Annals of Mathematical Logic 20 (1981) 77–108
North-Holland Publishing Company

INFINITE GAMES AND REDUCED PRODUCTS*

Wilfrid HODGES
Bedford College, London, England

Saharon SHELAH
Hebrew University, Jerusalem, Israel

Received 26 November 1979

We study reduced products over κ-complete filters. For such products the structures can carry relations and functions of any arity $<\kappa$, though nearly all our results allow the relations and functions to be finitary. Most of our theorems come in two forms: (H) for κ-complete filters and Horn logic, where κ is regular; and (S) for κ-complete ultrafilters and $L_{\kappa\kappa}$ with game quantifiers, where κ is strongly compact. Generally the (S) version is more elegant, but the (H) version applies in more situations.

Our main result appears as Theorem 4 in Section 2. In the (H) version, it says that under certain set-theoretic assumptions, two structures have isomorphic reduced powers over κ-complete filters if and only if their κ-Horn theories are consistent with each other. (Unpublished work of Laver shows that the set-theoretic assumption is consistent if there is a proper class of measurable cardinals.) We give several variants of this result. Both the (H) and the (S) versions are known to be true absolutely when $\kappa=\omega$; the (S) version is the Keisler–Shelah theorem on isomorphism of ultrapowers [16] and the (H) version appears as Exercise 6.2.6 in Chang and Keisler [2]. For uncountable κ the theorem is new.

In Section 3 we prove the same theorems for *limit* reduced products where both filters are required to be κ-complete. This time no special set-theoretic assumptions are needed. We also characterise limit reduced powers over κ-complete filters as the most general operation which commutes with taking reducts and gives elementary extensions for certain languages; the (S) version generalises a result of Keisler for $\kappa=\omega$.

Section 4 uses the results of Section 3 to deduce some infinitary model theory. We give interpolation and preservation theorems for Horn logic; these were originally proved [6] by using a more conventional proof-theoretic argument. When $\kappa>\omega$, the formulae which are preserved by reduced products over κ-complete filters are not necessarily Horn, even up to logical equivalence; we give

* The main work for this paper was done when the second author visited the first in London in August 1978 under Science Research Council Senior Visiting Fellowship GR/A/52997.

examples, some of which also answer a related question of Kueker. We show that the amalgamation property fails badly for logics with non-homogeneous infinitary quantifiers. Finally we describe an incompact logic which satisfies the Craig interpolation theorem and has the Feferman–Vaught property (but it has poor substitution properties).

In Sections 2 and 3 we handle the infinitary quantifiers by making two players play a large number of infinitary games simultaneously on separate boards. The idea in Section 3 is to make sure that the right player wins by making the players play elements of a limit product constructed from a family of independent functions; this idea was due to Shelah. The corresponding devices for reduced products in Section 2 are an amalgam of several people's ideas: we thank Fred Galvin and Richard Laver for letting us have their contributions.

1. Preliminaries

Throughout, κ is a regular cardinal. If A is a structure, then we always assume that the relations R_A and functions F_A of A have arities $<\kappa$. By a *κ-filter* we mean a filter which is proper and κ-complete (i.e. closed under infs of $<\kappa$ elements); a *κ-ultrafilter* is a κ-complete ultrafilter. The letters κ, λ, μ, ν are reserved for cardinals, while α, β, i, j etc. are ordinals.

1.1. Languages and games

A *pre-game* is an incompletely specified game, where we know what counts as playing it, but not necessarily what counts as winning it. If players play a pre-game, the result is a sequence of moves called a *play*.

By a *quantifier of length* α we mean a map $Q:\alpha\twoheadrightarrow\{\forall,\exists\}$. A quantifier Q and a structure A define a pre-game $G(Q,A)$ as follows: players $\forall$ and $\exists$ pick elements a_β of dom A (=domain of A) for each $\beta<\alpha=\text{length}(Q)$. Player $Q(\beta)$ picks a_β, and he is allowed to know what a_γ $(\gamma<\beta)$ have been chosen. Thus the play is a sequence $\bar{a}$ of length α from dom A. We sometimes generalise this definition a little by allowing the domain of Q to consist of any increasing sequence of ordinals of length α. The pre-game $G(Q,A)$ is said to have *length* α.

Let L be a language. We shall define a language PL. The formulae of PL are the pairs $\langle Q,\varphi\rangle$, written $Q\varphi$, where Q is a quantifier and φ is a quantifier-free formula of L. If A is a structure of the similarity type of L, then $G(Q\varphi,A)$ is the following game: players $\forall$ and $\exists$ play $G(Q,A)$, and if the resulting play is $\bar{a}$, then $\exists$ wins iff $A\models\varphi[\bar{a}]$. $A\models Q\varphi$ means that player $\exists$ has a winning strategy in $G(Q\varphi,A)$.

For example, if L is $L_{\kappa\kappa}$, then PL essentially consists of the prenex formulae of $L_{\kappa\kappa G}$. In this case PL is an extension of first-order logic, but note that it is not closed under negation (even up to logical equivalence) when $\kappa>\omega$.

Another example we shall often use is *Horn logic*. A quantifier-free *Horn* formula of $L_{\kappa\kappa}$ is a conjunction of $<\kappa$ formulae of form

$$\bigwedge \Phi \rightarrow \bigwedge \Psi$$

where Φ, Ψ are sets of $<\kappa$ atomic formulae, and Ψ is non-empty; $\perp$ (falsehood) counts as atomic. $H_{\kappa\kappa}$ consists of the quantifier-free Horn formulae of $L_{\kappa\kappa}$. When we speak of *Horn logic* we shall mean $PH_{\kappa\kappa}$.

1.2. Limit reduced products

The following definitions are taken from Keisler [10, 11] or Chang and Keisler [2], with some slight changes.

Let I be a non-empty set, and for each $i \in I$ let A_i be a structure, all of the same similarity type. Then we may form the *product* structure $\prod_{i \in I} A_i$ or more briefly $\prod_I A_i$.

Reduced products are homomorphic images of products, got by factoring out proper filters D on I. Precisely, let D be a proper filter on I (i.e. a proper filter on the boolean algebra $\mathcal{P}I$). For $f, g \in \prod_I A_i$, put $f \sim_D g$ iff $\{i \in I: f(i) = g(i)\} \in D$. Then $\sim_D$ is an equivalence relation on $\prod_I A_i$. If moreover D is κ-complete, then $\sim_D$ is a congruence with respect to the functions and relations of $\prod_I A_i$, and by factoring out $\sim_D$ we get a homomorphic image $\prod_D A_i$ of $\prod_I A_i$. $\prod_D A_i$ is called a *κ-reduced product* of the A_i.

Limit reduced products are substructures of reduced products, got from filters on the set of partitions of the index set I. Let $\text{Part}(I)$ be the set of partitions of I, and for $\pi, \rho \in \text{Part}(I)$ write $\pi \leqslant \rho$ when every partition class of π is included in a partition class of ρ. Then $(\text{Part}(I), \leqslant)$ is a complete lattice. For each element f of $\prod_I A_i$, write $[f]$ for the partition π of I such that i, j lie in the same class of π iff $A_i = A_j$ and $f(i) = f(j)$. Let F be a filter on $\text{Part}(I)$ and $f \in \prod_D A_i$; then we say f is an *F-element* iff $f = g/{\sim_D}$ for some $g \in \prod_I A_i$ with $[g] \in F$. Provided F is κ-complete, the F-elements form a substructure of $\prod_D A_i$, which we write as $\prod_D A_i \mid F$. This structure is called a *limit κ-reduced product* of the A_i. Note that $\prod_D A_i \mid F$ can be formed by first forming a limit product $\prod_I A_i \mid F$ and then factoring out by $\sim_D$.

If $A_i = A$ for each $i \in I$, we say 'power' instead of 'product', A^I instead of $\prod_I A_i$, and A^I_D for $\prod_D A_i$. Constant functions in A^I are F-elements for every filter F, so there is a natural embedding of A into $A^I_D \mid F$ which takes each element of dom A to the corresponding constant function.

When D is an ultrafilter, we say 'ultra-' for 'reduced'.

Note that every κ-reduced product is also a limit κ-reduced product, by taking F to be the whole of $\text{Part}(I)$.

1.3. Preservation and compactness

We say that a sentence φ is *preserved in* limit κ-reduced products iff every limit κ-reduced product of models of φ is again a model of φ; and likewise with other operations on models.

Lemma 1 (Preservation theorem—easy direction). *Every sentence of* $PH_{\kappa\kappa}$ ($PL_{\kappa\kappa}$) *is preserved in limit* κ*-reduced products (limit* κ*-ultraproducts).*

Proof. We prove only the $PH_{\kappa\kappa}$ case. Let $\prod_D A_i \mid F$ be a limit κ-reduced product of models of the sentence $Q\varphi$ of $PH_{\kappa\kappa}$. Then for each $i \in I$, player $\exists$ has a winning strategy σ_i in the game $G(Q\varphi, A_i)$; we may assume $\sigma_i = \sigma_j$ whenever $A_i = A_j$. For $G(Q, \prod_I A_i)$ he has the following strategy σ: play σ_i at the ith coordinate, for each $i \in I$. We claim that σ is also a strategy for player $\exists$ in $G(Q, \prod_I A_i \mid F)$. For suppose $\bar{a}_\gamma$ ($\gamma < \beta$) have been chosen from $\prod_I A_i \mid F$, and $Q(\beta) = \exists$. Writing $\bar{a}_\gamma(i)$ for the ith element of $\bar{a}_\gamma$, we want $\exists$ to choose $\bar{a}_\beta$ so that $\bar{a}_\beta(i) = \sigma_i\langle \bar{a}_\gamma(i)\colon \gamma < \beta\rangle$. But $[\bar{a}_\beta] \leqslant \bigwedge_{\gamma<\beta} [\bar{a}_\gamma]$, $\rho < \kappa$ and F is κ-complete; so he can choose this way.

Now let player $\exists$ play $G(Q\varphi, \prod_D A_i \mid F)$ by choosing representatives of equivalence classes and using σ in $G(Q\varphi, \prod_I A_i \mid F)$. Suppose the play on $\prod_I A_i \mid F$ is $\langle \bar{a}_\beta\colon \beta < \alpha\rangle$. Then for each i, $A_i \models \varphi[\bar{a}_\beta(i)]_{\beta<\alpha}$ by choice of σ_i. Then since φ is quantifier-free Horn and D is a κ-complete filter it follows easily that $\prod_D A_i \models \varphi[\bar{a}_\beta/\sim_D]_{\beta<\alpha}$. But since φ is quantifier-free, this implies that $\prod_D A_i \mid F \models \varphi[\bar{a}_\beta/\sim_D]_{\beta<\alpha}$. Hence player $\exists$ wins $G(Q\varphi, \prod_D A_i \mid F)$, and so $\prod_D A_i \mid F$ is a model of $Q\varphi$ as required.

The converse of Theorem 1 is to characterise those sentences which are preserved in all limit κ-reduced products. We give some positive and some negative results on this in Section 4.1 below.

Lemma 2 (Compactness). *If T is a set of sentences of* $PH_{\kappa\kappa}$ ($PL_{\kappa\kappa}$, *where κ is strongly compact), such that every subset of T of cardinality $<\kappa$ has a model, then T has a model.*

Proof. By Lemma 1 and the usual ultraproduct proof of the compactness theorem [2, Corollary 4.1.11], it suffices to find a κ-filter D on a set I and a family S of card(T) elements of D, such that every element of I is in fewer than κ elements of S. Take I to be the set of all sets of fewer than κ elements of T, and let S be the family of sets of form $\{t \in I\colon s \subseteq t\}$ for $s \in I$. Let D be the filter generated by S; then D is κ-complete since S is closed under intersections of fewer than κ elements (by the regularity of κ), and D is obviously proper.

Lemma 2 has another proof: set up a complete cut-free proof calculus for $PH_{\kappa\kappa}$ ($PL_{\kappa\kappa}$), and show that any proof of a contradiction from sentences in $PH_{\kappa\kappa}$ ($PL_{\kappa\kappa}$) has fewer than κ premises. (For the Horn case one can use Takeuti's completeness theorem for negative sequents with heterogeneous quantifiers; see [17, Proposition 24.19].) Hence the notion of a *consistent theory* of $PH_{\kappa\kappa}$ or $PL_{\kappa\kappa}$ is quite unambiguous.

2. Reduced products

In this section we give necessary and sufficient conditions for two structures (or families of structures) to have isomorphic κ-reduced products or κ-ultraproducts.

We can vary the question, for example by requiring the two filters to be the same, or asking for just one of them to be an ultrafilter. Some set-theoretic assumptions seem to be needed.

2.1. *The quantifier-free part*

Throughout this paragraph the setting is as follows. For each $i \in I$ a structure A_i is given, and for each $j \in J$ a structure B_j. μ is a cardinal, and sequences $\langle \bar{a}_\gamma : \gamma < \mu \rangle$, $\langle \bar{b}_\gamma : \gamma < \mu \rangle$ of elements of $\prod_I A_i$, $\prod_J B_j$ respectively are given. For each $i \in I$ the sequence

$$\langle \bar{a}_\gamma(i) : \gamma < \mu \rangle$$

is called the ith *thread* of $\langle \bar{a}_\gamma : \gamma < \mu \rangle$, and written $\bar{a}^i$; similarly $\bar{b}^j$. For the moment we assume that the variables of $\mathrm{PH}_{\kappa\kappa}$ ($\mathrm{PL}_{\kappa\kappa}$) are v_γ ($\gamma < \mu$).

Lemma 3a. *There is a family Δ of ordered pairs $\langle \theta, \eta \rangle$ of quantifier-free formulae of* $\mathrm{PH}_{\kappa\kappa}$, *not depending on $\langle \bar{a}_\gamma : \gamma < \mu \rangle$, $\langle \bar{b}_\gamma : \gamma < \mu \rangle$, such that the following are equivalent:*

(i) *for every pair $\langle \theta, \eta \rangle$ in Δ, either there is $i \in I$ such that $A_i \models \neg \theta[\bar{a}^i]$ or there is $j \in J$ such that $B_j \models \neg \eta[\bar{b}^j]$;*

(ii) *there are κ-filters D, E on I, J respectively, such that for every atomic formula ψ, $\prod_D A_i \models \psi[\bar{a}_\gamma]_{\gamma<\mu}$ iff $\prod_E B_j \models \psi[\bar{b}_\gamma]_{\gamma<\mu}$.*

Proof. To make later variants easier, we shall be a little more formal than we need. Consider the following notion of proof. By a *proof-scheme* we mean a tree P such that (1) there is a single bottom node, (2) each node has $<\kappa$ nodes immediately above it, (3) every node has finite height, (4) each node has one atomic formula attached to it, (5) each node is labelled A or B. The formula attached to the bottom node is called the *conclusion*. A node N labelled A will be said to be *correct at i* iff, writing ψ for the formula attached at N and Φ for the set of formulae attached immediately above N,

$$A_i \models (\bigwedge \Phi \rightarrow \psi)[\bar{a}^i].$$

Similarly for nodes labelled B. N is *everywhere correct* iff N is correct at every $i \in I$ or every $j \in J$ (according as N is labelled A or B). P is *valid* iff all its nodes are everywhere correct. We write $\vdash^* \psi$ iff there is a valid proof-scheme with conclusion ψ.

For each formula φ write

$$A(\varphi) = \{ i \in I : A_i \models \varphi[\bar{a}^i] \},$$

and similarly $B(\varphi)$. Define D to be the filter on I generated by all intersections of fewer than κ sets $A(\psi)$ such that $\vdash^* \psi$, and E the filter on J generated by intersections of fewer than κ sets $B(\psi)$ such that $\vdash^* \psi$. Clearly both D and E are κ-complete.

Claim. *For each atomic formula ψ, $A(\psi) \in D$ iff $\vdash^* \psi$ iff $B(\psi) \in E$.*

By symmetry we only need prove the first equivalence. Right to left is by definition of D. Conversely, suppose $A(\psi)\in D$; then there is a set Φ of fewer than κ atomic formulae φ such that $\vdash^* \varphi$, for which

$$\bigcap_{\varphi\in\Phi} A(\varphi)\subseteq A(\psi). \tag{I}$$

For each $\varphi\in\Phi$ there is a valid proof-scheme P_φ with conclusion φ. Let P be the proof-scheme with conclusion ψ labelled A, such that when the bottom node of P is removed, the segments which remain are precisely the P_φ $(\varphi\in\Phi)$. Then P is valid, because the P_φ are valid and (I) says that the bottom node of P is everywhere correct. Hence $\vdash^* \psi$, and the claim is proved.

The claim implies that for every atomic formula ψ, $\prod_D A_i\models\psi[\bar a_\gamma]_{\gamma<\mu}$ iff $\prod_E B_j\models\psi[\bar b_\gamma]_{\gamma<\mu}$.

Now for each proof-scheme P we can write down a pair of quantifier-free Horn formulae $\langle\theta_P,\eta_P\rangle$, not depending on the $\bar a_\gamma$ and the $\bar b_\gamma$, so that P is valid iff

$$\text{for all } i\in I,\quad A_i\models\theta_P[\bar a^i],$$

and

$$\text{for all } j\in J,\quad B_j\models\eta_P[\bar b^j].$$

Take Δ to be the set of all pairs $\langle\theta_P,\eta_P\rangle$ such that P is a proof-scheme with conclusion $\perp$. Then clause (i) of the lemma says that no such P is valid, or in other words: not $\vdash^*\perp$.

If not $\vdash^*\perp$, then by the claim, the empty set $A(\perp)=B(\perp)$ is not in D or E, and hence D and E are proper. By what we have already proved about D and E, it follows that they are κ-filters satisfying (ii). Thus (i) implies (ii).

Conversely suppose that D', E' are κ-filters, on I, J as in (ii). Then for every atomic formula ψ, $A(\psi)\in D'$ iff $B(\psi)\in E'$. Now let P be a valid proof-scheme and N a node of P labelled A; let Φ be the set of atomic formulae immediately above N, and ψ the formula attached at N. The condition that N is everywhere correct implies that (I) holds. But D' is κ-complete, so that if $A(\varphi)\in D'$ for every $\varphi\in\Phi$, then $A(\psi)\in D'$. By induction down valid proof-schemes, it follows that for every atomic formula ψ, if $\vdash^*\psi$, then $A(\psi)\in D'$. Since D' is a κ-filter it is proper, and hence not $\vdash^*\perp$. Hence (ii) implies (i).

Here follow some variants of Lemma 3a. The first is the obvious adaptation to κ-ultrafilters:

Lemma 3b. *Let κ be strongly compact. There is a family Δ of ordered pairs $\langle\theta,\eta\rangle$ of quantifier-free formulae of $\mathrm{PL}_{\kappa\kappa}$, not depending on $\langle\bar a_\gamma:\gamma<\mu\rangle$, $\langle\bar b_\gamma:\gamma<\mu\rangle$, such that the following are equivalent:*

(i) *for every pair $\langle\theta,\eta\rangle$ in Δ, either there is $i\in I$ such that $A_i\models\neg\theta[\bar a^i]$ or there is $j\in J$ such that $B_j\models\neg\eta[\bar b^j]$;*

(ii) *there are κ-ultrafilters D, E on I, J respectively, such that for every atomic formula ψ, $\prod_D A_i\models\psi[\bar a_\gamma]_{\gamma<\mu}$ iff $\prod_E B_j\models[\bar b_\gamma]_{\gamma<\mu}$.*

Proof. Define $A(\varphi)$, $B(\varphi)$ as in the proof of Lemma 3a. Let Φ be any set of fewer than κ atomic formulae. Then (ii) implies the following:

(iii) there are Φ_1, Φ_2 such that $\Phi = \Phi_1 \cup \Phi_2$, and there are κ-filters D', E' on I, J such that for every $\psi \in \Phi_1$, $A(\psi) \in D'$ and $B(\psi) \in E'$, while for every $\psi \in \Phi_2$, $A(\neg\psi) \in D'$ and $B(\neg\psi) \in E'$.

Conversely if (iii) holds for every set Φ of fewer than κ atomic formulae, then the strong compactness of κ allows us to deduce (ii). So it will be enough if we fix Φ and define Δ to make (i) and (iii) equivalent; the union of the Δ defined for the different Φ will work for the lemma.

Henceforth Φ is a fixed set of fewer than κ atomic formulae. We define proof-scheme as in the proof of Lemma 3a, with the following changes. In clause (4), 'atomic' becomes 'atomic or negated atomic'. We add a new clause: (6) zero or more maximal nodes of P are designated as *premises*. P is *valid* iff all its nodes which are not premises are everywhere correct. We write $\Psi \vdash^* \psi$ iff there is a valid proof-scheme whose bottom formula is ψ and whose premise-formulae are elements of Ψ. For each proof-scheme P we can write down a pair of quantifier-free formulae $\langle \theta_P, \eta_P \rangle$, not depending on the $\tilde{a}_\gamma$ and the $\bar{b}_\gamma$, so that P is valid iff

$$\text{for all } i \in I, \quad A_i \models \theta_P[\bar{a}^i],$$

and

$$\text{for all } j \in J, \quad B_j \models \eta_P[\bar{b}^j].$$

Next, for each $\Psi \subseteq \Phi$, define D_Ψ to be the filter on I generated by all intersections of fewer than κ sets $A(\psi)$ such that $\Psi \cup \{\neg\varphi : \varphi \in \Phi - \Psi\} \vdash^* \psi$; likewise E_Ψ on J. Then D_Ψ and E_Ψ are obviously κ-complete, and the same argument as for Lemma 3a shows:

Claim. *For each atomic or negated atomic formula* ψ, $A(\psi) \in D_\Psi$ *iff* $\Psi \cup \{\neg\varphi : \varphi \in \Phi - \Psi\} \vdash^* \psi$ *iff* $B(\psi) \in E_\Psi$.

Now by the claim, (iii) above holds provided that we have:

(iv) there is $\Psi \subseteq \Phi$ such that not $\Psi \cup \{\neg\varphi : \varphi \in \Phi - \Psi\} \vdash^* \bot$.

(Put $D' = D_\Psi$ and $E' = E_\Psi$.) Conversely if (iii) holds, then the argument of the last part of the proof of Lemma 3a shows that (iv) holds too. So we have to find a Δ which makes (i) equivalent to (iv). Define a *proof-system* to be a map $p : \mathcal{P}(\Phi) \to$ (proof-schemes) such that for each $\Psi \subseteq \Phi$, $p(\Psi)$ is a proof-scheme with premises $\subseteq \Psi \cup \{\neg\varphi : \varphi \in \Phi - \Psi\}$ and conclusion $\bot$. Then (iv) is equivalent to:

(v) for every proof-system p there is $\Psi \subseteq \Phi$ such that $p(\Psi)$ is not valid.

For each proof-system p define formulae

$$\theta_p = \bigwedge_{\Psi \subseteq \Phi} \theta_{p(\Psi)}, \qquad \eta_p = \bigwedge_{\Psi \subseteq \Phi} \eta_{p}(\Psi).$$

Then θ_p, η_p are in $L_{\kappa\kappa}$ because κ is strongly inaccessible. Let Δ be the set $\{\langle\theta_p, \eta_p\rangle$: p is a proof-system$\}$. For this Δ, (i) says precisely that for every proof-system p there is some invalid $p(\Psi)$; so (i) is equivalent to (v) as required.

The next two variants of Lemma 3 need no new ideas:

Lemma 3c. *There is a family Δ of ordered pairs $\langle\theta, \eta\rangle$ of quantifier-free formulae of* $\mathrm{PH}_{\kappa\kappa}$, *not depending on $\langle\bar{a}_\gamma: \gamma<\mu\rangle$, $\langle\bar{b}_\gamma: \gamma<\mu\rangle$, such that each θ is a conjunction of atomic formulae, and the following are equivalent:*

(i) *for every pair $\langle\theta, \eta\rangle$ in Δ, either there is $i\in I$ such that $A_i\models\neg\theta[\bar{a}^i]$ or there is $j\in J$ such that $B_j\models\neg\eta[\bar{b}^j]$;*

(ii) *there are κ-filters D, E on I, J respectively, such that for every atomic formula ψ, if $\prod_D A_i\models\psi[\bar{a}_\gamma]_{\gamma<\mu}$, then $\prod_E B_j\models\psi[\bar{b}_\gamma]_{\gamma<\mu}$.*

Lemma 3d. *Let κ be strongly compact. There is a family Δ of ordered pairs $\langle\theta, \eta\rangle$ of quantifier-free formulae of* $\mathrm{PL}_{\kappa\kappa}$, *not depending on $\langle\bar{a}_\gamma: \gamma<\mu\rangle$, $\langle\bar{b}_\gamma: \gamma<\mu\rangle$, such that each η is in* $\mathrm{PH}_{\kappa\kappa}$, *and the following are equivalent:*

(i) *for every pair $\langle\theta, \eta\rangle$ in Δ, either there is $i\in I$ such that $A_i\models\neg\theta[\bar{a}^i]$ or there is $j\in J$ such that $B_j\models\neg\eta[\bar{b}^j]$;*

(ii) *there are a κ-ultrafilter D on I and a κ-filter E on J such that for every atomic formula ψ, $\prod_D A_i\models\psi[\bar{a}_\gamma]_{\gamma<\mu}$ iff $\prod_E B_j\models\psi[\bar{b}_\gamma]_{\gamma<\mu}$.*

2.2. *Isomorphic reduced products*

We shall give a necessary and a sufficient condition for two structures to have isomorphic κ-reduced powers, assuming only the GCH. The conditions are both local in the sense that they involve only games of length less than κ on the two structures, but only one of them is straightforwardly syntactic. If there is a proper class of measurable cardinals, then it is consistent that the two conditions are equivalent. There are analogous results for κ-reduced products of sets of structures, for κ-ultrapowers when κ is strongly compact, for surjective homomorphisms between κ-reduced powers, and so on.

First we must weaken the notion of a winning strategy. Suppose two players $\forall$ and $\exists$ play a game G; imagine also that they imitate the chess wizards and play G simultaneously on μ boards. (If player $\forall$ makes the αth move in G and player $\exists$ makes the $(\alpha+1)$th, then player $\forall$ must make his αth move on all boards before player $\exists$ makes his $(\alpha+1)$th on any; and vice versa.) We can define a new game G^μ by declaring that player $\forall$ wins iff he wins G on at least one board. Obviously if player $\exists$ has a winning strategy in either of G or G^μ, then he wins the other too. But in general it is possible for player $\forall$ to have a winning strategy for G^μ and not for G. (Examples can be found along the lines of (4) in Section 4.3 below.)

Theorem 4a. *Let A and B be structures of the same similarity type. Then* (i)$\Rightarrow$(ii)$\Rightarrow$(iii), *where* (*i*)–(*iii*) are as follows:

(i) *there is a regular cardinal* $\mu \geq \max(\text{card } A, \text{card } B, \text{card(type of } A, B))$ *such that* $\mu^{<\kappa} = \mu$ *and* $2^\mu = \mu^+$, *and for every pair* $\langle \theta, \eta \rangle$ *of mutually inconsistent sentences of* $\mathrm{PH}_{\kappa\kappa}$, *player* $\forall$ *has a winning strategy for at least one of* $G^\mu(\theta, A)$ *and* $G^\mu(\eta, B)$;

(ii) *A and B have isomorphic* κ*-reduced powers*;

(iii) *if* T_A, T_B *are respectively the sets of* $\mathrm{PH}_{\kappa\kappa}$*-sentences true in A, B, then* $T_A \cup T_B$ *is consistent.*

(Note that if the GCH holds, then in (i) we can omit all the conditions on μ except the last one.)

Proof. (ii)$\Rightarrow$(iii) is by Lemma 1. Now we assume (i) and prove (ii). Referring back to the beginning of 2.1, we put $I = J = \mu$, $A_i = A$ and $B_j = B$ for each $i, j < \mu$. The cardinal μ of 2.1 now becomes $2^\mu = \mu^+$. By Lemma 3a, (ii) above is proved provided we can find $\bar{a}_\gamma$ $(\gamma < \mu^+)$ in A^μ and $\bar{b}_\gamma$ $(\gamma < \mu^+)$ in B^μ such that

(iv) for every pair $\langle \theta, \eta \rangle$ in Δ (from Lemma 3a), either there is $i < \mu$ such that $A \models \neg\theta[\bar{a}^i]$ or there is $j < \mu$ such that $B \models \neg\eta[\bar{b}^j]$; and $\langle \bar{a}_\gamma : \gamma < \mu^+ \rangle$, $\langle \bar{b}_\gamma : \gamma < \mu^+ \rangle$ list the whole of A^μ, B^μ respectively.

We shall make players $\forall$ and $\exists$ choose the $\bar{a}_\gamma$ in sequence; player $\forall$ chooses $\bar{a}_\gamma$ when γ is even. Independently of this, the two players will choose the $\bar{b}_\gamma$ in sequence, but here player $\forall$ will choose the $\bar{b}_\gamma$ with odd γ. The only requirement on player $\exists$ is that he chooses so that every element of A^μ, B^μ is chosen at some point. This guarantees the last part of (iv).

Player $\forall$ is going to have to splice together μ different strategies for μ^+ different games of length $<\kappa$ on μ different boards. To show how he can do it, we shall use an unpublished lemma of Galvin. We thank Galvin for permission to include this result. (He proved it in 1973 in answer to a question of Laver, whether $\in \cap (\omega_1 \times \omega_1)$ can be written as the union of an increasing ω-sequence of tree-orderings without branches of length ω_1.)

Lemma (Galvin). *Let* μ *be a regular cardinal. Then there is a sequence* $\langle R_j : j < \mu \rangle$ *such that*

(1) $\bigcup_{j<\mu} R_j = \{\langle \alpha, \beta \rangle : \alpha < \beta < \mu^+\}$;

(2) $j < k < \mu \Rightarrow R_j \subseteq R_k$;

(3) *for every* $j < \mu$, R_j *is a tree whose branches all have length* $\leq \mu$;

(4) *if* $\mu = \omega$, *then every branch of* R_j *has length* $\leq j + 1$.

Proof. By induction on α, construct for each $\alpha < \mu^+$ a function $f_\alpha : \mu \to \mu$, so that

$$\alpha < \beta \Rightarrow |\{j < \mu : f_\alpha(j) \geq f_\beta(j)\}| < \mu.$$

(This is possible since μ is regular.) Next, choose for each $\alpha<\mu^+$ a set $C_\alpha\subseteq\alpha$ as follows: $C_0=0$; $C_{\alpha+1}=\{\alpha\}$; if α is a limit ordinal, C_α is a cofinal subset of α of order-type $\leqslant\mu$.

Now we define the relation

$$\langle\alpha,\beta\rangle\in R_j \qquad (*)$$

by induction on β, simultaneously for all α and j, as follows. We define (*) to hold iff there is $\gamma\in C_\beta$ such that

(i) $\langle\alpha,\gamma\rangle\in R_j$ or $\alpha=\gamma$;

(ii) $f_\gamma(k)<f_\beta(k)$ for all $k\geqslant j$;

(iii) $\langle\delta,\gamma\rangle\in R_j$ whenever $\delta\in C_\beta\cap\gamma$.

Then (1), (2) are easily verified. For (3), observe that if $\langle\alpha,\beta\rangle\in R_j$, then $f_\alpha(j)<f_\beta(j)$. For (4), choose the f_α so that for each α and j, $f_\alpha(j)\leqslant j$.

Using Galvin's lemma and the fact that $\mu^{<\kappa}=\mu$, we can find a family $<_j$ $(j<\mu)$ of partial orderings with field μ^+, such that

(a) if $\alpha<_j\beta$ then $\alpha<\beta$;

(b) each $<_j$ is a tree whose branches all have length $\leqslant\kappa$;

(c) if S is a subset of μ^+ with cardinality $<\kappa$, then for some $j<\mu$, S is an initial part of a branch of $<_j$.

Now suppose the quantifier-free formula θ of $\mathrm{PH}_{\kappa\kappa}$ occurs in some pair $\langle\theta,\eta\rangle$ in the set Δ of Lemma 3a. Choose a quantifier prefix Q with domain a set of ordinals $<\mu^+$ of order-type $<\kappa$, such that if γ is an even ordinal and either v_γ or $v_{\gamma+1}$ occurs in θ, then $\mathsf{Q}(\gamma)=\forall$ and $\mathsf{Q}(\gamma+1)=\exists$, and no ordinals occur in the domain of Q except as just indicated. Write θ' for $\mathsf{Q}\theta$. Similarly for each η choose a prefix Q by the same rules but with $\mathsf{Q}(\gamma)=\exists$ and $\mathsf{Q}(\gamma+1)=\forall$, and write η' for $\mathsf{Q}\eta$. Write Δ' for the set of pairs $\langle\theta',\eta'\rangle$ such that $\langle\theta,\eta\rangle$ is in Δ.

Claim. *For each pair $\langle\theta',\eta'\rangle$ from Δ', the sentences θ' and η' are mutually contradictory.*

For suppose not; then both are true in some structure C. Now we can consider θ' as defining a game of length μ^+ on C in which player $\forall$ moves at even-numbered steps (by adding vacuous pairs $\forall v_\gamma\exists v_{\gamma+1}$ to the quantifier); likewise η' defines a game of length μ^+ in which player $\forall$ moves at the odd steps. By assumption, player $\exists$ has winning strategies σ, τ for these two games. Let the players now play $G(\theta',C)$; let player $\exists$ use his winning strategy σ, and let player $\forall$ use the strategy τ. Suppose $\bar c$ is the resulting sequence of elements of C. Then $C\models\theta\wedge\eta[\bar c]$. Applying Lemma 3a to the situation where I, J are singletons and $A_i=B_j=C$, this implies that for some atomic formula ψ, $C\models\psi[\bar c]$ iff not $C\models\psi[\bar c]$, which is absurd. The claim is proved.

By the claim and (i) of the theorem, player $\forall$ has a winning strategy for either $G^\mu(\theta',A)$ or $G^\mu(\eta',B)$, for each pair $\langle\theta,\eta\rangle$ in Δ.

Sh:109

List as φ_i $(i<\mu)$, possibly with repetitions, all the sentences of $\mathrm{PH}_{\kappa\kappa}$ whose variables have indices $<\kappa$, such that player $\forall$ moves at even steps in $G(\varphi_i, A)$ and has a winning strategy for $G^\mu(\varphi_i, A)$. For each $G^\mu(\varphi_i, A)$ choose a winning strategy σ_i. Take a bijection $g: \mu^3 \to \mu$. Player $\forall$ will now choose the elements $\bar{a}_\gamma$ (γ even) of A^μ as follows. For each $i, j<\mu$ the indices $g(i, j, k)$ $(k<\mu)$ form a set of μ boards on which he can play $G^\mu(\varphi_i, A)$. At move α (even), he plays on these indices using strategy σ_i and assuming that the previous moves on board k are $\langle \bar{a}_\delta(g(i, j, k)) : \delta <_j \gamma\rangle$, if this is a sequence whose length is an even ordinal less than the length of $G^\mu(\varphi_i, A)$; otherwise he plays as he likes. Player $\forall$ chooses the elements $\bar{b}_\gamma$ (γ odd) according to the same rubrics, but with odd and even reversed.

Now we can prove (iv). Let $\langle\theta, \eta\rangle$ be a pair from Δ. We have seen that player $\forall$ has a winning strategy for either $G^\mu(\theta', A)$ or $G^\mu(\eta', B)$; say he has one for $G^\mu(\theta', A)$. Let S be the set of all indices of variables which occur in θ'. By collapsing S down to an initial segment of the ordinals, we get a sentence θ^* of $\mathrm{PH}_{\kappa\kappa}$ whose variables all have indices $<\kappa$, such that player $\forall$ has a winning strategy for $G^\mu(\theta^*, A)$. Then θ^* is φ_i for some $i<\mu$. Also there is $j<\mu$ such that S is an initial segment of the partial ordering $<_j$. From the choice of the $\bar{a}_\gamma$ it follows that for some $k<\mu$,

$$A \models \neg\theta[\bar{a}^{g(i,j,k)}].$$

Thus (iv) is proved.

It was only to avoid a plethora of indices that we did not straight away prove:

Theorem 4b. *Let H and K be classes of structures, all of the same similarity type. Then* (i) $\Rightarrow$ (ii) $\Rightarrow$ (iii), where (i)–(iii) *are as follows*:

(i) *there are arbitrarily large regular cardinals* μ such that $\mu^{<\kappa} = \mu$ and $2^\mu = \mu^+$, *and for every pair $\langle\theta, \eta\rangle$ of mutually inconsistent sentences of* $\mathrm{PH}_{\kappa\kappa}$, *player $\forall$ has a winning strategy for at least one of the games $G^\mu(\theta, A)$ with $A \in H$ or $G^\mu(\eta, B)$ with $B \in K$ (the proof shows that a 'large enough' μ will do)*;

(ii) *some κ-reduced product of structures in H is isomorphic to a κ-reduced product of structures in K*;

(iii) *if T_H, T_K are respectively the sets of* $\mathrm{PH}_{\kappa\kappa}$*-sentences true throughout H, K, then $T_H \cup T_K$ is consistent.*

Proof. Again (ii) $\Rightarrow$ (iii) by Lemma 1. For (i) $\Rightarrow$ (ii), we can assume without loss that H and K are sets, since $\mathrm{PH}_{\kappa\kappa}$ has only a set of non-equivalent sentences. Then the argument proceeds very much as before. The cardinal μ is chosen larger than the cardinalities of all structures in $H \cup K$. Instead of listing the winning strategies σ_i, we list pairs (σ_i, A) such that σ_i is a winning strategy for player $\forall$ in $G^\mu(\varphi_i, A)$, and we take A to be the $g(i, j, k)$th element (for all $j, k<\mu$) in the product.

Since clause (i) has exactly the same form in all the versions of Lemma 3, the corresponding versions of Theorem 4 can be read off automatically from Theorems 4a and 4b above. For example we have conditions for two structures to have isomorphic κ-ultrapowers.

Next we ask when two structures have isomorphic κ-reduced powers over the *same* filter. This turns out to be a surprisingly strong condition. Since the filter is the same on both sides, it is not possible to separate out the syntactic conditions on the two structures. So the work in Section 2.1 is no help. Instead we introduce the following game $G(\Phi, A, B)$.

Let A and B be structures and Φ a set of atomic formulae. Then $G(\Phi, A, B)$ is played as follows. At the γth move, when γ is even, player $\forall$ chooses an element of A and then player $\exists$ chooses an element of B. When γ is odd, player $\forall$ first chooses an element of B and then player $\exists$ chooses from A. The players play thus until they have constructed sequences $\bar{a}$, $\bar{b}$ from A, B respectively which are long enough to cover the variables of the formulae in Φ. Player $\exists$ wins iff for every $\varphi \in \Phi$,

$$A \models \varphi[\bar{a}] \quad \text{iff } B \models \varphi[\bar{b}].$$

$G^\lambda(\Phi, A, B)$ is $G(\Phi, A, B)$ played on λ boards; player $\exists$ wins iff he wins on all boards.

Theorem 4c (GCH). *Let A and B be structures. Then* (i) $\Rightarrow$ (ii) $\Rightarrow$ (iii) $\Rightarrow$ (iv), *where* (i)–(iv) *are as follows:*

(i) *there is λ such that for every set Φ of fewer than κ atomic formulae player $\exists$ has a winning strategy in $G^\lambda(\Phi, A, B)$,*

(ii) *there are a set I and a κ-filter D on I such that for every κ-filter D' on I which extends D, $A^I/D' \cong B^I/D'$;*

(iii) *there are a set I and a κ-filter D on I such that $A^I/D \cong B^I/D$;*

(iv) *for every set Φ of fewer than κ atomic formulae, player $\forall$ has no winning strategy in $G(\Phi, A, B)$.*

Proof. (i) $\Rightarrow$ (ii): by GCH, choose a regular cardinal μ such that $\lambda \leq \mu^{<\kappa} = \mu$ and $2^\mu = \mu^+ \geq \text{card}(A) + \text{card}(B) + \text{card}(\text{type of } A, B)$. Players $\forall$ and $\exists$ shall play a pre-game of length μ^+ as follows. At the γth move, when γ is even, player $\forall$ chooses an element of A^μ and then player $\exists$ chooses an element of B^μ. When γ is odd, player $\forall$ first chooses from B^μ and then player $\exists$ from A^μ. Player $\forall$ plays so that his moves exhaust $A^\mu \cup B^\mu$. Let $\langle \bar{a}_\gamma : \gamma < \mu^+ \rangle$, $\langle \bar{b}_\gamma : \gamma < \mu^+ \rangle$ be the resulting sequences of elements of A^μ, B^μ. If (i) holds, then player $\exists$ can use Galvin's lemma as in the proof of Theorem 4a to ensure that for every set Φ of fewer than κ atomic formulae whose variables have indices $<\mu^+$, the set

$$X(\Phi) = \{i < \mu : A \models \varphi[\bar{a}^i] \Leftrightarrow B \models \varphi[\bar{b}^i] \text{ for all } \varphi \in \Phi\}$$

is non-empty. The filter D generated by the sets $X(\Phi)$ will satisfy (ii).

(ii) ⇒ (iii) is trivial. For (iii) ⇒ (iv), let D be a κ-filter on I and suppose (iv) fails for some Φ; then player ∀ has a winning strategy σ for $G(\Phi, A, B)$, and so he can win $G(\Phi, A^I/D, B^I/D)$ by playing σ at every coordinate. Clearly this implies that A^I/D and B^I/D are not isomorphic.

We make two remarks on Theorem 4c. First, when $\kappa = \omega$ the games $G(\Phi, A, B)$ are the familiar Ehrenfeucht–Fraissé games, which are all determinate. So in this case conditions (i) and (iv) are equivalent, and they simply say that A and B are elementarily equivalent. When $\kappa > \omega$, conditions (i) and (iv) are (possibly equivalent) generalisations of elementary equivalence. It is known that in general the Ehrenfeucht–Fraissé games of length κ do not characterise equivalence in $L_{\kappa\kappa G}$.

Second, there is no need to prove a separate version of Theorem 4c for the case where κ is strongly compact, because in this case a κ-ultrafilter can be got straight away from condition (ii).

It remains to ask about the gap between the sufficient and necessary conditions in Theorem 4; say (i) and (iii) in Theorem 4a. There is no situation where (i) and (iii) are known not to be equivalent. Nevertheless we conjecture that (i) is in general much stronger than (iii), and that it is almost never true in the constructible universe.

There are two situations in which we can show that (i) is equivalent to (iii). The first is when $\kappa = \omega$. The second situation is as follows. When I is an ideal on a boolean algebra B, we write I^+ for the set of elements of B which are not in I. Let $I(\lambda)$ be the statement:

There is an ideal I on λ^{++} which is λ^{++}-complete, normal and such that I^+ has a dense subset K with the property that every descending sequence of length $<\lambda^+$ in K has a lower bound in K.

A theorem of Richard Laver (due also in part to Menachem Magidor) states that if M is a model of ZFC containing a cardinal λ and a measurable greater than λ, then M has a boolean extension in which λ remains a cardinal, $I(\lambda)$ holds, and the GCH holds above λ. (The proof can be inferred from the case $\lambda = \omega$ which is described in Galvin et al. [5, Section 4].) In fact Laver shows that if there is a proper class of measurables in M, then there is a boolean model in which GCH holds and $I(\lambda)$ is true for arbitrarily large λ; if M has a proper class of supercompact cardinals, then the boolean model can have $I(\lambda)$ true for all infinite cardinals λ.

The next theorem is essentially a remark of Fred Galvin:

Theorem 5. *Suppose* $I(\lambda)$ *and the* GCH *hold, where* $\lambda^+ \geq \max(\kappa, \text{card } A, \text{card } B, \text{card(type of } A, B))$; *then in Theorem* 4*a*, (i) *and* (iii) *are equivalent.* (*The same argument shows that in Theorem* 4*c*, (i) *and* (iv) *are equivalent.*)

Proof. We show that if (iii) holds, then (i) holds with $\mu=\lambda^{++}$. Write I for the ideal given by $I(\lambda)$, and K for the dense subset of I^+. We can assume that K is closed under intersections of descending chains of length $<\lambda^+$. By Ulam matrices and the λ^{++}-completeness of I, every set in I^+ can be split into λ^{++} pairwise disjoint sets which are also in I^+. By the λ^{++}-completeness of I, every partition of a set in I^+ into at most λ^+ subsets includes at least one set in I^+.

Let ξ be an ordinal $<\kappa$, and let G be a game of length ξ in which players $\forall$ and $\exists$ alternately choose elements of a set C of cardinality $\leqslant\lambda^+$, and player $\exists$ wins iff the resulting sequence lies in a given set S. (The games $G(\theta, A)$ and $G(\eta, B)$ of Theorem 4a have this form.) Suppose that player $\exists$ has no winning strategy for G. We show that player $\forall$ has a winning strategy for G^μ.

First observe that any move $\bar{a}$ in G^μ determines a partition of each set $X\subseteq\mu$, by putting i, j in the same partition class iff $\bar{a}(i)=\bar{a}(j)$. Now before the αth move in G^μ, player $\forall$ should choose a family F_α of pairwise disjoint elements of K, so that if $\beta<\alpha$, then each set $X\in F_\alpha$ is included in some $Y\in F_\beta$. F_0 is chosen arbitrarily, and at limit δ the player should choose F_δ to be the set of all minimal non-empty intersections of sets from previous F_β. At successor moves $\alpha+1$ where $\exists$ has just moved and chosen $\bar{a}\in C^\mu$, player $\forall$ should choose $F_{\alpha+1}$ by replacing each $X\in F_\alpha$ by some subset of X which is in K and which lies inside one partition class of the partition determined by $\bar{a}$. Finally if player $\forall$ is to move at stage α, then after choosing F_α he should split each $X\in F_\alpha$ into card(C) disjoint parts which are all in I^+; now he chooses $\bar{a}$ so as to play a different element of C on each part of X, exhausting C. For $F_{\alpha+1}$ he replaces each of these parts of each X by an element of K which is included in it, and then makes $F_{\alpha+1}$ the set of these elements of K.

Now the sets in $\bigcup_{\alpha<\xi}F_\alpha$ form a downwards tree in I^+. By $I(\lambda)$, each branch b of this tree has non-empty intersection; pick an element i_b in the intersection. On these selected indices i_b, player $\exists$ is playing a constant strategy and player $\forall$ is trying every possible move against him. Since player $\exists$'s strategy was not winning, player $\forall$ wins on at least one index.

To make everything explicit, Theorems 4 and 5 together show that it is consistent (granted enough measurable cardinals) that two structures have isomorphic reduced powers over κ-complete filters iff their Horn theories in $\mathrm{PH}_{\kappa\kappa}$ are consistent with each other.

Does the statement 'Any two structures with consistent $\mathrm{PH}_{\kappa\kappa}$ theories have isomorphic κ-reduced powers' imply the existence of precipitous ideals?

3. Limit reduced products

In this section was assume only ZFC and prove the analogue of Theorem 4 for *limit* reduced products. We also give a characterisation of limit κ-reduced powers which generalises a theorem of Keisler for $\kappa=\omega$.

3.1. Combinatorial lemmas

The following lemma occurs under the name of Remark 3 in Engelking and Karłowicz [3].

Engelking–Karłowicz Lemma. *Suppose μ is regular and $\mu^{<\kappa} = \mu \geq \mathrm{card}(X)$. Then there is a family $(f_i: i<2^\mu)$ of maps $f_i: \mu \to X$ such that for every strictly increasing sequence $\langle \alpha_i: i<\xi \rangle$ of ordinals $\alpha_i < 2^\mu$, with $\xi < \kappa$, and every family $\langle x_i: i<\xi \rangle$ of elements of X, there is a $j<\mu$ such that*

$$f_{\alpha_i}(j) = x_i \quad \text{for all } i<\xi.$$

A family of functions $(f_i: i<2^\mu)$ as in this lemma will be said to be (μ, κ, X)-*independent.*

Let A be a structure and choose $\mu \geq \mathrm{card}(A)$ such that $\mu^{<\kappa} = \mu$. Let $(f_i: i<2^\mu)$ be a $(\mu, \kappa, \mathrm{dom}\, A)$-independent family. Then each f_i is an element of A^μ. Let F be the κ-complete filter on $\mathrm{Part}(\mu)$ generated by the partitions $[f_i]$, $i<2^\mu$. We call $A^\mu \mid F$ the (μ, κ)-*independent limit power of A generated by* $(f_i: i<2^\mu)$. If g is an element of $A^\mu \mid F$, then there is a unique smallest set Z of generators f_i such that $[g] \geq \bigwedge \{[f_i]: f_i \in Z\}$. Z has cardinality $<\kappa$. We call Z the *support* $\mathrm{supp}(g)$ of g.

If $A^\mu \mid F$ is as above, and Q is a quantifier of length $\leq 2^\mu$, then either player in $G(\mathrm{Q}, A^\mu \mid F)$ can paly by the following strategy: at the βth move, play the first f_i which is not in $\bigcup \{\mathrm{supp}(g)$: g was the γth move, $\gamma<\beta\}$. We call this the *independent strategy.*

Lemma 6. *Let Q be a quantifier of length $\alpha \leq 2^\mu$, and let $A^\mu \mid F$ be the (μ, κ)-independent limit power of A generated by $(f_i: i<2^\mu)$. Suppose that in a play of $G(\mathrm{Q}, A^\mu \mid F)$, player $\forall$ has followed his independent strategy, and let X be a subset of α of power $<\kappa$. Then there are $Z \subseteq \mu$ and a strategy σ for player $\exists$ in $G(\mathrm{Q} \restriction X, A)$ such that*

(i) *player $\exists$ has played σ in $G(\mathrm{Q} \restriction X, A)$ at each coordinate $i \in Z$;*

(ii) *each possible play of player $\forall$ against σ in $G(\mathrm{Q} \restriction X, A)$ occurs at some $i \in Z$.*

Proof. Let $\langle \bar{a}_\beta: \beta<\alpha \rangle$ be the play. Let $K_\forall$ be $\{\bar{a}_\beta: \beta \in X \cap \mathrm{Q}^{-1}(\forall)\}$ and let $K_\exists$ be $\bigcup \{\mathrm{supp}(\bar{a}_\beta): \beta \in X \cap \mathrm{Q}^{-1}(\exists)\}$. Pick any element $c \in \mathrm{dom}\, A$, and put $Z = \{i<\mu: f(i) = c$ for each $f \in K_\exists - K_\forall\}$. Now for each $\beta \in X \cap \mathrm{Q}^{-1}(\exists)$, $\bar{a}_\beta(i) \neq \bar{a}_\beta(j)$ implies $f(i) \neq f(j)$ for some $f \in \mathrm{supp}(\bar{a}_\beta)$; assuming $i, j \in Z$, this implies $f(i) \neq f(j)$ where f is $\bar{a}_\gamma$ for some $\gamma \in X \cap \mathrm{Q}^{-1}(\exists) \cap \beta$. It follows that player $\exists$ adopts a uniform strategy σ throughout Z, proving (i). Part (ii) then follows from the conclusion of the Engelking–Karłowicz lemma, since $\mathrm{card}(X) < \kappa$.

3.2. Characterisation of limit κ-reduced powers

We generalise Keisler's characterisation of limit ultrapowers as the most general operation which commutes with formation of reducts and preserves elementary equivalence. (See Theorem 6.4.10 of Chang and Keisler [2].)

Let us say that a structure A is *κ-complete* iff every function of arity $<\kappa$ which can be defined on dom A is of form F_A for some function symbol F of the language of A.

Theorem 7a. *If A and B are structures of the same similarity type, and A is κ-complete, then the following are equivalent:*

(i) *every Horn sentence of $L_{\kappa\kappa}$ which is true in A is true in B;*
(ii) *every universal Horn sentence of $L_{\kappa\kappa}$ which is true in A is true in B;*
(iii) *B is isomorphic to a limit κ-reduced power of A.*

Proof. (iii)$\Rightarrow$(i) is by Lemma 1 and (i)$\Rightarrow$(ii) is trivial. For (ii)$\Rightarrow$(iii), choose $\mu \geq \text{card}(A)$ such that $\mu^{<\kappa} = \mu$ and $2^\mu \geq \text{card}(B)$. Let $A^\mu \mid F$ be the (μ, κ)-independent limit power of A generated by $(f_\alpha : \alpha < 2^\mu)$. Choose any surjective map $\theta : \{f_\alpha : \alpha < 2^\mu\} \to \text{dom } B$. Define D to be the filter on μ generated by all intersections of fewer than κ sets of form

$$A(\varphi) = \{i < \mu : A \models \varphi[f_\alpha(i)]_{\alpha<2^\mu}\}$$

such that φ is atomic and $B \models \varphi[\theta f_\alpha]_{\alpha<2^\mu}$.

We claim that for every atomic ψ, if $A(\psi) \in D$, then $B \models \psi[\theta f_\alpha]_{\alpha<2^\mu}$. For suppose $A(\psi) \in D$. Then there are atomic φ_k $(k < \gamma < \kappa)$ such that $B \models \varphi_k[\theta f_\alpha]_{\alpha<2^\mu}$ for each $\alpha < \gamma$ and $\bigcap_{k<\gamma} A(\varphi_k) \subseteq A(\psi)$. If

$$A \models \forall \bar{x}\Big(\bigwedge_{k<\gamma} \varphi_k \to \psi\Big), \tag{1}$$

then by assumption (i) the same sentence holds in B, and so $B \models \psi[\theta f_\alpha]_{\alpha<2^\mu}$ as required. So we prove (1), as follows. If (1) fails, then there is a sequence $\bar{a}$ in A such that $A \models \varphi_k[\bar{a}]$ for each $k < \gamma$, but $A \models \neg\psi[\bar{a}]$. Since the f_α are $(\mu, \kappa, \text{dom } A)$-independent, and fewer than κ variables occur in ψ and the φ_k, we can find $i < \mu$ such that the sequence $(f_\alpha(i))_{\alpha<2^\mu}$ agrees with $\bar{a}$ at the relevant places. This contradicts the fact that $\bigcap_{k<\gamma} A(\varphi_k) \subseteq A(\psi)$. The claim is proved.

In particular D is proper, since $\bot$ is atomic. The claim also shows that θ induces an isomorphism between B and C/D where C is the substructure of A^μ which is generated by the f_α $(\alpha < 2^\mu)$. Now if $c \in \text{dom } A^\mu \mid F$, then for some $Y \subseteq 2^\mu$ of cardinality $<\kappa$, $[c] \geq \bigwedge_{\alpha\in Y} [f_\alpha]$; since A is κ-complete it follows that $c \in \text{dom } C$. Hence $A^\mu \mid F = C$ and $A^\mu_D \mid F = C/D \cong B$.

Essentially the same proof gives the following variant. Keisler's theorem was the case $\kappa = \omega$.

Theorem 7b. *Let κ be strongly compact, and let A and B be structures of the same similarity type, with A κ-complete. Then the following are equivalent:*

(i) *every (universal) sentence of $\text{PL}_{\kappa\kappa}$ which is true in A is true in B;*
(ii) *B is isomorphic to a limit κ-ultrapower of A.*

There is one further variant which we have not developed elsewhere in this paper, but it should be mentioned somewhere. Let κ and λ be regular cardinals with $\kappa \geq \lambda$. Define $PL_{\kappa\lambda}$ as $PL_{\kappa\kappa}$ but with the quantifier prefixes restricted to be of length $<\lambda$; we assume that all relations and functions have arity $<\lambda$.

Let F be a λ-complete filter on Part(I). For each partition $\pi \in F$, let B_π be the boolean algebra of all sets of form $\bigcup X$ where $X \subseteq \pi$; then B_π is a complete subalgebra of $\mathcal{P}(I)$. By a (κ, F)*-ultrafilter* on I we mean a subset D of $\bigcup_{\pi\in F} B_\pi$ whose restriction to each B_π is a κ-complete ultrafilter. Let A_i $(i \in I)$ be a family of structures of the same similarity type. By a λ*-limit* κ*-ultraproduct* of the A_i we mean a structure $\prod_D A_i \mid F$ where F is a λ-complete filter on Part(I) and D is a (κ, F)-ultrafilter on I. The reader can easily verify that this definition makes sense, and that every sentence of $PL_{\kappa\lambda}$ is preserved in λ-limit κ-ultraproducts.

Theorem 7c. *Let $\kappa \geq \lambda$ be regular cardinals. Let A and B be structures of the same similarity type, with A λ-complete. Then the following are equivalent:*

(i) *every sentence of* $PL_{\kappa\lambda}$ *which is true in A is true also in B*;

(ii) *B is isomorphic to a λ-limit κ-ultrapower of A.*

Proof. We proceed as in the proof of Theorem 7a up to the choice of the function $\theta : \{f_\alpha : \alpha < 2^\mu\} \to \operatorname{dom} B$. Then we take D to be the set of all sets of form

$$A(\varphi) = \{i \in \mu : A \models \varphi[f_\alpha(i)]_{\alpha<2^\mu}\}$$

such that φ is quantifier-free and has fewer than λ variables, and $B \models \varphi[\theta f_\alpha]_{\alpha<2^\mu}$. Let π be an element of F; then there is a set $J \subseteq 2^\mu$ of cardinality $<\lambda$ such that $\pi \geq \bigwedge_{\alpha\in J} [f_\alpha]$. Since A is λ-complete it follows that every set in B_π is of form $A(\varphi)$ for some quantifier-free formula φ with variables from v_α $(\alpha \in J)$. Hence D is a (κ, F)-ultrafilter. The proof that $A^\mu_D \mid F \cong B$ is as before.

3.3. *Isomorphic limit reduced products*

Now we shall prove analogues of Theorem 4 for limit κ-reduced products. The conditions for isomorphism which were proved necessary in Theorem 4 are now both necessary and sufficient, and we do not have to import any peculiar set-theoretic assumptions; ZFC alone is enough. (The analogue of Theorem 4c is a little more restricted.)

Theorem 8a. *Let A and B be structures of the same similarity type. Then the following are equivalent:*

(i) *A and B have isomorphic limit κ-reduced powers*:

(ii) *if T_A, T_B are respectively the sets of* $PH_{\kappa\kappa}$*-sentences true in A, B, then $T_A \cup T_B$ is consistent.*

Proof. (i)$\Rightarrow$(ii) is by Lemma 1. Now we assume (ii) and prove (i). Choose a cardinal μ such that $\mu=\mu^{<\kappa}\geqslant\text{card}(A)+\text{card}(B)$. Referring back to the beginning of Section 2.1, we put $I=J=\mu$, $A_i=A$ and $B_j=B$ for each $i,j<\mu$. The cardinal μ of 2.1 now becomes 2^μ. Let A^*, B^* be respectively (μ,κ)-independent limit powers of A, B; then A^*, B^* are substructures of A^μ, B^μ respectively. By Lemma 3a, (i) above is proved provided we can find $\bar{a}_\gamma(\gamma<2^\mu)$ in A^μ and $\bar{b}_\gamma$ $(\gamma<2^\mu)$ in B^μ such that

(iii) for every pair $\langle\theta,\eta\rangle$ in Δ (from Lemma 3a), either there is $i<\mu$ such that $A\models\neg\theta[\bar{a}^i]$ or there is $j<\mu$ such that $B\models\neg\eta[\bar{b}^j]$; and $\langle\bar{a}_\gamma:\gamma<2^\mu\rangle$, $\langle\bar{b}_\gamma:\gamma<2^\mu\rangle$ list the whole of A^*, B^* respectively.

We shall make players $\forall$ and $\exists$ choose the $\bar{a}_\gamma$ in sequence from A^*. Player $\forall$ chooses $\bar{a}_\gamma$ when γ is even, and he uses the independent strategy (cf. before Lemma 6). Player $\exists$ chooses at odd γ, and he plays so as to exhaust A^*. Likewise the players choose the $\bar{b}_\gamma$ from B^*; player $\forall$ uses the independent strategy at odd γ and player $\exists$ chooses so as to exhaust B^* at even γ. Player $\exists$'s choices guarantee the last part of (iii).

Now suppose the quantifier-free formula θ of $\text{PH}_{\kappa\kappa}$ occurs in some pair $\langle\theta,\eta\rangle$ in the set Δ of Lemma 3a. Let Q be the quantifier whose domain is the set of indices of variables which occur in θ, such that $\mathsf{Q}(\gamma)=\forall$ iff γ is even. Write θ' for $\mathsf{Q}\theta$. Similarly for each η choose a prefix Q by the same rules but with $\mathsf{Q}(\gamma)=\forall$ iff γ is odd. Write Δ' for the set of pairs $\langle\theta',\eta'\rangle$ such that $\langle\theta,\eta\rangle$ is in Δ. The following claim is proved exactly like the claim in the proof of Theorem 4a.

Claim. *For each pair $\langle\theta',\eta'\rangle$ from Δ', the sentences θ' and η' are mutually contradictory.*

Now we can prove (iii). Let $\langle\theta,\eta\rangle$ be a pair from Δ. By the claim and (ii) of the theorem, player $\exists$ does not have winning strategies for both $G(\theta',A)$ and $G(\eta',B)$; suppose he lacks one for $G(\theta',A)$. Then by Lemma 6 there are a set $Z\subseteq\mu$ and a strategy σ for player $\exists$ in $G(\theta',A)$ such that player $\exists$ is playing σ in $G(\theta',A)$ at each coordinate $i\in Z$, and each possible play of player $\forall$ in $G(\theta',A)$ occurs at some $i\in Z$. Since σ is not winning for $\exists$, player $\forall$ must win at some i, and hence $A\models\neg\theta[\bar{a}^i]$. Thus (iii) is proved.

The corresponding theorem with classes of structures is proved analogously, using products of (μ,κ)-independent limit powers:

Theorem 8b. *Let H and K be classes of structures, all of the same similarity type. Then the following are equivalent:*

(i) *some limit κ-reduced product of structures in H is isomorphic to a limit κ-reduced product of structures in K;*

(ii) *if T_H, T_K are respectively the sets of $\text{PH}_{\kappa\kappa}$-sentences true throughout H, K, then $T_H\cup T_K$ is consistent.*

We leave it to the reader to supply the remaining variants of Theorem 8. For the counterpart of Theorem 4c (where the filter D has to be the same on both sides) it seems to be necessary to assume that A and B have the same cardinality.

4. Applications

4.1. *Interpolation and definability for Horn logic*

From Lemma 2 (compactness) and Theorem 8b the usual argument gives:

Theorem 9 (Interpolation theorem). *Let φ, ψ be sentences of $\mathrm{PH}_{\kappa\kappa}$ such that $\varphi, \psi \vdash \bot$. Then there is a sentence θ of $\mathrm{PH}_{\kappa\kappa}$ containing only relation and function symbols which occur in both φ and ψ, such that $\varphi \vdash \theta$ and $\theta, \psi \vdash \bot$.*

Then once more the usual argument gives:

Corollary 10 (Beth definability). *Let T be a theory in $\mathrm{PH}_{\kappa\kappa}$, R a relation or function symbol occurring in T, and L a sublanguage of the language of T in which R does not occur. Suppose that if A, B are any two models of T for which $A \upharpoonright L = B \upharpoonright L$, then $R_A = R_B$. Then T entails an explicit definition $\forall \bar{v}(\varphi \leftrightarrow R\bar{v})$, where φ is a formula of $\mathrm{PH}_{\kappa\kappa}$ using only symbols from L.*

Isbell [8] raised the question whether there is a Beth definability theorem for equational theories in $\mathrm{PH}_{\kappa\kappa}$, and proved such a theorem assuming an extra hypothesis about functoriality. (See Hodges [7] for a model-theoretic result which generalises Isbell's.) Corollary 10 shows that there is a reasonable definability theorem without the functoriality condition. Under Isbell's assumption, θ is existential. In general, can we bound the number of quantifier alternations in θ when T is equational? Friedman [4] showed that already when $\kappa = \omega$ there is no finite bound. The example below was as near as we could get to showing that for uncountable κ there is no bound $<\kappa$.

Example 11. Sentences φ, ψ of $L_{\omega_1\omega_1}$ which are conjunctions of equational theories, such that $\varphi, \psi \vdash \bot$ but there is no interpolant (as in Theorem 9) in $L_{\infty\omega_1}^{<\omega_1}$.

Let T_L be a set of axioms for the variety of lattices with top and bottom elements 1, 0; we write $a \leqslant b$ for $a \wedge b = a$. Let $\mathrm{WF}_L(f)$ be the set of universal closures of the equations:

$$f(x, y) \wedge f(y, z) \leqslant f(x, z),$$

$$f(x, x) = 0,$$

$$s(1, 1, 1, \ldots) = 0 \quad (s \text{ is of arity } \omega),$$

$$s(f(x_0, x_1), f(x_1, x_2), f(x_2, x_3), \ldots) = 1.$$

Then $T_L \cup \mathrm{WF}_L(f)$ expresses that the relation $f(y, x) = 1$ is a well-founded irreflexive partial ordering. Let $\mathrm{Is}_{L,g}(f, c, d)$ consist of the identities

$$\forall x\, \forall y\, f(c, y) \wedge f(y, x) \leqslant f(g(c), g(y)) \wedge f(g(y), g(x)), \quad f(d, g(c)) = 1.$$

Let $\mathrm{Init}(x)$ be the set of all elements y such that $f(x, y) = 1$. Let φ be the conjunction of T_L, $\mathrm{WF}_L(f)$ and $\mathrm{Is}_{L,g}(f, c, d)$. Then φ expresses that g is an isomorphism from $\mathrm{Init}(c)$ to a proper initial part of $\mathrm{Init}(d)$, and hence that c has lower rank than d in the partial ordering $f(y, x) = 1$. Let $T_{L'}$, etc. be as T_L etc., but with the lattice operations except 1 replaced by new symbols. Let ψ be the conjunction of $T_{L'}$, $\mathrm{WF}_{L'}(f)$ and $\mathrm{Is}_{L',h}(f, d, c)$. Then clearly $\varphi, \psi \vdash \perp$.

Suppose now that θ is a sentence of $L^{<\omega_1}_{\infty\omega_1}$ in the language with symbols f, 1, c, d, and $\varphi \vdash \theta$. The following argument shows that θ is consistent with ψ, and hence θ is not an interpolant between φ and ψ. For any two transfinite ordinals α and β we can construct a model A of φ by taking the disjoint union of α and β, and choosing $c \in \alpha$ and $c < d \in \beta$. Then $A \vDash \theta$, since $\varphi \vdash \theta$. But a result of Chang [1] shows that the $L^{<\omega_1}_{\infty\omega_1}$ theory of the ordinal λ^λ is the same for all uncountable cardinals λ of cofinality $>\omega$, and the Feferman–Vaught theorem for disjoint unions of structures holds for $L^{<\omega_1}_{\infty\omega_1}$. Hence θ also has a model of the same form as A but with c of higher rank than d.

One of the unnumbered variants of Theorem 8 says that if K is a class of structures and B is a structure, then there is a surjective homomorphism from a limit κ-reduced product of structures in K to a limit κ-reduced power of B if and only if the $\mathrm{PH}_{\kappa\kappa}$ theory of B is consistent with the set of positive $\mathrm{PH}_{\kappa\kappa}$ sentences true throughout K. (Cf. Lemma 3c.) From this it is easy to deduce:

Theorem 12. *Let T be a theory in $\mathrm{PH}_{\kappa\kappa}$ and φ a sentence of $\mathrm{PH}_{\kappa\kappa}$ which is preserved in surjective homomorphisms between models of T, and suppose that $\neg\varphi$ is preserved in limit κ-reduced powers of models of T. Then φ is equivalent in T to a positive sentence of $\mathrm{PH}_{\kappa\kappa}$.*

We also have the analogue of Theorem 12 when κ is strongly compact, $\mathrm{PH}_{\kappa\kappa}$ is replaced by $\mathrm{PL}_{\kappa\kappa}$ and limit κ-reduced powers become limit κ-ultrapowers. Note that every sentence of $L_{\kappa\kappa}$ is preserved in limit κ-ultrapowers. The case $\kappa = \omega$ reduces to a well-known theorem of Lyndon. Curiously there do not seem to be any known counterexamples to Lyndon's theorem for any interesting infinitary language.

4.2. Sentences preserved in reduced products

When $\kappa = \omega$, a theorem of Keisler [12] says that a sentence of $L_{\kappa\kappa}$ is preserved in reduced products iff it is logically equivalent to a Horn sentence. (Cf. Chang and Keisler [2, Theorem 6.2.5]; Galvin eliminated Keisler's use of the continuum hypothesis.) The next two theorems partially generalise Keisler's theorem.

Theorem 13. (a) *Let φ be a sentence preserved in limit κ-reduced products, such that $\neg\varphi$ is preserved in limit κ-reduced powers. Then φ is logically equivalent to a set of sentences of* $\mathrm{PH}_{\kappa\kappa}$.

(b) *Suppose κ is strongly compact, and let φ be a sentence which is preserved in limit κ-reduced products, such that $\neg\varphi$ is preserved in limit κ-ultrapowers. Then φ is equivalent to a set of sentences of* $\mathrm{PH}_{\kappa\kappa}$.

(c) *Suppose κ is strongly compact, and let φ be a sentence of $L_{\kappa\kappa}$ which is preserved in limit κ-reduced products. Then φ is logically equivalent to a sentence of* $\mathrm{PH}_{\kappa\kappa}$.

Proof. (a) and (b) follow straightforwardly from Theorem 8. For (c) we use (b) and then compactness (Lemma 2) to reduce to a single sentence.

We conjecture that in case (c) of Theorem 13, the logically equivalent sentence cannot in general be chosen in $\mathrm{PH}_{\kappa\kappa} \cap L_{\kappa\kappa}$.

Theorem 14. *Let D be a proper κ-complete non-κ-saturated filter on the set I. If φ is a quantifier-free sentence of $L_{\kappa\kappa}$ which is preserved in reduced products modulo D, then φ is equivalent to a conjunction of (possibly $\geqslant\kappa$) quantifier-free Horn sentences of $L_{\kappa\kappa}$.*

Proof. By the hypothesis on D, there are pairwise disjoint subsets X_i $(i<\kappa)$ of I such that no X_i is zero (mod D).

Bring φ to conjunctive normal form in $L_{\kappa\kappa}$. Then φ is equivalent to the conjunction of a set of sentences χ of form

$$\bigwedge \Phi \rightarrow \bigvee \Psi$$

where Φ, Ψ are sets of $<\kappa$ atomic sentences. It suffices to show that for each such χ there is some $\psi \in \Psi$ such that φ entails $\bigwedge \Phi \rightarrow \psi$. Suppose χ is a counterexample to this, so that for each $\psi \in \Psi$ there is a model A_ψ of $\varphi \wedge \bigwedge \Phi \wedge \neg\psi$. Let $f: \Psi \rightarrow \kappa$ be an injection, and choose structures A_i $(i \in I)$ in such a way that $A_j = A_\psi$ whenever $j \in X_{f(\psi)}$. Our assumption on φ implies that $\prod_D A_i$ is a model of φ, and clearly it is also a model of $\bigwedge \Phi$. Now for each $\psi \in \Psi$, $\{i \in I\colon A_i \models \psi\} \subseteq I - X_{f(\psi)}$, so that $\{i \in I\colon A_i \models \psi\} \notin D$ and hence $\prod_D A_i \models \neg\psi$. Hence $\varphi, \bigwedge \Phi \not\models \psi$, a contradiction.

Proper κ-complete non-κ-saturated filters always exist: take the filter $\{\kappa\}$ on κ. When κ is strongly inaccessible, the proof of Theorem 14 does actually give a Horn sentence of $L_{\kappa\kappa}$ which is equivalent to φ. Example 15 will show that this is not always possible. Examples 15–17 illustrate three different failures of infinitary analogues of Keisler's theorem on sentences preserved in reduced products.

Example 15. A quantifier-free sentence of $L_{\kappa\kappa}$ which is preserved in κ-reduced products but is not equivalent to any sentence of $L_{\kappa\kappa}$, when $\kappa = \mu^+$ and $\mu^{<\mu} = \mu$.

For each $i<\mu$ let P_i be a distinct propositional letter. Let φ be the sentence of $L_{\kappa\kappa}$ which says:

Either fewer than μ of the P_i are true, or all of them are.

Then φ is preserved in all κ-reduced products, since it is equivalent to the conjunction of the set T of all sentences

$$\bigwedge \Phi \to P_j \quad (j<\mu, \Phi \subseteq \{P_i : i<\mu\}, \text{card}(\Phi)=\mu).$$

Every quantifier-free Horn consequence of φ in $L_{\kappa\kappa}$ is equivalent to a conjunction of fewer than κ sentences of T. But suppose $S \subseteq T$, card$(S) \leqslant \mu$. Then there is a proper subset X of $\{P_i : i<\mu\}$, with cardinality μ, such that for each formula $\bigwedge \Phi \to P_j$ in S, $\Phi \subseteq X$. (List the formulae of S in order-type μ, and for each formula $\bigwedge \Phi \to P_j$ in turn, put one element of Φ inside X and one outside.) Choose A so that $A \models P_i$ iff $P_i \in X$; then $A \models S$ but $A \models \neg\varphi$. Hence φ is not equivalent to any quantifier-free Horn sentence of $L_{\kappa\kappa}$.

Example 16. A sentence of $L_{\kappa\kappa}$ which is preserved in κ-reduced products but is not equivalent to any set of sentences of $\mathrm{PH}_{\kappa\kappa}$, when $\kappa=\mu^+$.

Let φ_0 be the Horn sentence of $L_{\kappa\kappa}$ which defines the class of κ-complete boolean algebras. In any boolean algebra B let $I_\mu(B)$ be the ideal generated by all sums of at most μ atoms of B, and let φ_1 be the statement which holds in B iff:

Either $B=I_\mu(B)$ or $B/I_\mu(B)$ is infinite.

φ_1 can be written as a sentence of $L_{\kappa\kappa}$. Our example φ is $\varphi_0 \wedge \varphi_1$.

We show that φ is preserved in κ-reduced products. Suppose that B_i is a model of φ for each $i \in I$, and D is a κ-complete filter on I. Write B for $\prod_D B_i$. We evidently have $B \models \varphi_0$ since φ_0 is Horn. There are now four cases to consider.

Case i: $\mathcal{P}(I)/D$ is not atomic. Then B is not atomic, and $B/I_\mu(B)$ is infinite.

Case ii: $\mathcal{P}(I)/D$ is atomic but contains some atom $X \subseteq I$ such that $B_i/I_\mu(B_i)$ is infinite for each $i \in X$. Write D' for the restriction of D to X; then (cf. Chang and Keisler [2, Proposition 6.2.1]) B is a product with a factor $B'=\prod_{D'} B_i$. Since X was an atom, D' is a κ-complete ultrafilter, and it follows by Łoś's theorem that $B'/I_\mu(B')$ is infinite. Hence the same holds also for B.

Case iii: $\mathcal{P}(I)/D$ is atomic with at most μ atoms, and there is no atom throughout which $B_i/I_\mu(B_i)$ is infinite. Then again we may write B as a product of at most μ algebras of form $B'=\prod_{D'} B_i$ where D' is a κ-complete ultrafilter on a subset I' of I and $B_i=I_\mu(B_i)$ for all $i \in I'$. Again by Łoś's theorem $B'=I_\mu(B')$ for each factor B', and hence the same holds for the product B.

Case iv: As Case iii, but $\mathcal{P}(I)/D$ is atomic with more than μ atoms. Then we can pick out κ atoms represented by sets $X_j \subseteq I$ $(j<\kappa)$ which are pairwise disjoint, and such that $B_i=I_\mu(B_i)$ for all $i \in \bigcup_{j<\kappa} X_j$. Partition κ into sets J_α $(\alpha<\kappa)$, each of cardinality κ. For each $i \in \bigcup_{j<\kappa} X_j$, choose an atom x_i of B. Let b_α be the element of $\prod_I B_i$ such that $b_\alpha(i) = x_i$ when $i \in \bigcup \{X_j : J \in J_\alpha\}$, and $b_\alpha(i) = 0$ otherwise. Then each b_α/D has μ^+ distinct atoms below it, and $(b_\alpha/D) \wedge (b_\beta/D) = 0$ whenever $\alpha \neq \beta$. Hence $B/I_\mu(B)$ has cardinality at least κ.

Hence φ is preserved in κ-reduced products. To prove that φ is not equivalent to any theory in $\mathrm{PH}_{\kappa\kappa}$, it will be enough (by Lemma 1) to construct a limit κ-reduced power which fails to preserve κ.

Let **2** be the two-element boolean algebra and let $\mathbf{2}^*$ be the expansion of **2** with all possible functions of arity $<\kappa$. By Theorem 7a, every reduct to the language of boolean algebras of a subalgebra of a κ-reduced power of $\mathbf{2}^*$ is a limit κ-reduced power of **2**. This makes it easy to construct limit κ-reduced powers of **2**.

Observe that **2** is a model of φ. Now partition κ into sets J, K of cardinality κ, and let B^* be $(\mathbf{2}^*)^\kappa$, and x the element of B^* such that $x(i)$ is 1 iff $i \in J$. Let B be the reduct to the language of boolean algebras of the subalgebra of B^* generated by the atoms and x. Then we have just seen that B is a limit κ-reduced power of **2**. Also B is the product of two copies of the subalgebra of **2** generated by all sums of $\leq\mu$ atoms. Hence $B/I_\mu(B)$ is the four-element algebra, and so φ_1 fails in B.

Example 17. A first-order sentence which is not equivalent to a sentence in any $\mathrm{PH}_{\lambda\lambda}$, but is preserved in all κ-reduced products with $\kappa > \omega$, provided there is no measurable cardinal.

For this we take the sentence φ which says: The structure is a boolean algebra and there is a maximal atomless element. We remark that Mansfield [15] used this sentence as an example of a non-Horn sentence which is preserved in direct products and what he calls normal submodels. We shall need the fact that φ is preserved in products.

First we show that φ is preserved in κ-reduced products provided there is no measurable cardinal and $\kappa > \omega$. Let each B_i $(i \in I)$ be a model of φ with maximal atomless element b_i, and let D be a κ-complete filter on I. Write $B = \prod_D B_i$. If Z is an atom of $\mathcal{P}(I)/D$, then the restriction of D to Z is a κ-complete ultrafilter on Z, and this ultrafilter must be principal since there are no measurable cardinals. Hence the atoms of $\mathcal{P}(I)/D$ are represented by singletons, and we can partition I into $X \cup Y$ where X is a union of singleton atoms and Y is atomless. Write D' for the restriction of D to Y, and B' for $\prod_{D'} B_i$. Then B' is atomless and B is the product of B' and all the B_i with $i \in X$. Since all these algebras satisfy φ, so does B.

It remains to show that φ is not equivalent to any set of sentences of any $\mathrm{PH}_{\kappa\kappa}$. We use the approach and notation of Example 16. Let J, K be disjoint sets of cardinality κ^+, and put $I = J \cup K$. Write D for the κ-complete filter on I consisting of those subsets X of I such that $I - X$ consists of fewer than κ elements of K. Let B^* be the subalgebra of $\prod_D \mathbf{2}^*$ consisting of those elements b/D such that b is constant on all but $\leq\kappa$ elements of I, and let B be the reduct of B^*

to the language of boolean algebras. The atomless elements of B are those which are zero throughout J, and among these elements none is maximal. Hence B is not a model of φ. But **2** is a model of φ, and as in Example 16, B is a limit κ-reduced power of **2**.

Kueker [14] introduces the closed unbounded filter D on $\mathcal{P}_{\omega_1}(\kappa)$, and defines countable approximations φ^s for sentences φ of $L_{\kappa^+\omega}$ and sets $s \in \mathcal{P}_{\omega_1}(\kappa)$. He shows (Theorem 4.6) that if φ is Horn in $L_{\kappa^+\omega}$, and structures A_s are given so that $\{s \in \mathcal{P}_{\omega_1}(\kappa): A_s \models \varphi^s\} \in D$, then $\prod_D A_s \models \varphi$. He asks whether every sentence of $L_{\kappa^+\omega}$ with this property is Horn up to logical equivalence. (In our notation, his Horn sentences are those sentences of $L_{\kappa^+\omega}$ which are in $\mathrm{PH}_{\kappa\kappa}$ when they are made prenex.) If $\kappa = \omega$, then D is principal and all φ^s can be taken to be equal to φ; so the question is only interesting when $\kappa > \omega$.

Now by Jech [9, Theorem 3.4], the filter D is atomless. It follows that if φ is as in Example 17, reduced products over D preserve φ regardless of whether there are measurable cardinals. Hence φ gives a negative answer to Kueker's question. If $\kappa^{<\kappa} = \kappa$, then the quantifier-free sentence of Example 15 gives another negative answer.

4.3. *Elementary extensions satisfying given sentences*

We investigate the following problem, which arises because neither $\mathrm{PL}_{\kappa\kappa}$ nor $\mathrm{PH}_{\kappa\kappa}$ is closed under negation. Write $A \preccurlyeq_L B$ to mean that A is a substructure of B and every formula in L which is true in A of elements of A is true in B too. For a given sentence $\mathrm{Q}\varphi$ of $\mathrm{PL}_{\kappa\kappa}$ ($\mathrm{PH}_{\kappa\kappa}$) and structure A, when is there B such that

$$A \preccurlyeq_{\mathrm{PL}_{\kappa\kappa}(\mathrm{PH}_{\kappa\kappa})} B \quad \text{and} \quad B \models \mathrm{Q}\varphi?$$

When L is fixed, let us write $A \models \Diamond \mathrm{Q}\varphi$ to mean that there is B such that $A \preccurlyeq_L B$ and $B \models \mathrm{Q}\varphi$. We write $\bar{\mathrm{Q}}$ for the dual of Q, got by replacing $\forall$ and $\exists$ and *vice versa* throughout. For any structure A and sentence $\mathrm{Q}\varphi$, consider these six possibilities:

(1) $A \models \mathrm{Q}\varphi$.

(2) $A \models \bar{\mathrm{Q}} \neg \varphi$.

(3) $A \models \Diamond \mathrm{Q}\varphi$ and $A \models \Diamond \bar{\mathrm{Q}} \neg \varphi$.

(4) Not (1), but $A \models \Diamond \mathrm{Q}\varphi$ and $A \models \neg \Diamond \bar{\mathrm{Q}} \neg \varphi$.

(5) Not (2), but $A \models \Diamond \bar{\mathrm{Q}} \neg \varphi$ and $A \models \neg \Diamond \mathrm{Q}\varphi$.

(6) $A \models \neg \Diamond \mathrm{Q}\varphi$ and $A \models \neg \Diamond \bar{\mathrm{Q}} \neg \varphi$.

Theorem 18. *Suppose φ above is required to be atomic, and L is $\mathrm{PL}_{\kappa\kappa}$ or $\mathrm{PH}_{\kappa\kappa}$. Then (1)–(6) are mutually exclusive. If κ is strongly compact and L is $\mathrm{PL}_{\kappa\kappa}$, then all of (1)–(6) do occur. If L is $\mathrm{PH}_{\kappa\kappa}$ and $\kappa > \omega$, then only possibilities (1), (2), (3), (5) occur.*

Proof. Lemma 1 shows that (1)–(6) are mutually exclusive. If L *is* $\mathit{PH}_{\kappa\kappa}$, we show as follows that $\mathrm{Q}\varphi$ is equivalent to $\neg \Diamond \bar{\mathrm{Q}} \neg \varphi$ (which eliminates (4) and (6)). $\mathrm{Q}\varphi$

entails $\neg\Diamond\tilde{\mathsf{Q}}\neg\varphi$ by Lemma 1. Conversely, suppose $A \models \neg\mathsf{Q}\varphi$; then player $\exists$ has no winning strategy for $G(\mathsf{Q}\varphi, A)$. For some $\mu \geqslant \kappa$, let $A^\mu \mid F$ be a (μ,κ)-independent limit power of A, and let players $\forall$ and $\exists$ play $G(\mathsf{Q}\varphi, A^\mu \mid F)$ with player $\forall$ using his independent strategy. (Cf. Section 3.1.) Then by Lemma 6, player $\forall$ wins on at least one coordinate. But since φ is atomic, this means that player $\forall$ wins on $A^\mu \mid F$. Hence A has an elementary extension $A^\mu \mid F$ which satisfies $\tilde{\mathsf{Q}}\neg\varphi$, and so $A \models \Diamond\tilde{\mathsf{Q}}\neg\varphi$.

It remains to construct situations in which (1)–(6) do occur. We shall treat the case where κ is strongly compact and L is $\mathrm{PL}_{\kappa\kappa}$, and leave the case of $\mathrm{PH}_{\kappa\kappa}$ to the reader. Note that ω is not strongly compact. (Our results for strongly compact κ in earlier sections did not use the assumption that $\kappa > \omega$.)

Consider a quantifier Q of length ξ and a cardinal $\nu > 0$. By a *ν-shuffle* of Q we mean an ordinal α together with a family $(x_i : i < \nu)$ of maps $x_i : \xi \to \alpha$ such that (i) α is the union of the images of the x_i, (ii) each x_i is order-preserving, and (iii) if $x_i(\beta) = x_j(\beta')$, then $\beta = \beta'$ and $x_i(\gamma) = x_j(\gamma)$ for all $\gamma < \beta$. If \$ is this ν-shuffle, then $\mathsf{Q}^\$$ is defined to be the quantifier of length α such that $\mathsf{Q}^\$(x_i(\beta)) = \mathsf{Q}(\beta)$ for each $i < \nu$ and $\beta < \xi$. If $\langle a_\beta : \beta < \alpha\rangle$ is a play of the pre-game $G(\mathsf{Q}^\$, A)$, then the ith *thread* $\bar{a}^i$ of this play is defined to be $\langle a_{x_i(\beta)} : \beta < \xi\rangle$, for each $i < \nu$.

Let G be the game $G(\mathsf{Q}\varphi, A)$. Then we define a game $G^\$ = G^\$(\mathsf{Q}\varphi, A)$: $G^\$$ is played as $G(\mathsf{Q}^\$, A)$, and player $\forall$ wins iff for at least one $i < \nu$, $A \models \neg\varphi[\bar{a}^i]$. We call $G^\$$ a *κ-derived* game of G iff $\nu < \kappa$.

The statement that player $\forall$ has a winning strategy for $G^\$(\mathsf{Q}\varphi, A)$ can be written in the form: Player $\exists$ has a winning strategy for $G(\mathsf{Q}'\varphi', A)$ where Q' is a certain quantifier of length α and φ' is a certain disjunction of instances of $\neg\varphi$.

Lemma (κ strongly compact). *Let* $\mathsf{Q}\varphi$ *be a sentence of* $\mathrm{PL}_{\kappa\kappa}$. *Then for any structure* A *the following are equivalent*:

(i) $A \models \Diamond\mathsf{Q}\varphi$.

(ii) *player* $\forall$ *has no winning strategy for any κ-derived game of* $G(\mathsf{Q}\varphi, A)$.

Proof of lemma. By Lemma 1 and the remark before the present lemma, if player $\forall$ has winning strategies for all κ-derived games of $G(\mathsf{Q}\varphi, A)$ and B is an elementary extension of A, then player $\forall$ has winning strategies for all κ-derived games of $G(\mathsf{Q}\varphi, B)$, and so $B \models \neg\mathsf{Q}\varphi$. This proves the implication (i) $\Rightarrow$ (ii).

For the converse, assume (ii). Choose $\mu \geqslant \mathrm{card}(A)$ such that $\mu^{<\kappa} = \mu$, and let $A^\mu \mid F$ be the (μ,κ)-independent limit power of A generated by $(f_i : i < 2^\mu)$. The number of sequences of length $<\kappa$ of elements of $A^\mu \mid F$ is at most 2^μ; we partition 2^μ into disjoint cofinal sets $X_{\bar{a}}$ indexed by such sequences $\bar{a}$. Let σ be the strategy for player $\exists$ in $G(\mathsf{Q}\varphi, A^\mu \mid F)$ which is the same as his independent strategy defined in Section 3.1, except that for 'the first f_i' we read 'the first f_i in $X_{\bar{a}}$ (where $\bar{a}$ is the sequence of moves played so far)'. Thus the whole preceding play can be inferred from each move of player $\exists$.

If player $\forall$ plays strategy τ against player $\exists$'s σ, the resulti g play of

$G(Q\varphi, A^\mu \mid F)$ is a sequence $\{\bar{a}_j : j < \xi\}$; write

$$A(\tau) = \{i < \mu : A \models \varphi[\bar{a}_j(i)]_{j<\xi}\}.$$

Let D be the filter on μ generated by all intersections of fewer than κ sets $A(\tau)$ where τ ranges over the possible strategies of player $\forall$ in $G(Q\varphi, A^\mu \mid F)$.

Clearly D is κ-complete. We claim that D is proper. For this it suffices to show that if $\{\tau_\gamma : \gamma < \nu\}$ is a set of $<\kappa$ strategies for player $\forall$ in $G(Q\varphi, A^\mu \mid F)$, then $\bigcap_{\gamma<\nu} A(\tau_\gamma)$ is not empty.

Write $\langle \bar{a}_{\gamma j} : j < \xi\rangle$ for the play when player $\forall$ plays τ_γ against σ. The choice of σ implies that there is a ν-shuffle \$ of Q with ordinal α and maps $(x_\gamma : \gamma < \nu)$, such that if $x_\gamma(h) < x_\delta(j)$ and $Q(j) = \exists$, then $\bar{a}_{\delta j} \notin \operatorname{supp}(\bar{a}_{\gamma h})$. Let $\bar{b}$ be the sequence of length α whose $x_\gamma(h)$th element is $\bar{a}_{\gamma h}$. After renumbering the f_i to match the order in which they appear in $\bar{b}$, $\bar{b}$ becomes a play of $G^\$(Q\varphi, A^\mu \mid F)$ in which player $\exists$ uses the independent strategy. By Lemma 6 with $X = \alpha$, there is $Z \subseteq \mu$ on which player $\forall$ uses a fixed strategy. Now $G^\$(Q\varphi, A)$ is a κ-derived game of $G(Q\varphi, A)$, so by assumption player $\forall$ has no winning strategy for it. By Lemma 6 again, player $\exists$ plays in every possible way against player $\forall$'s fixed strategy, so that player $\exists$ wins at some coordinate i. On i, each $\langle \bar{b}(x_\gamma(j))(i) : j < \xi\rangle = \langle \bar{a}_{\gamma j}(i) : j < \xi\rangle$ is winning for $\exists$ in $G(Q\varphi, A)$. In short, $i \in \bigcap_{\gamma<\nu} A(\tau_\gamma)$. The claim is proved.

Hence D can be extended to a κ-ultrafilter D' on μ. Then $A \preccurlyeq_{PL_{\kappa\kappa}} A^\mu_{D'} \mid F$. Since each $A(\tau)$ is in D', $A^\mu_{D'} \mid F \models Q\varphi$. This proves (i).

Now we return to the theorem. Examples of (1) and (2) are no trouble to find. We shall construct an example of (3). $Q\varphi$ will be of form

$$\forall v_0 \exists v_1 \forall v_2 \exists v_3 \cdots R(v_0 v_1 \cdots)$$

where Q has length ω and R is an ω-ary relation symbol. The structure A will be of form (λ, R_A) where $\lambda \geq \kappa$ and $2^\lambda = \lambda^\omega$. (For example, λ is the first strong limit number $>\kappa$.) By the choice of λ, we can list as $\langle \tau_i, G_i\rangle$ $(i < \lambda^\omega)$ all the pairs such that G_i is a κ-derived game of $G(Q\varphi, A)$ or $G(\tilde{Q}\neg\varphi, A)$, and τ_i is a strategy for player $\forall$ in G_i. (We know how these games are played, but since R_A is not yet defined we do not know what counts as winning them.) For each $i < \lambda^\omega$, let σ_i be the following strategy for player $\exists$ in G_i, defined by induction on i: at each thread, play a sequence from ${}^\omega\lambda$ which is distinct from the sequences played by either player at any thread of the game when player $\exists$ plays σ_j against player $\forall$'s τ_j, for all $j < i$.

Now by induction on i, we can ensure that σ_i always wins against τ_i, by putting some sequences from ${}^\omega\lambda$ into R_A (when G_i is a κ-derived game of $G(Q\varphi, A)$) or excluding some sequences from R_A (otherwise). The definition of the σ_i ensures that the sequences to be put in at one stage are all different from those to be excluded at another. R_A is otherwise chosen arbitrarily. Then no τ_i is a winning strategy for player $\forall$ in G_i. Hence (3) holds by the lemma.

Sh:109

Next we construct an example of (4). The sentence $\mathsf{Q}\varphi$ and the structure A will have the same form as for (3); but this time we must construct R_A so that (i) player $\exists$ has no winning strategy in $G(\mathsf{Q}\varphi, A)$, (ii) player $\forall$ has no winning strategy in any κ-derived game of $G(\mathsf{Q}\varphi, A)$, and (iii) player $\forall$ has a winning strategy for some κ-derived game of $G(\bar{\mathsf{Q}}\neg\varphi, A)$.

List without repetition all finite sequences of elements of λ as c_α $(\alpha<\lambda)$. Let \$ be the 2-shuffle of Q with ordinal ω, such that for each n, $x_0(n)=2n$ and $x_1(n)=2n+1$. We shall ensure that player $\forall$ wins $G^{\$}(\mathsf{Q}\neg\varphi, A)$ if he plays the following strategy ρ: when $a_0,\ldots,a_n$ is the play so far, $\forall$ shall play α where $\langle a_0,\ldots,a_n\rangle=c_\alpha$. Thus each move of player $\forall$ records the entire previous history of the play. List as Y_i $(i<\lambda^\omega)$ all pairs $\{\bar{a}^0,\bar{a}^1\}$ where $\bar{a}^0$, $\bar{a}^1$ are respectively the 0th and 1th thread of a play of $G^{\$}(\bar{\mathsf{Q}}\neg\varphi, A)$ in which player $\forall$ uses ρ. By choice of ρ, the Y_i are pairwise disjoint and all of cardinality 2. To ensure that ρ wins $G^{\$}(\bar{\mathsf{Q}}\neg\varphi, A)$ for player $\forall$, it suffices that for each $i<\lambda^\omega$, $Y_i\cap R_A$ is not empty.

Let $\langle\tau_i, G_i\rangle$ $(i<\lambda^\omega)$ list all pairs such that either G_i is $G(\mathsf{Q}\varphi, A)$ and τ_i is a strategy for player $\exists$ in G_i, or G_i is a κ-derived game of $G(\mathsf{Q}\varphi, A)$ and τ_i is a strategy for player $\forall$ in G_i. For each $i<\lambda^\omega$, we shall define sets M_i, $N_i\subseteq{}^\omega\lambda$, both of cardinality $<\kappa$, and a strategy σ_i for the player opposed to τ_i in G_i. The definition is by induction on i, as follows.

Case 1: G_i is a derived game of $G(\mathsf{Q}\varphi, A)$. Then σ_i shall be the following strategy for player $\exists$ in G_i: at each thread, play a sequence $\in{}^\omega\lambda$ which is distinct from every sequence played at any $j<i$ when σ_j is played against τ_j. Player $\exists$ wins G_i by playing σ_i against τ_i iff a certain subset M of ${}^\omega\lambda$ is in R_A; put $M_i=M$, $N_i=\emptyset$. M_i has cardinality $<\kappa<\lambda^\omega$.

Case 2: G_i is $G(\mathsf{Q}\varphi, A)$. Then σ_i shall be the following strategy for player $\forall$ in G_i: play a sequence $\in{}^\omega\lambda$ which is distinct from all sequences that are either in $\bigcup_{j<i}M_j$ or in any Y_k $(k<\lambda^\omega)$ such that $Y_k\cap(\bigcup_{j<i}N_j)\neq\emptyset$. Player $\forall$ wins $G(\mathsf{Q}\varphi, A)$ by playing σ_i against τ_i iff the resulting play $\bar{a}$ is not in R_A; put $N_i=\{\bar{a}\}$, $M_i=\emptyset$.

By construction, $M^*=\bigcup_{j<\lambda}M_j$ is disjoint from $N^*=\bigcup_{j<\lambda}N_j$, and N^* does not include any Y_i. We define R_A to be ${}^\omega\lambda-N^*$. Then $M^*\subseteq R_A$, so (ii) holds; N^* is disjoint from R_A, so (i) holds; each Y_i meets R_A, so (iii) holds. This makes (4) true. We get an example of (5) by dualising Q and taking the complement of R_A.

It remains to give an example of (6). We choose $\mathsf{Q}\varphi$ and A as for (3), but with R_A chosen as follows. Let \$ be any ω-shuffle of Q with ordinal ω such that the x_i $(i<\omega)$ have pairwise disjoint images. It suffices to choose R_A so that player $\forall$ has winning strategies for $G^{\$}(\mathsf{Q}\varphi, A)$ and $G^{\$}(\bar{\mathsf{Q}}\neg\varphi, A)$. As in the construction of (4), we can give player $\forall$ strategies ρ and σ for these two games respectively, so that at every move he codes up the preceding play. List as Y_α $(\alpha<\lambda^\omega)$ the sets of form $\{n\text{th thread of }\bar{a}: n<\omega\}$ where $\bar{a}$ is some play of $G^{\$}(\mathsf{Q}\varphi, A)$ in which player uses ρ; let Z_α $(\alpha<\lambda^\omega)$ be a corresponding list for $G^{\$}(\bar{\mathsf{Q}}\neg\varphi, A)$ and σ. The Y_α are pairwise disjoint and of cardinality ω; likewise the Z_α. For ρ and σ to be winning, it suffices that (i) for each $\alpha<\lambda^\omega$, $Y_\alpha\not\subseteq R_A$, and (ii) for each $\alpha<\lambda^\omega$, Z_α meets R_A.

Say that Y_α, Z_β are *close* iff $Y_\alpha \cap Z_\beta \neq \emptyset$. The transitive closure of closeness is an equivalence relation, and each equivalence class is countable. We procure (i) and (ii) on each equivalence class separately, by listing the class in order-type ω and then inductively putting one element of each Y_α or Z_α outside or inside R_A as required.

The *amalgamation property* for $\leqslant_L$ would say (if it were true): If $A \leqslant_L B_i$ $(i = 1, 2)$, then there is C such that $B_i \leqslant_L C$ $(i = 1, 2)$, up to isomorphism over A.

Corollary 19. *If L is $\mathrm{PH}_{\kappa\kappa}$ with $\kappa > \omega$, or if L is $\mathrm{PL}_{\kappa\kappa}$ with κ strongly compact, the amalgamation property fails for L.*

Proof. This follows from possibility (3).

4.4. Ultralimits and a logic with Craig and Feferman–Vaught properties

From now on we assume that κ is strongly compact and (except where stated) all relations and functions in structures are finitary.

Recall Kochen's notion of ultralimits [13]: if A_i $(i < \omega)$ are structures such that each A_{i+1} is an ultrapower $A_{iD_i}^{I_i}$ of A_i, and A_ω is the direct limit of the A_i under the natural embeddings, then we say that A_ω is an *ultralimit* of A_0. We shall write $A_\omega = \mathrm{Ult}\, A_0/D_i$. If the D_i are κ-ultrafilters, we call A_ω a κ*-ultralimit* of A_0.

The construction can be iterated beyond ω. Suppose for each ordinal α we have a set I_α and an ultrafilter D_α on I_α; then for every structure A we can define structures $A^{(\alpha)}$ by induction:

$$A^{(0)} = A,$$

$$A^{(\alpha+1)} = A^{(\alpha)I_\alpha}_{D_\alpha},$$

$$A^{(\delta)} = \varinjlim_{\alpha<\delta} A^{(\alpha)} \quad \text{when } \delta \text{ is a limit ordinal.}$$

For all $\alpha < \beta$ there are canonical elementary embeddings $h_{\alpha\beta} : A^{(\alpha)} \to A^{(\beta)}$, which we use to define the limits at limit ordinals. If $\alpha < \beta < \gamma$, then $h_{\alpha\gamma} = h_{\beta\gamma} h_{\alpha\beta}$. If $A^{(\alpha)} \cong B^{(\alpha)}$, then $A^{(\beta)} \cong B^{(\beta)}$ for all $\beta \geqslant \alpha$.

We shall apply this idea in the case where all the ultrafilters D_α are κ-ultrafilters. Let $D = \langle D_\alpha : \alpha \text{ an ordinal}\rangle$ be a sequence of κ-ultrafilters on sets I_α. Then we define an equivalence relation $\sim_{(D)}$ on structures by:

$$A \sim_{(D)} B \quad \text{iff for some } \alpha,\ A^{(\alpha)} \cong B^{(\alpha)}.$$

It is easy to see that if κ was ω, then we could choose D so that $\sim_{(D)}$ coincides with elementary equivalence. Theorem 20 generalises this fact:

Theorem 20. *There is a sequence D of κ-ultrafilters such that in each similarity type, $\sim_{(D)}$ has only a set of equivalence classes. Moreover there is a proper class C of ordinals such that $A \cong A^{(\gamma)}$ for every structure A and every $\gamma \in C$.*

The proof will rest on the following rather technical lemma:

Lemma 21. *Let λ be any cardinal. Then there is a sequence of κ-ultrafilters D_i on sets I_i $(i<\omega)$ such that if A, B are any two $L_{\kappa\kappa}$-equivalent structures of cardinality $\leqslant\lambda$, in any language, then*

$$\text{Ult } A/D_i \cong \text{Ult } B/D_i.$$

Proof. For the first part of this proof, we allow structures to carry relations of any arity $<\kappa$.

For some cardinal λ' there exists a family of pairs of structures, $\langle (A_\gamma, B_\gamma): \gamma<\lambda'\rangle$, such that each A_γ and each B_γ has a subset of λ as its domain, each A_γ is $L_{\kappa\kappa}$-equivalent to B_γ, and every pair of $L_{\kappa\kappa}$-equivalent structures of cardinality $\leqslant\lambda$ differs from some pair (A_γ, B_γ) in at most the choice of language. Choose a cardinal μ such that $\mu^{<\kappa}=\mu$ and $2^\mu\geqslant\lambda'$, and partition 2^μ into sets Y_γ $(\gamma<\lambda')$ of cardinality $\geqslant\lambda$. Choose a (μ,κ,λ)-independent family $\langle f_\alpha: \alpha<2^\mu\rangle$. For each $\gamma<\lambda'$ and $\alpha\in Y_\gamma$, let $g_{\alpha\gamma}$ be a map from μ to dom B_γ, such that if $f_\alpha(i)\in \text{dom } B_\gamma$, then $g_{\alpha\gamma}(i)=f_\alpha(i)$. Then $\langle g_{\alpha\gamma}: \alpha\in Y_\gamma\rangle$ forms a $(\mu,\kappa,\text{dom } B_\gamma)$-independent family of cardinality $\geqslant\lambda$. Choose an injection $\theta_\gamma: \text{dom } A_\gamma\to\{g_{\alpha\gamma}: \alpha\in Y_\gamma\}$. List the elements of A_γ in order-type λ (possibly with repetitions) as $\bar a_\gamma$, and write $\theta\bar a_\gamma$ for the sequence $\langle\theta(\bar a_\gamma(j)): j<\lambda\rangle$.

Now for each $\gamma<\lambda'$ and each atomic or negated atomic formula φ with variables v_α $(\alpha<\lambda)$, define

$$B_\gamma(\varphi)=\{i<\mu: B_\gamma\models\varphi[\theta\bar a_\gamma]\}.$$

Let D be the filter on μ generated by all intersections of fewer than κ sets of form $B_\gamma(\varphi)$ such that $A_\gamma\models\varphi[\bar a_\gamma]$.

Claim. *D is a κ-filter.*

As usual, the burden is to show that D is proper. Suppose φ_β $(\beta<\nu<\kappa)$ are formulae such that $A_\gamma\models\varphi_\beta[\bar a_\gamma]$ for all $\beta<\nu$. Then $A_\gamma\models\exists\bar v \bigwedge_{\beta<\nu}\varphi_\beta$ $(\bar v)$. Since A_γ and B_γ are $L_{\kappa\kappa}$-equivalent, there is a sequence $\bar b_\gamma$ in B so that $B_\gamma\models\bigwedge_{\beta<\nu}\varphi_\beta[\bar b_\gamma]$. Then since the $g_{\alpha\gamma}$ are $(\mu,\kappa,\text{dom } B_\gamma)$-independent, there is $i<\mu$ such that $(\theta\bar a_\gamma)_i$ agrees with $\bar b_\gamma$ at the relevant places. Of course most generators of D involve sets $B_\gamma(\varphi)$ from several different $\gamma<\lambda'$; but the fact that the f_α were (μ,κ,λ)-independent ensures that we can find an $i<\mu$ which works for all these γ at once. Hence the claim is proved.

Let D_0 be a κ-ultrafilter on μ which extends D. Then for each $\gamma<\lambda'$, θ_γ induces an embedding of A_γ into $B^\mu_{\gamma D_0}$. Hence if A, B are any two $L_{\kappa\kappa}$-equivalent structures of cardinality $\leqslant\lambda$, then A is embeddable in $B^\mu_{D_0}$.

Iterating this construction, we can find $\mu_1\geqslant\mu$ and a κ-ultrafilter D_1 on μ_1 such that if A, B are any two $L_{\kappa\kappa}$-equivalent structures of cardinality $\leqslant\lambda^\mu$, then A is embeddable in $B^{\mu_1}_{D_1}$. Then we can find $\mu_2 \geq \mu_1$ and a κ-ultrafilter D_2 on μ_2 which serves for structures of cardinality $\leq \lambda^{\mu_2}$; and so on for ω steps.

Now let A, B be any two $L_{\kappa\kappa}$-equivalent structures of cardinality $\leq\lambda$. Form A^*, B^* from A, B by adding relations for all formulae of $L_{\kappa\kappa}$. Then any embedding of A^* into an $L_{\kappa\kappa}$-equivalent structure is $L_{\kappa\kappa}$-elementary, so by the construction above we have an $L_{\kappa\kappa}$-elementary embedding $e_0: A^* \to B^{*\mu}_{D_0}$. Then likewise we have an $L_{\kappa\kappa}$-elementary embedding $e_1: B^{*\mu}_{D_0} \to A^{*\mu_1}_{D_1}$ so that the diagram

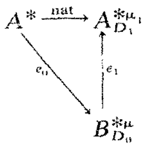

commutes. Continuing the diagram to the right ,for ω steps (cf. [13]), we eventually reach isomorphic κ-ultralimits Ult A/D_{2i+1} and Ult B/D_{2i}. (The non-finitary relations have to be dropped when we take limits.) Since the same proof shows that Ult $A/D_{2i+1} \cong$ Ult A/D_{2i}, it follows that

$$\text{Ult } A/D_{2i} \cong \text{Ult } B/D_{2i},$$

and so $\langle D_0, D_2, \ldots\rangle$ is the required sequence of κ-ultrafilters.

We remark that since κ is strongly compact, Lemma 21 implies that if K is any class of structures such that both K and its complement are closed under κ-ultralimits, then K is defined by a sentence of $L_{\kappa\kappa}$.

Proof of Theorem 20. We define the κ-ultrafilters D_α by induction, using global choice. Let λ be 0 or an uncountable cardinal, and suppose that for each $\alpha<\lambda$ the κ-ultrafilter D_α on μ_α has been defined. Let μ be $\sup\{\mu_\alpha : \alpha<\lambda\}$. By Lemma 21, choose $D_{\lambda+i}$ $(i<\omega)$ so that if A, B are any two $L_{\kappa\kappa}$-equivalent structures of cardinality $\leq 2^\lambda$, then Ult $A/D_{\lambda+i} \cong$ Ult $B/D_{\lambda+i}$. Then for each $\alpha<\lambda^+$, put $D_\alpha = D_\beta$ where $\alpha = (\lambda+\omega)\cdot\gamma+\beta$ and $0\leq\beta<\lambda+\omega$.

To show that this definition of D works for the theorem, consider a language with ν symbols. In this language there are at most $2^{\kappa+\nu}$ pairwise non-$L_{\kappa\kappa}$-equivalent structures. We claim that there are at most $2^{\kappa+\nu}$ equivalence classes of $\sim_{(D)}$ in this language. For if not, we can choose a set K of $(2^{\kappa+\nu})^+$ structures which are pairwise non-equivalent with respect to $\sim_{(D)}$. Let λ be any cardinal greater than the cardinalities of all the structures in K. Then there are distinct $A, B \in K$ such that $A^{(\lambda)}$ is $L_{\kappa\kappa}$-equivalent to $B^{(\lambda)}$, and both have cardinality $\leq 2^\lambda$. But then $A^{(\lambda+\omega)} \cong B^{(\lambda+\omega)}$ and hence $A \sim_{(D)} B$ contradicting the choice of K. This shows that $\sim_{(D)}$ has only a set of equivalence classes in each language.

Finally let C be the class of all transfinite ordinals of the form $(\lambda+\omega)\cdot\gamma$ with $\gamma<\lambda^+$. Then for each $\beta\in C$, the sequences $\langle D_\alpha : \alpha$ an ordinal$\rangle$ and $\langle D_{\beta+\alpha} : \alpha$ an ordinal$\rangle$ are identical, and so for any structure A, $A^{(\beta)(\lambda)} = A^{(\beta+\lambda)} = A^{(\lambda)}$ for all large enough λ, proving that $A \sim_{(D)} A^{(\beta)}$.

Given D as in Theorem 20, we can define a non-standard logic L_D as follows. Let S be a set of representatives (up to isomorphism) of all similarity types of cardinality $<\kappa$. For each $s \in S$ and each set Σ of $\sim_{(D)}$-equivalence classes of s-structures, we introduce a quantifier $\mathsf{Q}_{s\Sigma}$ by the rule

$$A \models \mathsf{Q}_{s\Sigma}(\varphi_1, \ldots) \text{ iff the } \sim_{(D)}\text{-class of the structure } \langle \text{dom } A, (\varphi_1)_A, \ldots \rangle \text{ is in } \Sigma.$$

The sentences of L_D in a given similarity type will be the expressions of form $\mathsf{Q}_{s\Sigma}(R_1\bar{v}, R_2\bar{v}, \ldots)$ where $R_1, R_2, \ldots$ are relation symbols of the similarity type. (For simplicity we are ignoring functions and constants.) Note that there is only a set of such sentences.

Sentences of $L_{\kappa\omega}$ are preserved in κ-ultralimits, and so any two $\sim_{(D)}$-equivalent structures must have the same $L_{\kappa\omega}$-theory. It follows that the logic L_D is not λ-compact for any $\lambda < \kappa$. Against this bad property, it has two good ones:

Theorem 22. *Let L_D be as defined above. Then:*

(i) *the Craig interpolation theorem holds for L_D;*

(ii) *(Feferman–Vaught property) the L_D-theory of a sum or product of two structures is determined by the L_D-theories of the structures.*

Proof. (i) Let s_1, s_2 be similarity types with intersection s. Let φ_1, φ_2 be sentences of $L_D(s_1)$, $L_D(s_2)$ respectively, so that φ_1 entails φ_2. We can suppose without loss that s was in the set S defined earlier, and so we can define Σ to be the set of all $\sim_{(D)}$-equivalence classes of s-structures which contain reducts of models of φ_1. We claim that the sentence $\mathsf{Q}_{s\Sigma}$ is an interpolant in $L_D(s)$ between φ_1 and φ_2.

If A is any model of φ_1, then the s-reduct of A is in the class Σ, and so $A \models \mathsf{Q}_{s\Sigma}$. This shows that φ_1 entails $\mathsf{Q}_{s\Sigma}$.

Suppose B is an s_2-structure which is a model of $\mathsf{Q}_{s\Sigma}$. Then for some model A of φ_1, $A \mid s \sim_{(D)} B \mid s$. It follows that for all large enough ordinals γ, $(A \mid s)^{(\gamma)} \cong (B \mid s)^{(\gamma)}$. In particular this holds for some γ in C. Since γ is in C, $A \sim_{(D)} A^{(\gamma)}$ and hence $A^{(\gamma)}$ is a model of φ_1. So $B^{(\gamma)}$ can be expanded to a model of φ_1, and hence $B^{(\gamma)}$ is a model of φ_2. But then B was also a model of φ_2, because $B \sim_{(D)} B^{(\gamma)}$. This shows that $\mathsf{Q}_{s\Sigma}$ entails φ_2, and so (i) is proved.

(ii) is proved similarly, using the facts that $A^{(\gamma)} \times B^{(\gamma)} \cong (A \times B)^{(\gamma)}$ and $A^{(\gamma)} + B^{(\gamma)} \cong (A+B)^{(\gamma)}$.

References

[1] C.C. Chang, Infinitary properties of models generated from indiscernibles, in: B. van Rootselaar and J.F. Staal; eds., Logic, Methodology and Philosophy of Science (North-Holland, Amsterdam, 1968).

[2] C.C. Chang and H.J. Keisler, Model Theory (North-Holland, Amsterdam, 1973).

[3] R. Engelking and M. Karłowicz, Some theorems of set theory, Fund. Math. 57 (1965) 275–285.

[4] H. Friedman, The complexity of explicit definitions, Advances in Math. 20 (1976) 18–29.

[5] F. Galvin, T. Jech and M. Magidor, An ideal game, J. Symbolic Logic 43 (1978) 284–292.
[6] W. Hodges, Compactness and interpolation for infinitary Horn sentences, Notices Am. Math. Soc. 23 (1976) A651.
[7] W. Hodges, Functorial uniform reducibility, Fund. Math. 108 (1980) 77–81.
[8] J.R. Isbell, Functorial implicit operations, Israel J. Math. 15 (1973) 185–188.
[9] T.J. Jech, Some combinatorial problems concerning uncountable cardinals, Ann. Math. Logic 5 (1973) 165–198.
[10] H.J. Keisler, Limit ultrapowers, Trans. Am. Math. Soc. 107 (1963) 382–408.
[11] H.J. Keisler, Limit ultraproducts, J. Symbolic Logic 30 (1965) 212–234.
[12] H.J. Keisler, Reduced products and Horn classes, Trans. Am. Math. Soc. 117 (1965) 307–328.
[13] S. Kochen, Ultraproducts in the theory of models, Ann. Math. 74 (1961) 221–261.
[14] D.W. Kueker, Countable approximations and Löwenheim–Skolem theorems, Ann. Math. Logic 11 (1977) 57–103.
[15] R. Mansfield, Sheaves and normal submodels, J. Symbolic Logic 42 (1977) 241–250.
[16] S. Shelah, Every two elementarily equivalent models have isomorphic ultrapowers, Israel J. Math. 10 (1971) 224–233.
[17] G. Takeuti, Proof Theory (North-Holland, Amsterdam, 1975).