THE EXTREMALS OF THE ALEXANDROV-FENCHEL INEQUALITY FOR CONVEX POLYTOPES

YAIR SHENFELD AND RAMON VAN HANDEL

Abstract. The Alexandrov-Fenchel inequality, a far-reaching generalization of the classical isoperimetric inequality to arbitrary mixed volumes, lies at the heart of convex geometry. The characterization of its extremal bodies is a long-standing open problem that dates back to Alexandrov's original 1937 paper. The known extremals already form a very rich family, and even the fundamental conjectures on their general structure, due to Schneider, are incomplete. In this paper, we completely settle the extremals of the Alexandrov-Fenchel inequality for convex polytopes. In particular, we show that the extremals arise from the combination of three distinct mechanisms: translation, support, and dimensionality. The characterization of these mechanisms requires the development of a diverse range of techniques that shed new light on the geometry of mixed volumes of nonsmooth convex bodies. Our main result extends further beyond polytopes in a number of ways, including to the setting of quermassintegrals of arbitrary convex bodies. As an application of our main result, we settle a question of Stanley on the extremal behavior of certain log-concave sequences that arise in the combinatorics of partially ordered sets.

Contents

2010 *Mathematics Subject Classification.* 52A39; 52A40; 52B05; 05B25.

Key words and phrases. Mixed volumes; Alexandrov-Fenchel inequality; convex polytopes; extremum problems in geometry and combinatorics.

1. Introduction

1.1. **The Alexandrov-Fenchel inequality and the extremal problem.** Let $C_1,\ldots,C_m$ be convex bodies (that is, nonempty compact convex sets) in $\mathbb{R}^n$. One of the most basic facts of convex geometry, due to Minkowski [22], is that the volume of convex bodies behaves as a homogeneous polynomial under addition $\lambda C+\mu C' := \{\lambda x+\mu y : x\in C, y\in C'\}$: that is, for all $\lambda_1,\ldots,\lambda_m\geqslant 0$

$$\mathrm{Vol}_n(\lambda_1C_1+\cdots+\lambda_mC_m) = \sum_{i_1,\ldots,i_n=1}^{m} \mathsf{V}_n(C_{i_1},\ldots,C_{i_n})\,\lambda_{i_1}\cdots\lambda_{i_n}. \tag{1.1}$$

The coefficients $\mathsf{V}_n(C_{i_1},\ldots,C_{i_n})$ of this polynomial, called *mixed volumes*, form a large family of natural geometric parameters associated to convex bodies. For example, the special cases $\mathsf{V}_n(C,\ldots,C,B,\ldots,B)$, called quermassintegrals, already capture familiar notions such as the volume, surface area, and mean width of C, and the average volume of the projections of C onto a random k-dimensional subspace.[1] In view of these and numerous other important examples, mixed volumes play a central role in convex geometry [3, 5, 24, 30].

When the convex bodies are polytopes, mixed volumes may also be viewed as belonging to combinatorial geometry. In this setting, striking connections arise between the theory of mixed volumes and other areas of mathematics. For example, in algebraic geometry, mixed volumes compute the number of solutions to systems of polynomial equations [5, §27] and intersection numbers of divisors on toric varieties [13, 9]; and in combinatorics, mixed volumes compute quantities associated to objects such as matroids, partial orders, and permanents [33, 17].

Given the central nature of mixed volumes, it is natural to expect that inequalities between mixed volumes capture important mathematical phenomena. The most fundamental result of this kind is the Alexandrov-Fenchel inequality, which expresses the fact that mixed volumes are log-concave.

Theorem 1.1 (Alexandrov-Fenchel inequality)**.** *We have*

$$\mathsf{V}_n(K,L,C_1,\ldots,C_{n-2})^2 \geqslant \mathsf{V}_n(K,K,C_1,\ldots,C_{n-2})\,\mathsf{V}_n(L,L,C_1,\ldots,C_{n-2})$$

for any convex bodies $K,L,C_1,\ldots,C_{n-2}$ in $\mathbb{R}^n$.

Theorem 1.1 was first proved by Minkowski in 1903 in dimension $n=3$ [22], and in full generality by Alexandrov in 1937 [1, 2]. (Fenchel independently announced a proof [12], but it was never published.) It lies at the heart of many applications of mixed volumes in convexity and in other areas of mathematics. This paper is concerned with a classical open problem surrounding the Alexandrov-Fenchel inequality that dates back to Alexandrov's original paper [2, p. 80].

To provide context for the problem studied in this paper, let us recall the original setting of Minkowski [22]. Minkowski viewed Theorem 1.1 as a far-reaching generalization of the isoperimetric inequality between volume and surface area, which are merely two special cases of mixed volumes. For example, the special case

$$\mathsf{V}_3(B,C,C)^2 \geqslant \mathsf{V}_3(C,C,C)\,\mathsf{V}_3(B,B,C)$$

states that the surface area of a three-dimensional convex body C is lower bounded by the product of its volume and mean width, a kind of isoperimetric inequality

[1] Throughout this paper B denotes the Euclidean unit ball in $\mathbb{R}^n$.

involving three geometric parameters. From this viewpoint, a complete understanding of Theorem 1.1 should capture not only the inequality but also the associated extremum problem: which bodies minimize surface area when the volume and mean width are fixed? This question is equivalent to the study of the cases of equality in the above inequality. Remarkably, it turns out that the extremals in this example possess highly unusual properties: they consist of cap bodies ("spiky balls") which are both non-unique and non-smooth, in sharp contrast with the situation in the classical isoperimetric problem (cf. [32] and the references therein).

The above example suggests that the extremum problems associated to more general cases of the Alexandrov-Fenchel inequality are likely to possess a rich and intricate structure. The problem of characterizing these extremals was raised in the original papers of Minkowski [22] and Alexandrov [1], but progress toward the resolution of this problem has proved to be elusive. None of the known proofs of the Alexandrov-Fenchel inequality provides information on its cases of equality. The geometric proofs (cf. [2, 31]) impose restrictions, such as smooth bodies or polytopes with identical face directions, under which only trivial extremals arise, and deduce the general result by approximation; nontrivial extremals arise only in the limit, and are thus invisible in the proofs of the inequality. The algebraic proofs (cf. [5, 13]) perform a reduction to a certain (non-toric) algebraic surface, which causes the convex geometric structure of the problem to be lost.

It was long believed that the extremals of the Alexandrov-Fenchel inequality are too numerous to admit a meaningful geometric characterization, cf. [5, §20.5] or [11, p. 248]. However, detailed conjectures on the structure of the extremals (attributed in part to Loritz) were published in 1985 by Schneider [26], breathing new life into the problem. Schneider's conjectures need not hold when some of the bodies have empty interior [8], and no conjectures have been formulated to date about this setting (which, as we will see, is of special importance in applications). However, the validity of Schneider's conjectures for full-dimensional bodies has remained open, except in a few special cases that are reviewed in [30, §7.6], [24]. Very recently, significant new progress was made in [32], which enabled the proof of Schneider's conjectures in the case that dates back to Minkowski [22]. The general case is however much richer, and entirely new ideas are needed.

1.2. **Main results.** In this paper, we completely settle the extremal problem *in the combinatorial setting.* Our main result characterizes all equality cases

$$\mathsf{V}_n(K, L, P_1, \ldots, P_{n-2})^2 = \mathsf{V}_n(K, K, P_1, \ldots, P_{n-2})\,\mathsf{V}_n(L, L, P_1, \ldots, P_{n-2})$$

when $P_1, \ldots, P_{n-2}$ are arbitrary convex polytopes in $\mathbb{R}^n$ and K, L are convex bodies. The characterization of the extremal bodies is described in section 2. In particular, we will show that the extremals of the Alexandrov-Fenchel inequality arise from the combination of three distinct mechanisms: translation, support, and dimensionality. The first two mechanisms were anticipated by Schneider's conjectures, while the third is responsible for the new extremals that arise when the polytopes P_i may have empty interior. The proof of our main result (Theorem 2.13), which is contained in sections 5–12, will in fact give considerably more precise information on the structure of the extremals than is provided by the characterization in section 2; the most detailed form of our main result will be formulated in section 13.

Aside from its intrinsic place in the foundation of convex geometry, the problem of characterizing the extremals of the Alexandrov-Fenchel inequality may be thought of in a broader context: the limited progress on this problem to date stems from major gaps in the understanding of the geometry of mixed volumes of non-smooth convex bodies. The fundamental issues that arise are both of a combinatorial and of an analytic nature, as we will explain presently.

As will become clear in section 2, the extremals of the Alexandrov-Fenchel inequality are controlled by the boundary structure of the bodies $C_1, \ldots, C_{n-2}$ in Theorem 1.1. In the case that was settled in [32], only the boundary structure of a single body plays a role. In general, however, each of the bodies $C_1, \ldots, C_{n-2}$ has an arbitrary boundary structure, and the interactions between the different bodies conspire to give rise to the extremals. This interaction already arises in its full complexity in the combinatorial setting considered in this paper. In settling the problem, we develop a theory that explains these interactions: this includes, among other ingredients, a local Alexandrov-Fenchel inequality for mixed area measures, strong gluing principles for projections from limited data, and new geometric structures ("propellers") of mixed area measures of bodies with empty interior. An overview of the proof of our main result will be given in section 4.

The main contribution of this paper is the complete solution of these combinatorial aspects of the problem. In contrast, the obstacle to going beyond polytopes stems from unresolved analytic problems in the theory of mixed volumes, which are largely independent of the problems studied in this paper. These analytic problems arise because the boundary of a general convex body may be almost arbitrarily irregular (for example, consider the convex hull of an arbitrary closed subset of the unit sphere), so that mixed volumes of general convex bodies give rise to analytic objects that live on highly irregular sets. The treatment of general bodies therefore requires the development of an appropriate functional-analytic framework, which has only been partially accomplished to date [32] (see section 16 for discussion). The main ideas of this paper are not specific to polytopes, however, and may be expected to apply more generally when placed in an suitable analytic framework.

1.3. **Extensions and applications.** While this paper is primarily concerned with the combinatorial setting, our methods already admit a number of extensions beyond the setting of convex polytopes. In particular, we will show in section 14 that our main result extends to the setting where the convex bodies $C_1, \ldots, C_{n-2}$ in Theorem 1.1 are a combination of polytopes, zonoids, and smooth bodies. By combining the present methods with [32], we will also fully characterize the extremals of the Alexandrov-Fenchel inequality for quermassintegrals of arbitrary convex bodies, a special case that arises frequently in applications.

Section 15 develops an application in combinatorics. It was noticed long ago that various combinatorially defined sequences (N_i) appear to be log-concave, that is, they satisfy $N_i^2 \geqslant N_{i-1}N_{i+1}$. Such phenomena have received much attention in recent years [17]. One of the earliest advances in this area is due to Stanley [33], who observed that if one can represent the relevant combinatorial quantities in terms of mixed volumes, log-concavity is explained by the Alexandrov-Fenchel inequality. Stanley further raises the following question: in cases where (N_i) is log-concave, can one characterize the associated extremum problem, that is, explain what combinatorial objects achieve equality $N_i^2 = N_{i-1}N_{i+1}$? As an illustration of our main result, we will settle this problem in one of the settings considered by Stanley, where N_i is the number of linear extensions of a partially ordered set for which a distinguished element has rank i. Such extremal problems appear to be inaccessible by currently known methods of enumerative or algebraic combinatorics. This example highlights the significance of the questions considered in this paper to extremal problems in other areas of mathematics, and hints at the possibility that the structures developed here might have analogues outside convexity; a brief discussion of algebraic analogues of our results is given in section 16.

Let us note that, far from being esoteric, it is precisely the case of convex bodies with empty interior (which is not covered by previous conjectures) that arises in combinatorial applications [33]. This reinforces the importance of a complete characterization of the extremals, whose formulation we turn to presently.

2. Three extremal mechanisms

The aim of this section is to formulate and explain the main result of this paper. We first recall some key facts on mixed volumes and mixed area measures. We will subsequently describe three distinct mechanisms that give rise to extremals of the Alexandrov-Fenchel inequality, and state our main result. Here and throughout the paper, our standard reference on convexity is the monograph [30].

2.1. **Basic facts.**

2.1.1. **Convex bodies, mixed volumes, mixed area measures.** Fix $n \geqslant 3$. A *convex body* is a nonempty compact convex set in $\mathbb{R}^n$. A (convex) *polytope* is the convex hull of a finite number of points.

To each convex body K, we associate its *support function*

$$h_K(u) := \sup_{y \in K} \langle y, u \rangle.$$

We think of h_K either as a function on S^{n-1} or as a 1-homogeneous function on $\mathbb{R}^n$. Geometrically, if $u \in S^{n-1}$, then $h_K(u)$ is the (signed) distance to the origin of the supporting hyperplane of K with outer normal u; thus $h_K : S^{n-1} \to \mathbb{R}$ uniquely determines K, as any convex body is the intersection of its supporting halfspaces. The key property of support functions is that they behave naturally under addition, that is, $h_{\lambda K + \mu L} = \lambda h_K + \mu h_L$ for any bodies K, L and $\lambda, \mu \geqslant 0$.

The *mixed volume* $\mathsf{V}_n(C_1, \ldots, C_n)$ of n convex bodies $C_1, \ldots, C_n$ in $\mathbb{R}^n$ is defined by (1.1). Mixed volumes are nonnegative, and are symmetric and multilinear in their arguments. Moreover, there exists a nonnegative measure $S_{C_1,\ldots,C_{n-1}}$ on S^{n-1}, called the *mixed area measure* of $C_1, \ldots, C_{n-1}$, such that

$$\mathsf{V}_n(K, C_1, \ldots, C_{n-1}) = \frac{1}{n} \int h_K(u)\, S_{C_1,\ldots,C_{n-1}}(du). \tag{2.1}$$

Like mixed volume, $S_{C_1,\ldots,C_{n-1}}$ is symmetric and multilinear in $C_1, \ldots, C_{n-1}$.

Consider a function $f = h_K - h_L$ that is a difference of support functions. As mixed volumes and mixed area measures are multilinear as functions of the underlying bodies (and hence of their support functions), we may uniquely extend their definitions to differences of support functions [30, §5.2]. That is, we will write

$$\mathsf{V}_n(f, C_1, \ldots, C_{n-1}) := \mathsf{V}_n(K, C_1, \ldots, C_{n-1}) - \mathsf{V}_n(L, C_1, \ldots, C_{n-1}),$$
$$S_{f,C_1,\ldots,C_{n-2}} := S_{K,C_1,\ldots,C_{n-2}} - S_{L,C_1,\ldots,C_{n-2}}.$$

We may analogously define $\mathsf{V}_n(f, g, C_1, \ldots, C_{n-2})$ when f, g are differences of support functions, etc. The extended definitions are still symmetric and multilinear, but are not necessarily nonnegative. Differences of support functions form a large class of functions on S^{n-1}: in particular, we have the following [30, Lemma 1.7.8].

Lemma 2.1. *Any $f \in C^2(S^{n-1})$ is a difference of support functions.*

2.1.2. **Positivity.** While mixed volumes and mixed area measures of convex bodies are always nonnegative, they need not be strictly positive. Positivity of mixed volumes and mixed area measures will play an important role throughout this paper. We presently state two key facts in this direction. First, we recall that positivity of mixed volumes is characterized by dimensionality conditions [30, Theorem 5.1.8]. Throughout this paper, we denote by $[n] := \{1, \ldots, n\}$.

Lemma 2.2. *For convex bodies $C_1, \ldots, C_n$ in $\mathbb{R}^n$, the following are equivalent:*

a. $\mathsf{V}_n(C_1, \ldots, C_n) > 0$.

b. There are segments $I_i \subseteq C_i$, $i \in [n]$ with linearly independent directions.

c. $\dim(C_{i_1} + \cdots + C_{i_k}) \geqslant k$ *for all* $k \in [n]$, $1 \leqslant i_1 < \cdots < i_k \leqslant n$.

Similarly, the mixed area measure $S_{C_1,\ldots,C_{n-1}}$ need not be supported on the entire sphere S^{n-1}. Unlike the positivity of mixed volumes, the problem of characterizing geometrically the support of mixed area measures of arbitrary convex bodies is not yet fully settled, cf. [30, Conjecture 7.6.14]. However, for the present purposes we require only the following special case. For any vector $u \in \mathbb{R}^n$, let

$$F(K, u) := \{x \in K : \langle u, x\rangle = h_K(u)\} \tag{2.2}$$

be the unique face of K with outer normal direction u. The following result states that when $P_1, \ldots, P_{n-2}$ are polytopes, the support of the mixed area measure $S_{B,P_1,\ldots,P_{n-2}}$ is characterized by dimensionality conditions on faces of $P_1, \ldots, P_{n-2}$. This result is essentially known; we will provide a proof in section 5.2.

Lemma 2.3. *Let $P_1, \ldots, P_{n-2}$ be any convex polytopes in $\mathbb{R}^n$, and let $u \in S^{n-1}$. Then the following conditions are equivalent:*

a. $u \in \operatorname{supp} S_{B,P_1,\ldots,P_{n-2}}$.

b. There are segments $I_i \subseteq F(P_i, u)$, $i \in [n-2]$ with linearly independent directions.

c. $\dim(F(P_{i_1}, u) + \cdots + F(P_{i_k}, u)) \geqslant k$ *for all* $k \in [n-2]$, $1 \leqslant i_1 < \cdots < i_k \leqslant n-2$.

*When a–c hold, $u \in S^{n-1}$ is called a $(B, P_1, \ldots, P_{n-2})$-*extreme *normal direction.*

The appearance the Euclidean ball B in Lemma 2.3 may appear rather arbitrary: we did not assume B appears as one of the bodies in Theorem 1.1. Its significance is that the associated mixed area measure has maximal support [30, Lemma 7.6.15] (an alternative proof may be given along the lines of Lemma 8.11 below).

Lemma 2.4. *For any convex bodies $M, C_1, \ldots, C_{n-2}$, we have*

$$\operatorname{supp} S_{M,C_1,\ldots,C_{n-2}} \subseteq \operatorname{supp} S_{B,C_1,\ldots,C_{n-2}}.$$

Let us note that Lemma 2.4 remains valid if B is replaced by any sufficiently smooth convex body; there is nothing uniquely special about B. However, the choice of Euclidean ball will prove to be particularly convenient in our proofs.

2.1.3. **Equality.** We finally recall a basic fact about equality in the Alexandrov-Fenchel inequality. It is evident that there is equality in Theorem 1.1 if and only if the difference between the left- and right-hand sides of the inequality is minimized. The first-order optimality condition associated to this minimum problem gives rise to an equivalent formulation of the equality cases of the Alexandrov-Fenchel inequality, due to Alexandrov [2, p. 80] (cf. section 3.3 or [30, Theorem 7.4.2]).

Lemma 2.5. *Let $K, L, C_1, \ldots, C_{n-2}$ be convex bodies in $\mathbb{R}^n$ such that*

$$\mathsf{V}_n(K, L, C_1, \ldots, C_{n-2}) > 0.$$

Then the following are equivalent:

a. $\mathsf{V}_n(K, L, C_1, \ldots, C_{n-2})^2 = \mathsf{V}_n(K, K, C_1, \ldots, C_{n-2})\, \mathsf{V}_n(L, L, C_1, \ldots, C_{n-2})$.

b. $S_{h_K - a h_L, C_1, \ldots, C_{n-2}} = 0$ *for some* $a > 0$.

Let us emphasize that this result provides essentially no information on the geometry of the extremal bodies $K, L, C_1, \ldots, C_{n-2}$: it is merely a reformulation of the equality condition. The main problem that will be addressed in this paper is to develop a geometric characterization of the extremals.

Remark 2.6. When $\mathsf{V}_n(K, L, C_1, \ldots, C_{n-2}) = 0$, there is automatically equality in Theorem 1.1. These *trivial* equality cases are fully characterized by Lemma 2.2. Nontrivial equality cases arise only when $\mathsf{V}_n(K, L, C_1, \ldots, C_{n-2}) > 0$, as is assumed in Lemma 2.5. This is the setting that will concern us in the rest of this paper.

2.2. **Extremal mechanisms.** What convex bodies yield equality in Theorem 1.1? We will now describe three mechanisms that yield extremals of the Alexandrov-Fenchel inequality, each capturing a different geometric phenomenon: translation (section 2.2.1), support (section 2.2.2), and dimensionality (section 2.2.3).

It is important to note that the bodies K, L and $C_1, \ldots, C_{n-2}$ play very different roles in Theorem 1.1: K, L vary, while $C_1, \ldots, C_{n-2}$ are the same in each term. We therefore consider the reference bodies $C_1, \ldots, C_{n-2}$ as fixed, and aim to characterize which K, L yield equality in Theorem 1.1. By Lemma 2.5, the problem can be formulated equivalently as follows: *given $C_1, \ldots, C_{n-2}$, we aim to characterize what differences of support functions f satisfy $S_{f, C_1, \ldots, C_{n-2}} = 0$.*

2.2.1. **Translation.** The simplest mechanism for equality in Theorem 1.1 stems from the most basic invariance property of mixed volumes: as volume is translation-invariant, (1.1) implies that mixed volumes are as well, that is,

$$\mathsf{V}_n(K, C_1, \ldots, C_{n-1}) = \mathsf{V}_n(K + v, C_1, \ldots, C_{n-1})$$

for all $v \in \mathbb{R}^n$. In terms of support functions, we have $h_{K+v}(u) = h_K(u) + \langle v, u \rangle$, that is, the support function of a convex body and its translate differ by a linear function. This gives rise to the following equality case.

Lemma 2.7. $S_{f, C_1, \ldots, C_{n-2}} = 0$ *whenever* $f = \langle v, \cdot \rangle$ *is a linear function.*

Proof. Let $f = \langle v, \cdot \rangle$ be any linear function. Then $f = h_{K+v} - h_K$ for any convex body K. Therefore, by translation-invariance of mixed volumes,

$$\frac{1}{n} \int g \, dS_{f, C_1, \ldots, C_{n-2}} = \mathsf{V}_n(g, f, C_1, \ldots, C_{n-2}) = 0$$

for any difference of support functions g, and thus *a fortiori* for any $g \in C^2(S^{n-1})$ by Lemma 2.1. It follows immediately that $S_{f, C_1, \ldots, C_{n-2}} = 0$. $\square$

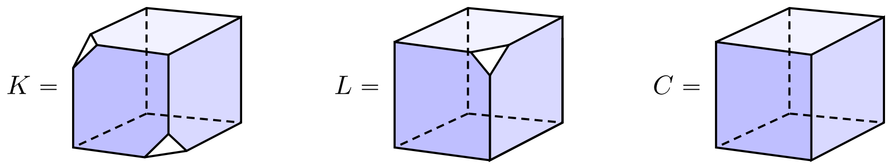

FIGURE 2.1. Example of an equality case described by Lemma 2.8.

Lemma 2.7 and Lemma 2.5 imply, for example, that equality occurs in the Alexandrov-Fenchel inequality whenever $h_K - ah_L = \langle v, \cdot \rangle$ for some $a > 0$ and $v \in \mathbb{R}^n$, which simply means that $K = aL + v$ (that is, K and L are homothetic). Of course, this also follows immediately from Theorem 1.1.

2.2.2. **Support.** A much more subtle invariance property of mixed volumes stems from the fact that mixed area measures need not be supported on the entire sphere S^{n-1}. Indeed, it follows immediately from (2.1) that

$$\mathsf{V}_n(K, C_1, \ldots, C_{n-1}) = \mathsf{V}_n(L, C_1, \ldots, C_{n-1})$$

whenever

$$h_K(u) = h_L(u) \text{ for all } u \in \operatorname{supp} S_{C_1,\ldots,C_{n-1}}.$$

That this phenomenon gives rise to new extremals of the Alexandrov-Fenchel inequality dates back essentially to the work of Minkowski, and has been put forward systematically by Schneider. Let us give a precise formulation [30, p. 430].

Lemma 2.8. $S_{f,C_1,\ldots,C_{n-2}} = 0$ *whenever* $f(u) = 0$ *for all* $u \in \operatorname{supp} S_{B,C_1,\ldots,C_{n-2}}$.

Proof. Suppose f vanishes on $\operatorname{supp} S_{B,C_1,\ldots,C_{n-2}}$. Then

$$\frac{1}{n}\int g\, dS_{f,C_1,\ldots,C_{n-2}} = \mathsf{V}_n(g, f, C_1, \ldots, C_{n-2}) = \frac{1}{n}\int f\, dS_{g,C_1,\ldots,C_{n-2}} = 0,$$

for any difference of support functions g, where we used the symmetry of mixed volumes and that $\operatorname{supp} S_{g,C_1,\ldots,C_{n-2}} \subseteq \operatorname{supp} S_{B,C_1,\ldots,C_{n-2}}$ by Lemma 2.4. The conclusion follows as we may choose any $g \in C^2(S^{n-1})$ by Lemma 2.1. $\square$

In the case that $C_1, \ldots, C_{n-2}$ are polytopes, we have given a geometric characterization of the support of $S_{B,C_1,\ldots,C_{n-2}}$ in Lemma 2.3. This yields a fully geometric interpretation of the situation described by Lemma 2.8: that $f = h_K - h_L$ vanishes on $\operatorname{supp} S_{B,C_1,\ldots,C_{n-2}}$ means precisely that the convex bodies K and L have the same supporting hyperplanes in all $(B, C_1, \ldots, C_{n-2})$-extreme normal directions.

Example 2.9. Let $C = [0,1]^3$ be a cube in $\mathbb{R}^3$, and let the bodies K and L be derived from C by slicing off some of its corners. This construction is illustrated in Figure 2.1. We claim that $h_K - h_L$ vanishes on $\operatorname{supp} S_{B,C}$, so that in particular

$$\mathsf{V}_3(K, L, C)^2 = \mathsf{V}_3(K, K, C)\,\mathsf{V}_3(L, L, C)$$

in this example by Lemmas 2.8 and 2.5.

To verify the claim, note that by part c of Lemma 2.3, we have $u \in \operatorname{supp} S_{B,C}$ if and only if u is a normal direction of a face of C of dimension at least one, that is, if u is the outer normal of a supporting hyperplane of one of the edges of the unit cube. But it is readily seen in Figure 2.1 that any such hyperplane also supports both K and L, so that $h_K(u) = h_L(u)$ for every $u \in \operatorname{supp} S_{B,C}$. There are of course many other directions in which the supporting hyperplanes of K, L differ, but these are all normal to a corner of the cube C and are therefore not in $\operatorname{supp} S_{B,C}$.

2.2.3. **Dimensionality.** We now describe yet another mechanism that gives rise to extremals of the Alexandrov-Fenchel inequality, which arises from the fact that mixed volumes may vanish for dimensionality reasons (Lemma 2.2). To make this idea precise, we introduce the following definition; recall that we are interested in extremals for given reference bodies $\mathcal{C} := (C_1, \ldots, C_{n-2})$.

Definition 2.10. Let (M, N) be a pair of convex bodies, and let $f : S^{n-1} \to \mathbb{R}$.

a. (M, N) is called a *$\mathcal{C}$-degenerate pair* if M is not a translate of N,

$$\mathsf{V}_n(M, N, C_1, \ldots, C_{n-2}) = 0, \tag{2.3}$$

$$\text{and} \quad \mathsf{V}_n(M, B, C_1, \ldots, C_{n-2}) = \mathsf{V}_n(N, B, C_1, \ldots, C_{n-2}). \tag{2.4}$$

b. f is a *$\mathcal{C}$-degenerate function* if $f = h_M - h_N$ for some $\mathcal{C}$-degenerate pair (M, N).

By Lemma 2.2, condition (2.3) is of a purely geometric nature: it is characterized by the dimensions of the relevant bodies. Condition (2.4) should be viewed merely as a normalization; for any pair (M, N) satisfying the first condition, the second condition can always be made to hold by rescaling M or N. We assume M is not a translate of N to exclude the trivial case that $f = h_M - h_N$ is a linear function.

Lemma 2.11. *$S_{f,C_1,\ldots,C_{n-2}} = 0$ whenever f is a $\mathcal{C}$-degenerate function.*

Proof. Let (M, N) be a $\mathcal{C}$-degenerate pair. The main observation is that we obtain equality in Theorem 1.1 for $K = B + M$ and $L = B + N$. Indeed, as

$$\mathsf{V}_n(K, K, C_1, \ldots, C_{n-2}) = \mathsf{V}_n(K, L, C_1, \ldots, C_{n-2}) + \mathsf{V}_n(M, M, C_1, \ldots, C_{n-2}),$$

$$\mathsf{V}_n(L, L, C_1, \ldots, C_{n-2}) = \mathsf{V}_n(K, L, C_1, \ldots, C_{n-2}) + \mathsf{V}_n(N, N, C_1, \ldots, C_{n-2})$$

by (2.3) and (2.4), we obtain

$$\mathsf{V}_n(K, L, C_1, \ldots, C_{n-2})^2 \leqslant \mathsf{V}_n(K, K, C_1, \ldots, C_{n-2})\, \mathsf{V}_n(L, L, C_1, \ldots, C_{n-2}).$$

As the reverse inequality holds by Theorem 1.1, we must in fact have equality.

Now note that if $\mathsf{V}_n(B, B, C_1, \ldots, C_{n-2}) > 0$, then $S_{h_K - a h_L, C_1, \ldots, C_{n-2}} = 0$ for some a by Lemma 2.5. Integrating against h_B and applying (2.1) and (2.4) yields $a = 1$. Thus $S_{f,C_1,\ldots,C_{n-2}} = 0$ for $f = h_K - h_L = h_M - h_N$.

If $\mathsf{V}_n(B, B, C_1, \ldots, C_{n-2}) = 0$, however, then we have $S_{B,C_1,\ldots,C_{n-2}} = 0$ by (2.1) as $h_B = 1$ on S^{n-1}. Thus in this case $S_{f,C_1,\ldots,C_{n-2}} = 0$ for any f by Lemma 2.4. $\square$

The geometric phenomena captured by Lemmas 2.8 and 2.11 are quite different: the former captures the facial structure of the bodies in $\mathcal{C}$, while the latter captures the dimensions of the bodies. Let us illustrate the distinction in a concrete example.

Example 2.12. Let $C_1 = [0, 1]^4$ be a cube in $\mathbb{R}^4$, and let $C_2 = [0, e_1] + [0, e_2]$ be a two-dimensional square in the plane spanned by the first two coordinate directions e_1, e_2. Let $M = [0, e_1]$ and $N = [0, e_2]$ be segments in the same plane.

We claim that (M, N) is a degenerate pair. Indeed, as $\dim(M + N + C_2) = 2$, Lemma 2.2 verifies (2.3). On the other hand, it is clear that (2.4) must hold, as this example is symmetric under exchanging the e_1 and e_2 directions. This gives rise, for example, to the following equality case of the Alexandrov-Fenchel inequality: if we choose $K = C_1 + M$ and $L = C_1 + N$, then Lemmas 2.11 and 2.5 yield

$$\mathsf{V}_4(K, L, C_1, C_2)^2 = \mathsf{V}_4(K, K, C_1, C_2)\, \mathsf{V}_4(L, L, C_1, C_2).$$

We now aim to show that the present example cannot be explained by a combination of Lemmas 2.7 and 2.8, confirming that Lemma 2.11 captures a genuinely distinct phenomenon. That is, we aim to show that $f = h_M - h_N$ does not coincide with a linear function on the support of S_{B,C_1,C_2}. To this end, note that

$$f(u) = \max(u_1, 0) - \max(u_2, 0).$$

On the other hand, for any $u \in S^3 \cap \mathrm{span}\{e_1, e_3\}$ we have $\dim F(C_1, u) \geqslant 2$ and $\dim F(C_2, u) \geqslant 1$, so that $S^3 \cap \mathrm{span}\{e_1, e_3\} \subset \mathrm{supp}\, S_{B,C_1,C_2}$ by Lemma 2.3. Thus f cannot coincide with any linear function on $\mathrm{supp}\, S_{B,C_1,C_2}$, as the restriction of f to the unit circle in $\mathrm{span}\{e_1, e_3\}$ is not a smooth function.

2.3. **Main result.** In the previous section, we described three distinct mechanisms for equality in the Alexandrov-Fenchel inequality in Lemmas 2.7, 2.8, and 2.11. However, these three mechanisms may all appear simultaneously by linearity: if $S_{f,C_1,\ldots,C_{n-2}} = 0$ and $S_{g,C_1,\ldots,C_{n-2}} = 0$, then $S_{f+g,C_1,\ldots,C_{n-2}} = 0$ as well. Thus any linear combination of the functions that appear in Lemmas 2.7, 2.8, and 2.11 gives rise to an extremal case of the Alexandrov-Fenchel inequality.

As no other mechanism for equality is known, one may conjecture that these are the *only* extremal cases of the Alexandrov-Fenchel inequality. The main result of this paper is a complete proof of this conjecture in the combinatorial setting. In geometric terms, we prove the following. (Recall that $\mathcal{P}$-degenerate pairs and $(B, \mathcal{P})$-extreme directions are defined in Definition 2.10 and Lemma 2.3.)

Theorem 2.13. *Let $\mathcal{P} := (P_1, \ldots, P_{n-2})$ be polytopes in $\mathbb{R}^n$, and let K, L be convex bodies such that $\mathsf{V}_n(K, L, P_1, \ldots, P_{n-2}) > 0$.*[2] *Then*

$$\mathsf{V}_n(K, L, P_1, \ldots, P_{n-2})^2 = \mathsf{V}_n(K, K, P_1, \ldots, P_{n-2})\, \mathsf{V}_n(L, L, P_1, \ldots, P_{n-2})$$

if and only if there exist $a > 0$, $v \in \mathbb{R}^n$, and a number $0 \leqslant m < \infty$ of $\mathcal{P}$-degenerate pairs $(M_1, N_1), \ldots, (M_m, N_m)$, so that $K + N_1 + \cdots + N_m$ and $aL + v + M_1 + \cdots + M_m$ have the same supporting hyperplanes in all $(B, \mathcal{P})$-extreme normal directions.

The *if* direction of Theorem 2.13 follows from Lemmas 2.5, 2.7, 2.8, and 2.11, so it is the *only if* part that requires proof. Some key ideas in the proof are described in section 4; the proof itself is contained in sections 5–12.

Schneider has conjectured [26] that equality in the Alexandrov-Fenchel inequality holds if and only if K and $aL + v$ have the same supporting hyperplanes in all $(B, \mathcal{P})$-extreme normal directions. That this is not always the case was illustrated in Example 2.12 (the existence of counterexamples was first noted in [8]). Nonetheless, no counterexample has been found to Schneider's conjecture in the case where all bodies in $\mathcal{P}$ are full-dimensional. This suggests that in the full-dimensional situation, degenerate pairs may not exist. Not only does this turn out to be the case, but in fact a much weaker condition suffices.

Definition 2.14. A collection of convex bodies $\mathcal{C} = (C_1, \ldots, C_{n-2})$ is *supercritical* if $\dim(C_{i_1} + \cdots + C_{i_k}) \geqslant k + 2$ for all $k \in [n - 2]$, $1 \leqslant i_1 < \cdots < i_k \leqslant n - 2$.

Lemma 2.15. *If $\mathcal{C}$ is supercritical, $\mathcal{C}$-degenerate functions do not exist.*

[2] As was noted in Remark 2.6, the trivial extremals $\mathsf{V}_n(K, L, P_1, \ldots, P_{n-2}) = 0$ are already fully characterized geometrically by Lemma 2.2, so we do not consider them further.

Proof. Suppose (M, N) is a $\mathcal{C}$-degenerate pair. By Lemma 2.2 and the supercriticality assumption, (2.3) implies that $\dim(M) = 0$, $\dim(N) = 0$, or $\dim(M + N) \leqslant 1$.

Assume first that $\dim(M) = 0$. Then $\mathsf{V}_n(N, B, C_1, \ldots, C_{n-2}) = 0$ by (2.4). But then by Lemma 2.2 and the supercriticality assumption, $\dim(N) = 0$ as well.

Thus there are two possibilities: $\dim(M) = \dim(N) = 0$, or $\dim(M) = \dim(N) = \dim(N + M) = 1$. In the first case M and N are singletons, while in the second case M and N are segments with parallel directions. Moreover, in the latter case $\mathsf{V}_n(N, B, C_1, \ldots, C_{n-2}) > 0$ by Lemma 2.2 and the supercriticality assumption, so (2.4) implies that M and N have the same length. Thus in either case M and N are translates of one another, which violates the definition of a degenerate pair. □

In other words, Lemma 2.15 yields:

Corollary 2.16. *Let $\mathcal{P} := (P_1, \ldots, P_{n-2})$ be a supercritical collection of polytopes in $\mathbb{R}^n$, and let K, L be convex bodies such that $\mathsf{V}_n(K, L, P_1, \ldots, P_{n-2}) > 0$. Then*

$$\mathsf{V}_n(K, L, P_1, \ldots, P_{n-2})^2 = \mathsf{V}_n(K, K, P_1, \ldots, P_{n-2})\, \mathsf{V}_n(L, L, P_1, \ldots, P_{n-2})$$

if and only if there exist $a > 0$ and $v \in \mathbb{R}^n$ so that K and $aL + v$ have the same supporting hyperplanes in all $(B, \mathcal{P})$-extreme normal directions.

Corollary 2.16 highlights that even though Theorem 2.13 provides a complete characterization of the extremals of the Alexandrov-Fenchel inequality for arbitrary polytopes $\mathcal{P}$, its formulation leaves key questions open: it does not explain how many degenerate pairs can appear, what they look like, or whether the decomposition into degenerate pairs is unique. A complete understanding of these questions will emerge from the proof of Theorem 2.13. As the requisite notions will only be introduced as we progress through the proof, we postpone formulating the definitive form of our main result until section 13.

While we have presented Corollary 2.16 as a special case of Theorem 2.13, the supercritical case will prove to be fundamental to the proof. We will first give a self-contained proof of Corollary 2.16 in sections 5–8, and then characterize the degenerate equality cases in sections 9–12 by a separate argument that requires the introduction of additional techniques. In particular, the proof of Corollary 2.16 may be read independently from the rest of the paper.

2.4. **Prior work.** Let us briefly review what was known prior to this paper. Three cases of Corollary 2.16 were previously verified: when $\mathcal{P}$ consists of strongly isomorphic simple polytopes [30, Theorem 7.6.21], when $P_1 = \cdots = P_{n-2}$ [28, 32] (in this case P_i need not be simple), or when $\mathcal{P}$ consists of full-dimensional zonotopes and K, L are symmetric [27]. All these results make crucial use of the special features that appear in these settings. In addition, one very special example of a degenerate equality case was previously known, when all the bodies $\mathcal{P}$ lie in a hyperplane [10, 29]. This example sheds little light on more general cases, however, as it is essentially amenable to explicit computation, cf. [32, §8].

The characterization of lower-dimensional extremals in terms of degenerate pairs was conjectured by the authors during initial work on this paper. We subsequently realized, however, that an analogous phenomenon appears in work of Panov [23] on Alexandrov's mixed discriminant inequality, which may be viewed as an analogue of the Alexandrov-Fenchel inequality in linear algebra. Despite tantalizing similarities between these inequalities, the main feature the Alexandrov-Fenchel inequality does not arise here: dimensionality is the only extremal mechanism in the mixed discriminant inequality, while the central difficulty in the analysis of the Alexandrov-Fenchel inequality stems from degeneration of the support of mixed area measures. While most of our analysis has little in common with [23], we will use a basic lemma of [23] to organize the collection of degenerate pairs (Lemma 9.2).

3. Preliminaries

The aim of this section is to recall some general background from convex geometry that will be needed in the remainder of the paper.

The following conventions will be in force throughout the paper. We always denote by B the Euclidean unit ball in $\mathbb{R}^n$. For any collection $\mathcal{C} := (C_1, \ldots, C_{n-2})$ of convex bodies in $\mathbb{R}^n$, we will often use the abbreviated notation

$$\mathsf{V}_n(K, L, \mathcal{C}) := \mathsf{V}_n(K, L, C_1, \ldots, C_{n-2}), \qquad S_{L,\mathcal{C}} := S_{L,C_1,\ldots,C_{n-2}}.$$

For $I \subseteq [n-2]$, we denote by $\mathcal{C}_I := (C_i)_{i\in I}$ and by $\mathcal{C}_{\backslash I} := (C_i)_{i\in[n-2]\backslash I}$.

We will also encounter mixed volumes of convex bodies $C_1, \ldots, C_m$ that lie in a subspace $E \subset \mathbb{R}^n$ with $\dim(E) = m$. Such mixed volumes will be denoted as $\mathsf{V}_E(C_1, \ldots, C_m)$, or as $\mathsf{V}_m(C_1, \ldots, C_m)$ when the subspace is clear from context.

3.1. Mixed volumes and mixed area measures.

Mixed volumes and mixed area measures were introduced in section 2.1. For future reference, we begin by spelling out their basic properties more carefully.

Mixed volumes are defined by (1.1). They satisfy the following [30, §5.1]. (Here and in the sequel, we use the notation $[\![A]\!] := \sqrt{\det A^*A}$ for a linear map A.)

Lemma 3.1. *Let $C, C', C_1, \ldots, C_n$ be convex bodies in $\mathbb{R}^n$.*

a. $\mathsf{V}_n(C, \ldots, C) = \mathrm{Vol}_n(C)$.
b. $\mathsf{V}_n(C_1, \ldots, C_n)$ *is symmetric and multilinear in its arguments.*
c. $\mathsf{V}_n(C_1, \ldots, C_n) \geqslant 0$.
d. $\mathsf{V}_n(C, C_2, \ldots, C_n) \geqslant \mathsf{V}_n(C', C_2, \ldots, C_n)$ *if* $C \supseteq C'$.
e. $\mathsf{V}_n(C_1, \ldots, C_n)$ *is invariant under translation* $C_i \leftarrow C_i + v_i$ *for* $v_i \in \mathbb{R}^n$.
f. $\mathsf{V}_n(AC_1, \ldots, AC_n) = [\![A]\!]\,\mathsf{V}_n(C_1, \ldots, C_n)$ *for any linear map* $A : \mathbb{R}^n \to \mathbb{R}^n$.

The identity (2.1) may be viewed as the definition of mixed area measures. The following basic properties are analogous to those of mixed volumes [30, §5.1].

Lemma 3.2. *Let $C, C_1, \ldots, C_{n-1}$ be convex bodies in $\mathbb{R}^n$.*

a. $S_{C_1,\ldots,C_{n-1}}$ *is symmetric and multilinear in its arguments.*
b. $S_{C_1,\ldots,C_{n-1}} \geqslant 0$.
c. $S_{C_1,\ldots,C_{n-1}}$ *is invariant under translation* $C_i \leftarrow C_i + v_i$.
d. $\int\langle v, x\rangle\, S_{C_1,\ldots,C_{n-1}}(dx) = 0$ *for all* $v \in \mathbb{R}^n$.

We now recall the basic continuity property of mixed volumes and mixed area measures. Recall that convex bodies $C^{(l)}$ converge to a convex body C in the sense of Hausdorff convergence if and only if $\|h_{C^{(l)}} - h_C\|_\infty \to 0$, cf. [30, Lemma 1.8.14]. Then we have the following result [30, pp. 280–281].

Lemma 3.3. *Suppose that $C_1^{(l)}, \ldots, C_n^{(l)}$ are convex bodies in $\mathbb{R}^n$ such that $C_i^{(l)} \to C_i$ as $l \to \infty$ in the sense of Hausdorff convergence. Then*

$$\mathsf{V}_n(C_1^{(l)}, \ldots, C_n^{(l)}) \to \mathsf{V}_n(C_1, \ldots, C_n), \qquad S_{C_1^{(l)},\ldots,C_{n-1}^{(l)}} \xrightarrow{w} S_{C_1,\ldots,C_{n-1}}$$

as $l \to \infty$, where the limit of measures is in the sense of weak convergence.

In the case that all the convex bodies are polytopes, mixed area measures take a particularly simple form [30, p. 279].

Lemma 3.4. *Let $P_1, \ldots, P_{n-1}$ be polytopes in $\mathbb{R}^n$. Then $S_{P_1,\ldots,P_{n-1}}$ is atomic, that is,* $\operatorname{supp} S_{P_1,\ldots,P_{n-1}} = \{u \in S^{n-1} : S_{P_1,\ldots,P_{n-1}}(\{u\}) > 0\}$*, with*

$$S_{P_1,\ldots,P_{n-1}}(\{u\}) = \mathsf{V}_{n-1}(F(P_1,u),\ldots,F(P_{n-1},u)).$$

Remark 3.5. In Lemma 3.4 we have made a slight abuse of notation: the faces $F(P_i,u)$, $i = 1,\ldots,n-1$ need not lie in a single $(n-1)$-dimensional subspace. However, by definition all these faces have u as a normal direction, so that each face may be translated to lie in $u^\perp$. We implicitly define $\mathsf{V}_{n-1}(F(P_1,u),\ldots,F(P_{n-1},u))$ as the mixed volume in $u^\perp$ of the translated faces; this convenient notation is consistent with the translation-invariance of mixed volumes.

As the faces $F(P,u)$ play a fundamental role in what follows, let us briefly recall at this stage some associated notions. A *facet* of a convex body C in $\mathbb{R}^n$ is an $(n-1)$-dimensional face of C. We recall that every polytope has a finite number of facets. We also recall the following basic property [30, §1.7].

Lemma 3.6. *Let C, C' be any convex bodies in $\mathbb{R}^n$ and $u, x \in \mathbb{R}^n$. Then*

$$h_{F(C,u)}(x) = \nabla_x h_C(u),$$

where ∇_x denotes the directional derivative in direction x. In particular,

$$F(C + C', u) = F(C,u) + F(C',u).$$

Consequently, we may observe that the mixed area measure in Lemma 3.4 is in fact supported on a finite number of points. Indeed, Lemmas 3.4 and 2.2 imply that every $u \in \operatorname{supp} S_{P_1,\ldots,P_{n-1}}$ must satisfy $\dim F(P_1 + \cdots + P_{n-1}, u) \geqslant n-1$, that is, each such u must be a facet normal of $P_1 + \cdots + P_{n-1}$. As the Minkowski sum of polytopes is a polytope, $\operatorname{supp} S_{P_1,\ldots,P_{n-1}}$ must be finite.

Finally, the following basic property of faces will be useful. Here and in the sequel, we denote by $\mathbf{P}_E$ the orthogonal projection onto a subspace E of $\mathbb{R}^n$.

Lemma 3.7. *For any convex body C in $\mathbb{R}^n$, linear subspace $E \subseteq \mathbb{R}^n$, and $u \in \mathbb{R}^n$,*

$$F(\mathbf{P}_E C, u) = \mathbf{P}_E F(C, \mathbf{P}_E u).$$

Proof. Using Lemma 3.6, we can compute

$$\begin{aligned} h_{F(\mathbf{P}_E C,u)}(x) &= \nabla_x h_{\mathbf{P}_E C}(u) = \nabla_{\mathbf{P}_E x} h_C(\mathbf{P}_E u) \\ &= h_{F(C,\mathbf{P}_E u)}(\mathbf{P}_E x) = h_{\mathbf{P}_E F(C,\mathbf{P}_E u)}(x) \end{aligned}$$

for every $x \in \mathbb{R}^n$, where we used $h_{\mathbf{P}_E C}(u) = h_C(\mathbf{P}_E u)$. □

3.2. **Projection formulae.** The relation between mixed volumes of convex bodies and their projections will play a recurring role in this paper. The following result captures this connection in a general setting [30, Theorem 5.3.1].

Lemma 3.8. *Let E be an m-dimensional subspace of $\mathbb{R}^n$, let $C_1, \ldots, C_m$ be convex bodies in E, and let $C_{m+1}, \ldots, C_n$ be convex bodies in $\mathbb{R}^n$. Then*

$$\binom{n}{m} \mathsf{V}_n(C_1,\ldots,C_n) = \mathsf{V}_E(C_1,\ldots,C_m)\, \mathsf{V}_{E^\perp}(\mathbf{P}_{E^\perp} C_{m+1},\ldots,\mathbf{P}_{E^\perp} C_n).$$

We will use Lemma 3.8 in its full force many times. The special case $m = 1$ is particularly important, however, so we highlight it separately.

Corollary 3.9. *Let $C_1, \ldots, C_{n-1}$ be convex bodies in $\mathbb{R}^n$, and let $u \in S^{n-1}$. Then*

$$n\,\mathsf{V}_n([0,u], C_1, \ldots, C_{n-1}) = \mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} C_1, \ldots, \mathbf{P}_{u^\perp} C_{n-1}).$$

When combined with Corollary 3.9, the following observation expresses certain n-dimensional mixed volumes in terms of $(n-1)$-dimensional projections.

Lemma 3.10. *Let $C_1, \ldots, C_{n-1}$ be convex bodies in $\mathbb{R}^n$. Then*

$$\int_{S^{n-1}} \mathsf{V}_n([0,u], C_1, \ldots, C_{n-1})\,\omega(du) = \kappa_{n-1} \mathsf{V}_n(B, C_1, \ldots, C_{n-1}),$$

where ω denotes the Lebesgue measure on S^{n-1} and κ_{n-1} denotes the volume of the Euclidean unit ball in $\mathbb{R}^{n-1}$.

Proof. Apply (2.1) and $\int h_{[0,u]}(x)\,\omega(du) = \int \langle u, x\rangle_+\,\omega(du) = \kappa_{n-1}\, h_B(x)$. ☐

3.3. **Alexandrov-Fenchel inequality and equality.** The classical formulation of the Alexandrov-Fenchel inequality given in Theorem 1.1 is not the most general one: as was emphasized by Alexandrov [1, 2], the convex body K may be replaced by any difference of support functions f. We will often require this more general inequality and its equality cases. We presently make precise the connection between these formulations. The results of this section could be deduced from [30, §7.4], but we find it more insightful to give direct proofs.

We begin by spelling out three equivalent formulations of Theorem 1.1.

Lemma 3.11. *Let $\mathcal{C} = (C_1, \ldots, C_{n-2})$ be convex bodies in $\mathbb{R}^n$. The following are three equivalent formulations of the Alexandrov-Fenchel inequality:*

a. For any convex bodies K, L,

$$\mathsf{V}_n(K, L, \mathcal{C})^2 \geqslant \mathsf{V}_n(K, K, \mathcal{C})\,\mathsf{V}_n(L, L, \mathcal{C}).$$

b. For any difference of support functions g and convex body L,

$$\mathsf{V}_n(g, L, \mathcal{C})^2 \geqslant \mathsf{V}_n(g, g, \mathcal{C})\,\mathsf{V}_n(L, L, \mathcal{C}).$$

c. For any difference of support functions f and convex body L with $\mathsf{V}_n(L, L, \mathcal{C}) > 0$,

$$\mathsf{V}_n(f, L, \mathcal{C}) = 0 \quad \textit{implies} \quad \mathsf{V}_n(f, f, \mathcal{C}) \leqslant 0.$$

Moreover, if $\mathsf{V}_n(L, L, \mathcal{C}) > 0$, then equality holds in part b if and only if there exists $a \in \mathbb{R}$ such that equality holds in part c with $f = g - a h_L$.

Proof. The implications $b \Rightarrow a$, $b \Rightarrow c$, and $c \Rightarrow b$ follow readily by choosing, respectively, $g = h_K$, $g = f$, and $f = g - a h_L$ with $a = \mathsf{V}_n(g, L, \mathcal{C})/\mathsf{V}_n(L, L, \mathcal{C})$ (we may assume $\mathsf{V}_n(L, L, \mathcal{C}) > 0$ in the latter case, as otherwise b is trivial.)

To prove $a \Rightarrow b$, note first that if $g = h_K - a h_L$ for some $a \in \mathbb{R}$, the $a h_L$ term cancels on both sides of the inequality in b by expanding the square, so that $a \Rightarrow b$ follows trivially. But if g and L are sufficiently smooth, then we may always write $g = h_K - a h_L$ for some $a > 0$ and convex body K [31, Corollary 2.2]; thus the implication $a \Rightarrow b$ follows under smoothness assumptions, and consequently in general by a standard approximation argument [30, §3.4].

Finally, suppose $\mathsf{V}_n(L, L, \mathcal{C}) > 0$. Then it is immediate that b holds with equality if and only if c holds with equality for $f = g - a h_L$ with $a = \mathsf{V}_n(g, L, \mathcal{C})/\mathsf{V}_n(L, L, \mathcal{C})$. It remains to note that if c holds with equality with $f = g - a h_L$ for some $a \in \mathbb{R}$, then it follows from $\mathsf{V}_n(f, L, \mathcal{C}) = 0$ that necessarily $a = \mathsf{V}_n(g, L, \mathcal{C})/\mathsf{V}_n(L, L, \mathcal{C})$. ☐

In view of Lemma 3.11, to study the equality cases of the Alexandrov-Fenchel inequality it suffices to consider the formulation of part c of Lemma 3.11. We presently reformulate the equality condition

$$\mathsf{V}_n(f, L, \mathcal{C}) = 0 \quad \text{and} \quad \mathsf{V}_n(f, f, \mathcal{C}) = 0 \tag{3.1}$$

using the first-order condition of optimality, following [2, p. 80]. For future reference, we consider a slightly more general situation than arises in Lemma 3.11.

Lemma 3.12. *Let f be a difference of support functions, and let L and $\mathcal{C} = (C_1, \ldots, C_{n-2})$ be convex bodies in $\mathbb{R}^n$.*

a. Suppose $\mathsf{V}_n(L, L, \mathcal{C}) > 0$. Then (3.1) holds if and only if $S_{f,\mathcal{C}} = 0$.

b. Suppose $\mathsf{V}_n(L, L, \mathcal{C}) = 0$ and $S_{L,\mathcal{C}} \neq 0$. Then (3.1) holds if and only if there exists $a \in \mathbb{R}$ such that $S_{f-ah_L,\mathcal{C}} = 0$.

Proof. We first prove part a. If $S_{f,\mathcal{C}} = 0$, then $\int h_L dS_{f,\mathcal{C}} = \int f dS_{f,\mathcal{C}} = 0$ and (2.1) yields (3.1). Conversely, suppose (3.1) holds, and let g be any difference of support functions. As $\mathsf{V}_n(L, L, \mathcal{C}) > 0$, we can choose a so that $\mathsf{V}_n(g - ah_L, L, \mathcal{C}) = 0$. Then

$$\varphi(\lambda) := \mathsf{V}_n(f + \lambda[g - ah_L], f + \lambda[g - ah_L], \mathcal{C})$$

satisfies $\varphi(\lambda) \leqslant 0$ by Lemma 3.11(c) and $\varphi(0) = 0$ by (3.1). Thus φ is a quadratic function with maximum at 0, so $\varphi'(0) = 0$. Using $\mathsf{V}_n(f, L, \mathcal{C}) = 0$, this yields

$$0 = \mathsf{V}_n(g, f, \mathcal{C}) = \frac{1}{n} \int g \, dS_{f,\mathcal{C}}.$$

As we may choose g to be any C^2 function by Lemma 2.1, we have $S_{f,\mathcal{C}} = 0$.

We now prove part b. If $S_{f-ah_L,\mathcal{C}} = 0$, then $n\mathsf{V}_n(f, L, \mathcal{C}) = \int h_L dS_{f-ah_L,\mathcal{C}} = 0$ as $\mathsf{V}_n(L, L, \mathcal{C}) = 0$; consequently, $n\mathsf{V}_n(f, f, \mathcal{C}) = \int f dS_{f-ah_L,\mathcal{C}} = 0$, proving (3.1). Conversely, suppose (3.1) holds. As $S_{L,\mathcal{C}} \neq 0$ we have $\mathsf{V}_n(B, L, \mathcal{C}) > 0$. Therefore:

- We may choose $a \in \mathbb{R}$ so that $\mathsf{V}_n(f - ah_L, B, \mathcal{C}) = 0$.
- $\mathsf{V}_n(f - ah_L, f - ah_L, \mathcal{C}) = 0$ by (3.1) and $\mathsf{V}_n(L, L, \mathcal{C}) = 0$.
- $\mathsf{V}_n(B, B, \mathcal{C}) > 0$ as $\mathsf{V}_n(B, L, \mathcal{C}) > 0$ and $L \subseteq cB$ for some $c > 0$.

We can now apply part a with $L \leftarrow B$, $f \leftarrow f - ah_L$ to conclude. □

For completeness, we conclude with a proof of Lemma 2.5.

Proof of Lemma 2.5. Let K, L and $\mathcal{C} = (C_1, \ldots, C_{n-2})$ be as in the statement of Lemma 2.5. To prove $b \Rightarrow a$, it suffices to note that integrating condition b against h_K and h_L yields $\mathsf{V}_n(K, K, \mathcal{C}) = a\mathsf{V}_n(K, L, \mathcal{C}) = a^2\mathsf{V}_n(L, L, \mathcal{C})$ by (2.1). To prove $a \Rightarrow b$, note that the assumption $\mathsf{V}_n(K, L, \mathcal{C}) > 0$ and condition a imply $\mathsf{V}_n(L, L, \mathcal{C}) > 0$. Thus Lemmas 3.11 and 3.12 imply $S_{h_K - ah_L,\mathcal{C}} = 0$ for some $a \in \mathbb{R}$. But integrating against h_L yields $\mathsf{V}_n(K, L, \mathcal{C}) = a\mathsf{V}_n(L, L, \mathcal{C})$ by (2.1), so $a > 0$. □

4. Overview of the proof

The main result of this paper, Theorem 2.13, is proved in sections 5–12 below. Before we proceed to the details, however, we aim to give a high-level overview of the proof in order to help the reader navigate the following sections. At the most basic level, the proof proceeds by induction on the dimension n. The argument splits into two parts that require completely different ideas and techniques.

Throughout the proof of Theorem 2.13, we will fix $n \geqslant 3$ and polytopes $\mathcal{P} = (P_1, \ldots, P_{n-2})$ in $\mathbb{R}^n$. Let us introduce at the outset a minimal dimensionality condition that will be assumed throughout most of this paper.

Definition 4.1. A collection of convex bodies $\mathcal{C} = (C_1, \ldots, C_{n-2})$ is *critical* if $\dim(C_{i_1} + \cdots + C_{i_k}) \geqslant k+1$ for all $k \in [n-2]$, $1 \leqslant i_1 < \cdots < i_k \leqslant n-2$.

Note that if there exist $i_1 < \cdots < i_k$ with $\dim(P_{i_1} + \cdots + P_{i_k}) \leqslant k$, then the bodies $(P_{i_1}, \ldots, P_{i_k})$ factor on both sides of the Alexandrov-Fenchel inequality by Lemma 3.8, and the problem reduces to a lower-dimensional one. For this reason, we may focus our attention on the case that $\mathcal{P}$ is critical, and the remaining cases will be easily dispensed with at the very end of the proof.

4.1. **The local Alexandrov-Fenchel inequality.** In order to perform induction on the dimension, we must understand how the extremals of the Alexandrov-Fenchel inequality in dimensions n and $n-1$ are related. The purpose of the first part of the proof of Theorem 2.13 is to make this connection. To explain how this is done, we begin by discussing an apparently unrelated question.

In view of their definition (2.1), it is natural to think of mixed area measures as local analogues of mixed volumes: they describe the behavior of mixed volumes in different normal directions. The analogy is even more explicit in the polytope case, cf. Lemma 3.4. One might therefore wonder whether there exists an analogue of the Alexandrov-Fenchel inequality for mixed area measures. This question makes little sense in the formulation of Theorem 1.1, of course, as one cannot square a measure. However, the question can be meaningfully formulated in the form of Lemma 3.11(*c*): given convex bodies $L, C_1, \ldots, C_{n-3}$, is it true that

$$S_{f,L,C_1,\ldots,C_{n-3}} = 0 \quad \overset{?}{\Longrightarrow} \quad S_{f,f,C_1,\ldots,C_{n-3}} \leqslant 0 \tag{4.1}$$

for any difference of support functions f? We will refer to any statement of the form (4.1) as a *local Alexandrov-Fenchel inequality.*

Let us first explain why such an inequality would enable an induction argument, at least in the full-dimensional case. To this end, we make a simple observation.

Lemma 4.2. *Let $\mathcal{C} = (C_1, \ldots, C_{n-2})$ be convex bodies in $\mathbb{R}^n$, let $r \in [n-2]$, and suppose C_r is full-dimensional. If $S_{f,\mathcal{C}} = 0$ and $S_{f,f,\mathcal{C}_{\backslash r}} \leqslant 0$, then $S_{f,f,\mathcal{C}_{\backslash r}} = 0$.*

Proof. By translation-invariance we may assume $0 \in \operatorname{int} C_r$, so that $h_{C_r} > 0$. Now note that as $S_{f,\mathcal{C}} = 0$, using (2.1) and the symmetry of mixed volumes yields

$$0 = \int f \, dS_{f,\mathcal{C}} = \int h_{C_r} \, dS_{f,f,\mathcal{C}_{\backslash r}}.$$

The conclusion follows as $S_{f,f,\mathcal{C}_{\backslash r}} \leqslant 0$ and $h_{C_r} > 0$. □

Now suppose we have equality in Theorem 1.1, and assume for simplicity that C_r is full-dimensional for some $r \in [n-2]$. Then by Lemma 2.5, we have

$$S_{f,\mathcal{C}} = 0 \quad \text{with} \quad f = h_K - a h_L$$

for some $a > 0$. If the local Alexandrov-Fenchel inequality (4.1) were to hold, we would obtain $S_{f,f,\mathcal{C}_{\backslash r}} \leqslant 0$, and thus $S_{f,f,\mathcal{C}_{\backslash r}} = 0$ by Lemma 4.2. Integrating both measures against $h_{[0,u]}$ (for any $u \in S^{n-1}$) yields, by (2.1) and Corollary 3.9,

$$\mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} f, \mathbf{P}_{u^\perp} \mathcal{C}) = 0 \qquad \text{and} \qquad \mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} f, \mathbf{P}_{u^\perp} f, \mathbf{P}_{u^\perp} \mathcal{C}_{\backslash r}) = 0,$$

where $\mathbf{P}_E f := h_{\mathbf{P}_E K} - a h_{\mathbf{P}_E L}$ and $\mathbf{P}_E \mathcal{C} := (\mathbf{P}_E C_1, \ldots, \mathbf{P}_E C_{n-2})$. But the latter is nothing other than an equality case (3.1) of the Alexandrov-Fenchel inequality in $u^\perp$. Thus *a local Alexandrov-Fenchel inequality would imply that extremality for the Alexandrov-Fenchel inequality in dimension n is inherited by projection onto any $(n-1)$-dimensional subspace*, opening the door to induction.

Unfortunately, it turns out that this approach breaks down precisely when the Alexandrov-Fenchel inequality has nontrivial extremals. That the above conclusion must fail in this case is immediately evident from the classical fact that equality $\mathsf{V}_2(K, L)^2 = \mathsf{V}_2(K, K)\,\mathsf{V}_2(L, L)$ holds in dimension $n = 2$ if and only if K, L are homothetic (cf. Remark 8.2 and the proof of Theorem 8.1). Thus it cannot be the case that the projections of a nontrivial equality case of the Alexandrov-Fenchel inequality in dimension $n \geqslant 3$ yield equality in dimension 2, as convex bodies in dimension $n \geqslant 3$ whose projections onto every hyperplane are homothetic must themselves be homothetic [36] (this is illustrated, for example, by Figure 2.1). In particular, it follows that the validity of the local Alexendrov-Fenchel inequality (4.1) is contradicted by the presence of nontrivial extremals.

At first sight, the failure of (4.1) appears to render the above approach useless for the study of the extremals. Remarkably, however, this turns out not to be the case. Recall that by Lemma 2.8, the measure $S_{f,\mathcal{C}}$ is unchanged if we modify f outside the support of $S_{B,\mathcal{C}}$; in particular, as we characterize extremals only up to $S_{B,\mathcal{C}}$-a.e. equivalence, we are free to modify f outside $\operatorname{supp} S_{B,\mathcal{C}}$ in the proof. On the other hand, the same property does *not* hold for $S_{f,f,\mathcal{C}}$: this measure may change drastically if we modify f outside the support of $S_{B,\mathcal{C}}$. One of the central ideas of this paper is that we can exploit the resulting degrees of freedom to force the validity of (4.1). More precisely, we will prove the following.

Theorem 4.3 (Local Alexandrov-Fenchel inequality)**.** *Let $\mathcal{P} = (P_1, \ldots, P_{n-2})$ be a critical collection of polytopes in $\mathbb{R}^n$, and fix $r \in [n-2]$. Then for any difference of support functions f so that $S_{f,\mathcal{P}} = 0$, there exists a difference of support functions g so that $S_{g,\mathcal{P}} = 0$, $S_{g,g,\mathcal{P}_{\backslash r}} \leqslant 0$, and $g(x) = f(x)$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$.*

The proof of 4.3 is the main part of this paper in which we exploit the assumption that the reference bodies are polytopes (see section 16 for discussion). The simplification provided by this setting is that it enables us to reduce Theorem 4.3 to a finite-dimensional problem, which will be accomplished in sections 5–6 by adapting ideas from Alexandrov's original proof of the Alexandrov-Fenchel inequality using strongly isomorphic polytopes [1] to the setting of arbitrary polytopes. It should be emphasized, however, that this reduction is merely a technical device: the entire difficulty of the proof lies in section 7, where we prove the existence of a function g with the requisite properties. We will ultimately reduce this problem to a system of linear equations, and the heart of the matter is to rule out the presence of degeneracies that would obstruct the existence of a solution.

Remark 4.4. The simple argument in the proof of Lemma 4.2 is due to Weyl [40]. It is used in classical proofs of the Alexandrov-Fenchel inequality precisely to *rule out* the existence of nontrivial extremals; see, e.g., [3, p. 110] or [30, p. 396]. It therefore appears rather surprising that such an argument provides a starting point for the study of nontrivial extremals. That this is in fact the case relies crucially on Theorem 4.3, which is a central new idea of this paper that opens the door to the analysis of the extremals by induction on the dimension.

A different induction argument was exploited by Schneider [27] to investigate extremals of the Alexandrov-Fenchel inequality for zonoids, that is, limits of Minkowski sums of segments. In this setting, the relation between the extremals and their projections arises from Corollary 3.9, but this appears as a very special property of this class of bodies. A notable aspect of our approach is that we are able to perform induction by projection in the absence of such special structure.

4.2. **The gluing argument.** Once we have shown that extremality is preserved by projection onto hyperplanes, we must combine the information contained in the $(n-1)$-dimensional projections in order to characterize the n-dimensional extremals. This is the purpose of the second part of the proof of Theorem 2.13.

At first sight, it seems evident that we may reconstruct a convex body from its projections, as $h_{\mathsf{P}_E K}(x) = h_K(\mathsf{P}_E x)$ for all x by the definition of support functions. Thus if the function $\mathsf{P}_{u^\perp} f$ were known for every u, the function f would be uniquely determined. The situation we encounter is much more delicate, however, as only very limited information about the projections will follow from the induction hypothesis that Theorem 2.13 holds in dimension $n-1$.

To illustrate the difficulty, suppose for simplicity that all polytopes in $\mathcal{P}$ are full-dimensional, and let f be an equality case of the Alexandrov-Fenchel inequality in dimension n, that is, $S_{f,\mathcal{P}} = 0$. We aim to prove the conclusion of Corollary 2.16, that is, there exists $s \in \mathbb{R}^n$ so that $f(x) = \langle s, x\rangle$ for $x \in \operatorname{supp} S_{B,\mathcal{P}}$. If we assume Corollary 2.16 holds in dimension $n-1$, then Theorem 4.3 and the arguments of the previous section show that there exists $s(u) \in \mathbb{R}^n$ such that

$$f(x) = \langle s(u), x\rangle \qquad \text{for all } x \in \operatorname{supp} S_{\mathsf{P}_{u^\perp} B, \mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}} \cap \operatorname{supp} S_{B,\mathcal{P}}$$

for every $u \in S^{n-1}$. We now face two problems: the linear function $\langle s(u), x\rangle$ may *a priori* depend on u; and we have only very limited information for any given u, as $\operatorname{supp} S_{\mathsf{P}_{u^\perp} B, \mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}} \cap \operatorname{supp} S_{B,\mathcal{P}}$ may only cover a very small part of $S^{n-1} \cap u^\perp$. We must therefore rule out, for example, that f is piecewise linear on disjoint parts of the supports of the mixed area measures that arise for different u.

In the supercritical case (Definition 2.14), these issues will be resolved in section 8, where we will glue together the linear functions $\langle s(u), x\rangle$ to form a single linear function $\langle s, x\rangle$. The idea behind the gluing argument is to show that there is sufficient overlap between the supports of the measures $S_{\mathsf{P}_{u^\perp} B, \mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}}$ for different u so that all the vectors $s(u)$ must be consistent with a single vector s. It will turn out that the supercriticality assumption is preserved by the induction, so that a self-contained proof of Corollary 2.16 will already be achieved in section 8.

To complete the proof of Theorem 2.13 it remains to consider the critical case, that is, when $\dim(C_{i_1} + \cdots + C_{i_k}) = k + 1$ for some *critical set* $i_1 < \cdots < i_k$. It is in this situation that nontrivial degenerate functions (Definition 2.10) appear. The problem of gluing together these degenerate functions in dimension $n-1$ to form degenerate functions in dimension n gives rise to numerous complications. We begin in section 9 by characterizing what degenerate functions look like; they will turn out to be intimately connected to the critical sets. In section 10, we will show that in the critical case, the supports of the relevant mixed area measures exhibit a striking phenomenon: they form geometric structures that we call *propellers*, which are responsible for the formation of the degenerate extremals. We exploit these insights in section 11 to solve the gluing problem for degenerate functions. The proof of Theorem 2.13 is finally completed in section 12.

Part 1. The local Alexandrov-Fenchel inequality

5. Polytopes, graphs, and extremals

The aim of this section is to give a concrete formulation of the equality condition $S_{f,\mathcal{P}} = 0$ in the case that $\mathcal{P} = (P_1, \dots, P_{n-2})$ are polytopes. In particular, we will describe the underlying combinatorial structure, and introduce the basic objects and notations that will be used in the following sections.

5.1. Basic constructions.

We fix at the outset $n \geqslant 3$ and an arbitrary collection of polytopes $\mathcal{P} = (P_1, \dots, P_{n-2})$ in $\mathbb{R}^n$. The notations and definitions that are introduced below will be in force throughout sections 5–7.

5.1.1. The background polytope.

We begin by introducing a certain background polytope P that will be fixed throughout the following constructions.

Recall that a polytope in $\mathbb{R}^n$ is called *simple* if it has nonempty interior and each of its extreme points meets exactly n facets.

Lemma 5.1. *There exists a polytope P_0 in $\mathbb{R}^n$ so that $P_0 + P_1 + \cdots + P_n$ is simple.*

Proof. Let R be any polytope in $\mathbb{R}^n$ with nonempty interior, and define $Q := R + P_1 + \cdots + P_n$. Then by [30, Lemma 2.4.14], there exists a simple polytope Q' so that each normal cone of an extreme point of Q' is contained in the normal cone of an extreme point of Q. As the normal cones of $Q' + Q$ are intersections of normal cones of Q' and of Q [30, Theorem 2.2.1], it follows that the normal cones of the extreme points of $Q' + Q$ coincide with the normal cones of the extreme points of Q'. Thus $Q' + Q$ is also simple. The proof is concluded by choosing $P_0 = Q' + R$. □

In the sequel, we fix a polytope P_0 as in Lemma 5.1, and define

$$P := P_0 + P_1 + \cdots + P_{n-2}.$$

We will use P to construct a certain graph structure, on which the various objects that will be encountered in the sequel are defined.

Remark 5.2. In this section we will only use the fact that P is full-dimensional. The reason for choosing P to be simple will become apparent in section 6.

5.1.2. The background graph.

Let $P^1, \dots, P^N$ be the facets of P. We will frequently identify a facet P^i by its index $i \in [N]$. For each $i \in [N]$, we denote by $u_i \in S^{n-1}$ the outer unit normal vector of facet P^i.

Two facets $i, j \in [N]$ of P are said to be *neighboring* if they intersect in an $(n-2)$-dimensional face of P. We denote the set of such pairs as

$$E_P := \{(i,j) \in [N]^2 : \dim(P^i \cap P^j) = n-2\}.$$

For any $i \in [N]$, we will denote by

$$E_P^i := \{j \in [N] : (i,j) \in E_P\}$$

the set of facets that are neighbors of facet i. One should view $([N], E_P)$ as a graph whose vertices are facets of P and whose edges are neighboring facets.

As P is full-dimensional, the angle θ_{ij} between the vectors u_i and u_j must satisfy $0 < \theta_{ij} < \pi$ for any $(i,j) \in E_P$. Thus there is a unique shortest geodesic in the sphere between u_i and u_j, which we denote as $e_{ij} \subset S^{n-1}$; note that the length of e_{ij} is precisely θ_{ij}. Geometrically, e_{ij} is precisely the set of outer unit normal vectors of the $(n-2)$-dimensional face $P^i \cap P^j$ of P.

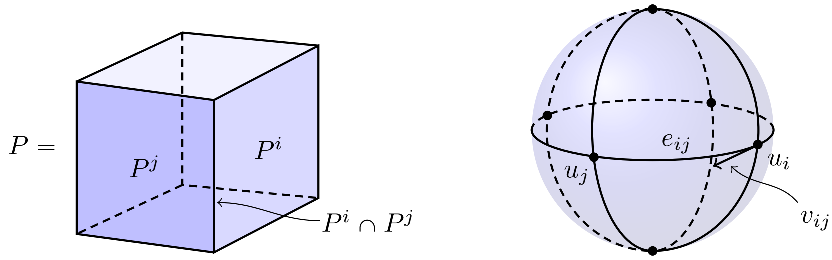

FIGURE 5.1. A polytope P in $\mathbb{R}^3$ and the associated geometric graph.

We further define for each $(i,j) \in E_P$ a vector $v_{ij} \in S^{n-1}$ such that $v_{ij} \perp u_i$ by

$$u_j =: u_i \cos\theta_{ij} + v_{ij} \sin\theta_{ij}.$$

Then v_{ij} is the unit tangent vector to e_{ij} at u_i, pointing toward u_j. Geometrically, if we view P^i as an $(n-1)$-dimensional convex body in aff P^i, then v_{ij} is precisely the outer unit normal vector of its facet $P^i \cap P^j = F(P^i, v_{ij})$.

The above definitions are illustrated in Figure 5.1. As is evident from the figure, one may naturally view these definitions as a geometric realization of the combinatorial graph $([N], E_P)$, whose vertices are the vectors $\{u_i\}_{i\in[N]}$ and whose edges are the geodesics $\{e_{ij}\}_{(i,j)\in E_P}$. We will often implicitly identify these viewpoints: we refer to both $i \in [N]$ and the associated vector u_i as a vertex, and to $(i,j) \in E_P$ and the associated geodesic e_{ij} as an edge, of the graph defined by P.

5.1.3. **Faces.** For any convex body C in $\mathbb{R}^n$ and $i \in [N]$, $j \in E_P^i$, we will denote

$$C^i := F(C, u_i), \qquad C^{ij} := F(C^i, v_{ij}).$$

We will frequently write $\mathcal{P}^i := (P_1^i, \ldots, P_{n-2}^i)$ and $\mathcal{P}^{ij} := (P_1^{ij}, \ldots, P_{n-2}^{ij})$, and analogously for other collections of bodies.

The notation C^i is consistent with the notation P^i for the facets of P, and we have $P^{ij} = P^i \cap P^j$. In particular, it follows from Lemma 3.6 that we can express the facets and $(n-2)$-faces of P in terms of faces of the polytopes P_r as

$$P^i = P_0^i + \cdots + P_{n-2}^i \qquad P^{ij} = P_0^{ij} + \cdots + P_{n-2}^{ij}.$$

In the sequel, we will apply these and similar consequences of the linearity of faces under Minkowski addition (Lemma 3.6) without further comment.

Note that P_r^i and P_r^{ij} are faces of the polytope P_r by definition. However, in contrast to the analogous faces of P, it is not necessarily the case that P_r^i is a facet and P_r^{ij} is an $(n-2)$-face of P_r. Nonetheless, the following lemma shows that the normal cone of the face P_r^{ij} of P_r always contains e_{ij}. In particular, it follows that $P_r^{ij} = P_r^{ji}$, which is not entirely obvious from the definition.

Lemma 5.3. *For every r, $i \in [N]$, $j \in E_P^i$, and $u \in \operatorname{relint} e_{ij}$, we have*

$$P_r^{ij} = F(P_r, u).$$

Proof. Recall that e_{ij} is the set of outer unit normal vectors of the face $P^i \cap P^j$ of P. But as any normal cone of a Minkowski sum $P = P_0 + \cdots + P_{n-2}$ of polytopes is contained in some normal cone of P_r for each r [30, Theorem 2.2.1], it follows that $F(P_r, u) = F(P_r, v) \subseteq F(P_r, w)$ for all $u, v \in \operatorname{relint} e_{ij}$ and $w \in e_{ij}$.

Choosing $w = u_i$, it follows that $F(P_r, u) \subseteq F(P_r, u_i) = P_r^i$. Thus $F(P_r, u)$ must be a face of P_r^i that has u as an outer normal vector, so that $F(P_r, u) \subseteq F(P_r^i, u)$. On the other hand, as $F(P_r^i, u)$ is a face of P_r with outer normal vector u, we must have $F(P_r^i, u) \subseteq F(P_r, u)$ as well. Thus we have shown $F(P_r, u) = F(P_r^i, u)$.

To conclude, note that as $u \in \operatorname{relint} e_{ij}$, we may write $u = au_i + bu_j$ for some $a, b > 0$, so that $\mathsf{P}_{u_i^\perp} u = cv_{ij}$ with $c = b \sin\theta_{ij} > 0$. It therefore follows from Lemma 3.7 that $F(P_r^i, u) = F(P_r^i, v_{ij}) = P_r^{ij}$, concluding the proof. $\square$

5.2. **The quantum graph.** The aim of this section is to describe the structure of the mixed area measure $S_{B,\mathcal{P}}$; this will be used in the next section to describe the extremal functions f such that $S_{f,\mathcal{P}} = 0$. It turns out that these objects are supported on a certain subgraph of the background graph defined by P in the previous section. We will rely on the formulation developed in [32], where the construction that arises here was called the "quantum graph". Related representations of mixed volumes and mixed area measures may be found in [8] and [30, p. 437].

Let us begin by describing the measure $S_{B,\mathcal{P}}$.

Lemma 5.4. *For every continuous function $f : S^{n-1} \to \mathbb{R}$, we have*

$$\int f\, dS_{B,\mathcal{P}} = \frac{1}{n-1} \sum_{(i,j)\in E_P : i<j} \mathsf{V}_{n-2}(P_1^{ij}, \ldots, P_{n-2}^{ij}) \int_{e_{ij}} f\, d\mathcal{H}^1,$$

where $\mathcal{H}^1$ is the 1-dimensional Hausdorff measure.

Proof. This is an immediate consequence of [32, Remark 5.11]. $\square$

Lemma 5.4 shows that $S_{B,\mathcal{P}}$ is supported on the edges $\{e_{ij}\}_{(i,j)\in E_P}$ of the geometric graph defined in the previous section. However, not every edge appears in the support: some of the weights $\mathsf{V}_{n-2}(\mathcal{P}^{ij})$ may be zero. Thus the collection of reference polytopes $\mathcal{P}$ defines a subgraph of the graph defined by P. Let us define some notation to describe this subgraph. In the sequel, we will write

$$\omega_{ij} := \mathsf{V}_{n-2}(P_1^{ij}, \ldots, P_{n-2}^{ij}).$$

The *active edges* of the graph defined by $\mathcal{P}$ are

$$E := \{(i,j) \in E_P : \omega_{ij} > 0\}.$$

Similarly, the *active vertices* of the graph defined by $\mathcal{P}$ are

$$V := \{i \in [N] : \textstyle\sum_{j\in[N]} \omega_{ij} > 0\},$$

that is, $i \in V$ when i is incident to at least one active edge $(i,j) \in E$. Denote by

$$E^i := \{j \in V : (i,j) \in E\}$$

the neighbors of $i \in V$ in the graph defined by $\mathcal{P}$.

We can now characterize the support of $S_{B,\mathcal{P}}$ as announced in Lemma 2.3.

Lemma 5.5. *The following are equivalent for any $u \in S^{n-1}$:*

a. $u \in \operatorname{supp} S_{B,\mathcal{P}}$.

b. $u \in e_{ij}$ *for some* $(i,j) \in E$.

c. There are segments $I_i \subseteq F(P_i, u)$, $i \in [n-2]$ with linearly independent directions.

d. $\dim(F(P_{i_1}, u) + \cdots + F(P_{i_k}, u)) \geqslant k$ *for all* $k \in [n-2]$, $1 \leqslant i_1 < \cdots < i_k \leqslant n-2$.

Proof. That $a \Leftrightarrow b$ is immediate by Lemma 5.4, while $c \Rightarrow d$ is trivial.

We now prove $b \Rightarrow c$. Suppose first that $u \in \operatorname{relint} e_{ij}$ for $(i,j) \in E$. Then $\mathsf{V}_{n-2}(F(P_1,u),\ldots,F(P_{n-2},u)) > 0$ by the definition of E and Lemma 5.3, which implies c by Lemma 2.2. It remains to consider the case where $u = u_i$ for some $i \in V$, so that $F(P_r,u) = P^i_r$. But by the definition of V, there exists j so that $\mathsf{V}_{n-2}(P^{ij}_1,\ldots,P^{ij}_{n-2}) > 0$, so c follows by Lemma 2.2 and $P^{ij}_r \subseteq P^i_r$.

It remains to prove $d \Rightarrow a$. To this end, suppose that d holds, and let Q be any polytope in $\mathbb{R}^n$ that has a facet with outer normal direction u. Then by Lemma 2.2, we have $\mathsf{V}_{n-1}(F(Q,u),F(P_1,u),\ldots,F(P_{n-2},u)) > 0$. Thus a follows as $u \in \operatorname{supp} S_{Q,\mathcal{P}} \subseteq \operatorname{supp} S_{B,\mathcal{P}}$ by Lemmas 3.4 and 2.4. $\square$

We now provide a useful description of the active vertices.

Lemma 5.6. *Let $i \in [N]$. Then the following hold:*

a. $i \in V$ if and only if $\mathsf{V}_{n-1}(P^i,P^i_1,\ldots,P^i_{n-2}) > 0$.

b. If $i \notin V$, then $\mathsf{V}_{n-1}(Q^i,P^i_1,\ldots,P^i_{n-2}) = 0$ for every polytope Q.

Proof. We may assume without loss of generality (by translation) that $P^i_1,\ldots,P^i_{n-2}$ are convex bodies in $u_i^\perp$ and that $0 \in \operatorname{relint} P^i$. As the facet normals of P^i in $u_i^\perp$ are precisely $\{v_{ij}\}_{j\in E_P}$, it follows from Lemma 3.4 that the mixed area measure $S_{P^i_1,\ldots,P^i_{n-2}}$ (computed in $u_i^\perp$) is supported on $\{v_{ij}\}_{j\in E_P}$, and that

$$S_{P^i_1,\ldots,P^i_{n-2}}(\{v_{ij}\}) = \mathsf{V}_{n-2}(F(P^i_1,v_{ij}),\ldots,F(P^i_{n-2},v_{ij})) = \omega_{ij}.$$

Thus (2.1) implies

$$\mathsf{V}_{n-1}(Q^i,P^i_1,\ldots,P^i_{n-2}) = \frac{1}{n-1}\sum_{j\in E^i_P} h_{Q^i}(v_{ij})\,\omega_{ij}$$

for any polytope Q, from which part b follows immediately. To prove part a, recall that P^i is a facet of P by definition, so our assumptions imply that P^i is a full-dimensional polytope in $u_i^\perp$ containing the origin in its interior. Thus $h_{P^i}(v_{ij}) > 0$ for all $j \in E^i_P$, and the conclusion of part a follows. $\square$

5.3. **The Alexandrov matrix.** The aim of this section is to give a combinatorial description of the equality cases of the Alexandrov-Fenchel inequality: we will show that the equality condition $S_{f,\mathcal{P}} = 0$ can be equivalently formulated in terms of separate conditions on the edges and vertices of the graph defined by $\mathcal{P}$. Such a characterization appears in the proof of [30, Theorem 7.6.21] for the special case that $P_1,\ldots,P_{n-2}$ are strongly isomorphic, which we do not assume here.

Define a symmetric matrix $\mathrm{A} \in \mathbb{R}^{N\times N}$ by

$$\mathrm{A}_{ij} := 1_{(i,j)\in E_P}\,\omega_{ij}\csc\theta_{ij} - 1_{i=j}\sum_{k\in E^i_P}\omega_{ik}\cot\theta_{ik}.$$

We will refer to A as the *Alexandrov matrix*, as a special case of this matrix arises in a much more restrictive setting (of strongly isomorphic polytopes) in Alexandrov's original proof of the Alexandrov-Fenchel inequality [1]. We now show that $S_{f,\mathcal{P}} = 0$ is equivalent to two conditions: f is (piecewise) linear on each edge in E, and the values of f on the vertices define a vector in the kernel of A.

Proposition 5.7. *Let $f : S^{n-1} \to \mathbb{R}$ be a difference of support functions. Then $S_{f,\mathcal{P}} = 0$ if and only if the following two conditions both hold:*

1. For every $(i,j) \in E$, there exists $t_{ij} \in \mathbb{R}^n$ such that $f(x) = \langle t_{ij}, x\rangle$ for $x \in e_{ij}$.
2. The vector $z := (f(u_i))_{i\in[N]} \in \mathbb{R}^N$ satisfies $z \in \ker \mathrm{A}$.

Proof. That $S_{f,\mathcal{P}} = 0$ may be equivalently stated as $\mathsf{V}_n(g,f,\mathcal{P}) = 0$ for every difference of support functions g. By [32, Theorem 5.1], this is equivalent to the statement that f lies in the kernel of the self-adjoint operator defined in [32, Theorem 5.7 and Remark 5.11], which is characterized by the following two conditions:

1. f is (piecewise) linear on each edge e_{ij} for $(i,j) \in E$.

2. f satisfies
$$\sum_{j\in E_P^i} \omega_{ij} \nabla_{v_{ij}} f(u_i) = 0 \quad \text{for every } i \in V.$$

It remains to show that the second condition is equivalent to $z \in \ker \mathrm{A}$. To this end, let us parametrize the edge e_{ij} as

$$e_{ij} = \{x(\theta) : 0 \leqslant \theta \leqslant \theta_{ij}\}, \qquad x(\theta) := u_i \cos\theta + v_{ij}\sin\theta.$$

By the first condition we can write $f(x) = \langle t, x\rangle$ on e_{ij} for some vector t, so

$$\begin{aligned} f(x(\theta)) &= \langle t, u_i\rangle \cos\theta + \langle t, v_{ij}\rangle \sin\theta \\ &= f(u_i)\cos\theta + \frac{f(u_j) - f(u_i)\cos\theta_{ij}}{\sin\theta_{ij}} \sin\theta. \end{aligned}$$

Consequently

$$\nabla_{v_{ij}} f(u_i) = \frac{d}{d\theta} f(x(\theta))\bigg|_{\theta=0} = \frac{f(u_j) - f(u_i)\cos\theta_{ij}}{\sin\theta_{ij}}.$$

Thus the second condition may be expressed equivalently as

$$0 = \sum_{j\in E_P^i} \omega_{ij} \frac{f(u_j) - f(u_i)\cos\theta_{ij}}{\sin\theta_{ij}} = (\mathrm{A}z)_i \quad \text{for all } i \in V.$$

But $(\mathrm{A}z)_i = 0$ always holds for $i \notin V$ by the definition of V, so we have shown that the second condition above is equivalent to $z \in \ker \mathrm{A}$. ☐

Remark 5.8. Instead of using the analytic theory of [32] as we have done here, one can give a more geometric proof by adapting the first part of the proof of [30, Theorem 7.6.21] to the present setting. Conditions 1 and 2 in the proof of Proposition 5.7 appear in [30] as (7.177) and (7.178), respectively.

Let us emphasize that the ith row and column of A are zero for every $i \notin V$. Thus the values $f(u_i)$ for $i \notin V$ never actually appear in Proposition 5.7. This simply reflects the fact that $\{u_i\}_{i\notin V}$ lie outside the support of $S_{B,\mathcal{P}}$, so these points play no role in the equality condition. Recall, however, that our ultimate aim is to prove the local Alexandrov-Fenchel inequality of Theorem 4.3, in which points outside the support of $S_{B,\mathcal{P}}$ play a crucial role. We therefore resist the temptation to simply remove the zero rows and columns from the definition of A at this stage.

6. Finite-dimensional reduction

The previous section introduced a combinatorial formulation of the equality condition $S_{f,\mathcal{P}} = 0$. In particular, Proposition 5.7 shows that an extremal function f is fully specified by its values $f(u_i)$ on the vertices u_i of the graph defined by $\mathcal{P}$: its values on the rest of the support of $S_{B,\mathcal{P}}$ are then uniquely determined by linearity. In order to prove the local Alexandrov-Fenchel inequality, however, we will also need to reason about the measure $S_{f,f,\mathcal{P}_{\backslash r}}$, and there is no reason to expect that only the directions $\{u_i\}_{i\in[N]}$ will appear in its description.

The aim of this section is to introduce a basic geometric construction that will enable us to surmount this issue. This construction will simultaneously serve two purposes: it will enable us to reduce the local Alexandrov-Fenchel inequality to a finite-dimensional problem, and it will furnish the objects that appear in Proposition 5.7 with a geometric interpretation that will be key to their analysis.

In this section, all assumptions and definitions of section 5 will be in force.

6.1. Strongly isomorphic polytopes.

We begin by recalling the definition.

Definition 6.1. Polytopes Q, Q' are said to be *strongly isomorphic* if

$$\dim F(Q,u) = \dim F(Q',u) \quad \text{for all } u \in S^{n-1}.$$

The key feature of strongly isomorphic polytopes Q, Q' is that they have identical facial structures: there is a bijection between the faces of Q and Q' such that each pair of identified faces has the same normal cone [30, §2.4]. Consequently, no new faces are created when one takes Minkowski sums of strongly isomorphic polytopes. Let us record this basic fact for future reference [30, Corollary 2.4.12].

Lemma 6.2. *Let Q, Q' be polytopes. Then all the polytopes $\lambda Q + \lambda' Q'$ with $\lambda, \lambda' > 0$ are strongly isomorphic. If Q, Q' are themselves strongly isomorphic, then all the polytopes $\lambda Q + \lambda' Q'$ with $\lambda, \lambda' \geqslant 0$ are strongly isomorphic.*

The following simple observation will play an important role in the sequel.

Lemma 6.3. *Let Q be a polytope that is strongly isomorphic to P. Then for every $(i,j) \in E_P$, there exists $t_{ij} \in \mathbb{R}^n$ such that $h_Q(x) = \langle t_{ij}, x\rangle$ for $x \in e_{ij}$.*

Proof. Let $(i,j) \in E_P$. As Q and P are strongly isomorphic, e_{ij} is the set of unit normal vectors to the face Q^{ij} of Q. Thus $Q^{ij} = F(Q,u)$ for any $u \in \operatorname{relint} e_{ij}$. If we therefore fix any $t_{ij} \in Q^{ij}$, then $h_Q(u) = \langle t_{ij}, u\rangle$ for all $u \in \operatorname{relint} e_{ij}$ by (2.2), and the conclusion extends to the endpoints of e_{ij} by continuity. □

The significance of Lemma 6.3 is immediately evident from Proposition 5.7: when Q is strongly isomorphic to P, the function $f = h_Q - h_P$ automatically satisfies the piecewise linearity condition that characterizes the extremals of the Alexandrov-Fenchel inequality on the edges of the graph defined by $\mathcal{P}$ (note that as P is strongly isomorphic to itself, Lemma 6.3 also applies to $Q = P$). We will shortly prove a strong converse to this statement: any extremal function f of the Alexandrov-Fenchel inequality may be represented in such a form.

6.2. Support vectors.

As is already anticipated by Proposition 5.7, we will frequently work with the restriction of support functions of convex bodies to the finite collection of directions $\{u_i\}_{i\in[N]}$. It will be convenient to introduce the following notation: for any convex body C in $\mathbb{R}^n$, we define its *support vector* $\mathrm{h}_C \in \mathbb{R}^N$ by

$$(\mathrm{h}_C)_i := h_C(u_i), \qquad i \in [N].$$

The following result shows that any vector $z \in \mathbb{R}^N$ can be expressed in terms of the support vector of a polytope that is strongly isomorphic to P. It is here that we make crucial use of the fact that P was chosen to be a simple polytope.

Lemma 6.4. *For any vector $z \in \mathbb{R}^N$, there exists a polytope Q that is strongly isomorphic to P and a scalar $a \in \mathbb{R}$ such that $z = \mathrm{h}_Q - a\mathrm{h}_P$.*

Proof. For any $y \in \mathbb{R}^N$, define

$$Q_y := \bigcap_{i\in[N]} \{x \in \mathbb{R}^n : \langle u_i, x\rangle \leqslant h_P(u_i) + y_i\}.$$

As P is a simple polytope, it follows from [30, Lemma 2.4.13] that there exists $\varepsilon > 0$ such that Q_y is strongly isomorphic to P whenever $\|y\|_\infty \leqslant \varepsilon$. In particular, we then have $\mathrm{h}_{Q_y} = \mathrm{h}_P + y$ as Q_y and P have the same facet normals. The conclusion follows by choosing $Q := aQ_{z/a}$ with $a := \varepsilon^{-1}(1 + \|z\|_\infty)$. □

We can now explain a key implication of the above construction: it enables us to modify any equality case of the Alexandrov-Fenchel inequality outside the support of $S_{B,\mathcal{P}}$ in such a way that the relevant mixed area measures are supported in the finite set $\{u_i\}_{i\in[N]}$. It is by virtue of this procedure that we will be able to reduce Theorem 4.3 to a finite-dimensional problem.

Corollary 6.5. *Let f be a difference of support functions so that $S_{f,\mathcal{P}} = 0$. Then there is a polytope Q that is strongly isomorphic to P and $a \in \mathbb{R}$ so that $g = h_Q - ah_P$ satisfies $g = f$ $S_{B,\mathcal{P}}$-a.e., $S_{g,\mathcal{P}} = 0$, and $S_{g,g,\mathcal{P}_{\backslash r}}$ is supported on $\{u_i\}_{i\in[N]}$ for all r.*

Proof. Choose any $z \in \mathbb{R}^N$ such that $z_i = f(u_i)$ for $i \in V$. Applying Lemma 6.4, we find a polytope Q that is strongly isomorphic to P and $a \in \mathbb{R}$ so that $g = h_Q - ah_P$ satisfies $g(u_i) = f(u_i)$ for all $i \in V$. Moreover, f is linear on e_{ij} for every $(i,j) \in E$ by Proposition 5.7, while g satisfies the same property by Lemma 6.3. Thus $f = g$ $S_{B,\mathcal{P}}$-a.e. by Lemma 5.5. That $S_{g,\mathcal{P}} = S_{f,\mathcal{P}} = 0$ now follows by Lemma 2.8. Finally, as the facet normals of $Q + P$ are $\{u_i\}_{i\in[N]}$ by Lemma 6.2, we can conclude that $S_{g,g,\mathcal{P}_{\backslash r}}$ is supported in this set for any r by Lemma 3.4. □

It should be emphasized that Corollary 6.5 does not in itself capture any aspect of the phenomenon described by Theorem 4.3: it merely reduces the problem to a finite universe $\{u_i\}_{i\in[N]}$ of normal directions, but does not otherwise guarantee any properties of the measure $S_{g,g,\mathcal{P}_{\backslash r}}$. On the other hand, we have considerable freedom in the construction of g in Corollary 6.5: we have only specified $g(u_i)$ for $i \in V$ in the proof, and we are therefore free to choose arbitrary values of $g(u_i)$ for $i \notin V$. What we must show in the proof of the local Alexandrov-Fenchel inequality is that there exists a choice of the latter values which ensures that $S_{g,g,\mathcal{P}_{\backslash r}} \leqslant 0$.

6.3. **The Alexandrov matrix revisited.** We now show that strongly isomorphic polytopes enable us to furnish the Alexandrov matrix of Proposition 5.7 with a geometric interpretation. To this end, it will be useful to introduce the following notation. For any $i \in [N]$, define a linear map $D_i : \mathbb{R}^N \to \mathbb{R}^{E_P^i}$ by

$$(D_i z)_j := z_j \csc\theta_{ij} - z_i \cot\theta_{ij}, \qquad j \in E_P^i,\ i \in [N],\ z \in \mathbb{R}^N.$$

The significance of this definition is the following.

Lemma 6.6. *Let Q be a polytope that is strongly isomorphic to P. Then*

$$(D_i \mathrm{h}_Q)_j = h_{Q^i}(v_{ij}) \qquad \text{for all } (i,j) \in E_P.$$

Proof. Fix $(i,j) \in E_P$. As Q is strongly isomorphic to P, it must be the case that $F(Q^i, v_{ij}) = Q^i \cap Q^j$. Thus for fixed $x \in Q^i \cap Q^j$, we have

$$\langle x, u_i \rangle = h_Q(u_i), \qquad \langle x, u_j \rangle = h_Q(u_j), \qquad \langle x, v_{ij} \rangle = h_{Q^i}(v_{ij}).$$

Taking the inner product with x in the definition of v_{ij} yields

$$h_Q(u_j) = h_Q(u_i) \cos\theta_{ij} + h_{Q^i}(v_{ij}) \sin\theta_{ij},$$

and the conclusion follows by rearranging this expression. □

As a consequence, we obtain the following geometric interpretation.

Corollary 6.7. *Let Q be strongly isomorphic to P and let $a \in \mathbb{R}$. Denote*

$$z := \mathrm{h}_Q - a\mathrm{h}_P, \qquad f := h_Q - a h_P, \qquad f^i := h_{Q^i} - a h_{P^i}.$$

Then for every $i \in [N]$

$$(\mathrm{A}z)_i = (n-1)\,\mathsf{V}_{n-1}(f^i, P^i_1, \ldots, P^i_{n-2}),$$

and for any convex body C

$$\langle \mathrm{h}_C, \mathrm{A}z \rangle = n(n-1)\,\mathsf{V}_n(C, f, P_1, \ldots, P_{n-2}).$$

Proof. It was shown in the proof of Lemma 5.6 that $\omega_{ij} = S_{P^i_1,\ldots,P^i_{n-2}}(\{v_{ij}\})$. We may therefore rewrite the definition of A as

$$(\mathrm{A}z)_i = \sum_{j \in E^i_P} (D_i z)_j \omega_{ij} = \sum_{j \in E^i_P} f^i(v_{ij})\, S_{P^i_1,\ldots,P^i_{n-2}}(\{v_{ij}\})$$

using Lemma 6.6. But as $\{v_{ij}\}_{j \in E^i_P}$ are the facet normals of P^i in aff P^i, we obtain

$$(\mathrm{A}z)_i = \int f^i\, dS_{P^i_1,\ldots,P^i_{n-2}} = (n-1)\,\mathsf{V}_{n-1}(f^i, P^i_1, \ldots, P^i_{n-2})$$

by Lemma 3.4 and (2.1). Now note that as Q is strongly isomorphic to P, the facet normals of $Q + P$ are $\{u_i\}_{i \in [N]}$ by Lemma 6.2. Thus

$$\begin{aligned}\langle \mathrm{h}_C, \mathrm{A}z \rangle &= (n-1) \sum_{i \in [N]} h_C(u_i)\,\mathsf{V}_{n-1}(f^i, P^i_1, \ldots, P^i_{n-2}) \\ &= (n-1) \int h_C\, dS_{f, P_1, \ldots, P_{n-2}} = n(n-1)\,\mathsf{V}_n(C, f, P_1, \ldots, P_{n-2})\end{aligned}$$

by Lemma 3.4 and (2.1). □

7. Proof of the local Alexandrov-Fenchel inequality

Now that the requisite machinery is in place, we proceed to the main part of the proof of Theorem 4.3. Throughout this section, all assumptions and definitions of sections 5 and 6 will be assumed without further comment.

Let us begin by reformulating Theorem 4.3 in a more combinatorial manner.

Theorem 7.1. *Assume that $\mathcal{P} = (P_1, \ldots, P_{n-2})$ is a critical collection of polytopes. Fix $r \in [n-2]$ and $z \in \ker \mathrm{A}$. Then there exist a polytope Q that is strongly isomorphic to P and $a \in \mathbb{R}$ such that the following hold:*

1. $(\mathrm{h}_Q - a\mathrm{h}_P)_i = z_i$ *for every* $i \in V$.
2. $\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r}) \leqslant 0$ *for every* $i \in [N]$.

With this result in hand, the conclusion of Theorem 4.3 follows readily:

Proof of Theorem 4.3. Fix $r \in [n-2]$ and a difference of support functions f so that $S_{f,\mathcal{P}} = 0$. Then $z := (f(u_i))_{i\in[N]}$ satisfies $z \in \ker \mathrm{A}$ by Proposition 5.7. We can therefore apply Theorem 7.1 to construct an associated polytope Q and $a \in \mathbb{R}$. We claim that $g := h_Q - ah_P$ satisfies the conclusion of Theorem 4.3.

To show this, note first that it follows exactly as in the proof of Corollary 6.5 that $g = f$ $S_{B,\mathcal{P}}$-a.e., that $S_{g,\mathcal{P}} = 0$, and that $S_{g,g,\mathcal{P}_{\backslash r}}$ is supported on $\{u_i\}_{i\in[N]}$. On the other hand, Lemma 3.4 implies that

$$S_{g,g,\mathcal{P}_{\backslash r}}(\{u_i\}) = \mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r}).$$

Thus the second property of Theorem 7.1 implies $S_{g,g,\mathcal{P}_{\backslash r}} \leqslant 0$. $\square$

The rest of this section is devoted to the proof of Theorem 7.1. The proof consists of two parts. First, we will show that the second property of Theorem 7.1 holds automatically for $i \in V$ by the Alexandrov-Fenchel inequality. We then show that Q, a can be chosen in such a way that this property holds also for $i \notin V$.

7.1. **The active vertices.** The first observation of the proof of Theorem 7.1 is that its second condition is automatically satisfied for the active vertices $i \in V$ whenever the first condition is satisfied, regardless of how Q is chosen.

Lemma 7.2. *Let Q be a polytope that is strongly isomorphic to P and let $a \in \mathbb{R}$. Suppose that $\mathrm{h}_Q - a\mathrm{h}_P \in \ker \mathrm{A}$. Then for any $r \in [n-2]$, we have*

$$\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r}) \leqslant 0 \quad \textit{for all } i \in V.$$

Proof. By Corollary 6.7, the assumption $\mathrm{h}_Q - a\mathrm{h}_P \in \ker \mathrm{A}$ implies

$$\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, \mathcal{P}^i) = 0 \quad \text{for all } i \in [N].$$

By the Alexandrov-Fenchel inequality

$$0 \geqslant \mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r})\, \mathsf{V}_{n-1}(P^i_r, P^i_r, \mathcal{P}^i_{\backslash r}).$$

The conclusion follows immediately for any i such that $\mathsf{V}_{n-1}(P^i_r, P^i_r, \mathcal{P}^i_{\backslash r}) > 0$.

Now suppose $i \in V$ but $\mathsf{V}_{n-1}(P^i_r, P^i_r, \mathcal{P}^i_{\backslash r}) = 0$. Then we can argue in a similar manner as in the proof of the second part of Lemma 3.12. As $i \in V$, Lemma 5.6 states that $\mathsf{V}_{n-1}(P^i, P^i_r, \mathcal{P}^i_{\backslash r}) > 0$. Thus we may choose $b \in \mathbb{R}$ so that

$$\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i} - bh_{P^i_r}, P^i, \mathcal{P}^i_{\backslash r}) = 0.$$

The Alexandrov-Fenchel inequality now yields

$$0 \geqslant \mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i} - bh_{P^i_r}, h_{Q^i} - ah_{P^i} - bh_{P^i_r}, \mathcal{P}^i_{\backslash r})\, \mathsf{V}_{n-1}(P^i, P^i, \mathcal{P}^i_{\backslash r}).$$

But $i \in V$ implies $\mathsf{V}_{n-1}(P^i, P^i, \mathcal{P}^i_{\backslash r}) \geqslant \mathsf{V}_{n-1}(P^i, \mathcal{P}^i) > 0$ by Lemma 5.6. Thus

$$\begin{aligned} 0 &\geqslant \mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i} - bh_{P^i_r}, h_{Q^i} - ah_{P^i} - bh_{P^i_r}, \mathcal{P}^i_{\backslash r}) \\ &= \mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r}), \end{aligned}$$

where we used that $\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, \mathcal{P}^i) = \mathsf{V}_{n-1}(P^i_r, P^i_r, \mathcal{P}^i_{\backslash r}) = 0$. $\square$

Now consider the setting of Theorem 7.1 for a given $z \in \ker \mathrm{A}$. As the ith row and column of A are zero for $i \notin V$, we have $z' \in \ker \mathrm{A}$ whenever $z_i = z'_i$ for $i \in V$. To any such choice of z', we can apply Lemma 6.4 to obtain a polytope Q that is strongly isomorphic to P and $a \in \mathbb{R}$ so that $z' = \mathrm{h}_Q - a\mathrm{h}_P$. Then:

1. $(\mathrm{h}_Q - a\mathrm{h}_P)_i = z_i$ for every $i \in V$ (as $z_i = z'_i$ for $i \in V$).
2. $\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r}) \leqslant 0$ for every $i \in V$ (by Lemma 7.2).

Thus the only part of the proof of Theorem 7.1 that remains is to ensure that the second condition holds for $i \notin V$. On the other hand, in the above construction, the choice of z'_i for $i \notin V$ is completely arbitrary.

The present discussion provides us with a key intuition about why the local Alexandrov-Fenchel inequality has any hope of being true: we aim to satisfy $N - |V|$ nontrivial equations, but we are free to choose $N - |V|$ parameters. In other words, *the number of degrees of freedom equals the number of equations we aim to satisfy.* This fact is not at all obvious from the formulation of Theorem 4.3.

On the other hand, this idea alone cannot suffice to complete the proof: it is possible that the system of equations we aim to solve is degenerate, in which case no solution may exist. It is far from obvious, *a priori*, why this situation cannot occur for some special choices of polytopes: had that been the case, there would have likely existed additional extremals of the Alexandrov-Fenchel inequality beyond the ones discussed in section 2. The main difficulty in the remainder of the proof of Theorem 7.1 is to rule out the existence of such degeneracies.

7.2. **Reduction to a linear system.** As was explained above, we now aim to choose the polytope Q in such a way that the second condition of Theorem 7.1 holds for $i \notin V$. In essence, this requires us to find a solution to a system of quadratic inequalities. The manipulation of these inequalities is somewhat awkward, however, so we begin by introducing a simplification: we will reduce the problem to solving a system of linear equations, which are formulated in the following result.

Proposition 7.3. *Assume that $\mathcal{P} = (P_1, \ldots, P_{n-2})$ is a critical collection of polytopes. Fix $r \in [n-2]$ and $z \in \mathbb{R}^N$. Then there exist a polytope Q that is strongly isomorphic to P and $a \in \mathbb{R}$ such that the following hold:*

1. $(\mathrm{h}_Q - a\mathrm{h}_P)_i = z_i$ *for every* $i \in V$*.*

2. $\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, P^i, \mathcal{P}^i_{\backslash r}) = 0$ *for every* $i \notin V$*.*

Proposition 7.3 will be proved in the next section. Before we do so, let us show that it implies Theorem 7.1. As in Lemma 7.2, the transition from linear equations to quadratic inequalities is a consequence of the Alexandrov-Fenchel inequality.

Proof of Theorem 7.1. Fix $r \in [n-2]$ and $z \in \ker \mathrm{A}$, and construct the polytope Q as in Proposition 7.3. Then the first condition of Theorem 7.1 holds by construction. Moreover, as the ith column of A vanishes for $i \notin V$, it follows that $\mathrm{h}_Q - a\mathrm{h}_P \in \ker \mathrm{A}$. Thus the second condition of Theorem 7.1 holds for $i \in V$ by Lemma 7.2.

Now let $i \notin V$. Then by Proposition 7.3 and the Alexandrov-Fenchel inequality

$$\begin{aligned} 0 &= \mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, P^i, \mathcal{P}^i_{\backslash r})^2 \\ &\geqslant \mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r})\,\mathsf{V}_{n-1}(P^i, P^i, \mathcal{P}^i_{\backslash r}). \end{aligned}$$

Thus the second condition of Theorem 7.1 holds provided $\mathsf{V}_{n-1}(P^i, P^i, \mathcal{P}^i_{\backslash r}) > 0$.

It remains to consider $i \in [N]$ such that $\mathsf{V}_{n-1}(P^i, P^i, \mathcal{P}^i_{\backslash r}) = 0$. By definition, P^i are the facets of P, so $\dim P^i = n-1$. It therefore follows from Lemma 2.2 that $\mathsf{V}_{n-1}(K^i, L^i, \mathcal{P}^i_{\backslash r}) = 0$ for any convex bodies K, L. In particular, for such i

$$\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, h_{Q^i} - ah_{P^i}, \mathcal{P}^i_{\backslash r}) = 0.$$

Thus the second condition of Theorem 7.1 is established for every $i \in [N]$. $\square$

To clarify the computations in the next section, let us further express the linear system of Proposition 7.3 explicitly in a finite-dimensional form. To this end, we would like to represent the mixed volume $\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, P^i, \mathcal{P}^i_{\backslash r})$ in terms of an Alexandrov matrix. In the present setting, however, the reference body P_r has been replaced by P, so that Corollary 6.7 does not directly apply.

Note, however, that $P + \sum_{i \notin r} P_i$ is strongly isomorphic to P by Lemma 6.2. Therefore, if we replace the reference bodies $\mathcal{P}$ by $(P, \mathcal{P}_{\backslash r})$, then the background graph defined in section 5.1 remains unchanged, and all the subsequent constructions in sections 5–6 extend *verbatim* to this setting up to a change of notation. In particular, if we define the Alexandrov matrix associated to $(P, \mathcal{P}_{\backslash r})$ as

$$\bar{\mathrm{A}}_{ij} := 1_{(i,j)\in E_P} \mathsf{V}_{n-2}(P^{ij}, \mathcal{P}^{ij}_{\backslash r}) \csc\theta_{ij} - 1_{i=j} \sum_{k \in E^i_P} \mathsf{V}_{n-2}(P^{ik}, \mathcal{P}^{ik}_{\backslash r}) \cot\theta_{ik},$$

then Corollary 6.7 extends immediately to show that whenever $z' = \mathrm{h}_Q - a\mathrm{h}_P$ for a polytope Q that is strongly isomorphic to P and $a \in \mathbb{R}$, we have

$$(\bar{\mathrm{A}}z')_i = (n-1)\,\mathsf{V}_{n-1}(h_{Q^i} - ah_{P^i}, P^i, \mathcal{P}^i_{\backslash r}),$$
$$\langle \mathrm{h}_C, \bar{\mathrm{A}}z' \rangle = n(n-1)\,\mathsf{V}_n(C, h_Q - ah_P, P, \mathcal{P}_{\backslash r})$$

for any $i \in [N]$ and convex body C. If we can therefore show that the linear system

$$\begin{cases} z'_i = z_i & \text{for } i \in V, \\ (\bar{\mathrm{A}}z')_i = 0 & \text{for } i \notin V \end{cases} \tag{7.1}$$

has a solution $z' \in \mathbb{R}^N$, the proof of Proposition 7.3 would follow from Lemma 6.4.

7.3. **The Fredholm alternative.** We are now ready to complete the proof of Proposition 7.3. To show that the linear system (7.1) has a solution, we will verify the dual condition provided by the Fredholm alternative $\operatorname{ran}\mathrm{M} = (\ker \mathrm{M}^*)^\perp$ of linear algebra. Surprisingly, it will turn out that the validity of this dual condition is itself a consequence of the equality condition of the Alexandrov-Fenchel inequality.

Proof of Proposition 7.3. We fix $r \in [n-2]$ and $z \in \mathbb{R}^N$ throughout the proof. Let us begin by rewriting the linear system (7.1) as a single equation. Let $V^c := [N] \backslash V$, and denote by P_V and P_{V^c} the orthogonal projections onto the subspaces of vectors supported on the coordinates V and V^c, respectively. Then clearly (7.1) has a solution $z' \in \mathbb{R}^N$ if and only if there exists $y \in \mathbb{R}^N$ such that

$$\mathrm{P}_{V^c}\bar{\mathrm{A}}\mathrm{P}_{V^c}y = -\mathrm{P}_{V^c}\bar{\mathrm{A}}\mathrm{P}_V z \tag{7.2}$$

(as then $z' = \mathrm{P}_{V^c}y + \mathrm{P}_V z$ is a solution to (7.1)).

To show there exists a solution to (7.2), we will prove the following claim:

$$\mathrm{P}_{V^c}w \in \ker\bar{\mathrm{A}} \quad \text{for every} \quad w \in \ker \mathrm{P}_{V^c}\bar{\mathrm{A}}\mathrm{P}_{V^c}. \tag{7.3}$$

Let us first argue that this suffices to conclude the proof. If (7.3) holds, then we clearly have $\langle w, \mathrm{P}_{V^c}\bar{\mathrm{A}}\mathrm{P}_V z\rangle = \langle \bar{\mathrm{A}}\mathrm{P}_{V^c}w, \mathrm{P}_V z\rangle = 0$ for every $w \in \ker \mathrm{P}_{V^c}\bar{\mathrm{A}}\mathrm{P}_{V^c}$.

The latter is precisely the dual condition for the existence of a solution y to (7.2). It therefore follows that there exists $z' \in \mathbb{R}^N$ satisfying (7.1), and the proof of Proposition 7.3 is concluded as explained at the end of the previous section.

It therefore remains to prove (7.3). To this end, let us fix $w \in \ker \mathrm{P}_{V^c}\bar{\mathrm{A}}\mathrm{P}_{V^c}$. By Lemma 6.4, there exists a polytope R that is strongly isomorphic to P and $b \in \mathbb{R}$ such that $\mathrm{P}_{V^c}w = \mathrm{h}_R - b\mathrm{h}_P$. We can therefore compute

$$\mathsf{V}_n(h_R - bh_P, h_R - bh_P, P, \mathcal{P}_{\backslash r}) = \frac{\langle \mathrm{P}_{V^c}w, \bar{\mathrm{A}}\mathrm{P}_{V^c}w\rangle}{n(n-1)} = 0.$$

On the other hand, we have

$$\begin{aligned}\mathsf{V}_n(h_R - bh_P, P_r, P, \mathcal{P}_{\backslash r}) &= \frac{1}{n}\sum_{i\in[N]} (\mathrm{h}_R - b\mathrm{h}_P)_i \,\mathsf{V}_{n-1}(P_r^i, P^i, \mathcal{P}^i_{\backslash r}) \\ &= \frac{1}{n}\sum_{i\in V^c} (\mathrm{h}_R - b\mathrm{h}_P)_i \,\mathsf{V}_{n-1}(P^i, \mathcal{P}^i) = 0,\end{aligned}$$

where the first equality follows from Lemma 3.4 and (2.1), the second equality follows as $(\mathrm{h}_R - b\mathrm{h}_P)_i = (\mathrm{P}_{V^c}w)_i = 0$ for $i \in V$, and the third equality follows as $\mathsf{V}_{n-1}(P^i, \mathcal{P}^i) = 0$ for $i \in V^c$ by Lemma 5.6. Finally, we have

$$\mathsf{V}_n(P_r, P_r, P, \mathcal{P}_{\backslash r}) > 0$$

using that $\mathcal{P}$ is critical (Definition 4.1) and Lemma 2.2. Thus Lemma 3.12 yields

$$0 = S_{h_R - bh_P, P, \mathcal{P}_{\backslash r}}(\{u_i\}) = \mathsf{V}_{n-1}(h_{R^i} - bh_{P^i}, P^i, \mathcal{P}^i_{\backslash r}) = \frac{(\bar{\mathrm{A}}\mathrm{P}_{V^c}w)_i}{n-1}$$

for every $i \in [N]$, where we used Lemma 3.4 in the second equality. In other words, we have shown that $\mathrm{P}_{V^c}w \in \ker \bar{\mathrm{A}}$, concluding the proof of (7.3). □

Remark 7.4. Let us emphasize that the definition of the matrix $\bar{\mathrm{A}}$ depends on the choice of r, so that the polytope Q and $a \in \mathbb{R}$ that are constructed in the proof of Proposition 7.3 will generally depend on r. This will not be a problem for our purposes, however, as we will fix r when we implement the induction argument.

Part 2. **Gluing**

8. The supercritical case

The aim of this section is to complete our characterization of the extremals of the Alexandrov-Fenchel inequality in the supercritical case (Definition 2.14).

Theorem 8.1. *Let $\mathcal{P} := (P_1, \ldots, P_{n-2})$ be a supercritical collection of polytopes in $\mathbb{R}^n$ $(n \geqslant 2)$. For any difference of support functions $f : S^{n-1} \to \mathbb{R}$, we have $S_{f,\mathcal{P}} = 0$ if and only if there exists $s \in \mathbb{R}^n$ so that $f(x) = \langle s, x\rangle$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$.*

Let us note that Theorem 8.1 is simply a reformulation of Corollary 2.16.

Proof of Corollary 2.16. By Lemma 2.5, the equality condition in Corollary 2.16 holds if and only if there exists $a > 0$ such that $S_{f,\mathcal{P}} = 0$ for $f = h_K - ah_L$. The conclusion now follows immediately from Theorem 8.1. □

Remark 8.2. We fixed at the beginning of this paper (section 2.1) $n \geqslant 3$, which has been assumed throughout without further comment. In dimension $n = 2$, the collection $\mathcal{P}$ is empty and the Alexandrov-Fenchel inequality reduces to Minkowski's first inequality [30, Theorem 7.2.1] whose equality cases are elementary. The case $n = 2$ does play a role in this paper, however, as it will be used as the base case for our induction arguments. For this reason, we have formulated Theorem 8.1 for $n \geqslant 2$. Note that the $n = 2$ case is always supercritical by definition.

Most of this section will be devoted to the proof of the induction step. We therefore fix until further notice $n \geqslant 3$ and a supercritical collection of polytopes in $\mathbb{R}^n$. By translation-invariance of mixed area measures, the equality condition $S_{f,\mathcal{P}} = 0$ is invariant under translation of the polytopes in $\mathcal{P}$, so there is no loss of generality in assuming that P_i *contains the origin in its relative interior for every* $i \in [n-2]$. Consequently, if we define for every $\alpha \subseteq [n-2]$ the linear space

$$\mathcal{L}_\alpha := \operatorname{span}\{P_i : i \in \alpha\} = \operatorname{span} \textstyle\sum_{i\in\alpha} P_i,$$

then $\dim \sum_{i\in\alpha} P_i = \dim \mathcal{L}_\alpha$ for any $\alpha \subseteq [n-2]$. We will denote by B_α the Euclidean unit ball in $\mathcal{L}_\alpha$, and we write $\mathcal{L}_r := \mathcal{L}_{\{r\}}$, $B_r := B_{\{r\}}$.

The above assumptions and notation will be assumed in the sequel without further comment. In particular, note that the supercriticality assumption may now be formulated as $\dim \mathcal{L}_\alpha \geqslant |\alpha| + 2$ for every $\alpha \subseteq [n-2]$, $\alpha \neq \varnothing$. Let us also note the simple identity $\mathcal{L}_{\alpha\cup\beta} = \mathcal{L}_\alpha + \mathcal{L}_\beta$ that will be used many times.

8.1. **The induction hypothesis.** The proof of Theorem 8.1 proceeds by induction on n: in the induction step, we will assume the theorem has been proved in dimension $n-1$, and deduce its validity in dimension n. The aim of this section is to formulate the resulting induction hypothesis. To this end, let us begin by stating a consequence of the local Alexandrov-Fenchel inequality.

Lemma 8.3. *Fix $r \in [n-2]$ and a difference of support functions f with $S_{f,\mathcal{P}} = 0$. Then there exists a difference of support functions g with the following properties:*

1. *$g(x) = f(x)$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$.*
2. *$\mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} P_r, \mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}) = 0$ for all $u \in S^{n-1}$.*
3. *$\mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}) = 0$ for all $u \in S^{n-1} \cap \mathcal{L}_r$.*
4. *$\mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} P_r, \mathbf{P}_{u^\perp} P_r, \mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}) > 0$ for all $u \in S^{n-1}$.*

Here the projections $\mathbf{P}_{u^\perp} g$, $\mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}$ are as defined in section 4.1.

Proof. By Theorem 4.3, there exists $g = f$ $S_{B,\mathcal{P}}$-a.e. such that $S_{g,\mathcal{P}} = 0$ and $S_{g,g,\mathcal{P}_{\backslash r}} \leqslant 0$. Let us check that each of the claimed properties holds for g. The first property holds by construction. To prove the second property, note that

$$0 = \int h_{[0,u]}\, dS_{g,\mathcal{P}} = \mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} P_r, \mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}),$$

where we have used Corollary 3.9 and (2.1).

The third property is analogous to Lemma 4.2, but in the present case we cannot assume that P_r is full-dimensional. We first note that as $S_{g,\mathcal{P}} = 0$, we have

$$0 = \int g\, dS_{g,\mathcal{P}} = \int h_{P_r}\, dS_{g,g,\mathcal{P}_{\backslash r}}$$

using (2.1) and the symmetry of mixed volumes. On the other hand, as $S_{g,g,\mathcal{P}_{\backslash r}} \leqslant 0$ by construction, it follows that $1_{h_{P_r}>0} dS_{g,g,\mathcal{P}_{\backslash r}} = 0$. Now note that as we assumed $0 \in \operatorname{relint} P_r$, there exists $\varepsilon > 0$ so that $\varepsilon[0,u] \subseteq P_r$ for every $u \in S^{n-1} \cap \mathcal{L}_r$. In particular, this implies $\varepsilon h_{[0,u]} \leqslant h_{P_r}$ and thus $\{x : h_{[0,u]}(x) > 0\} \subseteq \{x : h_{P_r}(x) > 0\}$ whenever $u \in S^{n-1} \cap \mathcal{L}_r$. We can therefore conclude that

$$0 = \int h_{[0,u]}\, dS_{g,g,\mathcal{P}_{\backslash r}} = \mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r})$$

for any $u \in S^{n-1} \cap \mathcal{L}_r$ using Corollary 3.9 and (2.1).

It remains to verify the fourth property, which is a consequence of the supercriticality assumption. As $\dim(\sum_{i\in\alpha} P_i) \geqslant |\alpha| + 2$ for all $\alpha \neq \varnothing$, it follows readily that $\dim(\sum_{i\in\alpha} \mathbf{P}_{u^\perp} P_i) \geqslant |\alpha| + 1$ for $\alpha \neq \varnothing$, and thus also $\dim(\mathbf{P}_{u^\perp} P_r + \sum_{i\in\alpha} \mathbf{P}_{u^\perp} P_i) \geqslant |\alpha| + 1$ for all α. The fourth property now follows from Lemma 2.2. □

From now on, we will fix $r \in [n-2]$ and a difference of support functions f with $S_{f,\mathcal{P}} = 0$, and construct the difference of support functions g as in Lemma 8.3. In particular, Lemma 8.3 ensures that the projection $\mathbf{P}_{u^\perp} g$ yields an equality case (3.1) of the Alexandrov-Fenchel inequality in dimension $n-1$ for any $u \in S^{n-1} \cap \mathcal{L}_r$. If we now assume that the conclusion of Theorem 8.1 is valid in dimension $n-1$, this will yield an explicit characterization of $\mathbf{P}_{u^\perp} g$ that will serve as the induction hypothesis for the proof of Theorem 8.1 in dimension n.

Theorem 8.1 is only valid, however, if the supercriticality assumption is satisfied. In order to implement the above program, we must therefore show that the supercriticality assumption on $\mathcal{P}$ is inherited by $\mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}$. We will presently show that this is in fact the case for almost all directions u, which will suffice for our purposes. More precisely, let us define the sets

$$N := \bigcup_{\substack{\alpha \subseteq [n-2]\backslash\{r\}: \\ \dim \mathcal{L}_\alpha = |\alpha|+2}} S^{n-1} \cap \mathcal{L}_r \cap \mathcal{L}_\alpha, \qquad U := (S^{n-1} \cap \mathcal{L}_r)\backslash N.$$

Then we have the following lemma. Here and in the remainder of this paper, we will frequently use the following simple linear algebra fact without further comment: for any linear subspace $E \subseteq \mathbb{R}^n$ and $u \in S^{n-1}$, we have $\dim(\mathbf{P}_{u^\perp} E) = \dim E$ if $u \notin E$, whereas $\dim(\mathbf{P}_{u^\perp} E) = \dim E - 1$ if $u \in E$.

Lemma 8.4. *The following hold:*

a. $\mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}$ *is supercritical for every* $u \in U$.

b. U *has full measure with respect to the uniform measure on* $S^{n-1} \cap \mathcal{L}_r$.

Proof. To prove part a, consider $u \in S^{n-1} \cap \mathcal{L}_r$ such that $\mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}$ is not supercritical. Then $\dim(\sum_{i\in\alpha} \mathbf{P}_{u^\perp} P_i) < |\alpha| + 2$ for some $\alpha \subseteq [n-2]\backslash\{r\}$, $\alpha \neq \varnothing$. On the other hand, as $\mathcal{P}$ is supercritical, we have $\dim(\sum_{i\in\alpha} P_i) \geqslant |\alpha| + 2$. By the above linear algebra fact, this can only occur if $\dim(\sum_{i\in\alpha} P_i) = |\alpha| + 2$ and $u \in \mathcal{L}_\alpha$, so that $u \in N$. Thus if $u \in U$, then $\mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}$ must be supercritical.

To prove part b, it suffices to show that N is the intersection of $S^{n-1} \cap \mathcal{L}_r$ with hyperplanes of codimension at least one. That is, for any $\alpha \subseteq [n-2]\backslash\{r\}$ such that $\dim \mathcal{L}_\alpha = |\alpha| + 2$, we claim that $\dim(\mathcal{L}_\alpha \cap \mathcal{L}_r) < \dim \mathcal{L}_r$. Indeed, if this is not the case, we must have $\mathcal{L}_r \subseteq \mathcal{L}_\alpha$. But that would imply that $\dim \mathcal{L}_{\alpha\cup\{r\}} = \dim \mathcal{L}_\alpha = |\alpha| + 2$, contradicting the supercriticality assumption on $\mathcal{P}$. □

Combining the above observations, we can now formally state the induction hypothesis (recall that r, f, g have been fixed in the remainder of this section).

Corollary 8.5. *Suppose that Theorem 8.1 has been proved in dimension $n-1$. Then for any $u \in U$, there exists $s(u) \in u^\perp$ such that*

$$g(x) = \langle s(u), x\rangle \quad \textit{for all } x \in \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}.$$

Proof. Applying Lemma 3.12 in $u^\perp$ and Lemma 8.3, we obtain $S_{\mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}} = 0$. As $\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}$ is supercritical by Lemma 8.4, applying Theorem 8.1 in $u^\perp$ yields

$$\mathbf{P}_{u^\perp} g(x) = \langle s(u), x\rangle \quad \text{for all } x \in \operatorname{supp} S_{\mathbf{P}_{u^\perp}B, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}}.$$

But as $\mathbf{P}_{u^\perp} g(x) = g(\mathbf{P}_{u^\perp} x)$ and as $S_{\mathbf{P}_{u^\perp}B, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}}$ is supported in $u^\perp$ by definition, we may remove $\mathbf{P}_{u^\perp}$ on the left-hand side. The conclusion now follows as $\operatorname{supp} S_{\mathbf{P}_{u^\perp}B, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}} = \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}$ by Corollary 3.9 (see Remark 8.6 below). ☐

Remark 8.6. In the proof of Corollary 8.5, we encountered a mixed area measure of the form $S_{\mathbf{P}_{u^\perp}C_1,\ldots,\mathbf{P}_{u^\perp}C_{n-2}}$ for convex bodies $C_1,\ldots,C_{n-2}$ in $\mathbb{R}^n$. By convention, this notation will be taken to mean that $\mathbf{P}_{u^\perp}C_1,\ldots,\mathbf{P}_{u^\perp}C_{n-2}$ are viewed as convex bodies in $u^\perp$, and that the mixed area measure is computed in this space. Even though we do not specify explicitly in the notation in which space the mixed area measure is computed, this will always be clear from context. For example, note that the collection $\mathbf{P}_{u^\perp}C_1,\ldots,\mathbf{P}_{u^\perp}C_{n-2}$ consists of $n-2$ bodies, so its mixed area measure only makes sense in an $(n-1)$-dimensional space.

Projected mixed area measures may be equivalently expressed as mixed area measures in $\mathbb{R}^n$ by Corollary 3.9. Indeed, note that

$$\int h\, dS_{[0,u],C_1,\ldots,C_{n-2}} = \frac{1}{n-1}\int h\, dS_{\mathbf{P}_{u^\perp}C_1,\ldots,\mathbf{P}_{u^\perp}C_{n-2}}$$

for any convex bodies $C_1,\ldots,C_{n-2}$ in $\mathbb{R}^n$ and any difference of support functions h by Corollary 3.9 and (2.1), where we used again that $\mathbf{P}_{u^\perp}h(x) = h(\mathbf{P}_{u^\perp}x)$. As we may choose h to be any C^2 function by Lemma 2.1, it follows that

$$(n-1)\, S_{[0,u],C_1,\ldots,C_{n-2}} = S_{\mathbf{P}_{u^\perp}C_1,\ldots,\mathbf{P}_{u^\perp}C_{n-2}}.$$

This is, of course, the direct counterpart of Corollary 3.9 for mixed area measures. Let us note, in particular, that $\operatorname{supp} S_{[0,u],C_1,\ldots,C_{n-2}} \subset u^\perp$.

8.2. **The gluing argument.** We now aim to show that the induction hypothesis of Corollary 8.5 implies the conclusion of Theorem 8.1 in dimension n. To this end, the main issue we encounter is to show that $s(u)$ may be replaced by a single vector $s \in \mathbb{R}^n$ that is independent of u. That is, we must "glue" together the linear functions obtained for different $u \in U$ to obtain a single linear function.

As a first step, we observe that supports of the measures $S_{[0,u],B,\mathcal{P}_{\backslash r}}$ for different $u \in U$ have a small but nontrivial overlap.

Lemma 8.7. *Let $u, v \in U$ be linearly independent. Then*

$$\operatorname{supp} S_{[0,u],[0,v],\mathcal{P}_{\backslash r}} \subseteq \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}} \cap \operatorname{supp} S_{[0,v],B,\mathcal{P}_{\backslash r}},$$

and

$$\operatorname{span}\operatorname{supp} S_{[0,u],[0,v],\mathcal{P}_{\backslash r}} = \{u,v\}^\perp.$$

Proof. The first claim is immediate by Lemma 2.4. To prove the second claim, note first that $\operatorname{span}\operatorname{supp} S_{[0,u],[0,v],\mathcal{P}_{\backslash r}} \subseteq \{u,v\}^\perp$ by Remark 8.6. Now suppose the inclusion is strict. Then $\operatorname{supp} S_{[0,u],[0,v],\mathcal{P}_{\backslash r}} \subset w^\perp$ for some $w \in S^{n-1} \cap \{u,v\}^\perp$, so

$$0 = \int \langle w, x\rangle_+\, S_{[0,u],[0,v],\mathcal{P}_{\backslash r}}(dx) = n\,\mathsf{V}_n([0,w],[0,u],[0,v],\mathcal{P}_{\backslash r})$$

using $h_{[0,w]}(x) = \langle w, x\rangle_+$ and (2.1). Now note that

$$\dim(\textstyle\sum_{i\in\alpha} P_i) \geqslant |\alpha| + 2 \qquad \text{and} \qquad \dim([0,v] + \textstyle\sum_{i\in\alpha} P_i) \geqslant |\alpha| + 3$$

for every $\alpha \subseteq [n-2]\backslash\{r\}$, $\alpha \neq \varnothing$ by the supercriticality assumption and the definition of U. As u, v, w are linearly independent, it follows from Lemma 2.2 that $\mathsf{V}_n([0,w],[0,u],[0,v],\mathcal{P}_{\backslash r}) > 0$, which entails a contradiction. □

We can now conclude the following.

Corollary 8.8. *Suppose the conclusion of Corollary 8.5 holds. Then there exists a function $a : U \times U \to \mathbb{R}$ such that $s(u) - s(v) = a(u,v)u - a(v,u)v$ whenever $u, v \in U$ are linearly independent.*

Proof. Let $u, v \in U$ be linearly independent. By Corollary 8.5 and Lemma 8.7,

$$\langle s(u), x\rangle = g(x) = \langle s(v), x\rangle \quad \text{for all } x \in \operatorname{supp} S_{[0,u],[0,v],\mathcal{P}_{\backslash r}}.$$

Thus Lemma 8.7 implies $s(u) - s(v) \perp \{u, v\}^\perp$, so that

$$s(u) - s(v) = a(u,v)u + b(u,v)v$$

for some functions a, b. But exchanging the roles of u, v, we obtain

$$a(u,v)u + b(u,v)v = s(u) - s(v) = -(s(v) - s(u)) = -a(v,u)v - b(v,u)u,$$

which implies $b(u,v) = -a(v,u)$ as u, v are linearly independent. □

Next, we show that the function $a(u,v)$ may be chosen to be independent of v.

Lemma 8.9. *Suppose the conclusion of Corollary 8.5 holds, and let $v, w \in U$ be linearly independent. Then there is a function $b : U \to \mathbb{R}$ such that the function $u \mapsto s(u) - b(u)u$ is constant on $U\backslash \operatorname{span}\{v, w\}$.*

Proof. Let the function a be as in Corollary 8.8. Consider first any linearly independent $u, v, w \in U$. Then we obtain by Corollary 8.8

$$\begin{aligned} 0 &= s(u) - s(v) + s(v) - s(w) + s(w) - s(u) \\ &= (a(u,v) - a(u,w))u + (a(v,w) - a(v,u))v + (a(w,u) - a(w,v))w. \end{aligned}$$

Thus $a(u,v) = a(u,w)$ by linear independence of u, v, w.

Let us now fix any linearly independent $v, w \in U$, and let $b(u) := a(u,v)$ for $u \in U$. As u, v, w are linearly independent for any $u \in U\backslash \operatorname{span}\{v, w\}$, we have $b(u) = a(u,v) = a(u,w)$ and $b(w) = a(w,v) = a(w,u)$ for all such u. Therefore

$$s(u) - b(u)u = s(w) + s(u) - s(w) - b(u)u = s(w) - b(w)w$$

for every $u \in U\backslash \operatorname{span}\{v, w\}$ by Corollary 8.8. □

Putting together the preceding arguments, we obtain the following.

Lemma 8.10. *Suppose that Theorem 8.1 has been proved in dimension $n-1$. Then there exists $s \in \mathbb{R}^n$ such that*

$$g(x) = \langle s, x\rangle \quad \textit{for all } x \in \operatorname{supp} S_{B,B_r,\mathcal{P}_{\backslash r}}.$$

Proof. We begin by noting that $\dim P_r \geqslant 3$ by the supercriticality assumption. Therefore, as U has full measure in $S^{n-1} \cap \mathcal{L}_r$ by Lemma 8.4, we may choose linearly independent $v, w \in U$. Moreover, as $\dim P_r \geqslant 3$ and $\dim \operatorname{span}\{u, v\} = 2$, it follows that $U\backslash \operatorname{span}\{v, w\}$ still has full measure.

By Corollary 8.5 and Lemma 8.9, there exists a function $b : U \to \mathbb{R}$ and $s \in \mathbb{R}^n$ so that $s(u) - b(u)u = s$ for all $u \in U \backslash \operatorname{span}\{v, w\}$. Thus Corollary 8.5 yields

$$g(x) = \langle s, x\rangle \quad \text{for all } x \in \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}} \text{ and } u \in U \backslash \operatorname{span}\{v, w\},$$

where we used that $\langle s, x\rangle = \langle s(u), x\rangle$ for $x \in u^\perp$.

Now note that it follows as in the proof of Lemma 3.10 and Remark 8.6 that $\int S_{[0,u],B,\mathcal{P}_{\backslash r}}\, \omega_r(du) = \kappa_{\dim P_r - 1}\, S_{B_r,B,\mathcal{P}_{\backslash r}}$, where ω_r denotes the uniform measure on $S^{n-1} \cap \mathcal{L}_r$. As $U \backslash \operatorname{span}\{v, w\}$ has full ω_r-measure, we can compute

$$\begin{aligned} 0 &= \int_{U \backslash \operatorname{span}\{v,w\}} \left(\int |g(x) - \langle s, x\rangle|\, S_{[0,u],B,\mathcal{P}_{\backslash r}}(dx) \right) \omega_r(du) \\ &= \kappa_{\dim P_r - 1} \int |g(x) - \langle s, x\rangle|\, S_{B_r,B,\mathcal{P}_{\backslash r}}(dx). \end{aligned}$$

The conclusion follows by the continuity of $g(x) - \langle s, x\rangle$. □

We have now almost concluded the induction step in the proof of Theorem 8.1, but there is a remaining subtlety: in Lemma 8.10 we have shown that $g(x) = \langle s, x\rangle$ for $x \in \operatorname{supp} S_{B,B_r,\mathcal{P}_{\backslash r}}$, while the conclusion of Theorem 8.1 states that this holds for $x \in \operatorname{supp} S_{B,P_r,\mathcal{P}_{\backslash r}}$. That the latter follows from the former is an immediate consequence of the following lower-dimensional analogue of Lemma 2.4.

Lemma 8.11. *For any convex bodies $\mathcal{C} = (C_1, \ldots, C_{n-2})$ in $\mathbb{R}^n$, we have*

$$\operatorname{supp} S_{P_r,\mathcal{C}} \subseteq \operatorname{supp} S_{B_r,\mathcal{C}}.$$

Proof. Let K be any convex body in $\mathbb{R}^n$ such that h_K is a C^2 function on S^{n-1}. It is shown in [32, Lemma 5.4] that we have

$$\int h\, dS_{K,\mathcal{C}} \leqslant \|\nabla^2 h_K\|_{L^\infty(S^{n-1})} \int h\, dS_{B,\mathcal{C}}$$

for any difference of support functions $h : S^{n-1} \to \mathbb{R}_+$. Let us now define $\Pi_\varepsilon := \mathsf{P}_{\mathcal{L}_r} + \varepsilon \mathsf{P}_{\mathcal{L}_r^\perp}$. Replacing $\mathcal{C} \leftarrow \Pi_\varepsilon^{-1}\mathcal{C}$ and $h \leftarrow h \circ \Pi_\varepsilon^{-1}$ in the above inequality yields

$$\int h\, dS_{\Pi_\varepsilon K,\mathcal{C}} \leqslant \|\nabla^2 h_K\|_{L^\infty(S^{n-1})} \int h\, dS_{\Pi_\varepsilon B,\mathcal{C}},$$

where we have used (2.1) and part f of Lemma 3.1. Letting $\varepsilon \to 0$ yields

$$\int h\, dS_{\mathsf{P}_{\mathcal{L}_r} K,\mathcal{C}} \leqslant \|\nabla^2 h_K\|_{L^\infty(S^{n-1})} \int h\, dS_{B_r,\mathcal{C}}$$

by Lemma 3.3. In particular, using Lemma 2.1, this implies that

$$\operatorname{supp} S_{\mathsf{P}_{\mathcal{L}_r} K,\mathcal{C}} \subseteq \operatorname{supp} S_{B_r,\mathcal{C}}$$

for any convex body K in $\mathbb{R}^n$ such that h_K is C^2 on S^{n-1}.

By a classical approximation argument [30, Theorem 3.4.1], we can find a sequence of convex bodies $K^{(l)}$ so that $h_{K^{(l)}}$ is C^2 for each l, and $K^{(l)} \to P_r$ in Hausdorff distance. Thus $S_{\mathsf{P}_{\mathcal{L}_r} K^{(l)},\mathcal{C}} \xrightarrow{w} S_{P_r,\mathcal{C}}$ by Lemma 3.3. But as each $S_{\mathsf{P}_{\mathcal{L}_r} K^{(l)},\mathcal{C}}$ is supported in $\operatorname{supp} S_{B_r,\mathcal{C}}$, this must be the case for the limiting measure as well. □

We can now conclude the proof of Theorem 8.1.

Proof of Theorem 8.1. The *if* direction of Theorem 8.1 follows directly from Lemmas 2.7 and 2.8, so it suffices to consider the *only if* direction.

Suppose first that Theorem 8.1 has been proved in dimension $n-1$ for some $n \geqslant 3$. Then we claim that Theorem 8.1 holds also in dimension n. Indeed, let $f : S^{n-1} \to \mathbb{R}$ be a difference of support functions such that $S_{f,\mathcal{P}} = 0$, and let g be the function constructed in Lemma 8.3 (for any $r \in [n-2]$ that is fixed throughout the proof). By Lemmas 8.10 and 8.11, there exists $s \in \mathbb{R}^n$ so that

$$g(x) = \langle s, x\rangle \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}.$$

The claim follows as $f(x) = g(x)$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$ by Lemma 8.3.

It remains to prove the base case $n = 2$. More precisely, we claim the following: for any difference of support functions $f : S^1 \to \mathbb{R}$ such that $S_f = 0$, there must exist $s \in \mathbb{R}^2$ so that $f(x) = \langle s, x\rangle$ for all $x \in S^1$. This is a classical fact; for example, it may be deduced from the equality case of the Brunn-Minkowski inequality as in [30, Theorem 7.2.1]. Let us give another proof here in order to illustrate a method that will be used again in section 10.3 in an essential manner.

Suppose f does not satisfy $f = \langle s, \cdot\rangle$ for any s. Then the Hahn-Banach theorem implies [7, Corollary IV.3.15] that there is a finite signed measure σ on S^1 so that

$$\int f \, d\sigma > 0 \qquad \text{and} \qquad \int x \, \sigma(dx) = 0.$$

Let $\sigma = \sigma^+ - \sigma^-$ be the Hahn-Jordan decomposition of σ, and let $m := \int x \, \sigma^{\pm}(dx)$ and $\mu^{\pm} := \sigma^{\pm} + \|m\| \delta_{-m/\|m\|} + S_B$. Then $\mu^{\pm}$ are nonnegative measures on S^1 so that $\int x \, \mu^{\pm}(dx) = 0$ and $\operatorname{span} \operatorname{supp} \mu^{\pm} = \mathbb{R}^2$. By the Minkowski existence theorem [30, Theorem 8.2.2], there exist convex bodies $C^{\pm}$ in $\mathbb{R}^2$ so that $\mu^{\pm} = S_{C^{\pm}}$. But then we obtain using (2.1) and the symmetry of mixed volumes

$$\int f \, d\sigma = \int f \, dS_{C^+} - \int f \, dS_{C^-} = \int h_{C^+} \, dS_f - \int h_{C^-} \, dS_f = 0,$$

which entails the desired contradiction. □

Remark 8.12. Let us highlight a surprising aspect of the proof of Theorem 8.1. By Lemma 2.8, the equality condition $S_{f,\mathcal{P}} = 0$ can only determine f on the support of $S_{B,\mathcal{P}}$. However, in Lemma 8.10 we have characterized the function g on the support of $S_{B,B_r,\mathcal{P}_{\backslash r}}$. The latter set is often much larger than the former. For example, if $P_1 = \cdots = P_{n-2} = P$ is a full-dimensional polytope, then $\operatorname{supp} S_{B,\mathcal{P}}$ is the set of normal directions of $(n-2)$-dimensional faces of P, but $\operatorname{supp} S_{B,B,\mathcal{P}_{\backslash r}}$ is the set of normal directions of $(n-3)$-dimensional faces of P (cf. [30, Theorem 4.5.3]).

Nonetheless, there is no contradiction, as Theorem 4.3 only ensures that $f = g$ on the smaller set $\operatorname{supp} S_{B,\mathcal{P}}$. The phenomenon exhibited here should be viewed as another manifestation of the fact that the local Alexandrov-Fenchel inequality fixes many degrees of freedom of the extremal functions.

9. Structure of critical sets

We now turn to the study of the extremals of the Alexandrov-Fenchel inequality in the critical case (Definition 4.1). The new feature that arises when $\mathcal{P}$ is critical is the appearance of $\mathcal{P}$-degenerate functions (Definition 2.10). Their analysis requires several new ideas, whose development will occupy us throughout sections 9–11.

The definition of the critical case differs from the supercritical case only in that there may now exist indices $i_1 < \cdots < i_k$ so that $\dim(P_{i_1} + \cdots + P_{i_k}) = k + 1$.

Such *critical sets* of indices will prove to be intimately connected to the structure of $\mathcal{P}$-degenerate pairs and functions. For example, we will show that for any $\mathcal{P}$-degenerate pair (M,N), the bodies M,N must be contained (up to translation) in the affine hull of $P_{i_1}+\cdots+P_{i_k}$ for some critical set $i_1<\cdots<i_k$.

In this section, we begin the analysis of the critical case by obtaining a classification of the critical sets, which will be used to give an explicit description of the structure of $\mathcal{P}$-degenerate functions. In section 10, we undertake a detailed study of the geometric structure of critical mixed area measures. These results will be employed in section 11 to prove Theorem 2.13 in the critical case.

Throughout this section, we fix $n\geqslant 3$ and a critical collection $\mathcal{P}=(P_1,\ldots,P_{n-2})$ of polytopes in $\mathbb{R}^n$. As in section 8, we will assume without loss of generality that P_i contains the origin in its relative interior for every $i\in[n-2]$, and we define the spaces $\mathcal{L}_\alpha$ and balls B_α as in the supercritical case. The criticality assumption may then be formulated as $\dim\mathcal{L}_\alpha\geqslant|\alpha|+1$ for every $\alpha\subseteq[n-2]$, $\alpha\neq\varnothing$.

9.1. **Critical sets.** The following definition will play a central role in the sequel.

Definition 9.1. Let $\mathcal{C}=(C_1,\ldots,C_m)$ be any collection of convex bodies.

a. $\alpha\subseteq[m]$ is called *$\mathcal{C}$-critical* if $\dim(\sum_{i\in\alpha}C_i)=|\alpha|+1$.

b. $\alpha\subseteq[m]$ is called *$\mathcal{C}$-maximal* if it is $\mathcal{C}$-critical, and there is no $\mathcal{C}$-critical set $\beta\supsetneq\alpha$.

A $\mathcal{P}$-critical ($\mathcal{P}$-maximal) set $\alpha\subseteq[n-2]$ will simply be called *critical* (*maximal*).

The analysis of degenerate functions will be greatly facilitated by the fact that the family of critical sets is organized in a very simple manner. The following lemma and its corollary are due to Panov [23, Lemma 6].

Lemma 9.2. *Let α,α' be critical sets. If $\alpha\cap\alpha'\neq\varnothing$, then $\alpha\cup\alpha'$ is a critical set.*

Proof. For any $\beta,\beta'\subset[n-2]$, we have $\mathcal{L}_{\beta\cup\beta'}=\mathcal{L}_\beta+\mathcal{L}_{\beta'}$ and $\mathcal{L}_{\beta\cap\beta'}\subseteq\mathcal{L}_\beta\cap\mathcal{L}_{\beta'}$ by the definition of $\mathcal{L}_\beta$. On the other hand, as we assumed $\mathcal{P}$ is critical and $\alpha\cap\alpha'\neq\varnothing$, we have $\dim\mathcal{L}_{\alpha\cup\alpha'}\geqslant|\alpha\cup\alpha'|+1$ and $\dim\mathcal{L}_{\alpha\cap\alpha'}\geqslant|\alpha\cap\alpha'|+1$. Therefore

$$\begin{aligned}|\alpha\cup\alpha'|+1&\leqslant\dim\mathcal{L}_{\alpha\cup\alpha'}=\dim\mathcal{L}_\alpha+\dim\mathcal{L}_{\alpha'}-\dim(\mathcal{L}_\alpha\cap\mathcal{L}_{\alpha'})\\&\leqslant\dim\mathcal{L}_\alpha+\dim\mathcal{L}_{\alpha'}-\dim\mathcal{L}_{\alpha\cap\alpha'}\\&\leqslant(|\alpha|+1)+(|\alpha'|+1)-(|\alpha\cap\alpha'|+1)\\&=|\alpha\cup\alpha'|+1,\end{aligned}$$

where we used that $\dim\mathcal{L}_\alpha=|\alpha|+1$ and $\dim\mathcal{L}_{\alpha'}=|\alpha'|+1$ as α,α' are critical. It follows that $\dim\mathcal{L}_{\alpha\cup\alpha'}=|\alpha\cup\alpha'|+1$, so $\alpha\cup\alpha'$ is critical. ☐

The key consequence of Lemma 9.2 is that distinct *maximal* sets α,α' must be disjoint. This structure is also reflected in the associated linear spaces: if α,α' are distinct maximal sets, then $\mathcal{L}_\alpha,\mathcal{L}_{\alpha'}$ are linearly independent.

Corollary 9.3. *Let $\alpha\neq\alpha'$ be maximal sets. Then $\alpha\cap\alpha'=\varnothing$ and $\mathcal{L}_\alpha\cap\mathcal{L}_{\alpha'}=\{0\}$.*

Proof. Let α,α' be distinct maximal sets. Then $\alpha\cup\alpha'$ cannot be a critical set: as either $\alpha\cup\alpha'\supsetneq\alpha$ or $\alpha\cup\alpha'\supsetneq\alpha'$, this would contradict maximality of α,α'. Thus $\alpha\cap\alpha'=\varnothing$, as otherwise $\alpha\cup\alpha'$ would be a critical set by Lemma 9.2.

Now note that as $\alpha\cup\alpha'$ is not a critical set and $\mathcal{P}$ is critical, we have

$$|\alpha\cup\alpha'|+2\leqslant\dim\mathcal{L}_{\alpha\cup\alpha'}\leqslant\dim\mathcal{L}_\alpha+\dim\mathcal{L}_{\alpha'}=|\alpha|+|\alpha'|+2=|\alpha\cup\alpha'|+2,$$

where we used that α, α' are critical sets and $\alpha \cap \alpha' = \varnothing$. Thus

$$\dim(\mathcal{L}_\alpha \cap \mathcal{L}_{\alpha'}) = \dim \mathcal{L}_{\alpha \cup \alpha'} - \dim \mathcal{L}_\alpha - \dim \mathcal{L}_{\alpha'} = 0,$$

completing the proof. □

In view of Corollary 9.3, we obtain the following picture. Associated to the critical collection $\mathcal{P}$ of polytopes is its collection $\{\alpha_1, \ldots, \alpha_\ell\}$ of disjoint maximal sets. Any critical set β is contained in exactly one of the maximal sets α_i. Moreover, the linear spaces $\mathcal{L}_{\alpha_i}$ are pairwise (but not jointly) linearly independent. The same properties extend *verbatim* to any critical collection $\mathcal{C}$ of convex bodies.

Let us finally record a simple observation.

Lemma 9.4. *Let $\alpha \subseteq [n-2]$ be a critical set and $\beta \subseteq [n-2]$ be arbitrary. Then*

$$\mathcal{L}_\beta \subseteq \mathcal{L}_\alpha \quad \text{if and only if} \quad \beta \subseteq \alpha.$$

Proof. If $\beta \subseteq \alpha$, then $\mathcal{L}_\beta \subseteq \mathcal{L}_\alpha$ by definition. Conversely, if $\mathcal{L}_\beta \subseteq \mathcal{L}_\alpha$, then

$$|\alpha| + 1 \leqslant |\alpha \cup \beta| + 1 \leqslant \dim \mathcal{L}_{\alpha \cup \beta} = \dim \mathcal{L}_\alpha = |\alpha| + 1,$$

where we used that $\mathcal{P}$ is critical, that $\mathcal{L}_\beta \subseteq \mathcal{L}_\alpha$, and that α is a critical set, respectively. Thus $|\alpha| = |\alpha \cup \beta|$, which implies $\beta \subseteq \alpha$. □

9.2. **Degenerate pairs and functions.** We now use the above classification of critical sets to obtain a better understanding of Definition 2.10. For simplicity, $\mathcal{P}$-degenerate pairs and functions will henceforth be called *degenerate pairs* and *degenerate functions*, respectively. However, the same structure will apply *verbatim* to $\mathcal{C}$-degenerate pairs and functions for any critical collection $\mathcal{C}$ of convex bodies.

Let us begin by introducing a more precise definition.

Definition 9.5. Let α be a maximal set and M, N be convex bodies in $\mathbb{R}^n$.

a. (M, N) is called an *α-degenerate pair* if

$$M, N \subset \mathcal{L}_\alpha \qquad \text{and} \qquad \mathsf{V}_{\mathcal{L}_\alpha}(M, \mathcal{P}_\alpha) = \mathsf{V}_{\mathcal{L}_\alpha}(N, \mathcal{P}_\alpha).$$

b. A function $f : S^{n-1} \to \mathbb{R}$ is called an *α-degenerate function* if $f = h_M - h_N$ for some α-degenerate pair (M, N).

If $\mathcal{C}$ is a critical collection of convex bodies and α is $\mathcal{C}$-maximal, the analogous definitions will be referred to as $(\mathcal{C}, \alpha)$-degenerate pairs and functions.

As a first step towards understanding Definition 9.5, let us note for any maximal (hence also critical) set α, we have $\dim \mathcal{L}_\alpha = |\alpha| + 1$ and $P_i \subset \mathcal{L}_\alpha$ for every $i \in \alpha$. Thus the mixed volume $\mathsf{V}_{\mathcal{L}_\alpha}(M, \mathcal{P}_\alpha)$ is indeed well defined: this is the mixed volume of $|\alpha| + 1$ convex bodies in the $(|\alpha| + 1)$-dimensional space $\mathcal{L}_\alpha$.

We will now show that in the present setting ($\mathcal{P}$ is critical), any degenerate pair or function in the sense of Definition 2.10 is in fact an α-degenerate pair or function up to translation. In other words, degenerate pairs must always be contained in translates of $\mathcal{L}_\alpha$ for some maximal set α, which provides an explicit geometric description of the dimensionality property that is implicit in Definition 2.10.

Lemma 9.6. *(M, N) is a degenerate pair if and only if M is not a translate of N and $(M + v, N + w)$ is an α-degenerate pair for some maximal set α and $v, w \in \mathbb{R}^n$. Thus f is a degenerate function if and only if f is nonlinear and $f - \langle v, \cdot \rangle$ is an α-degenerate function for some maximal set α and $v \in \mathbb{R}^n$.*

Proof. We begin by noting that for any critical set α and convex body $K \subset \mathcal{L}_\alpha$, Lemma 3.8 implies the projection formula

$$\binom{n}{|\alpha|+1} \mathsf{V}_n(K, B, \mathcal{P}) = \mathsf{V}_{\mathcal{L}_\alpha}(K, \mathcal{P}_\alpha)\, \mathsf{V}_{\mathcal{L}_\alpha^\perp}(\mathbf{P}_{\mathcal{L}_\alpha^\perp} B, \mathbf{P}_{\mathcal{L}_\alpha^\perp} \mathcal{P}_{\backslash\alpha}).$$

Moreover, as $\mathsf{V}_n(B_\alpha, B, \mathcal{P}) > 0$ by Lemma 2.2 and the assumption that $\mathcal{P}$ is critical, it follows that $\mathsf{V}_{\mathcal{L}_\alpha^\perp}(\mathbf{P}_{\mathcal{L}_\alpha^\perp} B, \mathbf{P}_{\mathcal{L}_\alpha^\perp} \mathcal{P}_{\backslash\alpha}) > 0$.

Consider first an α-degenerate pair (M, N) for some maximal set α, where M, N are not translates. We claim that (M, N) is a degenerate pair. Indeed, condition (2.3) follows from Lemma 2.2 as $\dim(M + N + \sum_{i\in\alpha} P_i) = \dim \mathcal{L}_\alpha = |\alpha| + 1$, while condition (2.4) follows from the projection formula and Definition 9.5.

Now consider a degenerate pair (M, N). As M is not a translate of N, at least one of M, N must have nonzero dimension. But as $\mathcal{P}$ is critical, $\mathsf{V}_n(K, B, \mathcal{P}) > 0$ whenever $\dim K \geqslant 1$ by Lemma 2.2. Thus (2.4) implies that $\dim(M) \geqslant 1$ and $\dim(N) \geqslant 1$. On the other hand, it cannot be the case that $\dim(M + N) = 1$. Indeed, if that were the case, then M, N must be segements with parallel directions; moreover, (2.4) then implies that M, N have equal length, so that M, N are translates. This case is therefore ruled out by the definition of a degenerate pair.

We have now shown that any degenerate pair (M, N) must satisfy

$$\dim(M) \geqslant 1, \qquad \dim(N) \geqslant 1, \qquad \dim(M + N) \geqslant 2.$$

Together with the assumption that $\mathcal{P}$ is critical, it follows from Lemma 2.2 and (2.3) that there must exist $\alpha' \subseteq [n-2]$, $\alpha' \neq \varnothing$ such that

$$\dim(M + N + \textstyle\sum_{i\in\alpha'} P_i) \leqslant |\alpha'| + 1.$$

On the other hand, as $\mathcal{P}$ is critical we have $\dim(\sum_{i\in\alpha'} P_i) \geqslant |\alpha'| + 1$. The only way this can happen is if $\dim \mathcal{L}_{\alpha'} = \dim(\sum_{i\in\alpha'} P_i) = |\alpha'| + 1$ (that is, α' is critical) and there exist $v, w \in \mathbb{R}^n$ so that $M + v$ and $N + w$ lie in $\mathcal{L}_{\alpha'}$.

Now let α be the maximal set containing α'. Then $M, N \subset \mathcal{L}_{\alpha'} \subseteq \mathcal{L}_\alpha$. Moreover, by the projection formula, the normalization condition of Definition 9.5 follows from (2.4). Thus we have shown that $(M + v, N + w)$ is an α-degenerate pair.

Finally, the equivalence between degenerate and α-degenerate functions is an immediate consequence of the corresponding equivalence for pairs. □

Lemma 9.6 explains the basic structure of the extremals of the Alexandrov-Fenchel inequality that appears in Theorem 2.13. Note that for a given maximal set α, any linear combination of α-degenerate functions is again an α-degenerate function by definition. On the other hand, if f is an α-degenerate function and f' is an α'-degenerate function for distinct maximal sets α, α', then linear combinations of f, f' need not be degenerate. Each maximal set α will therefore give rise to (at most) one α-degenerate pair in the statement of Theorem 2.13.

9.3. **An intrinsic description.** So far we have defined degenerate functions as differences of support functions of degenerate pairs of convex bodies. However, in the proof of Theorem 2.13, it will be necessary to construct degenerate functions directly by gluing together lower-dimensional degenerate functions. To this end, we now introduce a more intrinsic perspective on degenerate functions that does not require the auxiliary construction of a degenerate pair.

Before we proceed, we state a variant of the projection formula of Lemma 3.8 in terms of mixed area measures, which will be needed below.

Lemma 9.7. *Let $C_1, \ldots, C_{n-1}$ be convex bodies in $\mathbb{R}^n$, and suppose that $C_1, \ldots, C_k$ lie in a subspace E with $\dim E = k+1$. Then*

$$\binom{n-1}{k} \int \varphi(\mathsf{P}_E x)\, S_{C_1,\ldots,C_{n-1}}(dx) = \mathsf{V}_{E^\perp}(\mathsf{P}_{E^\perp} C_{k+1}, \ldots, \mathsf{P}_{E^\perp} C_n) \int \varphi\, dS_{C_1,\ldots,C_k}$$

for any 1-homogeneous function $\varphi : E \to \mathbb{R}$ that is $S_{C_1,\ldots,C_k}$-integrable.

Proof. Suppose first that the restriction of φ to $S^{n-1} \cap E$ is a C^2 function. Then we may write $\varphi = h_K - h_L$ for convex bodies K, L in E by Lemma 2.1. Moreover, by the definition of support functions, $\varphi(\mathsf{P}_E x) = h_K(x) - h_L(x)$ for any $x \in S^{n-1}$ as $K, L \subset E$. The conclusion now follows from Lemma 3.8 and (2.1).

Now define the map $\iota : S^{n-1} \backslash E^\perp \to S^{n-1} \cap E$ as $\iota(x) := \mathsf{P}_E x / \|\mathsf{P}_E x\|$. By 1-homogeneity of φ, the identity in the statement of the lemma may be written as

$$\binom{n-1}{k} \int \varphi \circ \iota\, d\mu = \mathsf{V}_{E^\perp}(\mathsf{P}_{E^\perp} C_{k+1}, \ldots, \mathsf{P}_{E^\perp} C_n) \int \varphi\, dS_{C_1,\ldots,C_k},$$

where the measure $\mu(dx) := \|\mathsf{P}_E x\|\, S_{C_1,\ldots,C_{n-1}}(dx)$ is supported on $S^{n-1} \backslash E^\perp$. As we have shown this identity holds for any φ of class C^2, it follows that

$$\binom{n-1}{k} \mu \circ \iota^{-1} = \mathsf{V}_{E^\perp}(\mathsf{P}_{E^\perp} C_{k+1}, \ldots, \mathsf{P}_{E^\perp} C_n)\, S_{C_1,\ldots,C_k} \tag{9.1}$$

as measures on $S^{n-1} \cap E$. The conclusion follows for any integrable 1-homogeneous function $\varphi : E \to \mathbb{R}$ by integrating this identity. $\square$

Remark 9.8. Suppose $C_1, \ldots, C_k$ are polytopes in Lemma 9.7. Then $S_{C_1,\ldots,C_k}$ has finite support by Lemma 3.4. Thus (9.1) shows that the measure $S_{C_1,\ldots,C_{n-1}} \circ \mathsf{P}_E^{-1}$ is supported on a finite union of rays emanating from the origin with directions in $\operatorname{supp} S_{C_1,\ldots,C_k}$. We now observe that any 1-homogeneous function φ is continuous on such a set: it is linear on each ray and zero at the origin. This implies that in the polytope setting, the function $x \mapsto \varphi(\mathsf{P}_E x)$ is continuous on $\operatorname{supp} S_{C_1,\ldots,C_{n-1}}$ for any 1-homogeneous function φ. This observation will be used below.

We can now introduce the main idea of this section: α-degenerate functions may be intrinsically described in terms of 1-homogeneous functions on $\mathcal{L}_\alpha$.

Lemma 9.9. *Let α be a maximal set.*

a. For any α-degenerate function f, there exists a 1-homogeneous function $\varphi : \mathcal{L}_\alpha \to \mathbb{R}$ with $\int \varphi\, dS_{\mathcal{P}_\alpha} = 0$ so that $f(x) = \varphi(\mathsf{P}_{\mathcal{L}_\alpha} x)$ for all $x \in S^{n-1}$.

b. For any 1-homogeneous function $\varphi : \mathcal{L}_\alpha \to \mathbb{R}$ with $\int \varphi\, dS_{\mathcal{P}_\alpha} = 0$, there exists an α-degenerate function f so that $f(x) = \varphi(\mathsf{P}_{\mathcal{L}_\alpha} x)$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$.

Proof. To prove part a, write $f = h_M - h_N$ for some α-degenerate pair (M, N). As $M, N \subset \mathcal{L}_\alpha$, we have $h_M(x) = h_M(\mathsf{P}_{\mathcal{L}_\alpha} x)$ and $h_N(x) = h_N(\mathsf{P}_{\mathcal{L}_\alpha} x)$ for all $x \in S^{n-1}$ by the definition of support functions. Now define φ to be the restriction of $h_M - h_N$ to $\mathcal{L}_\alpha$. Then φ is 1-homogeneous, $f(x) = \varphi(\mathsf{P}_{\mathcal{L}_\alpha} x)$ for all $x \in S^{n-1}$, and

$$\frac{1}{|\alpha| + 1} \int \varphi\, dS_{\mathcal{P}_\alpha} = \mathsf{V}_{\mathcal{L}_\alpha}(M, \mathcal{P}_\alpha) - \mathsf{V}_{\mathcal{L}_\alpha}(N, \mathcal{P}_\alpha) = 0$$

by (2.1) and the definition of an α-critical pair.

The same argument would apply *verbatim* in the converse direction if φ can be written as a difference of support functions. This is not clear, however, as we did not make any regularity assumption on φ. To work around this issue, we will exploit that $\mathcal{P}$ are polytopes to create a modification of φ with the requisite property.

More precisely, part b is proved as follows. As $\mathcal{P}$ are polytopes, $\operatorname{supp} S_{\mathcal{P}_\alpha}$ is a finite subset of $S^{n-1} \cap \mathcal{L}_\alpha$ by Lemma 3.4. Thus we can choose a C^2 function $\eta : S^{n-1} \cap \mathcal{L}_\alpha \to \mathbb{R}$ so that $\varphi(x) = \eta(x)$ for all $x \in \operatorname{supp} S_{\mathcal{P}_\alpha}$. By Lemma 2.1, there exist convex bodies $M, N \subset \mathcal{L}_\alpha$ so that $\eta(x) = h_M(x) - h_N(x)$ for all $x \in S^{n-1} \cap \mathcal{L}_\alpha$. We claim that $f := h_M - h_N$ has the properties stated in part b. Indeed, note that

$$\mathsf{V}_{\mathcal{L}_\alpha}(M, \mathcal{P}_\alpha) - \mathsf{V}_{\mathcal{L}_\alpha}(N, \mathcal{P}_\alpha) = \frac{1}{|\alpha| + 1} \int f \, dS_{\mathcal{P}_\alpha} = \frac{1}{|\alpha| + 1} \int \varphi \, dS_{\mathcal{P}_\alpha} = 0,$$

where we used (2.1) in the first equality and $f = \varphi$ $S_{\mathcal{P}_\alpha}$-a.e. in the second equality. Thus (M, N) is an α-degenerate pair and f is an α-degenerate function. On the other hand, as $f = \varphi$ on $\operatorname{supp} S_{\mathcal{P}_\alpha}$, we obtain

$$\begin{aligned} 0 &= \mathsf{V}_{\mathcal{L}_\alpha^\perp}(\mathbf{P}_{\mathcal{L}_\alpha^\perp} B, \mathbf{P}_{\mathcal{L}_\alpha^\perp} \mathcal{P}_{\backslash\alpha}) \int |f - \varphi| \, dS_{\mathcal{P}_\alpha} \\ &= \binom{n-1}{|\alpha|} \int |f(x) - \varphi(\mathbf{P}_{\mathcal{L}_\alpha} x)| \, S_{B,\mathcal{P}}(dx) \end{aligned}$$

by Lemma 9.7, where we used that $f(x) = f(\mathbf{P}_{\mathcal{L}_\alpha} x)$ as $M, N \subset \mathcal{L}_\alpha$. Thus $f(x) = \varphi(\mathbf{P}_{\mathcal{L}_\alpha} x)$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$ by Remark 9.8, completing the proof. $\square$

10. Propeller geometry

We have seen in the previous section that the appearance of degenerate functions is intimately connected to the critical sets of the reference bodies $\mathcal{P}$. In this section, we will develop a new geometric phenomenon that explains the origin of this behavior: we will show that the supports of critical mixed area measures exhibit certain geometric structures that we call *propellers*, in view of their resemblance to the propeller of a Mississippi steamboat. These propellers will play a crucial role in the proof of Theorem 2.13 in the critical case.

This section is organized as follows. We first introduce the propeller structure in section 10.1. In the proof of Theorem 2.13, this structure will be exploited in two nontrivial ways: to glue together lower-dimensional degenerate functions, and to decouple the contributions arising from distinct maximal sets. We develop both these methods in an abstract setting in sections 10.2 and 10.3, respectively. While the basic principles can be understood independently of the rest of the paper, their power will become clear when they are applied in section 11.

10.1. **The propeller.**

The following theorem describes the propeller structure.

Theorem 10.1. *Let $C_1, \ldots, C_{n-1}$ be convex bodies in $\mathbb{R}^n$, and suppose $C_1, \ldots, C_k$ lie in a subspace E with $\dim E = k + 1$. Define the space $F_z := \operatorname{span}\{E^\perp, z\}$ and halfspace $F_z^+ := \{x \in F_z : \langle z, x\rangle > 0\}$ for $z \in E$. Then*

$$\operatorname{supp} S_{C_1,\ldots,C_{n-1}} \subset E^\perp \cup \bigcup_{z \in \operatorname{supp} S_{C_1,\ldots,C_k}} F_z^+,$$

and

$$\begin{aligned} &\binom{n-1}{k} \int f(x) \, 1_{x \notin E^\perp} \, S_{C_1,\ldots,C_{n-1}}(dx) \\ &= \int \left(\int f(x) \, 1_{\langle z, x\rangle > 0} \, S_{\mathbf{P}_{F_z} C_{k+1},\ldots,\mathbf{P}_{F_z} C_{n-1}}(dx) \right) S_{C_1,\ldots,C_k}(dz) \end{aligned} \tag{10.1}$$

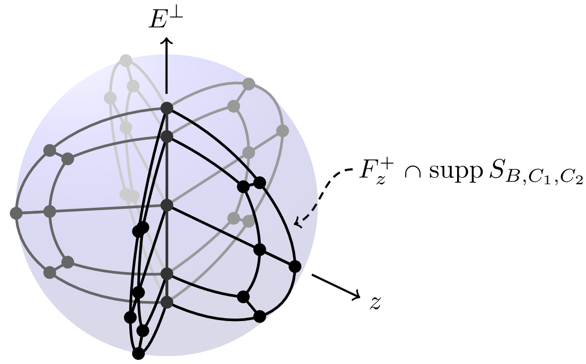

FIGURE 10.1. Illustration of a propeller structure in $\mathbb{R}^4$.

for any bounded measurable function $f : S^{n-1} \to \mathbb{R}$.

Informally, Theorem 10.1 states that $S_{C_1,\ldots,C_{n-1}}$ is supported in a union of half-spaces ("blades") centered around $E^\perp$ with orthogonal direction in $\operatorname{supp} S_{C_1,\ldots,C_k} \subset E$. Moreover, $S_{C_1,\ldots,C_{n-1}}$ agrees on each blade with the mixed area measure of the projections of $C_{k+1},\ldots,C_{n-1}$ onto the subspace in which that blade lies.

Example 10.2. The propeller structure is illustrated in Figure 10.1. Here C_1, C_2 are polytopes in $\mathbb{R}^4$, where C_1 is a pentagon contained in the plane $E = \operatorname{span}\{e_1, e_2\}$ spanned by the first two coordinate directions, and C_2 is a full-dimensional polytope (in the figure, we chose the Minkowski sum of a cube and an octahedron). We have visualized the support of S_{B,C_1,C_2} by projecting it onto $\operatorname{span}\{e_1, e_2, e_3\}$, which yields a geometric graph in the unit ball of $\mathbb{R}^3$ (cf. section 5.2). The propeller structure is immediately evident in the picture: the "blades" of the propeller lie in the halfspaces F_z^+, while the "shaft" of the propeller lies in $E^\perp$. There are five blades, corresponding to the five facet normals of the pentagon C_1.

Example 10.3. Suppose $C_1,\ldots,C_{n-2}$ all lie in a subspace $E = w^\perp$. Then we may apply Theorem 10.1 with $k = n-1$ to investigate the measure $S_{B,C_1,\ldots,C_{n-2}}$. In this case $S^{n-1} \cap F_z^+$ is merely a semicircular arc from w to $-w$ passing through z, and $S_{\mathsf{P}_{F_z} B}$ is the uniform measure on this arc. The propeller then takes the explicit "striped watermelon" form that was studied in [32, §8] and implicitly in [29].

Let us now discuss the interaction between degenerate functions and the propeller structure. In the setting of Theorem 10.1, a degenerate function may be expressed as $f(x) = \varphi(\mathsf{P}_E x)$ for a 1-homogeneous function $\varphi : E \to \mathbb{R}$ (cf. Lemma 9.9). Now note that for any $z \in S^{n-1} \cap E$, we have $\mathsf{P}_E F_z = \operatorname{span}\{z\}$ by definition, so that $\mathsf{P}_E x = \langle x, z\rangle z$ for any $x \in F_z$. We therefore obtain

$$f(x) = \varphi(z)\langle z, x\rangle \quad \text{for all } x \in F_z^+,\ z \in S^{n-1} \cap E$$

by the homogeneity of φ. On the other hand, clearly $f(x) = 0$ for $x \in E^\perp$. Thus we have shown that *a degenerate function is linear on each blade of the propeller, and vanishes on the shaft.* This provides a geometric explanation for why degenerate functions are extremals of the Alexandrov-Fenchel inequality: the conditions of Proposition 5.7 are satisfied for degenerate functions precisely because the propeller structure creates a geometric mechanism for this to happen.

Remark 10.4. The propeller structure was already hinted at by the observation of Remark 9.8 in the previous section: it is evident from the propeller structure that the projection of $\operatorname{supp} S_{C_1,\ldots,C_{n-1}}$ on E is supported on rays emanating from the origin in the directions in $\operatorname{supp} S_{C_1,\ldots,C_k}$. More generally, the reader may verify that Lemma 9.7 can be deduced directly from Theorem 10.1 and Corollary 3.9.

Remark 10.5. In Theorem 10.1 we only considered the effect of a single critical set on the geometry of the mixed area measure. However, many critical sets may coexist for the same collection of bodies: this is not ruled out by the assumptions of Theorem 10.1, where we singled out one critical set for analysis. When distinct maximal sets are present, the geometry of the mixed area measure will feature several propellers that are superimposed in different directions. Such "propellers within propellers" are hard to visualize, and we will not attempt to do so. Nonetheless, this situation must be addressed in the proof of Theorem 2.13, which will be done using a technique that is developed in section 10.3 below.

We now turn to the proof of Theorem 10.1.

Proof of Theorem 10.1. The statement about the support of $S_{C_1,\ldots,C_{n-1}}$ follows immediately from (10.1). We will first prove (10.1) in the case that $C_1,\ldots,C_{n-1}$ are polytopes, and then derive the general case by approximation.

Step 1. Suppose that $C_1,\ldots,C_{n-1}$ are polytopes. Fix any $x \in S^{n-1}\backslash E^\perp$ and let $z := \mathsf{P}_E x/\|\mathsf{P}_E x\|$. Then $x \in F_z$ by definition. As $C_1+\cdots+C_k \subset E$, we have

$$F(C_1+\cdots+C_k,x) = F(C_1+\cdots+C_k,z) \subset az + E\cap z^\perp$$

for some constant a by Lemma 3.7. In particular, $\dim F(C_1+\cdots+C_k,x) \leqslant k$. We can therefore write using Lemmas 3.4 and 3.8

$$\begin{aligned}\binom{n-1}{k} S_{C_1,\ldots,C_{n-1}}(\{x\}) &= \binom{n-1}{k}\mathsf{V}_{n-1}(F(C_1,x),\ldots,F(C_{n-1},x))\\ &= \mathsf{V}_{E\cap z^\perp}(F(C_1,z),\ldots,F(C_k,z))\,\mathsf{V}_{F_z}(\mathsf{P}_{F_z}F(C_{k+1},x),\ldots,\mathsf{P}_{F_z}F(C_{n-1},x))\\ &= S_{C_1,\ldots,C_k}(\{z\})\,S_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}}(\{x\}),\end{aligned}$$

where we used in the last line that $\mathsf{P}_{F_z}F(C_i,x) = F(\mathsf{P}_{F_z}C_i,x)$ by Lemma 3.7.

Now note that for any $u \in E$, we have $x \in F_u$ if and only if $u = z$ or $u = -z$. In particular, as $\langle z,x\rangle > 0$ and $\operatorname{supp} S_{\mathsf{P}_{F_u}C_{k+1},\ldots,\mathsf{P}_{F_u}C_{n-1}} \subset F_u$, we have

$$1_{\langle u,x\rangle>0}\,S_{\mathsf{P}_{F_u}C_{k+1},\ldots,\mathsf{P}_{F_u}C_{n-1}}(\{x\}) = 0 \quad \text{for all } u\in E,\ u\neq z.$$

We therefore obtain

$$\binom{n-1}{k} S_{C_1,\ldots,C_{n-1}}(\{x\}) = \sum_u 1_{\langle u,x\rangle>0}\,S_{\mathsf{P}_{F_u}C_{k+1},\ldots,\mathsf{P}_{F_u}C_{n-1}}(\{x\})\,S_{C_1,\ldots,C_k}(\{u\}).$$

As this identity holds for any $x\in S^{n-1}\backslash E^\perp$, (10.1) follows from Lemma 3.4.

Step 2. We now aim to show that (10.1) remains valid when $C_1,\ldots,C_{n-1}$ are arbitrary convex bodies. We first claim that the result is equivalent to

$$\begin{aligned}&\binom{n-1}{k}\int g(x)\,\|\mathsf{P}_E x\|\,S_{C_1,\ldots,C_{n-1}}(dx)\\ &= \int\left(\int g(x)\,\langle z,x\rangle_+\,S_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}}(dx)\right)S_{C_1,\ldots,C_k}(dz)\end{aligned} \tag{10.2}$$

for every continuous function $g : S^{n-1} \to \mathbb{R}$. That (10.1) implies (10.2) follows by choosing $f(x) = g(x)\|\mathsf{P}_E x\|$ and using that $\|\mathsf{P}_E x\| = |\langle z, x\rangle|$ on F_z. Conversely, suppose (10.2) holds; by a standard approximation argument, it extends to any nonnegative measurable function g. Choosing $g(x) = f(x)1_{x\notin E^\perp}\|\mathsf{P}_E x\|^{-1}$ yields (10.1) for nonnegative f, and the conclusion follows by linearity.

The advantage of (10.2) is that the integrands are continuous, so we may use weak convergence. Fix a continuous function $g : S^{n-1} \to \mathbb{R}$, and choose polytopes $C_1^{(l)}, \ldots, C_{n-1}^{(l)}$ so that $C_1^{(l)}, \ldots, C_k^{(l)} \subset E$ and $C_r^{(l)} \to C_r$ in Hausdorff distance for all r (the existence of such approximations is elementary [30, Theorem 1.8.16]). We have already shown in the first part of the proof that (10.2) holds for the polytopes $C_1^{(l)}, \ldots, C_{n-1}^{(l)}$. We would like to show the identity remains valid as $l \to \infty$. As mixed area measures are continuous by Lemma 3.3, it suffices by a standard weak convergence argument [18, Theorem 4.27] to show that

$$\begin{aligned}\int g(x)\,\langle z_l, x\rangle_+\, S_{\mathsf{P}_{F_{z_l}} C_{k+1}^{(l)}, \ldots, \mathsf{P}_{F_{z_l}} C_{n-1}^{(l)}}(dx)& \\ \xrightarrow{l\to\infty} \int g(x)\,\langle z, x\rangle_+\, S_{\mathsf{P}_{F_z} C_{k+1}, \ldots, \mathsf{P}_{F_z} C_{n-1}}(dx)&\end{aligned}$$

for any sequence $z_l \to z \in S^{n-1} \cap E$. But this follows readily from Lemma 3.3 as

$$\begin{aligned}\|h_{\mathsf{P}_{F_{z_l}} C_r^{(l)}} - h_{\mathsf{P}_{F_z} C_r}\|_\infty &\leqslant \|h_{\mathsf{P}_{F_{z_l}} C_r^{(l)}} - h_{\mathsf{P}_{F_{z_l}} C_r}\|_\infty + \|h_{\mathsf{P}_{F_{z_l}} C_r} - h_{\mathsf{P}_{F_z} C_r}\|_\infty \\ &\leqslant \|h_{C_r^{(l)}} - h_{C_r}\|_\infty + \|h_{\mathsf{P}_{F_{z_l}} C_r} - h_{\mathsf{P}_{F_z} C_r}\|_\infty \xrightarrow{l\to\infty} 0\end{aligned}$$

(that is, $\mathsf{P}_{F_{z_l}} C_r^{(l)} \to \mathsf{P}_{F_z} C_r$ in Hausdorff distance) and $\|\langle z_l, \cdot\rangle - \langle z, \cdot\rangle\|_\infty \to 0$. $\square$

10.2. **The gluing principle.** One of the difficulties we will encounter in the proof of Theorem 2.13 is that we must glue together degenerate functions in $(n-1)$-dimensional hyperplanes to form a degenerate function in dimension n. This will be accomplished using the following application of the propeller structure.

Lemma 10.6. *Let $C_{k+1}, \ldots, C_{n-1}$ be convex bodies in $\mathbb{R}^n$, and let E be a subspace of dimension $\dim E = k+1$. Assume that*

$$\mathsf{V}_{E^\perp}(\mathsf{P}_{E^\perp} C_{k+1}, \ldots, \mathsf{P}_{E^\perp} C_{n-1}) > 0.$$

Then for any 1-homogeneous function $h : \mathbb{R}^n \to \mathbb{R}$, there exists a 1-homogeneous function $\varphi : E \to \mathbb{R}$ with the following property: for every collection of convex bodies $K_1, \ldots, K_k$ in E such that we have

$$h(x) = \tilde{\varphi}(\mathsf{P}_E x) \quad \textit{for all } x \in \operatorname{supp} S_{K_1, \ldots, K_k, C_{k+1}, \ldots, C_{n-1}}$$

for some 1-homogeneous function $\tilde{\varphi} : E \to \mathbb{R}$, we have in fact

$$h(x) = \varphi(\mathsf{P}_E x) \quad \textit{for all } x \in \operatorname{supp} S_{K_1, \ldots, K_k, C_{k+1}, \ldots, C_{n-1}}.$$

The point of Lemma 10.6 is that the function $\tilde{\varphi}$ depends on the choice of bodies $K_1, \ldots, K_k$, while φ does not. Thus φ may be viewed as having "glued together" the functions $\tilde{\varphi}$ over all choices of $K_1, \ldots, K_k$ in E for which $\tilde{\varphi}$ exists. In the proof of Theorem 2.13, $\tilde{\varphi}$ will be degenerate functions of the $(n-1)$-dimensional projections, and φ will be the degenerate function in dimension n.

Proof of Lemma 10.6. Throughout the proof we adopt the same notation as in Theorem 10.1. First, we define $\varphi(z)$ for $z \in S^{n-1} \cap E$ as

$$\varphi(z) := \frac{\int h(x)\, 1_{\langle z,x\rangle>0}\, S_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}}(dx)}{\mathsf{V}_{E^\perp}(\mathsf{P}_{E^\perp}C_{k+1},\ldots,\mathsf{P}_{E^\perp}C_{n-1})}.$$

We may extend $\varphi : E \to \mathbb{R}$ to a 1-homogeneous function as $\varphi(z) := \|z\|\varphi(z/\|z\|)$ for $z \in E\backslash\{0\}$. We now show that φ satisfies the requisite property.

To this end, let $\tilde{\varphi} : E \to \mathbb{R}$ be any 1-homogeneous function, and fix any $z \in S^{n-1} \cap E$. As $\mathsf{P}_E x = \langle x, z\rangle z$ for any $x \in F_z$, we obtain

$$\tilde{\varphi}(\mathsf{P}_E x) = \tilde{\varphi}(z)\langle z, x\rangle \qquad \text{for any } x \in F_z^+.$$

Integrating this identity against $1_{\langle z,x\rangle>0}\, S_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}}(dx)$ yields

$$\begin{aligned}
&\int \tilde{\varphi}(\mathsf{P}_E x)\, 1_{\langle z,x\rangle>0}\, S_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}}(dx)\\
&= \tilde{\varphi}(z)\int \langle z, x\rangle_+\, S_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}}(dx)\\
&= \tilde{\varphi}(z)\,\mathsf{V}_{E^\perp}(\mathsf{P}_{E^\perp}C_{k+1},\ldots,\mathsf{P}_{E^\perp}C_{n-1}),
\end{aligned} \tag{10.3}$$

where we used Corollary 3.9 in the last line.

Now let $K_1,\ldots,K_k$ be convex bodies in E such that

$$h(x) = \tilde{\varphi}(\mathsf{P}_E x) \quad \text{for all } x \in \operatorname{supp} S_{K_1,\ldots,K_k,C_{k+1},\ldots,C_{n-1}}.$$

Then we also have

$$h(x) = \tilde{\varphi}(\mathsf{P}_E x) \quad \text{for all } x \in \operatorname{supp}(1_{\langle z,\cdot\rangle>0}\, dS_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}})$$

for any $z \in \operatorname{supp} S_{K_1,\ldots,K_k}$ by Theorem 10.1. Thus

$$\tilde{\varphi}(z) = \varphi(z) \quad \text{for all } z \in \operatorname{supp} S_{K_1,\ldots,K_k}$$

by (10.3) and the definition of φ. But then we may also conclude that

$$h(x) = \tilde{\varphi}(\mathsf{P}_E x) = \varphi(\mathsf{P}_E x) \quad \text{for all } x \in \operatorname{supp} S_{K_1,\ldots,K_k,C_{k+1},\ldots,C_{n-1}},$$

because $\mathsf{P}_E x \in \|\mathsf{P}_E x\| \operatorname{supp} S_{K_1,\ldots,K_k}$ for every $x \in \operatorname{supp} S_{K_1,\ldots,K_k,C_{k+1},\ldots,C_{n-1}}$ by Theorem 10.1, and as $\tilde{\varphi}$ and φ are both 1-homogeneous. $\square$

10.3. **Linear relations.** Lemma 10.6 is only applicable when a single degenerate function appears. In general there may be multiple degenerate functions corresponding to different maximal sets, and we will need a way to decouple their analysis. The technique that will be used for this purpose is developed in this section. The utility of the following result will be far from obvious at this point, but we will see in section 11.2 that it plays a key role in our proofs.

Unlike the other results of this section, we formulate the following result only for polytopes, which will suffice for our purposes. The polytope assumption is convenient in the proof, but does not appear to be of fundamental importance.

Proposition 10.7. *Let $C_1,\ldots,C_{n-1}$ be polytopes in $\mathbb{R}^n$. Suppose that $C_1,\ldots,C_k$ lie in a subspace E with $\dim E = k+1$ and satisfy the criticality condition of Definition 4.1. Let $h : S^{n-1} \to \mathbb{R}$ be a function such that*

$$h(x) = 0 \quad \textit{for all } x \in E^\perp \cap \operatorname{supp} S_{C_1,\ldots,C_{n-1}},$$

and such that

$$\int h\, dS_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}} = 0$$

for every full-dimensional polytope Q in E. Then there exists $w \in E$ so that

$$\int h(x)\, 1_{\langle z,x\rangle>0}\, S_{\mathsf{P}_{F_z}C_{k+1},\ldots,\mathsf{P}_{F_z}C_{n-1}}(dx) = \langle w, z\rangle$$

for all $z \in S^{n-1} \cap E$, where F_z is as defined in Theorem 10.1.

The reader should keep in mind the case where h is a degenerate function corresponding to a critical set disjoint from $[k]$. Then the bodies $C_1, \ldots, C_k$ factor out of the integral $\int h\, dS_{C_1,\ldots,C_{n-1}}$ by Lemma 9.7, and may thus be replaced by any other body Q in E. This motivates the assumption of Proposition 10.7. The conclusion of Proposition 10.7 then states that the average of h over each blade of the propeller generated by $C_1, \ldots, C_k$ must be linearly related across the blades. This is not at all clear from Theorem 10.1, which specifies the mixed area measure on each blade but does not explain the relations between different blades.

The proof of Proposition 10.7 is based on a duality argument that is similar to the one used at the end of the proof of Theorem 8.1. Let us begin by formulating a simple consequence of the Minkowski existence theorem.

Lemma 10.8. *Let σ be a signed measure on $S^{n-1} \cap E$ that is supported on a finite number of points and satisfies $\int x\, \sigma(dx) = 0$. Then there exist full-dimensional polytopes Q, Q' in E so that $\sigma = S_{Q,\ldots,Q} - S_{Q',\ldots,Q'}$.*

Proof. Let $\sigma = \sigma^+ - \sigma^-$ be the Hahn-Jordan decomposition of σ and $m := \int x\, \sigma^{\pm}(dx)$. Let R be any full-dimensional polytope in E, and define

$$\mu^{\pm} := \sigma^{\pm} + \|m\|\delta_{-m/\|m\|} + S_{R,\ldots,R}.$$

Then $\mu^{\pm}$ are finitely supported measures, $\int x\, d\mu^{\pm} = 0$, and $\operatorname{span}\operatorname{supp}\mu^{\pm} = E$. The Minkowski existence theorem [30, Theorem 8.2.1] therefore yields the existence of full-dimensional polytopes Q, Q' in E so that $\mu^+ = S_{Q,\ldots,Q}$ and $\mu^- = S_{Q',\ldots,Q'}$. The conclusion now follows as $\sigma = \mu^+ - \mu^-$. □

We will exploit Lemma 10.8 through a duality argument.

Corollary 10.9. *Let $\varrho : S^{n-1} \cap E \to \mathbb{R}$ be any function such that $\int \varrho\, dS_{Q,\ldots,Q} = 0$ for every full-dimensional polytope Q in E. Then $\varrho = \langle w, \cdot\rangle$ for some $w \in E$.*

Proof. We first claim that for any finite set $\Omega \subset S^{n-1} \cap E$, there exists $w_\Omega \in E$ so that $\varrho = \langle w_\Omega, \cdot\rangle$ on Ω. Indeed, suppose this is not the case; then by the Hahn-Banach theorem, there is a signed measure with support in Ω so that

$$\int \varrho\, d\sigma > 0 \qquad \text{and} \qquad \int x\, \sigma(dx) = 0.$$

This is contradicted by Lemma 10.8 and the assumption.

Now let $w := w_{\{v_1,\ldots,v_{k+1}\}}$, where $\{v_1, \ldots, v_{k+1}\}$ is a basis of E. Then we have

$$\langle w, v_i\rangle = \varrho(v_i) = \langle w_{\{x,v_1,\ldots,v_{k+1}\}}, v_i\rangle$$

for every $x \in S^{n-1} \cap E$ and i. Thus $w_{\{x,v_1,\ldots,v_{k+1}\}} = w$, so that

$$\varrho(x) = \langle w_{\{x,v_1,\ldots,v_{k+1}\}}, x\rangle = \langle w, x\rangle$$

for every $x \in S^{n-1} \cap E$. □

Remark 10.10. The reason for the somewhat roundabout finite-dimensional argument is that we did not assume any regularity (for example, continuity) of ϱ, so the Hahn-Banach theorem cannot be applied directly in infinite dimension.

We now formulate a useful consequence of the criticality condition. It is this part of the proof that is facilitated by the polytope assumption.

Lemma 10.11. *Let $C_1, \ldots, C_{n-1}$ be polytopes in $\mathbb{R}^n$. Suppose that $C_1, \ldots, C_k$ lie in a subspace E with $\dim E = k+1$ and satisfy the criticality condition of Definition 4.1. Then for any full-dimensional polytope Q in E, we have*

$$E^\perp \cap \operatorname{supp} S_{C_1,\ldots,C_{n-1}} = E^\perp \cap \operatorname{supp} S_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}}.$$

Proof. Fix $u \in E^\perp$. As $C_1, \ldots, C_k, Q \subset E$, Lemma 3.4 yields

$$\begin{aligned} S_{C_1,\ldots,C_{n-1}}(\{u\}) &= \mathsf{V}_{u^\perp}(C_1, \ldots, C_k, F(C_{k+1}, u), \ldots, F(C_{n-1}, u)), \\ S_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}}(\{u\}) &= \mathsf{V}_{u^\perp}(Q, \ldots, Q, F(C_{k+1}, u), \ldots, F(C_{n-1}, u)). \end{aligned}$$

Thus $S_{C_1,\ldots,C_{n-1}}(\{u\}) > 0$ implies $S_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}}(\{u\}) > 0$ by Lemma 2.2, as Q is full-dimensional. It remains to prove the converse implication.

To this end, suppose $S_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}}(\{u\}) > 0$. Then by Lemma 2.2, there exist segments $I_1, \ldots, I_k \subseteq Q$ and $I_r \subseteq F(C_r, u)$ for $r = k+1, \ldots, n-1$ so that

$$\mathsf{V}_{u^\perp}(I_1, \ldots, I_{n-1}) > 0.$$

Thus Lemma 3.8 yields

$$0 < \binom{n-1}{k} \mathsf{V}_{u^\perp}(Q, \ldots, Q, I_{k+1}, \ldots, I_{n-1}) = \mathrm{Vol}_G(\mathbf{P}_G Q)\, \mathsf{V}_H(I_{k+1}, \ldots, I_{n-1}),$$

where H is the linear span of the directions of the segments $I_{k+1}, \ldots, I_{n-1}$ and $G := H^\perp \cap u^\perp$. This implies that $\dim(\mathbf{P}_G Q) = k$. But as $\dim(Q) = k+1$ by assumption, the map $\mathbf{P}_G|_E$ must have a one-dimensional kernel. Therefore

$$\dim\left(\sum_{i\in\alpha} \mathbf{P}_G C_i\right) \geqslant \dim\left(\sum_{i\in\alpha} C_i\right) - 1 \geqslant |\alpha| \quad \text{for all } \alpha \subseteq [k],$$

where we used the criticality assumption in the second inequality. Therefore

$$\begin{aligned} 0 &< \mathsf{V}_G(\mathbf{P}_G C_1, \ldots, \mathbf{P}_G C_k)\, \mathsf{V}_H(I_{k+1}, \ldots, I_{n-1}) \\ &= \binom{n-1}{k} \mathsf{V}_{u^\perp}(C_1, \ldots, C_k, I_{k+1}, \ldots, I_{n-1}) \\ &\leqslant \binom{n-1}{k} \mathsf{V}_{u^\perp}(C_1, \ldots, C_k, F(C_{k+1}, u), \ldots, F(C_{n-1}, u)) \end{aligned}$$

by Lemmas 2.2 and 3.8. Thus we have shown that $S_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}}(\{u\}) > 0$ implies $S_{C_1,\ldots,C_{n-1}}(\{u\}) > 0$, completing the proof. □

We can now complete the proof of Proposition 10.7.

Proof of Proposition 10.7. Define the function $\varrho : S^{n-1} \cap E \to \mathbb{R}$ as

$$\varrho(z) := \int h(x)\, 1_{\langle z,x\rangle > 0}\, S_{\mathbf{P}_{F_z} C_{k+1},\ldots,\mathbf{P}_{F_z} C_{n-1}}(dx).$$

Then we have for any convex body Q in E

$$\int \varrho\, dS_{Q,\ldots,Q} = \binom{n-1}{k} \int h(x)\, 1_{x\notin E^\perp}\, S_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}}(dx)$$

by Theorem 10.1. But by the first assumption on h and Lemma 10.11, we have $h(x) = 0$ for $x \in E^\perp \cap \operatorname{supp} S_{Q,\ldots,Q,C_{k+1},\ldots,C_{n-1}}$ when Q is a full-dimensional polytope in E. Thus the second assumption on h shows that $\int \varrho\, dS_{Q,\ldots,Q} = 0$ for every full-dimensional polytope Q in E. The conclusion follows from Corollary 10.9. $\square$

11. The critical case

In this section, we complete the extremal characterization of the Alexandrov-Fenchel inequality in the critical case. More precisely, we will prove the following.

Theorem 11.1. *Let $\mathcal{P} = (P_1, \ldots, P_{n-2})$ be polytopes in $\mathbb{R}^n$ that contain the origin in their relative interior. Assume that $\mathcal{P}$ is critical but not supercritical, and denote by $\alpha_0, \ldots, \alpha_\ell$ the associated maximal sets. Then for any difference of support functions $f : S^{n-1} \to \mathbb{R}$, we have $S_{f,\mathcal{P}} = 0$ if and only if*

$$f(x) = \langle s, x\rangle + \sum_{j=0}^{\ell} g_j(x) \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}$$

holds for some $s \in \mathbb{R}^n$ and α_j-degenerate function g_j, $j = 0, \ldots, \ell$.

The proof of Theorem 11.1 proceeds by induction on n. Just as in the supercritical case, it will turn out that the criticality assumption (Definition 4.1) is preserved by the induction. The induction hypothesis may therefore give rise to a supercritical case, which is already covered by Theorem 8.1, or to a critical case, to which we may apply Theorem 11.1 in lower dimension.

The following setting will be assumed throughout this section. We fix $n \geqslant 3$ and a collection of polytopes $\mathcal{P} = (P_1, \ldots, P_{n-2})$ in $\mathbb{R}^n$ that contain the origin in their relative interior. We assume that $\mathcal{P}$ is critical but not supercritical, that is, there exists at least one critical set. We denote the maximal sets by $\alpha_0, \ldots, \alpha_\ell$. The spaces $\mathcal{L}_\alpha$ and the balls B_α are defined as in section 8. In particular, the criticality assumption may be formulated as $\dim \mathcal{L}_\alpha \geqslant |\alpha| + 1$ for every $\alpha \subseteq [n-2]$, $\alpha \neq \varnothing$.

11.1. The induction hypothesis.

In the induction step, we will assume that Theorem 11.1 has been proved in dimension $n-1$, and deduce its validity in dimension n. The aim of this section is to formulate the resulting induction hypothesis.

As in section 8, we will begin by applying the local Alexandrov-Fenchel inequality. To this end, we must choose an index $r \in [n-2]$ to which Theorem 4.3 will be applied. Unlike in the supercitical case, however, we may not choose r arbitrarily: the entire argument will be based on the fact that we will choose r to lie inside one of the maximal sets. We therefore fix an element $r \in \alpha_0$ at the outset, which will be used throughout the proof without further comment. As α_0 will play a special role throughout the proof, we will define henceforth $\gamma := \alpha_0$ in order to distinguish it in the notation from the remaining maximal sets $\alpha_1, \ldots, \alpha_\ell$.

Our starting point is the following direct analogue of Lemma 8.3.

Lemma 11.2. *Let f be a difference of support functions such that $S_{f,\mathcal{P}} = 0$. Then there exists a difference of support functions g with the following properties:*

1. *$g(x) = f(x)$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$.*
2. *$\mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} P_r, \mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}) = 0$ for all $u \in S^{n-1}$.*
3. *$\mathsf{V}_{n-1}(\mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} g, \mathbf{P}_{u^\perp} \mathcal{P}_{\backslash r}) = 0$ for all $u \in S^{n-1} \cap \mathcal{L}_r$.*

4. $S_{\mathsf{P}_{u^\perp}P_r,\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}} \neq 0$ *for all* $u \in S^{n-1}$.

Proof. By Theorem 4.3, there exists $g = f$ $S_{B,\mathcal{P}}$-a.e. such that $S_{g,\mathcal{P}} = 0$ and $S_{g,g,\mathcal{P}_{\backslash r}} \leqslant 0$. We must show that each of the claimed properties holds for g. The proof of properties 1–3 is identical to the proof of these properties in Lemma 8.3. To prove property 4, recall that $S_{\mathsf{P}_{u^\perp}P_r,\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}} = (n-1)\,S_{[0,u],\mathcal{P}}$ (cf. Remark 8.6). If there were to exist $u \in S^{n-1}$ so that $S_{[0,u],\mathcal{P}} = 0$, then integrating against h_B and using (2.1) would yield $\mathsf{V}_n([0,u],B,\mathcal{P}) = 0$. But this contradicts the criticality assumption by Lemma 2.2, concluding the proof. □

We would like to exploit Lemma 11.2 by applying Theorem 11.1 (or Theorem 8.1) in $u^\perp$. In order to do this, we must understand what happens to the criticality assumption under projection onto $u^\perp$. We will presently show that such a projection preserves not just the criticality assumption, but even the collection of maximal sets, for almost every choice of u. To this end, we define

$$N := \bigcup_{\substack{\alpha\subseteq[n-2]\backslash\{r\}:\\ \dim\mathcal{L}_\alpha\leqslant|\alpha|+2,\ \mathcal{L}_r\nsubseteq\mathcal{L}_\alpha}} S^{n-1}\cap\mathcal{L}_r\cap\mathcal{L}_\alpha, \qquad U := (S^{n-1}\cap\mathcal{L}_r)\backslash N$$

in the remainder of this section. Then we have the following.

Lemma 11.3. *The following hold for every* $u \in U$*:*

a. $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$ *is critical.*

b. The $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$*-maximal sets are precisely* $\gamma\backslash\{r\},\alpha_1,\ldots,\alpha_\ell$.

c. The map $\mathsf{P}_{u^\perp}|_{\mathcal{L}_{\alpha_i}} : \mathcal{L}_{\alpha_i} \to \mathsf{P}_{u^\perp}\mathcal{L}_{\alpha_i}$ *is a bijection for* $i = 1,\ldots,\ell$.

d. U *has full measure with respect to the uniform measure on* $S^{n-1}\cap\mathcal{L}_r$.

If $\gamma = \{r\}$ *is a singleton, then part b should be understood to say that the* $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$*-maximal sets are precisely* $\alpha_1,\ldots,\alpha_\ell$.

Proof. To prove part a, suppose that $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$ is not critical. Then we must have $\dim(\sum_{i\in\alpha}\mathsf{P}_{u^\perp}P_i) < |\alpha|+1$ for some $\alpha\subseteq[n-2]\backslash\{r\}$, $\alpha\neq\varnothing$. But as $\mathcal{P}$ is critical, we have $\dim(\sum_{i\in\alpha}P_i) \geqslant |\alpha|+1$. This can happen only if $\dim(\sum_{i\in\alpha}P_i) = |\alpha|+1$ and $u\in\mathcal{L}_\alpha$. Thus α is a critical set with $r\notin\alpha$, so that $\mathcal{L}_r\nsubseteq\mathcal{L}_\alpha$ by Lemma 9.4. Therefore $u\notin U$ by the definition of U, which entails a contradiction.

To prove part b, we must prove two distinct properties: that $\gamma\backslash\{r\},\alpha_1,\ldots,\alpha_\ell$ are $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-maximal sets, and that no other $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-maximal sets exist.

Claim. α_i is $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-maximal for all $i = 1,\ldots,\ell$.

Proof. Fix a maximal set $\alpha = \alpha_i$ for some $i = 1,\ldots,\ell$. Then α is disjoint from $\gamma\ni r$ by Corollary 9.3, so $\alpha\subseteq[n-2]\backslash\{r\}$. Therefore

$$|\alpha|+1 \leqslant \dim(\textstyle\sum_{i\in\alpha}\mathsf{P}_{u^\perp}P_i) \leqslant \dim(\sum_{i\in\alpha}P_i) = |\alpha|+1,$$

where the first inequality holds by part a, and the equality holds as α is critical. Thus we have shown that α is $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-critical.

To show α is $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-maximal, we must show there does not exist a $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-critical set $\beta\subseteq[n-2]\backslash\{r\}$ so that $\alpha\subsetneq\beta$. Indeed, suppose such a set β does exist. Then $\dim(\sum_{i\in\beta}\mathsf{P}_{u^\perp}P_i) = |\beta|+1$, so there are two possibilities:

i. $\dim(\sum_{i\in\beta}P_i) = |\beta|+1$ and $u\notin\mathcal{L}_\beta$; or
ii. $\dim(\sum_{i\in\beta}P_i) = |\beta|+2$ and $u\in\mathcal{L}_\beta$.

In case i, we have $\alpha \subsetneq \beta$ for a critical set β, which contradicts the maximality of α. On the other hand, in case ii, we must have $\mathcal{L}_r \subseteq \mathcal{L}_\beta$ by the definition of $U \ni u$, so that $\dim \mathcal{L}_{\beta\cup\{r\}} = \dim \mathcal{L}_\beta = |\beta| + 2$. Thus in this case $\alpha \subsetneq \beta \cup \{r\}$ and $\beta \cup \{r\}$ is a critical set, which contradicts again the maximality of α. □

Claim. If $\gamma\backslash\{r\} \neq \varnothing$, then $\gamma\backslash\{r\}$ is $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-maximal.

Proof. As $r \in \gamma$, we have $u \in \mathcal{L}_r \subseteq \mathcal{L}_\gamma$ by the definition of U. Therefore

$$|\gamma| = |\gamma\backslash\{r\}| + 1 \leqslant \dim(\textstyle\sum_{i\in\gamma\backslash\{r\}} \mathsf{P}_{u^\perp} P_i) \leqslant \dim \mathcal{L}_\gamma - 1 = |\gamma|,$$

where the first inequality holds by part a, the second inequality holds as $u \in \mathcal{L}_\gamma$, and the last equality holds as γ is critical. Thus $\gamma\backslash\{r\}$ is $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-critical.

Now suppose $\gamma\backslash\{r\} \subsetneq \beta \subseteq [n-2]\backslash\{r\}$ for some $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-critical set β. Then one of the possibilities i or ii in the proof of the previous claim must hold.

In case i, let β' be the maximal set containing the critical set β. As $\beta\backslash\gamma \neq \varnothing$, we must have $\gamma \neq \beta'$. But as $\gamma \cap \beta' \supseteq \gamma\backslash\{r\} \neq \varnothing$, this contradicts maximality of γ by Corollary 9.3. On the other hand, in case ii, we must have $\mathcal{L}_r \subseteq \mathcal{L}_\beta$ by the definition of $U \ni u$, so that $\dim \mathcal{L}_{\beta\cup\{r\}} = \dim \mathcal{L}_\beta = |\beta| + 2$. Thus $\gamma \subsetneq \beta \cup \{r\}$ and $\beta \cup \{r\}$ is a critical set, contradicting again the maximality of γ. □

Claim. There exists no $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-critical set $\beta \subseteq [n-2]\backslash(\gamma \cup \alpha_1 \cup \cdots \cup \alpha_\ell)$.

Proof. Suppose $\beta \subseteq [n-2]\backslash(\gamma \cup \alpha_1 \cup \cdots \cup \alpha_\ell)$ is a $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-critical set. Then one of the possibilities i or ii in the proofs of the previous claims must hold. Case i is impossible, as it would imply that β is a critical set that is disjoint from all maximal sets. In case ii, we must have $\mathcal{L}_r \subseteq \mathcal{L}_\beta$ by the definition of $U \ni u$, so that $\dim \mathcal{L}_{\beta\cup\{r\}} = \dim \mathcal{L}_\beta = |\beta| + 2$. Thus $\beta \cup \{r\}$ is a critical set containing r. But this would imply by Corollary 9.3 that $\beta \cup \{r\} \subseteq \gamma$, which is impossible as β and γ are disjoint. Thus we have shown the desired contradiction. □

Now recall that distinct maximal sets must be disjoint by Corollary 9.3. Thus the combination of the above three claims concludes the proof of part b.

To prove part c, it suffices to note that for any $i = 1, \ldots, \ell$,

$$\dim \mathcal{L}_{\alpha_i} = |\alpha_i| + 1 = \dim \mathsf{P}_{u^\perp}\mathcal{L}_{\alpha_i},$$

where the first equality holds as α_i is a critical set and the second equality holds as α_i is a $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-critical set by part b. The conclusion follows immediately.

Finally, we prove part d. To this end, note that by definition, N lies in a union of subspaces $\mathcal{L}_\alpha \not\supseteq \mathcal{L}_r$, so that $\dim(\mathcal{L}_\alpha \cap \mathcal{L}_r) < \dim \mathcal{L}_r$ for each of these spaces. In other words, N is the intersection of $S^{n-1} \cap \mathcal{L}_r$ with hyperplanes of codimension at least one, and is therefore a set of zero measure. □

In the rest of this section, we fix a difference of support functions f with $S_{f,\mathcal{P}} = 0$, and construct the difference of support functions g as in Lemma 11.2. Then Lemmas 11.2 and 11.3 ensure that the projection $\mathsf{P}_{u^\perp}g$ yields a critical equality case (3.1) of the Alexandrov-Fenchel inequality in dimension $n-1$ for every $u \in U$.

We would like to combine this fact with Theorem 11.1 in dimension $n-1$ to create the induction hypothesis for its proof in dimension n. In the critical case, however, there is a new subtlety: applying Theorem 11.1 in $u^\perp$ yields $\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}$-degenerate functions, while we must construct $\mathcal{P}$-degenerate functions to prove Theorem 11.1 in dimension n. We will address this problem by using condition c of Lemma 11.3 to map between the two types of degenerate functions. In addition, we must handle the new term that may now arise from part b of Lemma 3.12, which did not appear in the supercritical case. Both issues will be resolved presently.

Lemma 11.4. *Suppose that Theorem 11.1 has been proved in dimension $n-1$. Then for every $u \in U$, we have*

$$g(x) = \langle s(u), x\rangle + \sum_{j=0}^{\ell} g_{j,u}(x) \quad \text{for all } x \in \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}$$

for some $s(u) \in u^\perp$ and α_j-degenerate function $g_{j,u}$, $j = 0, \ldots, \ell$.

Proof. Fix $u \in U$. Applying Lemma 3.12 in $u^\perp$ and Lemma 11.2, we find that

$$S_{\mathbf{P}_{u^\perp} g - a_u h_{\mathbf{P}_{u^\perp} P_r}, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}} = 0 \quad \text{for some } a_u \in \mathbb{R}.$$

Now note that $\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}$ is critical with critical sets $\gamma\backslash\{r\}, \alpha_1, \ldots, \alpha_\ell$ by Lemma 11.3. If $\ell = 0$ and $\gamma = \{r\}$, then $\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}$ is supercritical and we may apply Theorem 8.1 in $u^\perp$; otherwise we may apply Theorem 11.1 in $u^\perp$. In either case

$$\mathbf{P}_{u^\perp} g(x) = \langle s(u), x\rangle + a_u h_{\mathbf{P}_{u^\perp} P_r}(x) + \sum_{j=0}^{\ell} \tilde g_{j,u}(x) \quad \text{for all } x \in \operatorname{supp} S_{\mathbf{P}_{u^\perp} B, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}}$$

holds for some $s(u) \in u^\perp$, $(\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}, \alpha_j)$-degenerate functions $\tilde g_{j,u}$, $j = 1, \ldots, \ell$, and $(\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}, \gamma\backslash\{r\})$-degenerate function $\tilde g_{0,u}$ (if $\gamma = \{r\}$, we simply set $\tilde g_{0,u} = 0$). But as $S_{\mathbf{P}_{u^\perp} B, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}}$ is supported in $u^\perp$ by definition, we may remove $\mathbf{P}_{u^\perp}$ on the left-hand side of this identity. We therefore obtain

$$g(x) = \langle s(u), x\rangle + a_u h_{P_r}(x) + \sum_{j=0}^{\ell} \tilde g_{j,u}(x) \quad \text{for all } x \in \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}} \tag{11.1}$$

as $\operatorname{supp} S_{\mathbf{P}_{u^\perp} B, \mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}} = \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}$ by Remark 8.6.

It remains to reformulate the above identity in terms of α_j-degenerate (as opposed to $(\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}, \alpha_j)$-degenerate) functions. To this end, we consider two cases.

Claim. For every $j = 1, \ldots, \ell$, there exists an α_j-degenerate function $g_{j,u}$ so that

$$g_{j,u}(x) = \tilde g_{j,u}(x) \quad \text{for all } x \in u^\perp.$$

Proof. Let $\tilde g_{j,u} = h_{\tilde M} - h_{\tilde N}$ for a $(\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}, \alpha_j)$-degenerate pair $(\tilde M, \tilde N)$. In particular, $\tilde M, \tilde N \subset \mathbf{P}_{u^\perp}\mathcal{L}_{\alpha_j}$. Therefore, by part c of Lemma 11.3, we can uniquely define convex bodies $M, N \subset \mathcal{L}_{\alpha_j}$ so that $\tilde M = \mathbf{P}_{u^\perp} M$ and $\tilde N = \mathbf{P}_{u^\perp} N$. We claim that the function $g_{j,u} := h_M - h_N$ satisfies the requisite properties.

Indeed, note that $g_{j,u}(\mathbf{P}_{u^\perp} x) = \tilde g_{j,u}(x)$ by construction, so $g_{j,u}(x) = \tilde g_{j,u}(x)$ for all $x \in u^\perp$. On the other hand, note that for any convex body $K \subset \mathcal{L}_{\alpha_j}$

$$\mathsf{V}_{\mathbf{P}_{u^\perp}\mathcal{L}_{\alpha_j}}(\mathbf{P}_{u^\perp} K, \mathbf{P}_{u^\perp}\mathcal{P}_{\alpha_j}) = [\![\mathbf{P}_{u^\perp}|_{\mathcal{L}_{\alpha_j}}]\!]\, \mathsf{V}_{\mathcal{L}_{\alpha_j}}(K, \mathcal{P}_{\alpha_j}),$$

by part f of Lemma 3.1. As $(\tilde M, \tilde N)$ is a $(\mathbf{P}_{u^\perp}\mathcal{P}_{\backslash r}, \alpha_j)$-degenerate pair, we obtain

$$[\![\mathbf{P}_{u^\perp}|_{\mathcal{L}_{\alpha_j}}]\!]\, \mathsf{V}_{\mathcal{L}_{\alpha_j}}(h_M - h_N, \mathcal{P}_{\alpha_j}) = \mathsf{V}_{\mathbf{P}_{u^\perp}\mathcal{L}_{\alpha_j}}(h_{\tilde M} - h_{\tilde N}, \mathbf{P}_{u^\perp}\mathcal{P}_{\alpha_j}) = 0.$$

As $[\![\mathbf{P}_{u^\perp}|_{\mathcal{L}_{\alpha_j}}]\!] > 0$ by part c of Lemma 11.3, and as $M, N \subset \mathcal{L}_{\alpha_j}$, we have verified that (M, N) is an α_j-degenerate pair. Thus $g_{j,u}$ is an α_j-degenerate function. □

Claim. There exists a γ-degenerate function $g_{0,u}$ so that

$$g_{0,u}(x) = a_u h_{P_r}(x) + \tilde g_{0,u}(x) \quad \text{for all } x \in u^\perp.$$

Proof. Fix $b \in \mathbb{R}$, and define $g_{0,u}$ as

$$g_{0,u} := a_u h_{P_r} + \tilde{g}_{0,u} + b h_{[0,u]}.$$

Clearly $h_{[0,u]}(x) = \langle u, x\rangle_+ = 0$ and thus $g_{0,u}(x) = a_u h_{P_r}(x) + \tilde{g}_{0,u}(x)$ for $x \in u^\perp$. We claim $g_{0,u}$ is γ-degenerate for a suitable choice of $b \in \mathbb{R}$.

If $\tilde{g}_{0,u} \neq 0$, then $\tilde{g}_{0,u} = h_M - h_N$ for a $(\mathsf{P}_{u^\perp}\mathcal{P}_{\backslash r}, \gamma\backslash\{r\})$-degenerate pair (M, N). As $r \in \gamma$, we have $u \in \mathcal{L}_r \subseteq \mathcal{L}_\gamma$ by the definition of U, so $M, N \subset \mathsf{P}_{u^\perp}\mathcal{L}_{\gamma\backslash\{r\}} \subset \mathcal{L}_\gamma$. But as $r \in \gamma$, we always have $[0, u], P_r \subset \mathcal{L}_\gamma$. Thus $g_{0,u}$ is a difference of support functions of convex bodies in $\mathcal{L}_\gamma$. On the other hand, we can choose $b \in \mathbb{R}$ so that

$$\mathsf{V}_{\mathcal{L}_\gamma}(g_{0,u}, \mathcal{P}_\gamma) = \mathsf{V}_{\mathcal{L}_\gamma}(a_u h_{P_r} + \tilde{g}_{0,u}, \mathcal{P}_\gamma) + b\,\mathsf{V}_{\mathcal{L}_\gamma}([0, u], \mathcal{P}_\gamma) = 0$$

as $\mathsf{V}_{\mathcal{L}_\gamma}([0, u], \mathcal{P}_\gamma) > 0$ by the criticality assumption and Lemma 2.2. Thus we have shown that $g_{0,u}$ is γ-degenerate, completing the proof. □

As $\operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}} \subset u^\perp$, the conclusion of Lemma 11.4 follows readily by combining the above two claims with the identity (11.1). □

11.2. **The decoupling argument.** With Lemma 11.4 in hand, the main difficulty we face is to remove the dependence of $s(u)$ and $g_{j,u}$ on u. We therefore begin, as in the supercritical case, by investigating the overlap between the supports of the measures $S_{[0,u],B,\mathcal{P}_{\backslash r}}$ for different $u \in U$. Surprisingly, the situation in the critical case proves to be completely different than in the supercritical case: the overlap between the supports does not depend on the choice of u.

In the remainder of this section, we define the polytope

$$P_\gamma := \sum_{i\in\gamma} P_i.$$

We state at the outset a simple technical lemma that will be used several times.

Lemma 11.5. $\mathsf{V}_n(K, P_\gamma, \mathcal{P}) > 0$ *whenever* $\dim(K) \geqslant 1$ *and* $0 \in K \not\subset \mathcal{L}_\gamma$.

Proof. Suppose $\mathsf{V}_n(K, P_\gamma, \mathcal{P}) = 0$. As $\mathcal{P}$ is critical (which implies $\dim P_\gamma \geqslant 2$), we must have $\dim(K + P_\gamma + \sum_{i\in\alpha} P_i) \leqslant |\alpha| + 1$ for some $\alpha \subseteq [n-2]$ by Lemma 2.2. As $\dim(\sum_{i\in\alpha} P_i) \geqslant |\alpha| + 1$ by the criticality assumption, this can happen only if α is a critical set and $K, P_\gamma \subset \mathcal{L}_\alpha$. But $P_\gamma \subset \mathcal{L}_\alpha$ implies $\gamma \subseteq \alpha$ by Lemma 9.4. Thus $\alpha = \gamma$ by the maximality of γ, which contradicts the assumption $K \not\subset \mathcal{L}_\gamma$. □

We can now formulate the key property of $\operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}$.

Lemma 11.6. *For every* $u \in U$*, we have*

$$\operatorname{supp} S_{P_\gamma,\mathcal{P}} \subseteq \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}},$$

and

$$\operatorname{span}\operatorname{supp} S_{P_\gamma,\mathcal{P}} = \mathcal{L}_\gamma^\perp.$$

Proof. We begin by noting that $\operatorname{supp} S_{[0,u],P_\gamma,\mathcal{P}_{\backslash r}} \subseteq \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}$ by Lemma 2.4. As $u \in \mathcal{L}_r \subseteq \mathcal{L}_\gamma$ by the definition of U, we have $[0, u], P_\gamma \subset \mathcal{L}_\gamma$, so that

$$\dim([0, u] + P_\gamma + \textstyle\sum_{i\in\gamma\backslash\{r\}} P_i) = \dim \mathcal{L}_\gamma = |\gamma| + 1.$$

Applying Lemma 3.8 as in Remark 8.6 yields

$$\binom{n-1}{|\gamma|+1} S_{[0,u],P_\gamma,\mathcal{P}_{\backslash r}} = \mathsf{V}_{\mathcal{L}_\gamma}([0,u],P_\gamma,\mathcal{P}_{\gamma\backslash\{r\}})\, S_{\mathbf{P}_{\mathcal{L}_\gamma^\perp}\mathcal{P}_{\backslash\gamma}},$$
$$\binom{n-1}{|\gamma|+1} S_{P_\gamma,\mathcal{P}} = \mathsf{V}_{\mathcal{L}_\gamma}(P_\gamma,\mathcal{P}_\gamma)\, S_{\mathbf{P}_{\mathcal{L}_\gamma^\perp}\mathcal{P}_{\backslash\gamma}}.$$

But as $\mathcal{P}$ is critical, $(P_\gamma,\mathcal{P}_{\backslash r})$ is critical as well, so $\mathsf{V}_{\mathcal{L}_\gamma}([0,u],P_\gamma,\mathcal{P}_{\gamma\backslash\{r\}})>0$ and $\mathsf{V}_{\mathcal{L}_\gamma}(P_\gamma,\mathcal{P}_\gamma)>0$ by Lemma 2.2. Thus we have proved the first claim.

Now note that the second identity above implies $\operatorname{span}\operatorname{supp} S_{P_\gamma,\mathcal{P}}\subseteq\mathcal{L}_\gamma^\perp$. If the inclusion were to be strict, then $\operatorname{supp} S_{P_\gamma,\mathcal{P}}\subset w^\perp$ for some $w\in S^{n-1}\cap\mathcal{L}_\gamma^\perp$, so

$$0=\int\langle w,x\rangle_+\,S_{P_\gamma,\mathcal{P}}(dx)=n\,\mathsf{V}_n([0,w],P_\gamma,\mathcal{P})$$

using $h_{[0,w]}(x)=\langle w,x\rangle_+$ and (2.1). As $[0,w]\not\subseteq\mathcal{L}_\gamma$, this entails a contradiction by Lemma 11.5. Thus the second claim is proved. □

Combining the above results, we conclude the following.

Corollary 11.7. *If the conclusion of Lemma 11.4 holds, then for any* $u,v\in U$

$$\langle s(u)-s(v),x\rangle+\sum_{j=1}^{\ell}\{g_{j,u}(x)-g_{j,v}(x)\}=0\quad\textit{for all } x\in\operatorname{supp}S_{P_\gamma,\mathcal{P}}.$$

Proof. As $g_{0,u}$ is a γ-degenerate function, $g_{0,u}(x)=g_{0,u}(\mathbf{P}_{\mathcal{L}_\gamma}x)=0$ for all $x\in\mathcal{L}_\gamma^\perp$ by Lemma 9.9. The claim follows immedately from Lemmas 11.4 and 11.6. □

Let us emphasize that the conclusion of Corollary 11.7 is of a fundamentally different nature than the analogous statement in the supercritical case: as the mixed area measure that appears here does not depend on u,v, there is no need to "glue" the functions $\langle s(u),\cdot\rangle+\sum_{j=1}^{\ell}g_{j,u}$ for different u. (We will still have to solve a gluing problem for $g_{0,u}$, which is addressed in the next section.) The problem we face here is that Corollary 11.7 only provides information on $\operatorname{supp}S_{P_\gamma,\mathcal{P}}$, while we must characterize these functions on $\operatorname{supp}S_{B,\mathcal{P}}$ to prove Theorem 11.1. As a first step towards this goal, let us make the following basic observation.

Lemma 11.8. *Let h be any α_i-degenerate function for some $i=1,\ldots,\ell$. Then the following statements are equivalent:*

a. $h(x)=0$ *for all* $x\in\operatorname{supp}S_{P_\gamma,\mathcal{P}}$.
b. $h(x)=0$ *for all* $x\in\operatorname{supp}S_{\mathcal{P}_{\alpha_i}}$.
c. $h(x)=0$ *for all* $x\in\operatorname{supp}S_{B,B,\mathcal{P}_{\backslash r}}$.

Proof. Lemmas 9.7 and 9.9 imply that

$$\binom{n-1}{|\alpha_i|}\int|h|\,dS_{B,B,\mathcal{P}_{\backslash r}}=\mathsf{V}_{\mathcal{L}_{\alpha_i}^\perp}(\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}B,\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}B,\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}\mathcal{P}_{\backslash\{\alpha_i,r\}})\int|h|\,dS_{\mathcal{P}_{\alpha_i}},$$
$$\binom{n-1}{|\alpha_i|}\int|h|\,dS_{P_\gamma,\mathcal{P}}=\mathsf{V}_{\mathcal{L}_{\alpha_i}^\perp}(\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}P_\gamma,\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}\mathcal{P}_{\backslash\alpha_i})\int|h|\,dS_{\mathcal{P}_{\alpha_i}}.$$

Thus the conclusion follows provided the two mixed volumes in the above identities are positive. To show that this is in fact the case, we note that

$$\mathsf{V}_{\mathcal{L}_{\alpha_i}}(B_{\alpha_i},\mathcal{P}_{\alpha_i})\,\mathsf{V}_{\mathcal{L}_{\alpha_i}^\perp}(\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}P_\gamma,\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}\mathcal{P}_{\backslash\alpha_i})=\binom{n}{|\alpha_i|+1}\mathsf{V}_n(B_{\alpha_i},P_\gamma,\mathcal{P})>0$$

by Lemmas 3.8 and 11.5 (here we used $B_{\alpha_i} \not\subset \mathcal{L}_\gamma$ by Corollary 9.3). Therefore

$$c\,\mathsf{V}_{\mathcal{L}_{\alpha_i}^\perp}(\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}B, \mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}B, \mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}\mathcal{P}_{\backslash\{\alpha_i,r\}}) \geqslant \mathsf{V}_{\mathcal{L}_{\alpha_i}^\perp}(\mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}P_\gamma, \mathbf{P}_{\mathcal{L}_{\alpha_i}^\perp}\mathcal{P}_{\backslash\alpha_i}) > 0,$$

for some $c > 0$, where we used $P_\gamma \subset \mathrm{diam}(P_\gamma)B$ and $P_r \subset \mathrm{diam}(P_r)B$. $\square$

The difficulty in the application of Lemma 11.8 is that it applies only to an individual α_i-degenerate function, which is used crucially in its proof. On the other hand, Corollary 11.7 involves a sum of degenerate functions for different sets α_i. The problem we must now address is therefore to decouple the different terms of the sum in Corollary 11.7. This will be accomplished using Proposition 10.7.

Proposition 11.9. *Let $t \in \mathbb{R}^n$, and let f_j be an α_j-degenerate function for every $j = 1, \ldots, \ell$. Suppose that we have*

$$\langle t, x\rangle + \sum_{j=1}^{\ell} f_j(x) = 0 \quad \textit{for all } x \in \mathrm{supp}\, S_{P_\gamma,\mathcal{P}}. \tag{11.2}$$

Then for every $j = 1, \ldots, \ell$, there exists $w_j \in \mathcal{L}_{\alpha_j}$ so that

$$f_j(x) = \langle w_j, x\rangle \quad \textit{for all } x \in \mathrm{supp}\, S_{B,B,\mathcal{P}_{\backslash r}}$$

and $t + w_1 + \cdots + w_\ell \in \mathcal{L}_\gamma$.

Proof. Fix $i \in [\ell]$ until further notice.

Step 1. We aim to apply Proposition 10.7 with $(C_1, \ldots, C_{n-1}) \leftarrow (P_\gamma, \mathcal{P})$, $(C_1, \ldots, C_k) \leftarrow \mathcal{P}_{\alpha_i}$, and $E \leftarrow \mathcal{L}_{\alpha_i}$ to the function

$$h(x) := \langle t, x\rangle + \sum_{j \neq i} f_j(x).$$

To this end, let us verify the assumptions of Proposition 10.7 are satisfied. The requisite criticality assumptions follow immediately as $\mathcal{P}$ is critical and α_i is a critical set. Now note that as f_i is α_i-degenerate, it follows from Lemma 9.9 that $f_i(x) = 0$ for $x \in \mathcal{L}_{\alpha_i}^\perp$. Thus (11.2) implies that

$$h(x) = 0 \quad \text{for all } x \in \mathcal{L}_{\alpha_i}^\perp \cap \mathrm{supp}\, S_{P_\gamma,\mathcal{P}}.$$

On the other hand, note that by Lemmas 9.7 and 9.9, we have

$$\binom{n-1}{|\alpha_j|} \int f_j\, dS_{\mathcal{Q}_{\alpha_i},P_\gamma,\mathcal{P}_{\backslash\alpha_i}} = \mathsf{V}_{\mathcal{L}_{\alpha_j}^\perp}(\mathbf{P}_{\mathcal{L}_{\alpha_j}^\perp}\mathcal{Q}_{\alpha_i}, \mathbf{P}_{\mathcal{L}_{\alpha_j}^\perp}P_\gamma, \mathbf{P}_{\mathcal{L}_{\alpha_j}^\perp}\mathcal{P}_{\backslash\{\alpha_i,\alpha_j\}}) \int f_j\, dS_{\mathcal{P}_{\alpha_j}}$$

for any $j \neq i$ and convex bodies $\mathcal{Q}_{\alpha_i} = (Q_l)_{l\in\alpha_i}$ in $\mathcal{L}_{\alpha_i}$. But the right-hand side vanishes as $\mathsf{V}_{\mathcal{L}_{\alpha_j}}(f_j, \mathcal{P}_{\alpha_j}) = 0$ by the definition of an α_j-degenerate function; thus

$$\int h\, dS_{Q,\ldots,Q,P_\gamma,\mathcal{P}_{\backslash\alpha_i}} = 0$$

for any full-dimensional polytope Q in $\mathcal{L}_{\alpha_i}$, where we used that the integral of the linear part of h vanishes by Lemma 3.2. The assumptions of Proposition 10.7 are therefore satisfied. Consequently, there exists $w_i' \in \mathcal{L}_{\alpha_i}$ so that

$$\int h(x)\, 1_{\langle z,x\rangle>0}\, S_{\mathbf{P}_{F_z}P_\gamma,\mathbf{P}_{F_z}\mathcal{P}_{\backslash\alpha_i}}(dx) = \langle w_i', z\rangle \tag{11.3}$$

for all $z \in S^{n-1} \cap \mathcal{L}_{\alpha_i}$, where $F_z := \mathrm{span}\{z, \mathcal{L}_{\alpha_i}^\perp\}$.

Step 2. Now note that $\mathbf{P}_{\mathcal{L}_{\alpha_i}}x = \langle x, z\rangle z$ for every $x \in F_z$. Thus

$$f_i(x) = f_i(z)\langle z, x\rangle \quad \text{for all } x \in F_z^+$$

by Lemma 9.9. On the other hand, (11.2) and Theorem 10.1 imply that

$$h(x) + f_i(x) = 0 \quad \text{for all } x \in \operatorname{supp}(1_{\langle z,\cdot\rangle>0}\, dS_{\mathsf{P}_{F_z} P_\gamma, \mathsf{P}_{F_z}\mathcal{P}_{\backslash\alpha_i}})$$

holds for every $z \in \operatorname{supp} S_{\mathcal{P}_{\alpha_i}}$. Substituting these identities in (11.3) yields

$$\mathsf{V}_{\mathcal{L}^\perp_{\alpha_i}}(\mathsf{P}_{\mathcal{L}^\perp_{\alpha_i}} P_\gamma, \mathsf{P}_{\mathcal{L}^\perp_{\alpha_i}} \mathcal{P}_{\backslash\alpha_i})\, f_i(z) = -\langle w_i', z\rangle$$

for every $z \in \operatorname{supp} S_{\mathcal{P}_{\alpha_i}}$, where we used Corollary 3.9. But we already showed in the proof of Lemma 11.8 that $\mathsf{V}_{\mathcal{L}^\perp_{\alpha_i}}(\mathsf{P}_{\mathcal{L}^\perp_{\alpha_i}} P_\gamma, \mathsf{P}_{\mathcal{L}^\perp_{\alpha_i}} \mathcal{P}_{\backslash\alpha_i}) > 0$. Therefore

$$f_i(z) = \langle w_i, z\rangle \quad \text{for all } z \in \operatorname{supp} S_{\mathcal{P}_{\alpha_i}},$$

where $w_i := -\mathsf{V}_{\mathcal{L}^\perp_{\alpha_i}}(\mathsf{P}_{\mathcal{L}^\perp_{\alpha_i}} P_\gamma, \mathsf{P}_{\mathcal{L}^\perp_{\alpha_i}} \mathcal{P}_{\backslash\alpha_i})^{-1} w_i' \in \mathcal{L}_{\alpha_i}$.

Step 3. As $i \in [\ell]$ was arbitrary, we have now constructed for every $j = 1, \ldots, \ell$ a vector $w_j \in \mathcal{L}_{\alpha_j}$ such that $f_j(z) - \langle w_j, z\rangle = 0$ for $z \in \operatorname{supp} S_{\mathcal{P}_{\alpha_j}}$. As $f_j - \langle w_j, \cdot\rangle$ is an α_j-degenerate function, we conclude by Lemma 11.8 that

$$f_j(x) = \langle w_j, x\rangle \quad \text{for all } x \in \operatorname{supp} S_{B,B,\mathcal{P}_{\backslash r}} \text{ and } x \in \operatorname{supp} S_{P_\gamma,\mathcal{P}}.$$

In particular, we obtain using (11.2)

$$\langle t + w_1 + \cdots + w_\ell, x\rangle = 0 \quad \text{for all } x \in \operatorname{supp} S_{P_\gamma,\mathcal{P}}.$$

Thus $\mathsf{P}_{\mathcal{L}^\perp_\gamma}(t + w_1 + \cdots + w_\ell) = 0$ by Lemma 11.6, completing the proof. ☐

We can now put together all the ideas of this section.

Corollary 11.10. *Suppose that Theorem 11.1 has been proved in dimension $n-1$. Then there exist $s \in \mathbb{R}^n$, an α_j-degenerate function g_j for every $j = 1, \ldots, \ell$, and $t(u) \in \mathcal{L}_\gamma$ for every $u \in U$, so that the following holds for all $u \in U$:*

$$g(x) - \langle s, x\rangle - \sum_{j=1}^{\ell} g_j(x) = \langle t(u), x\rangle + g_{0,u}(x) \quad \textit{for all } x \in \operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}.$$

Proof. Fix any $v \in U$, and define $s := s(v)$ and $g_j := g_{j,v}$ for $j = 1, \ldots, \ell$. Then Lemma 11.4, Corollary 11.7 and Proposition 11.9 yield $w_j(u) \in \mathcal{L}_{\alpha_j}$ so that

$$g_{j,u}(x) = g_j(x) + \langle w_j(u), x\rangle \quad \text{for all } x \in \operatorname{supp} S_{B,B,\mathcal{P}_{\backslash r}}$$

and

$$t(u) := s(u) - s + w_1(u) + \cdots + w_\ell(u) \in \mathcal{L}_\gamma$$

for every $u \in U$. As $\operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}} \subseteq \operatorname{supp} S_{B,B,\mathcal{P}_{\backslash r}}$ for every u by Lemma 2.4, the conclusion follows immediately from Lemma 11.4. ☐

11.3. **The gluing argument.** With Corollary 11.10 in hand, it remains to glue the γ-degenerate functions $g_{0,u}$ for different $u \in U$. This will be accomplished using Lemma 10.6. We remind the reader that the function f with $S_{f,\mathcal{P}} = 0$ was fixed at the beginning of the proof, and that g was constructed from f by Lemma 11.2.

Lemma 11.11. *Suppose that Theorem 11.1 has been proved in dimension $n-1$. Then there exist $s \in \mathbb{R}^n$ and an α_j-degenerate function g_j for $j = 0, \ldots, \ell$ so that*

$$f(x) = \langle s, x\rangle + \sum_{j=0}^{\ell} g_j(x) \quad \textit{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}.$$

Proof. We use the notations of Corollary 11.10 throughout the proof.

Step 1. We begin by applying Lemma 10.6 with $(C_{k+1},\ldots,C_{n-1}) \leftarrow (B,\mathcal{P}_{\backslash\gamma})$ and $E \leftarrow \mathcal{L}_\gamma$ to the function

$$h(x) := g(x) - \langle s,x\rangle - \sum_{j=1}^{\ell} g_j(x).$$

Note that

$$\mathsf{V}_{\mathcal{L}_\gamma}(P_\gamma,\mathcal{P}_\gamma)\,\mathsf{V}_{\mathcal{L}_\gamma^\perp}(\mathbf{P}_{\mathcal{L}_\gamma^\perp}B,\mathbf{P}_{\mathcal{L}_\gamma^\perp}\mathcal{P}_{\backslash\gamma}) = \binom{n}{|\gamma|+1}\mathsf{V}_n(B,P_\gamma,\mathcal{P}) > 0 \tag{11.4}$$

by Lemmas 3.8 and 11.5, so the assumptions of Lemma 10.6 are satisfied. Denote by $\varphi:\mathcal{L}_\gamma\to\mathbb{R}$ the 1-homogeneous function constructed from h by Lemma 10.6.

Step 2. Now fix any $u\in U$, and choose

$$\tilde\varphi(x) := \langle t(u),x\rangle + g_{0,u}(x).$$

As $g_{0,u}$ is a γ-degenerate function, it follows from Lemma 9.9 and $t(u)\in\mathcal{L}_\gamma$ that $\tilde\varphi(x)=\tilde\varphi(\mathbf{P}_{\mathcal{L}_\gamma}x)$. Corollary 11.10 therefore states that

$$h(x) = \tilde\varphi(\mathbf{P}_{\mathcal{L}_\gamma}x)\quad\text{for all } x\in\operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}.$$

Recalling that $[0,u]\subset\mathcal{L}_\gamma$ by the definition of U, we can apply the conclusion of Lemma 10.6 with $(K_1,\ldots,K_k)\leftarrow([0,u],\mathcal{P}_{\gamma\backslash\{r\}})$ to obtain

$$g(x) = \langle s,x\rangle + \sum_{j=1}^{\ell} g_j(x) + \varphi(\mathbf{P}_{\mathcal{L}_\gamma}x)\quad\text{for all } x\in\operatorname{supp} S_{[0,u],B,\mathcal{P}_{\backslash r}}$$

for any $u\in U$. Moreover, as U has full measure in $S^{n-1}\cap\mathcal{L}_r$ by Lemma 11.3, we may further integrate over $u\in U$ as in the proof of Lemma 8.10 to conclude that the previous identity remains valid for all $x\in\operatorname{supp} S_{B,B_r,\mathcal{P}_{\backslash r}}$. It follows that

$$f(x) = g(x) = \langle s,x\rangle + \sum_{j=1}^{\ell} g_j(x) + \varphi(\mathbf{P}_{\mathcal{L}_\gamma}x)\quad\text{for all } x\in\operatorname{supp} S_{B,\mathcal{P}} \tag{11.5}$$

by Lemma 8.11 and property 1 of Lemma 11.2.

Step 3. It remains to show that $\varphi(\mathbf{P}_{\mathcal{L}_\gamma}x)$ defines a γ-degenerate function. To this end, recall that $S_{f,\mathcal{P}}=0$ by assumption, and $S_{g_j,\mathcal{P}}=0$ for $j=1,\ldots,\ell$ by Lemmas 2.11 and 9.6. Thus $\int f\,dS_{B,\mathcal{P}} = \int g_j\,dS_{B,\mathcal{P}} = \int\langle s,\cdot\rangle\,dS_{B,\mathcal{P}} = 0$ by (2.1), the symmetry of mixed volumes, and Lemma 3.2. We therefore have

$$0 = \binom{n-1}{|\gamma|}\int\varphi(\mathbf{P}_{\mathcal{L}_\gamma}x)\,S_{B,\mathcal{P}}(dx) = \mathsf{V}_{\mathcal{L}_\gamma^\perp}(\mathbf{P}_{\mathcal{L}_\gamma^\perp}B,\mathbf{P}_{\mathcal{L}_\gamma^\perp}\mathcal{P}_{\backslash\gamma})\int\varphi\,dS_{\mathcal{P}_\gamma}$$

using (11.5) and Lemma 9.7. But as $\mathsf{V}_{\mathcal{L}_\gamma^\perp}(\mathbf{P}_{\mathcal{L}_\gamma^\perp}B,\mathbf{P}_{\mathcal{L}_\gamma^\perp}\mathcal{P}_{\backslash\gamma})>0$ by (11.4), we can apply Lemma 9.9 to construct a γ-degenerate function g_0 so that

$$g_0(x) = \varphi(\mathbf{P}_{\mathcal{L}_\gamma}x)\quad\text{for all } x\in\operatorname{supp} S_{B,\mathcal{P}}.$$

Subsituting this identity into (11.5) completes the proof. □

The proof of Theorem 11.1 is now readily completed.

Proof of Theorem 11.1. The *if* direction is an immediate consequence of Lemmas 2.7, 2.8, 2.11, and 9.6, so it suffices to consider the *only if* direction.

To this end, note first that the case $n = 2$ is always supercritical, so Theorem 11.1 is trivial in this case. For the induction step, it remains to prove that the validity of Theorem 11.1 in dimension $n-1$ implies its validity in dimension n for any $n \geqslant 3$. The latter is precisely the statement of Lemma 11.11. □

12. Proof of the main result

The aim of this section is to complete the proof of Theorem 2.13. While the results that we have proved in the supercritical and critial cases provide a lot more information on the extremals than can be read off from Theorem 2.13 (which will be described in full detail in section 13 below), the advantage of the formulation of Theorem 2.13 is that it unifies all the different cases that arise in our analysis in one simple and universal statement. What remains is to verify that all possible cases are in fact captured by the formulation of Theorem 2.13.

So far we have only considered the extremals under the criticality assumption on $\mathcal{P}$. The remaining cases turn out to be either trivial, or to reduce readily to a critical case in lower dimension. We will first investigate the latter phenomenon, and then put everything together to complete the proof of Theorem 2.13.

12.1. **Subcritical sets.** Let us begin by formally defining the cases that have not yet been considered in the previous sections.

Definition 12.1. A collection of convex bodies $\mathcal{C} = (C_1, \ldots, C_{n-2})$ is said to be *subcritical* if $\dim(C_{i_1} + \cdots + C_{i_k}) \geqslant k$ for all $k \in [n-2]$, $1 \leqslant i_1 < \cdots < i_k \leqslant n-2$. A collection of convex bodies that is not subcritical is called *null*.

In view of the following lemma, the null case is trivial.

Lemma 12.2. *Let $\mathcal{C} = (C_1, \ldots, C_{n-2})$ be a null collection of convex bodies in $\mathbb{R}^n$. Then $\mathsf{V}_n(K, L, \mathcal{C}) = 0$ for any convex bodies K, L.*

Proof. This is an immediate consequence of Lemma 2.2. □

In the remainder of this section, we fix $n \geqslant 3$ and a collection $\mathcal{P} = (P_1, \ldots, P_{n-2})$ of convex bodies in $\mathbb{R}^n$ that is subcritical but not critical. We will assume that P_i contains the origin in its relative interior for each $i \in [n-2]$, and define the spaces $\mathcal{L}_\alpha$ as in section 8. Thus the subcriticality assumption states that $\dim \mathcal{L}_\alpha \geqslant |\alpha|$ for every α, and the lack of criticality implies that $\dim \mathcal{L}_\alpha = |\alpha|$ for at least one set α. In analogy with the critical case, we introduce the following terminology.

Definition 12.3. $\alpha \subseteq [n-2]$ is called a *subcritical set* if $\dim(\sum_{i\in\alpha} P_i) = |\alpha|$.

The point of this definition is the following.

Lemma 12.4. *We have $\mathsf{V}_{\mathcal{L}_\alpha}(\mathcal{P}_\alpha) > 0$ and*

$$\binom{n}{|\alpha|} \mathsf{V}_n(K, L, \mathcal{P}) = \mathsf{V}_{\mathcal{L}_\alpha}(\mathcal{P}_\alpha)\, \mathsf{V}_{\mathcal{L}_\alpha^\perp}(\mathbf{P}_{\mathcal{L}_\alpha^\perp} K, \mathbf{P}_{\mathcal{L}_\alpha^\perp} L, \mathbf{P}_{\mathcal{L}_\alpha^\perp} \mathcal{P}_{\backslash\alpha})$$

for any convex bodies K, L and subcritical set $\alpha \subseteq [n-2]$.

Proof. That $\mathsf{V}_{\mathcal{L}_\alpha}(\mathcal{P}_\alpha) > 0$ follows from Lemma 2.2 and the assumption that $\mathcal{P}$ is subcritical. The second statement is an immediate consequence of Lemma 3.8. □

In other words, any subcritical set will factor out of all mixed volumes that appear in the Alexandrov-Fenchel inequality, reducing it to a lower-dimensional case. However, there may *a priori* be many subcritical sets, and it is also not clear what properties $\mathsf{P}_{\mathcal{L}_\alpha^\perp}\mathcal{P}_{\backslash\alpha}$ may have. The main aim of this section is to show that there is a special choice of subcritical set η so that $\mathsf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}$ is a *critical* collection of convex bodies, which reduces the study of extremals of the Alexandrov-Fenchel inequality in the subcritical case to the setting of Theorem 11.1.

To this end, we first prove a subcritical analogue of Lemma 9.2.

Lemma 12.5. *Let α, α' be subcritical sets. Then $\alpha \cup \alpha'$ is also a subcritical set.*

Proof. As $\mathcal{P}$ is subcritical, we have $\dim \mathcal{L}_{\alpha\cup\alpha'} \geqslant |\alpha \cup \alpha'|$ and $\dim \mathcal{L}_{\alpha\cap\alpha'} \geqslant |\alpha \cap \alpha'|$ (note that the latter holds even when $\alpha \cap \alpha' = \varnothing$, unlike in Lemma 9.2). Thus

$$\begin{aligned}|\alpha \cup \alpha'| &\leqslant \dim \mathcal{L}_{\alpha\cup\alpha'} = \dim \mathcal{L}_\alpha + \dim \mathcal{L}_{\alpha'} - \dim(\mathcal{L}_\alpha \cap \mathcal{L}_{\alpha'}) \\ &\leqslant \dim \mathcal{L}_\alpha + \dim \mathcal{L}_{\alpha'} - \dim \mathcal{L}_{\alpha\cap\alpha'} \\ &\leqslant |\alpha| + |\alpha'| - |\alpha \cap \alpha'| = |\alpha \cup \alpha'|,\end{aligned}$$

where we used that $\dim \mathcal{L}_\alpha = |\alpha|$ and $\dim \mathcal{L}_{\alpha'} = |\alpha'|$ as α, α' are subcritical. It follows that $\dim \mathcal{L}_{\alpha\cup\alpha'} = |\alpha \cup \alpha'|$, so $\alpha \cup \alpha'$ is subcritical. □

Corollary 12.6. *There is a unique maximal subcritical set η, that is, a subcritical set η such that $\alpha \subseteq \eta$ for every subcritical set α.*

Proof. By Lemma 12.5, we may choose η to be the union of all subcritical sets. □

We now claim that applying Lemma 12.4 to the set η of Corollary 12.6 reduces the Alexandrov-Fenchel inequlity to the critical case.

Lemma 12.7. *Let η be as in Corollary 12.6. Then $\mathsf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}$ is critical.*

Proof. Suppose the conclusion is false; then there must exist $\alpha \subseteq [n-2]\backslash\eta$, $\alpha \neq \varnothing$ such that $\dim \mathsf{P}_{\mathcal{L}_\eta^\perp}\mathcal{L}_\alpha \leqslant |\alpha|$. Now note that

$$\dim \mathsf{P}_{\mathcal{L}_\eta^\perp}\mathcal{L}_\alpha = \dim \mathcal{L}_\alpha - \dim(\mathcal{L}_\alpha \cap \ker \mathsf{P}_{\mathcal{L}_\eta^\perp}) = \dim \mathcal{L}_\alpha - \dim(\mathcal{L}_\alpha \cap \mathcal{L}_\eta).$$

Therefore

$$\dim \mathcal{L}_{\eta\cup\alpha} = \dim \mathcal{L}_\eta + \dim \mathcal{L}_\alpha - \dim(\mathcal{L}_\alpha \cap \mathcal{L}_\eta) \leqslant |\eta| + |\alpha| = |\eta \cup \alpha|,$$

where we used that η is subcritical and $\eta \cap \alpha = \varnothing$. Thus we have shown that $\eta \cup \alpha$ is subcritical, which contradicts the maximality of η. □

12.2. Proof of Theorem 2.13.

We are now finally ready to conclude the proof.

Proof of Theorem 2.13. The *if* direction of Theorem 2.13 follows from Lemmas 2.5, 2.7, 2.8, and 2.11, so only the *only if* part requires proof.

By translation-invariance, we may assume without loss of generality that each P_i contains the origin in its relative interior. The assumption $\mathsf{V}_n(K, L, \mathcal{P}) > 0$ implies that $\mathcal{P}$ cannot be null by Lemma 12.2. We now consider the remaining cases.

Suppose first that $\mathcal{P}$ is supercritical. Then the conclusion follows immediately from Corollary 2.16 (which was proved in section 8) and Lemma 2.15.

Now suppose that $\mathcal{P}$ is critical but not supercritical, and denote by $\alpha_0, \ldots, \alpha_\ell$ the $\mathcal{P}$-maximal sets. By Lemma 2.5, the equality condition in Theorem 2.13 implies that there exists $a > 0$ such that $S_{f,\mathcal{P}} = 0$ for $f = h_K - ah_L$. Thus

$$h_{K+N_0+\cdots+N_\ell}(x) = h_{aL+s+M_0+\cdots+M_\ell}(x) \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}$$

by Theorem 11.1, where $s \in \mathbb{R}^n$ and (M_j, N_j) is an α_j-degenerate pair for every $j = 0, \ldots, \ell$. The conclusion follows readily from Lemma 9.6.

Finally, suppose that $\mathcal{P}$ is subcritical but not critical, and let η be the maximal subcritical set of Corollary 12.6. Then the assumption and equality condition of Theorem 2.13 imply that $\mathsf{V}_{\mathcal{L}_\eta^\perp}(\mathbf{P}_{\mathcal{L}_\eta^\perp}K, \mathbf{P}_{\mathcal{L}_\eta^\perp}L, \mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}) > 0$ and

$$\begin{aligned}&\mathsf{V}_{\mathcal{L}_\eta^\perp}(\mathbf{P}_{\mathcal{L}_\eta^\perp}K, \mathbf{P}_{\mathcal{L}_\eta^\perp}L, \mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta})^2 = \\ &\mathsf{V}_{\mathcal{L}_\eta^\perp}(\mathbf{P}_{\mathcal{L}_\eta^\perp}K, \mathbf{P}_{\mathcal{L}_\eta^\perp}K, \mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta})\,\mathsf{V}_{\mathcal{L}_\eta^\perp}(\mathbf{P}_{\mathcal{L}_\eta^\perp}L, \mathbf{P}_{\mathcal{L}_\eta^\perp}L, \mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta})\end{aligned}$$

by Lemma 12.4. As $\mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}$ is critical by Lemma 12.7, we can apply the critical case of Theorem 2.13 in $\mathcal{L}_\eta^\perp$ to conclude that we have

$$h_{K+N_0+\cdots+N_\ell}(x) = h_{aL+s+M_0+\cdots+M_\ell}(x) \quad \text{for all } x \in \operatorname{supp} S_{\mathbf{P}_{\mathcal{L}_\eta^\perp}B, \mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}} \tag{12.1}$$

for some $a > 0$, $s \in \mathcal{L}_\eta^\perp$, and $\mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}$-degenerate pairs (M_j, N_j) for $j = 0, \ldots, \ell$, where we used that $h_{\mathbf{P}_{\mathcal{L}_\eta^\perp}K}(x) = h_K(x)$ and $h_{\mathbf{P}_{\mathcal{L}_\eta^\perp}L}(x) = h_L(x)$ for $x \in \mathcal{L}_\eta^\perp$. But as

$$\binom{n-1}{|\eta|} S_{B,\mathcal{P}} = \mathsf{V}_{\mathcal{L}_\eta}(\mathcal{P}_\eta)\, S_{\mathbf{P}_{\mathcal{L}_\eta^\perp}B, \mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}} \tag{12.2}$$

by applying Lemma 12.4 as in Remark 8.6, and as $\mathsf{V}_{\mathcal{L}_\eta}(\mathcal{P}_\eta) > 0$ by Lemma 12.4, it follows that (12.1) remains valid for $x \in \operatorname{supp} S_{B,\mathcal{P}}$. Moreover, it follows readily from Definition 2.10 and Lemma 12.4 that any $\mathbf{P}_{\mathcal{L}_\eta^\perp}\mathcal{P}_{\backslash\eta}$-degenerate pair is also a $\mathcal{P}$-degenerate pair. Thus the conclusion of Theorem 2.13 is proved. □

Part 3. Complements and applications

13. The extremal decomposition

Theorem 2.13 gives a very general description of the extremals of the Alexandrov-Fenchel inequality in terms of degenerate pairs. There is significant redundancy in this formulation, however: the same extremal bodies may be decomposed into degenerate pairs in different ways. A much more informative description of the extremals arises as a consequence of the theory developed in the previous sections. The aim of the present section is to extract from the proof of Theorem 2.13 a non-redundant characterization of the extremal decomposition. This formulation may be viewed as the definitive form of the main result of this paper.

13.1. **A unique extremal characterization.** Throughout this section, we fix polytopes $\mathcal{P} = (P_1, \ldots, P_{n-2})$ in $\mathbb{R}^n$. By translation-invariance, we may assume without loss of generality that each P_i contains the origin in its relative interior. We also assume without loss of generality that $\mathcal{P}$ is subcritical (Definition 12.1), as otherwise the extremal problem is vacuous by Lemma 12.2.

The following structural properties of $\mathcal{P}$ were introduced in the previous sections. Recall that subcritical and maximal sets are defined in Definitions 12.3 and 9.1.

- $\mathcal{P}$ has a unique maximal subcritical set $\eta \subseteq [n-2]$ by Corollary 12.6; we define

$$\mathcal{L} := \operatorname{span} \sum_{i \in \eta} P_i.$$

If $\mathcal{P}$ is critical (Definition 4.1), then $\eta = \varnothing$ and $\mathcal{L} = \{0\}$.

- The collection $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$ is critical by Lemma 12.7. Thus by Corollary 9.3, there are $c \geqslant 0$ disjoint $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$-maximal sets $\beta_1, \ldots, \beta_c \subseteq [n-2]\backslash\eta$; we define

$$\mathcal{L}_j := \operatorname{span} \sum_{i\in\beta_j} P_i.$$

If $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$ is supercritical (Definition 2.14), then $c = 0$.

Let us now introduce the spaces of degenerate functions (Definition 9.5)

$$\begin{aligned}\mathbb{D}_j := \{f : S^{n-1} \to \mathbb{R} \ : f \text{ is a } (\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}, \beta_j)\text{-degenerate function},\\ f \perp \langle v, \cdot\rangle \text{ in } L^2(S_{B,\mathcal{P}}) \text{ for every } v \in \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j\}\end{aligned}$$

for $j = 1, \ldots, c$. By introducing the orthogonality condition with respect to linear functions, we eliminate one source of redundancy: that part of the linear term in Theorem 11.1 may be absorbed in the definitions of the degenerate functions. We will presently show that this is in fact the *only* source of redundancy; once it is eliminated, the extremal decomposition is uniquely determined on $\operatorname{supp} S_{B,\mathcal{P}}$. That is, we have the following unique characterization of the extremals of the Alexandrov-Fenchel inequality (in the formulation of part b of Lemma 2.5).

Theorem 13.1. *Let $f : S^{n-1} \to \mathbb{R}$ be any difference of support functions. Then*

$$S_{f,\mathcal{P}} = 0$$

holds if and only if

$$f(x) = \langle s, x\rangle + \sum_{j=1}^{c} f_j(x) \quad \textit{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}$$

holds for some $s \in \mathcal{L}^\perp$ and $f_j \in \mathbb{D}_j$, $j = 1, \ldots, c$. Moreover, in this representation, s is unique and $f_1, \ldots, f_c$ are uniquely determined on $\operatorname{supp} S_{B,\mathcal{P}}$.

Before we prove Theorem 13.1, let us first record a basic property. Analogous arguments appeared already several times in the proof of Theorem 2.13.

Lemma 13.2. $\operatorname{span}\operatorname{supp} S_{B,\mathcal{P}} = \mathcal{L}^\perp$.

Proof. We first observe that $\operatorname{span}\operatorname{supp} S_{B,\mathcal{P}} \subseteq \mathcal{L}^\perp$ (this is trivial if $\eta = \varnothing$, and follows directly from (12.2) otherwise). Now suppose this inclusion is strict. Then we must have $\operatorname{supp} S_{B,\mathcal{P}} \subseteq w^\perp$ for some $w \in \mathcal{L}^\perp$, so that

$$0 = \binom{n}{|\eta|}\frac{1}{n}\int \langle w, x\rangle_+ \, S_{B,\mathcal{P}}(dx) = \mathsf{V}_{\mathcal{L}}(\mathcal{P}_\eta)\,\mathsf{V}_{\mathcal{L}^\perp}([0,w], \mathsf{P}_{\mathcal{L}^\perp}B, \mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta})$$

by (2.1) and Lemma 12.4 (if $\eta = \varnothing$, this expression remains valid with $\mathsf{V}_{\mathcal{L}}(\mathcal{P}_\eta) \equiv 1$). This contradicts criticality of $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$ by Lemma 2.2, establishing the claim. $\square$

We now turn to the proof of Thereom 13.1.

Proof of Theorem 13.1. We first note that $S_{f,\mathcal{P}} = 0$ if and only if $S_{\mathsf{P}_{\mathcal{L}^\perp}f,\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}} = 0$ by applying Lemma 12.4 as in Remark 8.6. Thus as $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$ is critical, we can use either Theorem 8.1 or Theorem 11.1 in $\mathcal{L}^\perp$ to show that $S_{f,\mathcal{P}} = 0$ if and only if

$$f(x) = \langle \bar{s}, x\rangle + \sum_{j=1}^{c} g_j(x) \quad \text{for all } x \in \operatorname{supp} S_{\mathsf{P}_{\mathcal{L}^\perp}B,\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}}$$

for some $\bar{s} \in \mathcal{L}^\perp$ and $(\mathbf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}, \beta_j)$-degenerate function g_j, $j = 1, \ldots, c$. The conclusion remains valid for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$ as $\operatorname{supp} S_{B,\mathcal{P}} = \operatorname{supp} S_{\mathbf{P}_{\mathcal{L}^\perp} B, \mathbf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}}$ (this is trivial when $\eta = \varnothing$, and follows from (12.2) otherwise).

Now note that for every $j = 1, \ldots, c$, there exists $s_j \in \mathbf{P}_{\mathcal{L}^\perp}\mathcal{L}_j$ so that

$$f_j := g_j - \langle s_j, \cdot\rangle \;\perp\; \operatorname{span}\{\langle v, \cdot\rangle : v \in \mathbf{P}_{\mathcal{L}^\perp}\mathcal{L}_j\} \quad \text{in } L^2(S_{B,\mathcal{P}}).$$

Then $f_j \in \mathbb{D}_j$ by construction. Moreover, let us write $s := \bar{s} + s_1 + \cdots + s_c \in \mathcal{L}^\perp$. Then we have shown that $S_{f,\mathcal{P}} = 0$ holds if and only if

$$f(x) = \langle s, x\rangle + \sum_{j=1}^{c} f_j(x) \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}$$

for some $s \in \mathcal{L}^\perp$ and $f_j \in \mathbb{D}_j$, $j = 1, \ldots, c$. It remains to show uniqueness.

To this end, let us suppose that

$$f(x) = \langle s, x\rangle + \sum_{j=1}^{c} f_j(x) = \langle s', x\rangle + \sum_{j=1}^{c} f'_j(x) \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}$$

holds for $s, s' \in \mathcal{L}^\perp$ and $f_j, f'_j \in \mathbb{D}_j$, $j = 1, \ldots, c$. Then we certainly have

$$\langle s - s', x\rangle + \sum_{j=1}^{c} \{f_j(x) - f'_j(x)\} = 0 \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{P}} = \operatorname{supp} S_{\mathbf{P}_{\mathcal{L}^\perp} B, \mathbf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}}.$$

Suppose first that $c \geqslant 1$. Then by Lemma 2.4, we may apply Proposition 11.9 in $\mathcal{L}^\perp$ to show that for every $j = 1, \ldots, c$, there exists $w_j \in \mathbf{P}_{\mathcal{L}^\perp}\mathcal{L}_j$ so that

$$f_j(x) - f'_j(x) = \langle w_j, x\rangle \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{P}}.$$

But the definition of $\mathbb{D}_j$ then implies that $w_j = 0$, so that each f_j is uniquely determined on $\operatorname{supp} S_{B,\mathcal{P}}$. Moreover, for any $c \geqslant 0$, we now obtain $\langle s - s', x\rangle = 0$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$. Thus $s = s'$ by Lemma 13.2, so that s is unique as well. $\square$

13.2. **The space of extremals.** In addition to the spaces $\mathbb{D}_j$, let us denote by

$$\mathbb{X} := \{f : S^{n-1} \to \mathbb{R} \;:\; f \text{ is a difference of support functions such that } S_{f,\mathcal{P}} = 0\}$$

the space of extremals of the Alexandrov-Fenchel inequality, and by

$$\mathbb{L} := \{f : S^{n-1} \to \mathbb{R} \;:\; f = \langle v, \cdot\rangle \text{ for some } v \in \mathcal{L}^\perp\}$$

the space of linear functions. In the present section, we will view $\mathbb{X}, \mathbb{L}, \mathbb{D}_j$ as subspaces of $L^2(S_{B,\mathcal{P}})$; in particular, we identify the elements of these spaces that agree $S_{B,\mathcal{P}}$-a.e. In these terms, we may reformulate Theorem 13.1 as

$$\mathbb{X} = \mathbb{L} \oplus \mathbb{D}_1 \oplus \cdots \oplus \mathbb{D}_c \quad \text{in } L^2(S_{B,\mathcal{P}}).$$

That is, $\mathbb{L}, \mathbb{D}_1, \ldots, \mathbb{D}_c$ are linearly independent subspaces of $L^2(S_{B,\mathcal{P}})$ that span $\mathbb{X}$.

Despite the definitive form of this result, its continuous formulation belies the essentially combinatorial nature of the extremals of the Alexandrov-Fenchel inequality for polytopes $\mathcal{P}$. For example, Proposition 5.7 implies that the extremals are fully described by the kernel of a matrix, so that $\mathbb{X}$ must be finite-dimensional; this fact is not evident above. To illustrate that this kind of information is indeed contained in the above characterization, let us presently compute the dimensions of the subspaces $\mathbb{L}, \mathbb{D}_j, \mathbb{X}$ of $L^2(S_{B,\mathcal{P}})$ in terms of the geometry of $\mathcal{P}$.

Proposition 13.3. *Let* $\Omega_j := \operatorname{supp} S_{\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\beta_j}}$. *Then*

$$\dim \mathbb{L} = n - |\eta|, \qquad \dim \mathbb{D}_j = |\Omega_j| - |\beta_j| - 2$$

for $j = 1, \ldots, c$. *In particular,* $\dim \mathbb{X} = n - |\eta| + \sum_{j=1}^{c}\{|\Omega_j| - |\beta_j| - 2\}$.

Remark 13.4. Note that $|\Omega_j| < \infty$ and that Ω_j may be computed by Lemma 3.4.

To prove Proposition 13.3, we begin by characterizing the affine hull of Ω_j.

Lemma 13.5. $\operatorname{aff}\Omega_j = \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j$.

Proof. That $\operatorname{span}\Omega_j = \operatorname{span}\operatorname{supp} S_{\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\beta_j}} = \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j$ follows as $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$ is critical by exactly the same argument as in the proof of Lemma 13.2. On the other hand, by part d of Lemma 3.2, there is a vanishing linear combination of the elements of Ω_j with positive coefficients, so $0 \in \operatorname{aff}\Omega_j$. Thus $\operatorname{aff}\Omega_j = \operatorname{span}\Omega_j$. □

Next, we observe that the support functions of convex bodies in $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j$ are uniquely determined up to $S_{B,\mathcal{P}}$-a.e. equivalence by their values on Ω_j.

Lemma 13.6. *Let* $f = h_M - h_N$ *for convex bodies* $M, N \subset \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j$. *Then* $f = 0$ $S_{B,\mathcal{P}}$*-a.e. if and only if* $f(x) = 0$ *for all* $x \in \Omega_j$. *Moreover, for any* $z \in \mathbb{R}^{\Omega_j}$, *there exists a function* f *of this form such that* $f(x) = z_x$ *for* $x \in \Omega_j$.

Proof. We first apply (12.2) and Lemma 9.7 to write

$$\begin{aligned}
&\binom{n-1}{|\eta|}\binom{n-|\eta|-1}{|\beta_j|}\int |f|\, dS_{B,\mathcal{P}} \\
&= \binom{n-|\eta|-1}{|\beta_j|}\mathsf{V}_{\mathcal{L}}(\mathcal{P}_\eta)\int |f|\, dS_{\mathsf{P}_{\mathcal{L}^\perp}B,\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}} \\
&= \mathsf{V}_{\mathcal{L}}(\mathcal{P}_\eta)\,\mathsf{V}_{\mathcal{L}^\perp\cap\mathcal{L}_j^\perp}(\mathsf{P}_{\mathcal{L}^\perp\cap\mathcal{L}_j^\perp}B, \mathsf{P}_{\mathcal{L}^\perp\cap\mathcal{L}_j^\perp}\mathcal{P}_{\backslash\{\eta,\beta_j\}})\int |f|\, dS_{\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\beta_j}}.
\end{aligned} \tag{13.1}$$

Now note that Lemma 12.4 implies that $\mathsf{V}_{\mathcal{L}}(\mathcal{P}_\eta) > 0$. On the other hand, we have $\mathsf{V}_{\mathcal{L}^\perp}(\mathsf{P}_{\mathcal{L}^\perp}B_j, \mathsf{P}_{\mathcal{L}^\perp}B, \mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}) > 0$ by Lemma 2.2 as $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$ is critical, where B_j denotes the unit ball in $\mathcal{L}_j$. Thus applying (13.1) with $f = h_{\mathsf{P}_{\mathcal{L}^\perp}B_j}$ shows that the mixed volumes on the last line of (13.1) are positive. We have therefore shown that $\int |f|\, dS_{B,\mathcal{P}} > 0$ if and only if $\int |f|\, dS_{\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\beta_j}} > 0$.

To prove the second claim, it suffices to note that as Ω_j is a finite set, there exists for any $z \in \mathbb{R}^{\Omega_j}$ a C^2 function $g : S^{n-1} \cap \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j \to \mathbb{R}$ such that $g(x) = z_x$ for all $x \in \Omega_j$; then $f(x) := g(\mathsf{P}_{\mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j}x)$ has the requisite form by Lemma 2.1. □

We can now conclude the proof of Proposition 13.3.

Proof of Proposition 13.3. We first note that for any $v \in \mathcal{L}^\perp$, we have $\langle v, \cdot\rangle = 0$ $S_{B,\mathcal{P}}$-a.e. if and only if $v = 0$ by Lemma 13.2. Thus $\dim \mathbb{L} = \dim \mathcal{L}^\perp = n - |\eta|$, where we used that $\dim \mathcal{L} = |\eta|$ as η is a subcritical set.

Now note that by Definition 9.5, a $(\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}, \beta_j)$-degenerate function f is defined by $f = h_M - h_N$ for convex bodies $M, N \subset \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j$ satisfying one linear constraint $\int f\, dS_{\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\beta_j}} = 0$. Thus by Lemma 13.6, the subspce of $L^2(S_{B,\mathcal{P}})$ defined by

$$\tilde{\mathbb{D}}_j := \{f : S^{n-1} \to \mathbb{R} \;:\; f \text{ is a } (\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}, \beta_j)\text{-degenerate function}\}$$

has dimension $\dim \tilde{\mathbb{D}}_j = |\Omega_j| - 1$. Moreover, it follows from Lemmas 13.5 and 13.6 that $\operatorname{span}\{\langle v, \cdot\rangle : v \in \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j\}$ is a subspace of $\tilde{\mathbb{D}}_j$ of dimension $\dim \mathsf{P}_{\mathcal{L}^\perp}\mathcal{L}_j = |\beta_j| + 1$, where we used that β_j is a $\mathsf{P}_{\mathcal{L}^\perp}\mathcal{P}_{\backslash\eta}$-critical set. Thus $\dim \mathbb{D}_j = |\Omega_j| - |\beta_j| - 2$ by the definition of $\mathbb{D}_j$. The remaining claim follows as $\mathbb{X} = \mathbb{L} \oplus \mathbb{D}_1 \oplus \cdots \oplus \mathbb{D}_c$. □

Remark 13.7. It has been emphasized throughout this paper that the distinction between the supercritical and critical cases of Theorem 2.13 is that nonlinear extremals appear in the latter case. It is sometimes possible, however, that no nonlinear extremals exist even in the critical case, because it may be the case that $|\Omega_j| = |\beta_j| + 2$ and thus $\mathbb{D}_j = \{0\}$ as a subspace of $L^2(S_{B,\mathcal{P}})$ by Proposition 13.3. For example, this is the case when $\mathcal{P}_{\beta_j} = (Q, \ldots, Q)$, where Q is a simplex in $\mathcal{L}^\perp$ of dimension $|\beta_j| + 1$. By the same token, however, Proposition 13.3 shows that this situation can occur only in very special cases: as soon as enough normal directions are in play, nonlinear extremals will always appear in the critical case.

14. Extensions to quermassintegrals, smooth bodies, and zonoids

In the proof of Theorem 2.13, the assumption that $\mathcal{P}$ are polytopes was used extensively in the proof of the local Alexandrov-Fenchel inequality. However, most other arguments of this paper are not specific to polytopes. As a byproduct of our methods, we will presently extend our main result to more general situations. In particular, we will characterize the extremals of Theorem 1.1 in the following cases:

- $C_1 = \cdots = C_m$ is arbitrary and $C_{m+1}, \ldots, C_{n-2}$ are smooth (Theorem 14.6). In particular, this settles the case of quermassintegrals (Corollary 14.8).
- $\mathcal{C}$ is any combination of polytopes, smooth bodies, and zonoids (Theorem 14.9).

In the interest of space, we will consider in this section only the supercritical case. Results may also be obtained for the critical case with additional work.

It should be emphasized that the analysis of this section will rely on the special structure of smooth bodies and zonoids; it does not address the main missing ingredient for extending our main results to general convex bodies, which is a general form of the local Alexandrov-Fenchel inequality (see section 16). However, the results of this section further illustrate the methods developed in this paper, and capture a number of cases that are important in applications.

14.1. **A smooth local Alexandrov-Fenchel principle.** A convex body C in $\mathbb{R}^n$ is called *smooth* if it has a unique normal vector at every point of its boundary. The main observation behind the extension of our results to cases that include smooth bodies is that an analogue of the local Alexandrov-Fenchel inequality may be obtained in this case by a direct argument. That this situation is rather special will be evident from the statement of the following result: in the smooth case, the naive inequality (4.1) holds, and none of the subtleties of Theorem 4.3 arise.

Proposition 14.1. *Let $\mathcal{C} = (C_1, \ldots, C_{n-2})$ be any convex bodies in $\mathbb{R}^n$ such that $\mathsf{V}_n(B, B, \mathcal{C}) > 0$, and suppose that C_r is smooth for a given $r \in [n-2]$. Then for any difference of support functions f so that $S_{f,\mathcal{C}} = 0$, we have $S_{f,f,\mathcal{C}_{\backslash r}} = 0$.*

The proof is based on an idea due to Schneider, whose basic step is the following.

Lemma 14.2. *Let C, M be smooth convex bodies in $\mathbb{R}^n$. Then there exist $\varepsilon, \delta > 0$ and a family $\{C^\tau\}_{\tau\in[-\varepsilon,\varepsilon]}$ of convex bodies in $\mathbb{R}^n$ so that $\|h_{C^\tau} - h_C\|_\infty \leqslant \delta|\tau|$ and*

$$\lim_{\tau\to 0} \frac{h_{C^\tau}(u) - h_C(u)}{\tau} = h_M(u) \quad \textit{for all } u \in S^{n-1}.$$

Proof. We define C^τ as $C^\tau := C + \tau M$ for $\tau \geqslant 0$ and $C^\tau := C \div (-\tau)M$ for $\tau < 0$, where $C \div A := \{x \in \mathbb{R}^n : x + A \subseteq C\}$ denotes Minkowski subtraction. That $\|h_{C^\tau} - h_C\|_\infty \leqslant \delta|\tau|$ is trivial for $\tau \geqslant 0$, and is shown in [30, p. 425] for $\tau < 0$. The remaining statement follows from [30, Lemma 7.5.4]. □

We can now conclude the proof of Proposition 14.1.

Proof of Proposition 14.1. Let M be any smooth body, and define the family C_r^τ as in Lemma 14.2 (with $C \leftarrow C_r$). Define the function

$$\varphi(\tau) := \mathsf{V}_n(f, B, \mathcal{C}_{\backslash r}, C_r^\tau)^2 - \mathsf{V}_n(f, f, \mathcal{C}_{\backslash r}, C_r^\tau)\,\mathsf{V}_n(B, B, \mathcal{C}_{\backslash r}, C_r^\tau).$$

Then $\varphi(\tau) \geqslant 0$ by Lemma 3.11, and $\varphi(0) = 0$ as $C_r^0 = C_r$ and $S_{f,\mathcal{C}} = 0$. Thus $\tau = 0$ is a local minimum of φ. It follows that

$$0 = \frac{d\varphi(\tau)}{d\tau}\bigg|_{\tau=0} = -\mathsf{V}_n(f, f, \mathcal{C}_{\backslash r}, M)\,\mathsf{V}_n(B, B, \mathcal{C}),$$

where we used Lemma 14.2, (2.1) and $S_{f,\mathcal{C}} = 0$ to compute the derivative. As M was an arbitrary smooth body and $\mathsf{V}_n(B, B, \mathcal{C}) > 0$, we have shown that

$$0 = \mathsf{V}_n(f, f, \mathcal{C}_{\backslash r}, g) = \frac{1}{n}\int g\, dS_{f,f,\mathcal{C}_{\backslash r}}$$

for every function g that is a difference of support functions of smooth bodies. As any $g \in C^2$ may be written in this manner by Lemma 2.1, the conclusion follows. □

One of the consequences of Proposition 14.1 is that for the characterization of extremals of the Alexandrov-Fenchel inequalities, all smooth bodies are indistinguishable. This conclusion is a variant of [30, Theorem 7.6.7].

Corollary 14.3. *Let $\mathcal{C} = (C_1, \ldots, C_n)$ be convex bodies in $\mathbb{R}^n$ so that $C_1, \ldots, C_m$ are smooth and $\mathsf{V}_n(B, B, \mathcal{C}) > 0$. Let $\mathcal{C}' = (C_1', \ldots, C_m', C_{m+1}, \ldots, C_n)$ for smooth bodies $C_1', \ldots, C_m'$. Then $\operatorname{supp} S_{B,\mathcal{C}} = \operatorname{supp} S_{B,\mathcal{C}'}$, and for any difference of support functions f we have $S_{f,\mathcal{C}} = 0$ if and only if $S_{f,\mathcal{C}'} = 0$.*

Proof. It suffices to prove the case $m = 1$, as the general case then follows by applying the result repeatedly. Note also that as smooth bodies are full-dimensional, $\mathsf{V}_n(B, B, \mathcal{C}) > 0$ if and only if $\mathsf{V}_n(K, K, \mathcal{C}') > 0$ for any smooth body K.

Let f be a difference of support functions such that $S_{f,\mathcal{C}} = 0$. By integrating the mixed area measures in Proposition 14.1 (with $r = 1$) against $h_{C_1'}$, we obtain $\mathsf{V}_n(f, C_1, \mathcal{C}') = 0$ and $\mathsf{V}_n(f, f, \mathcal{C}') = 0$. Thus Lemma 3.12 implies $S_{f,\mathcal{C}'} = 0$. The converse implication follows by reversing the roles of C_1, C_1'.

Now note that $x \notin \operatorname{supp} S_{B,\mathcal{C}}$ holds if and only if there is a nonnegative C^2 function f such that $f(x) > 0$ and $f(u) = 0$ for all $u \in \operatorname{supp} S_{B,\mathcal{C}}$. Suppose this is the case. Then $S_{f,\mathcal{C}} = 0$ by Lemma 2.8, so $S_{f,\mathcal{C}'} = 0$ as well. Integrating against h_B and using the symmetry of mixed volumes yields $\int f\, dS_{B,\mathcal{C}'} = 0$, and thus $f(u) = 0$ for $u \in \operatorname{supp} S_{B,\mathcal{C}'}$ as f is nonnegative and continuous. It follows that $x \notin \operatorname{supp} S_{B,\mathcal{C}'}$. The converse implication follows again by reversing the roles of C_1, C_1'. □

14.2. **Quermassintegrals.** *Quermassintegrals* of a convex body K, defined by

$$W_i(K) := \mathsf{V}_n(\underbrace{K, \ldots, K}_{n-i}, \underbrace{B, \ldots, B}_{i}),$$

play a special role in convexity and in integral geometry; see, e.g., [16, §6.4]. The Alexandrov-Fenchel inequality implies that quermassintegrals form a log-concave sequence, that is, $W_i(K)^2 \geqslant W_{i-1}(K)\, W_{i+1}(K)$. Even in this very special case, the extremal bodies K have been characterized only in the presence of symmetry assumptions [30, Theorem 7.6.20]. In this section, we will settle this problem as a special case of a much more general result.

In the remainder of this section, we fix the following setting. Let $m \in [n-2]$, and let M be any convex body in $\mathbb{R}^n$ such that $\dim M \geqslant m+2$. We further let $C_{m+1}, \ldots, C_{n-2}$ be any smooth convex bodies in $\mathbb{R}^n$, and denote by

$$\mathcal{C} = (\underbrace{M, \ldots, M}_{m}, C_{m+1}, \ldots, C_{n-2}). \tag{14.1}$$

The main result of this section will characterize the extremal bodies K, L such that $\mathsf{V}_n(K, L, \mathcal{C})^2 = \mathsf{V}_n(K, K, \mathcal{C})\, \mathsf{V}_n(L, L, \mathcal{C})$. The reason we are able to do this for a general convex body M (rather than a polytope) relies on two observations. First, note that the gluing argument of section 8 works *verbatim* for general convex bodies, as long as a local Alexandrov-Fenchel inequality is available. We may therefore use the gluing argument together with Proposition 14.1 to reduce to the case $m = n-2$. The latter case, known as *Minkowski's quadratic inequality*, was settled in complete generality in [32], which enables us to conclude the proof.

Before we formulate this result precisely, let us provide a geometric characterization of the support of $S_{B,\mathcal{C}}$ in the present setting.

Definition 14.4. A vector $u \in S^{n-1}$ is called an *r-extreme normal vector* of a convex body M in $\mathbb{R}^n$ if there do not exist linearly independent normal vectors $u_1, \ldots, u_{r+2}$ at a boundary point of M such that $u = u_1 + \cdots + u_{r+2}$.

For example, if M is a polytope, then u is an r-extreme normal vector of M if and only if it is an outer normal of a face of K of dimension at least $n-1-r$.

Lemma 14.5. *Let $m \in [n-2]$, let M be a convex body in $\mathbb{R}^n$ with $\dim M \geqslant m$, let $C_{m+1}, \ldots, C_{n-2}$ be smooth convex bodies in $\mathbb{R}^n$, and let $\mathcal{C}$ be as in (14.1). Then*

$$\operatorname{supp} S_{B,\mathcal{C}} = \operatorname{cl}\{u \in S^{n-1} : u \text{ is an } (n-1-m)\text{-extreme normal vector of } M\}.$$

Proof. That $\mathsf{V}_n(B, B, \mathcal{C}) > 0$ follows from $\dim M \geqslant m$ and Lemma 2.2. Therefore, by Corollary 14.3, we may assume without loss of generality that $C_{m+1} = \cdots = C_{n-2} = B$. In the latter case, the result was proved in [25, Satz 4]. □

We can now formulate the main result of this section.

Theorem 14.6. *Let $m \in [n-2]$, let M be a convex body in $\mathbb{R}^n$ with $\dim M \geqslant m+2$, let $C_{m+1}, \ldots, C_{n-2}$ be smooth convex bodies in $\mathbb{R}^n$, and let $\mathcal{C}$ be as defined in (14.1). Then for any convex bodies K, L in $\mathbb{R}^n$ so that $\mathsf{V}_n(K, L, \mathcal{C}) > 0$, we have*

$$\mathsf{V}_n(K, L, \mathcal{C})^2 = \mathsf{V}_n(K, K, \mathcal{C})\, \mathsf{V}_n(L, L, \mathcal{C})$$

if and only if there exist $a > 0$ and $v \in \mathbb{R}^n$ so that K and $aL + v$ have the same supporting hyperplanes in all $(n-1-m)$-extreme normal directions of M.

Proof. By Lemmas 2.5 and 14.5, the conclusion is equivalent to the statement that for any differences of support functions f, we have $S_{f,\mathcal{C}} = 0$ if and only if there exists $s \in \mathbb{R}^n$ so that $f(x) = \langle s, x\rangle$ for all $s \in \operatorname{supp} S_{B,\mathcal{C}}$. The *if* direction follows from Lemmas 2.7 and 2.8, so it remains to prove the *only if* direction.

We will fix $m \geqslant 1$, and prove this statement by induction on n. The base case of the induction, $n = m+2$, is the main result of [32]. We now suppose $n > m+2$, and assume the induction hypothesis that the conclusion has been proved in dimension $n-1$. We aim to show that the conclusion then also holds in dimension n.

To this end, let f be a difference of support functions such that $S_{f,\mathcal{C}} = 0$, and let $r = n-2$. Then C_r is smooth, so $S_{f,f,\mathcal{C}_{\backslash r}} = 0$ by Proposition 14.1. Moreover, $\mathcal{C}$ is supercritical as $\dim M \geqslant m+2$. The argument of section 8 now applies *verbatim* in the present setting with $\mathcal{P} \leftarrow \mathcal{C}$ and $g \equiv f$ (indeed, that $\mathcal{P}$ are polytopes was used in section 8 only to apply the local Alexandrov-Fenchel inequality). In particular, it follows from Lemma 8.10 that there exists $s \in \mathbb{R}^n$ so that $f(x) = \langle s, x\rangle$ for all $x \in \operatorname{supp} S_{B,B,\mathcal{C}_{\backslash r}}$, and the conclusion now follows from Lemma 2.4. □

When specialized to quermassintegrals, we obtain the following. In the case that K is centrally symmetric, this result was proved in [30, Theorem 7.6.20].

Definition 14.7. A convex body K in $\mathbb{R}^n$ is a $(n-1-i)$-*tangential body of a ball* if there exist $a > 0$, $v \in \mathbb{R}^n$ so that $aB + v \subseteq K$, and such that $aB + v$ and K have the same supporting hyperplanes in all i-extreme normal directions of K.

Corollary 14.8. *Let K be any convex body in $\mathbb{R}^n$, and let $i \in [n-1]$. Then we have equality $W_i(K)^2 = W_{i-1}(K)\, W_{i+1}(K)$ if and only if either $\dim K < n-i$, or K is an $(n-1-i)$-tangential body of a ball.*

Proof. By Lemma 2.2, we have $W_i(K) = 0$ if and only if $\dim K < n-i$; in this case equality always holds. On the other hand, if $\dim K = n-i$, then $W_i(K) > 0$ and $W_{i-1}(K) = 0$, so equality cannot hold; and K cannot be a tangential body, as $aB + v \subseteq K$ implies that any tangential body satisfies $\dim K = n$. Finally, if $\dim K \geqslant n-i+1$, Theorem 14.6 implies that equality $W_i(K)^2 = W_{i-1}(K)\, W_{i+1}(K)$ holds if and only if there exist $a > 0$ and $v \in \mathbb{R}^n$ so that $aB + v$ and K have the same supporting hyperplanes in all i-extreme normal directions of K.

It remains to show that the latter condition implies *a fortiori* that $aB + v \subseteq K$. Indeed, if $K \subset w^\perp$ for some $w \in S^{n-1}$, then both w and $-w$ are i-extreme, so $aB + v \subset w^\perp$ as well. As that cannot be, we must have $\dim K = n$. But then [30, Theorem 2.2.6] implies that K is the intersection of its regular (and thus i-extreme) supporting halfspaces. As these also support $aB+v$, it follows that $aB+v \subseteq K$. □

14.3. **Zonoids.** A *zonotope* is a polytope that is the Minkowski sum of a finite number of segments. A convex body Z is called a *zonoid* if it is a limit of zonotopes. For simplicity, we assume by convention that all zonoids are symmetric $Z = -Z$ (this entails no loss of generality for our purposes, as any zonoid is symmetric up to translation). Then Z is a zonoid if and only if [30, Theorem 3.5.3]

$$h_Z(x) = \int |\langle u, x\rangle|\, \rho(du)$$

for some even finite measure on ρ on S^{n-1}, called the *generating measure* of Z.

When the reference bodies in the Alexandrov-Fenchel inequality are zonoids, their additive structure enables an inductive approach to the analysis of the extremals that is very special to this case. Such arguments were exploited by Schneider [27] to characterize the extremals for full-dimensional zonoids under additional symmetry assumptions. In this section, we will fully characterize the extremals for any supercritical collection of zonoids (the analysis of the critical case is more delicate, and is omitted here in the interest of space). In fact, the proof of the following more general result will present no additional difficulties.

Theorem 14.9. *Let $m \in [n-2]$, let $C_1, \ldots, C_m$ be convex bodies in $\mathbb{R}^n$ such that each C_i is either a zonoid or a smooth body, and let $P_{m+1}, \ldots, P_{n-2}$ be any polytopes in $\mathbb{R}^n$. Assume that $\mathcal{C} := (C_1, \ldots, C_m, P_{m+1}, \ldots, P_{n-2})$ is supercritical. Then for any convex bodies K, L in $\mathbb{R}^n$ so that $\mathsf{V}_n(K, L, \mathcal{C}) > 0$, we have*

$$\mathsf{V}_n(K, L, \mathcal{C})^2 = \mathsf{V}_n(K, K, \mathcal{C})\,\mathsf{V}_n(L, L, \mathcal{C})$$

if and only if there exist $a > 0$ and $v \in \mathbb{R}^n$ so that K and $aL + v$ have the same supporting hyperplanes in all directions in $\operatorname{supp} S_{B,\mathcal{C}}$.

It must be emphasized that in contrast to the settings of Theorems 2.13 and 14.6, we have not given a geometric characterization of $\operatorname{supp} S_{B,\mathcal{C}}$ in the general setting of Theorem 14.9. This problem is in fact not yet fully settled [30, Conjecture 7.6.14], and its analysis is outside the scope of this paper. However, we will provide such a characterization in Proposition 14.13 for the case $m = n-2$ that all reference bodies are zonoids (or smooth), completing an analysis due to Schneider [27].

The basis for the proof of Theorem 14.9 is a kind of analogue of Corollary 14.3. Analogous arguments may be found in [30, Lemma 7.4.7] and in [27, §4].

Lemma 14.10. *Let $\mathcal{K} = (K_1, \ldots, K_{n-3})$ be convex bodies in $\mathbb{R}^n$ and Z be a zonoid with generating measure ρ. Assume $(Z, \mathcal{K})$ is supercritical. Then any difference of support functions f so that $S_{f,Z,\mathcal{K}} = 0$ satisfies $S_{f,[-u,u],\mathcal{K}} = 0$ for all $u \in \operatorname{supp} \rho$.*

Proof. Let f be a difference of support functions such that $S_{f,Z,\mathcal{K}} = 0$. Then

$$\int \mathsf{V}_n(f, f, [-u, u], \mathcal{K})\,\rho(du) = \mathsf{V}_n(f, f, Z, \mathcal{K}) = 0,$$

where we used $h_{[-u,u]}(x) = |\langle u, x\rangle|$ and the definition of ρ. On the other hand, for any $u \in S^{n-1}$, we have $\mathsf{V}_n(f, Z, [-u, u], \mathcal{K}) = 0$ by (2.1) and $\mathsf{V}_n(Z, Z, [-u, u], \mathcal{K}) > 0$ by supercriticality and Lemma 2.2. Thus $\mathsf{V}_n(f, f, [-u, u], \mathcal{K}) \leqslant 0$ for any $u \in S^{n-1}$ by Lemma 3.11. We may therefore conclude that

$$\mathsf{V}_n(f, f, [-u, u], \mathcal{K}) = 0 \quad \text{for all } u \in \operatorname{supp} \rho,$$

where we used that $u \mapsto \mathsf{V}_n(f, f, [-u, u], \mathcal{K})$ is continuous by Lemma 3.3. In particular, it follows from Lemma 3.12 that $S_{f,[-u,u],\mathcal{K}} = 0$ for every $u \in \operatorname{supp} \rho$. □

We can now complete the proof of Theorem 14.9.

Proof of Theorem 14.9. By Corollary 14.3, we may assume without loss of generality that each C_i that is a smooth body satisfies $C_i = B$. But as B is a zonoid, we can assume in the remainder of the proof that $C_1, \ldots, C_m$ are zonoids with generating measures $\rho_1, \ldots, \rho_m$, respectively. We also fix a difference of support functions f so that $S_{f,\mathcal{C}} = 0$; by Lemmas 2.5, 2.7, and 2.8, it suffices to prove that there exists $s \in \mathbb{R}^n$ so that $f(x) = \langle s, x\rangle$ for all $s \in \operatorname{supp} S_{B,\mathcal{C}}$.

For each $i \in [m]$, we may choose (for example, using the law of large numbers) a sequence $\{u_i^j\}_{j \geqslant 1} \subseteq \operatorname{supp} \rho_i$ so that the zonotopes

$$C_i^j := j^{-1}\{[-u_i^1, u_i^1] + \cdots + [-u_i^j, u_i^j]\}$$

satisfy $C_i^j \to C_i$ as $j \to \infty$ in Hausdorff distance. In particular, $\operatorname{aff} C_i^j = \operatorname{aff} C_i$ for sufficiently large j, so that $\mathcal{C}^j := (C_1^j, \ldots, C_m^j, P_{m+1}, \ldots, P_{n-2})$ is supercritical for sufficiently large j. Using linearity of mixed area measures and applying Lemma 14.10 repeatedly, it follows that $S_{f,\mathcal{C}^j} = 0$ for sufficiently large j. Therefore, as $\mathcal{C}^j$ consists entirely of polytopes, Theorem 8.1 yields for every sufficiently large j a vector $s_j \in \mathbb{R}^n$ so that $f(x) = \langle s_j, x\rangle$ for all $x \in \operatorname{supp} S_{B,\mathcal{C}^j}$.

Now note that by linearity of mixed area measures, $\operatorname{supp} S_{B,\mathcal{C}^j}$ is increasing in j. In particular, we have $\langle s_j, x\rangle = f(x) = \langle s_{j+1}, x\rangle$ for all $x \in \operatorname{supp} S_{B,\mathcal{C}^j}$ and sufficiently large j. It follows from Lemma 13.2 and supercriticality that $s_j = s_{j+1}$ for all sufficiently large j. Thus we have shown that there exists $s \in \mathbb{R}^n$ such that

$$f(x) = \langle s, x\rangle \quad \text{for all } x \in \operatorname{supp} S_{B,\mathcal{C}^j}$$

holds for all sufficiently large j. To conclude, note that

$$\int |f(x) - \langle s, x\rangle|\, S_{B,\mathcal{C}}(dx) = \lim_{j\to\infty} \int |f(x) - \langle s, x\rangle|\, S_{B,\mathcal{C}^j}(dx) = 0$$

by Lemma 3.3, so that $f(x) = \langle s, x\rangle$ for all $x \in \operatorname{supp} S_{B,\mathcal{C}}$ as well. □

Remark 14.11. We exploited Lemma 14.10 above to approximate zonoids C_i by zonotopes C_i^j. However, one could also attempt to use Lemma 14.10 as a replacement for the local Alexandrov-Fenchel inequality: by Remark 8.6, it implies that for any extremal f of the Alexandrov-Fenchel inequality with reference bodies $(Z, \mathcal{K})$ and $u \in \operatorname{supp} \rho$, the projection $\mathsf{P}_{u^\perp} f$ is extremal in $u^\perp$ with reference bodies $\mathsf{P}_{u^\perp}\mathcal{K}$. Such an argument was used by Schneider in [27]. The difficulty with this approach is that $\mathsf{P}_{u^\perp}\mathcal{K}$ need not be supercritical for ρ-a.e. u if we only assume that $\mathcal{K}$ is supercritical. On the other hand, this method works well when all the bodies in $\mathcal{K}$ are full-dimensional, and yields some more general results in this case (for example, an analogue of Theorem 14.6 where some of the bodies C_i are zonoids).

We now revisit the problem of characterizing $\operatorname{supp} S_{B,\mathcal{C}}$ geometrically. When $\mathcal{C}$ are polytopes, such a characterization is given in Lemma 2.3 in terms of their faces. It has been conjectured by Schneider that $\operatorname{supp} S_{B,\mathcal{C}}$ is characterized in general by a local analogue of Lemma 2.3, in which the faces are replaced by certain "tangent spaces" of the convex bodies $\mathcal{C}$. Let us recall the relevant notions.

A convex body C in $\mathbb{R}^n$ associates to each of its boundary points a cone of outer normal vectors. These cones generate a partition of $\mathbb{R}^n$ into relatively open convex cones, which are the *touching cones* of C. We denote by $T(C, u)$ the unique touching cone of C that contains $u \in \mathbb{R}^n$. One may think of $T(C, u)^\perp$ as the "tangent space" of C with outer normal vector u. In analogy with Lemma 2.3, we define:

Definition 14.12. Let $\mathcal{C} = (C_1, \ldots, C_{n-1})$ be convex bodies in $\mathbb{R}^n$. Then $u \in S^{n-1}$ is a $\mathcal{C}$-*extreme normal direction* if there are segments $I_i \subset T(C_i, u)^\perp$ for $i \in [n-1]$ with linearly independent directions.

The above notions are due to Schneider (see [30, §2.2] for equivalent definitions). In particular, Schneider has conjectured [30, Conjecture 7.6.14] that $\operatorname{supp} S_{\mathcal{C}}$ always coincides with the closure of the set of $\mathcal{C}$-extreme normal directions. It is readily verified that this conjecture agrees with the special cases of Lemmas 2.3 and 14.5. We will presently verify this conjecture for the case $m = n - 2$ of Theorem 14.9. This is essentially proved in [27], up to a minor observation.

Proposition 14.13. *Let $\mathcal{C} = (C_1, \ldots, C_{n-2})$ be convex bodies in $\mathbb{R}^n$ such that each C_i is either a zonoid or a smooth body. Then*

$$\operatorname{supp} S_{B,\mathcal{C}} = \operatorname{cl}\{u \in S^{n-1} : u \text{ is a } (B, \mathcal{C})\text{-extreme normal direction}\}.$$

In the proof we will need the following simple measure-theoretic fact.

Lemma 14.14. *Let X, Y be Polish spaces, ρ_x be a finite measure on Y for each $x \in X$, and η be a finite measure on X. Assume $x \mapsto \rho_x$ is weakly continuous. Then the measure $\mu := \int \rho_x \, \eta(dx)$ on Y satisfies* $\operatorname{supp} \rho_x \subseteq \operatorname{supp} \mu$ *for every* $x \in \operatorname{supp} \eta$.

Proof. Let $x \in \operatorname{supp} \eta$ and $y \in \operatorname{supp} \rho_x$. Let $f : Y \to [0, 1]$ be a continuous function so that $f(y) > 0$. Then $\varepsilon := \int f \, d\rho_x > 0$. As $x \mapsto \int f \, d\rho_x$ is continuous, there is an open neighborhood $W \ni x$ so that $\int f \, d\rho_z \geqslant \frac{\varepsilon}{2}$ for $z \in W$. Thus $\int f \, d\mu \geqslant \frac{\varepsilon}{2} \eta(W) > 0$. As this holds for any function f as above, it follows that $y \in \operatorname{supp} \mu$. $\square$

We can now complete the proof of Proposition 14.13.

Proof of Proposition 14.13. When $\mathsf{V}_n(B, B, \mathcal{C}) = 0$, there exist no $(B, \mathcal{C})$-extreme directions by Lemma 2.2, and the conclusion is trivial. When $\mathsf{V}_n(B, B, \mathcal{C}) > 0$, Corollary 14.3 shows that $\operatorname{supp} S_{B,\mathcal{C}}$ is unchanged if each smooth C_i is replaced by B. On the other hand, for smooth C_i we have $T(C_i, u) = T(B, u) = \operatorname{pos} u$ for all u, so the $(B, \mathcal{C})$-extreme directions are also unchanged by this replacement. Thus we may assume that $C_1, \ldots, C_{n-2}$ are zonoids with generating measures $\rho_1, \ldots, \rho_{n-2}$.

That $\operatorname{supp} S_{B,\mathcal{C}} \subseteq \operatorname{cl}\{(B, \mathcal{C})\text{-extreme directions}\}$ is shown in [27, Proposition 3.8]. We will prove the converse inclusion by induction on n. For $n = 3$, the conclusion is a special case of Lemma 14.5. From now on, we assume that the claim has been proved in dimension $n - 1$, and show that the result follows in dimension n.

Let $v \in S^{n-1}$ be a $(B, \mathcal{C})$-extreme normal direction. Then it is shown in [27, p. 125] that there exists $j \in [n - 2]$ and $u \in v^\perp \cap \operatorname{supp} \rho_j$ so that v is also a $(\mathsf{P}_{u^\perp} B, \mathsf{P}_{u^\perp} \mathcal{C}_{\setminus j})$-extreme normal direction (as defined in $u^\perp$). It remains to show that the latter are included in $\operatorname{supp} S_{B,\mathcal{C}}$. To this end, note that

$$S_{B,\mathcal{C}} = \int S_{B,[-u,u],\mathcal{C}_{\setminus j}} \, \rho_j(du),$$

and that $u \mapsto S_{B,[-u,u],\mathcal{C}_{\setminus j}}$ is continuous by Lemma 3.3. Thus $\operatorname{supp} S_{B,[-u,u],\mathcal{C}_{\setminus j}} \subseteq \operatorname{supp} S_{B,\mathcal{C}}$ for every $u \in \operatorname{supp} \rho_j$ by Lemma 14.14. But the induction hypothesis and Remark 8.6 imply that $v \in \operatorname{supp} S_{\mathsf{P}_{u^\perp} B, \mathsf{P}_{u^\perp} \mathcal{C}_{\setminus j}} = \operatorname{supp} S_{B,[-u,u],\mathcal{C}_{\setminus j}}$. Thus we have shown that any $(B, \mathcal{C})$-extreme normal direction is contained in $\operatorname{supp} S_{B,\mathcal{C}}$, and the conclusion follows as the latter is a closed set. $\square$

15. Application to combinatorics of partially ordered sets

A sequence $N_1, \ldots, N_n$ of positive numbers is *log-concave* if $N_i^2 \geqslant N_{i-1} N_{i+1}$ for all $i \in \{2, \ldots, n - 1\}$. It was noticed long ago that log-concave sequences arise in a surprisingly broad range of combinatorial problems [35]. One explanation for this phenomenon appears in the work of Stanley [33], who observed that if one can represent the relevant combinatorial quantities in terms of mixed volumes, log-concavity arises as a consequence of the Alexandrov-Fenchel inequality. This provides a common mechanism for the emergence of log-concavity in several combinatorial problems that appear to be otherwise unrelated. In recent years, it has been realized that this idea extends to a much broader setting: even for combinatorial problems that may not be represented in terms of classical convexity, one may often develop algebraic analogues of the Alexandrov-Fenchel inequality that explain the emergence of log-concavity. Such ideas have led to a series of recent breakthroughs in combinatorics due to Huh et al. [17].

As was explained in the introduction, one may view the Alexandrov-Fenchel inequality as a generalized isoperimetric inequality. In particular, associated to any instance of the Alexandrov-Fenchel inequality is a corresponding extremal problem: what bodies minimize the left-hand side in Theorem 1.1 when the right-hand side is fixed? One may analogously associate to any log-concave sequence of combinatorial quantities (N_i) a corresponding extremal problem: what combinatorial objects achieve equality $N_i = N_{i+1}N_{i-1}$ for a given i? This question was already posed by Stanley in [33]. Despite deep advances in understanding log-concavity through Alexandrov-Fenchel type inequalities, the analysis of the associated extremal problems appears to be inaccessible by currently known methods.

In this section, we will show how the theory developed in this paper makes it possible to settle such extremal problems in Stanley's original setting. For sake of illustration, we focus on one particular example from [33] that arises in the combinatorics of partially ordered sets; other combinatorial applications of the Alexandrov-Fenchel inequality may be investigated analogously. Whether the theory of this paper has analogues outside convexity is an intriguing question (cf. section 16).

15.1. **Linear extensions and extremal posets.** Let $\alpha := \{x, y_1, \ldots, y_{n-1}\}$ be a partially ordered set (poset) that will be fixed throughout this section. We denote by N_i the number of order-preserving bijections $\sigma : \alpha \to [n]$ such that $\sigma(x) = i$; that is, N_i is the number of linear extensions of the partial order of α for which x has rank i. The following was conjectured by Chung, Fishburn, and Graham [6].

Theorem 15.1 (Stanley [33]). *The sequence $N_1, \ldots, N_n$ is log-concave.*

Proof. The poset α defines polytopes K, L in $\mathbb{R}^{n-1}$ by

$$K := \{t \in [0,1]^{n-1} : t_j \leqslant t_k \text{ if } y_j \leqslant y_k,\ t_j = 1 \text{ if } y_j > x\},$$
$$L := \{t \in [0,1]^{n-1} : t_j \leqslant t_k \text{ if } y_j \leqslant y_k,\ t_j = 0 \text{ if } y_j < x\}.$$

Let us denote

$$\mathcal{K}_l := (\underbrace{K, \ldots, K}_{l}), \qquad \mathcal{L}_m := (\underbrace{L, \ldots, L}_{m}).$$

Then it is shown in [33, Theorem 3.2] that

$$N_i = (n-1)!\, \mathsf{V}_{n-1}(\mathcal{K}_{i-1}, \mathcal{L}_{n-i}).$$

The conclusion is now immediate by the Alexandrov-Fenchel inequality. ∎

The extremal question associated to Theorem 15.1 is: given $i \in \{2, \ldots, n-1\}$, which posets α attain equality $N_i^2 = N_{i+1}N_{i-1}$? The proof of Theorem 15.1 reduces this question to a special case of Theorem 2.13. To obtain a result of combinatorial interest, however, one must deduce from the geometric conditions of Theorem 2.13 a combinatorial characterization of the corresponding poset α. We presently state the resulting theorem, whose proof will occupy the remainder of this section.

As in Theorem 2.13, we must distinguish between trivial and nontrivial extremals. When $N_i = 0$, log-concavity implies that we always have $N_i^2 = N_{i+1}N_{i-1}$ for trivial reasons. Let us first characterize when this happens. In the following, we denote by $\alpha_{\mathrm{R}z} := \{y \in \alpha : y\mathrm{R}z\}$ for any relation $\mathrm{R} \in \{<, \leqslant, >, \geqslant\}$ of α.

Lemma 15.2 (Trivial extremals). *For any $i \in [n]$, we have*

$$N_i = 0 \qquad \text{if and only if} \qquad |\alpha_{<x}| > i - 1 \quad \text{or} \quad |\alpha_{>x}| > n - i.$$

This simple lemma, which will be proved in section 15.2, is intuitively obvious: it states that no linear extension of α can give x rank i if there are i elements of α that are smaller than x (as these elements must have smaller rank than x), or analogously if there are $n - i + 1$ elements larger than x. This statement has an easy direct proof. In contrast, the characterization of the nontrivial extremals is not obvious. The following theorem is the main result of this section.

Theorem 15.3 (Nontrivial extremals)**.** *Let $i \in \{2, \ldots, n-1\}$ be such that $N_i > 0$. Then the following are equivalent:*

a. $N_i^2 = N_{i+1}N_{i-1}$.
b. $N_i = N_{i+1} = N_{i-1}$.
c. Every linear extension $\sigma : \alpha \to [n]$ with $\sigma(x) = i$ assigns ranks $i-1$ and $i+1$ to elements of α that are incomparable to x.
d. $|\alpha_{<y}| > i$ *for all* $y \in \alpha_{>x}$, *and* $|\alpha_{>y}| > n - i + 1$ *for all* $y \in \alpha_{<x}$.

The formulation of condition d of Theorem 15.3 was derived by the authors from the analysis of the associated Alexandrov-Fenchel inequality (section 15.3). Once the correct statement has been realized, however, it is straightforward to find a direct proof of the easy directions $d \Rightarrow c \Rightarrow b \Rightarrow a$ of Theorem 15.3.

Proof of Theorem 15.3, $d \Rightarrow c \Rightarrow b \Rightarrow a$. Fix $i \in \{2, \ldots, n-1\}$ such that $N_i > 0$, and assume that condition d holds. We first show that this implies condition c. Indeed, suppose to the contrary that $\sigma(y) = i - 1$ for some $y \in \alpha_{<x}$; then we have $|\alpha_{>y}| \leqslant |\{z \in \alpha : \sigma(z) > \sigma(y)\}| = n - i + 1$, contradicting condition d. We can analogously rule out that $\sigma(y) = i + 1$ for some $y \in \alpha_{>x}$.

We now show that condition c implies b. Denote by $\mathcal{N}_i$ the set of linear extensions $\sigma : \alpha \to [n]$ with $\sigma(x) = i$ (so that $N_i = |\mathcal{N}_i|$). We further denote by $\mathcal{N}_i^{\pm} := \{\Pi_{i,i\pm1}\sigma : \sigma \in \mathcal{N}_i\}$, where $\Pi_{i,j}$ denotes the permutation of $[n]$ that exchanges i, j. Condition c now implies $\mathcal{N}_i^{\pm} \subseteq \mathcal{N}_{i\pm1}$, so that $N_i = |\mathcal{N}_i^{\pm}| \leqslant |\mathcal{N}_{i\pm1}| = N_{i\pm1}$. Thus

$$N_i^2 \geqslant N_{i+1}N_{i-1} \geqslant N_iN_{i-1} \geqslant N_i^2,$$

where the first inequality follows from Theorem 15.1. As $N_i > 0$, condition b follows readily. This completes the proof, as the implication $b \Rightarrow a$ is trivial. □

Condition c of Theorem 15.3 identifies one particular combinatorial mechanism that gives rise to equality in Theorem 15.1. It is far from obvious, however, why this suffices to yield a complete solution to the extremal problem. The hard part of Theorem 15.3 is to show that this is the *only* mechanism that gives rise to equality, which will be proved in section 15.3 below.

Let us emphasize that the combinatorial characterization of the equality cases of Stanley's inequality has strong consequences beyond the solution of the extrmal problem itself: it provides detailed information on the shape of the log-concave sequences $N_1, \ldots, N_n$ that can arise in Theorem 15.1. For example, the implication $a \Rightarrow b$ of Theorem 15.3 shows that such sequences cannot contain any 3-term geometric progressions. As any positive log-concave sequence is unimodal, it follows from Lemma 15.2 and Theorem 15.3 that one can decompose $[n] = I_1 \cup \cdots \cup I_5$ into consecutive (possibly empty) intervals I_k so that the sequence $N_1, \ldots, N_n$ has the form that is illustrated in Figure 15.1. Furthermore, Lemma 15.2 and Theorem 15.3 enable us to compute the length of each interval explicitly for any given poset.

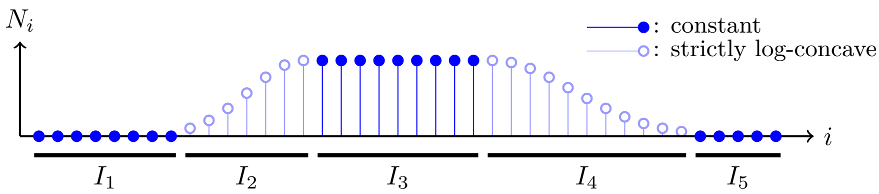

FIGURE 15.1. Structure of the sequence $N_1, \ldots, N_n$.

Example 15.4. For any $k, l, r, s, t \geqslant 1$, consider the poset

$$\alpha = \{x, (y_i)_{i\in[k]}, (z_i)_{i\in[l]}, (u_i)_{i\in[r-1]}, (v_i)_{i\in[s-1]}, (w_i)_{i\in[t+1]}\}$$

defined by the following relations:

$$y_1 < \cdots < y_k < x < z_1 < \cdots < z_l, \qquad y_k < w_1 < \cdots < w_{t+1} < z_1,$$
$$y_k < u_1 < \cdots < u_{r-1}, \qquad v_1 < \cdots < v_{s-1} < z_1.$$

Then $|\alpha_{<x}| = k$, $|\alpha_{>x}| = l$, $\min_{y\in\alpha_{<x}} |\alpha_{>y}| = l + r + t + 1$, and $\min_{y\in\alpha_{>x}} |\alpha_{<y}| = k + s + t + 1$. Therefore, Lemma 15.2 and Theorem 15.3 imply that the sequence $N_1, \ldots, N_{|\alpha|}$ has the form that is illustrated in Figure 15.1 with

$$|I_1| = k, \qquad |I_2| = s, \qquad |I_3| = t, \qquad |I_4| = r, \qquad |I_5| = l.$$

Thus any decomposition $[n] = I_1 \cup \cdots \cup I_5$ with $I_1, \ldots, I_5 \neq \varnothing$ is achievable. This example is readily modified to construct situations where some I_k may be empty.

15.2. **Order polytopes.** The convex bodies K, L that appear in the proof of Theorem 15.1 are examples of order polytopes. Before we proceed to the proof of Theorem 15.3, we must recall some basic properties of such polytopes.

We fix once and for all the poset $\alpha := \{x, y_1, \ldots, y_{n-1}\}$, and define $\bar{\alpha} := \alpha\backslash\{x\}$. For any $\beta \subseteq \alpha$ (with the induced partial order) and $z \in \alpha$, we will denote by $\beta_{\mathrm{R}z} := \{y \in \beta : y\mathrm{R}z\}$ for $\mathrm{R} \in \{<, \leqslant, >, \geqslant\}$, by $\beta_{\nsim z}$ the set of elements of β that are not comparable to z, and by $\beta^{\uparrow}$ ($\beta^{\downarrow}$) the set of maximal (minimal) elements of β.

For $y, z \in \beta$, we say that z *covers* y in β if $z \in (\beta_{>y})^{\downarrow}$. Moreover, $\beta \subseteq \alpha$ is called an *upper* (*lower*) set in α if $\alpha_{>y} \subseteq \beta$ ($\alpha_{<y} \subseteq \beta$) for every $y \in \beta$.

Now consider $\beta \subseteq \bar{\alpha}$. By a slight abuse of notation, we will define the subspace $\mathbb{R}^\beta := \{t \in \mathbb{R}^{n-1} : t_i = 0 \text{ for } y_i \notin \beta\}$. The *order polytope* O_β is defined as

$$O_\beta := \{t \in \mathbb{R}^\beta : t_j \in [0,1], \text{ and } t_j \leqslant t_k \text{ if } y_j \leqslant y_k, \text{ for all } y_j, y_k \in \beta\}.$$

The following basic facts may be found in [34, §1].

Lemma 15.5. *For any $\beta \subseteq \bar{\alpha}$, we have $\dim O_\beta = |\beta|$. The $(|\beta| - 1)$-dimensional faces of O_β are precisely the following subsets of O_β:*

1. $O_\beta \cap \{t_j = 0\}$ for $y_j \in \beta^{\downarrow}$.
2. $O_\beta \cap \{t_j = 1\}$ for $y_j \in \beta^{\uparrow}$.
3. $O_\beta \cap \{t_j = t_k\}$ for $y_j, y_k \in \beta$ such that y_k covers y_j in β.

In the sequel, we denote by $e_1, \ldots, e_{n-1}$ the coordinate basis of $\mathbb{R}^{n-1}$, and we let $1_\beta := \sum_{y_j\in\beta} e_j$ for any $\beta \subseteq \bar{\alpha}$. We can now formulate the following basic fact.

Lemma 15.6. *If $\beta \subseteq \bar{\alpha}$ is a lower set in $\bar{\alpha}$, then*

$$O_{\bar{\alpha}\backslash\beta} = \{t \in [0,1]^{n-1} : t_j \leqslant t_k \text{ if } y_j \leqslant y_k,\ t_j = 0 \text{ if } y_j \in \beta\}.$$

Analogously, if $\beta \subseteq \bar{\alpha}$ is an upper set in $\bar{\alpha}$, then

$$O_{\bar{\alpha}\backslash\beta} + 1_\beta = \{t \in [0,1]^{n-1} : t_j \leqslant t_k \text{ if } y_j \leqslant y_k,\ t_j = 1 \text{ if } y_j \in \beta\}.$$

Proof. Suppose β is a lower set and $t \in [0,1]^{n-1}$ satisfies $t_j = 0$ for all $y_j \in \beta$. Then $t_j \leqslant t_k$ holds trivially for $y_j \leqslant y_k$ with $y_j \in \beta$. On the other hand, as β is a lower set, $y_j \leqslant y_k$ with $y_j \notin \beta$ implies $y_k \notin \beta$. Thus the only nontrivial constraints $t_j \leqslant t_k$ for $y_j \leqslant y_k$ are those that appear in the definition of $O_{\bar{\alpha}\backslash\beta}$, which concludes the proof when β is a lower set. The proof when β is an upper set is analogous. □

In the remainder of this section, we will define the polytopes K, L as in the proof of Theorem 15.1. By Lemma 15.6, we can write equivalently

$$K = O_{\bar{\alpha}\backslash\alpha_{>x}} + 1_{\alpha_{>x}}, \qquad L = O_{\bar{\alpha}\backslash\alpha_{<x}},$$

where we used that $\alpha_{>x}$ is an upper set and $\alpha_{<x}$ is a lower set in $\bar{\alpha}$.

Lemma 15.7. $\dim K = n-1-|\alpha_{>x}|$, $\dim L = n-1-|\alpha_{<x}|$, *and* $\dim(K+L) = n-1$.

Proof. Lemma 15.5 shows that $K - 1_{\alpha_{>x}}$ is a full-dimensional polytope in $\mathbb{R}^{\bar{\alpha}\backslash\alpha_{>x}}$, and that L is a full-dimensional polytope in $\mathbb{R}^{\bar{\alpha}\backslash\alpha_{<x}}$. It follows that $\dim K = |\bar{\alpha}\backslash\alpha_{>x}|$, $\dim L = |\bar{\alpha}\backslash\alpha_{<x}|$, and $\dim(K+L) = |(\bar{\alpha}\backslash\alpha_{>x}) \cup (\bar{\alpha}\backslash\alpha_{<x})| = |\bar{\alpha}|$. □

Let us finally prove Lemma 15.2. While a direct combinatorial proof is a simple exercise, we find it instructive to show how it arises from the mixed volumes.

Proof of Lemma 15.2. The representation in the proof of Theorem 15.1 shows that $N_i = 0$ if and only if $\mathsf{V}_{n-1}(\mathcal{K}_{i-1}, \mathcal{L}_{n-i}) = 0$. By Lemma 2.2, this is the case if and only if either $\dim(K+L) < n-1$, $\dim(K) < i-1$, or $\dim(L) < n-i$. The conclusion now follows immediately from Lemma 15.7. □

15.3. **Combinatorial characterization of the extremals.** We now turn to the proof of Theorem 15.3. We have already proved the implications $d \Rightarrow c \Rightarrow b \Rightarrow a$; the remainder of this section is devoted to the proof of the implication $a \Rightarrow d$.

We begin by reducing the problem to an extremal case of the Alexandrov-Fenchel inequality in $\mathbb{R}^{n-1}$. While the polytopes $K, L \subset \mathbb{R}^{n-1}$ generally have empty interior, it turns out that nontrivial extremals can only arise in the present setting from a supercritical case of the Alexandrov-Fenchel inequality.

Lemma 15.8. *Let $i \in \{2, \ldots, n-1\}$ be such that $N_i > 0$ and $N_i^2 = N_{i+1}N_{i-1}$. Then $|\alpha_{<x}| + 1 < i < n - |\alpha_{>x}|$, and there exist $a > 0$ and $v \in \mathbb{R}^{n-1}$ so that*

$$h_K(x) = h_{aL+v}(x) \quad \text{for all} \quad x \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}.$$

Proof. We note first that $N_i > 0$ and $N_i^2 = N_{i+1}N_{i-1}$ imply that $N_{i-1} > 0$ and $N_{i+1} > 0$. It therefore follows from Lemma 15.2 that $|\alpha_{<x}| + 1 < i < n - |\alpha_{>x}|$. Moreover, this implies by Lemma 15.7 that $\dim K \geqslant i$ and $\dim L \geqslant n-i+1$, and $\dim(K+L) = n-1$. Thus the collection $(\mathcal{K}_{i-2}, \mathcal{L}_{n-i-1})$ is supercritical.

Now note that by the mixed volume representation in the proof of Theorem 15.1, $N_i > 0$ and $N_i^2 = N_{i+1}N_{i-1}$ imply $\mathsf{V}_{n-1}(K, L, \mathcal{K}_{i-2}, \mathcal{L}_{n-i-1}) > 0$ and

$$\mathsf{V}_{n-1}(K, L, \mathcal{K}_{i-2}, \mathcal{L}_{n-i-1})^2 = \mathsf{V}_{n-1}(K, K, \mathcal{K}_{i-2}, \mathcal{L}_{n-i-1})\,\mathsf{V}_{n-1}(L, L, \mathcal{K}_{i-2}, \mathcal{L}_{n-i-1}).$$

The conclusion therefore follows from Corollary 2.16. □

To exploit Lemma 15.8, there are two distinct difficulties: we must gain some understanding of which vectors lie in $\operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$, and we must understand how to exploit the fact that the supporting hyperplanes of K and $aL+v$ coincide in these directions. We begin by addressing the first issue.

Lemma 15.9. *Let $i \in \{2,\ldots,n-1\}$ be such that $N_i > 0$. Then the following hold:*

a. $-e_j \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ *for* $y_j \in \bar{\alpha}^{\downarrow}$.

b. $-e_j \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ *for* $y_j \in (\alpha_{>x})^{\downarrow}$ *such that* $|\alpha_{<y_j}| \leqslant i$.

c. $e_j \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ *for* $y_j \in \bar{\alpha}^{\uparrow}$.

d. $e_j \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ *for* $y_j \in (\alpha_{<x})^{\uparrow}$ *such that* $|\alpha_{>y_j}| \leqslant n-i+1$.

e. $e_{jk} := \frac{e_j-e_k}{\sqrt{2}} \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ *when the following conditions are all satisfied:*

(i) y_k *covers* y_j *in* $\bar{\alpha}$;

(ii) *if* $y_j \in \alpha_{<x}$, $y_k \in \bar{\alpha}\backslash\alpha_{<x}$ *then* $y_k \in (\bar{\alpha}\backslash\alpha_{<x})^{\downarrow}$;

(iii) *if* $y_k \in \alpha_{>x}$, $y_j \in \bar{\alpha}\backslash\alpha_{>x}$, *then* $y_j \in (\bar{\alpha}\backslash\alpha_{>x})^{\uparrow}$.

Proof. As $N_i > 0$, Lemma 15.2 implies $|\alpha_{<x}|+1 \leqslant i \leqslant n-|\alpha_{>x}|$. Moreover, by Lemma 2.3, we have $u \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ if and only if

$$\dim F(K,u) \geqslant i-2, \quad \dim F(L,u) \geqslant n-i-1, \quad \dim F(K+L,u) \geqslant n-3.$$

The latter condition will be verified in each part of the lemma.

Part *a*. We begin by noting that for any $j \in [n-1]$,

$$F(K,-e_j) = K \cap \{t_j = 1_{y_j\in\alpha_{>x}}\}, \qquad F(L,-e_j) = L \cap \{t_j = 0\}.$$

Indeed, it is readily seen by the definitions of K, L that $h_K(-e_j) = -\inf_{t\in K} t_j = -1_{y_j\in\alpha_{>x}}$ and $h_L(-e_j) = -\inf_{t\in L} t_j = 0$, so the claim follows from (2.2).

Now let $y_j \in \bar{\alpha}^{\downarrow}$. We claim that the following hold:

1. $F(K,-e_j)$ is a full-dimensional polytope in $\mathbb{R}^{\bar{\alpha}\backslash(\alpha_{>x}\cup\{y_j\})} + 1_{\alpha_{>x}}$; and
2. $F(L,-e_j)$ is a full-dimensional polytope in $\mathbb{R}^{\bar{\alpha}\backslash(\alpha_{<x}\cup\{y_j\})}$.

Indeed, recall first that $K = O_{\bar{\alpha}\backslash\alpha_{>x}} + 1_{\alpha_{>x}}$ by Lemma 15.6. If $y_j \in \alpha_{>x}$, then $F(K,-e_j) = K$ and the first claim follows directly from Lemma 15.5. On the other hand, if $y_j \in \bar{\alpha}\backslash\alpha_{>x}$, then $F(K,-e_j) = K \cap \{t_j = 0\} = O_{\bar{\alpha}\backslash\alpha_{>x}} \cap \{t_j = 0\} + 1_{\alpha_{>x}}$. But then $y_j \in \bar{\alpha}^{\downarrow}$ implies $y_j \in (\bar{\alpha}\backslash\alpha_{>x})^{\downarrow}$, and the first claim follows again from Lemma 15.5. The proof of the second claim is completely analogous.

To conclude the proof of part a, it suffices to note that the above claims imply

$$\begin{aligned}\dim F(K,-e_j) &= |\bar{\alpha}\backslash(\alpha_{>x}\cup\{y_j\})| \geqslant n-2-|\alpha_{>x}| \geqslant i-2,\\ \dim F(L,-e_j) &= |\bar{\alpha}\backslash(\alpha_{<x}\cup\{y_j\})| \geqslant n-2-|\alpha_{<x}| \geqslant n-i-1,\\ \dim F(K+L,-e_j) &= |\bar{\alpha}\backslash(\alpha_{>x}\cup\{y_j\})\cup\bar{\alpha}\backslash(\alpha_{<x}\cup\{y_j\})| = |\bar{\alpha}\backslash\{y_j\}| = n-2,\end{aligned}$$

where we used $|\alpha_{<x}|+1 \leqslant i \leqslant n-|\alpha_{>x}|$ and $\alpha_{>x}\cap\alpha_{<x} = \varnothing$.

Part *b*. Let $y_j \in (\alpha_{>x})^{\downarrow}$ with $|\alpha_{<y_j}| \leqslant i$. We already showed in part a that

$$F(K,-e_j) = K, \qquad F(L,-e_j) = L \cap \{t_j = 0\}$$

and that $F(K,-e_j)$ is a full-dimensional polytope in $\mathbb{R}^{\bar{\alpha}\backslash\alpha_{>x}} + 1_{\alpha_{>x}}$ (the proofs of these facts for $y_j \in \alpha_{>x}$ did not use the assumption of part a).

Now note that as $y_j \in \alpha_{>x}$, we have $y_i \leqslant y_j$ for all $y_i \in \alpha_{<x}$. It follows that

$$L \cap \{t_j = 0\} = \{t \in [0,1]^{n-1} : t_i \leqslant t_k \text{ if } y_i \leqslant y_k,\ t_i = 0 \text{ if } y_i \leqslant y_j\} = O_{\bar{\alpha}\backslash\bar{\alpha}_{\leqslant y_j}},$$

where we used Lemma 15.6 and that $\bar{\alpha}_{\leqslant y_j}$ is a lower set in $\bar{\alpha}$. It therefore follows from Lemma 15.5 that $F(L, -e_j)$ is a full-dimensional polytope in $\mathbb{R}^{\bar{\alpha}\backslash\bar{\alpha}_{\leqslant y_j}}$.

To conclude the proof of part b, note that as $i \leqslant n - |\alpha_{>x}|$, we have

$$\dim F(K, -e_j) = |\bar{\alpha}\backslash\alpha_{>x}| = n - 1 - |\alpha_{>x}| \geqslant i - 1.$$

On the other hand, as $y_j \in \alpha_{>x}$, we have $|\bar{\alpha}_{\leqslant y_j}| = |\alpha_{<y_j}| \leqslant i$, so that

$$\dim F(L, -e_j) = |\bar{\alpha}\backslash\bar{\alpha}_{\leqslant y_j}| = n - 1 - |\bar{\alpha}_{\leqslant y_j}| \geqslant n - i - 1.$$

Finally, we note that

$$\dim F(K + L, -e_j) = |\bar{\alpha}\backslash\alpha_{>x} \cup \bar{\alpha}\backslash\bar{\alpha}_{\leqslant y_j}| = |\bar{\alpha}\backslash\{y_j\}| = n - 2,$$

where we used that $\alpha_{>x} \cap \bar{\alpha}_{\leqslant y_j} = \{y_j\}$ as $y_j \in (\alpha_{>x})^{\downarrow}$.

Parts *c* and *d*. The proofs are completely analogous to those of parts a and b.

Part *e*. We begin by noting that

$$F(K, e_{jk}) = K \cap \{t_j = t_k\}, \qquad F(L, e_{jk}) = L \cap \{t_j = t_k\}$$

whenever $y_j < y_k$. Indeed, as $y_j < y_k$ implies $t_j \leqslant t_k$ for any $t \in K$, it follows readily from the definition of K that $h_K(e_{jk}) = 2^{-1/2} \sup_{t\in K}(t_j - t_k) = 0$. That $h_L(e_{jk}) = 0$ follows analogously, and the claim now follows from (2.2).

Now suppose y_j, y_k satisfy conditions (i)–(iii) of part e. We claim the following.

1. $F(K, e_{jk}) \subset \mathbb{R}^{\bar{\alpha}\backslash\alpha_{>x}} + 1_{\alpha_{>x}}$ with $\dim F(K, e_{jk}) \geqslant |\bar{\alpha}\backslash\alpha_{>x}| - 1$.
2. $F(L, e_{jk}) \subset \mathbb{R}^{\bar{\alpha}\backslash\alpha_{<x}}$ with $\dim F(L, e_{jk}) \geqslant |\bar{\alpha}\backslash\alpha_{<x}| - 1$.

Indeed, as $K = O_{\bar{\alpha}\backslash\alpha_{>x}} + 1_{\alpha_{>x}}$ by Lemma 15.6, the first part of the first claim is immediate. For the second part of the first claim, we consider three cases.

- If $y_j, y_k \in \alpha_{>x}$, then $F(K, e_{jk}) = K$, so $\dim F(K, e_{jk}) = |\bar{\alpha}\backslash\alpha_{>x}|$ by Lemma 15.5.
- If $y_j \in \bar{\alpha}\backslash\alpha_{>x}$ and $y_k \in \alpha_{>x}$, then $F(K, e_{jk}) = O_{\bar{\alpha}\backslash\alpha_{>x}} \cap \{t_j = 1\} + 1_{\alpha_{>x}}$, and (iii) states that $y_j \in (\bar{\alpha}\backslash\alpha_{>x})^{\uparrow}$. Thus $\dim F(K, e_{jk}) = |\bar{\alpha}\backslash\alpha_{>x}| - 1$ by Lemma 15.5.
- If $y_j, y_k \in \bar{\alpha}\backslash\alpha_{>x}$, then $F(K, e_{jk}) = O_{\bar{\alpha}\backslash\alpha_{>x}} \cap \{t_j = t_k\} + 1_{\alpha_{>x}}$, and (i) implies that y_k covers y_j in $\bar{\alpha}\backslash\alpha_{>x}$. Thus $\dim F(K, e_{jk}) = |\bar{\alpha}\backslash\alpha_{>x}| - 1$ by Lemma 15.5.

This proves the first claim. The proof of the second claim is completely analogous (using condition (ii) rather than condition (iii)).

To conclude the proof of part e, note that the above claims imply

$$\dim F(K, e_{jk}) \geqslant n - 2 - |\alpha_{>x}| \geqslant i - 2,$$
$$\dim F(L, e_{jk}) \geqslant n - 2 - |\alpha_{<x}| \geqslant n - i - 1$$

as $|\alpha_{<x}| + 1 \leqslant i \leqslant n - |\alpha_{>x}|$. On the other hand, we have

$$\begin{aligned}\dim F(K + L, e_{jk}) &\geqslant \dim F(K, e_{jk}) + \dim F(L, e_{jk}) - |\bar{\alpha}\backslash\alpha_{>x} \cap \bar{\alpha}\backslash\alpha_{<x}| \\ &\geqslant |\bar{\alpha}\backslash\alpha_{>x} \cup \bar{\alpha}\backslash\alpha_{<x}| - 2 = n - 3,\end{aligned}$$

where we used that $\alpha_{>x} \cap \alpha_{<x} = \varnothing$. The proof is complete. □

From this point onwards we place ourselves in the setting of Lemma 15.8. In particular, we will assume without further comment that $i \in \{2, \ldots, n-1\}$ with $N_i > 0$ and $N_i^2 = N_{i+1}N_{i-1}$, and that $a > 0$, $v \in \mathbb{R}^{n-1}$ have been fixed so that

$$h_K(x) = h_{aL+v}(x) \quad \text{for all} \quad x \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}.$$

Let us begin by formulating a first consequence of Lemma 15.9.

Lemma 15.10. *The following hold.*

a. $v_j = 0$ *for* $y_j \in \bar{\alpha}^{\downarrow}\backslash\alpha_{>x}$.
b. $v_j = 1 - a$ *for* $y_j \in \bar{\alpha}^{\uparrow}\backslash\alpha_{<x}$.
c. $v_j = v_k$ *whenever conditions* (i)–(iii) *of Lemma* 15.9 *are all satisfied.*

Proof. For part a, note that $-e_j \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ by Lemma 15.9, so that $h_K(-e_j) = ah_L(-e_j) - v_j$. But as $y_j \notin \alpha_{>x}$, we have $h_K(-e_j) = h_L(-e_j) = 0$ as in the proof of part a of Lemma 15.9, so the conclusion follows.

The argument for part b is analogous: we have $e_j \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ by Lemma 15.9, so that $h_K(e_j) = ah_L(e_j) + v_j$. But as $y_j \notin \alpha_{<x}$, it follows readily that $h_K(e_j) = h_L(e_j) = 1$, and the conclusion follows.

Finally, for part c, we have $e_{jk} \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ by Lemma 15.9, so that $h_K(e_{jk}) = ah_L(e_{jk}) + 2^{-1/2}(v_j - v_k)$. But condition (i) implies $h_K(e_{jk}) = h_L(e_{jk}) = 0$ as in the proof of part e of Lemma 15.9, so the conclusion follows. □

We can now use Lemma 15.10 as a basic step to compute a and v.

Corollary 15.11. $a = 1$.

Proof. We first note that $\alpha_{\nsim x} \neq \varnothing$. Indeed, if every element of α were comparable to x, then $|\alpha_{<x}| + |\alpha_{>x}| = n - 1$, which implies by Lemma 15.2 that $N_i = 0$ whenever $i \neq |\alpha_{<x}| + 1$. The latter contradicts $N_i^2 = N_{i+1}N_{i-1} > 0$.

Fix any $y_{j_0} \in \alpha_{\nsim x}$. We construct a chain $y_{j_{-s}} < \cdots < y_{j_i} < y_{j_{i+1}} < \cdots < y_{j_t}$ according to the following algorithm. For the upper part of the chain, we iteratively choose $y_{j_{i+1}} \in (\alpha_{>y_{j_i}})^{\downarrow}$ for $i \geqslant 0$ under the constraint that we select $y_{j_{i+1}} \in \alpha_{\nsim x}$ whenever possible. The chain is extended until a maximal element of α is reached. The lower part of the chain is constructed by iteratively choosing $y_{j_{i-1}} \in (\alpha_{<y_{j_i}})^{\uparrow}$ for $i \leqslant 0$ under the constraint that we select $y_{j_{i-1}} \in \alpha_{\nsim x}$ whenever possible. The chain is extended until a minimal element of α is reached.

We claim that this construction ensures the following properties:

1. $y_{j_{-s}} \in \bar{\alpha}^{\downarrow}\backslash\alpha_{>x}$ and $y_{j_t} \in \bar{\alpha}^{\uparrow}\backslash\alpha_{<x}$.
2. Conditions (i)–(iii) of Lemma 15.9 hold with $j = j_i$, $k = j_{i+1}$ for all i.

Indeed, it is impossible that $y_{j_i} \in \alpha_{\leqslant x}$ for some $i \geqslant 0$, or that $y_{j_i} \in \alpha_{\geqslant x}$ for some $i \leqslant 0$, as that would violate $y_{j_0} \in \alpha_{\nsim x}$. Thus the first claim follows as $y_{j_{-s}}$ is minimal and y_{j_t} is maximal by construction. To prove the second claim, consider first $i \geqslant 0$, so that $y_{j_i}, y_{j_{i+1}} \notin \alpha_{\leqslant x}$. Then (i) holds as $y_{j_{i+1}}$ covers y_{j_i} in $\bar{\alpha}$ by construction, and (ii) holds automatically. Finally, as we chose $y_{j_{i+1}} \in \alpha_{\nsim x}$ whenever possible, we can only have $y_{j_i} \in \bar{\alpha}\backslash\alpha_{>x}$ and $y_{j_{i+1}} \in \alpha_{>x}$ if $y_{j_i} \in (\bar{\alpha}\backslash\alpha_{>x})^{\uparrow}$, which establishes condition (iii). The proof of the second claim for $i \leqslant 0$ is completely analogous.

To conclude the proof, we observe that the above claims and Lemma 15.10 imply that $v_{j_{-s}} = 0$, $v_{j_t} = 1 - a$, and $v_{j_i} = v_{j_{i+1}}$ for all i. Thus $a = 1$. □

Corollary 15.12. $v_j = 0$ *for all* $y_j \in \alpha_{<x} \cup \alpha_{>x}$.

Proof. Fix any $y_{j_0} \in \alpha_{>x}$, and construct an increasing chain $y_{j_0} < \cdots < y_{j_t}$ by iteratively choosing $y_{j_{i+1}} \in (\alpha_{>y_{j_i}})^{\downarrow}$ until a maximal element is reached. Then $y_{j_i} \in \alpha_{>x}$ for all i, so $y_{j_t} \in \bar{\alpha}^{\uparrow}\backslash\alpha_{<x}$ and conditions (i)—(iii) of Lemma 15.9 hold with $j = j_i$, $k = j_{i+1}$ for all i. Thus Lemma 15.10 and Corollary 15.11 imply that $v_{j_t} = 0$ and $v_{j_i} = v_{j_{i+1}}$ for all i. We have therefore shown that $v_j = 0$ for any $y_j \in \alpha_{>x}$. The proof for the case $y_j \in \alpha_{<x}$ follows in a completely analogous manner by constructing a decreasing chain from any $y_{j_0} \in \alpha_{<x}$. □

We are now ready to complete the proof of Theorem 15.3.

Proof of Theorem 15.3, $a \Rightarrow d$. Fix $i \in \{2,\ldots,n-1\}$ such that $N_i > 0$ and $N_i^2 = N_{i+1}N_{i-1}$. Then Lemma 15.8 and Corollaries 15.11 and 15.12 imply that

$$h_K(x) = h_L(x) + \langle v, x\rangle \quad \text{for all} \quad x \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$$

holds for a vector $v \in \mathbb{R}^{n-1}$ with $v_j = 0$ for all $y_j \in \alpha_{<x} \cup \alpha_{>x}$.

Now consider any $y_j \in \alpha_{>x}$. Then we claim that $-e_j \notin \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$. Indeed, if $-e_j$ did lie in the support, then we would have $h_K(-e_j) = h_L(-e_j)$. But this entails a contradiction, as we showed in the proof of part a of Lemma 15.9 that $h_K(-e_j) = -1$ and $h_L(-e_j) = 0$. It follows by a completely analogous argument that $e_j \notin \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ when $y_j \in \alpha_{<x}$.

On the other hand, Lemma 15.9 implies that $-e_j \in \operatorname{supp} S_{B,\mathcal{K}_{i-2},\mathcal{L}_{n-i-1}}$ when $y_j \in (\alpha_{>x})^{\downarrow}$ such that $|\alpha_{<y_j}| \leqslant i$. Consequently, we have shown that y_j satisfying the latter condition cannot exist, that is, $|\alpha_{<y}| > i$ whenever $y \in (\alpha_{>x})^{\downarrow}$. As $\alpha_{<y}$ can only increase if we increase y, it follows that $|\alpha_{<y}| > i$ for all $y \in \alpha_{>x}$. By applying the same reasoning to e_j for $y_j \in (\alpha_{<x})^{\uparrow}$, it follows in a completely analogous manner that $|\alpha_{>y}| > n-i+1$ for all $y \in \alpha_{<x}$. $\square$

16. Discussion and open questions

The main results of this paper completely settle the extremals of the Alexandrov-Fenchel inequality for convex polytopes. Analogous extremal problems also arise, however, in other situations where Alexandrov-Fenchel type inequalities appear. The aim of the final section of this paper is to briefly discuss a number of basic open questions in this direction that arise from our results.

16.1. **General convex bodies.** The Alexandrov-Fenchel inequality (Theorem 1.1) applies to any reference bodies $C_1,\ldots,C_{n-2}$. While the main results of this paper require that the reference bodies are polytopes, the statements of our main results (Theorems 2.13 and 13.1) make sense for general convex bodies. One may therefore conjecture that the statements of our main results extend *verbatim* to the general setting. This conjecture is due to Schneider [26] for full-dimensional bodies, to which our results add detailed predictions on the lower-dimensional cases.

While we have made essential use of the polytope assumption in this paper, it should be emphasized that many of our arguments are already completely general: neither the gluing arguments nor the statement of the local Alexandrov-Fenchel inequality (Theorem 4.3) rely fundamentally on the polytope assumption, a fact that we already exploited in section 14. The main obstacle to extending the theory of this paper to general convex bodies therefore lies in the proof of Theorem 4.3: if such a result could be proved in the general setting, this would essentially complete the extremal characterization for general convex bodies.

16.1.1. **The local Alexandrov-Fenchel inequality.** Let us now briefly recall how the polytope structure was used in the proof of Theorem 4.3.

In sections 5–6, we used the combinatorial structure of polytopes to reduce Theorem 4.3 to a finite-dimensional problem. However, the actual proof of the local Alexandrov-Fenchel inequality in section 7 does not make direct use of the geometry of polytopes. What is really exploited here is that the mixed area measures $S_{g,\mathcal{P}}$ and $S_{g,g,\mathcal{P}_{\backslash r}}$ are supported in a finite set $\{u_i\}_{i\in[N]}$, which has two key consequences: the existence problem of Theorem 4.3 can be formulated as the solution of a system of linear equations; and the masses $S_{g,\mathcal{P}}(\{u_i\})$ satisfy Alexandrov-Fenchel inequalities $S_{g,\mathcal{P}}(\{u_i\})^2 \geqslant S_{g,g,\mathcal{P}_{\setminus r}}(\{u_i\})\, S_{P_r,P_r,\mathcal{P}_{\setminus r}}(\{u_i\})$ by Lemma 3.4.

In principle, however, one may conjecture that similar objects could be defined directly for general convex bodies in suitable functional-analytic framework. For example, it was shown in [32] that for any convex bodies $\mathcal{C} = (C_1, \ldots, C_{n-2})$ in $\mathbb{R}^n$, there exists a self-adjoint operator $\mathscr{A}_\mathcal{C}$ on the Hilbert space $L^2(S_{B,\mathcal{C}})$ so that

$$\mathsf{V}_n(K, L, \mathcal{C}) = \langle h_K, \mathscr{A}_\mathcal{C} h_L\rangle_{L^2(S_{B,\mathcal{C}})}$$

for any convex bodies K, L in $\mathbb{R}^n$. The operator $\mathscr{A}_\mathcal{C}$ may be viewed as an infinite-dimensional analogue of the Alexandrov matrix A (in the sense of Corollary 6.7), and one might therefore attempt to use such objects as a replacement for the finite-dimensional computations in the proof of Theorem 4.3. The problem with this construction, however, is that not only $\mathscr{A}_\mathcal{C}$ but also the underlying space $L^2(S_{B,\mathcal{C}})$ depends on $\mathcal{C}$, while the proof of Theorem 4.3 requires us to consider several such objects simultaneously. For example, an analogue "$(\mathscr{A}_\mathcal{C} f)^2 \geqslant \mathscr{A}_{f,\mathcal{C}_{\setminus r}} f \cdot \mathscr{A}_\mathcal{C} h_{C_r}$" of the Alexandrov-Fenchel inequality for $S_{g,\mathcal{P}}(\{u_i\})$ does not make sense in this form, as the operators $\mathscr{A}_\mathcal{C}$ and $\mathscr{A}_{f,\mathcal{C}_{\setminus r}}$ are defined on different spaces.

A basic question in this context is therefore whether one may construct analogues of the self-adjoint operators $\mathscr{A}_\mathcal{C}$ for different choices of $\mathcal{C}$ on the same space, and whether these operators satisfy an appropriate analogue of the Alexandrov-Fenchel inequality. Such analytic questions on the structure of mixed volumes are in principle completely independent from the study of the extremals. The development of such a functional-analytic framework could, however, provide a foundation for the extension of the proof of Theorem 4.3 to the general setting.

16.1.2. **Supports of mixed area measures.** The main results of this paper characterize the extremal functions f such that $S_{f,\mathcal{C}} = 0$ on the support of $S_{B,\mathcal{C}}$. To obtain a fully geometric interpretation of these results, however, they must be combined with a geometric characterization of $\operatorname{supp} S_{B,\mathcal{C}}$. The latter is elementary in the setting of polytopes (Lemma 2.3), but remains open in general.

The fundamental conjecture in this direction, due to Schneider [26], is that a local analogue of Lemma 2.3 (as formulated after Definition 14.12) remains valid in the general setting. Let us emphasize, however, that the characterization of $\operatorname{supp} S_{B,\mathcal{C}}$ played essentially no role in the proofs of our main results: this problem appears to be essentially orthogonal to the theory developed in this paper.

16.2. **Algebraic analogues.** Beyond its fundamental role in convex geometry, the Alexandrov-Fenchel inequality has deep connections with other areas of mathematics. It was realized in the 1970s by Teissier and Khovanskii (cf. [5, 13, 15]) that the Alexandrov-Fenchel inequality has natural analogues in algebraic and complex geometry. More recently, it has been realized that such connections extend even further: various combinatorial problems, which cannot be expressed directly in terms of convex or algebraic geometry, nonetheless fit within a general algebraic framework in which analogues of the Alexandrov-Fenchel inequality hold [17].

The rich algebraic theory surrounding the Alexandrov-Fenchel inequality raises the intriguing question whether our results might extend to a broader context. This question arises, for example, if we aim to develop combinatorial applications as in section 15 in situations that cannot be formulated in convex geometric terms. Related questions in algebraic geometry date back to Teissier [37, 4].

It may not be entirely obvious, however, how to even formulate algebraic analogues of the main results of this paper. The aim of this section is to sketch how our main results may be expressed in algebraic terms, which could (conjecturally) carry over to analogues of the Alexandrov-Fenchel inequality outside convexity. To the best of our knowledge, such problems are at present almost entirely open.

16.2.1. **Polytope algebra.** In order to describe our main results algebraically, we must first recall some aspects of the polytope algebra due to McMullen [21, 39].

Let us fix as in section 5 a simple polytope P in $\mathbb{R}^n$. Define the linear space

$$\mathrm{D} := \{h_Q - h_R : Q, R \text{ strongly isomorphic to } P\}.$$

For $f, g \in \mathrm{D}$, we will write $f \sim g$ if $f - g$ is a linear function, and denote by $[f]$ the equivalence class of f with respect to this equivalence relation.

The polytope algebra generated by P is a graded algebra

$$\mathrm{A}(P) = \bigoplus_{k=0}^{n} \mathrm{A}^k$$

with $\mathrm{A}^1 \simeq \mathrm{D}/{\sim}$ and $\mathrm{A}^n \simeq \mathbb{R}$. Moreover, the (commutative) multiplication of $\mathrm{A}(P)$ has the property that $F \cdot G \in \mathrm{A}^{k+l}$ for $F \in \mathrm{A}^k$, $G \in \mathrm{A}^l$, and

$$[f_1] \cdot \ldots \cdot [f_n] = n!\, \mathsf{V}_n(f_1, \ldots, f_n)$$

for any $f_1, \ldots, f_n \in \mathrm{D}$. By (2.1), one may therefore view $[f_1] \cdot \ldots \cdot [f_{n-1}] \in \mathrm{A}^{n-1}$ as the algebraic formulation of the mixed area measure $(n-1)!\, S_{f_1,\ldots,f_{n-1}}$.

Several different notions of positivity are defined by convex cones in A^1:

- $\mathrm{Amp} := \{[h_Q] : Q \text{ is strongly isomorphic to } P\}$ (the ample cone).
- $\mathrm{Nef} := \mathrm{cl}(\mathrm{Amp}) = \{[h_Q] : Q \text{ is homothetic to a summand of } P\}$ (the nef cone).
- $\mathrm{Big} := \{[f] : f \in \mathrm{D},\ f > 0\}$ (the big cone).
- $\overline{\mathrm{Eff}} := \mathrm{cl}(\mathrm{Big}) = \{[f] : f \in \mathrm{D},\ f \geqslant 0\}$ (the pseudoeffective cone).

The terminology used here is borrowed from algebraic geometry [19]. In these terms, two classical algebraic properties admit a familiar interpretation in convexity [17]. The *Hodge-Riemann relation* of degree one states that

$$(\eta \cdot L_0 \cdot \ldots \cdot L_{n-2})^2 \geqslant (\eta \cdot \eta \cdot L_1 \cdot \ldots \cdot L_{n-2})\,(L_0 \cdot L_0 \cdot \ldots \cdot L_{n-2})$$

for any $\eta \in \mathrm{A}^1$ and $L_0, \ldots, L_{n-2} \in \mathrm{Amp}$. This is nothing other than the Alexandrov-Fenchel inequality for polytopes strongly isomorphic to P. On the other hand, the *hard Lefschetz theorem* of degree one states that

$$\eta \cdot L_1 \cdot \ldots \cdot L_{n-2} = 0 \quad \text{if and only if} \quad \eta = 0$$

for any $\eta \in \mathrm{A}^1$ and $L_0, \ldots, L_{n-2} \in \mathrm{Amp}$. This statement is equivalent to the fact (which was proved in the original work of Alexandrov [1]) that $S_{f,\mathcal{P}} = 0$ if and only if f is a linear function when $f = h_K - h_L$ and $K, L, P_1, \ldots, P_{n-2}$ are strongly isomorphic polytopes; in other words, it states that the Alexandrov-Fenchel inequality has no nontrivial extremals in the strongly isomorphic setting.

With this algebraic language in hand, one may describe various notions of convexity in algebraic terms. To give one further example, if we associate to any $\eta \in \mathrm{Nef}$ a numerical dimension $\dim \eta := \max\{k : \eta^{\cdot k} \neq 0\}$ (cf. [20]), then it follows from Lemma 2.2 that $[h_Q] \in \mathrm{Nef}$ satisfies $\dim[h_Q] = \dim Q$. One can therefore readily describe notions such as critical sets, supercriticality, etc. algebraically.

16.2.2. **Extremals.** While the Hodge-Riemann inequality extends readily to the setting that $L_1, \ldots, L_{n-2} \in \mathrm{Nef}$ by continuity, this is not the case for the hard Lefschetz theorem. Indeed, the main results of this paper are concerned with the case that $\mathcal{P} = (P_1, \ldots, P_{n-2})$ are summands of P (see section 5), and it is precisely in this case that nontrivial extremals appear. Our results may therefore be viewed as refined forms of the hard Lefschetz theorem for nef classes.

To express our results algebraically, we must understand the algebraic meaning of the condition $f(x) = 0$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$. To this end, let us make the following observation (we freely use the notation of sections 5–6 in the proof).

Lemma 16.1. *$f \in \mathrm{D}$ satisfies $f(x) = 0$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$ if and only if there exist $f_1, f_2 \in \mathrm{D}$, $f_1, f_2 \geqslant 0$ so that $f = f_1 - f_2$ and $S_{f_1,\mathcal{P}} = S_{f_2,\mathcal{P}} = 0$.*

Proof. For the *if* direction, note that $S_{f_i,\mathcal{P}} = 0$ implies $\int f_i \, dS_{B,\mathcal{P}} = \int h_B \, dS_{f_i,\mathcal{P}} = 0$. As $f_i \geqslant 0$, it follows that $f_i(x) = 0$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$ and $i = 1, 2$.

For the *only if* direction, suppose $f(x) = 0$ for all $x \in \operatorname{supp} S_{B,\mathcal{P}}$. Applying Lemma 6.4 to the vectors $z_1 = (f(u_i)_+)_{i\in[N]}$ and $z_2 = (f(u_i)_-)_{i\in[N]}$, we obtain $f_1, f_2 \in \mathrm{D}$ so that $f_1(u_i), f_2(u_i) \geqslant 0$ and $f(u_i) = f_1(u_i) - f_2(u_i)$ for all $i \in [N]$, and $f_1(u_i) = f_2(u_i) = 0$ for $i \in V$. As any function in D is linear on the normal cones of P, it follows that $f_1, f_2 \geqslant 0$ and $f = f_1 - f_2$ everywhere, and that $f_1 = f_2 = 0$ on $\operatorname{supp} S_{B,\mathcal{P}}$ (by Lemma 5.5). The conclusion follows from Lemma 2.8. ☐

For sake of illustration, let us consider the simplest setting where $L_1, \ldots, L_{n-2}$ are big and nef, that is, $L_i = [h_{P_i}]$ for full-dimensional polytopes $P_1, \ldots, P_{n-2}$. The conclusion of Theorem 8.1 may then be reformulated as follows.

Corollary 16.2. *Let $L_1, \ldots, L_{n-2} \in \mathrm{Nef} \cap \mathrm{Big}$ and $\eta \in \mathrm{A}^1$. Then $\eta \cdot L_1 \cdot \ldots \cdot L_{n-2} = 0$ if and only if $\eta = \eta_1 - \eta_2$ for some $\eta_1, \eta_2 \in \overline{\mathrm{Eff}}$ so that $\eta_i \cdot L_1 \cdot \ldots \cdot L_{n-2} = 0$, $i = 1, 2$.*

One may analogously reformulate the result of Theorem 2.13 in algebraic terms to characterize any $\eta \in \mathrm{A}^1$ and $L_1, \ldots, L_{n-2} \in \mathrm{Nef}$ such that $\eta \cdot L_1 \cdot \ldots \cdot L_{n-2} = 0$ (in Definition 2.10, one may then replace B by any ample class).

From the perspective of convex geometry, there is of course nothing new in the present formulation. The point of the algebraic formulation is, however, that the same algebraic structures carry over to other mathematical problems [17]. The statement of Corollary 16.2 (for example) therefore gives rise to natural conjectures on what analogues of the results of this paper might look like in other contexts.

Example 16.3. A structure similar to $\mathrm{A}(P)$ arises in algebraic geometry: here D is the space of divisors, A^k is the space of k-cycles modulo numerical equivalence, and $\cdot$ is the intersection product on a projective variety [14]. The ample, nef, big, and pseudoeffective cones are described in [19]. One might therefore ask whether a result such as Corollary 16.2 carries over to this setting, at least in sufficiently nice situations. To the best of our knowledge this question is entirely open, except for toric varieties which admit a precise correspondence with convex geometry [13, 9] (for which such a conclusion follows from the results of this paper).

Example 16.4. A structure similar to $\mathrm{A}(P)$ arises in the theory of mixed discriminants [38]. This is a much simpler setting: for example, here $\mathrm{Amp} = \mathrm{Big}$ is the cone of positive definite matrices, so Corollary 16.2 reduces to the hard Lefschetz theorem. On the other hand, it was shown by Panov [23] that this setting admits degenerate extremals in complete analogy to Theorem 2.13. This setting therefore provides an example outside convexity in which analogous structures appear.

Acknowledgments. This work was supported in part by NSF grant DMS-1811735 and by the Simons Collaboration on Algorithms & Geometry. We thank Karim Adiprasito, Swee Hong Chan, June Huh, János Kollár, Igor Pak, Greta Panova, Rolf Schneider, and Amir Yehudayoff for helpful comments.

References

[1] A. D. Alexandrov. Zur Theorie der gemischten Volumina von konvexen Körpern II. *Mat. Sbornik N.S.*, 2:1205–1238, 1937.

[2] A. D. Alexandrov. *Selected works. Part I.* Gordon and Breach Publishers, Amsterdam, 1996.

[3] T. Bonnesen and W. Fenchel. *Theory of convex bodies.* BCS Associates, Moscow, ID, 1987.

[4] S. Boucksom, C. Favre, and M. Jonsson. Differentiability of volumes of divisors and a problem of Teissier. *J. Algebraic Geom.*, 18(2):279–308, 2009.

[5] Y. D. Burago and V. A. Zalgaller. *Geometric inequalities.* Springer-Verlag, Berlin, 1988.

[6] F. R. K. Chung, P. C. Fishburn, and R. L. Graham. On unimodality for linear extensions of partial orders. *SIAM J. Algebraic Discrete Methods*, 1(4):405–410, 1980.

[7] J. B. Conway. *A course in functional analysis*, volume 96 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, second edition, 1990.

[8] G. Ewald. On the equality case in Alexandrov-Fenchel's inequality for convex bodies. *Geom. Dedicata*, 28(2):213–220, 1988.

[9] G. Ewald. *Combinatorial convexity and algebraic geometry*, volume 168 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, 1996.

[10] G. Ewald and E. Tondorf. A contribution to equality in Alexandrov-Fenchel's inequality. *Geom. Dedicata*, 50(3):217–233, 1994.

[11] J. Favard. Sur les corps convexes. *J. Math. Pures Appl. (9)*, 12:219–282, 1933.

[12] W. Fenchel. Inégalités quadratiques entre les volumes mixtes des corps convexes. *C. R. Acad. Sci. Paris*, 203:647–650, 1936.

[13] W. Fulton. *Introduction to toric varieties*, volume 131 of *Annals of Mathematics Studies.* Princeton University Press, Princeton, NJ, 1993.

[14] W. Fulton. *Intersection theory.* Springer-Verlag, Berlin, second edition, 1998.

[15] M. Gromov. Convex sets and Kähler manifolds. In *Advances in differential geometry and topology*, pages 1–38. World Sci. Publ., Teaneck, NJ, 1990.

[16] P. M. Gruber. *Convex and discrete geometry*, volume 336 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences].* Springer, Berlin, 2007.

[17] J. Huh. Combinatorial applications of the Hodge-Riemann relations. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures*, pages 3093–3111. World Sci. Publ., Hackensack, NJ, 2018.

[18] O. Kallenberg. *Foundations of modern probability.* Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002.

[19] R. Lazarsfeld. *Positivity in algebraic geometry. I.* Springer-Verlag, Berlin, 2004. Classical setting: line bundles and linear series.

[20] B. Lehmann. Comparing numerical dimensions. *Algebra Number Theory*, 7(5):1065–1100, 2013.

[21] P. McMullen. On simple polytopes. *Invent. Math.*, 113(2):419–444, 1993.

[22] H. Minkowski. Volumen und Oberfläche. *Math. Ann.*, 57(4):447–495, 1903.

[23] A. A. Panov. Some properties of mixed discriminants. *Mat. Sb. (N.S.)*, 128(170)(3):291–305, 446, 1985.

[24] J. R. Sangwine-Yager. Mixed volumes. In *Handbook of convex geometry, Vol. A, B*, pages 43–71. North-Holland, Amsterdam, 1993.

[25] R. Schneider. Kinematische Berührmaße für konvexe Körper und Integralrelationen für Oberflächenmaße. *Math. Ann.*, 218(3):253–267, 1975.

[26] R. Schneider. On the Aleksandrov-Fenchel inequality. In *Discrete geometry and convexity*, volume 440 of *Ann. New York Acad. Sci.*, pages 132–141. New York Acad. Sci., 1985.

[27] R. Schneider. On the Aleksandrov-Fenchel inequality involving zonoids. *Geom. Dedicata*, 27(1):113–126, 1988.

[28] R. Schneider. Equality in the Aleksandrov-Fenchel inequality—present state and new results. In *Intuitive geometry (Szeged, 1991)*, volume 63 of *Colloq. Math. Soc. János Bolyai*, pages 425–438. North-Holland, Amsterdam, 1994.

[29] R. Schneider. Polytopes and Brunn-Minkowski theory. In *Polytopes: abstract, convex and computational*, volume 440 of *NATO Adv. Sci. Inst. Ser. C*, pages 273–299. Kluwer, 1994.

[30] R. Schneider. *Convex bodies: the Brunn-Minkowski theory*. Cambridge University Press, expanded edition, 2014.

[31] Y. Shenfeld and R. van Handel. Mixed volumes and the Bochner method. *Proc. Amer. Math. Soc.*, 147(12):5385–5402, 2019.

[32] Y. Shenfeld and R. Van Handel. The extremals of Minkowski's quadratic inequality. *Duke Math. J.*, 2022. To appear.

[33] R. P. Stanley. Two combinatorial applications of the Aleksandrov-Fenchel inequalities. *J. Combin. Theory Ser. A*, 31(1):56–65, 1981.

[34] R. P. Stanley. Two poset polytopes. *Discrete Comput. Geom.*, 1(1):9–23, 1986.

[35] R. P. Stanley. Log-concave and unimodal sequences in algebra, combinatorics, and geometry. In *Graph theory and its applications: East and West (Jinan, 1986)*, volume 576 of *Ann. New York Acad. Sci.*, pages 500–535. New York Acad. Sci., New York, 1989.

[36] W. Süss. Zusammensetzung von eikörpern und homothetische eiflächen. *Tôhoku Math. J.*, 35:47–50, 1932.

[37] B. Teissier. Bonnesen-type inequalities in algebraic geometry. I. Introduction to the problem. In *Seminar on Differential Geometry*, volume 102 of *Ann. of Math. Stud.*, pages 85–105. Princeton Univ. Press, Princeton, N.J., 1982.

[38] V. A. Timorin. The mixed Hodge-Riemann bilinear relations in the linear situation. *Funct. Anal. Appl.*, 32(4):268–272, 1998.

[39] V. A. Timorin. An analogue of the Hodge-Riemann relations for simple convex polytopes. *Russ. Math. Surv.*, 54(2):381–426, 1999.

[40] H. Weyl. Über die Starrheit der Eiflächen und konvexen Polyeder. *Sitzungsber. Preuss. Akad. Wiss. Berlin*, pages 250–266, 1917.

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA
Email address: shenfeld@mit.edu

Fine Hall 207, Princeton University, Princeton, NJ 08544, USA
Email address: rvan@math.princeton.edu