Generalising the weak compactness of ω

Andrew Brooke-Taylor

UNIVERSITY OF LEEDS

Generalised Baire Spaces Masterclass
Royal Netherlands Academy of Arts and Sciences
22 August 2018

Overview

One obstacle to generalising results about cardinal characteristics of the continuum is if the arguments use compactness properties of ω.

Overview

One obstacle to generalising results about cardinal characteristics of the continuum is if the arguments use compactness properties of ω. In these cases, assuming your cardinal is weakly compact will often allow the argument to generalise.

Overview

One obstacle to generalising results about cardinal characteristics of the continuum is if the arguments use compactness properties of ω. In these cases, assuming your cardinal is weakly compact will often allow the argument to generalise.

In this tutorial I want to give a couple of examples of this, digging into the necessary preliminaries on the way.

Weak compactness

There are many equivalent formulations of weak compactness; we will use a couple of different ones.

Weak compactness

There are many equivalent formulations of weak compactness; we will use a couple of different ones.

Notice in each case that, if we didn't simply decree that weakly compact cardinals (and inaccessible cardinals) must be uncountable, then ω would fit the definition.

Weak compactness

There are many equivalent formulations of weak compactness; we will use a couple of different ones.

Notice in each case that, if we didn't simply decree that weakly compact cardinals (and inaccessible cardinals) must be uncountable, then ω would fit the definition.

Recommended Reference:

The Exercises for Section 4.2 of Chang & Keisler's *Model Theory* (but note that their definition of weakly compact needs to have inaccessibility added to it).

Infinitary languages

For any vocabulary Σ (i.e. set of function and relation symbols) and for any cardinal κ the language $\mathcal{L}_{\kappa,\kappa}$ consists of formulas built via the usual construction rules along with:

- Conjunctions and disjunctions of less than κ many formulas: if $\delta < \kappa$ and φ_γ is a formula for every $\gamma < \delta$, then

$$\bigvee_{\gamma<\delta} \varphi_\gamma \qquad \text{and} \qquad \bigwedge_{\gamma<\delta} \varphi_\gamma$$

  are formulas.

- Less than κ-fold quantifications: if $\mathbf{x} = (x_\gamma : \gamma < \delta)$ is an δ-tuple of variables for some $\delta < \kappa$ and φ is a formula, then

$$\exists \mathbf{x} \varphi \qquad \text{and} \qquad \forall \mathbf{x} \varphi$$

  are formulas.

Satisfaction of these formulas is defined as you would expect.

Infinitary languages

E.g.

Being well-ordered can be expressed by a sentence of $\mathcal{L}_{\omega_1,\omega_1}$:

$$\neg\exists(x_i : i \in \omega)\left(\bigwedge_{i\in\omega} x_i > x_{i+1}\right)$$

Infinitary languages

E.g.

Being well-ordered can be expressed by a sentence of $\mathcal{L}_{\omega_1,\omega_1}$:

$$\neg\exists(x_i : i \in \omega)\left(\bigwedge_{i\in\omega} x_i > x_{i+1}\right)$$

Weak compactness, 1st formulation

An uncountable cardinal κ is *weakly compact* if and only if, for every set of T of $\mathcal{L}_{\kappa,\kappa}$ sentences over a vocabulary Σ of cardinality at most κ, if every subset of T of cardinality less than κ has a model then T itself has a model.

Infinitary languages

E.g.

Being well-ordered can be expressed by a sentence of $\mathcal{L}_{\omega_1,\omega_1}$:

$$\neg\exists(x_i : i \in \omega) \left(\bigwedge_{i\in\omega} x_i > x_{i+1} \right)$$

Weak compactness, 1st formulation

An uncountable cardinal κ is *weakly compact* if and only if, for every set of T of $\mathcal{L}_{\kappa,\kappa}$ sentences over a vocabulary Σ of cardinality at most κ, if every subset of T of cardinality less than κ has a model then T itself has a model.

Note that if the constraint on $|\Sigma|$ is dropped, then this defines *strongly compact cardinals*.

Elementary embeddings

Weak compactness, 2nd formulation

An uncountable cardinal κ is weakly compact if and only if, for any structure M of size κ for a vocabulary Σ of cardinality κ, there is a Σ-structure N such that M is a proper elementary substructure of N in the $\mathcal{L}_{\kappa,\kappa}$ sense.

Elementary embeddings

Weak compactness, 2nd formulation

An uncountable cardinal κ is weakly compact if and only if, for any structure M of size κ for a vocabulary Σ of cardinality κ, there is a Σ-structure N such that M is a proper elementary substructure of N in the $\mathcal{L}_{\kappa,\kappa}$ sense.

Proof sketch that formulation 1 $\Rightarrow$ formulation 2

Add to the vocabulary a constant c_m for each element m of M, and consider the the complete $\mathcal{L}_{\kappa,\kappa}$ theory T of M for this language, with c_m interpreted as m.

Elementary embeddings

Weak compactness, 2nd formulation

An uncountable cardinal κ is weakly compact if and only if, for any structure M of size κ for a vocabulary Σ of cardinality κ, there is a Σ-structure N such that M is a proper elementary substructure of N in the $\mathcal{L}_{\kappa,\kappa}$ sense.

Proof sketch that formulation 1 $\Rightarrow$ formulation 2

Add to the vocabulary a constant c_m for each element m of M, and consider the the complete $\mathcal{L}_{\kappa,\kappa}$ theory T of M for this language, with c_m interpreted as m. Now add another constant c to the vocabulary and add to the theory all of the sentences $c \neq c_m$. Every subset A of this extended theory T' with cardinality $< \kappa$ has a model (M with a suitable choice of c), so by formulation 1, T' has a model; this will be N.

Elementary embeddings

Weak compactness, 2nd formulation

An uncountable cardinal κ is weakly compact if and only if, for any structure M of size κ for a vocabulary Σ of cardinality κ, there is a Σ-structure N such that M is a proper elementary substructure of N in the $\mathcal{L}_{\kappa,\kappa}$ sense.

Proof sketch that formulation 1 $\Rightarrow$ formulation 2

Add to the vocabulary a constant c_m for each element m of M, and consider the the complete $\mathcal{L}_{\kappa,\kappa}$ theory T of M for this language, with c_m interpreted as m. Now add another constant c to the vocabulary and add to the theory all of the sentences $c \neq c_m$. Every subset A of this extended theory T' with cardinality $< \kappa$ has a model (M with a suitable choice of c), so by formulation 1, T' has a model; this will be N.

Remember that well-foundedness is definable in $\mathcal{L}_{\kappa,\kappa}$ for $\kappa > \omega$, so these embeddings can be nice from a set theory point of view.

Trees

A κ*-tree* is a tree T of height κ such for every $\alpha < \kappa$ there are fewer than κ many nodes of T of height α.

Trees

A κ*-tree* is a tree T of height κ such for every $\alpha < \kappa$ there are fewer than κ many nodes of T of height α. A cardinal κ has the *tree* property if and only if every κ-tree has a cofinal branch (i.e. a branch of height κ).

Trees

A κ-*tree* is a tree T of height κ such for every $\alpha < \kappa$ there are fewer than κ many nodes of T of height α. A cardinal κ has the *tree* property if and only if every κ-tree has a cofinal branch (i.e. a branch of height κ).

Weak compactness, 3rd formulation

A cardinal is weakly compact if and only if it is inaccessible and has the tree property.

Trees

A κ*-tree* is a tree T of height κ such for every $\alpha < \kappa$ there are fewer than κ many nodes of T of height α. A cardinal κ has the *tree* property if and only if every κ-tree has a cofinal branch (i.e. a branch of height κ).

Weak compactness, 3rd formulation

A cardinal is weakly compact if and only if it is inaccessible and has the tree property.

Proof sketch that formulation 2 $\Rightarrow$ formulation 3

Code your κ-tree T as a subset of V_κ, and take

$$M = \langle V_\kappa, \in, T, \text{ enough extra stuff to make } V_\kappa \text{ "rigid"} \rangle.$$

κ is in the model $N = \langle X, E, T', \ldots \rangle$ given by formulation 2, and T' below level κ is just T. Choosing any node t of T' at level κ, the set of nodes below t is then a cofinal branch through T.

Warm-up example: $\mathfrak{s}_\kappa$

For $A, B \subseteq \kappa$ of cardinality κ, say that A *splits* B if $|B \cap A| = |B \setminus A| = \kappa$.

Warm-up example: $\mathfrak{s}_\kappa$

For $A, B \subseteq \kappa$ of cardinality κ, say that A *splits* B if $|B \cap A| = |B \setminus A| = \kappa$.

A family $\mathcal{A} \subseteq [\kappa]^\kappa$ is a *splitting family* if for every $B \subseteq \kappa$ with $|B| = \kappa$, there is an $A \in \mathcal{A}$ which splits B.

Warm-up example: $\mathfrak{s}_\kappa$

For $A, B \subseteq \kappa$ of cardinality κ, say that A *splits* B if $|B \cap A| = |B \setminus A| = \kappa$.

A family $\mathcal{A} \subseteq [\kappa]^\kappa$ is a *splitting family* if for every $B \subseteq \kappa$ with $|B| = \kappa$, there is an $A \in \mathcal{A}$ which splits B.

The *splitting number* $\mathfrak{s}_\kappa$ is the least cardinality of a splitting family.

Proposition

$\mathfrak{s}_\omega \geq \omega_1$

Proposition

$\mathfrak{s}_\omega \geq \omega_1$

Proof.

Proposition

$\mathfrak{s}_\omega \geq \omega_1$

Proof.

Diagonalise!

Proposition

$\mathfrak{s}_\omega \geq \omega_1$

Proof.

Diagonalise!

Consider any $\mathcal{A} = \{A_i : i \in \omega\} \subseteq [\omega]^\omega$. We will inductively define a sequence of infinite subsets B_i of ω and a sequence of elements c_i of ω such that no A in $\mathcal{A}$ splits $C = \{c_i : i \in \omega\}$.

Proposition

$\mathfrak{s}_\omega \geq \omega_1$

Proof.

Diagonalise!

Consider any $\mathcal{A} = \{A_i : i \in \omega\} \subseteq [\omega]^\omega$. We will inductively define a sequence of infinite subsets B_i of ω and a sequence of elements c_i of ω such that no A in $\mathcal{A}$ splits $C = \{c_i : i \in \omega\}$.

For the base case, let $B_0 = \omega$ and $c_0 = 0$. Having defined B_i, at least one of $B_i \cap A_i$ and $B_i \smallsetminus A_i$ is infinite, so pick one that is infinite, and take that to be B_{i+1}. Then let c_{i+1} be the least element of B_{i+1} that is greater than c_i.

Note that for each $i \in \omega$, $\{c_j : j > i\} \subseteq B_{i+1}$, and so is disjoint from or contained in A_i. So no A_i splits C.

□

Generalising to higher κ

Generalising to higher κ

Potential problems

To generalise this inductive argument to higher κ, we have to be able to deal with limit stages.

- At limit stages α along the way, the natural choice is to take $B_\alpha = \bigcap_{\gamma<\alpha} B_\gamma$. What if this intersection is empty, or even just of size $< \kappa$?

Generalising to higher κ

Potential problems

To generalise this inductive argument to higher κ, we have to be able to deal with limit stages.

- At limit stages α along the way, the natural choice is to take $B_\alpha = \bigcap_{\gamma<\alpha} B_\gamma$. What if this intersection is empty, or even just of size $< \kappa$?
  - The tree of possible choices we could have made for B_γ is a binary tree. If κ is inaccessible, then at any stage α, since there are only $2^\alpha < \kappa$ possible nodes, and they between them partition κ, at least one of them must have cardinality κ.

Generalising to higher κ

Potential problems

To generalise this inductive argument to higher κ, we have to be able to deal with limit stages.

- At limit stages α along the way, the natural choice is to take $B_\alpha = \bigcap_{\gamma<\alpha} B_\gamma$. What if this intersection is empty, or even just of size $< \kappa$?
  - The tree of possible choices we could have made for B_γ is a binary tree. If κ is inaccessible, then at any stage α, since there are only $2^\alpha < \kappa$ possible nodes, and they between them partition κ, at least one of them must have cardinality κ.
- Does this tree of possibilities have a cofinal branch, allowing us to define C?

Generalising to higher κ

Potential problems

To generalise this inductive argument to higher κ, we have to be able to deal with limit stages.

- At limit stages α along the way, the natural choice is to take $B_\alpha = \bigcap_{\gamma<\alpha} B_\gamma$. What if this intersection is empty, or even just of size $< \kappa$?
  - The tree of possible choices we could have made for B_γ is a binary tree. If κ is inaccessible, then at any stage α, since there are only $2^\alpha < \kappa$ possible nodes, and they between them partition κ, at least one of them must have cardinality κ.
- Does this tree of possibilities have a cofinal branch, allowing us to define C?
  - If κ also has the tree property, then yes!

So:

Proposition (Kamo; see Zapletal [4])

If κ is weakly compact, then $\mathfrak{s}_\kappa > \kappa$.

So:

Proposition (Kamo; see Zapletal [4])

If κ is weakly compact, then $\mathfrak{s}_\kappa > \kappa$. Actually, this is if and only if.

So:

Proposition (Kamo; see Zapletal [4])

If κ is weakly compact, then $\mathfrak{s}_\kappa > \kappa$. Actually, this is if and only if.

To get $\mathfrak{s}_\kappa > \kappa^+$ one requires even more large cardinal strength — see Zapletal [4].

More involved example: $\mathfrak{e}_\kappa$

Evasion and prediction were introduced by Blass in a paper motivated by group theory [1].

More involved example: $\mathfrak{e}_\kappa$

Evasion and prediction were introduced by Blass in a paper motivated by group theory [1].

Definition

A *predictor* is a a function π such that

- $\mathrm{dom}(\pi) \subseteq \kappa$ and $|\mathrm{dom}(\pi)| = \kappa$, and
- for each $\alpha \in \mathrm{dom}(\pi)$, $\pi(\alpha)$ is a function from κ^α to κ.

More involved example: $\mathfrak{e}_\kappa$

Evasion and prediction were introduced by Blass in a paper motivated by group theory [1].

Definition

A *predictor* is a a function π such that

- $\mathrm{dom}(\pi) \subseteq \kappa$ and $|\mathrm{dom}(\pi)| = \kappa$, and
- for each $\alpha \in \mathrm{dom}(\pi)$, $\pi(\alpha)$ is a function from κ^α to κ.

Definition

Given a predictor π and a function $f : \kappa \to \kappa$, we say π *predicts* f if there is some $\alpha < \kappa$ such that for all $\beta > \alpha$, $\pi(\beta)(f \restriction \beta) = f(\beta)$.

Definition

The *evasion number* $\mathfrak{e}_\kappa$ is the "bounding number for prediction":

$$\mathfrak{e}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists \text{ predictor } \pi \forall f \in \mathcal{F}(\pi \text{ predicts } f)\}.$$

Definition

The *evasion number* $\mathfrak{e}_\kappa$ is the "bounding number for prediction":

$$\mathfrak{e}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists \text{ predictor } \pi \forall f \in \mathcal{F}(\pi \text{ predicts } f)\}.$$

In the framework of relations introduced in the first talk today, $\mathfrak{e}_\kappa$ is the norm of the dual relation to $(\kappa^\kappa, \text{predictors}, \text{is predicted by})$.

$\mathfrak{e}$ vs $\mathfrak{b}$

$$\mathfrak{e}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists \text{ predictor } \pi \forall f \in \mathcal{F}(\pi \text{ predicts } f)\}.$$

$$\mathfrak{b}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists g \in \kappa^\kappa \forall f \in \mathcal{F}(g \geq^* f)\}.$$

(Recall $g \geq^* f$ means there is some $\alpha < \kappa$ such that for all $\beta \geq \alpha$, $g(\beta) \geq f(\beta)$.)

$\mathfrak{e}$ vs $\mathfrak{b}$

$$\mathfrak{e}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists \text{ predictor } \pi \forall f \in \mathcal{F}(\pi \text{ predicts } f)\}.$$

$$\mathfrak{b}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists g \in \kappa^\kappa \forall f \in \mathcal{F}(g \geq^* f)\}.$$

(Recall $g \geq^* f$ means there is some $\alpha < \kappa$ such that for all $\beta \geq \alpha$, $g(\beta) \geq f(\beta)$.)

$\mathfrak{e}_\omega$ is independent of $\mathfrak{b}_\omega$: there is a model in which $\mathfrak{e}_\omega > \mathfrak{b}_\omega$, a model in which $\mathfrak{e}_\omega < \mathfrak{b}_\omega$, and a model in which $\mathfrak{e}_\omega = \mathfrak{b}_\omega$.

$\mathfrak{e}$ vs $\mathfrak{b}$

$$\mathfrak{e}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists \text{ predictor } \pi \forall f \in \mathcal{F}(\pi \text{ predicts } f)\}.$$

$$\mathfrak{b}_\kappa = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \kappa^\kappa \wedge \neg\exists g \in \kappa^\kappa \forall f \in \mathcal{F}(g \geq^* f)\}.$$

(Recall $g \geq^* f$ means there is some $\alpha < \kappa$ such that for all $\beta \geq \alpha$, $g(\beta) \geq f(\beta)$.)

$\mathfrak{e}_\omega$ is independent of $\mathfrak{b}_\omega$: there is a model in which $\mathfrak{e}_\omega > \mathfrak{b}_\omega$, a model in which $\mathfrak{e}_\omega < \mathfrak{b}_\omega$, and a model in which $\mathfrak{e}_\omega = \mathfrak{b}_\omega$.

I will focus on Brendle and Shelah's construction of a model with $\mathfrak{e}_\omega > \mathfrak{b}_\omega$ [2].

Getting $\mathfrak{e} > \mathfrak{b}$

How to get a model of $\mathfrak{e}_\omega > \mathfrak{b}_\omega$? The natural approach:

- start with a model of CH, where $\mathfrak{e}_\omega = \mathfrak{b}_\omega = \omega_1$,

Getting $\mathfrak{e} > \mathfrak{b}$

How to get a model of $\mathfrak{e}_\omega > \mathfrak{b}_\omega$? The natural approach:

- start with a model of CH, where $\mathfrak{e}_\omega = \mathfrak{b}_\omega = \omega_1$,
- force to make $\mathfrak{e}_\omega$ large, whilst keeping $\mathfrak{b}_\omega$ small.

Getting $\mathfrak{e} > \mathfrak{b}$

How to get a model of $\mathfrak{e}_\omega > \mathfrak{b}_\omega$? The natural approach:

- start with a model of CH, where $\mathfrak{e}_\omega = \mathfrak{b}_\omega = \omega_1$,
- force to make $\mathfrak{e}_\omega$ large, whilst keeping $\mathfrak{b}_\omega$ small.

A natural approach to this latter:

- do the most obvious forcing to make $\mathfrak{e}_\omega$ large, and
- hope for the best.

Getting $\mathfrak{e} > \mathfrak{b}$

How to get a model of $\mathfrak{e}_\omega > \mathfrak{b}_\omega$? The natural approach:

- start with a model of CH, where $\mathfrak{e}_\omega = \mathfrak{b}_\omega = \omega_1$,
- force to make $\mathfrak{e}_\omega$ large, whilst keeping $\mathfrak{b}_\omega$ small.

A natural approach to this latter:

- do the most obvious forcing to make $\mathfrak{e}_\omega$ large, and
- use a clever argument to show that $\mathfrak{b}_\omega$ remains small.

Making $\mathfrak{c}$ large

There is a standard way to make a "bounding number"-type cardinal characteristic large:

- force to kill all ground model unbounded sets: add a new bound for them all,

Making $\mathfrak{c}$ large

There is a standard way to make a "bounding number"-type cardinal characteristic large:

- force to kill all ground model unbounded sets: add a new bound for them all,
- do a long (however large you want the cardinal) iteration of these forcings with finite support.

Making $\mathfrak{c}$ large

There is a standard way to make a "bounding number"-type cardinal characteristic large:

- force to kill all ground model unbounded sets: add a new bound for them all,
- do a long (however large you want the cardinal) iteration of these forcings with finite support. Any small potentially unbounded set appears after an initial fragment of the iteration, and then is forced to be bounded at the next step.

Making $\mathfrak{e}$ large

There is a standard way to make a "bounding number"-type cardinal characteristic large:

- force to kill all ground model unbounded sets: add a new bound for them all,
- do a long (however large you want the cardinal) iteration of these forcings with finite support. Any small potentially unbounded set appears after an initial fragment of the iteration, and then is forced to be bounded at the next step.

In particular, to make $\mathfrak{e}_\omega$ equal to some regular cardinal $\lambda > \omega_1$, we:

- force to add a new predictor that predicts all ground model functions,

Making $\mathfrak{e}$ large

There is a standard way to make a "bounding number"-type cardinal characteristic large:

- force to kill all ground model unbounded sets: add a new bound for them all,
- do a long (however large you want the cardinal) iteration of these forcings with finite support. Any small potentially unbounded set appears after an initial fragment of the iteration, and then is forced to be bounded at the next step.

In particular, to make $\mathfrak{e}_\omega$ equal to some regular cardinal $\lambda > \omega_1$, we:

- force to add a new predictor that predicts all ground model functions,
- do a length λ iteration with finite supports of this forcing. Any set $\mathcal{F} \subseteq \omega^\omega$ in the extension of cardinality less than λ must have appeared by some initial stage of the forcing, and then the predictor added at the next step predicts it.

Adding a predictor

How do we force to add a predictor that predicts all ground model functions?

Adding a predictor

How do we force to add a predictor that predicts all ground model functions?
By finite approximations!

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F \rangle$ such that

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F \rangle$ such that

1. $d \in 2^{<\omega}$

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F \rangle$ such that

1. $d \in 2^{<\omega}$ (a finite approximation to the characteristic function of the domain of the predictor),

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F \rangle$ such that

1. $d \in 2^{<\omega}$ (a finite approximation to the characteristic function of the domain of the predictor),
2. π is a function with domain $d^{-1}\{1\}$ such that $\forall n \in d^{-1}\{1\}$, $\pi(n)$ is a finite partial function from ω^n to ω

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F\rangle$ such that

1. $d \in 2^{<\omega}$ (a finite approximation to the characteristic function of the domain of the predictor),
2. π is a function with domain $d^{-1}\{1\}$ such that $\forall n \in d^{-1}\{1\}$, $\pi(n)$ is a finite partial function from ω^n to ω (a finite approximation to the predictor),

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F\rangle$ such that

1. $d \in 2^{<\omega}$ (a finite approximation to the characteristic function of the domain of the predictor),
2. π is a function with domain $d^{-1}\{1\}$ such that $\forall n \in d^{-1}\{1\}$, $\pi(n)$ is a finite partial function from ω^n to ω (a finite approximation to the predictor),
3. $F \subset \omega^\omega$ is finite and for $f \neq g \in F$, $\max(\{n : f \restriction n = g \restriction n\}) < \text{dom}(d)$

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F \rangle$ such that

1. $d \in 2^{<\omega}$ (a finite approximation to the characteristic function of the domain of the predictor),
2. π is a function with domain $d^{-1}\{1\}$ such that $\forall n \in d^{-1}\{1\}$, $\pi(n)$ is a finite partial function from ω^n to ω (a finite approximation to the predictor),
3. $F \subset \omega^\omega$ is finite and for $f \neq g \in F$, $\max(\{n : f \restriction n = g \restriction n\}) < \mathrm{dom}(d)$ (a promise to predict the functions in F from now on).

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F \rangle$ such that

1. $d \in 2^{<\omega}$ (a finite approximation to the characteristic function of the domain of the predictor),
2. π is a function with domain $d^{-1}\{1\}$ such that $\forall n \in d^{-1}\{1\}$, $\pi(n)$ is a finite partial function from ω^n to ω (a finite approximation to the predictor),
3. $F \subset \omega^\omega$ is finite and for $f \neq g \in F$, $\max(\{n : f \restriction n = g \restriction n\}) < \mathrm{dom}(d)$ (a promise to predict the functions in F from now on).

We say $\langle d', \pi', F' \rangle \leq \langle d, \pi, F \rangle$ if and only if

- $d' \supseteq d$, $\pi' \supseteq \pi$, and $F' \supseteq F$, and

Adding a predictor

Definition (Brendle-Shelah [2])

We define the single step predictor forcing $\mathbb{P}$. Conditions are triples $\langle d, \pi, F\rangle$ such that

1. $d \in 2^{<\omega}$ (a finite approximation to the characteristic function of the domain of the predictor),
2. π is a function with domain $d^{-1}\{1\}$ such that $\forall n \in d^{-1}\{1\}$, $\pi(n)$ is a finite partial function from ω^n to ω (a finite approximation to the predictor),
3. $F \subset \omega^\omega$ is finite and for $f \neq g \in F$, $\max(\{n : f \restriction n = g \restriction n\}) < \mathrm{dom}(d)$ (a promise to predict the functions in F from now on).

We say $\langle d', \pi', F'\rangle \leq \langle d, \pi, F\rangle$ if and only if

- $d' \supseteq d$, $\pi' \supseteq \pi$, and $F' \supseteq F$, and
- for all $f \in F$ and $n \in (d')^{-1}\{1\} \setminus d^{-1}\{1\}$, $\pi'(n)(f \restriction n) = f(n)$ (and in particular is defined).

Centred-ness

Let P be a partial order.

- A subset X of P is $(1, < \omega)$-centred if any finitely many conditions in X have a common extension in P.
- A subset Y of P is $(\lambda, < \omega)$-centred if Y may be decomposed as $Y = \bigcup_{\gamma < \lambda} Y_\gamma$ where each Y_γ is $(1, < \omega)$-centred in P.
- P is said to be σ-centred if it is $(\omega, < \omega)$-centred (as a subset of itself).

Centred-ness

Let P be a partial order.

- A subset X of P is $(1, < \omega)$-centred if any finitely many conditions in X have a common extension in P.
- A subset Y of P is $(\lambda, < \omega)$-centred if Y may be decomposed as $Y = \bigcup_{\gamma < \lambda} Y_\gamma$ where each Y_γ is $(1, < \omega)$-centred in P.
- P is said to be σ-centred if it is $(\omega, < \omega)$-centred (as a subset of itself).

Clearly any σ-centred forcing is ccc.

Centred-ness

Let P be a partial order.

- A subset X of P is $(1, <\omega)$-centred if any finitely many conditions in X have a common extension in P.
- A subset Y of P is $(\lambda, <\omega)$-centred if Y may be decomposed as $Y = \bigcup_{\gamma<\lambda} Y_\gamma$ where each Y_γ is $(1, <\omega)$-centred in P.
- P is said to be σ-centred if it is $(\omega, <\omega)$-centred (as a subset of itself).

Clearly any σ-centred forcing is ccc.

Note that the predictor forcing $\mathbb{P}$ is σ-centred: any set of conditions with the same d and π components are compatible: take the union of the F components, and extend d to take the value 0 for long enough to satisfy requirement 3 on conditions.

So $\mathbb{P}$ is ccc, and hence preserves cardinals and cofinalities.

So $\mathbb{P}$ is ccc, and hence preserves cardinals and cofinalities.

Given a generic filter G for $\mathbb{P}$, the union of the π components of the conditions in G is a predictor, with domain the union of the d components, which predicts every ground model function from ω to ω.

So $\mathbb{P}$ is ccc, and hence preserves cardinals and cofinalities.

Given a generic filter G for $\mathbb{P}$, the union of the π components of the conditions in G is a predictor, with domain the union of the d components, which predicts every ground model function from ω to ω. Iterating $\mathbb{P}$ with finite support for length λ a regular cardinal makes $\mathfrak{e}_\omega = \lambda$ in the generic extension, as outlined above.

So $\mathbb{P}$ is ccc, and hence preserves cardinals and cofinalities.

Given a generic filter G for $\mathbb{P}$, the union of the π components of the conditions in G is a predictor, with domain the union of the d components, which predicts every ground model function from ω to ω. Iterating $\mathbb{P}$ with finite support for length λ a regular cardinal makes $\mathfrak{e}_\omega = \lambda$ in the generic extension, as outlined above.

To show that $\mathfrak{b}_\omega = \omega_1$ in the extension, we use the following lemma:

Lemma

If $\mathcal{F}$ is an unbounded (with respect to $\leq^*$) family of functions from ω to ω, then

$$\Vdash_{\mathbb{P}} \check{\mathcal{F}} \text{ is unbounded.}$$

To prove the Lemma, we use the following definitions.

- For a condition $p = \langle d, \pi, F \rangle \in \mathbb{P}$ define

$$I_p = \{f \restriction \mathrm{dom}(d) : f \in F\}.$$

- For $\dot{h}$ a $\mathbb{P}$-name for a function in ω^ω define $h_{d,\pi,I} \in (\omega+1)^\omega$ by

$$h_{d,\pi,I}(n) = \min\{m \leq \omega : \text{there is no } p \in \mathbb{P} \text{ of the form } p = \langle d, \pi, F \rangle \text{ with } I_p = I \text{ such that } p \Vdash \dot{h}(n) > m\}.$$

To prove the Lemma, we use the following definitions.

- For a condition $p = \langle d, \pi, F \rangle \in \mathbb{P}$ define

$$I_p = \{f \restriction \mathrm{dom}(d) : f \in F\}.$$

- For $\dot{h}$ a $\mathbb{P}$-name for a function in ω^ω define $h_{d,\pi,I} \in (\omega+1)^\omega$ by

$$h_{d,\pi,I}(n) = \min\{m \leq \omega : \text{there is no } p \in \mathbb{P} \text{ of the form } p = \langle d, \pi, F \rangle \text{ with } I_p = I \text{ such that } p \Vdash \dot{h}(n) > m\}.$$

Main Claim

Actually, $h_{d,\pi,I} \in \omega^\omega$.

Proof.

Compactness of ω. □

Part II

Recap

- π predicts f if for all large enough n in $\mathrm{dom}(\pi)$, $\pi(n)(f \upharpoonright n) = f(n)$.
- $\mathfrak{e}_\omega = \min\{|\mathcal{F}| : \mathcal{F} \subseteq \omega^\omega \wedge \forall$ predictor $\pi \exists f \in \mathcal{F}(\pi$ does not predict $f)\}$
- Want to force $\mathfrak{e}_\omega > \mathfrak{b}_\omega$ with a long iteration.
- Conditions in the individual forcing step: $\langle d, \pi, F\rangle$ — finite approximations to the generic predictor and its domain, and a promise to predict the functions in F.
- Lemma: $\mathcal{F}$ unbounded before the forcing $\Rightarrow \mathcal{F}$ unbounded after.
- Towards proof of the Lemma, let $I_p = \{f \upharpoonright \mathrm{dom}(d) : f \in F\}$, and given a $\mathbb{P}$-name $\dot{h}$ for a funtion $\omega \to \omega$, define

$$h_{d,\pi,I}(n) = \min\{m \leq \omega : \text{there is no } p \in \mathbb{P} \text{ of the form } p = \langle d, \pi, F\rangle \text{ with } I_p = I \text{ such that } p \Vdash \dot{h}(n) > m\}.$$

- Main Claim: $h_{d,\pi,I}(n)$ is finite for all n.

How the Main Claim proves the Lemma

Given $\dot{h}$, there are only countably many functions $h_{d,\pi,I}$ (as d, π and I vary), so there is some $h^*\colon \omega \to \omega$ such that $h_{d,\pi,I} \leq^* h^*$. Since $\mathcal{F}$ is unbounded there is some $f \in \mathcal{F}$ such that $f(n) > h^*(n)$ for infinitely many n.

How the Main Claim proves the Lemma

Given $\dot{h}$, there are only countably many functions $h_{d,\pi,I}$ (as d, π and I vary), so there is some $h^*\colon \omega \to \omega$ such that $h_{d,\pi,I} \leq^* h^*$. Since $\mathcal{F}$ is unbounded there is some $f \in \mathcal{F}$ such that $f(n) > h^*(n)$ for infinitely many n. Then also

$$\Vdash_{\mathbb{P}} \check{f}(n) > \dot{h}(n) \text{ for infinitely many } n. \tag{$\dagger$}$$

Otherwise, for some m and some p we would have $p \Vdash_{\mathbb{P}} \forall n \geq \check{m}(\check{f}(n) \leq \dot{h}(n))$. But taking d, π and I corresponding to p and $n \geq m$ such that $f(n) > h^*(n) \geq h_{d,\pi,I}(n)$, we'd have $p \Vdash_{\mathbb{P}} \check{h}_{d,\pi,I}(\check{n}) < \check{f}(\check{n}) \leq \dot{h}(\check{n})$, contradicting the definition of $h_{d,\pi,I}$. So ($\dagger$) holds, and so $\mathcal{F}$ is not bounded by the function named by $\dot{h}$, which was arbitrary; so $\mathcal{F}$ remains unbounded.

Proving the Main Claim

Main Claim: $h_{d,\pi,I}(n)$ is finite for all n, where

$$h_{d,\pi,I}(n) = \min\{m \leq \omega : \text{there is no } p \in \mathbb{P} \text{ of the form } p = \langle d, \pi, F\rangle \text{ with } I_p = I \text{ such that } p \Vdash \dot{h}(n) > m\}.$$

Suppose not. Then there are d, π, I and p_i for $i \in \omega$ with $p_i = \langle d, \pi, F_i\rangle$ and $I_{p_i} = I$ such that

$$p_i \Vdash_{\mathbb{P}} \dot{h}(\check{n}) > \check{i}.$$

Proving the Main Claim

Main Claim: $h_{d,\pi,I}(n)$ is finite for all n, where

$$h_{d,\pi,I}(n) = \min\{m \leq \omega : \text{there is no } p \in \mathbb{P} \text{ of the form } p = \langle d, \pi, F\rangle \text{ with } I_p = I \text{ such that } p \Vdash \dot{h}(n) > m\}.$$

Suppose not. Then there are d, π, I and p_i for $i \in \omega$ with $p_i = \langle d, \pi, F_i\rangle$ and $I_{p_i} = I$ such that

$$p_i \Vdash_{\mathbb{P}} \dot{h}(\check{n}) > \check{i}.$$

Let's index I as $I = \{\bar{f}_\ell : \ell < |I|\}$ and each F_i as $F_i = \{f^i_\ell : \ell < |I|\}$ such that $f^i_\ell \restriction \mathrm{dom}(d) = \bar{f}_\ell$.

We may thin out the sequence of p_i so that for each $\ell < |I|$, either

- $\exists g_\ell \in \omega^\omega \forall i (f^i_\ell \restriction i = g_\ell \restriction i)$, or
- $\exists i_\ell \in \omega \exists \hat{g}_\ell \in \omega^{i_\ell} \forall i (f^i_\ell \restriction i_\ell = \hat{g}_\ell \wedge f^i_\ell(i_\ell) > i)$.

Proving the Main Claim

Main Claim: $h_{d,\pi,I}(n)$ is finite for all n, where

$$h_{d,\pi,I}(n) = \min\{m \leq \omega : \text{there is no } p \in \mathbb{P} \text{ of the form } p = \langle d, \pi, F\rangle \text{ with } I_p = I \text{ such that } p \Vdash \dot{h}(n) > m\}.$$

Suppose not. Then there are d, π, I and p_i for $i \in \omega$ with $p_i = \langle d, \pi, F_i\rangle$ and $I_{p_i} = I$ such that

$$p_i \Vdash_{\mathbb{P}} \dot{h}(\check{n}) > \check{i}.$$

Let's index I as $I = \{\bar{f}_\ell : \ell < |I|\}$ and each F_i as $F_i = \{f^i_\ell : \ell < |I|\}$ such that $f^i_\ell \restriction \mathrm{dom}(d) = \bar{f}_\ell$.

We may thin out the sequence of p_i so that for each $\ell < |I|$, either

- $\exists g_\ell \in \omega^\omega \forall i (f^i_\ell \restriction i = g_\ell \restriction i)$, or
- $\exists i_\ell \in \omega \exists \hat{g}_\ell \in \omega^{i_\ell} \forall i (f^i_\ell \restriction i_\ell = \hat{g}_\ell \wedge f^i_\ell(i_\ell) > i)$.

Compactness of ω has been used here!

With the p_i thinned out to get this dichotomy, we have enough hands-on control to build a contradiction. □

Generalising to κ

(Joint work with Jörg Brendle)

To carry over the proof for the single step of the iteration, the main obstacle is to generalise this thinning out.

Generalising to κ

(Joint work with Jörg Brendle)

To carry over the proof for the single step of the iteration, the main obstacle is to generalise this thinning out.

There is a natural tree to consider — the set of all restrictions $f^i_\ell \restriction k$. But it would still be helpful to think of "climbing up through the tree" rather than just using a branch that the tree property hands down to use.

Generalising to κ

(Joint work with Jörg Brendle)

To carry over the proof for the single step of the iteration, the main obstacle is to generalise this thinning out.

There is a natural tree to consider — the set of all restrictions $f^i_\ell \restriction k$. But it would still be helpful to think of "climbing up through the tree" rather than just using a branch that the tree property hands down to use.

To do this we can use the embedding form of weak compactness to give us an ultrafilter to follow.

Proposition

Let κ be a weakly compact cardinal. Then for suitable structures M of size κ for a vocabulary of size at most κ, there is an M-normal ultrafilter: a set $\mathcal{U}$ such that

$$\langle M, \in, \mathcal{U} \models \mathcal{U} \text{ is a } \kappa\text{-complete normal ultrafilter on } \kappa.$$

In particular, if $M^{<\kappa} \subset M$, then $\mathcal{U}$ really is closed under $< \kappa$-fold intersections.

Proof sketch.

Use the embedding formulation of weak compactness, define for $X \subset \kappa$ in M,

$$X \in \mathcal{U} \Leftrightarrow j(X) \ni \kappa.$$

□

So working in a κ-sized model containing everything needed (such as the tree), we can use this ultrafilter to guide our way up the tree, and at the end, use normality to get the final thinned out sequence. We get:

Lemma

Suppose κ is a weakly compact cardinal, γ is a cardinal less than κ, and for each $\beta \in \gamma$, $\langle f_\beta^\delta : \delta \in \kappa \rangle$ is a sequence of functions in κ^κ. Then there is a strictly increasing sequence of ordinals less than κ, $\langle \delta_\eta : \eta \in \kappa \rangle$, such that for every $\beta \in \gamma$,

$$\text{either } (a)_\beta : \exists g_\beta \in \kappa^\kappa \forall \eta < \kappa (f_\beta^{\delta_\eta} \restriction \eta = g_\beta \restriction \eta)$$

$$\text{or } (b)_\beta : \exists \iota_\beta < \kappa \exists \hat{g}_\beta \in \kappa^{\iota_\beta} \forall \eta < \kappa (f_\beta^{\delta_\eta} \restriction \iota_\beta = \hat{g}_\beta \wedge f_\beta^{\delta_\eta}(\iota_\beta) \geq \eta).$$

Finishing the generalisation

The above Lemma is enough to generalise the single forcing step. How about the rest of the proof?

Finishing the generalisation

The above Lemma is enough to generalise the single forcing step. How about the rest of the proof?

More problems:

- We want κ to remain weakly compact for later stages of the forcing.

Finishing the generalisation

The above Lemma is enough to generalise the single forcing step. How about the rest of the proof?

More problems:

- We want κ to remain weakly compact for later stages of the forcing.

  Solution: Johnstone [3] showed that strongly unfoldable cardinals, which are somewhat stronger than weakly compact cardinals but still far below supercompact, can be made indestructible to a class of forcings including these. So we just assume this stronger large cardinals.

Finishing the generalisation

The above Lemma is enough to generalise the single forcing step. How about the rest of the proof?

More problems:

- We want κ to remain weakly compact for later stages of the forcing.

  Solution: Johnstone [3] showed that strongly unfoldable cardinals, which are somewhat stronger than weakly compact cardinals but still far below supercompact, can be made indestructible to a class of forcings including these. So we just assume this stronger large cardinals.
- The iteration theorems that deal with limit stages of the iteration in the ω case don't carry over for small cofinality limit stages.

Finishing the generalisation

The above Lemma is enough to generalise the single forcing step. How about the rest of the proof?

More problems:

- We want κ to remain weakly compact for later stages of the forcing.

  Solution: Johnstone [3] showed that strongly unfoldable cardinals, which are somewhat stronger than weakly compact cardinals but still far below supercompact, can be made indestructible to a class of forcings including these. So we just assume this stronger large cardinals.

- The iteration theorems that deal with limit stages of the iteration in the ω case don't carry over for small cofinality limit stages.

  Solution: Work with the fact that $\mathbb{P}$ is κ *centred with canonical lower bounds* (caution: not written down yet...).

Questions

- What large cardinal assumption is really needed for $\mathfrak{e}_\kappa > \mathfrak{b}_\kappa$? Strong unfoldability? Weak compactness? None?

Questions

- What large cardinal assumption is really needed for $\mathfrak{e}_\kappa > \mathfrak{b}_\kappa$? Strong unfoldability? Weak compactness? None?
- What other cardinal characteristics of the continuum results use compactness & need weak compactness to generalise?

Questions

- What large cardinal assumption is really needed for $\mathfrak{e}_\kappa > \mathfrak{b}_\kappa$? Strong unfoldability? Weak compactness? None?
- What other cardinal characteristics of the continuum results use compactness & need weak compactness to generalise?
- This is saying something about the *necessity* of compactness for these arguments from the ω case. Is there an interesting way to view this from a reverse mathematics perspective?

Andreas Blass.
Cardinal characteristics and the product of countably many infinite cyclic groups.
Journal of Algebra, 169:512–540, 1994.

Jörg Brendle and Saharon Shelah.
Evasion and prediction II.
Journal of the London Mathematical Society, 53(1):19–27, 1996.

Thomas Johnstone.
Strongly unfoldable cardinals made indestructible.
Journal of Symbolic Logic, 73(4):1215–1248, December 2008.

Jindřich Zapletal.
Splitting number at uncountable cardinals.
Journal of Symbolic Logic, 62(1):35–42, 1997.