An. Şt. Univ. Ovidius Constanţa Vol. 19(2), 2011, 101–130

Proof of Gaifman's conjecture for relatively categorical abelian groups

Wilfrid Hodges

Abstract

Abstract. Haim Gaifman conjectured in 1974 that if T is a complete first-order theory which is relatively categorical over its relativisation T^P to a predicate P, then every model B of T^P can be extended to a model A of T with $A^P = B$. He proved this when T is rigid over P, and it holds also for some cases of relative categoricity in uncountably categorical theories. In what may be the first example going significantly beyond these two types, we show that the conjecture is true when the relatively categorical theory T is the theory of an abelian group with P selecting a subgroup.

1 Gaifman's conjecture

In [5] Haim Gaifman introduced some notions equivalent to the following. Let L be a first-order language, P a 1-ary relation symbol not in L, and $L(P)$ the language which results from adding P to L. Let T be a complete theory in $L(P)$, with the property that if A is any model of T, then the L-reduct $A{\restriction}L$ of A has a substructure A^P whose domain is the interpretation of P in A. Note that the complete first order $\mathrm{Th}(A^P)$ of A^P depends only on T; we write it T^P.

Definition 1.1. *(a) We say that such a theory T is* relatively categorical *if whenever A, C are models of T and $i : A^P \to C^P$ is an isomorphism, then i extends to an isomorphism $j : A \to C$.*

(b) Gaifman's conjecture *states that if T is relatively categorical, then for every model B of T^P there is a model A of T with $B = A^P$.*

For Gaifman's own statement see [5], the turn of pages 31 and 32. Gaifman's conjecture is true when L is countable and T is rigid (which means that the isomorphism j in Definition 1.1(a) above is unique). Gaifman himself conjectured a proof of this in a footnote on page 32 of [5], and confirmed it later (cf. [6] Theorem 12.5.8, p. 645). As Gaifman noted, that proof doesn't generalise to the case where T is not rigid. We now know that there is a cohomological obstruction. If T is relatively categorical, then for every model A of T, restriction to A^P is a natural surjective group homomorphism $\nu : \text{Aut}(A) \to \text{Aut}(A^P)$ (where $\text{Aut}(A)$ is the automorphism group of A). Gaifman's construction of A from A^P implies that ν is a split surjection, in the sense that there is a group embedding $\iota : \text{Aut}(A^P) \to \text{Aut}(A)$ such that $\nu\iota = 1_{\text{Aut}(A^P)}$. So the construction fails when ν is not split.

> **Example**. Let the group A be $\mathbb{Z}(4)^{(\omega)}$. Write a_i for a generator of the i-th direct summand ($i < \omega$). Make A into an $L(P)$-structure by taking A^P to be the subgroup generated by $\{2a_i : i < \omega\}$. Then the theory of A is relatively categorical. The natural restriction map $\nu : \text{Aut}(A) \to \text{Aut}(A^P)$ is not a split surjection; this is proved in Evans, Hodges and Hodkinson [2].

Two facts about this example are worth noting. First, the same example appears in Ahlbrandt and Ziegler [1] as an example of an ω_1-categorical theory, where P selects a strongly minimal set. Every ω_1-categorical theory with a strongly minimal set selected by P is an example of a relatively categorical theory. Theories of this kind do satisfy Gaifman's conjecture. But they are hardly typical. For example if the $L(P)$-structure A is $\mathbb{Z}$ with $A^P = 2\mathbb{Z}$, then $\text{Th}(A)$ is relatively categorical but $\text{Th}(A^P)$ is not even superstable.

Second, let B be A^P and let B' be the subgroup of B generated by $\{2a_i : 0 < i < \omega\}$. Then $B' \preccurlyeq B$, and we can find a substructure A' of A with $A' \equiv A$ and $(A')^P = B'$ by taking A' to be generated by $\{a_i : 0 < i < \omega\}$. This is a special case of Gaifman's conjecture. But A' is not unique with these properties; taking a_1 to $a_1 + 2a_0$ and keeping the other generators of A fixed, we get an automorphism of A which fixes B' pointwise but moves A' to a different subgroup of A. In this context, we can't hope to prove the conjecture by finding the required A' as the set of all elements of A that satisfy some condition.

In 1986 Saharon Shelah published [9], which generalises relative categoricity by allowing a limited number of isomorphism types of A over a given A^P. The paper is famously difficult; its contents have never been reworked and brought into the mainstream. Very possibly the paper contains ideas sufficient for a proof of Gaifman's conjecture; though if it does, they are not yet in a form that logicians in general can use. It will be sensible to try to correlate features of the proof below with notions that Shelah uses.

At the meeting in Constanţa for Şerban Basarab's seventieth birthday, I sketched a proof of Gaifman's conjecture for theories of abelian group pairs that are relatively categorical for a pair of cardinals (as defined in [7]), and a way of generalising this proof to a larger class of theories. Writing this paper, I found that a detailed report of reasonable length would have to cover less ground. So below I restrict to the case of theories that are relatively categorical absolutely. There was no space here to discuss the pure model-theoretic core.

My warm thanks to the energetic and helpful organisers of the Constanţa meeting, and to Şerban Basarab himself for being both an inspiration to workers in the model theory of algebra, and also seventy years old.

2 Preliminaries

2.1 Models of relatively categorical theories

Henceforth the language L is the first-order language of abelian groups. By an *abelian group pair* we mean an $L(P)$-structure A which is an abelian group with A^P a subgroup. We write T for a complete theory of abelian group pairs in $L(P)$.

Fact 2.1. *If T is relatively categorical, then for every formula $\phi(v_0, \ldots, v_{n-1})$ of $L(P)$ there is a formula $(\phi)^\circ$ such that for every model A of T and every n-tuple $\bar{b}$ of elements of A^P,*

$$A \models \phi(\bar{b}) \quad \Leftrightarrow \quad A^P \models (\phi(\bar{b}))^\circ.$$

The property ascribed to T by Fact 2.1 is known as the *(Uniform) Reduction Property*. In fact this is a property of all relatively categorical first-order theories, without any restriction to abelian groups; cf. [6] p. 641, Lemma 12.5.1.

Fact 2.2. *Let T be relatively categorical. Then for every model A of T, A/A^P is a group of finite exponent. Moreover there is a finite abelian group pair D_0 with $D_0^P = \{0\}$, such that every model A of T has the form $A = C \oplus D$ where C, D as groups are subgroups of the group A, $A^P = C^P$, D is isomorphic to D_0 and C is tight over C^P. The theory Th(C) is relatively categorical and is uniquely determined by T.*

This fact draws together a number of results in [7]; see in particular Theorem 14.1 there. The notion of 'tight' is explained in [7]; we will not need it, but we will need a special case of it which is explained in subsection 2.2 below. Fact 2.2 reduces Gaifman's conjecture for T to the case where all models A^P of T are tight over A^P.

Fact 2.3. *Suppose T is relatively categorical and every model A of T is tight over A^P. Then there are finitely many primes $p_0, \ldots, p_{n-1}$ and relatively categorical theories $T_0, \ldots, T_{n-1}$ such that every model A of T is a pushout over A^P of models A_i of T_i ($i < n$) with $A_i^P = A^P$, and for each $i < n$, A_i/A_i^P is a p_i-group of finite exponent.*

This fact again draws together several results from [7]; see in particular sections 5 and 12 in [7]. The notion of pushout is standard universal algebra, but in Section 5 below we will explain what we need of it. The effect of Fact 2.3 is to reduce Gaifman's conjecture for T to the case where there are a prime p and a finite h_0 such that all models A of T are tight over A^P and A/A^P is a p-group of exponent h_0.

2.2 Algebraic facts and definitions

Suppose A is an $L(P)$-structure and X a set of elements of A. We write $\langle X \rangle$ for the subgroup of A generated by X. We extend the language $L(P)$ of A to a first-order language $L(P, X)$ by introducing each element of X as a constant. We write (A, X) for the expansion of A to an $L(P, X)$-structure by taking each element of X to name itself.

Let p be a prime. We write $p^k A$ for the subgroup of elements a of A such that for some $c \in A$, $p^k c = a$, and $p^k A[p]$ for the subgroup of $p^k A$ consisting of the elements a such that $pa = 0$. We say that A is *p-bounded over* its subgroup B if for some finite h_0, $p^{h_0} A \subseteq B$.

Suppose A is an abelian group and B a subgroup over which A is p-bounded. Then we say that A is *tight over B* if for each $k < \omega$,

$$p^k A[p] \;\subseteq\; p^{k+1} A + B.$$

(This is not the definition of 'tight' in [7], but by the argument of [7] Lemma 8.8 it is equivalent to that definition under the assumption that A is p-bounded over B.) Note that if A is tight over B then A is tight over every group C with $B \subseteq C \subseteq A$.

Let A be an abelian group and a an element of A. For any natural number h we say that a has *p-height at least h in A* if $a \in p^h A$. If we count a as having p-height ∞ when it has p-height at least h for each $h < \omega$, then each element a of A has a unique p-height in A, which is either a natural number or ∞. We write this p-height as $\mathrm{ht}_p^A(a)$; when the context allows, we sometimes omit the superscript ^{A}. Note that when h is finite, there is a first-order formula expressing that $\mathrm{ht}_p(x) \geqslant h$, and hence also a first-order formula expressing that $\mathrm{ht}_p(x) = h$.

Suppose the abelian group A is p-bounded over its subgroup B, and a is an element of A. Then a is said to be *proper over B in A* if $\mathrm{ht}_p^A(a)$ is maximal among the p-heights $\mathrm{ht}_p^A(a+b)$ with $b \in B$. Note that if C is a subgroup of B and a is proper over B, then a is proper over C.

Lemma 2.4. *Suppose B is a subgroup of A over which A is p-bounded and tight, and a is an element of A which is proper over B. Then*

$$ht_p^A(pa) = ht_p^A(a) + 1.$$

Proof. Put $h = \mathrm{ht}_p^A(a)$. Then for some element d of A, $p^h d = a$. It follows that pa has p-height at least $h+1$; we must show that it has p-height at most $h+1$. For contradiction suppose there is an element c of A such that $p^{h+2}c = pa$. Then

$$p(p^{h+1}c - p^h d) \;=\; pa - pa \;=\; 0.$$

Now A was assumed tight over B, so

$$p^{h+1}c - p^h d \in p^h A[p] \;\subseteq\; p^{h+1}A + B.$$

Hence there is $b \in B$ such that $a + b = p^h d + b$ has p-height $\geqslant h+1$, contradicting that p is proper over B. □

Suppose the abelian group A has a subgroup B. We say that *p-heights are preserved between A and B* if for each element b of B, $\mathrm{ht}_p^A(b) = \mathrm{ht}_p^B(b)$. Note that this is automatically true if $B \preccurlyeq A$. A key part of our argument in section 4 below will be that under certain conditions the converse holds too.

Much of the next section relies on the following assumption:

(†) A is an abelian group which is tight and p-bounded over its subgroup B, and p^{h_0} is the exponent of A/B.

3 Exploiting Kaplansky and Mackey

3.1 KM sequences, algebraic description

An essential tool in [7] was the Kaplansky-Mackey back-and-forth proof [8] of Ulm's Theorem, as described in Fuchs [4] section 77. The main task of the present section will be to extract the first-order content of the Kaplansky-Mackey procedure in the case of tight extensions. We convert the procedure into a recipe for constructing sequences of elements within a single structure. The original back-and-forth idea will make a brief but crucial reappearance in Lemma 3.11 below.

Definition 3.1. *Let A and B be abelian groups, with B a subgroup of A. A* KM(A, B) sequence *is a sequence $\bar{a} = (a_i : i < \alpha)$ of elements of A, such that for each $i < \alpha$,* *(i)* $pa_i \in \langle \bar{a}{\restriction}i\rangle + B$.

(ii) a_i *has finite* p*-height in* A.

(iii) a_i *is proper over* $\langle \bar{a}{\restriction}i\rangle + B$ *in* A.

The subgroup generated by $\bar{a}$ over B *is the subgroup* $\langle \bar{a}\rangle + B$.

Lemma 3.2. *If* $(a_i : i < \alpha)$ *is a KM(*A, B*) sequence, then for all* $i < \alpha$, $a_i \notin \langle \bar{a}{\restriction} i\rangle + B$.

Proof. By (iii) in Definition 3.1, $\mathrm{ht}_p(a_i - c) \leqslant \mathrm{ht}_p(a_i)$ for each $c \in \langle \bar{a}{\restriction}i\rangle + B$. If a_i was such a c, we would have

$$\infty \;=\; \mathrm{ht}_p(0) \;=\; \mathrm{ht}_p(a_i - a_i) \;\leqslant\; \mathrm{ht}_p(a_i),$$

which is impossible since $\mathrm{ht}_p(a_i)$ is finite by (ii) in Definition 3.1. □

Lemma 3.3. *Assume* (†). *Then every KM(*A, B*) sequence can be extended to a KM(*A, B*) sequence which generates* A *over* B.

Proof. Suppose $\bar{a} = (a_i : i < \alpha)$ is a KM(A, B) sequence. Write C for $\langle \bar{a}\rangle + B$, and suppose $C \neq A$, so that there is some element $a \in C \setminus A$. Since A/B has exponent p^{h_0}, there is some least $h \leqslant h_0$ such that $p^h a \in C$, and so $p^{h-1}a \notin C$. If c is any element of C then again $p^{h-1}a + c \notin C$, so $p^{h-1}a + c$ has p-height $< h_0$. Therefore among all $c \in C$ there is one c' such that $p^{h-1}a + c'$ has maximal p-height h' in A, i.e. is proper over C. Putting $a_\alpha = p^{h-1}a + c'$ extends the KM(A, B) sequence.

Now we proceed by induction on the ordinals. Starting with $\bar{a}$, we use the recipe of the previous paragraph to add elements at the end for so long as the resulting KM(A, B) sequence fails to generate A over B. The process must halt with a sequence $\bar{a'}$ of length $< |A|^+$, since by Lemma 3.2 the a'_i are distinct for distinct i. Since $\bar{a'}$ can't be extended to a longer KM(A, B) sequence, it must already generate A over B. □

Lemma 3.4. *Assume* (†), *and let* $\bar{a}$ *be a KM(*A, B*) sequence. Suppose* $\bar{c}$ *is a finite subsequence of* $\bar{a}$. *Then there is a finite subsequence* $\bar{d}$ *of* $\bar{a}$ *which includes* $\bar{c}$ *and is a KM(*A, B*) sequence.*

Proof. Let $\bar{a}$ be $(a_i : i < \alpha)$, and recall from Definition 3.1 that for each $i < \alpha$, $pa_i \in \langle \bar{a}{\restriction}i\rangle + B$. We say that a subset W of α is *closed* if for every i in W there is a finite subsequence $\bar{e}$ of $\bar{a}{\restriction}i$ such that each item in $\bar{e}$ is a_j for some $j < i$ with $j \in W$, and $pa_i \in \langle \bar{e}\rangle + B$. An application of König's tree lemma shows that each finite subset of α is contained in a finite closed set.

The lemma follows if we show that for each closed subset W of α, the restriction $\bar{d}$ of $\bar{a}$ to those a_i with $i \in W$ is a KM(A, B) sequence. List W in increasing order as $(i_j : j < \ell)$, so that $\bar{d} = (a_{i_j} : j < \ell)$. By Definition 3.1 we need to show that for each $j < \ell$,

(i) $pa_{i_j} \in \langle \bar{d}\restriction j\rangle + B$.

(ii) a_{i_j} has finite p-height in A.

(iii) a_{i_j} is proper over $\langle \bar{d}\restriction j\rangle + B$ in A.

Here (i) follows from the definition of closed set, and (ii) follows from (ii) for $\bar{a}$. For (iii) it suffices to note that $\langle \bar{d}\restriction j\rangle \subseteq \langle \bar{a}\restriction i_j\rangle$. □

Suppose B is a subgroup of N. We say that a KM(M, N)-sequence $\bar{a}$ *has support in* B if for each i, $pa_i \in \langle \bar{a}\restriction i\rangle + B$.

3.2 KM sequences, model-theoretic description

There are two main gaps to be bridged between the ideas of Kaplansky and Mackey and first-order logic. One is that the notion of being proper over a subgoup is hard to express with first-order formulas, and this accounts for some unnaturalness in the conditions (i:c) and (i:c)$^-$ in Lemma 3.6 below; maybe I missed some clearer way of doing it. The other is that Kaplansky and Mackey operate with equations and heights; it's only when we have equations and heights under good enough control to generate automorphisms that we can start to handle general first-order properties. (Cf. Lemma 3.12 below.) I strongly suspect that some kind of 'quantifier elimination over P' is at work here, but I don't yet have a good formulation of it.

Definition 3.5. *Assuming* (†), *we define* $Th(A, B)_{KM}$ *to consist of the following sentences of* $L(P, B)$*:*

(a) The first-order theory of abelian group pairs.

(b) First-order statements expressing that A is tight over B and A/B has exponent p^{h_0}.

(c) All true equations and negated equations in $Th(A, B)$.

(d) All true sentences in $Th(A, B)$ of the form $P(t)$ or $\neg P(t)$, where t is a closed term of $L(P, B)$.

(e) All true sentences in $Th(A, B)$ of the form $ht_p(t) \geqslant k$ or $ht_p(t) < k$ where t is a closed term of $L(P, B)$ and k is a positive integer.

We want a description of KM(A, B) sequences in terms of first-order properties of the structure (A, B). In short, we want to describe the type tp($\bar{a}/B$), in A and over B, of the KM(A, B) sequence $\bar{a}$. For this we introduce the following notation.

We fix a sequence $\bar{v} = (v_i : i < \xi)$ of variables indexed by ordinals (where ξ is taken large enough for the purpose in hand), and another sequence of variables $(z_i : i < \omega)$. The idea will be that the variables v_i range over elements of A while the variables z_i range over elements of B. The type tp($\bar{a}/B$) will be written in the language $L(P, B)$ with the variables v_i; writing $\bar{a} = (a_i : i < \alpha)$, the variable v_i will stand for the element a_i.

A term t (always in the language of abelian groups) will have variables from the v_i's and other variables from the z_i's. We write Term(i) for the set of all terms $t(\bar{v} \restriction i; \bar{z})$; here $\bar{z}$ is some tuple from the z_i's, and of course only finitely many of the variables in $\bar{v} \restriction i$ will actually occur in $t(\bar{v} \restriction i; \bar{z})$. We write Term($i, B$) for the set of terms of the form $t(\bar{v} \restriction i; \bar{b})$ such that $t \in$ Term(i) and $\bar{b}$ is from B.

Lemma 3.6. *Assume (†). Let $\bar{a} = (a_i : i < \alpha)$ be a sequence of elements of A. Then the following are equivalent:*

(1) $\bar{a}$ is a KM(A, B) sequence.

(2) For each $i < \alpha$ there are $h < h_0$ and a term $t(\bar{v}) \in$ Term(i, B) such that tp($\bar{a}/B$) contains the formulas

(i:a) $pv_i = t$.

(i:b) $ht_p(v_i) = h$.

(i:c) For each term $t'(\bar{v} \restriction i; \bar{z}) \in$ Term(i),

$$\forall \bar{z}\ (P(\bar{z}) \wedge ht_p(t'(\bar{v} \restriction i; \bar{z})) \geqslant h \rightarrow ht_p(t + pt'(\bar{v} \restriction i; \bar{z})) \leqslant h + 1).$$

(3) As (2), but with (i:b) and (i:c) replaced by

(i:b)$^-$ $ht_p(v_i) \geqslant h$.

(i:c)$^-$ For each term $t'(\bar{v} \restriction i; \bar{z}) \in$ Term(i) and each $\bar{b}$ in B,

$$ht_p(t'(\bar{v} \restriction i; \bar{b})) \geqslant h \rightarrow ht_p(t + pt'(\bar{v} \restriction i; \bar{b})) \leqslant h + 1.$$

Proof. First we show that (2) and (3) are equivalent. Clearly (2) entails (3), given that by Definition 3.5(d), Th$(A, B)_{KM}$ identifies the tuples in B. For the converse, note first that if $\bar{a} \restriction i$ satisfies each formula in (i:c)$^-$, then it also satisfies those in (i:c) by the truth definition for universal quantifiers. Second, (i:a) and (i:c)$^-$ together imply that $\mathrm{ht}_p(v_i) \leqslant h$; for otherwise t has p-height at least $h+2$, which contradicts (i:c)$^-$ when $t'(\bar{v}{\restriction}i;\bar{b})$ is taken to be 0.

It remains to show that (1) and (2) are equivalent.

(1) ⇒ (2): Let $i < \alpha$ and let h be the p-height of a_i in A. We prove (i:a), (i:b), (i:c) when the variables v_i are taken to stand for the elements a_i.

(i:a) By (i) in Definition 3.1, $pa_i \in \langle \bar{a}{\restriction}i\rangle + B$. But each element of $\langle \bar{a}{\restriction}i\rangle + B$ has the form $t(\bar{a}{\restriction}i;\bar{b})$ for some term t in $\mathrm{Term}(i)$.

(i:b) is by (†), Definition 3.1(ii) and the choice of h.

(i:c) As t' ranges over $\mathrm{Term}(i)$ and $\bar{z}$ ranges over tuples in B, $t'(\bar{v}';\bar{z})$ ranges over $\langle \bar{a}{\restriction}i\rangle + B$. So it suffices to consider any element c of $\langle \bar{a}{\restriction}i\rangle + B$ of p-height $\geqslant h$, and show that $t(\bar{v};\bar{b}) + pc$, i.e. $pv_i + pc$, has p-height $\leqslant h+1$. Now since c has p-height $\geqslant h = \mathrm{ht}_p(v_i)$, the element $v_i + c$ has p-height $\geqslant h$, and hence by (iii) in Definition 3.1 is proper over $\langle \bar{a}{\restriction}i\rangle + B$ and has p-height exactly h. It follows by Lemma 2.4 that $pv_i + pc$ has p-height $h+1$ in A.

(2) ⇒ (1): Let $i < \alpha$. We prove (i), (ii), (iii) in Definition 3.1.

(i) By (i:a) the element pa_i is $t(\bar{a}';\bar{b})$ for some subsequence $\bar{a}'$ of $\bar{a}{\restriction}i$ and some $\bar{b}$ in B.

(ii) follows at once from (i:b).

(iii) Suppose c is in $\langle \bar{a}{\restriction}i\rangle + B$, and hence is $t'(\bar{a}{\restriction}i;\bar{b}')$ for some term $t' \in \mathrm{Term}(i)$. We have to show that $a_i + c$ has p-height at most $\mathrm{ht}_p(a_i)$, which is h by (i:b). The p-height of $a_i + c$ can be $> h$ only if $\mathrm{ht}_p(c) = h$, so without loss we can assume this. Then by (i:c), $pa_i + pc$ has p-height at most $h+1$, so that $a_i + c$ has p-height at most h as required. □

Definition 3.7. *Assume* (†).

(a) A KM(A,B) preschedule *(of length α) is a set Γ of formulas of $L(P,B)$ that consists of, for each $i < \alpha$, a formula of the form (i:a) of Lemma 3.6 for some term $t \in$ Term(i,B), a formula of the form (i:b) for some $h < h_0$, and the formulas (i:c).*

(b) A KM(A,B) schedule *is a KM(A,B) preschedule Γ such that Th(A,B)$_{KM} \cup \Gamma$ is finitely consistent.*

(c) If $\bar{a}$ is a KM(A,B) sequence of length α, Γ is a KM(A,B) preschedule and $\Gamma \subseteq$ tp($\bar{a}/B$), then we describe Γ as the schedule of $\bar{a}$. *(It meets the definition of KM(A,B) schedule in (b) above by Lemma 3.6.)*

(d) If Γ is a KM(A,B) preschedule, we write $\Gamma{\restriction}i$ for the set of formulas (j:a), (j:b) and (j:c) in Γ with $j < i$.

(e) *The* support *of the KM(A, B) preschedule Γ is the set of constants $b \in B$ which occur in formulas in Γ.*

These definitions raise some obvious questions. (1) Does a KM(A, B) sequence determine its schedule (as defined in Definition 3.7(c)) uniquely? (2) Given the structure (A, B), does the schedule of the KM(A, B) sequence $\bar{a}$ determine tp($\bar{a}/B$)? (3) Is every KM(A, B) schedule the schedule of some KM(A, B) sequence? Here are some brief comments on these questions.

(1) The answer is strictly No: for example the term t in (i:a) could be replaced by $t + 0$, or it could have redundant variables. But this is the only indeterminacy, since when t is fixed, both (i:b) and (i:c) are determined by the height ht$_p(a_i)$. Also the element pa_i named by $t(\bar{a}\restriction i)$ is determined by $\bar{a}$; and if t_1 is another term in Term(i) such that $t_1(\bar{a}\restriction i)$ names pa_i, then Lemma 3.8(b) below will show that

$$\text{Th}(A, B)_{KM} \cup \Gamma\restriction i \vdash t(\bar{a}\restriction i) = t_1(\bar{a}\restriction i).$$

So the indeterminacy is slight.

(2) The answer is Yes, by Remark 3.14 below.

(3) The answer is Yes; this is Corollary 3.10 below.

3.3 Technical properties of Kaplansky-Mackey

Lemma 3.8. *Assume* (†). *Let Γ be a KM(A, B) schedule of length α, and suppose $i < \alpha$. Then from $\Gamma\restriction i$ and Th$(A, B)_{KM}$ we can deduce the following formulas:*

(a) *(For each term $t(\bar{v}) \in$ Term(α, B))*
Either the formula ($t(\bar{v}; \bar{b}) = 0$) or the formula ($t(\bar{v}; \bar{b}) \neq 0$).

(b) *(For each t as in (a))*
Either the formula $P(t(\bar{v}))$ or the formula $\neg P(t(\bar{v}))$.

(c) *(For each t as in (a))*
Either a formula (ht$_p(t(\bar{v})) = h$) for some $h < h_0$, or the formula (ht$_p(t(\bar{v})) \geqslant h_0$).

(d) *Each formula in (i:c)$^-$ (as in Lemma 3.6 — and note that we really do mean (i:c) and not ($i + 1$:c)).*

Proof. We prove (a)–(c) of the lemma for each term $t(\bar{v}; \bar{z})$ in Term(i) and each $\bar{b}$) in B, by induction on $i \leqslant \alpha$.

If $i = 0$ then terms in Term(0) contain no variables, so that (a)–(c) are determined by Th(A, B). There is nothing to prove at limit ordinals.

Assume then that (a)–(c) of the lemma are proved for terms in Term(i, B), and $t(\bar{v})$ is in Term$(i+1, B) \setminus$ Term(i, B). Then v_i is the highest-numbered variable in $t(\bar{v})$, so that $t(\bar{v})$ can be written as $mv_i + t'$ where $t' \in$ Term(i, B).

If p divides m then $t(\bar{v})$ is provably equal to a term t_1 in Term(i, B), so the induction hypothesis applies. If p doesn't divide m then m is prime to p, so that for some n prime to p, $nm \equiv 1 \pmod{p}$. The answers to (a)–(c) are not affected by multiplying $t(\bar{v})$ by n, so we can assume henceforth that $t(\bar{v})$ has the form $v_i + t'$ with $t' \in$ Term(i, B).

The answer to (a) is $\neq$, by the proof of Lemma 3.2. The answer to (b) is $\neg P$; for if $P(v_i + t')$ then $v_i + t' = b$ for some $b \in B$, and hence $v_i = b - t'$, contradicting Lemma 3.2 again.

We show (c). By (i:b) in $\Gamma \restriction i + 1$, ht$_p(v_i) = h$ for some $h < h_0$. Let h' be ht$_p(t')$; by induction hypothesis $\Gamma\restriction i$ and Th$(A, B)_{KM}$ determine h'. By first-order properties of height, ht$_p(v_i + t')$ is $\min\{h, h'\}$ and hence is determined by Th$(A, B)_{KM} \cup (\Gamma \restriction i + 1)$, unless possibly when $h = h$. If $h = h'$ then certainly ht$_p(v_i + t') \geqslant h$; moreover by (i:a) and one of the sentences (i:c) in $\Gamma \restriction i + 1$, ht$_p(t + pt') \leqslant h + 1$ and hence ht$_p(v_i + t') \leqslant h$. This deduces ht$_p(v_i + t') = h$ from Th$(A, B)_{KM} \cup (\Gamma\restriction i+1)$. We infer ht$_p(v_i + t'(\bar{v}'; \bar{b})) = h$.

Finally we show (d). Since t is in Term(i, B) and t' is in Term(i), (c) in the induction hypothesis shows that the p-heights of $pv_i + pt'(\bar{v}'; \bar{b}))$ and $t + pt'(\bar{v}'; \bar{b})$ are determined by Th$(A, B)_{KM} \cup \Gamma \restriction i$. Since Th$(A, B)_{KM} \cup \Gamma$ is assumed finitely consistent, these heights are determined consistently with the formulas (i:c) in Γ. The conclusion follows. □

Lemma 3.9. *Suppose Γ is a KM(A, B) schedule of length α, and for some $i < \alpha$, $\bar{a}$ is a KM(A, B) sequence of length i which satisfies $\Gamma\restriction i$. Then there is $\bar{c}$ such that $\bar{a}^\frown\bar{c}$ is a KM(A, B) sequence satisfying Γ.*

Proof. We begin by changing notation; we write $\bar{a}_0$ for the sequence $\bar{a}$ of the lemma. We will construct a KM(A, B) sequence $\bar{a}$ of length α in A so that

(a) $\bar{a}\restriction i = \bar{a}_0$ and

(b) for every $j < \alpha$, $\bar{a}\restriction j$ satisfies $\Gamma\restriction j$.

We proceed by induction on $j \geqslant i$. When $j = i$ we choose $\bar{a}\restriction i$ to be $\bar{a}_0$; then (a) holds, and (b) follows from (a) and the lemma assumption. There is nothing to do at limit ordinals. So we assume (b) when j is k and $i \leqslant k < \alpha$, and we find a_k to make (b) true when $j = k + 1$.

More precisely we assume that

$$A \models \left(\bigwedge \Gamma\restriction k\right)(\bar{a}\restriction k)$$

and we find an element a_k so that

$$A \models \left(\bigwedge \Delta\right)(\bar{a}{\restriction}k+1)$$

where $\bigwedge \Delta$ consists of the formulas (k:a), (k:b) and (k:c) of Lemma 3.6. By the equivalence of (2) and (3) in Lemma 3.6, it will suffice to show that $(\bar{a}{\restriction}k+1)$ satisfies

$$(k\text{:a}) \quad pv_i = t(\bar{v}{\restriction}k),$$
$$(k\text{:b})^- \quad \mathrm{ht}_p(v_k) \geqslant h$$

(where t and h are as determined by Γ) and the formulas of $(k\text{:c})^-$ from Lemma 3.6.

Now the formulas of $(k\text{:c})^-$ have their free variables in $\bar{v}{\restriction}k$, so we already know from Lemma 3.8(d) that $\bar{a} \restriction k+1$ will satisfy the formulas of $(k\text{:c})^-$, regardless of the choice of a_k.

The formulas (k:a) and (k:b) together entail that t has p-height at least $h+1$. But by Lemma 3.8(c), the set $\mathrm{Th}(A,B)_{KM} \cup \Gamma{\restriction}k$ determines the p-height of t, so by the finite consistency of $\mathrm{Th}(A,B)_{KM} \cup \Gamma$ it follows that

$$\mathrm{Th}(A,B)_{KM} \cup \Gamma{\restriction}i \;\vdash\; \mathrm{ht}_p(t) \geqslant h+1.$$

Thus by induction assumption

$$A \models \mathrm{ht}_p(t(\bar{a}{\restriction}k)) \geqslant h+1$$

and hence there is an element d in A such that $p^{h+1}d = t(\bar{a}{\restriction}k)$. Put $a_k = p^h d$. Then $pa_k = t(\bar{a}{\restriction}k)$ and $\mathrm{ht}_p(a_k) \geqslant h$, so that $(\bar{a}{\restriction}k+1)$ satisfies (k:a) and (k:b) as required. □

The following corollary shows the sense in which the model-theoretic notion of a KM(A, B) schedule captures the algebraic notion of a KM(A, B) sequence.

Corollary 3.10. *Assume (†), and let Γ be a set of formulas of $L(P,B)$. Then the following are equivalent:*

(a) Γ is a KM(A, B) schedule (i.e. a KM(A, B) preschedule which is finitely consistent with Th$(A,B)_{KM}$).

(b) Γ is the schedule of some KM(A, B) sequence.

(c) Γ is a KM(A, B) preschedule which is satisfied in A by some sequence.

Moreover any sequence of length α which satisfies a KM(A, B) preschedule of length α in A is a KM(A, B) sequence.

Proof. (a) ⇒ (b) is by Lemma 3.9. (b) ⇒ (c) is immediate. (c) ⇒ (a) is because A is a model of $\mathrm{Th}(A,B)_{KM}$. The last sentence is by (2) ⇒ (1) in Lemma 3.6. □

3.4 Construction of automorphisms over P

Lemma 3.11. *Assume* $(\dagger)$. *If $\bar{a}$ and $\bar{c}$ are KM(A, B) sequences with the same schedule of length i, then there is an automorphism of A which takes $\bar{a}$ to $\bar{c}$ and fixes B pointwise.*

Proof. We extend $\bar{a}$ and $\bar{c}$ back and forth. More precisely, rename the two sequences in the lemma $\bar{a}_0$ and $\bar{c}_0$. We define, by induction on $j \geqslant i$, sequences $\bar{a}$ and $\bar{c}$ so that

(a) $\bar{a}{\restriction}i = \bar{a}_0$ and $\bar{c}{\restriction}i = \bar{c}$,

(b) for each j, $\bar{a}{\restriction}j$ and $\bar{c}{\restriction}j$ are KM(A, B) sequences with the same schedule.

When $j = i$, we ensure (a) by defining $\bar{a}{\restriction}i$ and $\bar{c}{\restriction}i$ to be $\bar{a}_0$ and $\bar{c}_0$ respectively, and then (b) holds by the lemma assumption. At limit ordinals there is nothing to do. The procedure will halt when we reach a j such that each of $\bar{a}{\restriction}j$ and $\bar{c}{\restriction}j$ generates A over B.

Assume then that we have defined $\bar{a}{\restriction}k$ and $\bar{c}{\restriction}k$ so that (b) above holds when $j = k$. By assumption at least one of $\bar{a}{\restriction}k$ and $\bar{c}{\restriction}k$ fails to generate A over B; by symmetry we can assume it is $\bar{a}{\restriction}k$. Then as in the proof of Lemma 3.3, there is an element a_k of A which is proper over $\langle\bar{a}{\restriction}k\rangle + B$, and $(\bar{a}{\restriction}k+1)$ is a KM(A, B) sequence. Let Γ be the schedule of $(\bar{a}{\restriction}k+1)$. Then by (b), $\bar{c}{\restriction}k$ satisfies the schedule $\Gamma{\restriction}k$. So by Lemma 3.9 there is c_k so that $(\bar{c}{\restriction}k+1)$ satisfies Γ. Hence we have (b) for $j = k+1$ as required.

By Lemma 3.2 the procedure must eventually halt, say when $\bar{a}$ and $\bar{c}$ have length α. At this point each of $\bar{a}$ and $\bar{c}$ generates A over B. Again by Lemma 3.2 we can define a set-theoretic map f by putting $f(a_j) = c_j$ for each $j < \alpha$ and $f(b) = b$ for each $b \in B$. To show that f generates an automorphism g of A, it suffices to show that for each term $t(\bar{v})$ in Term(α, B),

$$A \models t(\bar{a}) = 0 \leftrightarrow t(\bar{c}) = 0.$$

This follows from Lemma 3.8(a), since $\bar{a}$ and $\bar{c}$ satisfy the same schedule. Then $g\bar{a}_0 = \bar{c}_0$ and g fixes B pointwise, since these are both true for f, and g extends f. □

3.5 Canonical definitions

Assume $(\dagger)$ throughout this subsection.

If $\bar{a}$ is a finite KM(A, B) sequence in A, then the schedule of $\bar{a}$ is a finite set of formulas. Writing their conjunction as a formula $\gamma(\bar{v})$, we call γ a *schedule formula*.

Lemma 3.12. *Let $\bar{a}$ be a finite KM(A, B) sequence with schedule formula $\gamma(\bar{v})$. Then*

$$Th(A, B), \gamma \vdash \bigwedge tp(\bar{a}/B).$$

Proof. Suppose $\phi(\bar{v})$ is a formula in tp($\bar{a}/B$). If

$$\mathrm{Th}(A, B) \cup \{\gamma(\bar{v}), \neg\phi(\bar{v})\}$$

is consistent, then $\exists\bar{v}(\gamma(\bar{v}) \wedge \neg\phi(\bar{v}))$ is consistent with the complete theory Th(A, B) and hence is a consequence of Th(A, B). So there is a tuple $\bar{c}$ in A such that

$$A \models \gamma(\bar{c}) \wedge \neg\phi(\bar{c}).$$

Since both $\bar{a}$ and $\bar{c}$ satisfy the same schedule formula γ, by Lemma 3.11 there is an automorphism of A fixing B pointwise and taking $\bar{c}$ to $\bar{a}$. But then $\bar{a}$ satisfies $\neg\phi$, contradicting the choice of ϕ as a formula in tp($\bar{a}/B$). □

Remark 3.13. Suppose $\gamma(\bar{v}, \bar{b})$ is a KM(A, B) schedule formula, where we have exhibited as $\bar{b}$ the elements of B in the support of the schedule. If $\bar{b}'$ is another tuple from B, when is $\gamma(\bar{v}, \bar{b}')$ also a schedule formula? The answer is in Corollary 3.10: when $\exists\bar{v}\gamma(\bar{x}, \bar{b}')$ is consistent with Th$(A, B)_{KM}$. This holds for example whenever $M \models \exists\bar{v}\gamma(\bar{x}, \bar{b}')$.

For the rest of this subsection, suppose $\bar{a}$ is a KM(A, B) sequence of length α, and Γ is the schedule of $\bar{a}$. We can define *closed subsets* W of α as in the proof of Lemma 3.4. In fact they can be defined in terms of Γ rather than $\bar{a}$: we require that for each $i \in W$, if t is the term in (i:a) then the indices of the variables in t are in W. (This could introduce some redundant ordinals, but no matter.) Then if W is a nonempty closed subset of α, we can define a restriction Γ_W of Γ to W: Γ_W contains, for each $i \in W$, the formulas (i:a) and (i:b) of Γ together with those formulas (i:c) where the term t' uses only indices in $W \cap i$. Then collapse the indices of variables in these formulas, in an order-preserving way, to an initial segment of the ordinals. The result is again a KM(A, B) schedule, and it is the schedule of the KM(A, B) sequence got in the proof of Lemma 3.4 by restricting $\bar{a}$ to the indices in W. In this sense, if X is any finite subset of α, then Γ contains a finite KM(A, B) schedule which includes the information about all the elements v_i with $i \in X$.

Remark 3.14. The construction above allows us to extend Lemma 3.12 to arbitrary KM(A, B) sequences. If $\bar{a}$ is such a sequence with schedule Γ, then to show that Th$(A, B) \cup \Gamma$ entails tp($\bar{a}/B$), it suffices to show that Th$(A, B) \cup \Gamma$ entails tp($\bar{a}'/B$) for each finite subsequence $\bar{a}'$ of $\bar{a}$. But each such subsequence is covered by a sub-schedule of Γ as above, and so Lemma 3.12 gives the result for $\bar{a}'$. Note that in general it would be impossible to apply the same argument directly to Γ itself, since we have made no saturation assumptions and Γ may have infinitely many variables.

Suppose next that $\bar{a}$ generates A over B. Let $\bar{c} = (c_0, \ldots, c_{n-1})$ be an n-tuple of elements of A. Then (for example by Lemma 3.4) there are a finite subsequence $\bar{d}$ of $\bar{a}$ and a tuple $\bar{b}_1$ of elements of B, and terms $t_0, \ldots, t_{n-1}$ in the language of abelian groups such that

(a) $\bar{d}$ is a KM(A, B) sequence, and

(b) $(c_0 = t_0(\bar{d}, \bar{b}_1)) \wedge \ldots \wedge (c_{n-1} = t_{n-1}(\bar{d}, \bar{b}_1))$.

We abbreviate the formula in (b) to $(\bar{c} = \bar{t}(\bar{d}, \bar{b}_1))$. Let $\gamma(\bar{v}, \bar{b}_2)$ be a schedule formula of $\bar{d}$, and let $\theta(\bar{x})$ be the formula

$$\exists \bar{v}(\gamma(\bar{v}, \bar{b}_2) \wedge (\bar{x} = \bar{t}(\bar{d}, \bar{b}_1))).$$

In practice we can generally amalgamate $\bar{b}_1$ and $\bar{b}_2$ into a single tuple $\bar{b}$ and write θ as $\theta(\bar{x}, \bar{b})$. We describe a formula $\theta(\bar{x}, \bar{b})$ constructed according to this recipe as a *canonical definition* of $\bar{c}$, or a *canonical $\bar{a}$-definition* of $\bar{c}$ to indicate the KM(A, B) sequence that was used.

Lemma 3.15. *Assume (†). Let $\theta(\bar{x}, \bar{b})$ be a canonical definition of a tuple in A. If $\bar{c}_0$ and $\bar{c}_1$ are two sequences in A that satisfy θ, then there is an automorphism of A taking $\bar{c}_0$ to $\bar{c}_1$ and fixing B pointwise.*

Proof. Unpacking θ, there are $\bar{d}_0$ and $\bar{d}_1$ in A so that

$$A \models \gamma(\bar{d}_k, \bar{b}_2) \wedge (\bar{c}_k = \bar{t}(\bar{d}, \bar{b}_1)) \;\; (k < 2).$$

Then $\bar{d}_0$ and $\bar{d}_1$ have the same schedule, so by Lemma 3.11 there is an automorphism f of A taking $\bar{d}_0$ to $\bar{d}_1$ and fixing B pointwise. Then

$$\bar{c}_1 = \bar{t}(\bar{d}_1, \bar{b}_1) = \bar{t}(f(\bar{d}_0), f(\bar{b}_1)) = f(\bar{t}(\bar{d}_0, \bar{b}_1)) = f(\bar{c}_0).$$

□

4 Proof of Gaifman's conjecture for a tight extension and one prime

Throughout this section, the complete $L(P)$-theory T is a relatively categorical theory of abelian group pairs, and there are a prime p and a natural number h_0 such that for every model M of T, the exponent of M/M^P is p^{h_0} and

M is tight over M^P. At the end of this section we will show that for every model B of T^P there is a model A of T with $A^P = B$. Up to but not including the main theorem, Theorem 4.9 we will assume the following:

Convention 4.1. The abelian group B is a model of T^P. Moreover there is a model M of T with $N = M^P$, such that B is an elementary substructure of N. We can construct a KM(M, N) sequence $\bar{a}_B$ which has support in B and is maximal with this property (since limit steps in the construction add no new support). Let α be the length of $\bar{a}_B$. Then by Lemma 3.3 we can extend $\bar{a}_B$ to a KM(M, N) sequence $\bar{a}$ which generates M over N. Let β be the length of $\bar{a}$. Let A be the subgroup of M generated by $\bar{a}_B$ and B, i.e. $A = \langle \bar{a}_B \rangle + B$. We make A into an abelian group pair by putting $A^P = B$. For the rest of this section the following items remain fixed: M, N, B, $\bar{a}$, α, β, A.

Note that Convention 4.1 includes (†) from the end of section 2 above, but with M, N in place of A, B. It will be a major step to prove in Corollary 4.6 that (†) holds for A, B as well.

Lemma 4.2. *Assuming Convention 4.1, we have $A^P = B$.*

Proof. Since A inherits its $L(P)$-structure from M, $A^P = A \cap N$. Then Convention 4.1 gives at once that $B \subseteq A \cap N$. It remains to show $A \cap N \subseteq B$. If this fails, there is some first $i < \alpha$ such that some $b \in N \setminus B$ can be written in the form $ma_i + c$ where c is in $\langle \bar{a} \restriction i \rangle + B$. As usual, m can be assumed to be 1. But then $a_i = b - c$, where $b - c$ lies in $\langle \bar{a} \restriction i \rangle + N$. Since $\bar{a}$ is assumed to be a KM(M, N) sequence, this contradicts Lemma 3.2. □

In the following lemma we introduce an important technique, using the fact that $B \preccurlyeq N$ to slide tuples of elements from N into B. To lead the eye, I underline the elements from N that need to be slid into B. So an element $\underline{b} \in N$ becomes replaced by $b \in B$.

Lemma 4.3. *Assume Convention 4.1. Let $\bar{c}$ be elements of A and $\underline{\bar{b}}'$ elements of N, and ϕ a first-order formula. If*

$$M \models \phi(\bar{c}, \underline{\bar{b}})$$

then there are $\bar{b}$ in B such that

$$M \models \phi(\bar{c}, \bar{b}).$$

Proof. Let $\theta(\bar{x})$ be a canonical $\bar{a}$-definition of $\bar{c}$ in M. Since $\bar{a}$ is in A, we can write θ as $\theta(\bar{x}, \bar{b}_0)$ where $\bar{b}_0$ is in B. Hence

$$M \models \exists \bar{x}\, (\theta(\bar{x}, \bar{b}_0) \wedge \phi(\bar{x}, \underline{\bar{b}})).$$

Since T is relatively categorical, we can apply the Reduction Property from Fact 2.1:

$$N \models (\exists \bar{x}\,(\theta(\bar{x}, \bar{b}_0) \wedge \phi(\bar{x}, \underline{b})))^{\circ}.$$

Since $B \preccurlyeq N$, there is $\bar{b}$ in B such that

$$N \models (\exists \bar{x}\,(\theta(\bar{x}, \bar{b}_0) \wedge \phi(\bar{x}, \bar{b})))^{\circ}$$

and so by the Reduction Property again

$$M \models \exists \bar{x}\,(\theta(\bar{x}, \bar{b}_0) \wedge \phi(\bar{x}, \bar{b})).$$

This tells us that there is $\bar{a}_1$ in M such that

$$M \models \theta(\bar{a}_1, \bar{b}_0) \wedge \phi(\bar{a}_1, \bar{b}).$$

Since $\bar{c}$ and $\bar{a}_1$ satisfy the same canonical definition in M, by Lemma 3.15 there is an automorphism of M which takes $\bar{a}_1$ to $\bar{c}$ and pointwise fixes B. Then

$$M \models \phi(\bar{c}, \bar{b})$$

as required. □

Lemma 4.4. *Assume Convention 4.1. Let c be an element of A, h a finite number and a an element of M of p-height h such that $pa = c$. Then there is an element $a' \in A$ of p-height $\geqslant h$ such that $pa' = c$. Moreover if a is not in N then a' can be found not in B.*

Proof. **Case One**: $a \in A + N$. Put $a = a_1 + \underline{b}$ with $a_1 \in A$ and $\underline{b} \in N$. Let $\phi(x, y, z)$ be the formula

$$p(y + z) = x \wedge \mathrm{ht}_p(y + z) \geqslant h.$$

Then

$$M \models \phi(c, a_1, \underline{b}).$$

By Lemma 4.3 there is $b \in B$ such that

$$M \models \phi(c, a_1, b)$$

and so $a' = a_1 + b$ meets our requirements. If a is not in N, then add to ϕ the conjunct

$$\neg P(y + z).$$

Case Two: $a \notin A+N$. Consider the set of pairs $(a^\star, e^\star)$ where $\mathrm{ht}_p(a^\star) \geqslant h$, $a^\star \notin A + N$, $e^\star \in A + N$ and $pa^\star = c$. The set is not empty, since it contains

$(a, 0)$. For each such pair, $a^\star + e^\star$ is not in N, so that its p-height is at most $h_0 - 1$. So there is such a pair $(a^\star, e^\star)$ for which the p-height of $a^\star + e^\star$ has maximal value, say the finite value $h^\star \geqslant h$.

Fixing this pair $(a^\star, e^\star)$, write $e^\star$ as $a_1 + \underline{b}$ where a_1 is in A and $\underline{b}$ is in N. Let $\psi(x, y, z, w)$ be the formula

$$pw = x \wedge \text{ht}_p(w) \geqslant h \wedge \text{ht}_p(w + y + z) = h^\star.$$

Then

$$M \models \exists w\, \psi(c, a_1, \underline{b}, w).$$

Hence by Lemma 4.3 there are b in B and a'' in M such that

$$M \models \psi(c, a_1, b, a'').$$

If a'' is not in $A + N$ then neither is $a'' + a_1 + b$, and moreover $a'' + a_1 + b$ is proper over $A + N$ by the choice of h'. So since $p(a'' + a_1 + b) = c + pa_1 + pb \in A + B$, we can extend $\bar{a}{\restriction}\alpha$ to a longer KM(M, N)-sequence over B by putting $a_\alpha = a'' + a_1 + b$. This contradicts the choice of $\bar{a}{\restriction}\alpha = \bar{a}_B$ in Convention 4.1. Therefore $a'' \in A + N$ and we can revert to Case One. If $a \notin N$, then again we can ensure that $a'' \notin N$ by adding a suitable conjunct to ψ. □

Lemma 4.5. *Assume Convention 4.1. Then p-heights are preserved between M and A.*

Proof. Let h be a positive integer and a an element of A with p-height $\geqslant h$ in M. We claim:

> For every $i \leqslant h$ there is $d \in M$ such that $p^h d = a$ and $p^{h-i} d \in A$.

The proof is by induction on i.

For $i = 0$ the claim holds by assumption. Assuming the claim for $i = k < h$ we prove it for $i = k + 1$ as follows. By assumption there are $d \in M$ and $c \in A$ such that $c = p^{h-k} d$ and $p^k c = a$. Then $p^{h-k-1} d$, which exists since $k < h$, is an element a' such that $pa' = c$ and a' has p-height $\geqslant h - k - 1$ in M. By Lemma 4.4 there is such an element a'' in A'; so for some $d' \in M$, $p^{h-(k+1)} d' = a'' \in A$ and $p^{k+1} a'' = p^k c = a$. This proves the claim for $i = k + 1$.

Hence the claim holds. When $i = h$ it says that a has p-height $\geqslant h$ in A. □

Corollary 4.6. *Assume Convention 4.1. Then:*

(a) A/A^P has exponent p^{h_0}.

(b) A is tight over A^P.

(c) *(†) holds for both* (M, N) *and* (A, B)*, with the same value of* h_0*.*

(d) $Th(A, B)_{KM} = Th(M, N)_{KM} \cap L(P, B)$*.*

Proof. (a) Since $A \subseteq M$, A/A^P certainly has exponent at most p^{h_0}. To show that the exponent of A/A^P is at least p^{h_0}, consider an element d_1 in M such that $p^{h_0}d_1 \in N$ but $p^{h_0-1}d_1 \notin N$. By Lemma 4.3 there is an element d_2 in M such that $p^{h_0}d_2 \in B$ but $p^{h_0-1}d_2 \notin N$. Then by Lemma 4.4 (putting $a = p^{h_0-1}d_2$) there is an element a' in $A \notin B$ such that $pa' \in B$ and $\mathrm{ht}_p^M(a') \geqslant h_0 - 1$. But by Lemma 4.5, $\mathrm{ht}_p^M(a') = \mathrm{ht}_p^A(a')$. This shows that A/A^P has exponent at least p^{h_0}.

(b) Suppose $a \in A$ and $a \in p^k A[p]$. Then $a \in p^k M[p]$, and so by the tightness of M over N, a can be written as $p^{k+1}c + \underline{b}$ with $c \in M$ and $\underline{b} \in N$. Hence

$$M \models \mathrm{ht}_p(a - \underline{b}) \geqslant k + 1.$$

By Lemma 4.3 there is $b \in B$ such that

$$M \models \mathrm{ht}_p(a - b) \geqslant k + 1,$$

so by Lemma 4.5

$$\mathrm{ht}_p^A(a - b) = \mathrm{ht}_p^M(a - b) \geqslant k + 1$$

and thus $a \in p^{k+1}A + B$ as required.

(c) is checked from (a), (b) and the text of (†) at the end of section 2..

For (d) we check the clauses of Definition 3.5. Clause (a) is clear. Clause (b) is by (a) and (b) above. Clauses (c) and (d) are because the formulas involved are quantifier-free. Clause (e) is by Lemma 4.5. □

Corollary 4.6(d) tells us that

$$M \models \bigwedge \mathrm{Th}(A, B)_{KM}.$$

If we can raise this result to

$$M \models \bigwedge \mathrm{Th}(A, B)$$

then we are home. The next two lemmas carry out this raising.

Lemma 4.7. *Assume Convention 4.1. Then:*

(a) *Let* Γ *be a KM(A, B) preschedule and* $\bar{a}$ *a sequence in A. Then*

$$A \models \bigwedge \Gamma(\bar{a}) \;\Leftrightarrow\; M \models \bigwedge \Gamma(\bar{a}).$$

(b) *A set of formulas of $L(P,B)$ is a KM(A,B) schedule if and only if it is a KM(M,N) schedule.*

(c) *Every KM(A,B) sequence is a KM(M,N) sequence.*

Proof. (a) We show that for every formula ϕ in Γ,

$$A \models \phi(\bar{a}) \;\Leftrightarrow\; M \models \phi(\bar{a}).$$

If ϕ is a formula (i:a) then this holds because ϕ is atomic. If ϕ is a formula (i:b) then it holds because of Lemma 4.5. If ϕ is a formula in (i:c) then the implication holds from right to left because ϕ is a $\forall$ formula; the implication from left to right holds by Lemma 4.3.

(b) By Lemma 4.6 the exponent p^{h_0} is the same for M/N as it is for A/B. It follows by Definition 3.7(a) that Γ is a KM(A,B) preschedule if and only if it is a KM(M,N) preschedule. If Γ is a KM(A,B) schedule, then by Corollary 3.10 it is satisfied in A by some sequence $\bar{a}$, and hence by (a) above, $\bar{a}$ satisfies it in M too, so that by Lemma 3.6 $\bar{a}$ is a KM(M,N) sequence with KM(M,N) schedule Γ. Conversely if Γ is a KM(M,N) schedule, then Γ is finitely consistent with Th$(M,N)_{KM}$, and hence with Th$(A,B)_{KM}$ since Th$(A,B)_{KM} \subseteq$ Th$(K,M)_{KM}$ by Corollary 4.6(d). Hence Γ is a KM(A,B) schedule.

(c) If $\bar{c}$ is a KM(A,B) sequence, the schedule Γ of $\bar{c}$ in A is a KM(A,B) schedule in $L(P,B)$, so by (b) it is a KM(M,N) schedule. By (a), $\bar{c}$ satisfies Γ in M too, so $\bar{c}$ is a KM(M,N) sequence. □

Lemma 4.8. *Assume Convention 4.1. Then $A \preccurlyeq M$.*

Proof. To apply the Tarski-Vaught criterion for elementary substructures, let $\bar{c}$ be a tuple of elements of A and ϕ a formula of $L(P)$ such that for some $a \in M$,

$$M \models \phi(\bar{c}, a).$$

It suffices to find an element $a' \in A$ such that

$$M \models \phi(\bar{c}, a')$$

We convert the question into one about KM sequences. First, since $\bar{c}$ is in A there is a finite subsequence $\bar{a}_1$ of $\bar{a}{\restriction}\alpha$ such that $\bar{c}$ lies in $\langle \bar{a}_1 \rangle + B$ and $\bar{a}_1$ is a KM(M,N) sequence. Second, there is a finite subsequence $\bar{a}_2^\star$ of $\bar{a}$ such that a lies in $\langle \bar{a}_2^\star \rangle + N$ and $\bar{a}_2^\star$ is a KM(M,N) sequence. We can choose $\bar{a}_2^\star$ so that it includes $\bar{a}_1$, and then without loss we can expand $\bar{a}_1$ to be the part of $\bar{a}_2^\star$ that lies in $\bar{a}{\restriction}\alpha$. Write $\bar{a}_2^\star$ as $\bar{a}_1 {}^\frown \bar{a}_2$. Let $\gamma_2(\bar{v}_1\bar{v}_2, \bar{b}_1, \underline{\bar{b}}_2)$ be the schedule formula of

$\bar{a}_1{}^\frown\bar{a}_2$, where $\bar{b}_1$ is in B and $\underline{b}_2$ is in $N \setminus B$. A part of $\gamma_2(\bar{v}_1\bar{v}_2, \bar{b}_1, \underline{b}_2)$ forms the schedule formula $\gamma_1(\bar{v}_1, \bar{b}_1)$ of $\bar{a}_1$, in such a way that

$$\vdash \ \forall\bar{v}_1\forall\bar{v}_2\forall\bar{z}_1\forall\bar{z}_2\ (\gamma_2(\bar{v}_1, \bar{v}_2, \bar{z}_1, \bar{z}_2) \to \gamma_1(\bar{v}_1, \bar{z}_1)).$$

(See the discussion before Remark 3.14 above.) Finally there are abelian group terms $\bar{t}_1$, t_2 such that

$$M \ \models\ (\bar{c} = \bar{t}_1(\bar{a}_1)) \wedge (a = t_2(\bar{a}_1, \bar{a}_2).)$$

We assemble all this information and quantify:

$$M \ \models\ \exists\bar{v}_1\bar{v}_2\ (\phi(\bar{t}_1(\bar{v}_1), t_2(\bar{v}_1, \bar{v}_2)) \wedge \gamma_2(\bar{v}_1, \bar{v}_2, \bar{b}_1, \underline{b}_2)).$$

From this point onwards we proceed by a series of seven claims.

Claim One. There are $\bar{b}_2$ in B and $\bar{a}_1'', \bar{a}_2''$ in M such that

$$M \ \models\ \phi(\bar{t}_1(\bar{a}_1''), t_2(\bar{a}_1'', \bar{a}_2'')) \wedge \gamma_2(\bar{a}_1'', \bar{a}_2'', \bar{b}_1, \bar{b}_2).$$

This is by Lemma 4.3.

Claim Two. $\gamma_2(\bar{v}_1, \bar{v}_2, \bar{b}_1, \bar{b}_2)$ and and $\gamma_1(\bar{v}_1, \bar{b}_1)$ are KM(M, N) schedule formulas. This claim holds for γ_2 by Remark 3.13 and Claim One, and for γ_1 by its definition.

Claim Three. $\gamma_2(\bar{v}_1, \bar{v}_2, \bar{b}_1, \bar{b}_2)$ and $\gamma_1(\bar{v}_1, \bar{b}_1)$ are KM(A, B) schedule formulas. This claim holds by Lemma 4.7(b).

Claim Four. $\bar{a}_1$ satisfies $\gamma_1(\bar{v}_1, \bar{b}_1)$ in A. To show this, note first that

$$M \ \models \gamma_1(\bar{a}_1, \bar{b}_1)$$

by the choice of γ_1 and $\bar{b}_1$. Then apply Lemma 4.7(a).

Claim Five. There is $\bar{a}_2'$ in A such that

$$A \ \models\ \gamma_2(\bar{a}_1, \bar{a}_2', \bar{b}_1, \bar{b}_2).$$

so that $\bar{a}_1{}^\frown\bar{a}_2$ is a KM(A, B) sequence. This is by Lemma 3.9, Claims Three and Four and the relationship between γ_1 and γ_2.

Claim Six. There are a KM(M, N) sequence $\bar{a}_3'$ which generates M and extends $\bar{a}_1{}^\frown\bar{a}_2'$ and a KM(M, N) sequence $\bar{a}_3''$ which has the same schedule as $\bar{a}_3'$ and extends $\bar{a}_1''{}^\frown\bar{a}_2''$. To prove the first part, $\bar{a}_1{}^\frown\bar{a}_2'$ is a KM(M, N) sequence by Claim Five and Lemma 4.7(c), so we can find $\bar{a}'_3$ by Lemma 3.3. Then we can find $\bar{a}''_3$ by Lemma 3.9.

Claim Seven. There is an automorphism f of M which takes $\bar{a}''_1{}^\frown\bar{a}''_2$ to $\bar{a}_1{}^\frown\bar{a}'_2$.

This is by Claim Six and Lemma 3.11.

With these claims we are home. Let f be as in Claim Seven, and put $a' = t_2(\bar{a}_1, \bar{a}'_2) = f(t_2(\bar{a}''_1, \bar{a}''_2))$. By choice of $\bar{a}_1$ and Claim Five, $\bar{a}_1$ and $\bar{a}'_2$ are in A, so a' is in A. Applying the automorphism f to Claim One, and recalling that $\bar{c} = \bar{t}_1(\bar{a}_1)$ in M,

$$M \models \phi(\bar{c}, a').$$

□

Theorem 4.9. *Let T be a relatively categorical theory of abelian group pairs, such that for some prime p, if A is any model of T then A is tight and p-bounded over A^P. Then for every model B of T^P there is a model A of T with $A^P = B$.*

Proof. Take $\kappa > |B|$. Let M be a κ-saturated model of T, and define N to be M^P. Since M was κ-saturated, so is N, and hence we can elementarily embed B in N. With the usual adjustments we can assume $B \preccurlyeq N$. Now we can define $\bar{a}, \alpha, \beta, A$ as in Convention 4.1. Then by Lemma 4.8, $A \preccurlyeq M$ and so the theorem is proved. □

5 Proof of Gaifman's conjecture for relatively categorical abelian group pairs

5.1 Module-like structures and their pushouts

Definition 5.1.

(a) Given a first-order language L, an L-structure and a formula $\phi(v_0, \ldots, v_{n-1})$ of L, we write $\phi(A^n)$ for the set of all n-tuples $\bar{a}$ in A such that $A \models \phi(\bar{a})$.

(b) A formula is positive primitive, *for short* pp, *if it has the form*

$$\exists \bar{x}\ (\phi_0 \wedge \ldots \wedge \phi_{n-1})$$

where each ϕ_i is atomic. (Here atomic means equations and relational formulas $R(\bar{v})$, not including $\perp$.)

(c) We say that an L-structure A is module-like *if L has a binary function symbol $+$, A forms an abelian group under $+$, and for every pp formula $\phi(v_0, \ldots, v_{n-1})$ of L without parameters, $\phi(A^n)$ is a subgroup of the product abelian group A^n.*

(d) A first-order theory T is module-like *if all its models are module-like (with respect to the same symbol $+$).*

(e) The Baur-Monk invariants *of a module-like structure A are the numbers $|\phi(A)/\psi(A)\cap\phi(A)|$, where ϕ and ψ are pp formulas with one free variable, and the numbers are taken as either finite or ∞. Since the conjunction of two pp formulas is logically equivalent to a single pp formula, there is no loss in considering only invariants $|\phi(A)/\psi(A)|$ where $\vdash \psi \rightarrow \phi$.*

We note that every abelian group pair is a module-like structure, so that any theory of abelian group pairs is module-like. Module-like theories satisfy the Baur-Monk quantifier elimination theorem, cf. [6] Theorem A.1.1, page 654. One corollary of this fact is the following criterion for elementary embedding.

Fact 5.2. *Let T be a module-like theory in a language L and $f : A \rightarrow M$ a homomorphism between models of T. Then f is an elementary embedding of A in M if and only if:*

(a) For every pp formula $\phi(\bar{v})$ of L and tuple $\bar{a}$ in A, if $M \models \phi(f\bar{a})$ then $A \models \phi(\bar{a})$, and

(b) for each Baur-Monk invariant $|\phi/\psi|$,

$$\left|\frac{\phi(A)}{\psi(A)}\right| \geqslant \left|\frac{\phi(M)}{\psi(M)}\right|.$$

Proof. If f is an elementary embedding then (a) and (b) are easily checked. For the converse we use a result essentially due to Gabriel Sabbagh in the case of modules (cf. [6] Corollary A.1.3, page 656). Namely, if f is an embedding and (a) holds and (b) holds with $=$ in place of $\geqslant$, then f is elementary. Now if (a) holds then f is an embedding. Also if (a) holds then (b) holds with $\leqslant$ for $\geqslant$. For suppose $|\phi(A)/\psi(A)| \geqslant m$. Then there are $a_0, \ldots, a_{m-1}$ in $\phi(A)$ such that for all $i < j < m$, $a_i - a_j \notin \psi(A)$. But by (a) the elements $fa_0, \ldots, fa_{m-1}$ have the same properties in M, so that $|\phi(M)/\psi(M)|$ too. So to satisfy Sabbagh's criterion it suffices to prove (a) and (b). □

Let T be a module-like theory in a language L. We define pushouts of models of T as follows. (The essentials are in Fuchs [3] section 10. There may be fuller accounts in some recent theoretical computer science texts.)

Definition 5.3. *Let $A_0, \ldots, A_{n-1}$ be models of T, and suppose that some L-structure B is a substructure of each A_i ($i < n$). Let Π be the direct product*

$$\Pi = A_0 \times \ldots \times A_{n-1}.$$

Let $\mathbb{K}$ be the subgroup of Π generated by the elements $(b_0, \ldots, b_{n-1})$ of B^n such that $b_0 + \ldots + b_{n-1} = 0$. The pushout *of $A_0, \ldots, A_{n-1}$ over B is the structure $\Pi/\mathbb{K}$. For each $i < n$ we write $\iota_i : A_i \to A$ for the induced homomorphism taking a_i to $(0, \ldots, 0, a_i, 0, \ldots, 0) + \mathbb{K}$. Since B is a substructure of each A_i, each ι_i is injective (cf. [3] Theorem 10.2).*

Fact 5.4. *Pushouts have the following universal mapping property. Let A be the pushout of $A_0, \ldots, A_{n-1}$ over B. Suppose that D is a module-like structure for the same language as A, and for each $i < n$ there is a homomorphism $g_i : A_i \to D$ such that if $i, j < n$ and $b \in B$ then $g_i(b) = g_j(b)$. Then there is a unique homomorphism $f : A \to D$ such that for each $i < n$, $g_i = f.\iota_i$. (Again see [3] Theorem 10.2.)*

We consider the following situation. Abelian group pairs $M_0, \ldots, M_{n-1}$ are given, with an abelian group $N = M_i^P$ for each $i < n$, and M is the pushout of the M_i over N. Also abelian group pairs $A_0, \ldots, A_{n-1}$ are given, with an abelian group $B = A_i^P$ for each $i < n$, and A is the pushout of the A_i over B. We have $B \preccurlyeq N$ and $A_i \preccurlyeq M_i$ for each $i < n$. We illustrate with $n = 2$ in the following diagram:

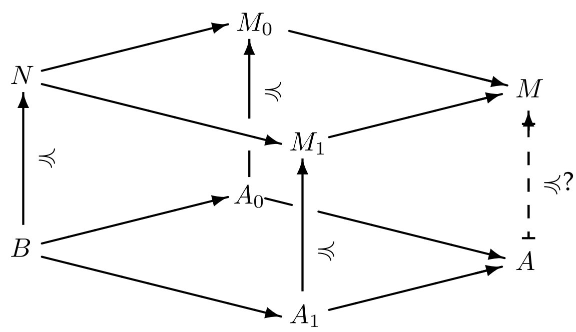

By the universal mapping property in Fact 5.4 there is a homomorphism $f : A \to M$. We will prove, under suitable assumptions, that f is an elementary embedding.

In view of the importance of pp formulas for applying Fact 5.2, we give necessary and sufficient conditions for a tuple to satisfy a pp formula in the pushout of a finite family of extensions of a module-like structure. For a matrix C we write $\rho_i(C)$ for the i-th row of C and $\kappa_j(C)$ for the j-th column of C. We write $\Sigma(C)$ for the row matrix which is the sum of the rows of C, and $\bar{0}$ for the row matrix consisting of zeros.

Lemma 5.5. *Let A be the pushout of $A_0, \ldots, A_{n-1}$ over B, where all these structures are models of a module-like theory in a language L. Let $\phi(x_0, \ldots, x_{m-1})$ be a pp formula of L. Let C be an $n \times m$ matrix (c_{ij}) such that for each $i < n$ the i-th row $\rho_i(C) = (c_{i0}, \ldots, c_{i,m-1})$ is an m-tuple of elements of A_i. As above, we write $\mathbb{K}$ for the substructure of $A_0 \times \ldots \times A_{n-1}$ consisting of all n-tuples of elements of B that sum to zero. Then the following are equivalent:*

(1)
$$A \models \phi(\kappa_0(C) + \mathbb{K}, \ldots, \kappa_{m-1}(C) + \mathbb{K}).$$

(2) There is an $n \times m$ matrix H of elements of B so that
$$B \models \Sigma(H) = \bar{0}$$
and for each $i < n$
$$A_i \models \phi(\rho_i(C) + \rho_i(H)).$$

(3) (a) There is an $n \times m$ matrix H of elements of B so that for each $i < n$,
$$A_i \models \phi(\rho_i(C) + \rho_i(H)),$$
and

(b) if G is any matrix of elements of B so that for each $i < n$,
$$A_i \models \phi(\rho_i(C) + \rho_i(G))$$
then
$$B \models \phi(\Sigma(G)).$$

Proof. (1) ⇔ (2): (1) holds if and only if $(\kappa_0(C) + \mathbb{K}, \ldots, \kappa_{m-1}(C) + \mathbb{K})$ is the image under $\Pi \mapsto \Pi/\mathbb{K}$ of an m-tuple of elements of Π which satisfy ϕ in Π, i.e. such that it satisfies ϕ at each of the n factors. This is what (2) says.

(2) ⇒ (3): Assuming (2), H certainly satisfies (3)(a). If G is a matrix over B and for each $i < n$
$$A_i \models \phi(\rho_i(C) + \rho_i(G))$$
then since ϕ defines an abelian group, for each $i < n$
$$B \models \phi(\rho_i(H) - \rho_i(G))$$
and so
$$B \models \phi(\Sigma(H) - \Sigma(G)).$$
But by (1), $\Sigma(H) = \bar{0}$ and so $\Sigma(G) \in \phi(B^m)$, proving (3)(b).

(3) ⇒ (2): If H is as in (3)(a), then taking H for G in (3)(b) gives that

$$B \models \phi(\Sigma(H)).$$

Form the matrix H' by subtracting $\Sigma(H)$ from the first row of H. Then $\Sigma(H') = \bar{0}$, and by the additivity of ϕ,

$$A_i \models \phi(\rho_i(C) + \rho_i(H'))$$

for each $i < n$. This proves (2) with H' for H. □

5.2 Proof of Gaifman's conjecture

Theorem 5.6. *Let T be a relatively categorical theory of abelian group pairs. If B is a model of T^P then there is a model A of T with $A^P = B$.*

Proof. By Fact 2.2 there are a relatively categorical theory T' and a finite abelian group pair D_0 with $D_0^P = \{0\}$, such that every model A of T is the direct sum of a model C of T' and an isomorphic copy of D_0, and moreover for any such C, $A^P = C^P$ and C is tight over C^P. So it will suffice to find a model C of T' with $C^P = B$ and put $A = C \oplus D_0$. Then by Feferman-Vaught, A is a model of T and $A^P = C^P = B$. The effect of this reduction is that we need only prove the case of the theorem where every model A of T is tight over A^P. Henceforth we assume this is the case.

By Fact 2.3 there are primes $p_0, \ldots, p_{n-1}$ and relatively categorical theories $T_0, \ldots, T_{n-1}$ such that every model A of T is a pushout over A^P of models A_i of T_i with $A_i^P = A^P$, each A_i/A^P is a p_i-group of finite exponent, and each A_i is tight over A^P.

Let B be a model of T^P. Then by Theorem 4.9, for each $i < n$ there is a model A_i of T_i with $B = A_i^P$. In fact the earlier arguments showed that if we elementarily embed B in $N = M^P$ for some model M of T, then each A_i can be found as an elementary substructure of the component M_i of M. Let A be the pushout of the A_i's over B. Then the universal mapping property of pushouts gives us a homomorphism $f : A \to M$ as in the previous section. To show that A is a model of T, and hence to prove Gaifman's conjecture, it suffices to show that f is an elementary embedding.

We will apply Fact 5.2. We write Π for the product $A_0 \times \ldots \times A_{n-1}$ and $\mathbb{K}$ for the kernel consisting of all m-tuples of elements of B which sum to zero. We write Π^+ for the product $M_0 \times \ldots \times M_{n-1}$ and $\mathbb{K}^+$ for the set of all m-tuples of elements of N which sum to zero. So $A = \Pi/\mathbb{K}$ and $M = \Pi^+/\mathbb{K}^+$.

First we prove (a) in Fact 5.2. Let $\phi(x_0, \ldots, x_{n-1})$ be a pp formula and C an $n \times m$ matrix of elements of C such that

$$M \models \phi(f(\kappa_0(C) + \mathbb{K}), \ldots, f(\kappa_0(C) + \mathbb{K})). \tag{1}$$

We need to show that

$$A \models \phi(\kappa_0(C) + \mathbb{K}, \ldots, \kappa_0(C) + \mathbb{K}). \tag{2}$$

We can eliminate the f as follows. The pushout was constructed as a quotient of the direct product, which happens also to be a direct sum when (as here) the number of factors is finite. There is a universal mapping theorem for direct sums, which maps the product of the A_i's into the product of the M_i's, and we can assume it is an inclusion map. Because of this, the homomorphism $f : A \to M$ can be read as the map

$$\bar{a} + \mathbb{K} \mapsto \bar{a} + \mathbb{K}^+,$$

so that in place of (1) we can write

$$M \models \phi(\kappa_0(C) + \mathbb{K}^+, \ldots, \kappa_0(C) + \mathbb{K}^+). \tag{3}$$

By (3) and (1) $\Rightarrow$ (2) in Lemma 5.5 there is an $n \times m$ matrix $\underline{H} = (h_{ij})$ of elements of N so that

$$M \models \Sigma(\underline{H}) = \bar{0}$$

and for each i,

$$M_i \models \phi(\rho_i(C) + \rho_i(\underline{H})).$$

Now by the theory in earlier parts of this paper, each M_i is generated over N by a KM(M_i, N) sequence with an initial segment that generates A_i over B. Also the Reduction Property from Fact 2.1 holds for each T_i; we write $()^\circ$ as $()^\circ_i$ to show that we are dealing with T_i. Using canonical definitions θ_i of $\rho_i(C)$ in each M_i gives

$$N \models \exists \bar{z}(\Sigma \bar{z} = 0 \wedge \bigwedge_{i<n} (\exists \bar{w}\ (\phi(\bar{w} + \bar{z}(i)) \wedge \theta_i(\bar{w}, \bar{b}_i))^\circ_i) \tag{4}$$

where $\bar{z}$ represents an $n \times m$ matrix. Since the i-th row of C is already in A_i, the tuple $\bar{b}_i$ inside the canonical definition of $\rho_i(C)$ can be chosen inside B. Since $B \preccurlyeq N$, we infer that for some matrix H in B whose columns each sum to zero,

$$B \models \bigwedge_{i<n} (\exists \bar{w}\ \left(\phi(\bar{w} + \rho_i(H)) \wedge \theta_i(\bar{w}, \bar{b}_i)\right)^\circ_i \tag{5}$$

and hence for each $i < n$ there is $\bar{c}'_i$ in A_i that satisfies the canonical definition of $\rho_i(C)$ and

$$A_i \models \phi(\bar{c}'_i + \rho_i(H)) \wedge \theta_i(\bar{c}'_i, \bar{b}_i). \tag{6}$$

So there is an automorphism of A_i taking $\bar{c}'_i$ to $\rho_i(C)$ and hence

$$A_i \models \phi(\rho_i(C) + \rho_i(H)). \tag{7}$$

Given the requirement $\Sigma(H) = \bar{0}$, this is exactly the condition for

$$A \models \phi(\kappa_0(C) + \mathbb{K}, \dots, \kappa_0(C) + \mathbb{K}). \tag{8}$$

Thus (a) in Fact 5.2 holds.

Next we prove (b) in Fact 5.2. Let ϕ/ψ be a Baur-Monk invariant and m a positive integer. We assume $|\phi(M)/\psi(M)| \geqslant m$ and we deduce $|\phi(A)/\psi(A)| \geqslant m$. By assumption there are elements $c_0, \dots, c_{m-1}$ of M such that

$$M \models \bigwedge_{j<m} \phi(c_k) \wedge \bigwedge_{j<k<m} \neg\psi(c_j - c_k). \tag{9}$$

We translate this into a statement about the M_i by using (2) of Lemma 5.5 for ϕ and (3) of Lemma 5.5 for ψ. Thus, there is an $n \times m$ matrix $C = (c_{ij})$ such that each row consists of elements of M_i, and there is an $n \times m$ matrix $\underline{H} = (\underline{h}_{ij})$ of elements of N with $\Sigma(\underline{H}) = \bar{0}$, such that for each $i < n$,

$$M_i \models \bigwedge_{j<m} \phi(c_{ij} + \underline{h}_{ij}) \tag{10}$$

but for every pair of indices (j,k) with $j < k < m$, **either** there is no $n \times 1$ column matrix $\underline{G} = (\underline{g}_i)$ over N such that for each $i < n$,

$$M_i \models \psi(c_{ij} - c_{ik} + \underline{g}_i)$$

or there is such a $\underline{G}$, but it satisfies

$$B \models \neg\psi(\Sigma(\underline{G})).$$

Let $\mathbb{G}$ be the set of pairs (j,k) where the 'either' case holds, and $\mathbb{H}$ the set of the remaining pairs (j,k) with $j < k < m$ (so that the 'or' case holds for these pairs). For each $(j,k) \in \mathbb{G}$ there is $i(j,k) < n$ such that

$$M_{i(j,k)} \models \forall z\, (P(z) \to \neg\psi(c_{i(j,k),j} - c_{i(j,k),k} + z)),$$

since otherwise we could refute the 'either' case by choosing each $\underline{g}_i$ separately.

To summarise the facts so far: There are a matrix $\underline{H}$ of elements of N, and for each pair $(j,k) \in \mathbb{H}$ there is a column matrix $\underline{G}^{jk} = (\underline{g}_i^{jk})$ of elements of N, such that the following hold:

(i) For each $i < n$

$$M_i \models \bigwedge_{j<m} \phi(c_{ij} + \underline{h}_{ij}),$$

(ii)
$$N \models \Sigma(\underline{H}) = \bar{0},$$

(iii) for each $(j,k) \in \mathbb{G}$,
$$M_{i(j,k)} \models \forall z\, (P(z) \to \neg\psi(c_{i(j,k),j} - c_{i(j,k),k} + z)),$$

(iv) for each $(j,k) \in \mathbb{H}$ and each $i < n$,
$$M_i \models \psi(c_{ij} - c_{ik} + \underline{g}_i^{jk})$$
and
$$N \models \neg\psi(\Sigma(\underline{G}^{jk}).$$

We can group the conditions here into $n+1$ conjuncts: one for each M_i and one for N. There is an implied quantifier $\exists C$ over all of this. But the requirements on the elements of C are now separate for each $i < n$, so we can distribute them between the separate conjuncts. By the Reduction Property we can state the conditions on each M_i as conditions on N. If for example 0 is $i(j,k)$ for just one pair in $(j,k) \in \mathbb{G}$, namely $(1,2)$, then the condition on M_0 will have the form:

$$M_0 \models \Big(\exists \bar{v}\, \Big(\bigwedge_{j<m} \phi(v_j + \underline{h}_{ij}) \wedge \forall z\, (P(z) \to \neg\psi(v_1 - v_2 + z)) \wedge \bigwedge_{(j,k)\in\mathbb{H}} \psi(v_j - v_k + \underline{g}_i^{jk})\Big)\Big)_i^{\circ}.$$

Writing χ_i for the conjunct that deals with M_i $(i < n)$ and χ_n for the conjunction of the statements in (ii) and (iv) that refer to N, the whole statement takes the form
$$N \models \exists \underline{H} (\exists \underline{G}^{jk})_{(j,k)\in\mathbb{H}}\ (\chi_0 \wedge \ldots \wedge \chi_n).$$
Since $B \preccurlyeq N$,
$$B \models \exists H (\exists G^{jk})_{(j,k)\in\mathbb{H}}\ (\chi_0 \wedge \ldots \wedge \chi_n).$$
Running the entire argument backwards, we find m elements in $\phi(A)$ such that none of their differences are in $\psi(A)$, and hence $|\phi(A)/\psi(A)| \geqslant m$ as required. □

References

[1] Gisela Ahlbrandt and Martin Ziegler, 'What's so special about $(\mathbb{Z}/4\mathbb{Z})^\omega$?', *Archive for Mathematical Logic* 31 (1991) 115–132.

[2] David Evans, Wilfrid Hodges and Ian Hodkinson, Automorphisms of bounded abelian groups, *Forum Mathematicum* 3 (1991) 523–541.

[3] Laszlo Fuchs, *Infinite Abelian Groups I*, Academic Press, New York 1970.

[4] Laszlo Fuchs, *Infinite Abelian Groups II*, Academic Press, New York 1973.

[5] Haim Gaifman, 'Operations on relational structures, functors and classes I', in *Proceedings of the Tarski Symposium*, ed. Leon Henkin et al., American Mathematical Society, Providence RI 1974, pp. 21–39.

[6] Wilfrid Hodges, *Model Theory*, Cambridge University Press, Cambridge 1993.

[7] Wilfrid Hodges and Anatoly Yakovlev, 'Relative categoricity in abelian groups II', *Annals of Pure and Applied Logic* 158 (2009) 203–231.

[8] I. Kaplansky and G. W. Mackey, 'A generalization of Ulm's theorem', *Summa Brasil. Math.* 2 (1951) 195–202.

[9] Saharon Shelah, Classification over a predicate II, in *Around Classification Theory of Models*, Lecture Notes in Mathematics 1182, Springer, Berlin 1986, pp. 47–90.

Herons Brook,
Sticklepath, Okehampton EX20 2PY