The price of unreachable perfection: dimension laws at every regular scale

Seth Douglas Nidhal Mghirbi

Version 1.0.0, October 2026

Abstract

How expensive is an optimum that a finite quantum system can approach but never attain? For a fixed finite rational nonlocal game of quantum value one, write $A_G(n) = \log_2 D_G(2^{-n})$, where D_G is the least local dimension at loss 2^{-n}. We prescribe $\log_2 A_G(n)$ at every logarithmic-space, nondecreasing unit-increment scale F with $F(n)/\log n \to \infty$, up to an additive $O(\log n + \log^2(F(n) + 2))$ error, and match the one-way distillable entanglement of the original mixed strategy state. At every asymptotically flat such scale, the positive finite lower and upper normalized dimension endpoints are exactly the ordered pairs with lower endpoint in Σ_2^0 and upper endpoint in Π_2^0; existing positive limits are exactly Δ_2^0. A single uniform spectral-tail construction supplies the density-free profile law behind these results. It also classifies stretched-scale pairs in the unit square and gives quasi-polynomial dimension windows of width three in the exponent. A computable unit-increment target outside every rational game's profile shows that computability alone is insufficient. We locate the boundaries through a quantitative locality cap, an honest-linear compression cap, and an explicitly conditional route to finite exponents above one. A separate exact-marginal correlation theorem identifies a Farey law and its Diophantine endpoint curve. The dimension arithmetic necessity applies to all rational games; the operational realization statements concern the constructed games.

The authors developed this work with AI models from more than one developer and take responsibility for every claim. How the work was done, and what the machine checks cover, is described at the end of the paper.

1 A theory of precision scales

A finite nonlocal game can have quantum value one while every finite-dimensional strategy loses with positive probability. The absence of a perfect strategy then has a quantitative price. How large must a system be to lose at most ε? Which growth functions can occur when the game is fixed once and for all? We study this whole-profile question together with an operational version: how many entangled bits can be distilled, with classical communication from Alice to Bob, from the very mixed state used by a near-perfect strategy?

For a finite rational game G of quantum value one let $D_G(\varepsilon)$ be the least d for which a strategy on $\mathbb{C}^d \otimes \mathbb{C}^d$, with arbitrary state and POVMs, has loss at most ε. Local spaces of smaller sizes may be padded. Write

$$A_G(n) = \log_2 D_G(2^{-n}), \qquad d_G^{\rightarrow}(\varepsilon) = \inf_{\text{loss}\le\varepsilon} E_D^{A\to B}(\rho_{AB}). \tag{1.1}$$

The infimum in the second expression ranges over all finite-dimensional strategies, with no dimension restriction and no access to a purifying environment. Entanglement is measured in bits per copy. We also write p_G for the infimum of pure-state entanglement entropy over pure original strategies and f_G for the infimum of single-copy entanglement of formation over all original strategies. Our lower estimates apply before taking these infima. The maximally entangled upper strategies simultaneously bound all four demands. We abbreviate $d_G^{\rightarrow}(2^{-n})$ by $d_G(n)$.

All logarithms are base two unless ln is written. When an iterated logarithm is used for an arbitrary game, we regularize it at a fixed positive value, for example by replacing $\log_2 A_G$ with $\log_2 \max\{2, A_G\}$. This changes none of the positive endpoints below. If a game of value one has no finite perfect strategy, compactness at each fixed dimension implies $D_G(2^{-n}) \to \infty$, so the unregularized logarithms are eventually defined.

Definition 1.1. A regular scale is a function $F\colon \mathbb{N} \to \mathbb{N}$, computable in deterministic space $O(\log(n+2))$, such that F is nondecreasing, its adjacent increments belong to $\{0,1\}$, and $F(n)/\log(n+2) \to \infty$. It is asymptotically flat if $F(n+1) - F(n)$ has average slope tending uniformly to zero on long intervals: for every $\eta > 0$ there is C_η with

$$F(v) - F(u) \le \eta(v-u) + C_\eta \quad (v \ge u \ge 0).$$

This explicit interval condition is the meaning of flatness throughout the paper. It implies $F(n) = o(n)$. Smooth sublinear targets, with a finite initial prefix change and integer rounding, provide examples.

We use the real-number hierarchy convention that a finite lower limit of a computable rational sequence is a Σ_2^0 real, a finite upper limit is a Π_2^0 real, and a limit is a Δ_2^0 real; see Zheng and Weihrauch [36]. The hierarchy here classifies cuts of real numbers.

Theorem 1.2 (Theory of scales)**.** *The following statements hold for fixed finite rational games.*

1. *For every regular scale F, a rational Boolean constraint-system game G of quantum value one, with no finite-dimensional perfect strategy, satisfies*

$$\log_2 A_G(n) = F(n) + O\big(\log(n+2) + \log^2(F(n)+2)\big),$$

   and the same estimate holds for $\log_2 d_G(n)$, $\log_2 p_G(2^{-n})$ and $\log_2 f_G(2^{-n})$.

2. *At every asymptotically flat regular scale F, the positive finite dimension endpoint pairs of $\log_2 A_G(n)/F(n)$ are exactly*

$$\{(L,U) : 0 < L \le U < \infty,\ L \in \Sigma_2^0,\ U \in \Pi_2^0\}. \tag{1.2}$$

Every pair is realized by a BCS game with the same operational endpoint pair for $\log_2 d_G(n)/F(n)$. *In particular the positive existing dimension limits are exactly the positive* Δ_2^0 *reals.*

3. *The endpoint pairs in* $[0,1]^2$ *of*

$$\frac{\log_2 \log_2 A_G(n)}{\log_2 n}$$

are exactly the ordered (Σ_2^0, Π_2^0) *pairs in that square. They have matched original-state usable realizations.*

4. *There is a computable nondecreasing unit-increment* F *with*

$$\lfloor n^{1/3} \rfloor \le F(n) \le \lfloor n^{2/3} \rfloor + 1$$

for all sufficiently large n, *for which no finite rational game satisfies* $\log_2 A_G(n) = F(n)(1 + o(1))$.

The arithmetic necessity in this theorem is about dimension. The matched usable laws are realization statements for the constructed games. Our proof does not compute distillable entanglement of arbitrary mixed states. The main input beyond the compiler, controller and extraction theorems of [13, 12] is a single source model with a tail estimate at every spectral threshold. It yields the pointwise profile

$$\log_2 A_G(n),\ \log_2 d_G(n) = 2k(\lfloor n/2 \rfloor) + O\big(\log(n+2) + \log^2(k(\lfloor n/2 \rfloor) + 2)\big),$$

where k is the cumulative one-count of a programmed pulse word. No positive pulse-density assumption is needed. Sections 3–5 prove the theorem and the quasi-polynomial windows.

Table 1 separates dimension from its logarithm. Put $n = \log_2(1/\varepsilon)$. The first row is a Bell-functional deficit problem, and the exact-marginal example in Section 7 has a third contract. Neither is silently converted to a Boolean value-one game.

Dimension scale	What is established	Boundary of the statement
$D = \Theta(n)$	I_{3322} Bell deficit [11]	Non-attainment also proved by Pauwels and Coladangelo [29, 8]; distinct contract
$D = 2^{\alpha n + o(n)}$	Positive coefficients and endpoint pairs classified in [13, Theorem Q4.2]	Matched usable realizations
$D = 2^{n^{c+o(1)}}$	Quasi-polynomial windows of bounded width	Exact exponents in these windows unresolved
$D = 2^{2^{F(n)(1+o(1))}}$	Every regular F; arithmetic law for flat F	Computability alone is insufficient
$D = 2^{2^{n^{c+o(1)}}}$	Ordered stretched pairs in $[0,1]^2$	Statement is restricted to this square
$D = 2^{2^{\beta n + o(n)}}$	$0 < \beta \le 1$ classified in [13, Theorems P1.1, Q6.1]	A finite existing $\beta > 1$ is unresolved
General upper power exponent	$+\infty$ occurs in [13, Theorem P6.11]	No usable-entanglement conclusion from this argument

Table 1: Dimension prices, with precision $\varepsilon = 2^{-n}$. Classification refers to dimension; operational matching refers to the realized games. Rows 1, 2, 6 and 7 cite earlier results; rows 3–5 are proved here.

1.1 Relation to earlier work

Slofstra established nonclosure through finite nonlocal games [33]. Slofstra and Vidick gave quantitative dimension and Schmidt-rank bounds [34]; their polynomial lower bound is a dimension bound, not a lower bound on distillable ebits. Slofstra's later two-player construction has a bound $\exp(c\varepsilon^{-a})$ for some $0 < a < 1$ and commuting-operator value one; quantum value one was left open [32, Corollary 1.4]. Ji, Leung and Vidick obtained a three-player coherent-state embezzlement game [18]. Coladangelo then obtained a two-player near-optimal dimension bound $2^{\Omega(\varepsilon^{-1/8})}$ [7]. Binary constraint-system and linear-system games have the operator formulations of Cleve–Mittal and Cleve–Liu–Slofstra [6, 5]. The undecidability and recursive-compression background includes $\mathrm{MIP}^* = \mathrm{RE}$ [19]. Mousavi, Nezhadi and Yuen classified the value-one decision problem at the second universal level of the arithmetical hierarchy [26]. We use their hardness for the upper-exponent obstruction and their compression construction for the cap in Section 6; the regular-scale realizations are direct programmed constructions.

The two base-scale classifications are results of the companion paper *The uncomputable price of precision*, version 4.0.0: Theorem 16.6 and Corollary 16.7 give matching polynomial dimension laws and their positive coefficients; Theorem Q4.2 gives their endpoint pairs; Theorems P1.1 and Q6.1 give the power scale and its pairs in $(0, 1]$ [13]. All theorem and section numbers citing that paper refer to version 4.0.0, identified by its concept DOI in the bibliography. The locality and unbounded-upper-exponent results below are also restatements of that paper, with their quantitative hypotheses; we do not reproduce their proofs. The new contributions here are the uniform-tail source, its density-free profile and regular-scale consequences, the stretched and window estimates, the diagonal in the iterated-logarithm normalization, the stopping cap and conditional interface, and the Farey comparison and endpoint curve.

Bartholdi–Erschler realize intermediate word-growth functions under their dilation conditions [1]; Kassabov–Pak construct oscillating intermediate growth between widely separated prescribed bounds [21]. These are close precedents for prescribed speeds and oscillations, with different invariants and equivalence relations. Hochman–Meyerovitch classify entropies of multidimensional shifts of finite type [17]. As stated, Meyerovitch's Theorem 5.1 classifies lower, upper and existing entropy dimensions of $\mathbb{Z}^d$ shifts of finite type, for $d \geq 2$, as the Σ_2, Π_3 and Δ_2 reals in $[0, d]$, respectively [24]. The upper class there is Π_3, whereas finite-dimensional feasibility here is decidable and gives a Π_2^0 upper cut; global pattern extension in that setting is a different computational problem. Here we also classify the joint ordered endpoint pair. This comparison uses his three stated classes without identifying the invariants. Sapir–Birget–Rips show how computational restrictions enter realization of isoperimetric functions [31]. Brady–Bridson prove density, with closure $\{1\} \cup [2, \infty)$, for isoperimetric exponents [2]; a closure theorem is different from realizing every real exponent. Gromov's polynomial-growth theorem and Grigorchuk's gap conjecture [16, 15] give further contrasts. No group-growth theorem is transferred to the present game invariant.

2 Effective dimension and the imported interfaces

Proposition 2.1 (Arithmetic necessity at every effective scale)**.** *For rational finite games G of quantum value one, there is one uniform partial algorithm taking G and a positive rational tolerance as input and returning its least feasible local dimension whenever its search terminates. It terminates on every value-one game. If ϕ is an effectively evaluable real function defined at the relevant $A_G(n)$ and $s(n)$ is an effectively evaluable, eventually positive sequence tending to infinity, every finite lower endpoint of $\phi(A_G(n))/s(n)$ belongs to Σ_2^0 and every*

finite upper endpoint belongs to Π_2^0*. An existing finite limit belongs to* Δ_2^0*.*

Proof. At fixed d, write a state and each POVM using real and imaginary matrix entries. Hermiticity, positivity, normalization and loss at most a positive rational threshold form a first-order sentence over the real field with rational polynomial coefficients. Positivity may be expressed by principal minors, or by a factorization BB^*. Real quantifier elimination decides this sentence, including equality at the threshold. Test $d = 1, 2, \ldots$. The value-one promise makes the search finite for every positive tolerance. It supplies neither a recognizer for the promise nor a total bound on all input descriptions.

Effective evaluation gives computable rationals within $1/n$ of the normalized profile. Their lower and upper limits are unchanged. For a computable rational sequence, its finite lower limit is left computably enumerable relative to the halting set, and its finite upper limit is right computably enumerable relative to that set. One can see this by taking tail infima and suprema, uniformly computable as extended reals relative to the halting set, then increasing or decreasing the tail index. These are the asserted cut classes. A finite limit is a limit of a computable rational sequence. Monotonicity of ϕ is unnecessary; effective evaluation, its domain, and positive denominators are necessary. □

We now state precisely the interfaces used from [13, 12]. This records dependencies rather than new assumptions about arbitrary games. A weighted model has $\xi = \sqrt{\rho}$ and norm $\|R\|_\rho = \|R\xi\|_2$. Domains (unitaries, projections and bounded self-adjoint contractions) are exact; each tested polynomial residual and each required square-root-state commutator has weighted norm at most its precision.

Compiler. The polynomial compiler [13, Theorem 3.1] sends a finite bounded rational $*$-polynomial presentation to a fixed finite rational BCS game. The game-to-model transfer of [12, Theorem 4.9 and Proposition 6.1], together with that compiler, decodes loss ε on the original state at precision $C\sqrt{\varepsilon}$, including actual mirrors, inverse mirrors and square-root-state commutators. A flat precision-δ model of dimension D gives a maximally entangled strategy in dimension at most cD and loss at most $C\delta^2$. The constants depend on the fixed presentation.

Clock and controller. The unguarded two-clock controller and normalized stopped-product construction of [13, Sections P3–P5] simulate a fixed deterministic INC/DECJZ/NOP program with two output controls. If t_J is the physical time immediately after pulse J and $k(J)$ counts its one pulses, then $k(0) = 0$ and $k(J+1) - k(J) \in \{0, 1\}$. At horizon H the localized errors are $C(H+1)^b\delta$. Initialized-corner amplitude is charged by a polynomial in H. Exact scalar upper histories allow cyclic reset before a programmed stop. No tensor factor proportional to t_J is required.

For clarity the complete-column estimate used below is

$$\|T\varphi\| \le 2^{-k(J)} + 2^{J-k(J)+1}\big(\|S\varphi\| + \sqrt{H}\, r_H\big) + C\sqrt{H+1}\, r_H, \qquad H = t_J, \tag{2.1}$$

where $r_H = C(H+1)^b\delta$, $S = 0$ is tested, and charge back to $\|T\|_\rho$ costs another polynomial in H. The proof uses ordered measurement columns: $(PK)^*(PK) + ((I-P)K)^*((I-P)K) = K^*K$ exactly. Normalizing the two counters divides by their actual products 2^k and 2^J. The largest stopped ratio is $2^{J-k(J)+1}$ because $J - k(J)$ is nondecreasing. This explains why no positive density enters (2.1). This is inequality (P5.3) in the proof of [13, Proposition P5.1]: Proposition P3.1 supplies localization, Proposition P4.2 supplies counter stability, the normalized identities (P5.2) give products $M_h = 2^{k(h)}$, $N_h = 2^{j(h)}$, and Lemma P3.2 sums the stopped columns. That derivation uses polynomial delay but not pulse density. Lemma Q2.2 there is the analogous estimate for the separate carry checker.

Original-state extraction. The proof of [12, Theorem 12.4], via its signed regular extraction Theorem 6.4 and displayed budget (12.4), gives the following source-level estimate, also recorded in [13, Section P1]:

$$\log_2 E_D^{A\to B}(\rho_{AB}) \geq 2\log_2(1/e) - C\big(\log_2\log_2(1/e) + C\big)^2 - C. \tag{2.2}$$

Here the common tolerance e includes all source residuals, mirrors, inverse mirrors and square-root-state commutators, and the marked central sector has fixed visibility. Its explicit budget is $2^q\Lambda_q^4 e^2 \leq c$ with $\Lambda_q \leq 2^{(\log_2 q+23)^2}$, giving rate at least $c'2^q$. The Alice filter probability is included in that rate. Thus (2.2) is a statement about the original mixed state. These are the exact quantitative inputs needed below; replacing extraction by an entropy bound would not suffice.

3 One source model with every spectral tail

The marked affine lamp presentation in [13, Section P2] has six unitary letters p, q, t, d, x, z, 27 certificate words of [25, Definition 7.1] and the marker $J = (xz)^2$. Write $R_w = (w - I)/2$ and $R_J = (J + I)/2$. We retain that fixed presentation and change only the finite list used in its Clifford construction. The certificate separates 23 host words from four lamp words. The host verification is independent of the list; the four lamp residuals are controlled by its character sums.

Lemma 3.1 (Uniform tail)**.** *For every integer $N \geq 1$ the marked presentation has a flat unitary assignment with $x^2 = z^2 = I$, $J = -I$, and*

$$\log_2 D_N \leq 3N + (4M+4)^2 + 2\log_2 M, \qquad M = \max\{48, \lceil N^{1/4}\rceil\}. \tag{3.1}$$

Every host residual has operator norm at most $C_h N^{-1/2}$, where $C_h = 24\pi$. Each lamp residual satisfies, simultaneously for every $\kappa > 0$,

$$\operatorname{tr}(1_{(\kappa,1]}(|R_w|)) \leq \sqrt{2}\exp(-N\kappa^2/16). \tag{3.2}$$

The same assignment, for every real $\nu > 0$, satisfies

$$\operatorname{tr}(|R_w|^{2\nu}) \leq \sqrt{2}\left(\frac{16\nu}{eN}\right)^{\nu} + \left(\frac{C_h^2}{N}\right)^{\nu}. \tag{3.3}$$

For nondecreasing $f : [0,1] \to [0,\infty)$ and $C_h/\sqrt{N} \leq \kappa \leq 1$, every residual obeys

$$\operatorname{tr}(f(|R_w|)^2) \leq f(\kappa)^2 + \sqrt{2}f(1)^2\exp(-N\kappa^2/16). \tag{3.4}$$

The assignments can be selected effectively.

Proof. Put $r = 4M+4$. For an arbitrary list $\zeta_1, \ldots, \zeta_N \in \mathbb{F}_2^r$, with repetitions and zero entries allowed, define

$$v_u = N^{-1/2}\big((-1)^{\langle\zeta_j,u\rangle}\big)_{j=1}^N, \qquad D_u = \operatorname{diag}\big((-1)^{\langle\zeta_j,u\rangle}\big).$$

Every v_u is a real unit vector and $D_u v_w = v_{u+w}$. On the full exterior space of $\mathbb{C}^{3N}$, second quantization gives a genuine representation

$$U_u = \Lambda(D_u \otimes I_3), \qquad U_u U_w = U_{u+w}.$$

Use the overlapping-qubit Majorana construction of Chao–Reichardt–Sutherland–Vidick [4], with three copies A_u, B_u, C_u associated with v_u, and set $X_u = iA_uB_u$, $Z_u = iA_uC_u$. The canonical anticommutation relations give reflections, exact onsite anticommutation, $(X_uZ_u)^2 = -I$, and affine covariance under U_u. These identities use only unit norm and $D_uv_w = v_{u+w}$. They require no independence or distinctness of the list. Using the full exterior space also avoids any projective lifting issue on an irreducible Clifford module.

Induce over $\mathrm{GL}(r,2)$ and use the clock and holonomy correction of [13, Section P2]. Their certificate calculation gives zero residual on host rows 5, 10–22 and 27; the other host rows have norm at most $3h_M$, where $h_M \le 8\pi M^{-2}$. Therefore every host row is bounded by $24\pi N^{-1/2}$. The only dependence on the list is in the four lamp words, involving t, x, z. On a block indexed by $g \in \mathrm{GL}(r,2)$ their two labels are 0 and $u = g^{-1}e_1$. Put

$$c(u) = \langle v_0, v_u \rangle = N^{-1} \sum_{j=1}^{N} (-1)^{\langle \zeta_j, u \rangle}.$$

For the two mixed words, the additive commutator is $2c(u)B_0C_u$; multiplication by a unitary converts it to the normalized group residual. Its absolute value is $|c(u)|I$. For the two same-type words, the CAR identity gives residual norm $2|c(u)|\sqrt{1-c(u)^2} \le 2|c(u)|$. Clock coordinates repeat these blocks with equal multiplicity. Transitivity makes $g^{-1}e_1$ exactly uniform over the nonzero vectors. Thus each tail is bounded by $\Pr_u(2|c(u)| > \kappa)$, with $u \neq 0$ uniform.

Choose the list independently and uniformly. For fixed nonzero u, $S = Nc(u)$ is a sum of N independent Rademacher signs. With Z standard Gaussian,

$$\mathbb{E} \exp(S^2/(4N)) = \mathbb{E}_Z \cosh(Z/\sqrt{2N})^N < \mathbb{E}_Z e^{Z^2/4} = \sqrt{2}.$$

The strict inequality follows from $\cosh y < e^{y^2/2}$ for $y \neq 0$. Average over nonzero u and then over lists. Some list has

$$\Phi = \frac{1}{2^r - 1} \sum_{u \neq 0} e^{Nc(u)^2/4} < \sqrt{2}.$$

Fix this one list. Markov's inequality gives (3.2) for every threshold at once. This is an average over affine blocks, not a maximum over labels. Independence between different u is unused. Effectiveness follows by dovetailing certified interval comparisons for the finitely many lists until a strict comparison $\Phi < \sqrt{2}$ is certified. We do not ask an algorithm to decide equality of computable reals.

The full space has dimension $M^2|\mathrm{GL}(r,2)|2^{3N}$. Since $\log_2 |\mathrm{GL}(r,2)| \le r^2$, this proves (3.1), namely logarithmic dimension $3N + O(\sqrt{N})$. For lamp words, apply $y^\nu \le (\nu/e)^\nu e^y$ with $y = Nc(u)^2/4$. This gives the first term in (3.3); the host estimate gives the second. The list is unchanged as ν varies. Finally split the spectral measure at κ. The host tail vanishes above $C_h/\sqrt{N}$ and the lamp tail is bounded by (3.2), proving (3.4). If the interval for κ is empty, the split assertion is unused. □

The lemma is stronger than a separate fixed-order moment family. A list chosen for bounded independence can have an exceptional norm-one block. Uniform control in the moment order cannot be inferred from that earlier family. Here the exponential-moment average supplies the additional information directly, without an additional dimension exponent.

4 The density-free profile law

Fix a pulse program π with unbounded one-count k and polynomial delay $t_J \le C(J+1)^p$. Let

$$h(\ell) = \min\{J : k(J) = \ell\}, \quad h(0) = 0, \qquad s(t) = 2^{-h(\lfloor \log_2(1/t) \rfloor)-1} \quad (0 < t \le 1), \quad s(0) = 0.$$

The first hit, including the beginning of a zero plateau, is essential. The function s is nondecreasing and at most $1/2$.

For each of the 28 normalized source and marker residuals take its own two-clock controller and stopped calibrator from Section 2. Replace the original test $R_i = 0$ by $R_i = V_i T_i$ and the penalty test $S_i = 0$. The domains make V_i unitary and T_i a positive contraction. Use the same finite program table in every copy. Different residual workspaces have no imposed commutations with one another's moduli or with the source letters; the polar factors are excluded from the instantaneous commuting lists. This is exactly the source-calibration presentation of [13, Section P5] for the chosen program, compiled by its Theorem 3.1. Denote the fixed rational BCS game by G_π.

Theorem 4.1 (Pointwise profile)**.** *There are constants C, C_*, n_0 depending on the fixed program such that for $n \ge n_0$, with $k_n = k(\lfloor n/2 \rfloor)$,*

$$\log_2 A_{G_\pi}(n) \le 2k_n + \log_2 n + C, \tag{4.1}$$

$$\log_2 E_D^{A\to B}(\rho_{AB}) \ge 2k_n - C\log(n+2) - C\log^2(k_n+2) - C, \tag{4.2}$$

where the second estimate holds for every original mixed state of every loss-2^{-n} strategy whenever $k_n \ge C_* \log(n+2)$. The upper estimate is attained as an upper bound by maximally entangled strategies. If $k(J)/\log(J+2) \to \infty$, then, for $X = A_{G_\pi}, d_{G_\pi}, p_{G_\pi}, f_{G_\pi}$ at dyadic precision,

$$\log_2 X(n) = 2k_n + O\big(\log(n+2) + \log^2(k_n+2)\big). \tag{4.3}$$

The game has value one and no finite-dimensional perfect strategy.

Proof. Let $C_c \ge 1$ be the completeness loss constant and choose flat precision η by $C_c\eta^2 = 2^{-n}$. Put

$$x = \frac{n + \log_2 C_c + 1}{2}, \qquad \ell = k(\lceil x \rceil - 1) + 1, \qquad \kappa = 2^{-\ell},$$

and

$$N = \left\lceil 16\big(n\ln 2 + \ln(4C_c) + 1\big)4^\ell \right\rceil.$$

The first-hit identity $h(\ell) \ge x$ is equivalent to $k(\lceil x \rceil - 1) < \ell$. Hence $s(\kappa)^2 \le \eta^2/8$. The choice of N gives

$$\sqrt{2}\, s(1)^2 e^{-N\kappa^2/16} < \eta^2/8.$$

Also $\kappa \ge C_h/\sqrt{N}$ eventually. Lemma 3.1, applied with $f = s$, makes every $S_i = s(|R_i|)$ have flat norm at most $\eta/2$.

For each residual choose its unitary polar extension $R_i = V_i|R_i|$ and set $T_i = |R_i|$. On a spectral value $t > 0$, run the exact scalar controller, stop at the first physical age at which the one-count is $\lfloor \log_2(1/t) \rfloor$, and initialize $S = s(t)$. Until stopping, the two calibrator coordinates are $2^{k(j)}t$ and $2^j s(t)$; at the stop the second equals $1/2$. Afterwards they freeze. At $t = 0$ set both to zero and do not stop. A cyclic reset reinitializes the coordinates and counters even when this first stop lies beyond the reset period. These are the exact upper histories of the imported construction.

The reciprocal-age coordinates $1/(a+1)$, counter coordinates $n_i(a)/(a+1)$, witnesses $\max\{n_i(a)-1, 0\}/(a+1)$ and exact zero tests remain in their domains because each instruction changes at most one counter by one. All non-clock equations other than $S_i = 0$ are exact. Spectral histories are assigned separately for each T_i, so simultaneous assignments require no commutation of different source residuals. Use $n + O(1)$ bits on each of the 56 clock factors. Their unguarded increment norm is $\sqrt{2}\, 2^{-N_c/2}$ and their Cantor defects are $O(4^{-N_c})$; all are at most $\eta/2$ after increasing the fixed offset. No register of size $t_{h(\ell)}$ is added. The total logarithmic dimension is

$$3N + O(\sqrt{N}) + 56(n + O(1)).$$

Completeness now gives loss at most 2^{-n}. Taking a logarithm once more gives $2\ell+\log_2 n+O(1)$. The unit-increment property absorbs the fixed shift between $\lceil x\rceil-1$ and $\lfloor n/2\rfloor$, proving (4.1).

For the lower estimate start with any original strategy and decode at the common precision $\delta = C\,2^{-n/2}$. Set

$$J = \lfloor n/2 - K\log(n+2)\rfloor - O(1), \qquad H = t_J.$$

Choose K larger than every fixed localization and controller degree. Polynomial delay ensures $2^J(H+2)^{a_1}\delta \le 1$. In (2.1), the tested $S=0$ bounds $\|S\varphi\|$ by δ times a polynomial in H. The potentially large term is

$$2^{J-k(J)}\delta\ \mathrm{poly}(H) = 2^{-k(J)}\big[2^J\delta\ \mathrm{poly}(H)\big].$$

The bracket is bounded by the choice of J. Charge the initialized corner back to the state. Since $k(J)\le J$, all additive errors are absorbed, yielding

$$\|T_i\|_\rho \le C(n+2)^b 2^{-k(J)}.$$

The decoded factorization is approximate: $\|R_i - V_iT_i\|_\rho \le \delta$. Thus $\|R_i\|_\rho \le \|T_i\|_\rho + \delta$, rather than an exact equality. Unit increments give $k(J)\ge k_n - O(\log(n+2))$. All 28 source and marker residuals are therefore at most

$$e_n = C(n+2)^c 2^{-k_n}.$$

Actual mirrors, inverse mirrors and square-root-state commutators are $O(2^{-n/2})$, and fit the same bound because $k_n \le n/2$. For C_* large the tolerance tends to zero and gives fixed marker visibility. Apply (2.2) with $\log_2(1/e_n) = k_n - O(\log n)$. This is (4.2), on the original mixed state before taking any infimum.

Distillable entanglement is bounded by formation and by the logarithm of local dimension; on a pure state it equals its entropy. The maximally entangled upper strategies therefore squeeze all four demands into (4.3). Their losses tend to zero, proving value one. To exclude finite perfection for every unbounded count, including counts outside the lower-estimate regime, suppose a finite perfect strategy exists. Zero-error decoding gives an exact model on the support of its marginal state. Exact state commutators make that support invariant and the restricted state a faithful trace on the decoded algebra. The complete Cantor clock has 2^M mutually orthogonal binary cylinders of trace 2^{-M} at every depth M, by [13, Lemma 4.1, fair-cylinder identity (4.3)]. All these projections are nonzero, forcing dimension at least 2^M for every M, a contradiction. □

Corollary 4.2 (Prescribed regular speeds)**.** *For every regular F there is a game satisfying part 1 of Theorem 1.2. In particular $A_G(n) = 2^{F(n)(1+o(1))}$, with the same law for $d_G(n)$.*

Proof. Set $c = \lfloor F(0)/2\rfloor$ and

$$k(J) = \lfloor F(2J)/2\rfloor - c.$$

This starts at zero, is unbounded and has increments in $\{0,1\}$. A finite program stores the count and at the next pulse compares it with $\lfloor F(2J+2)/2\rfloor - c$, retaining the same initial offset. The logarithmic-space computation has only polynomially many complete configurations, including its input head, bounded output and bookkeeping. Since it halts, its time is polynomial. Stack encodings of its work tapes have polynomial values and can be simulated by unary counters in polynomially many elementary instructions. Recomputing each target and emitting its difference therefore gives polynomial total pulse delay. Apply Theorem 4.1. The main term differs from $F(n)$ by $O(1)$. Both $\log n = o(F(n))$ and $\log^2(F(n)+2) = o(F(n))$. □

Examples include integer-rounded rational powers n^a for $0 < a < 1$, integer powers of $\log(n+2)$ above one, and suitable effective versions of $n/\log(n+2)$ or $2^{\sqrt{\log n}}$, after a finite initial prefix change. For an example $\lfloor n/g(n) \rfloor$, monotonicity of g alone is insufficient: one must also verify monotonicity and unit increments of the quotient. Indeed $g(n) = 2^{\lfloor \lfloor \log_2 n \rfloor / 2 \rfloor}$ makes that quotient drop at powers of four.

Corollary 4.3 (Subscale space)**.** *The relative conclusion* $\log_2 X(n) = F(n)(1+o(1))$ *remains valid when the regular target's space bound is replaced by a total computation in space* $S(n) = o(F(n))$*, including input and output bookkeeping. The sharper additive error of Corollary 4.2 is not asserted.*

Proof. Apply the prefix-maximum and time-charge argument of [13, Corollary Q5.1] to the density-free estimate of Theorem 4.1. The changed normalization requires $S_*(n) = \max_{j \le n} S(j) + O(\log(n+2)) = o(F(n))$, which follows from monotonicity of F. A horizon shift $O(S_*)$ then loses only $o(F)$ unit increments, and the logarithms of the computation and localization costs are $o(F)$. The new source error $O(\log^2 F)$ is also $o(F)$. These substitutions give the stated relative law; no sharper additive bound is imported. □

Corollary 4.4 (The intermediate band and a matched root scale)**.** *There is one fixed game for which each* $X = A_G, d_G, p_G, f_G$ *obeys*

$$\frac{X(n)}{n^m} \longrightarrow \infty \quad (m > 0), \qquad \frac{\log_2 X(n)}{n} \longrightarrow 0.$$

The root-count game $k(J) = \lfloor J^{3/4} \rfloor$ *has*

$$\log_2 X(n) = 2\lfloor \lfloor n/2 \rfloor^{3/4} \rfloor + O(\log^2(n+2)).$$

The count $k(J) \sim J/\log_2 J$ *yields* $\log_2 X(n) = (1+o(1))n/\log_2 n$.

Proof. The root count is legal. Its finite counter implementation can find the next increment by integer comparisons $K^4 \le J^3$, using polynomially many unary instructions at pulse J. For the near-exponential example, stage r has 2^r pulses and emits one at zero-based positions divisible by r. A prefix of length v in that stage contributes $\lceil v/r \rceil$ ones. The completed-stage sum satisfies

$$\sum_{t<r} 2^t/t = 2^r/r + O(2^r/r^2).$$

To see this, put $t = r - s$, expand $1/(r-s) = 1/r + O(s/r^2)$ for $s \le r/2$, and bound the remaining geometric tail. Ceiling errors total $O(r)$. Thus throughout the stage $k(J) = J/r + O(2^r/r^2 + r) = J/\log_2 J\,(1 + O(1/\log J))$. Copying and doubling the epoch length costs $O(2^r)$ instructions; reloading the spacing countdown costs $O(r\lceil 2^r/r \rceil)$, and the remaining work is linear in the stage length. Even a partial stage costs $O(2^r) = O(J+1)$, giving linear total pulse delay. In either example $k/\log J \to \infty$ and $k = o(J)$, so (4.3) proves the formulas and the two band limits. Monotonicity between adjacent dyadic tolerances extends these assertions to general ε. □

The root example gives $\log D_G(\varepsilon) = \exp((\log(1/\varepsilon))^{3/4+o(1)})$. In dimension itself, the band lies above 2^{n^m} for every fixed m and below $2^{2^{an}}$ for every $a > 0$, eventually. The operational quantity is measured in bits and is matched to A_G, not to D_G.

5 Arithmetic laws, windows and a time obstruction

Lemma 5.1 (Endpoint names)**.** *For $0 \le L \le U < \infty$, with $L \in \Sigma^0_2$ and $U \in \Pi^0_2$, there is a bounded computable rational sequence with those lower and upper limits. If the endpoints lie in $[0,1]$, it can be chosen in $[0,1]$.*

Proof. For $L < U$ choose a rational separator r and rational bounds $a < L < U < b$. If $L = 0$, use the constant lower-limit name zero instead of u_s below. Interleave a computable lower-limit name u_s clipped to $[a, r]$ and a computable upper-limit name v_s clipped to $[r, b]$. The first has lower limit L and is at most r; the second has upper limit U and is at least r. Their interleaving has the desired endpoints. For equality use a convergent computable name, which exists because membership in both cut classes is Δ^0_2. Clipping into $[0,1]$ preserves endpoints there. This is the elementary naming fact of [13, Lemma Q4.1]. □

Theorem 5.2 (The arithmetic law at a regular scale)**.** *At every asymptotically flat regular scale F, the positive finite dimension pairs are exactly (1.2). Each is realized with the same pair for original-state usable demand, pure entropy and formation.*

Proof. Necessity is Proposition 2.1, with $\phi(x) = \log_2 \max\{2, x\}$ and denominator $F(n)$.

For sufficiency let a_s be a computable rational endpoint name, clipped into a fixed rational compact interval $[c, C] \subset (0, \infty)$ containing $[L, U]$ in its interior. Put $G(J) = F(2J)/2$. Run a foreground tracker: at pulse $J+1$ emit one if $k(J) < a_s G(J+1)$, and otherwise zero. A background counter simulation computes the next rational one elementary instruction per pulse. Pad the new denominator q by q^2 instructions before declaring it ready, and install a ready target only when the foreground count is within a fixed bounded distance of its current target. The program keeps emitting pulses during arbitrary background computations.

We give the tracking details. For a fixed target a, $aG(J)$ is unbounded and $o(J)$. A perpetual run of zeros is impossible because the target is unbounded. A perpetual run of ones is impossible because its count would be linear while the target is sublinear. Changes of run therefore recur; at a change from ones to zeros the count is within a fixed error of aG, since adjacent increments of G are at most one. Thus a ready stage is installed in finite time, and every background stage is eventually completed and installed. Stage-end ratio errors tend to zero.

No extra endpoints arise inside a stage. During a zero run k/G decreases, with its final ratio at least $a - O(1/G)$. During a one run write $k(J+h) = k(J) + h$. For any $\eta < 1/(2C+2)$, flatness gives $G(J+h) \le G(J) + \eta h + C_\eta$. It follows that this ratio cannot cross below the smaller of its starting ratio and a, except by $O(C_\eta/G)$; at each emitted one it is at most $a + O(1/G)$ whenever it has reached the active target range. Equivalently, split the run before and after its first entrance into that range and apply the displayed inequality to the initial segment. During zeros it cannot cross below $a - O(1/G)$, and during ones it cannot cross above the larger of the starting ratio and a, up to the same vanishing error. These estimates are uniform for $a \in [c, C]$. Hence all intermediate ratios remain between the stage's initial ratio and its active target up to $o(1)$. Every stage ends, so

$$\liminf_J k(J)/G(J) = L, \qquad \limsup_J k(J)/G(J) = U.$$

Denominator padding bounds both the active numerator and denominator polynomially in J, since a_s stays in $[c, C]$. Computing F and making the integer comparison have polynomial foreground cost in J. Background instructions and target installations also have polynomial pulse cost. Thus the stream has polynomial total delay. The ratio is eventually bounded below, so $k/\log J \to \infty$. Equation (4.3) has error $o(F(n))$, because $k \asymp F$ and $\log^2 F = o(F)$. Since $2G(\lfloor n/2 \rfloor) = F(n) + O(1)$, it transfers precisely these endpoints to all four demands. □

The diagonal $L = U$ proves the existing-limit assertion in Theorem 1.2. Endpoint selection is existential: rational compact bounds and separators need not be extracted uniformly from arbitrary names.

Remark 5.3 (Zero and infinite endpoints: sketch, not used). The same tracking construction permits ordered extended endpoints. For a zero lower endpoint use a baseline

$$b(J) = \left\lfloor \sqrt{G(J)\lfloor \log_2(J+2) \rfloor} \right\rfloor$$

and track $\max\{a_s G(J), b(J)\}$. Then $b/G \to 0$, $b/\log J \to \infty$ and $b = o(J)$. The target can jump at dyadic boundaries; the emitted count still has unit increments. On a preceding dyadic block there are enough pulses to fill its initial baseline, giving $k(J) \geq b(\lfloor J/4 \rfloor) - O(1)$ eventually. Flatness gives $G(J/4) \gg \log J$, so this is enough for extraction. Every fixed-target stage ends because its target is unbounded and sublinear. During zeros its ratio decreases toward the active target; at an emitted one its upper ratio is $a_s + o(1)$ or $b/G + o(1)$. This sketches the zero-endpoint extension without using an unsuitable fixed $\log^2 J$ baseline.

For unbounded rational targets, pad $(p_s + q_s)^2$, rather than only q_s^2, and require $G(J) \geq s(a_s + 1)$ before ending stage s. Each wait is finite. Numerators and denominators are polynomially bounded in the pulse index, and stage-end ratio errors are $O(1/s) + o(1)$. For any fixed finite lower barrier, choose η small enough for that barrier in the one-run argument above. This prevents upward transitions to arbitrarily large finite targets from introducing a smaller lower endpoint. Applying the argument to each finite barrier also handles both endpoints infinite. These observations describe possible extensions; no extended-endpoint theorem is used or claimed here beyond the positive finite classification and the separately proved stretched-scale theorem.

Theorem 5.4 (Stretched-scale pairs)**.** *The dimension pairs in* $[0,1]^2$ *of* $\log_2 \log_2 A_G(n)/\log_2 n$ *are precisely the ordered* (Σ_2^0, Π_2^0) *pairs there, with matched usable, pure and formation realizations by value-one BCS games without finite perfection.*

Proof. Necessity again follows from Proposition 2.1 with an eventually defined iterated logarithm. For sufficiency take an endpoint name $a_s \in [0,1]$. At $J \geq 2$ put $b = \lfloor \log_2 J \rfloor$ and

$$\tau_s(J) = \max\{2^{\lfloor a_s b \rfloor}, b^2\},$$

capped by J on a finite initial segment. Track this target greedily with zeros and ones. A ready stage ends when $\tau_s/2 \leq k \leq 2\tau_s$. Use the same pulse-triggered background and rational padding as above. Foreground arithmetic is polynomial in J.

For a fixed target, τ_s is constant in each dyadic block and increases by a factor at most two at its next boundary, eventually. If k is too large, holding it lets the unbounded target catch it. If it is too small, a block of 2^b pulses has enough positions to fill at least half a target of size at most 2^b. The factor-two stopping window therefore recurs, including $a_s = 1$. Every stage is installed and ends.

At a stage end,

$$\frac{\log_2 k(J)}{\log_2 J} = a_s + O\left(\frac{\log(b+2)}{b}\right),$$

uniformly in a_s. During zeros this exponent decreases without crossing below the active target exponent beyond that error. During ones, $k = k_0 + (J - J_0)$; for $1 \leq k \leq J$, $\log k/\log J$ is nondecreasing along such a run. At an emitted one, $k \leq \tau_s + 1$, giving the upper target bound. Thus intermediate exponents stay between adjacent stage endpoints up to a vanishing error. Their lower and upper limits are the prescribed ones.

The baseline is not asserted pointwise as $k \geq b^2$ immediately after a dyadic jump. What holds is $k \geq (b-1)^2$ eventually, because the preceding complete dyadic block has enough pulses to fill that baseline even while targets change. Consequently $k/\log J \to \infty$. Theorem 4.1 gives $\log_2 X = 2k(1+o(1))$. Taking one more logarithm and using $J = \lfloor n/2 \rfloor$ transfers the endpoints to the stated normalization. □

This theorem classifies pairs within the square. It makes no assertion excluding larger stretched endpoints for other games.

Theorem 5.5 (Quasi-polynomial windows)**.** *For every rational $m > 1$, a fixed value-one BCS game, without finite perfection, has*

$$(2m-2)\log_2 n - O_m((\log\log(n+4))^2) \leq \log_2 d_G(n),$$
$$\log_2 d_G(n) \leq \log_2 A_G(n) \leq (2m+1)\log_2 n + O_m(1).$$

The same outer bounds hold for pure entropy and formation. These are windows; no exact exponent is claimed.

Proof. Write $m = A/B > 0$. Choose a fixed r_0 with $2^r \geq \lceil m \rceil + 1$ for $r \geq r_0$ and emit a finite initial zero prefix. In each epoch $[2^r, 2^{r+1})$ emit

$$a_r = \lfloor m(r+1) \rfloor - \lfloor mr \rfloor$$

ones at its first positions and zeros elsewhere. Completed epochs telescope, and a partial epoch contributes at most $\lceil m \rceil$. Hence $k(J) = m\log_2(J+1) + O_m(1)$. This is a legal unit-increment count even for large m.

A finite unary-counter program stores the epoch length, copies and doubles it in $O(2^r)$ instructions, and counts it down with bounded work per pulse. The residue r mod B is finite control, and the finitely many possible a_r are hard-coded. Total setup work through a partial epoch is $O_m(J)$, even when all setup occurs at its beginning, because its starting pulse index is already comparable to 2^r. Thus $t_J \leq C_m(J+1)$, with the same delay exponent for every m.

For the lower bound we use the explicit initialized-corner charge, rather than asserting that controller degrees are uniform across programs. Fix this program. All its localization and counter errors are bounded by $P_m(H)\delta$ for some fixed polynomial P_m, by [13, Propositions P3.1, P4.2 and P5.1]. At decoded precision $\delta = C_m 2^{-n/2}$ choose

$$J = \lfloor n/2 - K_m \log_2(n+2) \rfloor - O_m(1), \qquad H = t_J.$$

The constant K_m may depend on every degree in P_m. Since $H = O_m(J+1)$, choose it so that $2^J P_m(H)\delta \leq 1$, enlarging P_m to include the $\sqrt{H}$ factors and the localized penalty norm in (2.1). That estimate then gives $\|T_i Q\xi\|_2 \leq C_m 2^{-k(J)} + P_m(H)\delta$. Proposition P3.1's input charge is explicitly

$$\|T_i\|_\rho \leq 2^{b+1/2}\|T_i Q\xi\|_2 + P_m(H)\delta, \qquad H < 2^b \leq 2(H+1).$$

It has a linear, not an unspecified polynomial, multiplicative cost $O(H+1)$. Moreover $J/n \to 1/2$, so the logarithmic count satisfies $k(J) = m\log_2 n + O_m(1)$, even though K_m depends on the program. All remaining additive errors are exponentially small times a fixed polynomial. The factorization error and actual source mirrors fit the same tolerance, giving

$$e_n \leq C_m n^{1-m}.$$

For every fixed $m > 1$ this tends to zero and gives fixed marker visibility. Equation (2.2) yields the lower coefficient $2(m-1)$ and error $O_m((\log\log(n+4))^2)$. The upper proof of Theorem 4.1

gives $2k_n + \log_2 n + O_m(1) = (2m+1)\log_2 n + O_m(1)$. Its maximally entangled models bound all four demands. Thus the exponent window has width three; no program-uniform localization degree is assumed. □

Theorem 5.6 (A computable speed outside every game)**.** *Part 4 of Theorem 1.2 holds, with the integer-root envelope stated there.*

Proof. The logarithmic-scale analogue is [13, Theorem Q7.2]. The argument here is a direct diagonal against dimension feasibility in the new iterated-logarithm normalization, with integer-root envelopes; it does not reproduce that theorem or use its time bound. Enumerate all finite rational games, repeating every index infinitely often. Fixed-dimension feasibility is a total decision even outside the value-one promise. Build F on successive blocks $[b, b^4)$, starting at $b \geq 64$, and test the next game at $m = b^3$. Write

$$\ell(n) = \lfloor n^{1/3} \rfloor, \quad u(n) = \lfloor n^{2/3} \rfloor, \quad h = \left\lfloor \frac{\ell(m) + u(m)}{2} \right\rfloor.$$

Search feasible dimensions up to the integer 2^{2^m}. If none works, or if the least feasible D is at least 2^{2^h}, make F follow the lower envelope on this block; otherwise follow the upper one. These are finite integer comparisons, requiring neither a logarithmic equality decision nor a definition of $\log \log 1$.

To follow the lower envelope, hold the inherited value until ℓ catches it, then copy ℓ. The inherited value is at most $u(b) + 1$, so catch-up occurs by $(u(b) + 2)^3 < b^3$. To follow the upper envelope from below, increment by one until it is caught, then copy it; an inherited one-unit excess is first held until caught. The upper envelope has unit increments, and upward catch-up takes $O(b^{2/3})$ steps. Thus before m the chosen envelope is reached exactly. The construction preserves monotonicity, unit increments and $\ell \leq F \leq u + 1$.

In the high case $\log_2 A_G(m) \geq h$, or the dimension exceeds the search bound and gives an even stronger inequality. Since $F(m) = \ell(m)$, the normalized ratio tends to infinity along such tests. In the low case $\log_2 A_G(m) < h$ and $F(m) = u(m)$, so its ratio is at most $1/2 + o(1)$. Regularize logs for bounded-dimensional games. Each game is tested infinitely often, excluding convergence to one. Each block choice is a finite computation, so F is total computable, although no efficient time bound is asserted. □

Together the realizations and the obstruction distinguish a regular band with no gaps in the stated sense from the much larger class of arbitrary computable targets. Endpoint classification does not classify whole profiles.

6 Caps and the boundary above one

For this section a power exponent means an endpoint of $\log_2 A_G(n)/n$. The results of Section 5 prescribe intermediate profiles, but do not supply a finite existing power exponent greater than one. The locality cap and unbounded-upper-exponent statements below are recalled from [13]; the stopping cap and conditional interface that follow are deductions developed here.

6.1 A quantitative locality cap

Let W be a finite set of words, each representing the identity in the marked Pauli source group, or its indicated six-letter restriction. The Pauli hash models of Mghirbi–Douglas [25, Proposition 9.4] decompose into equal-dimensional blocks permuted by the generators. Every word in this relator set fixes the same good blocks pointwise; the complementary fraction is

at most ν, the marker $J = -I$ is exact, and logarithmic dimension is $O_W(\nu^{-1}\log(2/\nu))$. This simultaneous version is [13, Lemma P6.6]. The marker is tested separately by $(J + I)/2 = 0$; J is not required to represent the identity.

Definition 6.1. Following [13, Definition P6.5], fix a radius L and let W_L contain all source-trivial words of length at most L. A bounded rational presentation extending the source is quantitatively local if there are constants K, c such that every equal-block model with generators permuting its blocks, exact marker $J = -I$, and common good blocks for W_L admits exact-domain auxiliary assignments with logarithmic auxiliary dimension at most $c(1 + \log(1/\nu))$ such that each defining residual satisfies

$$Q_j = Q_j(\Pi \otimes I) + E_j, \qquad \|Q_j\| \le K, \quad \|E_j\|_{2,\mathrm{tr}} \le K\sqrt{\nu},$$

where Π is its common bad-block projection, of trace at most ν. All constants and L are fixed before precision tends to zero.

Theorem 6.2 (Locality cap, recalled from [13])**.** *The compiled game of every quantitatively local extension has quantum value one and satisfies*

$$\log D_G(\varepsilon) \le C_G \varepsilon^{-1} \log(2/\varepsilon).$$

The same upper bound holds for usable, pure and formation demands. All four upper power exponents are at most one.

This is [13, Theorem P6.7], using its Lemma P6.6; its proof is not repeated here. The displayed stopped calibrators, finite mixing and chaining, and fixed residual powers belong to this class by its Lemma P6.8. Fixed-radius invariance tests are covered by Lemma P6.9 there, using neighborhoods of actual relator failures and finitely many source-trivial conjugates, rather than an arbitrarily padded bad set. Thus the relator hypothesis above is essential. The carry-checker presentation of [13, Section Q1] is a separate construction.

For integer stopped multipliers $m \ge r \ge 2$, [13, Corollary P6.10] gives the saturated exponent one, including the dimension and matched usable laws. In particular $m > r$ does not realize the nominal exponent $\log m/\log r > 1$. These cited conclusions concern the quantitative extension class, rather than all finite extensions containing source letters.

6.2 General games have unbounded upper exponents

Theorem 6.3 (Recalled from [13])**.** *For each total computable function $B\colon \mathbb{N}^2 \to \mathbb{N}$, there is a finite rational value-one game G such that, for every C, $D_G(2^{-n}) > B(C, n)$ for infinitely many n. Some such game has upper power exponent $+\infty$, as do its pure-entropy and formation demands. These existence statements also hold within rational BCS games.*

This is [13, Theorem P6.11]. Its BCS conclusion uses the repaired Lemma P6.12 there: a synchronous game must have positive probability on every diagonal question pair for the stated value-one equivalence. The hard synchronous instances have the uniform question distribution, as in [26, Theorem 1.6 and the proof of Lemma 6.7], so they satisfy this hypothesis. We use that precise reduction, including its sequential-POVM reverse transfer, without reproducing it. The quantifier order is one game for each two-argument bound B, escaping every fixed parameter C infinitely often. Each fixed value-one game's dimension profile remains computable. No infinite-upper usable demand or finite existing exponent above one is inferred.

6.3 Cheap stopping and honest-linear compression

Lemma 6.4 (A cheap deterministic answer-reduction strategy)**.** *In the literal answer-reduced game of [26, Table 5], with formula length L and underlying question count m, there is a one-dimensional synchronous oracularizable strategy of loss at most $6L^{-3}$.*

Proof. Choose the same fixed legal padded answer $a_0(x)$ for either player on each underlying question x. For each oracle pair (x, y) choose an arbitrary proof string whose two input portions are $a_0(x), a_0(y)$. Answer every bit or triple query by reading these strings. All equal-question, proof-consistency and isolated-input consistency tests pass, in both orientations. Only clause tests may fail. We allow every clause triple to reject; no accepting-tableau property is assumed.

The question alphabet has size $M = |X^{\mathrm{orac}}|(L + L^2 + L^3)$ with $|X^{\mathrm{orac}}| \geq m^2$. A clause test compares an oracle single index against a matching triple containing it. There are at most $2m^2 \cdot 3L^3$ losing ordered pairs, including both orientations. Uniform sampling on all ordered question pairs gives loss at most $6m^2L^3/M^2 \leq 6L^{-3}$. Repeated indices reduce this count; malformed padded questions are trivial. The strategy reads one fixed string in either role and is oracularizable. □

The ladder interface. Precisely, an honest-linear stopping ladder consists of finite games G_i, sizes $L_i \geq 1$, and loss weights $0 \leq \lambda_i \leq \lambda < 1$. Fix an admissible strategy class preserved by the following maps. For every finite admissible strategy T of G_{i+1} there is a strategy $\mathcal{C}_i[T]$ of G_i with

$$\operatorname{loss}(\mathcal{C}_i[T]) \leq \lambda_i \operatorname{loss}(T), \qquad \log \dim \mathcal{C}_i[T] \leq HL_i + \log \dim T.$$

At each level there is an admissible one-dimensional stopping strategy of loss at most cL_i^{-a}, with fixed $a, c, H > 0$. A fixed top wrapper may add logarithmic dimension C_0 and no loss. These bounds use the same size L_i and have no additive loss. Set $S_1 = 1$ and $S_i = \prod_{j<i} \lambda_j$.

Theorem 6.5 (Honest-linear stopping cap)**.** *The top game of an honest-linear stopping ladder satisfies*

$$\log D_G(\varepsilon) \leq C_0 + \frac{H(c/\varepsilon)^{1/a}}{1 - \lambda^{1/a}}.$$

In particular, literal MNY answer-reduction ladders satisfying this interface, with honest-linear direct parts measured in whole-level formula length, have $\log D_G(\varepsilon) \leq C_G \varepsilon^{-1/3}$.

Proof. Stop at the first level K with $S_K c L_K^{-a} \leq \varepsilon$, and compose the maps from that one-dimensional terminal strategy. This gives loss at most ε and cost at most $C_0 + H \sum_{i<K} L_i$. Such a level exists because $L_i \geq 1$ and $S_i \leq \lambda^{i-1}$. For $i < K$,

$$L_i < (c/\varepsilon)^{1/a} S_i^{1/a} \leq (c/\varepsilon)^{1/a} \lambda^{(i-1)/a}.$$

Sum the geometric series. If a weight is zero, stop immediately after it; its terminal loss is annihilated. No monotonicity of the sizes is required.

For the indicated MNY specialization, Lemma 6.4 supplies $a = 3, c = 6$ and an oracularizable terminal strategy. Introspection adds logarithmic dimension linear in the question length of the introspected game, which is bounded by whole-level formula length; answer reduction adds no dimension [26, Propositions 4.2 and 5.3]. Their completeness maps retain oracularizability. The recursive tagged mixture contributes a factor at most $1/2$ as in [26, Theorem 6.1(2)]. Tensoring the direct and recursive strategies supplies compatible tagged measurements when the direct parts have the stipulated exact honest strategies and linear cost. Under precisely those hypotheses the displayed interface applies. No assertion that every MNY recursion has honest-linear direct parts is needed. □

Both caps have a size-detection interpretation within their proved interfaces: a defect invisible outside a fraction ν of blocks costs loss $O(\nu)$, and a rejecting event sampled with probability L^{-a} permits stopping at size $O(\varepsilon^{-1/a})$. A uniform clause tester with only $O(1)$ violated clauses out of $\Theta(L)$ would give $a = 1$ under the same honest-linear and completeness assumptions. This interpretation is not a theorem constraining every nonlocal game or every compression construction. The 1/3 cap is an upper bound, with no matching sharpness assertion.

6.4 A conditional route above one

We state a sufficient hypothesis with all interfaces included. Let *linear closed compression* mean an effective compressor for polynomial-time verifiers describing finite rational synchronous games, including efficiently represented joint question distributions. Its class must contain the correlated Pauli tests of [9], the rational tagged mixtures below, and its own outputs, and must support the effective recursion construction. For each fixed verifier assume constants C, γ and a starting index such that its compressed game at index n has time and question length $O((\log n)^\gamma)$; every finite full-matrix oracularizable tracial PVM strategy lifts with loss at most C times the old loss and logarithmic dimension increased by $O((\log n)^\gamma)$, remaining in the same strategy class; every sufficiently low-loss finite tracial PVM strategy extracts, with loss multiplied by at most C, to a nonzero corner of its original algebra with normalized restricted trace. Constants and loss ranges are uniform in n. This is a hypothesis, not a consequence of existing gapless compression. In particular, independent uniform question sampling does not automatically implement a correlated direct test with uniform quantitative soundness.

Write $D^s_G(\varepsilon)$ for the least representation dimension of a finite tracial PVM strategy with loss at most ε; a finite direct sum is counted by the sum of its matrix sizes. This is the synchronous dimension demand.

Theorem 6.6 (Conditional positive power spectrum)**.** *Under linear closed compression, every positive Δ^0_2 real β is an existing power exponent for synchronous dimension demand. For general strategies a fixed synchronization test and known fourth-root corner rounding give lower exponent at least $\beta/4$ and upper exponent at most β. With dimension-compatible linear rounding, the general exponent is exactly β.*

Proof. Fix a convergent computable rational name for $\beta > 0$ and an integer $B > 2\beta + 2$. At levels $n_k = 2^{Bk^2}$ simulate a finite sequential prefix of that name within a bounded polynomial-in-n_k time budget. Clip the last completed output into $[1/(k+1), B/2]$ and call it t_k. Every fixed prefix eventually completes, so $t_k \to \beta$ without a modulus of convergence. Set $N_k = 2^{\lfloor t_k k^2 \rfloor}$, with a harmless finite initial prefix change. Then $\log_2 N_k = \beta k^2 + o(k^2)$, and the direct Pauli verifier has polynomial time in n_k. The next compressed level has time polynomial in k; its index has $O(k^2)$ bits. Define trivial games off schedule and use the effective recursion theorem within the hypothesized closed format.

At level k play the direct Pauli test on N_k qubits with probability 1/2, the compressed successor with probability $s_k = 2^{-(2k+1)}$, and a trivial branch otherwise. The cumulative recursive weight is $W_k = 2^{-k^2}$. Rational joint mixtures are explicitly part of the hypothesis. Honest play through level k followed by an arbitrary admissible terminal strategy has loss at most $C^k 2^{-(k+1)^2}$ and logarithmic dimension at most a constant times

$$\sum_{j \le k} (N_j + \mathrm{poly}(j)).$$

For each $\eta > 0$, all late terms lie between $2^{(\beta-\eta)j^2}$ and $2^{(\beta+\eta)j^2+O(1)}$. The final term lower-bounds the sum; at most $k+1$ terms plus a finite prefix upper-bound it. Hence the logarithm

of this sum is $\beta k^2 + o(k^2)$, even for oscillating approximations. No comparison of the sum with $(k+1)N_k$ is used. The adjacent logarithmic precision depths $x_k = \log_2(1/\varepsilon_k) = (k+1)^2 - k\log_2 C$, where $\varepsilon_k = C^k 2^{-(k+1)^2}$, satisfy $x_{k+1}/x_k \to 1$. Monotonicity therefore interpolates the upper exponent to every sufficiently small loss and proves value one.

At loss 2^{-x}, successive restrictions and soundness extractions give level-j loss at most $2^{-x}C^j 2^{j^2}$. When $x \geq j^2 + j\log_2 C + O(1)$, the direct Pauli rigidity theorem [9, Theorem 3.11] gives flexible closeness parameter at most $1/2$ to its 2^{N_j}-dimensional ideal algebra. We spell out the dimension point because flexible closeness permits amplification. If a finite algebra $\bigoplus_i M_{d_i}$ is a corner P in the amplified ideal algebra, the partial isometry of closeness gives $\tau_\infty(P) \geq 1 - \delta$. On a nonzero finite type-I central component with auxiliary size r_i and central weight p_i, a rank-d_i projection has semifinite trace $p_i d_i/(2^{N_j} r_i) \leq d_i/2^{N_j}$. A type-II component has no nonzero finite-dimensional full corner. Summing gives

$$\sum_i d_i \geq (1-\delta)2^{N_j}.$$

Every intervening corner extraction is dimension-nonincreasing by hypothesis. Thus the original faithful representation dimension has logarithm at least $N_j - 1$. Adjacent extraction thresholds again have ratio tending to one; this yields the synchronous lower exponent β.

Add fixed diagonal synchronization mass covering every question marginal if needed. Vidick's finite-dimensional almost-synchronous rounding [35], in the corner formulation of [23, Corollaries 3.2 and 3.4] gives, after a fixed-game Naimark dilation for POVMs,

$$D_G^s(K_G \varepsilon^{1/4}) \leq c_G D_G(\varepsilon).$$

It implies a general lower exponent $\beta/4$; the synchronous honest strategies give upper exponent β. A dimension-compatible linear version gives equality. The dimension arithmetic necessity of Proposition 2.1 then supplies the converse for existing positive dimension exponents under that additional hypothesis. No original-state entanglement theorem is inferred from this corner rounding. □

The Pauli rigidity and rounding inputs do not establish the missing compression interface. Quantum-sound low-degree testing already exists [20], and other compression frameworks include granular and hypercompression for free games [28]. The open hypothesis here is the uniform linear, additive-free, dimension-compatible, iterably closed interface just stated. It should not be identified with a quantum PCP conjecture without a reduction.

7 An exact-marginal Diophantine law

This section starts from the K_5 graph-correlation nonclosure witness of Dykema, Paulsen and Prakash [14] and the scalar-sum and Coxeter framework of Kruglyak, Rabanovich and Samoilenko [22]; see also the related nonclosure construction of Musat–Rørdam [27]. We use the quantitative exact-marginal carrier bounds of [10, Lemmas 4.4 and 4.6], version 1.0.2. The new conclusions here are the uniform Farey comparison and its Diophantine endpoint curve. Its carrier is a finite direct sum of matrix algebras with a faithful tracial state. Cost is the largest matrix-block dimension. Fix an irrational

$$\lambda \in \left(\frac{5-\sqrt{5}}{2}, \frac{5+\sqrt{5}}{2}\right),$$

and five projections with $\tau(P_i) = \lambda/5$. Write $S = \sum_i P_i$. The deficit is $\mathrm{Var}_\tau(S)$, by the identity $\sum_{i \neq j} \tau(P_i P_j) = \tau(S^2) - \tau(S)$. Let $D_\lambda(\varepsilon)$ be the least block cost at deficit at most ε.

Let $b_N < \lambda < a_N$ be the nearest rationals with denominator at most N, put

$$g(N) = (a_N - \lambda)(\lambda - b_N), \qquad H(\varepsilon) = \min\{N : g(N) \le \varepsilon\},$$

and let $\mu = \mu(\lambda) \in [2, \infty]$ be the irrationality exponent.

Theorem 7.1 (Farey law and endpoint curve)**.** *For every $\varepsilon > 0$,*

$$H(\varepsilon) \le D_\lambda(\varepsilon) \le 10H(\varepsilon).$$

Its lower and upper dimension exponents are respectively

$$\liminf_{\varepsilon \downarrow 0} \frac{\log D_\lambda(\varepsilon)}{\log(1/\varepsilon)} = \frac{1}{\mu + 2}, \qquad \limsup_{\varepsilon \downarrow 0} \frac{\log D_\lambda(\varepsilon)}{\log(1/\varepsilon)} = \frac{\mu - 1}{2\mu}, \tag{7.1}$$

with endpoint pair $(0, 1/2)$ when $\mu = \infty$. Uniformly over the parameter, $D_\lambda(\varepsilon) \le 10\lceil 1/(2\sqrt{\varepsilon})\rceil$.

Proof. Suppose the carrier has block dimensions at most N. On block l, the mean $s_l = \mathrm{tr}_{d_l}(S_l)$ is rational with reduced denominator at most N. Total variance gives

$$\mathrm{Var}_\tau(S) \ge \sum_l w_l (s_l - \lambda)^2.$$

No s_l lies in (b_N, a_N) and $\sum_l w_l s_l = \lambda$. Taking expectations of $(s_l - a_N)(s_l - b_N) \ge 0$ shows that the right side is at least $g(N)$. Nonscalar blocks and arbitrarily many central summands therefore cannot improve the lower bound.

For the upper bound, Lemmas 4.4 and 4.6 of [10] realize any rational p/q in the closed window as a scalar sum of five projections in dimension at most $2q$, and cyclic symmetrization makes ranks equal at a factor-five cost. We check coarse Farey levels as well. Put $\alpha = (5 - \sqrt{5})/2 = [1; 2, 1, 1, \ldots]$. Its best lower approximants under a denominator bound are $1, 4/3, 11/8, 29/21, \ldots$, generated by

$$r_0 = 1, \qquad r_{j+1} = 1 + \frac{1}{4 - r_j}.$$

These chain points lie outside the closed scalar-sum window, so Lemma 4.6 of [10] does not supply them. The first is realized in dimension one. The Coxeter transformation of [22] sends p/q in dimension q to $(5q - p)/(4q - p)$ in dimension $4q - p$. Its determinant is one, so there is no cancellation. Thus every exterior lower approximant is realized in its denominator dimension. Complementation supplies the upper exterior chain. If a Farey neighbor of the interior λ lies below α, it is necessarily the best lower approximant to α at that denominator bound; a better one would contradict its being the neighbor of λ. The upper case is symmetric. Consequently every needed neighbor, including coarse ones, has the same realization bound.

At $N = H(\varepsilon)$ realize both neighbors, symmetrize each, and put weight $(\lambda - b_N)/(a_N - b_N)$ on the upper block. The marginals are exactly $\lambda/5$ and the variance is exactly $g(N)$. Maximum block dimension is at most $10N$. Total carrier dimension is at most $20N$, so its exponents agree, although the literal factor ten belongs to the block convention.

We prove the endpoints, including irregular continued fractions. Write p_j/q_j for the convergents and

$$M = \limsup_j \frac{\log q_{j+1}}{\log q_j} = \mu - 1.$$

The identity follows from the convergent error bounds and best approximation. For $q = q_j$, $r = q_{j-1}/q$, and complete quotient $t = [a_{j+1}; a_{j+2}, \ldots]$, the convergent and the opposite intermediate convergent with index $0 \le h \le a_{j+1}$ have error product

$$g_h = \frac{t - h}{q^4 (h + r)(t + r)^2}. \tag{7.2}$$

For $h < a_{j+1}$ they are the Farey neighbors from level $\max\{q, q_{j-1}+hq\}$ until the next mediant denominator becomes available. This follows by substituting $\lambda = (p_j t + p_{j-1})/(q_j t + q_{j-1})$ and the determinant identity for consecutive convergents.

For finite μ, fix $\nu > \mu$. Every reduced fraction obeys $|\lambda - p/q| \geq c_\nu q^{-\nu}$. Farey adjacency gives gap $1/(q_a q_b) \geq N^{-2}$, so one error is at least $1/(2N^2)$ and the other at least $c_\nu N^{-\nu}$. Thus $g(N) \geq (c_\nu/2)N^{-(\nu+2)}$, implying lower endpoint at least $1/(\mu+2)$. Conversely at $N = q_j + q_{j-1}$ the convergent and first opposite intermediate convergent give

$$g(N) \leq C/(q_j^3 q_{j+1}), \qquad N \leq 2q_j.$$

The consecutive pair gives the same bound if the next quotient is one. Select a subsequence on which $\log q_{j+1}/\log q_j \to M$ and take the actual deficit $g(N)$. Since $H(g(N)) \leq N$, the lower endpoint is at most $1/(3+M)$. The same argument gives zero when $M = \infty$.

For the upper endpoint choose the largest convergent denominator $q \leq N/2$ and let Q be the next. Then $Q > N/2$ and $Q \leq Cq^{\nu-1}$, hence $q \geq cN^{1/(\nu-1)}$. If $Q \leq N$, consecutive convergents give product at most $1/(q^2Q^2)$. If $Q > N$, the last opposite intermediate denominator below N is at least $N/2$, giving product at most $C/(q^2N^2)$. Therefore

$$g(N) \leq C' N^{-2-2/(\nu-1)}.$$

Inversion and $\nu \downarrow \mu$ bound the upper endpoint by $(\mu-1)/(2\mu)$. For $M > 1$, use pairs q, Q approaching M and $N = \lfloor Q/2 \rfloor$. Their partial quotient tends to infinity, and (7.2) gives $g(N) \asymp (q^2Q^2)^{-1}$. At $\varepsilon = g(N)/2$, $H(\varepsilon) > N$, giving upper endpoint at least $M/(2(M+1))$. When $M = 1$, the lower endpoint and the upper bound already agree at $1/4$.

Finally $q_a + q_b > N$ implies $q_a q_b \geq N$. The product of two nonnegative errors with sum $1/(q_a q_b)$ is at most $1/(4q_a^2 q_b^2) \leq 1/(4N^2)$. This proves the universal bound and the missing upper bound $1/2$ when $M = \infty$. □

Badly approximable parameters have bounded partial quotients and $g(N) \asymp N^{-4}$, hence $D_\lambda(\varepsilon) = \Theta(\varepsilon^{-1/4})$. Roth's theorem [30] gives $\mu = 2$ for algebraic irrationals, yielding $\varepsilon^{-1/4+o(1)}$, without asserting a uniform Θ law for every algebraic parameter. The standard continued-fraction construction [3, Chapter 1], choosing partial quotients with prescribed $\limsup \log q_{j+1}/\log q_j$ realizes every $\mu \in [2, \infty]$, so the entire curve (7.1) occurs.

For any eventually positive unbounded nondecreasing f, there is a Liouville parameter in the window with $\liminf_{\varepsilon \downarrow 0} D_\lambda(\varepsilon)/f(1/\varepsilon) = 0$. Indeed choose a continued-fraction prefix in the window, then at stage j, with present denominator q, choose the next denominator Q so large that

$$f(q^3 Q/C) \geq 20jq, \qquad Q \geq q^j.$$

Later tails preserve the first-intermediate product bound up to the same absolute constant. At $N = q + q_{\text{prev}} \leq 2q$ and actual deficit $\varepsilon_j = g(N)$, the dimension is at most $20q$ while $1/\varepsilon_j \geq q^3 Q/C$. The ratio is at most $1/j$. The second condition makes the parameter Liouville. The parameter depends on f, and this is a subsequence conclusion; its upper exponent remains $1/2$.

The exact-marginal problem does not by itself decide any dimension floor for Boolean value-one games. In fact for a fixed Boolean predicate and finite support, any strictly positive real question distribution p is comparable to the rational uniform distribution u: if $c = |T| \min p_t$ and $C = |T| \max p_t$, then $cL_u \leq L_p \leq CL_u$. Constant precision rescaling preserves polynomial exponents. Thus a claim that the probability rationality alone explains a floor would not follow from Theorem 7.1.

8 Open problems

The results locate a band of prescribed speeds, with a computability obstruction to extending it without a complexity condition. Several boundaries remain substantive.

First, the least possible divergence for rational value-one games without finite perfection is unresolved. The logarithmic I_{3322} dimension law belongs to a Bell-functional optimum, and the slow lower subsequences of Section 7 belong to exact-marginal correlations. Neither changes that game question.

Second, no result here realizes a finite existing power exponent greater than one. The locality cap and honest-linear compression cap explain why their respective methods cannot supply one. Infinite upper exponents for general games do not interpolate to finite limits. The conditional interface of Section 6.4 is one sufficient route, whose existence remains open.

Third, the seam $A_G(n) = n^{1+o(1)}$ with $A_G(n)/n \to \infty$ is not settled by either a linear logarithmic profile or a regular $F \gg \log n$ in $A_G = 2^{F(1+o(1))}$. Quasi-polynomial windows give coarse control at a neighboring normalization, but leave exact quasi-polynomial exponents unresolved. Removing the polynomial localization loss and the $+\log n$ upper overhead would require new estimates; no claim that either overhead is intrinsic follows from the numerical checks.

Finally, what computational condition exactly characterizes whole dimension profiles? Logarithmic-space regular targets are realized; some computable targets with the same simple integer-root envelopes are excluded. A sharp characterization must account for this distinction as well as arithmetic endpoint classes. The original-state operational matching in this paper is proved for the constructed games.

9 How this work was done

The authors developed the mathematical direction, selected the claims and sources, and are responsible for the statements and proofs. AI models from more than one developer assisted with derivations, attempted counterexamples, source comparisons, finite computations and manuscript preparation. The generic process was to state a quantitative interface, test its weak points against adversarial examples, and retain deductions only with the hypotheses and error terms justified by that interface.

The finite checks exercise scalar recurrences, instruction tables, small matrix identities, spectral-tail selection, tracking schedules, integer arithmetic and combinatorial question-pair counts. They do not verify asymptotic theorems by numerical extrapolation, formalize the full compiler or extraction proof, or establish the conditional compression hypothesis. Those dependencies are cited mathematical inputs, and the new deductions are proved in the text. Reports of the checks are available from the authors.

A Finite checks and reproducibility

The ancillary directory contains the construction checks and independently written checks for the uniform-tail and scale arguments, sparse programs, Farey arithmetic, the recalled locality cap and compression calculations. A single runner invokes them sequentially, with numerical thread counts fixed to one. Recorded outputs and a SHA-256 manifest accompany the code. The source distribution is self-contained for these finite checks; Python with NumPy is required for its small matrix computations.

The uniform-tail checks enumerate small list multisets by direct character sums, compare their weighted expectation with the binomial expectation, and check a selected list at every spectral atom and several noninteger moment orders. Full exterior-space matrices test affine

covariance, onsite markers and the four lamp residual bounds for repeated and zero list entries. The scale checks include the initial-offset identity, a nonmonotone quotient counterexample, legal logarithmic-count epochs, dyadic endpoint tracking, survival-mass differences and exact integer-root catch-up inequalities. The sparse-program checks independently interpret finite counter tables and use noncommuting ordered-column examples. The Farey checks compare mediant and brute-force neighbors and verify the intermediate-convergent and variance identities in exact rational arithmetic.

The cap checks test the parameter and weighted-norm estimates in the recalled locality and stopped-multiplier statements. Compression checks enumerate small synthetic instances of the literal question formats, pessimistically rejecting every matching clause, and test jagged honest costs, corner-rank inequalities and the nonmonotone cumulative-cost example. Counts from these finite models support implementation and algebra checks only. They do not construct the huge source matrices used asymptotically, implement a full recursive verifier, or prove a game theorem by fitting a graph.

In particular the spectral-band diagnostic adds overlapping tail upper bounds. Its computed sufficient size is an upper-envelope diagnostic, not an exact extremizer over spectral laws. Neither a necessary $+\log n$ cost nor intrinsic window width follows from it. The analytic threshold split in Lemma 3.1 is the upper-bound proof used here.

References

[1] Laurent Bartholdi and Anna Erschler. Groups of given intermediate word growth. *Annales de l'Institut Fourier*, 64(5):2003–2036, 2014. https://arxiv.org/abs/1110.3650.

[2] Noel Brady and Martin R. Bridson. There is only one gap in the isoperimetric spectrum. *Geometric and Functional Analysis*, 10(5):1053–1070, 2000. https://doi.org/10.1007/PL00001646.

[3] Yann Bugeaud. *Approximation by Algebraic Numbers*, volume 160 of *Cambridge Tracts in Mathematics*. Cambridge University Press, 2004. https://doi.org/10.1017/CBO9780511542886.

[4] Rui Chao, Ben W. Reichardt, Chris Sutherland, and Thomas Vidick. Overlapping qubits. In *8th Innovations in Theoretical Computer Science Conference*, volume 67 of *Leibniz International Proceedings in Informatics*, pages 48:1–48:21, 2017. https://doi.org/10.4230/LIPIcs.ITCS.2017.48. https://arxiv.org/abs/1701.01062.

[5] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58(1):012202, 2017. https://doi.org/10.1063/1.4973422. https://arxiv.org/abs/1606.02278.

[6] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *Automata, Languages, and Programming*, volume 8572 of *Lecture Notes in Computer Science*, pages 320–331, 2014. https://doi.org/10.1007/978-3-662-43948-7_27. https://arxiv.org/abs/1209.2729.

[7] Andrea Coladangelo. A two-player dimension witness based on embezzlement, and an elementary proof of the non-closure of the set of quantum correlations. *Quantum*, 4:282, 2020. https://doi.org/10.22331/q-2020-06-18-282. https://arxiv.org/abs/1904.02350.

[8] Andrea Coladangelo. The I_{3322} Bell inequality requires infinite dimensions, 2026. https://arxiv.org/abs/2609.06038.

[9] Mikael de la Salle. Spectral gap and stability for groups and non-local games. *Journal de l'École polytechnique – Mathématiques*, 12:1417–1444, 2025. https://doi.org/10.5802/jep.314. Theorem numbers refer to https://arxiv.org/abs/2204.07084v2.

[10] Seth Douglas. How much dimension does an epsilon-approximation to a non-attained quantum correlation cost? a Diophantine answer for K_5, 2026. Version 1.0.2; lemma numbers refer to this version. https://doi.org/10.5281/zenodo.22944181. Concept DOI: https://doi.org/10.5281/zenodo.21876896.

[11] Seth Douglas. The I_{3322} quantum value is attained spatially but not in finite dimension, 2026. https://arxiv.org/abs/2609.05555.

[12] Seth Douglas and Nidhal Mghirbi. The price of precision: usable entanglement in nonlocal games and the geometry of groups, 2026. Version 1.0.0, unrefereed preprint. https://doi.org/10.5281/zenodo.23697346.

[13] Seth Douglas and Nidhal Mghirbi. The uncomputable price of precision, 2026. Version 4.0.0; theorem and section numbers refer to this version. Concept DOI: https://doi.org/10.5281/zenodo.23701733.

[14] Ken Dykema, Vern I. Paulsen, and Jitendra Prakash. Non-closure of the set of quantum correlations via graphs. *Communications in Mathematical Physics*, 365:1125–1142, 2019. https://doi.org/10.1007/s00220-019-03301-1. https://arxiv.org/abs/1709.05032.

[15] Rostislav Grigorchuk. On the gap conjecture concerning group growth. *Bulletin of Mathematical Sciences*, 4:113–128, 2014. https://doi.org/10.1007/s13373-012-0029-4.

[16] Michael Gromov. Groups of polynomial growth and expanding maps. *Publications Mathématiques de l'IHÉS*, 53:53–78, 1981. With an appendix by Jacques Tits. https://doi.org/10.1007/BF02698687.

[17] Michael Hochman and Tom Meyerovitch. A characterization of the entropies of multidimensional shifts of finite type. *Annals of Mathematics*, 171(3):2011–2038, 2010. https://annals.math.princeton.edu/2010/171-3/p12.

[18] Zhengfeng Ji, Debbie Leung, and Thomas Vidick. A three-player coherent state embezzlement game. *Quantum*, 4:349, 2020. https://arxiv.org/abs/1802.04926.

[19] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP* = RE, 2020. https://arxiv.org/abs/2001.04383.

[20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Quantum soundness of the classical low individual degree test, 2020. https://arxiv.org/abs/2009.12982.

[21] Martin Kassabov and Igor Pak. Groups of oscillating intermediate growth. *Annals of Mathematics*, 177(3):1113–1145, 2013. https://doi.org/10.4007/annals.2013.177.3.7.

[22] S. A. Kruglyak, V. I. Rabanovich, and Yu. S. Samoilenko. On sums of projections. *Functional Analysis and Its Applications*, 36(3):182–195, 2002. https://doi.org/10.1023/A:1020193804109.

[23] Amine Marrakchi and Mikael de la Salle. Almost synchronous correlations and Tomita–Takesaki theory, 2023. https://arxiv.org/abs/2307.08129.

[24] Tom Meyerovitch. Growth-type invariants for $\mathbb{Z}^d$ subshifts of finite type and arithmetical classes of real numbers. *Inventiones Mathematicae*, 184(3):567–589, 2011. https://doi.org/10.1007/s00222-010-0296-1. Theorem 5.1 is cited as stated in https://arxiv.org/abs/0902.0223.

[25] Nidhal Mghirbi and Seth Douglas. An explicit nonlocal game of quantum value one that needs exponential entanglement, 2026. Version 1.0.0, unrefereed preprint. https://doi.org/10.5281/zenodo.23697255.

[26] Hamoon Mousavi, Seyed Sajjad Nezhadi, and Henry Yuen. Nonlocal games, compression theorems, and the arithmetical hierarchy. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1–11, 2022. https://doi.org/10.1145/3519935.3519949. Preprint 2021; theorem numbers refer to https://arxiv.org/abs/2110.04651v2.

[27] Magdalena Musat and Mikael Rørdam. Non-closure of quantum correlation matrices and factorizable channels that require infinite dimensional ancilla. *Communications in Mathematical Physics*, 375(3):1761–1776, 2020. With an appendix by Narutaka Ozawa. https://arxiv.org/abs/1806.10242.

[28] Anand Natarajan and Tina Zhang. Quantum free games, 2023. https://arxiv.org/abs/2302.04322.

[29] Jef Pauwels. The quantum supremum of the I_{3322} Bell inequality is not attained in finite dimension, 2026. https://arxiv.org/abs/2608.29734.

[30] Klaus F. Roth. Rational approximations to algebraic numbers. *Mathematika*, 2(1):1–20, 1955. https://doi.org/10.1112/S0025579300000644.

[31] Mark V. Sapir, Jean-Camille Birget, and Eliyahu Rips. Isoperimetric and isodiametric functions of groups. *Annals of Mathematics*, 156(2):345–466, 2002. https://annals.math.princeton.edu/2002/156-2/p01.

[32] William Slofstra. A group with at least subexponential hyperlinear profile, 2018. https://arxiv.org/abs/1806.05267.

[33] William Slofstra. The set of quantum correlations is not closed. *Forum of Mathematics, Pi*, 7:e1, 2019. https://arxiv.org/abs/1703.08618.

[34] William Slofstra and Thomas Vidick. Entanglement in non-local games and the hyperlinear profile of groups. *Annales Henri Poincaré*, 19:2979–3005, 2018. https://arxiv.org/abs/1711.10676.

[35] Thomas Vidick. Almost synchronous quantum correlations. *Journal of Mathematical Physics*, 63(2):022201, 2022. https://doi.org/10.1063/5.0056512. https://arxiv.org/abs/2103.02468.

[36] Xizhong Zheng and Klaus Weihrauch. The arithmetical hierarchy of real numbers. *Mathematical Logic Quarterly*, 47(1):51–65, 2001.