

UC Berkeley

UC Berkeley Previously Published Works

Title

On Holder-Brascamp-Lieb inequalities for torsion-free discrete Abelian groups

Permalink

<https://escholarship.org/uc/item/9kx6s2c7>

Authors

Christ, Michael
Demmel, James
Knight, Nicholas
et al.

Publication Date

2015-10-14

Peer reviewed

On Hölder-Brascamp-Lieb Inequalities for Torsion-Free Discrete Abelian Groups

Michael Christ*, James Demmel†, Nicholas Knight‡, Thomas Scanlon§ and Katherine Yelick¶

October 6, 2015

Abstract

Hölder-Brascamp-Lieb inequalities provide upper bounds for a class of multilinear expressions, in terms of L^p norms of the functions involved. They have been extensively studied for functions defined on Euclidean spaces. Bennett-Carbery-Christ-Tao have initiated the study of these inequalities for discrete Abelian groups and, in terms of suitable data, have characterized the set of all tuples of exponents for which such an inequality holds for specified data, as the convex polyhedron defined by a particular finite set of affine inequalities.

In this paper we advance the theory of such inequalities for torsion-free discrete Abelian groups in three respects. The optimal constant in any such inequality is shown to equal 1 whenever it is finite. An algorithm that computes the admissible polyhedron of exponents is developed. It is shown that nonetheless, existence of an algorithm that computes the full list of inequalities in the Bennett-Carbery-Christ-Tao description of the admissible polyhedron for all data, is equivalent to an affirmative solution of Hilbert's Tenth Problem over the rationals. That problem remains open.

1 Main Results

By a Hölder-Brascamp-Lieb inequality is typically meant an inequality of the form

$$\int_{\mathbb{R}^d} \prod_{j=1}^m f_j(L_j(x)) dm(x) \leq C \prod_{j=1}^m \|f_j\|_{L^{p_j}(\mathbb{R}^{d_j})} \quad (1.1)$$

where $L_j : \mathbb{R}^d \rightarrow \mathbb{R}^{d_j}$ are surjective linear mappings, the functions f_j are Lebesgue measurable and nonnegative, m denotes Lebesgue measure on \mathbb{R}^d , the exponents p_j belong to the natural range $[1, \infty]$, and C is some finite constant which is independent of the functions f_j but may depend on the dimensions d, d_j , the exponents p_j , and the mappings L_j . Fundamental examples include Hölder's inequality, Young's convolution inequality, and the Loomis-Whitney inequality. A substantial literature concerning this class of inequalities has developed.

In this paper we are concerned with analogous inequalities in which $\mathbb{R}^d, \mathbb{R}^{d_j}$ are replaced by discrete Abelian groups G, G_j respectively, where G is torsion-free. Lebesgue measure is replaced by counting measure. This situation was studied, without the restriction that G be torsion-free, in [3].

We have found these inequalities for $G = \mathbb{Z}^d$ to be relevant to the analysis of lower bounds for communication in a class of algorithms that arise, or may potentially arise, in computer science. The present paper, motivated by this connection, serves to develop the underlying analytic theory. The connection and its applications are explored in a companion paper [6].

*Mathematics Dept., Univ. of California, Berkeley (mchrist@berkeley.edu).

†Computer Science Div. and Mathematics Dept., Univ. of California, Berkeley (demmel@cs.berkeley.edu).

‡Computer Science Div., Univ. of California, Berkeley (knight@cs.berkeley.edu).

§Mathematics Dept., Univ. of California, Berkeley (scanlon@math.berkeley.edu).

¶Computer Science Div., Univ. of California, Berkeley and Lawrence Berkeley Natl. Lab. (yelick@cs.berkeley.edu).

Definition 1.1. An Abelian group HBL datum \mathfrak{G} is a 3-tuple

$$\mathfrak{G} = (G, (G_j), (\phi_j)) = (G, (G_1, \dots, G_m), (\phi_1, \dots, \phi_m))$$

where G and each G_j are finitely generated Abelian groups, G is torsion-free, and each $\phi_j: G \rightarrow G_j$ is a group homomorphism.

When this notation is employed, it is always implicitly assumed that $j \in \{1, 2, \dots, m\}$. By a Hölder-Brascamp-Lieb inequality associated to an Abelian group HBL datum, we mean one of the form

$$\sum_{x \in G} \prod_{j=1}^m f_j(\phi_j(x)) \leq C \prod_{j=1}^m \|f_j\|_{\ell^{1/s_j}} \quad \text{for all functions } 0 \leq f_j \in \ell^{1/s_j}(G_j) \quad (1.2)$$

where each $s = (s_1, \dots, s_m) \in [0, 1]^m$, and where $C < \infty$ may depend on $(G, (G_j), (\phi_j))$ and on s , but is independent of the functions f_j . The ℓ^p norm is $\|f\|_{\ell^p(G)} = (\sum_{x \in G} |f(x)|^p)^{1/p}$ for $p < \infty$, while $\|f\|_{\ell^\infty(G)} = \sup_{x \in G} |f(x)|$. We will often write $\|\cdot\|_p$ as shorthand for $\|\cdot\|_{\ell^p}$.

A certain convex polytope $\mathcal{P}(G, (G_j), (\phi_j)) \subseteq [0, 1]^m$ plays a central role in the theory.

Definition 1.2. For any Abelian group HBL datum $\mathfrak{G} = (G, (G_j), (\phi_j))$, $\mathcal{P}(\mathfrak{G}) = \mathcal{P}(G, (G_j), (\phi_j))$ denotes the set of all $s \in [0, 1]^m$ that satisfy

$$\text{rank}(H) \leq \sum_{j=1}^m s_j \text{rank}(\phi_j(H)) \quad \text{for every subgroup } H \leq G. \quad (1.3)$$

$\mathcal{P}(\mathfrak{G}) = \mathcal{P}(G, (G_j), (\phi_j))$ is the set of all $s \in \mathbb{R}^m$ specified by the inequalities $0 \leq s_j \leq 1$ for all j together with all inequalities $\sum_{j=1}^m s_j r_j \geq r$, where (r, r_1, \dots, r_m) ranges over all elements of $\{1, 2, \dots, \text{rank}(G)\} \times \prod_{j=1}^m \{0, 1, \dots, \text{rank}(\phi_j(G))\}$ for which there exists a finitely generated subgroup $H \leq G$ that satisfies $\text{rank}(H) = r$ and $\text{rank}(\phi_j(H)) = r_j$ for all j . Although infinitely many candidate subgroups H must potentially be considered in any calculation of $\mathcal{P}(\mathfrak{G})$, there exists a collection of fewer than $(\text{rank}(G) + 1)^{m+1}$ tuples (r, r_1, \dots, r_m) that suffices to generate all of the inequalities defining $\mathcal{P}(\mathfrak{G})$. Thus \mathcal{P} is a convex polytope with finitely many extreme points. It is equal to the convex hull of this finite set.

The first of our three main results states that the optimal constant C in such an inequality equals 1.

Theorem 1.3. For any Abelian group HBL datum $\mathfrak{G} = (G, (G_j), (\phi_j))$ and any $s \in \mathcal{P}(\mathfrak{G})$,

$$\sum_{x \in G} \prod_{j=1}^m f_j(\phi_j(x)) \leq \prod_{j=1}^m \|f_j\|_{1/s_j} \quad \text{for all functions } 0 \leq f_j \in \ell^{1/s_j}(G_j). \quad (1.4)$$

In combination with results previously known, this implies a more comprehensive statement.

Theorem 1.4. The following conditions are mutually equivalent, for any Abelian group HBL datum $\mathfrak{G} = (G, (G_j), (\phi_j))$ and any $s \in [0, 1]^m$.

1. $s \in \mathcal{P}(\mathfrak{G})$.
2. There exists a constant $C < \infty$ for which the Hölder-Brascamp-Lieb inequality (1.2) holds.
3. There exists $A < \infty$ such that

$$|E| \leq A \prod_{j=1}^m |\phi_j(E)|^{s_j} \quad \text{for all nonempty finite sets } E \subseteq G. \quad (1.5)$$

4. The Hölder-Brascamp-Lieb inequality (1.2) holds with $C = 1$.

5. Inequality (1.5) holds with $A = 1$.

Indeed, the first three of these conditions are proved to be equivalent in [3], in the more general setting where G is an arbitrary finitely generated Abelian group, possibly with torsion. In combination with this equivalence, Theorem 1.3 implies the mutual equivalence of all five conditions when G is torsion-free. \square

If G has torsion then the optimal constant is no longer equal to 1. Its value is determined for all finite Abelian group HBL data in [5]. In combination with Theorem 1.3, this allows determination of the optimal constant for all finitely generated Abelian group HBL data.

That (1.4) cannot hold for any $C < 1$ may be seen by fixing any point $x_0 \in G$ and defining f_j to be the indicator function of the singleton set $\{\phi_j(x_0)\}$.

Our second main theorem establishes an algorithm for the computation of $\mathcal{P}(\mathfrak{G})$. To have such an algorithm is desirable for applications [6] to computer science. While $\mathcal{P}(\mathfrak{G})$ is defined in terms of a finite list of inequalities, its definition refers to each of the subgroups of G in order to specify this list. Therefore the definition does not directly provide an effective way to calculate $\mathcal{P}(\mathfrak{G})$, unless G has rank 0 or 1.

Theorem 1.5. *There exists an algorithm that takes as input any Abelian group HBL datum $\mathfrak{G} = (G, (G_j), (\phi_j))$ and returns as output both a list of finitely many linear inequalities over \mathbb{Z} that jointly specify the associated polytope $\mathcal{P}(\mathfrak{G})$, and a list of all extreme points of $\mathcal{P}(\mathfrak{G})$.*

This algorithm is specified, and proved to be correct, in §3.

The polytope $\mathcal{P}(\mathfrak{G})$ is specified in Definition 1.2 by finitely many inequalities, all of which lie in a finite set that can be read off directly from the datum; for any subgroup $H \leq G$, $(r, r_1, \dots, r_m) = (\text{rank}(H), \text{rank}(\phi_1(H)), \dots, \text{rank}(\phi_m(H)))$ is an element of $\{0, 1, \dots, \text{rank}(G)\}^{m+1}$. But only some tuples $(r, r_1, \dots, r_m) \in \{0, 1, \dots, \text{rank}(G)\}^{m+1}$ are actually realized in this way by some subgroup. An algorithm that determines for each (r, r_1, \dots, r_m) whether there exists a subgroup $H \leq G$ satisfying $\text{rank}(H) = r$ and $\text{rank}(\phi_j(H)) = r_j$ for each $j \in \{1, 2, \dots, m\}$, would accomplish the first task of the algorithm of Theorem 1.5. The algorithm of Theorem 1.5 does not do this. Instead, it produces a finite sublist of certain (H, r, r_1, \dots, r_m) satisfying these conditions, along with a proof that the polytope specified by all the inequalities (1.3) specified by this sublist coincides with \mathcal{P} . Any inequalities (1.3) missing from the computed sublist are guaranteed to be redundant for the definition of $\mathcal{P}(\mathfrak{G})$.

We have not succeeded in devising an algorithm that computes the entire list of constraints (1.3) in the definition of $\mathcal{P}(\mathfrak{G})$. Our third main theorem asserts that the existence of such an algorithm is equivalent to an affirmative answer to a longstanding open question, Hilbert's Tenth Problem over the rationals.

Theorem 1.6. *There exists an effective algorithm for computing the set of constraints (1.3) defining $\mathcal{P}(\mathfrak{G})$ for any arbitrary Abelian group HBL datum \mathfrak{G} , if and only if there exists an effective algorithm to decide whether any given system of polynomial equations with rational coefficients has a rational solution.*

A simply computable alternative description of $\mathcal{P}(\mathfrak{G})$ is given by Carlen, Lieb, and Loss [4] for the case in which each target group G_j has rank equal to one.¹ Theorem 1.6 makes it clear that the theory is more complex in the general situation of arbitrary rank.

In a subsequent paper we will refine the results obtained here. Consider any Abelian group HBL datum $\mathfrak{G} = (G, (G_j), (\phi_j))$. Define $\mathcal{L}(\mathfrak{G})$ to be the smallest collection of subgroups of G that contains $\ker(\phi_j)$ for all indices j and is closed with respect to intersections and formation of sums. Define the polytope $\mathcal{P}_* = \mathcal{P}_*(\mathfrak{G})$ to be the set of all $s \in [0, 1]^m$ that satisfy (1.3) for every subgroup $H \in \mathcal{L}$. Then $\mathcal{P}_*(\mathfrak{G}) = \mathcal{P}(\mathfrak{G})$. Moreover, whereas the algorithm promised by Theorem 1.5 and specified in its proof below requires searching a list of all subgroups of G until halting, it suffices to search only a list of subspaces belonging to \mathcal{L} . These can be generated iteratively from the subspaces $\ker(\phi_j)$ by repeated formation of intersections and sums.

¹ [4] is concerned with continuum inequalities (1.1), but the description of $\mathcal{P}(\mathfrak{G})$ obtained carries over to the Abelian group context.

2 Upper bounds

In this section we prove Theorem 1.3. This entails the development of some auxiliary material. We have found it more convenient to work with vector spaces $G \otimes \mathbb{Q}$ than with general torsion-free Abelian groups. Therefore we formulate a version of the HBL inequalities in terms of vector spaces, and reduce Theorem 1.3 to its analogue in the vector space context.

Notation 2.1. *Let X be any set. For $p \in [1, \infty)$, $\ell^p(X)$ denotes the space of all p -power summable functions $f: X \rightarrow \mathbb{C}$, equipped with the norm*

$$\|f\|_p = \left(\sum_{x \in X} |f(x)|^p \right)^{1/p}.$$

$\ell^\infty(X)$ is the space of all bounded functions $f: X \rightarrow \mathbb{C}$, equipped with the norm $\|f\|_\infty = \sup_{x \in X} |f(x)|$.

These definitions apply even when X is uncountable: If $p \in [1, \infty)$, any function in $\ell^p(X)$ vanishes at all but countably many points in X ; and the ℓ^∞ -norm is still the supremum of $|f|$. In the discussion below, X will always be either a finitely generated Abelian group, or a finite-dimensional vector space over a field \mathbb{F} . If $s_j = 0$ then $1/s_j$ is interpreted as ∞ .

2.1 Preliminaries

The following result demonstrates there is no loss of generality in restricting attention to exponent tuples $s \in [0, 1]^m$, rather than $s \in [0, \infty)^m$, in Theorem 1.3. Its proof is simple, but is deferred to §2.4 where it is discussed in conjunction with a parallel result in which Abelian groups are replaced by vector spaces.

Proposition 2.2. *Let $s \in [0, \infty)^m$. Define $t \in [0, 1]^m$ by $t_j = \min(s_j, 1)$ for each index j . If (1.3) or (1.4) holds for s then it holds for t .*

If X is a measure space equipped with counting measure then whenever $0 < \alpha \leq \beta \leq \infty$, $\ell^\alpha(X) \subseteq \ell^\beta(X)$, and the inclusion mapping is a contraction. Therefore when s, t are related as above, inequality (1.4) for the exponent tuple t subsumes the inequality for s .

Proof of necessity of the hypothesis (1.3) for validity of the inequality (1.4). This necessity is proved in [3, Theorem 2.4]; the simple proof is included here for the sake of completeness. Indeed, consider any subgroup $H \leq G$. Let $r = \text{rank}(H)$. By definition of rank, there exists a set $\{e_i\}_{i=1}^r$ of elements of H such that for any coefficients $m_i, n_i \in \mathbb{Z}$, $\sum_{i=1}^r m_i e_i = \sum_{i=1}^r n_i e_i$ only if $m_i = n_i$ for all i . For any positive integer N define E_N to be the set of all elements of the form $\sum_{i=1}^r n_i e_i$, where each n_i ranges freely over $\{1, 2, \dots, N\}$. Then $|E_N| = N^r$.

On the other hand, for $j \in \{1, \dots, m\}$,

$$|\phi_j(E_N)| \leq A_j N^{\text{rank}(\phi_j(H))}, \tag{2.1}$$

where A_j is a finite constant which depends on ϕ_j , on the structure of H_j , and on the choice of $\{e_i\}$, but not on N . Indeed, it follows from the definition of rank that for each j it is possible to permute the indices i so that for each $i > \text{rank}(\phi_j(H))$ there exist integers k_i and $\kappa_{i,l}$ such that

$$k_i \phi_j(e_i) = \sum_{l=1}^{\text{rank}(\phi_j(H))} \kappa_{i,l} \phi_j(e_l).$$

The upper bound (2.1) follows from these relations.

Consider any large positive integer N . Define f_j to be the indicator function of $\phi_j(E_N)$. Then $\prod_j f_j \circ \phi_j \equiv 1$ on E_N . Therefore inequality (1.4), applied to these functions f_j , asserts that

$$N^{\text{rank}(H)} \leq \prod_{j=1}^m A_j^{s_j} N^{s_j \text{rank}(\phi_j(H))} = AN^{\sum_{j=1}^m s_j \text{rank}(\phi_j(H))}$$

where $A < \infty$ is independent of N . By letting N tend to infinity, we conclude that $\text{rank}(H) \leq \sum_{j=1}^m s_j \text{rank}(\phi_j(H))$, as was to be shown. \square

We show that Theorem 1.3 is a consequence of the special case in which all of the groups G_j are torsion-free.

Reduction of Theorem 1.3 to the case of torsion-free codomains. Let $\mathfrak{G} = (G, (G_j), (\phi_j))$ be an Abelian group HBL datum and let $s \in \mathcal{P}(\mathfrak{G})$. According to the structure theorem of finitely generated Abelian groups, each G_j is isomorphic to $\tilde{G}_j \oplus T_j$ where T_j is a finite group and \tilde{G}_j is torsion-free. Here \oplus denotes the direct sum of Abelian groups; $G' \oplus G''$ is the Cartesian product $G' \times G''$ equipped with the natural componentwise group law. Define $\pi_j: \tilde{G}_j \oplus T_j \rightarrow \tilde{G}_j$ to be the natural projection; thus $\pi_j(x, t) = x$ for $(x, t) \in \tilde{G}_j \times T_j$. Define $\tilde{\phi}_j = \pi_j \circ \phi_j: G \rightarrow \tilde{G}_j$.

If K is a subgroup of G_j , then $\text{rank}(\pi_j(K)) = \text{rank}(K)$ since the kernel T_j of π_j is a finite group. Therefore for any subgroup $H \leq G$, $\text{rank}(\tilde{\phi}_j(H)) = \text{rank}(\phi_j(H))$. Therefore whenever $(G, (G_j), (\phi_j))$ and s together satisfy the hypothesis (1.3), so do $(G, (\tilde{G}_j), (\tilde{\phi}_j))$ and s .

Under this hypothesis, consider any m -tuple $f = (f_j)$ of nonnegative functions with $f_j \in \ell^{1/s_j}(G_j)$, and for each j , define $\tilde{f}_j \in \ell^{1/s_j}(\tilde{G}_j)$ by

$$\tilde{f}_j(y) = \max_{t \in T_j} f_j(y, t), \text{ for } y \in \tilde{G}_j.$$

For any $x \in G$, $f_j(\phi_j(x)) \leq \tilde{f}_j(\tilde{\phi}_j(x))$. Consequently $\prod_{j=1}^m f_j(\phi_j(x)) \leq \prod_{j=1}^m \tilde{f}_j(\tilde{\phi}_j(x))$.

We are assuming validity of Theorem 1.3 in the torsion-free case. Its conclusion asserts that

$$\sum_{x \in G} \prod_{j=1}^m \tilde{f}_j(\tilde{\phi}_j(x)) \leq \prod_{j=1}^m \| \tilde{f}_j \|_{1/s_j}.$$

For each j , $\| \tilde{f}_j \|_{1/s_j} \leq \| f_j \|_{1/s_j}$. This is obvious when $s_j = 0$. If $s_j \neq 0$ then

$$\| \tilde{f}_j \|_{1/s_j}^{1/s_j} = \sum_{y \in \tilde{G}_j} \tilde{f}_j(y)^{1/s_j} = \sum_{y \in \tilde{G}_j} \max_{t \in T_j} f_j(y, t)^{1/s_j} \leq \sum_{y \in \tilde{G}_j} \sum_{t \in T_j} f_j(y, t)^{1/s_j} = \| f_j \|_{1/s_j}^{1/s_j}.$$

Combining these inequalities gives the HBL inequality (1.4). \square

2.2 Vector space HBL inequalities

In this subsection we formulate HBL inequalities for vector spaces over arbitrary fields. This vector space framework is more convenient for our method of proof, which involves quotients; these are awkward in the Abelian group context since quotients of torsion-free groups need not be torsion-free, while quotients of vector spaces suffer no such calamity. We continue to endow all measure spaces with counting measure, even vector spaces over \mathbb{R} or \mathbb{C} . Theorem 2.5, below, is the analogue of Theorem 1.3 for vector spaces. We will show how Theorem 2.5 for $\mathbb{F} = \mathbb{Q}$ implies Theorem 1.3, and then prove Theorem 2.5. The case of fields other than \mathbb{Q} is not the real thrust of our investigation, but since it is treated by our analysis with no extra effort we have formulated the results for general fields.

All vector spaces considered in this paper are assumed to be finite-dimensional.

Notation 2.3. $\dim(V)$ will denote the dimension of a vector space V over a field \mathbb{F} . The notation $W \leq V$ indicates that W is a subspace of V , and $W < V$ indicates that W is a proper subspace.

Definition 2.4. A vector space HBL datum is a 3-tuple

$$\mathbb{V} = (V, (V_j), (\phi_j)) = (V, (V_1, \dots, V_m), (\phi_1, \dots, \phi_m))$$

where V and V_j are finite-dimensional vector spaces over a field \mathbb{F} and $\phi_j: V \rightarrow V_j$ is an \mathbb{F} -linear map.

It will always be understood that $j \in \{1, 2, \dots, m\}$, although the quantity m will not usually be indicated.

To any Abelian group HBL datum $(G, (G_j), (\phi_j))$ with all G_j (as well as G) torsion-free we associate a vector space HBL datum, as follows. Firstly, G_j is isomorphic to \mathbb{Z}^{d_j} where $d_j = \text{rank}(G_j)$. Likewise G is isomorphic to \mathbb{Z}^d where $d = \text{rank}(G)$. Define homomorphisms $\tilde{\phi}_j: \mathbb{Z}^d \rightarrow \mathbb{Z}^{d_j}$ by composing ϕ_j with these isomorphisms. $(G, (G_j), (\phi_j))$ is thereby identified with $(\mathbb{Z}^d, (\mathbb{Z}^{d_j}), (\tilde{\phi}_j))$. Defining scalar multiplication in the natural way (i.e., treating \mathbb{Z}^d and \mathbb{Z}^{d_j} as \mathbb{Z} -modules), represent each \mathbb{Z} -linear map $\tilde{\phi}_j$ by a matrix with integer entries. Secondly, let $\mathbb{F} = \mathbb{Q}$. Regard \mathbb{Z}^d and \mathbb{Z}^{d_j} as subsets of \mathbb{Q}^d and of \mathbb{Q}^{d_j} , respectively, via the inclusion of \mathbb{Z} into \mathbb{Q} . Define the vector space HBL datum $(V, (V_j), (\psi_j))$ by setting $V = \mathbb{Q}^d$, $V_j = \mathbb{Q}^{d_j}$, and letting $\psi_j: \mathbb{Q}^d \rightarrow \mathbb{Q}^{d_j}$ be the \mathbb{Q} -linear map represented by the same integer matrix as $\tilde{\phi}_j$.

This construction requires certain arbitrary choices, so $(V, (V_j), (\psi_j))$ is properly referred to as *an* associated vector space HBL datum.

Theorem 2.5. Let $(V, (V_j), (\phi_j))$ be a vector space HBL datum, and let $s \in [0, 1]^m$. If

$$\dim(W) \leq \sum_{j=1}^m s_j \dim(\phi_j(W)) \quad \text{for all subspaces } W \leq V \quad (2.2)$$

then

$$\sum_{x \in V} \prod_{j=1}^m f_j(\phi_j(x)) \leq \prod_{j=1}^m \|f_j\|_{1/s_j} \quad \text{for all functions } 0 \leq f_j \in \ell^{1/s_j}(V_j). \quad (2.3)$$

In particular,

$$|E| \leq \prod_{j=1}^m |\phi_j(E)|^{s_j} \quad \text{for all nonempty finite sets } E \subseteq V. \quad (2.4)$$

Conversely, if (2.4) holds for $s \in [0, 1]^m$, and if $\mathbb{F} = \mathbb{Q}$ or if \mathbb{F} is finite, then s satisfies (2.2).

The conclusions (1.4) and (2.3) remain valid for functions f_j without the requirement that the ℓ^{1/s_j} norms are finite, under the convention that the product $0 \cdot \infty$ is interpreted as zero whenever it arises. For then if any f_j has infinite norm, then either the right-hand side is infinite, or both sides are zero; in either case, the inequality holds.

Reduction of Theorem 1.3 in the case of torsion-free codomains to Theorem 2.5. Let $\mathfrak{G} = (G, (G_j), (\phi_j))$ be an Abelian group HBL datum, and let $s \in \mathcal{P}(\mathfrak{G})$; furthermore suppose that each group G_j is torsion-free. Let $\mathbb{V} = (V, (V_j), (\psi_j))$ be a vector space HBL datum associated to \mathfrak{G} by the construction specified above. We claim that $s \in \mathcal{P}(\mathbb{V})$. Indeed, given any subspace $W \leq V$, there exists a basis S for W over \mathbb{Q} which consists of elements of \mathbb{Z}^d . Define $H \leq \mathbb{Z}^d$ to be the subgroup generated by S (over \mathbb{Z}). Then $\text{rank}(H) = \dim(W)$. Moreover, $\psi_j(W)$ equals the span over \mathbb{Q} of $\tilde{\phi}_j(H)$, and $\dim(\psi_j(W)) = \text{rank}(\tilde{\phi}_j(H))$. The hypothesis $\text{rank}(H) \leq \sum_{j=1}^m s_j \text{rank}(\tilde{\phi}_j(H))$ is therefore equivalently written as $\dim(W) \leq \sum_{j=1}^m s_j \dim(\psi_j(W))$, which is (2.2) for W . Since this holds for every subspace W , $s \in \mathcal{P}(\mathbb{V})$.

Consider the m -tuple $f = (f_j)$ corresponding to any inequality in (1.4) for $(\mathbb{Z}^d, (\mathbb{Z}^{d_j}), (\tilde{\phi}_j))$ and s , and for each j define F_j by $F_j(y) = f_j(y)$ if $y \in \mathbb{Z}^{d_j} \subseteq \mathbb{Q}^{d_j}$, and $F_j(y) = 0$ otherwise. By Theorem 2.5,

$$\sum_{x \in \mathbb{Z}^d} \prod_{j=1}^m f_j(\tilde{\phi}_j(x)) = \sum_{x \in \mathbb{Q}^d} \prod_{j=1}^m F_j(\psi_j(x)) \leq \prod_{j=1}^m \|F_j\|_{1/s_j} = \prod_{j=1}^m \|f_j\|_{1/s_j}.$$

Thus the conclusion (1.4) of Theorem 1.3 is satisfied. \square

Conversely, it is possible to derive Theorem 2.5 for $\mathbb{F} = \mathbb{Q}$ from the special case of Theorem 1.3 in which $G = \mathbb{Z}^d$ and $G_j = \mathbb{Z}^{d_j}$ by similar reasoning, using multiplication by large integers to clear denominators.

This reasoning in this reduction establishes

Lemma 2.6. *Let $(G, (G_j), (\phi_j))$ be an Abelian group HBL datum with torsion-free codomains G_j , let $(\mathbb{Z}^d, (\mathbb{Z}^{d_j}), (\tilde{\phi}_j))$ be the associated datum specified above, and let $(\mathbb{Q}^d, (\mathbb{Q}^{d_j}), (\psi_j))$ be an associated vector space HBL datum. Then*

$$\mathcal{P}(G, (G_j), (\phi_j)) = \mathcal{P}(\mathbb{Z}^d, (\mathbb{Z}^{d_j}), (\tilde{\phi}_j)) = \mathcal{P}(\mathbb{Q}^d, (\mathbb{Q}^{d_j}), (\psi_j)).$$

2.3 On groups G with torsion

In this subsection we temporarily relax the requirement in the definition of an Abelian group HBL datum that the finitely generated Abelian group G be torsion-free. We call such a more general $(G, (G_j), (\phi_j))$ an HBL datum with torsion.

It was shown in [3, Theorem 2.4] that for an HBL datum with torsion, the rank conditions (1.3) are necessary and sufficient for the existence of some finite constant C such that (1.2) holds. A consequence of Theorem 1.3 is a concrete upper bound for the constant C in these inequalities. The torsion subgroup $T(G)$ of G is the (finite) set of all elements $x \in G$ for which there exists $0 \neq n \in \mathbb{Z}$ such that $nx = 0$.

Theorem 2.7. *Consider $(G, (G_j), (\phi_j))$, an HBL datum with torsion, and $s \in [0, 1]^m$. If (1.3) holds, then (1.2) holds with $C = |T(G)|$. In particular,*

$$|E| \leq |T(G)| \cdot \prod_{j=1}^m |\phi_j(E)|^{s_j} \quad \text{for all nonempty finite sets } E \subseteq G. \quad (2.5)$$

Conversely, if (2.5) holds for $s \in [0, 1]^m$, then s satisfies (1.3).

Proof. To prove (1.2), express G isomorphically as $\tilde{G} \oplus T(G)$ where $\tilde{G} \leq G$ is torsion-free. Thus arbitrary elements $x \in G$ are expressed as $x = (\tilde{x}, t)$ with $\tilde{x} \in \tilde{G}$ and $t \in T(G)$. Then $(\tilde{G}, (G_j), (\phi_j|_{\tilde{G}}))$ is an Abelian group HBL datum (with \tilde{G} torsion-free) to which Theorem 1.3 can be applied. Consider any $t \in T(G)$. Define $g_j: G_j \rightarrow [0, \infty)$ by $g_j(x_j) = f_j(x_j + \phi_j(0, t))$. Then $f_j(\phi_j(y, t)) = f_j(\phi_j(y, 0) + \phi_j(0, t)) = g_j(\phi_j(y, 0))$, so

$$\sum_{y \in \tilde{G}} \prod_{j=1}^m f_j(\phi_j(y, t)) = \sum_{y \in \tilde{G}} \prod_{j=1}^m g_j(\phi_j(y, 0)) \leq \prod_{j=1}^m \|g_j|_{\phi_j(\tilde{G})}\|_{1/s_j} \leq \prod_{j=1}^m \|f_j\|_{1/s_j}.$$

The first inequality is an application of Theorem 1.3. Summation with respect to $t \in T(G)$ gives the required bound.

To show necessity, we consider just the inequalities (2.5) corresponding to the subsets $E \subseteq \tilde{G}$ and follow the proof of the converse of Theorem 1.3, except substituting $A|T(G)|$ for A . \square

The factor $|T(G)|$ cannot be improved if the groups G_j are torsion free, or more generally if $T(G)$ is contained in the intersection of the kernels of all the homomorphisms ϕ_j ; this is seen by considering $E = T(G)$. However, it is not optimal, in general. The optimal bound, for arbitrary finitely generated Abelian groups, is determined in a paper [5] that builds on the present one.

2.4 Polytopes \mathcal{P} for vector space HBL data

Definition 2.8. For any vector space HBL datum $\mathbb{V} = (V, (V_j), (\phi_j))$, $\mathcal{P}(\mathbb{V})$ denotes the set of all $s \in [0, 1]^m$ that satisfy (2.2).

Now we prove Proposition 2.2, which asserts that there is no loss of information in restricting the discussion to exponents $s_j \leq 1$.

Proof of Proposition 2.2. Proposition 2.2 was formulated in terms of Abelian group HBL data and inequalities. Here we show the corresponding result for vector space HBL data and inequalities, obtaining the version for Abelian groups as a direct consequence.

Suppose that a vector space HBL datum $(V, (V_j), (\phi_j))$ and $s \in [0, \infty)^m$ satisfy (2.2), and suppose that $s_k > 1$ for some k . Define $t \in [0, \infty)^m$ by $t_j = s_j$ for $j \neq k$, and $t_k = 1$. We claim that t continues to satisfy (2.2).

Consider any subspace $W \leq V$. In order to verify (2.2) for W and t , define $W' = W \cap \ker(\phi_k)$. Choose a supplement U for W' in W ; that is, $W = W' + U$ and $W' \cap U = \{0\}$. Then by (2.2) for s ,

$$\dim(W') \leq \sum_{j=1}^m s_j \dim(\phi_j(W')) = s_k \cdot 0 + \sum_{j \neq k} s_j \dim(\phi_j(W')) = \sum_{j \neq k} t_j \dim(\phi_j(W')).$$

Since ϕ_k is injective on U and $t_k = 1$, $\dim(U) = t_k \dim(\phi_k(U))$. Therefore

$$\begin{aligned} \dim(W) &= \dim(U) + \dim(W') \\ &= t_k \dim(\phi_k(U)) + \dim(W') \\ &\leq t_k \dim(\phi_k(U)) + \sum_{j \neq k} t_j \dim(\phi_j(W')) \\ &\leq t_k \dim(\phi_k(W)) + \sum_{j \neq k} t_j \dim(\phi_j(W)) \\ &= \sum_{j=1}^m t_j \dim(\phi_j(W)). \end{aligned}$$

Given an m -tuple s with multiple components $s_k > 1$, we argue by induction on the number of such indices k , with the result just proved serving as the induction step.

Our desired conclusion concerning (1.3) follows from Lemma 2.6, by considering the associated vector space HBL datum $(\mathbb{Q}^d, (\mathbb{Q}^{d_j}), (\psi_j))$ and noting that the lemma was established without assuming $s_j \leq 1$.

Next, we show the result concerning (1.4): If (1.4) holds for some $s \in [0, \infty)^m$, then it also holds for t , where $t_j = \min(s_j, 1)$ for all $j \in \{1, 2, \dots, m\}$. Consider the following less structured situation. Let X, X_1, \dots, X_m be sets and $\phi_j: X \rightarrow X_j$ be functions for $j \in \{1, \dots, m\}$. Let $s \in [0, \infty)^m$ with some $s_k > 1$, and suppose that $|E| \leq \prod_{j=1}^m |\phi_j(E)|^{s_j}$ for any finite nonempty subset $E \subseteq X$. Consider any such set E . For each $y \in \phi_k(E)$, let $E_y = \phi_k^{-1}(y) \cap E$, the preimage of y under $\phi_k|_E$; thus $|\phi_k(E_y)| = 1$. By assumption, $|E_y| \leq \prod_{j=1}^m |\phi_j(E_y)|^{s_j}$, so it follows that

$$|E_y| \leq \prod_{j \neq k} |\phi_j(E_y)|^{s_j} \leq \prod_{j \neq k} |\phi_j(E)|^{s_j} = \prod_{j \neq k} |\phi_j(E)|^{t_j}.$$

Since E can be written as the union of disjoint sets $\bigcup_{y \in \phi_k(E)} E_y$, we obtain

$$|E| = \sum_{y \in \phi_k(E)} |E_y| \leq \sum_{y \in \phi_k(E)} \prod_{j \neq k} |\phi_j(E)|^{t_j} = |\phi_k(E)| \cdot \prod_{j \neq k} |\phi_j(E)|^{t_j} = \prod_{j=1}^m |\phi_j(E)|^{t_j}.$$

Given an m -tuple s with multiple components $s_k > 1$, we can consider each separately and apply the same reasoning to obtain an m -tuple $t \in [0, 1]^m$ with $t_j = \min(s_j, 1)$ for all j . By picking $(X, (X_j), (\phi_j)) := (\mathbb{Z}^d, (\mathbb{Z}^{d_j}), (\phi_j))$, we now conclude that (1.5) holds for t . By Theorem 1.4, (1.5) implies (1.4); the proof of that theorem given below is independent of Proposition 2.2 so this concludes its proof.

Similar conclusions can be obtained for (1.5) and (2.4). \square

2.5 Interpolation between extreme points of \mathcal{P}

To complete the proof of Theorem 2.5, we must show that (2.3) holds for every $s \in \mathcal{P}$. We next prove that if (2.3) holds at each extreme point s of \mathcal{P} , then it holds for all $s \in \mathcal{P}$. In Section 2.6, we will show that for any extreme point s of \mathcal{P} , the hypothesis (2.2) can be restated in a special form. Finally, in Section 2.8 we will use this reformulation to prove (2.3) for extreme points.

Let $(X_j, \mathcal{A}_j, \mu_j)$ be measure spaces for $j \in \{1, 2, \dots, m\}$, where each μ_j is a nonnegative measure on the σ -algebra \mathcal{A}_j . Let S_j be the set of all simple functions $f_j: X_j \rightarrow \mathbb{C}$. Thus S_j is the set of all $f_j: X_j \rightarrow \mathbb{C}$ that can be expressed in the form $\sum_i c_i 1_{E_i}$ where $c_i \in \mathbb{C}$, $E_i \in \mathcal{A}_j$, $\mu_j(E_i) < \infty$, and the sum extends over finitely many indices i .

Let $T: \prod_{j=1}^m \mathbb{C}^{X_j} \rightarrow \mathbb{C}$ be a multilinear map; i.e., for any m -tuple $f \in \prod_{j=1}^m \mathbb{C}^{X_j}$ where $f_k = c_0 f_{k,0} + c_1 f_{k,1}$ for $c_0, c_1 \in \mathbb{C}$ and $f_{k,0}, f_{k,1} \in \mathbb{C}^{X_k}$,

$$T(f) = c_0 T(f_1, \dots, f_{k-1}, f_{k,0}, f_{k+1}, \dots, f_m) + c_1 T(f_1, \dots, f_{k-1}, f_{k,1}, f_{k+1}, \dots, f_m).$$

One multilinear extension of the Riesz-Thorin theorem states the following (see, e.g., [1]).

Proposition 2.9 (Multilinear Riesz-Thorin theorem). *Suppose that $p_0 = (p_{j,0}), p_1 = (p_{j,1}) \in [1, \infty]^m$. Suppose that there exist $A_0, A_1 \in [0, \infty)$ such that*

$$|T(f)| \leq A_0 \prod_{j=1}^m \|f_j\|_{p_{j,0}} \quad \text{and} \quad |T(f)| \leq A_1 \prod_{j=1}^m \|f_j\|_{p_{j,1}} \quad \text{for all } f \in \prod_{j=1}^m S_j.$$

For each $\theta \in (0, 1)$ define exponents $p_{j,\theta}$ by

$$\frac{1}{p_{j,\theta}} := \frac{\theta}{p_{j,0}} + \frac{1-\theta}{p_{j,1}}.$$

Then for each $\theta \in (0, 1)$,

$$|T(f)| \leq A_0^\theta A_1^{1-\theta} \prod_{j=1}^m \|f_j\|_{p_{j,\theta}} \quad \text{for all } f \in \prod_{j=1}^m S_j.$$

Here $\|f_j\|_p = \|f_j\|_{L^p(X_j, \mathcal{A}_j, \mu_j)}$.

In the context of Theorem 2.5 with vector space HBL datum $(V, (V_j), (\phi_j))$, we consider the multilinear map

$$T(f) := \sum_{x \in V} \prod_{j=1}^m f_j(\phi_j(x))$$

representing the left-hand side in (2.3).

Lemma 2.10. *If (2.3) holds for every extreme point of \mathcal{P} , then it holds for every $s \in \mathcal{P}$.*

Proof. For any $\tilde{f} \in \prod_{j=1}^m S_j$, we define another m -tuple f where for each j , $f_j = |\tilde{f}_j|$ is a nonnegative simple function. By hypothesis, the inequality in (2.3) corresponding to f holds at every extreme point s of \mathcal{P} , giving

$$\left|T(\tilde{f})\right| \leq \sum_{x \in V} \prod_{j=1}^m \left| \tilde{f}_j(\phi_j(x)) \right| = \sum_{x \in V} \prod_{j=1}^m f_j(\phi_j(x)) \leq \prod_{j=1}^m \|f_j\|_{1/s_j} = \prod_{j=1}^m \|\tilde{f}_j\|_{1/s_j}.$$

As a consequence of Proposition 2.9 (with constants $A_i = 1$), and the fact that any $s \in \mathcal{P}$ is a finite convex combination of the extreme points, this expression holds for any $s \in \mathcal{P}$. For any nonnegative function F_j (e.g., in $\ell^{1/s_j}(V_j)$), there is an increasing sequence of nonnegative simple functions f_j whose (pointwise) limit is F_j . Consider the m -tuple $F = (F_j)$ corresponding to any inequality in (2.3), and consider a sequence of m -tuples f which converge to F ; then $\prod_{j=1}^m f_j$ also converges to $\prod_{j=1}^m F_j$. So by the monotone convergence theorem, the summations on both sides of the inequality converge as well. \square

2.6 Critical subspaces and extreme points

Let $\mathbb{V} = (V, (V_j), (\phi_j))$ be an arbitrary vector space HBL datum. Let $\mathcal{P} = \mathcal{P}(\mathbb{V})$ continue to denote the set of all $s \in [0, 1]^m$ that satisfy (2.2). The following key definition appears in [2] and [3].

Definition 2.11. *Consider any $s \in [0, 1]^m$. A subspace $W \leq V$ satisfying $\dim(W) = \sum_{j=1}^m s_j \dim(\phi_j(W))$ is said to be a critical subspace with respect to s ; one satisfying $\dim(W) \leq \sum_{j=1}^m s_j \dim(\phi_j(W))$ is said to be subcritical with respect to s ; and a subspace satisfying $\dim(W) > \sum_{j=1}^m s_j \dim(\phi_j(W))$ is said to be supercritical with respect to s . W is said to be strictly subcritical with respect to s if $\dim(W) < \sum_{j=1}^m s_j \dim(\phi_j(W))$.*

In this language, the conditions (2.2) assert that that every subspace W of V , including $\{0\}$ and V itself, is subcritical; equivalently, there are no supercritical subspaces. We may sometimes omit the phrase “with respect to s ”, but these notions are always relative to some implicitly or explicitly specified tuple s of exponents.

The goal of Section 2.6 is to establish the following:

Proposition 2.12. *Let s be an extreme point of $\mathcal{P}(\mathbb{V})$. Then some subspace $\{0\} < W < V$ is critical with respect to s , or $s \in \{0, 1\}^m$.*

These two possibilities are not mutually exclusive.

Lemma 2.13. *If s is an extreme point of $\mathcal{P}(\mathbb{V})$, and if i is an index for which $s_i \notin \{0, 1\}$, then $\dim(\phi_i(V)) \neq 0$.*

Proof. Suppose $\dim(\phi_i(V)) = 0$. If $t \in [0, 1]^m$ satisfies $t_j = s_j$ for all $j \neq i$, then $\sum_{j=1}^m t_j \dim(\phi_j(W)) = \sum_{j=1}^m s_j \dim(\phi_j(W))$ for all subspaces $W \leq V$, so $t \in \mathcal{P}(\mathbb{V})$ as well. If $s_i \notin \{0, 1\}$, then this contradicts the assumption that s is an extreme point of $\mathcal{P}(\mathbb{V})$. \square

Lemma 2.14. *Let s be an extreme point of $\mathcal{P}(\mathbb{V})$. Suppose that no subspace $\{0\} < W \leq V$ is critical with respect to s . Then $s \in \{0, 1\}^m$.*

Proof. Membership in $\mathcal{P}(\mathbb{V})$ is decided by finitely many affine inequalities. The hypothesis of the lemma means that s satisfies each of these with strict inequality, except for the inequality $0 \leq \sum_j s_j \cdot 0$ that arises from $W = \{0\}$. Consequently any $t \in [0, 1]^m$ that is sufficiently close to s also belongs to $\mathcal{P}(\mathbb{V})$. If $s_i \notin \{0, 1\}$ for some index i , then any $t \in [0, 1]^m$ with $t_j = s_j$ for all $j \neq i$ and t_i sufficiently close to s_i belongs to $t \in \mathcal{P}(\mathbb{V})$. This contradicts the assumption that s is an extreme point. \square

Lemma 2.15. *Let s be an extreme point of $\mathcal{P}(\mathbb{V})$. Suppose that no subspace $\{0\} < W < V$ is critical with respect to s . Then there exists at most one index i for which $s_i \notin \{0, 1\}$.*

Proof. Suppose to the contrary that there were to exist distinct indices k, l such that neither of s_k, s_l belongs to $\{0, 1\}$. By Lemma 2.13, both $\phi_k(V)$ and $\phi_l(V)$ have positive dimensions. For $\varepsilon \in \mathbb{R}$ define t by $t_j = s_j$ for all $j \notin \{k, l\}$,

$$t_k = s_k + \varepsilon \dim(\phi_l(V)) \text{ and } t_l = s_l - \varepsilon \dim(\phi_k(V)).$$

Whenever $|\varepsilon|$ is sufficiently small, $t \in [0, 1]^m$. Moreover, V remains subcritical with respect to t . If $|\varepsilon|$ is sufficiently small, then every subspace $\{0\} < W < V$ remains strictly subcritical with respect to t , because the set of all parameters $(\dim(W), \dim(\phi_1(W)), \dots, \dim(\phi_m(W)))$ which arise, is finite. Thus $t \in \mathcal{P}(\mathbb{V})$ for all sufficiently small $|\varepsilon|$. Therefore s is not an extreme point of $\mathcal{P}(\mathbb{V})$. \square

Lemma 2.16. *Let $s \in [0, 1]^m$ be an extreme point of $\mathcal{P}(\mathbb{V})$. Suppose that V is critical with respect to s . Suppose that there exists exactly one index $i \in \{1, 2, \dots, m\}$ for which $s_i \notin \{0, 1\}$. Then V has a subspace that is supercritical with respect to s .*

Proof. By Lemma 2.13, $\dim(\phi_i(V)) > 0$. Let K be the set of all indices k for which $s_k = 1$. The hypothesis that V is critical means that

$$\dim(V) = s_i \dim(\phi_i(V)) + \sum_{k \in K} s_k \dim(\phi_k(V)).$$

Since $s_i > 0$ and $\dim(\phi_i(V)) > 0$,

$$\sum_{k \in K} \dim(\phi_k(V)) = \sum_{k \in K} s_k \dim(\phi_k(V)) = \dim(V) - s_i \dim(\phi_i(V)) < \dim(V).$$

Consider the subspace $W \leq V$ defined by

$$W = \bigcap_{k \in K} \ker(\phi_k);$$

this intersection is interpreted to be $W = V$ if the index set K is empty. W necessarily has positive dimension. Indeed, W is the kernel of the map $\psi: V \rightarrow \bigoplus_{k \in K} \phi_k(V)$, defined by $\psi(x) = (\phi_k(x) : k \in K)$, where \bigoplus denotes the direct sum of vector spaces. The image of ψ is a subspace of $\bigoplus_{k \in K} \phi_k(V)$, a vector space whose dimension $\sum_{k \in K} \dim(\phi_k(V))$ is strictly less than $\dim(V)$. Therefore $\ker(\psi) = W$ has dimension greater than or equal to $\dim(V) - \sum_{k \in K} \dim(\phi_k(V)) > 0$. Since $\phi_k(W) = \{0\}$ for all $k \in K$,

$$\begin{aligned} \sum_{j=1}^m s_j \dim(\phi_j(W)) &= s_i \dim(\phi_i(W)) + \sum_{k \in K} \dim(\phi_k(W)) \\ &= s_i \dim(\phi_i(W)) \\ &\leq s_i \dim(W). \end{aligned}$$

Since $s_i < 1$ and $\dim(W) > 0$, $s_i \dim(\phi_i(W))$ is strictly less than $\dim(W)$, whence W is supercritical. \square

Proof of Proposition 2.12. Suppose that there exists no critical subspace $\{0\} < W < V$. By Lemma 2.14, either $s \in \{0, 1\}^m$ — in which case the proof is complete — or V is critical. By Lemma 2.15, there can be at most one index i for which $s_i \notin \{0, 1\}$. By Lemma 2.16, for critical V , the existence of one single such index i implies the presence of some supercritical subspace, contradicting the main hypothesis of Proposition 2.12. Thus again, $s \in \{0, 1\}^m$. \square

2.7 Factorization of HBL data

Notation 2.17. Suppose V, V' are finite-dimensional vector spaces over a field \mathbb{F} , and $\phi: V \rightarrow V'$ is an \mathbb{F} -linear map. For any subspace $W \leq V$, $\phi|_W: W \rightarrow \phi(W)$ denotes the restriction of ϕ to W , also a \mathbb{F} -linear map. V/W denotes the quotient of V by W , a finite-dimensional vector space. Thus $x+W = x'+W$ if and only if $x - x' \in W$. Every subspace of V/W can be written as U/W for some $W \leq U \leq V$.

The quotient space $V'/\phi(W)$, and the quotient linear map $[\phi]: V/W \ni x + W \mapsto \phi(x) + \phi(W) \in V'/\phi(W)$, are likewise defined.

In this quotient situation, it is elementary that for any $U/W \leq V/W$, $[\phi](U/W) = \phi(U)/\phi(W)$.

Definition 2.18. Let $(V, (V_j), (\phi_j))$ be a vector space HBL datum. To any subspace $W \leq V$ are associated the two vector space HBL data

$$\begin{cases} \mathbb{V}_W = (W, (\phi_j(W)), (\phi_j|_W)) \\ \mathbb{V}_{V/W} = (V/W, (V_j/\phi_j(W)), ([\phi_j])). \end{cases} \quad (2.6)$$

Lemma 2.19. Let $\mathbb{V} = (V, (V_j), (\phi_j))$ be a vector space HBL datum. For any subspace $W \leq V$,

$$\mathcal{P}(\mathbb{V}_W) \cap \mathcal{P}(\mathbb{V}_{V/W}) \subseteq \mathcal{P}(\mathbb{V}). \quad (2.7)$$

Proof. Consider any subspace $U \leq V$ and any $s \in \mathcal{P}(\mathbb{V}_W) \cap \mathcal{P}(\mathbb{V}_{V/W})$. Then

$$\begin{aligned} \dim(U) &= \dim((U+W)/W) + \dim(U \cap W) \\ &\leq \sum_{j=1}^m s_j \dim([\phi_j]((U+W)/W)) + \sum_{j=1}^m s_j \dim(\phi_j(U \cap W)) \\ &= \sum_{j=1}^m s_j \dim(\phi_j(U+W)/\phi_j(W)) + \sum_{j=1}^m s_j \dim(\phi_j(U \cap W)) \\ &= \sum_{j=1}^m s_j (\dim(\phi_j(U+W)) - \dim(\phi_j(W))) + \sum_{j=1}^m s_j \dim(\phi_j(U \cap W)) \\ &= \sum_{j=1}^m s_j (\dim(\phi_j(U) + \phi_j(W)) + \dim(\phi_j(U \cap W)) - \dim(\phi_j(W))) \\ &\leq \sum_{j=1}^m s_j (\dim(\phi_j(U) + \phi_j(W)) + \dim(\phi_j(U) \cap \phi_j(W)) - \dim(\phi_j(W))) \\ &= \sum_{j=1}^m s_j \dim(\phi_j(U)). \end{aligned}$$

The last inequality is a consequence of the inclusions $\phi_j(U \cap W) \subseteq \phi_j(U) \cap \phi_j(W)$. The last equality uses the relation $\dim(A) + \dim(B) = \dim(A+B) + \dim(A \cap B)$, which holds for any subspaces A, B of a vector space. Thus U is subcritical with respect to s , so $s \in \mathcal{P}(\mathbb{V})$. \square

Lemma 2.20. Let $\mathbb{V} = (V, (V_j), (\phi_j))$ be a vector space HBL datum. Let $s \in [0, 1]^m$ and let $W \leq V$ be a subspace of V . If W is critical with respect to s then

$$s \in \mathcal{P}(\mathbb{V}) \iff s \in \mathcal{P}(\mathbb{V}_W) \cap \mathcal{P}(\mathbb{V}_{V/W}). \quad (2.8)$$

Proof. With Lemma 2.19 in hand, it remains to show that if $s \in \mathcal{P}(\mathbb{V})$ then $s \in \mathcal{P}(\mathbb{V}_W) \cap \mathcal{P}(\mathbb{V}_{V/W})$. Any subspace $U \leq W$ is also a subspace of V . U is subcritical with respect to s when regarded as a subspace of W , if and only if U is subcritical when regarded as a subspace of V . So $s \in \mathcal{P}(\mathbb{V}_W)$.

Next consider any subspace $U/W \leq V/W$. We have $W \leq U \leq V$ and $\dim(U/W) = \dim(U) - \dim(W)$. Moreover,

$$\dim([\phi_j](U/W)) = \dim(\phi_j(U)/\phi_j(W)) = \dim(\phi_j(U)) - \dim(\phi_j(W)).$$

Therefore since $\dim(W) = \sum_{j=1}^m s_j \dim(\phi_j(W))$,

$$\begin{aligned} \dim(U/W) &= \dim(U) - \dim(W) \leq \sum_{j=1}^m s_j \dim(\phi_j(U)) - \sum_{j=1}^m s_j \dim(\phi_j(W)) \\ &= \sum_{j=1}^m s_j (\dim(\phi_j(U)) - \dim(\phi_j(W))) = \sum_{j=1}^m s_j \dim([\phi_j](U/W)) \end{aligned}$$

by the subcriticality of U , which holds because $s \in \mathcal{P}(\mathbb{V})$. Thus any subspace $U/W \leq V/W$ is subcritical with respect to s , so $s \in \mathcal{P}(\mathbb{V}_{V/W})$, as well. \square

2.8 Conclusion of proof of Theorem 2.5

In order to prove Theorem 2.5 for a vector space HBL datum $(V, (V_j), (\phi_j))$, we argue by induction on the dimension of the ambient vector space V . If $\dim(V) = 0$ then V has a single element, and the inequality (2.3) is trivially valid.

To establish the inductive step, consider any extreme point s of \mathcal{P} . According to Proposition 2.12, there are two cases which must be analyzed. We begin with the case in which there exists a critical subspace $\{0\} < W < V$. We assume that Theorem 2.5 holds for all HBL data for which the ambient vector space has strictly smaller dimension than is the case for the given datum.

Lemma 2.21. *Let $\mathbb{V} = (V, (V_j), (\phi_j))$ be a vector space HBL datum, and let $s \in \mathcal{P}(\mathbb{V})$. Suppose that there exists a subspace $\{0\} < W < V$ that is critical with respect to s . Then the inequality (2.3) holds for this s .*

Proof. Consider the inequality (2.3) for some $s \in [0, 1]^m$. We may assume that none of the exponents s_j equals zero. For if $s_k = 0$, then $f_k(\phi_k(x)) \leq \|f_k\|_{1/s_k}$ for all x , and therefore

$$\sum_{x \in V} \prod_{j=1}^m f_j(\phi_j(x)) \leq \|f_k\|_{1/s_k} \cdot \sum_{x \in V} \prod_{j \neq k} f_j(\phi_j(x)).$$

If $\|f_k\|_{1/s_k} = 0$, then (2.3) holds with both sides 0. Otherwise we divide by $\|f_k\|_{1/s_k}$ to conclude that $s \in \mathcal{P}(\mathbb{V})$ if and only if $(s_j)_{j \neq k}$ belongs to the polytope associated to the vector space HBL datum $(V, (V_j)_{j \neq k}, (\phi_j)_{j \neq k})$. Thus the index k can be eliminated. This reduction can be repeated to remove all indices that equal zero.

Let $W_j := \phi_j(W)$. By Lemma 2.20, $s \in \mathcal{P}(W, (W_j), (\phi_j|_W))$. Therefore by the inductive hypothesis, one of the inequalities in (2.3) is

$$\sum_{x \in W} \prod_{j=1}^m f_j(\phi_j(x)) \leq \prod_{j=1}^m \|f_j|_{W_j}\|_{1/s_j}. \quad (2.9)$$

Define $F_j: V_j/W_j \rightarrow [0, \infty)$ to be the function

$$F_j(x + W_j) = \left(\sum_{y \in W_j} f_j(y + x)^{1/s_j} \right)^{s_j}.$$

This quantity is a function of the coset $x + W_j$ alone, rather than of x itself, because for any $z \in W_j$,

$$\sum_{y \in W_j} f_j(y + (x + z))^{1/s_j} = \sum_{y \in W_j} f_j(y + x)^{1/s_j}$$

by virtue of the substitution $y + z \mapsto y$. Moreover,

$$\|F_j\|_{1/s_j} = \|f_j\|_{1/s_j}. \quad (2.10)$$

To prove this, choose one element $x \in V_j$ for each coset $x + W_j \in V_j/W_j$. Denoting by X the set of all these representatives,

$$\|F_j\|_{1/s_j}^{1/s_j} = \sum_{x \in X} \sum_{y \in W_j} f_j(y + x)^{1/s_j} = \sum_{z \in V_j} f_j(z)^{1/s_j}$$

because the map $X \times W_j \ni (x, y) \mapsto x + y \in V_j$ is a bijection.

The inductive bound (2.9) can be equivalently written in the more general form

$$\sum_{x \in W} \prod_{j=1}^m f_j(\phi_j(x + y)) \leq \prod_{j=1}^m F_j([\phi_j](y + W)) \quad (2.11)$$

for any $y \in V$, by applying (2.9) to (\hat{f}_j) where $\hat{f}_j(z) = f_j(z + \phi_j(y))$.

Denote by $Y \subseteq V$ a set of representatives of the cosets $y + W \in V/W$, and identify V/W with Y . Then

$$\sum_{x \in V} \prod_{j=1}^m f_j(\phi_j(x)) = \sum_{y \in Y} \sum_{x \in W} \prod_{j=1}^m f_j(\phi_j(y + x)) \leq \sum_{y \in Y} \prod_{j=1}^m F_j([\phi_j](y + W))$$

by (2.11).

By (2.10), it suffices to show that

$$\sum_{y \in Y} \prod_j F_j([\phi_j](y + W)) \leq \prod_j \|F_j\|_{1/s_j} \quad \text{for all functions } 0 \leq F_j \in \ell^{1/s_j}(V_j/W_j). \quad (2.12)$$

This is a set of inequalities of exactly the form (2.3), with \mathbb{V} replaced by $\mathbb{V}_{V/W} = (V/W, (V_j/W_j), ([\phi_j]))$. By Lemma 2.20, $s \in \mathcal{P}(\mathbb{V}_{V/W})$, and since $\dim(V/W) < \dim(V)$, we conclude directly from the inductive hypothesis that (2.12) holds, concluding the proof of Lemma 2.21. \square

According to Proposition 2.12, in order to complete the proof of Theorem 2.5, it remains only to analyze the case in which the extreme point s is an element of $\{0, 1\}^m$. Let $K = \{k : s_k = 1\}$. Consider $W = \bigcap_{k \in K} \ker(\phi_k)$. Since W is subcritical by hypothesis,

$$\dim(W) \leq \sum_{j=1}^m s_j \dim(\phi_j(W)) = \sum_{k \in K} \dim(\phi_k(W)) = 0$$

so $\dim(W) = 0$, that is, $W = \{0\}$. Therefore the map $x \mapsto (\phi_k(x))_{k \in K}$ from V to the Cartesian product $\prod_{k \in K} V_k$ is injective.

For any $x \in V$,

$$\prod_{j=1}^m f_j(\phi_j(x)) \leq \prod_{k \in K} f_k(\phi_k(x)) \prod_{i \notin K} \|f_i\|_\infty = \prod_{k \in K} f_k(\phi_k(x)) \prod_{i \notin K} \|f_i\|_{1/s_i}$$

since $s_i = 0$ for all $i \notin K$. Thus it suffices to prove that

$$\sum_{x \in V} \prod_{k \in K} f_k(\phi_k(x)) \leq \prod_{k \in K} \|f_k\|_1.$$

This is a special case of the following result.

Lemma 2.22. *Let V be any finite-dimensional vector space over \mathbb{F} . Let K be a finite index set, and for each $k \in K$, let ϕ_k be an \mathbb{F} -linear map from V to a finite-dimensional vector space V_k . If $\bigcap_{k \in K} \ker(\phi_k) = \{0\}$ then for all functions $f_k: V_k \rightarrow [0, \infty)$,*

$$\sum_{x \in V} \prod_{k \in K} f_k(\phi_k(x)) \leq \prod_{k \in K} \|f_k\|_1.$$

Proof. Define $\Phi: V \rightarrow \prod_{k \in K} V_k$ by $\Phi(x) = (\phi_k(x))_{k \in K}$. The hypothesis $\bigcap_{k \in K} \ker(\phi_k) = \{0\}$ is equivalent to Φ being injective. The product $\prod_{k \in K} \|f_k\|_1$ can be expanded as the sum of products

$$\sum_y \prod_{k \in K} f_k(y_k)$$

where the sum is taken over all $y = (y_k)_{k \in K}$ belonging to the Cartesian product $\prod_{k \in K} V_k$. The quantity of interest,

$$\sum_{x \in V} \prod_{k \in K} f_k(\phi_k(x)),$$

is likewise a sum of such products. Each term of the latter sum appears as a term of the former sum, and by virtue of the injectivity of Φ , appears only once. Since all summands are nonnegative, the former sum is greater than or equal to the latter. Therefore

$$\prod_{k \in K} \|f_k\|_1 = \sum_y \prod_{k \in K} f_k(y_k) \geq \sum_{x \in V} \prod_{k \in K} f_k(\phi_k(x)).$$

□

As mentioned above, necessity of the condition (2.2) for the inequality (2.3) in the case $\mathbb{F} = \mathbb{Q}$ can be deduced from the corresponding necessity in Theorem 1.3, by clearing denominators. First, we identify V and V_j with \mathbb{Q}^d and \mathbb{Q}^{d_j} and let E be any nonempty finite subset of \mathbb{Q}^d . Let $\hat{\phi}_j: \mathbb{Q}^d \rightarrow \mathbb{Q}^{d_j}$ be the linear map represented by the matrix of ϕ_j multiplied by the lowest common denominator of its entries, i.e., an integer matrix. Likewise, let \hat{E} be the set obtained from E by multiplying each point by the lowest common denominator of the coordinates of all points in E . Then by linearity,

$$|\hat{E}| = |E| \leq \prod_{j=1}^m |\phi_j(E)|^{s_j} = \prod_{j=1}^m |\hat{\phi}_j(\hat{E})|^{s_j}.$$

Recognizing $(\mathbb{Z}^d, (\mathbb{Z}^{d_j}), (\hat{\phi}_j|_{\mathbb{Z}^d}))$ as an Abelian group HBL datum, we conclude (1.3) for this datum from the implication (1.4) \Rightarrow (1.3) of Theorem 1.3. According to Lemma 2.6, (2.2) holds for the vector space HBL datum $(\mathbb{Q}^d, (\mathbb{Q}^{d_j}), (\hat{\phi}_j))$; our conclusion follows since $\dim(\hat{\phi}_j(W)) = \dim(\phi_j(W))$ for any $W \leq \mathbb{Q}^d$.

It remains to show that (2.2) is necessary for (2.3) in the case of a finite field \mathbb{F} . Whereas the above reasoning required only the validity of (2.4) in the weakened form $|E| \leq C \prod_{j=1}^m |\phi_j(E)|^{s_j}$ for some constant $C < \infty$ independent of E (see proof of necessity for Theorem 1.3), now the assumption that this holds with $C = 1$ becomes essential. Let W be any subspace of V . Since $|\mathbb{F}| < \infty$ and W has finite dimension over \mathbb{F} , W is a finite set and the hypothesis (2.4) can be applied with $E = W$. Therefore $|W| \leq \prod_{j=1}^m |\phi_j(W)|^{s_j}$. This is equivalent to

$$|\mathbb{F}|^{\dim(W)} \leq \prod_{j=1}^m |\mathbb{F}|^{s_j \dim(\phi_j(W))},$$

so since $|\mathbb{F}| \geq 2$, by taking base- $|\mathbb{F}|$ logarithms of both sides, we obtain $\dim(W) \leq \sum_{j=1}^m s_j \dim(\phi_j(W))$, as was to be shown. □

3 An algorithm that computes the polytope \mathcal{P}

From the perspective of potential applications to communication bounds and the analysis of algorithms, it is desirable to compute the convex polytope $\mathcal{P} = \mathcal{P}(\mathfrak{G})$ associated to an Abelian group HBL datum.

We have already shown in Lemma 2.6 that \mathcal{P} is unchanged when the groups \mathbb{Z}^d and \mathbb{Z}^{d_j} are embedded in the natural way into the vector spaces \mathbb{Q}^d and \mathbb{Q}^{d_j} over \mathbb{Q} and the homomorphisms ϕ_j are viewed as \mathbb{Q} -linear maps. Thus \mathcal{P} is identical to the polytope defined by the inequalities

$$\dim(V) \leq \sum_{j=1}^m s_j \dim(\phi_j(V)) \quad \text{for all subspaces } V \leq \mathbb{Q}^d \quad (3.1)$$

and $0 \leq s_j \leq 1$ for all indices j . Indeed, (3.1) is the hypothesis (2.2) of Theorem 2.5 in the case $\mathbb{F} = \mathbb{Q}$.

We will show how to compute \mathcal{P} in the case $\mathbb{F} = \mathbb{Q}$. Throughout the remainder of this section, V and V_j denote finite-dimensional vector spaces over \mathbb{Q} , and ϕ_j denotes a \mathbb{Q} -linear map. The reasoning presented below applies to any countable field \mathbb{F} , provided that elements of \mathbb{F} and the field operations are computable.

Remark 3.1. *The algorithm described below relies on a search of a list of all subspaces of V . A similar algorithm was sketched, less formally, in [10] for computing the corresponding polytope in [3, Theorem 2.1]. That algorithm searches a smaller collection of subspaces, namely the lattice generated by the kernels of ϕ_1, \dots, ϕ_m under the operations of intersection and pairwise sum of subspaces. In a forthcoming sequel to this work, we will show that it suffices to search a corresponding lattice in our situation. Searching this lattice would make the algorithm below more efficient, in principle. But our goal here is to prove decidability, leaving issues of efficiency for future work. A minor point is that this modification also allows relaxation of the requirement that \mathbb{F} be countable.*

The proof of Theorem 1.5 is built upon several smaller results.

Lemma 3.2. *There exists an algorithm that takes as input a finite-dimensional vector space V over \mathbb{Q} , and returns a list of its subspaces. More precisely, this algorithm takes as input a finite-dimensional vector space V and a positive integer N , and returns as output the first N elements W_i of a list (W_1, W_2, \dots) of all subspaces of V . This list is independent of N . Each subspace W is expressed as a finite sequence $(d; w_1, \dots, w_d)$ where $d = \dim(W)$ and $\{w_i\}$ is a basis for W .*

Proof. Generate a list of all nonempty subsets of V having at most $\dim(V)$ elements. Test each subset for linear independence, and discard all that fail to be independent. Output a list of those that remain. \square

We do not require this list to be free of redundancies.

Lemma 3.3. *For any positive integer m , there exists an algorithm that takes as input a finite set of linear inequalities over \mathbb{Z} for $s \in [0, 1]^m$, and returns as output a list of all the extreme points of the convex subset $\mathcal{P} \subseteq [0, 1]^m$ specified by these inequalities.*

Proof. To the given family of inequalities, adjoin the $2m$ inequalities $s_j \geq 0$ and $-s_j \geq -1$. \mathcal{P} is the convex polytope defined by all inequalities in the resulting enlarged family. Express these inequalities as $\langle s, v_\alpha \rangle \geq c_\alpha$ for all $\alpha \in A$, where A is a finite nonempty index set.

An arbitrary point $\tau \in \mathbb{R}^m$ is an extreme point of \mathcal{P} if and only if (i) there exists a set B of indices α having cardinality m , such that $(v_\beta : \beta \in B)$ is linearly independent and $\langle \tau, v_\beta \rangle = c_\beta$ for all $\beta \in B$, and (ii) τ satisfies $\langle \tau, v_\alpha \rangle \geq c_\alpha$ for all $\alpha \in A$.

Create a list of all subsets $B \subset A$ with cardinality equal to m . There are finitely many such sets, since A itself is finite. Delete each one for which $(v_\beta : \beta \in B)$ is not linearly independent. For each subset B not deleted, compute the unique solution τ of the system of equations $\langle \tau, v_\beta \rangle = c_\beta$ for all $\beta \in B$. Include τ in the list of all extreme points, if and only if τ satisfies $\langle \tau, v_\alpha \rangle \geq c_\alpha$ for all $\alpha \in A \setminus B$. \square

Proposition 3.4. *There exists an algorithm that takes as input a vector space HBL datum $\mathbb{V} = (V, (V_j), (\phi_j))$, an element $t \in [0, 1]^m$, and a subspace $\{0\} < W < V$ which is critical with respect to t , and determines whether $t \in \mathcal{P}(\mathbb{V})$.*

Theorem 1.5 and Proposition 3.4 will be proved inductively in tandem, according to the following induction scheme. The proof of Theorem 1.5 for HBL data in which V has dimension n , will rely on Proposition 3.4 for HBL data in which V has dimension n . The proof of Proposition 3.4 for HBL data in which V has dimension n and there are m subspaces V_j , will rely on Proposition 3.4 for HBL data in which V has dimension strictly less than n , on Theorem 1.5 for HBL data in which V has dimension strictly less than n , and also on Theorem 1.5 for HBL data in which V has dimension n and the number of subspaces V_j is strictly less than m . Thus there is no circularity in the reasoning.

Proof of Proposition 3.4. Let $(V, (V_j), (\phi_j))$ and t, W be given. Following Notation 2.17, consider the two HBL data $\mathbb{V}_W = (W, (\phi_j(W)), (\phi_j|_W))$ and $\mathbb{V}_{V/W} = (V/W, (V_j/\phi_j(W)), ([\phi_j]))$, where $[\phi_j]: V/W \rightarrow V_j/\phi_j(W)$ are the quotient maps. From a basis for V , bases for V_j , a basis for W , and corresponding matrix representations of ϕ_j , it is possible to compute the dimensions of, and bases for, V/W and $V_j/\phi_j(W)$, via row operations on matrices. According to Lemma 2.20, $t \in \mathcal{P}(\mathbb{V})$ if and only if $t \in \mathcal{P}(\mathbb{V}_W) \cap \mathcal{P}(\mathbb{V}_{V/W})$.

Because $\{0\} < W < V$, both $W, V/W$ have dimensions strictly less than the dimension of V . Therefore by Theorem 1.5 and the induction scheme, there exists an algorithm which computes both a finite list of inequalities characterizing $\mathcal{P}(\mathbb{V}_W)$, and a finite list of inequalities characterizing $\mathcal{P}(\mathbb{V}_{V/W})$. Testing each of these inequalities on t determines whether t belongs to these two polytopes, hence whether t belongs to $\mathcal{P}(\mathbb{V})$. \square

Lemma 3.5. *Let an HBL datum $\mathbb{V} = (V, (V_j), (\phi_j))$ be given. Let $i \in \{1, 2, \dots, m\}$. Let $s \in [0, 1]^m$ and suppose that $s_i = 1$. Let $V' = \ker(\phi_i)$. Define $\hat{s} \in [0, 1]^{m-1}$ to be (s_1, \dots, s_m) with the i^{th} coordinate deleted. Then $s \in \mathcal{P}(V, (V_j), (\phi_j))$ if and only if $\hat{s} \in \mathcal{P}(V', (V_j)_{j \neq i}, (\phi_j|_{V'})_{j \neq i})$.*

Proof. Write $\mathbb{V}' = (V', (V_j)_{j \neq i}, (\phi_j|_{V'})_{j \neq i})$. For any subspace $W \leq V'$, since $\dim(\phi_i(W)) = 0$,

$$\sum_j s_j \dim(\phi_j(W)) = \sum_{j \neq i} s_j \dim(\phi_j(W)).$$

So if $s \in \mathcal{P}(\mathbb{V})$ then $\hat{s} \in \mathcal{P}(\mathbb{V}')$.

Conversely, suppose that $\hat{s} \in \mathcal{P}(\mathbb{V}')$. Let W be any subspace of V . Write $W = W'' + (W \cap V')$ where the subspace $W'' \leq V$ is a supplement to $W \cap V'$ in W , so that $\dim(W) = \dim(W'') + \dim(W \cap V')$. Then

$$\begin{aligned} \sum_j s_j \dim(\phi_j(W)) &= \dim(\phi_i(W)) + \sum_{j \neq i} s_j \dim(\phi_j(W)) \\ &\geq \dim(\phi_i(W'')) + \sum_{j \neq i} s_j \dim(\phi_j(W \cap V')) \\ &\geq \dim(W'') + \dim(W \cap V'); \end{aligned}$$

$\dim(\phi_i(W'')) = \dim(W'')$ because ϕ_i is injective on W'' . So $s \in \mathcal{P}(\mathbb{V})$. \square

To prepare for the proof of Theorem 1.5, let $\mathcal{P}(\mathbb{V}) = \mathcal{P}(V, (V_j), (\phi_j))$ be given. Let (W_1, W_2, W_3, \dots) be the list of subspaces of V produced by the algorithm of Lemma 3.2. Let $N \geq 1$. To each index $\alpha \in \{1, 2, \dots, N\}$ is associated a linear inequality $\sum_{j=1}^m s_j \dim(\phi_j(W_\alpha)) \geq \dim(W_\alpha)$ for elements $s \in [0, 1]^m$, which we encode by an $(m+1)$ -tuple $(v(W_\alpha), c(W_\alpha))$; the inequality is $\langle s, v(W_\alpha) \rangle \geq c(W_\alpha)$. Define $\mathcal{P}_N \subseteq [0, 1]^m$ to be the polytope defined by this set of inequalities.

Lemma 3.6.

$$\mathcal{P}_N \supseteq \mathcal{P}(\mathbb{V}) \quad \text{for all } N. \quad (3.2)$$

Moreover, there exists a positive integer N such that $\mathcal{P}_M = \mathcal{P}(\mathbb{V})$ for all $M \geq N$.

Proof. The inclusion holds for every N , because the set of inequalities defining \mathcal{P}_N is a subset of the set defining $\mathcal{P}(\mathbb{V})$.

$\mathcal{P}(\mathbb{V})$ is specified by some finite set of inequalities, each specified by some subspace of V . Choose one such subspace for each of these inequalities. Since (W_α) is a list of all subspaces of V , there exists M such that each of these chosen subspaces belongs to $(W_\alpha : \alpha \leq M)$. \square

Lemma 3.7. *Let $m \geq 2$. If s is an extreme point of \mathcal{P}_N , then either $s_j \in \{0, 1\}$ for some $j \in \{1, 2, \dots, m\}$, or there exists $\alpha \in \{1, 2, \dots, N\}$ for which W_α is critical with respect to s and $0 < \dim(W_\alpha) < \dim(V)$.*

In the following argument, we say that two inequalities $\langle s, v(W_\alpha) \rangle \geq c(W_\alpha)$, $\langle s, v(W_\beta) \rangle \geq c(W_\beta)$ are distinct if they specify different subsets of \mathbb{R}^m .

Proof. For any extreme point s , equality must hold in at least m distinct inequalities among those defining \mathcal{P}_N . These inequalities are of three kinds: $\langle s, v(W_\alpha) \rangle \geq c(W_\alpha)$ for $\alpha \in \{1, 2, \dots, N\}$, $s_j \geq 0$, and $-s_j \geq -1$, with $j \in \{1, 2, \dots, m\}$. If $W_\beta = \{0\}$ then W_β specifies the tautologous inequality $\sum_j s_j \cdot 0 = 0$, so that index β can be disregarded.

If none of the coordinates s_j are equal to 0 or 1, there must exist β such that equality holds in at least two distinct inequalities $\langle s, v(W_\beta) \rangle \geq c(W_\beta)$ associated to subspaces W_α among those which are used to define \mathcal{P}_N . We have already discarded the subspace $\{0\}$, so there must exist β such that W_β and V specify distinct inequalities. Thus $0 < \dim(W_\beta) < \dim(V)$. \square

Proof of Theorem 1.5. Set $\mathcal{P} = \mathcal{P}(V, (V_j), (\phi_j))$. Consider first the base case $m = 1$. The datum is a pair of finite-dimensional vector spaces V, V_1 with a \mathbb{Q} -linear map $\phi: V \rightarrow V_1$. The polytope \mathcal{P} is the set of all $s \in [0, 1]$ for which $s \dim(\phi(W)) \geq \dim(W)$ for every subspace $W \leq V$. If $\dim(V) = 0$ then $\mathcal{P}(V, (V_j), (\phi_j)) = [0, 1]$. If $\dim(V) > 0$ then since $\dim(\phi(W)) \leq \dim(W)$ for every subspace, the inequality can only hold if the kernel of ϕ has dimension 0, and then only for $s = 1$. The kernel of ϕ can be computed. So \mathcal{P} can be computed when $m = 1$.

Suppose that $m \geq 2$. Let $(V, (V_j), (\phi_j))$ be given. Let $N = 0$. Recursively apply the following procedure.

Replace N by $N + 1$. Consider \mathcal{P}_N . Apply Lemma 3.3 to obtain a list of all extreme points τ of \mathcal{P}_N , and for each such τ which belongs to $(0, 1)^m$, a nonzero proper subspace $W(\tau) \leq V$ which is critical with respect to τ .

Examine each of these extreme points τ , to determine whether $\tau \in \mathcal{P}(V, (V_j), (\phi_j))$. There are three cases. Firstly, if $\tau \in (0, 1)^m$, then Proposition 3.4 may be invoked, using the critical subspace $W(\tau)$, to determine whether $\tau \in \mathcal{P}$.

Secondly, if some component τ_i of τ equals 1, let V' be the kernel of ϕ_i . Set

$$\mathcal{P}' = \mathcal{P}(V', (V_j)_{j \neq i}, (\phi_j|_{V'})_{j \neq i}).$$

According to Lemma 3.5, $\tau \in \mathcal{P}$ if and only if $\hat{\tau} = (\tau_j)_{j \neq i} \in \mathcal{P}'$. This polytope \mathcal{P}' can be computed by the induction hypothesis, since the number of indices j has been reduced by one.

Finally, if some component τ_i of τ equals 0, then because the term $s_i \dim(\phi_i(W)) = 0$ contributes nothing to sums $\sum_{j=1}^m s_j \dim(\phi_j(W))$, $\tau \in \mathcal{P}$ if and only if $\hat{\tau}$ belongs to $\mathcal{P}(V, (V_j)_{j \neq i}, (\phi_j)_{j \neq i})$. To determine whether $\hat{\tau}$ belongs to this polytope requires again only an application of the induction hypothesis.

If every extreme point τ of \mathcal{P}_N belongs to \mathcal{P} , then because \mathcal{P}_N is the convex hull of its extreme points, $\mathcal{P}_N \subseteq \mathcal{P}$. The converse inclusion holds for every N , so in this case $\mathcal{P}_N = \mathcal{P}$. The algorithm halts, and returns the conclusion that $\mathcal{P} = \mathcal{P}_N$, along with information already computed: a list of the inequalities specified by all the subspaces W_1, \dots, W_N , and a list of extreme points of $\mathcal{P}_N = \mathcal{P}$.

On the other hand, if at least one extreme point of \mathcal{P}_N fails to belong to \mathcal{P} , then $\mathcal{P}_N \neq \mathcal{P}$. Then increment N by one, and repeat the above steps.

Lemma 3.6 guarantees that this procedure will halt after finitely many steps. \square

4 On (un)computability of the list of constraints defining \mathcal{P}

We have discussed the computation of \mathcal{P} . Now we turn to the problem of computing the set of inequalities (1.3). There may be many sets of inequalities which specify \mathcal{P} ; thus in order to compute \mathcal{P} , it suffices to compute any such set of inequalities, rather than the specific set (1.3). This distinction accounts for the coexistence of Theorems 1.5 and 1.6.

In order to compute the set of inequalities (1.3), we must compute the answer to the question: Given any group homomorphisms ϕ_1, \dots, ϕ_m and integers $0 \leq r, r_1, \dots, r_m \leq d$, does there exist a subgroup $H \leq \mathbb{Z}^d$ such that

$$\text{rank}(H) = r, \text{ and } \text{rank}(\phi_j(H)) = r_j \text{ for all } 1 \leq j \leq m?$$

Notation 4.1. For a natural number d and ring R , $M_d(R)$ denotes the ring of d -by- d matrices with entries from R . We identify $M_d(R)$ with the endomorphism ring of the R -module R^d (or R -vector space, if R is a field) and thus may write elements of $M_d(R)$ as R -linear maps rather than as matrices. Via the usual coordinates, we may identify $M_d(R)$ with R^{d^2} . We write $R[x_1, \dots, x_q]$ to denote the ring of polynomials over R in variables x_1, \dots, x_q .

Recall that there are given $d, d_j \in \mathbb{N}$ and \mathbb{Z} -linear maps $\phi_j: \mathbb{Z}^d \rightarrow \mathbb{Z}^{d_j}$, for $j \in \{1, 2, \dots, m\}$ for some positive integer m . Each ϕ_j can also be interpreted as a \mathbb{Q} -linear map (from \mathbb{Q}^d to \mathbb{Q}^{d_j}), represented by the same integer matrix. It is no loss of generality to assume that each $d_j = d$, so that each ϕ_j is an endomorphism of \mathbb{Z}^d .

Definition 4.2. Given $m, d \in \mathbb{N}$, and a finite sequence r, r_1, \dots, r_m of natural numbers each bounded by d , we define the sets

$$E_{d;r,r_1,\dots,r_m} := \{(\phi_1, \dots, \phi_m) \in (M_d(\mathbb{Z}))^m : (\exists H \leq \mathbb{Z}^d) \text{ rank}(H) = r \text{ and } \text{rank}(\phi_j(H)) = r_j, 1 \leq j \leq m\},$$

$$E_{d;r,r_1,\dots,r_m}^{\mathbb{Q}} := \{(\phi_1, \dots, \phi_m) \in (M_d(\mathbb{Q}))^m : (\exists V \leq \mathbb{Q}^d) \text{ dim}(V) = r \text{ and } \text{dim}(\phi_j(V)) = r_j, 1 \leq j \leq m\}.$$

Remark 4.3. The question of whether a given m -tuple $(\phi_1, \dots, \phi_m) \in (M_d(R))^m$ is a member of $E_{d;r,r_1,\dots,r_m}$ (when $R = \mathbb{Z}$) or $E_{d;r,r_1,\dots,r_m}^{\mathbb{Q}}$ (when $R = \mathbb{Q}$) is an instance of the problem of whether some system of polynomial equations has a solution over the ring R . We let B be a d -by- r matrix of variables, and construct a system of polynomial equations in the dr unknown entries of B and md^2 known entries of ϕ_1, \dots, ϕ_m that has a solution if and only if the aforementioned rank (or dimension) conditions are met. The condition $\text{rank}(M) = s$ for a matrix M is equivalent to all $(s+1)$ -by- $(s+1)$ minors of M equaling zero (i.e., the sum of their squares equaling zero), and at least one s -by- s minor being nonzero (i.e., the sum of their squares not equaling zero — see Remark 4.6). We construct two polynomial equations in this manner for $M = B$ (with $s = r$) and for each matrix $M = \phi_j B$ (with $s = r_j$).

Lemma 4.4. With the notation as in Definition 4.2, $E_{d;r,r_1,\dots,r_m} = E_{d;r,r_1,\dots,r_m}^{\mathbb{Q}} \cap (M_d(\mathbb{Z}))^m$.

Proof. This result was already established in Lemma 2.6; we restate it here using the present notation. For the left-to-right inclusion, observe that if $H \leq \mathbb{Z}^d$ witnesses that $(\phi_1, \dots, \phi_m) \in E_{d;r,r_1,\dots,r_m}$, then $H_{\mathbb{Q}}$ witnesses that $(\phi_1, \dots, \phi_m) \in E_{d;r,r_1,\dots,r_m}^{\mathbb{Q}}$. For the other inclusion, if $(\phi_1, \dots, \phi_m) \in E_{d;r,r_1,\dots,r_m}^{\mathbb{Q}} \cap (M_d(\mathbb{Z}))^m$ witnessed by $V \leq \mathbb{Q}^d$, then we may find a submodule $H = V \cap \mathbb{Z}^d$ of \mathbb{Z}^d , with $\text{rank}(H) = \text{dim}(V)$. Then $\text{dim}(\phi_j(V)) = \text{rank}(\phi_j(H)) = r_j$ showing that $(\phi_1, \dots, \phi_m) \in E_{d;r,r_1,\dots,r_m}$. \square

Definition 4.5. Hilbert's Tenth Problem for \mathbb{Q} is the question of whether there is an algorithm which given a finite set of polynomials $f_1(x_1, \dots, x_q), \dots, f_p(x_1, \dots, x_q) \in \mathbb{Q}[x_1, \dots, x_q]$ (correctly) determines whether or not there is some $a \in \mathbb{Q}^q$ for which $f_1(a) = \dots = f_p(a) = 0$.

Remark 4.6. One may modify the presentation of Hilbert's Tenth Problem for \mathbb{Q} in various ways without affecting its truth value. For example, one may allow a condition of the form $g(a) \neq 0$ as this is equivalent to $(\exists b)(g(a)b - 1 = 0)$. On the other hand, using the fact that $x^2 + y^2 = 0 \iff x = 0 = y$, one may replace the finite sequence of polynomial equations with a single equality.

Remark 4.7. Hilbert's Tenth Problem, proper, asks for an algorithm to determine solvability in integers of finite systems of equations over \mathbb{Z} . From such an algorithm one could positively resolve Hilbert's Tenth Problem for \mathbb{Q} . However, by the celebrated theorem of Matiyasevich-Davis-Putnam-Robinson [8], no such algorithm exists. The problem for the rationals remains open. The most natural approach would be to reduce from the problem over \mathbb{Q} to the problem over \mathbb{Z} , say, by showing that \mathbb{Z} may be defined by an expression of the form

$$a \in \mathbb{Z} \iff (\exists y_1) \cdots (\exists y_q) P(a; y_1, \dots, y_q) = 0$$

for some fixed polynomial P . Koenigsmann [7] has shown that there is in fact a universal definition of \mathbb{Z} in \mathbb{Q} , that is, a formula of the form

$$a \in \mathbb{Z} \iff (\forall y_1) \cdots (\forall y_q) \theta(a; y_1, \dots, y_q) = 0$$

where θ is a finite Boolean combination of polynomial equations, but he also demonstrated that the existence of an existential definition of \mathbb{Z} would violate the Bombieri-Lang conjecture. Koenigsmann's result shows that it is unlikely that Hilbert's Tenth Problem for \mathbb{Q} can be resolved by reducing to the problem over \mathbb{Z} using an existential definition of \mathbb{Z} in \mathbb{Q} . However, it is conceivable that this problem could be resolved without such a definition.

Proof of Theorem 1.6 (necessity). Evidently, if Hilbert's Tenth Problem for \mathbb{Q} has a positive solution, then there is an algorithm to (correctly) determine for $d \in \mathbb{N}$, $r, r_1, \dots, r_m \leq d$ also in \mathbb{N} , and $(\phi_1, \dots, \phi_m) \in (M_d(\mathbb{Z}))^m$ whether $(\phi_1, \dots, \phi_m) \in E_{d;r,r_1,\dots,r_m}$. By Lemma 4.4, $(\phi_1, \dots, \phi_m) \in E_{d;r,r_1,\dots,r_m}$ just in case $(\phi_1, \dots, \phi_m) \in E_{d;r,r_1,\dots,r_m}^{\mathbb{Q}}$. This last condition leads to an instance of Hilbert's Tenth Problem (for \mathbb{Q}) for the set of rational polynomial equations given in Remark 4.3. \square

Notation 4.8. Given a set $S \subseteq \mathbb{Q}[x_1, \dots, x_q]$ of polynomials, we denote the set of rational solutions to the equations $f = 0$ as f ranges through S by

$$V(S)(\mathbb{Q}) := \{a \in \mathbb{Q}^q : (\forall f \in S) f(a) = 0\}.$$

Definition 4.9. For any natural number t , we say that the set $D \subseteq \mathbb{Q}^t$ is Diophantine if there is some $q - t \in \mathbb{N}$ and a set $S \subseteq \mathbb{Q}[x_1, \dots, x_t; y_1, \dots, y_{q-t}]$ for which

$$D = \{a \in \mathbb{Q}^t : (\exists b \in \mathbb{Q}^{q-t})(a; b) \in V(S)(\mathbb{Q})\}.$$

We will show sufficiency in Theorem 1.6 by establishing a stronger result: namely, that an algorithm to decide membership in sets of the form $E_{d;r,r_1,\dots,r_m}$ could also be used to decide membership in any Diophantine set. (Hilbert's Tenth Problem for \mathbb{Q} concerns membership in the specific Diophantine set $V(S)(\mathbb{Q})$.)

With the next lemma, we use a standard trick of replacing composite terms with single applications of the basic operations to put a general Diophantine set in a standard form (see, e.g., [9]).

Lemma 4.10. Given any finite set of polynomials $S \subseteq \mathbb{Q}[x_1, \dots, x_q]$, let $d := \max_{f \in S} \max_{i=1}^q \deg_{x_i}(f)$ and $\mathcal{D} := \{0, 1, \dots, d\}^q$. There is another set of variables $\{u_\alpha\}_{\alpha \in \mathcal{D}}$ and another finite set $S' \subseteq \mathbb{Q}[\{u_\alpha\}_{\alpha \in \mathcal{D}}]$ consisting entirely of affine polynomials (polynomials of the form $c + \sum c_\alpha u_\alpha$ where not all c and c_α are zero) and polynomials of the form $u_\alpha u_\beta - u_\gamma$ with α, β , and γ distinct, so that $V(S)(\mathbb{Q}) = \pi(V(S')(\mathbb{Q}))$ where $\pi: \mathbb{Q}^{\mathcal{D}} \rightarrow \mathbb{Q}^q$ is given by

$$(u_\alpha)_{\alpha \in \mathcal{D}} \mapsto (u_{(1,0,\dots,0)}, u_{(0,1,0,\dots,0)}, \dots, u_{(0,\dots,0,1)}).$$

Proof. Denote by $\mathbf{0}$ the element $(0, 0, \dots, 0)$ of $\mathcal{D} = \{0, 1, \dots, d\}^q$. Let $T \subset \mathbb{Q}[\{u_\alpha\}_{\alpha \in \mathcal{D}}]$ consist of

- $u_{(0, \dots, 0)} - 1$ and
- $u_{\alpha+\beta} - u_\alpha u_\beta$ for $(\alpha + \beta) \in \mathcal{D}$, $\alpha \neq \mathbf{0}$ and $\beta \neq \mathbf{0}$.

Define $\chi: \mathbb{Q}^q \rightarrow \mathbb{Q}^{\mathcal{D}}$ by

$$(x_1, \dots, x_q) \mapsto (x^\alpha)_{\alpha \in \mathcal{D}}$$

where $x^\alpha := \prod_{j=1}^q x_j^{\alpha_j}$. One sees immediately that χ induces a bijection $\chi: \mathbb{Q}^q \rightarrow V(T)(\mathbb{Q})$ with inverse $\pi|_{V(T)(\mathbb{Q})}$.

Let S' be the set containing T and the polynomials $\sum_{\alpha \in \mathcal{D}} c_\alpha u_\alpha$ for which $\sum_{\alpha \in \mathcal{D}} c_\alpha x^\alpha \in S$.

One checks that if $a \in \mathbb{Q}^q$, then $\chi(a) \in V(S')(\mathbb{Q})$ if and only if $a \in V(S)(\mathbb{Q})$. Applying π , and noting that π is the inverse to χ on $V(T)(\mathbb{Q})$, the result follows. \square

Notation 4.11. *For the remainder of this argument, we call a set enjoying the properties identified for S' (namely that each polynomial is either affine or of the form $u_\alpha u_\beta - u_\gamma$) a basic set.*

Proof of Theorem 1.6 (sufficiency). It follows from Lemma 4.10 that the membership problem for a general finite Diophantine set may be reduced to the membership problem for a Diophantine set defined by a finite basic set of equations. Let $S \subseteq \mathbb{Q}[x_1, \dots, x_q]$ be a finite basic set of equations and let $t \leq q$ be some natural number, We now show that there are natural numbers $\mu, \nu, \rho, \rho_1, \dots, \rho_\mu$ and a computable function $f: \mathbb{Q}^t \rightarrow (M_\nu(\mathbb{Z}))^\mu$ so that for $a \in \mathbb{Q}^t$ one has that there is some $b \in \mathbb{Q}^{q-t}$ with $(a, b) \in V(S)(\mathbb{Q})$ if and only if $f(a) \in E_{\nu; \rho, \rho_1, \dots, \rho_\mu}$.

List the ℓ affine polynomials in S as

$$\lambda_{0,1} + \sum_{i=1}^q \lambda_{i,1} x_i, \quad \dots, \quad \lambda_{0,\ell} + \sum_{i=1}^q \lambda_{i,\ell} x_i$$

and the k polynomials expressing multiplicative relations in S as

$$x_{i_1,1} x_{i_2,1} - x_{i_3,1}, \quad \dots, \quad x_{i_1,k} x_{i_2,k} - x_{i_3,k}.$$

Note that by scaling, we may assume that all of the coefficients $\lambda_{i,j}$ are integers.

We shall take $\mu := 4 + q + t + |S|$, $\nu := 2q + 2$, $\rho := 2$ and the sequence ρ_1, \dots, ρ_μ to consist of $4 + q + k$ ones followed by $t + \ell$ zeros. Let us describe the map $f: \mathbb{Q}^t \rightarrow (M_\nu(\mathbb{Z}))^\mu$ by expressing each coordinate. For the sake of notation, our coordinates on \mathbb{Q}^ν are $(u; v) := (u_0, u_1, \dots, u_q; v_0, v_1, \dots, v_q)$.

- A. The map f_1 is constant taking the value $(u; v) \mapsto (u_0, 0, \dots, 0)$.
- B. The map f_2 is constant taking the value $(u; v) \mapsto (v_0, 0, \dots, 0)$.
- C. The map f_3 is constant taking the value $(u; v) \mapsto (u_0, \dots, u_q; 0, \dots, 0)$.
- D. The map f_4 is constant taking the value $(u; v) \mapsto (v_0, \dots, v_q; 0, \dots, 0)$.
- E. The map f_{4+j} (for $0 < j \leq q$) is constant taking the value $(u; v) \mapsto (u_0 - v_0, u_j - v_j, 0, \dots, 0)$.
- F. The map f_{4+q+j} (for $0 < j \leq k$) is constant taking the value $(u; v) \mapsto (u_{i_1,j} + v_0, u_{i_3,j} + v_{i_2,j}, 0, \dots, 0)$.
- G. The map $f_{4+q+k+j}$ (for $0 < j \leq t$) takes $a = (\frac{p_1}{q_1}, \dots, \frac{p_t}{q_t})$ (written in lowest terms) to the linear map $(u; v) \mapsto (p_j u_0 - q_j u_j, 0, \dots, 0)$.
- H. The map $f_{4+q+k+t+j}$ (for $0 < j \leq \ell$) takes the value $(u; v) \mapsto (\sum_{i=0}^q \lambda_{i,j} u_i, 0, \dots, 0)$.

Note that only the components $f_{4+q+k+j}$ for $0 < j \leq t$ actually depend on $(a_1, \dots, a_t) \in \mathbb{Q}^t$.

Let us check that this construction works. First, suppose that $a \in \mathbb{Q}^t$ and that there is some $b \in \mathbb{Q}^{q-t}$ for which $(a, b) \in V(S)(\mathbb{Q})$. For the sake of notation, we write $c = (c_1, \dots, c_q) := (a_1, \dots, a_t, b_1, \dots, b_{q-t})$.

Let $V := \mathbb{Q}(1, c_1, \dots, c_q, 0, \dots, 0) + \mathbb{Q}(0, \dots, 0, 1, c_1, \dots, c_q)$. We will check now that V witnesses that $f(a) \in E_{\nu, \rho, \rho_1, \dots, \rho_\mu}$. Note that a general element of V takes the form $(\alpha, \alpha c_1, \dots, \alpha c_q, \beta, \beta c_1, \dots, \beta c_q)$ for $(\alpha, \beta) \in \mathbb{Q}^2$. Throughout the rest of this proof, when we speak of a general element of some image of V , we shall write α and β as variables over \mathbb{Q} .

Visibly, $\dim(V) = 2 = \rho$.

Clearly, $f_1(a)(V) = \mathbb{Q}(1, 0, \dots, 0) = f_2(a)(V)$, so that $\rho_1 = \dim(f_1(a)(V)) = 1 = \dim(f_2(a)(V)) = \rho_2$ as required.

Likewise, $f_3(a)(V) = \mathbb{Q}(1, c_1, \dots, c_q, 0, \dots, 0) = f_4(a)(V)$, so that $\rho_3 = \dim(f_3(a)(V)) = 1 = \dim(f_4(a)(V)) = \rho_4$.

For $0 < j \leq q$ the general element of $f_{4+j}(a)(V)$ has the form $(\alpha - \beta, \alpha c_j - \beta c_j, 0, \dots, 0) = (\alpha - \beta)(1, c_j, 0, \dots, 0)$. Thus, $f_{4+j}(a)(V) = \mathbb{Q}(1, c_j, 0, \dots, 0)$ has dimension $\rho_{4+j} = 1$.

For $0 < j \leq k$ we have $c_{i_{3,j}} = c_{i_{1,j}} c_{i_{2,j}}$, the general element of $f_{4+q+j}(a)(V)$ has the form $(\alpha c_{i_{1,j}} + \beta, \alpha c_{i_{3,j}} + \beta c_{i_{2,j}}, 0, \dots, 0) = (\alpha c_{i_{1,j}} + \beta, \alpha c_{i_{1,j}} c_{i_{2,j}} + \beta c_{i_{2,j}}, 0, \dots, 0) = (\alpha c_{i_{1,j}} + \beta)(1, c_{i_{2,j}}, 0, \dots, 0)$ so we have that $\rho_{4+q+j} = \dim(f_{4+q+j}(a)(V)) = 1$.

For $0 < j \leq t$, the general element of $f_{4+q+k+j}(a)(V)$ has the form $(p_j \alpha - q_j \alpha a_j, 0, \dots, 0) = \mathbf{0}$. That is, $\rho_{4+q+k+j} = \dim(f_{4+q+k+j}(a)(V)) = 0$.

Finally, if $0 < j \leq \ell$, then the general element of $f_{4+q+k+t+j}(a)(V)$ has the form $(\lambda_{0,j} \alpha + \sum_{i=1}^q \lambda_{i,j} \alpha c_i, 0, \dots, 0) = \mathbf{0}$ since $\lambda_{0,j} + \sum_{i=1}^q \lambda_{i,j} c_i = 0$. So $\rho_{4+q+k+t+j} = \dim(f_{4+q+k+t+j}(a)(V)) = 0$.

Thus, we have verified that if $a \in \mathbb{Q}^t$ and there is some $b \in \mathbb{Q}^{d-t}$ with $(a, b) \in V(S)(\mathbb{Q})$, then $f(a) \in E_{\nu, \rho, \rho_1, \dots, \rho_\mu}$.

Conversely, suppose that $a = (\frac{p_1}{q_1}, \dots, \frac{p_t}{q_t}) \in \mathbb{Q}^t$ and that $f(a) \in E_{\nu, \rho, \rho_1, \dots, \rho_\mu}$, with the same $\mu, \nu, \rho, \rho_1, \dots, \rho_\mu$. Let $V \leq \mathbb{Q}^\nu$ have $\dim(V) = \rho$ witnessing that $f(a) \in E_{\nu, \rho, \rho_1, \dots, \rho_\mu}$.

Lemma 4.12. *There are elements g and h in V for which $g = (0, \dots, 0; g_0, \dots, g_q)$, $h = (h_0, \dots, h_q; 0, \dots, 0)$, $g_0 = h_0 = 1$, and $\mathbb{Q}g + \mathbb{Q}h = V$.*

Proof. The following is an implementation of row reduction. Let the elements $d = (d_0, \dots, d_q; d'_0, \dots, d'_q)$ and $e = (e_0, \dots, e_q; e'_0, \dots, e'_q)$ be a basis for V . Since $\dim(f_1(a)(V)) = 1$, at the cost of reversing d and e and multiplying by a scalar, we may assume that $d_0 = 1$. Since $\dim(f_3(a)(V)) = 1$, we may find a scalar γ for which $(\gamma d_0, \dots, \gamma d_q) = (e_0, \dots, e_q)$. Set $\tilde{g} := e - \gamma d$. Write $\tilde{g} = (0, \dots, 0, \tilde{g}_0, \dots, \tilde{g}_q)$. Since \tilde{g} is linearly independent from e and $\dim(f_4(a)(V)) = 1$, we see that there is some scalar δ for which $(\delta \tilde{g}_0, \dots, \delta \tilde{g}_q) = (d'_0, \dots, d'_q)$. Set $h := d - \delta \tilde{g}$. Using the fact that $\dim(f_2(a)(V)) = 1$ we see that $\tilde{g}_0 \neq 0$. Set $g := \tilde{g}_0^{-1} \tilde{g}$. \square

Lemma 4.13. *For $0 \leq j \leq q$ we have $g_j = h_j$.*

Proof. We arranged $g_0 = 1 = h_0$ in Lemma 4.12. The general element of V has the form $\alpha h + \beta g$ for some $(\alpha, \beta) \in \mathbb{Q}^2$. For $0 < j \leq q$, the general element of $f_{4+j}(a)(V)$ has the form $(\alpha h_0 - \beta h_0, \alpha h_j - \beta g_j, 0, \dots, 0) = ((\alpha - \beta)h_0, (\alpha - \beta)h_j + \beta(h_j - g_j), 0, \dots, 0)$. Since $h_0 \neq 0$, if $h_j \neq g_j$, then this vector space would have dimension two, contrary to the requirement that $f(a) \in E_{\nu, \rho, \rho_1, \dots, \rho_\mu}$. \square

Lemma 4.14. *For $0 < j \leq t$ we have $h_j = a_j$.*

Proof. The image of $\alpha h + \beta g$ under $f_{4+q+k+j}(a)$ is $(p_j \alpha - q_j \alpha h_j, 0, \dots, 0)$ where $a_j = \frac{p_j}{q_j}$ in lowest terms. Since $\dim(f_{4+q+k+j}(a)(V)) = 0$, we have $q_j h_j = p_j$. That is, $h_j = a_j$. \square

Lemma 4.15. *For any $F \in S$, we have $F(h_1, \dots, h_q) = 0$.*

Proof. If F is an affine polynomial, that is, if $F = \lambda_{0,j} + \sum_{i=1}^q \lambda_{i,j} x_i$ for some $0 < j \leq \ell$, then because $\dim(f_{4+q+k+t+j}(a)(V)) = 0$, we have $\lambda_{0,j} + \sum_{i=1}^q \lambda_{i,j} h_i = 0$. On the other hand, if F is a multiplicative relation, that is, if $F = x_{i_1,j} x_{i_2,j} - x_{i_3,j}$ for some $0 < j \leq k$, then because $\dim(f_{4+q+j}(a)(V)) = 1$ we see that there is some scalar γ so that for any pair (α, β) we have $\gamma(\alpha h_{i_1,j} + \beta) = \alpha h_{i_3,j} + \beta h_{i_2,j}$. Specializing to $\beta = 0$ and $\alpha = 1$, we have $\gamma = \frac{h_{i_3,j}}{h_{i_1,j}}$ (unless both are zero, in which case the equality $h_{i_1,j} h_{i_2,j} = h_{i_3,j}$ holds anyway), which we substitute to obtain $\frac{h_{i_3,j}}{h_{i_1,j}} = h_{i_2,j}$, or $h_{i_2,j} h_{i_1,j} = h_{i_3,j}$. \square

Taking $b := (h_{t+1}, \dots, h_q)$ we see that $(a, b) \in V(S)(\mathbb{Q})$. \square

Acknowledgements

We thank Bernd Sturmfels for suggesting the connection with Hilbert’s Tenth Problem over the rationals.

The research of various of the authors was supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under awards DE-SC0004938, DE-SC0003959, and DE-SC0010200; by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, X-Stack program under awards DE-SC0005136, DE-SC0008699, DE-SC0008700, and AC02-05CH11231; by DARPA award HR0011-12-2-0016 with contributions from Intel, Oracle, and MathWorks, and by NSF grants DMS-1363324 and DMS-1001550.

References

- [1] C. Bennett and R.C. Sharpley. *Interpolation of Operators*, volume 129 of *Pure and Applied Mathematics*. Academic Press, 1988.
- [2] J. Bennett, A. Carbery, M. Christ, and T. Tao. The Brascamp-Lieb inequalities: finiteness, structure and extremals. *Geom. Funct. Anal.*, 17(5):1343–1415, 2008.
- [3] J. Bennett, A. Carbery, M. Christ, and T. Tao. Finite bounds for Hölder-Brascamp-Lieb multilinear inequalities. *Mathematical Research Letters*, 17(4):647–666, 2010.
- [4] E. A. Carlen, E. H. Lieb, and M. Loss. A sharp analog of Young’s inequality on S^N and related entropy inequalities. *Jour. Geom. Anal.*, 14:487–520, 2004.
- [5] M. Christ. Optimal constants in Hölder-Brascamp-Lieb inequalities for discrete Abelian groups. 2013. arXiv preprint 1307.8442.
- [6] M. Christ, J. Demmel, N. Knight, T. Scanlon, and K. Yelick. Communication lower bounds and optimal algorithms for programs that reference arrays. *in preparation*, 2014.
- [7] J. Koenigsmann. Defining \mathbb{Z} in \mathbb{Q} . *Annals of Mathematics*, to appear. arXiv preprint 1011.3424.
- [8] Y. Matiyasevich. *Hilbert’s Tenth Problem*. MIT Press, 1993.
- [9] R. Vakil. Murphy’s law in algebraic geometry: badly-behaved deformation spaces. *Inventiones Mathematicae*, 164(3):569–590, 2006.
- [10] S.I. Valdimarsson. The Brascamp-Lieb polyhedron. *Canadian Journal of Mathematics*, 62(4):870–888, 2010.