measure algebra forcing) to any ground model, then in the extension $\mathfrak{d}$ will have the same value as in the ground model while $\mathbf{non}(\mathcal{L})$ will be $\aleph_1$. The former follows from the fact that all reals in a random extension are majorized by ground model reals. The latter follows from the fact that any $\aleph_1$ of the added random reals form a set of positive outer measure. (A measure-zero Borel set, or rather its code, depends on only countably many of the added random reals; all the rest of the added random reals, being random over an intermediate model containing the code, must be outside that Borel set.)

It follows that all the cardinals in the middle and bottom rows of Cichoń's diagram are $\aleph_1$, and therefore so are $\mathfrak{s}$, $\mathfrak{e}$, $\mathfrak{g}$, $\mathfrak{h}$, $\mathfrak{t}$, $\mathfrak{p}$, and $\mathfrak{m}$.

Finally, we show, adapting Kunen's proof for the Cohen model, that $\mathfrak{a} = \aleph_1$ in the random model. Since every real in the random model is in a submodel that can be generated by a single random real, it suffices to construct a family $\mathcal{A}$ in the ground model that is MAD and remains so when one random real is adjoined to the universe. We proceed as follows in the ground model. Because the forcing notion to adjoin one random real has cardinality $\mathfrak{c}$ and satisfies the ccc, there are only $\mathfrak{c} = \aleph_1$ essentially different names for subsets of ω; enumerate them as $\langle x_\alpha : \alpha < \omega_1 \rangle$. We construct $\mathcal{A}$ by a recursion of length $\aleph_1$, starting with a partition of ω into $\aleph_0$ infinite pieces, and adding one set a_α to $\mathcal{A}$ at each step. This set will be chosen so as to be almost disjoint from the previous a_β's and to have infinite intersection with the denotation (with respect to every generic set) of x_α (unless some earlier a_β already does or x_α is finite). That will ensure that $\mathcal{A} = \{a_\alpha : \alpha < \omega_1\}$ remains MAD in the random extension. Let $[B]$ be the Boolean truth value of "x_α is not almost included in the union of finitely many $\check{a}_\beta$ with $\beta < \check{\alpha}$." We shall make sure that the truth value of "$x_\alpha \cap \check{a}_\alpha$ is infinite" is at least $[B]$. Equivalently, since we are dealing with a measure algebra, we shall make sure that for every n the Boolean truth value of "$x_\alpha \cap \check{a}_\alpha$ has a member $> n$" intersected with $[B]$ has measure at least $\mu[B] - \frac{1}{n}$. And of course we must ensure that a_α is almost disjoint from the earlier a_β's. We define a_α as follows.

Let the earlier a_β's be enumerated in an ω-sequence as a'_n. We shall construct a_α by a recursion of length ω, adding finitely many elements at each stage, and ensuring at stage n that the measure requirement at the end of the last paragraph is satisfied for n. To ensure almost disjointness, we shall not add any elements of a'_k after stage k. We now describe stage n. Let $v = \bigcup_{k<n} a'_k$, whose elements are no longer to be added to a_α. With truth value at least $[B]$, $x_\alpha - v$ is infinite. So $[B]$ is the Boolean sum of (countably many) pairwise incompatible conditions $[B_i]$ each forcing a specific value for

the first element z of $x_\alpha - v$ that is $> n$. Since the measures of all the $[B_i]$ add up to the measure of $[B]$, finitely many of them come to within $\frac{1}{n}$ of that total. Put the corresponding finitely many z's into a_α. This completes the construction of a_α; we omit the routine verification that it does what was required.

11.5. Sacks Reals

The Sacks forcing notion, introduced in [95] and also called perfect set forcing, consists of perfect subtrees of ${}^{<\omega}2$, i.e., nonempty subtrees that have branching beyond each node; the partial ordering is inclusion. This is a proper forcing that adjoins a real s, namely the unique common path through all the trees in a generic set G.

The forcing extension $V[s]$ enjoys the *Sacks property*: For every function $f : \omega \to V$ in $V[s]$, there is a function $g : \omega \to V$ in V such that for all $n \in \omega$ we have $f(n) \in g(n)$ and $|g(n)| \leq 2^n$. (Although 2^n emerges naturally from the proof, we could, as in Remark 5.15, replace 2^n by any function tending to ∞.) To prove this, suppose we are given a name $\dot{f}$ for f and a condition p. Working in V, we prune the tree p in ω steps to produce a perfect subtree q forcing that a certain g is as required; by genericity, this will suffice. Begin by choosing $p_0 \leq p$ deciding a specific value for $\dot{f}(0)$. This value will be the unique element of $g(0)$. The first branching node a of p_0 will be the first branching node of the final q; i.e., neither a nor its immediate successors $a^\frown\langle 0\rangle$ and $a^\frown\langle 1\rangle$ will be pruned away later. Regard p_0 as the union of two perfect subtrees, one consisting of the nodes comparable with $a^\frown\langle 0\rangle$ and the other of the nodes comparable with $a^\frown\langle 1\rangle$. In each of these, find a perfect subtree deciding $\dot{f}(1)$ (possibly different decisions for the two subtrees). Reuniting these two subtrees, we get a perfect subtree p_1 of p_0, where a is still a branching node, and such that p_1 forces $\dot{f}(1)$ to have one of just two specific values. Those values will be the elements of $g(1)$. All later steps will preserve the two second-level branching nodes of p_1. Regard p_1 as the union of four perfect subtrees, one through each of the immediate successors of those nodes. Shrink each of the four to decide a (possibly different) value for $\dot{f}(2)$; and reunite them to get p_2. Continuing in this way, we finally obtain a tree $q = \bigcap_{n\in\omega} p_n$ that is perfect because we retain more and more branching as the construction progresses. q is an extension of p forcing each $\dot{f}(n)$ to have one of 2^n specific values known in V, so the desired g exists in V.

This sort of construction, repeatedly pruning a tree but retaining more and more branching, is referred to as *fusion*. It can also be used to prove that adjoining a Sacks real produces a minimal extension in the sense that if $x \in V[s] - V$ is a set of ordinals then $V[x] = V[s]$.

The usual way to iterate Sacks forcing is with countable support for $\aleph_2$ steps, starting with a model of GCH. The resulting model is often called

the *Sacks model*. Properness of Sacks forcing implies that cardinals are preserved. Furthermore, the Sacks model has the Sacks property, because this property is preserved by countable-support proper iterations; see [5, Section 6.3.F], [96, Sections VI.1–2], or [50]. It follows, by the dual of Theorem 5.14, that $\mathbf{cof}(\mathcal{L}) = \aleph_1$ in the Sacks model. Therefore, all cardinals in Cichoń's diagram as well as $\mathfrak{d}$, $\mathfrak{e}$, $\mathfrak{b}$, $\mathfrak{g}$, $\mathfrak{s}$, $\mathfrak{h}$, $\mathfrak{t}$, $\mathfrak{p}$, and $\mathfrak{m}$ are equal to $\aleph_1$ in this model. Baumgartner and Laver [11] showed that selective ultrafilters in the ground model, which exist since GCH holds there, generate ultrafilters in the Sacks model. (In fact, the same is true of P-points.) Therefore the Sacks model has $\mathfrak{u} = \mathfrak{r} = \aleph_1$.

Spinas has shown (private communication) that the Sacks model satisfies $\mathfrak{a} = \aleph_1$. In outline, his argument is as follows. By general properties of Souslin proper forcing (see [59], [50, Section 7], and [100]), it suffices to find, in the ground model, a MAD family $\mathcal{A}$ that remains MAD in the extension obtained by iterating Sacks forcing for ω_1 steps with countable support. List in an ω_1-sequence all pairs (τ, p) where p is a condition in this iteration and τ is a name forced by p to denote an infinite subset of ω. We define the desired $\mathcal{A} = \{A_\alpha : \alpha < \omega_1\}$ by induction in the ground model, ensuring at step α that for the α^{th} pair (τ, p) some extension of p either forces (a) "$\tau \cap \check{A}_\alpha$ is infinite" or forces (b) "τ is almost included in $\check{A}_{\beta_1} \cup \cdots \cup \check{A}_{\beta_r}$" for some finitely many $\beta_1, \ldots, \beta_r < \alpha$. Either way, p cannot force $\mathcal{A} \cup \{\tau\}$ to be almost disjoint with $\tau \notin \mathcal{A}$, so the maximality is preserved. To define A_α, assume the previous A_β's are already defined; modifying them finitely and re-numbering them (see the proof of Proposition 8.4), we can pretend that the ω we are working in is $\omega \times \omega$ and that these earlier A_β's are the columns $\{n\} \times \omega$. We can also assume that p forces τ to meet infinitely many of these columns, as otherwise we already have alternative (b) above. We shall take A_α to be $\{(a, b) : b < f(a)\}$ for a suitably large $f : \omega \to \omega$. Then clearly A_α is almost disjoint from the previous A_β's (the columns). To obtain alternative (a) and thus complete the proof, we need only choose f large enough. Specifically, use the name τ to produce a name D for the set of n such that the n^{th} column meets τ and a name g for a function $D \to \omega$ such that p forces "D is infinite and, for each $d \in D$, τ contains an element (d, b) with $b < g(d)$." Then, thanks to the Sacks property, p also forces "some ground model function $f : \omega \to \omega$ majorizes g." Choosing an extension of p that decides what f is, we obtain alternative (a), and the proof is complete.

Finally, Eisworth and Shelah (unpublished) have shown that $\mathfrak{i} = \aleph_1$ in the Sacks model.

For many cardinal characteristics, a recent result of Shelah gives a uniform reason why they are $\aleph_1$ in the Sacks model. Shelah has shown that a countable-support proper iteration of forcings that individually add no reals can, at limit stages of cofinality ω, introduce Sacks reals. But there are numerous iteration theorems (see [96, 50, 44]) saying that certain properties

of a ground model will be unchanged by a countable-support proper forcing iteration provided they are unchanged by the individual steps. These properties, then, are not changed by adding Sacks reals.

Another explanation for the smallness of many cardinal characteristics in the Sacks model is the fact that countable support iteration of Sacks forcing is the optimal forcing for increasing $\mathfrak{c}$, in the sense of Zapletal [114]; see Remark 11.1. Thus, all tame cardinal characteristics that can be forced to remains small when $\mathfrak{c}$is increased by some set forcing in fact remain small in the Sacks model, provided there is a proper class of measurable Woodin cardinals.

11.6. Hechler Reals

Introduced by Hechler [56] for his proof of Theorem 2.5, Hechler forcing, also called dominating forcing, is the set of pairs (s, f) where $s \in {}^{<\omega}\omega$ and $f \in {}^{\omega}\omega$. (The "meaning" of (s, f) is that the generic real in ${}^{\omega}\omega$ has s as an initial segment and thereafter majorizes f.) The ordering puts $(s', f') \leq (s, f)$ if s is an initial segment of s', $f \leq f'$, and $s'(n) \geq f(n)$ for all $n \in \mathrm{dom}(s') - \mathrm{dom}(s)$. This forcing satisfies ccc; in fact it is σ-centered, since any finitely many conditions with the same first component have a lower bound. A Hechler-generic set G determines a function $g : \omega \to \omega$, namely the union of the first components of the members of G. Such a g is called a Hechler real. Genericity implies that it dominates all ground model functions $\omega \to \omega$, i.e., g solves $\mathfrak{D}$. ("Dominating real" is sometimes used as a synonym for "Hechler real" and sometimes to mean any real that dominates all ground model reals.) Replacing each of the values of g by its parity, we obtain a Cohen real, g mod 2.

By "the Hechler model" we mean the result of a finite support iteration of Hechler forcing over a model of GCH, where the number of steps is some regular uncountable cardinal λ. One can also consider countable-support iterations (for up to ω_2 stages, as usual) but we shall not do so here. Hechler's original use of Hechler forcing [56] amounted to a combination of finite-support iteration and product constructions.

Since a Hechler real solves $\mathfrak{D}$ and its parity solves $\mathbf{Cov}(\mathcal{B})^{\perp}$, the Hechler model satisfies $\mathbf{cov}(\mathcal{B}) = \mathfrak{b} = \lambda = \mathfrak{c}$. By Theorem 5.6, it satisfies $\mathbf{add}(\mathcal{B}) = \mathfrak{c}$. Thus, in this model, the cardinals in the second through fourth columns of Cichoń's diagram equal $\mathfrak{c}$. Those in the first column, on the other hand, equal $\aleph_1$ since this forcing adds no random reals [5, second model in 7.6.9]. Since $\mathfrak{b}$ is large, so are $\mathfrak{r}$, $\mathfrak{u}$, $\mathfrak{a}$, and $\mathfrak{i}$. Baumgartner and Dordal showed in [10] that $\mathfrak{s}$ in the Hechler model is $\aleph_1$, and therefore so are $\mathfrak{h}$, $\mathfrak{t}$, $\mathfrak{p}$, and $\mathfrak{m}$. Brendle [30, Theorem 10.4] showed that $\mathfrak{e} = \aleph_1$ in the Hechler model.

The value of $\mathfrak{g}$ in the Hechler model should be $\aleph_1$. Brendle has shown (private communication) that it is $\aleph_1$ if Hechler forcing is iterated for only ω_2 steps. Shelah has sketched a proof that it is $\aleph_1$ in general, but so far as

I know this proof has yet to be written down carefully and checked (private communication from Eisworth).

Pawlikowski [83] showed that, although $\mathbf{add}(\mathcal{B})$ is large in the Hechler model, adjoining a single Hechler real to the ground model does not produce any real solving $\mathbf{Cof}(\mathcal{B})$. Such a real appears, however, when two Hechler reals are added iteratively. This last fact follows from part 1 of Theorem 5.6, which says that $\mathbf{Cof}(\mathcal{B})$ admits a morphism from a sequential composition of two relations each of which is solved when a single Hechler real is adjoined.

11.7. Laver Reals

Conditions in Laver forcing are trees $p \subseteq {}^{<\omega}\omega$ in which there is a node s, called the stem, such that all nodes are comparable with s and every node beyond s has infinitely many immediate successors. (So, starting at the root of p, one finds no branching until one reaches s and then infinite branching everywhere thereafter.) The ordering is inclusion. A generic set G determines a function $g : \omega \to \omega$ called a Laver real, namely the union of the stems of all the conditions in G, or equivalently the unique common path through all members of G. Laver forcing is proper. Genericity implies that a Laver real dominates all ground model functions $\omega \to \omega$.

The Laver model is obtained by an ω_2-stage countable-support iteration of Laver forcing over a model of GCH. (Historically, Laver forcing and countable-support iteration were introduced together in [71]. For the purpose of that paper, producing a model of the Borel conjecture, one needs to dominate all ground model reals, but one must not introduce Cohen reals, so neither Hechler forcing nor a finite support iteration can be used.) Since a Laver real solves $\mathfrak{D}$, the Laver model has $\mathfrak{b} = \aleph_2 = \mathfrak{c}$. It follows that the cardinals in all but the left column and bottom row of Cichon's diagram are $\aleph_2$, and so are $\mathfrak{r}$, $\mathfrak{i}$, $\mathfrak{u}$, and $\mathfrak{a}$.

Like Hechler forcing, Laver forcing even when iterated does not produce random reals, but unlike Hechler forcing it does not produce Cohen reals either. In fact, the set of ground model reals does not have measure zero in the extension. See [5, Section 7.3.D] for proofs of these facts. It follows that $\mathbf{cov}(\mathcal{L})$ and $\mathbf{non}(\mathcal{L})$ are both $\aleph_1$ in the Laver model, and therefore so are $\mathbf{add}(\mathcal{L})$, $\mathbf{add}(\mathcal{B})$, $\mathbf{cov}(\mathcal{B})$, $\mathfrak{e}$, $\mathfrak{s}$, $\mathfrak{h}$, $\mathfrak{t}$, $\mathfrak{p}$, and $\mathfrak{m}$.

Finally, Brendle has pointed out that the proof of $\mathfrak{g} = \mathfrak{c}$ for the Miller model [25, 18] applies also to the Laver model. The same argument was used for a slightly different purpose in [44, Lemma 4.3.5].

11.8. Mathias Reals

Mathias forcing was described in Remark 7.8. It consists of pairs (s, A) with $s \in [\omega]^{<\omega}$ and $A \in [\omega]^{\omega}$ ("meaning" that the generic subset of ω has s as an initial segment and otherwise is included in A). The ordering, defined

in Remark 7.8, is based on this meaning. A generic filter G determines an infinite subset X of ω called a Mathias real, namely the union of the first components of all the members of G. Mathias forcing was used in [73] and was studied in detail in [75].

The essential property of a Mathias real X is that, if $\mathcal{D} \subseteq [\omega]^\omega$ is any dense open family in the ground model, then X is included in some member of $\mathcal{D}$. To prove this, consider an arbitrary condition (s, A) and use the density of $\mathcal{D}$ to extend it to (s, A') with $A' \in \mathcal{D}$. Then (s, A') forces the generic real X to be almost included in A' and therefore included in some member of $\mathcal{D}$ since dense open families are closed under finite modifications.

By the Mathias model, we mean the result of an ω_2-stage countable-support iteration of Mathias forcing over a model of GCH. The preceding paragraph together with a reflection argument implies that $\mathfrak{h} = \aleph_2 = \mathfrak{c}$ in this model. Specifically, given any $\aleph_1$ dense open families $\mathcal{D}_\xi$, we can find a common member as follows. Using the $< \aleph_2$-chain condition, we obtain an $\alpha < \omega_2$ (in fact an ω_1-closed unbounded set of such α's) such that each $\mathcal{D}_\xi \cap V_\alpha$ is a member of V_α and is a dense open set in the sense of V_α. Then the Mathias real X adjoined in going from V_α to $V_{\alpha+1}$ has, by the preceding paragraph, supersets in each $\mathcal{D}_\xi \cap V_\alpha$ and therefore belongs to each $\mathcal{D}_\xi$.

Because $\mathfrak{h}$ is large, so are $\mathfrak{b}$, $\mathfrak{g}$, $\mathfrak{s}$, $\mathfrak{r}$, $\mathfrak{d}$, $\mathfrak{a}$, $\mathfrak{u}$, $\mathfrak{i}$, and **non** and **cof** of both category and measure.

On the other hand, both $\mathbf{cov}(\mathcal{B})$ and $\mathbf{cov}(\mathcal{L})$ are only $\aleph_1$ because neither Cohen nor random reals are added. See [5, Section 7.4.A] for the proof. It follows that $\mathbf{add}(\mathcal{L})$, $\mathbf{add}(\mathcal{B})$, $\mathfrak{e}$, $\mathfrak{t}$, $\mathfrak{p}$, and $\mathfrak{m}$ are also $\aleph_1$.

11.9. Miller Reals

The Miller forcing notion, introduced in [81], consists of superperfect trees (also called rational perfect trees), i.e., subtrees of ${}^{<\omega}\omega$ in which beyond every node there is one with infinitely many immediate successors. The order is inclusion. As with other such tree forcings, this is proper, and a generic set G determines a real $g : \omega \to \omega$, namely the union of the stems of the members of G or equivalently the unique path through all members of G. It is sometimes convenient to replace the Miller forcing notion with the isomorphic one in which the nodes of the trees are strictly increasing finite sequences from ω. Then the generic g is an increasing map $\omega \to \omega$, the enumeration of an infinite $X \subseteq \omega$. Either g or X can be called a Miller real or a superperfect real.

The Miller model is the result of an ω_2-stage countable-support iteration of Miller forcing over a model of GCH. It is shown in [25, 18] that a Miller real X has supersets in all groupwise dense families from the ground model. This and a reflection argument show, just as in the discussion of Mathias forcing above, that $\mathfrak{g} = \aleph_2 = \mathfrak{c}$ in the Miller model. It follows that $\mathfrak{d}$, $\mathfrak{i}$, $\mathbf{cof}(\mathcal{B})$, and $\mathbf{cof}(\mathcal{L})$ are also $\aleph_2$.

On the other hand, it is shown in [5, 7.3.E] that both $\mathbf{non}(\mathcal{L})$ and $\mathbf{non}(\mathcal{B})$ are $\aleph_1$ in the Miller model. Therefore so are $\mathfrak{s}$, $\mathfrak{e}$, $\mathfrak{b}$, $\mathfrak{h}$, $\mathfrak{t}$, $\mathfrak{p}$, $\mathfrak{m}$, and all the cardinals in Cichoń's diagram except $\mathfrak{d}$ and the two cofinalities.

It is also shown in [25] that every P-point in the ground model generates an ultrafilter in the Miller model. Therefore this model satisfies $\mathfrak{u} = \mathfrak{r} = \aleph_1$.

Finally, the proof that $\mathfrak{a} = \aleph_1$ in the Sacks model can, as Spinas pointed out, be transferred to the Miller model with only a minor modification. At the end of the proof, instead of using the Sacks property (which fails in the Miller model), one uses the fact that the ground model is an unbounded family in ${}^{\omega}\omega$ to show that p forces the function g in the extension to be majorized on an infinite subset of D by an f from the ground model. Another proof that $\mathfrak{a} = \aleph_1$ in the Miller model is given in [43, Prop. 8.24]. Eisworth pointed out (private communication) that the same argument applies to the Sacks model.

11.10. Summary of Iterated Forcing Results

Table 4 summarizes the preceding results concerning the values of cardinal characteristics in the iterated forcing models described above. Remember that in the countable-support models, i.e., in the Sacks, Laver, Mathias, and Miller columns of the table, $\mathfrak{c}$ is just $\aleph_2$.

Figure 1 is a Hasse diagram of the main cardinal characteristics discussed in this chapter, except for the characteristics of the measure and category ideals. A line joining two characteristics in the figure means that the lower one is provably $\leq$ the upper one.

11.11. Other Forcing Iterations

The preceding sections cover only a few of the many kinds of iterated forcing, over models of GCH, that have been used in the theory of cardinal characteristics. There are other kinds of reals that one can adjoin, for example infinitely equal reals, Prikry-Silver reals, Matet reals, Grigorieff reals. Except for Matet reals, which are defined in the last section of [18], these and many others can be found in [5] or [61]. Most of these forcing notions do not satisfy the ccc, so they are iterated with countable support and therefore one enlarges $\mathfrak{c}$ only to $\aleph_2$.

Two models constructed in [24, Sections 2 and 6] involve iterating a forcing that looks less natural than those discussed in the preceding sections or mentioned in the preceding paragraph, but we list their cardinal characteristics here because they are somewhat unusual, e.g., $\mathfrak{u} < \mathfrak{s}$. Both models have $\mathfrak{u} = \aleph_1$ and therefore all of $\mathfrak{r}$, $\mathfrak{e}$, $\mathfrak{b}$, $\mathfrak{h}$, $\mathfrak{t}$, $\mathfrak{p}$, $\mathfrak{m}$, and the covering numbers and additivities for both category and measure are $\aleph_1$. On the other hand, they have $\mathfrak{s} = \mathfrak{c} = \aleph_2$ and therefore all of $\mathfrak{d}$, $\mathfrak{i}$, and the uniformities and cofinalities of both measure and category are $\aleph_2$. (See [24, Theorem 5.2 and

	MA	Cohen	Random	Sacks	Hechler	Laver	Mathias	Miller
$\mathfrak{a}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$
$\mathfrak{b}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$
$\mathfrak{d}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$
$\mathfrak{e}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathfrak{g}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$
$\mathfrak{h}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$
$\mathfrak{i}$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$
$\mathfrak{m}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathfrak{p}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathfrak{r}$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$
$\mathfrak{s}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$
$\mathfrak{t}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathfrak{u}$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$
$\mathbf{add}(\mathcal{L})$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathbf{cov}(\mathcal{L})$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathbf{non}(\mathcal{L})$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$
$\mathbf{cof}(\mathcal{L})$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$
$\mathbf{add}(\mathcal{B})$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathbf{cov}(\mathcal{B})$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$	$\aleph_1$	$\aleph_1$
$\mathbf{non}(\mathcal{B})$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$
$\mathbf{cof}(\mathcal{B})$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\aleph_1$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$	$\mathfrak{c}$

Table 4: Cardinal Characteristics in Iterated Forcing Models

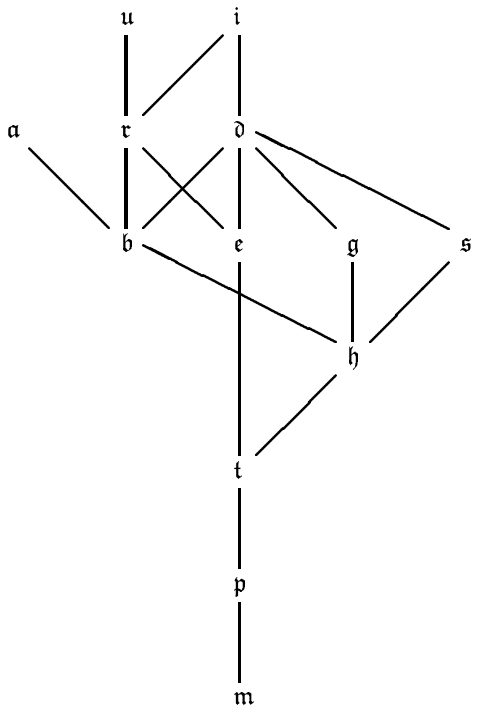

Figure 1: Hasse Diagram of Combinatorial Characteristics

the end of Section 6].) The first of the two models, the one designed to satisfy NCF, has $\mathfrak{g} = \aleph_2$, as was shown in [23, Theorem 2]. This model also has $\mathfrak{a} = \aleph_1$ by the same Souslin-forcing argument used above for Sacks and Miller reals. The second model, the one with simple $P_{\aleph_1}$-points and simple $P_{\aleph_2}$-points, does not satisfy NCF and therefore must have $\mathfrak{g} = \aleph_1$. I do not know the value of $\mathfrak{a}$ in this second model.

A frequently useful sort of iterated forcing is one where two or more different forcings are used alternately. Numerous examples of this can be found in [5, Chapter 7]. Dow's paper [42] describes, among other things, the models obtained by alternating Laver and Mathias forcings; it turns out to make a difference which forcing one uses at limit ordinals.

Dordal [40] uses a mixed-support iteration of Mathias forcings. Viewing Mathias forcing as a two-step iteration, where one first adjoins an ultrafilter generically and then does Mathias forcing with respect to this ultrafilter (see Remark 7.8), he defines an iteration in which the adjunctions of ultrafilters are done with countable support while the interleaved Mathias forcings with respect to these ultrafilters are done with finite support.

All the preceding forcing iterations began with a ground model satisfying GCH. Thus, all cardinal characteristics are $\aleph_1$ in the ground model, and the iterations are designed to raise some characteristics while leaving others small. An alternative approach is to begin with a model where $\mathfrak{c}$ and some other characteristics are already large (e.g., a model of MA) and to do an iteration, usually of small length, to lower some characteristics while leaving others large. We briefly describe two examples; many more can be found in [5, Chapter 7].

Start with a model of MA + $\neg$CH (so all the characteristics we have discussed are large) and adjoin $\aleph_1$ random reals. Since the ${}^{\omega}\omega$ of a random extension is dominated by that of the ground model, we obtain a model where $\mathfrak{b}$ has the same large value that it had in the ground model of MA. On the other hand, $\mathfrak{s}$ is only $\aleph_1$ in the extension, and in fact so is $\mathbf{non}(\mathcal{L})$, since the $\aleph_1$ random reals form a set of positive outer measure and thus a splitting family. This proof for the consistency of $\mathfrak{b} > \mathfrak{s}$, due to Balcar and Simon, is easier than either of the ones obtainable from Table 4 (the Hechler and Laver models).

Another application of forcing over a model with large continuum is the construction in [26] of a model where $\mathfrak{u} < \mathfrak{d}$. This model, which predates the ones in [24, 25] that establish the stronger $\mathfrak{u} < \mathfrak{g}$, has the advantage that $\mathfrak{u} < \mathfrak{d}$ can be any prescribed uncountable regular cardinals. It begins with a Cohen model, where $\mathfrak{d}$ has the desired value, and extends it by a finite-support iteration of Mathias forcings with respect to carefully chosen ultrafilters. The length of the iteration is the prescribed $\mathfrak{u}$. The easier part of "carefully chosen" is that each ultrafilter contains the previously adjoined Mathias reals, so that the sequence of Mathias reals is almost decreasing and generates an ultrafilter in the final model. Thus $\mathfrak{u}$ will be small. The

hard part of "carefully chosen," which we omit here, is to keep $\mathfrak{d}$ large.

11.12. Adding One Real

In this subsection, we briefly summarize some results about the effect on cardinal characteristics of adjoining one real to a model of ZFC. Here the ground model will not satisfy CH, for the single-real forcings we consider would preserve CH and leave all characteristics at $\aleph_1$. We consider situations where some characteristics are large in the ground model and we ask how adding a single real affects them. Most of what is known about this concerns the cardinals from Cichoń's diagram. The results summarized here are from [29, 33, 7, 38, 83].

Adding a Cohen real to any model of ZFC makes $\mathbf{add}(\mathcal{L}) = \mathbf{cov}(\mathcal{L}) = \aleph_1$ and $\mathbf{non}(\mathcal{L}) = \mathbf{cof}(\mathcal{L}) = \mathfrak{c}$. The values of $\mathbf{add}(\mathcal{B})$, $\mathbf{non}(\mathcal{B})$, and $\mathfrak{b}$ in the extension are the $\mathbf{add}(\mathcal{B})$ of the ground model, and dually the values of $\mathbf{cof}(\mathcal{B})$, $\mathbf{cov}(\mathcal{B})$, and $\mathfrak{d}$ in the extension are the $\mathbf{cof}(\mathcal{B})$ of the ground model.

Adding a random real produces a value for $\mathbf{cov}(\mathcal{L})$ that is at least $\max\{\mathbf{cov}(\mathcal{L}), \mathfrak{b}\}$ of the ground model, and may be strictly larger. Dually, the extension's $\mathbf{non}(\mathcal{L})$ is at most $\min\{\mathbf{non}(\mathcal{L}), \mathfrak{d}\}$ and may be strictly smaller. Except for $\mathbf{cov}(\mathcal{L})$ and $\mathbf{non}(\mathcal{L})$, the cardinals in Cichoń's diagram remain unchanged.

Adding one Hechler real makes all cardinals in the left half of Cichoń's diagram $\aleph_1$ and all those in the right half $\mathfrak{c}$. It also makes $\mathfrak{a} = \aleph_1$.

Adding one Laver or Mathias real makes the $\mathfrak{d}$ of the extension $\aleph_1$. These forcings also collapse $\mathfrak{c}$ to $\mathfrak{h}$. Since $\mathfrak{h} \leq \mathfrak{d}$, it follows that a two-step iteration of these forcings produces a model of CH.

Bibliography

[1] Bohuslav Balcar, Jan Pelant, and Petr Simon. The space of ultrafilters on N covered by nowhere dense sets. *Fund. Math.*, 110:11–24, 1980.

[2] Bohuslav Balcar and Petr Simon. On minimal π-character of points in extremally disconnected compact spaces. *Topology Appl.*, 41:133–145, 1991.

[3] Tomek Bartoszyński. Additivity of measure implies additivity of category. *Trans. Amer. Math. Soc.*, 281:209–213, 1984.

[4] Tomek Bartoszyński. Combinatorial aspects of measure and category. *Fund. Math.*, 127:225–239, 1987.

[5] Tomek Bartoszyński and Haim Judah. *Set Theory, On the Structure of the Real Line.* A K Peters, 1995.

[6] Tomek Bartoszyński, Haim Judah, and Saharon Shelah. The Cichoń diagram. *J. Symbolic Logic*, 58:401–423, 1993.

[7] Tomek Bartoszyński, Andrzej Rosłanowski, and Saharon Shelah. Adding one random real. *J. Symbolic Logic*, 61:80–90, 1996.

[8] James Baumgartner. Almost-disjoint sets, the dense-set problem, and the partition calculus. *Ann. Math. Logic*, 9:401–439, 1976.

[9] James Baumgartner. Iterated forcing. In A. R. D. Mathias, editor, *Surveys in Set Theory*, London Math. Soc. Lecture Notes, pages 1–59. Cambridge Univ. Press, 1983.

[10] James Baumgartner and Peter Dordal. Adjoining dominating functions. *J. Symbolic Logic*, 50:94–101, 1985.

[11] James Baumgartner and Richard Laver. Iterated perfect-set forcing. *Ann. Math. Logic*, 17:271–288, 1979.

[12] Murray Bell. On the combinatorial principle $P(c)$. *Fund. Math.*, 114:149–157, 1981.

[13] Andreas Blass. *Orderings of Ultrafilters*. PhD thesis, Harvard Univ., 1970.

[14] Andreas Blass. The Rudin-Keisler ordering of P-points. *Trans. Amer. Math. Soc.*, 179:145–166, 1973.

[15] Andreas Blass. Near coherence of filters, I: Cofinal equivalence of models of arithmetic. *Notre Dame J. Formal Logic*, 27:579–591, 1986.

[16] Andreas Blass. Ultrafilters related to Hindman's finite unions theorem and its extensions. In S. Simpson, editor, *Logic and Combinatorics*, volume 65 of *Contemp. Math.*, pages 89–124. Amer. Math. Soc., 1987.

[17] Andreas Blass. Selective ultrafilters and homogeneity. *Ann. Pure Appl. Logic*, 38:215–255, 1988.

[18] Andreas Blass. Applications of superperfect forcing and its relatives. In J. Steprāns and S. Watson, editors, *Set Theory and its Applications*, volume 1401 of *Lecture Notes in Math.*, pages 18–40. Springer-Verlag, 1989.

[19] Andreas Blass. Groupwise density and related cardinals. *Arch. Math. Logic*, 30:1–11, 1990.

[20] Andreas Blass. Simple cardinal characteristics of the continuum. In H. Judah, editor, *Set Theory of the Reals*, volume 6 of *Israel Math. Conf. Proc.*, pages 63–90. Amer. Math. Soc., 1993.

[21] Andreas Blass. Cardinal characteristics and the product of countably many infinite cyclic groups. *J. Algebra*, 169:512–540, 1994.

[22] Andreas Blass. Reductions between cardinal characteristics of the continuum. In T. Bartoszyński and M. Scheepers, editors, *Set Theory (Annual Boise Extravaganza in Set Theory conference, 1992–94)*, volume 192 of *Contemp. Math.*, pages 31–49. Amer. Math. Soc., 1996.

[23] Andreas Blass and Claude Laflamme. Consistency results about filters and the number of inequivalent growth types. *J. Symbolic Logic*, 54:50–56, 1989.

[24] Andreas Blass and Saharon Shelah. There may be simple $P_{\aleph_1}$ and $P_{\aleph_2}$ points and the Rudin-Keisler ordering may be downward directed. *Ann. Pure Appl. Logic*, 33:213–243, 1987.

[25] Andreas Blass and Saharon Shelah. Near coherence of filters III: A simplified consistency proof. *Notre Dame J. Formal Logic*, 30:530–538, 1989.

[26] Andreas Blass and Saharon Shelah. Ultrafilters with small generating sets. *Israel J. Math.*, 65:259–271, 1989.

[27] Andreas Blass and Saharon Shelah. Dominating, unsplit, and Ramsey reals. to appear.

[28] David Booth. Ultrafilters on a countable set. *Ann. Math. Logic*, 2:1–24, 1970.

[29] Jörg Brendle. Combinatorial properties of classical forcing notions. *Ann. Pure Appl. Logic*, 73:143–170, 1995.

[30] Jörg Brendle. Evasion and prediction — the Specker phenomenon and Gross spaces. *Forum Math.*, 7:513–541, 1995.

[31] Jörg Brendle. Strolling through paradise. *Fund. Math.*, 148:1–25, 1995.

[32] Jörg Brendle. Evasion and prediction. III. constant prediction and dominating reals. *J. Math. Soc. Japan*, 55:101–115, 2003.

[33] Jörg Brendle, Haim Judah, and Saharon Shelah. Combinatorial properties of Hechler forcing. *Ann. Pure Appl. Logic*, 58:185–199, 1992.

[34] Jörg Brendle and Saharon Shelah. Evasion and prediction. II. *J. London Math. Soc. (2)*, 53:19–27, 1996.

[35] Maxim Burke. A proof of Hechler's theorem on embedding $\aleph_1$-directed sets cofinally into $(\omega^\omega, <^*)$. *Arch. Math. Logic*, 36:399–403, 1997.

[36] R. Michael Canjar. On the generic existence of special ultrafilters. *Proc. Amer. Math. Soc.*, 110:233–241, 1990.

[37] Georg Cantor. Über eine Eigenschaft des Inbegriffes aller reellen algebraischen Zahlen. *Crelles Journal f. Mathematik*, 77:258–262, 1874.

[38] Jacek Cichoń and Janusz Pawlikowski. On ideals of subsets of the plane and on Cohen reals. *J. Symbolic Logic*, 51:560–569, 1986.

[39] Paul Cohen. The independence of the continuum hypothesis. *Proc. Nat. Acad. Sci. U. S. A.*, 50:1143–1148, 1963.

[40] Peter Dordal. A model in which the base-matrix tree cannot have cofinal branches. *J. Symbolic Logic*, 52:651–664, 1987.

[41] Eric van Douwen. The integers and topology. In K. Kunen and J. Vaughan, editors, *Handbook of Set Theoretic Topology*, pages 111–167. North-Holland, 1984.

[42] Alan Dow. Tree π-bases for $\beta\mathbb{N}-\mathbb{N}$ in various models. *Topology Appl.*, 33:3–19, 1989.

[43] Alan Dow. More set-theory for topologists. *Topology Appl.*, 64:243–300, 1995.

[44] Erick Todd Eisworth. *Contributions to the Theory of Proper Forcing.* PhD thesis, University of Michigan, 1994.

[45] Grigorii Fichtenholz and Leonid Kantorovitch. Sur les opérations linéaires dans l'espace des fonctions bornées. *Studia Math.*, 5:69–98, 1935.

[46] David Fremlin. Cichoń's diagram. In G. Choquet, M. Rogalski, and J. Saint Raymond, editors, *Séminaire Initiation à l'Analyse*, pages 5–01 – 5–13. Publications Mathématiques de l'Université Pierre et Marie Curie, Paris, 1984.

[47] David Fremlin. *Consequences of Martin's Axiom*, volume 84 of *Cambridge Tracts in Math.* Cambridge Univ. Press, 1984.

[48] Fred Galvin and Karel Prikry. Borel sets and Ramsey's theorem. *J. Symbolic Logic*, 38:193–198, 1973.

[49] Leonard Gillman and Meyer Jerison. *Rings of Continuous Functions.* Van Nostrand, 1960.

[50] Martin Goldstern. Tools for your forcing construction. In H. Judah, editor, *Set Theory of the Reals*, volume 6 of *Israel Math. Conf. Proc.*, pages 305–360, 1993.

[51] Martin Goldstern and Saharon Shelah. Ramsey ultrafilters and the reaping number—Con($\mathfrak{r} < \mathfrak{u}$). *Ann. Pure Appl. Logic*, 49:121–142, 1990.

[52] Martin Goldstern and Saharon Shelah. Many simple cardinal invariants. *Arch. Math. Logic*, 32:203–221, 1993.

[53] Felix Hausdorff. Über zwei Sätze von G. Fichtenholz und L. Kantorovitch. *Studia Math.*, 6:18–19, 1936.

[54] Stephen Hechler. Short complete nested sequences in $\beta N \setminus N$ and small maximal almost-disjoint families. *Gen. Topology Appl.*, 2:139–149, 1972.

[55] Stephen Hechler. A dozen small uncountable cardinals. In R. Alo, R. Heath, and J.-I. Nagata, editors, *TOPO 72. General Topology and its Applications*, volume 378 of *Lecture Notes in Math.*, pages 207–218. Springer-Verlag, 1974.

[56] Stephen Hechler. On the existence of certain cofinal subsets of ${}^{\omega}\omega$. In T. Jech, editor, *Axiomatic Set Theory, Part II*, volume 13(2) of *Proc. Symp. Pure Math.*, pages 155–173. Amer. Math. Soc., 1974.

[57] Stephen Hechler. On a ubiquitous cardinal. *Proc. Amer. Math. Soc.*, 52:348–352, 1975.

[58] Neil Hindman. Finite sums from sequences within cells of a partition of N. *J. Combin. Theory Ser. A*, 17:1–11, 1974.

[59] Jaime Ihoda (= Haim Judah) and Saharon Shelah. Souslin forcing. *J. Symbolic Logic*, 53:1188–1207, 1988.

[60] Jaime Ihoda (= Haim Judah) and Saharon Shelah. Δ^1_2 sets of reals. *Ann. Pure Appl. Logic*, 42:207–223, 1989.

[61] Thomas Jech. *Multiple Forcing*, volume 88 of *Cambridge Tracts in Math.* Cambridge Univ. Press, 1986.

[62] Masaru Kada. The Baire category theorem and the evasion number. *Proc. Amer. Math. Soc.*, 126:3381–3383, 1988.

[63] Anastasis Kamburelis and Bogdan Węglorz. Splittings. *Arch. Math. Logic*, 35:263–277, 1996.

[64] S. Kamo. Cardnal invariants associated with predictors. In P. Pudl'ak, S. Buss, and P. Hájek, editors, *Logic Colloquium '98*, Lecture Notes in Logic, pages 280–295. Associtaion for Symbolic Logic, 2000.

[65] Jussi Ketonen. On the existence of P-points in the Stone-Čech compactification of integers. *Fund. Math.*, 92:91–94, 1976.

[66] Kenneth Kunen. Some points in βN. *Math. Proc. Cambridge Philos. Soc.*, 80:385–398, 1976.

[67] Kenneth Kunen. *Set Theory: An Introduction to Independence Proofs.* North-Holland, 1980.

[68] Kenneth Kunen and Franklin Tall. Between Martin's axiom and Souslin's hypothesis. *Fund. Math.*, 102:173–181, 1979.

[69] Claude Laflamme. Equivalence of families of functions on natural numbers. *Trans. Amer. Math. Soc.*, 330:307–319, 1992.

[70] Claude Laflamme. Combinatorial aspects of F_σ filters with an application to $\mathcal{N}$-sets. *Proc. Amer. Math. Soc.*, 125:3019–3025, 1997.

[71] Richard Laver. On the consistency of Borel's conjecture. *Acta Math.*, 137:151–169, 1976.

[72] Viatcheslav I. Malyhin. Topological properties of Cohen generic extension. *Trans. Moscow Math. Soc.*, 52:1–32, 1979.

[73] Donald A. Martin and Robert Solovay. Internal Cohen extensions. *Ann. Math. Logic*, 2:143–178, 1970.

[74] Adrian Mathias. Solution of problems of Choquet and Puritz. In W. Hodges, editor, *Conference in Mathematical Logic—London '70*, volume 255 of *Lecture Notes in Math.*, pages 204–210. Springer-Verlag, 1972.

[75] Adrian Mathias. Happy families. *Ann, Math. Logic*, 12:59–111, 1977.

[76] Heike Mildenberger. Non-constructive Galois-Tukey connections. *J. Symbolic Logic*, 62:1179–1186, 1997.

[77] Heike Mildenberger. Changing cardinal invariants of the reals without changing cardinals or the reals. *J. Symbolic Logic*, to appear.

[78] Arnold Miller. There are no Q-points in Laver's model for the Borel conjecture. *Proc. Amer. Math. Soc.*, 78:103–106, 1980.

[79] Arnold Miller. Some properties of measure and category. *Trans. Amer. Math. Soc.*, 266:93–114, 1981.

[80] Arnold Miller. Additivity of measure implies dominating reals. *Proc. Amer. Math. Soc.*, 91:111–117, 1984.

[81] Arnold Miller. Rational perfect set forcing. In J. Baumgartner, D. A. Martin, and S. Shelah, editors, *Axiomatic Set Theory*, volume 31 of *Contemp. Math.*, pages 143–159. American Mathematical Society, 1984.

[82] John Oxtoby. *Measure and Category.* Springer-Verlag, second edition, 1980.

[83] Janusz Pawlikowski. Why Solovay real produces Cohen real. *J. Symbolic Logic*, 51:957–968, 1986.

[84] Janusz Pawlikowski and Ireneusz Recław. Parametrized Cichoń's diagram and small sets. *Fund. Math.*, 147:135 – 155, 1995.

[85] Zbigniew Piotrowski and Andrzej Szymański. Some remarks on category in topological spaces. *Proc. Amer. Math. Soc.*, 101:156–160, 1987.

[86] Szymon Plewik. Intersections and unions of ultrafilters without the baire property. *Bull. Polish Acad. Sci. Math.*, 35:805–808, 1987.

[87] Bedřich Pospíšil. On bicompact spaces. *Publ. Fac. Sci. Univ. Masaryk*, 270:1939, 3–16.

[88] Jean Raisonnier and Jacques Stern. The strength of measurability hypotheses. *Israel J. Math.*, 50:337–349, 1985.

[89] Fritz Rothberger. Eine Äquivalenz zwischen der Kontinuumhypothese und der Existenz der Lusinschen und Sierpińskischen Mengen. *Fund. Math.*, 30:215–217, 1938.

[90] Fritz Rothberger. Sur un ensemble toujours de première catégorie qui est dépourvu de la propriété λ. *Fund. Math.*, 32:294–300, 1939.

[91] Fritz Rothberger. On some problems of Hausdorff and Sierpiński. *Fund. Math.*, 35:29–46, 1948.

[92] Fritz Rothberger. On the property c and a problem of Hausdorff. *Canad. J. Math.*, 4:111–116, 1952.

[93] Mary Ellen Rudin. Partial orders on the types in βN. *Trans. Amer. Math. Soc.*, 155:353–362, 1971.

[94] Walter Rudin. Homogeneity problems in the theory of Čech compactifications. *Duke Math. J.*, 23:409–419, 1956.

[95] Gerald Sacks. Forcing with perfect closed sets. In D. Scott, editor, *Axiomatic Set Theory*, volume 13, Part 1 of *Proc. Symp. Pure Math*, pages 331–355. American Mathematical Society, 1971.

[96] Saharon Shelah. *Proper Forcing*, volume 940 of *Lecture Notes in Math.* Springer-Verlag, 1982.

[97] Saharon Shelah. Can you take Solovay's inaccessible away? *Israel J. Math.*, 48:1–47, 1984.

[98] Saharon Shelah. On cardinal invariants of the continuum. In J. Baumgartner, D. Martin, and S. Shelah, editors, *Axiomatic Set Theory*, volume 31 of *Contemp. Math.*, pages 183–207, 1984.

[99] Saharon Shelah. Are $\mathfrak{a}$ and $\mathfrak{d}$ your cup of tea? to appear.

[100] Saharon Shelah and Otmar Spinas. The distributivity numbers of $\mathcal{P}(\omega)/\mathrm{fin}$ and its square. to appear.

[101] Jack Silver. Every analytic set is Ramsey. *J. Symbolic Logic*, 35:60–64, 1970.

[102] R. C. Solomon. Families of sets and functions. *Czechoslovak Math. J.*, 27:556–559, 1977.

[103] Robert Solovay. A model of set theory in which every set of reals is Lebesgue measurable. *Ann. Math.*, 92:1–56, 1970.

[104] Robert Solovay. Real-valued measurable cardinals. In D. Scott, editor, *Axiomatic Set Theory*, volume 13, Part 1 of *Proc. Symp. Pure Math*, pages 331–355. American Mathematical Society, 1971.

[105] Robert Solovay and Stanley Tennenbaum. Iterated Cohen extensions and Souslin's problem. *Ann. Math.*, 94:201–245, 1971.

[106] Otmar Spinas. Analytic countably splitting families. *J. Symbolic Logic*, to appear.

[107] Michel Talagrand. Compacts de fonctions mesurables et filtres non mesurables. *Studia Math.*, 67:13–43, 1980.

[108] Daniel Talayco. *Applications of Homological Algebra to Questions in Set Theory: Gaps and Trees.* PhD thesis, University of Michigan, 1993.

[109] John Truss. Sets having calibre $\aleph_1$. In R. O. Gandy and J. M. E. Hyland, editors, *Logic Colloquium 76*, volume 87 of *Stud. Logic Found. Math.*, pages 595–612. North-Holland, 1977.

[110] Jerry Vaughan. Small uncountable cardinals and topology. In J. van Mill and G. Reed, editors, *Open Problems in Topology*, pages 195–218. North-Holland, 1990.

[111] Peter Vojtáš. Generalized Galois-Tukey connections between explicit relations on classical objects of real analysis. In H. Judah, editor, *Set Theory of the Reals*, volume 6 of *Israel Math. Conf. Proc.*, pages 619–643. Amer. Math. Soc., 1993.

[112] Ed Wimmers. The Shelah P-point independence theorem. *Israel J. Math.*, 43:28–48, 1982.

[113] Olga Yiparaki. *On Some Tree Partitions.* PhD thesis, University of Michigan, 1994.

[114] Jindřich Zapletal. Isolating cardinal invariants. *J. Math. Log.*, 3:143–162, 2003.