CONJURA · SOLUTION

Signed Low-Entropy Reweightings to Rank One on the Sphere

Resolves. Conjecture 1 (see `main.tex` in this folder) = Question 8.1 of Barak, Kothari and Steurer [BKS16] — proven true, and with a strictly stronger conclusion than the conjecture asks for: exact rank one (Frobenius error 0, not merely $\le \varepsilon\|L\|_F$), at a constant $C(\varepsilon,\delta) = 2/(e\delta)$ that does not depend on ε at all.

Provenance. This solution was produced by the *conjecture-harness* campaign framework used throughout this site, not by a single AI pass. The chain transcribed below (Lemmas 1–4, the Theorem, the Corollary, and Corollary 2) was fixed at the initial construction (campaign artifact `0048-RKHS-1`) and then re-verified, unchanged, across six revision rounds by twenty independent blind adversarial referee passes (five referees per audited round, across four audited rounds, drawn from Claude 3.5 Sonnet, Claude Opus 4 and Claude Haiku 3.5) and four independent triage rulings. No upheld referee finding, in any round, has ever touched this load-bearing chain; every upheld finding concerned only the surrounding scope commentary (Remark 2) or citation hygiene. **No human has reviewed the mathematics below.** The full audit trail – every referee report and every triage ruling – is recorded at `c/0048/campaign/LEDGER.md`, which this document treats as its record of verification.

An unresolved fork. A separate, concurrent session forked this campaign at round 3 and closed the same scope-commentary gap addressed in Remark 2 by *deleting* the disputed material, rather than by adding the argument this document reproduces; that fork is recorded as artifact `0048-RKHS-1-r5` (built on an earlier, superseded triage ruling, `0048-RKHS-1-r4`). Both branches leave the load-bearing mathematics – Lemmas 1–4, the Theorem, the Corollary, and Corollary 2 – byte-identical and mathematically complete; the difference between them is editorial only (addition versus deletion of scope commentary). That fork is, as of this writing, **unreconciled** in the campaign ledger and awaits a human decision between the two sessions; this document follows the main branch, frozen artifact `0048-RKHS-1-r6`. Date: 25 August 2026.

The conjecture, as posed

The definition and the conjecture below are reproduced verbatim from the statement note *Signed Low-Entropy Reweightings to Rank One on the Sphere* (`main.tex` in this folder), whose numbering is preserved: Definition 1 and Conjecture 1 here *are* Definition 1 and Conjecture 1 there. That note's "Known boundary" paragraph records that the statement restricted to nonnegative r is established, and tight, for $\delta \ge 1/2$ (this is Theorem 2.3 of [BKS16], specialised to honest distributions), and that the regime $\delta < 1/2$, which requires cancellation, is the content of the question; both facts are used in Remark 2 below and are not re-derived here.

Definition 1 (Signed reweighting and its entropy cost). Let μ be a Borel probability measure on $\mathbb{S}^{n-1}$ and let $K \ge 0$. A Borel measurable function $r : \mathbb{S}^{n-1} \to \mathbb{R}$ is called a *signed reweighting of μ with entropy cost at most K* if r is

μ-integrable and

$$\mathop{\mathbb{E}}_{v\sim\mu} \big|r(v)\big| \;=\; 1,$$

$$\mathop{\mathbb{E}}_{v\sim\mu} \Big[\big|r(v)\big| \log \big|r(v)\big| \Big] \;\le\; K.$$

If in addition $r \ge 0$ holds μ-almost everywhere, then r is a *nonnegative* reweighting; in that case $d\mu' := r\,d\mu$ is a probability measure and the entropy cost is exactly the Kullback–Leibler divergence $\Delta_{\mathrm{KL}}(\mu' \,\|\, \mu) = \mathbb{E}_{v\sim\mu'} \log\big(d\mu'/d\mu\big)$. Throughout, $0\log 0 := 0$.

Conjecture 1 (signed low-entropy reweighting to rank one). *For every $\varepsilon > 0$ and every $\delta > 0$ there exists a finite constant $C = C(\varepsilon,\delta)$ such that the following holds for every $n \in \mathbb{N}$ and every Borel probability measure μ on $\mathbb{S}^{n-1}$: there exist a signed reweighting r of μ with entropy cost at most $C \cdot n^{\delta}$ (in the sense of Definition 1) and a nonzero rank one matrix $L \in \mathbb{R}^{n\times n}$ such that*

$$\Big\| \mathop{\mathbb{E}}_{v\sim\mu} \big[\, r(v)\, vv^{\top} \,\big] \;-\; L \Big\|_F \;\le\; \varepsilon\, \|L\|_F.$$

1 Introduction

This note transcribes and typesets, without alteration, the proof recorded in the frozen campaign artifact `0048-RKHS-1-r6`, which resolves Conjecture 1 affirmatively – and more than affirmatively. For every n and every Borel probability measure μ on $\mathbb{S}^{n-1}$, there is a signed reweighting of μ, of entropy cost at most $\log\big(n(n+1)/2\big)$, whose reweighted second moment equals a nonzero rank one matrix *exactly*: Frobenius error 0, not merely $\le \varepsilon\|L\|_F$ for the given ε (Theorem 1). Consequently the conjectured constant can be taken $C(\varepsilon,\delta) = 2/(e\delta)$, independent of ε (Corollary 1). The construction generalises verbatim from the sphere to any distribution over rank one matrices of unit Frobenius norm (Corollary 2), which is the domain the source paper's own Theorem 2.3 and its intended application actually concern.

Two things this does *not* establish are stated as plainly as the source artifact states them (Remark 2). First, no consequence for the log rank conjecture is claimed: the submatrix dictionary the source paper uses to convert a rank one approximation into a monochromatic rectangle is stated for *flat* distributions only, and nothing in the construction below guarantees flatness of r. Second, the pseudo-distribution extension that the source paper's own running-time application actually needs is out of scope: every use of $\operatorname{supp}\mu$ in the proof below – Lemma 1 (b),(c), Lemma 2, Lemma 3 and the Theorem – presumes an honest Borel probability distribution, not a sum-of-squares pseudo-distribution.

Nor does Theorem 1 conflict with the source paper's own tightness claim for *nonnegative* reweightings. That claim (its Theorem 2.3, tight as stated) is a worst-case statement: it says some μ forces $\delta \ge 1/2$ if $r \ge 0$ is required. It is not a claim about every μ, and Remark 2 shows, by two parallel arguments internal to the artifact – one for the sphere case, one for Corollary 2's generalisation – that whenever *this* construction happens to return a nonnegative r, that r is itself already a nonnegative reweighting of entropy cost at most $\log(n(n+1)/2)$ (respectively $\log(n^2+1)$), so such a μ was never a hard instance for the nonnegative question to begin with. No gap remains open in the mathematics transcribed below; the only outstanding item is the editorial fork disclosed above the title.

2 The kernel construction

Fix $n \geq 1$ and a Borel probability measure μ on $\mathbb{S}^{n-1}$. Write Sym_n for the real symmetric $n \times n$ matrices and, for $Q \in \mathrm{Sym}_n$, let $q_Q(v) := v^\top Q v$ be the associated quadratic form, regarded as a function on $\mathbb{S}^{n-1}$. Let

$$F_\mu \;:=\; \big\{\, [q_Q] \in L^2(\mu) \;:\; Q \in \mathrm{Sym}_n \,\big\}, \qquad d \;:=\; \dim F_\mu .$$

Lemma 1 (the quadratic space, evaluation on the support, and its kernel). *(a) F_μ is a subspace of $L^\infty(\mu) \cap L^2(\mu)$ with $d \leq n(n+1)/2$ and $\|q_Q\|_\infty \leq \|Q\|_F$;*

(b) if $q_Q = q_{Q'}$ μ-a.e. then $q_Q(v) = q_{Q'}(v)$ for every $v \in \operatorname{supp}\mu$; hence $\mathrm{ev}_v : F_\mu \to \mathbb{R}$, $[q_Q] \mapsto q_Q(v)$, is well defined and linear for $v \in \operatorname{supp}\mu$;

(c) for each $v \in \operatorname{supp}\mu$ there is a unique $k_v \in F_\mu$ with $\langle k_v, f\rangle_{L^2(\mu)} = \mathrm{ev}_v(f)$ for all $f \in F_\mu$; set $K(v,v) := \|k_v\|^2_{L^2(\mu)} = \mathrm{ev}_v(k_v)$;

(d) $[1] = [q_I] \in F_\mu$, and $[v \mapsto v_a v_b] = [q_{E_{ab}}] \in F_\mu$ where $E_{ab} = (e_a e_b^\top + e_b e_a^\top)/2$.

Preliminaries used in the proof. (P1) $\mu(\operatorname{supp}\mu) = 1$: the complement of the support is a union of open null sets, and $\mathbb{S}^{n-1}$ is second countable, so that union is a countable union of null sets. (P2) if $v \in \operatorname{supp}\mu$ and $U \ni v$ is open then $\mu(U) > 0$ (definition of support).

Proof. (a) $Q \mapsto q_Q$ is linear, so F_μ is the image of Sym_n and $d \leq \dim \mathrm{Sym}_n = n(n+1)/2$. For $v \in \mathbb{S}^{n-1}$, $|v^\top Q v| \leq \|Q\|_{\mathrm{op}}\|v\|_2^2 = \|Q\|_{\mathrm{op}} \leq \|Q\|_F$. Since μ is a probability measure, $L^\infty(\mu) \subseteq L^2(\mu)$.

(b) Put $h := q_{Q-Q'}$, continuous, with $\mu(\{h \neq 0\}) = 0$. If $h(v_0) \neq 0$ for some $v_0 \in \operatorname{supp}\mu$, then $U := \{|h| > |h(v_0)|/2\}$ is open, contains v_0, so $\mu(U) > 0$ by (P2); but $U \subseteq \{h \neq 0\}$ is null. Contradiction. Well-definedness of ev_v follows: ev_v is defined through quadratic-form representatives only, and by (b) any two quadratic-form representatives of the same class of F_μ take the same value at every $v \in \operatorname{supp}\mu$; linearity is inherited from $Q \mapsto q_Q(v)$.

(c) F_μ is a finite-dimensional subspace of $L^2(\mu)$, hence a Hilbert space; ev_v is a linear functional on it, hence bounded, hence has a unique Riesz representer $k_v \in F_\mu$. Explicitly $k_v = \sum_j \mathrm{ev}_v(f_j) f_j$ for any orthonormal basis (f_j). Taking $f = k_v$ gives $K(v,v) = \mathrm{ev}_v(k_v)$.

(d) $q_I(v) = \|v\|_2^2 = 1$ identically on $\mathbb{S}^{n-1}$; $q_{E_{ab}}(v) = v_a v_b$ identically. □

Note. Part (b) is the only place where the general Borel case differs from the finitely supported case, and it removes the need for any approximation argument; for finitely supported μ it is trivial, since every support point is then an atom.

Lemma 2 (three reproducing identities). *Let $v_0 \in \operatorname{supp}\mu$ and $k := k_{v_0}$. Then*

(i) $\mathbb{E}_\mu\left[k(v)\, vv^\top\right] = v_0 v_0^\top$*;*

(ii) $\mathbb{E}_\mu[k] = 1$*; hence* $\mathbb{E}_\mu|k| \geq 1$, $K(v_0, v_0) \geq 1$, *and* $k \neq 0$*;*

(iii) $\mathbb{E}_\mu[k^2] = K(v_0, v_0)$ *and* $1 \leq \mathbb{E}_\mu|k| \leq K(v_0, v_0)^{1/2} < \infty$.

Proof. All integrals converge absolutely: k is bounded (Lemma 1(a)) and the other factors are bounded by 1 on $\mathbb{S}^{n-1}$.

(i) Fix a, b. By Lemma 1(d), $g_{ab} := [v \mapsto v_a v_b] \in F_\mu$, so the reproducing property gives $\mathbb{E}_\mu[k\, v_a v_b] = \langle k, g_{ab}\rangle = \mathrm{ev}_{v_0}(g_{ab}) = (v_0)_a (v_0)_b$. The (a, b) entries of the two matrices therefore agree, for all a, b.

(ii) By Lemma 1(d), $[1] \in F_\mu$ with $\mathrm{ev}_{v_0}([1]) = q_I(v_0) = 1$, so $\mathbb{E}_\mu[k] = \langle k, [1]\rangle = 1$. Then $\mathbb{E}_\mu|k| \geq |\mathbb{E}_\mu k| = 1$ and, by Cauchy–Schwarz, $K(v_0, v_0) = \mathbb{E}_\mu[k^2] \geq (\mathbb{E}_\mu|k|)^2 \geq 1$. And $\mathbb{E}_\mu[k] = 1 \neq 0$ forces $k \neq 0$.

(iii) $K(v_0, v_0) = \|k\|^2 = \mathbb{E}_\mu[k^2]$ by definition; finiteness is Lemma 1(a); the upper bound on $\mathbb{E}_\mu|k|$ is Cauchy–Schwarz against 1. □

Structural note. Part (ii) is exactly the trace of part (i), since $\operatorname{tr}\mathbb{E}_\mu[k\, vv^\top] = \mathbb{E}_\mu[k\, \|v\|_2^2] = \mathbb{E}_\mu[k]$ and $\operatorname{tr}(v_0 v_0^\top) = 1$. This locates where the construction gets its normalisation for free: the sphere constraint $\|v\|_2 = 1$ is exactly what makes the constant function a quadratic form.

Lemma 3 (a cheap base point exists). $\int K(v, v)\, d\mu(v) = d$, *and consequently there exists* $v_0 \in \operatorname{supp}\mu$ *with* $K(v_0, v_0) \leq d \leq n(n+1)/2$.

Proof. Fix an orthonormal basis $f_1, \ldots, f_d$ of F_μ and quadratic-form representatives q_j of f_j. For $v \in \operatorname{supp}\mu$, Lemma 1(c) and Parseval give $k_v = \sum_j q_j(v) f_j$ and hence

$$K(v, v) \;=\; \sum_{j=1}^{d} q_j(v)^2 \;=:\; g(v) \qquad \text{for all } v \in \operatorname{supp}\mu. \tag{1}$$

g is a polynomial, hence Borel, and by (P1) it represents $v \mapsto K(v, v)$ up to a null set, so the integral is well defined and

$$\int K(v, v)\, d\mu \;=\; \sum_j \int q_j^2\, d\mu \;=\; \sum_j \|f_j\|^2 \;=\; d. \tag{2}$$

(For $v \in \operatorname{supp}\mu$ the value $g(v) = \|k_v\|^2$ is independent of the choices, by uniqueness in Lemma 1(c); off the support nothing below uses g.)

Let $A := \{g \leq d\}$, Borel. If $\mu(A) = 0$ then $h := g - d$ satisfies $h > 0$ μ-a.e. while $\int h\, d\mu = 0$ by (2); a nonnegative function with vanishing integral vanishes a.e., so $h = 0$ μ-a.e., which cannot hold simultaneously with $h > 0$ μ-a.e. under a probability measure. So $\mu(A) > 0$, hence by (P1) $A \cap \operatorname{supp}\mu \neq \emptyset$; any v_0 in it satisfies $K(v_0, v_0) = g(v_0) \leq d$. □

Note. Identity (2) is $\int K(v,v)\,d\mu = \operatorname{tr} P$ for P the orthogonal projection onto F_μ, i.e. the mean of the reciprocal Christoffel function; it is proved here from scratch and used nowhere else in this form.

Lemma 4 (entropy at most the log of the second moment). *If $r \in L^2(\mu)$ and $\mathbb{E}_\mu |r| = 1$, then $\mathbb{E}_\mu[\,|r|\log|r|\,] \le \log \mathbb{E}_\mu[r^2]$.*

Proof. Write $\varphi(t) = t\log t$, $\varphi(0) = 0$. Then $\varphi(t) \ge -1/e$ for $t \ge 0$, and $\varphi(t) \le t^2$ for all $t \ge 0$ (for $t \ge 1$ because $\log t \le t$; for $t < 1$ because $\varphi \le 0$). Since $r \in L^2(\mu)$ and μ is a probability measure, $\varphi(|r|) \in L^1(\mu)$, so the entropy is a finite real number.

Let $d\nu := |r|\,d\mu$; then $\nu(\mathbb{S}^{n-1}) = 1$, so ν is a probability measure, and $\nu(\{r = 0\}) = 0$, so $\log|r|$ is ν-a.e. finite and $\mathbb{E}_\nu[\log|r|] = \mathbb{E}_\mu[\,|r|\log|r|\,]$ (the convention $0\log 0 = 0$ matching the ν-null set $\{r = 0\}$). Also $\mathbb{E}_\nu[|r|] = \mathbb{E}_\mu[r^2] < \infty$. Jensen's inequality for the concave function log against ν gives $\mathbb{E}_\nu[\log|r|] \le \log \mathbb{E}_\nu[|r|] = \log \mathbb{E}_\mu[r^2]$. □

Sharpness. Equality holds iff $|r| \in \{0, c\}$ μ-a.e., i.e. exactly on the indicator-type ("conditioning on an event") reweightings, where the bound reads $\log(1/\mu(\{|r| = c\}))$. The bound is lossy precisely to the extent that $|r|$ is spread out.

3 The main theorem

Theorem 1. *For every $n \ge 1$ and every Borel probability measure μ on $\mathbb{S}^{n-1}$ there exist a point $v_0 \in \operatorname{supp}\mu$, a constant c with $\big(n(n+1)/2\big)^{-1/2} \le c \le 1$, and a signed reweighting r of μ – indeed the restriction to $\mathbb{S}^{n-1}$ of a single quadratic form, hence a bounded Borel function – such that*

$$\mathop{\mathbb{E}}_{v\sim\mu}\left[\, r(v)\, vv^\top \,\right] \;=\; c \cdot v_0 v_0^\top \qquad \text{(exactly)}$$

and

$$\mathop{\mathbb{E}}_{\mu}\left[\, |r|\log|r| \,\right] \;\le\; \log K(v_0,v_0) \;\le\; \log\left(\frac{n(n+1)}{2}\right).$$

Here $L := c\, v_0 v_0^\top$ is rank one, nonzero and positive semidefinite, with $\|L\|_F = c$.

Proof outline. $F_\mu \subseteq L^2(\mu)$ is the space of restrictions to $\mathbb{S}^{n-1}$ of quadratic forms, of dimension $d \le n(n+1)/2$. On the sphere the constant function 1 *is* the quadratic form q_I, and each coordinate product $v_a v_b$ is the quadratic form $q_{E_{ab}}$. Hence the reproducing kernel $k = k_{v_0}$ of F_μ at a support point v_0 satisfies simultaneously $\mathbb{E}_\mu[k] = 1$ (free normalisation) and $\mathbb{E}_\mu[k(v)vv^\top] = v_0 v_0^\top$ (exact rank one). Its squared L^2 norm averages to d over μ, so some base point is cheap; and entropy is bounded by the log of the second moment. Cancellation is available through k, which is negative somewhere whenever $\mathbb{E}_\mu|k| > \mathbb{E}_\mu[k] = 1$.

Proof of the Theorem. Lemma 3 supplies $v_0 \in \operatorname{supp}\mu$ with $K(v_0,v_0) \le d \le n(n+1)/2$; let $k := k_{v_0}$. By Lemma 2(ii),(iii), $1 \le \mathbb{E}_\mu|k| \le K(v_0,v_0)^{1/2} < \infty$ and $k \ne 0$, so

$$c \;:=\; \frac{1}{\mathbb{E}_\mu|k|} \;\in\; \left[\, K(v_0,v_0)^{-1/2}, 1 \,\right] \;\subseteq\; \left[\, \big(n(n+1)/2\big)^{-1/2}, 1 \,\right].$$

Set $r := c \cdot k$, taking as representative the quadratic form q_{cQ_k}; then r is a bounded Borel function on $\mathbb{S}^{n-1}$, hence μ-integrable, and $\mathbb{E}_\mu |r| = c\, \mathbb{E}_\mu |k| = 1$. So r is admissible for Definition 1.

Second moment. By Lemma 2(i), $\mathbb{E}_\mu[r(v)vv^\top] = c \cdot v_0 v_0^\top =: L$. As $c > 0$ and $\|v_0\|_2 = 1$, $L = (cv_0)v_0^\top$ is rank one, nonzero, positive semidefinite, with $\|L\|_F = c > 0$. The Frobenius error is $0 \le \varepsilon \|L\|_F$ for every $\varepsilon > 0$.

Entropy. $r \in L^2(\mu)$ with $\mathbb{E}_\mu |r| = 1$, so Lemma 4 and Lemma 2(iii) give

$$\mathop{\mathbb{E}}_{\mu}\bigl[\, |r| \log |r| \,\bigr] \;\le\; \log \mathop{\mathbb{E}}_{\mu}[r^2] \;=\; \log\Bigl(c^2 K(v_0, v_0)\Bigr) \;=\; \log\left(\frac{K(v_0,v_0)}{(\mathbb{E}_\mu |k|)^2}\right) \;\le\; \log K(v_0,v_0) \;\le\; \log \Bigl(\frac{n(n+1}{2}$$

using $\mathbb{E}_\mu |k| \ge 1$ and Lemma 3. □

Corollary 1 (Conjecture 1). *Conjecture 1 holds with $C(\varepsilon,\delta) := 2/(e\delta)$, independent of ε.*

Proof. For $n = 1$, $\log(n(n+1)/2) = \log 1 = 0$. For $n \ge 2$, $n(n+1)/2 \le n^2$, so $\log(n(n+1)/2) \le 2 \log n$; and $\max_{x>0} (\log x)/x^\delta = 1/(e\delta)$ (attained at $x = e^{1/\delta}$), so $2 \log n \le (2/(e\delta))\, n^\delta$. Combine with Theorem 1, taking $L = c\, v_0 v_0^\top$ for the nonzero rank one matrix required by Conjecture 1: the Frobenius error is exactly $0 \le \varepsilon \|L\|_F$ for every $\varepsilon > 0$, so the bound holds with a single $C(\varepsilon,\delta)$ that does not use ε at all. □

This is strictly more than Conjecture 1 asks for. The conjecture only demands, for each fixed $\varepsilon > 0$, some finite $C(\varepsilon,\delta)$; Corollary 1 exhibits a single $C(\delta) = 2/(e\delta)$ that works simultaneously at *every* $\varepsilon > 0$, because the construction's Frobenius error is exactly 0 rather than merely $\le \varepsilon \|L\|_F$.

4 Further remarks

Remark 1 (the conclusion is not obtainable by cancelling to zero). This remark is recorded because it shows that the relative normalisation in Conjecture 1 does real work, so Theorem 1 is not exploiting a degenerate reading of the statement; it is not used elsewhere in this document. Let $M \in \mathrm{Sym}_n$ have eigenvalues indexed so that $|\lambda_1| \ge |\lambda_2| \ge \cdots$, and let $\varepsilon \in (0,1)$.

- If $M = 0$, no admissible L exists: $\|0 - L\|_F = \|L\|_F > \varepsilon \|L\|_F$ for every $L \ne 0$.
- If $M \ne 0$, a nonzero rank one L with $\|M - L\|_F \le \varepsilon \|L\|_F$ exists *iff* $\sum_{i \ge 2} \lambda_i^2 \le \bigl(\varepsilon^2/(1-\varepsilon^2)\bigr)\lambda_1^2$.

Proof. Write $L = t\, uw^\top$, $\|u\| = \|w\| = 1$, $t > 0$, so $\|L\|_F = t$ and $\|M - L\|_F^2 = \|M\|_F^2 - 2t \cdot u^\top M w + t^2$. The requirement is $(1-\varepsilon^2)t^2 - 2(u^\top M w)t + \|M\|_F^2 \le 0$, and increasing $u^\top M w$ only helps, so the optimal choice is $u^\top M w = s_1(M) = |\lambda_1|$ (singular value decomposition). The quadratic in t has a real root iff $s_1^2 \ge (1-\varepsilon^2)\|M\|_F^2 = (1-\varepsilon^2)\sum_i \lambda_i^2$, which rearranges to the stated inequality; when $M \ne 0$ both roots are positive (their product is $\|M\|_F^2/(1-\varepsilon^2) > 0$, their sum $2s_1/(1-\varepsilon^2) > 0$), so an admissible $t > 0$ exists exactly then. □

Corollary 2 (the same for distributions over rank one matrices of unit Frobenius norm). *Let* $\mathcal{X} := \{X \in \mathbb{R}^{n\times n} : \operatorname{rank} X = 1,\ \|X\|_F = 1\}$ *– equivalently* $X = ab^\top$ *with* $\|a\|_2 = \|b\|_2 = 1$*, the normalisation used in the source paper's intended application (§2.3 of [BKS16]: "we will restrict our attention to the case that all the columns of U and V are of unit norm. (This restriction is easy to lift and anyway holds automatically in our intended application.)"). Let* μ *be a Borel probability measure on* $\mathcal{X}$*. Then there exist* $X_0 \in \operatorname{supp}\mu$*,* $c \in (0,1]$*, and a signed reweighting* r *of* μ *with* $\mathbb{E}_\mu[\,|r|\log|r|\,] \le \log(n^2+1)$ *such that* $\mathbb{E}_{X\sim\mu}[\,r(X)\cdot X\,] = c\cdot X_0$*, exactly rank one and nonzero.*

Proof. Run Lemmas 1–4 verbatim with $\mathbb{S}^{n-1}$ replaced by $\mathcal{X}$ (compact) and the space of quadratic forms replaced by the affine space

$$G \;:=\; \left\{X \mapsto \alpha + \langle A, X\rangle \;:\; \alpha \in \mathbb{R},\ A \in \mathbb{R}^{n\times n}\right\},$$

of dimension at most n^2+1. Only two properties of the function space were used: it contains the constant 1 (here trivially, $\alpha = 1$, $A = 0$) and it contains each coordinate function $X \mapsto X_{ab}$ (here $\alpha = 0$, $A = e_a e_b^\top$). Continuity of the elements of G gives the analogue of Lemma 1(b) unchanged; boundedness on $\mathcal{X}$ gives the analogue of Lemma 1(a). □

Note. The mechanism is therefore not an artefact of specialising to $X = vv^\top$: it is the presence of constants and coordinates in one finite-dimensional function space. On the sphere the homogeneous quadratic forms suffice because $\|v\|_2 = 1$ makes the constant a quadratic form; that only buys the sharper bound $\log(n(n+1)/2)$ over $\log(n^2+1)$.

Remark 2 (scope). Theorem 1 answers Question 8.1 of [BKS16] as literally posed, and nothing more.

(i) No conflict with the source's negative results. The statement "it can be shown that as stated, Theorem 2.3 is tight" ([BKS16], p. 6) is about *deficient* reweightings, which are probability distributions and hence nonnegative by Definition 2.2 of [BKS16]; and the statement "the answer to this question is No if one does not allow negative reweighting functions" ([BKS16], p. 19) is an existential statement over μ. Neither is contradicted here, because whenever the construction returns a nonnegative r, that r is itself a nonnegative reweighting of entropy cost at most $\log(n(n+1)/2) = O(\log n)$ – equivalently, $k \ge 0$ forces $\mu(\{\pm v_0\}) = 1/K(v_0, v_0) \ge 2/(n(n+1))$ – so such a μ is not a hard instance for the nonnegative question.

(ii) No log-rank consequence is claimed. The submatrix dictionary of [BKS16], §2.2, is stated for *flat* distributions, and the Theorem supplies no flatness guarantee, so that route is unavailable; no other route is addressed.

(iii) The pseudo-distribution version is out of scope. Every use of $\operatorname{supp}\mu$ in the proof – Lemma 1(b),(c), Lemma 2, Lemma 3, the Theorem – presumes an honest distribution. The source paper attaches its pseudo-distribution condition to its *running-time* consequence (Question 8.1, p. 19), and this document addresses neither.

Derivation of the "equivalently" clause of (i), sphere case. Let v_0 and $k = k_{v_0}$ be as in the proof of the Theorem, and suppose $k \geq 0$ μ-a.e. By Lemma 2(ii), $d\mu' := k\,d\mu$ is a probability measure, and by Lemma 2(i), $\mathbb{E}_{\mu'}[vv^\top] = v_0v_0^\top$. For a unit $w \perp v_0$ this gives $\mathbb{E}_{\mu'}\langle v, w\rangle^2 = w^\top v_0 v_0^\top w = 0$, so $\langle v, w\rangle = 0$ μ'-a.s.; ranging w over an orthonormal basis of $v_0^\perp$ and using $\|v\|_2 = 1$ forces $\mu'(\{\pm v_0\}) = 1$. Since k is (the class of) a quadratic form, $k(-v_0) = k(v_0) = K(v_0, v_0)$ by Lemma 1(c), so $1 = \mu'(\{\pm v_0\}) = K(v_0, v_0)\cdot\mu(\{\pm v_0\})$, i.e. $\mu(\{\pm v_0\}) = 1/K(v_0, v_0)$; and $K(v_0, v_0) \leq n(n+1)/2$ by Lemma 3 gives $\mu(\{\pm v_0\}) \geq 2/(n(n+1))$. Conditioning on that atom, $r' := \mathbf{1}_{\{\pm v_0\}}/\mu(\{\pm v_0\})$, is then a nonnegative reweighting reaching $v_0v_0^\top$ at entropy cost $\log(1/\mu(\{\pm v_0\})) = \log K(v_0, v_0) \leq \log(n(n+1)/2)$, the same bound the Theorem gives for r itself. Every quantity used here is Lemma 1(c), Lemma 2(i), (ii), Lemma 3 or the Theorem; nothing is taken on the authority of a document outside this construction.

The same dichotomy for Corollary 2's construction. This second argument is fully internal to the construction above and strictly simpler than the sphere-case argument just given, because $\mathcal{X}$ carries no antipodal identification and there is no $-X_0$ to account for. Let X_0 and $k = k_{X_0}$ be as in the proof of Corollary 2, and suppose $k \geq 0$ μ-a.e. By Corollary 2's own transfer of Lemma 2(ii) (using $[1] \in G$), $d\mu' := k\,d\mu$ is a probability measure, and by the transfer of Lemma 2(i) (using $[X \mapsto X_{ab}] \in G$ for every a, b), $\mathbb{E}_{\mu'}[X] = X_0$. Every X in the domain $\mathcal{X}$ has $\|X\|_F = 1$, and $\|X_0\|_F = 1$ since $X_0 \in \mathcal{X}$, so, using only the bilinear expansion of the Frobenius inner product exactly as in Lemma 2 and Remark 1,

$$\mathop{\mathbb{E}}_{\mu'} \left\|X - X_0\right\|_F^2 = \mathop{\mathbb{E}}_{\mu'} \|X\|_F^2 - 2\langle X_0, \mathop{\mathbb{E}}_{\mu'} X\rangle + \|X_0\|_F^2 = 1 - 2\langle X_0, X_0\rangle + 1 = 0.$$

A nonnegative integrand with vanishing integral vanishes a.e., so $X = X_0$ μ'-a.s., i.e. $\mu'(\{X_0\}) = 1$ – a point mass, as promised, simpler than the sphere case. Hence $1 = \int_{\{X_0\}} k\,d\mu = K(X_0, X_0)\cdot\mu(\{X_0\})$, so $\mu(\{X_0\}) = 1/K(X_0, X_0) \geq 1/(n^2+1)$ (Corollary 2's transfer of Lemma 3, with $\dim G \leq n^2 + 1$), and conditioning on that atom, $r' := \mathbf{1}_{\{X_0\}}/\mu(\{X_0\})$, is a nonnegative reweighting reaching X_0 exactly at entropy cost $\log(1/\mu(\{X_0\})) = \log K(X_0, X_0) \leq \log(n^2 + 1)$ – the same bound Corollary 2 already proves for its own (possibly signed) r. So whenever Corollary 2's construction returns a nonnegative r, that μ is, by exactly this argument, not a hard instance for the nonnegative version of Theorem 2.3 of [BKS16] either. Every quantity used here is Corollary 2's own statement and proof, which states the transfer of Lemmas 1–4 verbatim, together with the Frobenius inner product already used in Lemma 2 and Remark 1; nothing is taken on the authority of any document outside this construction.

5 A worked example

The following hand computation, independent of the proof above, illustrates Theorem 1 at $n = 2$. Let μ be uniform on the four points

$$v_1 = (1, 0), \quad v_2 = (0, 1), \quad v_3 = \tfrac{1}{\sqrt{2}}(1, 1), \quad v_4 = \tfrac{1}{\sqrt{2}}(1, -1).$$

Identifying F_μ with a 4-dimensional coordinate space via evaluation at $v_1, \ldots, v_4$, one finds $F_\mu = \{x \in \mathbb{R}^4 : x_1 + x_2 = x_3 + x_4\}$, so $d = 3$. The kernel at $v_0 = v_1$ is $k = (3, -1, 1, 1)$

(i.e. $k(v_1) = 3$, $k(v_2) = -1$, $k(v_3) = k(v_4) = 1$). Then

$$\mathop{\mathbb{E}}_{\mu}[k] = \tfrac{1}{4}(3 - 1 + 1 + 1) = 1 \quad \checkmark, \qquad K(v_0, v_0) = \|k\|^2 = \tfrac{1}{4}(9 + 1 + 1 + 1) = 3 = d \quad \checkmark,$$

$$\mathop{\mathbb{E}}_{\mu}[k\, vv^\top] = \tfrac{1}{4}\Big(3v_1v_1^\top - v_2v_2^\top + v_3v_3^\top + v_4v_4^\top\Big) = v_1v_1^\top \quad \checkmark$$

(exactly rank one, as Lemma 2(i) predicts). Normalising, $\mathbb{E}_\mu |k| = \frac{1}{4}(3+1+1+1) = \frac{3}{2}$, so $c = 2/3$ and $r = c\,k = \big(2, -\frac{2}{3}, \frac{2}{3}, \frac{2}{3}\big)$, with $\mathbb{E}_\mu |r| = 1$ as required, and

$$\mathop{\mathbb{E}}_{\mu}\big[\,|r| \log |r|\,\big] \;=\; 0.1438\ldots \;\le\; \log \mathop{\mathbb{E}}_{\mu}[r^2] \;=\; 0.2877\ldots \;\le\; \log 3 \;=\; 1.0986\ldots \qquad \checkmark,$$

matching Lemma 4 and Theorem 1. Note that r is genuinely signed here (it is negative at v_2), illustrating that the cheap base point v_1 need not make k nonnegative even in this small example.

6 Bibliography

[BKS16] B. Barak, P. Kothari, D. Steurer. *Quantum entanglement, sum of squares, and the log rank conjecture.* arXiv:1701.06321v2 [quant-ph], 2017.