GOST SEMINAR NOTES

AN INTRODUCTION TO MEASURES AND ULTRAPOWERS

James Holland
math.rutgers.edu/~jch258
jch258@math.rutgers.edu

Start Date: 2020-02-06
End Date: 2020-02-13
Updated: 2022-10-11

Section 1. *Ultrafilters and Logic*

Ultrafilters can be seen in a variety of places around mathematics, especially within set theory. For our purposes, ultrafilters give rise to ultraproducts, and certain ultraproducts result in inner models. The existence and structure of these inner models give rise to deep results about the original universe we start in, and they present important connections to large cardinal assumptions and consistency strength.

§ 1 A. Filters

The notion of a filter over a set makes precise the notion of largeness as well as "almost every". Its association with measure also leads to saying a set is "measure one" to mean that it is in the filter, alluding to measuring subsets of $[0,1] \subseteq \mathbb{R}$ similar to probability. This way of referring to sets in a filter $F \subseteq \wp(X)$ is motivated by the idea that if $x \in \wp(X)$ is "large" and $x \subseteq y \in \wp(X)$, then y is "large" too. This leads to the following definition.

1 A • 1. Definition

Let $A \neq \emptyset$ be a set. A *filter* over A is a non-empty subset $F \subsetneq \wp(A)$ such that the following hold: for all $x, y \in \wp(A)$,

1. If $x \in F$ and $x \subseteq y$, then $y \in F$; and
2. If $x, y \in F$, then $x \cap y \in F$.

An *ultrafilter* over A is a $\subseteq$-maximal filter $U \subseteq \wp(A)$.

Other references will often require $A \in F \subseteq \wp(A)$ and $\emptyset \notin F$, but these are implied by (1) and that $F \neq \emptyset$ is a *proper* subset $F \subsetneq \wp(A)$. Without this requirement, we'd have trivial filters like all of $\wp(A)$, or just $\emptyset$. We wouldn't want to allow such sets to be filters, because it would muck with the definition of ultrafilters.

To help grasp the concept a bit more, we have some relatively easy examples of filters.

1 A • 2. Example

1. Let $A \neq \emptyset$ be any set with $a \in A$. Therefore $\{x \in \wp(A) : a \in x\}$ is a filter, and in fact an ultrafilter.
2. Let $A \neq \emptyset$ be any set with $\emptyset \neq x \subsetneq A$. Therefore $\{y \in \wp(A) : x \subseteq y\}$ is a filter, but not an ultrafilter unless x is a singleton.
3. Suppose A is infinite. Therefore $\{x \in \wp(A) : A \setminus x \text{ is finite}\}$ is a filter, but not an ultrafilter.
4. Let κ be an uncountable, regular cardinal. Call $x \subseteq \kappa$ a *club* iff $\sup x = \kappa$, and for all bounded $y \subseteq x$, $\sup y \in x$. Therefore $\{x \in \wp(\kappa) : x \text{ contains a club of } \kappa\}$ is a filter called the *club filter*, but it is not an ultrafilter.

The first is in effect the most trivial kind of filter, and it is something we will try to avoid. Note that we can come up with all sorts of filters. First we just start with a family of non-pairwise-disjoint sets X, and then we close under intersections and supersets. This yields a filter containing X. So this is the process by which we construct[i] ultrafilters: just keep adding a set or its complement until we can't anymore.

For now, let's try to generate more examples of filters. To do this, we need to be somewhat careful. The general idea is to simply close a given set under finite intersections, and then add all supersets. The issue with this is that we need to ensure that we don't accidentally end up with $\emptyset$ after intersecting a bunch of elements. Otherwise $\emptyset$ would be in our filter, and after closing upwards under $\subseteq$, we'd end up with the full powerset. Luckily, this is the only obstruction to generating a filter.

1 A • 3. Definition

A set X has the *finite intersection property* iff for all finite subsets $\{x_0, \cdots, x_n\} \subseteq X$, $\bigcap_{i \leq n} x_i \neq \emptyset$.

1 A • 4. Result

Let $A \neq \emptyset$ be a set, and let $X \subseteq \wp(A)$ have the finite intersection property. Therefore there is a filter $F \supseteq X$.

Proof .:.

Consider the closure Y of X under pairwise intersections. By the finite intersection property, $\emptyset \notin Y$. Now define $F = \{x \in \wp(A) : \exists y \in Y (y \subseteq x)\}$. As $\emptyset \notin Y$, $\emptyset \notin F$ and hence $F \subsetneq \wp(A)$. F is clearly closed under supersets and pairwise intersection because Y is. Hence F is a filter with $X \subseteq Y \subseteq F$. ⊣

The filter given in the proof is generated by X not just in the sense that the construction is given by X, but also in the sense that it is the $\subseteq$-minimal filter containing X. Now the question becomes how to generate an *ultrafilter*. Without AC, the situation is a bit odd and different[ii], but in our case, *every* filter can be extended to an ultrafilter. The proof of this can be easily shown through Zorn's lemma: consider the set of filters containing F, and for each $\subseteq$-chain, just take the union to get another filter, and end up with a $\subseteq$-maximal filter $U \supseteq F$.

The characterization of ultrafilters just as maximal filters is useful to prove their existence[iii], but for the most part, it doesn't help one understand properties of ultrafilters. A much more useful characterization is the following.

1 A • 5. Result

Let $U \subseteq \wp(A)$ be a filter. Therefore U is an ultrafilter iff for all $x \in \wp(A)$, either $x \in U$ or $A \setminus x \in U$.

Proof .:.

Suppose U contains every subset of A or its complement, but there is some other filter with $U \subsetneq F \subsetneq \wp(A)$. Take $x \in F \setminus U$ and note that we must have $A \setminus x \in U \subseteq F$. Since F is a filter, $\emptyset = x \cap (A \setminus x) \in F$ which implies $F = \wp(A)$, a contradiction.

Now suppose U is an ultrafilter. Let $x \subseteq A$ be such that $x, A \setminus x \notin U$. Consider the set $X = \{u \setminus x : u \in U\}$ which contains $A \setminus x$, for example. Note $\emptyset \notin X$ since otherwise $u \setminus x = \emptyset$ for some $u \in U$, meaning $u \subseteq x \in U$. Therefore X has the finite intersection property because U does. So let F be the filter generated by X: $F = \{y \subseteq X : \exists z \in F\ (z \subseteq y)\}$. This contains U, contradicting that U is maximal: $x \in F \setminus U$. ⊣

We will only be interested in ultrafilters over infinite sets, since the only ultrafilters over finite sets are principal: for U an ultrafilter over $N \in \omega$, each $m \in N$ has some $X_m \in U$ with $m \notin X_m$. Therefore, intersecting these finitely many sets yields $\bigcap_{m<N} X_m = \emptyset \in U$, contradicting that U is a filter.

§ 1 B. Background on Clubs

Without loss of generality, we will consider filters on infinite cardinals. This somewhat simplifies the notation and situation, but it is a primer for later ideas which work directly with concepts related to cardinals. Example 1 A • 2 (4) already contains an example of how working cardinals can give additional information. This example also includes the notion of a *club*, being a closed and unbounded subset. These sets have further properties that can be connected to

[i] Choice is needed to show the existence of (nonprincipal) ultrafilters, so it's not exactly an explicit construction.
[ii] Things being odd is usually the case without AC. In particular, the existence of non-principal ultrafilters cannot be proven without AC.
[iii] and it's the only definition that works if you define filters for posets in general rather than just for the poset $\langle \wp(A), \subseteq \rangle$

ultrafilters later. First we repeat a definition.

1 B • 1. Definition

Let κ be an uncountable, regular cardinal.
A subset $x \subseteq \kappa$ is *club* in κ or *a club* iff $x \cup \{\kappa\}$ is closed under supremum of subsets and $\sup x = \kappa$.
Let $\{x_\alpha : \alpha < \lambda\}$ be a family of sets indexed by $\lambda \in \text{Ord}$. The *diagonal intersection* of this family is

$$\mathop{\triangle}_{\alpha<\kappa} x_\alpha := \left\{\alpha < \kappa : \alpha \in \bigcap_{\beta<\alpha} x_\alpha\right\}.$$

The diagonal intersection is important because the set of clubs is closed under diagonal intersections of length κ. This cannot be strengthened to full intersections, however. To see this, for each $\alpha < \kappa$, take the club $C_\alpha = \{\beta < \kappa : \alpha < \beta\}$. This gives that $\bigcap_{\alpha<\kappa} C_\alpha = \emptyset$. The diagonal intersection, however, will still be a club, and in fact will be κ itself: for every $\alpha < \kappa$, $\alpha \in \bigcap_{\beta<\alpha} C_\beta$.

1 B • 2. Result

Let κ be an uncountable, regular cardinal. Let $\{C_\alpha : \alpha < \kappa\}$ be a collection of clubs. Therefore

1. For each $\lambda < \kappa$, $\bigcap_{\alpha<\lambda} C_\alpha$ is a club.
2. $\triangle_{\alpha<\kappa} C_\alpha$ is a club.

Proof .:.

1. Let $\lambda < \kappa$ be given. First we will show that $\bigcap_{\alpha<\lambda} C_\alpha$ is unbounded. So let $\gamma < \kappa$ be arbitrary. Since κ is regular, choose an increasing sequence of x_αs such that each $x_\alpha \in C_\alpha$ and $\gamma < x_0$. Now we have a sequence $\langle x_\alpha : \alpha < \lambda\rangle = \langle x_{0+\alpha} : \alpha < \lambda\rangle$. Set $x_{\lambda+0} > \sup_{\alpha<\lambda} x_{0+\alpha}$, and again choose an increasing sequence as before: $x_{\lambda+\alpha} \in C_\alpha$. In the end, we get an interlaced, increasing sequence $X = \langle x_{\lambda\cdot n+\alpha} : n < \omega \wedge \alpha < \lambda\rangle$ where $x_{\lambda\cdot n+\alpha} \in C_\alpha$ for each $\alpha < \lambda$. Notice that as the sequence was interlaced and increasing, each C_α slice for $\alpha < \lambda$ has the same supremum:
$$\sup(X \cap C_\alpha) = \sup\{x_{\lambda\cdot n+\alpha} : n \in \omega\} = \sup X.$$
Hence this supremum is in $\bigcap_{\alpha<\lambda} C_\alpha$, and is bigger than γ. Thus the intersection is unbounded.

   To see that $\bigcap_{\alpha<\lambda} C_\alpha$ is closed, any bounded subset $Y \subseteq \bigcap_{\alpha<\lambda} C_\alpha$ has $Y \subseteq C_\alpha$ for each $\alpha < \lambda$. Yet $\sup Y \in C_\alpha$ by the hypothesis, for each $\alpha < \lambda$, implying that $\sup Y \in \bigcap_{\alpha<\lambda} C_\alpha$ as desired.

2. We will again show that $\triangle_{\alpha<\kappa} C_\alpha$ is unbounded, as closure is the easier of the two. Let $\gamma < \kappa$ be arbitrary. Choose an increasing sequence $\langle x_n > \gamma : n < \omega\rangle$ with $x_0 \in C_0 \setminus \gamma$ and $x_{n+1} \in \bigcap_{\alpha<x_n} C_\alpha$. This can be done since each $\bigcap_{\alpha<x_n} C_\alpha$ is club by (1). Now write $X = \{x_n : n \in \omega\}$ with $x = \sup X$.

   To see that $x \in \triangle_{\alpha<\kappa} C_\alpha$, we just need to see that $x \in \bigcap_{\alpha<x} C_\alpha$. For each $\alpha < x$, $\alpha \leq x_m$ for some m, which means the tail of X is contained in C_α:
$$\{x_n : n > m\} \subseteq \bigcap_{\beta<x_m} C_\beta \subseteq C_\alpha.$$
Therefore $x = \sup X \in C_\alpha$ and hence $x \in \bigcap_{\alpha<x} C_\alpha$.

   To see that $\triangle_{\alpha<\kappa} C_\alpha$ is closed, let $X \subseteq \gamma$ be a bounded subset of it with $x = \sup X$. Note that for any $\alpha < \kappa$, we have $X \setminus \alpha \subseteq \bigcap_{\beta<\alpha} C_\beta$. In particular, for $\alpha < x$, the tail of X is a subset of C_α and hence $x = \sup X \in C_\alpha$. Therefore $x \in \bigcap_{\alpha<x} C_\alpha$ and so $x \in \triangle_{\alpha<\kappa} C_\alpha$. ⊣

The importance of the diagonal intersection is primarily for the purpose of Fodor's lemma, which motivates an important property for filters. Fodor's lemma talks about stationary sets: sets which intersect every club set, but which might not be clubs themselves.

1 B • 3. Definition

Let κ be an uncountable, regular cardinal. A subset $X \subseteq \kappa$ is *stationary* iff $C \cap X \neq \emptyset$ for every club $C \subseteq \kappa$.

The existence of stationary sets is easy to see just from the fact that every club set is stationary: κ itself is trivially a stationary subset of κ. The existence of stationary, co-stationary subsets—i.e. stationary subsets that do not contain a club—can be shown through direct example. Since $\kappa > \aleph_0$, we can consider $S^\kappa_\omega = \{\alpha < \kappa : \text{cof}(\alpha) = \omega\}$. It's clear that S^κ_ω is stationary, since each club contains a sequence of length ω, whose supremum is then in S^κ_ω since $\kappa > \aleph_0$

is regular. More generally, the set S^κ_λ of ordinals with cofinality λ will be stationary whenever $\lambda < \kappa$ is regular for precisely the same reason as with ω.

1 B • 4. Definition

Let $X \subseteq \mathrm{Ord}$ and let $f : X \to \mathrm{Ord}$. f is *regressive* iff $f(\alpha) < \alpha$ for all $\alpha \in X$.

1 B • 5. Lemma (Fodor's Lemma)

Let κ be an uncountable, regular cardinal. Let $S \subseteq \kappa$ be stationary, and let $f : S \to \kappa$ be regressive. Therefore f is constant on a stationary set: $f^{-1}``\{\beta\}$ is stationary for some $\beta \in S$.

Proof .:.

Otherwise for each $\beta < \kappa$, let $C_\beta \cap f^{-1}``\{\beta\} = \emptyset$ with C_β a club. Consider $\triangle_{\beta<\kappa} C_\beta$, which is a club by Result 1 B • 2, and hence $S \cap \triangle_{\beta<\kappa} C_\beta \neq \emptyset$. Taking $\alpha \in S \cap \triangle_{\beta<\kappa} C_\beta$ requires that $f(\alpha) < \alpha$. But note that

$$f^{-1}``\{\beta\} \cap C_\beta \subseteq f^{-1}``\{\beta\} \cap \bigcap_{\gamma<\alpha} C_\gamma = \emptyset$$

for each $\beta < \alpha$. In particular, for $\beta = f(\alpha) < \alpha$, $\alpha \in f^{-1}``\{f(\alpha)\}$ has $\alpha \notin \bigcap_{\gamma<\alpha} C_\gamma$, contradicting that $\alpha \in \triangle_{\beta<\kappa} C_\beta$. Hence there must be some $\beta$ with $f^{-1}``\{\beta\}$ stationary, meaning that f is constant on a stationary set. ⊣

Such a result is extremely useful for combinatorial parts of set theory, being used to prove statements like the generalized Δ-system lemma, tremendously useful in methods of forcing. Stated in terms of filters, any ultrafilter extending the club filter will necessarily contain only stationary sets, and thus will abide by Fodor's Lemma (1 B • 5). This is a nice property of ultrafilters for various reasons, as will be covered later.

§ 1 C. Logic and filters

Now as stated above, filters and ultrafilters give a notion of "size" or "largeness" to subsets, but they also then give a notion of "how often" something is in a given subset. In this way, as with a probability measure, ultrafilters give a notion of how often something is true. To make this connection a little more apparent, the following notation will be used extensively.

1 C • 1. Definition

Let F be a filter over a set K. Let $\varphi(x, \vec{w})$ be a formula. Write "$\forall^* x\ \varphi(x, \vec{w})$" to say that $\{x \in K : \varphi(x, \vec{w})\} \in F$. We write $\exists^* x\ \varphi(x, \vec{w})$ to say that $K \setminus \{x \in K : \varphi(x, \vec{w})\} \notin F$, i.e. $\neg\forall^* x\ \neg\varphi(x, \vec{w})$.

$\forall^*$ should be read as "for almost every", and $\exists^*$ doesn't have a standard phrase, but one can read it as "there is a positive set", analogous to measure on the real numbers as if to say it's not measure 0. If we need to specify the ultrafilter, we write $\forall^*_U$ for "for U-almost every". In everyday language, words and phrases like "almost every", "by-and-large", and "many" come into play to gloss over details. These words are usually vague or ambiguous, but the notion of an ultrafilter makes them precise in a way that is consistent with ordinary usage. Moreover, the new quantifiers have their own sort of logic to them based just on Definition 1 A • 1. This new vocabulary dramatically simplifies some proofs, and is overall a better way of thinking about ultrafilters, as well as their properties. Definitions that may seem unmotivated or hard to understand can become more intuitive and natural with the new logical framework.

It's useful to present some easy results about how this quantifier interacts with the other connectives of first order logic. Note that the two properties of Definition 1 A • 1 can be restated as

1. If $\forall^*\alpha\ \varphi$ and $\forall\alpha\ (\varphi \to \psi)$, then $\forall^*\alpha\ \psi$; and
2. If $\forall^*\alpha\ \varphi$ and $\forall^*\alpha\ \psi$, then $\forall^*\alpha\ (\varphi \wedge \psi)$.

This isn't difficult to see if we allow parameters, since $\varphi(\alpha, x)$ might just be $\alpha \in x$ and $\psi(\alpha, y)$ might just be $\alpha \in y$. Regardless, these immediately give the following. As a result of the results to follow, (1) above can be weakened so that if $\forall^*\alpha\ \varphi$ and $\forall^*\alpha\ (\varphi \to \psi)$, then $\forall^*\alpha\ \psi$.

1 C • 2. Result

Let U be a filter over set K. Let φ and ψ be formulas, possibly with parameters. Therefore

1. $\forall^* x\ \varphi \to \exists^* x\ \varphi$. The two are equivalent for U an ultrafilter.
2. $\neg\forall^* x\ \neg\varphi \leftrightarrow \exists^* x\ \varphi$.

3. $(\forall^* x\ \varphi \wedge \forall^* x\ \psi) \leftrightarrow \forall^* x\ (\varphi \wedge \psi)$;
4. $(\exists^* x\ \varphi \vee \exists^* x\ \psi) \leftrightarrow \exists^* x\ (\varphi \vee \psi)$;
5. $\exists y\ \forall^* x\ \varphi$ implies $\forall^* x\ \exists y\ \varphi$;
6. $\forall^* x\ \forall y\ \varphi$ implies $\forall y\ \forall^* x\ \varphi$; and
7. $\forall x\ \varphi$ implies $\forall^* x\ \varphi$, which implies $\exists^* x\ \varphi$, which implies $\exists x\ \varphi$.

Proof .:.

1. Suppose $\{x \in K : \varphi(x)\} \in U$. If $\neg\exists^* x\ \varphi$ then $K \setminus \{x \in K : \varphi(x)\} \in U$. By closure under intersections, this would imply $\emptyset \in U$, a contradiction. For the other direction, suppose $\exists^* x\ \varphi$, meaning $K \setminus \{x \in K : \varphi(x)\} \notin U$. By Result 1 A • 5, this means the complement $\{x \in K : \varphi(x)\} \in U$, meaning $\forall^* x\ \varphi$.
2. We have that $\neg\forall^* x\ \neg\varphi$ iff $\{x \in K : \neg\varphi(x)\} = K \setminus \{x \in K : \varphi(x)\} \notin U$ iff $\exists^* x\ \varphi$.
3. The '←' direction is immediate since filters are closed under supersets: $\forall x\ (\varphi \wedge \psi \to \varphi)$ and $\forall^* x\ (\varphi \wedge \psi)$ implies $\forall^* x\ \varphi$ and similarly for ψ. For the '→' direction, use that filters are closed under intersections.
4. This is (3) used with the fact that $\neg(\varphi \wedge \psi)$ is equivalent to $\neg\varphi \vee \neg\psi)$.
5. This is just from basic first-order logic: if there is a y such that $\{x \in K : \varphi(x, y)\} \in U$ then as a superset, $\{x \in K : \exists y\ \varphi(x, y)\} \in U$.
6. If $\forall^* x\ \forall y\ \varphi$, then $\{x \in K : \forall y\ \varphi(x, y)\} \in U$. This set is contained in $\{x \in K : \varphi(x, y)\}$ for any given y, meaning that for every y, $\{x \in K : \varphi(x, y)\} \in U$. Hence $\forall y\ \forall^* x\ \varphi$.
7. These implications are is clear: $K \in U$ so $\forall x\ \varphi$ implies $\{x \in K : \varphi(x)\} = K \in U$. The second implication is from (1). The third implication follows from $K \in U$: if $\exists^* x\ \varphi$ then $K \setminus \{x \in K : \varphi(x)\} \notin U$. But $\neg\exists x\ \varphi$ would imply $K = K \setminus \{x \in K : \varphi(x)\} \notin U$, a contradiction. ⊣

The weakness of (5) and (6) cannot be improved in general, since

- $\forall^* \alpha\ \exists x\ (x = \alpha)$ doesn't satisfy $\exists x\ \forall^* \alpha\ (x = \alpha)$ unless U is principal; and
- Often $\forall \beta\ \forall^* \alpha\ (\alpha > \beta)$—i.e. almost everything is bigger than any particular β—but we likely won't have $\forall^* \alpha\ \forall \beta\ (\alpha > \beta)$—i.e. almost every α is bigger than *everything*.

To find specific examples where this happens, we need to consider some particular properties of ultrafilters. Note that when considering ultrafilters, (1) implies that we don't need the notation of $\exists^*$. But the two are distinct for filters. For example, the filter of measure one subsets of $[0, 1] \subseteq \mathbb{R}$ has that $\exists^* x\ (0 \le x \le 1/2)$ but clearly $\neg\forall^* x\ (0 \le x \le 1/2)$.

§ 1 D. Ultrafilter properties

With all of this logical notation at our disposal, we can more easily state some definitions.

1 D • 1. Definition

Let κ be a cardinal, and let U be a filter over κ.

- U is *uniform* iff $|x| = \kappa$ for every $x \in U$.
- U is *unbounded* iff for every $\beta < \kappa$, $\forall^* \alpha\ (\alpha > \beta)$.
- U is *normal* iff for every f such that $\forall^* \alpha\ (f(\alpha) < \alpha)$ there is some $\beta < \kappa$ with $\exists^* \alpha\ (f(\alpha) = \beta)$.
- U is *η-complete* iff U is closed under $< \eta$-intersections: for $\gamma < \eta$ and formulas φ_ξ for $\xi < \gamma$, $\bigwedge_{\xi<\gamma}(\forall^* \alpha\ \varphi_\xi)$ iff $\forall^* \alpha\ (\bigwedge_{\xi<\gamma} \varphi_\xi)$. Equivalently, for $\gamma < \eta$, $\{X_\alpha : \alpha < \gamma\} \subseteq U$ implies $\bigcap_{\alpha<\gamma} X_\alpha \in U$.

The club filter over a regular, uncountable cardinal will have all of these properties, for example, although it isn't an ultrafilter.

1 D • 2. Example

Let κ be an uncountable, regular cardinal. Let Club_κ be the club filter over κ (the filter generated by closed, unbounded [i.e. club] subsets of κ). Therefore Club_κ is uniform, unbounded, normal, and κ-complete.

Proof .:.

Obviously all club sets are unbounded, and so by regularity of κ, each club has cardinality κ. So Club_κ is uniform. That Club_κ is normal is just Fodor's lemma, since each element of Club_κ contains a club and is thus stationary. In other words, Fodor's Lemma (1 B • 5) tells us a regressive f is constant on a stationary set which is therefore Club_κ-positive: the complement doesn't contain a club. κ-completeness follows from facts about club sets (closure of the intersection is immediate, and for unboundedness, interlace a λ-length sequence of members in the $\lambda < \kappa$ clubs and take the supremum which is in all the clubs). ⊣

Many of the properites of Definition 1 D • 1 are connected as shown below for ultrafilters over a cardinal κ. Implications are denoted with arrows (a dashed arrow denotes the implication merely of the *existence* of an ultrafilter with both properties). If κ is regular, unboundedness is equivalent to uniformity, which is otherwise stronger.

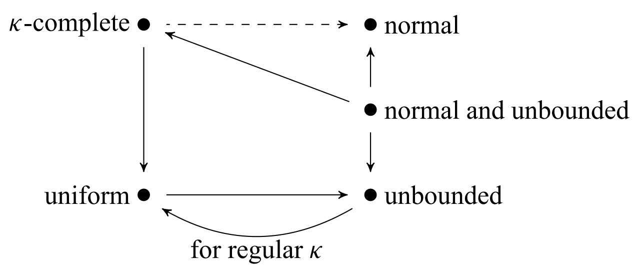

1 D • 3. Figure: Properties of non-principal ultrafilters over $\kappa \geq \aleph_0$

Arguably the most difficult of the properites in Figure 1 D • 3 to achieve is normality, which isn't directly implied by any combination of the other properties. However the strongest two properties here are clearly κ-completeness and normality. This combination is important enough to get its own name.

1 D • 4. Definition

Let κ be an uncountable cardinal with U a non-principal ultrafilter over κ. We say that U is a *measure* iff U is κ-complete and normal.

We call this a *measure* as motivated from the fact that the function

$$\mu(X) := \begin{cases} 1 & \text{if } X \in U \\ 0 & \text{if } X \notin U \end{cases}$$

is a κ-additive, two-valued, probability measure over κ. Other authors often drop the requirement of normality in the definition of a measure (which makes sense with this motivation), but then always work with normal measures.

What this definition tells us is that the dashed arrowof Figure 1 D • 3 means that if κ has a κ-complete, non-principal ultrafilter, then it also has a measure. Another notable property of κ-completeness is that it implies that κ is regular. To see this, since it implies unboundedness, if $\langle \gamma_\beta : \beta < \text{cof}(\kappa) \rangle$ is unbounded with $\text{cof}(\kappa) < \kappa$, then the infinitary conjunction $\bigwedge_{\beta < \text{cof}(\kappa)} \forall^* \alpha\ (\alpha > \gamma_\beta)$ implies $\forall^* \alpha\ (\alpha > \gamma_\beta$ for all $\beta < \text{cof}(\kappa))$, which contradicts that the sequence of γ_βs is unbounded in κ.

The two difficult arrows in Figure 1 D • 3 are that κ-complete, non-principal ultrafilters yield measures, and that unbounded, normal, non-principal ultrafilters are measures. The second of these is easier.

1 D • 5. Result

Let κ be an uncountable cardinal with U an unbounded, normal ultrafilter over κ. Therefore U is κ-complete, and hence a measure.

Proof .:.

Let $\{X_\alpha : \alpha < \lambda\} \in \mathcal{P}(U)$. If $\bigcap_{\alpha<\lambda} X_\alpha \notin U$, we may assume without loss of generality that $\bigcap_{\alpha<\lambda} X_\alpha = \emptyset$. So define $f : \kappa \to \lambda$ to be such that $f(\alpha)$ is the least $\xi < \lambda$ with $\alpha \notin X_\xi$. As $\lambda < \kappa$, $f``\kappa$ is bounded in κ and thus (as U is unbounded),

$$\forall^* \alpha\ (\alpha > \sup(f``\kappa)) \quad \text{implies} \quad \forall^* \alpha\ (\alpha > f(\alpha)).$$

So f is regressive on a set in U. By normality, there is then a $\beta < \kappa$ with $\forall^* \alpha\ (f(\alpha) = \beta)$. But this means

$\forall^* \alpha\ (\alpha \notin X_\beta)$, contradicting that $X_\beta \in U$. ⊣

Normality is kind of a weird definition, but its usefulness will become more apparent as we investigate ultrapowers and elementary embeddings. Really, one should think of the dashed arrow of Figure 1 D • 3 as being a property of κ rather than of the ultrafilters. We could still prove now, without reference to later material, the dashed arrow: we can get normal ultrafilters through *possibly different* κ-complete ultrafilters. However, the proof of this with our current understanding is not the best proof as it is fairly technical without additional concepts. But with later material, the idea becomes much more natural.

It should be noted that the existence of measures isn't provable just from ZFC. The reason for this is that any κ which admits such a measure, called a *measurable cardinal*, will be quite large. We will see later that they will be inaccessible, for example, and thus can't be shown to exist just from ZFC. But they will be much more and much larger than mere inaccessibles. To further investigate measurable cardinals, it is useful to take a look at ultrapowers and elementary embeddings.

Section 2. Ultrapowers and Elementary Embeddings

At this point, we can get to our first true application of filters. Ultraproducts are a model-theoretic notion which serve two purposes. Firstly, they are a sort of average of the starting models: what's true in the ultraproduct is what's "almost always" true in the models. Secondly, they are a way to enlarge the universe: the ultraproduct using just one universe yields an elementary embedding. The usefulness for set theory comes when the ultraproduct is well-founded so that it may be collapsed down into an inner model.

In model theory, there is a general concept of a *reduced product* where you kind of "average out" a set of models over a filter. We will not be so concerned with general reduced products, since we will be focused on ultrafilters. The idea is that your objects are now sequences of elements in these models, and a statement $\varphi(\vec{f})$ is true iff for almost every α, $\varphi(\vec{f}(\alpha))$ is true in the corresponding model. So in particular, a sentence is true iff it is true in "most" of the models. The result is an ultraproduct instead of a mere reduced product. Even still, we will not be concerned with ultraproducts in general, but ultraproducts where the models we're "averaging" are all the same model.

2•1. Definition

Let σ be a signature. Let A be an first-order logic model for σ. Let U be an ultrafilter over a set K.
For $f, g : K \to A$, say $f \approx g$ iff $\forall^* x\ (f(x) = g(x))$. The *ultrapower* of A by U is the structure $\mathsf{Ult}(\mathsf{A}, U)$

- with universe $[f]_U = \{g : f \approx g\}$;
- σ-relation interpretations $R^{\mathsf{Ult}(\mathsf{A},U)}([\vec{f}])$ iff $\forall^* x\ R^{\mathsf{A}}(\vec{f}(x))$; and
- σ-function interpretations $F^{\mathsf{Ult}(\mathsf{A},U)}([\vec{f}]) = [g]$ iff $\forall^* x\ (F^{\mathsf{A}}(\vec{f}(x)) = g(x))$.

One should check that these are actually well-defined, but this is easy given that U is an ultrafilter. This, however, is not very difficult as filters are closed under finite intersections.

What's happening here is that the choice of what is true at the level of atomic formulas is left up to the ultrafilter: what happens often enough in the factors happens in the ultraproduct. This goes through to all levels of first-order formula complexity, as shown in the following indispensable theorem known as Łoś's Theorem. The theorem fully characterizes first-order truth in ultrapowers based on first-order truth in the factors. As the name "Łoś" is Polish, it is pronounced [ˈwɔɕ]. Also note that the statement of the theorem here uses only one parameter x in $\varphi(x)$, but the result actually allows for arbitrarily many: $\varphi(\vec{x})$. This would clutter notation for the proof, which essentially the same either way.

2•2. Theorem (Łoś's Theorem)

Let σ be a signature. Let A be an first-order logic model for σ. Let U be an ultrafilter over K, and write Ult for $\mathsf{Ult}(\mathsf{A}, U)$.
Let $\varphi(x)$ be an first-order formula in the signature σ, and let $[f] \in \mathrm{Ult}$ be a parameter.
Therefore $\mathsf{Ult} \vDash \varphi([f])$ iff $\forall^* x[\mathsf{A} \vDash \varphi(f(x))]$.

Proof .:.

Write Ult for $\mathsf{Ult}(\mathsf{A}, U)$. This is a proof by structural induction on φ. For atomic formulas, this is just by definition. Sentential connectives $\wedge$ and $\neg$ follow easily as well, since U is an ultrafilter:

$$\begin{aligned} \mathsf{Ult} \vDash \text{“}\neg\varphi([f])\text{”} &\text{ iff } \mathsf{Ult} \nvDash \text{“}\varphi([f])\text{”} \text{ iff } \neg\forall^* x\ (\mathsf{A} \vDash \text{“}\varphi(f(x))\text{”}) \\ &\text{ iff } \forall^* x\ (\mathsf{A} \nvDash \text{“}\varphi(f(x))\text{”}) \text{ iff } \forall^* x\ (\mathsf{A} \vDash \text{“}\neg\varphi(f(x))\text{”}). \\ \mathsf{Ult} \vDash \text{“}\varphi([f]) \wedge \psi([f])\text{”} &\text{ iff } \mathsf{Ult} \vDash \text{“}\varphi([f])\text{” and } \mathsf{Ult} \vDash \text{“}\psi([f])\text{”} \\ &\text{ iff } \forall^* x\ \big[\mathsf{A} \vDash \text{“}\varphi(f(x))\text{”}\big] \wedge \forall^* x\ \big[\mathsf{A} \vDash \text{“}\psi(f(x))\text{”}\big] \\ &\text{ iff } \forall^* x\ \big[\mathsf{A} \vDash \text{“}\varphi(f(x)) \wedge \psi(f(x))\text{”}\big]. \end{aligned}$$

For existential quantification, suppose $\mathsf{Ult} \vDash$ "$\exists y\ \varphi(y, [f])$". Thus there is some $[g] \in \mathrm{Ult}$ where $\mathsf{Ult} \vDash$

"$\varphi([g],[f])$". By (a modified version of) the inductive hypothesis, this happens iff $\forall^* x$ ($\mathsf{A} \vDash$ "$\varphi(g(x), f(x))$"), and so clearly it follows that $\forall^* x$ ($\mathsf{A} \vDash$ "$\exists y\ \varphi(y, f(x))$").

For the other direction, we need AC: for each $x \in K$ such that $\mathsf{A} \vDash$ "$\exists y\ \varphi(y, f(x))$", let $g(x)$ witness this. Otherwise, let $g(x)$ be any arbitrary element of A_x. The resulting function $g = \{\langle x, g(x)\rangle : x \in K\}$ witnesses that $\mathsf{Ult} \vDash$ "$\varphi([g],[f])$" and thus that $\mathsf{Ult} \vDash$ "$\exists y\ \varphi(y,[f])$". ⊣

2 • 3. Corollary

Any model is elementarily equivalent to any of its ultrapowers.

§ 2 A. Elementary embeddings

We have actually a much stronger correspondence between truth in A and its ultrapowers, but to talk further about this relation, we need the concept of an elementary embedding. The issue is that the two models of universes composed of fundamentally different things, and so we can't just compare them outright. Instead, we translate by a function.

2 A • 1. Definition

Let σ be a signature. Let A and B be first-order logic models for σ. Let $f : A \to B$ be a function. f is an *elementary embedding* iff for all formulas φ in the signature σ and parameters $a_0, \cdots, a_n \in A$,

$$\mathsf{A} \vDash \text{“}\varphi(a_0, \cdots, a_n)\text{”} \quad \text{iff} \quad \mathsf{B} \vDash \text{“}\varphi(f(a_0), \cdots, f(a_n))\text{”}.$$

Any elementary embedding will be an embedding just by considering the atomic formulas: $x \in^{\mathsf{A}} y$ iff $f(x) \in^{\mathsf{B}} f(y)$. It should be obvious from this that any elementary embedding is injective: for $x, y \in A$,

$$x \neq y \quad \text{iff} \quad \mathsf{A} \vDash \text{“}x \neq y\text{”} \quad \text{iff} \quad \mathsf{B} \vDash \text{“}f(x) \neq f(y)\text{”} \quad \text{iff} \quad f(x) \neq f(y).$$

Elementary embeddings aren't necessarily surjective, however, meaning that they are stronger than a mere embedding, but weaker than a full isomorphism.

Elementary embeddings are crucial to the understanding of ultrapowers and inner models. Some of the basic facts are not recorded because they are seen to be obvious. To better familiarize the reader with some of these basic facts, the following extensively used results will be given explicit proofs.

2 A • 2. Lemma

Let $j : \mathrm{V} \to \mathrm{M}$ be elementary with $\mathrm{M} \subseteq \mathrm{V}$ a transitive class. Therefore the following hold for all formulas φ.

1. If φ is absolute between V and M, then $\varphi(x)$ iff $\varphi(j(x))$. Hence if x is defined by an absolute formula—meaning $y = x$ iff $\varphi(y)$—then $j(x) = x$.
2. If f is a function, then $j(f)$ is a function, and $j(f(x)) = j(f)(j(x))$.
3. If f, g are functions, then $j(f \circ g) = j(f) \circ j(g)$.

Proof .:.

1. By elementarity, $\varphi(x)$ iff $\mathsf{M} \vDash$ "$\varphi(j(x))$". By absoluteness, this happens iff $\varphi(j(x))$. Now if x is defined by φ—i.e. $x = y$ iff $\varphi(y)$—for some absolute formula φ, then $\mathsf{M} \vDash$ "$\varphi(j(x))$". By absoluteness, $\varphi(j(x))$ so that $x = j(x)$.

2. Being an ordered pair is definable by a formula absolute between transitive models. So being a set of ordered pairs, $x \in f$ implies x is an ordered pair. Elementarity then gives that every $x \in j(f)$ has that x is an ordered pair. Moreover, f is a function iff $\forall x\ (\exists y\ \langle x, y\rangle \in f \to \exists! y\ \langle x, y\rangle \in f)$, which is absolute between transitive models. By elementarity,

$$\mathsf{M} \vDash \text{“}\forall x\ (\exists y\ \langle x, y\rangle \in j(f) \to \exists! y\ \langle x, y\rangle \in j(f))\text{”}$$

By absoluteness, this holds in V so that $j(f)$ is then a function.

Let $x \in \mathrm{dom}(f)$ be arbitrary. $f(x)$ is the unique y such that $\langle x, y\rangle \in f$. Hence $j(f(x))$ is the unique y such that $\langle j(x), y\rangle \in j(f)$. Hence $j(f)$ is a function, and it obeys $j(f)(j(x)) = j(f(x))$ whenever $x \in \mathrm{dom}(f)$.

3. Following easily from (1), for all x and z, $\langle x,z\rangle \in f\circ g$ iff there is some $y\in f$ where $\langle x,y\rangle \in g$ and $\langle y,z\rangle\in f$. By elementarity, for all x,z, $\langle x,z\rangle\in j(f\circ g)$ iff $\exists y\ (\langle x,y\rangle\in j(g)\wedge\langle y,z\rangle\in j(f))$, meaning $j(f\circ g)=j(f)\circ j(g)$. ⊣

Often arguments like these will be written in shorthand, so it's important to know what will be moved by j or won't be moved by j. For example, consider $\gamma=\{\alpha:\alpha<\gamma\}$. Note that $j(\gamma)$ will generally not be $\{j(\alpha):j(\alpha)<j(\gamma)\}$: α is a dummy variable that plays no role here, but γ is still a parameter. So $j(\gamma)=\{\alpha:\alpha<j(\gamma)\}$ (as expected). Similarly, $j(\{x:f(x)=\alpha\})=\{x:j(f)(x)=j(\alpha)\}$, again, x is just a dummy variable, but f and α aren't.

Before heading too deep into this, however, we need to think about how we regard these elementary embeddings. As functions from perhaps the entire universe of sets, they will not be sets. And a priori, there's no reason to think they need to be definable. To counter this issue, we will work in a relatively simple class theory, like $\mathsf{NBG}-\mathsf{GC}+\mathsf{AC}$—von Neumann–Bernays–Gödel class theory with choice for sets. In other words, we have **V** being the usual set-theoretic universe adjoining predicates for classes of V (whatever those happen to be, but at least including the definable classes). One can show that this is a conservative extension of ZFC, meaning that no new theorems with quantifiers ranging over sets are proven by NBG.

Obviously there is a kind of trivial elementary embeddings from V into an inner model: the identity map. This map isn't exactly interesting, however, and so we will be interested with maps that actually move sets. It turns out that if an elementary embedding moves a set, then it moves an ordinal.

2 A • 3. Result

Let $j:\mathrm{V}\to\mathrm{M}$ be elementary with $\mathrm{M}\subseteq\mathrm{V}$ a transitive class. Let $\alpha\in\mathrm{Ord}$. Therefore $j\restriction\alpha=\mathrm{id}\restriction\alpha$ iff $j\restriction\mathrm{V}_\alpha=\mathrm{id}\restriction\mathrm{V}_\alpha$.

Proof .:.

Obviously $j\restriction\mathrm{V}_\alpha=\mathrm{id}\restriction\mathrm{V}_\alpha$ implies $j\restriction\alpha=\mathrm{id}\restriction\alpha$ since $\alpha\subseteq\mathrm{V}_\alpha$.

For the other direction, in essence, the rank of $x\in\mathrm{V}_\alpha$ is still preserved. For $\alpha=\emptyset$ and α a limit, the result clearly holds. For the successor case, we assume $j\restriction\alpha=\mathrm{id}\restriction\alpha$ and that $j(\alpha)=\alpha$. Let $x\in\mathrm{V}_{\alpha+1}$. Hence $x\subseteq\mathrm{V}_\alpha$ so by elementarity, $j(x)\subseteq\mathrm{V}^{\mathsf{M}}_{j(\alpha)}=\mathrm{V}_\alpha\cap\mathrm{M}$ since M is transitive. Now for any fixed $y\in\mathrm{V}_\alpha$, $y\in j(x)$ inductively is equivalent to $y=j(y)\in j(x)$. So by elementarity, this is equivalent to $y\in x$. Thus $j(x)=x$, as desired. ⊣

2 A • 4. Corollary

If $j:\mathrm{V}\to\mathrm{M}$ is elementary into $\mathrm{M}\subseteq\mathrm{V}$ a transitive class and $j\neq\mathrm{id}$, then the least α with $j(\alpha)\neq\alpha$ is also the least rank of a set moved by j.

This motivates the following definition of a critical point, below which j is just the identity, and which is moved by j.

2 A • 5. Definition

Let $\mathrm{M}\subseteq\mathrm{V}$ be a transitive class of **V** with $j:\mathrm{V}\to\mathrm{M}$ elementary such that $j\neq\mathrm{id}$.
An ordinal α is a *critical point* of j—denoted $\mathrm{cp}(j)$—iff α is the least ordinal where $\alpha\neq j(\alpha)$.

If $\alpha<\kappa=\mathrm{cp}(j)$ then by elementarity $\alpha=j(\alpha)<j(\kappa)$ so that $j(\kappa)>\kappa$. So the first ordinal moved is always moved up. This implies that nontrivial elementary embeddings will never be surjective: no ordinal α with between $j(\kappa)$ and κ (more precisely, $\kappa\le\alpha<j(\kappa)$) is in the image of j.

On the topic of the identity embedding, there is a kind of ceiling to how close $\mathrm{M}\subseteq\mathrm{V}$ can be to V when $j:\mathrm{V}\to\mathrm{M}$ is elementary. The following theorem tells us that in particular, M cannot be V. This is important for ruling out the existence of reinhardt cardinals—the critical points of elementary embeddings from V into itself, rather than merely an inner model. The proof of this theorem can only be given after we introduce the concept of measurable cardinals.

2 A • 6. Theorem (Kunen's Inconsistency Theorem)

Let $j:\mathrm{V}\to\mathrm{V}$ elementary. Therefore $j=\mathrm{id}$.
More precisely, suppose $\mathrm{M}\subseteq\mathrm{V}$ is a transitive class and $j:\mathrm{V}\to\mathrm{M}$, also a class of V, is elementary. Therefore $j\neq\mathrm{id}$ implies $\mathrm{M}\neq\mathrm{V}$.

A nice property of elementary j from V into classes of V is that they will preserve $\wp(\mathrm{cp}(j))$. In general, if $j : N \to M$ is elementary between two transitive classes, there's no guarantee that the powerset is preserved, and we'd only get $\wp(\mathrm{cp}(j)) \cap N \subseteq \wp(\mathrm{cp}(j)) \cap M$.

2 A • 7. Result

Let $\mathrm{M} \subseteq \mathrm{V}$ be a transitive class. Let $j : \mathrm{V} \to \mathrm{M}$ be elementary with $\mathrm{cp}(j) = \kappa$. Therefore $\wp(\kappa)^{\mathsf{M}} = \wp(\kappa)$. In fact, $\mathrm{V}_{\kappa+1} \subseteq \mathrm{M}$ so that $x \subseteq \mathrm{V}_\kappa$ is just $x = j(x) \cap \mathrm{V}_\kappa$.

Proof .:.

As M is transitive, $\wp(\kappa)^{\mathsf{M}} = \wp(\kappa) \cap \mathrm{M}$. Let $x \in \wp(\kappa) \cap V$. For every α, since $x \subseteq \kappa$,

$$\alpha \in x \quad \text{iff} \quad j(\alpha) \in j(x) \wedge \alpha = j(\alpha) < \kappa,$$

Hence $x = j(x) \cap \kappa \in \wp(\kappa) \cap \mathrm{M}$, and thus $\wp(\kappa) \subseteq \wp(\kappa) \cap \mathrm{M}$. And since $\mathrm{M} \subseteq \mathrm{V}$, $\wp(\kappa) \cap \mathrm{M} \subseteq \wp(\kappa) \cap \mathrm{V}$.

By Result 2 A • 3, looking at $j \upharpoonright \mathrm{V}_{\alpha+1}$, $j(\mathrm{V}_\alpha) = \mathrm{V}_\alpha$ for each $\alpha < \kappa$ so that

$$\mathrm{V}_\kappa^{\mathsf{M}} = \bigcup_{\alpha<\kappa} \mathrm{V}_\alpha^{\mathsf{M}} = \bigcup_{\alpha<\kappa} \mathrm{V}_\alpha = \mathrm{V}_\kappa.$$

Thus $\mathrm{V}_\kappa \in \mathrm{M}$. Now consider $x \subseteq \mathrm{V}_\kappa$. Since $j \upharpoonright \mathrm{V}_\kappa = \mathrm{id} \upharpoonright \mathrm{V}_\kappa$ again follows from Result 2 A • 3, we have by elementarity that $y \in j(x) \cap \mathrm{V}_\kappa$ iff $y \in x$, which means that $j(x) \cap \mathrm{V}_\kappa = x \in \mathrm{M}$ and thus $\mathrm{V}_{\kappa+1} \subseteq \mathrm{M}$. ⊣

Hence the "strength" of a non-trivial, elementary embedding $j : \mathrm{V} \to \mathrm{M}$ is at least $\mathrm{cp}(j) + 1$ in the sense that we always have $\mathrm{V}_{\mathrm{cp}(j)+1} \subseteq \mathrm{M}$. It may be possible[iv] for j to have a larger strength, but to do this, we would need *extenders* rather than mere ultrafilters (this also motivates the notion of a *strong* cardinal). But now that we have thought about elementary embeddings in general, let us return to the notion of an ultrapower.

§ 2 B. Characterizing ultrapowers

With all of this talk about elementary embeddings, we should perhaps note that we always have an elementary embedding from a model into its ultrapower.

2 B • 1. Theorem

Let σ be a signature. Let $\mathbf{A}$ be a first-order logic model for σ, and let U be an ultrafilter over a set K. Therefore $\mathbf{A}$ is elementarily embedded in $\mathbf{Ult}(\mathbf{A}, U)$ by $x \mapsto [\mathrm{const}_x]_U$, where $\mathrm{const}_x : K \to A$ is the constant x map.

Proof .:.

By Łoś's Theorem (2 • 2), $\mathbf{Ult}(\mathbf{A}, U) \vDash$ "$\varphi([\mathrm{const}_x])$" iff for almost every $k \in K$, $\mathbf{A} \vDash$ "$\varphi(\mathrm{const}_x(k))$" (i.e. $\mathbf{A} \vDash$ "$\varphi(x)$") which is just to say that the following set is in U:

$$\{k \in K : \mathbf{A} \vDash \text{“}\varphi(x)\text{”}\} = \begin{cases} K & \text{if } \mathbf{A} \vDash \text{“}\varphi(x)\text{”} \\ \emptyset & \text{otherwise.} \end{cases}$$

As an ultrafilter over K, $\emptyset \notin U$ and $K \in U$ so that $\mathbf{Ult}(\mathbf{A}, U) \vDash$ "$\varphi([\mathrm{const}_x])$" iff $\mathbf{A} \vDash$ "$\varphi(x)$". ⊣

Note that for a proper class like V, each equivalence class $[f] \in \mathrm{Ult}(\mathrm{V}, U)$ will be a proper class as well.[v] This can be rectified if we just consider $[f] = \{g : g \approx f \wedge \mathrm{rank}(g) \text{ is minimal}\}$. Doing this, we get the usual equivalence class of f just intersected with some V_α for α least. Doing this, one still has that $[f] = [g]$ for all $f \approx g$, and thus $x \mapsto [\mathrm{const}_x]$ is a legitimate (class) function. So the result above also holds with proper classes too under this variant definition.

The existence of such an elementary embedding, however, doesn't tell you that it's nontrivial.

[iv] This consistency of this can't be proven in ZFC alone, as such embeddings yield the existence of certain large cardinals, which in turn imply the consistency of ZFC.

[v] To see this, let $v \in \mathrm{V}$ be arbitrary. As U is an ultrafilter over K, let $X \in U$ be such that $X \neq K$ so that there is some $x \in K \setminus X$. Now for any $v \in \mathrm{V}$, consider f_v to be the map sending every $y \in K$ to x except for x itself, which is sent to v, i.e. $f_v = (\mathrm{const}_x \setminus \{\langle x, x\rangle\}) \cup \{\langle x, v\rangle\}$. Note that $\forall y \in X\ (f_v(y) = \mathrm{const}_x(y))$ so that $f_v \in [\mathrm{const}_x]$. Also note that $f_v \neq f_{v'}$ for $v \neq v' \in \mathrm{V}$ so that $[\mathrm{const}_x]$ is a proper class.

2 B • 2. Result

Let U be a *principal* ultrafilter over a set K. Therefore $\mathsf{A} \cong \mathsf{Ult}(\mathsf{A}, U)$ by the canonical embedding.

Proof .:.

It suffices to show that the canonical embedding of Theorem 2 B • 1 is surjective. Let $u \in \mathrm{Ult}(A, U)$ be arbitrary. We know that $u = [f]$ for some $f : K \to A$. As U is principal, there is some $a \in A$ where $\{a\} \in U$. Hence $g \approx f$ iff $g(a) = f(a)$. In particular, $u = [\mathrm{const}_{f(a)}]$. Hence $x \mapsto [\mathrm{const}_x]$ is a bijective embedding, and thus an isomorphism. ⊣

As the transitive collapse of a well-founded structure is unique, it follows that the collapsed version of the ultrapower $\mathrm{Ult}(\mathrm{V}, U)$ (if well-founded) is precisely V, and it's not difficult to show inductively that in this case, $[\mathrm{const}_x]$ is collapsed to x.

The importance of using ultrapowers, however, is when they are well-founded, because set theory in general is more concerned with transitive models. Transitive models are easier to work with, and are related to the actual universe of sets. Hence studying transitive models allows us to learn about the actual set-theoretic universe—though perhaps only conditional on large cardinal assumptions or other hypotheses. Well-founded models that satisfy extensionality are just an isomorphism away from transitive models by a mostowski collapse. In general, not all ultrapowers will be well-founded. Crucially, if we can take countably many conjunctions, then the ultrapower must be well-founded: we can collect together the countable amount of information that $f_0 \ni f_1$, and $f_1 \ni f_2$, and so on all at once. It's nice that we then have a characterization of the ultrapower being well-founded.

2 B • 3. Theorem

Let U be an ultrafilter in V. Therefore $\mathsf{Ult}(\mathsf{V}, U)$ is well-founded iff U is $\aleph_1$-complete.

Proof .:.

($\leftarrow$) Suppose U is $\aleph_1$-complete, but $\mathsf{Ult} = \mathsf{Ult}(\mathsf{V}, U)$ is ill-founded. Let $\langle f_n : n \in \omega \rangle \in \mathrm{V}$ be one such descending $\in^{\mathsf{Ult}}$-sequence in Ult: for every $n \in \omega$, $\mathsf{Ult} \models$ "$[f_{n+1}] \in [f_n]$". As U is $\aleph_1$-complete in V,

$$\mathsf{V} \models \text{“}\bigwedge_{n\in\omega} \forall^*\alpha\ (f_{n+1}(\alpha) \in f_n(\alpha))\text{”} \quad \text{iff} \quad \mathsf{V} \models \text{“}\forall^*\alpha \left(\bigwedge_{n\in\omega} f_{n+1}(\alpha) \in f_n(\alpha)\right)\text{”}.$$

But any such α yields an infinite, decreasing sequence $\langle f_n(\alpha) : n \in \omega \rangle$ in V, contradicting foundation.

($\rightarrow$) Now suppose $\mathsf{Ult}(\mathsf{V}, U)$ is not $\aleph_1$-complete. Let $\{X_n : n \in \omega\} \in \mathcal{P}(U)$ with $\bigcap_{n\in\omega} X_n \notin U$. From this, we will construct an infinite, decreasing $\in^{\mathsf{Ult}}$-sequence in Ult, contradicting foundation. Without loss of generality, assume $X_n \supseteq X_{n+1}$ just by replacing each X_n with $\bigcap_{i\leq n} X_i$. Without loss of generality, U is an ultrafilter over a cardinal κ.

For each $\alpha < \kappa$, let $\mathrm{index}(\alpha)$ be the least n for which $\alpha \notin X_n$. If there is no such n, then write $\mathrm{index}(\alpha) = 0$. For each $n \in \omega$, define the function $f_n : \kappa \to \omega$ by taking, for $\alpha < \kappa$,

$$f_n(\alpha) = \begin{cases} \mathrm{index}(\alpha) - n & \text{if } \mathrm{index}(\alpha) \geq n \\ 0 & \text{otherwise.} \end{cases}$$

So in essence, $\langle f_n(\alpha) : n \in \omega \rangle$ will start at $\mathrm{index}(\alpha)$ and decrease by 1 until it is eventually, constantly 0. As a result, if $\alpha \in X_n$, then $\mathrm{index}(\alpha) > n$ and so $f_n(\alpha) > f_{n+1}(\alpha)$. As almost every α is in X_n, it follows that $\forall^*\alpha\ (f_n(\alpha) > f_{n+1}(\alpha))$. So for each $n \in \omega$, $\mathsf{Ult} \models$ "$[f_n] \in [f_{n+1}]$". Consequently, $\langle [f_n] : n \in \omega \rangle$ is a decreasing $\in^{\mathsf{Ult}}$-sequence, meaning Ult is ill-founded. ⊣

Of course, this doesn't say that $\mathsf{Ult}(\mathsf{V}, U)$ or V is necessarily *actually* well-founded, just that if V is well-founded—if we start from a well-founded class—then we still remain well-founded after taking the ultrapower. You might think that this result is obvious, since if $\mathrm{Ult}(\mathrm{V}, U)$ is a class of V, and V thinks itself is well-founded, surely it must think this class is too. But the issue is the difference in interpretation of '$\in$'. If $\mathsf{Ult}(\mathsf{V}, U)$ is well-founded, then we can identify it with a transitive class of V, but otherwise, it's just some structure whose universe is a class of V.

2 B • 4. Definition

Let U be an $\aleph_1$-complete ultrafilter. Define $\mathbf{cUlt}(\mathsf{V}, U)$ be transitive collapse of $\mathbf{Ult}(\mathsf{V}, U)$ via π_U. Set $j_U : \mathrm{V} \to \mathrm{cUlt}(\mathrm{V}, U)$ to be the canonical embedding: $j_U(x) = \pi_U([\mathrm{const}_x])$.

This doesn't inherently tell us that this (collapsed) ultrapower is different from V, however, which was more obviously the case when U was principal. If V has no measurable cardinals, it will turn out that V has no $\aleph_1$-complete ultrafilters, as such ultrafilters will actually be λ-complete for some (maximal) λ that turns out to be a measurable cardinal.[vi]

The notion of completeness is also important as it determines the critical point of the canonical embedding.

2 B • 5. Theorem

Let U be a non-principal ultrafilter over K that is κ-complete, but not κ^+-complete for some cardinal $\kappa > \aleph_0$. Let $j : \mathrm{V} \to \mathrm{cUlt}(\mathrm{V}, U)$ be the canonical embedding. Therefore, $j \neq \mathrm{id}$ and $\mathrm{cp}(j) = \kappa$.

Proof .:.

Let $\pi : \mathrm{Ult}(\mathrm{V}, U) \to \mathrm{cUlt}(\mathrm{V}, U)$ be the collapsing isomorphism. First we show that $j \restriction \kappa = \mathrm{id} \restriction \kappa$. To see this, suppose inductively that $j \restriction \xi = \mathrm{id} \restriction \xi$ for some $\xi < \kappa$. We aim to show $j(\xi) = \xi$. By elementarity and the inductive hypothesis, $\alpha \in \xi$ iff $j(\alpha) = \alpha \in j(\xi)$ so that $j(\xi) \supseteq \xi$. So it suffices to show $\pi([\mathrm{const}_\xi]) = j(\xi) \leq \xi$.

So let $\zeta < j(\xi)$ be arbitrary. ζ can be represented in the ultrapower by some $f : K \to \mathrm{V}$: $\zeta = \pi([f])$. Since $\mathbf{Ult}(\mathsf{V}, U) \vDash$ "$[f] < [\mathrm{const}_\xi]$", it follows that

$$\forall^* x\ (f(x) < \mathrm{const}_\xi(x)) \quad \text{iff} \quad \forall^* x\ (f(x) < \xi) \quad \text{iff} \quad \forall^* x \left(\bigvee_{\varepsilon < \xi} f(x) = \varepsilon \right).$$

Suppose for each $\varepsilon < \xi$ that $\neg\forall^* x\ (f(x) = \varepsilon)$ iff $\forall^* x\ (f(x) \neq \varepsilon)$. As $\xi < \kappa$, by κ-completeness,

$$\forall^* x \left(\bigwedge_{\varepsilon < \xi} f(x) \neq \varepsilon \right) \quad \text{iff} \quad \forall^* x \left(\neg \bigvee_{\varepsilon < \xi} f(x) = \varepsilon \right) \quad \text{iff} \quad \forall^* x\ (f(x) \not< \xi),$$

a contradiction. Hence there must be some $\varepsilon < \xi$ where $\forall^* x\ (f(x) = \varepsilon = \mathrm{const}_\varepsilon(x))$. For this ε, we then have $\mathbf{Ult}(\mathsf{V}, U) \vDash$ "$[f] = [\mathrm{const}_\varepsilon]$" so after collapsing, $\zeta = \pi([f]) = j(\varepsilon) = \varepsilon < \xi$. This shows $j(\xi) \subseteq \xi$ and thus equality.

To see that $j \neq \mathrm{id}$, it suffices to show $j(\kappa) > \kappa$. This also shows that $\mathrm{cp}(j) = \kappa$, since we already know $j \restriction \kappa = \mathrm{id} \restriction \kappa$. To do this, we find a function $f : K \to \kappa$ sitting between every $[\mathrm{const}_\alpha]$ and $[\mathrm{const}_\kappa]$ in the ultrapower. To construct f, proceed as follows. As U is not κ^+-complete, let $\langle X_\alpha : \alpha < \kappa \rangle$ witness this: $X_\alpha \in U$ for each $\alpha < \kappa$, but $\bigcap_{\alpha<\kappa} X_\alpha \notin U$. By subtracting this intersection we can assume without loss of generality that $\bigcap_{\alpha<\kappa} X_\alpha = \emptyset$. Furthermore, by using κ-completeness, each $\bigcap_{\xi<\alpha} X_\xi \in U$ so we can without loss of generality obtain a sequence where $X_\alpha \subseteq X_\beta \in U$ for $\beta < \alpha < \kappa$. Consider the map $f : K \to \kappa$ sending $x \in K$ to the least $\alpha < \kappa$ with $x \notin X_\alpha$. Now consider $[f]$ in the ultrapower.

Note that for each α, almost every $x \in K$ is in X_α. In particular, for any fixed α, almost every $x \in K$ has $f(x) > \alpha$. So in the ultrapower, for each $\alpha < \kappa$, $\mathbf{Ult}(\mathsf{V}, U) \vDash$ "$[f] > [\mathrm{const}_\alpha]$" so in taking the transitive collapse, $\pi([f]) > j(\alpha) = \alpha$. In particular, $\pi([f]) \geq \kappa$. But clearly, as f is a function from K to κ, $\forall^* x\ (f(x) < \mathrm{const}_\kappa(x))$ and therefore $\mathbf{Ult}(\mathsf{V}, U) \vDash$ "$[f] < [\mathrm{const}_\kappa]$", meaning $\pi([f]) < j(\kappa)$. Hence $\kappa \leq \pi([f]) < j(\kappa)$ so that $\mathrm{cp}(j) = \kappa$. ⊣

We now aim to prove two main theorems about ultrapowers dealing with how we can factor embeddings through ultrapowers, and how we may represent the elements of ultrapowers. The idea is that an arbitrary elementary embedding $j : \mathrm{V} \to \mathrm{M}$ can be coded through an ultrafilter U_j over $\mathrm{cp}(j)$ and thus through the ultrapower.

To derive an ultrafilter from j, note that for $\kappa = \mathrm{cp}(j)$, most subsets of κ will be shot up beyond κ in the sense that $A \subseteq \kappa$ will likely have $j(A)$ be unbounded in $j(\kappa) > \kappa$. In this sense, $j(A)$ will have many more elements above those in A. The key thing for us is whether κ is in this stretched version of A, $j(A)$. This clearly is answerable for any

[vi] This can be proven just by translating the ultrafilter U to a separate ultrafilter on λ according to how elements disappear from a λ-length $\subseteq$-decreasing sequence of elements of U.

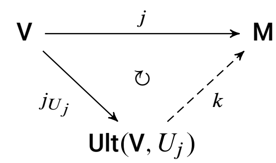

2 B • 6. Figure: Factoring through the ultrapower embedding

given subset of A, and by elementarity, will be preserved under the necessary operations.

2 B • 7. Definition

Let $\mathrm{M} \subseteq \mathrm{V}$ be a transitive class. Let $j : \mathrm{V} \to \mathrm{M}$ be elementary with $\mathrm{cp}(j) = \kappa$. Define the *ultrafilter derived from* j to be $U_j = \{A \subseteq \kappa : \kappa \in j(A)\}$.

As described above, it's not difficult to see that U_j is an ultrafilter. More importantly, U_j is actually a measure over κ.

2 B • 8. Result

Let $\mathrm{M} \subseteq \mathrm{V}$ be a transitive class. Let $j : \mathrm{V} \to \mathrm{M}$ be elementary with $\mathrm{cp}(j) = \kappa$. Thus U_j is a measure in **V** over κ. (Moreover, the club filter $\mathrm{Club}_\kappa \subseteq U_j$.)

Proof .:.

That U_j is an ultrafilter is easy enough to see as j is elementary: for $A \subseteq \kappa$ in V, $A \notin U_j$ iff $\kappa \notin j(A)$ iff $\kappa \in j(\kappa) \setminus j(A) = j(\kappa \setminus A)$ iff $\kappa \setminus A \in U_j$.

- U_j is easily seen to be non-principal. Otherwise, for some $\alpha < \kappa$, we'd have $\{\alpha\} \in U_j$ and hence $\kappa \in j(\{\alpha\})$. By elementarity of j, $j(\{\alpha\})$ has just one element: $j(\alpha) = \alpha \neq \kappa$, a contradiction.

- For κ-completeness, consider $\{A_\alpha : \alpha < \lambda\} \subseteq U_j$ in V for $\lambda < \kappa$. Since $\kappa \in j(A_\alpha)$ for each $\alpha < \lambda$, $\kappa \in \bigcap_{\alpha<\lambda} j(A_\alpha)$. Now since $\lambda < \kappa$, $j(\lambda) = \lambda$ and hence
$$\kappa \in \bigcap_{\alpha<\lambda} j(A_\alpha) = \bigcap_{\alpha<j(\lambda)} j(A_\alpha) = j\left(\bigcap_{\alpha<\lambda} A_\alpha\right).$$

- To show that U_j is normal, let $f : \kappa \to \kappa$ be such that $\forall^*\alpha\ (f(\alpha) < \alpha)$. This means
$$\kappa \in j(\{\alpha < \kappa : f(\alpha) < \alpha\}) = \{\alpha < j(\kappa) : j(f)(\alpha) < \alpha\}$$
and thus $j(f)(\kappa) < \kappa$. So there is some $\beta < \kappa$ with $j(f)(\kappa) = \beta = j(\beta)$ and hence $\forall^*\alpha\ (f(\alpha) = \beta)$.

U_j extends the club filter, since being a club is a first-order property. Hence $C \in \mathrm{Club}_\kappa$ has $j(C)$ containing a club of $j(\kappa)$. Since $C \subseteq \kappa$, $C = j(C) \cap \kappa$, which contains a club of κ. As $j(C)$ is closed, $\kappa = \sup C \in j(C)$. This means that $C \in U_j$ and thus $\mathrm{Club}_\kappa \subseteq U_j$. ⊣

This actually proves the earlier claim of Figure 1 D • 3: if U is a κ-complete ultrafilter over κ, then $j : \mathrm{V} \to \mathrm{cUlt}(\mathrm{V}, U)$ has $\mathrm{cp}(j) = \kappa$ and thus its derived ultrafilter U_j is a measure over κ.

2 B • 9. Theorem (Factoring)

Let $\mathrm{M} \subseteq \mathrm{V}$ be a transitive class. Let $j : \mathrm{V} \to \mathrm{M}$ be elementary with $\mathrm{cp}(j) = \kappa$.
Let U_j be the derived ultrafilter, and let $j_{\mathrm{ult}} : \mathrm{V} \to \mathrm{Ult}(\mathrm{V}, U_j)$ be the canonical ultrapower embedding.
Therefore there is a (unique) elementary $k : \mathrm{Ult}(\mathrm{V}, U_j) \to \mathrm{M}$ such that $j = k \circ j_{\mathrm{ult}}$ and $k([f]_{U_j}) = j(f)(\kappa)$.

Proof .:.

Write **Ult** for **Ult**(**V**, U_j). For each $[f]$, define $k([f]) = j(f)(\kappa)$. Note that this is independent on the choice of f, since if $\forall^*\alpha\ (f(\alpha) = g(\alpha))$, then by definition of U_j, $\kappa \in \{\alpha < j(\kappa) : j(f)(\alpha) = j(g)(\alpha)\}$ and so $k([f]) = j(f)(\kappa) = j(g)(\kappa) = k([g])$. Note also that $j = k \circ j_{\mathrm{ult}}$, since $k \circ j_{\mathrm{ult}}(x) = k([\mathrm{const}_x]) = j(\mathrm{const}_x)(\kappa) = \mathrm{const}_{j(x)}(\kappa) = j(x)$.

To see that k as defined is elementary, let $\varphi(x)$ be a formula and suppose **Ult** ⊨ "$\varphi([f])$" for some $[f] \in \mathrm{Ult}$. By Łoś's Theorem (2 • 2), this happens iff $\forall^*\alpha\ \varphi(f(\alpha))$. By definition of U_j, this means $\kappa \in j(\{\alpha < \kappa : \varphi(f(\alpha))\})$, i.e. **M** ⊨ "$\varphi(j(f)(\kappa))$". Rewritten, this says **M** ⊨ "$\varphi(k([f]))$". Thus k is elementary. ⊣

There are a number of corollaries to this. Firstly, we have a nice theorem of how we can break down ultrapowers.

2 B • 10. Lemma

Let U be a measure over κ. Therefore $\pi([\mathrm{id} \restriction \kappa]) = \kappa$ and $U_{j_U} = U$.

Proof .:.

- We know from Łoś's Theorem (2 • 2) that $\mathbf{Ult}(\mathrm{V}, U) \vDash$ "$[\mathrm{id} \restriction \kappa] > [\mathrm{const}_\alpha]$" for each $\alpha < \kappa$. Hence in the collapse (as κ is the critical point), $\mathbf{cUlt}(\mathrm{V}, U) \vDash$ "$\pi([\mathrm{id} \restriction \kappa]) > \alpha$" for each $\alpha < \kappa$ and thus $\pi([\mathrm{id} \restriction \kappa]) \geq \kappa$. To show that $\pi([\mathrm{id} \restriction \kappa]) \leq \kappa$, we appeal to normality.

  Let $\alpha < \pi([\mathrm{id} \restriction \kappa])$ be arbitrary. Therefore, by Łoś's Theorem (2 • 2), $\forall^* \beta\ (\mathrm{const}_\alpha(\beta) < \beta)$. So applying normality to const_α, we get that there must be some particular $\gamma < \kappa$ where $\forall^* \beta\ (\mathrm{const}_\alpha(\beta) = \gamma = \mathrm{const}_\gamma(\beta))$ so that $[\mathrm{const}_\alpha] = [\mathrm{const}_\gamma]$ and thus $\alpha = \gamma < \kappa$. Therefore $\pi([\mathrm{id} \restriction \kappa]) \subseteq \kappa$, and hence equal.

- It suffices to show that $A \in U$ iff $\kappa \in j_U(A)$. Rewritten, $A \in U$ says $\forall^* \alpha\ (\alpha \in \mathrm{const}_A(\alpha))$ which is equivalent to $\mathbf{Ult}(\mathrm{V}, U) \vDash$ "$[\mathrm{id} \restriction \kappa] \in [\mathrm{const}_A]$", meaning $\pi([\mathrm{id} \restriction \kappa]) = \kappa \in j_U(A)$. ⊣

This has the consequence of showing a trivial version of the factor lemma when M is the ultrapower by a measure. But this allows us to think about the ultrapower and the "M" in the same way.

2 B • 11. Corollary

Let U be a measure over κ. Therefore $\mathrm{cUlt}(\mathrm{V}, U) = \{j_U(f)(\kappa) : f \in {}^\kappa\mathrm{V}\}$.

Proof .:.

Let $j : \mathrm{V} \to \mathrm{cUlt}(\mathrm{V}, U)$ be the canonical ultrapower embedding. By Factoring (2 B • 9), there is a unique, elementary $k : \mathrm{Ult}(\mathrm{V}, U_j) \to \mathrm{cUlt}(\mathrm{V}, U)$ which obeys $k([f]_{U_j}) = j(f)(\kappa)$. Since $U_j = U$ by Lemma 2 B • 10, $\mathrm{cUlt}(\mathrm{V}, U_j) = \mathrm{cUlt}(\mathrm{V}, U)$ so that k must just be the collapsing isomorphism. Hence every element of $\mathrm{cUlt}(\mathrm{V}, U)$ can be represented in this way. ⊣

A slight generalization of this can be used for U that are merely κ-complete and not actually measures. The argument just replaces κ with $[\mathrm{id} \restriction \kappa]$. In fact, Corollary 2 B • 11 is equivalent to a κ-complete ultrafilter U being normal.

2 B • 12. Theorem

Let $\mathrm{M} \subseteq \mathrm{V}$ be a transitive class. Let $j : \mathrm{V} \to \mathrm{M}$ be elementary with $\mathrm{cp}(j) = \kappa$. Therefore, there is some ultrafilter U where $\mathrm{M} = \mathrm{Ult}(\mathrm{V}, U)$ with j as the canonical embedding iff there is some $s \in \mathrm{M}$ where

$$\mathrm{M} = \{j(f)(s) : f \in {}^\kappa\mathrm{V}\}.$$

Proof .:.

Suppose $\mathrm{M} = \mathrm{cUlt}(\mathrm{V}, U)$ with j as the canonical embedding. Set s (the *seed*) to be $\pi_U([\mathrm{id} \restriction \kappa])$ where $\pi_U : \mathrm{Ult}(\mathrm{V}, U) \to \mathrm{cUlt}(\mathrm{V}, U)$ is the collapsing map. We know already that $\mathrm{M} = \{\pi_U([f]_U) : f \in {}^\kappa\mathrm{V}\}$ so for $\pi_U([f]_U) \in \mathrm{M}$ arbitrary, it suffices to show that $\pi_U([f]_U) = j(f)(s)$. Note that $\forall^* \alpha\ (f(\alpha) = f(\alpha))$. We can think of $f(\alpha)$ as coming from the map $\alpha \mapsto f(\alpha)$ or coming from the map $\alpha \mapsto (\mathrm{const}_f(\alpha))(\mathrm{id}(\alpha))$. Using these two interpretations, by Łoś's Theorem (2 • 2), we have that $\mathbf{Ult}(\mathrm{V}, U) \vDash$ "$[f] = [\mathrm{const}_f]([\mathrm{id} \restriction \kappa])$", meaning that in the collapse, recallling that $j(x) = \pi_U([\mathrm{const}_x])$,

$$\pi_U([f]) = \pi_U\big([\mathrm{const}_f]([\mathrm{id} \restriction \kappa])\big) = \pi_U([\mathrm{const}_f])\big(\pi_U([\mathrm{id} \restriction \kappa])\big) = j(f)(s).$$

Now suppose there is some $s \in \mathrm{M}$ with $\mathrm{M} = \{j(f)(s) : f \in {}^\kappa\mathrm{V}\}$. Consider the ultrafilter $U = \{A \subseteq \kappa : s \in j(A)\}$. As in Result 2 B • 8, U can be easily shown to be an ultrafilter (by elementarity), and κ-complete (by elementarity and that $\mathrm{cp}(j) = \kappa$). As in Factoring (2 B • 9), consider the map $k : \mathrm{Ult}(\mathrm{V}, U) \to \mathrm{M}$ defined by $k([f]_U) = j(f)(s)$. To see that this is well defined, note that $[f]_U = [g]_U$ means $\forall^* \alpha\ (f(\alpha) = g(\alpha))$ implying $s \in j(\{\alpha < \kappa : f(\alpha) = g(\alpha)\})$ so that $s \in \{\alpha < j(\kappa) : j(f)(\alpha) = j(g)(\alpha)\}$ and thus $j(f)(s) = j(g)(s)$.

This k is elementary by the same reasoning as in Factoring (2 B • 9):

$$\begin{aligned}\mathbf{Ult}(\mathrm{V}, U) \vDash \text{“}\varphi([f]_U\text{”} \quad &\text{iff} \quad \forall^* \alpha\ \varphi(f(\alpha)) \\ &\text{iff} \quad s \in j(\{\alpha < \kappa : \varphi(f(\alpha))\} = \{\alpha < j(\kappa) : \mathbf{M} \vDash \text{“}\varphi(j(f)(\alpha))\text{”}\}\end{aligned}$$

$$\text{iff} \quad \mathbf{M} \vDash \text{"}\varphi(j(f)(s))\text{"}.$$

This implies k is injective and an embedding. So k is actually an ismorphism since it's clearly surjective: $\mathrm{M} = \{j(f)(s) : f \in {}^{\kappa}\mathrm{V}\}$ allow us to merely consider $k([f]_U) = j(f)(s)$ for any $f : \kappa \to \mathrm{V}$. Hence $\mathrm{cUlt}(\mathrm{V}, U) = \mathrm{M}$ and by uniqueness, k is just the collapsing isomorphism, meaning j is the canonical ultrapower embedding. ⊣

As stated before, the $s = \pi_U([\mathrm{id} \restriction \kappa])$ being κ is equivalent to normality for κ-complete ultrafilters over $\kappa > \aleph_0$.

2 B • 13. Result

Let U be a non-principal, κ-comlete ultrafilter over $\kappa > \aleph_1$. Let $\pi_U : \mathrm{Ult}(\mathrm{V}, U) \to \mathrm{cUlt}(\mathrm{V}, U)$ be the collapsing isomorphism. Therefore U is normal iff $\pi_U([\mathrm{id} \restriction \kappa]) = \kappa$.

Proof .:.

Let $j : \mathrm{V} \to \mathrm{Ult}(\mathrm{V}, U)$ be the canonical embedding which then has $\mathrm{cp}(j) = \kappa$. Since U is κ-complete, it's clearly unbounded and thus $\forall^*\alpha\ (\alpha > \beta)$ for each β. Restated, this says $\mathbf{Ult}(\mathbf{V}, U) \vDash$ "$[\mathrm{id} \restriction \kappa] > [\mathrm{const}_\beta]$" for each $\beta < \kappa$. So after collapsing,

$$\pi_U([\mathrm{id} \restriction \kappa]) \supseteq \{\pi_U([\mathrm{const}_\beta]) : \beta < \kappa\} = \{j(\beta) : \beta < \kappa\} = \kappa.$$

Note that a function f being regressive on a set in U is equivalent to $\forall^*\alpha\ (f(\alpha) < \alpha)$, meaning $\mathbf{Ult}(\mathbf{V}, U) \vDash$ "$[f] < [\mathrm{id} \restriction \kappa]$". Thus $\pi_U([\mathrm{id} \restriction \kappa]) = \{\pi_U([f]) : f \text{ is regressive on a set in } U\}$.

So suppose U is normal. Note that every regressive function $f : \kappa \to \kappa$ has some $\beta < \kappa$ where $[f] = [\mathrm{const}_\beta]$ and thus $\pi_U([f]) = \pi_U([\mathrm{const}_\beta]) = j(\beta) = \beta$. Therefore $\pi_U([\mathrm{id} \restriction \kappa]) \subseteq \kappa$, and so we have equality.

Now suppose $\pi_U([\mathrm{id} \restriction \kappa]) = \kappa$. Thus every element of $\pi_U([\mathrm{id} \restriction \kappa])$ is an ordinal less than κ, meaning that every regressive function $f : \kappa \to \kappa$ has $\pi_U([f]) = \beta = \pi_U([\mathrm{const}_\beta])$ for some $\beta < \kappa$. So $[f] = [\mathrm{const}_\beta]$ and thus $\forall^*\alpha\ (f(\alpha) = \beta)$, meaning U is normal. ⊣

§ 2 C. Properties of ultrapowers

So we have ultrapowers, and we know what they look like thanks to Theorem 2 B • 12. What are some of their properties, however? The main goal of this subsection is now to look at what happens with the critical point of the ultrapower embedding: Where is it sent? How close can the ultrapower be to V? What are the combinatorial effects of taking an ultrapower? A complete answer to these questions won't be given here (if there even is such an answer). Instead, we will consider the following results.

2 C • 1. Result

Let U be a measure on κ. Let $j : \mathrm{V} \to \mathrm{cUlt}(\mathrm{V}, U)$ be the canonical embedding. Therefore,

1. $\mathbf{cUlt}(\mathbf{V}, U)$ is closed under κ-length sequences, meaning ${}^{\kappa}\,\mathrm{cUlt}(\mathrm{V}, U) \subseteq \mathrm{cUlt}(\mathrm{V}, U)$.
2. $\mathbf{cUlt}(\mathbf{V}, U)$ and $\mathbf{V}$ *agree up to* $\kappa + 1$, meaning $\mathrm{V}_{\kappa+1} \subseteq \mathrm{cUlt}(\mathrm{V}, U)$ but $\mathrm{V}_{\kappa+2} \nsubseteq \mathrm{cUlt}(\mathrm{V}, U)$;
3. $\wp(\kappa) = \wp(\kappa) \cap \mathrm{cUlt}(\mathrm{V}, U)$;
4. $U \notin \mathrm{cUlt}(\mathrm{V}, U)$; and
5. $j(\kappa)$ is not a cardinal of $\mathbf{V}$: $\kappa < 2^\kappa \leq (2^\kappa)^{\mathbf{cUlt}(\mathbf{V}, U)} < j(\kappa) < (2^\kappa)^+$.

To prove these from the ground up, we need some results about *measurable* cardinals which we have not introduced yet. Instead, just assume the following lemma.

2 C • 2. Lemma

Let κ have a measure U over it. Therefore, κ is strongly inaccessible.

Proof .:.

κ is regular by κ-completeness of its measure. κ is uncountable by elementarity of j. To show that κ is a strong limit, suppose not, and let $\lambda < \kappa$ have $2^\lambda \geq \kappa$.

So consider family $\Lambda \subseteq \wp(\lambda)$ be of size $|\Lambda| = \kappa$. Take a corresponding ultrafilter $W \subseteq \wp(\Lambda)$ with U and the bijection with κ. This W, however, will not be κ-complete, contradicting that U is. To see this, for each $\alpha < \lambda$, consider

$$X_\alpha = \begin{cases} \{x \in \Lambda : \alpha \in x\} & \text{if this is in } U \\ \{x \in \Lambda : \alpha \notin x\} & \text{otherwise.} \end{cases}$$

By construction, $X_\alpha \in U$ for all $\alpha < \lambda$. The intersection of all these, by κ-completeness of U, is in W. But $\bigcap_{\alpha<\lambda} X_\alpha$ is a single subset of λ, contradicting nonprincipality. ⊣

Proof of Result 2 C • 1 .:.

To save space, write $\mathrm{M} = \mathrm{cUlt}(\mathrm{V}, U)$ and $\mathrm{Ult} = \mathrm{Ult}(\mathrm{V}, U)$ with $\pi : \mathrm{Ult} \to \mathrm{M}$ the collapsing isomorphism.

1. Let $\vec{x} = \langle x_\alpha \in \mathrm{M} : \alpha < \kappa\rangle$ be a κ-length sequence (in V). Represent $x_\alpha = \pi([f_\alpha])$ for $f_\alpha : \kappa \to \mathrm{V}$. Consider the sequence (also in V) $\vec{f} = \langle f_\alpha : \alpha < \kappa\rangle$. Now we consider $j(\vec{f})$. By elementarity, $j(\vec{f})$ is a sequence of length $j(\kappa)$. Moreover, for every $\beta < \kappa$, $\forall^*\alpha\ (\vec{f}(\beta)(\alpha) = f_\beta)$ so by Łoś's Theorem (2 • 2),
$$\mathsf{Ult} \vDash \text{“}[\mathrm{const}_{\vec{f}}]([\mathrm{const}_\beta]) = [f_\beta]\text{”} \quad \text{iff} \quad j(\vec{f})(\beta) = \pi([\mathrm{const}_{\vec{f}}])(\pi([\mathrm{const}_\beta])) = \pi([f_\beta]) = x_\beta.$$
Thus $j(\vec{f}) \upharpoonright \kappa = \vec{x}$. As $\kappa, j(\vec{f}) \in \mathrm{M}$, it then follows that $\vec{x} \in \mathrm{M}$.
2. This follows by Result 2 A • 7 and (4) below.
3. This follows by Result 2 A • 7.
4. Every $\alpha < j(\kappa)$ has a represenation $[f]$ in Ult which then obeys $\forall^*\beta\ (f(\beta) < \mathrm{const}_\kappa(\beta))$, meaning we can assume without loss of generality that $f : \kappa \to \kappa$. So let $F : {}^\kappa\kappa \to j(\kappa)$ be the surjective map $f \mapsto \pi([f])$. Suppose $U \in \mathrm{M}$ so that for any $f \in ({}^\kappa\kappa)^{\mathsf{M}} = {}^\kappa\kappa$, we can form $[f]$ and thus the map F within M. Hence $\mathsf{M} \vDash$ “$\kappa < j(\kappa) \leq \kappa^\kappa = 2^\kappa$”, contradicting Lemma 2 C • 2 since by elementarity, $j(\kappa)$ is also strongly inaccessible.
5. By (3), $2^\kappa \leq (2^\kappa)^{\mathsf{M}}$. We of course know $\kappa < 2^\kappa$ by Cantor's theorem. We have $j(\kappa) > (2^\kappa)^{\mathsf{M}}$ because κ is a strong limit in V so that $j(\kappa)$ is a strong limit in M. We have $j(\kappa) < (2^\kappa)^+$ since the argument given in (4) tells us that there's a surjection from $\kappa^\kappa = 2^\kappa$ to $j(\kappa)$ in V. ⊣

Now all of this has been a kind of coded way of talking about measurable cardinals by way of their measures.

§ 2 D. Measurable cardinals

Although we have mentioned measurable cardinals before, they should be given a formal introduction. Measurable cardinals are important for their two equivalent characaterizations: having a measure, and being the critical point of an elementary embedding. Measurable cardinals will be quite large, and their importance is partly for the ultrapowers mentioned in the rest of this section, but also in motivating a canonical inner model $\mathrm{L}[U]$ to be introduced later.

2 D • 1. Definition

A cardinal $\kappa > \aleph_0$ is *measurable* iff there is a non-principal, κ-complete ultrafilter over κ.

Note that by the results above, there are several different characterizations of this.

2 D • 2. Result

Let $\kappa \geq \aleph_0$ be a cardinal. Therefore, the following are equivalent:

1. κ is measurable, i.e. $\kappa > \aleph_0$ has a non-principal, κ-complete ultrafilter over it.
2. κ has a measure over it.
3. κ is the critical point of an elementary $j : \mathrm{V} \to \mathrm{M}$, where M is a transitive class of V.

Proof .:.

Clearly (2) implies (1) with the only thing to check being that κ is uncountable. But normality implies this: suppose $\kappa = \aleph_0$ with U a measure over $\aleph_0$. Consider $f : \omega \to \omega$ defined by $f(0) = 0$ and $f(n) = n - 1$ for $n > 0$. As $f(n) \geq n$ iff $n = 0$, it follows by uniformity that $\forall^*n\ (f(n) < n)$. So by normality, there is some $m < \omega$ where $\forall^*n\ (f(n) = m)$. But $f^{-1}(m) \subseteq \{m, m+1\} \notin U$ by uniformity. Therefore $\kappa \neq \aleph_0$.

So suppose (1) holds: κ is measurable as witnessed by U. Therefore $\mathsf{Ult}(\mathsf{V}, U)$ is well-founded since $\kappa \geq \aleph_1$: κ-completeness implies $\aleph_1$-completeness. Hence the canonical embedding $j : \mathrm{V} \to \mathrm{cUlt}(\mathrm{V}, U)$ has $\mathrm{cp}(j) = \kappa$ by Theorem 2 B • 5 showing (3).

If (3) holds, the derived ultrafilter U_j is a measure on κ by Result 2 B • 8, yielding (2). ⊣

This equivalence of measurability and being a critical point is an important one in the sense that each characterization has various corollaries, and when combined they give a clearer picture of measurable cardinals. Consider the following consequences, for example, showing just how large measurables need to be. We already know that just one inaccessible goes beyond what ZFC can prove. In fact, the consistency of just any number of inaccessibles can't be proven relative to the consistency of any fewer number of them. Now consider how strong the existence of measurables is.

2 D • 3. Corollary

Let κ be measurable. A cardinal λ is *mahlo* iff $\{\theta < \lambda : \theta = |\theta| \text{ is inaccessible}\}$ is a stationary subset of λ. Therefore

1. κ is strongly inaccessible by Lemma 2 C • 2;
2. κ is the κth (strongly) inaccessible cardinal;
3. κ is the κth mahlo cardinal;
4. κ has a measure by Result 2 D • 2; and
5. κ has a measure that extends the club filter Club_κ by Result 2 B • 8.

Proof .:.

Let U be a measure on κ, and let $j : \mathrm{V} \to \mathrm{M}$ be elementary with $\mathrm{M} \subseteq \mathrm{V}$ a transitive class.

2. Note that a cardinal λ being strongly inaccessible is downward absolute. So if κ is strongly inaccessible in V, then it is in M, meaning that M thinks that $j(\kappa)$ has an inaccessible below it: κ. So for each $\alpha < \kappa$, $\mathsf{M} \models$ "$\exists x(x \text{ is inaccessible and } \alpha < x < j(\kappa))$". So by elementarity, for each $\alpha < \kappa$, $\mathsf{V} \models$ "$\exists x(x \text{ is inaccessible and } \alpha < x < \kappa)$". So the set of inaccessible cardinals below κ is unbounded in κ. As κ is regular, κ is the κth inaccessible.

3. Firstly, to see that κ is mahlo, take $j : \mathrm{V} \to \mathrm{M} \subseteq \mathrm{V}$ elementary with $\mathrm{cp}(j) = \kappa$. For any club $C \subseteq \kappa$, $j(C) \subseteq j(\kappa)$ is also club, and since $C = j(C) \cap \kappa$, it follows that $\kappa \in j(C)$ and thus $\mathsf{M} \models$ "$j(C)$ has an inaccessible member". By elementarity and absoluteness, C has an inaccessible member so that the set of inaccessibles below κ is stationary and κ is mahlo.

   κ is still mahlo in M, since $\wp(\kappa) = \wp(\kappa) \cap \mathrm{M}$ meaning that M contains every club of κ as well as the stationary set of inaccessibles above. Hence being a stationary subset of κ is absolute between M and V. Thus the above $j(C)$ contains a mahlo cardinal in M. By elementarity, C contains a mahlo cardinal in V, and thus the set of mahlos below κ is stationary, and thus κ is the κth mahlo cardinal. ⊣

One might be tempted to apply the reasoning of Corollary 2 D • 3 to the property of being measurable, which would seem to indicate that any measurable cardinal κ would need to be the κth measurable cardinal, or it seems at least there can't be a least measurable. To simplify the issue, let κ be the least measurable cardinal, and let $j : \mathrm{V} \to \mathrm{M} \subseteq \mathrm{V}$ be elementary. It would seem that $j(\kappa)$ has a measurable below it, and thus κ does too, contradicting that κ is the least measurable. The issue is that κ might not be measurable in M, because we've thinned out the universe to $\mathrm{M} \subseteq \mathrm{V}$ such that it no longer contains a measure, as seen in Result 2 C • 1.

Moreover, M, being the collapsed ultrapower, has further properties that present limitations on the kinds of embeddings that can be realized by ultrapowers. The properties of being inaccessible, mahlo, and so forth could be used with the above reasoning, since they deal only at the level of V_κ and $\mathrm{V}_{\kappa+1}$, but issues creep in if we try going beyond this, like the statement of being measurable. This is again a result of the agreement between the ultrapower and V as seen in Result 2 C • 1.

Now despite the fact that the reasoning of Corollary 2 D • 3 breaks down when we try to apply them to the property of, for example, being measurable, the reasoning *does* apply when $\mathsf{V} = \mathsf{L}$. This is because of L being the smallest inner model: $\mathsf{cUlt}(\mathsf{V}, U) = \mathsf{V} = \mathsf{L}$ which forces M to still recognize κ as measurable.

2 D • 4. Theorem (L Has No Measurable Cardinals)

Let κ be measurable. Therefore $V \neq L$.

Proof .:.

Without loss of generality, let κ be the least measurable cardinal, and assume $V = L$. By Result 2 D • 2, there is an elementary embedding $j : L \to M$ with a transitive class $M \subseteq L$. By elementarity,

$$M \vDash \text{ZFC} + \text{“}V = L\text{”} + \text{“}j(\kappa) \text{ is the least measurable”}.$$

Condensation implies $M = L$, and thus the two agree on κ: $M \vDash$ “κ is the least measurable”, which contradicts that κ is the critical point of j: $\kappa \neq j(\kappa)$. ⊣

This is a relatively easy proof due to condensation, but there is a more complicated proof due to a more general result.

2 D • 5. Theorem (Kunen's Inconsistency Theorem)

Let $j : V \to V$ be elementary (as a class of V). Therefore $j = \text{id}$.

Proof .:.

Assume $j \neq \text{id}$. By Result 2 A • 3, there is some critical point $\kappa = \text{cp}(j)$. By repeatedly applying j, we get the sequence $\langle j^n(\kappa) : n \in \omega \rangle$. Let $\theta = \sup_{n\in\omega} \kappa_n$. By applying j to the sequence, by elementarity, we get that $j(\langle j^n(\kappa) : n \in \omega\rangle) = \langle j^{n+1}(\kappa) : n \in \omega\rangle$, and that $j(\theta) = \sup_{n\in\omega} j^{n+1}(\kappa) = \theta$. As a fixed point of j, this is good. Unfortunately, θ isn't regular. So instead consider the next cardinal, which by elementarity is also fixed: $j(\theta^+) = j(\theta)^+ = \theta^+$.

As θ^+ is regular, consider the stationary subset of ordinals with cofinality ω: $S = S^{\theta^+}_\omega = \{\alpha < \theta^+ : \text{cof}(\alpha) = \omega\}$. This can be closed under fix-points of j, since $j(\alpha) = \sup(j\text{"}\alpha)$ for $\text{cof}(\alpha) < \text{cp}(j)$. The resulting set is also unbounded since $j\text{"}\theta^+ = \theta^+$. What this means is that

$$C = \{\alpha < \theta^+ : \text{cof}(\alpha) = \omega \wedge j(\alpha) = \alpha\}$$

is almost a club. In particular, C^+—the closure of C under all sequences—is club in θ^+ with no new elements of cofinality ω. As a result, any stationary subset of S will intersect C.

But any stationary set of θ^+ may be partitioned into θ^+ stationary subsets. In particular, we can consider subsets $S_\alpha \subseteq S$ for $\alpha < \kappa$—just take the union of S_0 with the guaranteed S_α for $\kappa \leq \alpha < \theta^+$ and make this the new S_0—where all the S_αs are stationary and pairwise disjoint. Applying j, we get another sequence, this time of length $j(\kappa)$, of pairwise disjoint, stationary subsets of θ^+: $\langle Z_\alpha : \alpha < j(\kappa)\rangle = j(\langle S_\alpha : \alpha < \kappa\rangle)$. By the above ideas on C^+, $Z_\kappa \cap C^+ \neq \emptyset$. So there is some element $\zeta \in Z_\kappa \cap C^+$. As the S_αs partition S, there is also some $\alpha < \kappa$ with $\zeta \in S_\alpha \cap C^+$. But then $\zeta = j(\zeta) \in j(S_\alpha) = Z_{j(\alpha)}$. As $\alpha < \kappa = \text{cp}(j)$, $j(\alpha) = \alpha$, yielding that $Z_\alpha \cap Z_\kappa \neq \emptyset$, a contradiction. ⊣

It's a good exercise to see where this proof breaks down for elementary $j : V \to M$ for $M \subsetneq V$ a transitive class. Note that this doesn't say that there can be no (non-trivial) $j : V \to V$ for $V = L$, for example, just that no j can exist as a class of V in this case.