Sharp universal Rényi threshold for iid sums on finite Sidon sets

Dev D. Goyal

28 October 2026; revised 1 November 2026

Abstract

Let X, X' be independent random variables with the same probability law on a nonempty finite Sidon set, where sums are unique up to permutation, including repeated summands. The result determines exactly which Rényi orders $\alpha > 1$ have the following universal property: uniform input maximizes the entropy of $X + X'$ on every such support. The necessary and sufficient condition is $2^{\alpha-1} \le 2\alpha - 1$, equivalently $1 < \alpha \le \alpha_c$, where $\alpha_c \in (3, 4)$ is the unique nontrivial solution of $2^\alpha - 4\alpha + 2 = 0$ and is approximately 3.65986. The sufficient direction follows from a sharp two-coordinate power inequality and a gradient comparison. A binary obstruction proves necessity. The proof includes zero input probabilities and the critical endpoint. The supplied source context credits Gu [1] with an earlier sufficient bound, the threshold equation and the binary obstruction; that attribution has not been verified against the primary text, and no priority is claimed. A Lean formalization covers the classification and its convolution-law bridge.

1 The model and source context

A finite set $A = \{v_1, \dots, v_n\}$ in an abelian group is a *Sidon set* in the strong B_2 sense if

$$v_i + v_j = v_k + v_\ell \quad \Longrightarrow \quad (i = k \text{ and } j = \ell) \text{ or } (i = \ell \text{ and } j = k).$$

Indices may repeat. Thus diagonal sums $2v_i$ are included in the uniqueness requirement. This displayed condition defines the support convention in this note. The formal proof works more generally in a commutative additive semigroup under exactly the displayed condition.

Let $p = (p_1, \dots, p_n)$ be any probability vector: $p_i \ge 0$ and $\sum_i p_i = 1$. Let X and X' be independent with $\mathbb{P}(X = v_i) = \mathbb{P}(X' = v_i) = p_i$. All logarithms in this note are natural, so entropies are measured in nats. For $\alpha > 1$,

$$H_\alpha(X + X') = \frac{1}{1-\alpha} \log\left(\sum_z \mathbb{P}(X + X' = z)^\alpha \right).$$

Zero probabilities are allowed, with $0^\alpha = 0$.

The question is universal in the support and its cardinality:

> For which $\alpha > 1$ does uniform input maximize $H_\alpha(X + X')$ for every nonempty finite Sidon support?

Gu [1] is credited in the supplied source context for an earlier sufficient bound, the threshold equation, and a binary obstruction. Hillion and Johnson [2] are credited for related binary Tsallis curvature and gradient-comparison methods. The primary pages and current versions were not verified for this note. Exact proposition numbers, support conventions, and the extent

of earlier sufficient results therefore remain unverified. If those conventions and attributions agree, the theorem below gives the matching sufficient bound in this model. It is presented as a self-contained classification, without a claim to close a verified published gap or to have first priority. The theorem concerns identical iid input laws and orders strictly above one; it makes no Shannon-entropy or different-input-law claim. The Lean formalization described in Section 9 proves Theorem 1 for every order above one.

2 Complete classification

Theorem 1 (Sharp universal threshold)**.** *Fix a real number $\alpha > 1$. The following are equivalent:*

(i) For every nonempty finite Sidon support and every probability law on it, uniform input maximizes $H_\alpha(X + X')$ for two iid draws.

(ii) $2^{\alpha-1} \le 2\alpha - 1$.

(iii) $\alpha \le \alpha_c$, where α_c is the unique real number greater than 1 satisfying $2^{\alpha_c - 1} = 2\alpha_c - 1$.

The critical order satisfies $3 < \alpha_c < 4$. In the valid range, the maximum on any Sidon support of cardinality n is

$$\frac{1}{1-\alpha} \log \left(\frac{n + 2^{\alpha-1} n(n-1)}{n^{2\alpha}} \right). \tag{1}$$

For every $\alpha > \alpha_c$, a strictly positive nonuniform law on the concrete Sidon set $\{0, 1\}$ has strictly larger entropy than uniform input.

This theorem classifies *universal uniform optimality*. For orders above the threshold, the required negative conclusion is existence of a support where uniform input fails. The theorem does not assert failure for every cardinality, or classify all maximizing laws for a fixed larger cardinality. No equality-case uniqueness statement about the optimizing law is required or claimed in the formal classification.

3 The exact output moment

The Sidon condition gives one sum atom of weight p_i^2 for each diagonal pair (i, i) and one of weight $2p_ip_j$ for every unordered distinct pair $\{i, j\}$. Consequently, if $c = 2^{\alpha-1}$ and $S = \sum_i p_i^\alpha$, the output power moment is

$$\begin{aligned} F_\alpha(p) &= \sum_i p_i^{2\alpha} + \sum_{i<j} (2p_ip_j)^\alpha \\ &= c \left(\sum_i p_i^\alpha \right)^2 - (c-1) \sum_i p_i^{2\alpha}. \end{aligned} \tag{2}$$

This is strictly positive for every probability vector: at least one p_i is positive, giving a positive diagonal term. Since $1 - \alpha < 0$ and the logarithm is increasing, maximizing entropy is equivalent to minimizing F_α.

At $u = (1/n, \ldots, 1/n)$, direct substitution into the first line gives

$$F_\alpha(u) = n\, n^{-2\alpha} + \frac{n(n-1)}{2} (2n^{-2})^\alpha = (n + c\, n(n-1)) n^{-2\alpha}.$$

This proves the asserted value of the maximum once uniform optimality is established.

4 A sharp power-difference inequality

Lemma 2. *If $r > 1$ and $x > y > 0$, then*

$$x^r - y^r > r(xy)^{(r-1)/2}(x - y). \tag{3}$$

The corresponding non-strict inequality holds for $x \geq y > 0$.

Proof. Write $x = e^{m+t}$ and $y = e^{m-t}$ with $t > 0$. The difference between the two sides of (3) is

$$2e^{rm}(\sinh(rt) - r\sinh t).$$

As a function of $t \geq 0$, the expression in parentheses vanishes at zero and has derivative $r(\cosh(rt) - \cosh t) > 0$ for $t > 0$. It is therefore strictly positive. The case $x = y$ is equality. The Lean proof uses the equivalent exponential expressions rather than introducing hyperbolic functions. □

The choice $r = 2\alpha - 1 > 1$ gives

$$x^{2\alpha-1} - y^{2\alpha-1} \geq (2\alpha - 1)(xy)^{\alpha-1}(x - y). \tag{4}$$

The inequality is strict if $x > y > 0$. This strictness persists even when the eventual threshold condition is an equality.

5 Ordering the gradient

For a strictly positive vector p, differentiating (2) gives

$$g_i(p) := \frac{\partial F_\alpha}{\partial p_i} = 2\alpha p_i^{\alpha-1}(cS - (c-1)p_i^\alpha). \tag{5}$$

Lemma 3. *Suppose $\alpha > 1$ and $c = 2^{\alpha-1} \leq 2\alpha - 1$. For every strictly positive finite vector p, $p_i \leq p_j$ implies $g_i(p) \leq g_j(p)$. The implication is strict when $p_i < p_j$.*

Proof. Equality of the two coordinates gives equality of their gradient components. For the strict case put $x = p_j > y = p_i > 0$ and $D = x^{\alpha-1} - y^{\alpha-1} > 0$. Since all summands of S are nonnegative, $S \geq x^\alpha + y^\alpha$. Equation (5) yields

$$\begin{aligned}
\frac{g_j - g_i}{2\alpha} &= cSD - (c-1)(x^{2\alpha-1} - y^{2\alpha-1}) \\
&\geq c(x^\alpha + y^\alpha)D - (c-1)(x^{2\alpha-1} - y^{2\alpha-1}) \\
&= x^{2\alpha-1} - y^{2\alpha-1} - c(xy)^{\alpha-1}(x - y) \\
&> (2\alpha - 1 - c)(xy)^{\alpha-1}(x - y) \geq 0.
\end{aligned}$$

The exact identity in the third line follows by expanding $(x^\alpha + y^\alpha)(x^{\alpha-1} - y^{\alpha-1})$. The strict inequality is Lemma 2. Finally $2\alpha > 0$, so $g_j > g_i$. □

The use of *both* coordinates in $S \geq x^\alpha + y^\alpha$ is what retains the sharp coefficient. A bound using only one coordinate loses information needed to reach the binary threshold.

6 From the gradient to a global minimum

Proposition 4. *If $\alpha > 1$ and $2^{\alpha-1} \le 2\alpha - 1$, then $F_\alpha(u) \le F_\alpha(p)$ for every nonnegative probability vector p of every finite length $n \ge 1$.*

Proof. Let $u_i = 1/n$ and interpolate from uniform to p by

$$q_i(t) = (1-t)/n + tp_i, \qquad 0 \le t \le 1.$$

For $0 < t < 1$ every $q_i(t)$ is strictly positive, including when some $p_i = 0$. The ordering of the coordinates of $q(t)$ agrees with that of p. Lemma 3 therefore says that p and $g(q(t))$ have the same coordinate order. Chebyshev's finite sum inequality gives

$$\left(\sum_i p_i\right)\left(\sum_i g_i(q(t))\right) \le n\sum_i p_i g_i(q(t)).$$

For completeness, the difference of the right and left sides equals

$$\frac{1}{2}\sum_{i,j}(p_i - p_j)(g_i(q(t)) - g_j(q(t))),$$

whose summands are nonnegative by the common ordering. Since $\sum_i p_i = 1$, the chain rule implies

$$\begin{aligned}\frac{d}{dt}F_\alpha(q(t)) &= \sum_i (p_i - 1/n)g_i(q(t)) \\ &= \sum_i p_i g_i(q(t)) - \frac{1}{n}\sum_i g_i(q(t)) \ge 0.\end{aligned}$$

The function $F_\alpha(q(t))$ is continuous on $[0,1]$, because the real powers involved have positive exponents. The derivative is needed only on $(0,1)$. The mean-value theorem therefore makes it nondecreasing on the closed interval. Evaluating at the endpoints gives $F_\alpha(u) \le F_\alpha(p)$. □

Combining Proposition 4 with positivity of the moment and the sign of $1-\alpha$ proves the entropy bound for every finite Sidon support. The argument covers $n = 1$, arbitrary zero coordinates of p, and the critical equality $2^{\alpha-1} = 2\alpha - 1$.

7 The matching obstruction above the threshold

The supplied source context attributes this binary obstruction to Gu; that exact source comparison remains unverified. The argument is included and formalized here to make both directions of Theorem 1 self-contained. Take $A = \{0,1\} \subset \mathbb{Z}$ and input weights $(t, 1-t)$. The three output probabilities are t^2, $2t(1-t)$, and $(1-t)^2$. Thus

$$b_\alpha(t) = t^{2\alpha} + (2t(1-t))^\alpha + (1-t)^{2\alpha}.$$

Direct differentiation at $t = 1/2$ gives

$$\begin{aligned} b'_\alpha(1/2) &= 0, \\ b''_\alpha(1/2) &= 2^{4-2\alpha}\alpha(2\alpha - 1 - 2^{\alpha-1}). \end{aligned} \tag{6}$$

If $2^{\alpha-1} > 2\alpha - 1$, the second derivative is negative. All bases of the powers are positive in a neighborhood of $1/2$, so the function is twice differentiable there. Hence, for all sufficiently small nonzero ε,

$$b_\alpha(1/2 + \varepsilon) < b_\alpha(1/2).$$

Choosing $|\varepsilon| < 1/2$ gives a strictly positive normalized input law. Its output moment is smaller, so its Rényi entropy is strictly larger. The formal proof derives the existence of such a point directly from derivatives, a one-sided neighborhood, and the mean-value theorem. It does not rely on a floating-point search.

Related Tsallis curvature arguments are attributed to Hillion and Johnson [2]; the literal lemma comparison remains unverified. Sections 4–6 prove the universal sufficient bound used here.

8 The critical order

Put $\phi(a) = 2^{a-1} - (2a-1)$. It is continuous and strictly convex on $\mathbb{R}$. Its values at $1, 3, 4$ are respectively $0, -1, 1$. The intermediate-value theorem gives a root $\alpha_c \in (3,4)$. Strict convexity gives $\phi(a) < 0$ for $1 < a < \alpha_c$. For $a > \alpha_c$, the strict ordering of adjacent secant slopes gives

$$0 = \frac{\phi(\alpha_c) - \phi(1)}{\alpha_c - 1} < \frac{\phi(a) - \phi(\alpha_c)}{a - \alpha_c},$$

so $\phi(a) > 0$. These observations also prove uniqueness of the root above one and the equivalence

$$2^{a-1} \le 2a - 1 \quad \Longleftrightarrow \quad a \le \alpha_c \qquad (a > 1).$$

Together with the preceding sufficient and necessary directions, this completes Theorem 1. The familiar decimal 3.65986 is only a readable approximation; the formal statement uses the exact root.

9 Lean verification and reproducibility

The accompanying formalization targets Lean 4.19.0 [3] and the following pinned mathlib commit:

`c44e0c8ee63ca166450922a373c7409c5d26b00b`.

The top-level theorem is

```
theorem universalUniformOptimal_iff_criticalOrder
    {a : Real} (ha : 1 < a) :
    UniversalUniformOptimal a <-> a <= criticalOrder
```

Here `UniversalUniformOptimal` is explicitly quantified over every nonempty finite support, every strong Sidon enumeration into an abelian group, and every nonnegative normalized input probability vector. The separate forward theorem is universe-polymorphic. The entropy definition is the logarithm of the power sum of the actual finite convolution law, divided by $1 - a$.

Module	Verified content
`SidonLaw`	Product law, pushforward by addition, normalization, zero outside the finite sum support, exact power-moment identity.
`PowerDifference`	Sharp scalar power inequality and cross-term identity.
`Core`, `Gradient`	Derivative formula, continuity, strict coordinate ordering of the gradient.
`Optimization`	Global minimum at uniform input, including boundary laws.
`AboveThreshold`, `Binary`	Analytic binary improvement and a genuine probability law on the concrete Sidon support $\{0, 1\}$.
`Entropy`	Moment positivity, entropy comparisons, uniform moment value.
`Threshold`	Exact critical root, existence, uniqueness, and interval sign.
`Main`	Both directions of the entropy classification and exact uniform value.

The complete source was freshly rebuilt with Lean 4.19.0 on 28 October 2026. A separate audit with `--trust=0` checked all twelve named declarations in `AxiomAudit.lean`. Both checks succeeded, and the audited declarations depend only on the standard Lean foundations `propext`, `Classical.choice`, and `Quot.sound`. All twelve Lean source files were unchanged during these checks. The Lean kernel and the stated foundational axioms remain the logical trust basis.

The source package contains the pinned toolchain, package manifest, complete Lean sources, and `AxiomAudit.lean`. Reproduction commands and the map from informal results to formal declarations are in the accompanying README and formalization guide.

The verified classification is for orders $a > 1$. Lower orders are outside the stated formal classification. The expression $\log F_a/(1-a)$ is not used to claim a Shannon-entropy theorem at $a = 1$.

10 Scope and attribution

The theorem identifies the exact universal threshold for orders above one. It applies to every finite cardinality and the whole probability simplex. Any comparison with earlier sufficient bounds is conditional on the unverified primary-source conventions and results described above. The threshold equation, its approximate value, the binary obstruction, and the gradient comparison method are not claimed as new. A formal proof establishes the stated mathematical implication; it does not establish historical priority.

The problem of classifying maximizing laws for each fixed cardinality above the universal threshold is outside this theorem. So is the mixed-input Shannon-entropy question at order one.

11 Research process and AI disclosure

This work is fully autonomous agentic research. Autonomous AI research agents produced the research, proofs, formalization and manuscript preparation.

Opus 5.5 by Anthropic used with Claude Code revised this manuscript editorially on 1 November 2026: it moved this disclosure from the title page into its own section and revised the abstract and introduction. It did not change the mathematics.

Dev D. Goyal is responsible for publication. No expert validation or peer review is claimed, and the work has not been externally peer reviewed. Section 9 states the formal scope, and Section 10 states the limits of the result and of the source attributions.

References

[1] Yuzhou Gu. Zero-error communication over adder MAC. arXiv:1809.07364v1, https://arxiv.org/abs/1809.07364, 2018. Supplied source record; primary-text comparison unverified.

[2] Erwan Hillion and Oliver Johnson. A proof of the Shepp–Olkin entropy concavity conjecture. *Bernoulli*, 23(4B):3638–3649, 2017. Supplied source record; primary-text comparison unverified.

[3] Lean 4. Release 4.19.0. https://github.com/leanprover/lean4/releases/tag/v4.19.0.