Sharp stability and exact cap optima for Rademacher sums

Dev D. Goyal

1 November 2026

Abstract

For a normalized finite Rademacher sum, we study the central support gap and normal-approximation discrepancy under an upper bound on every coefficient. Our principal result gives sharp quantitative geometry for large gaps: if $|b_i| \le a$, every signed sum has magnitude at least $x > a/2$, and $C = \{i : |b_i| \ge x\}$, then C has positive odd size k and the squared Euclidean distance to the flat vector on C is at most $2 - 2\sqrt{1-(a-x)^2}$. Equality is exactly an odd saturated core with one residual coefficient. Combining the finite-slice geometry with a sharp tail-compression argument gives a uniform stability theorem for the exact support-gap optimum at every cap. Its optimal linear constant is $\sqrt{2+\sqrt{2}}$, attained by splitting a residual into two equal parts. We determine the exact local nonlinear modulus on the saturated branch and the exact fixed-core flat modulus, including all equality cases and the admissibility restrictions imposed by the cap.

We extend the exact full Kolmogorov cap optimum from the previously proved range $a \ge 1/\sqrt{3}$ to every $a \ge 1/\sqrt{5}$, with all optimizing coefficient vectors, thresholds, and transitions. The two new cap intervals alternate between a saturated core with one residual and the next flat vector. The probability proofs combine two- and four-coordinate localization with a four-shift lower-discrepancy envelope, certified on 27972 rational boxes by validated interval arithmetic. They use explicitly stated Keller–Klein, Hollom–Portier, and Tomaszewski inputs. The deterministic stability results are independent of those external theorems and certificates. Lean 4 checks the actual finite-vector inverse, all-cap stability and optimizer selection, exact saturated modulus and equality, fixed-flat inequalities, and concrete sharpness examples. Writing $c(t) = \Phi(t) - 1/2$ and $\mathfrak{G}(a)$ for the optimal gap, the lower-discrepancy optimum equals $c(\mathfrak{G}(a))$ throughout the new range, so every lower optimizer has an empty strict central interval. The probability optimization and its interval certificates remain outside this kernel-checked scope.

1 Problem, contributions, and prior work

Let ε_i be independent uniform signs and let

$$X_b = \sum_{i=1}^{N} b_i \varepsilon_i, \qquad \sum_{i=1}^{N} b_i^2 = 1, \qquad |b_i| \le a \le 1.$$

The number of coefficients is arbitrary and finite. A cap a need not be attained by a coefficient. Write

$$g(b) = \min_{\varepsilon \in \{-1,1\}^N} \left| \sum_i b_i \varepsilon_i \right|, \qquad \mathfrak{G}(a) = \sup_{N,b} g(b).$$

Signs of coefficients may always be absorbed into the random signs. Vectors of different lengths are compared after arranging their absolute coefficients in decreasing order and padding with zeros. We

write $\operatorname{dist}(b, v)$ for the resulting Euclidean distance, and occasionally abbreviate it by $\|b - v\|_2$ when this convention is understood. Thus (a^k, r) denotes k copies of a followed by one residual r, with zero residuals omitted. The flat vector of size k has all its nonzero coefficients equal to $k^{-1/2}$.

The earlier research package [10] established the exact value

$$\mathfrak{G}(a) = \max\left\{\frac{1}{\sqrt{n+2}},\, a - \sqrt{1-na^2}\right\}, \qquad \frac{1}{\sqrt{n+2}} \le a \le \frac{1}{\sqrt{n}}, \quad n \text{ odd}. \tag{1}$$

It also classified the maximizers: the winning type is either the next odd-flat vector or an odd saturated core with one residual, and both types maximize at the unique transition. We include the complete short proof in Section 2; this exact value is an established starting point, not a new theorem of the present extension.

1.1 New structural results

The question addressed here is quantitative: how much coefficient structure does a nearly maximal gap force? A variance estimate alone would typically yield a square-root dependence on the extremal deficit. The relevant majority slice instead has a sharp linear geometric bound. Theorem 3.1 proves, for every certified gap $x > a/2$ and the actual core $C = \{i : |b_i| \ge x\}$,

$$\left\|(|b_i|)_i - |C|^{-1/2}\mathbf{1}_C\right\|_2^2 \le 2 - 2\sqrt{1-(a-x)^2}.$$

The cap can be unattained and the certified gap can be smaller than the actual one. Equality determines the coefficients completely. The strict threshold $1/2$ is optimal: the normalized vector $(2,1,1,1)/\sqrt{7}$ has gap equal to half its largest coefficient but has an even core.

Theorem 4.1 turns this geometry into stability for the exact optimum (1). If $x = g(b) > \max\{a/2, (n+4)^{-1/2}\}$, the core has size n or $n+2$, and the appropriate candidate lies within

$$\sqrt{2+\sqrt{2}}\,(\mathfrak{G}(a) - g(b))$$

of b. The constant is best possible. Below the difference between the two competing values, the candidate must be the actual unique optimizer; at a transition both candidates are necessary. The fixed-flat curve in Proposition 3.3 and the exact saturated modulus in Theorem 5.1 further describe the equality geometry. In particular, the optimal local coefficient on a fixed saturated branch is 1; the larger universal constant appears when the entire residual splits into two equal coefficients.

1.2 Exact normal discrepancy down to the five-flat cap

Let Φ and φ denote the standard normal distribution and density, and put $c(x) = \Phi(x) - 1/2$. For $x \ge 0$ define

$$U_b(x) = \tfrac{1}{2}\mathbb{P}(|X_b| \le x) - c(x).$$

For $x > 0$, define $L_b(x) = c(x) - \frac{1}{2}\mathbb{P}(|X_b| < x)$, and set $L_b(0) = 0$. The distinction between the closed and open central events is essential at atoms. Symmetry and continuity of Φ give

$$d_K(X_b, N(0,1)) = \max\{\sup_{x\ge 0} U_b(x), \sup_{x\ge 0} L_b(x)\}.$$

Write $\mathfrak{D}(a)$, $\mathfrak{L}(a)$, and $\mathfrak{K}(a)$ for the cap suprema of U, L, and the full Kolmogorov distance respectively.

The new probability results determine the full optimum on the next two cap intervals. Combining them with the established large-cap formulas gives the following explicit profile.

Theorem 1.1 (The full cap optimum for $a \geq 1/\sqrt{5}$)**.** *For* $1/\sqrt{5} \leq a \leq 1$,

$$\mathfrak{K}(a) = \begin{cases} \max\{c(1/\sqrt{5}),\, 3/16 - c(\sqrt{1-4a^2})\}, & 1/\sqrt{5} \leq a \leq 1/2, \\ \max\{3/16,\, c(a - \sqrt{1-3a^2})\}, & 1/2 \leq a \leq 1/\sqrt{3}, \\ \max\{c(1/\sqrt{3}),\, 1/4 - c(\sqrt{1-2a^2})\}, & 1/\sqrt{3} \leq a \leq 1/\sqrt{2}, \\ \max\{1/4,\, c(a - \sqrt{1-a^2})\}, & 1/\sqrt{2} \leq a \leq 1. \end{cases} \tag{2}$$

When the constant branch is strictly larger, the unique maximizing type is the next flat vector, of size $5, 4, 3, 2$ *respectively. When the varying branch is strictly larger, it is the saturated vector* $(a^n, \sqrt{1-na^2})$ *with* $n = 4, 3, 2, 1$ *respectively. Each interval has one transition; precisely both types maximize there. Winning even cores attain their value in* U *at the residual threshold, odd cores in* L *at the support gap. A flat vector attains its value in* U *at zero when its size is even, and in* L *at its support gap when its size is odd.*

The first two rows and their complete equality statements are new: they are proved in Theorems 8.1 and 7.2. The last two rows are the earlier exact large-cap theorem [10], whose complete proofs are in that unpublished companion and are not repeated here. The formulas agree at shared endpoints.

Theorem 6.1 first determines the positive discrepancy for every $a \geq 1/2$. A competitive threshold forces two large coefficients and an explicit residual variance obstruction. A four-coordinate argument settles the corresponding positive side on $[1/\sqrt{5}, 1/2]$. The new lower-discrepancy argument expresses the law as a four-shift mixture and bounds its strict central interval. Two finite rational-box partitions and one three-box partition certify the remaining scalar inequalities; their full analytic interface and reproducible enclosure algorithm are given in Section 7. In fact the lower result is stronger than the full-distance formula requires: $\mathfrak{L}(a) = c(\mathfrak{G}(a))$ for every $a \geq 1/\sqrt{5}$, and its optimizers are exactly the support-gap optimizers, including all ties. Thus the deterministic classification identifies every lower-discrepancy optimizer throughout that range.

The preceding package also established $\mathfrak{K}(a)$ for every sufficiently small cap, with an unevaluated positive cutoff [10]. That asymptotic proof remains in the companion. The deterministic gap formula alone does not exclude competitive thresholds containing positive central mass; the new probability reductions supply this missing control on the two additional finite-cap intervals.

1.3 Comparison with the literature

Keller and Klein's work on Tomaszewski's conjecture includes sharp normal-approximation questions for Rademacher sums [1]. The original small-coefficient research package [9] supplies the missing upper-concentration regime relative to Ivanisvili's October 2026 manuscript [8]. Claims using the completed comparison retain the external verification boundary described in Section 6. Hollom and Portier determine the uncapped strict-tail profile used in our new probability argument [2]; our coefficient-cap optimization is a different extremal problem. Their method already uses capped tail bounds and conditioning on leading signs; those devices are not claimed as new here.

Sharp coefficient stability has substantial precedents. De, Diakonikolas, and Servedio bound the excess of $\mathbb{E}|X_b|$ above $1/\sqrt{2}$ by distance from its two-flat minimizers [3]. Fang and Wang give a sharp nonlinear profile near that set [4]; Xie determines the optimal global linear constant $(2-\sqrt{2})^{3/2}/8$, attained at four-flat coefficients [5]. These are moment inequalities. Our objective is the minimum absolute signed sum, with odd cores and coefficient caps; these displayed sharpness vectors have zero support gap. Xie's Fourier–Haagerup argument differs from the finite-slice and tail-repair geometry

used here. Bell studies near-maximal Berry–Esseen ratios for general independent arrays and their common-lattice and two-point structure [7]; that objective and normalization also differ from the capped central-gap problem.

The equal-core-plus-residual construction itself is not novel; it occurs in earlier Rademacher extremal investigations, including Pinelis [6]. Nor do we claim independent novelty for the elementary hypersimplex or ℓ_1/ℓ_2 inequalities used below. The contribution is their sharp consequence for actual signed-sum gaps, across every cap, with a complete near-optimizer and equality analysis, together with the exact probability cap profile on the two additional intervals. The primary-source comparison found no identical prior theorem, but it is not a guarantee of priority.

2 Majority structure and the exact support-gap optimum

The extremal value and classification in this section are from the earlier research package [10]; their short proof is included because it supplies the geometry used in all the new stability theorems.

For a finite positive coefficient vector $b = (b_1, \dots, b_N)$ define

$$g(b) = \min_{\varepsilon \in \{-1,1\}^N} \left| \sum_{i=1}^{N} b_i \varepsilon_i \right|,$$

and, for $0 < a \le 1$, put

$$\mathfrak{G}(a) = \sup_{\substack{N<\infty,\ \sum_i b_i^2 = 1 \\ 0 < b_i \le a}} g(b).$$

Zero coefficients may be added or removed freely; signs of coefficients may be absorbed into the signs in the minimum.

Theorem 2.1 (Exact support-gap optimum for every cap)**.** *Let $n \ge 1$ be odd and let*

$$\frac{1}{\sqrt{n+2}} \le a \le \frac{1}{\sqrt{n}}.$$

Then

$$\mathfrak{G}(a) = \max\left\{ \frac{1}{\sqrt{n+2}},\ a - \sqrt{1 - na^2} \right\}. \tag{3}$$

The supremum is attained. Up to permutations, changes of signs, and zero coefficients, all maximizing vectors are as follows. If the first candidate is strictly larger, the only maximizer consists of $n+2$ equal coefficients $1/\sqrt{n+2}$. If the second candidate is strictly larger, the only maximizer consists of n coefficients equal to a and one coefficient $\sqrt{1-na^2}$, with a zero residual omitted. At a tie, precisely both types maximize.

There is exactly one tie in the open interval $(1/\sqrt{n+2}, 1/\sqrt{n})$. Writing $\lambda = 1/\sqrt{n+2}$, its location is

$$a_n = \frac{\lambda + \sqrt{n+1-n\lambda^2}}{n+1} = \frac{1 + \sqrt{n^2+2n+2}}{(n+1)\sqrt{n+2}}. \tag{4}$$

Thus the first type maximizes below a_n and the second above a_n. The formulas agree at shared endpoints of adjacent odd-n intervals, so they specify the optimum and every maximizing vector for every $0 < a \le 1$.

2.1 A high-gap sum has a majority core

The following structural fact does not require variance normalization. For positive coefficients, a Hamming layer consists of the sign vectors having a fixed number of positive signs.

Lemma 2.2 (Majority core). *Let $0 < b_i \le a$ be finitely many positive coefficients. Suppose $x > a/2$ and*

$$\left|\sum_i b_i \varepsilon_i\right| \ge x \quad \text{for every sign vector } \varepsilon. \tag{5}$$

Put

$$C = \{i : b_i \ge x\}, \qquad k = |C|, \qquad \ell = \sum_{i \notin C} b_i.$$

Then k is positive and odd. The sign of the complete sum is the majority sign of the k coordinates in C and is independent of every coordinate outside C. Moreover, for every core sign pattern $\eta \in \{-1, 1\}^C$,

$$\left|\sum_{i \in C} b_i \eta_i\right| \ge x + \ell. \tag{6}$$

In particular,

$$x + \ell \le \frac{1}{k} \sum_{i \in C} b_i. \tag{7}$$

Proof. Flipping a coordinate outside C changes the full sum by $2b_i < 2x$. Such an edge cannot join a sum at least x to one at most $-x$. Hence its endpoints have the same sign. The cube of outside coordinates is connected, so the sign is constant on every fiber with the core coordinates fixed.

For a fixed core outcome $S = \sum_{i \in C} b_i \eta_i$, the two residual endpoints are $S - \ell$ and $S + \ell$. They have the same sign, and each has magnitude at least x. Thus either $S - \ell \ge x$ or $S + \ell \le -x$, proving (6). In particular $k = 0$ is impossible.

Swap a positive and a negative sign at two core coordinates. The core sum changes by at most $2|b_i - b_j| \le 2(a - x) < 2x$. By (6), this change cannot reverse its sign. Any two sign patterns in the same Hamming layer are connected by such swaps. Therefore the sign is constant on each layer. As the number of positive signs increases, these layer signs are nondecreasing: one may pass from any layer to the next by changing a negative sign to a positive one, which strictly increases the sum. Negating all signs reverses the sign and exchanges layers j and $k - j$.

If k were even, its middle layer would have a constant sign equal to its own negative, which is impossible. Thus k is odd, and monotonicity together with sign reversal forces the majority rule: layers with at least $(k + 1)/2$ positive signs are positive, and the others are negative.

Average the core sum over the layer with exactly $(k + 1)/2$ positive signs. Each coordinate has average sign $1/k$, since it is positive in the fraction $(k + 1)/(2k)$ of the layer. The average sum is therefore $k^{-1} \sum_{i \in C} b_i$. Every sum on this layer is at least $x + \ell$ by (6), proving (7). □

A pair-deletion proof used in the formalization. The core margin also gives the stronger coordinatewise conclusion $x + \ell \le b_i$ for every $i \in C$. Indeed, let $y = x + \ell$ and choose two core coefficients u, v. For every signing of the remaining coordinates, the two opposite assignments to this pair give $S + (u - v)$ and $S - (u - v)$. Both have absolute value at least y, whereas $|u - v| \le a - x < x \le y$. They therefore have the same sign and force $|S| \ge y + |u - v| \ge y$. Deleting any pair preserves the gap. Repeated deletion would give $0 \ge y > 0$ if the core size were even; if it is odd, one may retain any designated coefficient and conclude that it is at least y. This proves

oddness and the coordinatewise bound directly, without an assumed majority rule. Summing their squares gives $k(x+\ell)^2 \le 1-V$ in the normalized case.

Corollary 2.3 (Variance form)**.** *Under the hypotheses of Lemma 2.2, assume in addition* $\sum_i b_i^2 = 1$ *and put* $V = \sum_{i\notin C} b_i^2$*. Then*

$$x \le \sqrt{\frac{1-V}{k}} - \sqrt{V}. \tag{8}$$

If $ka^2 \le 1$*, this further implies*

$$x \le a - \sqrt{1-ka^2}. \tag{9}$$

Proof. One has $\ell \ge \sqrt{V}$, and Cauchy–Schwarz gives $k^{-1}\sum_{i\in C} b_i \le \sqrt{(1-V)/k}$. Combine these inequalities with (7). Also $V \ge 1-ka^2$, and the function $v \mapsto \sqrt{(1-v)/k} - \sqrt{v}$ decreases strictly on $[0,1]$. Substituting its smallest permitted argument gives (9). □

2.2 Exact optimization and equality

Proof of Theorem 2.1. Put $\lambda = 1/\sqrt{n+2}$ and $g = a - \sqrt{1-na^2}$. The $(n+2)$-flat vector is admissible and has gap λ. Whenever $g \ge \lambda$, the residual $r = \sqrt{1-na^2}$ satisfies $0 \le r \le a-\lambda < a$; hence the saturated vector of n coefficients a and the residual r is admissible. Since n is odd and $r < a$, its minimum absolute signed sum is $a - r = g$. Thus the right side of (3) is attained.

To prove the upper bound and classify equality, take an arbitrary admissible vector and write $x = g(b)$. If $x < \lambda$ there is nothing to prove. Otherwise

$$x \ge \lambda > \frac{a}{2},$$

because $a \le 1/\sqrt{n}$ and $4n > n+2$ for $n \ge 1$. Apply Lemma 2.2 with the threshold x. The number of core coefficients satisfies $kx^2 \le \sum_i b_i^2 = 1$, so $k \le n+2$. If $k = n+2$, equality is forced throughout

$$1 \ge \sum_{i\in C} b_i^2 \ge (n+2)x^2 \ge (n+2)\lambda^2 = 1.$$

Thus $x = \lambda$, all core coefficients equal λ, and there is no nonzero residual. This is exactly the $(n+2)$-flat type.

In the remaining case $k < n+2$. Since k is odd, $k \le n$. In particular $ka^2 \le na^2 \le 1$, so Corollary 2.3 gives

$$x \le a - \sqrt{1-ka^2} \le a - \sqrt{1-na^2} = g. \tag{10}$$

This proves (3).

If a maximizer is not the $(n+2)$-flat type, then it belongs to the second case and $x = \max\{\lambda, g\} \le g$. Hence $x = g$ and all inequalities in (10) and its proof are equalities. The last inequality is strict for $k < n$, so $k = n$. Strict decrease in the variance substitution gives $V = 1-na^2$. Since all n core coefficients are at most a, their squared sum na^2 forces all of them to equal a. Equality $\ell = \sqrt{V}$ forces at most one nonzero residual coefficient: for nonnegative numbers, $(\sum b_i)^2 = \sum b_i^2$ holds only when at most one is nonzero. That residual is $\sqrt{1-na^2}$, or is absent when this value is zero. Conversely this vector has already been shown to attain the active second candidate. The asserted classification follows, including at a tie.

Finally $a \mapsto a - \sqrt{1-na^2}$ increases strictly on the interval in the theorem. At its left endpoint it equals $(1-\sqrt{2})\lambda < \lambda$, and at its right endpoint it is $1/\sqrt{n} > \lambda$. Solving $a - \sqrt{1-na^2} = \lambda$ gives (4); the positive solution is the unique one in the interval. At a shared endpoint $a = 1/\sqrt{n}$, the result is the n-flat vector, matching the neighboring interval's constant branch. This also covers $n = 1$ and $a = 1$. □

3 A sharp global inverse theorem for large support gaps

For a nonzero finite real vector b, write

$$g(b) = \min_{\varepsilon_i \in \{-1,1\}} \left| \sum_i b_i \varepsilon_i \right|, \qquad m(b) = \max_i |b_i|.$$

Signs of coefficients may be absorbed into the independent signs, so the proofs below use nonnegative coefficients and allow zeros. For a nonempty set C of cardinality k, let $e_C = k^{-1/2}\mathbf{1}_C$.

Theorem 3.1 (Global sharp gap inverse)**.** *Suppose $\sum_i b_i^2 = 1$, $|b_i| \le a \le 1$, and $a/2 < x \le g(b)$. The cap need not be attained and x may be any certified lower bound for the actual gap. Let*

$$C = \{i : |b_i| \ge x\}, \quad k = |C|, \quad \delta = a - x, \quad \rho = \big\|(|b_i|)_i - e_C\big\|_2.$$

Then k is positive and odd, the sign of the signed sum is the majority sign on C, and

$$\rho^2 \le 2 - 2\sqrt{1-\delta^2}. \tag{11}$$

Equivalently, $\rho\sqrt{1-\rho^2/4} \le \delta$. Equality holds if and only if, after signs, permutations, and zeros are discarded, the vector consists of k copies of a and one residual coefficient δ; the residual is omitted when $\delta = 0$. Necessarily $ka^2 + \delta^2 = 1$ and $\delta < a/2$; equality also forces $a = m(b)$ and $x = g(b)$.

For every odd $k \ge 1$ and every $0 \le r < 1/\sqrt{4k+1}$, the vector with k coefficients $\sqrt{(1-r^2)/k}$ and one coefficient r satisfies the hypotheses and attains equality. In particular the leading linear coefficient in $\rho \le (1 + o(1))\delta$ cannot be reduced.

The inequality is valid throughout the high-gap regime, with no asymptotic smallness assumption and no numerical integration. Its equality family is specified explicitly; the assertion is not that this right side is an attained optimum at every arbitrarily prescribed value of δ.

3.1 A sharp finite-slice inequality

Lemma 3.2. *Let $k = 2h + 1 \ge 3$, let $d_1, \dots, d_k$ be real numbers with $\sum_i d_i = 0$, and put*

$$s = \left(\sum_i d_i^2 \right)^{1/2}, \quad M = \max_i d_i, \quad A = \max_{|J|=h} \sum_{i \in J} d_i, \quad \gamma = \sqrt{\frac{k-1}{k}}.$$

Then

$$2A \ge \gamma s, \qquad M \ge \frac{s}{k\gamma}, \qquad s \le \gamma(M + 2A). \tag{12}$$

For $s > 0$, equality in either of the first two inequalities occurs precisely when $k-1$ entries are equal and positive, and the last entry is minus their sum. The same statement holds for equality in the third.

Proof. First, whenever $0 \le z_i \le 1$ and $\sum_i z_i = h$, one has

$$\sum_i d_i z_i \le A. \tag{13}$$

Indeed, a nonintegral such vector has at least two fractional coordinates. Choose two, say i, j, with $d_i \ge d_j$, and increase z_i while decreasing z_j by the same amount until at least one becomes integral.

This preserves feasibility and does not decrease the objective. Repeating terminates at the indicator of an h-subset and proves the claim.

Assume $s > 0$. Since $d_i = -\sum_{j \neq i} d_j$, Cauchy–Schwarz gives $d_i^2 \leq (k-1)(s^2 - d_i^2)$, hence $|d_i| \leq \gamma s$. Set

$$r = \frac{h}{\sqrt{k(k-1)}} = \frac{h}{k\gamma}, \qquad z_i = \frac{h}{k} + r\frac{d_i}{s}.$$

Then $0 \leq z_i \leq 2h/k < 1$ and $\sum_i z_i = h$. Equation (13) yields $A \geq rs$, so $2A \geq \gamma s$.

If equality holds, some z_j must be zero. Otherwise all coordinates are strictly between zero and one, and a small transfer from a coordinate with smaller d_i to one with larger d_i strictly increases the objective, contradicting equality in (13). Thus $d_j = -\gamma s$, and equality in the displayed Cauchy–Schwarz inequality forces the other $k-1$ entries to be equal and positive. Conversely that vector attains the bound.

For the second inequality put $u_i = M - d_i \geq 0$. Then

$$s^2 + kM^2 = \sum_i u_i^2 \leq \left(\sum_i u_i\right)^2 = k^2 M^2.$$

Thus $M \geq s/\sqrt{k(k-1)}$. Equality for $s > 0$ requires at most one u_i to be nonzero, which is the same vector shape. Finally $\gamma + 1/(k\gamma) = 1/\gamma$, so adding the first two inequalities proves the third and its equality statement. The case $s = 0$ is immediate. □

3.2 From a genuine support gap to the coefficient distance

Proof of Theorem 3.1. Replace b_i by $|b_i|$. Put

$$\ell = \sum_{i \notin C} b_i, \qquad V = \sum_{i \notin C} b_i^2.$$

Lemma 2.2 and its pair-deletion argument give positive odd k, the majority rule, and $b_i \geq x + \ell$ for every core coordinate. In particular the least positive core sum lies on the nearest majority layer.

Let $\beta = k^{-1}\sum_{i \in C} b_i$ and $T = \sum_{i \in C}(b_i - \beta)^2$. The actual maximum lies in C. The least positive core sum occurs on a nearest majority layer, so, with $h = (k-1)/2$ and $A = \max_{|J|=h} \sum_{i \in J}(b_i - \beta)$,

$$x + \ell \leq \beta - 2A, \qquad (a - \beta) + 2A \leq \delta - \ell.$$

In particular $0 \leq \ell \leq \delta$. If $k \geq 3$, Lemma 3.2, with $\max_{i \in C}(b_i - \beta) \leq a - \beta$, gives

$$T \leq \frac{k-1}{k}(\delta - \ell)^2. \tag{14}$$

For $k = 1$, $T = 0$, so the same formula remains true. Since $V \leq \ell^2$,

$$T + V \leq \frac{k-1}{k}(\delta - \ell)^2 + \ell^2 \leq \delta^2. \tag{15}$$

The last inequality follows for $0 \leq \ell \leq \delta$ by expansion; if $\delta > 0$ it is strict unless $\ell = \delta$.

Normalization gives $k\beta^2 + T + V = 1$. Since $\beta > 0$,

$$\rho^2 = 2 - 2\sqrt{k}\,\beta = 2 - 2\sqrt{1 - (T+V)} \leq 2 - 2\sqrt{1 - \delta^2}.$$

Here $0 \le \delta < a/2 \le 1/2$, so every square root is well defined. The identity $T + V = \rho^2 - \rho^4/4$ also gives the equivalent form.

If $\delta > 0$, equality forces $\ell = \delta$, $T = 0$, and $V = \ell^2$. Hence the core is flat. Its value β satisfies $a = x + \delta = x + \ell \le \beta \le a$, so the core value is a, and precisely one nonzero tail coefficient equals δ. If $\delta = 0$, the same inequalities give $T = V = 0$ and $\beta = a$, so the vector is odd flat. Conversely, for an odd flat core of coefficient a and one residual $r < a/2$, the actual gap is $a - r$, the core selected by its gap is exactly the flat core, and $\rho^2 = 2 - 2\sqrt{1 - r^2}$. Choosing $a = \sqrt{(1 - r^2)/k}$ proves the asserted equality families and sharpness. □

Sharpness of the structural threshold. The strict threshold $g > m/2$ cannot be relaxed to $g \ge m/2$. For the normalized vector $(2, 1, 1, 1)/\sqrt{7}$, the gap is $1/\sqrt{7}$ and the maximum is $2/\sqrt{7}$, while its gap-defined core has even cardinality four. The large-gap core theorem therefore has the exact threshold $1/2$.

3.3 A complementary sharp bound relative to a fixed flat dimension

Proposition 3.3. *In the actual-gap case $x = g(b) = g$ of Theorem 3.1, if $k \ge 3$,*

$$g \le \frac{1 - \rho^2/2}{\sqrt{k}} - \sqrt{\frac{k-1}{k}}\,\rho\sqrt{1 - \rho^2/4}.$$

Equivalently,

$$\rho^2 \le 2 - \frac{2}{\sqrt{k}}\left(g + \sqrt{(k-1)(1-g^2)}\right).$$

Equality holds precisely for $k - 1$ equal high coefficients and one low coefficient, with no tail. These vectors are

$$\left(\underbrace{\sqrt{\frac{1-g^2}{k-1}}, \dots, \sqrt{\frac{1-g^2}{k-1}}}_{k-1}, g\right), \qquad \frac{1}{\sqrt{4k-3}} < g \le \frac{1}{\sqrt{k}}.$$

For a prescribed cap a, the original conditions $\sqrt{(1-g^2)/(k-1)} \le a$ and $g > a/2$ must also hold. At $a = 1/\sqrt{k}$ the nonflat equality family is inadmissible; on an interior flat phase it is admissible for all sufficiently small positive deficits.

Proof. The first slice inequality gives $g \le \beta - \gamma\sqrt{T} - \ell$. Since $\ell \ge \sqrt{V}$ and $0 < \gamma < 1$, $\gamma\sqrt{T} + \sqrt{V} \ge \gamma\sqrt{T + V}$, with equality only if $V = 0$. Substitute $\beta = (1 - \rho^2/2)/\sqrt{k}$ and $T + V = \rho^2 - \rho^4/4$. For the equivalent form put $\cos\theta = \sqrt{k}\beta = 1 - \rho^2/2$ and $\cos\psi = 1/\sqrt{k}$, with $\theta, \psi \in [0, \pi/2)$. The first inequality reads $g \le \cos(\theta + \psi)$; its positive right side implies $\theta + \psi < \pi/2$, so it can be inverted. This yields $\cos\theta \ge [g + \sqrt{(k-1)(1-g^2)}]/\sqrt{k}$. Equality follows from the exact equality cases in Lemma 3.2; the high-gap condition is precisely $g > \frac{1}{2}\sqrt{(1-g^2)/(k-1)}$. □

4 Uniform sharp stability of the exact cap optimum

Let $n \ge 1$ be odd and $1/\sqrt{n+2} \le a \le 1/\sqrt{n}$. Put

$$\lambda = (n+2)^{-1/2}, \quad \mu = (n+4)^{-1/2}, \quad r = \sqrt{1 - na^2}, \quad s = a - r, \quad G = \max\{\lambda, s\}, \quad C_* = \sqrt{2 + \sqrt{2}}.$$

Let v_f denote the $(n+2)$-flat vector, and, when $r \le a$, let $v_s = (a^n, r)$, with zero entries omitted. All coefficient vectors in distances are sorted in decreasing magnitude and padded with zeros.

Theorem 4.1 (Sharp linear stability across all cap bands)**.** *Let b be any finite normalized coefficient vector with $|b_i| \le a$. Write $x = g(b)$ and $\Delta = G - x$. If*

$$x > \max\{a/2, \mu\}, \tag{16}$$

then the actual core $\{i : |b_i| \ge x\}$ has size n or $n+2$. More precisely:

1. *If its size is $n+2$, then $x \le \lambda$ and*

$$\|b - v_f\|_2 \le \sqrt{\frac{n+2}{n+1}}(\lambda - x) \le C_*\Delta.$$

2. *If its size is n, then $r < a/2$, the saturated candidate is admissible, $x \le s$, and*

$$\|b - v_s\|_2 \le C_*(s - x) \le C_*\Delta.$$

Consequently the distance to the admissible members of these two candidate types is at most $C_*\Delta$. If the branches are unequal and $\Delta < |\lambda - s|$, the candidate selected above is necessarily the unique maximizing type. At a branch tie both candidates are maximizers.

The constant C_* is optimal, even on a subinterval with a unique maximizer and under the strict trigger (16). The endpoint μ in the core-count trigger cannot be included in a uniform linear estimate of this form.

Unlike a bound that chooses a unique branch continuously in a, this statement remains valid through the switching caps. The need to keep both competing types near a switch is real: on either side of a switch, the losing candidate has deficit tending to zero while its distance to the unique winner tends to the nonzero distance between the two shapes.

4.1 Sharp concentration of a nonnegative tail

Lemma 4.2. *Let t be a finite nonnegative vector, $\sigma = \|t\|_2$, $\ell = \|t\|_1$, and let j be a largest coordinate, after adding a zero coordinate if the index set is empty. Then*

$$\|t - \sigma e_j\|_2 \le C_*(\ell - \sigma).$$

If $\ell > \sigma$, equality holds if and only if exactly two coordinates are nonzero and they are equal. If $\ell = \sigma$, both sides vanish.

Proof. Only $\sigma > 0$ and $\ell > \sigma$ need consideration. Normalize $\sigma = 1$, and put $z = t_j \in (0, 1)$. The squared distance is $2(1 - z)$. If $z \le 1/\sqrt{2}$, then $1 = \sum_i t_i^2 \le z\ell$, and hence

$$\frac{2(1-z)}{(\ell-1)^2} \le \frac{2z^2}{1-z} \le 2 + \sqrt{2}.$$

The last expression increases strictly with $z \in (0, 1)$, and its value at $1/\sqrt{2}$ is $2 + \sqrt{2}$. If $z \ge 1/\sqrt{2}$, then $\ell \ge z + \sqrt{1 - z^2}$. Put $u = \sqrt{(1-z)/(1+z)} \in (0, \sqrt{2} - 1]$. Direct substitution gives

$$\frac{\sqrt{2(1-z)}}{z + \sqrt{1-z^2} - 1} = \frac{\sqrt{1+u^2}}{1-u} \le \sqrt{2+\sqrt{2}}.$$

The middle expression increases strictly with u, as is seen by differentiation. In either case equality forces $z = 1/\sqrt{2}$. The first case further requires every nonzero t_i to equal z; the second requires at most one further nonzero coordinate. Both give exactly two equal nonzero coordinates. The converse is direct. □

4.2 The saturated branch

Lemma 4.3 (Repair of an almost saturated core)**.** *Let b be normalized and nonnegative, with coefficients at most a, actual gap $x > a/2$, and core size k satisfying $ka^2 \le 1$. Write*

$$r = \sqrt{1 - ka^2}, \quad e = a - r - x, \quad \ell = \|b_{C^c}\|_1, \quad \sigma = \|b_{C^c}\|_2,$$

and put $q = \ell - \sigma$, $D = \sigma - r$, and $\kappa = a - x - \ell$. Then $q, D, \kappa \ge 0$ and $e = q + D + \kappa$. If j is a largest tail coordinate, with a zero coordinate added when necessary, then

$$\operatorname{dist}(b, (a^k, r)) \le \|b_{C^c} - \sigma e_j\|_2 + \kappa + D. \tag{17}$$

Proof. The cap and normalization give $\sigma \ge r$, and the core lemma gives $\ell \le a - x$ and $x + \ell \le b_i \le a$ on the core. These prove the asserted signs and identity. Put $d_i = a - b_i$ on the core, $S = \sum d_i$, and $Q = \sum d_i^2$. Then $0 \le d_i \le \kappa$, so $Q \le \kappa S$ and

$$\sigma^2 - r^2 = 2aS - Q \ge (2a - \kappa)S.$$

If $\sigma + r = 0$, normalization forces every core coefficient to equal a, and the assertion is immediate. Otherwise

$$D = \frac{\sigma^2 - r^2}{\sigma + r} \ge S, \qquad \sqrt{Q + D^2} \le \sqrt{\kappa D + D^2} \le \kappa + D,$$

because $\sigma + r \le 2\ell = 2(a - \kappa - x) < 2a - \kappa$. First merge the tail to a coefficient σ at j. Changing the core to k copies of a and this residual from σ to r then costs $\sqrt{Q + D^2}$ in Euclidean distance. The triangle inequality proves (17). Since core coefficients exceed every tail coefficient, this alignment gives the distance under our sorting convention. □

Lemma 4.4. *Under the preceding hypotheses,*

$$\operatorname{dist}(b, (a^k, r)) \le C_* e.$$

For $e > 0$, equality holds precisely when the core consists of k copies of a and the tail consists of two equal coefficients $r/\sqrt{2}$; necessarily $e = (\sqrt{2} - 1)r$.

Proof. The tail concentration lemma and (17) give

$$\operatorname{dist}(b, (a^k, r)) \le C_* q + (e - q) \le C_* e.$$

Since $C_* > 1$, equality with $e > 0$ forces $q = e$, hence $D = \kappa = 0$. The core is therefore saturated and $\sigma = r$. Equality in the tail lemma gives exactly two equal residual coefficients. Their sum is $\sqrt{2}r$, yielding the stated deficit. Conversely this vector attains the bound whenever its actual gap $a - \sqrt{2}r$ exceeds $a/2$. When $e = 0$, all the nonnegative quantities vanish and the vector is the saturated comparator itself. □

4.3 The flat branch and the exact core-count trigger

Proof of Theorem 4.1. Absorb signs and discard zeros. Let k be the positive odd core size, and let ℓ, V be the tail sum and variance. The core lemma gives $\ell \le a - x < a/2$. Hence

$$1 = \sum_C b_i^2 + V \le ka^2 + \ell^2 < (k + 1/4)a^2 \le (k + 1/4)/n.$$

Thus $k \geq n$. Also $kx^2 \leq 1$ and $x > \mu$ imply $k < n+4$. Oddness gives exactly $k = n$ or $k = n+2$.

For $k = n$, use Lemma 4.4. It also gives $r \leq \sigma \leq \ell < a/2$, so the saturated candidate is admissible.

For $k = n+2$, put $\gamma = \sqrt{(k-1)/k}$ and $\lambda = 1/\sqrt{k}$. Proposition 3.3 yields, with $\rho = \|b - v_f\|_2$ and $S = \rho^2 - \rho^4/4$,

$$x \leq \lambda\sqrt{1-S} - \gamma\sqrt{S}.$$

Since $x > 0$, this implies $S < \lambda^2$. Moreover $\rho^2 = 2 - 2\sqrt{1-S} \leq 2S$, so $\gamma\rho/2 < \lambda$. Using $\sqrt{1-\rho^2/4} \geq 1 - \rho^2/4$ gives

$$\lambda - x \geq \lambda\rho^2/2 + \gamma\rho\sqrt{1-\rho^2/4} \geq \gamma\rho + \frac{\rho^2}{2}(\lambda - \gamma\rho/2) \geq \gamma\rho.$$

This proves the stated flat estimate. Its leading coefficient $1/\gamma$ is sharp on the interior of a flat phase by the admissible one-low-coordinate equality family in Proposition 3.3.

Every core-n vector has gap at most s, and every core-$(n+2)$ vector has gap at most λ. Therefore a deficit below $|\lambda - s|$ excludes the losing core size. The exact cap formula identifies the remaining candidate as the actual optimizer.

To prove sharpness of C_*, take $n = 1$, let $r > 0$ tend to zero, and set $a = \sqrt{1-r^2}$. Then the unique optimizer is (a, r). For $b = (a, r/\sqrt{2}, r/\sqrt{2})$,

$$x = a - \sqrt{2}r, \quad \Delta = (\sqrt{2}-1)r, \quad \operatorname{dist}(b,(a,r)) = r\sqrt{2-\sqrt{2}} = C_*\Delta.$$

For sufficiently small r, the strict trigger holds, the losing three-flat vector stays a positive distance away, and $\Delta < |\lambda - s|$. Thus the constant is sharp even for distance to the unique global optimizer.

Finally take $a = 1/\sqrt{n+2}$ and let b be $(n+4)$-flat. Its gap is exactly μ, the saturated candidate is inadmissible, and $G = \lambda$. Writing $k = n+2 \geq 3$, the squared ratio of its distance from the k-flat optimizer to its deficit is

$$\frac{2 - 2\sqrt{k/(k+2)}}{(k^{-1/2} - (k+2)^{-1/2})^2} = k(k+2+\sqrt{k(k+2)}) > C_*^2.$$

The ratio is unbounded as $k \to \infty$. This proves the endpoint obstruction and completes the proof. □

Interpretation and limit. For $n \geq 3$, $\mu > a/2$ throughout the band, and the trigger is simply $\Delta < G - \mu$. On the flat branch this allowable deficit is $\lambda - \mu \asymp a^3$. The theorem is a sharp quantitative extension of the all-cap support-gap classification, valid without an unevaluated small-cap cutoff. It does not resolve the remaining probability optimization at thresholds containing positive central mass.

5 The exact local stability modulus on the saturated branch

The optimal linear constant is the endpoint of an exact nonlinear stability modulus. For $r > 0$ and $0 \leq e \leq (\sqrt{2}-1)r$, define

$$F_r(e) = \sqrt{r\left(r - e - \sqrt{r^2 - 2re - e^2}\right)}.$$

Both square roots are nonnegative on this interval. Indeed $(r-e)^2 - (r^2 - 2re - e^2) = 2e^2$ and $e < r$.

Theorem 5.1 (Exact saturated stability modulus)**.** *Let b be a finite normalized nonnegative vector with coefficients at most $a \le 1$, actual gap $x > a/2$, and actual core of odd size k. Suppose $ka^2 < 1$, put*

$$r = \sqrt{1 - ka^2} > 0, \qquad e = a - r - x,$$

and assume $e \le (\sqrt{2} - 1)r$. Then $e \ge 0$ and

$$\operatorname{dist}(b, (a^k, r)) \le F_r(e). \tag{18}$$

Equality holds precisely when the core consists of k copies of a and the tail has the two coefficients

$$t_\pm = \frac{r + e \pm \sqrt{r^2 - 2re - e^2}}{2}, \tag{19}$$

with the zero coefficient omitted at $e = 0$. Conversely these coefficients attain the bound whenever the high-gap condition $a - r - e > a/2$ holds.

In particular $F_r(e)/e$ increases strictly from 1 to $\sqrt{2 + \sqrt{2}}$ as e/r increases from 0 to $\sqrt{2} - 1$.

When the saturated type is the unique cap optimizer and the deficit is below the branch separation, this gives the exact coefficient-distance modulus for near optimizers in the stated deficit range. The two small residuals in (19) are forced at equality; the second residual cannot be replaced by an arbitrary small diffuse tail.

Lemma 5.2 (Exact near-axis tail modulus)**.** *Let t be a finite nonnegative vector with $\sigma = \|t\|_2 > 0$, $\ell = \|t\|_1$, and $q = \ell - \sigma \le (\sqrt{2} - 1)\sigma$. If j is a largest coordinate, then*

$$\|t - \sigma e_j\|_2 \le F_\sigma(q).$$

For $q > 0$, equality holds precisely when at most two coordinates are nonzero. Their values are obtained by replacing (r, e) with (σ, q) in (19). For $q = 0$ the vector is one-sparse.

Proof. Put $M = t_j$. Since $\sigma^2 \le M\ell$ and $\ell \le \sqrt{2}\sigma$, one has $M \ge \ell/2$. Also

$$\sigma^2 - M^2 \le (\ell - M)^2,$$

with equality exactly when at most one coordinate besides j is nonzero. Solving this quadratic inequality on $M \ge \ell/2$ gives

$$M \ge \frac{\ell + \sqrt{2\sigma^2 - \ell^2}}{2}.$$

The squared distance is $2\sigma(\sigma - M)$. Substituting the last bound yields $F_\sigma(q)^2$, with exactly the stated equality cases. □

Lemma 5.3 (Monotonicity of the modulus)**.** *The function $R(\alpha) = F_1(\alpha)/\alpha$ is strictly increasing on $0 < \alpha \le \sqrt{2} - 1$, has limit 1 at zero, and has endpoint value $\sqrt{2 + \sqrt{2}}$. Consequently, whenever $0 \le q \le e \le (\sqrt{2} - 1)r$ and $\sigma \ge r > 0$,*

$$F_\sigma(q) \le F_r(q), \qquad F_r(q) + (e - q) \le F_r(e).$$

If $e > 0$, the second inequality is strict unless $q = e$.

Proof. For $0 < \alpha \le \sqrt{2} - 1$, rationalization gives

$$R(\alpha)^2 = \frac{2}{1 - \alpha + \sqrt{1 - 2\alpha - \alpha^2}}.$$

The denominator decreases strictly from 2 to $2 - \sqrt{2}$, proving the monotonicity and the stated endpoint values. Since $F_\sigma(q) = qR(q/\sigma)$, the first consequence follows. For the second, if $q > 0$,

$$F_r(e) - F_r(q) = (e - q)R(e/r) + q(R(e/r) - R(q/r)) \ge e - q.$$

The same conclusion for $q = 0$ follows from $R(e/r) \ge 1$. If $e > q$ and $e > 0$, then $R(e/r) > 1$, so the inequality is strict. □

Proof of Theorem 5.1. Use the nonnegative quantities q, D, κ and the identity $e = q + D + \kappa$ from Lemma 4.3. Since $r > 0$, the tail norm σ is positive. The exact tail lemma and the modulus lemma yield

$$\operatorname{dist}(b, (a^k, r)) \le F_\sigma(q) + (e - q) \le F_r(q) + (e - q) \le F_r(e).$$

Their hypotheses hold because $q \le e \le (\sqrt{2} - 1)r \le (\sqrt{2} - 1)\sigma$.

If $e > 0$, equality forces $q = e$ by strictness in the modulus lemma. The identity $e = q + D + \kappa$ then gives $D = \kappa = 0$. Every core coefficient equals a, the tail variance is r^2, and equality in the exact tail lemma gives (19). If $e = 0$, the same nonnegative decomposition gives $q = D = \kappa = 0$, so the vector is the saturated optimizer itself. Conversely, the displayed two residuals have sum $r + e$ and squared sum r^2. As k is odd and $r + e < a/2$, the actual gap is $a - r - e$; both residuals lie below this gap. The exact tail calculation gives equality in (18). □

Local expansion. For fixed $r > 0$ and $e/r \to 0$,

$$F_r(e) = e + \frac{e^2}{2r} + O(e^3/r^2).$$

Thus the sharp leading linear coefficient on every fixed saturated branch is 1. The larger universal coefficient $\sqrt{2 + \sqrt{2}}$ is needed when the entire residual splits into two equal pieces, at $e = (\sqrt{2} - 1)r$.

6 Exact positive discrepancy for every cap at least one half

Let $X = \sum_i b_i \varepsilon_i$, $\sum_i b_i^2 = 1$, and $|b_i| \le a$, with finitely many independent uniform signs. Write $c(x) = \Phi(x) - 1/2$, $U_X(x) = \frac{1}{2}\mathbb{P}(|X| \le x) - c(x)$ for $x \ge 0$, and

$$\mathfrak{D}(a) = \sup_{N, b} \sup_{x \ge 0} U_X(x).$$

This section uses two stated analytic inputs. First, the completed Keller–Klein comparison for normalized Rademacher sums is

$$d_K(Z, N(0,1)) \le c\left(\max_i |b_i|\right). \tag{20}$$

Its small-coefficient part is proved in the original package [9]; the passage to every coefficient size uses the external computer-assisted theorem of Ivanisvili [8]. The expanded companion [10] gives

another reconstruction with its own explicitly recorded external inputs. We use (20) as a theorem input here and do not claim to reprove its external certificates.

Second, the published theorem of Hollom and Portier [2] gives the following strict-tail lower bounds:

$$\mathbb{P}(|Z| > t) \geq f(t), \qquad \begin{array}{c|ccccc} t & [0,1/\sqrt{7}) & [1/\sqrt{7},1/\sqrt{5}) & [1/\sqrt{5},1/\sqrt{3}) & [1/\sqrt{3},\sqrt{2/3}) & [\sqrt{2/3},1) \\ \hline f(t) & 1/2 & 29/64 & 3/8 & 1/4 & 7/32. \end{array} \tag{21}$$

Their theorem is stated for non-strict tails. For a fixed t in a listed half-open interval, apply it at $t+\eta$ inside that interval and let $\eta \downarrow 0$. Continuity of probability from below gives exactly (21), including the left endpoints. The external computer-assisted part of their proof is not replayed in this package.

The probability theorem below is thus conditional on these precise external analytic inputs in the verification sense. The deterministic theorems in the previous sections have no such dependencies.

Theorem 6.1. *For $1/2 \leq a \leq 1/\sqrt{2}$, put $r = \sqrt{1-2a^2}$. Then*

$$\mathfrak{D}(a) = \max\{3/16,\ 1/4 - c(r)\}.$$

Let $q > 0$ be the unique solution of $c(q) = 1/16$, and put $a_0 = \sqrt{(1-q^2)/2}$. For $1/2 \leq a < a_0$, the only optimizer is the four-flat vector, at threshold zero. For $a_0 < a \leq 1/\sqrt{2}$, the only optimizer is (a, a, r), at threshold r, omitting a zero residual. At $a = a_0$, precisely both types optimize. These statements are up to permutations, sign changes, and zero padding. For $a \geq 1/\sqrt{2}$, $\mathfrak{D}(a) = 1/4$, with only the equal pair, at threshold zero.

6.1 Elementary localization estimates

We use the scalar inequalities

$$\frac{199}{500} < \varphi(0) < \frac{2}{5}, \quad c(1/\sqrt{7}) > \frac{9}{64}, \quad c(1/\sqrt{3}) > \frac{43}{200}, \quad 1 - \Phi(1) < \frac{7}{40}, \quad q < \frac{4}{25}. \tag{22}$$

Here are elementary proofs. The classical bounds $223/71 < \pi < 22/7$ imply the first pair. Since $e^{-u} \geq 1 - u$ for $u \geq 0$,

$$c(t) > \frac{199}{500}\left(t - \frac{t^3}{6}\right) \quad (0 < t \leq 1).$$

For $t = 1/\sqrt{7}$ use $t > 37/100$ and $t^2 = 1/7$; for $t = 1/\sqrt{3}$ use $t > 577/1000$ and $t^2 = 1/3$. The resulting rational lower bounds exceed $9/64$ and $43/200$. At $t = 1$, the lower bound exceeds $13/40$, proving the tail assertion. Finally

$$c(4/25) > \frac{199}{500}\left(\frac{4}{25} - \frac{(4/25)^3}{6}\right) > \frac{1}{16},$$

so strict monotonicity of c gives $q < 4/25$. All remaining numerical margins below are rational inequalities; the supplementary script `check_two_core_upper.py` checks them using exact fractions.

Lemma 6.2 (The zero threshold)**.** *If at least three coefficients are nonzero, then $\mathbb{P}(X = 0) \leq 3/8$. Equality forces four equal nonzero coefficients.*

Proof. Absorb signs and remove zeros. Condition on all but three coefficients. The signs of three positive coefficients giving any fixed value form an antichain, so have probability at most 3/8. For completeness, a random maximal chain of the three-dimensional Boolean lattice meets an antichain at most once, and each antichain member is met with probability at least 1/3, proving the count bound three.

For three positive weights $u \geq v \geq w$, an atom has multiplicity three only if $u = v = w$, and its value is then $\pm u$. Indeed the extreme values $\pm(u+v+w)$ are singletons; the remaining values are $\pm(u+v-w), \pm(u-v+w), \pm(-u+v+w)$. Three coincide only when the three positive expressions coincide, which gives $u = v = w$; if $-u+v+w \leq 0$, a triple equality would force $v = 0$ or $w = 0$. Thus equality in the conditional bound forces the three selected weights to equal some $b > 0$ and the residual R to be supported on $\{-b, b\}$. Its largest outcome is its coefficient sum $\ell = b$, while $\mathbb{E}R^2 = b^2$. Hence $\ell^2 = \sum_{i \in R} b_i^2$, forcing exactly one nonzero residual coefficient. Normalization gives four coefficients $b = 1/2$. □

Lemma 6.3 (Translated residual intervals)**.** *If a finite Rademacher sum R has variance $V > 0$ and every coefficient has magnitude at most $x > 0$, then every interval of radius x has probability less than $(8/5)x/\sqrt{V}$.*

Proof. After normalization, Keller–Klein bounds the difference of each CDF (and its left limits) from the Gaussian by $c(x/\sqrt{V})$. The Gaussian mass of an interval of radius $x/\sqrt{V}$ is at most $2c(x/\sqrt{V})$. Subtracting the two CDF values therefore gives at most $4c(x/\sqrt{V})$, which is strictly less than $(8/5)x/\sqrt{V}$. □

Lemma 6.4 (Exact two-core radius)**.** *Suppose $a \leq 1/\sqrt{2}$, $x > 0$, and*

$$U_X(x) \geq T := \max\{3/16, 1/4 - c(r)\}, \qquad r = \sqrt{1-2a^2}.$$

Then $x \geq r$ and $\mathbb{P}(|X| \leq x) \leq 1/2$.

Proof. For $1/\sqrt{7} \leq x < 1/\sqrt{3}$, the strict tail is at least 3/8, so $U_X(x) \leq 5/16 - c(1/\sqrt{7}) < 11/64 < 3/16$. For $1/\sqrt{3} \leq x < 1$, it is at least 7/32, giving $U_X(x) \leq 25/64 - c(1/\sqrt{3}) < 3/16$. For $x \geq 1$, $U_X(x) \leq 1 - \Phi(1) < 3/16$. Consequently $x < 1/\sqrt{7}$, where the first tail phase gives $p := \mathbb{P}(|X| \leq x) \leq 1/2$. Thus $U_X(x) \geq T \geq 3/16$ implies $x \leq q < 4/25$.

Also $p = 2U_X(x) + 2c(x) > 3/8$. If three coefficients were greater than x, conditioning on the others and applying the same antichain argument would give $p \leq 3/8$. Hence, after sorting, $b_3 \leq x$. If there are at most two nonzero coefficients, normalization and $a \leq 1/\sqrt{2}$ force the equal pair and $r = 0$, proving the claim. Otherwise put $V = \sum_{i \geq 3} b_i^2 > 0$, $s = b_1 + b_2$, $h = b_1 - b_2$. Lemma 6.3, applied after conditioning the first two signs, gives

$$V < \left(\frac{(8/5)x}{p}\right)^2.$$

The law of X is the equal mixture of the Rademacher laws $h\varepsilon + R$ and $s\varepsilon + R$, with the displayed sign independent of R. The former has variance

$$h^2 + V = 1 - 2b_1b_2 \geq r^2.$$

Assume for contradiction that $x < r$, and put $u = \min\{r, q\}$ and $t = x/u$. Then $0 < t \leq 1$; if $t = 1$, we necessarily have $u = q < r$. Since $T = 1/4 - c(u)$,

$$p \geq \frac{1}{2} - 2(c(u) - c(x)) \geq \frac{93 + 32t}{250} =: P(t).$$

Here $u \le 4/25$ and $c' \le 2/5$ were used. In particular

$$V < v(t) := \left(\frac{64t}{93+32t}\right)^2, \qquad x \le \frac{4t}{25}.$$

The core sum satisfies $s \ge (1-V)/a \ge \sqrt{2}(1-V)$, so

$$s - x > \frac{141}{100}(1-v(t)) - \frac{4t}{25} =: d(t) > 0.$$

The positivity follows from $v(t) \le v(1) = 4096/15625$ and $d(1) = 1375589/1562500 > 0$.

The balanced component has central probability at most $1 - f(t)$, using the appropriate phase above. In the last phase the same bound $1 - 7/32$ applies also when $t = 1$, because its actual normalized threshold is strictly less than one by $x < r$. For the other component, symmetry and one-sided Chebyshev give

$$\mathbb{P}(|s\varepsilon + R| \le x) = \mathbb{P}(s - x \le R \le s + x) \le \frac{V}{2(s-x)^2}.$$

Therefore

$$p < \frac{1-f(t)}{2} + F(t), \qquad F(t) := \frac{v(t)}{4d(t)^2}.$$

The function v increases and d decreases on $[0,1]$, so F increases. In each phase, replace t in F by the rational upper endpoint in the next table, and replace t in P by the rational lower endpoint. Each displayed positive margin is an exact rational lower bound for $P(t) - (1-f(t))/2 - F(t)$.

$(1-f(t))/2$	t lower	t upper	margin greater than
1/4	0	38/100	11/100
35/128	37/100	45/100	13/100
5/16	44/100	58/100	9/100
3/8	57/100	82/100	2/100
25/64	81/100	1	49/100000

The square-root phase endpoints lie inside the listed rational bounds, as is verified by squaring. Thus the upper bound for p is strictly below $P(t)$ in every phase, a contradiction. Hence $x \ge r$. □

Proof of Theorem 6.1. For $a < 1/\sqrt{2}$, there must be at least three nonzero coefficients, so Lemma 6.2 gives $U_X(0) \le 3/16$, with equality only for four-flat. For positive thresholds competitive with T, Lemma 6.4 gives

$$U_X(x) \le 1/4 - c(x) \le 1/4 - c(r) \le T.$$

The endpoint $a = 1/\sqrt{2}$ also permits the equal pair, whose value is $1/4$ at zero. For any larger cap, the same value remains maximal: the localization estimates above give $U_X(x) < 3/16$ for $x \ge 1/\sqrt{7}$, while the first strict-tail phase gives $U_X(x) \le 1/4 - c(x) < 1/4$ for $0 < x < 1/\sqrt{7}$. At zero, conditioning on one nonzero sign gives an atom of mass at most $1/2$. Equality requires exactly two equal nonzero coefficients, because three or more are covered by Lemma 6.2.

Four-flat is admissible for every $a \ge 1/2$ and gives $3/16$. Whenever $1/4 - c(r) \ge 3/16$, one has $r \le q < 4/25 < a$. Thus (a, a, r) is admissible and its central mass at $x = r$ is $1/2$; it gives the other branch. The branch is strictly increasing in a, yielding the unique transition a_0.

It remains to classify positive-threshold equality. It requires $x = r$ and $p = 1/2$. The two-coefficient endpoint is settled, so assume $V > 0$. By the preceding conditioning argument, $b_3 \le x$ and

$$V < \big((8/5)x/(1/2)\big)^2 < 4096/15625.$$

Consequently $b_2^2 = 1 - V - b_1^2 > 1/2 - 4096/15625 > (4/25)^2 > x^2$. For each residual outcome z, an interval of radius x can contain at most two of the four values $-s, -h, h, s$; if it contains two, they must be the central pair $-h, h$, because $s - h = 2b_2 > 2x$. The average conditional central probability equals $1/2$, so each conditional probability equals $1/2$. Therefore $|z| + h \le r$ for every residual outcome. Taking $\ell = \sum_{i\ge 3} b_i$ gives

$$r^2 \le V \le \ell^2 \le (r - h)^2 \le r^2.$$

Equality throughout forces $h = 0$, $b_1 = b_2 = a$, and one nonzero residual coefficient r. Together with Lemma 6.2, this proves all equality and tie cases. □

A general balanced-core observation. The decisive variance inequality extends to every even core size M. Generate the uniform balanced sign slice by a uniform perfect matching, then independently orient each matched pair with opposite signs. Conditional on a matching, the balanced core plus residual is a Rademacher sum of variance

$$V + \sum_{\{i,j\}\ \text{matched}} (b_i - b_j)^2 = 1 - 2 \sum_{\{i,j\}\ \text{matched}} b_i b_j \ge 1 - Ma^2.$$

Thus the same Hollom–Portier bound applies to every matching. For $M \ge 4$, the missing global step is threshold localization and control of the other core layers. No all-cap upper formula is inferred from this observation alone.

7 Lower discrepancy: two finite certificates

Write $Q(t) = 1 - \Phi(t)$ and

$$L_X(x) = c(x) - \tfrac{1}{2}\mathbb{P}(|X| < x) = \mathbb{P}(X \ge x) - Q(x), \qquad x > 0.$$

The strict central interval is essential. Besides the Keller–Klein comparison and the Hollom–Portier profile f in (21), this section uses Tomaszewski's theorem $\mathbb{P}(|Z| \le 1) \ge 1/2$ for a normalized Rademacher sum [1]. All finite certificates below concern explicit scalar inequalities; their use of validated interval arithmetic is separate from Lean.

Theorem 7.1 (The five-flat lower plateau)**.** *If* $\max_i |b_i| \le 1/2$*, then*

$$\sup_{x>0} L_{X_b}(x) \le c(1/\sqrt{5}).$$

Equality holds only for the five-flat vector, at $x = 1/\sqrt{5}$*. Consequently the lower-discrepancy cap optimum equals* $c(1/\sqrt{5})$ *for every* $a \in [1/\sqrt{5}, 1/2]$*.*

Theorem 7.2 (Exact lower discrepancy in the first missing band)**.** *For* $1/\sqrt{5} \le a \le 1/\sqrt{3}$*, put*

$$G_3(a) = \max\{1/\sqrt{5}, a - \sqrt{1 - 3a^2}\} = \mathfrak{G}(a).$$

Then

$$\sup_{\substack{N<\infty,\ \sum_i b_i^2=1\\ |b_i|\le a}} \sup_{x>0} L_{X_b}(x) = c(G_3(a)). \tag{23}$$

Equality occurs precisely at $x = G_3(a)$ for a maximizing support-gap vector: five-flat on the flat branch, $(a, a, a, \sqrt{1-3a^2})$ on the saturated branch, and both at the unique transition. Together with the earlier high-cap result [10], the identity $\mathfrak{L}(a) = c(\mathfrak{G}(a))$ and this equivalence of lower- discrepancy and support-gap optimizers now hold for every $a \ge 1/\sqrt{5}$.

The positive-discrepancy theorem consequently gives

$$\mathfrak{K}(a) = \max\{3/16, c(a - \sqrt{1-3a^2})\}, \qquad 1/2 \le a \le 1/\sqrt{3}. \tag{24}$$

Put $h = c^{-1}(3/16)$ and $a_3 = (h + \sqrt{4-3h^2})/4$. Below a_3 the unique maximizing type is four-flat, at positive-discrepancy threshold zero; above a_3 it is $(a, a, a, \sqrt{1-3a^2})$, at lower-discrepancy threshold $a - \sqrt{1-3a^2}$. At a_3 precisely both types maximize. Permutations, coefficient sign changes and zero padding are immaterial.

7.1 The four-shift envelope

Absorb coefficient signs and order $m = b_1 \ge b_2 \ge b_3 \ge \cdots$. Write $R = \sum_{i\ge 4} b_i\varepsilon_i$ and $\sigma^2 = 1 - m^2 - b_2^2 - b_3^2 > 0$. Set

$$\rho = \min\{1, b_3/\sigma\}, \quad C = c(\rho), \qquad (q_0, q_1, q_2, q_3) = (m+b_2+b_3,\ |b_2+b_3-m|,\ m-b_2+b_3,\ m+b_2-b_3).$$

The distribution of X is the uniform mixture of the four symmetric laws $q_j\varepsilon + R$. For $t > 0$ define

$$T_C(t) = \min\left\{\frac{1}{2}, \frac{1}{2t^2}, e^{-t^2/2}, Q(t) + C, \quad \frac{1}{4} \text{ if } t > 1\right\}. \tag{25}$$

An inapplicable last entry is omitted. For $u \in \mathbb{R}$ set

$$A_C(u) = \begin{cases} \max\{1 - T_C(-u), Q(u) - C, 0\}, & u < 0, \\ \max\{f(u)/2, Q(u) - C, 0\}, & u \ge 0, \end{cases} \tag{26}$$

where $f(u) = 0$ for $u \ge 1$ and otherwise has the stated strict-tail values, with the smaller value at every jump. For each j, put

$$u_j = (q_j - x)/\sigma, \quad v_j = (q_j + x)/\sigma, \qquad H_j = \max\{0, A_C(u_j) - T_C(v_j), Q(u_j) - Q(v_j) - 2C\}.$$

Additional valid lower bounds on $\mathbb{P}(R/\sigma > u_j)$ may be included in the maximum defining $A_C(u_j)$. We use only the explicit additional bounds described below. The resulting scalar envelope is

$$E(m, b_2, b_3, x) = \min\left\{c(m),\ c(x) - \frac{1}{8}\sum_{j=0}^{3} H_j,\ \frac{1}{2x^2} - Q(x)\right\}. \tag{27}$$

Lemma 7.3. *With these definitions, $L_X(x) \le E(m, b_2, b_3, x)$ for every $x > 0$.*

Proof. For the normalized residual $Z = R/\sigma$, symmetry and Chebyshev give

$$\mathbb{P}(Z \ge t) \le \min\{1/2, 1/(2t^2)\}.$$

The exponential bound follows from $\mathbb{E}e^{sZ} = \prod_i \cosh(sb_i/\sigma) \le e^{s^2/2}$ and Markov, optimized at $s = t$. Keller–Klein, including CDF left limits, gives $\mathbb{P}(Z \ge t) \le Q(t) + C$ and $\mathbb{P}(Z > u) \ge Q(u) - C$. For $t > 1$, Tomaszewski gives $\mathbb{P}(Z \ge t) \le \mathbb{P}(Z > 1) \le 1/4$. These prove (25). If $u < 0$, symmetry gives $\mathbb{P}(Z > u) = 1 - \mathbb{P}(Z \ge -u)$; if $u \ge 0$, Hollom–Portier gives $\mathbb{P}(Z > u) = \frac{1}{2}\mathbb{P}(|Z| > u) \ge f(u)/2$. Thus (26) is valid, with the stated endpoints. Finally,

$$\mathbb{P}(|q_j\varepsilon + R| < x) = \mathbb{P}(q_j - x < R < q_j + x) = \mathbb{P}(Z > u_j) - \mathbb{P}(Z \ge v_j) \ge H_j.$$

Averaging the four laws proves the middle bound in (27); the other two are the global Keller–Klein bound and one-sided Chebyshev for X. □

7.2 Certified scalar inequalities

The following is the exact interface between the analysis and the computation. Every box coordinate and every bisection is rational. Only the evaluation of square roots, exponentials and Gaussian integrals uses validated real intervals.

Lemma 7.4 (Finite scalar certificates)**.** *The envelope (27) has the following strict bounds, on the full indicated domains.*

1. *If $1/2 \le m \le 1/\sqrt{3}$, $0 \le b_3 \le b_2 \le m$, $447/1000 \le x \le 9/5$, $\sigma > 0$, and $x > G_3(m)$, then $E < c(G_3(m))$, where the extra bounds $A_C(u_1), A_C(u_2) \ge 7/64$ are included, and $A_C(u_3) \ge 7/64$ is included whenever $b_3 \ge \sqrt{1 - 3m^2}$.*
2. *If $1/\sqrt{5} \le m \le 1/2$, $0 \le b_3 \le b_2 \le m$, $447/1000 \le x \le 9/5$, and $q_1 < x$, then $E < c(1/\sqrt{5})$, where the extra bound $A_C(u_1) \ge 1/2$ is included.*
3. *If $1/\sqrt{5} \le m \le 1/2$, $0 \le b_3 \le b_2 \le m/8$, and $447/1000 \le x \le 9/5$, then $E < c(1/\sqrt{5})$ without extra bounds.*

Proof and reproducibility of the scalar certificates. The respective partitions have 13666, 14303 and 3 leaves. They are supplied, compressed with `xz`, as `lower_exact_gap_pilot.json.xz`, `lower_five_flat_pilot_B.json.xz`, and `lower_five_flat_pilot_C.json.xz`. Their names record their origin in a nonrigorous partition search; none of the floating point values in those files is trusted by the verifiers. The verifiers read only the finite bisection paths, reconstruct every endpoint with rational arithmetic, verify complete coverage, and then prove a strictly positive interval margin on every leaf.

Here are complete details of the enclosure algorithm. Parameterize $w = m^2$, $b_2 = m\alpha$, $b_3 = m\alpha\beta$. In case (i),

$$(w, \alpha, \beta, x) \in [1/4, 1/3] \times [0, 1]^2 \times [447/1000, 9/5];$$

in cases (ii),(iii), the w interval is $[1/5, 1/4]$, the x interval is $[447/1000, 9/5]$, and the α interval is respectively $[0, 1]$ or $[0, 1/8]$. A digit $2k + e$ in a leaf path means bisect coordinate $k \in \{0, 1, 2, 3\}$ and choose half $e \in \{0, 1\}$. The exact prefix-tree check requires that every nonleaf has precisely the two children $2k, 2k + 1$ for one common k; it also checks absence of duplicate or prefix leaves and exact total normalized volume one. Thus no region can be omitted.

For a box write subscripts $-$ and $+$ for its coordinate endpoints, and $m_\pm = \sqrt{w_\pm}$. Valid residual bounds are

$$V_- = 1 - w_+(1 + \alpha_+^2(1 + \beta_+^2)), \qquad V_+ = 1 - w_-(1 + \alpha_-^2(1 + \beta_-^2)).$$

Both are nonnegative on the specified domains, and $V_+ > 0$. Set $s_\pm = \sqrt{V_\pm}$ and $C_+ = c(\min\{1, m_+\alpha_+\beta_+/s_-\})$, taking $C_+ = c(1)$ if $s_- = 0$. The elementary interval bounds for the four shifts are

j	$q_{j,-}$	$q_{j,+}$
0	$m_-(1+\alpha_-(1+\beta_-))$	$m_+(1+\alpha_+(1+\beta_+))$
2	$m_-(1-\alpha_+(1-\beta_-))$	$m_+(1-\alpha_-(1-\beta_+))$
3	$m_-(1+\alpha_-(1-\beta_+))$	$m_+(1+\alpha_+(1-\beta_-))$.

For $j = 1$, take the minimum and maximum of the absolute value on $[\alpha_-(1+\beta_-)-1, \alpha_+(1+\beta_+)-1]$, then multiply by m_- and m_+ respectively. In case (i) replace x_- by $\max\{x_-, 1/\sqrt{5}, m_- - \sqrt{1-3w_-}\}$; this only restricts to a larger set than the required condition $x > G_3(m)$. The target is bounded below by $c(G_3(m_-))$ using monotonicity.

An upper bound for u_j is $(q_{j,+} - x_-)/s_-$ when its numerator is nonnegative, and $(q_{j,+} - x_-)/s_+$ when negative; use $+\infty$ if the first case has $s_- = 0$. A lower bound for v_j is $(q_{j,-} + x_-)/s_+ > 0$. Apply the formulas (25)–(27) to these endpoints, using C_+, $c(m_+)$, $c(x_+)$ and $1/(2x_-^2) - Q(x_+)$. Monotonicity gives an upper bound for E: T_C decreases with its argument and increases with C, while A_C decreases with its argument and with C. For extra bounds, cases (i),(ii) use exactly their stipulated entries; the optional $j = 3$ entry in (i) is used only when the exact rational comparison $w_-\alpha_-^2\beta_-^2 \geq 1 - 3w_-$ holds. If an interval comparison cannot determine a sign, the verifier uses a weaker bound, never the unproved branch. At a phase boundary it uses only a phase whose endpoint is proved strictly larger. The exceptional forced 7/64 entries are justified by a strict analytic inequality below, not by rounding.

The programs `verify_lower_exact_gap_arb.py` and `verify_lower_five_flat_arb.py` implement these formulas with `Fraction` for exact rational data and `python-flint` 0.9.0 (Arb) for real enclosures [11, 12]. A floating midpoint is used only to choose one among already valid upper or lower bounds; correctness never relies on that choice selecting the best bound. The acceptance test is an interval comparison proving the entire margin strictly positive. All 27972 leaves pass at both 128 and 192 bits. The supplied JSON reports contain each leaf margin, partition and verifier SHA-256 hashes, coverage checks, precision and execution metadata. The commands are

```
python verify_lower_exact_gap_arb.py 192
python verify_lower_five_flat_arb.py B 192
python verify_lower_five_flat_arb.py C 192
```

This proves the stated scalar inequalities with validated arithmetic. It is not a Lean kernel check or a replay of the external Hollom–Portier computation. □

7.3 Proofs of the probability conclusions

Proof of Theorem 7.2. The five-flat plateau covers actual maxima $m \leq 1/2$; its only equality is five-flat at $x = 1/\sqrt{5}$. Assume $1/2 < m \leq a \leq 1/\sqrt{3}$. For $x \leq G_3(m)$, trivially $L_X(x) \leq c(x) \leq c(G_3(m))$. We show that $x > G_3(m)$ gives strict inequality. If $x \geq 9/5$, one-sided Chebyshev gives $L_X(x) \leq 25/162 < c(1/\sqrt{5}) \leq c(G_3(m))$. Otherwise $x > G_3(m) \geq 1/\sqrt{5} > 447/1000$, so the variables lie in the domain of Lemma 7.4(i). If $\sigma = 0$, normalization and $m \leq 1/\sqrt{3}$ force the three-flat vector. For $x > G_3(m) = 1/\sqrt{3}$, its strict central probability is at least 3/4, hence $L_X(x) \leq 1/8 < c(1/\sqrt{5})$.

To justify the additional tail entries in (i), set $r_3 = \sqrt{1-3m^2}$. Since $q_1, q_2 \leq m$ and $\sigma \geq r_3$, the strict inequality $x > G_3(m) \geq m - r_3$ gives $q_1 - x < \sigma$ and $q_2 - x < \sigma$. Thus $u_1, u_2 < 1$;

Hollom–Portier gives $A_C(u_1), A_C(u_2) \geq 7/64$ if these arguments are nonnegative, and symmetry gives at least $1/2$ otherwise. For $j = 3$, let $d = b_2 - b_3 \geq 0$. If $b_3 \geq r_3$, then

$$\sigma^2 - r_3^2 = 2m^2 - b_2^2 - b_3^2 \geq d^2 + 2b_3 d \geq d^2 + 2r_3 d.$$

The first inequality follows by subtraction, leaving $2m^2 - 2b_2^2 \geq 0$. Thus $\sigma \geq r_3 + d$, and $q_3 - x = m + d - x < r_3 + d \leq \sigma$. This proves the optional last entry. The scalar certificate and Lemma 7.3 now give the required strict bound.

The function G_3 is nondecreasing. Therefore every admissible law has lower discrepancy at most $c(G_3(a))$, and equality must have $x = G_3(a)$ and an empty strict central interval. Equivalently, $g(b) = \mathfrak{G}(a)$, so the support-gap theorem proves both attainment and the complete optimizer classification. This establishes (23). The earlier theorem for $a \geq 1/\sqrt{3}$ gives the claimed combined range.

On $[1/2, 1/\sqrt{3}]$, the positive optimum is $3/16$ and $c(1/\sqrt{5}) < 3/16$. Taking the maximum of the positive and lower optima proves (24), with the equality cases inherited from their separate classifications. Solving $a - \sqrt{1 - 3a^2} = h$ gives the displayed transition a_3. □

Proof of Theorem 7.1. Write $\lambda = 1/\sqrt{5}$ and $H = c(\lambda)$. If $m < \lambda$, Keller–Klein gives $L_X \leq c(m) < H$. If $x < \lambda$, the trivial bound $L_X(x) \leq c(x)$ is strict. At $x = \lambda$, equality means $\mathbb{P}(|X| < \lambda) = 0$, hence $g(b) \geq \lambda$. The exact support-gap theorem gives $\mathfrak{G}(1/2) = \lambda$ with only five-flat attaining it, so this case has precisely the asserted equality type. If $x \geq 9/5$, then $L_X(x) \leq 25/162 < H$, a scalar comparison in the certificates. We may therefore assume $\lambda \leq m \leq 1/2$ and $\lambda < x < 9/5$. Here $\sigma^2 \geq 1 - 3m^2 \geq 1/4$, so the four-shift envelope is always defined.

If $q_1 < x$, symmetry gives $\mathbb{P}(Z > u_1) \geq \mathbb{P}(Z \geq 0) \geq 1/2$. Lemma 7.4(ii) gives $L_X(x) < H$. If $q_1 \geq x$ and $b_2 + b_3 < m$, then $b_2 + b_3 \leq m - x \leq m - \lambda$ and hence

$$b_2/m \leq 1 - \lambda/m \leq 1 - 2\lambda < 1/8.$$

Lemma 7.4(iii) again gives $L_X(x) < H$.

It remains to treat the triangle case $b_2 + b_3 \geq m$ and $q_1 \geq x$. Then $b_2 + b_3 \geq m + x$, so

$$x \leq q_1 \leq q_2 \leq q_3, \qquad x \leq m \leq 1/2, \qquad q_3 - x \leq 2(m - x) \leq 1 - 2\lambda < \sigma/\sqrt{7}.$$

The last comparison follows from $\sigma \geq 1/2$ and $2 - 4\lambda < 1/\sqrt{7}$, also verified in the certificates. For $j = 1, 2, 3$, the first Hollom–Portier phase and one-sided Chebyshev therefore give

$$\mathbb{P}(q_j - x < R < q_j + x) \geq \frac{1}{4} - \frac{\sigma^2}{2(q_j + x)^2} \geq \frac{1}{4} - \frac{\sigma^2}{8x^2}.$$

Also $b_2^2 + b_3^2 \geq (b_2 + b_3)^2/2 \geq (m + x)^2/2$, whence $\sigma^2 \leq 1 - m^2 - (m + x)^2/2 \leq 1 - 3x^2$, with equality in the last inequality only if $m = x$. Therefore

$$L_X(x) \leq c(x) - \frac{3}{32} + \frac{3\sigma^2}{64x^2} \leq c(x) - \frac{15}{64} + \frac{3}{64x^2} =: J(x).$$

On $[\lambda, 1/2]$, $J'(x) = \varphi(x) - 3/(32x^3) < 2/5 - 3/4 < 0$, while $J(\lambda) = c(\lambda)$ exactly. Thus $L_X(x) < H$ when $x > \lambda$. This exhausts every remaining case and proves both the bound and its equality assertion. □

8 The exact full discrepancy down to the five-flat cap

The preceding lower certificate combines with a four-coordinate localization argument to determine the full discrepancy throughout the next cap interval. No additional numerical optimization is used here.

Theorem 8.1. *Let $\lambda = 1/\sqrt{5}$, $H = c(\lambda)$, and $r_4(a) = \sqrt{1-4a^2}$. Then*

$$\mathfrak{K}(a) = \max\{H,\, 3/16 - c(r_4(a))\}, \qquad \lambda \le a \le 1/2. \tag{28}$$

Put $q_4 = c^{-1}(3/16 - H)$ and $a_4 = \frac{1}{2}\sqrt{1-q_4^2}$. For $\lambda \le a < a_4$, the unique maximizing type is five-flat, at lower-discrepancy threshold λ. For $a_4 < a \le 1/2$, it is $(a,a,a,a,r_4(a))$, at positive-discrepancy threshold $r_4(a)$. At a_4, precisely both types maximize. At $a = 1/2$ the latter vector is four-flat and its threshold is zero.

Lemma 8.2 (Four-coordinate localization)**.** *Suppose $\lambda \le a \le 1/2$, $x > 0$, and $U_X(x) \ge H$. Then exactly four coefficients are greater than x, $\mathbb{P}(|X| \le x) \le 3/8$, and $x \ge r_4(a)$.*

Proof. The scalar estimates needed below are

$$H > 11/64 > 17/100, \quad c(1/\sqrt{7}) > 9/64, \quad c(1/\sqrt{3}) > 43/200, \quad c(4/5) > 7/25, \quad Q(1) < 17/100,$$

$$c(1/5) > 5/64, \qquad c(1/25) > 1/64, \qquad c(x) > 39x/100 \quad (0 < x \le 1/5).$$

They follow from $199/500 < \varphi(0) < 2/5$, $e^{-t^2/2} \ge 1 - t^2/2$, the elementary square-root bounds $1/\sqrt{5} > 447/1000$, $1/\sqrt{7} > 377/1000$, $1/\sqrt{3} > 577/1000$, and the exact rational comparisons in `check_four_core_upper.py`. For example, $H > (199/500)(29/30)(447/1000) > 11/64$, and on $[0, 1/5]$ we have $\varphi(x) > (199/500)(49/50) > 39/100$.

The strict-tail profile first forces $x < 1/\sqrt{7}$. Indeed, in $[1/\sqrt{7}, 1/\sqrt{3})$ it gives $U_X(x) \le 5/16 - c(1/\sqrt{7}) < 11/64 < H$; in $[1/\sqrt{3}, \sqrt{2/3})$ it gives $U_X(x) \le 3/8 - c(1/\sqrt{3}) < 4/25 < H$; and in $[\sqrt{2/3}, 1)$ it gives $U_X(x) \le 25/64 - c(4/5) < 17/100 < H$. Finally $x \ge 1$ gives $U_X(x) \le Q(1) < H$. For $x < 1/\sqrt{7}$ the first tail phase gives $p := \mathbb{P}(|X| \le x) \le 1/2$. Consequently $c(x) \le 1/4 - H < 5/64$, so $x < 1/5$, and

$$p = 2U_X(x) + 2c(x) > 17/50 + (39/50)x. \tag{29}$$

If at most two coefficients exceed x, condition on the first two signs. The remaining variance is at least $1 - 2a^2 \ge 1/2$, and all remaining coefficients are at most x. The translated interval bound gives $p < (8\sqrt{2}/5)x < (12/5)x$. This contradicts (29), since $(12/5 - 39/50)(1/5) = 81/250 < 17/50$. Thus at least three coefficients exceed x. The sign patterns of any three such coefficients falling in an interval of length $2x$ form an antichain, so $p \le 3/8$ by Sperner's bound. It follows that $c(x) \le 3/16 - H < 1/64$, hence $x < 1/25$.

Exactly three coefficients cannot exceed x: conditioning on them leaves variance at least $1 - 3a^2 \ge 1/4$, and the translated interval bound gives $p < (16/5)x < 16/125 < 17/50$. Nor can five coefficients exceed x: conditioning on any five gives the antichain bound $p \le \binom{5}{2}/2^5 = 5/16 < 17/50$. Thus exactly four exceed x.

Sort them as $b_1 \ge b_2 \ge b_3 \ge b_4 > x$, and let $R = \sum_{i \ge 5} b_i \varepsilon_i$ have variance V. If $V = 0$, then normalization and $a \le 1/2$ force four-flat, and $r_4(a) = 0$, proving the radius assertion. Otherwise, the translated interval bound and $p > 17/50$, $x < 1/25$ give

$$V < \left(\frac{(8/5)(1/25)}{17/50}\right)^2 = \left(\frac{16}{85}\right)^2 < 9/250. \tag{30}$$

Suppose $x < r_4(a)$. The six balanced sign patterns have total weight $6/16$. Their uniform law can be generated by choosing one of the three perfect matchings of four indices uniformly, and independently orienting each matched pair with opposite signs. Every balanced sign pattern is then counted twice among the twelve equally likely outcomes. Conditional on a matching, the balanced core plus R is a Rademacher sum of variance

$$V + \sum_{\{i,j\} \text{ matched}} (b_i - b_j)^2 = 1 - 2 \sum_{\{i,j\} \text{ matched}} b_i b_j \geq 1 - 4a^2 = r_4(a)^2.$$

Its normalized threshold is strictly below one. Thus the last Hollom–Portier phase, or any stronger preceding phase, bounds its central probability by $25/32$.

For every nonbalanced sign pattern the core magnitude is at least $S - 2a$, where $S = \sum_{i=1}^4 b_i$. Since $S \geq (1 - V)/a$, we have

$$S - 2a \geq 2a - V/a > \frac{894}{1000} - \frac{9/250}{447/1000} > \frac{813}{1000}.$$

Here $a \leq 1/2$ gives $1/a - 2a \geq 2a$, and $a \geq \lambda > 447/1000$. For a nonbalanced magnitude s, symmetry and one-sided Chebyshev therefore give

$$\mathbb{P}(|s\varepsilon + R| \leq x) = \mathbb{P}(s - x \leq R \leq s + x) \leq \frac{V}{2(s-x)^2} < \frac{9/250}{2(773/1000)^2} < \frac{31}{1000}.$$

Combining the six balanced and ten nonbalanced patterns yields

$$p < \frac{6}{16}\frac{25}{32} + \frac{10}{16}\frac{31}{1000} < \frac{17}{50},$$

contrary to (29). Hence $x \geq r_4(a)$. □

Proof of Theorem 8.1. The lower discrepancy is at most H by Theorem 7.1, with only five-flat attaining it. For positive x competitive with H, the localization lemma gives

$$U_X(x) \leq 3/16 - c(x) \leq 3/16 - c(r_4(a)).$$

At zero, if $a < 1/2$ there are at least five nonzero coefficients; the five-sign atom bound gives $U_X(0) \leq 5/32 < H$. If $a = 1/2$, then $U_X(0) \leq 3/16$, with equality only for four-flat. This proves the upper bound in (28).

Five-flat is admissible throughout the interval and attains H. The other branch can compete only when $c(r_4(a)) \leq 3/16 - H < 1/64$, which implies $r_4(a) < 1/25 < a$. The vector $(a, a, a, a, r_4(a))$ is therefore admissible; precisely its six balanced core patterns fall in the closed central interval of radius $r_4(a)$, giving $U = 3/16 - c(r_4(a))$. This branch increases strictly with a, so the displayed transition is unique.

For positive-threshold equality on that branch, the localization lemma forces $x = r_4(a)$, exactly four coefficients greater than x, and $p = 6/16$. For each residual outcome, the qualifying four-sign patterns form an antichain of size at most six. An antichain of size six in the four-dimensional Boolean lattice is exactly the middle layer: the maximal-chain identity $\sum_{A \in \mathcal{A}} \binom{4}{|A|}^{-1} \leq 1$ forces every member to have size two. Since the average conditional probability is $6/16$, every residual outcome must include all six balanced patterns. Their extreme values are $\pm d$, where $d = b_1 + b_2 - b_3 - b_4 \geq 0$. Writing $\ell = \sum_{i \geq 5} b_i$ gives $\ell + d \leq r_4(a)$, and hence

$$r_4(a)^2 \leq V \leq \ell^2 \leq (r_4(a) - d)^2 \leq r_4(a)^2.$$

All inequalities are equalities. The core energy is $4a^2$, so all four core coefficients equal a; the tail satisfies $\ell^2 = V$, so it has one nonzero coefficient, of magnitude $r_4(a)$. The zero-threshold endpoint and the lower equality classification already proved give exactly the stated optimizer and tie cases. □

9 Verification scope

The written mathematical proofs, external theorem inputs, exact-arithmetic checks, and Lean kernel checks are distinct parts of the verification record. Successful compilation of a selected module does not certify unformalized analytic statements elsewhere in the paper.

Kernel-checked scope. The inherited support-gap development proves the exact all-dimensional value, actual attainment, and optimizer classification from finite signed sums. The new module `GlobalGapRigidity.lean` proves the sharp global inverse and its equality classification on the original coefficient indices. It derives the actual core from the given vector; the majority structure, normalized variance identities, and finite-slice inequality are proved rather than assumed. The wrappers `signed_supportGap_inverse` and `signed_supportGap_inverse_equality_iff` accept arbitrary signed coefficients and their actual minimum absolute signed sum.

The wrapper in `ActualGlobalGapStability.lean` derives the odd cap band and actual core, constructs a normalized comparison vector obeying the cap, proves the sharp $\sqrt{2+\sqrt{2}}$ distance bound, and proves that the comparison is an actual optimizer under the stated phase separation or at a tie. The comparison is given on the original indices with one zero coordinate added when necessary. This is a concrete finite-dimensional distance statement; sorting magnitudes gives the distance convention used here.

`FixedFlatTradeoff.lean` proves both displayed fixed-flat inequalities, including their exact radical form, from the actual input vector. The equality-shape analysis of Proposition 3.3 is supplied by its written proof. `ExactSaturatedModulus.lean` proves the exact nonlinear saturated modulus and its complete equality equivalence: every core coordinate equals a, the largest tail is t_+, and the remaining tail has at most one nonzero coefficient with sum t_-. The module `ActualSaturatedModulus.lean` joins the bound to arbitrary signed input, the actual support gap, zero padding, and global phase selection. The separate `ExactModulusAttainment.lean` realizes the entire two-residual equality family for every odd core size and every feasible local deficit, including both endpoints. It proves the actual gap, cap, normalization, core size, and distance from explicit coefficients.

The examples in `GapRigidityExamples.lean` prove the actual saturated equality family and the normalized endpoint counterexample $(2,1,1,1)/\sqrt{7}$, including its signed minimum and even core. The separate module `SharpStabilityExamples.lean` proves the optimal-constant witness (a,t,t), with $a^2+2t^2=1$ and $0<t\le a/10$, strictly inside the unique saturated phase. It verifies the actual cap optimum, actual gap, distance equality, and a finite-dimensional lower bound against every saturated comparison, including after zero padding. Exact feasible rational parameters are supplied, so this family is not an empty conditional implication.

The new modules compile with Lean 4.19.0, `-trust=0`, and kernel type checking enabled. The printed dependency lists for the stated results contain only the standard logical axioms `propext`, `Classical.choice`, and `Quot.sound`. The accompanying fresh replay record supplies the exact source hashes, dependencies, and checked declarations. The nine inherited support-gap modules imported from the unpublished companion [10] and the 62 modules of the original package [9] are

included and recompiled in the same replay. No claim of a kernel-checked entire manuscript is made.

The separate `FourShiftLaw.lean` and `FourShiftEnvelope.lean` derive the actual three-head/four-shift mixture, the strict-interval identity as open tail minus closed tail, residual normalization and cap, and the scalar max/min envelope. Their Gaussian comparisons, strict-tail profile, exponential estimate, above-one tail estimate, and additional open-tail bounds are explicit named hypotheses. This checks the finite-law interface conditionally; the full probability optimization and validated interval certificates remain outside Lean.

Computational scope. The finite phase comparisons in Sections 6 and 8 are reproduced with Python's exact rational arithmetic. Section 7 gives the full interface and enclosure algorithm for the validated real interval certificates. Their acceptance proves the explicit scalar inequalities on complete rational-box partitions; it does not prove the external Keller–Klein, Hollom–Portier, or Tomaszewski inputs. These analytic probability reductions and interval computations are not formalized end to end in Lean. Floating enumeration and randomized searches were used to falsify candidate statements during research; they are diagnostics and supply no proof of the theorems.

10 Remaining probability questions

The support-gap problem is fully optimized for every cap, and its near-maximizers now have the sharp quantitative structure proved above. The full distributional problem requires an additional step: at a threshold containing positive central mass, its probability cost must be compared with the increase in Gaussian mass. No deterministic support-gap theorem alone supplies that comparison.

The full Kolmogorov optimum is now determined for every $a \geq 1/\sqrt{5}$. The small-cap theorem in the companion has an ineffective cutoff, leaving a genuine intermediate range below $1/\sqrt{5}$. The balanced matching construction used in the four-core proof extends to larger even cores, but localization and control of the remaining core layers become additional obligations. The current finite certificates establish exactly their stated domains; they are not evidence of a proved all-cap probability formula.

The exact flat modulus is attained under the stated cap constraints; at a flat endpoint $a = 1/\sqrt{k}$, its nonflat equality family is inadmissible. The saturated modulus is exact in the stated local range $e \leq (\sqrt{2}-1)r$. No assertion is made that these separate curves already solve the full coefficient-distance optimization at every cap and every deficit.

Research process and AI disclosure

This is fully autonomous agentic research. Dev D. Goyal is responsible for publication. No expert validation or external peer review is claimed, and the novelty assessment is provisional.

The results, their proofs, the Lean development, the interval certificates and the literature comparison were produced by autonomous AI research agents in an earlier session. Its records describe an autonomous GPT-6 Astra Pro session with parallel research agents. The proof audits, counterexample searches and source comparisons recorded in that session are internal AI-assisted reviews; they are not included here.

Opus 5.5 by Anthropic used with Claude Code assessed the package for this release. It checked the written proofs, compared the Lean statements with the theorems, ran fresh Lean replays and axiom audits, reran the exact rational checks and the interval certificates at 128 and 192 bits, and

revised this manuscript: it removed references to files not included here, stated the dependence on the unpublished companion [10], and gathered the disclosure into this section. AI checks are not human peer review.

References

[1] N. Keller and O. Klein. *Proof of Tomaszewski's conjecture on randomly signed sums*. Advances in Mathematics 407 (2022), 108558. arXiv:2006.16834.

[2] L. Hollom and J. Portier. *Tight anti-concentration of Rademacher sums*. Random Structures & Algorithms 67 (2025), e70024. doi:10.1002/rsa.70024.

[3] A. De, I. Diakonikolas, and R. A. Servedio. *A robust Khintchine inequality, and algorithms for computing optimal constants in Fourier analysis and high-dimensional geometry*. SIAM Journal on Discrete Mathematics 30(2) (2016), 1058–1094. doi:10.1137/130919143.

[4] X. Fang and T. Wang. *A two-regime Khintchine inequality and an improved bound on the degree-1 Fourier weight for linear threshold functions*. Preprint, 28 August 2026. arXiv:2608.27908.

[5] X. Xie. *Stability of Szarek's inequality with best constant*. Preprint, 5 October 2026. arXiv:2610.07526.

[6] I. Pinelis. *On the supremum of the tails of normalized sums of independent Rademacher random variables*. Statistics & Probability Letters 99 (2015), 131–134. arXiv:1204.1761.

[7] L. Bell. *Small Lyapunov ratios and the structure of Berry–Esseen maximizers*. Preprint, 6 October 2026, in the independent Berry–Esseen bounds collection. `https://github.com/epicycloids/independent-berry-esseen-bounds`. Used for comparison only, not as a proof input.

[8] P. Ivanisvili. *The Keller–Klein Berry–Esseen conjecture: the reachable half for all Rademacher sums, and the full conjecture for $a_1 \geq 0.07$*. Computer-assisted AI-generated manuscript, 21 October 2026. `https://paatai.github.io/ai-proofs/keller-klein/keller-klein.pdf`.

[9] D. D. Goyal. *The small-coefficient upper half of the Keller–Klein conjecture: analytic completion and Lean 4 formalization*. Autonomous AI research package, 30 October 2026. `https://github.com/dvgyl/math-and-theory/tree/main/papers/mathematics/keller-klein`.

[10] D. D. Goyal. *Exact extremizers and sharp stability in Rademacher normal approximation*. Unpublished expanded research package preceding the present extension, 1 November 2026. It is not part of this repository; the nine Lean modules of it that the present development imports are included with this manuscript.

[11] F. Johansson. *Arb: efficient arbitrary-precision midpoint-radius interval arithmetic*. IEEE Transactions on Computers 66(8) (2017), 1281–1292. doi:10.1109/TC.2017.2690633.

[12] The python-flint developers. *python-flint documentation: real ball arithmetic*. `https://python-flint.readthedocs.io/en/latest/arb.html`. Version 0.9.0 used for the accompanying interval certificates.