THE BRASCAMP–LIEB INEQUALITIES: FINITENESS, STRUCTURE AND EXTREMALS

JONATHAN BENNETT, ANTHONY CARBERY, MICHAEL CHRIST, AND TERENCE TAO

Abstract. We consider the Brascamp–Lieb inequalities concerning multilinear integrals of products of functions in several dimensions. We give a complete treatment of the issues of finiteness of the constant, and of the existence and uniqueness of centred gaussian extremals. For arbitrary extremals we completely address the issue of existence, and partly address the issue of uniqueness. We also analyse the inequalities from a structural perspective. Our main tool is a monotonicity formula for positive solutions to heat equations in linear and multilinear settings, which was first used in this type of setting by Carlen, Lieb, and Loss [CLL]. In that paper, the heat flow method was used to obtain the rank one case of Lieb's fundamental theorem concerning exhaustion by gaussians; we extend the technique to the higher rank case, giving two new proofs of the general rank case of Lieb's theorem.

1. Introduction

Important inequalities such as the multilinear Hölder inequality, the sharp Young convolution inequality and the Loomis–Whitney inequality find their natural generalisation in the *Brascamp–Lieb inequalities*, which we now describe.

Definition 1.1 (Brascamp–Lieb constant)**.** Define a *Euclidean space* to be a finite-dimensional real Hilbert space[1] H, endowed with the usual Lebesgue measure dx; for instance, $\mathbf{R}^n$ is a Euclidean space for any n. If $m \geq 0$ is an integer, we define an *m-transformation* to be a triple

$$\mathbf{B} := (H, (H_j)_{1\leq j\leq m}, (B_j)_{1\leq j\leq m})$$

where $H, H_1, \dots, H_m$ are Euclidean spaces and for each j, $B_j : H \to H_j$ is a linear transformation. We refer to H as the *domain* of the m-transformation $\mathbf{B}$. We say that an m-transformation is *non-degenerate* if all the B_j are surjective (thus $H_j = B_jH$) and the common kernel is trivial (thus $\bigcap_{j=1}^m \ker(B_j) = \{0\}$). We define an *$m$-exponent* to be an m-tuple $\mathbf{p} = (p_j)_{1\leq j\leq m} \in \mathbf{R}_+^m$ of non-negative real numbers. We define a *Brascamp–Lieb datum* to be a pair $(\mathbf{B}, \mathbf{p})$, where $\mathbf{B}$ is an

The third author was supported in part by NSF grant DMS-040126.

[1]It is convenient to work with arbitrary finite-dimensional Hilbert spaces instead of just copies of $\mathbf{R}^n$ in order to take advantage of invariance under Hilbert space isometries, as well as such operations as restriction of a Hilbert space to a subspace, or quotienting one Hilbert space by another. In fact one could dispense with the inner product structure altogether and work with finite-dimensional vector spaces with a Haar measure dx, but as the notation in that setting is less familiar, especially when regarding heat equations on such domains, we shall retain the inner product structure for notational convenience.

m-transformation and $\mathbf{p}$ is an m-exponent for some integer $m \geq 0$. When we are in a situation which involves a Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p})$, it is always understood that the objects H, H_j, B_j, p_j denote the relevant components of this Brascamp–Lieb datum. If $(\mathbf{B}, \mathbf{p})$ is a Brascamp–Lieb datum, we define an *input* for $(\mathbf{B}, \mathbf{p})$ to be an m-tuple $\mathbf{f} := (f_j)_{1 \leq j \leq m}$ of nonnegative measurable functions $f_j : H_j \to \mathbf{R}^+$ such that $0 < \int_{H_j} f_j < \infty$, and then define the quantity $0 \leq \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}) \leq +\infty$ by the formula

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}) := \frac{\int_H \prod_{j=1}^m (f_j \circ B_j)^{p_j}}{\prod_{j=1}^m (\int_{H_j} f_j)^{p_j}}.$$

Note that if $(\mathbf{B}, \mathbf{p})$ is non-degenerate and the f_j are bounded with compact support, then $\mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}) < +\infty$. We then define the *Brascamp–Lieb constant* $\mathrm{BL}(\mathbf{B}, \mathbf{p}) \in (0, +\infty]$ to be the supremum of $\mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f})$ over all inputs $\mathbf{f}$. Equivalently, $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ is the smallest constant for which the m-linear Brascamp–Lieb inequality

$$\int_H \prod_{j=1}^m (f_j \circ B_j)^{p_j} \leq \mathrm{BL}(\mathbf{B}, \mathbf{p}) \prod_{j=1}^m (\int_{H_j} f_j)^{p_j} \tag{1}$$

holds for nonnegative measurable functions $f_j : H_j \to \mathbf{R}^+$.

Remark 1.2. By testing (1) on functions which are strictly positive near the origin, one can easily verify that the Brascamp–Lieb constant must be strictly positive, though it can of course be infinite. We give this definition assuming only that the inputs f_j are non-negative measurable, but it is easy to see (using Fatou's lemma) that one could just as easily work with strictly positive Schwartz functions with no change in the Brascamp–Lieb constant. One can of course define $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ when $\mathbf{B}$ is degenerate but it is easily seen that this constant is infinite in that case (see also Lemma 4.1). Thus we shall often restrict our attention to non-degenerate Brascamp–Lieb data.

We now give some standard examples of Brascamp–Lieb data and their associated Brascamp–Lieb constants.

Example 1.3 (Hölder's inequality)**.** *If* $\mathbf{B}$ *is the non-degenerate* m*-transformation*

$$\mathbf{B} := (H, (H)_{1 \leq j \leq m}, (\mathrm{id}_H)_{1 \leq j \leq m})$$

for some Euclidean space H *and some* $m \geq 1$*, where* $\mathrm{id}_H : H \to H$ *denotes the identity on* H*, then the multilinear Hölder inequality asserts that* $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ *is equal to 1 when* $p_1 + \ldots + p_m = 1$*, and is equal to* $+\infty$ *otherwise.*

Example 1.4 (Loomis–Whitney inequality)**.** *If* $\mathbf{B}$ *is the non-degenerate* n *- transformation*

$$B := (\mathbf{R}^n, (e_j^\perp)_{1 \leq j \leq n}, (P_j)_{1 \leq j \leq n})$$

where $e_1, \ldots, e_n$ *is the standard basis of* $\mathbf{R}^n$*,* $e_j^\perp \subset \mathbf{R}^n$ *is the orthogonal complement of* e_j*, and* $P_j : \mathbf{R}^n \to e_j^\perp$ *is the orthogonal projection onto* e_j*, then the Loomis–Whitney inequality* [LW] *can be interpreted as an assertion that* $\mathrm{BL}(\mathbf{B}, \mathbf{p}) = 1$ *when* $\mathbf{p} = (\frac{1}{n-1}, \ldots, \frac{1}{n-1})$*, and is infinite for any other value of* $\mathbf{p}$*. For instance, when* $n = 3$ *this inequality asserts that*

$$\int\int\int f(y,z)^{1/2} g(x,z)^{1/2} h(x,y)^{1/2}\ dxdydz \leq \|f\|_{L^1(\mathbf{R}^2)}^{1/2} \|g\|_{L^1(\mathbf{R}^2)}^{1/2} \|h\|_{L^1(\mathbf{R}^2)}^{1/2} \tag{2}$$

whenever f, g, h are non-negative measurable functions on $\mathbf{R}^2$. More generally, Finner [F] *established multilinear inequalities of Loomis-Whitney type involving orthogonal projections to co-ordinate subspaces.*

Example 1.5 (Sharp Young inequality)**.** *The sharp Young inequality (*[Be], [BL]*) can be viewed as an assertion that if $\mathbf{B}$ is the non-degenerate 3-transformation*

$$\mathbf{B} := (\mathbf{R}^d \times \mathbf{R}^d, (\mathbf{R}^d)_{1\le j\le 3}, (B_j)_{1\le j\le 3})$$

where $d \ge 1$ is an integer and the maps $B_j : \mathbf{R}^d \times \mathbf{R}^d \to \mathbf{R}^d$ are defined for $j = 1, 2, 3$ by

$$B_1(x, y) := x; \quad B_2(x, y) := y; \quad B_3(x, y) := x - y$$

then we have

$$\mathrm{BL}(\mathbf{B}, (p_1, p_2, p_3)) = \left(\prod_{j=1}^{3} \frac{(1 - p_j)^{1-p_j}}{p_j^{p_j}} \right)^{d/2}$$

if $p_1 + p_2 + p_3 = 2$ and $0 \le p_1, p_2, p_3 \le 1$, with $\mathrm{BL}(\mathbf{B}, (p_1, p_2, p_3)) = +\infty$ for any other values of (p_1, p_2, p_3). (See Example 3.8.)

Example 1.6 (Geometric Brascamp–Lieb inequality)**.** *Let $\mathbf{B}$ be the non-degenerate m-transformation*

$$\mathbf{B} := (H, (H_j)_{1\le j\le m}, (B_j)_{1\le j\le m})$$

where H is a Euclidean space, $H_1, \dots, H_m$ are subspaces of H, and $B_j : H \to H_j$ are orthogonal projections, thus the adjoint $B_j^* : H_j \to H$ is the inclusion map. The geometric Brascamp–Lieb inequality of Ball [Ball] *(and later generalised by Barthe* [Bar2]*) asserts that $\mathrm{BL}(\mathbf{B}, \mathbf{p}) = 1$ whenever $\mathbf{p} = (p_1, \dots, p_m) \in \mathbf{R}_+^m$ obeys the identity $\sum_{j=1}^m p_j B_j^* B_j = 1$. Note that equality is obtained when we have $f_j(x) = \exp(-\pi \|x\|_{H_j}^2)$ for all j. This significantly generalises the Loomis–Whitney inequality (Example 1.4) and the Hölder inequality (Example 1.3).*

Example 1.7 (Rank-one Brascamp–Lieb inequality)**.** *Let $\mathbf{B}$ be the non-degenerate m-transformation*

$$\mathbf{B} := (H, (\mathbf{R})_{1\le j\le m}, (v_j^*)_{1\le j\le m})$$

where H is a Euclidean space, $v_1, \dots, v_m$ are non-zero vectors in H which span H, and $v_j^* : H \to \mathbf{R}$ is the corresponding linear functional $v_j^*(x) := \langle v_j, x\rangle_H$. Then the work of Barthe [Bar2] *(see also* [CLL]*) shows that $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ is finite if and only if $\mathbf{p} \in \mathbf{R}_+^m$ lies in the convex polytope whose vertices are the points $(1_{j\in I})_{1\le j\le m}$, where I is a subset of $\{1, \dots, m\}$ such that the vectors $(v_j)_{j\in I}$ form a basis of H (in particular, this forces $|I| = \dim(H)$). Furthermore, if $\mathbf{p}$ lies in the $(m-1)$-dimensional interior of this polytope, then equality in* (1) *can be attained. We shall reprove these statements as Theorem 5.5.*

Remark 1.8. The rank one case can differ dramatically from the general case. In particular, rearrangement inequalities such as that of Brascamp, Lieb, and Luttinger [BLL] apply for rank one, that is, when all H_j have dimensions one, but for higher rank are only very rarely applicable.

It is thus of interest to compute the Brascamp–Lieb constants $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ explicitly, or at least to determine under what conditions these constants are finite. A fundamental theorem of Lieb [L] shows that these constants are exhausted by centred gaussians. More precisely, given any positive definite transformation[2] $A : H \to H$, we consider the associated gaussian $\exp(-\pi\langle Ax, x\rangle_H)$. As is well-known we have the formula

$$\int_H \exp(-\pi\langle Ax, x\rangle)\, dx = (\det_H A)^{-1/2}, \tag{3}$$

where $\det_H$ is the determinant associated to transformations on the Euclidean space H. Define a *gaussian input* for $(\mathbf{B}, \mathbf{p})$ to be any m-tuple $\mathbf{A} = (A_j)_{1\le j\le m}$ of positive definite linear transformations $A_j : H_j \to H_j$. If we now test the inequality (1) with the input $(\exp(-\pi\langle A_j x, x\rangle_H))_{1\le j\le m}$ for some arbitrary positive-definite transformations $A_j : H_j \to H_j$, we conclude that

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) \ge \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A})$$

for arbitrary gaussian input $\mathbf{A} = (A_j)_{1\le j\le m}$, where $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}) \in (0, +\infty)$ is the quantity

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}) := \left(\frac{\prod_{j=1}^m (\det_{H_j} A_j)^{p_j}}{\det_H(\sum_{j=1}^m p_j B_j^* A_j B_j)} \right)^{1/2}, \tag{4}$$

and $B_j^* : H_j \to H$ is the adjoint of H. In particular, if we define

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = \sup\{\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}) : \mathbf{A} \text{ is a gaussian input for } (\mathbf{B}, \mathbf{p})\} \tag{5}$$

then we have

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) \ge \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) \tag{6}$$

for arbitrary Brascamp–Lieb data $(\mathbf{B}, \mathbf{p})$.

Lieb [L] showed that the above inequality is in fact an equality:

Theorem 1.9 (Lieb's theorem)**.** [L] *For any Brascamp–Lieb datum* $(\mathbf{B}, \mathbf{p})$*, we have* $\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$.

Remark 1.10. Lieb's theorem was proved in the rank-one and certain intermediate cases by Brascamp and Lieb [BL] in part using rearrangement techniques from [BLL]. Lieb gave the first full proof in [L] using, *inter alia*, a clever $O(2)$-invariance and arguments related to the central limit theorem. Barthe gave an alternative proof using deep ideas of transportation of mass due to Brenier, McCann and Caffarelli; see [Bar2] and the references therein. More relevantly to our own approach, Carlen, Lieb and Loss [CLL] gave a proof of Lieb's theorem in the rank one case using heat flow methods. Our own arguments, though rediscovered independently, can be viewed as an extension of the arguments in [CLL] to the higher rank case.

The question of understanding the Brascamp–Lieb constants is solved by Theorem 1.9, in the sense that the task is reduced to the simpler task of understanding the (in principle computable) gaussian Brascamp–Lieb constants. However, there are still a number of issues that are not easily resolved just from this theorem

[2]By positive definite, we mean that the transformation is self-adjoint and that the associated quadratic form $\langle Ax, x\rangle$ is positive definite. It is in fact slightly more natural to view A as a transformation from H to H^*, the dual of H, rather than H itself, but of course since H and H^* are canonically identifiable using the Hilbert space structure we will usually not bother to enforce the distinction between H and H^*.

alone; for instance, it is not immediately obvious what the necessary and sufficient conditions are for either $\mathrm{BL}(\mathbf{B},\mathbf{p})$ or $\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p})$ to be finite. Also, this theorem does not characterise the extremals for (1) or for (5), or even address whether such extremals exist at all.

More precisely, we formulate

Definition 1.11 (Extremisability)**.** A Brascamp–Lieb datum is said to be *extremisable* if $\mathrm{BL}(\mathbf{B},\mathbf{p})$ is finite and there exists an input $\mathbf{f}$ for which $\mathrm{BL}(\mathbf{B},\mathbf{p}) = \mathrm{BL}(\mathbf{B},\mathbf{p};\mathbf{f})$.

A Brascamp–Lieb datum is said to be *gaussian-extremisable* if there exists a gaussian input $\mathbf{A}$ for which $\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p};\mathbf{A})$ (so in particular, $\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p})$ is finite).

Note that gaussian-extremisability does not formally imply extremisability. Nonetheless, Lieb's theorem asserts that every gaussian-extremisable datum is also extremisable. However it does not give the converse (which turns out to be true, see Proposition 6.5 below).

In this paper we shall address these issues as follows. Firstly, we shall use the multilinear heat flow monotonicity formula technique to give two fully self-contained proofs of the general-rank case of Lieb's theorem (in Section 5 and Section 8 respectively). This technique was first employed in this setting [CLL] to establish the analogues of most of the results here in the rank-one case $\dim(H_j) = 1$ (and also for more general domains than Euclidean spaces), and our arguments are thus an extension of those in [CLL]. (See also [BarC] for another recent application of the heat flow method to Brascamp–Lieb type inequalities.) Secondly, we establish a characterisation of when the Brascamp–Lieb constants are finite; in fact this will be achieved simultaneously with the first of our two new proofs of Lieb's theorem. Thirdly, we shall give necessary and sufficient conditions for the Brascamp–Lieb data $(\mathbf{B},\mathbf{p})$ to be gaussian-extremisable. Fourthly we shall give necessary and sufficient conditions for the Brascamp–Lieb data $(\mathbf{B},\mathbf{p})$ to possess unique gaussian extremisers (up to trivial symmetries). We shall achieve these results and others partly with the aid of a two-stage structural perspective on the Brascamp–Lieb inequalities (1).

In order to describe our criterion for finiteness we first need a definition. This definition (or more precisely an equivalent formulation of this definition) was first introduced in the context of the rank one problem in [CLL].

Definition 1.12 (Critical subspace and simplicity)**.** [CLL] Let $(\mathbf{B},\mathbf{p})$ be a Brascamp–Lieb datum. A *critical subspace* V for $(\mathbf{B},\mathbf{p})$ is a non-zero proper subspace of H such that

$$\dim(V) = \sum_{j=1}^{m} p_j \dim(B_j V).$$

The datum $(\mathbf{B},\mathbf{p})$ is *simple* if it has no critical subspaces.

Theorem 1.13. *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum. Then* $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ *is finite if and only if we have the scaling condition*

$$\dim(H) = \sum_j p_j \dim(H_j) \tag{7}$$

and the dimension condition

$$\dim(V) \leq \sum_{j=1}^m p_j \dim(B_j V) \text{ for all subspaces } V \subseteq H. \tag{8}$$

Furthermore, if $(\mathbf{B}, \mathbf{p})$ *is simple, then it is gaussian-extremisable.*

Remark 1.14. In the rank-one case, a finiteness criterion for $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ and $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ equivalent to Theorem 1.13 was given by Barthe [Bar2], see Theorem 5.5 below.

By Lieb's theorem we have as an immediate corollary:

Theorem 1.15 (Finiteness of Brascamp–Lieb constant)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum. Then the following three statements are equivalent.*

(a) $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ *is finite.*
(b) $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ *is finite.*
(c) (7) *and* (8) *hold.*

Furthermore, if any of (a)-(c) hold, and $(\mathbf{B}, \mathbf{p})$ *is simple, then* $(\mathbf{B}, \mathbf{p})$ *is extremisable.*

Remarks 1.16. The conditions (7), (8) imply in particular that $\mathbf{B}$ is non-degenerate, as can be seen by testing on $V := H$ and on $V := \bigcap_{j=1}^m \ker(B_j)$. The deduction of (c) from (a) or (b) is very easy, see Lemma 4.1. The more difficult part of the theorem is to establish the reverse implication in Theorem 1.13; we shall do so in Proposition 5.2. One can also easily show that Theorem 1.15 implies all the negative results in Examples 1.3-1.6.

Remark 1.17. A key element of our analysis is a certain factorisation of the inequality through critical subspaces. This factorisation method played a similarly key role in the work of [CLL] in the rank one case, and also in the work of Finner [F], who analysed the case of orthogonal projections to co-ordinate subspaces. One can view this factorization method as generalizing the arguments of Loomis and Whitney [LW]. In the companion paper [BCCT] we give an alternative proof of the equivalence (c) $\iff$ (a) in Theorem 1.15, without recourse to Lieb's theorem, heat flow deformation, or related techniques, based on this factorisation together with multilinear interpolation. Our first proof of Theorem 1.9 (given in Section 5) combines heat flows with the factorisation, while a second proof (in Section 8) is a pure heat flow argument. The factorisation also plays a central role in our structural perspective.

Remark 1.18. Theorem 1.15 implies in particular that for any fixed $\mathbf{B}$, the set of all $\mathbf{p}$ for which $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ or $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ is finite is a convex polytope, with the faces determined by the natural numbers $n, n_1, \ldots, n_m$ for which an n-dimensional subspace V of H exists with n_j-dimensional images $B_j V$. This is a purely geometric condition, which can in principle be computed algebraically. However, the problem of determining which $n, n_1, \dots, n_m$ are attainable seems to be a problem in Schubert calculus, and given the algebraic richness and complexity of this calculus (see e.g. [Bel]), a fully explicit and easily computable description of these numbers (and hence of the above polytope) may be too ambitious to hope for in general. However in the rank-one case (Example 1.7) a concrete description of the polytope has been given by Barthe [Bar2]. As already observed in [CLL], Barthe's description of the polytope is equivalent to that given by Theorem 1.15; we recall this equivalence in Section 5.

Our proofs of Lieb's theorem and Theorem 1.15 will occupy Sections 2-5. However our methods provide a quite short proof of the geometric Brascamp–Lieb inequality (Example 1.6) which does not rely on Lieb's theorem; this is Proposition 2.8. Moreover, the geometric Brascamp–Lieb inequality, combined with a simple linear change of variables argument, is already strong enough to obtain many of the standard applications of the Brascamp–Lieb inequality, such as the sharp Young inequality (Example 1.5), as well as the special case of Theorem 1.15 and Lieb's theorem when $(\mathbf{B}, \mathbf{p})$ is simple. It will also be one of two pillars of our first proof of Lieb's theorem, the other pillar being the aforementioned factorisation argument.

In the work of Barthe [Bar2] (Example 1.7), a structural analysis of the rank-one Brascamp–Lieb functional (4) was given, involving in particular a decomposition of the rank-one Brascamp–Lieb data into *indecomposable* components. Barthe also considered the question of when extremals to (5) exist, and when they are unique. The uniqueness question was answered based upon this decomposition and upon an explicit algebraic description of the gaussian Brascamp–Lieb constant (5) in the rank-one case. Part of the purpose of this paper is to extend (albeit by different methods) Barthe's programme to the higher rank case. Barthe's decomposition of Brascamp–Lieb data is different from the factorisation method used here (and also in [CLL], [F]); the difference is analogous to the distinction between direct product and semi-direct product in group theory. Barthe's decomposition depends only on the data $\mathbf{B}$ while ours depends also upon $\mathbf{p}$. We will combine a higher-rank analogue of Barthe's decomposition of Brascamp–Lieb data into indecomposable components with our notion of factorisation in order to answer the question of when Brascamp–Lieb data is gaussian-extremisable. For precise definitions see Section 7 below.

Theorem 1.19. *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum for which* (7) *and* (8) *hold. Then* $(\mathbf{B}, \mathbf{p})$ *is gaussian-extremisable if and only if each indecomposable component for* $\mathbf{B}$ *is simple.*

For the full story see Theorem 7.13 below. It is also shown in Proposition 6.5 that extremisability is equivalent to gaussian-extremisability. This latter equivalence uses Lieb's theorem and an iterated convolution idea of Ball. Thus extremisability can be viewed as a kind of *semisimplicity* for the Brascamp–Lieb datum.

Theorem 7.13 gives a satisfactory description of the Brascamp–Lieb data for which extremisers or gaussian extremisers exist. We can also show that these gaussian extremisers are *unique* (up to scaling) if and only if the data is simple:

Theorem 1.20. *Let* $(\mathbf{B}, \mathbf{p})$ *be Brascamp–Lieb datum with* $p_j > 0$ *for all* j*. Then gaussian extremizers for* $(\mathbf{B}, \mathbf{p})$ *exist and are unique (up to scaling) if and only if* $(\mathbf{B}, \mathbf{p})$ *is simple.*

See Corollary 9.2. As for the uniqueness (up to trivial symmetries) of general extremisers, this problem seems to be more difficult, and we have only a partial result (Theorem 9.3).

As mentioned above, Carlen, Lieb and Loss in [CLL] have introduced the idea of using heat flow in the Brascamp–Lieb context in the rank-one case (and also on the sphere $\mathbf{S}^n$ instead of a Euclidean space). We remark also that prior to the work in [CLL], the known proofs of versions of Lieb's theorem relied on methods such as rearrangement inequalities [BLL] and mass transfer inequalities [Bar2]. While such methods are similar in spirit to that of heat flows – in that they are all ways of deforming non-extremal solutions to extremal gaussians – they are not identical. For instance the heat flow method will continuously deform sums of gaussians to other sums of gaussians, whereas continuous mass transfer does not achieve this; and it is not obvious how to perform rearrangement in a continuous manner.

Guide to the paper. In Section 2 we introduce the heat flow method and use it to prove the geometric Brascamp–Lieb inequality. In Section 3 we consider the gaussian-extremisable case and give a characterisation of gaussian extremisers, recovering the sharp Young inequality Example 1.5 as an immediate application. Our first approach to the structure of the Brascamp–Lieb inequalities, via factorisation, is taken up in Section 4 where we also establish the necessity of (8) and (7) for $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ or $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ to be finite. In Section 5 we establish the sufficiency of these conditions and prove Theorem 1.9 and Theorem 1.15 in the general case. In Section 6 we prove that extremisability and gaussian-extremisability are equivalent. Our second approach to structural issues is made in Section 7 where the main characterisation of extremisability is also given. In Section 8 we examine variants of the Brascamp–Lieb inequalities for regularised inputs in a gaussian setting, and this leads to our second proof of Lieb's theorem via a purely heat-flow method. As a further consequence of our analysis, we give versions of Theorem 1.9 and Theorem 1.15 in a gaussian-localised setting in Corollary 8.16 and Theorem 8.17 respectively. Corollary 8.16 is very much in the spirit of Lieb's original paper [L]. In Section 9 we discuss uniqueness of extremals. Finally, in Section 10 we make some remarks on the heat flow method in so far as it applies in non-gaussian contexts. We give a general monotonicity formula for log-concave kernels (the class of which of course includes gaussians) in Lemma 10.4.

Our work in this paper was motivated by the forthcoming article [BCT] in which we obtain almost optimal results for the multilinear Kakeya and restriction problems. It was in this context that we rediscovered the applicability of heat flows in multilinear inequalities.

Acknowledgements. We are grateful to Assaf Naor for bringing [CLL] to our attention. We also thank Rodrigo Bañuelos for some corrections, and Eric Carlen for help with the references.

2. The geometric Brascamp–Lieb inequality

In this section we establish Lieb's theorem in the "geometric" case (see Proposition 2.8). We begin with a definition.

Definition 2.1 (Geometric Brascamp–Lieb data)**.** A Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p})$ is said to be *geometric* if we have $B_j B_j^* = \mathrm{id}_{H_j}$ (thus B_j^* is an isometry) and

$$\sum_{j=1}^{m} p_j B_j^* B_j = \mathrm{id}_H . \tag{9}$$

Remark 2.2. The data in Example 1.6 is of this type. Conversely, if $(\mathbf{B}, \mathbf{p})$ is geometric, and we let E_j be the range of the isometry B_j^*, then we may identify H_j with E_j, and B_j with the orthogonal projection from H to E_j, thus placing ourselves in the situation of Example 1.6. The "geometry" here refers to Euclidean or Hilbert space geometry, since the inner product structure of H is clearly being used (for instance, to define the notion of orthogonal projection). We remark that the condition $B_j B_j^* = \mathrm{id}_{H_j}$ forces B_j to be surjective, and (9) implies that $\bigcap_{j=1}^m \ker(B_j) = \{0\}$, and thus geometric Brascamp–Lieb data is always non-degenerate. Furthermore, by taking traces of (9) we conclude (7). More generally, if we let V be any subspace of H, and let $\Pi : H \to H$ be the orthogonal projection onto V, then by multiplying (9) by Π and taking traces we conclude

$$\sum_{j=1}^{m} p_j \,\mathrm{tr}(\Pi B_j^* B_j) = \dim(V).$$

Since $\Pi B_j^* B_j$ is a contraction and has range $B_j V$, we conclude (8) also. Thus the assertion that geometric Brascamp–Lieb data have finite Brascamp–Lieb constants is a special case of Theorem 1.15; we establish this special case in Proposition 2.8 below.

Remark 2.3. Let $(\mathbf{B}, \mathbf{p})$ be a Brascamp–Lieb datum with $p_j > 0$ for all j, and let $\bigoplus_{j=1}^m p_j H_j$ be the Cartesian product $\prod_{j=1}^m H_j$ endowed with the inner product $\langle x, y\rangle_{\bigoplus_{j=1}^m p_j H_j} := \sum_{j=1}^m p_j \langle x_j, y_j\rangle_{H_j}$; this is thus a Euclidean space. Let V be the image of H in $\bigoplus_{j=1}^m p_j H_j$ under the map $x \mapsto (B_j x)_{1\le j\le m}$, and let $\pi_j : V \to H_j$ be the projection maps induced by restricting the projection maps on $\bigoplus_{j=1}^m p_j H_j$ to V. Then the statement that $(\mathbf{B}, \mathbf{p})$ is geometric is equivalent to the assertion that the map $x \mapsto (B_j x)_{1\le j\le m}$ is a Euclidean isomorphism from H to V, and that the projection maps are co-isometries (i.e. surjective partial isometries, or equivalently their adjoints are isometries). Thus geometric Brascamp–Lieb data can be thought of as subspaces in $\bigoplus_{j=1}^m p_j H_j$ which lie "co-isometrically" above each of the factor spaces H_j.

Our approach to the geometric Brascamp–Lieb inequality is via monotonicity formulae, which as observed in [CLL] seems to be especially well suited for the geometric Brascamp–Lieb setting. Abstractly speaking, if one wishes to prove an inequality of the form $A \le B$, one can do so by constructing a monotone non-decreasing quantity $Q(t)$ which equals A when $t = 0$ (say) and equals B when $t = \infty$ (say). It is thus of interest to have a general scheme for generating such monotone quantities, which we will also use later in this paper to deduce some variants of the Brascamp–Lieb inequalities. We begin with the following simple lemma.

Lemma 2.4 (Monotonicity for transport equations)**.** *Let $I \subseteq \mathbf{R}$ be a time interval, let H be a Euclidean space, let $u : I \times H \to \mathbf{R}^+$ be a smooth non-negative function, and $\vec{v} : I \times H \to H$ be a smooth vector field, such that $\vec{v}u$ is rapidly decreasing at spatial infinity $x \to \infty$ locally uniformly on I. Let $\alpha \in \mathbf{R}$ be fixed. Suppose that we have the transport inequality*

$$\partial_t u(t,x) + \operatorname{div}(\vec{v}(t,x)u(t,x)) \geq \alpha u/t \tag{10}$$

for all $(t,x) \in I \times H$, where div *of course is the spatial divergence on the Euclidean space H. Then the quantity*

$$Q(t) := t^{-\alpha} \int_H u(t,x)\ dx \in [0,+\infty]$$

is non-decreasing in time. Furthermore, if (10) *holds with strict inequality for all x,t, then Q is strictly increasing in time. Similarly, if the signs are reversed in* (10)*, then $Q(t)$ is now non-increasing in time.*

Remark 2.5. We allow the velocity field $\vec{v}$ to vary in both space and time. This is important as we will typically be dealing with functions u which solve a heat equation such as $\partial_t u = \Delta u$, in which case the natural velocity vield $\vec{v}$ is given by $\vec{v} = -\nabla \log u$.

Proof By multiplying (10) by t^α and replacing u by $\tilde{u} := t^{-\alpha}u$, we may reduce to the case $\alpha = 0$. Let $t_1 < t_2$ be two times in I. From Stokes' theorem we have

$$\int_H u(t_2,x)\psi(x)\ dx - \int_H u(t_1,x)\psi(x)\ dx \\ = \int_{t_1}^{t_2} \int_H (\partial_t u(t,x)\psi(x) + \operatorname{div}(\psi(x)\vec{v}(t,x)u(t,x))\ dxdt$$

for any non-negative smooth cutoff function ψ. Using the product rule and (10) we conclude

$$\int_H u(t_2,x)\psi(x)\ dx - \int_H u(t_1,x)\psi(x)\ dx \geq \int_{t_1}^{t_2} \int_H \langle \nabla\psi(x), \vec{v}(t,x)u(t,x)\rangle\ dx.$$

Letting ψ approach the constant function 1 and using the hypothesis that $\vec{v}u$ is rapidly decreasing at spatial infinity, uniformly in $[t_1,t_2]$, we obtain the claim. ■

We now generalise this result substantially.

Lemma 2.6 (Multilinear monotonicity for transport equations)**.** *Let $p_1, \ldots, p_m > 0$ be real exponents, let H be a Euclidean space, and for each $1 \leq j \leq m$ let $u_j : \mathbf{R}^+ \times H \to \mathbf{R}^+$ be a smooth strictly positive function, and $\vec{v}_j : \mathbf{R}^+ \times H \to H$ be a smooth vector field. Let $\alpha \in \mathbf{R}$ be fixed. Suppose we also have an additional smooth vector field $\vec{v} : \mathbf{R}^+ \times H \to H$, such that $\vec{v}\prod_{j=1}^m u_j^{p_j}$ is rapidly decreasing in* *space locally uniformly on I, and we have the inequalities*

$$\partial_t u_j(t,x) + \operatorname{div}(\vec{v}_j u_j(t,x)) \geq 0 \text{ for all } 1 \leq j \leq m \tag{11}$$

$$\operatorname{div}(\vec{v} - \sum_{j=1}^m p_j \vec{v}_j) \geq \alpha/t \tag{12}$$

$$\sum_{j=1}^m p_j \langle \vec{v} - \vec{v}_j, \nabla \log u_j \rangle_H \geq 0. \tag{13}$$

Then the quantity

$$Q(t) := t^{-\alpha} \int_H \prod_{j=1}^m u_j(t,x)^{p_j}\ dx \tag{14}$$

is non-decreasing in time. Furthermore, if at least one of (11), (12), (13) *holds with strict inequality for all x, t, then Q is strictly increasing in time. If all the signs in* (11), (12), (13) *are reversed, then $Q(t)$ is now non-increasing in time.*

Proof We begin with the first claim. By Lemma 2.4, it suffices to show that

$$\partial_t \prod_{j=1}^m u_j^{p_j} + \operatorname{div}(\vec{v} \prod_{j=1}^m u_j^{p_j}) \geq \alpha \prod_{j=1}^m u_j^{p_j}.$$

We divide both sides by the positive quantity $\prod_{j=1}^m u_j^{p_j}$ and then use the product rule, to reduce to showing that

$$\sum_{j=1}^m p_j \frac{\partial_t u_j}{u_j} + \operatorname{div}(\vec{v}) + \langle \vec{v}, \sum_{j=1}^m p_j \frac{\nabla u_j}{u_j} \rangle_H \geq \alpha.$$

But the left-hand side can be rearranged as

$$\begin{aligned} &\sum_{j=1}^m \frac{p_j}{u_j} (\partial_t u_j(t,x) + \operatorname{div}(\vec{v}_j u_j(t,x))) \\ &+ \operatorname{div}(\vec{v} - \sum_{j=1}^m p_j \vec{v}_j) \\ &+ \sum_{j=1}^m p_j \langle \vec{v} - \vec{v}_j, \nabla \log u_j \rangle_H \end{aligned}$$

and the claim follows from (11), (12), (13). A similar argument gives the second claim. ∎

Remark 2.7. The above lemma shows that in order to show that the quantity (14) synthesised from various spacetime functions u_j is monotone, one needs to locate velocity fields $\vec{v}_j$, $\vec{v}$ obeying the pointwise inequalities (11), (12), (13). In practice the u_j will be chosen to obey a transport equation $\partial_t u_j + \operatorname{div}(\vec{v}_j u_j) = 0$, so that (11) is automatically satisfied. As for the other two inequalities, in the linear case $j = 1$ one can obtain (13) automatically by setting $\vec{v} = \vec{v}_j$, leaving only (12) to be verified. In the multilinear case, one would have to set $\vec{v}$ to be a suitable average of the $\vec{v}_j$. This latter strategy seems to only work well in the case when the u_j solve a heat equation $\partial_t u_j = \operatorname{div}(G_j \nabla u_j)$, since in this case the relevant velocity field $\vec{v}_j = -G_j \nabla \log u_j$ is related to $\nabla \log u_j$ by a positive definite matrix and one is more likely to ensure (13) is positive.

We now give the first major application of the above abstract machinery, by reproving K. Ball's geometric Brascamp–Lieb inequality (Example 1.6). We give some further applications in Section 8 and Section 10.

Proposition 2.8 (Geometric Brascamp–Lieb inequality)**.** [Ball], [Bar2] *Let* $(\mathbf{B}, \mathbf{p})$ *be geometric Brascamp–Lieb data. Then*

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = 1. \tag{15}$$

Furthermore, $(\mathbf{B}, \mathbf{p})$ *is both extremisable and gaussian-extremisable.*

Proof We may drop those exponents p_j for which $p_j = 0$ as being irrelevant. By Remark 2.2 we may assume without loss of generality that H_j is a subspace of H and that B_j is the orthogonal projection from H to H_j. By considering (4) with the gaussian input $(\mathrm{id}_{H_j})_{1 \leq j \leq m}$ we observe that $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) \geq 1$. By (6) we see that it suffices to show that $\mathrm{BL}(\mathbf{B}, \mathbf{p}) \leq 1$; note that this will imply that $(\mathrm{id}_{H_j})_{1 \leq j \leq m}$ is a gaussian extremiser and $(\exp(-\pi \|x\|_{H_j}^2))_{1 \leq j \leq m}$ is an extremiser.

By Definition 1.1, it thus suffices to show that for any non-negative measurable functions $f_j : H_j \to \mathbf{R}^+$ we have

$$\int_H \prod_{j=1}^m (f_j \circ B_j)^{p_j} \leq \prod_{j=1}^m (\int_{H_j} f_j)^{p_j}. \tag{16}$$

By Fatou's lemma it suffices to verify this inequality when the f_j are smooth, rapidly decreasing, and strictly positive. Now let $u_j : \mathbf{R}^+ \times H \to \mathbf{R}^+$ be the solution to the heat equation Cauchy problem

$$\begin{aligned} \partial_t u_j(t, x) &= \Delta_H u_j(t, x) \\ u_j(0, x) &= f_j \circ B_j(x) \end{aligned}$$

where $\Delta_H := \mathrm{div}\, \nabla$ is the usual Laplacian on H. More explicitly, we have

$$u_j(t, x) = \frac{1}{(4\pi t)^{\dim(H)/2}} \int_H e^{-\|x-y\|_H^2/4t} f_j(B_j y)\, dy.$$

We can split y into components in H_j and in the orthogonal complement $H_j^\perp$, using Pythagoras' theorem to split $\|x - y\|_H^2$ correspondingly. The contribution from the orthogonal complement can be evaluated explicitly by (3), to obtain

$$u_j(t, x) = \frac{1}{(4\pi t)^{\dim(H_j)/2}} \int_{H_j} e^{-\|B_j x - z\|_{H_j}^2/4t} f_j(z)\, dz. \tag{17}$$

Alternatively, one can verify that u_j also solves the above Cauchy problem.

In order to apply Lemma 2.6, we rewrite the heat equation as a transport equation

$$\partial_t u_j + \mathrm{div}(\vec{v}_j u_j) = 0$$

where $\vec{v}_j := -\nabla \log u_j$; thus (11) is trivially satisfied. Next we set $\alpha := 0$ and

$$\vec{v} := \sum_{j=1}^m p_j \vec{v}_j$$

so that (12) is also trivially satisfied.

Next, we verify the technical condition that $\vec{v}\prod_{j=1}^m u_j^{p_j}$ is rapidly decreasing in space. From (17) and the hypothesis that f_j is smooth and rapidly decreasing, we see that for any t in a compact interval in $(0, \infty)$, $\vec{v}_j(t, x) = -\nabla u_j / u_j(t, x)$ grows at most polynomially in space. On the other hand, u_j is bounded and rapidly decreasing in the directions $\|B_j x\|_{H_j} \to \infty$. Since $\mathbf{B}$ is non-degenerate, $\prod_{j=1}^m u_j^{p_j}$ decays rapidly in all spatial directions. The claim follows.

Now we verify (13). Observe that $u_j(t, x)$ depends only on $B_j x$ and not on x itself, which shows that $\vec{v}_j$ lies in the range H_j of the projection $B_j^* B_j$. In particular we have

$$\nabla \log u_j = -\vec{v}_j = -B_j^* B_j \vec{v}_j.$$

Hence we have

$$\sum_{j=1}^m p_j \langle \vec{v} - \vec{v}_j, \nabla \log u_j \rangle_H = \sum_{j=1}^m p_j \langle B_j^* B_j(\vec{v} - \vec{v}_j), -\vec{v}_j \rangle_H.$$

Also, from (9) we have

$$\sum_{j=1}^m p_j B_j^* B_j(\vec{v} - \vec{v}_j) = \vec{v} - \sum_{j=1}^m p_j B_j^* B_j \vec{v}_j = \vec{v} - \sum_{j=1}^m p_j \vec{v}_j = 0$$

and hence

$$\sum_{j=1}^m p_j \langle \vec{v} - \vec{v}_j, \nabla \log u_j \rangle_H = \sum_{j=1}^m p_j \langle B_j^* B_j(\vec{v} - \vec{v}_j), (\vec{v} - \vec{v}_j) \rangle_H. \tag{18}$$

Since the orthogonal projection $B_j^* B_j$ is positive semi-definite, we obtain (13). We may now invoke Lemma 2.6 and conclude that the quantity

$$Q(t) := \int_H \prod_{j=1}^m u_j^{p_j}(t, x)\ dx$$

is non-decreasing for $0 < t < \infty$. In particular we have

$$\limsup_{t \to 0^+} Q(t) \le \liminf_{t \to \infty} Q(t).$$

From Fatou's lemma we have

$$\int_H \prod_{j=1}^m (f_j \circ B_j)^{p_j} \le \limsup_{t \to 0^+} Q(t)$$

(in fact we have equality, but we will not need this), so it will suffice to show that

$$\liminf_{t \to \infty} Q(t) \le \prod_{j=1}^m \left(\int_{H_j} f_j \right)^{p_j}. \tag{19}$$

(Again, we will have equality, but we do not need this.) Using (17) we can write

$$Q(t) = \frac{1}{(4\pi t)^{\sum_{j=1}^m p_j \dim(H_j)/2}} \int_H \prod_{j=1}^m \left(\int_{H_j} e^{-\|B_j x - z\|_{H_j}^2/4t} f_j(z)\ dz \right)^{p_j}\ dx.$$

By taking traces of the hypothesis (9) we obtain (7). Thus by making the change of variables $x = t^{1/2} w$ we obtain

$$Q(t) = \frac{1}{(4\pi)^{\dim(H)/2}} \int_H \prod_{j=1}^m \left(\int_{H_j} e^{-\|B_j w - t^{-1/2} z\|_{H_j}^2/4} f_j(z)\ dz \right)^{p_j}\ dw.$$

Since the f_j are rapidly decreasing and $\bigcap_{j=1}^m \ker(P_j) = \{0\}$, we may then use dominated convergence to conclude

$$\begin{aligned}\liminf_{t\to\infty} Q(t) &= \frac{1}{(4\pi)^{\dim(H)/2}} \int_H \prod_{j=1}^m \left(\int_{H_j} e^{-\|B_j w\|_{H_j}^2/4} f_j(z)\ dz \right)^{p_j}\ dw \\ &= \frac{1}{(4\pi)^{\dim(H)/2}} \prod_{j=1}^m \left(\int_{H_j} f_j \right)^{p_j} \int_H e^{-\langle \sum_{j=1}^m p_j B_j^* B_j w, w\rangle_H/4}\ dw.\end{aligned}$$

Using (9) and (3), the claim (19) follows. ■

Remark 2.9. It should be clear from the above argument that the method should also extend to the case when m is countably or even uncountably infinite; in the latter case the exponents p_j will be replaced by some positive measure on the index set that j ranges over. We will however not pursue such generalisations here.

3. The gaussian-extremisable case

We can now prove Theorem 1.9 and Theorem 1.15 in the gaussian-extremisable case.

We begin with a notion of equivalence between Brascamp–Lieb data.

Definition 3.1 (Equivalence)**.** Two m-transformations $\mathbf{B} = (H, (H_j)_{1\le j\le m}, (B_j)_{1\le j\le m})$ and $\mathbf{B}' = (H', (H_j')_{1\le j\le m}, (B_j')_{1\le j\le m})$ are said to be *equivalent* if there exist invertible linear transformations $C : H' \to H$ and $C_j : H_j' \to H_j$ such that $B_j' = C_j^{-1} B_j C$ for all j; we refer to C and C_j as the *intertwining transformations.* Note in particular that this forces $\dim(H) = \dim(H')$ and $\dim(H_j) = \dim(H_j')$. We say that two Brascamp–Lieb data $(\mathbf{B}, \mathbf{p})$ and $(\mathbf{B}', \mathbf{p}')$ are equivalent if $\mathbf{B}, \mathbf{B}'$ are equivalent and $\mathbf{p} = \mathbf{p}'$.

Remark 3.2. This is clearly an equivalence relation. Up to equivalence, the only relevant features of a non-degenerate m-transformation are the kernels $\ker(B_j)$ and how they are situated inside H. More precisely, if one fixes the dimensions $n = \dim(H)$ and $n_j = \dim(H_j)$, then the moduli space of non-degenerate m-transformations with these dimensions, quotiented out by equivalence, can be identified with the moduli space of m-tuples of subspaces V_j of $\mathbf{R}^n$, with $\dim(V_j) = n - n_j$ and $\bigcap_{j=1}^m V_j = \{0\}$, quotiented out by the general linear group $GL(\mathbf{R}^n)$

of $\mathbf{R}^n$. The problem of understanding this moduli space is part of the more general question of understanding quiver representations, which is a rich and complex subject (see [DW] for a recent survey).

The Brascamp–Lieb constants of two equivalent Brascamp–Lieb data are closely related:

Lemma 3.3 (Equivalence of Brascamp–Lieb constants)**.** *Suppose that* $(\mathbf{B}, \mathbf{p})$, $(\mathbf{B}', \mathbf{p}')$ *are two equivalent Brascamp–Lieb data, with intertwining transformations* $C : H' \to H$ *and* $C_j : H'_j \to H_j$. *Then we have*

$$\mathrm{BL}(\mathbf{B}', \mathbf{p}') = \frac{\prod_{j=1}^m |\det_{H'_j \to H_j} C_j|^{p_j}}{|\det_{H' \to H} C|} \mathrm{BL}(\mathbf{B}, \mathbf{p}) \tag{20}$$

and

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}', \mathbf{p}') = \frac{\prod_{j=1}^m |\det_{H'_j \to H_j} C_j|^{p_j}}{|\det_{H' \to H} C|} \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}). \tag{21}$$

Here of course $\det_{H' \to H} C$ *denotes the determinant of the transformation* $C : H' \to H$ *with respect to the Lebesgue measures on* H', H, *and similarly for* $\det_{H'_j \to H_j} C_j$. *In particular* $\mathrm{BL}(\mathbf{B}', \mathbf{p}')$ *(or* $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}', \mathbf{p}')$*) is finite if and only if* $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ *(or* $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$*) is finite.*

Furthermore, $(\mathbf{B}, \mathbf{p})$ *is extremisable if and only if* $(\mathbf{B}', \mathbf{p}')$ *is extremisable, and* $(\mathbf{B}, \mathbf{p})$ *is gaussian-extremisable if and only if* $(\mathbf{B}', \mathbf{p}')$ *are gaussian-extremisable.*

Proof To see (20) we simply apply the change of variables $x = C^{-1}y$ on $\mathbf{R}^n$, and replace an input $(f_j)_{1 \le j \le m}$ for $(\mathbf{B}, \mathbf{p})$ with the corresponding input $(f_j \circ C_j)_{1 \le j \le m}$ for $(\mathbf{B}', \mathbf{p}')$. To see the second identity we similarly replace a gaussian input $(A_j)_{1 \le j \le m}$ for $(\mathbf{B}, \mathbf{p})$ with the gaussian input $(C_j^* A_j C_j)_{1 \le j \le m}$ for $(\mathbf{B}', \mathbf{p}')$. The corresponding claims about extremisers are proven similarly. ∎

Remark 3.4. The transformation between A_j and A'_j can be described using the commutative diagram

$$\begin{array}{ccccccc} H & \xrightarrow{B_j} & H_j & \xrightarrow{A_j} & H_j^* & \xrightarrow{B_j^*} & H^* \\ \uparrow C & & \uparrow C_j & & \downarrow C_j^* & & \downarrow C^* \\ H' & \xrightarrow{B'_j} & H'_j & \xrightarrow{A'_j} & (H'_j)^* & \xrightarrow{(B'_j)^*} & (H')^* \end{array}$$

where we have suppressed the identification between a Hilbert space H and its dual H^* to emphasize the self-adjointness of the above diagram, and the horizontal rows do not denote exact sequences. The above lemma shows that the inner product structure of H is not truly relevant for the analysis of Brascamp–Lieb constants, however we retain this structure in order to take advantage of convenient notions such as orthogonal complement or induced Lebesgue measure.

We are now ready to give a satisfactory algebraic characterisation of gaussian extremisable Brascamp–Lieb data.

Definition 3.5 (Ordering of self-adjoint transformations)**.** If $A : H \to H$ and $B : H \to H$ are two self-adjoint linear transformations on a Euclidean space H, we write $A \geq B$ if $A - B$ is positive semi-definite, and $A > B$ if $A - B$ is positive definite. (We recall that $A \geq B$ and $A \neq B$ do not together imply that $A > B$.)

Proposition 3.6 (Characterisation of gaussian-extremisers)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum with* $p_j > 0$ *for all* j*, and let* $\mathbf{A} = (A_j)_{1 \leq j \leq m}$ *be a gaussian input for* $(\mathbf{B}, \mathbf{p})$*. Let* $M : H \to H$ *be the positive semi-definite transformation* $M := \sum_{j=1}^m p_j B_j^* A_j B_j$*. Then the following seven statements are equivalent.*

- *(a)* $\mathbf{A}$ *is a global extremiser to* (4) *(in particular,* $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A})$ *is finite and* $(\mathbf{B}, \mathbf{p})$ *is gaussian-extremisable).*
- *(b)* $\mathbf{A}$ *is a local extremiser to* (4)*.*
- *(c)* M *is invertible, and we have*
$$A_j^{-1} - B_j M^{-1} B_j^* = 0 \text{ for all } 1 \leq j \leq m. \tag{22}$$
- *(d) The scaling condition* (7) *holds,* $\mathbf{B}$ *is non-degenerate, and*
$$A_j^{-1} \geq B_j M^{-1} B_j^* \text{ for all } 1 \leq j \leq m. \tag{23}$$
- *(e) The scaling condition* (7) *holds,* $\mathbf{B}$ *is non-degenerate, and*
$$B_j^* A_j B_j \leq M \text{ for all } 1 \leq j \leq m. \tag{24}$$
- *(f)* $(\mathbf{B}, \mathbf{p})$ *is equivalent to a geometric Brascamp–Lieb datum* $(\mathbf{B}', \mathbf{p}')$ *with intertwining operators* $C := M^{-1/2}$ *and* $C_j := A_j^{-1/2}$*, and*
$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}(\mathbf{B}', \mathbf{p}')\, \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A});$$
$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}', \mathbf{p}')\, \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}).$$
- *(g)* $(\mathbf{B}, \mathbf{p})$ *is equivalent to a geometric Brascamp–Lieb datum* $(\mathbf{B}', \mathbf{p}')$ *with intertwining operators* $C := M^{-1/2}$ *and* $C_j := A_j^{-1/2}$*, and* H_j *equal to the range of* C_j*, and*
$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}).$$

Remark 3.7. One may wish to view (22) as a commutative diagram

$$\begin{array}{ccc} H & \xrightarrow{M^{-1}} & H \\ \uparrow{\scriptstyle B_j^*} & & \downarrow{\scriptstyle B_j} \\ H_j & \xrightarrow{A_j^{-1}} & H_j \end{array}$$

We shall shortly expand this diagram in the proof below as

$$\begin{array}{ccccccc} H & \xrightarrow{M^{-1/2}} & H & \xrightarrow{id} & H & \xrightarrow{M^{-1/2}} & H \\ \uparrow{\scriptstyle B_j^*} & & \uparrow{\scriptstyle T_j^*} & & \downarrow{\scriptstyle T_j} & & \downarrow{\scriptstyle B_j} \\ H_j & \xrightarrow{A_j^{-1/2}} & H_j & \xrightarrow{id} & H_j & \xrightarrow{A_j^{-1/2}} & H_j \end{array}$$

Proof The implication (a) $\implies$ (b) is trivial. Now we verify that (b) $\implies$ (c). From Lemma 4.1 we see that $\bigcap_{j=1}^m \ker(B_j) = \{0\}$. Since the A_j are positive definite, this implies that M is also.

Taking logarithms in (4), we see that $\mathbf{A} > 0$ is a local maximiser for the quantity

$$(\sum_{j=1}^m p_j \log \det_{H_j} A_j) - \log \det_H \sum_{j=1}^m p_j B_j^* A_j B_j.$$

Let us fix a j in $\{1, \dots m\}$, and let $Q_j : H_j \to H_j$ be an arbitrary self-adjoint transformation. By perturbing A_j by a small multiple of Q_j, we conclude that

$$\frac{d}{d\varepsilon} \left(p_j \log \det_{H_j}(A_j + \varepsilon Q_j) - \log \det_H (M + \varepsilon p_j B_j^* Q_j B_j)\right)\Big|_{\varepsilon=0} = 0.$$

We can thus subtract off the term $p_j \log \det_{H_j} A_j - \log \det_H M$, which does not depend on ε, from the above equation and obtain

$$\frac{d}{d\varepsilon} \left(p_j \log \det_{H_j}(\mathrm{id}_{H_j} + \varepsilon A_j^{-1} Q_j) - \log \det_H (\mathrm{id}_H + \varepsilon p_j M^{-1} B_j^* Q_j B_j)\right)\Big|_{\varepsilon=0} = 0.$$

A simple Taylor expansion shows

$$\frac{d}{d\varepsilon} \log \det_H (\mathrm{id}_H + A)|_{\varepsilon=0} = \mathrm{tr}_H(A), \tag{25}$$

hence

$$p_j \,\mathrm{tr}_{H_j}(A_j^{-1} Q_j) - \mathrm{tr}_H(p_j M^{-1} B_j^* Q_j B_j) = 0.$$

We rearrange this using the cyclic properties of the trace as

$$\mathrm{tr}_{H_j}((A_j^{-1} - B_j M^{-1} B_j^*) Q_j) = 0.$$

Since Q_j was an arbitrary self-adjoint transformation, and $A_j^{-1} - B_j M^{-1} B_j^*$ is also self-adjoint, we conclude (22).

Now we verify that (c) $\implies$ (d). The implication of (23) from (22) is trivial. From the invertibility of M we conclude $\bigcap_{j=1}^m \ker(B_j) = \{0\}$, and from the surjectivity of A_j^{-1} and (22) we conclude that B_j is surjective. Hence $\mathbf{B}$ is non-degenerate. Finally we compute

$$\mathrm{tr}_H(\mathrm{id}_H) = \sum_{j=1}^m p_j \,\mathrm{tr}_H(M^{-1} B_j^* A_j B_j) = \sum_{j=1}^m p_j \,\mathrm{tr}_{H_j}(A_j B_j M^{-1} B_j^*) = \sum_{j=1}^m p_j \,\mathrm{tr}_{H_j}(\mathrm{id}_{H_j}) \tag{26}$$

which is (7).

Now we verify (d) $\implies$ (c). From the non-degeneracy of $\mathbf{B}$ we see that M is positive definite, hence invertible. From (23) we have $\mathrm{tr}_{H_j}(A_j B_j M^{-1} B_j^*) \leq \mathrm{tr}_{H_j}(\mathrm{id}_{H_j})$. Using (7) and reversing the argument in (26) we conclude

$$\mathrm{tr}_{H_j}(A_j B_j M^{-1} B_j^*) = \mathrm{tr}_{H_j}(\mathrm{id}_{H_j}),$$

which together with (23) and the positive definiteness of A_j yields (22).

Now we verify that (d) $\iff$ (e). If we introduce the operators $T_j : H \to H_j$ defined by $T_j := A_j^{1/2} B_j M^{-1/2}$ then (23) is equivalent to $T_j T_j^* \leq \mathrm{id}_{H_j}$, while (24) is equivalent to $T_j^* T_j \leq \mathrm{id}_H$. Since $T_j T_j^*$ and $T_j^* T_j$ have the same operator norm, the claim follows.

Now we verify that (c) $\implies$ (f). Let $\mathbf{B}' := (H, (H_j)_{1\leq j\leq m}, (T_j)_{1\leq j\leq m})$ where T_j was defined earlier, thus $(\mathbf{B}', \mathbf{p})$ is equivalent to $(\mathbf{B}, \mathbf{p})$ with intertwining maps $C = M^{-1/2}$ and $C_j := A_j^{-1/2}$. From the hypothesis (22) we have $T_j T_j^* = \mathrm{id}_{H_j}$, and from definition of M we have

$$\sum_{j=1}^m p_j T_j^* T_j = \sum_{j=1}^m M^{-1/2} p_j B_j^* A_j B_j M^{-1/2} = M^{-1/2} M M^{-1/2} = \mathrm{id}_H,$$

and hence $(\mathbf{B}', \mathbf{p})$ is geometric. The remaining claims in (f) then follow from Lemma 3.3.

Finally, the implication (f) $\implies$ (g) follows from Proposition 2.8, and the implication (g) $\implies$ (a) is trivial. ■

Example 3.8. *One can now deduce the sharp Young inequality (Example 1.5) from Proposition 3.6 by setting* $A_j := \frac{1}{p_j(1-p_j)} I_d$ *for* $j = 1, 2, 3$. *In block matrix notation, one has*

$$\sum_{i=1}^3 p_i B_i^* A_i B_i = \begin{pmatrix} \frac{p_2}{(1-p_1)(1-p_3)} I_d & \frac{1}{1-p_3} I_d \\ \frac{1}{1-p_3} I_d & \frac{p_1}{(1-p_2)(1-p_3)} I_d \end{pmatrix}$$

and one easily verifies any of the conditions (c)-(e).

From Proposition 3.6 we see in particular that we have proven Theorem 1.9 and Theorem 1.15 in the case when $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable. However, not every Brascamp–Lieb datum is gaussian-extremisable, even if we assume the Brascamp–Lieb constants to be finite, as the following example shows.

Example 3.9. *Consider the rank one case (Example 1.7) with* $H = \mathbf{R}^2$, $m = 3$, *with any two of* v_1, v_2, v_3 *being linearly independent; this is the case for instance with the one-dimensional version of Young's inequality (Example 1.5). Then one can verify that the quantity* (5) *is finite if and only if* (p_1, p_2, p_3) *lies in the solid triangle with vertices* $(1, 1, 0), (0, 1, 1), (1, 0, 1)$, *but if* (p_1, p_2, p_3) *lies on one of the open edges of this triangle then no gaussian extremiser will exist. However, gaussian extremisers do exist on the three vertices of the triangle and on the interior; this corresponds to the well-known fact that extremisers to Young's inequality* $\|f * g\|_r \leq \|f\|_p \|g\|_q$ *with* $1/r + 1 = 1/p + 1/q$ *and* $1 \leq p, q, r \leq \infty$ *exist when* $1 < p, q, r < \infty$, *or if all of* p, q, r *are equal to* 1 *or* ∞, *but do not exist in the remaining cases. Observe also that this data is simple if and only if* (p_1, p_2, p_3) *lies in the interior of this triangle.*

Fortunately, it turns out that when a Brascamp–Lieb datum is not gaussian-extremisable, then it can be factored into lower-dimensional Brascamp–Lieb data, with the corresponding constants also factoring accordingly; this was first observed by [CLL] in the rank one case. We will in fact have two means of factoring, which roughly speaking correspond to the notions of direct product $G = H \times K$ and semi-direct product $0 \to N \to G \to G/N \to 0$ in group theory; they also correspond to the notions of *decomposability* and *reducibility* respectively in quiver theory. We begin with the analogue of semi-direct product for Brascamp–Lieb data in the next section; the analogue of direct product will be studied in Section 7.

4. Structural theory of Brascamp–Lieb data I : simplicity

In this section we begin our analysis of the structure of general Brascamp–Lieb data, and in particular consider the questions of whether such data can be placed into a "normal form" or decomposed into "indecomposable" components.

Readers familiar with quivers (see e.g. [DW] for an introduction) will recognise an m-transformation as a special example of a quiver representation, and indeed much of the structural theory of m-transformations which we develop here can be viewed as a crude version of the basic representation theory for quivers. It is likely that the deeper theory of such representations is of relevance to this theory, but we do not pursue these connections here.

We first make a trivial remark, that if one of the exponents p_j in an m-exponent $\mathbf{p}$ is zero, then we may omit this exponent, as well as the corresponding components of an m-transformation $\mathbf{B}$, without affecting the Brascamp–Lieb constants $\mathrm{BL}(\mathbf{B},\mathbf{p})$ or $\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p})$ (or the conditions(7), (8)). Also, the omission of such exponents does not affect extremisability or gaussian-extremisability (though it does affect any issues regarding uniqueness of extremisers, albeit in a rather trivial way). We shall use this remark from time to time to reduce to the cases where the exponents p_j are strictly positive.

Next, we establish the (a) $\implies$ (c) and (b) $\implies$ (c) directions of Theorem 1.15:

Lemma 4.1 (Necessary conditions for finiteness)**.** *Let* $(\mathbf{B},\mathbf{p})$ *be Brascamp–Lieb data such that* $\mathrm{BL}(\mathbf{B},\mathbf{p})$ *or* $\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p})$ *is finite. Then we have the scaling condition* (7) *and the dimension inequalities* (8)*. In particular this implies that* $\mathbf{B}$ *is non-degenerate.*

Proof By (6) it suffices to verify the claim assuming that $\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p})$ is finite. Let $\lambda > 0$ be arbitrary. By applying (5) with the gaussian input $(\lambda\,\mathrm{id}_{H_j})_{1\le j\le m}$ we see that

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p}) \geq \lambda^{\frac{1}{2}[\sum_{j=1}^m p_j \dim(H_j) - \dim(H)]} / \det(\sum_{j=1}^m B_j^* B_j)^{1/2}.$$

Since λ is arbitrary, we see that $\mathrm{BL}_{\mathbf{g}}(\mathbf{B},\mathbf{p})$ can only be finite if (7) holds.

Next, let V be any subspace in H, and let $0 < \varepsilon < 1$ be a small parameter. Let $\mathbf{A}^{(\varepsilon)} = (A_j)_{1\le j\le m}$ be the gaussian input $A_j := \varepsilon\,\mathrm{id}_{B_jV} \oplus \mathrm{id}_{(B_jV)^\perp}$. Then $\det_{H_j}(A_j)$ decays like $\varepsilon^{\dim(B_jV)}$ as $\varepsilon \to 0$. Also, we see that $\sum_{j=1}^m B_j^* A_j B_j$ is bounded uniformly in ε, and when restricted to V decays linearly in ε. Thus $\det_H(\sum_{j=1}^m B_j^* A_j B_j)$ decays at least as fast as $\varepsilon^{\dim(V)}$ as $\varepsilon \to 0$. By (4) we conclude that $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}^{(\varepsilon)})$ grows at least as fast as $\varepsilon^{\frac{1}{2}[\sum_{j=1}^m p_j \dim(B_j V) - \dim(V)]}$. Since we are assuming that $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ is finite, we obtain (8) as desired.

Comparing (8) for $V = H$ with (7) we conclude that the B_j must all be surjective, and applying (8) with $V = \bigcap_{j=1}^m \ker(B_j)$ we conclude that $\bigcap_{j=1}^m \ker(B_j)$ must equal $\{0\}$. Thus $\mathbf{B}$ is necessarily non-degenerate. ∎

We remark that by testing (8) on $V := \ker B_j$ we conclude that in order for $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ or $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ to be finite it is necessary that each $p_j \leq 1$.

To proceed further, we need some notation.

Definition 4.2 (Restriction and quotient of m-transformations)**.** Let $\mathbf{B} = (H, (H_j)_{1 \leq j \leq m}, (B_j)_{1 \leq j \leq m})$ be an m-transformation, and let V be a subspace of H. We define the *restriction* $\mathbf{B}_V$ *of* $\mathbf{B}$ *to* V *to be the* m*-transformation*

$$\mathbf{B}_V = (V, (B_j V)_{1 \leq j \leq m}, (B_{j,V})_{1 \leq j \leq m})$$

where $B_{j,V} : V \to B_j V$ *is the restriction of* $B_j : H \to H_j$ *to* V*, and we also define the quotient* $\mathbf{B}_{H/V}$ *of* $\mathbf{B}$ *to be the* m*-transformation*

$$\mathbf{B}_{H/V} = (H/V, (H_j/(B_j V))_{1 \leq j \leq m}, (B_{j,H/V})_{1 \leq j \leq m})$$

where $B_{j,H/V} : H/V \to H_j/(B_j V)$ *is defined by setting* $B_{j,H/V}(x + V)$ *to be the coset* $B_j x + B_j V$*. Equivalently,* $\mathbf{B}_V$ *and* $\mathbf{B}_{H/V}$ *are the unique* m*-transformations for which the diagram*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & V & \longrightarrow & H & \longrightarrow & H/V & \longrightarrow & 0 \\ & & \big\downarrow{\scriptstyle B_{j,V}} & & \big\downarrow{\scriptstyle B_j} & & \big\downarrow{\scriptstyle B_{j,H/V}} & & \\ 0 & \longrightarrow & B_j V & \longrightarrow & H_j & \longrightarrow & H_j/(B_j V) & \longrightarrow & 0 \end{array}$$

commutes.

Remark 4.3. Using the Hilbert structure, one can of course identify H/V with $V^\perp$, the orthogonal complement of V, and we will use this identification when convenient (particularly for computations). However we prefer the notation H/V as it reinforces the idea that the inner product structure is not truly essential to the Brascamp–Lieb analysis. One can schematically represent the above commutative diagram by the "short exact sequence" $0 \to \mathbf{B}_V \to \mathbf{B} \to \mathbf{B}_{H/V} \to 0$.

While one can restrict or quotient arbitrary m-transformations using arbitrary subspaces V of H, it turns out for the purposes of analysing Brascamp–Lieb constants that the only worthwhile subspaces V to restrict to or quotient by are the *critical* subspaces. Recall that a subspace $V \subset H$ is a *critical subspace* for a Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p})$ if

$$\dim(V) = \sum_{j=1}^m p_j \dim(B_j V), \tag{27}$$

and V is non-zero and proper. The trivial subspace $\{0\}$ satisfies (27), and the scaling condition (7) is simply the requirement that H itself satisfy (27). We remind the reader that we call a Brascamp–Lieb datum *simple* if there are no critical subspaces. Also observe that equivalent Brascamp–Lieb data have critical subspaces in 1-1 correspondence; indeed if V is a critical subspace for $(\mathbf{B}, \mathbf{p})$ and $(\mathbf{B}', \mathbf{p})$ is an equivalent Brascamp–Lieb datum with intertwining maps C and C_j, then $C^{-1}V$ is a critical subspace for $(\mathbf{B}', \mathbf{p})$.

Example 4.4. *For Hölder's inequality (Example 1.3), every non-zero proper subspace of $\mathbf{R}^n$ is critical. For the Loomis–Whitney inequality (Example 1.4), any co-ordinate plane (spanned by some subset of $\{e_1, \dots, e_n\}$) is critical. For Young's inequality (Example 1.5) with $d = 1$, there are no critical subspaces when p_1, p_2, p_3 lie strictly between 0 and 1, but if one of p_1, p_2, p_3 equals 1 then one of the lines $\{(x, 0) : x \in \mathbf{R}\}$, $\{(0, y) : y \in \mathbf{R}\}$ or $\{(x, -x) : x \in \mathbf{R}\}$ will be critical.*

Remark 4.5. Critical subspaces play a role in Brascamp–Lieb data analogous to the role normal subgroups play in group theory, as will hopefully become clearer from the other results in this section.

A crucial observation for our analysis is that the necessary conditions (7), (8) factor through critical subspaces:

Lemma 4.6 (Necessary conditions split)**.** *Let $(\mathbf{B}, \mathbf{p})$ be a Brascamp–Lieb datum, and let V_c be a critical subspace. Then $(\mathbf{B}, \mathbf{p})$ obeys the conditions* (7), (8) *if and only if both $(\mathbf{B}_{V_c}, \mathbf{p})$ and $(\mathbf{B}_{H/V_c}, \mathbf{p})$ obey the conditions* (7), (8) *for all subspaces V of V_c, H/V_c respectively.*

Proof The basic idea is to view $\mathbf{B}$ as the "semi-direct product" of $\mathbf{B}_{V_c}$ and $\mathbf{B}_{H/V_c}$, as the schematic diagram $0 \to \mathbf{B}_{V_c} \to \mathbf{B} \to \mathbf{B}_{H/V_c} \to 0$ already suggests. A similar idea will also underlie the proof of Lemma 4.8 below.

First suppose that $(\mathbf{B}, \mathbf{p})$ obeys (7), (8). Then it is clear that $(\mathbf{B}_{V_c}, \mathbf{p})$ obeys (8), simply by restricting the spaces $V \subseteq \mathbf{R}^n$ to be a subspace of V_c. Also, the scaling condition (7) for $(\mathbf{B}_{V_c}, \mathbf{p})$ is precisely (27). To verify that $(\mathbf{B}_{V_c}, \mathbf{p})$ obeys (8), we let V be a subspace of H/V_c. Applying (8) to the subspace $V + V_c$, followed by (27), we observe

$$\begin{aligned}
\dim(V) &= \dim(V + V_c) - \dim(V_c) \\
&\le \sum_{j=1}^{m} p_j \dim(B_j(V + V_c)) - \dim(V_c) \\
&= \sum_{j=1}^{m} p_j \dim(B_jV + B_jV_c) - \sum_{j=1}^{m} p_j \dim(B_jV_c) \\
&= \sum_{j=1}^{m} p_j \dim(B_{j,H/V_c}V)
\end{aligned}$$

as desired. A similar computation with $V = H/V_c$ shows that $(\mathbf{B}_{H/V_c}, \mathbf{p})$ also obeys (7). This proves the "only if" direction of the lemma.

Now suppose conversely that $(\mathbf{B}_{V_c}, \mathbf{p})$ and $(\mathbf{B}_{H/V_c}, \mathbf{p})$ both obey (7), (8). Then by adding together the two instances of (7) we see that $(\mathbf{B}, \mathbf{p})$ also obeys (7). Also, for any $V \subseteq H$, we write $U := V \cap V_c$ and $W := (V + V_c)/V_c$ and compute using both instances of (8),

$$\begin{aligned}
\dim(V) &= \dim(V \cap V_c) + \dim(V + V_c) - \dim(V_c) \\
&= \dim(U) + \dim(W) \\
&\le \sum_{j=1}^{m} p_j \dim(B_{j,V} U) + p_j \dim(B_{j,H/V} W) \\
&= \sum_{j=1}^{m} p_j \dim(B_j U) + p_j (\dim(B_j W + B_j V_c) - \dim(B_j V_c)) \\
&= \sum_{j=1}^{m} p_j [\dim(B_j U) + \dim(B_j V + B_j V_c) - \dim(B_j V_c)]
\end{aligned}$$

since $W + V_c = V + V_c$. Observe that

$$\dim(B_j U) = \dim(B_j (V \cap V_c)) \le \dim(B_j V \cap B_j V_c).$$

Since $\dim(B_j V \cap B_j V_c) + \dim(B_j V + B_j V_c) - \dim(B_j V_c) = \dim(B_j V)$, we see that $(\mathbf{B}, \mathbf{p})$ obeys (8) as desired. ■

Another key fact is the submultiplicativity of Brascamp–Lieb constants through critical subspaces; this was already observed in [CLL] (in the rank one case) and Finner (in the case of orthogonal projections to co-ordinate spaces).

Lemma 4.7 (Submultiplicativity of Brascamp–Lieb constants)**.** *If* $(\mathbf{B}, \mathbf{p})$ *is a Brascamp–Lieb datum and* V *is a critical subspace with respect to this datum, then*

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) \le \mathrm{BL}(\mathbf{B}_V, \mathbf{p}) \, \mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p}). \tag{28}$$

Proof We may assume that $\mathrm{BL}(\mathbf{B}_V, \mathbf{p})$ and $\mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p})$ are finite. Let $\mathbf{f} = (f_j)_{1 \le j \le m}$ be an arbitrary input for $(\mathbf{B}, \mathbf{p})$. We may normalise $\int_{H_j} f_j = 1$. We have to show that

$$\int_H \prod_{j=1}^{m} (f_j \circ B_j)^{p_j} \le \mathrm{BL}(\mathbf{B}_V, \mathbf{p}) \, \mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p}).$$

By the Fubini-Tonelli theorem, the left-hand side can be rewritten as

$$\int_{V^\perp} \Big(\int_V \prod_{j=1}^{m} (f_j \circ B_j)^{p_j} (v + w) \; dv \Big) \; dw.$$

But we can write $f_j \circ B_j(v + w) = f_{j,w} \circ B_{j,V}(v)$, where $f_{j,w} : B_j V \to \mathbf{R}^+$ is the function

$$f_{j,w}(v_j) := f_j(v_j + B_j w) \text{ for all } v_j \in B_j V.$$

Applying (1) we conclude that

$$\int_V \prod_{j=1}^{m} (f_j \circ B_j)^{p_j} (v + w) \; dv \le \mathrm{BL}(\mathbf{B}_V, \mathbf{p}) \prod_{j=1}^{m} \Big(\int_{B_j V} f_{j,w} \Big)^{p_j}$$

and we are reduced to showing that

$$\int_{V^\perp} \prod_{j=1}^m \Big(\int_{B_j V} f_{j,w}\Big)^{p_j}\ dw \leq \mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p}). \tag{29}$$

We then identify $V^\perp$ with H/V and observe that $\int_{B_j V} f_{j,w} = f_{j,H/V} \circ B_{j,H/V}(w)$, where $f_{j,H/V} : H_j/(B_j V) \to \mathbf{R}^+$ is the function

$$f_{j,H/V}(w_j) = \int_{w_j + B_j V} f_j.$$

Applying (1) again, we can bound the left-hand side of (29) by

$$\mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p}) \prod_{j=1}^m \Big(\int_{H_j/(B_j V)} f_{j,H/V}\Big)^{p_j}$$

and the claim follows from the Fubini-Tonelli theorem and the normalisation $\int_{H_j} f_j = 1$. Note the same argument also shows that if $(\mathbf{B}, \mathbf{p})$ is extremisable, then $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_{H/V}, \mathbf{p})$ are also extremisable, because if equality is attained in (1) then all the inequalities above must in fact be equalities. ∎

This result will be used to prove the finiteness theorem ((c) $\implies$ (a) in Theorem 1.15) in the next section. In fact, the Brascamp–Lieb constants and their gaussian versions are not just submultiplicative, but multiplicative:

Lemma 4.8 (Brascamp–Lieb constants split)**.** *If $(\mathbf{B}, \mathbf{p})$ is a Brascamp–Lieb datum and V is a critical subspace with respect to this datum, then*

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}(\mathbf{B}_V, \mathbf{p})\, \mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p})$$

and

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}_V, \mathbf{p})\, \mathrm{BL}_{\mathbf{g}}(\mathbf{B}_{H/V}, \mathbf{p}).$$

Moreover, if $(\mathbf{B}, \mathbf{p})$ is extremisable, then $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_{H/V}, \mathbf{p})$ are also extremisable. Similarly, if $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable, then $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_{H/V}, \mathbf{p})$ are also gaussian-extremisable.

Proof If $(\mathbf{B}, \mathbf{p})$ fails to obey (7) or (8), then Lemma 4.1 and Lemma 4.6 show that all expressions in this lemma are infinite and so the conclusion trivially holds. Thus we may assume that $(\mathbf{B}, \mathbf{p})$ obeys (7), (8). By Lemma 4.6 we see that $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_{H/V}, \mathbf{p})$ also obey (7), (8). In particular all of these data are non-degenerate.

Having proved (28) above, let us first establish the reverse inequality

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) \geq \mathrm{BL}(\mathbf{B}_V, \mathbf{p})\, \mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p}). \tag{30}$$

Let $0 < C_V < \mathrm{BL}(\mathbf{B}_V, \mathbf{p})$ and $0 < C_{H/V} < \mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p})$ be arbitrary constants. Then by definition of Brascamp–Lieb constant and homogeneity, we can find inputs $\mathbf{f}_V = (f_{j,V})_{1\leq j\leq m}$ and $\mathbf{f}_{H/V} = (f_{j,H/V})_{1\leq j\leq m}$ for $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_{H/V}, \mathbf{p})$ respectively, such that

$$\mathrm{BL}(\mathbf{B}_V, \mathbf{p}; \mathbf{f}_V) > C_V; \quad \mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p}; f_{H/V}) > C_{H/V}.$$

We may normalise

$$\int_{B_jV} f_{j,V} = \int_{H_j/(B_jV)} f_{j,H/V} = 1 \text{ for all } j.$$

Let $\lambda > 1$ be a large parameter. We define the input $\mathbf{f}^{(\lambda)} = (f_j)_{1\leq j\leq m}$ for $(\mathbf{B},\mathbf{p})$ by the formula

$$f_j(v+w) := f_{j,V}(v_j)f_{j,H/V}(\lambda w_j) \text{ whenever } v_j \in B_jV; w_j \in (B_jV)^\perp,$$

where we identify $H_j/(B_jV)$ with $(B_jV)^\perp$ in the usual manner. Then by the Fubini-Tonelli theorem and the normalisation of $\mathbf{f}_V, \mathbf{f}_{H/V}$, we observe that

$$\int_{H_j} f_j = \lambda^{\dim(B_jV)-n_j}. \tag{31}$$

Now we investigate the expression $\prod_{j=1}^m f_j \circ B_j(v+w)$ where $v \in V$ and $w \in V^\perp$, where we identify $V^\perp$ with H/V in the usual manner. Observe for any $v \in V$, $w \in V^\perp$ that

$$\prod_{j=1}^m f_j \circ B_j(v + \lambda^{-1}w) = \prod_{j=1}^m f_{j,V}(B_{j,V}v + \lambda^{-1}\pi_{B_jV}(B_jw))f_{j,H/V}(B_{j,H/V}w)$$

and thus by the Fubini-Tonelli theorem and rescaling

$$\begin{aligned}&\lambda^{\dim(V)-\dim(H)}\int_H \prod_{j=1}^m f_j\circ B_j\\ &= \int_{V^\perp}\int_V \prod_{j=1}^m f_{j,V}(B_{j,V}v + \lambda^{-1}\pi_{B_jV}(B_jw))\,dv \prod_{j=1}^m f_{j,H/V}(B_{j,H/V}w)\,dw.\end{aligned}$$

From this, (31), (1) and (27) we conclude

$$\mathrm{BL}(\mathbf{B},\mathbf{p}) \geq \int_{V^\perp}(\int_V \prod_{j=1}^m f_{j,V}(B_{j,V}v + \lambda^{-1}\pi_{B_jV}(B_jw))\,dv) \prod_{j=1}^m f_{j,H/V}(B_{j,H/V}w)\,dw.$$

Since $\mathbf{B}_V$ and $\mathbf{B}_{H/V}$ are non-degenerate, we have

$$\bigcap_{j=1}^m \ker(B_{j,H/V}) = \bigcap_{j=1}^m \ker(B_{j,V}) = \{0\}.$$

Thus in the above integrals, v and w range over a compact set uniformly in $\lambda > 1$. We may thus take limits as $\lambda \to \infty$ (using the smoothness of the $f_{j,V}$) to conclude

$$\mathrm{BL}(\mathbf{B},\mathbf{p}) \geq \int_{V^\perp}(\int_V \prod_{j=1}^m f_{j,V}(B_{j,V}v)dv) \prod_{j=1}^m f_{j,V^\perp}(B_{j,H/V}w)\,dw.$$

By construction of the inputs $\mathbf{f}_V, \mathbf{f}_{H/V}$ we thus have

$$\mathrm{BL}(\mathbf{B},\mathbf{p}) \geq C_V C_{H/V}$$

and upon taking suprema in C_V, $C_{H/V}$ we obtain (30) as desired.

We observe that the corresponding inequality for the gaussian Brascamp–Lieb constants follows by an identical argument, with the $f_{j,V}$, $f_{j,H/V}$ (and hence f_j) now being centred gaussians instead of test functions. Note that one can still justify the limit as $\lambda \to \infty$ using dominated convergence. Thus one has

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) \geq \mathrm{BL}_{\mathbf{g}}(\mathbf{B}_V, \mathbf{p})\, \mathrm{BL}_{\mathbf{g}}(\mathbf{B}_{H/V}, \mathbf{p}). \tag{32}$$

Turning now to the analogue of (28) for the gaussian case, we first observe that the argument given above does not quite work in the gaussian case because if f_j is a centred gaussian then $f_{j,w}$ is merely an uncentred gaussian except when $w = 0$. For this reason it is pertinent to introduce the *uncentred* gaussian Brascamp–Lieb constant $\mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$ defined as the best constant in (1) when the input is restricted to consist of uncentred gaussians. The argument now shows that the analogue of (28) holds for $\mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$. The proof of the inequality corresponding to (30) for $\mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$ follows by the identical argument, with the $f_{j,V}$, $f_{j,H/V}$ (and hence f_j) now being uncentred gaussians instead of test functions. Note that one can still justify the limit as $\lambda \to \infty$ using dominated convergence. Thus the constants $\mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$ are multiplicative; that is

$$\mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{u}}(\mathbf{B}_V, \mathbf{p})\, \mathrm{BL}_{\mathbf{u}}(\mathbf{B}_{H/V}, \mathbf{p}). \tag{33}$$

It therefore suffices to show that the centred and uncentred gaussian constants coincide, that is $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$. Clearly $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) \leq \mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$ so it is enough to show $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) \geq \mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$. While this is an obvious consequence of Lieb's theorem, an elementary argument is available.

Consider an m-tuple of gaussians $f_j(x) := \exp(-\pi\langle A_j(x - \xi_j), (x - \xi_j)\rangle)$ centred at $\xi_j \in \mathbf{R}^{n_j}$. Since B_j is surjective we may take $\xi_j = B_j w_j$, for some $w_j \in \mathbf{R}^n$. Let $\overline{w} = M^{-1}\sum_{j=1}^m p_j B_j^* A_j B_j w_j$, where $M = \sum_{j=1}^m p_j B_j^* A_j B_j$, so that $\sum_{j=1}^m p_j B_j^* A_j B_j (w_j - \overline{w}) = 0$. Now by a simple change of variables,

$$\begin{aligned}
\int \prod f_j(B_j(x))^{p_j}\ dx &= \int \exp\Big\{-\pi \sum_{j=1}^m p_j \langle A_j B_j(x - w_j), B_j(x - w_j)\rangle\Big\}\ dx \\
&= \int \exp\Big\{-\pi \sum_{j=1}^m p_j \langle A_j B_j(x - (w_j - \overline{w})), B_j(x - (w_j - \overline{w}))\rangle\Big\}\ dx \\
&= \exp\Big\{-\pi \sum_{j=1}^m p_j \langle A_j B_j(w_j - \overline{w}), B_j(w_j - \overline{w})\rangle\Big\} \int \exp\Big\{-\pi \sum_{j=1}^m p_j \langle A_j B_j x, B_j x\rangle\Big\}\ dx \\
&\leq \int \exp\Big\{-\pi \sum_{j=1}^m p_j \langle A_j B_j x, B_j x\rangle\Big\}\ dx.
\end{aligned}$$

Hence $\int \prod f_j(B_j(x))^{p_j}\ dx$ is maximised when all of the ξ_j's are zero, and thus $\mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p}) \leq \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$.

The assertion concerning extremisability of $(\mathbf{B}, \mathbf{p})$ is made at the end of the proof of Lemma 4.7 above. Similarly, if $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable, with $\mathbf{A}$ a gaussian input extremising (4), then with $f_j(x) = (\det A_j)^{1/2} \exp(-\pi\langle A_j x, x\rangle)$ we can follow the proof of (28), to see that for all w the uncentred gaussians $(f_{j,w})_{1\leq j\leq m}$

attain $\mathrm{BL}_{\mathbf{u}}(\mathbf{B}_V, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}_V, \mathbf{p})$, and the centred gaussians $(f_{j,H/V})_{1\le j\le m}$ attain $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}_{H/V}, \mathbf{p})$. In particular, taking $w = 0$, both $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_{H/V}, \mathbf{p})$ are gaussian-extremisable. ∎

Remark 4.9. We note that the above proof of $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) \ge \mathrm{BL}_{\mathbf{u}}(\mathbf{B}, \mathbf{p})$ shows that if a collection of uncentred gaussians are extremisers, then they are obtained from the corresponding centred gaussians by translation by $B_j\overline{w}$ for a fixed $\overline{w} \in H$, and that any such translations also furnish extremisers. This will complement Theorem 9.3.

Remark 4.10. Recall that a Brascamp–Lieb datum is *simple* if it has no critical subspaces. Lemma 4.6 and Lemma 4.8 allow us to reduce the problem of finiteness of Brascamp–Lieb constants (or of proving Lieb's theorem) to the case of simple Brascamp–Lieb data. As we shall see later, simple data obeying the necessary conditions (7), (8) are always gaussian-extremisable and hence equivalent to geometric Brascamp–Lieb data, and thus Lieb's theorem reduces entirely to checking the geometric Brascamp–Lieb case, which was already done in Proposition 2.8.

Remark 4.11. The question of whether gaussian extremisers exist in general is a little bit more complicated; the correct characterisation is not simplicity but *semi-simplicity* (the direct sum of simple data). We shall study this issue in depth in Section 7.

5. Sufficient conditions for finiteness, and Lieb's theorem

We are now ready to prove Theorem 1.15 and the general case on Theorem 1.9. We begin with a lemma.

Lemma 5.1. *Let $(\mathbf{B}, \mathbf{p})$ be a Brascamp–Lieb datum such that*

$$\dim(H/V) \ge \sum_{j=1}^m p_j \dim(H_j/(B_jV)) \tag{34}$$

for all subspaces V of H. (In particular, (34) is equivalent to (8) if (7) holds.) Then there exists a real number $c > 0$, such that for every orthonormal basis $e_1, \dots, e_n$ of H there exists a set $I_j \subseteq \{1, \dots, n\}$ for each $1 \le j \le m$ with $|I_j| = \dim(H_j)$ such that

$$\sum_{j=1}^m p_j |I_j \cap \{1, \dots, k\}| \le k \text{ for all } 0 \le k \le n \tag{35}$$

and

$$\| \bigwedge_{i\in I_j} B_j e_i \|_H \ge c \text{ for all } 1 \le j \le m. \tag{36}$$

Here $\bigwedge_{i\in I_j} B_j e_i$ denotes the wedge product of the vectors $B_j e_i$, and the $\|\|_H$ norm is the usual norm on forms induced from the Hilbert space structure. In particular, if (7) holds, we have

$$\sum_{j=1}^m p_j |I_j \cap \{k+1, \dots, n\}| \ge n-k \text{ for all } 0 \le k \le n. \tag{37}$$

If furthermore there are no critical subspaces, then we can enforce strict inequality in (37) *for all* $0 < k < n$.

Proof If (7) holds then $n = \sum_{j=1}^m p_j \dim(H_j)$, and (37) follows from (35) and the fact that $|I_j| = \dim(H_j)$. Thus we shall focus on proving the claim (35). The case $k = 0$ of (35) is trivial, and the $k = n$ case follows from (34), so we restrict attention to $0 < k < n$. From the hypotheses we know that the B_j are surjective.

The space of all orthonormal bases is compact, and the number of possible I_j is finite. Thus by continuity and compactness we may replace the conclusion (36) by the weaker statement

$$\bigwedge_{i \in I_j} B_j e_i \neq 0 \text{ for all } 1 \leq j \leq m.$$

In other words, we require the vectors $(B_j e_i)_{i \in I_j}$ to be linearly independent on H_j for each j.

We shall select the I_j by a backwards greedy algorithm. Namely, we set I_j equal to those indices i for which $B_j e_i$ is not in the linear span of $\{B_j e_{i'} : i < i' \leq n\}$ (thus for instance n will lie in I_j as long as $B_j e_n \neq 0$). Since the B_j are surjective, we see that $|I_j| = \dim(H_j)$. To prove (35), we apply the hypothesis (34) with V equal to the span of $\{e_{k+1}, \dots, e_n\}$, to obtain

$$\sum_j p_j \dim(H_j / B_j V) \leq k.$$

But by construction of I_j we see that $\dim(B_j V) = |I_j \cap \{k+1, \dots, n\}|$ and hence $\dim(H_j / B_j V) = |I_j \cap \{1, \dots, k\}|$. The claim (35) follows.

If (7) holds and there are no critical spaces, then one always has strict inequality in (8) or (34), hence in (35) or (37) when $0 < k < n$. ■

We can now prove the sufficiency of (7) and (8) in Theorem 1.15 for the finiteness of the gaussian Brascamp–Lieb constant.

Proposition 5.2. *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum such that* (7) *and* (8) *hold. Then* $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ *is finite. Furthermore, if* $(\mathbf{B}, \mathbf{p})$ *is simple (i.e., there is no critical subspace), then* $(\mathbf{B}, \mathbf{p})$ *is gaussian-extremisable.*

We remark that this proposition was established in the rank one case by [CLL].

Proof We may discard those j for which $p_j = 0$ or for which $H_j = \{0\}$, as these factors clearly give no contribution to the Brascamp–Lieb constants (or to (7) or (8)), nor does they affect whether $(\mathbf{B}, \mathbf{p})$ is simple or not. In particular $\mathbf{B}$ will still be non-degenerate after doing this.

Fix a gaussian input $\mathbf{A} = (A_j)_{1 \leq j \leq m}$, and let $M := \sum_j p_j B_j^* A_j B_j$. This transformation is self-adjoint; since $\mathbf{B}$ is non-degenerate and $p_j > 0$, we also see that it is positive definite. Thus by choosing an appropriate orthonormal basis $\{e_1, \dots, e_n\} \subset H$ we may assume that $M = \mathrm{diag}(\lambda_1, \dots, \lambda_n)$ for some $\lambda_1 \geq \dots \geq \lambda_n > 0$.

Applying Lemma 5.1, we can find $I_j \subseteq \{1, \dots, n\}$ for each $1 \leq j \leq m$ of cardinality $|I_j| = \dim(H_j)$ obeying (37) and (36). For each $i \in I_j$, we have

$$\langle A_j B_j e_i, B_j e_i \rangle_{H_j} = \langle e_i, B_j^* A_j B_j e_i \rangle_H \leq \frac{1}{p_j} \langle e_i, M e_i \rangle_H = \lambda_i / p_j.$$

On the other hand, from (36) we see that $(B_j e_i)_{i \in I_j}$ is a basis of H_j with a lower bound on the degeneracy. We thus conclude that

$$\det(A_j) \leq C \prod_{i \in I_j} \lambda_i$$

for some constant $C > 0$ depending on the Brascamp–Lieb data. Thus

$$\prod_{j=1}^m (\det A_j)^{p_j} \leq C \prod_{i=1}^n \lambda_i^{\sum_{j=1}^m p_j |I_j \cap \{i\}|}.$$

We can telescope the right-hand side (using (7)) and obtain

$$\prod_{j=1}^m (\det A_j)^{p_j} \leq C \lambda_1^n \prod_{0 \leq k \leq n-1} (\lambda_{k+1}/\lambda_k)^{\sum_{j=1}^m p_j |I_j \cap \{k+1, \dots, n\}|}$$

where we adopt the convention $\lambda_0 = \lambda_1$. Applying (37) we conclude

$$\prod_{j=1}^m (\det A_j)^{p_j} \leq C \lambda_1^n \prod_{0 \leq k \leq n-1} (\lambda_{k+1}/\lambda_k)^{n-k}$$

which by reversing the telescoping becomes

$$\prod_{j=1}^m (\det A_j)^{p_j} \leq C \lambda_1 \dots \lambda_k = C \det(M).$$

Comparing this with (4) we conclude that $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ is finite.

Now suppose that $(\mathbf{B}, \mathbf{p})$ is simple. Then we have strict inequality in (37). We may thus refine the above analysis and conclude that

$$\prod_{j=1}^m (\det A_j)^{p_j} \leq C \det(M) \prod_{1 \leq k \leq n-1} (\lambda_{k+1}/\lambda_k)^c$$

for some $c > 0$ depending on the Brascamp–Lieb data. This shows that the expression in (4) goes to zero whenever λ_n/λ_1 goes to zero. Thus to evaluate the supremum it suffices to do so in the region $\lambda_1 \leq C\lambda_n$. Also using the scaling hypothesis (7) we may normalise $\lambda_n = 1$. This means that M is now bounded above and below, which by surjectivity of B_j implies that A_j is also bounded. We may now also assume that A_j is bounded from below since otherwise the expression (4) in the supremum in (5) will be small. We have thus localised each the A_j to a compact set, and hence by continuity we see that an extremiser exists. Thus $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable as desired. ∎

We can now prove Theorem 1.9 and Theorem 1.15 simultaneously and quickly.

Proof [of Theorem 1.9 and Theorem 1.15] As in [CLL], we induct on the dimension $\dim(H)$. When $\dim(H) = 0$ the claim is trivial. Now suppose inductively that $\dim(H) > 0$ and the claim has already been proven for smaller values of $\dim(H)$. In light of Lemma 4.1 we may assume that (7) and (8) hold. From Proposition 5.2 we know that $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ is finite; our task is to show that $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ is equal to $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$.

We divide into two cases. First suppose that $(\mathbf{B}, \mathbf{p})$ is simple. Then by Proposition 5.2 the datum $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable, and the claim then follows from Proposition 3.6. Now consider the case when $(\mathbf{B}, \mathbf{p})$ is not simple, i.e., there is a proper critical subspace V. Then we can split the Brascamp–Lieb data $(\mathbf{B}, \mathbf{p})$ into $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_{H/V}, \mathbf{p})$. By Lemma 4.6 and the induction hypothesis we see that

$$\mathrm{BL}(\mathbf{B}_V, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}_V, \mathbf{p}) < \infty$$

and

$$\mathrm{BL}(\mathbf{B}_{H/V}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}_{H/V}, \mathbf{p}) < \infty$$

Applying Lemma 4.8 we conclude

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) < \infty$$

thus closing the induction. ■

Let us now fix the m-transformation $\mathbf{B}$, and let the m-exponent $\mathbf{p} = (p_1, \dots, p_m)$ vary. Let us define the following subsets of $\mathbf{R}^m$:

$$S(\mathbf{B}) := \{\mathbf{p} \in \mathbf{R}^m : p_j \geq 0 \text{ for all } j; \quad \sum_{j=1}^{m} p_j \dim(H_j) = \dim(H)\}$$

$$\Pi(\mathbf{B}) := \{\mathbf{p} \in \mathbf{R}^m : \sum_{j=1}^{m} p_j \dim(B_j V) \geq \dim(V) \text{ for all } \{0\} \subsetneq V \subseteq H\}$$

$$\Pi^\circ(\mathbf{B}) := \{\mathbf{p} \in \mathbf{R}^m : \sum_{j=1}^{m} p_j \dim(B_j V) > \dim(V) \text{ for all } \{0\} \subsetneq V \subseteq H\}$$

Observe that even though there are an infinite number of vector spaces V, there are only a finite number of possible values for the dimensions $\dim(V), \dim(B_j V)$. Thus $\Pi(\mathbf{B})$ is a closed convex cone with only finitely many faces, and $\Pi^\circ(\mathbf{B})$ is the m-dimensional interior of that cone.

Theorem 1.15 thus asserts that the Brascamp–Lieb constant $\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ is finite if and only if $\mathbf{p}$ lies in $S(\mathbf{B}) \cap \Pi(\mathbf{B})$, and furthermore if $\mathbf{p}$ lies in $S(\mathbf{B}) \cap \Pi(\mathbf{B})^\circ$ then $(\mathbf{B}, \mathbf{p})$ is also gaussian-extremisable. This of course implies that equality in (1) can also be attained.

Remark 5.3. It is perfectly possible for $S(\mathbf{B}) \cap \Pi(\mathbf{B})$ or $S(\mathbf{B}) \cap \Pi^\circ(\mathbf{B})$ to be empty. For instance for the Loomis–Whitney inequality (Example 1.4), $S(\mathbf{B}) \cap \Pi(\mathbf{B})$ consists of a single point $(\frac{1}{n-1}, \dots, \frac{1}{n-1})$, and $S(\mathbf{B}) \cap \Pi(\mathbf{B})^\circ$ is empty. If $\mathbf{B}$ is degenerate, then $S(\mathbf{B}) \cap \Pi(\mathbf{B})$ is empty.

An alternative proof that the hypotheses (7) and (8) characterise the Brascamp–Lieb constant $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ is given in the companion paper [BCCT]. That proof does not involve gaussians, extremisers, or monotonicity formulae, but instead uses multilinear interpolation, Hölder's inequality, and an induction argument on dimension based on factorisation through critical subspaces. Such an induction and factorisation argument had previously appeared in [CLL] (for the rank one case) and [F] (for the case of orthogonal projections to co-ordinate spaces).

We now connect the above results to the work of Barthe [Bar2], who considered the rank-one case (Example 1.7). This connection was also noted in [CLL], though our treatment here differs slightly from that in [CLL]. We are assuming $\mathbf{B}$ to be non-degenerate, thus the v_j are non-zero and span H. The condition (7) in this case simplifies to $\sum_{j=1}^m p_j = n$, while (8) in this case simplifies to

$$\dim(V) \leq \sum_{1 \leq j \leq m : v_j \notin V^\perp} p_j.$$

Subtracting this from the scaling identity $\sum_{j=1}^m p_j = n$, we see that (8) has now become

$$\sum_{1 \leq j \leq m : v_j \in V^\perp} p_j \leq \dim(V^\perp).$$

Since we can set $V^\perp$ to be any subspace, and in particular to be the span of any set of vectors v_j, we thus easily see that the constraints (8) are then (assuming (7)) equivalent to the assertion

$$\sum_{j \in I} p_j \leq d_I \text{ for all } I \subseteq \{1, \dots, n\}, \tag{38}$$

where d_I is the integer $d_I := \dim \operatorname{span}((v_j)_{j \in I})$. In particular we have $0 \leq p_j \leq 1$ for all j, (which, as we have noted above, is in general a necessary condition). Thus in the rank one case, we have characterised the polytope where the Brascamp–Lieb constant is finite as $S(\mathbf{B}) \cap \Pi(\mathbf{B})$

$$= \{(p_1, \dots, p_m) \in \mathbf{R}_+^m : \sum_{1 \leq j \leq m} p_j = n; \sum_{j \in I} p_j \leq d_I \text{ for all } I \subseteq \{1, \dots, m\}\}.$$

Let us now characterise the extreme points of this polytope.

Lemma 5.4. [CLL] *Suppose $H_j = \mathbf{R}$ for all j. Let $(p_1, \dots, p_m)$ be an extreme point of $S \cap \Pi$. Then all of the p_j are equal to 0 or 1.*

Proof We induct on $n + m$. When $n + m = 0$ the claim is trivial, so suppose that $n + m > 0$ and the claim has been proven for all smaller values of $n + m$. If one of the p_j is already equal to zero, say $p_m = 0$, then we can remove that index m (and the associated vector v_m) from $\mathbf{B}$ to form an $(m-1)$-transformation $\tilde{\mathbf{B}}$, and observe that $(p_1, \dots, p_{m-1})$ is an extreme point of $S(\tilde{\mathbf{B}}) \cap \Pi(\tilde{\mathbf{B}})$. Thus the claim follows from the induction hypothesis. Now suppose that one of the p_j is equal to one, say $p_m = 1$. Let $H' := H/\mathbf{R}v_m$ be the quotient of H by v_m, and let $v_1', \dots, v_{m-1}'$ be the image of $v_1, \dots, v_m$ under this quotient map. We let $\mathbf{B}'$ be the associated $(m-1)$-transformation. If we then let $d_I' := \dim \operatorname{span}((v_j')_{j \in I})$ for all $I \subseteq \{1, \dots, m-1\}$, then we have

$$S(\mathbf{B}') \cap \Pi(\mathbf{B}') = \{(q_1, \dots, q_{m-1}) \in \mathbf{R}_+^{m-1} : \sum_{1 \leq j \leq m-1} q_j = n-1;$$
$$\sum_{j \in I} q_j \leq d'_I \text{ for all } I \subseteq \{1, \dots, m\}\}.$$

Now

$$\sum_{j \in I} p_j = \sum_{j \in I \cup \{m\}} p_j - 1 \leq d_{I \cup \{m\}} - 1 = d'_I.$$

Thus $(p_1, \dots, p_{m-1}) \in S(\mathbf{B}') \cap \Pi(\mathbf{B}')$. Conversely, if $(q_1, \dots, q_{m-1}) \in S(\mathbf{B}') \cap \Pi(\mathbf{B}')$ then reversing the above argument shows that $(q_1, \dots, q_{m-1}, 1) \in S(\mathbf{B}) \cap \Pi(\mathbf{B})$. Since $(p_1, \dots, p_{m-1}, 1)$ was an extreme point of $S(\mathbf{B}) \cap \Pi(\mathbf{B})$, this implies that $(p_1, \dots, p_{m-1})$ is an extreme point of $S(\mathbf{B}') \cap \Pi(\mathbf{B}')$. Thus we can again apply the induction hypothesis to close the argument.

Finally, suppose that none of the p_j are equal to 0 or 1. Call a proper subset $\emptyset \subsetneq I \subseteq \{1, \dots, m\}$ of $\{1, \dots, m\}$ *critical* if $\sum_{j \in I} p_j = d_I$. Since d_I has to be an integer, we see that no singleton sets are critical. On the other hand, the entire set $\{1, \dots, m\}$ is critical. Thus if we let I_{min} be a non-empty critical set of minimal size, then I_{min} has at least two elements. Now suppose that I is another critical set which intersects I_{min}. Then

$$\begin{aligned}
d_{I_{min}} + d_I &= \sum_{j \in I_{min}} p_j + \sum_{j \in I} p_j \\
&= \sum_{j \in I_{min} \cap I} p_j + \sum_{j \in I_{min} \cup I} p_j \\
&\leq d_{I_{min} \cap I} + d_{I_{min} \cup I} \\
&\leq \dim[\operatorname{span}((v_j)_{j \in I_{min}}) \cap \operatorname{span}((v_j)_{j \in I})] \\
&\quad + \dim[\operatorname{span}((v_j)_{j \in I_{min}}) + \operatorname{span}((v_j)_{j \in I})] \\
&= d_{I_{min}} + d_I
\end{aligned}$$

and hence all the above inequalities must in fact be equality. In particular this implies that the non-empty set $I_{min} \cap I$ is critical, which by minimality of I_{min} forces $I_{min} \subseteq I$. Thus all the critical sets either contain I_{min} or are disjoint from it. Now recall that I_{min} has at least two elements, and that p_j lies strictly between 0 and 1 for all j in I_{min}. Thus one can reduce one of the p_j, $j \in I_{min}$ by an epsilon and increase another p_j, $j \in I_{min}$ by the same epsilon, and stay in $S(\mathbf{B}) \cap \Pi(\mathbf{B})$, or conversely. This shows that $(p_1, \dots, p_m)$ is not an extreme point of $S(\mathbf{B}) \cap \Pi(\mathbf{B})$, a contradiction. ■

Note that if $(p_1, \dots, p_m) \in S(\mathbf{B}) \cap \Pi(\mathbf{B})$ consists entirely of 0's and 1's, then the set $\{v_j : p_j = 1\}$ is a basis for H (because it has cardinality n by (7), and any subset of this set of cardinality k must span a space of dimension at least k by (38)). Conversely, if $I \subseteq \{1, \dots, m\}$ is such that $\{v_j : j \in I\}$ is a basis for H, then the m-tuple $(1_{j \in I})_{1 \leq j \leq m}$ is easily verified to lie in $S(\mathbf{B}) \cap \Pi(\mathbf{B})$. If one takes the convex hull of all these points, then any one of these points $(1_{j \in I})_{1 \leq j \leq m}$ will be vertices of these convex hull, as can be seen by maximising the linear functional $\sum_{j \in I} p_j$ on this convex hull. Combining all these facts, we have reproved the following theorem of Barthe.

Theorem 5.5. [Bar2, Section 2] *Let $v_1, \dots, v_m$ be non-zero vectors which span a Euclidean space $\mathbf{R}^n$, and let $B_j : \mathbf{R}^n \to \mathbf{R}$ be the maps $B_j x := \langle x, v_j \rangle$. Let $P \subset \mathbf{R}^m_+$ be the set of all exponents $\mathbf{p}$ for which the Brascamp–Lieb constant $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ is finite. Then P is a convex polytope, whose vertices are precisely those points $(1_{j \in I})_{1 \leq j \leq m})$ where $(v_j)_{j \in I}$ is a basis for $\mathbf{R}^n$. If $\mathbf{p}$ lies in the $(m-1)$-dimensional interior of this polytope, $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable.*

Remark 5.6. Barthe's proof of this theorem relies upon an analysis of the constant (5). The sup in (5) is now over $(0, \infty)^m$ and the Cauchy-Binet theorem gives an explicit formula for the denominator in (4). This facilitates a direct analysis of (5) in the rank-one case.

6. EXTREMISERS

In this section we analyse extremisers to the Brascamp–Lieb inequality, and connect these extremisers to the heat flow approach used previously.

We now recall an important observation of K. Ball (which can be found for instance in [Bar]) which can be used to motivate the monotonicity formula approach. If $\mathbf{f} = (f_j)_{1 \leq j \leq m}$ and $\mathbf{f}' = (f'_j)_{1 \leq j \leq m}$ are inputs for a Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p})$, we define their *convolution* $\mathbf{f} * \mathbf{f}'$ to be the m-tuple $\mathbf{f} * \mathbf{f}' := (f_j * f'_j)_{1 \leq j \leq m}$, where of course $f_j * g_j(x) := \int_{H_j} f_j(x-y) g_j(y)\, dy$. Observe from the Fubini-Tonelli theorem that $\mathbf{f} * \mathbf{f}'$ will also be an input for $(\mathbf{B}, \mathbf{p})$.

Lemma 6.1 (Convolution inequality)**.** [Bar] *Let $(\mathbf{B}, \mathbf{p})$ be a Brascamp–Lieb datum with $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ finite, and let $\mathbf{f}$, $\mathbf{f}'$ be inputs for $(\mathbf{B}, \mathbf{p})$. Then*

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f} * \mathbf{f}') \geq \frac{\mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f})\, \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}')}{\mathrm{BL}(\mathbf{B}, \mathbf{p})}.$$

Proof Write $\mathbf{f} = (f_j)_{1 \leq j \leq m}$ and $\mathbf{f}' = (f'_j)_{1 \leq j \leq m}$ After normalisation we can then assume $\int_{H_j} f_j = \int_{H_j} f'_j = 1$. Thus

$$\int_H \prod_{j=1}^m (f_j \circ B_j)^{p_j} = \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}); \quad \int_H \prod_{j=1}^m (f'_j \circ B_j)^{p_j} = \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}').$$

Convolving these together we obtain

$$\int_H \prod_{j=1}^m (f_j \circ B_j)^{p_j} * \prod_{j=1}^m (f'_j \circ B_j)^{p_j} = \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f})\, \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}').$$

The left-hand side can be rewritten using the Fubini-Tonelli theorem as

$$\int_H (\int_H \prod_{j=1}^m (g_j^x \circ B_j)^{p_j})\ dx$$

where for each $x \in H$, $g_j^x : H_j \to \mathbf{R}^+$ is the function $g_j^x(y) := f_j(B_j x - y) f_j'(y)$. Applying (1) to the inner integral we conclude that

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f})\, \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}') \leq \mathrm{BL}(\mathbf{B}, \mathbf{p}) \int_H \prod_{j=1}^m \Big(\int_{H_j} g_j^x\Big)^{p_j}\, dx.$$

But clearly $\int_{H_j} g_j^x = f_j * f_j'(B_j x)$, hence

$$\mathrm{BL}(\mathbf{B}, \mathbf{p})^2 \leq \mathrm{BL}(\mathbf{B}, \mathbf{p}) \int_H \prod_{j=1}^m ((f_j * f_j') \circ B_j)^{p_j}.$$

On the other hand, from (1) we have

$$\int_H \prod_{j=1}^m ((f_j * f_j') \circ B_j)^{p_j} = \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f} * \mathbf{f}') \prod_{j=1}^m \Big(\int_H f_j * f_j'\Big)^{p_j} = \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f} * \mathbf{f}')$$

since by the normalisation we have $\int_H f_j * f_j' = (\int_H f_j)(\int_H f_j') = 1$. Combining these inequalities we obtain the claim. ∎

Remark 6.2. Lemma 4.8 can be viewed as a degenerate version of this inequality, in which $\mathbf{f}'$ is Lebesgue measure on the subspace V.

By applying this lemma we can conclude the following closure properties of extremisers.

Lemma 6.3 (Closure properties of extremisers)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum with* $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ *finite.*

- *(Scale invariance) If* $\mathbf{f} = (f_j)_{1 \leq j \leq m}$ *is an extremising input, then so is* $(f_j(\lambda \cdot))_{1 \leq j \leq m}$ *for any non-zero real number* λ*.*
- *(Homogeneity) If* $\mathbf{f} = (f_j)_{1 \leq j \leq m}$ *is an extremising input, then so is* $(c_j f_j)_{1 \leq j \leq m}$ *for any non-zero real numbers* $c_1, \dots, c_m$*.*
- *(Translation invariance) If* $\mathbf{f} = (f_j)_{1 \leq j \leq m}$ *is an extremising input, then so is* $(f_j(\cdot - B_j x_0))_{1 \leq j \leq m}$ *for any* $x_0 \in H$*.*
- *(Closure in* L^1*) If* $\mathbf{f}^{(n)} = (f_j^{(n)})_{1 \leq j \leq m}$ *is a sequence of extremising inputs, which converges in the product* L^1 *sense to another input* $\mathbf{f} = (f_j)_{1 \leq j \leq m}$*, so* $\lim_{n \to \infty} \|f_j^{(n)} - f_j\|_{L^1(H_j)} = 0$ *for all* $1 \leq j \leq m$*, then* $\mathbf{f}$ *is also an extremising input.*
- *(Closure under convolution) If* $\mathbf{f}$ *and* $\mathbf{f}'$ *are extremisers, then so is* $\mathbf{f} * \mathbf{f}'$*.*
- *(Closure under multiplication) If* $\mathbf{f}$ *and* $\mathbf{f}'$ *are extremisers, then the input* $(f_j(\cdot - B_j x_0) f_j'(\cdot))_{1 \leq j \leq m}$ *is an extremiser for almost every* $x_0 \in H$ *for which* $\int_H f_j(x - B_j x_0) f_j'(x)\, dx > 0$ *for all* j*. If furthermore* $\mathbf{f}$ *and* $\mathbf{f}'$ *are bounded, then the "almost" in "almost every" can be removed.*

Proof From Lemma 4.1 we see that (7) holds, which guarantees the scale invariance property. The closure in L^1 follows from standard arguments since one can use (1) and the hypothesis that $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ is finite to show that $\int_H \prod_{j=1}^m (f_j^{(n)} \circ B_j)^{p_j}$ converges to $\int_H \prod_{j=1}^m (f_j \circ B_j)^{p_j}$. The homogeneity and translation invariance can be verified by direct computation. The closure under convolution follows from

Lemma 6.1. Now consider the closure under multiplication. Let us write $\tilde{f}_j(x) := f_j(-x)$, and let us repeat the proof of Lemma 6.1 with $\mathbf{f} = (f_j)_{1 \leq j \leq m}$ replaced by $\tilde{f} = (\tilde{f}_j)_{1 \leq j \leq m}$. Note from scale invariance that $\tilde{f}$ is still an extremiser. We then argue as before to obtain

$$\mathrm{BL}(\mathbf{B}, \mathbf{p})^2 = \int_H (\int_H \prod_{j=1}^m (g_j^x \circ B_j)^{p_j})\ dx \leq \mathrm{BL}(\mathbf{B}, \mathbf{p}) \int_H \prod_{j=1}^m (\int_{H_j} g_j^x)^{p_j}\ dx$$
$$= \mathrm{BL}(\mathbf{B}, \mathbf{p})\, \mathrm{BL}(\mathbf{B}, \mathbf{p}; \tilde{\mathbf{f}} * \mathbf{f}'),$$

where $g_j^x(y) := f_j(y - B_j x) f_j'(y)$. Since $\mathrm{BL}(\mathbf{B}, \mathbf{p}; \tilde{\mathbf{f}} * \mathbf{f}') \leq \mathrm{BL}(\mathbf{B}, \mathbf{p})$, the above inequality must in fact be equality, and thus

$$\int_H \prod_{j=1}^m (g_j^x \circ B_j)^{p_j}\ dx = \mathrm{BL}(\mathbf{B}, \mathbf{p}) \prod_{j=1}^m (\int_{H_j} g_j^x)^{p_j}$$

for almost every x. Also from the Fubini-Tonelli theorem we know that $\int_{H_j} g_j^x$ is finite for almost every x. This proves the first part of the closure under multiplication. Also, since f_j and f_j' are integrable and bounded, the convolution g_j^x lies in $L^1(H_j)$ continuously in x. Since we are assuming $\mathrm{BL}(\mathbf{B}, \mathbf{p})$ to be finite, this means that the left-hand side of the above expression also depends continuously on x. Thus in fact we have equality for every x, not just almost every x, and the second part of the closure under multiplication follows. ■

If we specialise the above lemma to the case of centred gaussian extremisers and use Theorem 1.9, observing that the convolution or product of two gaussians is again a gaussian, we conclude

Corollary 6.4 (Closure properties of gaussian extremisers)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum with* $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ *finite.*

- *(Scale invariance) If* $\mathbf{A}$ *is a gaussian extremiser, then so is* $\lambda \mathbf{A}$ *for any* $\lambda > 0$.
- *(Topological closure) If* $\mathbf{A}^{(n)}$ *is a sequence of extremising gaussian inputs which converge to a gaussian input* $\mathbf{A}$*, then* $\mathbf{A}$ *is also an extremiser.*
- *(Closure under harmonic addition) If* $\mathbf{A} = (A_j)_{1 \leq j \leq m}$ *and* $\mathbf{A}' = (A_j')_{1 \leq j \leq m}$ *are extremisers, then so is* $((A_j^{-1} + (A_j')^{-1})^{-1})_{1 \leq j \leq m}$.
- *(Closure under addition) If* $\mathbf{A}$ *and* $\mathbf{A}'$ *are gaussian extremisers, then so is* $\mathbf{A} + \mathbf{A}'$.

Note that these properties can also be deduced from Proposition 3.6 (or Proposition 9.1 below), although the closure under harmonic addition is not particularly obvious from these propositions. By applying all of these closure properties we can now conclude

Proposition 6.5. *Let* $(\mathbf{B}, \mathbf{p})$ *be an extremisable Brascamp–Lieb datum* $(\mathbf{B}, \mathbf{p})$. *Then* $\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ *and* $(\mathbf{B}, \mathbf{p})$ *is gaussian-extremisable.*

Carlen, Lieb and Loss have proved this in the rank-one case; see Theorem 5.4 of [CLL].

Proof Intuitively, the idea (which was first discussed in Barthe [Bar2]) is to start with an extremising input $\mathbf{f}$ and repeatedly convolve it using the closure under convolution property of extremisers, rescale these convolutions using the scale invariance, and then take limits using the central limit theorem and the closure under L^1 to obtain a gaussian extremiser. However there is a technical difficulty because the central limit theorem requires some moment conditions on the input $\mathbf{f}$ which are not obviously available, and so we will instead proceed in stages, starting with an arbitrary extremiser $\mathbf{f}$ and successively replacing $\mathbf{f}$ with increasingly more regular and well-behaved extremisers, until we end up with a centred gaussian extremiser.

Let $\mathbf{f} = (f_j)_{1\leq j\leq m}$ be an extremising input for a Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p})$. To begin with, all we know about the f_j are that they are non-negative, integrable, and that $\int_{H_j} f_j > 0$. However we can use the following trick to create a more regular extremiser. Observe that the reflection $\tilde{\mathbf{f}}(x) := (f_j(-x))_{1\leq j\leq m}$ is also an extremiser (by scale invariance). Using closure under convolution we conclude that $\mathbf{f} * \tilde{\mathbf{f}}$ is also an extremiser, and is also symmetric around the origin, and is strictly positive at zero. In fact it must be strictly positive on a small ball surrounding zero; this is because for each j there must exist a set E_j of positive finite Lebesgue measure such that $f_j > c1_{E_j}$ for some constant $c > 0$, and the convolution $c1_{E_j} * c1_{E_j}$ is continuous and positive at zero. Thus, replacing $\mathbf{f}$ with $\mathbf{f} * \tilde{\mathbf{f}}$ if necessary, we can (and shall) assume that $\mathbf{f}$ is symmetric, and strictly positive near the origin.

Now let $1 \leq \lambda \leq 2$ and $x_0 \in H$ be parameters with $\|x_0\|_H \leq \varepsilon$ for some small ε, and consider the m-tuple $\mathbf{f}_{\lambda,x_0} = (f_j(x)f_j(\lambda x - B_j x_0))_{1\leq j\leq m}$. From the Fubini-Tonelli theorem one verifies that

$$\int_{H_j} f_j(x)f_j(\lambda x - B_j x_0)\ dx < \infty \tag{39}$$

for almost every λ, x_0, and by previous hypotheses we know that $f_j(x)f_j(\lambda x - B_j x_0)$ is strictly positive near the origin if ε is sufficiently small. In particular $\mathbf{f}_{\lambda,x_0}$ is an input for almost every λ, x_0 with $\|x_0\|_H$ small. From closure under multiplication we thus see that $\mathbf{f}_{\lambda,x_0}$ is an extremiser for almost every λ, x_0. We now claim that $\mathbf{f}_{\lambda,x_0}$ obeys the moment condition

$$\int_{H_j} f_j(x)f_j(\lambda x - B_j x_0)(1 + \|x\|_{H_j})\ dx < \infty$$

for each j. To see this, we fix j. By (39) we only need to show that

$$\int_{H_j:\|x\|_{H_j}>1} f_j(x)f_j(\lambda x - B_j x_0)\|x\|_{H_j}\ dx < \infty.$$

By the Fubini-Tonelli theorem it then suffices to show that

$$\int_{x_0\in H:\|x_0\|_H\leq\varepsilon} \int_1^2 \int_{x\in H_j:\|x\|_{H_j}>1} f_j(x)f_j(\lambda x - B_j x_0)\|x\|_{H_j}\ dx d\lambda dx_0 < \infty.$$

In fact we claim that

$$\int_{x_0\in H:\|x_0\|_H\leq\varepsilon} \int_1^2 f_j(\lambda x - B_j x_0)\ d\lambda dx_0 \leq \frac{C}{\|x\|_{H_j}} \int_{H_j} f_j$$

for all $x \in H_j$ with $\|x\|_{H_j} > 1$, and some finite constant $C = C(\varepsilon, B_j, H_j)$ depending only on ε, B_j, and H_j. To see this, we divide the interval $\{1 \leq \lambda \leq 2\}$ into $O(1/\|x\|_{H_j})$ intervals I of radius $O(\|x\|_{H_j})$. For each interval I, we observe that

$$\int_{x_0 \in H : \|x_0\|_H \leq \varepsilon} f_j(\lambda x - B_j x_0)\ d\lambda dx_0 \leq C \int_{\|y - \lambda_I x\|_{H_j} \leq C(1+\varepsilon)} f_j(y)\ dy$$

where λ_I is the midpoint of I and the constants C can depend on ε, B_j, H_j. The claim then follows by integrating in $\lambda \in I$, then summing in I, observing that the balls $\{y \in H_j : \|y - \lambda_I x\|_{H_j} \leq C(1+\varepsilon)\}$ have an overlap of at most C.

To conclude, we have located an extremiser $\mathbf{f}$ which is positive near the origin, and obeys the moment condition $\int_{H_j}(1 + \|x\|_{H_j}) f_j(x)\ dx < \infty$. We can then use closure under multiplication again, replacing the f_j by $f_j(x) f_j(x - B_j x_0)$ for some small $x_0 \in H$ and arguing as before, to improve this moment condition to $\int_{H_j}(1 + \|x\|_{H_j}^2) f_j(x)\ dx < \infty$. Indeed by iterating this we can ensure that $\int_{H_j}(1 + \|x\|_{H_j}^N) f_j(x)\ dx < \infty$ for any specified N (e.g., $N = 100 \dim(H_j)$). In particular, each f_j is square integrable, which implies that the Fourier transforms $\hat{f}_j$ are also square-integrable. It is also bounded, since f_j is integrable. One can then replace $|\hat{f}_j(\xi)|^2$ with $|\hat{f}_j(\xi)|^2 |\hat{f}_j(\lambda\xi)|^2$ for any $1 \leq \lambda \leq 2$, by convolving each f_j with a rescaled version of itself; this may slightly reduce the amount of moment conditions available but this is of no concern since we can make N arbitrary. Note that the new input will still be an extremiser, thanks to Lemma 6.3. By arguing as before (but now in the Fourier domain) we see that we have the moment conditions $\int_{H_j} |\hat{f}_j(\xi)|^2 |\hat{f}_j(\lambda\xi)|^2 \|\xi\|_{H_j}\ d\xi$ for almost every λ. Thus by replacing $\mathbf{f}$ if necessary we can assume that the extremiser $\mathbf{f}$ obeys the regularity condition $\int_{H_j} |\hat{f}_j(\xi)|^2 (1 + \|\xi\|_{H_j})\ d\xi < \infty$ for each j. Indeed one can iterate this argument and obtain an extremiser for which $\int_{H_j} |\hat{f}_j(\xi)|^2 (1 + \|\xi\|_{H_j}^N)\ d\xi < \infty$. To conclude, we have now obtained an extremiser which has any specified amount of Sobolev regularity and decay.

We can now convolve $\mathbf{f}$ with its reflection $\tilde{\mathbf{f}}$ as before to recover the symmetry of $\mathbf{f}$. We may also normalise $\int_{H_j} f_j = 1$ for all j.

The input $\mathbf{f}$ now obeys enough regularity for the central limit theorem. If we set $\mathbf{f}^{(n)} = (f_j^{(n)})_{1 \leq j \leq m}$ to be the rescaled iterated convolution

$$f_j^{(n)}(x) := n^{(\dim H_j)/2} f_j * \ldots * f_j(\sqrt{n} x)$$

where $f_j * \ldots * f_j$ is the n-fold convolution of f_j, then each of the $\mathbf{f}^{(n)}$ are extremisers thanks to Lemma 6.3. Also, the central limit theorem shows that $f_j^{(n)}$ converges in the L^1 topology (for instance) to a centred gaussian, normalised to have total mass one. Applying Lemma 6.3 again we conclude that we can find an extremiser which consists entirely of centred gaussians, which thus implies that every extremisable Brascamp–Lieb datum is gaussian extremisable as desired. ■

Remark 6.6. Proposition 6.5 implies Lieb's theorem (Theorem 1.9) for extremisable data $(\mathbf{B}, \mathbf{p})$, though this is not directly interesting since it is difficult to verify that

$(\mathbf{B}, \mathbf{p})$ is extremisable without first using Lieb's theorem. However, if one assumes Lieb's theorem as a "black box", then the above results do shed some light on why the monotonicity formula approach worked in the gaussian-extremisable case, as follows. Suppose that $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable with some gaussian extremiser $\mathbf{A} = (A_j)_{1 \le j \le t}$. Using Lieb's theorem we see that $(\exp(-\pi \langle A_j x, x \rangle_{H_j}))_{1 \le j \le m}$ is then an extremiser. By Lemma 6.3, for any $t > 0$ the heat kernel

$$\mathbf{K}_{\mathbf{A}}(t)(x) := \Big(\det(A_j/t)^{1/2} \exp(-\pi \langle A_j x, x \rangle_{H_j}/t)\Big)_{1 \le j \le t}$$

is also an extremiser. Define the heat operators $U_{\mathbf{A}}(t)$ on inputs $\mathbf{f}$ by $U_{\mathbf{A}}(t)\mathbf{f} := \mathbf{f} * \mathbf{K}_{\mathbf{A}}(t)$. Applying Lemma 6.1 we see that

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}; U_{\mathbf{A}}(t)f) \ge \mathrm{BL}(\mathbf{B}, \mathbf{p}; f) \text{ for all } t > 0.$$

Indeed, thanks to the well-known semigroup law

$$U_{\mathbf{A}}(t+s) = U_{\mathbf{A}}(t) U_{\mathbf{A}}(s) \text{ for all } s, t > 0$$

for the heat equation, we conclude that $\mathrm{BL}(\mathbf{B}, \mathbf{p}; U_{\mathbf{A}}(t)f)$ is non-decreasing in time. Also, one can compute (using (7)) that $\lim_{t \to +\infty} \mathrm{BL}(\mathbf{B}, \mathbf{p}; U_{\mathbf{A}}(t)\mathbf{f}) \le \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A})$ for all sufficiently well-behaved $\mathbf{f}$ (e.g., rapidly decreasing $\mathbf{f}$ will suffice). From this we can recover the type of heat flow monotonicity that was so crucial in the proof of Proposition 2.8.

7. Structural theory II : semisimplicity and existence of extremisers

Let $(\mathbf{B}, \mathbf{p})$ be a Brascamp–Lieb datum obeying the conditions (7), (8), which we have shown to be necessary and sufficient for the Brascamp–Lieb constant $\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$ to be finite. We now explore further the question of when $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable. We already have Theorem 1.15, which shows that when there are no critical subspaces then $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable. However it is certainly possible to be gaussian-extremisable in the presence of critical spaces; consider for instance Hölder's inequality (Example 1.3), which has plenty of gaussian extremisers but for which every proper subspace is critical. To resolve this question more satisfactorily we need the notion of an *indecomposable m-transformation*.

Definition 7.1 (Direct sum)**.** If $\mathbf{B} = (H, (H_j)_{1 \le j \le m}, (B_j)_{1 \le j \le m})$ and $\mathbf{B}' = (H', (H'_j)_{1 \le j \le m}, (B'_j)_{1 \le j \le m})$ are m-transformations, we define the *direct sum* $\mathbf{B} \oplus \mathbf{B}'$ to be the m-transformation

$$\mathbf{B} \oplus \mathbf{B}' := (H \oplus H', (H_j \oplus H'_j)_{1 \le j \le m}, (B_j \oplus B'_j)_{1 \le j \le m})$$

where $H \oplus H' := \{(x, x') : x \in H, x' \in H'\}$ is the Hilbert space direct sum of H and H' with the usual direct sum inner product $\langle (x, x'), (y, y') \rangle_{H \oplus H'} := \langle x, y \rangle_H + \langle x', y' \rangle_{H'}$, and $B_j \oplus B'_j : H \oplus H' \to H_j \oplus H'_j$ is the direct sum $(B_j \oplus B'_j)(x, x') := (B_j x, B'_j x')$. We say that an m-transformation $\mathbf{B}$ is *decomposable* if it is equivalent to the direct sum $\mathbf{B}_1 \oplus \mathbf{B}_2$ of two m-transformations $\mathbf{B}_1, \mathbf{B}_2$ whose domains have strictly smaller dimension than that of $\mathbf{B}$, in which case we refer to $\mathbf{B}_1$ and $\mathbf{B}_2$ as *factors* of $\mathbf{B}$. We say that $\mathbf{B}$ is *indecomposable* if it is not decomposable.

It is easy to verify that if $(\mathbf{B}_1, \mathbf{p}), (\mathbf{B}_2, \mathbf{p})$ are Brascamp–Lieb data with domains H_1, H_2 which obey (7), and $(\mathbf{B}, \mathbf{p}) = (\mathbf{B}_1 \oplus \mathbf{B}_2, \mathbf{p})$ is the direct sum (with domain

$H_1 \oplus H_2$), then the subspaces $H_1 \oplus \{0\}$ and $\{0\} \oplus H_2$ of $H_1 \oplus H_2$ are critical subspaces with respect to $(\mathbf{B}, \mathbf{p})$. Furthermore, the restriction of $\mathbf{B}$ to $H_1 \oplus \{0\}$ is equivalent to $\mathbf{B}_1$, and the quotient of $\mathbf{B}$ by $H_1 \oplus \{0\}$ is equivalent to $\mathbf{B}_2$; similarly with the roles of 1 and 2 reversed. In particular this implies (from Lemma 4.8 and Lemma 3.3) that

$$\mathrm{BL}(\mathbf{B}_1 \oplus \mathbf{B}_2, \mathbf{p}) = \mathrm{BL}(\mathbf{B}_1, \mathbf{p}) \, \mathrm{BL}(\mathbf{B}_2, \mathbf{p}). \tag{40}$$

We of course have a similar assertion for the gaussian Brascamp–Lieb constants.

Remark 7.2. The Krull–Schmidt theorem for quiver representations (see e.g., [ARS]) implies (as a special case) that every m-transformation has a factorisation (up to equivalence) as the direct sum of indecomposable m-transformations, and furthermore that this factorisation is unique up to equivalence and permutations. The question of classifying what the indecomposable factors are of a given m-transformation, however, is quite difficult. In the rank one case (Example 1.7) a satisfactory description was given by Barthe [Bar2, Proposition 1]. In fact the indecomposable (and prime) components can be described explicitly in this case as follows. Let us introduce the relation $\bowtie$ on $\{1, \dots, m\}$ by requiring $i \bowtie j$ whenever there exists a collection $I \subset \{1, \dots, m\} \backslash \{i, j\}$ such that $(v_k)_{k \in I \cup \{i\}}$ and $(v_k)_{k \in I \cup \{j\}}$ are bases for H. Let $\sim$ be the transitive completion of $\bowtie$. Then the indecomposable components of H take the form $\mathrm{span}(v_j : j \in I)$, where I is any equivalence class for $\sim$. In the higher rank case we cannot expect such a completely explicit factorisation; at a minimum, we must allow for factorisation to only be unique up to equivalence. For instance, for Hölder's inequality (Example 1.3), any decomposition of H into $\dim(H)$ independent one-dimensional subspaces will induce a factorisation of the associated Brascamp–Lieb datum into indecomposables; these factorisations are only unique up to equivalence.

We now give a geometric criterion for indecomposability.

Definition 7.3 (Critical pair)**.** Let $\mathbf{B}$ be an m-transformation. A pair (V, W) of subspaces of H is said to be a *critical pair* for $\mathbf{B}$ if V and W are complementary in H (thus $V \cap W = \{0\}$ and $V + W = H$), and for each j, $B_j V$ and $B_j W$ are complementary in H_j. We say the critical pair (V, W) is *proper* if $\{0\} \subsetneq V, W \subsetneq H$.

Remarks 7.4. Observe that the exponents p_j play no role in this definition. If the B_j are surjective, then $(0, H)$ and $(H, 0)$ are of course critical pairs, though they are not proper. Furthermore, in this case one can use identities such as $\dim(V) = \dim(B_j V) + \dim(V \cap \ker(B_j))$ to formulate an equivalent definition of critical pair by requiring that V, W are complementary in H, and $V \cap \ker(B_j)$ and $W \cap \ker(B_j)$ are complementary in $\ker(B_j)$. In the rank one case $\dim(H_j) \equiv 1$, this condition was identified by Barthe [Bar2], who used it in obtaining a factorisation of Brascamp–Lieb data into indecomposable components. We shall extend that analysis here to the higher rank case.

Example 7.5. *For Hölder's inequality (Example 1.3), every complementary pair in $\mathbf{R}^n$ is a critical pair. For the Loomis–Whitney inequality (Example 1.4), any coordinate plane and its orthogonal complement will be a critical pair. For Young's inequality (Example 1.5) with $d = 1$, there are no proper critical pairs (compare with Example 4.4). In the rank-one case (Example 1.7), (V, W) is a critical pair if and only if V and W are complementary, and $\{v_1, \dots, v_m\} \subseteq V \cup W$. Also, if (V, W)*

is a critical pair for $\mathbf{B}$*, and* $\mathbf{B}'$ *is equivalent to* $\mathbf{B}$ *with intertwining transformations* C, C_j*, then* $(C^{-1}V, C^{-1}W)$ *is a critical pair for* $\mathbf{B}'$*.*

Example 7.6. *If* $\mathbf{B}$ *is an* m*-transformation with domain* H*, then the direct sum* $\mathbf{B} \oplus \mathbf{B}$ *on* $H \oplus H$ *has* $(H \oplus \{0\}, \{0\} \oplus H)$ *as an obvious critical pair. But there is also a "diagonal" critical pair*

$$(\{(x, x) : x \in H\}, \{(y, -y) : y \in H\}).$$

By taking tensor products of inputs in the first critical pair and then projecting onto the second critical pair (in the spirit of Lemma 4.8) one can recover K. Ball's convolution inequality Lemma 6.1, as well as several of the closure properties in Lemma 6.3. We omit the details.

Lemma 7.7 (Geometric criterion for indecomposability)**.** *An* m*-transformation* $\mathbf{B}$ *is indecomposable if and only if it has no proper critical pairs.*

Proof If $\mathbf{B}$ is decomposable, then it is equivalent to a direct sum $\mathbf{B}_1 \oplus \mathbf{B}_2$ on a domain $H_1 \oplus H_2$ where H_1 and H_2 are not $\{0\}$. The pair $(H_1 \oplus \{0\}, \{0\} \oplus H_2)$ can be easily seen to be a proper critical pair for $\mathbf{B}_1 \oplus \mathbf{B}_2$, and thus $\mathbf{B}$ also has a proper critial pair by the remarks in Example 7.5.

Now suppose conversely that $\mathbf{B}$ has a proper critical pair (V, W). Since V and W are complementary, we can apply an invertible linear transformation C on H to make V and W orthogonal complements, while replacing $\mathbf{B}$ with an equivalent m-transformation. Thus we may assume without loss of generality that V and W are orthogonal complements. Similarly we may assume that B_jV and B_jW are orthogonal complements. But then $\mathbf{B}$ is canonically equivalent to the direct sum $\mathbf{B}_V \oplus \mathbf{B}_W$ of its restrictions to V and W, and is thus decomposable. ■

Remark 7.8. The above proof in fact shows that if (V, W) is a critical pair for $\mathbf{B}$, then $\mathbf{B}$ is equivalent to $\mathbf{B}_V \oplus \mathbf{B}_W$.

One feature of critical pairs is that they are "universally critical", in the sense that they are critical for all admissible exponents $\mathbf{p}$:

Lemma 7.9 (Critical pairs are critical subspaces)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum obeying* (8), (7)*, and let* (V, W) *be a proper critical pair for* $\mathbf{B}$*. Then* V *and* W *are both critical subspaces for* $(\mathbf{B}, \mathbf{p})$*.*

Remark 7.10. This lemma implies that all simple Brascamp–Lieb data are indecomposable. The converse is not true however; consider for instance the indecomposable example in Example 3.9.

Proof From (8), (7), and the hypothesis that (V, W) is a critical pair, we have

$$\begin{aligned}\dim(H) &= \dim(V) + \dim(W)\\ &\leq \sum_{j=1}^m p_j \dim(B_j V) + \sum_{j=1}^m p_j \dim(B_j W)\\ &= \sum_{j=1}^m p_j \dim(H_j)\\ &= \dim(H)\end{aligned}$$

and hence the inequality above must be equality. This implies that V and W are critical, as claimed. ∎

Critical pairs are related to extremisability in two ways. First of all, extremisability of a product is equivalent to extremisability in the factors:

Lemma 7.11 (Extremisability factors)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum obeying the conditions* (7), (8), *and let* (V, W) *be a critical pair for* $\mathbf{B}$. *Then* $(\mathbf{B}, \mathbf{p})$ *is extremisable if and only if* $(\mathbf{B}_V, \mathbf{p})$ *and* $(\mathbf{B}_W, \mathbf{p})$ *are both extremisable. Similarly,* $(\mathbf{B}, \mathbf{p})$ *is gaussian-extremisable if and only if* $(\mathbf{B}_V, \mathbf{p})$ *and* $(\mathbf{B}_W, \mathbf{p})$ *are both gaussian-extremisable.*

Proof By Remark 7.8 we know that $\mathbf{B}$ is equivalent to $\mathbf{B}_V \oplus \mathbf{B}_W$. From Lemma 3.3 and Lemma 4.8 we thus see that if $(\mathbf{B}, \mathbf{p})$ is extremisable then $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_W, \mathbf{p})$ are extremisable, and if $(\mathbf{B}, \mathbf{p})$ is gaussian-extremisable then $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_W, \mathbf{p})$ are gaussian-extremisable.

In the converse direction, if equality is attained in (1) for the input $(f_{j,V})_{1 \leq j \leq m}$ for $(\mathbf{B}_V, \mathbf{p})$ and for the input $(f_{j,W})_{1 \leq j \leq m})$ for $(\mathbf{B}_W, \mathbf{p})$, then the direct sums $(f_{j,V} \oplus f_{j,W})_{1 \leq j \leq m}$ will attain equality in $(\mathbf{B}_V \oplus \mathbf{B}_W, \mathbf{p})$, thanks to (40) (or Lemma 4.8). The claim for the extremisation problem (5) is similar. ∎

Furthermore, in the presence of an extremiser one can obtain a converse to Lemma 7.9.

Lemma 7.12 (Critical complements for gaussian-extremisable data)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a gaussian-extremisable Brascamp–Lieb datum. Then for every critical subspace* V *of* $(\mathbf{B}, \mathbf{p})$ *there exists a complementary space* W *to* V *such that* (V, W) *is a critical pair for* $\mathbf{B}$. *If furthermore* $(\mathbf{B}, \mathbf{p})$ *is geometric, then* W *is the orthogonal complement of* V.

Proof We may drop exponents for which $p_j = 0$, and thus assume that $p_j > 0$ for all j. By Proposition 3.6 we may assume without loss of generality that $(\mathbf{B}, \mathbf{p})$ is a geometric Brascamp–Lieb datum, thus we may assume H_j is a subspace of H, $B_j : H \to H_j$ is the orthogonal projection, and (9) holds.

Now let $\pi : H \to V$ be the orthogonal projection onto V. Observe that $P_j \pi$ is a contraction from H to $P_j V$ and hence has trace at most $\dim(P_j V)$. Thus, since V

is a critical subspace,

$$\begin{aligned}\dim(V) &= \mathrm{tr}_H(V) \\ &= \mathrm{tr}_H(\sum_{j=1}^m p_j P_j V) \\ &\leq \sum_{j=1}^m p_j \dim(P_j V) \\ &= \dim(V)\end{aligned}$$

and hence we must in fact have $\mathrm{tr}_H(P_j\pi) = \dim(P_j V)$ for all j. Thus $P_j\pi$ is in fact a co-isometry (the adjoint is an isometry), which means that P_jV and $P_j(V^\perp)$ are orthogonal complements in H_j. This implies that $(V, V^\perp)$ form a critical pair, and the claim follows. ■

We can now give a characterisation of when extremisers exist.

Theorem 7.13 (Necessary and sufficient conditions for extremisers)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum with* $p_j > 0$ *for all* j*. Then the following nine statements are equivalent.*

- *(a)* $(\mathbf{B}, \mathbf{p})$ *is extremisable.*
- *(b)* $(\mathbf{B}, \mathbf{p})$ *is gaussian-extremisable.*
- *(c) A local extremiser exists to* (4)*.*
- *(d) There exists a gaussian input* $\mathbf{A} = (A_j)_{1\leq j\leq m}$ *such that the matrix* $M := \sum_{j=1}^m p_j B_j^* A_j B_j$ *obeys* $A_j^{-1} = B_j M^{-1} B_j^*$ *for all* j*.*
- *(e) The scaling condition* (7) *holds,* $\mathbf{B}$ *is non-degenerate, and there exists a gaussian input* $\mathbf{A} = (A_j)_{1\leq j\leq m}$ *such that the matrix* $M := \sum_{j=1}^m p_j B_j^* A_j B_j$ *obeys* $A_j^{-1} \geq B_j M^{-1} B_j^*$ *for all* j*.*
- *(f) The scaling condition* (7) *holds,* $\mathbf{B}$ *is non-degenerate, and there exists a gaussian input* $\mathbf{A} = (A_j)_{1\leq j\leq m}$ *such that* $B_j^* A_j B_j \leq \sum_{i=1}^m p_i B_i^* A_i B_i$*.*
- *(g)* $(\mathbf{B}, \mathbf{p})$ *is equivalent to a geometric Brascamp–Lieb datum.*
- *(h) The bounds* (8)*,* (7) *hold, and every critical subspace* V *is part of a critical pair* (V, W)*.*
- *(i) The bounds* (8)*,* (7) *hold, and every indecomposable factor of* $(\mathbf{B}, \mathbf{p})$ *is simple.*

One can view the equivalence (a) $\iff$ (i) as a statement that extremisability is equivalent to being "semisimple" (the direct sum of simple factors).

Proof The implication (a) $\implies$ (b) is given by Proposition 6.5, while the converse direction (b) $\implies$ (a) is given by Theorem 1.9. The equivalence of (b)-(g) follows from Proposition 3.6. The implication (b) $\implies$ (h) is given by Theorem 1.15 and Lemma 7.12. Now let us prove the converse implication (h) $\implies$ (b). We induct on the dimension of H. If there are no proper critical subspaces then the claim follows from Theorem 1.15. Now suppose that there is a proper critical space V, which is then part of a critical pair (V, W) of proper subspaces. We then pass to the factors $(\mathbf{B}_V, \mathbf{p})$ and $(\mathbf{B}_W, \mathbf{p})$. By Lemma 4.6 these factors also obey (8), (7). Also, we observe that every critical subspace V' of V is part of a critical pair (V', W') in V. To see this, observe from hypothesis that V' is part of a critical pair (V', W'') in H. Now set $W' := V \cap W''$; since V' and W'' are complementary in H, and $V' \subset V$, we see that V' and W' are complementary in V. In particular, $B_j V' + B_j W' = B_j V$. On the other hand, since $B_j V'$ and $B_j W''$ are complementary in H_j, and $B_j W'$ is contained in $B_j W''$, we have $B_j V' \cap B_j W' = \{0\}$. Thus $B_j V'$ and $B_j W'$ are complementary in $B_j V$, and hence (V', W') is a critical pair in V as claimed.

Applying the induction hypothesis, we conclude that the supremum in (5) for $(\mathbf{B}_V, \mathbf{p})$ is attained. Similarly for V replaced by W. The claim now follows from Lemma 7.12.

Now we show the implication (b) $\implies$ (i). Applying Lemma 4.8 we see that the supremum in (5) is attained for every indecomposable factor of $(\mathbf{B}, \mathbf{p})$. Applying the equivalence between (b) and (h) already proven, and noting that indecomposable factors cannot contain proper critical pairs by definition, we thus see that every indecomposable factor contains no proper critical subspaces as claimed.

Finally, to show that (h) $\implies$ (b), we see by the equivalence of (b) and (h) that the supremum in (5) is attained for every indecomposable factor of $(\mathbf{B}, \mathbf{p})$. Factoring $(\mathbf{B}, \mathbf{p})$ in some arbitrary manner as a direct sum of indecomposable components and using Lemma 7.11 we obtain the claim. ■

It should be mentioned that the above equivalences only hold with the assumption $p_j > 0$. Of course any exponent with $p_j = 0$ can be omitted without affecting most of the properties listed above (specifically, this does not affect (a)-(g)), but it does affect the factorisation of data into indecomposable components.

Example 7.14. *Let us return to Example 3.9. This* 3*-transformation* $\mathbf{B}$ *is indecomposable, and the datum* $(\mathbf{B}, \mathbf{p})$ *is simple if* (p_1, p_2, p_3) *lies in the interior of the triangle. However if* (p_1, p_2, p_3) *is on a vertex of this triangle, then one of the* p_j *vanishes, and on removing this exponent we are left with a* 2*-transformation which is now decomposable (with the two components clearly being simple). This is why extremisers exist on the interior and vertices of the triangles but not on the open edges.*

More generally, the above proposition yields an explicit test as to whether a given datum $(\mathbf{B}, \mathbf{p})$ is extremisable. First, one removes any exponents for which $p_j = 0$. Then one splits $\mathbf{B}$ into indecomposable components $(\mathbf{B}_i, \mathbf{p})$. If $\mathbf{p}$ lies in $S(\mathbf{B}_i) \cap \Pi(\mathbf{B}_i)^\circ$ for each component $\mathbf{B}_i$, then $(\mathbf{B}, \mathbf{p})$ is extremisable, but if $\mathbf{p}$ lies on the boundary of $\Pi(\mathbf{B}_i)$ for any i then $(\mathbf{B}, \mathbf{p})$ is not extremisable.

8. Regularised and localised Brascamp–Lieb inequalities

We have seen a number of connections between Brascamp–Lieb constants, gaussians, and the heat equation. We now pursue these connections further by introducing a generalisation of the Brascamp–Lieb constants. As a byproduct of this analysis we give an alternate proof of Lieb's theorem (Theorem 1.9) that does not rely on factorisation through critical subspaces.

We first examine the situation when the data are regularised at a certain scale; morally this corresponds to a discrete setting. This allows us to then examine the situation when the data are gaussian-localised, effected by appending a fixed gaussian multiplier to the input and then scaling. The gaussian-localised situation falls in principle under the framework of [L].

Our first definition formalises our notion of regularity.

Definition 8.1 (Type G functions)**.** Let H be a Euclidean space, and let $G : H \to H$ be positive definite. We define a *type G function* to be any function $u : H \to \mathbf{R}^+$ of the form

$$u(x) = \det_H(G)^{1/2} \int_H \exp(-\pi\langle G(x-y), (x-y)\rangle_H)\ d\mu(y)$$

where μ is a positive finite Borel measure on H with non-zero total mass. If μ is a point mass, we say that u is of *extreme type G*; thus the functions of extreme type G are simply the translates and positive scalar multiples of the function $\exp(-\pi\langle Gx, x\rangle_H)$.

Remark 8.2. Observe that type G functions are smooth and strictly positive. Also, if $G_1 \geq G_2 > 0$ then every function of type G_2 is also of type G_1, since a gaussian such as $\exp(-\pi\langle G_2 x, x\rangle_H)$ can be expressed as a convolution of $\exp(-\pi\langle G_1 x, x\rangle_H)$ with a positive finite measure (this can be seen for instance using the Fourier transform). Positive finite measures themselves can be informally viewed as "functions of type $+\infty$".

Remark 8.3. Type G functions arise naturally in the study of heat equations. More precisely, if $u : \mathbf{R}^+ \times H \to \mathbf{R}^+$ is a solution to the heat equation

$$u_t = \frac{1}{4\pi} \operatorname{div}(G^{-1}\nabla u)$$

and $u(0) = \mu$ is a positive finite measure, then from the explicit solution

$$u(t,x) = \det_H(G/t)^{1/2} \int_H \exp(-\pi\langle G(x-y), (x-y)\rangle_H/t)\ d\mu(y)$$

we see that $u(t)$ is of type G/t for all $t > 0$. More generally, if $u(s)$ is of type A for some $s \geq 0$ and $A > 0$, then $u(t)$ is of type $(A^{-1} + (t-s)G^{-1})^{-1}$ for all $t > s$, as can be seen for instance by using the Fourier transform, or alternatively by convolving one fundamental solution to a heat equation with another using (3). Note that this operation of harmonic addition has already appeared in Corollary 6.4.

Definition 8.4 (generalised Brascamp–Lieb constant)**.** If $(\mathbf{B}, \mathbf{p})$ is a Brascamp–Lieb datum, we define a *generalised Brascamp–Lieb datum* to be a triple $(\mathbf{B}, \mathbf{p}, \mathbf{G})$, where $(\mathbf{B}, \mathbf{p})$ is a Brascamp–Lieb datum and $\mathbf{G}$ is a gaussian input for $(\mathbf{B}, \mathbf{p})$. We shall refer to $\mathbf{G}$ as an *m-type*. We say that an input $\mathbf{f} = (f_j)_{1\leq j\leq m}$ is of *type* $\mathbf{G} = (G_j)_{1\leq j\leq m}$ if each f_j is of type G_j. We define the *generalised Brascamp–Lieb constant* $\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G})$ to be

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \sup_{\mathbf{f} \text{ of type } \mathbf{G}} \mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}). \tag{41}$$

We also define the generalised gaussian Brascamp–Lieb constant $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G})$ by the formula

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}) := \sup_{\mathbf{A} \leq \mathbf{G}} \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}), \tag{42}$$

where the supremum extends over all gaussian inputs $\mathbf{A} = (A_j)_{1 \leq j \leq m}$ such that $A_j \leq G_j$ for all j. We say that $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ is *extremisable* if $\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G})$ is finite and equality can be attained in (41) for some input $\mathbf{f}$ of type $\mathbf{G}$, and is *gaussian-extremisable* if $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G})$ is finite and the supremum in (42) can be attained for some gaussian input $\mathbf{A} \leq \mathbf{G}$.

Remark 8.5. One may restrict attention to non-degenerate $\mathbf{B}$ for the following reasons. If $\bigcap_{j=1}^m \ker(B_j) \neq \{0\}$ then it is easy to see that both constants will be infinite. If one or more of the B_j is not surjective, one can simply restrict H_j to $B_j V$ (and restrict the quadratic form associated to G_j to $B_j V$ also) to obtain an equivalent problem.

Clearly we have

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) \geq \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}); \tag{43}$$

we shall later show in Corollary 8.15 that this is in fact an equality, in analogy with Theorem 1.9. Clearly $\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) \leq \mathrm{BL}(\mathbf{B}, \mathbf{p})$ and $\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}) \leq \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p})$, but it is certainly possible for strict inequality to hold for either inequality. On the other hand, a simple regularisation argument shows that

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}) = \lim_{\lambda \to \infty} \mathrm{BL}(\mathbf{B}, \mathbf{p}, \lambda \mathbf{G}); \quad \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}) = \lim_{\lambda \to \infty} \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \lambda \mathbf{G}) \tag{44}$$

for any m-type $\mathbf{G}$. Furthermore, it turns out that there are analogues not only of Lieb's theorem but also of Theorem 1.15 in this generalised setting. We shall state these results presently, but let us first develop some preliminary lemmas. We begin with a basic log-convexity estimate.

Lemma 8.6 (Log-convexity lemma)**.** *Let $G : H \to H$ be positive definite, let u be of type G, and let v of be of extreme type G. Then u/v is log-convex (i.e., $\nabla^2 \log(u/v)$ is positive semi-definite). If u is not of extreme type G, then u/v is strictly log-convex (i.e., $\nabla^2 \log(u/v)$ is positive definite).*

Proof By applying an invertible linear transformation if necessary one can reduce to the case $H = \mathbf{R}^n$, $G = I_n$. By translation we may then assume that v is the extreme type I_n function

$$v(x) = \exp(-\pi|x|^2).$$

We can write the type I_n function $u(x)$ as

$$u(x) = \int_{\mathbf{R}^n} \exp(-\pi|x-y|^2)\ d\mu(y)$$

and hence

$$\frac{u(x)}{v(x)} = \int_{\mathbf{R}^n} e^{2\pi\langle x,y\rangle/2}\ e^{-\pi|y|^2} d\mu(y). \tag{45}$$

At this point we could argue that u/v is log-convex by observing that the exponential functions $e^{2\pi\langle x,y\rangle}$ are log-convex, and that the superposition of log-convex functions is again log-convex. However we shall give a more explicit argument which also yields the second claim. Taking gradients of the above equation we obtain

$$\frac{u(x)}{v(x)}\nabla \log \frac{u(x)}{v(x)} = \int_{\mathbf{R}^n} 2\pi y e^{2\pi\langle x,y\rangle}\ e^{-\pi|y|^2} d\mu(y)$$

and hence by (45) again

$$\int_{\mathbf{R}^n} (\nabla \log \frac{u(x)}{v(x)} - 2\pi y) e^{2\pi\langle x,y\rangle}\ e^{-\pi|y|^2} d\mu(y) = 0.$$

We multiply this by an arbitrary transformation $A : \mathbf{R}^n \to \mathbf{R}^n$ and take divergences, to conclude

$$\int_{\mathbf{R}^n} (\operatorname{div} A\nabla \log \frac{u(x)}{v(x)} + \langle 2\pi Ay, \nabla \log \frac{u(x)}{v(x)} - 2\pi y\rangle) e^{2\pi\langle x,y\rangle}\ e^{-\pi|y|^2} d\mu(y) = 0.$$

Rearranging this using (45) again, we conclude

$$\operatorname{div} A\nabla \log \frac{u(x)}{v(x)} = \frac{\int_{\mathbf{R}^n} \langle A(\nabla \log \frac{u(x)}{v(x)} - 2\pi y), \nabla \log \frac{u(x)}{v(x)} - 2\pi y\rangle e^{2\pi\langle x,y\rangle}\ e^{-\pi|y|^2} d\mu(y)}{\int_{\mathbf{R}^n} e^{2\pi\langle x,y\rangle}\ e^{-\pi|y|^2} d\mu(y)}.$$

Observe that if A is positive semi-definite, then the right-hand side is non-negative. This shows that $\nabla^2 \log \frac{u}{v}$ is log-convex as claimed. Also, if A is positive definite, then we see that for any fixed x there is at most one y for which the integrand in the numerator is zero. Thus if μ is not a point mass then the numerator is always strictly positive, and hence $\frac{u}{v}$ is strictly convex in this case. This concludes the proof. ∎

Corollary 8.7. *Let $G : H \to H$ be positive definite, let $u : H \to \mathbf{R}^+$ be of type G, and let $A : H \to H$ be positive semi-definite. Then*

$$\operatorname{div}(A\nabla \log u) \geq -2\pi \operatorname{tr}_H(AG).$$

If A is not identically zero, then equality occurs if and only if u is extreme type G.

Remark 8.8. Intuitively speaking, this corollary asserts that positive solutions to the heat equation cannot diverge too rapidly, and can be viewed as an explanation for the monotonicity for heat flows.

Proof Let $v : H \to \mathbf{R}^+$ be the extreme type G function $v(x) := \exp(-\pi\langle Gx, x\rangle_H\rangle)$. Observe that $\operatorname{div}(A\nabla \log v) = -2\pi \operatorname{tr}_H(AG)$, and from Lemma 8.6 we have $\operatorname{div}(A\nabla \log \frac{u}{v}) \geq 0$, with strict inequality if A is not identically zero and u is not extreme type G. The claim follows. ∎

We can now give the monotonicity formula for multiple heat flows. We begin with a rather general statement.

Proposition 8.9 (Multilinear heat flow monotonicity formula)**.** *Let $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ be a generalised Brascamp–Lieb datum with all the B_j surjective, and suppose that there is a gaussian input $\mathbf{A} = (A_j)_{1\leq j\leq m}$ with $\mathbf{A} \leq \mathbf{G}$, such that the transformation $M : H \to H$ defined by $M := \sum_{j=1}^m p_j B_j^* A_j B_j$ is invertible and obeys the inequality*

$$A_j^{-1} - B_j M^{-1} B_j^* \geq 0 \tag{46}$$

and the identity

$$(A_j^{-1} - B_j M^{-1} B_j^*)(G_j - A_j) = 0 \tag{47}$$

for each $1 \leq j \leq m$*. Also, for each* $1 \leq j \leq m$*, let* $\tilde{u}_j(1) : H_j \to \mathbf{R}^+$ *be of type* G_j*, and let* $\tilde{u}_j : [1, \infty) \times H_j \to \mathbf{R}^+$ *be the solution to the heat equation*

$$\partial_t \tilde{u}_j = \frac{1}{4\pi} \operatorname{div}(A_j^{-1} \nabla \tilde{u}_j)$$

with initial data $\tilde{u}_j(1)$ *at* $t = 1$*. Then the quantity*

$$Q(t) := t^{(\sum_{j=1}^m p_j \dim(H_j) - \dim(H))/2} \int_H \prod_{j=1}^m (\tilde{u}_j \circ B_j)^{p_j}$$

is monotone non-decreasing in time for $t > 1$*.*

Furthermore, if there is a j *for which equality in* (46) *does not hold, and* $\tilde{u}_j(1)$ *is not of extreme type* G_j*, then* $Q(t)$ *is strictly increasing for* $t > 1$*.*

Proof We have

$$\tilde{u}_j(1, x) = \det_{H_j}(G_j)^{1/2} \int_{H_j} \exp(-\pi \langle G_j(x - y), (x - y) \rangle_{H_j})\, d\mu_j(y) \tag{48}$$

for some positive finite Borel measure μ_j. A simple computation using the Fourier transform then shows that

$$\begin{aligned} \tilde{u}_j(t, x) &= \det_{H_j}(G_j^{-1} + (t-1)A_j^{-1})^{-1/2} \\ &\quad \times \int_{H_j} \exp(-\pi \langle (G_j^{-1} + (t-1)A_j^{-1})^{-1}(x - y), (x - y) \rangle_{H_j})\, d\mu_j(y). \end{aligned} \tag{49}$$

By a Fatou lemma argument we may assume without loss of generality that the measures μ_j are compactly supported. We may also assume $p_j > 0$ for all j as the degenerate exponents $p_j = 0$ can be easily omitted.

Let $D_j : H_j \to H$ be any linear right-inverse to B_j, thus $B_j D_j = \mathrm{id}_{H_j}$. Such an inverse exists since B_j is assumed to be surjective. The exact choice of D_j will not be relevant; for instance one could take $D_j = B_j^*(B_j B_j^*)^{-1}$. We will apply Lemma 2.6 with

$$\begin{aligned} I &:= (0, +\infty) \\ u_j &:= \tilde{u}_j \circ B_j \\ \vec{v}_j &:= -\frac{1}{4\pi} D_j A_j^{-1} (\nabla \log \tilde{u}_j) \circ B_j \\ \vec{v} &:= M^{-1} \sum_{j=1}^m p_j B_j^* A_j B_j \vec{v}_j \\ \alpha &:= \frac{1}{2}(\dim(H) - \sum_{j=1}^m p_j \dim(H_j)). \end{aligned}$$

Let us first verify the technical condition that $\vec{v} \prod_{j=1}^m u_j^{p_j}$ is rapidly decreasing in space. Because the μ_j are compactly supported, one can verify that $\nabla \log \tilde{u}_j$ grows at most polynomially in space, locally uniformly on I. Hence $\vec{v}$ also grows at most polynomially in space. Since M is invertible, we have $\bigcap_{j=1}^m \ker(B_j) = 0$. Since each of the $\tilde{u}_j$ are rapidly decreasing in space, locally uniformly on I, we see that $\prod_{j=1}^m u_j^{p_j}$ is also, and the claim follows.

By Lemma 2.6, we will now be done as soon as we verify the inequalities (11), (12), (13). We begin with (11). On the one hand, we have

$$\partial_t u_j = (\partial_t \tilde{u}_j) \circ B_j = \frac{1}{4\pi} \operatorname{div}(A_j^{-1} \nabla \tilde{u}_j) \circ B_j.$$

On the other hand, from the chain rule and the choice of D_j we have

$$\begin{aligned}
\operatorname{div}(\vec{v}_j u_j) &= -\operatorname{div}([\frac{1}{4\pi} D_j A_j^{-1} (\nabla \log \tilde{u}_j) \tilde{u}_j] \circ B_j) \\
&= -\frac{1}{4\pi} \operatorname{div}(D_j A_j^{-1} (\nabla \tilde{u}_j) \circ B_j) \\
&= -\frac{1}{4\pi} \operatorname{div}(B_j D_j A_j^{-1} \nabla \tilde{u}_j) \circ B_j \\
&= -\frac{1}{4\pi} \operatorname{div}(A_j^{-1} \nabla \tilde{u}_j) \circ B_j
\end{aligned}$$

and so the left-hand side of (11) is zero.

Next we verify (12). We expand the left-hand side as

$$\begin{aligned}
\operatorname{div}(\vec{v} - \sum_{j=1}^m p_j \vec{v}_j) &= \sum_{j=1}^m p_j \operatorname{div}(M^{-1} B_j^* A_j B_j \vec{v}_j - \vec{v}_j) \\
&= -\frac{1}{4\pi} \sum_{j=1}^m p_j \operatorname{div}((M^{-1} B_j^* A_j B_j - \mathrm{id}_{H_j}) D_j A_j^{-1} (\nabla \log \tilde{u}_j) \circ B_j) \\
&= -\frac{1}{4\pi} \sum_{j=1}^m p_j \operatorname{div}(B_j (M^{-1} B_j^* A_j B_j - \mathrm{id}_{H_j}) D_j A_j^{-1} (\nabla \log \tilde{u}_j)) \circ B_j \\
&= \frac{1}{4\pi} \sum_{j=1}^m p_j \operatorname{div}((A_j^{-1} - B_j M^{-1} B_j^*) \nabla \log \tilde{u}_j) \circ B_j.
\end{aligned}$$

From (49) we see that $\tilde{u}_j(t)$ is of type $(G_j^{-1} + (t-1)A_j^{-1})^{-1}$. Since $G_j \geq A_j > 0$ and $t > 1$, we have $(G_j^{-1} + (t-1)A_j^{-1})^{-1} \leq G_j/t$. Thus $\tilde{u}_j(t)$ is also of type G_j/t. Applying (46) and Corollary 8.7 we have

$$\operatorname{div}((A_j^{-1} - B_j M^{-1} B_j^*) \nabla \log \tilde{u}_j) \geq -2\pi \operatorname{tr}_{H_j}((A_j^{-1} - B_j M^{-1} B_j^*) G_j / t).$$

Using (47) we conclude that

$$\operatorname{div}((A_j^{-1} - B_j M^{-1} B_j^*) \nabla \log \tilde{u}_j) \geq -\frac{2\pi}{t} \operatorname{tr}_{H_j}((A_j^{-1} - B_j M^{-1} B_j^*) A_j).$$

Inserting this into the preceding computation, we conclude

$$\begin{aligned}
\operatorname{div}(\vec{v} - \sum_{j=1}^m p_j \vec{v}_j) &\geq -\frac{1}{2t} \sum_{j=1}^m p_j \operatorname{tr}_{H_j}((A_j^{-1} - B_j M^{-1} B_j^*) A_j) \\
&= \frac{1}{2t} \sum_{j=1}^m p_j (\operatorname{tr}_H(M^{-1} B_j^* A_j B_j) - \dim(H_j)) \\
&= \frac{1}{2t} (\operatorname{tr}_H(M^{-1} M) - \sum_{j=1}^m p_j \dim(H_j)) \\
&= \frac{1}{2t} (\dim(H) - \sum_{j=1}^m p_j \dim(H_j)) \\
&= \frac{\alpha}{t}
\end{aligned}$$

as desired.

Finally, we verify (13). From the chain rule we have

$$\nabla \log u_j = B_j^* (\nabla \log \tilde{u}_j) \circ B_j$$

and hence by definition of $\vec{v}_j$ and D_j

$$\nabla \log u_j = -B_j^* A_j B_j \vec{v}_j.$$

Thus we can write the left-hand side of (13) as

$$\sum_{j=1}^m p_j \langle B_j^* A_j B_j (\vec{v} - \vec{v}_j), -\vec{v}_j \rangle_H.$$

On the other hand, by definition of $\vec{v}$ and M we have

$$\sum_{j=1}^m p_j B_j^* A_j B_j (\vec{v} - \vec{v}_j) = M\vec{v} - \sum_{j=1}^m p_j B_j^* A_j B_j \vec{v}_j = 0.$$

Thus we can write the left-hand side of (13) as

$$\sum_{j=1}^m p_j \langle B_j^* A_j B_j (\vec{v} - \vec{v}_j), (\vec{v} - \vec{v}_j) \rangle_H.$$

Since A_j is positive definite, we see that this expression is non-negative as desired. This completes the proof of monotonicity. The proof of strict monotonicity when equality does not hold in (46) for some j, and $\tilde{u}_j(1)$ is not of extreme type G_j, follows by a minor variation of the above argument which we omit. ■

Let us compute the limiting behaviour of $Q(t)$ as $t \to \infty$, under the assumption that the μ_j are compactly supported. After a change of variables we have

$$Q(t) = t^{\sum_{j=1}^m p_j \dim(H_j)/2} \int_H \prod_{j=1}^m \tilde{u}_j(t, t^{1/2} B_j x)^{p_j} \, dx$$

and then after applying (49) we can write this as

$$\int_H \prod_{j=1}^m \Big[t^{\dim(H_j)/2} \det_{H_j}(G_j^{-1} + (t-1)A_j^{-1})^{-1/2}$$
$$\int_{H_j} \exp(-\pi\langle (G_j^{-1}+(t-1)A_j^{-1})^{-1}(t^{1/2}B_jx - y), (t^{1/2}B_jx-y)\rangle_{H_j})\ d\mu_j(y)\Big]^{p_j}\ dx.$$

Via another change of variables we may rewrite this as

$$\int_H \prod_{j=1}^m \Big[\det_{H_j}(t^{-1}G_j^{-1} + (1-t^{-1})A_j^{-1})^{-1/2}$$
$$\int_{H_j} \exp(-\pi\langle (t^{-1}G_j^{-1}+(1-t^{-1})A_j^{-1})^{-1}(B_jx - t^{-1/2}y), (B_jx-t^{-1/2}y)\rangle_{H_j})\ d\mu_j(y)\Big]^{p_j}\ dx.$$

Taking limits as $t \to \infty$ using dominated convergence (which can be justified since μ_j is compactly supported) we conclude

$$\lim_{t\to\infty} Q(t) = \int_H \prod_{j=1}^m \left[\det_{H_j}(A_j^{-1})^{-1/2} \int_{H_j} \exp(-\pi\langle A_jB_jx, B_jx\rangle_{H_j})\ d\mu_j(y)\right]^{p_j}\ dx,$$

which simplifies using (3) to

$$\begin{aligned}\lim_{t\to\infty} Q(t) &= \frac{\prod_{j=1}^m \det_{H_j}(A_j)^{p_j/2}}{\det_H(M)^{1/2}} \prod_{j=1}^m \mu_j(H_j)^{p_j} \\ &= \frac{\prod_{j=1}^m \det_{H_j}(A_j)^{p_j/2}}{\det_H(M)^{1/2}} \prod_{j=1}^m \left(\int_{H_j} \tilde u_j(1)\right)^{p_j}.\end{aligned}$$

In particular we have

$$Q(1) \le \frac{\prod_{j=1}^m \det_{H_j}(A_j)^{p_j/2}}{\det_H(M)^{1/2}} \prod_{j=1}^m \left(\int_{H_j} \tilde u_j(1)\right)^{p_j}.$$

We thus conclude (using a Fatou lemma argument to eliminate the hypothesis that μ_j is compactly supported)

Corollary 8.10 (Towards a generalised Brascamp–Lieb inequality)**.** *Let $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ be a generalised Brascamp–Lieb datum with all the B_j surjective, and let $\mathbf{A} = (A_j)_{1\le j\le m}$ be a gaussian input with $\mathbf{A} \le \mathbf{G}$, such that the transformation $M : H \to H$ defined by $M := \sum_{j=1}^m p_j B_j^* A_j B_j$ is invertible and obeys the inequality* (46) *and the identity* (47) *for all $1 \le j \le m$. Then*

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}).$$

One special case of this Corollary is when $\mathbf{G} = \mathbf{A}$, as in this case the condition (47) is automatic. Another special case is the limiting case $\mathbf{G} \to +\infty$, which recovers portions of Proposition 3.6. For instance:

Corollary 8.11 (Lieb's theorem generalised, gaussian-extremisable case)**.** *Let* $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ *be a generalised Brascamp–Lieb datum which is non-degenerate and gaussian-extremisable. Then* $\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G})$*. In particular* $\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G})$ *is finite in this case.*

Proof We may assume that $p_j > 0$ for all j since otherwise we can simply omit any exponents for which $p_j = 0$. Let $A_1, \ldots, A_m > 0$ be an extremiser to (42). Since $\mathbf{B}$ is non-degenerate, we have $\bigcap_{j=1}^{\infty} \ker(B_j) = \{0\}$. In particular if we set $M : H \to H$ to be the transformation $M := \sum_{j=1}^{m} p_j B_j^* A_j B_j$ then M is positive definite.

Taking logarithms in (42), we see that $\mathbf{A}$ is a local maximiser for the quantity

$$(\sum_{j=1}^{m} p_j \log \det_{H_j} A_j) - \log \det_H \sum_{j=1}^{m} p_j B_j^* A_j B_j$$

subject to the constraint $\mathbf{A} \leq \mathbf{G}$. Now let us fix a j. Let $V_j \subseteq H$ denote the kernel of the positive definite operator $G_j - A_j$, let $\iota_j : V_j \to H_j$ be the inclusion map, and let $Q_j : H_j \to H_j$ be an arbitrary self-adjoint transformation which is negative definite on V_j (i.e. $\iota_j^* Q_j \iota_j \leq 0$). Then $0 \leq A_j + \varepsilon Q_j \leq G_j$ for all sufficiently small $\varepsilon > 0$, and hence

$$\frac{d}{d\varepsilon} p_j \log \det_{H_j}(A_j + \varepsilon Q_j) - \log \det_H (M + \varepsilon p_j B_j^* Q_j B_j)|_{\varepsilon=0} \geq 0.$$

Arguing as in Proposition 3.6 we then conclude that

$$\mathrm{tr}_{H_j}((A_j^{-1} - B_j M^{-1} B_j^*) Q_j) \geq 0 \text{ whenever } Q_j \text{ is negative definite on } V_j;$$

In particular, this trace is positive whenever Q_j is negative definite on H_j, which implies (46), Also, by considering both Q_j and $-Q_j$ we see that

$$\mathrm{tr}_{H_j}((A_j^{-1} - B_j M^{-1} B_j^*) Q_j) = 0 \text{ whenever } \iota_j^* Q_j \iota_j = 0, Q_j = Q_j^*$$

which by negative definiteness of $A_j^{-1} - B_j M^{-1} B_j^*$ implies that

$$A_j^{-1} - B_j M^{-1} B_j^* = \iota_j N_j \iota_j^*$$

for some transformation $N_j : V_j \to V_j$. In particular, since $G_j - A_j$ vanishes on V_j, so $\iota_j (G_j - A_j) = 0$ and hence by self-adjointness $\iota_j^* (G_j - A_j) = 0$. Thus we obtain (47). Applying Corollary 8.10, and the hypothesis that $\mathbf{A}$ extremises (42), we conclude

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) \leq \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}; \mathbf{A}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G})$$

and the claim follows from (43). ■

Notice that if in addition to the hypotheses of Corollary 8.11 we impose (7), then any gaussian extremiser for $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ is also one for $(\mathbf{B}, \mathbf{p})$ since (46) and invertibility of M together imply gaussian-extremisability by Proposition 3.6.

Now we remove the condition that an extremiser exists. The analysis here is in fact somewhat simpler than in the non-regularised case, because it turns out there is a large class of generalised Brascamp–Lieb data, namely the *localised* generalised Brascamp–Lieb data, which are gaussian-extremisable and are in some sense "dense" in the space of all generalised Brascamp–Lieb data. (The reason for the nomenclature "localised" will become clear below.)

Definition 8.12 (Localised data)**.** A generalised Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ is said to be *localised* if there exists $1 \leq j \leq m$ such that $p_j = 1$, $H_j = H$, and $B_j = \mathrm{id}_H$.

Lemma 8.13. *Let* $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ *be a localised generalised Brascamp–Lieb datum. Then* $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ *is gaussian-extremisable, and thus (by Corollary 8.11) we have* $\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}) < \infty$.

Proof We may assume that $p_j > 0$ for all j, since we can drop all indices for which $p_j = 0$. We can also assume that $m > 1$ since the $m = 1$ case is trivial. Without loss of generality we may assume that m is the localising index, thus $p_m = 1$, $H_m = H$, and $B_m = \mathrm{id}_H$. We can rewrite (42) as

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \sup_{\mathbf{A} \leq \mathbf{G}} \det{}_H^{-1/2}(\mathrm{id}_H + M^{1/2} A_m^{-1} M^{1/2}) \prod_{j=1}^{m-1} (\det{}_{H_j} A_j)^{p_j/2}$$

where $\mathbf{A} = (A_j)_{1 \leq j \leq m}$ and $M := \sum_{j=1}^{m-1} p_j B_j^* A_j B_j$. Observe that for fixed $A_1, \ldots,$ A_{m-1}, the quantity $\det_H^{-1/2}(\mathrm{id}_H + M^{1/2} A_m^{-1} M^{1/2})$ with $0 < A_m \leq G_m$ is maximised when $A_m = G_m$. Thus we have

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \sup_{0 < A_j \leq G_j, (1 \leq j \leq m-1)} \det{}_H^{-1/2}(\mathrm{id}_H + M^{1/2} G_m M^{1/2}) \prod_{j=1}^{m-1} (\det{}_{H_j} A_j)^{p_j/2}.$$

Observe that this expression extends continuously to semi-definite A_j, and thus

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \sup_{0 \leq A_j \leq G_j, (1 \leq j \leq m-1)} \det{}_H^{-1/2}(\mathrm{id}_H + M^{1/2} G_m M^{1/2}) \prod_{j=1}^{m-1} (\det{}_{H_j} A_j)^{p_j/2}.$$

Thus an extremiser exists in the range $0 \leq A_j \leq G_j$. But the expression in the supremum vanishes when $\det(A_j) = 0$ for some j, and hence at the extremum we have $A_j > 0$ for all $1 \leq j \leq m-1$. The claim follows. ∎

Lemma 8.14 (Approximation by localised data)**.** *Let* $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ *be a generalised Brascamp–Lieb datum. Let* $(\mathbf{B}, \mathrm{id}_H)$ *be the* $(m+1)$*-transformation formed by appending the identity operator* $B_{m+1} = \mathrm{id}_H : H \to H$ *to the* m*-transformation* $\mathbf{B}$ *(thus* $H_{m+1} = H$*), and for any real number* $\lambda > 0$*, let* $(\mathbf{G}, \lambda \, \mathrm{id}_H)$ *be the* $(m+1)$*-type formed by appending the dilation operator* $\lambda \, \mathrm{id}_H : H \to H$ *to the* m*-type* $\mathbf{G}$*. Then*

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \lim_{\lambda \to 0} \lambda^{-\dim(H)/2} \, \mathrm{BL}((\mathbf{B}, \mathrm{id}_H), (\mathbf{p}, 1), (\mathbf{G}, \lambda \, \mathrm{id}_H))$$

and

$$\mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \lim_{\lambda \to 0} \lambda^{-\dim(H)/2} \, \mathrm{BL}_{\mathbf{g}}((\mathbf{B}, \mathrm{id}_H), (\mathbf{p}, 1), (\mathbf{G}, \lambda \, \mathrm{id}_H)).$$

Proof We begin with the first equality. Let $\mathbf{f} = (f_j)_{1 \leq j \leq m}$ be a normalised input of type $\mathbf{G}$. Then from (41) and (3) we see that

$$\int_H \exp(-\lambda \|x\|_H^2) \prod_{j=1}^m (f_j \circ B_j)^{p_j}(x) \, dx \leq \lambda^{-\dim(H)/2} \, \mathrm{BL}((\mathbf{B}, \mathrm{id}_H), (\mathbf{p}, 1), (\mathbf{G}, \lambda \, \mathrm{id}_H))$$

for all $\lambda > 0$. Letting $\lambda \to 0$ and using monotone convergence and (41) we conclude

$$\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) \leq \liminf_{\lambda \to 0} \lambda^{-\dim(H)/2} \mathrm{BL}((\mathbf{B}, \mathrm{id}_H), (\mathbf{p}, 1), (\mathbf{G}, \lambda \,\mathrm{id}_H)). \tag{50}$$

Conversely, let $\lambda > 0$ and $f_{m+1} : H \to \mathbf{R}^+$ be a function of type $\lambda \,\mathrm{id}_H$, then we can write

$$f_{m+1}(x) = \lambda^{\dim(H)/2} \int_H \exp(-\lambda \|x - y\|_H^2) \; d\mu(y).$$

By the Fubini-Tonelli theorem and (3) we have $\int_H f_{m+1} = \mu(H)$. Now if $(f_j)_{1 \leq j \leq m}$ is of type $\mathbf{G}$ and we set $B_{m+1} = \mathrm{id}_H$ and $p_{m+1} = 1$, then by the Fubini-Tonelli theorem again

$$\begin{aligned}
&\int_H \prod_{j=1}^{m+1} (f_j \circ B_j)^{p_j}(x) \; dx \\
&\qquad = \lambda^{\dim(H)/2} \int_H \int_H \exp(-\lambda \|x - y\|_H^2) \; d\mu(y) \prod_{j=1}^{m} (f_j \circ B_j)^{p_j}(x) \; dx \\
&\qquad \leq \lambda^{\dim(H)/2} \int_H \Big(\int_H \prod_{j=1}^{m} (f_j \circ B_j)^{p_j}(x) \; dx \Big) \; d\mu(y) \\
&\qquad \leq \lambda^{\dim(H)/2} \mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) \prod_{j=1}^{m} \Big(\int_{H_j} f_j \Big)^{p_j} \; \mu(H) \\
&\qquad = \lambda^{\dim(H)/2} \mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) \prod_{j=1}^{m+1} \Big(\int_{H_j} f_j \Big)^{p_j}.
\end{aligned}$$

We thus conclude that

$$\lambda^{-\dim(H)/2} \mathrm{BL}((\mathbf{B}, \mathrm{id}_H), (\mathbf{p}, 1), (\mathbf{G}, \lambda \,\mathrm{id}_H)) \leq \mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}).$$

Combining this with (50) we obtain the first equality of the lemma. The second equality is proven in exactly the same way but with all the f_j (and f_{m+1})) constrained to be centred gaussians. (The fact that $\exp(-\lambda \|x - y\|_H^2)$ is not centred is of no consequence as we are simply bounding it by 1). ■

Observe that the data $((\mathbf{B}, \mathrm{id}_H), (\mathbf{p}, 1), (\mathbf{G}, \lambda \,\mathrm{id}_H))$ is manifestly localised. Thus if we combine Lemma 8.14 with Lemma 8.13, we immediately obtain

Corollary 8.15 (Lieb's theorem generalised)**.** *Let* $(\mathbf{B}, \mathbf{p}, \mathbf{G})$ *be a generalised Brascamp–Lieb datum. Then* $\mathrm{BL}(\mathbf{B}, \mathbf{p}, \mathbf{G}) = \mathrm{BL}_{\mathbf{g}}(\mathbf{B}, \mathbf{p}, \mathbf{G})$.

Combining this corollary with (44) we obtain another proof of Lieb's theorem (Theorem 1.9). Another application – moving now to the setting of localised data – is the following result, which can also be found implicitly in [L].

Corollary 8.16 (Lieb's theorem localised)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum, and let* $G : H \to H$ *be positive definite. Then the best constant* $0 < K \leq \infty$ *in the* *estimate*

$$\int_H \exp(-\pi\langle Gx, x\rangle_H) \prod_{j=1}^m (f_j \circ B_j(x))^{p_j}\ dx \leq K \prod_{j=1}^m \int_H f_j \tag{51}$$

is given by

$$K := \sup_{A_1,\dots,A_m>0} \left(\frac{\prod_{j=1}^m (\det_{H_j} A_j)^{p_j}}{\det_H(G + \sum_{j=1}^m p_j B_j^* A_j B_j)} \right)^{1/2}.$$

Proof By testing (51) on centred gaussians we certainly see that we cannot replace K by any better constant. Thus it will suffice to prove (51) with the specified value of K.

Let $\lambda > 0$ be a large parameter. An inspection of the proof of Lemma 8.14 shows that if each f_j is of type $\lambda\,\mathrm{id}_{H_j}$, then we have

$$\int_H \exp(-\pi\langle Gx, x\rangle_H) \prod_{j=1}^m (f_j \circ B_j(x))^{p_j}\ dx \leq K_\lambda \prod_{j=1}^m \int_H f_j$$

where

$$K_\lambda := \sup_{0<A_j\leq\lambda\,\mathrm{id}_{H_j}} \left(\frac{\prod_{j=1}^m (\det_{H_j} A_j)^{p_j}}{\det_H(G + \sum_{j=1}^m p_j B_j^* A_j B_j)} \right)^{1/2}.$$

Taking limits as $\lambda \to \infty$, we obtain the claim. ■

We can also obtain a localised version of Theorem 1.15.

Theorem 8.17 (Localised necessary and sufficient conditions)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be a Brascamp–Lieb datum with* $\mathbf{B}$ *non-degenerate, and let* $G : H \to H$ *be positive definite. Then the estimate* (51) *holds for some finite constant* K *if and only if the inequalities* (34) *hold for all subspaces* $0 \subseteq V \subseteq H$.

Proof Recall that (34) asserts that $\dim(H/V) \geq \sum_{j=1}^m p_j \dim(H_j/(B_j V))$. We can discard those j for which $p_j = 0$ or $H_j = \{0\}$ as they have no impact on either claim in the theorem. By a linear transformation we may take $G = \mathrm{id}_H$.

The necessity of the conditions (34) can be seen by testing (51) for functions f_j which lie in an ε-neighbourhood of the unit ball in $B_j V$ and letting $\varepsilon \to 0$; we omit the easy computations. Now suppose conversely that (34) holds. Applying Corollary 8.16, we reduce to showing that

$$\prod_{j=1}^m (\det_{H_j} A_j)^{p_j} \leq K^2 \det_H(\mathrm{id}_H + M)$$

for all gaussian inputs $\mathbf{A} = (A_j)_{1\leq j\leq m}$ for some finite constant K, where $M := \sum_{j=1}^m p_j B_j^* A_j B_j$.

We now repeat the proof of Proposition 5.2. By choosing an appropriate orthonormal basis $e_1, \dots, e_n \in H$ we may assume that $M = \mathrm{diag}(\lambda_1, \dots, \lambda_n)$ for some

$\lambda_1 \geq \ldots \geq \lambda_n > 0$. Our task is thus to show that

$$\prod_{j=1}^m (\det_{H_j} A_j)^{p_j} = O(\prod_{i=1}^n (1+\lambda_i)).$$

Here we allow the constants in $O(\ )$ to depend on the Brascamp–Lieb data. Applying Lemma 5.1, we can find $I_j \subseteq \{1, \ldots, n\}$ for each $1 \leq j \leq m$ of cardinality $|I_j| = \dim(H_j)$ obeying (35) and (36). From the arguments in Proposition 5.2 we have

$$\prod_{j=1}^m (\det A_j)^{p_j} \leq C \prod_{i=1}^n \lambda_i^{\sum_{j=1}^m p_j |I_j \cap \{i\}|}.$$

Writing $\mu_i := 1 + \lambda_i$ and $c_i := \sum_{j=1}^m p_j |I_j \cap \{i\}|$, it then suffices to show that

$$\prod_{i=1}^n \mu_i^{c_i} \leq \prod_{i=1}^n \mu_i.$$

We can telescope the right-hand side as

$$(\prod_{i=1}^n \mu_i) \prod_{k=1}^n (\frac{\mu_{k+1}}{\mu_k})^{k-c_1-\ldots-c_k}$$

where we adopt the convention $\mu_{n+1} := 1$. But from (35) we have $k - c_1 - \ldots - c_k \geq 0$ for all k, and from construction we have $\frac{\mu_{k+1}}{\mu_k} \leq 1$ for all k. The claim follows. ■

9. Uniqueness of extremals

In Theorem 7.13 we gave necessary and sufficient conditions for gaussian extremisers or extremisers to exist. In this section we address the issue of whether these extremisers are unique.

One trivial source of non-uniqueness is if $p_j = 0$ for some j, then f_j can clearly be arbitrary. However one can simply omit these indices j to eliminate this source of non-uniqueness, and thus we shall only consider the case where $p_j > 0$ for all j.

From Theorem 7.13 (and Lemma 3.3) it now suffices to consider the case of geometric Brascamp–Lieb data $(\mathbf{B}, \mathbf{p})$, which has the obvious gaussian extremisers $(\lambda \, \mathrm{id}_{H_j})_{1 \leq j \leq m}$ for any $\lambda > 0$, and similarly the obvious extremisers $(c_j \exp(-\lambda \|x - B_j x_0\|_{H_j}^2))_{1 \leq j \leq m}$ for any $\lambda, c_1, \ldots, c_m > 0$ and $x_0 \in H$; compare with Lemma 6.3 and Lemma 6.4. The natural question is whether any other extremisers exist.

In the gaussian case, the answer is provided by the following proposition, which can be viewed as a continuation of Proposition 3.6 in the geometric case.

Proposition 9.1 (Characterisation of gaussian extremisers)**.** *Let $(\mathbf{B}, \mathbf{p})$ be a geometric Brascamp–Lieb datum with $p_j > 0$ for all $1 \leq j \leq m$. Let $\mathbf{A} = (A_j)_{1 \leq j \leq m}$ be a gaussian input, and let $M : H \to H$ be a positive definite transformation. Then the following seven statements are logically equivalent.*

(a) *$\mathbf{A}$ is a gaussian extremiser, and $M = \sum_{j=1}^m p_j B_j^* A_j B_j$.*

(b) $M = \sum_{j=1}^m p_j B_j^* A_j B_j$, *and* $A_j^{-1} = B_j M^{-1} B_j^*$ *for all* $1 \leq j \leq m$.
(c) $B_j M = A_j B_j$ *and* $A_j^{-1} = B_j M^{-1} B_j^*$ *for all* $1 \leq j \leq m$.
(d) $M B_j^* = B_j^* A_j$ *and* $A_j^{-1} = B_j M^{-1} B_j^*$ *for all* $1 \leq j \leq m$.
(e) M *leaves the space* $B_j^* H_j$ *invariant for all* $1 \leq j \leq m$, *and* $A_j^{-1} = B_j M^{-1} B_j^*$ *for all* $1 \leq j \leq m$.
(f) M *leaves the space* $\ker(B_j)$ *invariant for all* $1 \leq j \leq m$, *and* $A_j^{-1} = B_j M^{-1} B_j^*$ *for all* $1 \leq j \leq m$.
(g) *Each proper eigenspace of* M *is a critical subspace for* $(\mathbf{B}, \mathbf{p})$, *and* $A_j^{-1} = B_j M^{-1} B_j^*$ *for all* $1 \leq j \leq m$.

We remark that in (g), we use the term eigenspace to denote a *maximal* subspace consisting of eigenvectors of M with a common eigenvalue.

Proof For purposes of visualisation, the reader may wish to follow the arguments below in the case when the H_j are subspaces of H and B_j is an orthogonal projection, so that B_j^* is simply the inclusion map from H to H_j. (The general case is in fact equivalent to this case after some Euclidean isomorphisms.) However for notational reasons it is slightly more convenient for us to keep H_j and H separate from each other.

The equivalence of (a) and (b) follows from Proposition 3.6. Now to see that (c) $\Longrightarrow$ (b), we observe that if $B_j M = A_j B_j$ then $\sum_{j=1}^m p_j B_j^* A_j B_j = \sum_{j=1}^m p_j B_j^* B_j M = M$ by (9).

From duality we see that (c) and (d) are equivalent. Next we prove that (b) $\Longrightarrow$ (d). Using the hypothesis $M = \sum_{j=1}^m p_j B_j^* A_j B_j$ and (9), we compute

$$\begin{aligned}
&\sum_{j=1}^m p_j \operatorname{tr}_H(M^{-1}(MB_j^* - B_j^* A_j)(MB_j^* - B_j^* A_j)^*) \\
&\quad = \operatorname{tr}_H(\sum_{j=1}^m p_j M^{-1}(MB_j^* - B_j^* A_j)(B_j M - A_j B_j)) \\
&\quad = \operatorname{tr}_H(\sum_{j=1}^m p_j B_j^* B_j M - p_j M^{-1} B_j^* A_j B_j M - p_j B_j^* A_j B_j + p_j M^{-1} B_j^* A_j^2 B_j) \\
&\quad = \operatorname{tr}_H(M - M - M + M^{-1} \sum_{j=1}^m p_j B_j^* A_j^2 B_j) \\
&\quad = -\sum_{j=1}^m p_j \operatorname{tr}_H(B_j^* A_j B_j) + \sum_{j=1}^m p_j \operatorname{tr}_H(M^{-1} B_j^* A_j^2 B_j) \\
&\quad = -\sum_{j=1}^m p_j \operatorname{tr}_{H_j}(A_j B_j B_j^*) + \sum_{j=1}^m p_j \operatorname{tr}_{H_j}(B_j M^{-1} B_j^* A_j^2).
\end{aligned}$$

Since $B_j B_j^* = \mathrm{id}_{H_j}$ and $B_j M^{-1} B_j^* = A_j^{-1}$, we conclude

$$\begin{aligned}\sum_{j=1}^m p_j \operatorname{tr}_H(M^{-1}(MB_j^* - B_j^* A_j)(MB_j^* - B_j^* A_j)^*) \\ = -\sum_{j=1}^m p_j \operatorname{tr}_{H_j}(A_j) + \sum_{j=1}^m p_j \operatorname{tr}_{H_j}(A_j) \\ = 0.\end{aligned}$$

On the other hand, M^{-1} and $(MB_j^* - B_j^* A_j)(MB_j^* - B_j^* A_j)^*$ are positive semi-definite operators, and thus their product has non-negative trace. Since $p_j > 0$, we conclude that

$$\operatorname{tr}_H(M^{-1}(MB_j^* - B_j^* A_j)(MB_j^* - B_j^* A_j)^*) = 0 \text{ for all } j.$$

Since M^{-1} is positive definite, we conclude that

$$(MB_j^* - B_j^* A_j)(MB_j^* - B_j^* A_j)^* = 0 \text{ for all } j,$$

and hence $MB_j^* - B_j^* A_j = 0$. This gives (d).

The implication (d) $\implies$ (e) is immediate, since $MB_j^* H_j = B_j^* A_j H_j = B_j^* H_j$. Now we show that (e) $\implies$ (c); thus suppose that M leaves $B_j^* H_j$ invariant. Since B_j^* is an isometry, we have $(B_j^* H_j)^\perp = \ker(B_j)$; since M is self-adjoint we conclude that M also leaves $\ker(B_j)$ invariant. Since M is invertible, we see that M^{-1} will then also leave the spaces $B_j^* H_j$ and $\ker(B_j)$ invariant. From $A_j^{-1} = B_j M^{-1} B_j^*$ we conclude that

$$(A_j B_j - B_j M)(M^{-1} B_j^*) = A_j B_j M^{-1} B_j^* - B_j B_j^* = A_j A_j^{-1} - \mathrm{id}_{H_j} = 0$$

and hence $A_j B_j - B_j M$ vanishes on $M^{-1} B_j^* H_j = B_j^* H_j$. Also, since M^{-1} preserves $\ker(B_j)$ we see that $B_j M$ and $A_j B_j$ both annihilate $\ker(B_j)$. Thus $A_j B_j - B_j M$ vanishes identically, and we conclude (c).

By duality we see that (e) and (f) are equivalent. Now we show that (e) and (f) together imply (g). Let V be a proper eigenspace of M. Let $W_j = B_j^* H_j$. Since M leaves W_j invariant and V is maximal, we have $W_j = (W_j \cap V) \oplus (W_j \cap V^\perp)$. Similarly we have $W_j^\perp = (W_j^\perp \cap V) \oplus (W_j^\perp \cap V^\perp)$. Thus $H = W_j \oplus W_j^\perp = (W_j \cap V) \oplus (W_j \cap V^\perp) \oplus (W_j^\perp \cap V) \oplus (W_j^\perp \cap V^\perp)$. Now let $\Pi : H \to H$ be the orthogonal projection onto V; $B_j^* B_j : H \to H$ is the orthogonal projection onto W_j. Then Π and $B_j^* B_j$ have a common orthonormal set of eigenvectors and hence they commute. Consequently, $B_j^* B_j \Pi$ is the orthogonal projection onto $B_j^* B_j V$, and $\dim(B_j^* B_j V) = \operatorname{tr}_H(B_j^* B_j \Pi)$. Now B_j^*, being an isometry, is injective. So $\dim(B_j V) = \dim(B_j^* B_j V) = \operatorname{tr}_H(B_j^* B_j \Pi)$. Multiplying this by p_j and summing, we see from (9) That

$$\sum_{j=1}^m p_j \dim(B_j V) = \operatorname{tr}_H(\sum_{j=1}^m p_j B_j^* B_j \Pi) = \operatorname{tr}_H(\Pi) = \dim(V).$$

Thus V is a critical subspace as desired.

Finally, we show that (g) implies (e) and (f). Decomposing M as a direct sum of scalar dilation maps on the eigenspaces, we see that it suffices to show that each eigenspace V splits as the direct sum of a subspace of $\ker(B_j)$ and a subspace of the orthogonal complement $B_j^* H_j$. But by hypothesis the eigenspace V is a critical subspace, and hence by Lemma 7.12, $(V,\ V^\perp\ )$ is a critical pair. The claim follows. ■

The equivalence of (a) and (g) provides a means to construct all the gaussian extremisers $\mathbf{A}$ of a geometric Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p})$. Namely, we first decompose $\mathbf{B}$ into the direct sum of indecomposable mutually orthogonal spaces, and then let M be the direct sum of arbitrary positive scalar dilations on these indecomposable components. One then defines the A_j by the formula $A_j^{-1} = B_j M^{-1} B_j^*$. As a corollary we obtain

Corollary 9.2 (Uniqueness of gaussian extremisers)**.** *Let $(\mathbf{B}, \mathbf{p})$ be a gaussian-extremisable Brascamp–Lieb datum with $p_j > 0$ for all j. Then the following three statements are equivalent.*

(a) *The gaussian extremisers $\mathbf{A}$ of $(\mathbf{B}, \mathbf{p})$ are unique (up to scaling $\mathbf{A} \mapsto \lambda \mathbf{A}$).*
(b) *$\mathbf{B}$ is indecomposable.*
(c) *$(\mathbf{B}, \mathbf{p})$ is simple.*

We remark that this result was obtained in the rank-one case by Barthe [Bar2].

Proof The implication (b) $\implies$ (c) follows from Theorem 7.13 (or Lemma 7.12).

Now we prove (c) $\implies$ (a). If $(\mathbf{B}, \mathbf{p})$ is simple, then by statement (g) of Proposition 9.1 there are no proper eigenspaces of M. Consequently the unique eigenspace of M must be the whole of H, i.e., M is a scalar multiple of the identity. Proposition 9.1 now gives the desired uniqueness.

Now we prove (a) $\implies$ (b) Conversely, if $\mathbf{B}$ is decomposable, then by Lemma 7.7 we have a proper critical pair (V, W), which we can choose to be orthogonal complements, from Lemma 7.12. One can now define a two-parameter family of positive-definite operators M with eigenspaces V, W, and each one generates a gaussian extremiser by Proposition 9.1. The claim follows. ■

Now we address the issue of uniqueness of extremisers which are not gaussian. Here the situation appears to be significantly more complicated, especially in the decomposable case. For instance, for the Loomis–Whitney inequality (2) the extremisers take the form

$$f(y,z) = aG(y)H(z); \quad g(x,z) = bF(x)H(z); \quad h(x,y) = cF(x)G(y)$$

for arbitrary scalars $a, b, c > 0$ and non-negative integrable functions F, G, H of one variable. Even in the indecomposable case there can be plenty of extremisers. For instance, as is well known, for Hölder's inequality (Example 1.3) one has an extremiser whenever the functions f_j are scalar multiples of each other, even in the one-dimensional case which is irreducible and thus has unique gaussian extremisers up to scaling. However, in the rank one case it is known that for any irreducible Brascamp–Lieb datum whose domain H has dimension strictly larger than one, the only extremisers are the gaussian ones; see [Bar2, Theorem 4]. Further analysis of the extremisers in the rank one case was conducted to [CLL]. We cannot fully generalise this analysis to the higher rank case. We do however have the following partial result.

Theorem 9.3 (Uniqueness of extremisers)**.** *Let* $(\mathbf{B}, \mathbf{p})$ *be an extremisable Brascamp–Lieb datum with* $0 < p_j < 1$ *for all* j*. Suppose also that the spaces* $B_j^* H_j$ *are all disjoint except at 0, thus* $B_i^* H_i \cap B_j^* H_j = \{0\}$ *whenever* $1 \leq i < j \leq m$*. Then if* $\mathbf{f} = (f_j)$ *is an extremising input, then all the* f_j *are gaussians, thus there exist real numbers* $c_j > 0$*, positive definite transformations* $A_j : H_j \to H_j$*, and points* $x_j \in H_j$ *such that* $f_j(x) = c_j \exp(-\pi \langle A_j(x - x_j), (x - x_j)\rangle_{H_j})$*.*

Remark 9.4. Remark 4.9 tells us that the permitted x_j's in the conclusion of the theorem are precisely of the form $B_j \overline{w}$ as $\overline{w}$ varies over H.

Remark 9.5. The condition $p_j < 1$ is automatic if one also assumes that $(\mathbf{B}, \mathbf{p})$ is simple, $\dim(H) > 1$, and that $\dim(H_j) > 0$ for all j, as can be seen by testing (8) on one-dimensional subspaces of H not in the kernel of B_j. In the rank one case (Example 1.7), the condition that the $B_j^* H_j$ are all disjoint amounts to the assertion that no two of the vectors $v_1, \ldots, v_m$ are scalar multiples of each other. But in the case when two of the v_i are multiples of each other they can easily be concatenated using Hölder's inequality; if the concatenated extremiser is gaussian, then by the converse Hölder inequality we see that the original extremisers are also gaussian. Using this observation we can recover the result of [Bar2, Theorem 4] that in the simple rank one case in dimensions two or greater, the only extremisers are given by gaussians. We also recover the well-known fact that for the sharp Young's inequality (Example 1.5) with p_1, p_2, p_3 strictly less than 1, the only extremisers are gaussians.

Proof Our arguments here are based on a more careful analysis of the heat flow monotonicity argument, as in [CLL]. By Theorem 7.13 it suffices to consider the case when $(\mathbf{B}, \mathbf{p})$ is geometric. In particular we already have $\mathbf{f}_0 = (\exp(-\pi\|x\|_{H_j}^2))_{1 \leq j \leq m}$ as an extremising input. From Lemma 6.3 we see that $\mathbf{f}$ is any extremising input for $(\mathbf{B}, \mathbf{p})$, then so is $(\mathbf{f} * \mathbf{f}_0)\mathbf{f}_0$. But since $\mathbf{f}_0$ is Schwartz and $\mathbf{f}$ is integrable, we see that $(\mathbf{f} * \mathbf{f}_0)\mathbf{f}_0$ is Schwartz and strictly positive (in fact it is of type id_{H_j}). Also one can easily verify (using the Fourier transform) that $(\mathbf{f} * \mathbf{f}_0)\mathbf{f}_0$ consists of gaussians if and only if $\mathbf{f}$ consists of gaussians. Thus to prove the claim, it suffices to do so for inputs $\mathbf{f}$ which are Schwartz strictly positive functions of type id_{H_j}.

We now review the proof of Proposition 2.8, which among other things proves $\mathrm{BL}(\mathbf{B}, \mathbf{p}; \mathbf{f}) \leq 1 = \mathrm{BL}(\mathbf{B}, \mathbf{p})$. But since $\mathbf{f}$ is extremising, we must have equality at every stage of the proof of this Proposition. In particular, since this proof needed (18) to be non-negative, we now see that in fact (18) needs to be zero everywhere:

$$\sum_{j=1}^{m} p_j \langle B_j^* B_j(\vec{v} - \vec{v}_j), (\vec{v} - \vec{v}_j)\rangle_H = 0.$$

Since $p_j > 0$ for all j, we conclude that

$$\|B_j(\vec{v} - \vec{v}_j)\|_H = 0 \text{ for all } j.$$

Since $\vec{v} = \sum_{i=1}^m p_i \vec{v}_i$, we thus have

$$B_j \sum_{i=1}^m p_i \vec{v}_i = B_j \vec{v}_j \text{ for all } j.$$

Since p_j is strictly less than 1, and $\vec{v}_j$ lies in the range of $B_j^* B_j$, we conclude that

$$\vec{v}_j = B_j \sum_{i \neq j} \frac{p_i}{1 - p_j} B_j^* B_j \vec{v}_i.$$

Next, observe from the chain rule that $\vec{v}_i$ takes values in $B_i^* H_i \subseteq H$, and we can write $\vec{v}_i = w_i \circ B_i$ for some function $w_i : H_i \to B_i^* H_i$, thus

$$w_j \circ B_j = B_j \sum_{i \neq j} \frac{p_i}{1 - p_j} B_j^* B_j w_i \circ B_i. \tag{52}$$

The disjointness of the spaces $B_i^* H_i$ will now force a lot of structure on these functions w_i; this is easiest to establish using Fourier analysis and the theory of distributions[3]. First observe that since $\mathbf{f}$ is Schwartz and strictly positive, it is not hard to see from the definition of $\vec{v}$ that the w_i are continuous and grow at most linearly (thus we have $\|w_i(x)\|_H \leq C_{\mathbf{f},i}(1 + \|x\|_{H_j})$ for all $x \in H_j$ for some constant $C_{\mathbf{f},i}$ depending on the input $\mathbf{f}$ and the index i). In particular w_i is a tempered distribution. We can now take distributional Fourier transforms of (52) in the Euclidean space H. The left-hand side is a tempered distribution supported in $B_j^* H_j$, while the right-hand side is supported in $\bigcup_{i \neq j} B_i^* H_i$. Thus by hypothesis, both distributions are in fact supported on $\{0\}$. Inverting the Fourier transform, we conclude in particular that $w_j \circ B_j$ (and hence w_j) is a polynomial. But since w_j has at most linear growth, we conclude that w_j is a linear polynomial. We now specialise to the limiting time case $t = 0$, in which $v_j = -\nabla \log f_j$, taking advantage of the fact that f_j is of type id_{H_j}, to conclude that $\nabla \log f_j$ is a linear polynomial. Integrating this we see that $\log f_j$ must be a quadratic polynomial; since f_j is integrable, the leading term must be strictly negative definite. The claim follows by completing the square of the quadratic polynomial. ■

10. Sliding kernels

We now recast the monotonicity formulae obtained earlier as a monotonicity property of sliding gaussians. This leads naturally to the question of whether such monotonicity also holds for other kernels than gaussians; in the linear case, we will show that this is true for log-concave kernels.

---

[3]An alternate approach, which exploits the smoothness of the w_i, is to differentiate (52) in various directions, designed to eliminate the terms on the right-hand side but not on the left.

Let us first return to Proposition 2.8. A key component of the proof of that proposition was the claim that the quantity

$$Q(t) := \int_H \prod_{j=1}^m u_j^{p_j}(t,x)\ dx$$

was monotone non-decreasing in time for $t > 0$, where

$$u_j(t,x) = \frac{1}{(4\pi t)^{\dim(H_j)/2}} \int_{H_j} e^{-\|B_j x - z\|_{H_j}^2/4t} f_j(z)\ dz$$

and $(f_j)_{1\leq j\leq m}$ was an arbitrary input. Making the change of variables $s := (4\pi t)^{-1/2}$, $y := sx$, and $v := z$, and $d\mu_j(v) = f_j(v)d\mu(v)$, we thus see that the quantity

$$\int_H \prod_{j=1}^m \left(\int_{H_j} \exp(-\pi\|B_j y - vs\|_{H_j}^2)\ d\mu_j(v) \right)^{p_j}\ dy$$

is monotone non-increasing in s for $s > 0$. We view s as a time parameter, y as a position variable, v as a velocity variable, and μ_j as a velocity distribution. Each function $\exp(-\pi\|B_j y - vs\|_{H_j}^2)$ then becomes a "sliding gaussian" which equals $\exp(-\pi\|B_j y\|_{H_j}^2)$ at time $s = 0$ and then moves with velocity v (where we embed H_j into H using the isometry B_j^*). The above monotonicity then represents the intuitively plausible fact that the multilinear L^p norm of these sliding gaussians is maximised at time $s = 0$, at which time all the gaussians are centred at the origin.

In light of Proposition 3.6, there should be an analogous monotonicity of sliding gaussians for any Brascamp–Lieb datum $(\mathbf{B}, \mathbf{p})$ for which one can locate a gaussian input $\mathbf{A}$ obeying (7) and the inequalities $B_j^* A_j B_j \leq \sum_{i=1}^m p_i B_i^* A_i B_i$ for all j. Indeed, in light of Proposition 8.9 (in the special case $\mathbf{G} = \mathbf{A}$), the scaling condition can be dropped:

Proposition 10.1. *Let $(\mathbf{B}, \mathbf{p})$ be a Brascamp–Lieb datum with all the B_j surjective, and let $\mathbf{A} = (A_j)_{1\leq j\leq m}$ be a gaussian input such that $B_j^* A_j B_j \leq \sum_{i=1}^m p_i B_i^* A_i B_i$ for all j. Then for any positive finite Borel measures $d\mu_j$ on H_j, the quantity*

$$\int_H \prod_{j=1}^m \left(\int_{H_j} \exp(-\pi\langle A_j(B_j y - vs), (B_j y - vs)\rangle_{H_j})\ d\mu_j(v) \right)^{p_j}\ dy$$

is monotone non-increasing in s for $s > 0$.

Proof An easy scaling argument shows that it suffices to prove this when $0 < s < (4\pi)^{-1/2}$. The claim then follows from Proposition 8.9 in the case $\mathbf{G} = \mathbf{A}$, after the change of variables $s := (4\pi t)^{-1/2}$, $y := sx$, and $v := z$ as in the preceding discussion. ■

Now we turn to log-concave kernels. We begin with a divergence estimate, which is a weak analogue to Corollary 8.7.

Proposition 10.2 (Divergence estimate for log-concave kernels)**.** *Let $\psi : \mathbf{R}^n \to \mathbf{R}^+$ be a smooth, strictly positive, absolutely integrable function which is log-concave* *(thus $\nabla^2\psi(x) \geq 0$ for all x). Let μ be a non-zero positive finite Borel measure, and let $\vec{y} : \mathbf{R}^n \to \mathbf{R}^n$ be the centre-of-mass vector field*

$$\vec{y}(x) := \frac{\int_{\mathbf{R}^n} \psi(x-y)\ y\ d\mu(y)}{\int_{\mathbf{R}^n} \psi(x-y)\ d\mu(y)}.$$

Then $\operatorname{div} \vec{y} \geq 0$*, with equality if and only if μ is a point mass.*

Proof From the definition of $\vec{y}$ we have

$$\int_{\mathbf{R}^n} \psi(x-y)(\vec{y}(x)-y)\ d\mu(y) = 0 \text{ for all } x \in \mathbf{R}^n. \tag{53}$$

Taking divergences of both sides and writing $\nabla\psi = \psi\nabla\log\psi$, we obtain

$$\int_{\mathbf{R}^n} \psi(x-y)[\langle \nabla \log \psi(x-y), \vec{y}(x)-y\rangle + \operatorname{div}(\vec{y})(x)]\ d\mu(y) = 0.$$

On the other hand, from (53) again we have

$$\int_{\mathbf{R}^n} \psi(x-y)\langle \nabla \log \psi(x-\vec{y}), \vec{y}(x)-y\rangle\ d\mu(y) = 0.$$

Combining these two equations we obtain

$$\operatorname{div}(\vec{y})(x) = \frac{\int_{\mathbf{R}^n} \psi(x-y)\langle \nabla \log \psi(x-\vec{y}) - \nabla \log \psi(x-y), \vec{y}-y\rangle\ d\mu(y)}{\int_{\mathbf{R}^n} \psi(x-y) d\mu(y)}.$$

Since $\nabla^2 \log\psi \geq 0$, we see from the mean-value theorem that $\langle \nabla \log \psi(x-\vec{y}) - \nabla \log \psi(x-y), \vec{y}-y\rangle$ is non-negative, and the claim follows. ■

Remark 10.3. In the gaussian case $\psi(x) = \exp(-\pi\langle Gx, x\rangle)$ the above proposition follows from the $A = G^{-1}$ case of Corollary 8.7 after some simple algebraic manipulations which we omit here.

As a consequence of this divergence estimate we obtain the following monotonicity formula.

Lemma 10.4 (L^p monotonicity for log-concave kernels)**.** *Let $\psi : \mathbf{R}^n \to \mathbf{R}^+$ be a strictly positive log-concave function which vanishes at infinity. Then for any positive finite Borel measure μ on $\mathbf{R}^n$ and any $p \geq 1$ the quantity*

$$Q(t) = \int_{\mathbf{R}^n} (\int_{\mathbf{R}^n} \psi(x-vt)\ d\mu(v))^p\ dx$$

is non-increasing in time for $t \geq 0$.

Remark 10.5. The intuition here is that the travelling waves $\psi(x-vt)$ are diverging from each other as $t > 0$ increases, and the total mass of $\int_{\mathbf{R}^n} \psi(x-vt)\ d\mu(v)$ is constant, so the higher L^p norms should decrease.

Proof We may assume of course that μ is not identically zero. By a limiting argument we can also assume that μ is compactly supported and ψ is rapidly decreasing. We set

$$u(t,x) := \int_{\mathbf{R}^n} \psi(x-vt)\ d\mu(v)$$

and

$$\vec{v}(t,x) := \frac{\int_{\mathbf{R}^n} v\psi(x-vt)\ d\mu(v)}{\int_{\mathbf{R}^n} \psi(x-vt)\ d\mu(v)}$$

(note that the denominator is strictly positive by our assumptions on $\psi, d\mu$). A simple calculation then shows that we have the transport equation

$$\partial_t u + \operatorname{div}(\vec{v}u) = 0.$$

From Proposition 10.2 we also have

$$\operatorname{div}(\vec{v}) \geq 0.$$

Also by our assumptions on ψ, μ we see that $\vec{v}u^p$ is rapidly decreasing in space. The claim now follows from Lemma 2.6 with $m = 1$, $\alpha = 0$, $u_1 := u$, and $\vec{v}_1 := \vec{v}$, and with the reversed signs. ■

Remarks 10.6. A similar argument shows that when $0 < p < 1$ the quantity $Q(t)$ is non-increasing in time. Note that Q is constant in the boundary case $p = 1$. One can easily prove a similar result when f is merely a positive finite measure rather than a Schwartz function by a standard limiting argument which we omit here. In the case when ψ is strictly log-concave, a more refined analysis in the $p \neq 1$ case also shows that Q is strictly monotone unless f is a point mass; we again omit the details.

Remark 10.7. In the gaussian case $\psi(x) = \exp(-\pi\langle Ax, x\rangle)$ this proposition is a special case of Proposition 10.1. Indeed, the above arguments can be easily modified to give a direct proof of Proposition 10.1 from Lemma 2.6, using Corollary 8.7 instead of Proposition 10.2; we omit the details.

We can apply this lemma to concrete log-concave kernels such as the one-dimensional heat kernel $\psi(x) := e^{-\pi x^2}$.

Corollary 10.8. *Let $f : \mathbf{R} \to \mathbf{R}^+$ be a positive finite measure and $p \geq 1$. Similarly, if $u : \mathbf{R}^+ \times \mathbf{R} \to \mathbf{R}^+$ denotes the heat extension*

$$u(t,x) := \frac{1}{\sqrt{4\pi t}} \int_{\mathbf{R}} e^{-|x-y|^2/4t} f(y)\ dy$$

then $t^{1/2p'} \|\phi(\cdot,t)\|_{L^p(\mathbf{R})}$ is non-decreasing in t. If f is not a point mass, then these quantities are strictly increasing in t. Here p' is the dual exponent of p, defined by $1/p + 1/p' = 1$.

This innocuous-sounding result does not appear to be previously in the literature; it can be derived by explicit computation when $p \geq 1$ is an integer but is not trivial to prove for other values of p; one can also establish these results directly from Lemma 2.6 of course. It is an interesting question to ask whether an analogous result holds for the harmonic extension $\phi(t,x)$ of u (with $t^{1/2p'}$ replaced by $t^{1/p'}$); this corresponds to setting $\psi(x) = \frac{1}{\pi(1+x^2)}$, which is not log-concave. In this regard, it is a classical result, essentially due to Hardy and Littlewood [HL] (see also [D], Theorem 5.9) that if $p \geq 1$, then $t^{1/p'}\|\phi(\cdot,t)\|_p$ is bounded as a function of s, and in fact it tends to zero as s decreases to zero. A similar claim can also be proven easily for the heat extension. It is also amusing to note that there is a dyadic analogue of these monotonicity formulae: if for each integer k we let $E_k(f)$ denote the orthogonal projection of f onto functions which are constant on dyadic intervals of length 2^k, then it is easy to verify that the quantity $2^{k/p'}\|E_k(f)\|_{L^p(\mathbf{R})}$ is non-decreasing in k.

References

[ARS] M. Auslander, I. Reiten and S. Smalo, Representation theory of Artin algebras, Cambridge University Press, 1995.

[Ball] K. M. Ball, *Volumes of sections of cubes and related problems*, Geometric Aspects of Functional Analysis, ed. by J. Lindenstrauss and V. D. Milman, Lecture Notes in Math. 1376, Springer, Heidelberg, 1989, pp. 251 –260.

[Bar] F. Barthe, *Optimal Young's inequality and its converse: a simple proof*, Geom. Func. Anal. **8** (1998), 234–242.

[Bar2] F. Barthe, *On a reverse form of the Brascamp–Lieb inequality*, Invent. Math. **134** (1998), 335–361.

[BarC] F. Barthe and D. Cordero-Erausquin, *Inverse Brascamp–Lieb inequalities along the heat equation*, in *Geometric Aspects of Functional Analysis, 2002-2003* (eds. V. D. Milman and G. Schechtman) pp. 65–71, Lecture Notes in Mathematics **1850**, Springer-Verlag, Berlin 2004.

[Be] W. Beckner, *Inequalities in Fourier analysis*, Ann. of Math. **102** (1975), 159–182.

[BCCT] J. M. Bennett, A. Carbery, M. Christ and T. Tao, *Finite bounds for Hölder–Brascamp–Lieb multilinear inequalities*, preprint.

[BCT] J. M. Bennett, A. Carbery and T. Tao, *Kakeya rearrangement inequalities*, preprint.

[Bel] P. Belkale, *Geometric proofs of the Horn and saturation conjectures*, preprint.

[BL] H. J. Brascamp and E. H. Lieb, *Best constants in Young's inequality, its converse, and its generalization to more than three functions*, Adv. Math. **20** (1976), 151–173.

[BLL] H. J. Brascamp, E. H. Lieb and J. M. Luttinger, *A general rearrangement inequality for multiple integrals*, J. Funct. Anal. **17** (1974), 227–237.

[CLL] 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** (2004), 487 –520.

[DW] H. Derksen and J. Weyman, *Quiver representations*, Notices Amer. Math. Soc. **52** (2005), no. 2, 200–206.

[D] P. L. Duren, Theory of H^p spaces, Academic Press, 1970.

[F] H. Finner, *A generalization of Hölder's inequality and some probability inequalities*, Ann. Probab. **20** (1992), no 4., 1893–1901.

[HL] G. H. Hardy and J. E. Littlewood, *Some properties of fractional integrals, II*, Math. Z. **34** (1932), 403–439.

[L] E. H. Lieb, *Gaussian kernels have only Gaussian maximizers*, Invent. Math. **102** (1990), 179–208.

[LW] L. H. Loomis and H. Whitney, *An inequality related to the isoperimetric inequality*, Bull. Amer. Math. Soc **55**, (1949). 961–962.

School of Mathematics, University of Birmingham, Birmingham, B15 2TT, UK

E-mail address: J.Bennett@bham.ac.uk

School of Mathematics, University of Edinburgh, Edinburgh, EH9 3JZ, UK

E-mail address: A.Carbery@ed.ac.uk

Department of Mathematics, University of California, Berkeley, CA 94720-3840, USA

E-mail address: mchrist@math.berkeley.edu

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu