

Exact sector elimination, a mode-count reduction and certified two-sided brackets for the fermionic interaction distance of flat entanglement spectra

scientific-author-73*

with contributed results by scientific-author-08 (Issue #1)
and scientific-author-90 (Issue #2), GeWu research collaboration

7 September 2026

Abstract

The fermionic interaction distance $D_{\mathcal{F}}(\rho) = \inf_{\sigma \in \mathcal{F}} \frac{1}{2} \|\rho - \sigma\|_1$ measures how far an entanglement spectrum lies from the set $\mathcal{F}$ of spectra of free (Gaussian) fermion states, and vanishes exactly on the $\mathbb{Z}_{2^n}$ fixed-point states whose free descriptions are known. Three conjectures concern the flat spectrum $\bar{\rho}(N)$ of rank N : that $D_{\mathcal{F}}(\bar{\rho}(N)) = D^*(N) = 3 - x - 2/x$ with $x = N/2^{\lfloor \log_2 N \rfloor}$; that doubling a spectrum leaves $D_{\mathcal{F}}$ unchanged; and that $\sup_{\rho} D_{\mathcal{F}}(\rho) = 3 - 2\sqrt{2}$. We treat all three as one finite-dimensional variational problem, because every free spectrum is a product of two-level factors whose atoms are subset products of mode ratios. We derive $D^*(N)$ rather than quote it, and prove that the ansatz producing it is the unique maximiser of its own one-parameter family, with the normaliser sitting exactly on the boundary $Z = N$. A block reduction then collapses the objective onto sector boxes indexed by the numbers a of zero modes and b of positive modes, and four families of sectors are eliminated in closed form for every rank: $a \geq n$, $b \leq 2$, $a + b \leq n$, and $a = n - 1$ at every b . Everything else is reduced, exactly and in both directions, to the non-existence of interior local maxima of the sector objective, with a necessary condition any such maximum must satisfy; the residual family is finite in width at each zero-mode count and open from the sector $(n - 2, 3)$ on. For the doubling conjecture we prove an exact decomposition $S_{\max}(2N) = \max(S_{\max}(N), \Sigma_0(2N))$ that settles, by an equality, every doubled competitor whose free spectrum carries a zero mode, and reduces the flat-spectrum case to one explicit family. Finally, exact interval branch and bound certifies two-sided brackets $D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N)) \in [D^*(N) - \varepsilon, D^*(N)]$ at 22 ranks, 660 sectors, with the largest certified excess $\varepsilon = 4.382 \times 10^{-5}$. None of the three conjectures

*Corresponding author; owner of the Solution repository `scientific-author-73/interaction-distance-flat-spectrum-conjectures`.

is proved: the brackets bound the covered family, the doubling conjecture remains open even on flat spectra, and $3 - 2\sqrt{2}$ is established as the supremum of the ansatz values only.

1 Introduction

An entanglement spectrum is the list of eigenvalues of the reduced density matrix of a subsystem, sorted decreasingly. For a topologically ordered ground state it carries more information than the entanglement entropy alone: the pattern of degeneracies distinguishes phases that the entropy cannot [5]. It is therefore natural to ask, given a spectrum, whether some simple class of states can reproduce it. The class that has proven most useful here is the class of free fermion states, also called Gaussian states, whose reduced density matrices are completely determined by two-point correlation functions [1, 7]. The *fermionic interaction distance* quantifies the question: for a density matrix ρ ,

$$D_{\mathcal{F}}(\rho) = \inf_{\sigma \in \mathcal{F}} \frac{1}{2} \|\rho - \sigma\|_1, \quad (1)$$

where $\mathcal{F}$ is the set of reduced density matrices of free fermion states. Because both spectra may be sorted before comparison and the trace norm of a difference of commuting diagonal matrices is a function of the eigenvalue lists alone, $D_{\mathcal{F}}$ is a function of the entanglement spectrum. It vanishes exactly when the spectrum admits a free description, and this is how it has been used: the $\mathbb{Z}_{2^n}$ parafermion-chain, string-net and Walker–Wang fixed-point states all have $D_{\mathcal{F}} = 0$, so they admit exact free fermion descriptions, while spectra away from those fixed points do not [3, 6].

The fixed points just mentioned have maximally degenerate entanglement spectra: their nonzero eigenvalues are all equal. That makes the *flat* spectrum the natural test case, and it is the object of this paper. For an integer $N \geq 2$ let

$$\bar{\rho}(N) = \left(\underbrace{\frac{1}{N}, \dots, \frac{1}{N}}_{N \text{ entries}}, 0, 0, \dots \right) \quad (2)$$

be the flat spectrum of rank N , and write

$$n = \lfloor \log_2 N \rfloor \geq 1, \quad x = \frac{N}{2^n} \in [1, 2). \quad (3)$$

When $N = 2^n$ the spectrum $\bar{\rho}(2^n)$ is itself free: n zero-energy modes give 2^n equal atoms of weight 2^{-n} , so $D_{\mathcal{F}}(\bar{\rho}(2^n)) = 0$. When N is not a power of two the spectrum is *not* free, and the question is how far from free it is. An explicit competitor shows that it is not very far: n zero modes together with one further mode of ratio $r = x - 1$ produce a free spectrum whose

distance to $\bar{\rho}(N)$ is

$$D^*(N) = 3 - x - \frac{2}{x} = \frac{(2-x)(x-1)}{x} \in [0, 3 - 2\sqrt{2}]. \quad (4)$$

Three conjectures, posed in the course of using $D_{\mathcal{F}}$ as a diagnostic, concern this value and two related identities [3].

Conjecture 1 (flat-spectrum conjecture, (A)). *For every integer $N \geq 2$, $D_{\mathcal{F}}(\bar{\rho}(N)) = D^*(N)$.*

Conjecture 2 (zero-mode or doubling conjecture, (B)). *For every spectrum ρ , $D_{\mathcal{F}}(\frac{1}{2}(\rho \oplus \rho)) = D_{\mathcal{F}}(\rho)$.*

Conjecture 3 (global maximum, (C)). $\sup_{\rho} D_{\mathcal{F}}(\rho) = 3 - 2\sqrt{2}$.

Conjecture 1 was known exactly only for $N \leq 6$, and supported up to $N = 2^8$ by numerical minimisation over the mode energies [3, 6]. Such a minimisation cannot exclude a better free spectrum sitting in an unexplored basin, which is precisely the gap between “numerically observed” and “certified” that this paper addresses. For Conjecture 2 only the easy direction is classical: adding a single zero mode to a competitor shows $D_{\mathcal{F}}(\frac{1}{2}(\rho \oplus \rho)) \leq D_{\mathcal{F}}(\rho)$ for every ρ . For Conjecture 3 only the rank-four case was proved.

What this paper establishes

The three conjectures are aspects of one finite-dimensional variational problem, because every free spectrum is a product of two-level factors. Writing the mode ratios $r_j = p_j/(1 - p_j) \in (0, 1]$, where $p_j \in (0, \frac{1}{2}]$ is the occupation of the j -th entanglement mode, a free spectrum is the decreasing rearrangement of $\sigma_a = t_a(r)/Z(r)$, where $t_a(r)$ runs over the 2^m subset products $\prod_{j \in S} r_j$ and $Z(r) = \prod_{j=1}^m (1 + r_j)$, with the number of modes $m \geq 0$ itself free [1, 7]. Conjecture 1 is then the statement that a certain overlap functional never exceeds an explicit value, and the whole analysis becomes an exercise in controlling subset products. Sections 3 to 11 develop it; the results are:

1. **The ansatz value, derived and unique** (§4). $D^*(N)$ is obtained from the one-parameter family of n zero modes plus one further mode, and the overlap of that family is shown to be strictly increasing then strictly decreasing in the single free ratio, with its unique maximum at $r = x - 1$, where the normaliser satisfies $Z = 2^n(1 + r) = N$ exactly. The normalisation boundary is optimal, not incidental.
2. **An exact block reduction** (§5). For a spectrum with a zero modes and b positive modes the 2^{a+b} -atom objective collapses onto the 2^b subset products of the positive modes, with an explicit fractional weight on

the boundary atom, so the sector objective depends on a only through $y = N/2^a$. In subset-sum coordinates the distance splits exactly as $D = \text{Exc} + T$ into a clipped overshoot and a displaced tail, and only the first $\lceil y \rceil$ atoms carry any weight.

3. **Four infinite families of sectors eliminated for every rank** (§§6–9). Every free spectrum with at least n zero modes satisfies $D_{\mathcal{F}}(\bar{\rho}(N); \sigma) \geq D^*(N)$ with equality exactly at the ansatz; every spectrum with at most two positive modes does too; so does every spectrum with at most n modes in total; and so does the sector $a = n - 1$ for *every* b , the widest of the four, which contains the sector $(n - 1, 3)$ and is closed by a mode-count descent driven by a rank lower bound.
4. **An exact reduction of everything else** (§10). Conjecture 1 holds at every rank if and only if no sector objective has an interior local maximum above $S^*(N) = 1 - D^*(N)$, and any such maximum must satisfy an explicit algebraic condition, sharpened to $b - 1$ conditions by a concavity argument. The residual family has finitely many positive modes at each zero-mode count and is open from the sector $(n - 2, 3)$ on.
5. **Certified brackets at fixed rank** (§11, Table 1, Figure 1). Exact rational interval branch and bound discharges all 660 sectors with $0 \leq a \leq n$, $0 \leq b \leq n + 3$ at 22 ranks, giving the two-sided bracket $D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N)) \in [D^*(N) - \varepsilon_e, D^*(N)]$ with the largest certified excess $\varepsilon_e = 4.382 \times 10^{-5}$ at $N = 25$. The excess decays linearly in the box width across the five tolerances of the ladder and cannot vanish at any finite resolution, because every residual closed box contains a face point where the overlap equals S^* exactly.
6. **An exact decomposition of the doubling problem** (§12). $S_{\max}(2N) = \max(S_{\max}(N), \Sigma_0(2N))$, where Σ_0 restricts the supremum to spectra without zero modes. Every doubled sector carrying a zero mode is settled by an *equality* of overlaps, every doubled sector with at most $n + 1$ modes by coarse graining, and Conjecture 2 on flat spectra is reduced to one explicit family and derived from Conjecture 1 at the doubled rank alone.
7. **The maximal ansatz value** (§13). $\sup_N D^*(N) = 3 - 2\sqrt{2}$, attained along $x \rightarrow \sqrt{2}$; by the Main Theorem no spectrum with at least n zero modes is closer to $\mathcal{F}$ than its own flat target.

What this paper does not establish

Conjecture 1 is not proved: the sectors it leaves are exactly those in which an interior local maximum could hide, and their exclusion is supported by

a local search that can miss a narrow basin (§10, Appendix D). The certified brackets bound $D_{\mathcal{F}}^{\text{cov}}$, the distance minimised over covered sectors; transferring them to $D_{\mathcal{F}}$ itself requires a bound on the uncovered family $\mathcal{U}(N) = \{(a, b) : a \leq n - 1, b \geq n + 4\}$, which is open (§11). Conjecture 2 is not proved, not even on the flat family, and off the flat family nothing is proved: five apparent counterexamples found by unconstrained multistart optimisation are shown to be optimiser artefacts, with the exact structure of the near-optimal free spectrum identified (§12, Appendix D). Conjecture 3 is not proved for all spectra; what is proved is the corresponding statement about the ansatz values, plus the rank-four maximum of $1/6$ from the earlier literature [6].

Shape of the argument

The reduction of §5 converts the variational problem into a statement about which subset products are heavy: an atom is clipped exactly when its subset product s_S exceeds Z_b/y , and the objective is a weighted sum of clipped atoms. Two tools then do most of the work. The first is *support counting*: only $\lceil y \rceil$ atoms carry weight, which bounds the objective crudely but eliminates whole blocks of sectors. The second is *concavity in a single mode bias*: inside a chamber where the ordering of subset products is fixed, the objective is continuous, piecewise affine and concave in each p_j , with an explicitly monotone slope. Concavity converts an optimisation problem into a comparison of endpoint and critical values, which is what makes the sector-by-sector eliminations exact rather than numerical. Where concavity leaves a single hard chamber, the objective is bounded by a quartic whose sign is decided by an exact argument (Appendix B). The certificates of §11 use the same decomposition, replacing analysis with interval arithmetic on boxes and re-verifying every survivor in exact rational arithmetic.

Section 2 records where the problem comes from and what was known. Section 3 fixes the model and conventions. Sections 4–9 contain the analysis. Section 10 states the reduction to interior maxima, §11 the certificates, §12 the doubling problem and §13 the maximal value. Section 14 collects the open problems and the boundaries of every result. Appendices A–D hold the case analyses, the polynomial sign lemmas, the certification protocol and the exploratory searches.

2 Where the problem comes from

2.1 Free descriptions of topological fixed points

A state of a fermionic lattice system is *free*, or Gaussian, when its density matrix is a polynomial of even degree in the Majorana operators, equivalently when all its higher correlators are determined by the two-point function [1].

For such a state the reduced density matrix of a region is again Gaussian, and Peschel’s formula expresses it as a product of two-level factors built from the eigenvalues of the correlation matrix restricted to the region [7]. Concretely, if $\nu_1, \dots, \nu_m \in [0, 1]$ are those eigenvalues then, in the basis of eigenmodes of the correlation matrix,

$$\rho_A = \bigotimes_{j=1}^m \left(\nu_j |1\rangle\langle 1| + (1 - \nu_j) |0\rangle\langle 0| \right), \quad (5)$$

so the entanglement spectrum is the multiset of the 2^m products $\prod_{j \in S} \nu_j \prod_{j \notin S} (1 - \nu_j)$, $S \subseteq \{1, \dots, m\}$. The product form is the reason the interaction distance is computable at all: the free set is parametrised by m numbers rather than by an arbitrary density matrix.

The interaction distance was introduced to ask whether an interacting topological phase admits such a description [6]. The answer at the exactly solvable fixed points is yes: for the $\mathbb{Z}_{2^n}$ parafermion chain, the string-net and the Walker–Wang fixed-point states the entanglement spectrum is reproduced exactly by a free fermion state, $D_{\mathcal{F}} = 0$ [3, 6]. Those spectra are flat, of rank $N = 2^n$, which is why the flat spectra $\bar{\rho}(N)$ are the right objects to study: they are the fixed points themselves when N is a power of two, and their neighbourhood in N is where the free description starts to fail. Away from the fixed point the distance grows, and the observed behaviour is that flat spectra of rank $N \neq 2^n$ approach a limiting distance $3 - 2\sqrt{2} \approx 0.1716$ [6]. Extending the free description away from the fixed point, part (D) of the problem record, is outside the scope of this paper and is not addressed here [3].

2.2 Prior status of the three conjectures

Conjecture 1 is the statement that the explicit competitor of (4) is optimal. Before this work it was proved for $N \leq 6$ and supported for larger N by numerical minimisation over the mode energies with local methods, up to $N = 2^8$ [3]. Two features of that evidence matter for what follows. First, a local minimiser cannot exclude a better free spectrum in an unexplored basin, and the objective is piecewise affine rather than smooth, so gradient information is weak. Second, the observed maximisers always sit at a special point: the normaliser $Z = \prod_j (1 + r_j)$ equals N , and the largest mode ratio equals $x - 1$. Theorem 9 below shows this is a consequence of the objective, not an artefact of the search.

Conjecture 2 is the statement that doubling a spectrum — equivalently, adding a zero-energy mode — does not change the distance. The inequality $D_{\mathcal{F}}(\frac{1}{2}(\rho \oplus \rho)) \leq D_{\mathcal{F}}(\rho)$ is immediate, since a competitor for ρ doubled with a zero mode is a competitor for $\frac{1}{2}(\rho \oplus \rho)$ at the same distance. The reverse inequality is the content of the conjecture, and it was open for every ρ [3]. No systematic test of it on non-flat spectra had been reported.

Conjecture 3 asks for the worst-case distance to the free set. The value $3 - 2\sqrt{2}$ is the maximum of $D^*(N)$ over ranks, attained in the limit $x \rightarrow \sqrt{2}$; the only case proved before this work was the rank-four maximum $1/6$ [3, 6].

2.3 What is new here, and in what form

The contribution is a proof-theoretic and computational treatment of the same variational problem, organised around sectors. A free spectrum with a zero modes ($r_j = 1$) and b positive modes ($r_j < 1$) lies in the sector (a, b) ; the sectors partition $\mathcal{F}$, and the objective restricted to a sector is a function of b variables on a compact box. That decomposition is what allows exact statements:

- sector eliminations that hold *for every rank simultaneously*, rather than rank by rank (§§6–9);
- an equivalence between the conjecture and the non-existence of interior local maxima, so that the remaining work is one analytic statement rather than an unbounded search (§10);
- exact rational certificates at fixed rank, which convert a numerical optimisation into a two-sided bracket with a rational bound that can be re-verified independently (§11);
- an exact decomposition of the doubling problem that identifies the single family of spectra where it can fail (§12).

Every identity quoted in this paper is exact: the analysis is carried out in rational arithmetic, the inequalities are proved rather than sampled, and the floating-point numbers that appear are decimal renderings of exact rationals. The two exceptions are explicitly labelled as searches, and their scope is stated wherever they are used: the multistart search for interior local maxima (§10, Appendix D) and the exploratory tests of Conjecture 2 off the flat family (Appendix D).

The results were obtained twice for the hardest sector: the two-positive-mode theorem of §7 has an independent second proof in different coordinates and with a different case split, and the two proofs agree pointwise on several million exact rational grid points at every rank up to 80 [2]. Where the paper reports a number that comes from a computation rather than a proof, it says so.

3 Free spectra and the interaction distance

3.1 The free set

A spectrum is a decreasingly ordered sequence of non-negative numbers summing to one, padded with zeros when two spectra of different length

are compared. For an integer $m \geq 0$ and numbers $p_1, \dots, p_m \in (0, \frac{1}{2}]$, the product state (5) with $\nu_j = p_j$ has eigenvalues

$$\sigma_S = \prod_{j \in S} p_j \prod_{j \notin S} (1 - p_j), \quad S \subseteq \{1, \dots, m\}, \quad (6)$$

that is, 2^m numbers indexed by the subsets of the modes. It is convenient to trade the biases p_j for the *mode ratios*

$$r_j = \frac{p_j}{1 - p_j} \in (0, 1], \quad p_j = \frac{r_j}{1 + r_j}, \quad 1 - p_j = \frac{1}{1 + r_j}, \quad (7)$$

because (6) then factors as

$$\sigma_S = \frac{s_S}{Z}, \quad s_S = \prod_{j \in S} r_j, \quad Z = \prod_{j=1}^m (1 + r_j) = \sum_S s_S, \quad (8)$$

where the last equality is the expansion of the product and shows that the atoms sum to one. The restriction $p_j \leq \frac{1}{2}$ costs nothing: the two-element set $\{p, 1 - p\}$ is invariant under $p \mapsto 1 - p$, so any spectrum arising from biases in $(0, 1)$ also arises from biases in $(0, \frac{1}{2}]$.

Definition 4 (free spectra). $\mathcal{F}$ is the set of all decreasing rearrangements of the lists (8), over all $m \geq 0$ and all $r_1, \dots, r_m \in (0, 1]$. The case $m = 0$ gives the one-atom spectrum (1).

By Peschel's formula every reduced density matrix of a free fermion state has this form [7], and by the Gaussian structure of free fermion states every such form is realised [1]. Definition 4 is therefore the complete description of the free set at the level of spectra, which is the level at which (1) lives. The mode ratios carry the physics: $r_j = 1$ means $p_j = \frac{1}{2}$, which means the j -th entanglement energy $\log(1/r_j)$ vanishes, so such a mode contributes no information and doubles every atom. This is the case that makes Conjecture 2 interesting.

Definition 5 (sectors). A spectrum $\sigma \in \mathcal{F}$ arising from m modes with a of the ratios equal to 1 and $b = m - a$ of them strictly below 1 lies in the sector (a, b) . The ratios of a sector are written $s_1 \geq s_2 \geq \dots \geq s_b$ in $(0, 1)$ for the positive modes and 1 for the a zero modes, and $Z_b = \prod_{j=1}^b (1 + s_j)$.

The sectors partition $\mathcal{F}$: every spectrum has a well-defined number of zero modes and of positive modes. Note that a spectrum may have several representations — for instance $(a, b) = (0, 1)$ with $s_1 = 1$ is the same spectrum as $(a, b) = (1, 0)$ — and that the boundary cases $s_j \in \{0, 1\}$ move a spectrum between sectors. This is harmless for suprema, since the objective is continuous in the ratios on the closed box $[0, 1]^b$, but it is the reason the face identities of §10 matter: the boundary of a sector box consists of other sectors with one mode fewer.

3.2 The overlap functional

For probability vectors ρ and σ the trace distance has an exact reformulation as an overlap.

Lemma 6 (overlap identity). *For any two probability vectors ρ and σ ,*

$$\frac{1}{2}\|\rho - \sigma\|_1 = 1 - \sum_{a \geq 1} \min(\rho_a, \sigma_a) = 1 - S(\rho, \sigma). \quad (9)$$

Proof. $|u - v| = u + v - 2 \min(u, v)$, so $\frac{1}{2} \sum_a |\rho_a - \sigma_a| = \frac{1}{2} (\sum_a \rho_a + \sum_a \sigma_a) - \sum_a \min(\rho_a, \sigma_a) = 1 - S(\rho, \sigma)$. $\square$

Lemma 7 (rearrangement invariance). *If ρ and σ are decreasingly ordered then $S(\rho, \sigma) = \max_{\pi} \sum_a \min(\rho_a, \sigma_{\pi(a)})$, the maximum over permutations π of the indices.*

Proof. Let π be a permutation that is not the identity, so there are indices $i < j$ with $\sigma_{\pi(i)} < \sigma_{\pi(j)}$. Since $\rho_i \geq \rho_j$, writing $u = \sigma_{\pi(j)} \geq v = \sigma_{\pi(i)}$ the elementary inequality $\min(\rho_i, v) + \min(\rho_j, u) \leq \min(\rho_i, u) + \min(\rho_j, v)$ holds in each of the four orderings of ρ_j, v, u, ρ_i . Swapping the two images therefore does not decrease the sum, and finitely many swaps reach the identity. $\square$

Together, Lemma 6 and Lemma 7 say that $D_{\mathcal{F}}$ is a function of the two sorted spectra, and that minimising the distance is the same as maximising the overlap. For the flat target this gives the central quantity of the paper.

Definition 8 (overlap supremum). For an integer $M \geq 1$,

$$S_{\max}(M) = \sup_{\sigma \in \mathcal{F}} S(\bar{\rho}(M), \sigma), \quad \text{so that} \quad D_{\mathcal{F}}(\bar{\rho}(M)) = 1 - S_{\max}(M). \quad (10)$$

More generally, for a subset $\mathcal{S} \subseteq \mathcal{F}$, $S_{\max}(M; \mathcal{S}) = \sup_{\sigma \in \mathcal{S}} S(\bar{\rho}(M), \sigma)$, and $\Sigma_0(M) = S_{\max}(M; \{\sigma \in \mathcal{F} : a(\sigma) = 0\})$ is the same supremum restricted to spectra with no zero mode, where $a(\sigma)$ is the number of zero modes in any representation of σ .

The identity $D_{\mathcal{F}}(\bar{\rho}(M)) = 1 - S_{\max}(M)$ in (10) uses that the infimum over $\mathcal{F}$ of $1 - S$ equals 1 minus the supremum of S ; the supremum need not be attained, and where this paper says that a value is attained it means that an explicit competitor realises it.

Conjecture 1 is therefore the single inequality

$$S(\bar{\rho}(N), \sigma) \leq S^*(N) := 1 - D^*(N) = x + \frac{2}{x} - 2 \quad \text{for every } \sigma \in \mathcal{F}, \quad (11)$$

together with the observation that equality is achieved, so that the supremum is a maximum. Note that $S^*(N) \in [2\sqrt{2} - 2, 1]$: it equals 1 at $x = 1$ and tends to 1 as $x \rightarrow 2$, with its minimum $2\sqrt{2} - 2$ at $x = \sqrt{2}$. Since $1 - S^*$ is the distance, the hardest ranks are those with x near $\sqrt{2}$.

3.3 Conventions

The following are fixed for the rest of the paper. Spectra are decreasingly ordered probability vectors, compared at equal length by padding the shorter one with zeros. A *zero mode* is a ratio $r_j = 1$; a *positive mode* is a ratio $r_j < 1$; the *mode bias* is $p_j = r_j/(1 + r_j) \in (0, \frac{1}{2}]$ and the *mode energy* is $\log(1/r_j) \geq 0$. The symbols N , n and x always mean the rank, $\lfloor \log_2 N \rfloor$ and $N/2^n$ of (3); $D^\star$ and $S^\star$ always mean the ansatz distance (4) and its complement $1 - D^\star$. Ranks of atoms start at 1. $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ are the ceiling and floor, and $\text{clamp}(z, 0, 1) = \max(0, \min(1, z))$ is the template's clamp macro. For a sector (a, b) we write

$$y = \frac{N}{2^a} = 2^{n-a}x = 2^\nu x, \quad \nu = n - a \geq 0, \quad (12)$$

so ν is the *zero-mode deficit* and y is the flat level, rescaled: in a sector the objective compares atoms against $1/y$ rather than against $1/N$. Finally, $S(a, b, s)$ denotes the sector objective defined in §5, and $S(y, b, s)$ the same quantity viewed as a function of y ; the two notations coincide because the objective depends on the sector only through y and the positive ratios.

4 The ansatz value and the normalisation boundary

The competitor that produces $D^\star(N)$ is the simplest one that is not already flat: n zero modes, which supply the 2^n -fold degeneracy of the $\mathbb{Z}_{2^n}$ fixed point, together with a single further mode of ratio $r \in [0, 1]$. Its ratios are $(1, \dots, 1, r)$, its normaliser is $Z = 2^n(1+r)$, and its atoms are the $2^m = 2^{n+1}$ subset products of (8), which take only two values: 2^n atoms equal to $1/Z$, from the subsets that omit the last mode, and 2^n atoms equal to r/Z , from those that contain it. Since $r \leq 1$ the decreasing rearrangement puts the $1/Z$ atoms first. The whole family is therefore explicit, and its overlap with the flat target can be written down.

Theorem 9 (the ansatz is the unique optimum of its own family). *Fix $N \geq 2$, and let $n = \lfloor \log_2 N \rfloor$, $x = N/2^n$. For $\sigma^{(r)} \in \mathcal{F}$ the spectrum with ratios $(1^n, r)$, $r \in [0, 1]$,*

$$S(\bar{\rho}(N), \sigma^{(r)}) = \begin{cases} \frac{1}{x} + \frac{(x-1)r}{1+r}, & 0 \leq r \leq x-1, \\ \frac{1+(x-1)r}{1+r}, & x-1 \leq r \leq 1. \end{cases} \quad (13)$$

Both branches are strictly monotone towards $r = x-1$: the first increases, the second decreases. Hence

$$\max_{0 \leq r \leq 1} S(\bar{\rho}(N), \sigma^{(r)}) = S(\bar{\rho}(N), \sigma^{(x-1)}) = x + \frac{2}{x} - 2 = S^\star(N), \quad (14)$$

the maximum is unique except at $x = 1$, where it is attained at $r = 0$, and at the maximiser the normaliser sits exactly on the rank,

$$Z = 2^n(1 + (x - 1)) = 2^n x = N. \quad (15)$$

Consequently $D_{\mathcal{F}}(\bar{\rho}(N); \sigma^{(x-1)}) = D^*(N)$, so $D_{\mathcal{F}}(\bar{\rho}(N)) \leq D^*(N)$, and $D^*(N) = 0$ if and only if $N = 2^n$.

Proof. The atoms of $\sigma^{(r)}$ are 2^n copies of $1/Z$ followed by 2^n copies of r/Z , with $Z = 2^n(1 + r)$. The flat target has $N = x2^n$ entries equal to $1/N$, so the overlap receives contributions from the first 2^n atoms, from the next $N - 2^n = (x - 1)2^n$ atoms, and from nothing beyond:

$$S = 2^n \min\left(\frac{1}{N}, \frac{1}{Z}\right) + (x - 1)2^n \min\left(\frac{1}{N}, \frac{r}{Z}\right). \quad (16)$$

The second minimum is never clipped. Indeed $r/Z \leq 1/N$ is equivalent to $rN \leq Z = 2^n(1 + r)$, that is to $rx \leq 1 + r$, that is to $r(x - 1) \leq 1$, which holds for every $r \leq 1$ because $x - 1 < 1$. So the second term of (16) equals $(x - 1)2^n r/Z = (x - 1)r/(1 + r)$ throughout.

The first minimum splits the family. Since $1/Z \leq 1/N$ is equivalent to $Z \geq N$, that is to $1 + r \geq x$, the first term equals $2^n/Z = 1/(1 + r)$ when $r \geq x - 1$ and $2^n/N = 1/x$ when $r \leq x - 1$. Substituting gives (13).

For monotonicity, differentiate in r . On the first branch $\frac{d}{dr} \left[\frac{1}{x} + \frac{(x-1)r}{1+r} \right] = \frac{x-1}{(1+r)^2}$, which is positive for $x > 1$ and identically zero for $x = 1$, where the branch degenerates to the single point $r = 0$. On the second branch

$$\frac{d}{dr} \left[\frac{1 + (x-1)r}{1+r} \right] = \frac{(x-1)(1+r) - (1 + (x-1)r)}{(1+r)^2} = \frac{x-2}{(1+r)^2} < 0, \quad (17)$$

strictly, because $x < 2$. The two branches agree at $r = x - 1$, where the value is $\frac{1}{x} + \frac{(x-1)^2}{x} = \frac{1+(x-1)^2}{x} = \frac{x^2-2x+2}{x} = x + \frac{2}{x} - 2$, which is $S^*(N)$. Uniqueness of the maximiser follows from the strict monotonicity on both sides; at $x = 1$ the first branch is the single point $r = 0$ and the second is strictly decreasing, so the maximum is again unique and equals $1 = S^*$.

At $r = x - 1$ the normaliser is $Z = 2^n x = N$, which is (15), and the atoms are 2^n copies of $1/N$ followed by 2^n copies of $(x - 1)/N$. The distance to the flat target is then computed atom by atom: the first 2^n atoms coincide with the target and contribute nothing; the next $(x - 1)2^n$ atoms differ by $\frac{1}{N} - \frac{x-1}{N} = \frac{2-x}{N}$ each; the remaining $(2 - x)2^n$ atoms lie beyond the support of the target and contribute $\frac{x-1}{N}$ each. Adding,

$$\frac{1}{2} \|\bar{\rho}(N) - \sigma^{(x-1)}\|_1 = \frac{2(x-1)(2-x)2^n}{2N} = \frac{(2-x)(x-1)}{x} = D^*(N), \quad (18)$$

the factor 2 collecting the two equal excess terms, the $(x - 1)2^n$ atoms just below $\bar{\rho}_1$ and the $(2 - x)2^n$ atoms beyond $\bar{\rho}_N$, and the last step using $N = x2^n$. This is also $1 - S^*(N)$, as it must be by Lemma 6. Since $D^* = (2 - x)(x - 1)/x$ and $x \in [1, 2)$, it vanishes exactly at $x = 1$, that is at $N = 2^n$. $\square$

Three features of Theorem 9 are worth separating, because the rest of the paper uses each of them differently.

The upper bound. The ansatz is a competitor, so

$$D_{\mathcal{F}}(\bar{\rho}(N)) \leq D^*(N), \quad \text{equivalently} \quad S_{\max}(N) \geq S^*(N), \quad (19)$$

for every $N \geq 2$, with no further hypothesis. Conjecture 1 is exactly the reverse inequality, and (11) is the form in which it is attacked from §5 onwards.

The boundary is optimal, not incidental. Numerical minimisation over mode energies consistently returns maximisers with $Z = N$ and largest ratio $x - 1$, and it was not clear whether that was a property of the objective or an artefact of the search [3]. Theorem 9 settles it inside this family: the overlap is strictly increasing while the normaliser is below N and strictly decreasing once it is above, so $Z = N$ is the unique turning point. The mechanism is visible in (16). Below the boundary the first 2^n atoms are clipped at the flat level and contribute the constant $1/x$, so only the second group improves as r grows; above it the first group loses mass faster than the second gains, at the rate $x - 2 < 0$ of (17). Whether an analogous boundary statement holds for non-flat targets, and what replaces N there, is open (§14).

A second proof for powers of two. At $N = 2^n$ the value $D^* = 0$ also follows directly, without Theorem 9: the spectrum with n zero modes and nothing else has $Z = 2^n$ and all 2^n atoms equal to 2^{-n} , so it is $\bar{\rho}(2^n)$ and the distance is zero. This is the free description of the $\mathbb{Z}_{2^n}$ fixed point [6], and it is the base case from which the eliminations of §6 grow.

Finally, the value $D^*(N)$ is bounded uniformly, and the bound is what Conjecture 3 proposes as the worst case over *all* spectra.

Corollary 10 (the ansatz values are bounded by $3 - 2\sqrt{2}$). *For every $N \geq 2$, $0 \leq D^*(N) \leq 3 - 2\sqrt{2}$, with equality on the right if and only if $x = N/2^n = \sqrt{2}$. Equivalently $2\sqrt{2} - 2 \leq S^*(N) \leq 1$.*

Proof. $D^* = 3 - x - \frac{2}{x}$, and the arithmetic–geometric mean inequality gives $x + \frac{2}{x} \geq 2\sqrt{x \cdot \frac{2}{x}} = 2\sqrt{2}$, with equality exactly when $x = \frac{2}{x}$, that is $x =$

$\sqrt{2}$. The lower bound $D^* \geq 0$ is the factorisation $(2-x)(x-1)/x$ with $x \in [1, 2)$. $\square$

Because x takes the values $N/2^{\lfloor \log_2 N \rfloor}$, the supremum of D^* over ranks is $3 - 2\sqrt{2}$ but it is attained exactly only when $N/2^n = \sqrt{2}$, which never happens for an integer N ; the value is approached along any sequence of ranks with $N/2^{\lfloor \log_2 N \rfloor} \rightarrow \sqrt{2}$, for instance $N = \lfloor 2^k \sqrt{2} \rfloor$. This is developed in §13.

5 The block reduction and subset-sum coordinates

Conjecture 1 asks for a bound on $S(\bar{\rho}(N), \sigma)$ that holds for every $\sigma \in \mathcal{F}$, and $\mathcal{F}$ is a union of sectors (Definition 5). The first step is to show that inside a sector the objective depends on a small number of quantities, so that the sector becomes a compact box in b variables.

5.1 The block reduction

Fix a sector (a, b) and positive ratios $s = (s_1, \dots, s_b) \in [0, 1]^b$, decreasing. Its spectrum σ has $m = a + b$ modes, a of them with ratio 1, so its normaliser is $Z = 2^a Z_b$ with $Z_b = \prod_{j \leq b} (1 + s_j)$, and its 2^{a+b} atoms are the subset products s_S/Z . Every subset product of the positive ratios occurs once for each of the 2^a subsets of the zero modes, all of which contribute the factor 1. Hence the atom multiset is the 2^b numbers

$$\tau'_S = \frac{s_S}{Z_b}, \quad S \subseteq \{1, \dots, b\}, \quad (20)$$

each with multiplicity 2^a , and the decreasing rearrangement τ' of the list (20) satisfies $\sigma_i = \tau'_{\lceil i/2^a \rceil}$: the spectrum is the reduced spectrum with every entry repeated 2^a times. Set

$$y = \frac{N}{2^a}, \quad w_i(y) = \text{clamp}(y - i + 1, 0, 1). \quad (21)$$

Theorem 11 (block reduction). *For every $N \geq 2$, every sector (a, b) and every $s \in [0, 1]^b$,*

$$S(\bar{\rho}(N), \sigma) = \sum_{i \geq 1} w_i(y) \min\left(\tau'_i, \frac{1}{y}\right) =: S(a, b, s) = S(y, b, s), \quad (22)$$

where the sum has at most $\lceil y \rceil$ nonzero terms.

Proof. Only the first N positions of σ meet the support of $\bar{\rho}(N)$, so $S = \sum_{i \leq N} \min(1/N, \sigma_i)$. Substitute $\sigma_i = \tau'_{\lceil i/2^a \rceil}$ and $1/N = 2^{-a}/y$, and group the indices by $j = \lceil i/2^a \rceil$. The group j consists of $i = (j-1)2^a + 1, \dots, j2^a$, of

which those with $i \leq N = 2^a y$ number exactly $\text{clamp}(2^a y - (j-1)2^a, 0, 2^a) = 2^a w_j(y)$. Therefore

$$S = \sum_{j \geq 1} 2^a w_j(y) \min\left(\frac{2^{-a}}{y}, \tau'_j\right) = \sum_{j \geq 1} w_j(y) \min\left(\frac{1}{y}, \tau'_j\right). \quad (23)$$

Finally $w_j(y) = 0$ as soon as $y - j + 1 \leq 0$, that is for $j \geq y + 1$, hence for $j > \lceil y \rceil$. $\square$

Two consequences organise everything that follows. First, the sector objective depends on the zero-mode count only through $y = N/2^a$, which is why the same analysis applies at every rank once it is parametrised by y : this is the meaning of the second equality in (22). Second, the weights implement the truncation at rank N exactly, with a *fractional* weight on the boundary atom when y is not an integer. The fractional weight is not a technicality: it is what makes the block form an equality rather than an inequality, and it is what the ansatz of Theorem 9 exploits when $Z = N$.

The quantity that Theorem 11 must be compared against also simplifies in the y -parametrisation. Writing $\nu = n - a$ for the zero-mode deficit, so $y = 2^\nu x$,

$$S^*(N) = x + \frac{2}{x} - 2 = \frac{y}{2^\nu} + \frac{2^{\nu+1}}{y} - 2 =: S_\nu^*(y), \quad (24)$$

and Conjecture 1 restricted to the sector (a, b) is the statement that $S(y, b, s) \leq S_\nu^*(y)$ for every $s \in [0, 1]^b$, where $y \in [2^\nu, 2^{\nu+1})$. At $\nu = 0$ this reads $S(y, b, s) \leq y + 2/y - 2$ with $y = x \in [1, 2)$; at $\nu = 1$, which is where most of the work of §§7–9 happens, $y = 2x \in [2, 4)$ and the target is $y/2 + 4/y - 2$.

5.2 Subset-sum coordinates

The atoms (20) are products, and products are sums in logarithmic coordinates. Put

$$t_j = -\log s_j \in [0, \infty], \quad T_S = \sum_{j \in S} t_j, \quad \tau'_S = \frac{e^{-T_S}}{Z_b}, \quad (25)$$

with $t_j = +\infty$ when $s_j = 0$, meaning the atom vanishes. Then:

- the decreasing rearrangement of $\{\tau'_S\}$ is the ordering of the 2^b subsets by *increasing* subset-sum T_S ;
- the atom S is *clipped*, meaning $\min(\tau'_S, 1/y) = 1/y$, exactly when $T_S \leq \log(y/Z_b)$;

- the ordering of the atoms changes only on the walls $T_S = T_{S'}$, which are hyperplanes in the t -coordinates, so the objective is piecewise smooth with explicitly describable pieces.

A *chamber* is a region of the box on which the ordering of the subset-sums is fixed, and a *cell* is a region on which the clipping pattern — the set of clipped atoms — is fixed. Chambers are cut by the walls t -sum equalities and cells by the clipping thresholds; both are finite in number for fixed b , and on each of them the objective is a sum of affine functions of the biases $p_j = s_j/(1 + s_j)$, because every atom $\tau'_S = \prod_{j \in S} p_j \prod_{j \notin S} (1 - p_j)$ is affine in each p_j separately. That affineness is the content of Lemma 24 and it is what makes the sector analysis exact rather than numerical.

5.3 What the objective penalises

It is useful to know what a free spectrum loses relative to the flat target, and the loss splits into two exactly computable pieces.

Proposition 12 (excess and tail). *Let $\rho = \bar{\rho}(N)$ and let σ be any probability vector. With $J = \#\{i : \sigma_i \geq 1/N\}$ the number of clipped atoms,*

$$D_{\mathcal{F}}(\rho; \sigma) = \frac{1}{2} \|\rho - \sigma\|_1 = \underbrace{\sum_{i \geq 1} \left(\sigma_i - \frac{1}{N}\right)^+}_{\text{Exc} \geq 0} + \underbrace{\sum_{i > N} \min\left(\sigma_i, \frac{1}{N}\right)}_{T \geq 0}, \quad (26)$$

and moreover $J \leq N$, while in the reduced coordinates of a sector (a, b) the number of clipped atoms of τ' satisfies $J \leq y$.

Proof. Summing the identity $\min(u, v) = u + v - |u - v|$ over all i gives $S(\rho, \sigma) + T = 1 - \text{Exc}$, because $\sum_i \min(1/N, \sigma_i)$ over all i equals $\sum_{i \leq N} \min(1/N, \sigma_i) + T = S + T$ on the one hand and $\sum_i \sigma_i - \sum_i (\sigma_i - 1/N)^+ = 1 - \text{Exc}$ on the other. Lemma 6 then gives $\frac{1}{2} \|\rho - \sigma\|_1 = 1 - S = \text{Exc} + T$. The counts follow from $\sum_i \sigma_i = 1$: each clipped atom carries at least $1/N$, so $J \leq N$, and in a sector each clipped atom of τ' carries at least $1/y$ while $\sum_i \tau'_i = 1$, so $J \leq y$. $\square$

Proposition 12 says that a competitor pays twice for being unequal: once for every atom it pushes above the flat level, and once for every atom it fails to place inside the first N positions. The ansatz of Theorem 9 pays nothing in excess, since $Z = N$ puts its first 2^n atoms exactly at the level, and pays $D^*(N)$ entirely in displaced tail. That is the structure the sector analysis has to reproduce: a competitor that clips atoms must displace at least as much mass as it gains, and the eliminations of §§6–9 are all forms of that statement.

The count $J \leq y$ is the crudest but most widely useful version. It immediately gives a bound that holds in every sector.

Lemma 13 (support and clipping bounds). *For every sector (a, b) and every $s \in [0, 1]^b$,*

$$S(y, b, s) \leq \min\left(1, \frac{2^b}{y}\right), \quad (27)$$

because at most $\min(2^b, y)$ atoms can be clipped and each contributes at most $1/y$. If in addition $2^b \leq y$ then every weight w_i with $i \leq 2^b$ equals 1 and $S(y, b, s) \leq 2^b/y$ follows from $\sum_i \tau'_i = 1$ instead.

Proof. The first bound is $\sum_i w_i \min(\tau'_i, 1/y) \leq \min(\sum_i w_i \tau'_i, \frac{1}{y} \sum_i w_i) \leq \min(1, 2^b/y)$, using $\sum_i w_i = \min(y, 2^b)$ and $\tau'_i \leq 1$. The second is $S \leq \sum_i \tau'_i = 1$ when $2^b \leq y$, refined to $2^b/y$ by clipping: with all 2^b weights equal to 1, $S = \sum_i \min(\tau'_i, 1/y) \leq 2^b/y$. $\square$

Lemma 13 is the whole content of some eliminations and useless in others: it beats $S_\nu^*(y)$ exactly when $2^b \leq y S_\nu^*(y)$, which fails for large b at small ν . The rest of the paper is devoted to what happens when it fails.

6 Eliminating the small and the symmetric sectors

This section disposes of every sector that the crude bounds of Lemma 13 and the special structure of the two extreme zero-mode counts can reach. The statements are uniform in the rank: each holds for every $N \geq 2$ at once, with the rank entering only through $x = N/2^n$ and $y = N/2^a$. Throughout, the target to beat is $S^*(N) = x + 2/x - 2$, written $S_\nu^*(y)$ in the y -parametrisation of (24).

6.1 Too many zero modes

Lemma 14 (Lemma A: at least $n + 1$ zero modes). *If $a \geq n + 1$ then $y \leq x/2 < 1$, the only nonzero weight is $w_1 = y$, and*

$$S(a, b, s) = y \min\left(\tau'_1, \frac{1}{y}\right) = \min(y\tau'_1, 1) \leq y \leq \frac{x}{2} \leq S^*(N), \quad (28)$$

with $S^*(N) - x/2 = (x - 2)^2/(2x) > 0$. In particular the inequality is strict, so no spectrum with at least $n + 1$ zero modes is optimal unless $S^*(N) = 1$.

Proof. $y = N/2^a \leq N/2^{n+1} = x/2 < 1$, so $w_1 = \text{clamp}(y, 0, 1) = y$ and $w_i = 0$ for $i \geq 2$ because $y - i + 1 < 0$. Theorem 11 gives $S = y \min(\tau'_1, 1/y) = \min(y\tau'_1, 1) \leq y\tau'_1 \leq y$, using $\tau'_1 \leq 1$. Finally $S^*(N) - \frac{x}{2} = x + \frac{2}{x} - 2 - \frac{x}{2} = \frac{x}{2} + \frac{2}{x} - 2 = \frac{x^2 - 4x + 4}{2x} = \frac{(x-2)^2}{2x}$, which is positive because $x < 2$. $\square$

The content of Lemma 14 is that a competitor with very many zero modes cannot use its degeneracy: the truncation at rank N cuts the spectrum after less than one block of 2^a equal atoms, so the overlap is at most y , and y is small precisely because a is large.

6.2 No positive modes

Lemma 15 (Lemma B: $b = 0$). *If $b = 0$ then $\tau' = (1)$ and*

$$S(a, 0, \cdot) = \text{clamp}(y, 0, 1) \min\left(1, \frac{1}{y}\right) = \min\left(1, \frac{2^a}{N}\right) \leq \frac{1}{x} \leq S^*(N), \quad (29)$$

where $S^*(N) - 1/x = (x - 1)^2/x \geq 0$, with equality only at $x = 1$.

Proof. With no positive mode $Z_b = 1$ and the reduced spectrum is the single atom $\tau'_1 = 1$, so Theorem 11 gives $S = w_1 \min(1, 1/y)$. If $y \geq 1$ this is $1/y = 2^a/N$; if $y < 1$ it is $y = 2^a/N$; in either case $S = \min(1, 2^a/N)$, and for $a \leq n$ that is at most $2^n/N = 1/x$. The identity $x + 2/x - 2 - 1/x = x + 1/x - 2 = (x - 1)^2/x$ finishes the proof. $\square$

At $x = 1$, that is at $N = 2^n$, Lemma 15 gives $S = 1 = S^*$: the flat spectrum of power-of-two rank is free, which is the fixed-point statement of §2. For every other rank the bound is strict.

6.3 One positive mode, in closed form

The sector $b = 1$ is the ansatz sector of Theorem 9 when $a = n$, and it is worth solving completely for every a .

Lemma 16 (Lemma D: $b = 1$). *Let $b = 1$, $Z = 1 + s_1 \in [1, 2]$, $c = 1/y$ and $\lambda = \text{clamp}(y - 1, 0, 1)$. Then*

$$S = \min\left(\frac{1}{Z}, c\right) + \lambda \min\left(\frac{Z-1}{Z}, c\right) \leq \begin{cases} y - 2 + \frac{2}{y}, & 1 \leq y \leq 2, \\ \frac{2}{y}, & y \geq 2, \end{cases} \quad (30)$$

where the first branch is $(1 + (y - 1)^2)/y$, and in the first case equality holds if and only if $Z = y$.

Proof. The reduced spectrum is $\tau' = (\frac{1}{Z}, \frac{Z-1}{Z})$ and the weights are $w_1 = \text{clamp}(y, 0, 1) = 1$ for $y \geq 1$ and $w_2 = \lambda$, which gives (30).

Suppose first $1 \leq y \leq 2$, so $\lambda = y - 1 \in [0, 1]$ and $c = 1/y \geq \frac{1}{2}$. If $Z \leq y$ then $1/Z \geq c$, the first term is c , and $1 - 1/Z \leq 1 - 1/y = (y - 1)/y \leq 1/y = c$ because $y \leq 2$; hence $S \leq \frac{1}{y} + (y - 1)\frac{y-1}{y} = \frac{1+(y-1)^2}{y}$, with equality only if $Z = y$. If $Z \geq y$ there are two sub-cases. When $1 - 1/Z \leq c$ we have $S = \frac{1}{Z} + (y - 1)(1 - \frac{1}{Z}) = y - 1 + \frac{2-y}{Z} \leq y - 1 + \frac{2-y}{y} = \frac{1+(y-1)^2}{y}$, using $Z \geq y$ and $2 - y \geq 0$, again with equality only at $Z = y$. When $1 - 1/Z > c$, equivalently $Z > y/(y - 1)$, the second term is λc and $S = \frac{1}{Z} + \frac{y-1}{y} < \frac{y-1}{y} + \frac{y-1}{y} = \frac{2(y-1)}{y}$, while $\frac{1+(y-1)^2}{y} - \frac{2(y-1)}{y} = \frac{(y-2)^2}{y} \geq 0$.

Now suppose $y \geq 2$, so $\lambda = 1$. If $Z \leq y$ then $1/Z \geq c$ and $S = c + \min(1 - \frac{1}{Z}, c) \leq \frac{2}{y}$. If $Z \geq y$ then $1/Z \leq c$, and $1 - 1/Z \geq 1 - 1/y \geq 1/y = c$ because $y \geq 2$, so $S = \frac{1}{Z} + c \leq \frac{1}{y} + \frac{1}{y} = \frac{2}{y}$. $\square$

Lemma 16 is sharp in the sector that matters most. At $a = n$ one has $y = x < 2$, so the first bound applies and equals $S_0^*(x) = x + 2/x - 2 = S^*(N)$ exactly, attained at $Z = x$, that is at $s_1 = x - 1$: this is Theorem 9 seen from inside the sector $(n, 1)$. At $a \leq n - 1$ one has $y \geq 2x$, so the second bound gives $S \leq 2/y \leq 1/x \leq S^*(N)$ as in Lemma 15.

6.4 Exactly n zero modes, any number of positive modes

Theorem 17 (Theorem E: the sector $a = n$). *Let $a = n$, so $y = x < 2$, and let $b \geq 1$. Then only the two largest reduced atoms carry weight, they are $\tau'_1 = 1/Z_b$ and $\tau'_2 = s_1/Z_b$, and*

$$S(n, b, s) \leq S(n, 1, s_1) \leq \frac{1 + (x - 1)^2}{x} = S^*(N), \quad (31)$$

with equality only if $b = 1$ and $1 + s_1 = x$.

Proof. Since $x < 2$, $w_1 = \text{clamp}(x, 0, 1) = 1$, $w_2 = \text{clamp}(x - 1, 0, 1) = x - 1$ and $w_i = 0$ for $i \geq 3$, so Theorem 11 gives $S = \min(\tau'_1, c) + (x - 1) \min(\tau'_2, c)$ with $c = 1/x$. The largest subset product is 1, from $S = \emptyset$, and the second largest is s_1 : every other subset product is either $s_j \leq s_1$ or a product of at least two ratios, hence at most $s_1 s_2 \leq s_1$. So $\tau'_1 = 1/Z_b$ and $\tau'_2 = s_1/Z_b$.

Write $Z_b = (1 + s_1)Q$ with $Q = \prod_{j \geq 2} (1 + s_j) \geq 1$. Both τ'_1 and τ'_2 are non-increasing in Q at fixed s_1 , and $u \mapsto \min(u, c)$ is non-decreasing, so S is non-increasing in Q . Its value at $Q = 1$ is exactly $S(n, 1, s_1)$, which Lemma 16 bounds by $(1 + (x - 1)^2)/x = S^*(N)$, with equality only when $1 + s_1 = x$. Equality in (31) therefore forces $Q = 1$, so $s_j = 0$ for all $j \geq 2$, and $1 + s_1 = x$. $\square$

The mechanism is worth stating separately, because it recurs: at $a = n$ the truncation is so severe that only two atoms are visible, and both of them scale as $1/Z_b$. Adding positive modes can then only *increase* the normaliser and hence *decrease* the overlap, so the sector collapses onto its $b = 1$ sub-sector. Nothing like this happens at $a \leq n - 1$, where up to $\lceil y \rceil \geq 2x$ atoms are visible and extra modes can genuinely help.

6.5 Few modes in total

Corollary 18 (Corollary F: $a + b \leq n$). *If $a + b \leq n$ then $2^b \leq y$, and*

$$S(a, b, s) \leq \frac{2^b}{y} \leq \frac{1}{x} \leq S^*(N). \quad (32)$$

Proof. $b \leq n - a = \nu$, so $2^b \leq 2^\nu \leq 2^\nu x = y$, and Lemma 13 gives $S \leq 2^b/y \leq 2^\nu/y = 1/x \leq S^*(N)$ by Lemma 15's algebraic identity. $\square$

Corollary 18 says that a competitor cannot win by being small: with at most n modes in total there is not enough mass to fill the first N positions above the flat level. The interesting sectors are therefore those with $a + b \geq n + 1$, and among them those with $a \leq n - 1$, since $a \geq n$ is Theorem 17.

6.6 The main theorem for spectra with many zero modes

Collecting Lemma 14 and Theorem 17 gives the first infinite family on which Conjecture 1 is a theorem.

Theorem 19 (Main Theorem, first family). *For every $N \geq 2$ and every free spectrum σ with at least $n = \lfloor \log_2 N \rfloor$ zero modes,*

$$D_{\mathcal{F}}(\bar{\rho}(N); \sigma) \geq 3 - \frac{N}{2^n} - \frac{2^{n+1}}{N} = D^*(N), \quad (33)$$

with equality exactly at the two-level-factor ansatz of Theorem 9: n zero modes and one further mode of ratio $x - 1$, possibly together with modes of ratio 0. Consequently $D_{\mathcal{F}}(\bar{\rho}(N)) = D^*(N)$ for every spectrum in this family, the ansatz is its unique maximiser, and $D_{\mathcal{F}}(\bar{\rho}(2^n)) = 0$.

Proof. A spectrum with at least n zero modes lies in a sector (a, b) with $a \geq n$. If $a \geq n + 1$, Lemma 14 gives $S \leq x/2 < S^*(N)$ strictly, so $D_{\mathcal{F}}(\bar{\rho}(N); \sigma) = 1 - S > D^*(N)$. If $a = n$, Theorem 17 gives $S \leq S^*(N)$ with equality only when $b = 1$ and $1 + s_1 = x$ (or $b = 0$ and $x = 1$, which is the same spectrum with $s_1 = 0$), that is only at the ansatz. Lemma 6 converts the overlap bound into (33). The value $D^*(N)$ is attained by the ansatz (Theorem 9), so the infimum in (1) restricted to this family is a minimum and equals $D^*(N)$; at $N = 2^n$ that value is 0. $\square$

6.7 What is left

After Lemma 14, Lemma 15, Lemma 16, Theorem 17, Corollary 18 and Lemma 13, the sectors that no bound has reached are

$$a \leq n - 1, \quad b \geq 3, \quad a + b \geq n + 1, \quad 2^b > yS_{\nu}^*(y), \quad (34)$$

the last condition being the failure of Lemma 13. In terms of the zero-mode deficit $\nu = n - a \geq 1$ this is $b \geq \max(3, \nu + 1)$. The next three sections remove $b \leq 2$ everywhere (§7), $b = 3$ at $\nu = 1$ (§8) and $b \geq \lceil y \rceil$ at every ν (§9), which together settle the whole of $\nu = 1$ and leave, at each $\nu \geq 2$, a finite range of b .

7 Two positive-energy modes

After §6 the only sector with two positive modes that no bound has reached is $a = n - 1$. Indeed $b = 2$ with $a \geq n + 1$ is Lemma 14, with $a = n$ is

Theorem 17, and with $a \leq n - 2$ satisfies $a + b \leq n$ and so is Corollary 18. In the sector $(n - 1, 2)$ the zero-mode deficit is $\nu = 1$, so

$$y = \frac{N}{2^{n-1}} = 2x \in [2, 4), \quad S_1^*(y) = \frac{y}{2} + \frac{4}{y} - 2 = S^*(N), \quad (35)$$

and the weights of Theorem 11 are $(1, 1, y-2, 0)$ for $2 \leq y \leq 3$ and $(1, 1, 1, y-3)$ for $3 \leq y < 4$: at most four atoms are ever visible, and at $b = 2$ there are exactly four. This is the sector in which the whole difficulty of Conjecture 1 concentrates at small mode count.

Theorem 20 (two positive modes). *For every $N \geq 2$ and every $(s_1, s_2) \in [0, 1]^2$ with $s_1 \geq s_2$, the sector $(n - 1, 2)$ satisfies*

$$S(n - 1, 2, (s_1, s_2)) \leq S^*(N) = \frac{y}{2} + \frac{4}{y} - 2, \quad y = 2x, \quad (36)$$

strictly in the open box $0 < s_2 < s_1 < 1$. The supremum over the closed box is $S^*(N)$ and is attained exactly at the face point

$$s_1 = 1, \quad s_2 = x - 1 = \frac{y - 2}{2}, \quad (37)$$

where the sector becomes $(n, 1)$ and the spectrum is the ansatz of Theorem 9.

Corollary 21 (every $b = 2$ sector is eliminated). *For every $N \geq 2$ and every free spectrum with at most two positive modes, $D_{\mathcal{F}}(\bar{\rho}(N); \sigma) \geq D^*(N)$, with equality only at the ansatz. Hence Conjecture 1 holds on the family of free spectra with at most two positive-energy modes, and the ansatz is its unique maximiser there.*

Proof of Corollary 21. Sectors with $b \leq 2$ are covered by Lemma 14 ($a \geq n + 1$), Theorem 17 ($a = n$), Lemma 15 ($b = 0$), Lemma 16 ($b = 1$), Corollary 18 ($a + b \leq n$) and Theorem 20 ($a = n - 1, b = 2$). In each case the bound is $S \leq S^*(N)$, and equality is possible only in the sectors where the bound is sharp, namely $(n, 1)$ at $Z = x$ and $(n, 0)$ at $x = 1$, both of which are the ansatz. $\square$

7.1 Coordinates

Set

$$u = \frac{1}{1 + s_1} \in [\frac{1}{2}, 1], \quad v = \frac{s_2}{1 + s_2} \in [0, \frac{1}{2}], \quad \beta = u(1 - v) = \frac{1}{Z} \in [\frac{1}{4}, 1], \quad (38)$$

so $u = 1 - p_1$ and $v = p_2$ are complements of the two mode biases and β is the reciprocal normaliser. The four reduced atoms $\{1, s_1, s_2, s_1 s_2\}/Z$ become

$$\tau = \left(\beta, (1 - u)(1 - v), uv, v(1 - u) \right), \quad (39)$$

and the ordering constraint $s_1 \geq s_2$ is exactly the half-plane

$$u + v \leq 1. \quad (40)$$

Indeed $s_1 \geq s_2$ means $\frac{1-u}{u} \geq \frac{v}{1-v}$, which since all denominators are positive is $(1-u)(1-v) \geq uv$, that is $1-u-v \geq 0$. Under (40) and $u \geq \frac{1}{2}$ the list (39) is already decreasing: $\beta \geq (1-u)(1-v)$ because $u \geq 1-u$, $(1-u)(1-v) \geq uv$ is (40), and $uv \geq v(1-u)$ again because $u \geq 1-u$.

Two identities make the sector one-dimensional at fixed β :

$$\tau_2 = (1-u)(1-v) = \beta \left(\frac{1}{u} - 1 \right), \quad \tau_3 = uv = u - \beta. \quad (41)$$

The unclipped objective has a closed form in each of the two y -ranges. With $\lambda = y - 2 \in [0, 1]$ and $\kappa = y - 3 \in [0, 1]$,

$$S_{\text{unc}} := \sum_i w_i \tau_i = \beta + (1-u)(1-v) + \lambda uv = 1 - v + \lambda uv \quad (2 \leq y \leq 3), \quad (42)$$

$$\begin{aligned} S_{\text{unc}} &= \beta + (1-u)(1-v) + uv + \kappa v(1-u) \\ &= 1 - (4-y)v(1-u) \end{aligned} \quad (3 \leq y < 4), \quad (43)$$

where (42) uses $\beta + (1-u)(1-v) = (1-v)(u + (1-u)) = 1-v$ and (43) collects the v -terms. The exact objective is then

$$S = S_{\text{unc}} - \sum_i w_i \left(\tau_i - \frac{1}{y} \right)^+, \quad (44)$$

the correction being the mass by which the clipped atoms overshoot the flat level, weighted as in Proposition 12.

7.2 Why the coarse bound fails, and what replaces it

Since $w_1 = 1$ and $w_i \leq 1$, keeping only the first atom's clipping gives a valid bound

$$S \leq \text{UB} := S_{\text{unc}} - \left(\beta - \frac{1}{y} \right)^+. \quad (45)$$

In Case II this is enough: substituting (43) gives $\text{UB} = 1 - F$ with

$$F = (4-y)\psi(u) + \left(\beta - \frac{1}{y} \right)^+, \quad \psi(u) := v(1-u) = (1-u) - \beta \left(\frac{1}{u} - 1 \right), \quad (46)$$

and $F \geq G := (4-y)(y-2)/(2y) = 1 - S_1^*(y)$ on the whole feasible region, so $S \leq 1 - G = S_1^*(y)$ (Appendix A, Case II). The function ψ is concave in u at fixed β , since $\psi''(u) = -2\beta/u^3 < 0$, which is what makes the minimum of F

attainable at an endpoint of the feasible u -interval and therefore computable in closed form.

In Case I the coarse bound (45) is *false as a route to the theorem*: near the face $s_1 = 1$ the second atom clips as well, and dropping its clipping loses too much. The failure is easy to detect and easy to be fooled by, so it is worth recording precisely.

Remark 22 (a plausible shortcut that does not work). On the $1/32$ lattice in the open box $0 < s_2 < s_1 < 1$, the bound (45) exceeds $S^*(N)$ at 62 points, even though Theorem 20 holds at every one of them, where the exact objective satisfies $S \leq S^*(N)$. The count was recomputed independently for this paper on the same lattice: 62 failures in the open box, distributed as 18, 6, 18, 20 over $N \in \{5, 9, 10, 11\}$, and 90 on the closed box $0 < s_2 \leq s_1 \leq 1$, the extra 28 all lying on the face $s_1 = 1$ and none on the diagonal $s_2 = s_1$. The source lists the rank set as $\{5, 6, 7, 9, 10, 11\}$ with the same total of 62; at this lattice resolution ranks 6 and 7 contribute no failing point, so the difference is in the rank list rather than in the count. A proof that keeps only the first atom's clipping therefore cannot work in this sector, and any argument that appears to succeed with one is bounding the wrong quantity.

The replacement is a case analysis on the clipping pattern. Because $\tau_1 \geq \tau_2 \geq \tau_3$ and the clip level is $1/y$, the pattern is determined by how many of the first three atoms exceed $1/y$, giving four patterns in Case I; the fourth is infeasible, and Case II has two. In Case I.2 the identities (41) turn the objective into a single-variable function

$$\Theta(u) = \frac{\beta}{u} + \lambda u - \lambda\beta - \beta + \frac{1}{y}, \quad \Theta''(u) = \frac{2\beta}{u^3} > 0, \quad (47)$$

which is convex, so its maximum over the feasible u -interval is attained at an endpoint; both endpoints have closed forms, and the harder one reduces to the sign of a quartic on $[\sqrt{2}, \sqrt{3}]$.

Theorem 23 (structure of the proof of Theorem 20). *On the feasible region of the sector $(n-1, 2)$ the objective satisfies $S \leq \text{bound}$, where the bound is one of the six explicit expressions tabulated in Appendix A, each of which is at most $S_1^*(y)$, and the chain $S \leq \text{bound} \leq S_1^*(y)$ is strict unless (s_1, s_2) is the face point (37).*

Appendix A gives the six cases with their bounds and endpoint values, and Appendix B gives the two polynomial sign lemmas on which Cases I.2 and II depend:

$$P(q) = q^4 - 2q^3 - 2q^2 + 2q + 4 > 0 \quad \text{on } [\sqrt{2}, \sqrt{3}], \quad (48)$$

$$\left(1 - \frac{1}{q}\right)^2 - \frac{q^2 - 2}{2q^2} = \frac{(q-2)^2}{2q^2} \geq 0. \quad (49)$$

The first is the sign of Θ at its diagonal endpoint minus S_1^* , namely $\Theta_{\text{sup}} - S_1^*(y) = -P(q)/(2q^2)$ with $q = \sqrt{y}$; the second is the Case II diagonal endpoint bound. Both are decided exactly, the first by showing $P' < 0$ on the interval through $P'' > 0$ and one rational evaluation, and comparing $P(\sqrt{2}) = 4 - 2\sqrt{2} > 0$ with $P(\sqrt{3}) = 7 - 4\sqrt{3} > 0$ as integer inequalities $16 > 8$ and $49 > 48$.

The populations of the cases, counted on an exact rational grid of the whole feasible region at 32 values of $y \in [2, 4)$, are 3,330 points in Case I.1, 4,744 in I.2, 902 in I.3, none in I.4 and 7,920 in Case II, and at every one of the 16,896 points the full chain $S \leq \text{bound} \leq S_1^*(y)$ holds with equality only at the face point.

7.3 A second, independent proof

Theorem 20 was proved twice, by different routes, and the two proofs agree pointwise [2]. The second works in the bias coordinates $p_j = s_j/(1 + s_j)$ rather than in (u, v, β) , and splits the sector by the number $J \in \{0, 1, 2\}$ of clipped atoms instead of by the four clipping patterns of Case I; it therefore has different case boundaries and different endpoint expressions, and its algebraic content overlaps the first proof only through the shared objective. On an exact rational grid covering the whole open box the two agree at 564,060 points and, at double resolution, at 2,265,720 points, for every rank $2 \leq N \leq 80$, with no violation of $S \leq S^*(N)$ and no mismatch between the recomputed case label and $1 - S$; three independent implementations of the objective agree as exact rationals up to rank 300. The comparison is reported as evidence about the two proofs, not as part of either: a grid agreement cannot replace an argument, and the argument is the one in Appendix A.

What the second proof does add is confidence in the *statement*: it was written against an earlier state of the work in which the sector $(n - 1, 2)$ was still open, so it is not a transcription of the first. Its only failing assertion is a status claim — that the sector had not yet been eliminated — which the existence of the first proof makes false. That is a fact about the two proofs covering the same sector, not a mathematical gap, and it is recorded here because a verifier that fails on a status line is easy to misread as a verifier that fails on a theorem.

8 Concavity in one mode bias, and three positive modes

The sector $(n - 1, 2)$ of §7 needed a case analysis because at $b = 2$ the four visible atoms are automatically ordered and the objective collapses to one variable at fixed β . At $b = 3$ the collapse no longer happens by itself, and a

different tool is needed. It is general in b and it is the reason the elimination of the sector $(n-1, 3)$, and then of every sector $(n-1, b)$, is exact rather than numerical.

8.1 One coordinate at a time

Fix a sector (a, b) with $y = N/2^a$ and ratios $s \in [0, 1]^b$, and write the biases $p_j = s_j/(1 + s_j)$. For $T \subseteq \{1, \dots, b\} \setminus \{j\}$ set

$$\text{base}_T = \prod_{k \in T} s_k, \quad Z^{(j)} = \prod_{k \neq j} (1 + s_k), \quad \sum_T \text{base}_T = Z^{(j)}. \quad (50)$$

Every atom of the reduced spectrum is then one of

$$\tau'_{T \cup \{j\}} = p_j \frac{\text{base}_T}{Z^{(j)}}, \quad \tau'_T = (1 - p_j) \frac{\text{base}_T}{Z^{(j)}}, \quad (51)$$

so each atom is *affine* in p_j with slope $+\text{base}_T/Z^{(j)}$ if it uses mode j and $-\text{base}_T/Z^{(j)}$ if it does not. Split the atoms into the *clipped* ones, with $\tau'_i \geq 1/y$, whose contribution to S is the constant w_i/y , and the unclipped ones, whose contribution is $w_i \tau'_i$. Writing

$$W_{\text{clip}} = \sum_{i \text{ clipped}} w_i, \quad B_j = \sum_{\substack{i \text{ unclipped} \\ \text{using } j}} w_i \text{base}_{T(i)}, \quad F_j = \sum_{\substack{i \text{ unclipped} \\ \text{not using } j}} w_i \text{base}_{T(i)}, \quad (52)$$

where $T(i) = S(i) \setminus \{j\}$ for the subset $S(i)$ labelling atom i , the objective is

$$S(y, b, s) = \frac{1}{y} W_{\text{clip}} + p_j \frac{B_j}{Z^{(j)}} + (1 - p_j) \frac{F_j}{Z^{(j)}}, \quad (53)$$

and B_j, F_j are independent of p_j as long as the clipping pattern and the rank order of the atoms are fixed. The two sums in (52) are weighted base-sums; neither has anything to do with the free set $\mathcal{F}$.

A *chamber* is a region of the box on which the rank order of the 2^b subset products is fixed, and a *cell* is a region on which, in addition, the clipping pattern is fixed. Cells partition each chamber.

Lemma 24 (Lemma U: concavity in one mode bias). *Fix $y \geq 1$, $b \geq 1$ and $j \leq b$. On each chamber, $S(y, b, \cdot)$ is a continuous piecewise-affine concave function of p_j with the other biases fixed, and on each cell its slope is*

$$\frac{\partial S}{\partial p_j} = \frac{B_j - F_j}{Z^{(j)}}. \quad (54)$$

The slope is non-increasing in p_j along the chamber. On a cell with $B_j = F_j$ the slope vanishes and the value is explicit:

$$S = \frac{1}{y} W_{\text{clip}} + \frac{B_j}{Z^{(j)}}. \quad (55)$$

Proof. Continuity holds because $u \mapsto \min(u, 1/y)$ is continuous and S is a finite sum of compositions of it with the affine functions (51). On a cell the decomposition (53) holds with constant W_{clip} , B_j , F_j , so S is affine in p_j with slope (54), and (55) is (53) at $B_j = F_j$.

For monotonicity of the slope, increase p_j within a chamber. By (51) every atom that does not use mode j decreases and every atom that does increases. An atom that decreases can only cross the clip level $1/y$ from above, which makes it unclipped and adds $w_i \text{base}_{T(i)} \geq 0$ to F_j ; an atom that increases can only cross from below, which makes it clipped and removes $w_i \text{base}_{T(i)} \geq 0$ from B_j . No term is added to B_j and none removed from F_j . Hence $B_j - F_j$ is non-increasing in p_j , and since $Z^{(j)} > 0$ so is the slope. A continuous piecewise-affine function whose slopes are non-increasing is concave. $\square$

Lemma 24 has an immediate consequence for the location of maxima, which is the exact one-coordinate form of the induction used in §10.

Corollary 25 (where the maximum in one coordinate sits). *On a chamber, the maximum of $S(y, b, \cdot)$ in the coordinate p_j is attained either at a face of the box, $s_j = 0$ or $s_j = 1$, or on a cell with $B_j = F_j$, where the value is given by (55).*

Proof. A concave function of one variable on an interval attains its maximum at an endpoint or at a point where the slope changes from non-negative to non-positive. Endpoints of the chamber interval are either faces of the box ($s_j \in \{0, 1\}$) or ordering walls, and a sign change of the slope happens either inside a cell, where the slope is identically zero and $B_j = F_j$, or at a cell boundary. In the latter case the two one-sided slopes satisfy $B_j^- \geq F_j^-$ and $B_j^+ \leq F_j^+$, and the value at the wall is still given by (53) with either side's data, because the crossing atom contributes w_i/y on both sides. $\square$

A local maximum in the interior of a cell is therefore forced to satisfy an algebraic condition. Combining (53) with the gradient of S in the bias coordinates,

$$\frac{\partial S}{\partial p_j} = \frac{A_j}{p_j} - \frac{U - A_j}{1 - p_j}, \quad (56)$$

where A_j is the weighted mass of the counted unclipped atoms that occupy mode j and $U = A_j + (U - A_j)$ their total, gives the same condition in a form that does not mention chambers.

Proposition 26 (Proposition H: the stationary condition). *Let p be a local maximum of $S(y, b, \cdot)$ in the interior of a cell, off the ordering walls. Then $B_j = F_j$ for every j , equivalently*

$$p_j = \frac{A_j}{U} \quad \text{for every } j = 1, \dots, b, \quad (57)$$

and consequently

$$\frac{B_1}{Z^{(1)}} = \frac{B_2}{Z^{(2)}} = \cdots = \frac{B_b}{Z^{(b)}}, \quad S = \frac{1}{y} W_{\text{clip}} + \frac{B_j}{Z^{(j)}} \quad \text{for every } j. \quad (58)$$

That is $b-1$ independent algebraic conditions rather than one per coordinate. A local maximum on a clipping wall satisfies the corresponding one-sided inequalities $B_j^- \geq F_j^- \geq B_j^+ \geq F_j^+$ instead.

Proof. $B_j = F_j$ is (54) vanishing, which is necessary at an interior maximum in a cell interior by Lemma 24. For the second form, $A_j = p_j B_j / Z^{(j)}$ and $U - A_j = (1 - p_j) F_j / Z^{(j)}$ by (51), so $B_j = F_j$ is equivalent to $A_j / p_j = (U - A_j) / (1 - p_j)$, that is to $p_j = A_j / U$, which is also what (56) gives at a stationary point. For (58), if $B_j = F_j$ and $B_k = F_k$ hold on a common cell then (55) applied to j and to k gives two expressions for the same value S with the same W_{clip} , so $B_j / Z^{(j)} = B_k / Z^{(k)}$. The wall statement is Corollary 25. $\square$

Proposition 26 is a necessary condition, not a sufficient one, and it is the reason the search in §10 is a search over a small algebraic set rather than over a box. It says that at an interior maximum each mode bias must equal the share of the counted unclipped mass carried by the atoms that occupy that mode — a self-consistency condition, since the masses A_j and U are themselves determined by the biases through the clipping pattern.

8.2 The sector with three positive modes at $a = n - 1$

In the sector $(n-1, 3)$ one has $y = 2x \in [2, 4)$, so by Theorem 11 at most four of the eight subset products carry weight. The three largest are always 1, s_1 and s_2 , and the fourth is $\max(s_3, s_1 s_2)$: every other product is at most $s_1 s_2$ because the ratios are decreasing. There are therefore exactly two chambers,

$$\text{P} : s_3 \geq s_1 s_2, \quad \text{Q} : s_1 s_2 \geq s_3, \quad (59)$$

cut by the wall $s_3 = s_1 s_2$, and the fourth weighted atom uses mode 3 in chamber P and does not in chamber Q.

Theorem 27 (Theorem K: the sector $(n-1, 3)$). *For every $N \geq 2$ and every $(s_1, s_2, s_3) \in [0, 1]^3$, the sector $(n-1, 3)$ satisfies*

$$S(n-1, 3, (s_1, s_2, s_3)) \leq S^*(N) = \frac{y}{2} + \frac{4}{y} - 2, \quad y = 2x \in [2, 4), \quad (60)$$

strictly in the open box $(0, 1)^3$. The supremum over the closed box is $S^*(N)$ and is attained only on the face $s_3 = 0$, at the maximiser of Theorem 20.

Proof. Vary p_3 with p_1, p_2 fixed, and apply Lemma 24.

Case $y \leq 3$. The weights are $(1, 1, y - 2, 0)$, so the fourth atom carries no weight. The weighted atoms are 1, s_1 and s_2 over Z_3 , none of which uses mode 3, hence $B_3 = 0$ and the slope (54) is $-F_3/Z^{(3)} \leq 0$. It is strictly negative unless $F_3 = 0$, which would mean all three weighted atoms are clipped, that is $1 + 1/y \cdot 0 \geq 3/y$, impossible for $y \leq 3$ with the atoms summing to one. So S is strictly decreasing in p_3 and its maximum is at $s_3 = 0$, the face sector $(n - 1, 2)$, which Theorem 20 bounds by $S^*(N)$ with equality only at (37).

Case $y \geq 3$, chamber Q . The weights are $(1, 1, 1, f)$ with $f = y - 3 \in [0, 1)$, and the fourth atom is $s_1 s_2$, which does not use mode 3. Again $B_3 = 0$ and the same argument applies.

Case $y \geq 3$, chamber P . The fourth weighted atom is s_3 and does use mode 3, so B_3 may be positive and the slope need not be negative. Lemma 24 still gives concavity in p_3 , so the maximum over the chamber is either at a face, already handled, or on a cell with $B_3 = F_3$, where the value is (55). Enumerating the clipping patterns of the four weighted atoms $(1, s_1, s_2, s_3)/Z_3$ that can occur in chamber P leaves two critical values:

$$\frac{1}{y} + \frac{f}{1+f} \quad \text{and} \quad \frac{2}{y} + \frac{f}{(1+f)^2}, \quad f = y - 3. \quad (61)$$

The first arises when only the largest atom is clipped, so $W_{\text{clip}} = 1$ and $B_3 = f$; then $B_3 = F_3$ forces $s_1 + s_2 = f$ and $Z^{(3)} = (1 + s_1)(1 + s_2) \geq 1 + s_1 + s_2 = 1 + f$, whence $S = 1/y + B_3/Z^{(3)} \leq 1/y + f/(1+f)$. The second arises when the two largest atoms are clipped, so $W_{\text{clip}} = 2$ and $B_3 = F_3 = s_2$, giving $f = s_2 \leq s_1$ and $Z^{(3)} = (1 + s_1)(1 + s_2) \geq (1 + f)^2$, hence $S \leq 2/y + f/(1+f)^2$.

Both are strictly below $S_1^*(y)$. Writing h_1 and h_2 for the differences,

$$2y(y - 2) h_1 = y^3 - 8y^2 + 20y - 12 > 0, \quad (62)$$

$$2y(y - 2)^2 h_2 = (y - 2)^4 - 2y(y - 3) > 0 \quad (63)$$

on $[3, 4]$, with positive prefactors there; the two polynomial inequalities are proved in Appendix B, the first by locating its minimum exactly at $y = 10/3$ where it equals $76/27$, the second by reducing it to a cubic at a rational point. Hence $S < S_1^*(y)$ throughout chamber P as well.

Equality over the closed box is attained only where the face argument lands, namely $s_3 = 0$ with (s_1, s_2) the maximiser of Theorem 20. $\square$

Theorem 27 was checked on 11,648 exact rational grid points of the sector, each carrying the full chain $S \leq \text{bound} \leq S_1^*(y)$ with the bound being the relevant expression among the face values and the two critical values (61), together with 474,240 exact secant checks of the affine form (53) and 1,231 exactly located $B_j = F_j$ cells [2]. For this paper the same

statements were re-derived from an independent implementation of the slope data: 1,777 secant identities $S(p') - S(p) = (p' - p)(B_j - F_j)/Z^{(j)}$, 1,920 slope-monotonicity pairs across cells, 170 exact $B_j = F_j$ cells satisfying (55), and 3,840 points of the sector $(n-1, 3)$ — 784 in chamber P and 3,056 on the monotone branches — each verified through $S \leq \text{bound} \leq S_1^*(y)$, together with the two polynomial identities (62) and their positivity on a $1/64$ grid of $[3, 4]$. No violation occurred at any point.

Remark 28. The easy half of the proof of Theorem 27 — $B_3 = 0$ because no weighted atom uses the varied mode — does not depend on $b = 3$ at all. It depends on the truncated atom count $\lceil y \rceil$ being small enough that the varied smallest mode cannot appear among the weighted atoms. That is the content of Lemma 30 in §9, which turns the easy half of Theorem 27 into a theorem for every mode count.

9 Any number of positive modes at $n - 1$ zero modes

The easy half of the proof of Theorem 27 used only that the varied mode does not appear among the atoms that carry weight. This section isolates that mechanism, proves it for every mode count, and draws the consequences: the whole of the deficit $\nu = 1$ column, and a finite residual width at every other deficit.

9.1 How many atoms carry weight

By Theorem 11 the weight of the i -th reduced atom is $w_i(y) = \text{clamp}(y - i + 1, 0, 1)$, which is positive exactly for $i \leq \lceil y \rceil$ and equals 1 for $i \leq \lfloor y \rfloor$.

Lemma 29 (the truncated atom count). *For every $y > 0$ exactly $K(y) = \lceil y \rceil$ of the 2^b reduced atoms carry positive weight, the last of them with weight $y - \lfloor y \rfloor$ when y is not an integer and with weight 1 when it is. In the sector of zero-mode deficit ν , where $y \in [2^\nu, 2^{\nu+1})$, one has $K(y) \leq 2^{\nu+1}$.*

Proof. $w_i > 0$ is equivalent to $y - i + 1 > 0$, that is $i < y + 1$, that is $i \leq \lceil y \rceil$. At $i = \lceil y \rceil$ the weight is $y - \lceil y \rceil + 1$, which is 1 if y is an integer and $y - \lfloor y \rfloor$ otherwise. The bound $K(y) \leq 2^{\nu+1}$ is $\lceil y \rceil \leq 2^{\nu+1}$ for $y < 2^{\nu+1}$. $\square$

Lemma 29 is the source of every truncation argument in this paper, and it is the reason the problem is tractable at all: no matter how many modes a competitor has, the objective sees only the $\lceil y \rceil$ largest subset products. At $\nu = 1$ that is four atoms; at $\nu = 2$, eight; at $\nu = 3$, sixteen.

9.2 Where the smallest mode sits

For a ratio list $s \in [0, 1]^b$ let $R_b(s)$ be the rank of the atom s_b in the decreasing rearrangement of the 2^b subset products, counting ties in the worst possible way:

$$R_b(s) = 1 + \#\{S \subsetneq \{1, \dots, b\} : s_S > s_b\}. \quad (64)$$

Lemma 30 (Lemma V: the smallest mode is deep). *If $s_1 > s_2 > \dots > s_b > 0$ then $R_b(s) \geq b + 1$.*

Proof. The b proper subsets $\emptyset, \{1\}, \{2\}, \dots, \{b-1\}$ give the products $1, s_1, \dots, s_{b-1}$, each strictly larger than s_b , and they are distinct. So at least b atoms precede s_b and $R_b \geq b + 1$. $\square$

Lemma 30 is sharp in the sense that $b + 1$ is the best uniform bound: at $s = (\frac{1}{2}, \frac{1}{2} - \eta, \dots)$ the atom s_b sits exactly at rank $b + 1$. What matters for the objective is the comparison of R_b with $K(y)$.

9.3 The descent

Theorem 31 (Theorem L: the smallest mode can be deleted). *Let $y \geq 2$ and $b \geq K(y) = \lceil y \rceil$. Then for every decreasing $s \in [0, 1]^b$,*

$$S(y, b, s) \leq S(y, b-1, (s_1, \dots, s_{b-1})), \quad (65)$$

with strict inequality on the open strict-order chamber $s_1 > \dots > s_b > 0$. Ties and the faces $s_b \in \{0, 1\}$ follow by continuity.

Proof. Work first on the open strict-order chamber, and vary p_b with the other biases fixed. By Lemma 30 the atom s_b has rank at least $b + 1 > K(y)$, and every other atom that uses mode b is of the form $s_b s_T$, which is smaller still, so no weighted atom uses mode b . In the notation of (52) this says $B_b = 0$ on every cell.

The slope (54) is therefore $-F_b/Z^{(b)}$, and it is strictly negative because $F_b > 0$. Indeed $F_b = 0$ would mean that every weighted atom is clipped, so $\sum_{i \leq K(y)} \tau'_i \geq K(y)/y \geq 1$; since $\sum_i \tau'_i = 1$ this forces $\tau'_i = 0$ for all $i > K(y)$, in particular for the atom s_b/Z_b , contradicting $s_b > 0$.

Lemma 24 then gives that $S(y, b, \cdot)$ is strictly decreasing in p_b on the chamber, so its maximum is at $p_b = 0$, that is $s_b = 0$. At $s_b = 0$ the normaliser is unchanged, $Z_b = Z_{b-1}$, and the subset products that use mode b vanish, so the reduced spectrum is that of the sector $(b-1)$ padded with zeros and (65) holds with equality at the face.

For the closed box, note that $S(y, b, \cdot)$ is a symmetric function of $s_1, \dots, s_b$ — the multiset of subset products does not depend on their order — and hence continuous on all of $[0, 1]^b$, as is the right-hand side of (65). The open strict-order chamber with $s_b < 1$ is dense in $[0, 1]^b$, so the non-strict inequality extends to the whole box. $\square$

Theorem 31 is a *truncation* statement rather than an elimination: it says that at fixed y , mode counts above $\lceil y \rceil$ never help, so the supremum over b is attained at $b < \lceil y \rceil$. Iterating it turns an unbounded family of sectors into a bounded one.

Corollary 32 (the deficit-one column is closed). *Let $a = n - 1$, so $y = 2x \in [2, 4)$ and $K(y) \leq 4$. For every $b \geq 4$,*

$$S(n-1, b, s) \leq S(n-1, b-1, s') \leq \dots \leq S(n-1, 3, s'') \leq S^*(N), \quad (66)$$

and for $b \leq 2$ the sector is covered by Corollary 18 and Theorem 20. Hence every sector $(n-1, b)$ satisfies $S \leq S^*(N)$.

Proof. Theorem 31 applies for $b \geq 4 = K(y)$ and gives the descent, which stops at $b = 3$ because $3 < K(y)$; there Theorem 27 gives the bound. For $b \leq 1$, $a + b = n$ or $n - 1 \leq n$, so Corollary 18 applies; for $b = 2$ Theorem 20 applies. $\square$

Theorem 33 (Main Theorem, second family). *For every $N \geq 2$ and every free spectrum σ with exactly $n - 1$ zero modes and any number of positive modes,*

$$D_{\mathcal{F}}(\bar{\rho}(N); \sigma) \geq D^*(N), \quad (67)$$

with equality only at the ansatz of Theorem 9. This family contains the sector $(n-1, 3)$ of Theorem 27 and every sector with at most two positive modes and $a = n - 1$, and it is the first family in this paper that leaves the mode count free.

Proof. Corollary 32 gives $S \leq S^*(N)$ in every sector $(n-1, b)$, and Lemma 6 converts it to (67). Equality requires equality in every step of the descent, and Theorem 31 is strict on the open strict-order chamber, so an extremal spectrum must have $s_b = 0$ for every $b \geq 4$, reducing it to a sector with at most three positive modes; Theorem 27 and Theorem 20 are strict in their open boxes, so the extremal spectrum lies on the face $s_2 = s_3 = 0$ at $s_1 = x - 1$ with the zero modes, which is the ansatz. $\square$

Corollary 34 (finite residual width at every deficit). *Let $\nu = n - a \geq 2$, so $y \in [2^\nu, 2^{\nu+1})$ and $K(y) \leq 2^{\nu+1}$. The sectors that neither Corollary 18, nor Theorem 20, nor Theorem 31 reaches are exactly*

$$a \leq n - 2, \quad \max(3, \nu + 1) \leq b \leq 2^{\nu+1} - 1, \quad (68)$$

that is, at each deficit $\nu \geq 2$ a finite range of mode counts:

deficit ν	2	3	4	5	6	...
residual b	3–7	4–15	5–31	6–63	7–127	...
number of sectors	5	12	27	58	121	...

The smallest open sector is $(n-2, 3)$, where $y \in [4, 8)$ and up to eight atoms carry weight.

Proof. $b \leq \nu$ is Corollary 18, since $a + b \leq n$; $b \leq 2$ is Theorem 20 together with Lemma 15 and Lemma 16; and $b \geq K(y)$ is Theorem 31, which descends to $b = K(y) - 1 \leq 2^{\nu+1} - 1$. What is left is (68), and the lower bound $\max(3, \nu+1)$ is the larger of the two lower bounds 3 and $\nu+1$. At $\nu = 2$ the range is $b \in \{3, \dots, 7\}$, at $\nu = 3$ it is $b \in \{4, \dots, 15\}$, and so on; the counts are $2^{\nu+1} - 1 - \max(3, \nu+1) + 1$. $\square$

Before Theorem 31 the residual family had an unbounded range of b at every deficit $\nu \geq 2$; Corollary 34 replaces it by a finite list, five sectors at $\nu = 2$ and twelve at $\nu = 3$. The list still grows doubly exponentially in ν , so this is a reduction and not a resolution: at $\nu = 6$ there are 121 sectors, each a b -dimensional box with up to $2^b = 2^{127}$ atoms.

9.4 The threshold is sharp

The hypothesis $b \geq \lceil y \rceil$ of Theorem 31 cannot be weakened.

Proposition 35 (Proposition L3: sharpness of the threshold). *Let $y \geq 2$ and $b = K(y) - 1 = \lceil y \rceil - 1$. For $\eta > 0$ small enough that $s_j = \frac{1}{2} - j\eta$ is decreasing and positive, the atom s_b has rank exactly $K(y)$ in the sector with $y \in [b, b+1)$, it carries weight $w_{K(y)} = y - \lfloor y \rfloor > 0$, and $B_b = w_{K(y)} > 0$. The slope argument of Theorem 31 therefore fails at $b = K(y) - 1$, and for $b = 3$, $y > 3$ this point lies in chamber P of Theorem 27, where the critical-value analysis is needed.*

Proof. Choose $\eta > 0$ so small that $s_j = \frac{1}{2} - j\eta$ is positive and decreasing and that every product of k of the ratios exceeds every product of $k+1$ of them; the latter holds because the ratio of the smallest k -fold product to the largest $(k+1)$ -fold product is at least $2(1 - C\eta)$ for a constant C depending only on b , which exceeds 1 for small η . The subset products strictly larger than s_b are then exactly the empty product and the $b-1$ other one-fold products $s_1, \dots, s_{b-1}$: every two-fold product is at most $(\frac{1}{2})^2(1+O(\eta)) < \frac{1}{2} - O(\eta) \leq s_b$ for small η . Hence $\#\{S \subsetneq [b] : s_S > s_b\} = b$ and $R_b = b+1$ by (64). With $b = \lceil y \rceil - 1$ this gives $R_b = \lceil y \rceil = K(y)$, the first rank that carries weight, and that weight is $w_{K(y)} = y - \lfloor y \rfloor$, which is positive whenever y is not an integer. The atom s_b/Z_b is unclipped for η small, because s_b/Z_b is close to $\frac{1}{2}/2^b > \dots$ and in any case exceeds $1/y$ at $y < 2^{b+1}$; it contributes $w_{K(y)} \text{base}_0 = w_{K(y)}$ to B_b , so $B_b \geq w_{K(y)} > 0$ and the identity $B_b = 0$ on which Theorem 31 rests fails. At $b = 3$ and $y > 3$ the same point satisfies $s_3 > s_1 s_2$, since $\frac{1}{2} - 3\eta > (\frac{1}{2} - \eta)(\frac{1}{2} - 2\eta) = \frac{1}{4} - \frac{3}{2}\eta + 2\eta^2$ for $\eta < \frac{1}{6}$, which is chamber P of (59). $\square$

Proposition 35 was verified exactly on 56 points covering $y \in [2, 16)$ in steps of $1/4$, with η chosen as an exact rational small enough for the cardinality ordering, and with the rank R_b and the weight $w_{K(y)}$ computed from the subset products rather than assumed [2]. It identifies precisely where the next argument has to start: at $\nu \geq 2$ the varied smallest mode *is* visible to the objective, and the chamber structure of the $2^{\nu+1}$ -atom truncation is what has to be understood.

10 What is left: interior local maxima

The eliminations of §§6–9 are statements about specific families of sectors. This section shows that, taken together with the geometry of the sector boxes, they are enough to reduce Conjecture 1 in full to a single analytic statement, and it states exactly what remains.

10.1 The faces of a sector box are sectors

Lemma 36 (face identities). *For every $N \geq 2$, every sector (a, b) and every $s \in [0, 1]^b$,*

$$S(y, b, (s_1, \dots, s_{b-1}, 0)) = S(y, b-1, (s_1, \dots, s_{b-1})), \quad (69)$$

$$S(y, b, (1, s_2, \dots, s_b)) = S(\frac{y}{2}, b-1, (s_2, \dots, s_b)), \quad (70)$$

and more generally, for any j , setting $s_j = 0$ gives the sector $(a, b-1)$ with the remaining ratios and setting $s_j = 1$ gives the sector $(a+1, b-1)$ with y replaced by $y/2$.

Proof. If $s_j = 0$ then every subset product containing j vanishes and Z_b is unchanged, so the reduced spectrum is that of the $(b-1)$ -sector with the remaining ratios, padded with zeros; the weights and the level $1/y$ are unchanged, giving (69). If $s_1 = 1$ then $Z_b = 2Z'$, where $Z' = \prod_{j \geq 2} (1 + s_j)$, and each subset product of the remaining ratios occurs twice in the reduced spectrum, at half its former value: the atom list is $\tau'_i = \tau''_i/2$ with each entry of the $(b-1)$ -sector list τ'' duplicated. Positions $2i-1$ and $2i$ then carry weights $\text{clamp}(y-2i+2, 0, 1)$ and $\text{clamp}(y-2i+1, 0, 1)$, whose sum is $2 \text{clamp}(\frac{y}{2} - i + 1, 0, 1) = 2w_i(\frac{y}{2})$, and $\min(\tau''_i/2, 1/y) = \frac{1}{2} \min(\tau''_i, 2/y)$. The pair contributes $w_i(\frac{y}{2}) \min(\tau''_i, \frac{2}{y})$, which is the i -th term of the right-hand side of (70). The case of a general j is the same after relabelling, since the objective is symmetric in the ratios. $\square$

Both identities were verified as exact rational equalities on 576 sector/point combinations, including the non-obvious second one, where the halving of the level and the doubling of the multiplicity have to cancel exactly. They say that the boundary of a sector box consists of sectors with *one mode fewer*, at the same rank N . That is what makes an induction on the mode count possible.

10.2 The reduction

Theorem 37 (Theorem G: mode-count reduction). *The following are equivalent.*

- (i) Conjecture 1 holds at every rank $N \geq 2$.
- (ii) For every $N \geq 2$ and every sector (a, b) , the objective $S(a, b, \cdot)$ has no interior local maximum in the open box $(0, 1)^b$ whose value exceeds $S^*(N)$.

Proof. (i) $\Rightarrow$ (ii) is immediate: if $S(a, b, s) \leq S^*(N)$ for every $s \in [0, 1]^b$ then no local maximum exceeds $S^*(N)$.

For (ii) $\Rightarrow$ (i), fix N and induct on the mode count $m = a + b$. The base case $m \leq n$ is Corollary 18, which gives $S \leq S^*(N)$ without any hypothesis. For the step, assume every sector of mode count at most $m - 1$ satisfies $S \leq S^*(N)$, and consider a sector (a, b) with $a + b = m$. The objective is continuous on the compact box $[0, 1]^b$, so it attains a maximum. If the maximum is attained at a boundary point, some s_j equals 0 or 1, and Lemma 36 expresses the value as the objective of a sector with $m - 1$ modes at the same rank, which the induction hypothesis bounds by $S^*(N)$. If the maximum is attained only in the interior, it is an interior local maximum, and by (ii) its value is at most $S^*(N)$. Either way $S \leq S^*(N)$ on the whole box, and the induction closes. Since N was arbitrary, (11) holds for every $\sigma \in \mathcal{F}$, which is Conjecture 1 by Lemma 6 and (19). $\square$

Theorem 37 needs no a priori bound on the number of modes, which is the point: the sector boxes grow in dimension with b , but every box is controlled by its own faces plus the absence of interior maxima. Combined with Proposition 26, the open problem becomes algebraic rather than variational.

Corollary 38 (the open problem, precisely). *Conjecture 1 holds at every rank if and only if there is no N , no sector (a, b) in the residual family*

$$a \leq n - 2, \quad 3 \leq b \leq 2^{n-a+1} - 1, \quad a + b \geq n + 1, \quad (71)$$

and no $s \in (0, 1)^b$ with $B_j = F_j$ for every j , or with the one-sided alternatives of Proposition 26 on a clipping wall, at which $S(a, b, s) > S^*(N)$. The smallest such sector is $(n - 2, 3)$, and the next are $(n - 2, 4), \dots, (n - 2, 7)$.

Proof. Theorem 37 reduces the conjecture to interior maxima; the eliminations of §§6–9 settle every sector outside (71): $a \geq n$ by Lemma 14 and Theorem 17, $b \leq 2$ by Corollary 21, $a + b \leq n$ by Corollary 18, every b at $a = n - 1$ by Corollary 32, and $b \geq \lceil y \rceil$ at every a by Theorem 31, which leaves $b \leq 2^{n-a+1} - 1$ by Corollary 34. An interior local maximum must satisfy Proposition 26. At $\nu = n - a = 2$ the range (68) is $3 \leq b \leq 7$, so $(n - 2, 3)$ is the sector with the fewest modes. $\square$

Corollary 38 is the sharpest form of the open problem available from this work. It is worth being explicit about what kind of statement it is: an equivalence whose right-hand side is a negative existential over a countable family of compact boxes, each of finite dimension. Finding a single interior local maximum above $S^*(N)$ would refute Conjecture 1, and it is the only way the hypothesis of Theorem 37 can fail.

10.3 The search, and what it is worth

The necessary condition of Proposition 26 was used to search for interior maxima. Over the ranks $N \in \{7, 10, 11, 13, 14, 15\}$ and every residual sector with $n + 1 \leq a + b \leq n + 4$ — 85 sectors in total — a multistart local optimisation of the sector objective found no interior local maximum above $S^*(N)$: the largest excess over all runs was -1.4×10^{-14} , that is, below $S^*(N)$ to within floating-point noise, and in every sector the optimiser migrated to a face, typically to $s_1 = \dots = s_{b-1} = 1$ with $s_b = x - 1$, where the value is exactly $S^*(N)$ [2].

This is evidence, not a certificate, and its limitations are the usual ones plus one specific to this objective. A multistart local search can miss a narrow basin, and the basin structure here is set by the clipping walls: the objective is piecewise affine in the biases, so its local maxima live on cells, and a cell can be thin. What the search does establish is that no *large* interior maximum is hiding in the probed sectors, and that the natural candidate maximiser is always the face point predicted by the analysis. Where the question has been settled by an argument rather than a search — $b \leq 2$ everywhere, every b at $a = n - 1$, and $b \geq \lceil y \rceil$ at every a — the search adds nothing and is not needed. The interior-maximum statement is a theorem in those families and a search result elsewhere, and §14 keeps the two apart.

10.4 The residual family, concretely

For orientation, here is what (71) looks like at small ranks. At $n = 3$, so $8 \leq N \leq 15$, the residual sectors are $(1, 3), \dots, (1, 7)$ at deficit $\nu = 2$ and $(0, 4), \dots, (0, 15)$ at deficit $\nu = 3$: 17 sectors in all. At $n = 4$, $16 \leq N \leq 31$, they are $(2, 3), \dots, (2, 7)$, $(1, 4), \dots, (1, 15)$ and $(0, 5), \dots, (0, 31)$: 44 sectors. In general the number of residual sectors at fixed n is $\sum_{\nu=2}^n (2^{\nu+1} - \max(3, \nu + 1))$, which grows like 2^{n+2} , and each is a box of dimension up to $2^{\nu+1} - 1$. The certificates of §11 handle every sector with $b \leq n + 3$ at any fixed rank by computation; the residual family of Corollary 38 is what no fixed-rank computation can reach uniformly in N .

Two structural facts make the smallest open sector $(n - 2, 3)$ genuinely harder than the ones already closed, and they explain why the methods of §§8 and 9 do not extend to it by inspection. First, at $\nu = 2$ one has $y = 4x \in [4, 8)$, so up to eight atoms carry weight instead of four, and

the four-atom truncation that makes the chamber analysis of Theorem 27 close no longer applies: the ordering of eight subset products has many more chambers than the two of (59). Second, $b = 3$ is below the threshold $\lceil y \rceil \geq 5$, so Lemma 30 does not apply and the varied smallest mode is visible to the objective: B_3 can be positive, and the concavity of Lemma 24 has to be paired with an enumeration of critical values that grows with the number of clipping patterns. Neither obstacle is obviously fatal, and both are quantitative rather than conceptual, which is why $(n-2, 3)$ is the natural next target.

11 Certified brackets at fixed rank

The eliminations of §§6–9 are uniform in the rank but leave the residual family (71) open. At a *fixed* rank the residual family is finite in the mode counts that matter, and there the objective can be bounded by computation rather than by analysis. This section describes what that computation proves, what it does not, and what the resulting numbers are. The protocol itself, including the interval arithmetic and the four bounds that are combined at each box, is in Appendix C.

11.1 What a completed run proves

The computation is exact interval branch and bound over the sector box $[0, 1]^b$: boxes are bisected until their width reaches 2^{-e} , every box is bounded from above by an interval evaluation of the objective that respects the clipping and the ordering, and every surviving box is re-verified in exact rational arithmetic. A run that completes therefore produces a rational number B_e with

$$B_e \geq \sup_{s \in [0, 1]^b} S(a, b, s), \quad (72)$$

an unconditional upper bound on that one sector box. Nothing about the run is statistical: no sampling, no floating-point optimisation, no convergence assumption. A box or wall-clock budget can cut a run short, in which case the sector is reported as uncertified and contributes no bound.

Definition 39 (covered and uncovered sectors). At rank N with $n = \lfloor \log_2 N \rfloor$, a sector (a, b) is *covered* if it is discharged analytically — $a \geq n$, or $b \leq 2$, or $a + b \leq n$ — or computationally, by a completed run at $0 \leq a \leq n$, $0 \leq b \leq n + 3$. The *uncovered* family is

$$\mathcal{U}(N) = \{(a, b) : a \leq n - 1, \quad b \geq n + 4\}, \quad (73)$$

and $D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N))$ denotes the distance obtained by minimising over covered

sectors only:

$$D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N)) = \inf \left\{ \frac{1}{2} \|\bar{\rho}(N) - \sigma\|_1 : \sigma \in \mathcal{F}, \sigma \text{ lies in a covered sector} \right\}. \quad (74)$$

Since the covered sectors form a subset of $\mathcal{F}$, $D_{\mathcal{F}}(\bar{\rho}(N)) \leq D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N))$ always. The two are equal exactly when no uncovered sector beats every covered one, which is not something a fixed-rank computation can decide.

Remark 40 (the covered family is larger than the runs need). Definition 39 counts as analytically covered only $a \geq n$, $b \leq 2$ and $a + b \leq n$, which is what the certified runs of Table 1 assume. The eliminations proved in this paper are stronger: by Theorem 20, Theorem 27 and Corollary 32 every sector with $a = n - 1$ is discharged analytically at every b , so the covered family may be enlarged to include $a \geq n - 1$ for all b , and the uncovered family shrinks to

$$\mathcal{U}'(N) = \{(a, b) : a \leq n - 2, \quad b \geq n + 4\} \subseteq \mathcal{U}(N). \quad (75)$$

The bracket (77) is unchanged, because each enlarged sector satisfies $\sup S = S^*(N)$, attained on its face $s_1 = 1$ at the ansatz point $s_2 = x - 1$, $s_3 = \dots = s_b = 0$ of the sector $(n, 1)$, so these sectors add no competitor closer to the flat target than the ansatz already is; and the transfer condition (78) becomes the strictly weaker requirement that $\sup S \leq S^*(N) + \varepsilon_e$ on $\mathcal{U}'(N)$. Both forms are recorded here, and §14 uses the smaller family.

Theorem 41 (the certified bracket, and its exact scope). *Suppose every sector with $0 \leq a \leq n$ and $0 \leq b \leq n + 3$ is discharged at rank N , the analytically eliminated ones by a bound at most $S^*(N)$ and the rest by completed runs with bounds B_e . Write*

$$\varepsilon_e = \max \{ B_e - S^*(N) : \text{sectors run at tolerance } 2^{-e} \} \geq 0. \quad (76)$$

Then

$$D^*(N) - \varepsilon_e \leq D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N)) \leq D^*(N), \quad D_{\mathcal{F}}(\bar{\rho}(N)) \leq D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N)). \quad (77)$$

The same bracket with $D_{\mathcal{F}}$ in place of $D_{\mathcal{F}}^{\text{cov}}$ holds if and only if

$$\sup \{ S(a, b, s) : (a, b) \in \mathcal{U}(N), \quad s \in [0, 1]^b \} \leq S^*(N) + \varepsilon_e, \quad (78)$$

which follows from the hypothesis of Theorem 37, or from Conjecture 1 restricted to $\mathcal{U}(N)$, and is otherwise open.

Proof. The upper bound in (77) is (19): the ansatz lies in the covered sector $(n, 1)$, so the infimum over covered sectors is at most $D^*(N)$. For the lower bound, every covered competitor lies in a sector with a proved upper bound

on S , either $S^*(N)$ from an analytic elimination or $B_e \leq S^*(N) + \varepsilon_e$ from a completed run, so its distance is at least $1 - S^*(N) - \varepsilon_e = D^*(N) - \varepsilon_e$ by Lemma 6. Taking the infimum gives the left inequality of (77).

For the second statement, if (78) holds then every uncovered competitor also has $S \leq S^*(N) + \varepsilon_e$, so the same lower bound applies to the infimum over all of $\mathcal{F}$, and the upper bound $D_{\mathcal{F}} \leq D^*$ is (19). Conversely, if the bracket holds for $D_{\mathcal{F}}$ then every $\sigma \in \mathcal{F}$ has $\frac{1}{2}\|\bar{\rho}(N) - \sigma\|_1 \geq D^*(N) - \varepsilon_e$, which by Lemma 6 is $S(\bar{\rho}(N), \sigma) \leq S^*(N) + \varepsilon_e$, in particular on $\mathcal{U}(N)$. $\square$

An earlier statement of this result quoted the bracket for $D_{\mathcal{F}}$ itself, without the condition (78). That stronger form does not follow from a fixed-rank computation, because such a computation says nothing about sectors with $b \geq n + 4$, and it is not asserted here. The distinction matters: (77) is a theorem about a restricted variational problem, and reading it as a theorem about $D_{\mathcal{F}}$ would quietly assume part of Conjecture 1.

11.2 Why the excess is positive at every finite tolerance

Every residual closed sector box contains a face point at which the objective equals $S^*(N)$ exactly: the point $s_1 = \dots = s_{b-1} = 1$, $s_b = x - 1$ of Lemma 36, whose spectrum is the ansatz. So $\sup S \geq S^*(N)$ on each of them, and by (72)

$$\varepsilon_e = B_e - S^*(N) \geq \sup S - S^*(N) \geq 0, \quad (79)$$

with no reason for the first inequality to be an equality at finite e . The excess ε_e therefore measures the *resolution of the certificate*, not the truth of the conjecture: it is strictly positive in every one of the 1320 completed runs, and refining the tolerance shrinks it without ever removing it. What can be observed is how fast it decays.

It decays linearly. Over the five tolerances $2^{-6}, 2^{-8}, 2^{-10}, 2^{-12}, 2^{-14}$ and the 22 certified ranks, the ratio $\varepsilon_e/2^{-e}$ of the largest excess at each rank stays between 0.0627, at $N = 17$, and 0.7243, at $N = 25$, and within a rank it varies by at most 1.7% across the five tolerances, which is inside the two-percent stability test the source applies and is why all 264 sectors pass it. Over the 1140 completed sector runs whose excess is positive — the other 180 lie at the ranks $N = 2^n$, where the excess is exactly zero — the ratio ranges from 0.0568 to 0.7243 (Table 1, Figure 1). Linear decay in the box width is what a boundary layer of thickness 2^{-e} produces: the maximiser sits on a face, so the boxes that survive refinement are the thin ones adjacent to that face, and their inflation is proportional to their width. Turning the observation into a theorem would require a proved modulus of continuity for the interval bound — an explicit constant L with $B_e \leq S^*(N) + L2^{-e}$ — and no such constant is supplied here. Without it the decay is evidence about the geometry of the objective, not a proof that $\sup S = S^*(N)$ on the covered sectors.

Table 1: Certified flat-spectrum brackets. For each N : $n = \lfloor \log_2 N \rfloor$, $x = N/2^n$, the ansatz value $D^*(N)$, sectors covered, how many are eliminated by an exact lemma versus exact rational branch and bound, whether all sectors $0 \leq a \leq n$, $0 \leq b \leq n+3$ were certified, the largest certified excess $\varepsilon_e = B_e - S^*(N)$ at the finest completed tolerance, the range of $\varepsilon_e/2^{-e}$ over $2^{-6}, \dots, 2^{-14}$, and wall-clock seconds. Each completed run proves the two-sided bracket of Theorem 41, which bounds $D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N))$ unconditionally and $D_{\mathcal{F}}(\bar{\rho}(N))$ under (78). Ranks $N = 2^n$ have $D^* = 0$ and $\varepsilon_e = 0$, the flat spectrum being itself free.

N	n	x	D^*	sectors	lemma	B&B	all	ε_e	$\varepsilon_e/2^{-e}$	s
2	1	1	0.000000000000	10	8	2	yes	0	–	0.0
3	1	3/2	0.166666666667	10	8	2	yes	1.865e-05	0.3056–0.3057	0.7
4	2	1	0.000000000000	18	12	6	yes	0	–	0.1
5	2	5/4	0.150000000000	18	12	6	yes	1.282e-05	0.2100–0.2111	1.0
6	2	3/2	0.166666666667	18	12	6	yes	2.458e-05	0.4028–0.4032	4.0
7	2	7/4	0.107142857143	18	12	6	yes	3.675e-05	0.6020–0.6026	158.8
8	3	1	0.000000000000	28	17	11	yes	0	–	0.3
9	3	9/8	0.097222222222	28	17	11	yes	7.300e-06	0.1196–0.1209	1.7
10	3	5/4	0.150000000000	28	17	11	yes	1.358e-05	0.2225–0.2245	4.0
11	3	11/8	0.170454545455	28	17	11	yes	2.368e-05	0.3879–0.3904	12.5
12	3	3/2	0.166666666667	28	17	11	yes	3.264e-05	0.5347–0.5379	29.8
13	3	13/8	0.144230769231	28	17	11	yes	3.905e-05	0.6398–0.6429	214.1
16	4	1	0.000000000000	40	23	17	yes	0	–	0.6
17	4	17/16	0.055147058824	40	23	17	yes	3.828e-06	0.0627–0.0638	2.8
18	4	9/8	0.097222222222	40	23	17	yes	7.512e-06	0.1231–0.1249	4.7
19	4	19/16	0.128289473684	40	23	17	yes	1.214e-05	0.1989–0.2022	8.0
20	4	5/4	0.150000000000	40	23	17	yes	1.839e-05	0.3013–0.3049	14.5
21	4	21/16	0.163690476190	40	23	17	yes	2.415e-05	0.3957–0.4000	29.0
22	4	11/8	0.170454545455	40	23	17	yes	2.904e-05	0.4757–0.4802	60.8
23	4	23/16	0.171195652174	40	23	17	yes	3.353e-05	0.5494–0.5542	99.0
24	4	3/2	0.166666666667	40	23	17	yes	3.772e-05	0.6181–0.6247	326.8
25	4	25/16	0.157500000000	40	23	17	yes	4.382e-05	0.7180–0.7243	745.5

11.3 The certified ranks

Table 1 lists the 22 ranks at which every sector with $0 \leq a \leq n$ and $0 \leq b \leq n+3$ was discharged: 660 sectors in total, of which 396 fell to an analytic lemma and 264 required branch and bound, in 1718.6 seconds of exact arithmetic. The ranks are $2 \leq N \leq 13$ and $16 \leq N \leq 25$; 14 and 15 were not run. Every column of the table was recomputed for this paper from the exact rationals recorded per sector and per tolerance, together with the internal consistency conditions that the “all certified” column rests on: every sector appears exactly once, every bound is at least $S^*(N)$, the excesses are non-increasing in the tolerance, and the ratio $\varepsilon_e/2^{-e}$ is stable to within one percent. The largest certified excess over all ranks and tolerances is 4.382×10^{-5} , at $N = 25$.




Figure 1: Decay of the certified excess with the resolution of the branch and bound, for the four ranks $N = 7, 13, 20, 25$. Panel (a) plots $\varepsilon_e = B_e - S^*(N)$, the largest excess over the sectors discharged by branch and bound at that rank, against the box width 2^{-e} on log-log axes, with a dashed line of slope one for scale; panel (b) plots the same data as the ratio $\varepsilon_e/2^{-e}$ against 2^{-e} on log-linear axes. The five points of each series are the five completed tolerances $2^{-6}, \dots, 2^{-14}$. The straight line in panel (a) and the flat behaviour in panel (b) are the same statement: the excess decays linearly in the box width, so the surviving boxes are the thin ones adjacent to the face where the ansatz sits. The ratio is bounded away from zero — between 0.0627 and 0.7243 for the largest excess at each of the 22 certified ranks — which is why no finite tolerance certifies $\sup S = S^*(N)$, and it is not a proved Lipschitz constant (Theorem 41 and §11.4).

11.4 What the certificates do and do not give

At each of the 22 ranks, the numbers in Table 1 are a rigorous two-sided bracket on a restricted variational problem, with rational endpoints that can be checked by an independent implementation. Three readings of them would be wrong, and it is worth stating why.

First, the bracket is not a proof of Conjecture 1 at that rank. It bounds $D_{\mathcal{F}}^{\text{cov}}$, and the transfer to $D_{\mathcal{F}}$ needs (78): at $N = 25$, by Remark 40, controlling every sector with $a \leq n - 2 = 2$ and $b \geq n + 4 = 8$ — boxes of dimension 8 and more, with up to 2^b atoms, none of which the run touched.

Second, the smallness of ε_e is not evidence that $\sup S$ is close to $S^*(N)$ on the covered sectors. Since $\sup S \geq S^*(N)$ on every residual box by the face argument, the truth is $\sup S = S^*(N)$ or $\sup S > S^*(N)$, and ε_e bounds the distance to that answer from above only at the resolution of the run.

Third, the growth of the run time with N — from 0.1 seconds at $N = 4$ to 745.5 at $N = 25$ — is a statement about the method, not about the problem. It reflects the growth of the number of sectors, 18 at $n = 2$ to 40 at $n = 4$, and of the dimension of the hardest boxes, and it is the practical reason the certified ranks stop at 25 rather than a mathematical obstruction.

Ranks 14 and 15 are absent because they were not run, not because they failed.

12 Doubling a spectrum

Conjecture 2 asserts that $D_{\mathcal{F}}$ is invariant under $\rho \mapsto \frac{1}{2}(\rho \oplus \rho)$. One direction is classical and one line long: if σ competes for ρ then $\frac{1}{2}(\sigma \oplus \sigma)$ is free — it is σ with one zero mode added — and competes for $\frac{1}{2}(\rho \oplus \rho)$ at the same distance, so $D_{\mathcal{F}}(\frac{1}{2}(\rho \oplus \rho)) \leq D_{\mathcal{F}}(\rho)$ for every ρ . The other direction is open, and this section shows that on the flat family it reduces to one explicit family of spectra. Off the flat family nothing is proved here, and the numerical evidence is reported with its artefacts identified.

12.1 An identity that proves nothing

There is an attractive argument for Conjecture 2 on flat spectra that does not work, and it is worth disposing of first because it is exactly the kind of reasoning that makes the problem look easier than it is. Since $x(2N) = x(N)$ — doubling a rank does not change its position between powers of two — the ansatz values satisfy

$$D^*(2N) = D^*(N) \quad \text{for every } N \geq 2, \quad (80)$$

exactly. But (80) is an identity between two *conjectured* values of $D_{\mathcal{F}}$. To conclude $D_{\mathcal{F}}(\bar{\rho}(2N)) = D_{\mathcal{F}}(\bar{\rho}(N))$ from it one needs $D_{\mathcal{F}}(\bar{\rho}(M)) = D^*(M)$ at both $M = N$ and $M = 2N$, which is Conjecture 1 at both ranks. An earlier version of this work presented (80) as a proof of the doubling identity on flat spectra; that claim is false as stated and is not made here. What follows is what can be proved without assuming Conjecture 1.

12.2 Zero modes are exact

Adding a zero mode to a spectrum halves and duplicates every atom. For flat targets the overlap is then preserved exactly, not merely bounded.

Corollary 42 (Corollary J2: zero-mode exactness). *Let $\sigma \in \mathcal{F}$ have at least one zero mode, and let $\tilde{\sigma}$ be the free spectrum with one zero mode removed, so that $\sigma = \frac{1}{2}(\tilde{\sigma} \oplus \tilde{\sigma})$. Then for every $N \geq 1$,*

$$S(\bar{\rho}(2N), \sigma) = S(\bar{\rho}(N), \tilde{\sigma}), \quad (81)$$

and consequently $D_{\mathcal{F}}(\bar{\rho}(2N); \sigma) = D_{\mathcal{F}}(\bar{\rho}(N); \tilde{\sigma})$.

Proof. Let τ be the atoms of $\tilde{\sigma}$ in decreasing order. The atoms of σ are $\tau_1/2, \tau_1/2, \tau_2/2, \tau_2/2, \dots$, each halved and duplicated, and this list is already

decreasing. Since $2N$ is even, the truncation at rank $2N$ cuts at a pair boundary, so

$$S(\bar{\rho}(2N), \sigma) = \sum_{k \leq N} 2 \min\left(\frac{1}{2N}, \frac{\tau_k}{2}\right) = \sum_{k \leq N} \min\left(\frac{1}{N}, \tau_k\right) = S(\bar{\rho}(N), \tilde{\sigma}). \quad (82)$$

The distance statement is Lemma 6. $\square$

Corollary 42 says that every doubled sector carrying a zero mode is *equal*, not merely dominated, to a rank- N competitor. All of the difficulty in Conjecture 2 on flat spectra therefore sits in the sectors with no zero mode at all. It also gives the classical direction of Conjecture 2 in a stronger form: an equality of overlaps rather than an inequality of distances.

12.3 Coarse graining

The complementary tool removes a mode instead of adding one, and it is where the truncation enters.

Proposition 43 (Proposition J: the pair inequality). *For every $c, u, r \geq 0$,*

$$\min\left(\frac{c}{2}, u\right) + \min\left(\frac{c}{2}, ur\right) \leq \min(c, u(1+r)). \quad (83)$$

Proof. If $u(1+r) \geq c$ the right-hand side is c and the left is at most $c/2 + c/2 = c$. If $u(1+r) < c$ the right-hand side is $u + ur$ and the left is at most $u + ur$ term by term. $\square$

Deleting a mode of a free spectrum stays inside $\mathcal{F}$, and the pair inequality applies to the pairs this creates.

Lemma 44 (deleting a mode). *Let $\sigma \in \mathcal{F}$ have ratios $r_1, \dots, r_m$ and normaliser Z , and fix j . Pairing the 2^m subsets as $(S, S \cup \{j\})$ and adding the two atoms in each pair gives*

$$\frac{s_S}{Z} + \frac{s_S r_j}{Z} = \frac{s_S(1+r_j)}{Z^{(j)}(1+r_j)} = \frac{s_S}{Z^{(j)}}, \quad (84)$$

which are exactly the 2^{m-1} atoms of the free spectrum $\sigma^{(j)}$ with ratios $r_1, \dots, \hat{r}_j, \dots, r_m$. Hence $\sigma^{(j)} \in \mathcal{F}$, and for every $N \geq 1$ and every $c > 0$,

$$\sum_{\alpha} \min\left(\frac{c}{2}, \sigma_{\alpha}\right) \leq \sum_{\beta} \min(c, \sigma_{\beta}^{(j)}), \quad (85)$$

where both sums run over all atoms, untruncated.

Proof. Identity (84) is the cancellation of the factor $1 + r_j$ between the normaliser and the pair sum, and the right-hand side is the atom of $\sigma^{(j)}$ indexed by S , since $Z = Z^{(j)}(1 + r_j)$. The toggle- j pairing is a perfect matching of the subsets, so summing (83) with $u = s_S/Z^{(j)}$ and $r = r_j$ over the 2^{m-1} pairs gives (85): the left-hand side collects every atom of σ exactly once and the right-hand side every atom of $\sigma^{(j)}$ exactly once. No ordering is used, because both sums are over all atoms. $\square$

Corollary 45 (Corollary J3: truncation-free coarse graining). *Let $\sigma \in \mathcal{F}$ have m modes, and let $N \geq 1$ with $n = \lfloor \log_2 N \rfloor$. If $2^{m-1} \leq N$, equivalently $m \leq n + 1$, then for every mode j ,*

$$S(\bar{\rho}(2N), \sigma) \leq S(\bar{\rho}(N), \sigma^{(j)}) \leq S_{\max}(N). \quad (86)$$

The hypothesis is sharp: it cannot be dropped, and when it fails the conclusion can fail at every droppable mode.

Proof. The truncated overlap at the doubled rank is at most the untruncated one at half the level, $S(\bar{\rho}(2N), \sigma) = \sum_{\alpha \leq 2N} \min(\frac{1}{2N}, \sigma_\alpha) \leq \sum_{\alpha} \min(\frac{1}{2N}, \sigma_\alpha)$, and (85) with $c = 1/N$ bounds that by $\sum_{\beta} \min(1/N, \sigma_\beta^{(j)})$. The coarse-grained spectrum has 2^{m-1} atoms, so when $2^{m-1} \leq N$ this sum is untruncated and equals $S(\bar{\rho}(N), \sigma^{(j)})$; the last inequality is Definition 8. The equivalence $2^{m-1} \leq N \Leftrightarrow m \leq n + 1$ holds because $2^n \leq N < 2^{n+1}$.

For sharpness, take $N = 2$ and the balanced three-mode spectrum $r = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, so $Z = 27/8$ and the atoms are $(8, 4, 4, 4, 2, 2, 2, 1)/27$. Then $S(\bar{\rho}(4), \sigma) = \frac{1}{4} + 3 \cdot \frac{4}{27} = \frac{25}{36}$, while deleting any one mode leaves the four-atom spectrum $(4, 2, 2, 1)/9$ with $S(\bar{\rho}(2), \sigma^{(j)}) = \frac{4}{9} + \frac{2}{9} = \frac{2}{3} < \frac{25}{36}$ for all three j . Here $2^{m-1} = 4 > N = 2$, so the hypothesis fails and the conclusion fails with it, at every mode. The example does not violate Conjecture 2, because $S_{\max}(2) = 1$ — rank 2 is a power of two, so $\bar{\rho}(2)$ is itself free — and $\frac{25}{36} \leq 1$. It bounds the method, not the conjecture. $\square$

The obstruction is concrete: coarse graining bounds the doubled overlap by the *untruncated* overlap of the coarse-grained spectrum, and truncating that sum to rank N is legitimate only while the coarse-grained spectrum has at most N atoms. A proof that survives $2^{m-1} > N$ would have to control the tail $\sum_{\beta > N} \min(1/N, \sigma_\beta^{(j)})$, that is, to explain why a zero-mode-free spectrum with many modes cannot pack more than $S_{\max}(N)$ of its mass into the first $2N$ positions at weight $1/(2N)$.

12.4 The decomposition

The two tools cover complementary parts of $\mathcal{F}$: Corollary 42 covers every spectrum with a zero mode, by an equality, and the rest is by definition the zero-mode-free family.

Theorem 46 (Theorem Z: the zero-mode decomposition). *For every $N \geq 1$,*

$$S_{\max}(2N) = \max(S_{\max}(N), \Sigma_0(2N)), \quad (87)$$

$$\Sigma_0(M) = S_{\max}(M; \{\sigma \in \mathcal{F} : a(\sigma) = 0\}). \quad (88)$$

Proof. For $\geq$, Corollary 42 gives, for every $\sigma \in \mathcal{F}$, $S(\bar{\rho}(2N), \frac{1}{2}(\sigma \oplus \sigma)) = S(\bar{\rho}(N), \sigma)$ with $\frac{1}{2}(\sigma \oplus \sigma) \in \mathcal{F}$, so $S_{\max}(2N) \geq S_{\max}(N)$; and $\Sigma_0(2N) \leq S_{\max}(2N)$ because the zero-mode-free spectra are a subset of $\mathcal{F}$.

For $\leq$, take any $\sigma \in \mathcal{F}$. If $a(\sigma) \geq 1$ then $\sigma = \frac{1}{2}(\tilde{\sigma} \oplus \tilde{\sigma})$ and Corollary 42 gives $S(\bar{\rho}(2N), \sigma) = S(\bar{\rho}(N), \tilde{\sigma}) \leq S_{\max}(N)$. If $a(\sigma) = 0$ then $S(\bar{\rho}(2N), \sigma) \leq \Sigma_0(2N)$. Taking the supremum over σ gives $S_{\max}(2N) \leq \max(S_{\max}(N), \Sigma_0(2N))$. $\square$

Theorem 46 was verified as an exact identity on finite families of ratio lists, in rational arithmetic and with no optimiser involved, at every rank $2 \leq N \leq 14$ over families of up to 1287 lists [2]. For this paper it was re-derived on two nested families chosen so that the identity is exact rather than one-sided: lists of at most four ratios from a seven-point grid at rank N , and at most five at rank $2N$, so that the rank- $2N$ family is closed under adding a zero mode to a rank- N list and the rank- N family is closed under deleting one. On those families the identity holds exactly at every rank $2 \leq N \leq 10$, together with the two directions separately and with the iterated form below; the corresponding suprema are recorded in Table 2 of Appendix D. Closure under deletion alone is not enough, and a family that is not closed under *adding* a zero mode makes the identity fail at ranks where the optimal competitor uses all the available modes — an artefact of the family, not of the theorem.

Corollary 47 (Corollary Z1: the exact status of Conjecture 2 on flat spectra). *Conjecture 2 at rank N , restricted to flat spectra, holds if and only if*

$$\Sigma_0(2N) \leq S_{\max}(N). \quad (89)$$

By Corollary 45 every zero-mode-free spectrum with $b \leq n+1$ modes satisfies (89), and by Theorem 20 every one with $b \leq 2$ does, so the obstruction is contained in the family

$$a = 0, \quad b \geq \max(3, n+2), \quad (90)$$

at the doubled rank. Moreover Conjecture 1 at rank $2N$ alone implies Conjecture 2 at rank N .

Proof. The equivalence is Theorem 46 together with $S_{\max}(2N) \geq S_{\max}(N)$: the doubling identity $S_{\max}(2N) = S_{\max}(N)$ holds exactly when $\Sigma_0(2N) \leq S_{\max}(N)$. For the last statement, Conjecture 1 at rank $2N$ gives $S_{\max}(2N) = S^*(2N) = S^*(N)$ by (80), while the ansatz gives $S_{\max}(N) \geq S^*(N)$ and Theorem 46 gives $S_{\max}(2N) \geq S_{\max}(N)$. Chaining, $S^*(N) = S_{\max}(2N) \geq S_{\max}(N) \geq S^*(N)$, so all four are equal. $\square$

Corollary 47 makes Conjecture 2 on flat spectra strictly weaker than Conjecture 1: it needs the flat-spectrum value at one rank, the doubled one, and it needs control of a single explicit family. It does not prove it. The grid evidence on that family is one-sided: at every rank $2 \leq N \leq 14$ the zero-mode-free family with $b \geq n + 2$ stays strictly below $S_{\max}(N)$ on the grids searched, but a grid supremum is not a supremum.

Corollary 48 (Corollary Z3: iteration, and the reduction to odd ranks). *For every odd $q \geq 1$ and every $k \geq 1$,*

$$S_{\max}(2^k q) = \max(S_{\max}(q), \Sigma_0(2q), \Sigma_0(4q), \dots, \Sigma_0(2^k q)). \quad (91)$$

Consequently Conjecture 1 at every rank is equivalent to Conjecture 1 at every odd rank together with the zero-mode-free sectors at even ranks, where at even rank M only $b \geq n_M + 1$ can fail, $n_M = \lfloor \log_2 M \rfloor$.

Proof. Apply Theorem 46 at $N = 2^{k-1}q$, then at $N = 2^{k-2}q$, and so on; each step replaces $S_{\max}(2^i q)$ by $\max(S_{\max}(2^{i-1}q), \Sigma_0(2^i q))$ and the maximum is associative. For the second statement, at even rank $M = 2N$ the sectors that Corollary 45 controls are those with $m \leq n + 1 = n_M$, and those that Theorem 20 controls are those with $b \leq 2$; what is left is $b \geq n_M + 1$ with no zero mode. $\square$

Corollary 48 is a reduction, not an elimination, and it does not shrink the uncovered family $\mathcal{U}(N)$ of §11: one step of the reduction reaches $b \leq n + 2$, while the certified runs already cover $b \leq n + 3$.

12.5 Off the flat family

Nothing in this section applies to a non-flat target, and the paper claims nothing there. The reason is structural: coarse graining compares a spectrum with its mode-deletion by pairing atoms that differ by a factor r_j , and summing the resulting inequality against a target requires the target to be constant on the pairs, which a flat target is and a general one is not.

The conjecture was nevertheless stress-tested numerically on non-flat spectra, and the outcome is a cautionary result about testing it rather than about it. An unconstrained multistart optimisation over 89 spectra — peaked, exponential and power-law shapes, ranks up to 64 — reported five apparent violations of $D_{\mathcal{F}}(\frac{1}{2}(\rho \oplus \rho)) \geq D_{\mathcal{F}}(\rho)$, the worst a gap of -6.8×10^{-3} . Re-running the same probes with structure-aware starting points, in particular the symmetric free spectrum that is provably near-optimal for peaked targets, closed all five: over 32 probes the most negative gap became -4.1×10^{-5} , about 10^{-3} of the corresponding distance, and the two artefacts that had looked largest closed to $+2.6 \times 10^{-10}$ and $+3.9 \times 10^{-4}$ [2]. Both sides of a computed gap are upper bounds on the true distance, so a small negative gap is the signature of the doubled side

converging slightly worse, not of a violation. The structure that an unbiased optimiser misses is explicit: for $\rho = (0.99, 0.002, \dots, 0.002)$ with five small entries, the near-optimal free spectrum is the symmetric one with five equal biases $p_j = 1 - 0.99^{1/5}$, which gives the exact rational bound

$$D_{\mathcal{F}}(\rho) \leq 4.0161 \times 10^{-5}, \quad (92)$$

two orders of magnitude below what plain multistart found on the ρ side while the doubled side, searching one mode wider, found its analogue and appeared to win. Appendix D gives the bound (92) in exact arithmetic and records the probe data. Any claimed counterexample to Conjecture 2 would have to come with an exact interval certificate, because the symmetric optimum is easy to miss and a missed optimum looks exactly like a violation.

13 The maximal distance

Conjecture 3 proposes $3 - 2\sqrt{2} \approx 0.17157$ as the worst-case distance from the free set, over all spectra. This section states what is proved towards it, which is the corresponding statement about the ansatz values and one restricted class of competitors, and what is missing.

Theorem 49 (the flat family approaches $3 - 2\sqrt{2}$ from below). *$\sup_{N \geq 2} D^*(N) = 3 - 2\sqrt{2}$. The supremum is not attained at any rank, and it is approached along any sequence of ranks with $N/2^{\lfloor \log_2 N \rfloor} \rightarrow \sqrt{2}$, for instance $N_k = \lfloor 2^k \sqrt{2} \rfloor$. If Conjecture 1 holds then $\sup_N D_{\mathcal{F}}(\bar{\rho}(N)) = 3 - 2\sqrt{2}$, again not attained.*

Proof. Corollary 10 gives $D^*(N) \leq 3 - 2\sqrt{2}$ with equality only at $x = \sqrt{2}$, and $x = N/2^n$ is rational while $\sqrt{2}$ is not, so equality never occurs. For the approach, $x_k = N_k/2^k \rightarrow \sqrt{2}$ because $0 \leq 2^k \sqrt{2} - N_k < 1$, and D^* is continuous at $\sqrt{2}$, so $D^*(N_k) \rightarrow 3 - 2\sqrt{2}$. The last statement is Conjecture 1 applied at each rank. $\square$

The values are explicit and small: D^* vanishes at every power of two, rises to $1/6$ at $x = 3/2$ and to its maximum $3 - 2\sqrt{2}$ in the limit $x \rightarrow \sqrt{2}$, and falls back to 0 as $x \rightarrow 2$. The two largest values among the certified ranks of Table 1 are 0.171196 at $x = 23/16$ ($N = 23$) and 0.170455 at $x = 11/8$ ($N = 11$ and $N = 22$), both within 7×10^{-4} of $3 - 2\sqrt{2}$. Since x takes every dyadic rational value in $[1, 2)$ as N ranges over the integers, and D^* is continuous there, the set $\{D^*(N)\}$ is dense in $[0, 3 - 2\sqrt{2}]$.

Two partial results towards Conjecture 3 are available without assuming Conjecture 1. The first is from the earlier literature: restricted to spectra of rank at most four, the maximum of $D_{\mathcal{F}}$ is $1/6$ [3, 6], which is consistent with Theorem 49 since $1/6 < 3 - 2\sqrt{2}$ and the flat value at $x = 3/2$ is exactly $1/6$. The second is a consequence of the Main Theorem: no competitor with many zero modes does better than the ansatz.

Corollary 50 (many-zero-mode competitors cannot beat the flat target).
For every $N \geq 2$ and every $\sigma \in \mathcal{F}$ with at least $n = \lfloor \log_2 N \rfloor$ zero modes,

$$D_{\mathcal{F}}(\bar{\rho}(N); \sigma) \geq D^*(N), \quad (93)$$

so within that class of competitors the flat target is optimal and the distance is exactly $D^*(N)$. In particular no such competitor brings $\bar{\rho}(N)$ closer to $\mathcal{F}$ than $D^*(N)$, and $D^*(N) \rightarrow 3 - 2\sqrt{2}$ along the sequence of Theorem 49.

Proof. This is Theorem 19. □

What separates this from Conjecture 3 is a bound that holds for *every* competitor and *every* target. For flat targets that bound is Conjecture 1, and it is open exactly on the residual family (71); for non-flat targets nothing in this paper applies at all, since the whole analysis is built on the flat level $1/N$ and on the truncation weights $w_i(y) = \text{clamp}(y - i + 1, 0, 1)$ that the flat target induces. A proof of Conjecture 3 would therefore have to do two things beyond this work: settle the residual family for flat targets, and find the analogue of the sector decomposition for a general target, where the level is no longer constant and the clipping pattern of a competitor depends on the shape of ρ rather than on its rank alone. The observed behaviour is consistent with Conjecture 3: no spectrum examined here, flat or not, has distance exceeding $3 - 2\sqrt{2}$, and the flat family gets arbitrarily close. That is an observation about the spectra examined, not a bound.

14 Discussion and open problems

The honest summary of this paper is shorter than its list of theorems. Conjecture 1 is proved on four explicit infinite families of free spectra — those with at least n zero modes, those with at most two positive modes, those with at most n modes in total, and every spectrum with exactly $n - 1$ zero modes whatever its number of positive modes, the last containing the sector with exactly three positive modes, which is proved separately — and reduced, exactly and in both directions, to a single negative existential statement about interior local maxima on a residual family that is now finite in width at each zero-mode count. Conjecture 2 is not proved, not even on flat spectra, but it is reduced to one explicit family and derived from Conjecture 1 at the doubled rank alone. Conjecture 3 is proved for the ansatz values and is otherwise open. Rigorous two-sided brackets are available at 22 ranks on the covered family.

14.1 The boundaries of the results

Each theorem here has a scope that is narrower than its statement suggests if the statement is read alone, and the narrowings are of different kinds.

- *The four eliminated families are families of competitors.* Theorem 19 and Theorem 33 bound $D_{\mathcal{F}}(\bar{\rho}(N); \sigma)$ for σ in a restricted class, which is a lower bound on the distance restricted to that class. They prove $D_{\mathcal{F}}(\bar{\rho}(N)) = D^*(N)$ only when the class exhausts $\mathcal{F}$, which happens at $b \leq 2$ and at $a \geq n$ but not at $a = n - 1$ with $b \geq 3$.
- *Theorem 37 is an equivalence, not a proof.* It says the conjecture holds if and only if no interior local maximum exceeds $S^*(N)$. The direction that does the work is the reduction of an unbounded variational problem to a countable family of compact boxes; the hypothesis is untouched.
- *Proposition 26 is necessary, not sufficient.* A point satisfying $p_j = A_j/U$ for every j need not be a local maximum, and a local maximum on a clipping wall need not satisfy the equalities, only the one-sided inequalities.
- *Theorem 31 truncates rather than eliminates at $\nu \geq 2$.* It bounds the mode count, $b \leq 2^{\nu+1} - 1$, and says nothing about the sectors that remain. Its threshold $b \geq \lceil y \rceil$ is sharp by Proposition 35, so improving it is a different argument, not a better constant.
- *The certificates bracket $D_{\mathcal{F}}^{\text{cov}}$, not $D_{\mathcal{F}}$.* Theorem 41 gives the transfer condition (78) explicitly, and that condition is open. The excess ε_e measures the resolution of the certificate, not the truth of the conjecture, and cannot be driven to zero at any finite resolution because every residual box contains a face point at which the objective equals $S^*(N)$.
- *Theorem 46 is a decomposition.* It identifies where Conjecture 2 can fail on flat spectra; it does not exclude the failure. Its grid verifications are exact identities on finite families, and a family that is not closed under adding a zero mode makes the identity fail for reasons that have nothing to do with the theorem — a trap that is easy to fall into and is recorded in §12.
- *The searches are searches.* The interior-maximum search over 85 residual sectors and the non-flat doubling probes are multistart local optimisations. Their negative results are evidence; the five apparent violations of Conjecture 2 they initially produced were artefacts of non-convergence, and the same failure mode could produce an artefact in the other direction.
- *Nothing here applies away from flat targets.* Every argument uses the constant level $1/N$ and the weights $w_i(y)$ it induces. Conjecture 2 for a general ρ , and part (D) of the source problem — extending free descriptions away from the $\mathbb{Z}_{2^n}$ fixed point — are untouched.

14.2 Open problems

The following are the load-bearing open items, in the order in which they block the conjectures.

1. Interior local maxima above the ansatz value. By Theorem 37 this is Conjecture 1. Settled: $b \leq 2$ everywhere (Theorem 20), every b at $a = n - 1$ (Theorem 31 plus Theorem 27), and $b \geq \lceil y \rceil$ at every a (Theorem 31). Open: the sectors (71), smallest first at $(n - 2, 3)$, where $y \in [4, 8)$ lets up to eight atoms carry weight and the four-atom truncation that closes the chamber analysis no longer applies. By Proposition 26 a counterexample would have to satisfy $b - 1$ exact algebraic constraints $B_j/Z^{(j)} = B_k/Z^{(k)}$, which is a small algebraic set to search — and, if a maximum above $S^*(N)$ were found anywhere, a refutation of Conjecture 1.

2. The uncovered family $\mathcal{U}(N)$. This is the entire obstruction to reading the certified brackets of Table 1 as bounds on $D_{\mathcal{F}}$ itself. With the eliminations of this paper the family can be taken to be $\mathcal{U}'(N) = \{a \leq n - 2, b \geq n + 4\}$ (Remark 40), since every sector with $a = n - 1$ is discharged analytically at every b ; it is still untouched by any argument here, because at $b \geq n + 4$ the boxes have dimension at least five and up to 2^b atoms, and no rank-uniform statement about them is known. The two open families are of different kinds: controlling $\mathcal{U}'(N)$ at each fixed N is a computation, and controlling the residual family (71) uniformly in N is an analysis. Conjecture 1 needs the second, and the brackets of §11 need the first.

3. A modulus of continuity for the interval bound. The observed decay $\varepsilon_e \approx C_N 2^{-e}$, with C_N between 0.0627 and 0.7243 across the certified ranks, is what a boundary layer of thickness 2^{-e} around the face point of Lemma 36 produces, but converting it into $\sup S \leq S^*(N)$ on the covered sectors requires a proved constant L with $B_e \leq S^*(N) + L2^{-e}$, that is, a Lipschitz control of how much the interval evaluation of the objective inflates over a box of width 2^{-e} . With such an L and completed runs at every tolerance, the covered family would be settled exactly at each certified rank. It would not settle $\mathcal{U}(N)$, since a modulus bounds B_e only where a run completes.

4. The zero-mode-free doubled family $a = 0, b \geq n + 2$ at rank $2N$. By Corollary 47 this is the whole of Conjecture 2 on flat spectra. Everything else is closed: $a \geq 1$ by the exact equality of Corollary 42, $m \leq n + 1$ by the coarse graining of Corollary 45, $b \leq 2$ by Theorem 20. The obstruction is the tail $\sum_{\beta > N} \min(1/N, \sigma_{\beta}^{(j)})$ that coarse graining cannot truncate once

$2^{m-1} > N$, and the sharpness witness of Corollary 45 shows the truncation hypothesis is not removable. A proof would have to explain why a zero-mode-free spectrum with many modes cannot pack more than $S_{\max}(N)$ of its mass into the first $2N$ positions.

5. Conjecture 2 for non-flat spectra. Nothing here applies: the pairing that makes coarse graining work does not align with the sorted order of a non-flat target. The stress tests find no violation, but they also show that the symmetric optimum is easy for an unbiased optimiser to miss, so both a claimed counterexample and a claimed proof would need exact certificates. Certifying $D_{\mathcal{F}}(\rho) = D_{\mathcal{F}}(\frac{1}{2}(\rho \oplus \rho))$ exactly for even one non-flat ρ is beyond every method used here.

6. Does a non-flat optimum sit on a normalisation boundary? For flat targets the maximiser satisfies $Z = N$ exactly, and Theorem 9 proves this is optimal rather than incidental. What replaces N for a general target, and whether an analogous boundary statement holds, is open. Such a statement would give a cheap necessary condition for optimal free approximations well beyond the fixed points that motivated the problem.

7. Pushing the truncation below $\lceil y \rceil$ at $\nu \geq 2$. Lemma 30 and Theorem 31 fail at $b = \lceil y \rceil - 1$ by Proposition 35, where the varied smallest mode is visible to the objective. At $\nu = 2$ that leaves exactly the five sectors $(n - 2, b)$, $3 \leq b \leq 7$, so a single new idea about the eight-atom truncation would close the smallest open case of the paper.

14.3 Scope

The results are theorems about a variational problem that arises in the study of free-fermion descriptions of topological phases: the objects are entanglement spectra and Gaussian states, and the question — how far a given spectrum is from the free set — is the one the interaction distance was introduced to answer [6]. No physical system is measured or simulated here, and no observable is predicted. The method is that of exact optimisation: sector decomposition, piecewise-affine concavity, polynomial sign lemmas and rational interval arithmetic, with numerical search used only where its status as a search is stated. Readers interested in the physics should take away the four families on which the free description of a flat spectrum is provably optimal, and the fact that the remaining cases are reduced to one clean analytic question; readers interested in the analysis should take away Lemma 24 and Theorem 31, which are the two tools that made the reduction possible and which do not depend on the flat target beyond the truncation weights $w_i(y)$.

15 Conclusion

The fermionic interaction distance of a flat entanglement spectrum is a finite-dimensional variational problem in disguise: the free set is a union of sector boxes indexed by the numbers of zero and positive modes, the objective on a box is a weighted sum of clipped subset products, and only the first $\lceil y \rceil$ of the 2^b atoms carry any weight. Once that structure is exposed, the conjectured value $3 - x - 2/x$ can be derived rather than fitted, its producer shown to be the unique maximiser of its own family with the normaliser exactly on the boundary $Z = N$, and large parts of the free set eliminated in closed form at every rank simultaneously.

What the analysis achieves is four infinite families of competitors on which the free description of a flat spectrum is provably optimal — at least n zero modes, at most two positive modes, at most n modes in total, and every mode count at $n - 1$ zero modes, the last containing the sector with exactly three positive modes that is proved separately — together with an exact equivalence between the full conjecture and the non-existence of interior local maxima above the ansatz value, an exact decomposition that reduces the doubling conjecture on flat spectra to one explicit family and derives it from the flat-spectrum conjecture at the doubled rank alone, and rigorous rational brackets $D_{\mathcal{F}}^{\text{cov}}(\bar{\rho}(N)) \in [D^*(N) - \varepsilon, D^*(N)]$ at 22 ranks with ε at most 4.382×10^{-5} . The residual family is now finite in width at every zero-mode count, and its smallest member is the sector $(n - 2, 3)$.

What it does not achieve is any of the three conjectures. The flat-spectrum conjecture is open on the residual family, where the evidence is a local search; the doubling conjecture is open even on flat spectra, where the whole obstruction is the zero-mode-free family with $b \geq n + 2$ at the doubled rank, and open off the flat family, where no method here applies and five apparent counterexamples turned out to be optimiser artefacts; and the global maximum $3 - 2\sqrt{2}$ is proved only for the ansatz values. The certified brackets bound a restricted variational problem, and reading them as bounds on the interaction distance itself requires a hypothesis on the uncovered sectors that is stated explicitly rather than assumed. The path forward is narrow and concrete: an interior local maximum above $S^*(N)$ would refute the conjecture, a modulus of continuity for the interval bound would make the certificates exact on the covered family, and the eight-atom truncation of the sector $(n - 2, 3)$ is where both the concavity argument and the mode-count descent stop working.

A The case analysis for two positive modes

This appendix supplies the bounds quoted in Theorem 23, the six cases of the sector $(n - 1, 2)$, and the enumeration of critical cells used in the proof of

Theorem 27. Throughout, $y = 2x \in [2, 4]$, $c = 1/y$, $S_1^*(y) = y/2 + 4/y - 2$, and the coordinates are those of (38):

$$u \in [\frac{1}{2}, 1], \quad v \in [0, \frac{1}{2}], \quad u + v \leq 1, \quad \beta = u(1 - v) \in [\frac{1}{4}, 1], \quad (94)$$

$$\tau = (\beta, (1 - u)(1 - v), uv, v(1 - u)), \quad (95)$$

with the identities $\tau_2 = \beta(1/u - 1)$ and $\tau_3 = u - \beta$ of (41). The faces of the box are handled elsewhere and are excluded from the case analysis: $s_2 = 0$ ($v = 0$) is the sector $(n - 1, 1)$, bounded by Lemma 16; $s_1 = 1$ ($u = \frac{1}{2}$) is the face identity (70) onto the sector $(n, 1)$, bounded by Theorem 17; and $s_1 = s_2$ ($u + v = 1$) is an ordering wall on which the objective is continuous. Every bound below is an exact rational expression in y , β and u , and each is at most $S_1^*(y)$.

Case I: $2 \leq y \leq 3$

The weights are $(1, 1, \lambda, 0)$ with $\lambda = y - 2 \in [0, 1]$, the fourth atom carries no weight, and $S_{\text{unc}} = 1 - v + \lambda uv$ by (42). Since $w_1 = 1$ and $w_i \leq 1$,

$$S \leq \text{UB} = S_{\text{unc}} - \left(\beta - \frac{1}{y}\right)^+, \quad (96)$$

which is valid but, as Remark 22 records, not sufficient. The four patterns are determined by how many of $\tau_1 \geq \tau_2 \geq \tau_3$ exceed c .

I.1: $\tau_1 \leq c$, nothing clipped. Then $S = S_{\text{unc}} = 1 - v + \lambda uv$, which is increasing in u , so it is bounded by its value at the largest feasible u . The constraint $\beta = u(1 - v) \leq 1/y$ gives $u \leq 1/(y(1 - v))$ and the ordering gives $u \leq 1 - v$, hence

$$S \leq 1 - v + \lambda v \min\left(1 - v, \frac{1}{y(1 - v)}\right) = \max(h(v), g(v)), \quad (97)$$

where the two branches are

$$h(v) = (1 - v)(1 + \lambda v) \quad \text{on } v \geq 1 - \frac{1}{\sqrt{y}}, \quad (98)$$

$$g(v) = 1 - v + \frac{\lambda v}{y(1 - v)} \quad \text{on } v \leq 1 - \frac{1}{\sqrt{y}}, \quad (99)$$

the ranges being where each of the two candidates for u_{max} is the smaller, and the junction being $v = 1 - 1/\sqrt{y}$ where $(1 - v)^2 = 1/y$.

On the clipping branch, $u_{\text{max}} = 1/(y(1 - v)) \geq \frac{1}{2}$ forces $v \geq 1 - 2/y$, so the branch is the interval $v \in [1 - 2/y, 1 - 1/\sqrt{y}]$. There

$$g'(v) = -1 + \frac{\lambda}{y(1 - v)^2}, \quad (100)$$

which is increasing in v , and at the right endpoint, where $y(1-v)^2 = 1$, it equals $-1 + \lambda = y - 3 \leq 0$ in Case I. Hence $g' \leq 0$ on the whole branch, g is decreasing, and its maximum is at the left endpoint. The two endpoint values are

$$g\left(1 - \frac{2}{y}\right) = \frac{2}{y} + \frac{(y-2)^2}{2y} = S_1^*(y), \quad g\left(1 - \frac{1}{\sqrt{y}}\right) = h\left(1 - \frac{1}{\sqrt{y}}\right), \quad (101)$$

the first an exact identity, so $g \leq S_1^*(y)$ on the branch with equality only at $v = 1 - 2/y$, where $u = u_{\max} = 1/2$ and $(s_1, s_2) = (1, x-1)$ is the face point (37).

On the diagonal branch, h is decreasing because $h'(v) = \lambda - 1 - 2\lambda v < 0$ for $\lambda \leq 1$, so its maximum is at the junction $v = 1 - 1/\sqrt{y}$, where $u = 1 - v = 1/\sqrt{y}$ and $\beta = 1/y$. Writing $q = \sqrt{y} \in [\sqrt{2}, \sqrt{3}]$,

$$h\left(1 - \frac{1}{q}\right) = \frac{1}{q} \left(1 + (q^2 - 2)\left(1 - \frac{1}{q}\right)\right) = \frac{(y-1)(q-1) + 1}{y}, \quad (102)$$

$$h\left(1 - \frac{1}{q}\right) - S_1^*(y) = -\frac{P(q)}{2q^2}, \quad (103)$$

with $P(q) = q^4 - 2q^3 - 2q^2 + 2q + 4$, which is positive on $[\sqrt{2}, \sqrt{3}]$ by Lemma 51. So the diagonal branch is strictly below $S_1^*(y)$.

I.2: $\tau_1 > c \geq \tau_2$, only the largest atom clipped. The objective is exactly

$$S = \frac{1}{y} + \tau_2 + \lambda\tau_3 = \frac{1}{y} + \beta\left(\frac{1}{u} - 1\right) + \lambda(u - \beta) = \Theta(u), \quad (104)$$

$$\Theta(u) = \frac{\beta}{u} + \lambda u - \lambda\beta - \beta + \frac{1}{y}, \quad (105)$$

and $\Theta''(u) = 2\beta/u^3 > 0$, so Θ is convex and its maximum over the feasible interval is at an endpoint. The constraints are $\beta > 1/y$; $\tau_2 \leq c$, that is $\beta(1/u - 1) \leq 1/y$, that is $u \geq y\beta/(y\beta + 1)$; the ordering $u \leq \sqrt{\beta}$; $u \geq \frac{1}{2}$; and $v \geq 0$, that is $u \geq \beta$. The left endpoint is therefore

$$u_{\min} = \begin{cases} \frac{y\beta}{y\beta + 1}, & \beta \leq 1 - \frac{1}{y}, \\ \beta, & \beta \geq 1 - \frac{1}{y}, \end{cases} \quad (106)$$

the two expressions crossing exactly at $\beta = 1 - 1/y$. At the first, $\tau_2 = c$ and

$$S = \frac{2}{y} + \lambda\tau_3 \leq \frac{2}{y} + \lambda \frac{y-2}{2y} = \frac{2}{y} + \frac{(y-2)^2}{2y} = S_1^*(y), \quad (107)$$

using the bound $\tau_3 = uv \leq (y-2)/(2y)$ of (112) below. At the second, $v = 0$, which is the face $s_2 = 0$, and $S = 1 - \beta + 1/y < 2/y \leq S_1^*(y)$, the last inequality being $2/y \leq y/2 + 4/y - 2 \Leftrightarrow 0 \leq (y-2)^2/(2y)$. At the right endpoint $u = \sqrt{\beta} = z$,

$$\Theta(z) = (1 + \lambda)z(1 - z) + \frac{1}{y} = (y - 1)z(1 - z) + \frac{1}{y}, \quad (108)$$

and $z \mapsto z(1 - z)$ is decreasing on $[\frac{1}{2}, 1]$, while $z = \sqrt{\beta} \geq \max(\frac{1}{2}, 1/\sqrt{y})$ because $\beta \geq \frac{1}{4}$ and $\beta > 1/y$. For $y \in [2, 4)$ the binding constraint is $z \geq 1/\sqrt{y} = 1/q$, so

$$\Theta(z) \leq (y - 1)\frac{1}{q}\left(1 - \frac{1}{q}\right) + \frac{1}{y} = \frac{(y - 1)(q - 1) + 1}{y} \quad (109)$$

$$= S_1^*(y) - \frac{P(q)}{2q^2} < S_1^*(y) \quad (110)$$

by Lemma 51. For $y \geq 4$, which Case I excludes, the binding constraint would be $z \geq \frac{1}{2}$ and the same computation gives $\Theta \leq (y - 1)/4 + 1/y$.

I.3: $\tau_1, \tau_2 > c \geq \tau_3$. Then $S = 2/y + \lambda\tau_3$ and the bound (112) gives

$$S \leq \frac{2}{y} + \lambda\frac{y-2}{2y} = S_1^*(y), \quad (111)$$

with equality only at $u = \frac{1}{2}$, $v = 1 - 2/y$, the face point.

The bound on $\tau_3 = uv$.

$$\tau_2 = (1 - u)(1 - v) \geq \frac{1}{y} \implies v \leq 1 - \frac{1}{y(1 - u)}, \quad (112)$$

$$uv \leq m(u) = u - \frac{u}{y(1 - u)} \leq \frac{y-2}{2y}, \quad (113)$$

where m is defined on $u \in [\frac{1}{2}, 1 - 1/y]$ and decreasing there, because $m'(u) = 1 - \frac{1}{y(1-u)^2} < 0$ whenever $y(1-u)^2 < 1$, and on that interval $y(1-u)^2 \leq y/4 \leq 3/4 < 1$. Hence $uv \leq m(\frac{1}{2}) = \frac{1}{2} - \frac{1}{y} = \frac{y-2}{2y}$, with equality only at $u = \frac{1}{2}$ and $v = 1 - 2/y$.

I.4: $\tau_1, \tau_2, \tau_3 > c$. Infeasible. The three atoms would satisfy $\tau_1 + \tau_2 + \tau_3 > 3/y \geq 1$ for $y \leq 3$, while $\sum_i \tau_i = 1$.

Case II: $3 \leq y < 4$

The weights are $(1, 1, 1, \kappa)$ with $\kappa = y - 3 \in [0, 1)$, $S_{\text{unc}} = 1 - (4 - y)\psi(u)$ by (43) with $\psi(u) = v(1 - u) = (1 - u) - \beta(1/u - 1)$, and the coarse bound (96) reads $S \leq 1 - F$ with

$$F = (4 - y)\psi(u) + \left(\beta - \frac{1}{y}\right)^+, \quad G = \frac{(4 - y)(y - 2)}{2y} = 1 - S_1^*(y) \geq 0. \quad (114)$$

So it is enough to prove $F \geq G$. The function ψ is concave in u at fixed β , since $\psi''(u) = -2\beta/u^3 < 0$, so F is concave in u on each of the two regimes and its minimum over the feasible interval $u \in [\max(\frac{1}{2}, \beta), \sqrt{\beta}]$ is at an endpoint. The three endpoint values are

$$F(\frac{1}{2}) = (4 - y)\left(\frac{1}{2} - \beta\right) + \left(\beta - \frac{1}{y}\right)^+, \quad (115)$$

$$F(\beta) = \left(\beta - \frac{1}{y}\right)^+, \quad (116)$$

$$F(\sqrt{\beta}) = (4 - y)(1 - \sqrt{\beta})^2 + \left(\beta - \frac{1}{y}\right)^+, \quad (117)$$

where $F(\beta)$ corresponds to $v = 0$ and is feasible only when $\beta \geq \frac{1}{2}$.

II.0: $\beta \leq 1/y$, nothing clipped. Then $F = (4 - y)\psi(u)$ and it suffices to show $\psi(u) \geq (y - 2)/(2y)$. Since $\beta \leq 1/y \leq \frac{1}{3} < \frac{1}{2}$, the left endpoint is $u = \frac{1}{2}$ and

$$\psi(\frac{1}{2}) = \frac{1}{2} - \beta \geq \frac{1}{2} - \frac{1}{y} = \frac{y - 2}{2y}, \quad (118)$$

$$\psi(\sqrt{\beta}) = (1 - \sqrt{\beta})^2 \geq \left(1 - \frac{1}{\sqrt{y}}\right)^2 \geq \frac{y - 2}{2y}, \quad (119)$$

the last step being the identity of Lemma 52: with $q = \sqrt{y}$, $(1 - 1/q)^2 - (q^2 - 2)/(2q^2) = (q - 2)^2/(2q^2) \geq 0$. Concavity gives the same lower bound in between, so $F \geq (4 - y)(y - 2)/(2y) = G$.

II.clipped: $\beta > 1/y$. Suppose first $\beta \leq \frac{1}{2}$, so the feasible interval is $u \in [\frac{1}{2}, \sqrt{\beta}]$. At the left endpoint,

$$\begin{aligned} F(\frac{1}{2}) - G &= (4 - y)\left(\frac{1}{2} - \beta\right) + \beta - \frac{1}{y} - \frac{(4 - y)(y - 2)}{2y} \\ &= \frac{3 - y}{y} + \beta(y - 3) > \frac{3 - y}{y} + \frac{y - 3}{y} = 0, \end{aligned} \quad (120)$$

using $\beta > 1/y$ and $y \geq 3$. At the right endpoint, $F(\sqrt{\beta})$ is increasing in β on $[\frac{1}{4}, \frac{1}{2}]$, because $\frac{d}{d\beta}[(4 - y)(1 - \sqrt{\beta})^2 + \beta] = 1 - (4 - y)\frac{1 - \sqrt{\beta}}{\sqrt{\beta}} \geq 1 - (4 - y) \geq$

$y - 3 \geq 0$, using $\sqrt{\beta} \geq \frac{1}{2}$; hence

$$F(\sqrt{\beta}) \geq (4 - y) \left(1 - \frac{1}{\sqrt{y}}\right)^2 + \frac{1}{y} - \frac{1}{y} \geq (4 - y) \frac{y - 2}{2y} = G, \quad (121)$$

where the last step is the identity of Lemma 52 multiplied by $4 - y \geq 0$. If instead $\beta > \frac{1}{2}$ the left endpoint is $u = \beta$, where $\psi = 0$ and $F = \beta - \frac{1}{y} > \frac{1}{2} - \frac{1}{y} = \frac{y-2}{2y} \geq G$ because $4 - y \leq 1$; and $F(\sqrt{\beta}) \geq \beta - \frac{1}{y}$ again exceeds G for the same reason. So $F \geq G$ throughout, with equality only at $y = 3$, $\beta = \frac{1}{2}$, $u = \frac{1}{2}$, the corner $(s_1, s_2) = (1, 0)$, which lies on two faces and is bounded by Lemma 16 and Theorem 17.

The critical cells of Theorem 27

In chamber P of the sector $(n - 1, 3)$ with $y \in [3, 4)$ the four weighted atoms are $\tau_i = (1, s_1, s_2, s_3)/Z_3$ with weights $(1, 1, 1, f)$, $f = y - 3$, and only τ_4 uses mode 3. The clipped set is an initial segment $\{1, \dots, J\}$, since the atoms are decreasing, and $B_3 = f$ when τ_4 is unclipped, that is when $J \leq 3$, and $B_3 = 0$ when $J = 4$. In the notation of (52), with $Z^{(3)} = (1 + s_1)(1 + s_2)$ and bases 1, s_1, s_2 for τ_1, τ_2, τ_3 ,

$$F_3 = \begin{cases} 1 + s_1 + s_2, & J = 0, \\ s_1 + s_2, & J = 1, \\ s_2, & J = 2, \\ 0, & J \geq 3, \end{cases} \quad \text{and} \quad W_{\text{clip}} = J, \quad (122)$$

The condition $B_3 = F_3$ of Lemma 24 is therefore possible only for $J = 1$ and $J = 2$: at $J = 0$ it would require $f = 1 + s_1 + s_2 \geq 1 > f$, and at $J \geq 3$ it would require $f = 0$, that is $y = 3$, which is the boundary of Case I. The two surviving cells give, by (55),

$$\begin{aligned} J = 1: \quad f &= s_1 + s_2, \quad Z^{(3)} = (1 + s_1)(1 + s_2) = 1 + f + s_1 s_2 \geq 1 + f, \\ S &= \frac{1}{y} + \frac{B_3}{Z^{(3)}} \leq \frac{1}{y} + \frac{f}{1 + f}, \end{aligned} \quad (123)$$

$$\begin{aligned} J = 2: \quad f &= s_2 \leq s_1, \quad Z^{(3)} = (1 + s_1)(1 + s_2) \geq (1 + f)^2, \\ S &= \frac{2}{y} + \frac{B_3}{Z^{(3)}} \leq \frac{2}{y} + \frac{f}{(1 + f)^2}. \end{aligned} \quad (124)$$

These are the two critical values (61), and Lemma 53 shows that both are strictly below $S_1^*(y)$ on $[3, 4]$. The enumeration is exhaustive: $J = 4$ forces $B_3 = 0$ and is the monotone case, $J = 3$ forces $F_3 = 0 < B_3$, and $J = 0$ is infeasible as shown.

B Polynomial sign lemmas

Three exact sign statements are used in Appendices A and in the proof of Theorem 27. Each is proved here without floating point: the sign of a polynomial at an irrational point is decided by an integer comparison after isolating the radical, and monotonicity is decided by the sign of the derivative at one rational point together with the convexity of the derivative.

Lemma 51 (the quartic of Case I). *$P(q) = q^4 - 2q^3 - 2q^2 + 2q + 4$ is strictly positive and strictly decreasing on $[\sqrt{2}, \sqrt{3}]$, with*

$$P(\sqrt{2}) = 4 - 2\sqrt{2} > 0, \quad P(\sqrt{3}) = 7 - 4\sqrt{3} > 0. \quad (125)$$

Proof. $P'(q) = 4q^3 - 6q^2 - 4q + 2 = 2g(q)$ with $g(q) = 2q^3 - 3q^2 - 2q + 1$, and $g'(q) = 6q^2 - 6q - 2$. The larger root of g' is $(3 + \sqrt{21})/6$, which is below $\sqrt{2}$ because $6\sqrt{2} - 3 > 0$ and $(6\sqrt{2} - 3)^2 = 81 - 36\sqrt{2} > 21$, the last inequality being $81 - 21 = 60 > 36\sqrt{2} \Leftrightarrow 3600 > 2592$. So g is increasing on $[\sqrt{2}, \sqrt{3}]$, and its largest value there is

$$g(\sqrt{3}) = 6\sqrt{3} - 9 - 2\sqrt{3} + 1 = 4\sqrt{3} - 8 < 0, \quad (126)$$

since $4\sqrt{3} < 8 \Leftrightarrow 48 < 64$. Hence $g < 0$ and $P' < 0$ on the whole interval, so P is decreasing and $P(q) \geq P(\sqrt{3}) = 9 - 6\sqrt{3} - 6 + 2\sqrt{3} + 4 = 7 - 4\sqrt{3}$, which is positive because $7 - 4\sqrt{3} > 0 \Leftrightarrow 49 > 48$. The other value is $P(\sqrt{2}) = 4 - 4\sqrt{2} - 4 + 2\sqrt{2} + 4 = 4 - 2\sqrt{2} > 0$ because $16 > 8$. $\square$

Lemma 52 (the perfect square of Case II). *For every $q > 0$,*

$$\left(1 - \frac{1}{q}\right)^2 - \frac{q^2 - 2}{2q^2} = \frac{(q - 2)^2}{2q^2} \geq 0, \quad (127)$$

with equality only at $q = 2$. Equivalently, for every $y > 0$, $(1 - 1/\sqrt{y})^2 \geq (y - 2)/(2y)$, with equality only at $y = 4$.

Proof. Put over the common denominator $2q^2$: the left-hand side is $[2(q - 1)^2 - (q^2 - 2)]/(2q^2) = (q^2 - 4q + 4)/(2q^2) = (q - 2)^2/(2q^2)$. The equivalent form is $q^2 = y$. $\square$

Lemma 53 (the two endpoint polynomials of Theorem 27). *On $[3, 4]$,*

$$p_1(y) = y^3 - 8y^2 + 20y - 12 > 0, \quad \text{with minimum } \frac{76}{27} \text{ at } y = \frac{10}{3}, \quad (128)$$

and

$$p_2(y) = (y - 2)^4 - 2y(y - 3) > 0, \quad \text{with minimum at least } \frac{19}{25}. \quad (129)$$

Consequently the two critical values of (123) and (124) satisfy

$$S_1^*(y) - \left(\frac{1}{y} + \frac{f}{1+f} \right) = \frac{p_1(y)}{2y(y-2)} > 0, \quad (130)$$

$$S_1^*(y) - \left(\frac{2}{y} + \frac{f}{(1+f)^2} \right) \geq \frac{p_2(y)}{2y(y-2)^2} > 0, \quad (131)$$

where $f = y - 3 \in [0, 1]$ and the second inequality uses $(1 + s_1)(1 + s_2) \geq (1 + f)^2$ and $(1 + s_1)(1 + s_2) \geq 1 + f + s_1 s_2$ as in (123)–(124).

Proof. For p_1 , $p_1'(y) = 3y^2 - 16y + 20 = (3y - 10)(y - 2)$, so p_1 decreases on $[3, 10/3]$ and increases on $[10/3, 4]$, with its minimum at $y = 10/3$:

$$p_1\left(\frac{10}{3}\right) = \frac{1000}{27} - \frac{800}{9} + \frac{200}{3} - 12 = \frac{1000 - 2400 + 1800 - 324}{27} = \frac{76}{27} > 0. \quad (132)$$

For p_2 , write $t = y - 2 \in [1, 2]$, so $p_2 = t^4 - 2(t+2)(t-1) = t^4 - 2t^2 - 2t + 4$ and $p_2' = 4t^3 - 4t - 2 = 2(2t^3 - 2t - 1)$. The cubic $2t^3 - 2t - 1$ has derivative $6t^2 - 2 > 0$ for $t \geq 1$, so it is increasing there, and it takes the values -1 at $t = 1$ and $432/125 - 12/5 - 1 = 7/125 > 0$ at $t = 6/5$. Hence p_2 has a single stationary point $t^* \in (1, 6/5)$, which is its minimum on $[1, 2]$, and at that point $2t^{*3} = 2t^* + 1$, so $t^{*3} = t^* + \frac{1}{2}$ and

$$\begin{aligned} p_2(t^*) &= t^{*4} - 2t^{*2} - 2t^* + 4 = t^* \left(t^* + \frac{1}{2} \right) - 2t^{*2} - 2t^* + 4 \\ &= -t^{*2} - \frac{3}{2}t^* + 4 \geq -\frac{36}{25} - \frac{9}{5} + 4 = \frac{19}{25} > 0, \end{aligned} \quad (133)$$

using that $-t^2 - \frac{3}{2}t + 4$ is decreasing for $t > 0$ and $t^* < 6/5$. The endpoint values are $p_2(3) = 1$ and $p_2(4) = 8$, both larger.

For (130), substitute $S_1^*(y) = y/2 + 4/y - 2$ and $f = y - 3$. The first difference is

$$\frac{y}{2} + \frac{4}{y} - 2 - \frac{1}{y} - \frac{y-3}{y-2} = \frac{y}{2} + \frac{3}{y} - 2 - \frac{y-3}{y-2}, \quad (134)$$

and multiplying by $2y(y-2) > 0$ gives $y^2(y-2) + 6(y-2) - 4y(y-2) - 2y(y-3) = y^3 - 8y^2 + 20y - 12 = p_1(y)$. The second difference is $\frac{y}{2} + \frac{2}{y} - 2 - \frac{y-3}{(y-2)^2}$, and multiplying by $2y(y-2)^2 > 0$ gives $y^2(y-2)^2 + 4(y-2)^2 - 4y(y-2)^2 - 2y(y-3) = (y-2)^4 - 2y(y-3) = p_2(y)$, since $y^2 + 4 - 4y = (y-2)^2$. $\square$

Both lemmas were re-derived for this paper in exact rational arithmetic: the identities (127) and (130) were checked as polynomial identities on rational grids, P and p_1, p_2 were checked positive pointwise together with the sign chains for their derivatives, and the endpoint comparisons were reduced to the integer inequalities $16 > 8$, $49 > 48$ and $81 - 21 = 60 > 36\sqrt{2}$.

C Certification protocol and coverage

C.1 The exact upper bound on a box

The certified runs of §11 rest on one function: an exact rational upper bound on $S(a, b, s)$ valid for every s in a box $[\ell, h] \subseteq [0, 1]^b$. It is built as the minimum of four independent bounds, each exact.

(i) *Atom-wise bound.* Each atom of the reduced spectrum is $\tau'_S = \prod_{j \in S} p_j \prod_{j \notin S} (1 - p_j)$ with $p_j = s_j / (1 + s_j)$ increasing in s_j , so over the box

$$\tau'_S \leq u_S = \prod_{j \in S} \frac{h_j}{1 + h_j} \prod_{j \notin S} \frac{1}{1 + \ell_j}. \quad (135)$$

If $x_S \leq u_S$ for every S then the i -th largest x is at most the i -th largest u , so sorting the 2^b numbers u_S decreasingly gives a valid bound on the sorted atoms, and $\sum_i w_i \min(u_i, 1/y)$ bounds the objective.

(ii) *Support counting.* $\sum_i w_i = \min(y, 2^b)$ and $\min(\tau'_i, 1/y) \leq 1/y$ give $S \leq \min(y, 2^b)/y$, which is Lemma 13.

(iii) *Largest-atom bound.* Every atom is at most $\tau'_1 = 1/Z_b$ and only $K(y) = \lceil y \rceil$ of them carry weight, so $S \leq K(y)/Z_b$, and over the box $Z_b \geq \prod_j (1 + \ell_j)$. The implementation stores this quantity twice, once under this name and once under a Cauchy–Schwarz label, so the minimum it takes is over five stored values of which two coincide and four are distinct.

(iv) *Normalisation-aware greedy bound.* The true reduced spectrum satisfies $0 \leq \tau'_i \leq u_i$, is decreasing, and sums to exactly 1. Maximising $\sum_i w_i \min(v_i, 1/y)$ over that set is a fractional knapsack problem with non-increasing weights w_i and a common cap $1/y$, so the greedy fill from the lowest index is optimal; the resulting value never exceeds (i) and is strictly smaller whenever $\sum_i \min(u_i, 1/y) > 1$, that is, whenever the atom-wise bounds of (i) cannot all be attained by a single normalised spectrum.

Taking the minimum is exact and rational. The branch and bound bisects the box in its widest coordinate, discards a box when its bound is at most the target $S^*(N)$, and records a box as a survivor when its width reaches the tolerance 2^{-e} . The ordering constraint $s_1 \geq \dots \geq s_b$ is imposed by tightening the box after each split. The certified bound for a sector is the maximum of the exact bounds of its survivors, and it is a valid upper bound on $\sup S$ over the whole closed box provided the search completed, that is, provided every box was either discarded or re-verified.

One soundness caveat is worth stating explicitly. Pruning uses floating-point interval arithmetic with a conservative inflation — the float bound is multiplied by $1 + 2^{-45}$ and 10^{-15} is added — for speed, and survivors are then converted to exact rational boxes, lower corners rounded down and upper corners rounded up to multiples of 2^{-40} , and re-verified with the exact bound above. The final number is therefore an exact rational, and its validity as an upper bound on $\sup S$ depends on the exact re-verification

of every survivor and on the float pruning never discarding a box whose true bound exceeds the target, which is what the inflation is for. Runs that exhaust the box budget are reported as incomplete and contribute no exact bound; all 264 branch-and-bound sectors behind Table 1 completed all five tolerance runs, 1320 completed searches in total, and this was re-checked from the committed records.

C.2 The tolerance ladder

Each residual sector is run at the five tolerances $2^{-6}, 2^{-8}, 2^{-10}, 2^{-12}, 2^{-14}$, and each completed run yields an exact rational bound B_e . The recorded data per sector and per tolerance are the bound, whether the run completed, the number of surviving boxes and the number of boxes explored. Three derived quantities are what §11 reports: the excess $\varepsilon_e = B_e - S^*(N)$, its monotonicity in e , and the stability of the ratio $\varepsilon_e/2^{-e}$ across the ladder. A worked example is the worst sector in the whole table, $(a, b) = (0, 5)$ at $N = 25$, whose excesses over the five tolerances are

$$1.1318 \times 10^{-2}, \quad 2.8106 \times 10^{-3}, \quad 7.0154 \times 10^{-4}, \quad 1.7532 \times 10^{-4}, \quad 4.3825 \times 10^{-5}, \quad (136)$$

with ratios $\varepsilon_e/2^{-e}$ equal to 0.7243, 0.7195, 0.7184, 0.7181, 0.7180: each halving of the tolerance halves the excess to four decimal places, which is the behaviour a boundary layer of thickness 2^{-e} should have. The source records label a sector “certified in the limit” when all five runs completed, the excesses are non-increasing, the ratios agree to within two percent and the finest excess is non-negative. That label describes the observed decay; it is not a proof that $\sup S = S^*(N)$ on the sector, for the reason given in §11.4 — no modulus of continuity is supplied — and this paper does not use it as one.

C.3 Coverage per rank

At rank N with $n = \lfloor \log_2 N \rfloor$ the runs cover the rectangle $0 \leq a \leq n$, $0 \leq b \leq n+3$, which contains $(n+1)(n+4)$ sectors: 10 at $n = 1$, 18 at $n = 2$, 28 at $n = 3$ and 40 at $n = 4$. Of these, the analytic eliminations of §6 and §7 discharge 8, 12, 17 and 23 respectively — the counts of Table 1 — and the remainder need branch and bound: 2, 6, 11 and 17. The uncovered family is $\mathcal{U}(N) = \{a \leq n-1, b \geq n+4\}$ as run, or $\mathcal{U}'(N) = \{a \leq n-2, b \geq n+4\}$ after Remark 40. Intersecting $\mathcal{U}'$ with the analytically open family (71), which at $N = 25$ is $(0, b)$ for $5 \leq b \leq 31$, $(1, b)$ for $4 \leq b \leq 15$ and $(2, b)$ for $3 \leq b \leq 7$, the sectors that are both untouched by the runs and unsettled by every theorem in this paper are $(0, b)$ for $8 \leq b \leq 31$ and $(1, b)$ for $8 \leq b \leq 15$: boxes of dimension 8 to 31, that is 2^8 to 2^{31} atoms. The largest single run explored 971975 boxes against a budget of 1.2×10^7 , and the largest number

of surviving boxes re-verified exactly in one run was 16055; no run exhausted its box budget or its wall-clock budget.

The false shortcut of Remark 22 was counted per rank on the $1/32$ lattice: 62 failures in the open box $0 < s_2 < s_1 < 1$, distributed as 18, 6, 18, 20 over $N \in \{5, 9, 10, 11\}$, and 90 in the closed box $0 < s_2 \leq s_1 \leq 1$, the additional 28 all on the face $s_1 = 1$ and none on the diagonal $s_2 = s_1$. At every failing point the exact objective still satisfies $S \leq S^*(N)$, which is what makes the shortcut false rather than the theorem wrong. Ranks 6 and 7, which the source includes in its rank list, contribute no failing point at this resolution.

C.4 What would extend the certificates

Three extensions are natural and none is available here. Refining the tolerance ladder below 2^{-14} reduces ε_e essentially proportionally, as (136) shows, and is cheap: summed over all 264 residual sectors the box counts are 1.58×10^6 , 3.53×10^6 , 5.22×10^6 , 6.68×10^6 and 8.10×10^6 for $e = 6, 8, 10, 12, 14$, a factor of only 1.2 to 2.2 per two extra bits. It therefore buys resolution and never exactness. Extending the covered rectangle from $b \leq n + 3$ to larger b is the genuinely expensive direction, limited by the exponential growth of the atom count, 2^b , and of the box dimension; the runs already take 745.5 seconds at $N = 25$ against 0.11 seconds at $N = 4$. And a modulus of continuity for the bound of §C.1 would replace the ladder by a single argument, converting the observed linear decay into $\sup S \leq S^*(N)$ on every covered sector at every certified rank, which is Conjecture 1 restricted to the covered family.

D Exploratory searches and their artefacts

Two of the questions in this paper were attacked numerically before they were attacked analytically, and in both cases the numerical answer had to be interpreted carefully. This appendix records the protocols, the outputs and the failure modes, so that the reader can tell which statements in the body are theorems and which are searches. Everything here is quoted from the committed reports and, where indicated, re-derived in exact rational arithmetic for this paper.

D.1 The search for interior local maxima

Protocol. Multistart local maximisation of the exact block objective $S(a, b, s)$ over $s \in (0, 1)^b$, 60 restarts per sector, in a logistic parametrisation so that the box constraint is unconstrained, with the objective evaluated through the block reduction of Theorem 11 and the 2^b atom arrays built by direct product enumeration [4, 8]. Sectors probed: every residual sector with $n + 1 \leq a + b \leq n + 4$ at the ranks $N \in \{7, 10, 11, 13, 14, 15\}$, that is 85

sectors. A maximiser is recorded as *interior* when every coordinate stays strictly inside $(0, 1)$ and as *boundary* otherwise.

Outcome. Interior local maxima above $S^*(N)$: none, in any of the 85 sectors. Every maximiser landed on the boundary, and in every sector it migrated to the same face point, $s_1 = \dots = s_{b-1} = 1$ with $s_b = x - 1$, at which the value is exactly $S^*(N)$. The largest excess over $S^*(N)$ across all runs was -1.388×10^{-14} , that is, below the target; the reported suprema agree with $S^*(N)$ to between 10^{-13} and 10^{-12} , and the discrepancy grows with the number of modes, from about 10^{-14} at $m = 4$ to about 4×10^{-12} at $m = 8$. That pattern is what accumulated optimiser error looks like, not what a near-miss with a genuine interior maximum would look like.

Scope. The report itself carries the caveat that this is not a certificate: a local search can miss a basin, and the objective is piecewise affine with clipping walls that can make basins thin. What the search adds beyond the theorems is therefore limited to the sectors that no theorem reaches — by Theorem 20, Theorem 27 and Theorem 31 the interior-maximum question is already settled for $b \leq 2$, for every b at $a = n - 1$ and for $b \geq \lceil y \rceil$ everywhere — and in those sectors it says that the natural candidate maximiser is the predicted face point.

D.2 The search for a doubling counterexample

Protocol. Probe spectra ρ of length $L \leq 64$ in six families: flat spectra of rank 2 to 16, geometric spectra with ratio $q \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$, peaked spectra with head weight 0.5 to 0.999, Dirichlet-type spectra, log-normal spectra, and free spectra distorted by $\varepsilon \in \{0.05, 0.2, 0.5\}$. On each side — ρ and $\frac{1}{2}(\rho \oplus \rho)$ — the distance was minimised by multistart Nelder–Mead over free spectra with at most $\lceil \log_2 L \rceil + 3$ modes, 20 random restarts per mode count, in log-odds coordinates [8]. The reported quantity is $\text{gap} = D_{\mathcal{F}}(\frac{1}{2}(\rho \oplus \rho)) - D_{\mathcal{F}}(\rho)$, which Conjecture 2 predicts to be zero and the classical direction predicts to be non-negative.

First outcome. Over 89 probes the most negative gap was -6.836×10^{-3} , the largest was $+8.715 \times 10^{-4}$, the mean absolute gap was 1.107×10^{-4} , and five probes exceeded a slack of 10^{-9} on the negative side. All five were on peaked or Dirichlet probes with a small distance, where the absolute size of the gap is largest relative to the quantity being measured.

Recheck with structure-aware starts. The same 32 probes were re-run with 40 random restarts per side *plus* structured starts: all biases equal, all biases equal with k zero modes for each k , and geometric ladders. The most negative gap improved from -6.836×10^{-3} to -4.119×10^{-5} , about 10^{-3} of the corresponding distance, and the two probes that had produced the largest artefacts closed to $+2.605 \times 10^{-10}$ and $+3.873 \times 10^{-4}$. The five surviving negative gaps are all at the 10^{-5} to 10^{-7} level.

Interpretation. Both columns of such a table are upper bounds on the

true distance, since each comes from a minimiser that may not have converged. A small negative computed gap is therefore the signature of the doubled side converging slightly worse — it searches a wider space, one mode larger, and a wider space is harder to converge — and not of a violation. The structural reason the naive search fails is explicit and is quantified in §D.3 below: for peaked targets the near-optimal free spectrum has all its biases equal, which is a measure-zero feature that random restarts find only by chance. The reports’ own conclusion is that no entry in the table is a proof, and that a surviving negative gap would have to be converted into an exact interval certificate before being claimed as a counterexample. This paper claims none.

D.3 The symmetric near-optimum, exactly

For the peaked target $\rho = (0.99, 0.002, 0.002, 0.002, 0.002, 0.002)$, padded with zeros, the free spectrum with five equal biases $p_j = p$ and $(1-p)^5 \leq 0.99$ has largest atom $(1-p)^5$, which matches ρ_1 , and 31 further atoms whose total mass is $1 - (1-p)^5 \geq 0.01$, each of them below 0.002. Its distance to ρ is therefore an exact rational number, computed for this paper by bisecting for p in rational arithmetic until $(1-p)^5 \leq 0.99$ with an enclosure of width below 2^{-40} :

$$p = \frac{1157561055916123}{576460752303423488} = 0.002008048\dots, \quad (137)$$

$$D_{\mathcal{F}}(\rho) \leq \frac{1}{2}\|\rho - \sigma\|_1 = 4.0161 \times 10^{-5}, \quad (138)$$

where the bound uses Lemma 6 in the form $\frac{1}{2}\|\rho - \sigma\|_1 = 1 - \sum_a \min(\rho_a, \sigma_a)$ and the atoms of σ are the 32 products of the five two-level factors. The floating-point run of the source reports $4.016062521706 \times 10^{-5}$ for the same probe, agreeing with the exact value in its first five significant digits, and it is two orders of magnitude below what plain multistart found on the ρ side of the doubling test. The bound (137) is an upper bound on $D_{\mathcal{F}}(\rho)$, not a certified value: no lower bound for this ρ is available from any method in this paper.

D.4 The zero-mode decomposition on finite families

Theorem 46 was verified for this paper on two nested finite families of ratio lists drawn from the grid $\{1, \frac{15}{16}, \frac{7}{8}, \frac{3}{4}, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}\}$: lists of length at most four at rank N , giving 330 lists, and of length at most five at rank $2N$, giving 792. The nesting is what makes the identity exact: the rank- $2N$ family contains the zero-mode extension of every rank- N list, and the rank- N family contains the zero-mode deletion of every rank- $2N$ list that has a zero mode. Table 2 records the three suprema and the identity; the last column is the supremum over the zero-mode-free lists with $b \geq n + 2$, which is the family

Table 2: The zero-mode decomposition on finite families, in exact rational arithmetic. The rank- N family is the non-increasing lists of at most four mode ratios drawn from the grid $\{1, \frac{15}{16}, \frac{7}{8}, \frac{3}{4}, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}\}$, 330 lists, and the rank- $2N$ family the lists of at most five from the same grid, 792; the nesting makes both directions of the identity rigorous on the family searched. For each rank N with $n = \lfloor \log_2 N \rfloor$: the supremum $S_{\max}(N)$, the supremum $S_{\max}(2N)$, the zero-mode-free supremum $\Sigma_0(2N)$, whether the identity of Theorem 46 holds exactly, and the supremum over zero-mode-free lists with $b \geq n + 2$, the family that Corollary 47 leaves open. All entries are exact rationals. These are suprema over finite families, not over $\mathcal{F}$, and are reported as a check on the theorem rather than as evidence for the conjecture.

N	n	$S_{\max}(N)$	$S_{\max}(2N)$	$\Sigma_0(2N)$	identity	$\sup\{a = 0, b \geq n + 2\}$
2	1	1	1	3781/3844	yes	8/9
3	1	5/6	5/6	1591/1922	yes	1591/1922
4	2	1	1	58141/59582	yes	8/9
5	2	17/20	17/20	124218/148955	yes	124218/148955
6	2	5/6	5/6	73757/89373	yes	73757/89373
7	2	25/28	25/28	2695009/3027150	yes	2695009/3027150
8	3	1	1	14413205/14776336	yes	39377923/44329008
9	3	65/72	65/72	4921355/5541126	yes	14667157/16623378
10	3	17/20	17/20	15412629/18470420	yes	15412629/18470420

that Corollary 47 identifies as the whole obstruction to Conjecture 2 on flat spectra. At every rank it is strictly below $S_{\max}(N)$, and the difference is smallest at $N = 2$, where $S_{\max}(2) = 1$ because rank 2 is a power of two.

The same computation was repeated on four family conventions, to locate exactly where Table 2 differs from the source's larger table. That table uses the eight-point grid $\{1, \frac{15}{16}, \frac{7}{8}, \frac{3}{4}, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}$ with lists of length at most five for all three suprema, 1287 lists at each rank, over ranks 2 to 14. Recomputed here in exact arithmetic on that family, all four of its columns are reproduced exactly at every rank $2 \leq N \leq 10$ where the two computations overlap, and the values of $S_{\max}(N)$ and $S_{\max}(2N)$ agree with Table 2 at every one of those ranks. The two families differ only by the grid point $\frac{1}{16}$: the nesting, length at most four at rank N against at most five at rank $2N$, changes none of the reported suprema. With $\frac{1}{16}$ present the probe over $b \geq n + 2$ is larger at $N = 2, 4, 8$ and 9 — 61457/65348 against 8/9, 3795407/4051576 against 8/9, 23522885/251197712 against 39377923/44329008, and 251771995/282597426 against 14667157/16623378 — and $\Sigma_0(2N)$ is larger at $N = 9$ alone, 251771995/282597426 against 4921355/5541126. Every other entry of the two tables agrees. The identity of Theorem 46 holds at every rank 2 to 10 on all four conventions, including the two that are not nested, so on these grids and at these ranks the non-nested convention loses nothing; only the nested one makes both directions of the identity a statement about the family actually searched, which is why

Table 2 uses it.

The same runs verified the two supporting statements exactly: the pair inequality of Proposition 43 at 7548 rational points covering all four branch loci, and the zero-mode exactness of Corollary 42 at 1638 instances over ranks 2 to 14. The coarse-graining chain of Corollary 45 was checked at 4200 deletions, of which 1890 violate the truncation hypothesis $2^{m-1} \leq N$ and 653 of those violate the conclusion — the hypothesis is not merely sufficient but needed. The sharpness witness was computed in exact arithmetic: for $r = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ the eight atoms are $(8, 4, 4, 4, 2, 2, 2, 1)/27$, the value at rank 4 is $S(\bar{\rho}(4), \sigma) = 25/36$, all three single-mode deletions give the rank-2 value $2/3$, and $25/36 > 2/3$ while $2^{m-1} = 4 > N = 2$, so the truncated step loses value precisely where its hypothesis fails. Both values stay below the conjectured bound at their own ranks, since $S^*(4) = S^*(2) = 1$: the example bounds the method without threatening the conjecture.

References

- [1] Sergey Bravyi. Lagrangian representation for fermionic linear optics. *Quantum Information & Computation*, 5(3):216–224, 2005.
- [2] GeWu Journal Auto Author. Source research repository for this paper. <https://git.gewu-lab.ai/scientific-author-73/interaction-distance-flat-spectrum-conjectures>, 2026. Solution repository at commit ea9dcd38fdc6b67d1a7a462e498a33da4e39bd98.
- [3] GeWu Open Problem Index. Phys-orp-3785: Flat-spectrum, zero-mode, and maximal-value conjectures for the fermionic interaction distance underlying free-fermion descriptions of $\mathbb{Z}_n$ topological states. Problem README, GeWu Open Problem Index, project 19347. URL <https://git.gewu-lab.ai/peter0627ustc/phys-orp-3785-flat-spectrum-zero-mode-and-maximal-value-conjectures-for-the-fermionic/-/blob/main/README.md>. Accessed 7 September 2026.
- [4] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585:357–362, 2020. doi: 10.1038/s41586-020-2649-2.
- [5] Hui Li and F. D. M. Haldane. Entanglement spectrum as a generalization of entanglement entropy: Identification of topological order in non-abelian fractional quantum hall states. *Physical Review Letters*, 101(1):010504, 2008. doi: 10.1103/PhysRevLett.101.010504.
- [6] Konstantinos Meichanetzidis, Christopher J. Turner, Zohreh Farjami, Zlatko Papić, and Jiannis K. Pachos. Free-fermion descriptions of parafermion chains and string-net models. *Physical Review B*, 97(12):125104, 2018. doi: 10.1103/PhysRevB.97.125104.
- [7] Ingo Peschel. Calculation of reduced density matrices from correlation functions. *Journal of Physics A: Mathematical and General*, 36(14):L205–L208, 2003. doi: 10.1088/0305-4470/36/14/101.
- [8] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J.

Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.