The uncomputable price of precision

Seth Douglas Nidhal Mghirbi

Version 4.0.0, October 2026

Abstract

Which pairs can be the lower and upper entanglement price of perfection for one fixed game? For finite rational Boolean constraint-system games of quantum value one, the positive finite logarithmic dimension-exponent pairs are exactly $0 < L \leq U$ with $L \in \Sigma_2^0$ and $U \in \Pi_2^0$. Every pair has a matched one-way distillable-entanglement realization on the original mixed strategy state. The diagonal recovers exactly the positive Δ_2^0 coefficients and gives a proof using no $\mathrm{MIP}^* = \mathrm{RE}$ witness. For a counter program with polynomial pulse delay, positive lower one-density and one-count k, a carry-checker construction gives $\log_2 D_G(2^{-n}) = \Lambda + (m/2)k(\Lambda) + O(\log n)$, with the same usable profile, where $m, q \geq 1$ are fixed integers and $\Lambda = \lfloor n/(2q) \rfloor$. At the power scale, every admissible endpoint pair in $(0, 1]$ has matched log-dimension, usable, pure-entropy and formation realizations; the dimension endpoint conditions are also necessary. The full usable endpoint ceilings remain open. With linear slope bounds, target profiles computable in logarithmic space are realized to $O(\log n)$, and sublinear-space targets modulo $o(n)$. Every dimension profile with finite upper exponent is computable within doubly exponential time; some computable bounded-increment profiles with any admissible distinct endpoints are not realized by any game even modulo $o(n)$. This separates endpoint classification from whole-profile realization. The witness-based proof also yields the exact-trace barrier and the projection-minimum classification. In the companion reset model, local-heat endpoint pairs double while joint-reset heat stays bounded. Meyerovitch's entropy dimensions, Hochman-Meyerovitch's local-rule spectra and Sapir-Birget-Rips's time conditions are precedents in form; the invariants here are fixed-game precision costs.

The authors developed this work with current frontier AI models from more than one developer and take responsibility for every claim. How the work was done and checked, including refutation-first proof audits and finite computations, is described before the appendices.

Concept DOI: 10.5281/zenodo.23701733.

Contents

1 The price of perfection . . . 3
2 Weighted models and elementary estimates . . . 10
3 From bounded polynomials to Boolean constraints . . . 11
4 A finite Cantor clock . . . 14
5 The finite programmed presentation . . . 17
6 Every tracial limit pays for halting . . . 19
7 Simultaneous finite models . . . 21
8 Counting measurements on the actual state . . . 23
9 The matched lower bound . . . 25
10 The game and its exponent . . . 29
11 Every positive left-c.e. exponent . . . 30
12 Usable entanglement of the original state . . . 31
13 Heat depends on the order of reset and communication . . . 36
14 The exact-trace barrier and matrix approximation . . . 39
15 Escaping the barrier: a complementary halting flag . . . 44
16 Every positive limit-computable exponent . . . 48

Power-scale exponents 56
P1 Power spectra and the realization theorem . . . 56
P2 Finite-moment models for the lamp source . . . 57
P3 Polynomial localization and stopped products . . . 60
P4 Scaled counters and a stream with polynomial delay . . . 63
P5 Programmed calibration and matched game laws . . . 66
P6 Ceilings, explicit exponents and the remaining gap . . . 69

Precision profiles and exponent pairs 80
Q1 A carry checker for logarithmic profiles . . . 80
Q2 Hold streams and profile calibration . . . 83
Q3 The matched logarithmic profile . . . 84
Q4 Exponent pairs and the witness-free classification proof . . . 86
Q5 Profiles computable in small space . . . 88
Q6 Power-scale exponent pairs . . . 88
Q7 A time bound for whole profiles . . . 90
Q8 Group profiles and the trace obstruction . . . 92
Q9 Weak pure rounding of original strategies . . . 97
17 Open spectra and sharper resource laws . . . 100
How this work was done . . . 101
A Why exact periodic verification is insufficient . . . 102
B Finite computations and reproducibility . . . 104
C The finite reader construction . . . 107
D Complementary flag estimates and resource laws . . . 109

1 The price of perfection

Which numbers can be the entanglement price of perfection for one fixed finite game? The positive dimension exponents of finite rational Boolean constraint-system (BCS) games of quantum value one are exactly the positive Δ_2^0 reals. Every number in that class also occurs as the logarithmic coefficient of one-way distillable entanglement on the original mixed strategy state. The exponent measures how rapidly the least local dimension grows when one fixed device must play closer to perfection. It is the entanglement budget that precision forces on a device whose rules are fixed, so the class of possible exponents decides whether that budget can be read off from the finite rulebook. In general it cannot: a computable finite rulebook can have a noncomputable limiting resource rate. At the power scale, the positive exponents in $(0, 1]$ of log dimension, pure entropy and formation are also exactly the Δ_2^0 reals, with every such exponent realized for the original mixed-state usable demand. The full usable spectra remain open. The profile part classifies the positive finite lower/upper logarithmic dimension pairs and the power pairs in $(0, 1]$, with matched original-state usable realizations. It also separates endpoint classification from realization of a whole precision profile.

A real is *left-computably-enumerable*, or *left-c.e.*, if it is the limit of a computable non-decreasing sequence of rationals. Increasing approximations below are understood in this non-strict sense. All logarithms without a displayed base are base two, and entanglement is measured in ebits. For a finite game G of quantum value one, let $D_G(\varepsilon)$ be the least local dimension of a strategy with loss at most ε, padding the two local spaces equally. Define

$$d_G^{\rightarrow}(\varepsilon) = \inf_{\text{finite-dimensional strategies of loss } \leq \varepsilon} E_D^{A\to B}(\rho_{AB}). \tag{1.1}$$

The state may be mixed. Its one-way distillable entanglement is the asymptotic Bell-pair rate using local operations and classical communication from Alice to Bob. A real is *right-c.e.* if its negative is left-c.e.; it is Δ_2^0 if it is the limit of a computable sequence of rationals, equivalently, by Shoenfield's limit lemma [64], if it is computable relative to the halting set. Let $\mathcal{B}_{\mathrm{BCS}}$ be the set of positive limits $\lim_{\varepsilon\downarrow 0} \log_2 D_G(\varepsilon)/\log_2(1/\varepsilon)$ for finite BCS games with rational question probabilities and quantum value one.

Theorem 1.1 (Complete dimension classification and matched realization)**.** *The positive dimension spectrum is*

$$\mathcal{B}_{\mathrm{BCS}} = \{\alpha > 0 : \alpha \text{ is } \Delta_2^0\}.$$

For every positive Δ_2^0 real α there is a fixed finite two-player Boolean constraint-system game G_α, with rational question probabilities, quantum value one and no finite-dimensional perfect strategy, such that, as $\varepsilon \downarrow 0$,

$$\log_2 D_{G_\alpha}(\varepsilon) = \alpha \log_2(1/\varepsilon) + o(\log(1/\varepsilon)), \tag{1.2}$$

$$d_{G_\alpha}^{\rightarrow}(\varepsilon) = \alpha \log_2(1/\varepsilon) + o(\log(1/\varepsilon)). \tag{1.3}$$

Conversely, every finite dimension exponent of a rational finite game is Δ_2^0 whenever its defining limit exists. The usable-entanglement assertion is an inclusion, not a classification of its full spectrum.

Proof routes. Our main proof is Corollary Q4.4: a carry checker and programmed pulse counts establish the classification with no MIP* = RE witness. Section 16 gives a second, witness-based proof and additionally supplies the projection-minimum and exact-trace results.

Corollary Q4.4 proves the realization, with a second proof in Section 16; Lemma 2.4, with the limit argument in the proof of Corollary 16.7, gives the converse (compare [26, Proposition 11.13]). Corollary 13.1 extends the order contrast to every positive Δ_2^0 coefficient:

local-reset heat has coefficient 2α, while joint-reset heat stays bounded. This identifies limiting coefficients, without asserting computability or noncomputability of the finite-precision heat function.

1.1 The halting game

Fix a standard effective enumeration of programs, with halting set $K \subseteq \mathbb{N} = \{0, 1, \ldots\}$.

Theorem 1.2 (The halting game)**.** *There is a fixed finite two-player Boolean constraint-system game G_K, with rational question probabilities, such that $\omega_q(G_K) = 1$, no finite-dimensional strategy wins perfectly, and*

$$D_{G_K}(\varepsilon) = \varepsilon^{-(3+\theta_K)/4+o(1)}, \qquad \theta_K = \sum_{e \in K} 4^{-(e+1)}. \tag{1.4}$$

The exponent is Turing-equivalent to K. The sequence $\log D_{G_K}(2^{-n})/n$ has no computable certified rate of convergence.

The existence of increasing computable rational sequences with uncomputable limits is classical, beginning with Specker [69]; this construction places such a limit inside the resource law of one fixed game. Sections 2–10 prove its clock, compiler, finite-model and arbitrary-state estimates. Sections 11–13 give the left-c.e. reader, operational extraction, heat and recovery laws. Section 14 proves the exact-trace obstruction and classifies matrix-approximable projection minima. Section 15 gives a complementary non-left-c.e. example, and Section 16 proves the full logarithmic classification.

Theorem 1.3 (Power-scale classification)**.** *For every Δ_2^0 real $\beta \in (0, 1]$, one fixed finite rational BCS game of value one, without a finite-dimensional perfect strategy, has matching power exponents*

$$d_G^{\rightarrow}(\varepsilon) = \varepsilon^{-\beta+o(1)}, \qquad \log D_G(\varepsilon) = \varepsilon^{-\beta+o(1)}.$$

In $(0, 1]$, the existing power exponents of log dimension, pure original-state entropy and single-copy entanglement of formation are exactly the Δ_2^0 reals. The usable ceiling remains open; an exponent outside that class would force the polynomial formation/distillation gap of Proposition P6.2.

Proof. This is Theorem P1.1, proved in the power-scale part after Section 16. Its new ingredients are the finite-moment source models, polynomial same-state clock localization, scaled counters and stopped normalized products. Its finite program computes a rational approximation while an older rational continues to emit multiplier pulses; no convergence modulus or witness Hilbert space is needed. Section 17 records the remaining usable ceiling and the question of exponents above one. □

1.2 Dependency map

The finite presentations, stop-or-continue proof, simultaneous upper models and matched resource estimates are proved here. The Boolean correspondence is used as stated from [26, Theorem 4.9], and the actual-state interface from [26, Proposition 6.1]; Section 12 records how they feed an explicit extraction on the original mixed state. The main proof of Theorem 1.1 uses the carry-checker interface of Theorem Q1.2 and the pulse construction of Lemma Q2.1. The second proof additionally uses the finite perfect-completeness reduction of MIP* = RE [39, Theorem 12.6]. Both routes use Holevo's bound [37] in the measurement count, and Devetak–Winter hashing [23]. Classical results used without proof are counter-machine universality [63], real quantifier elimination and Shoenfield's limit lemma [64]. The bounded tracial moment

method [12, 42] enters only the exact-trace barrier of Section 14 and its uses in Proposition 15.9 and Section 16.6. The normal form [27, Theorem 5.1] enters only Theorem 14.5. The real simulation [47] and [26, Lemma 11.12] enter only the remark on real strategies and question weights. The heat paper [24] is needed only for Section 13 and the heat consequences. These dependencies are substantive: this is not a self-contained proof of the imported correspondences or complexity theorem.

Theorem (Exponent pairs and whole profiles)**.** *The positive finite logarithmic dimension pairs are exactly* $0 < L \le U < \infty$ *with* $L \in \Sigma_2^0$ *and* $U \in \Pi_2^0$*. One fixed rational BCS game realizes each pair with the same original-state usable pair. At the power scale every such pair in* $(0, 1]$ *is realized, and these conditions classify dimension pairs there. A polynomial-delay pulse word with positive lower one-density and one-count* k *gives*

$$\log_2 D_G(2^{-n}) = \Lambda + \tfrac{m}{2} k(\Lambda) + O(\log n), \qquad \Lambda = \lfloor n/(2q) \rfloor,$$

with the same usable profile, for fixed integers $m, q \ge 1$*. Under linear slope bounds, space-*$o(n)$ *target profiles are realized modulo* $o(n)$*, but some computable bounded-increment profiles with any admissible distinct endpoints are not realized by any game modulo* $o(n)$*.*

These assertions are proved in Theorems Q3.1, Q4.2, Q6.1 and Q7.2 and Corollary Q5.1. Proposition Q7.1 supplies the necessary doubly exponential time bound. The arbitrary-game usable endpoint ceiling and the intervening profile frontier remain open.

The local lamp class has an additional sharp ceiling: $\log D_G(\varepsilon) = O(\varepsilon^{-1} \log(1/\varepsilon))$, and its positive existing power spectra for log dimension, pure entropy and formation are exactly $\Delta_2^0 \cap (0, 1]$. General rational BCS games have no such upper-exponent ceiling: some fixed games have infinite upper power exponent for all three resources. Existing finite limits above one remain open. The usable-spectrum problem is reduced conditionally to weak pure rounding of original mixed strategies. For groups we prove effective upper classes and a trace obstruction for these counter presentations; neither supplies group realizations or excludes all analytic decoders. Turing-machine encodings into group rings and irrational analytic kernel projections are established precedents [33, 3, 61]. See Theorems P6.7 and P6.11 and Sections Q8 and Q9.

1.3 Related work and credit

Mousavi, Nezhadi and Yuen prove that deciding whether the quantum value is exactly one is Π_2^0-complete [53, 54]. Their theorem rules out uniform computable dimension bounds. Theorem P6.11 extracts a stronger consequence for individual games: given any computable two-parameter envelope, one fixed value-one game exceeds each of its parameter slices infinitely often. The BCS version uses Paddock's SynchBCS construction [59, Section 4.3 and Proposition 4.22] and the weighted sequential-POVM proof in Lemma P6.12. This hierarchy argument establishes unbounded upper exponents, without establishing an existing limit above one. Theorems 1.1 and 1.2 identify convergent coefficients for individual fixed games. The theorem $\mathrm{MIP}^* = \mathrm{RE}$ puts halting information in a game's value [39]; it does not by itself identify a resource exponent of a fixed value-one game.

The algebraic connection between constraint games and operator assignments goes back to Cleve and Mittal [18] and, for linear systems, Cleve, Liu and Slofstra [17]. Slofstra established an embedding theorem for solution groups [67] and the non-closure of quantum correlations [66]. Slofstra and Vidick introduced the hyperlinear profile as a quantitative bridge from approximate representations to dimension requirements [68]. Slofstra later obtained a finitely presented group with at least subexponential hyperlinear profile [65]. Mančinska and Vidick give a two-player game with two questions per player, but answers of unbounded length, whose optimal value one is reached only in the limit: some local-dimension-d strategies

succeed with probability $1 - O(d^{-c})$, and none exceeds $1 - \Omega(d^{-2})$ [46]. Ji, Leung and Vidick construct a fixed finite three-player game of value one for which near-perfect play requires at least $\Omega(\varepsilon^{-c})$ qubits for a fixed $c > 0$ [38]. Coladangelo constructs a fixed finite two-player score game for which every ε-optimal strategy has local dimension $2^{\Omega(\varepsilon^{-1/8})}$ [19]. These results supply the framework and important growth examples; they do not identify the matched finite exponent in (1.4).

Mehta, Slofstra and Zhao construct halting-dependent Pauli systems and embed them in finite presentations with polynomial bounds on relation decompositions [48, Proposition 4.8]. Those bounds transfer lower estimates; they do not assert small simultaneous finite-model extensions. The construction here supplies its own upper models. The conjugacy gadgets used in our compiler are those of Fanizza, Kroell, Mehta, Paddock, Rochette, Slofstra and Zhao [28, Lemma 5.8 and Theorem 5.5]. Their credit is essential: the passage from nested conjugacy to ordinary Boolean contexts is not a new gadget of this paper. We give its matrices and converse calculation because quantitative completeness is needed here.

The state-preserving Boolean correspondence was established in *The price of precision* [26, Theorem 4.9]; its linear-system predecessor and further dimension examples appear in the companion paper [50]. We use that correspondence as an existing theorem. Section 10 records the elementary square-root estimate needed for the present local-dimension application and checks the constants' effect on the exponent. Theorem 1.1 proves and strengthens the computable-coefficient inclusion conjectured in the discussion of [26, Problem 18.3(i)], for its original usable-entanglement demand as well as dimension. The Δ_2^0 upper bound for dimension exponents is [26, Proposition 11.13]; Lemma 2.4 records the rational case.

Uncomputable asymptotic data in finite systems have other precedents. Hochman and Meyerovitch characterize the entropies of multidimensional shifts of finite type as exactly the nonnegative right-c.e. reals [36]; their construction passes through a shift whose maximal symbol frequency is a prescribed right-c.e. real. Their theorem is a template for identifying the possible values of a physical quantity of a fixed finite system with a computability class. The exact precedent in form is Meyerovitch's entropy-dimension classification [49, Theorem 5.1(3), arXiv version], for every $d \geq 2$:

> The class of entropy dimensions of $\mathbb{Z}^d$ SFT's is $[0, d] \cap \Delta_2$.

The same theorem gives upper entropy dimensions Π_3 and lower entropy dimensions Σ_2. Our classification concerns nonlocal-game dimension exponents and matched entanglement coefficients. It is not the first Δ_2^0 classification of an exponent. Sapir, Birget and Rips relate group isoperimetry and computation [62]; a Dehn exponent is not a dimension exponent, and no identification between them is used here. Cubitt, Perez-Garcia and Wolf prove undecidability of the spectral gap over families of local Hamiltonians [21]. Bausch, Cubitt and Watson obtain uncomputable phase diagrams for a continuous one-parameter family [6]. These are precedents for asymptotic noncomputability, rather than proofs of (1.4).

The binary adding machine used below is the classical odometer on a Cantor set; its interpretation through operator algebras is standard [34]. The new point needed here is a finite polynomial presentation and a carry estimate that remain quantitative in a non-flat state. The final information bound uses Holevo's theorem [37] and the random-access-code principle exemplified by Nayak [57]. Its branching estimate is proved explicitly.

Jockusch and Schupp prove that the densities of computable sets are exactly the Δ_2^0 reals in $[0, 1]$ [40, Theorem 2.21], and Downey, Jockusch and Schupp classify the densities, lower densities and upper densities of computable and c.e. sets [27, Section 5]. The flag densities forced in Section 16 realize the same arithmetic class; what is new here is forcing such a density on every near-perfect strategy of one fixed game. The moving-window estimate in Section 11 uses the binary radical-inverse ordering of van der Corput [70], not a new

discrepancy principle. Vidick's *Almost synchronous quantum correlations* [72] supplies the closely related spectral-compression mechanism; Section 14 gives the word-moment form needed here. Tracial moment optimization [12, 42] supplies the exact-trace lower hierarchy. Its distinction from general commuting-model NPA [56] is the distinction used to escape the barrier in Section 15.

Devetak–Winter hashing [23] is the operational last step in Section 12. The new ingredient is the inactive-outcome extension of the triangular extraction in [26, Theorem 6.2] and its cyclic finite-alphabet form [26, Lemma 6.3], with random targets and an original-mixed-state bridge. The heat conclusion invokes the existing reset sandwich of [24, Theorem 2.1 and Corollary 2.2]. Kolchinsky and Wolpert [43] study noncomputable heat functions for universal Turing machines. Corollary 13.1 and Theorem 15.7 identify uncomputable logarithmic precision coefficients for fixed finite nonlocal games and compare their local and joint reset schedules. Section 15 additionally uses the finite perfect-completeness form of Ji–Natarajan–Vidick–Wright–Yuen, states its exact theorem numbers, and credits the Connes embedding and commuting-complexity context, including Lin [45].

Fixed-size quantum-correlation undecidability also has important precedents in Fu–Miller–Slofstra [30], and quantitative spectral-gap and representation-stability methods for Pauli-testing game families in de la Salle [22]. Their statements concern the specified correlation inputs or varying game families. They do not alone supply a matched fixed-game original-state usable spectrum. Finite-field independence and small-bias methods are credited to Naor–Naor and Alon–Goldreich–Håstad–Peralta [55, 1] at the source construction.

The power-scale comparison with group isoperimetry is precise. Brady–Bridson [10] realize the dense constant-scaling Dehn exponents $2\log_2(2p/q)$. Brady–Bridson–Forester–Shankar [11, Theorem A] obtain $2\log_\lambda r$ from an irreducible nonnegative integer matrix with Perron eigenvalue $\lambda > 1$ and rational r above every row sum. Sapir–Birget–Rips [62] program Dehn exponents under explicit digit-computation time hypotheses (double-exponential digit computation suffices for their stated $\alpha \geq 4$ Dehn laws, and a triple-exponential necessary bound is proved); this is not an unrestricted theorem for all computable exponents. Here the invariant is an original-state game resource, with no Dehn-to-precision reduction. Arbitrarily slow rational computation is interleaved with continued pulse generation, and scaled-counter errors remain polynomial in physical depth. Meyerovitch's limit-computable entropy-dimension theorem remains the direct precedent in form. Classical counter universality is credited to Shepherdson–Sturgis and Minsky [63, 51]; rotation coding and logarithmic independence to Morse–Hedlund and Baker [52, 4]. Clifford overlapping qubits [14, Theorem 3.1], the collective Houghton clock [26, Proposition 13.9], the source certificate [50, Definition 7.1] and the original-state extraction [26, Theorems 6.4 and 12.4] are the essential geometric and operational inputs. Their role is explicit in the power-scale proofs.

1.4 Proof outline

The main route starts with two compulsory carry-checker clocks. A polynomial-delay counter program emits a pulse word of positive lower one-density, and stopped calibrators force the source-clock depth at a precision-dependent scale. If $k(J)$ is the one-count, Theorem Q3.1 gives the whole profile

$$\log_2 D_G(2^{-n}) = \Lambda + \frac{m}{2}k(\Lambda) + O(\log n), \qquad \Lambda = \lfloor n/(2q) \rfloor,$$

with the same original-state usable law. Computable pulse densities realize the admissible lower and upper endpoint pairs by Theorem Q4.2; taking equal endpoints gives Theorem 1.1 through Corollary Q4.4. No fixed observable is assigned the limiting coefficient. The next paragraphs describe the second proof and its additional exact-trace information.

At source precision 4^{-N}, three compulsory clocks cost $3N - o(N)$ bits of logarithmic dimension. They hold an address, fuel and time. For G_K, a finite counter program recognizes halting inputs on separated base-four cylinders. Reset forces initialized computations at every finite fuel scale. Exact tracial models must flag every halting cylinder; compactness extends this to all weighted approximate models, without an effective modulus.

A conditional clock contributes $\theta_K N - o(N)$ further bits. Explicit finite histories are diagonal labels on the existing registers. A moving register of optional bits realizes the conditional clock in dimension $2^{\theta_N N + O(\log N)}$, with computable $\theta_N \uparrow \theta_K$. Source precision 4^{-N} becomes game loss of order 16^{-N}, giving the denominator four.

For an arbitrary left-c.e. $\theta \in [0, 1]$, an address reader accepts exactly the binary values below θ. Radical-inverse window counts match its finite flag density. Extra compulsory clocks add to the numerator, and replacing polynomial residuals by fixed powers divides the exponent. These two operations cover every positive left-c.e. α. Finally a coherent extraction with an inactive outcome converts the expected active count into distillable Bell pairs on the original mixed state. The cited heat sandwich then matches both reset laws.

The complementary construction reverses the flag rule: an unflagged decoded cylinder must carry a finite perfect witness for its halting game. A fixed counter transducer with two operator registers tests every relation of that witness's projection presentation, packed into three data generators. Nonhalting unflagged corners have zero weight in matrix-approximable traces, and moment compression transfers that bound to the actual strategy state. Finite perfect witnesses can be admitted so slowly that their combined logarithmic dimension is at most $\sqrt{N}$, while the unflagged density tends to θ_K. The flag density is therefore $1 - \theta_K$, giving exponent $(4 - \theta_K)/4$. The same extraction and heat sandwich apply. The exact-trace minimum is strictly smaller, as the barrier forces it to be.

The stop-or-continue construction replaces monotone admission by a stopping parameter R_*. Each continuation has one loaded finite perfect witness; each stop runs a scalar verifier against the entire remaining approximation. A bad address can neither stop safely nor continue forever. For the upper bound, any fixed threshold below the limit needs only finitely many witnesses. Letting the threshold advance sufficiently slowly makes their logarithmic overhead sublinear. This handles every computable convergent approximation, without a variation bound. Within this mechanism, even though both quantum and commuting values equal one, the exact-versus-matrix-approximable trace gap associated with the failure of Connes' embedding conjecture changes the attainable exponent class beyond the exact-trace left-c.e. barrier. The carry-checker proof of Corollary Q4.4 reaches the same class with no exact-versus-approximable gap, so this gap belongs to the flag mechanism and is not necessary for the classification.

The witness-based construction at a glance

The following witness-based pipeline separates the finite rulebook from the two asymptotic estimates. The imported Boolean correspondence is the final conversion, not a new complexity reduction.

Stage	Finite construction	Role in the estimates
Counter program	A program computes $a_s \to \theta$. At each depth it stops and verifies the tail, or loads one witness and continues.	Bad prefixes force flag mass at least $1 - \theta$ in matrix-approximable traces.
Polynomial presentation	Three compulsory clocks, one conditional clock, R_*, commutative workspace, per-gate data and two contraction registers.	Positive initialized corners test one tuple; fixed thresholds give exact new relations with a fixed finite witness factor.
Parameters	Add r clocks and take the qth power of every defining residual.	The resource coefficient is $(4 + r - \theta)/(4q)$; every positive Δ_2^0 real has this form.
Compiler	Rational bounded relations pass through conjugacy gadgets to finitely many Boolean contexts.	Upper residual $O(4^{-qN})$ gives loss $O(16^{-qN})$; loss ε gives unpowered precision $(C\sqrt{\varepsilon})^{1/q}$.
BCS game	A fixed finite rational question distribution tests those contexts.	Explicit flat models have $\log d = (4 + r - \theta)N + o(N)$; no game changes as precision improves.
Upper flow	Threshold $b < \theta \to$ finitely many witnesses $\to$ initialized histories $\to$ slowly advancing $b \uparrow \theta$.	Witness logarithmic cost is $o(N)$; window counts fix the optional-register cost.
Lower flow	Near-perfect strategy $\to$ actual-state interface $\to$ weighted flag bound $\to$ raw clock measurements.	Counting bounds dimension; coherent extraction and hashing bound one-way entanglement on the original mixed state.
Heat consequence	Apply the existing reset sandwich only after both bounds match.	Local reset costs $2\alpha \log_2(1/\varepsilon) + o(\log(1/\varepsilon))$; joint reset stays bounded.

The carry-checker route in Section Q1 uses two compulsory checker clocks and calibrated source clocks instead of witness admission. Its scale-dependent pulse count is not the exact-trace minimum of one fixed observable. The profile and time-bound comparisons with Meyerovitch, Hochman–Meyerovitch, Zheng–Weihrauch and Sapir–Birget–Rips are detailed in Section Q7.

The sharp lamp models use the hash-register mechanism of the companion construction [50, Proposition 9.4], building on Chao–Reichardt–Sutherland–Vidick's overlapping-qubit methods [14, 15]. The uniform canonical-trace rationality input is Zalesskii's theorem, with Kaplansky's positivity distinguished from rationality; Burger–Valette give a proof and historical account [13]. Hayden–Leung–Winter's high-entanglement subspaces [35] furnish the one effect-only obstruction to universal purification used here. None of these inputs identifies the general usable or group realization spectrum.

2 Weighted models and elementary estimates

Definition 2.1. A bounded polynomial presentation consists of finitely many variables, each declared to be a positive contraction, a projection, a self-adjoint unitary, or a unitary, and finitely many rational real-coefficient noncommutative $*$-polynomial equations $r = 0$. Its weighted precision-δ model on $\mathbb{C}^d$ is an assignment in the declared domains and a density matrix ρ such that, with $\xi = \sqrt{\rho}$,

$$\|r\xi\|_2 \le \delta, \qquad \|[s,\xi]\|_2 \le \delta \tag{2.1}$$

for each defining relation r and each variable s. The norm is the unnormalized Hilbert–Schmidt norm. A flat model has $\rho = I/d$ and $\|r\|_f = d^{-1/2}\|r\|_2$. Write $W(\delta)$ and $H(\delta)$ for the least weighted and flat model dimensions, with value ∞ if no model exists.

We write $|D|_\rho = \|D\xi\|_2$. Constants for a fixed presentation may depend on its finite polynomials, but never on d, ρ, δ or a growing digit depth.

Lemma 2.2 (Right multiplication)**.** *For bounded A, B,*

$$|AB|_\rho \le \|A\||B|_\rho, \qquad |BA|_\rho \le \|A\||B|_\rho + \|B\|\|[A,\xi]\|_2.$$

A word w of length l in contraction variables and unitary adjoints satisfies $\|[w,\xi]\|_2 \le l\delta$. A fixed finite algebraic derivation from the defining relations has error at most $C\delta$ on the same state.

Proof. Use $BA\xi = B\xi A + B[A,\xi]$. Expand the commutator of a product one factor at a time. Adjoint unitaries have the same state-commutator norm as the unitary. For adjoint relations,

$$|r^*|_\rho = \|\xi r\|_2 \le |r|_\rho + \|[r,\xi]\|_2.$$

All variables are bounded, so a fixed polynomial has a fixed operator-norm bound and a commutator bound $C\delta$. Expand a finite ideal identity as a finite sum $a_j r_j b_j$ and apply the first two inequalities. This proves the last assertion. □

Lemma 2.3 (Hilbert–Schmidt functional calculus)**.** *Let A, B be self-adjoint, Z an arbitrary rectangular Hilbert–Schmidt matrix, and g an L-Lipschitz function on their spectra. Then*

$$\|g(A)Z - Zg(B)\|_2 \le L\|AZ - ZB\|_2, \tag{2.2}$$

$$|g(A) - g(B)|_\rho \le L(|A - B|_\rho + 2\|[B,\xi]\|_2), \tag{2.3}$$

$$|[Y, g(A)]|_\rho \le L(|[Y,A]|_\rho + 2\|[A,\xi]\|_2) \tag{2.4}$$

in the last line for a contraction Y. Also $\|[g(A),\xi]\|_2 \le L\|[A,\xi]\|_2$.

Proof. Diagonalize A on the left and B on the right. Corresponding entries in the first inequality are $(g(a) - g(b))Z_{ab}$ and $(a - b)Z_{ab}$. Square the scalar Lipschitz inequality and sum. The state-commutator estimate is its case $A = B, Z = \xi$. Insert and subtract $\xi g(B)$ to obtain (2.3). For (2.4), apply (2.2) with $A = B, Z = Y\xi$; then insert and subtract $Y\xi g(A)$. The two resulting state-commutator terms are each bounded by $L\|[A, \xi]\|_2$. These are Schatten-two estimates, not operator-norm Lipschitz statements. □

Lemma 2.4 (Finite-precision computability)**.** *For a rational-data game promised to have quantum value one, $D_G(\varepsilon)$ is computable at every positive rational ε. If $a_n = \log D_G(2^{-n})/n$ converges to β and $|a_n - \beta| \le r(n)$ eventually for a computable $r(n) \to 0$, then β is computable.*

Proof. In a fixed local dimension d, introduce real variables for the real and imaginary matrix entries of the state and both families of POVMs. Positivity is a finite list of principal-minor inequalities. Normalization and winning probability at least $1 - \varepsilon$ are polynomial conditions with rational coefficients. Real quantifier elimination decides feasibility. Test $d = 1, 2, \ldots$; the value-one promise ensures termination. For the second assertion, hardcode a threshold beyond which the error bound holds. Given k, search above it for n with $r(n) < 2^{-k-1}$ and compute a_n within 2^{-k-1}. This gives a 2^{-k} approximation to β. □

3 From bounded polynomials to Boolean constraints

A BCS specifies finitely many Boolean variables and finitely many contexts. A context lists variables and its allowed sign assignments. In an operator model the context variables are commuting reflections and their joint spectrum is restricted to those assignments. Its game asks Alice for a valid context assignment and Bob for one sign, testing consistency on an incidence. We use the uniform distribution on the finite incidence set, hence rational probabilities. The algebraic interpretation is standard [18, 17].

Theorem 3.1 (Polynomial compiler)**.** *For each finite bounded polynomial presentation $\mathcal{P}$, there are a finite ordinary BCS $\mathcal{B}$ and constants $c, C > 0$ with the following properties.*

(a) Every weighted precision-η model of $\mathcal{B}$ reconstructs a precision-$C\eta$ model of $\mathcal{P}$ on the same space and with the same state.

(b) Every flat precision-δ model of $\mathcal{P}$ on $\mathbb{C}^d$ gives a strategy for its context-variable game in local dimension at most cd, with loss at most $C\delta^2$.

All constants depend only on the presentation. In particular, guarded equations such as $(V - W)(I - A) = 0$ are compiled with their guard intact.

3.1 Domains and selectors

Every positive contraction is

$$r = \frac{2I + u + u^*}{4}, \qquad u = 2r - I + 2i\sqrt{r(I - r)}. \tag{3.1}$$

The real and imaginary parts in this expression commute and their squared sum is I, so u is unitary. Conversely the first formula is a positive contraction for any unitary. For upper models realify the whole assignment by

$$\mathcal{R}(M) = \begin{pmatrix} \Re M & -\Im M \\ \Im M & \Re M \end{pmatrix}.$$

This doubles dimension, preserves real-coefficient equations and flat norms, and takes unitaries to real orthogonal matrices. A real orthogonal matrix is a product of two symmetric orthogonal involutions: on a rotation plane use

$$\begin{pmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} \cos t & -\sin t \\ -\sin t & -\cos t \end{pmatrix},$$

and use scalar blocks on the remaining eigenspaces. Orthogonal conjugation and direct sums finish the factorization. Replace each unitary by a product of two reflections. Lower reconstructions are automatically unitary, and (3.1) is automatically positive. Their state commutators cost a fixed factor. No commutation between the two factors is imposed. The upper models are flat, so no state estimate for the square root in (3.1) is needed.

A source equation is now $p = \sum_j a_j W_j = 0$, with rational a_j and reflection words W_j. Clear denominators with an integer c, duplicate signed terms, and add cancelling pairs until

$$\sum_{j=1}^{L} \widetilde{W}_j = 2cp, \qquad L = 2^q.$$

Introduce q Pauli pairs X_i, Z_i, commuting across pairs and with every source reflection. Let e_j be the joint Z spectral projections and E the joint $+1$ projection of the X_i. Then

$$Ee_jE = L^{-1}E, \qquad \sum_{v\in\{0,1\}^q} Z^v E Z^v = I. \tag{3.2}$$

For the first identity, expand e_j in Z monomials. Every nonconstant monomial anticommutes with some X_i and hence has zero compression by E. The second identity sums the joint X spectral projections.

The controlled reflection $(I - e_j) + e_j s$ is a Boolean truth table in commuting variables. Products of these controls and controlled signs form a reflection word whose value is the unitary

$$S = \sum_j e_j \widetilde{W}_j.$$

The projection E receives its own Boolean variable x_E. Its defining context is $\{X_1, \dots, X_q, x_E\}$, whose allowed assignments are those with $x_E = +1$ exactly when $X_1 = \cdots = X_q = +1$. In an operator model this truth table says $(I + x_E)/2 = \prod_i (I + X_i)/2 = E$. Introduce $x_F = S x_E S^*$ by successive reflection conjugacies and put $F = (I + x_F)/2 = SES^*$. Orthogonality is imposed on this pair alone: the context $\{x_E, x_F\}$ forbids exactly the assignment $(+1, +1)$, which in an operator model is the condition $EF = 0$. The exact identities are

$$EFS = ESE = \frac{2c}{L} pE, \qquad p = \frac{L}{2c} \sum_v Z^v EFS Z^v. \tag{3.3}$$

They show why a low-weight selector corner cannot conceal an equation. All of its translates are used.

3.2 The conjugacy gadgets

We use the gadgets of [28, Lemma 5.8]. Their formulas are included to establish the precise finite-dimensional extension needed here. For a general relation $x_1 x_2 x_1 = x_3$, introduce commuting parity contexts

$$x_i y_i z_i = I, \qquad x_i f w_i = I, \qquad g y_2 z_3 = I,$$

and the conjugacies $fy_if = z_i$ and $w_1y_2w_1 = z_3$. The targets of these remaining conjugacies commute by the parity contexts. If the original reflections are X_i, an exact doubled extension is

$$x_i = \mathrm{diag}(X_i, X_i), \quad y_i = \mathrm{diag}(X_i, I), \quad z_i = \mathrm{diag}(I, X_i),$$
$$f = \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix}, \quad w_i = \begin{pmatrix} 0 & X_i \\ X_i & 0 \end{pmatrix}, \quad g = \mathrm{diag}(X_2, X_3).$$

Conversely $w_1 = x_1f$, $[f, x_1] = 0$, and $fy_2f = z_2$ give $x_1z_2x_1 = z_3$. Conjugation by f gives $x_1y_2x_1 = y_3$. Multiplying gives the required conjugacy.

For $aba = c$ with commuting targets b, c, use six commuting parity contexts:

$$ad_1d_2 = bd_2d_3 = d_3d_4d_5 = ad_5d_6 = cd_6d_7 = d_1d_4d_7 = I. \tag{3.4}$$

They imply

$$aba = d_1d_2d_2d_3d_5d_6 = d_1d_4d_6 = d_7d_6 = c.$$

Writing $M(a) = \left(\begin{smallmatrix} 0 & a \\ a & 0 \end{smallmatrix}\right)$, an exact extension, with old variables repeated diagonally, is

$$d_1 = M(a), \quad d_2 = M(I), \quad d_3 = M(b), \quad d_4 = \begin{pmatrix} 0 & ba \\ ab & 0 \end{pmatrix},$$
$$d_5 = \mathrm{diag}(bab, a), \quad d_6 = \mathrm{diag}(bc, I), \quad d_7 = \mathrm{diag}(b, c).$$

For clarity, $d_4^* = d_4$ and $d_4^2 = I$. The equality $aba = c$ implies that a commutes with bc, while bc is a reflection. These facts verify the fourth context; the other five follow by direct multiplication. Each listed triple consists of commuting reflections with product I. Sequentially applying the two gadgets has constant dimension cost for a fixed presentation. Pauli anticommutation is $XZX = -Z$, with a Boolean variable for the negation of Z.

3.3 Weighted soundness and flat completeness

Proof of Theorem 3.1. All identities above are finite algebraic calculations in the defining relations. A Boolean truth table gives finite algebraic relations: expand in its assignment projections, and any polynomial vanishing on allowed assignments is the finite sum of its values on forbidden assignments times their projections. Thus the controlled-variable definitions, both gadget converses and (3.3) are finite ideal derivations. In an approximate model put $E_{\mathrm{pol}} = \prod_i (I + X_i)/2$, $E_0 = (I + x_E)/2$ and $F_0 = (I + x_F)/2$. The defining AND context of x_E and the conjugacy gadgets give $|E_{\mathrm{pol}} - E_0|_\rho = O(\eta)$ and $|SE_{\mathrm{pol}}S^* - F_0|_\rho = O(\eta)$, while the orthogonality context gives $|E_0F_0|_\rho = O(\eta)$. Interpret (3.3) as a polynomial identity modulo the Pauli and controlled-variable relations. Substitution of these estimates, using Lemma 2.2, gives $|p|_\rho = O(\eta)$. Only E_0 and F_0 are automatically projections here; the selector polynomial need not be a projection in an approximate model. Lemma 2.2 proves linear weighted soundness. This statement is stronger than mere exact implication; the explicit finite derivations are what justify the linear error.

For upper models choose the selectors on an exact tensor factor and define every controlled variable and successive conjugate projection canonically. Put $x_E = 2E - I$ with $E = \prod_i (I + X_i)/2$. It commutes with every X_i, and the joint spectrum of $(X_1, \ldots, X_q, x_E)$ lies in the allowed assignments, so the defining context of x_E is exact. The conjugacies defining x_F are exact even when p is nonzero. All gadget contexts extend exactly. The only inexact constraint is the final orthogonality, where

$$\|EF\|_f = \|ESE\|_f = \frac{2c}{L\sqrt{L}}\|p\|_f. \tag{3.5}$$

Let F' be the projection onto $(I-E)\operatorname{ran} F$. Then

$$EF' = 0, \qquad \|F - F'\|_f^2 \le 2\|EF\|_f^2. \tag{3.6}$$

To prove this, diagonalize FEF on $\operatorname{ran} F$ with eigenvalues λ_i and orthonormal eigenvectors v_i. For $\lambda_i < 1$, the vectors $(I-E)v_i/\sqrt{1-\lambda_i}$ are orthonormal and span the new range. Those with $\lambda_i = 1$ disappear. Hence

$$\|F - F'\|_2^2 = \#\{i : \lambda_i = 1\} + 2\sum_{\lambda_i<1} \lambda_i \le 2\sum_i \lambda_i = 2\|EF\|_2^2.$$

On the orthogonality context $\{x_E, x_F\}$, Alice answers with the PVM

$$E \mapsto (+1,-1), \qquad F' \mapsto (-1,+1), \qquad I - E - F' \mapsto (-1,-1).$$

These are orthogonal projections summing to I, because $F' \le I - E$, and every outcome is an allowed assignment. Her operator for x_E is therefore $2E - I$ and her operator for x_F is $2F' - I$. Bob uses $2E - I$ and $2F - I$. On a maximally entangled state there is no disagreement for x_E, and the disagreement probability for x_F is $\|F - F'\|_f^2$. This is why E carries its own variable: if the context listed $X_1, \dots, X_q$ together with x_F, the PVM above would not be a legal answer, since F' commutes with E but in general not with each X_i, and so does not fix their values. All other contexts, including the defining context of x_E, are exact. With incidence probability π, summing over the finitely many source equations and using (3.6) and (3.5) bounds the loss by

$$2\pi \sum \|EF\|_f^2 = 8\pi \sum \frac{c^2}{L^3} \|p\|_f^2 \le C\delta^2.$$

Here c and L are the integers of each source equation. The new variable adds no tensor factor. Realification, selectors and gadget extensions multiply the dimension by the fixed factor $2^{1+q_{\mathrm{tot}}+g}$, where q_{tot} is the total number of selector qubits and g the number of gadget doublings; this is the dimension constant in (b). This proves both assertions. □

4 A finite Cantor clock

The clock has a positive contraction A, a projection P, a unitary T and a self-adjoint unitary X, with relations

$$[A, P] = 0, \quad TAT^* = 4A - 3P, \quad XPX = I - P, \quad [X, TAT^*] = 0. \tag{4.1}$$

For $j \in \mathbb{N}$ put $P_j = T^j P T^{-j}$ and $X_j = T^j X T^{-j}$. The Cantor set used here is

$$\mathcal{K} = \left\{ \sum_{j\ge 0} 3b_j 4^{-j-1} : b_j \in \{0,1\} \right\}.$$

Let p be its first-digit function, extended to $[0,1]$ by

$$p(a) = \begin{cases} 0 & a \le 1/4, \\ 2a - 1/2 & 1/4 \le a \le 3/4, \\ 1 & a \ge 3/4, \end{cases} \qquad s(a) = 4a - 3p(a).$$

These functions are respectively 2- and 4-Lipschitz and map $[0,1]$ into itself.

Lemma 4.1 (Exact digits)**.** *In an exact model, A has spectrum in $\mathcal{K}$, all P_j are commuting spectral functions of A, and*

$$A = \sum_{j<M} 3 \cdot 4^{-j-1} P_j + 4^{-M} T^M A T^{-M}. \tag{4.2}$$

The flip X_j exchanges P_j and $I - P_j$ and preserves all P_l with $l > j$. In every tracial state,

$$\tau\left(\prod_{j<M} P_j^{b_j} (I - P_j)^{1-b_j}\right) = 2^{-M}. \tag{4.3}$$

The spectral projection of A at 1 is zero in the faithful tracial representation.

Proof. Since A, P commute and $0 \le 4A - 3P \le I$, the $I - P$ block of A has spectrum in $[0, 1/4]$ and the P block in $[3/4, 1]$. Hence $P = p(A)$ and $TAT^* = s(A)$. Conjugating this statement repeatedly proves $P_j = p(s^j(A))$ and (4.2). Its remainder has norm at most 4^{-M}, which places the spectrum in $\mathcal{K}$. The tail coordinate TAT^* determines every later digit. Its commutation with X, and $XPX = I - P$, give the flip assertions, then their translates. A trace assigns equal mass to a cylinder and its earliest-bit flip. Add the two and repeat on the remaining suffix to obtain (4.3). The spectral projection at 1 lies under every all-ones cylinder and has trace at most 2^{-M} for every M. Faithfulness makes it zero. □

Lemma 4.2 (Linear spectral repair)**.** *There is an absolute constant C such that any weighted precision-δ clock has a self-adjoint $\widehat{A}$ with spectrum in $\mathcal{K}$ satisfying*

$$|\widehat{A} - A|_\rho + \|[\widehat{A}, \xi]\|_2 \le C\delta, \tag{4.4}$$

$$|P_j - p(s^j(\widehat{A}))|_\rho \le C 4^j \delta. \tag{4.5}$$

The first M raw digit projections and triangular flips satisfy digit–digit commutations, own-bit flips, later-digit preservation and state commutations with error at most $C(M+1)^3 4^M \delta$. The spectral digit projections $p(s^j(\widehat{A}))$ commute exactly and differ from the raw ones by (4.5).

Proof. Set $B_0 = TAT^*$, $D_0 = PB_0P + (I-P)B_0(I-P)$ and $C_0 = (D_0 + 3P)/4$. The affine relation and $[A, P]$ give $|[B_0, P]|_\rho = O(\delta)$ by Lemma 2.2. Since $B_0 - D_0 = (2P - I)[P, B_0]$, it follows that $|D_0 - B_0|_\rho = O(\delta)$ and $|C_0 - A|_\rho = O(\delta)$. Their state commutators are $O(\delta)$ from their fixed polynomial definitions. All are positive contractions, and D_0 commutes with P.

Let $d_{\mathcal{K}}$ be distance to $\mathcal{K}$, a 1-Lipschitz function. Cantor self-similarity on the two P blocks gives $d_{\mathcal{K}}(C_0) = d_{\mathcal{K}}(D_0)/4$. Lemma 2.3 and unitary transport give

$$|d_{\mathcal{K}}(A)|_\rho \le \tfrac{1}{4} |d_{\mathcal{K}}(A)|_\rho + O(\delta).$$

Thus $|d_{\mathcal{K}}(A)|_\rho = O(\delta)$. Round each spectral value of A to a nearest point of $\mathcal{K}$. The resulting $\widehat{A}$ has $|\widehat{A} - A|_\rho = O(\delta)$. For any self-adjoint E, $\|[E, \xi]\|_2 \le 2|E|_\rho$; apply this to $E = \widehat{A} - A$ to prove (4.4). No continuity of the nearest-point map is needed.

Since $p(C_0) = P$, functional calculus gives $|P - p(\widehat{A})|_\rho = O(\delta)$, and then the affine relation gives $|T\widehat{A}T^* - s(\widehat{A})|_\rho = O(\delta)$. For $g = p \circ s^r$, the one-step error

$$|Tg(\widehat{A})T^* - g(s(\widehat{A}))|_\rho \le C 4^r \delta$$

follows from (2.3). Telescope these identities to depth j. The sum is $O(4^j\delta)$. Every digit function has norm at most one, so transporting the errors through the outer powers of T adds only $O(j^2\delta)$, absorbed by $O(4^j\delta)$. This proves (4.5).

For the flips, translate $XPX = I - P$ and replace each P_j using (4.5). For $l > j$, use the tail coordinate $T^{j+1}AT^{-(j+1)}$. The translated last relation of (4.1) makes its commutator with X_j

at most $C(j+1)\delta$ in weighted norm. Its digit at position $l-j-1$ is a function with Lipschitz constant at most $2\cdot 4^{l-j-1}$. Apply (2.4), replace the digit back, and use Lemma 2.2. The error is at most $C(M+1)^3 4^M\delta$. Raw digit commutators follow by replacing both with commuting functions of $\widehat{A}$; the two replacement costs add. Finally $\|[P_j,\xi]\|_2, \|[X_j,\xi]\|_2 \le (2j+1)\delta$. This proves all assertions. □

Remark 4.3. The spectral repair and digit estimate use only positivity, the affine equation, $[A,P]=0$, and state commutators. They apply without a flip. This will be used for the conditional coordinates.

Theorem 4.4 (Clock profile)**.** *For the presentation* (4.1)*,*

$$\log W(2^{-n}) = \frac{n}{2} + O(\log(n+2)), \qquad \log H(2^{-n}) = \frac{n}{2} + O(\log(n+2)).$$

There are flat models on N qubits of precision $3\cdot 4^{-N}$.

Proof. On N qubits let P_j read bit j, let T rotate contents from slot j to slot $j+1$ cyclically, and let X flip bit zero. Put

$$A_N = \sum_{j=0}^{N-1} 3\cdot 4^{-j-1} P_j. \tag{4.6}$$

Then $[A_N,P_0]=0$, $XP_0X = I-P_0$, and

$$TA_NT^* - 4A_N + 3P_0 = 3\cdot 4^{-N}P_0, \qquad [X, TA_NT^*] = 3\cdot 4^{-N}[X,P_0].$$

All state commutators vanish in the flat state, proving the upper bound. For the lower bound, use M digits and flips from Lemma 4.2 in the binary case of the counting lemma, Theorem 8.1, proved below. Choose $M = \lfloor \log_4(1/\delta) - 20\log_4(2+\log_4(1/\delta))\rfloor$. Its error terms tend to zero, so $\log d \ge M - o(1)$. At $\delta = 2^{-n}$ this is $n/2 - O(\log(n+2))$. Since $W \le H$, both conclusions follow. This forward reference uses no property of the program constructed later. □

4.1 Guarded increment

Add a unitary V and the relation

$$(V - X[(I-P) + (TVT^*)P])(I-A) = 0. \tag{4.7}$$

Proposition 4.5. *The finite clocks have increment models of dimension 2^N whose guarded defect has operator norm $\sqrt{2}\,4^{-N}$ and flat norm $\sqrt{2}\,4^{-N}2^{-N/2}$. In an exact faithful tracial model, V increments every finite binary digit prefix.*

Proof. Let V_N add one modulo 2^N, with bit zero least significant. If that bit is zero, the right-hand side before the guard is the correct first-bit flip. If it is one and some later bit is zero, TV_NT^* carries through slots $1,2,\dots$ to the first zero and X then clears slot zero. Only the all-ones input fails:

$$R_N := V_N - X[(I-P_0) + (TV_NT^*)P_0] = (|0^N\rangle - |10^{N-1}\rangle)\langle 1^N|.$$

The guard equals 4^{-N} there. The two norm statements follow from the rank-one formula. For exact tracial models, Lemma 4.1 makes $I-A$ have dense range, so its bounded left factor in (4.7) vanishes. Put $C = TVT^*$ and $D = (I-P) + CP$. Since $V = XD$, D is unitary. It fixes $\ker P$ pointwise, so unitarity forces $\operatorname{ran} P$ onto itself; there D is C. Hence $[C,P]=0$. The recursion now reads binary increment on the first bit followed by increment on the tail on its carry branch. Induction on prefix length proves the assertion. □

Proposition 4.6 (Mass near the vanishing guard)**.** *There is an absolute constant C such that every weighted precision-δ clock and every integer $M \geq 1$ satisfy*

$$\mathrm{Tr}(\rho\, 1_{[1-4^{-M},1]}(A)) \leq 2^{-M} + C(M+1)^5 4^M \delta. \tag{4.8}$$

If $\|R\| \leq L$ and $|R(I-A)|_\rho \leq \delta$, then

$$|R|_\rho \leq 4^M(1+L)\delta + L\sqrt{2^{-M} + C(M+1)^5 4^M \delta}. \tag{4.9}$$

In particular $|R|_\rho = O((1+\log(1/\delta))^{5/2}\delta^{1/6})$ for fixed L.

Proof. Here it is useful to avoid rounding A first. From the proof of Lemma 4.2, $|P - p(A)|_\rho = O(\delta)$ and $|TAT^* - s(A)|_\rho = O(\delta)$. The same telescoping gives $|P_j - p(s^j(A))|_\rho \leq C4^j\delta$. Threshold $p(s^j(A))$ at $1/2$ to commuting projections R_j. Against an exact projection the cross-spectrum Lipschitz bound for this threshold is two: for $b \in \{0,1\}$ and $a \in [0,1]$, $|1_{[1/2,1]}(a) - b| \leq 2|a-b|$. The eigenbasis proof of Lemma 2.3, followed by its state-commutator insertion, therefore gives $|R_j - P_j|_\rho \leq C(j+1)4^j\delta$. Their commutators with ξ and the triangular flip defects are at most $\eta = C(M+1)^3 4^M \delta$.

For completeness, an approximately invariant suffix law differs from its earliest-bit swap by at most $CM^2\eta$. Pinch ξ in the commuting R_j to ζ; each pinching costs at most η, so $\|\xi - \zeta\|_2 \leq M\eta$. Intertwining a flip with one recorded digit has error $O((M+1)\eta)$ on ζ. Inserting previous joint spectral projections does not enlarge the sum of squared branch errors, since they commute with ζ. Telescoping at most M digits and returning to ξ costs $CM^2\eta$. The initial state changes by trace norm at most 2η under its flip. Purification distance now bounds the claimed total variation.

Let q_i be the probability that $R_i, \ldots, R_{M-1}$ are all one, with $q_M = 1$. The swap bound gives $|2q_i - q_{i+1}| \leq CM^2\eta$, hence $q_0 \leq 2^{-M} + CM^2\eta$. If $a \geq 1 - 4^{-M}$, then $s^j(a) = 1 - 4^j(1-a) \geq 3/4$ for $j < M$. Thus all R_j are one on that spectral interval. This proves (4.8).

Put $G = I - A$ and $E = 1_{[0,4^{-M}]}(G)$. Since $R\xi G = RG\xi + R[\xi, G]$, its norm is at most $(1+L)\delta$. On $I - E$ the inverse of G has norm at most 4^M, while $\|R\xi E\|_2 \leq L\sqrt{\mathrm{Tr}(\rho E)}$. Add the two column bounds and use (4.8). Finally choose $M = \lfloor \frac{1}{3}\log_2(1/\delta) \rfloor$. Both 2^{-M} and $4^M\delta$ are of order $\delta^{1/3}$, proving the last assertion. The relaxed polynomial factor here is harmless; only exact guard cancellation is needed in the flag proof. □

5 The finite programmed presentation

Fix a finite deterministic counter program recognizing K. Counters take nonnegative integer values and instructions increment, decrement on a nonzero branch, or test zero. Such a recognizer follows from the classical universality of counter machines [63]. One direct construction stores the two halves of a Turing tape as integer-coded stacks. Pushing multiplies a stack code by a fixed base and adds a digit; popping computes quotient and remainder by that base. Increment, decrement and zero-test loops perform these operations using finitely many scratch counters. A unary input loop writes the integer input on the simulated tape. This gives one fixed finite instruction table. Add absorbing idle and halt controls.

The variables introduced below have the domains stated explicitly. Projection inclusions $Q \leq F$ mean the polynomial equation $(I - F)Q = 0$. A precision model satisfies every listed commutator and every displayed defining equation to precision δ in Definition 2.1.

5.1 Three compulsory clocks and a carry coordinate

Use three copies of (4.1), with indices a, f, t for address, fuel and time. Complete variable sets of different clocks commute. Add an address increment V and a time increment U,

satisfying (4.7) for their own clocks. The address and fuel sets commute with U, and the time and fuel sets commute with V. Add a positive contraction H with

$$H = (I - P_a) + \tfrac{1}{4} P_a T_a H T_a^*, \tag{5.1}$$

$$[P_a, T_a H T_a^*] = [H, A_a] = 0, \qquad V^* A_a V = A_a - I + \tfrac{7}{4} H. \tag{5.2}$$

Include H in the address set, commuting with all fuel and time variables. The last equality will control many address translations without an exponential precision cost.

5.2 Address decoding and reset

Put $P_{a,1} = T_a P_a T_a^*$, $E_{00} = (I - P_a)(I - P_{a,1})$ and $E_1 = P_a(I - P_{a,1})$. Add a projection L and a positive contraction R with

$$L = E_1 + E_{00} T_a^2 L T_a^{-2}, \tag{5.3}$$

$$R = (I - E_{00}) + \tfrac{1}{2} E_{00} T_a^2 R T_a^{-2}. \tag{5.4}$$

The products are polynomials even before the approximate digits commute. In exact models the cylinder

$$B_e = \left(\prod_{j<e} T_a^{2j} E_{00} T_a^{-2j} \right) T_a^{2e} E_1 T_a^{-2e} \tag{5.5}$$

selects first nonzero base-four address digit one in position e. Here each base-four address digit groups two successive binary clock bits, with the first bit least significant.

Add a reset projection Z, set $S = T_f T_t$, and impose

$$Z = (I - P_f) + P_f (I - P_t) S Z S^*. \tag{5.6}$$

In exact models write $Q_j^f = T_f^j P_f T_f^{-j}$, $P_{t,j} = T_t^j P_t T_t^{-j}$ and

$$C_k = \left(\prod_{j<k} Q_j^f \right) (I - Q_k^f), \qquad D_k = \prod_{j<k} (I - P_{t,j}). \tag{5.7}$$

The product defining C_k means $(\prod_{j<k} Q_j^f)(I - Q_k^f)$, with empty product I.

5.3 Counters and transitions

Encode counter value n by $r = 2^{-n}$. For each counter add a positive contraction r_i, a zero-test projection z_i and a positive contraction s_i with

$$r_i = z_i + \tfrac{1}{2}(I - z_i) s_i. \tag{5.8}$$

All instantaneous workspace variables commute. For each control state σ, add a projection q_σ, with $q_\sigma q_\tau = 0$ for $\sigma \neq \tau$ and $\sum_\sigma q_\sigma = I$. Each program branch selector is q_σ, $q_\sigma z_i$ or $q_\sigma(I - z_i)$; these selectors partition I. The next control is the sum of selectors entering that state. The next counter is the sum, over selectors, of that selector times $r_i/2$, r_i or $2r_i$ as its instruction dictates. The factor two is used only on a nonzero decrement branch. These prescriptions define a fixed rational polynomial map Step on counters and controls.

Initialization is

$$\mathrm{Init}(q_{\mathrm{start}}) = L, \quad \mathrm{Init}(q_{\mathrm{idle}}) = I - L, \quad \mathrm{Init}(r_1) = R, \quad \mathrm{Init}(r_i) = I \ (i > 1),$$

with other controls zero. For each counter or control W, impose

$$ZW = Z\,\mathrm{Init}(W), \tag{5.9}$$

$$U^*WU = Z_+\,\mathrm{Init}(W) + (I - Z_+)\,\mathrm{Step}(W), \qquad Z_+ = U^*ZU. \tag{5.10}$$

The witnesses and zero-tests satisfy (5.8) at every time; separate update equations for them are unnecessary. Add a flag projection F and require

$$q_{\mathrm{halt}} \le F. \tag{5.11}$$

5.4 Commutations and the conditional clock

The following is the complete additional commutation list.

- L, R, F commute with each other, with A_a, and with every fuel and time clock variable, including U.
- Z and all instantaneous workspace variables commute with one another and with A_a, A_f, A_t, L, R, F.
- The address set, including V, H, commutes with the complete fuel and time sets.

Workspace is not required to commute with U or the address shifts and flips.

For the conditional clock add positive contractions B, C, a projection Q and a self-adjoint unitary Y. Impose

$$Q \le F, \qquad [B,Q] = [B,F] = [C,Q] = [C,F] = [B,C] = 0, \tag{5.12}$$

$$VBV^* = 4B - 3Q, \qquad VCV^* = 4C - 3F, \tag{5.13}$$

$$YFY = F, \qquad YQY = F - Q, \tag{5.14}$$

$$[Y, VBV^*] = 0, \qquad [Y, C] = 0. \tag{5.15}$$

Each of B, C, Q, Y commutes with A_a and every fuel and time clock variable, including U. Define

$$F_i = V^i F V^{-i}, \qquad Q_i = V^i Q V^{-i}, \qquad Y_i = V^i Y V^{-i}. \tag{5.16}$$

The intended ternary measurement is $(I - F_i, F_i - Q_i, Q_i)$. Its inactive outcome contributes no required bit; its two active outcomes are exchanged by Y_i. The coordinate C records the flags. It is needed to control commutation of their translates, which commutation with A_a alone would not imply.

This finishes the finite presentation, denoted $\mathcal{P}_K$. Every branch and every variable used above occurs in a finite list; e, k, i in later arguments index consequences, not additional defining relations.

6 Every tracial limit pays for halting

Lemma 6.1 (Forced finite initialization)**.** *In an exact model of $\mathcal{P}_K$,*

$$B_e L = B_e, \quad B_e R = 2^{-e} B_e, \qquad C_k Z = C_k D_k. \tag{6.1}$$

The zero-test equation forces $z_i v = v$ when $r_i v = v$, and $z_i v = 0$ when $r_i v = 2^{-n} v$ for an integer $n > 0$.

Proof. By Lemma 4.1, every digit is a spectral function of its clock coordinate. The required commutations imply that L, R, Z commute with the relevant digits. Expand (5.3) and (5.4) along the e successive E_{00} factors of B_e. The nonrecursive terms vanish there. At the first E_1 factor, the nonrecursive term is I. The coefficient accumulated in the second expansion is 2^{-e}. These finite expansions prove the first two identities.

In (5.6), on C_k the first k fuel bits force the recursive branch. Each supplies the next factor $I - P_{t,j}$. At the next fuel bit the constant branch is forced. Thus $C_k Z = C_k D_k$. In particular the $k = 0$ identity forces reset on $I - P_f$, so $Z = 0$ cannot be a tracial solution.

For the zero-test, z_i commutes with s_i. On its range $r_i = I$; on its kernel $0 \le r_i \le I/2$. If $r_i v = v$, the latter component of v must vanish; if $r_i v = 2^{-n} v$ with $n > 0$, the former component must vanish. The argument applies to every vector in a given subspace, including after unitary conjugation. Although $r_i = 0$ is permitted, no finite sequence of the prescribed updates starting at an ordinary encoded integer reaches it. □

Theorem 6.2 (Flag density)**.** *Every exact tracial model of $\mathcal{P}_K$ satisfies $\tau(F) \ge \theta_K$. There is a function $\omega(\delta) \to 0$ such that every weighted precision-δ model satisfies*

$$\mathrm{Tr}(\rho F) \ge \theta_K - \omega(\delta). \tag{6.2}$$

No effective modulus is asserted.

Proof. Work first in the faithful tracial representation. The clock variables satisfy Lemma 4.1. By Proposition 4.5, U increments each finite time prefix; it commutes with fuel and address variables. Fix $e \in K$ with halting time h and choose k with $2^k > h$. The operator

$$Q_0 = B_e C_k D_k (I - F)$$

is a projection. Lemma 6.1 and (5.9) set all counters and controls on its range to the initialized values. Input counter one has value $r_1 = 2^{-e}$, while other counters encode zero.

Write $W_l = U^{-l} W U^l$ and $Z_l = U^{-l} Z U^l$. On Q_0, the first k time bits are initially zero. Increment does not return them to zero before step 2^k, so

$$Z_0 Q_0 = Q_0, \qquad Z_l Q_0 = 0 \quad (1 \le l \le h). \tag{6.3}$$

We prove by induction that every counter and control at time $l \le h$ has its classical scalar right action on Q_0. The assertion holds at zero by initialization. Given these right actions, the zero-test part of Lemma 6.1, applied to the conjugated workspace, forces the correct zero-test right actions as well. Every branch selector and every update polynomial therefore has its prescribed scalar right action. If $AQ_0 = aQ_0$ and $BQ_0 = bQ_0$, then $ABQ_0 = abQ_0$; no inter-time commutation is needed.

Conjugate (5.10) by U^{-l} and multiply by Q_0 on the right. The initialized polynomials are functions of L, R and constants, so retain their scalar actions. By (6.3) the reset term vanishes and the step term yields exactly the next configuration. This proves the induction. At time h,

$$U^{-h} q_{\mathrm{halt}} U^h Q_0 = Q_0.$$

Since $q_{\mathrm{halt}} \le F$ and $[F, U] = 0$, the left side is killed by $I - F$. Thus $Q_0 = 0$.

It remains to remove the auxiliary conditions, which must not be assumed to have uniform conditional mass. The operator $B_e(I - F)$ commutes with all fuel and time flips. Conjugate $B_e(I - F) C_k D_k = 0$ by the flip of the earliest specified fuel digit. It preserves every later specified digit. Adding the two equations removes that digit. Repeat in increasing order to remove all fuel conditions, then do the same for time. The result is $B_e(I - F) = 0$. The address cylinders B_e are orthogonal and have trace $4^{-(e+1)}$ by (4.3). Taking finite sums over halting inputs and then their supremum proves $\tau(F) \ge \theta_K$.

For the uniform approximate statement, suppose instead that a sequence of models has $\delta_j \to 0$ but $\mathrm{Tr}(\rho_j F) \le \theta_K - \gamma$ for some $\gamma > 0$. Form their states on the universal C^*-algebra of the finitely many bounded-domain variables before imposing equations. A weak-* limit of moment states exists. It is tracial, since

$$\|[\rho_j, s]\|_1 \le 2\|[\sqrt{\rho_j}, s]\|_2 \longrightarrow 0$$

for each generator and hence each fixed word. Here $\|\cdot\|_1$ denotes trace norm; the inequality follows by expanding $[\xi^2, s]$ and applying Hilbert–Schmidt Cauchy–Schwarz. For every relation r, the limit has $\tau(r^*r) = 0$. Traciality gives, for bounded a,

$$\tau(a^*r^*ra) = \tau(r^*raa^*) \le \|a\|^2\tau(r^*r) = 0.$$

Thus all relations are zero operators in the tracial GNS representation. This exact model has $\tau(F) \le \theta_K - \gamma$, contradicting the first part. Taking the supremum of the positive shortfall over precision-δ models defines a function $\omega(\delta) \to 0$ and proves (6.2). □

The proof treats only finite computations before taking compactness. It never tries to quantify how far a given finite precision sees into the halting set. This compactness argument supplies no effective modulus for the uniform flag-density estimate.

7 Simultaneous finite models

Theorem 7.1. *Let $h(e)$ be the recognizer's halting time, or ∞ if it does not halt. For every even integer $N \ge 2$, define*

$$K_N = \{e < N/2 : h(e) < 2^N\}, \qquad \theta_N = \sum_{e \in K_N} 4^{-(e+1)}.$$

There is a flat model of $\mathcal{P}_K$ on dimension

$$d_N = 2^{3N+k_N}, \qquad k_N = \theta_N N + O(\log(N+2)), \tag{7.1}$$

with every defining defect at most $3 \cdot 4^{-N}$ in operator norm. The coefficients θ_N increase to θ_K.

Proof. Take N-bit address, fuel and time registers, with the finite coordinate (4.6). Their bit zero is least significant. Let the address integer be $x \in \mathbb{Z}/2^N\mathbb{Z}$ and put

$$f_N(x) = \sum_{e \in K_N} 1_{\{x \equiv 4^e \pmod{4^{e+1}}\}}. \tag{7.2}$$

The residue classes are disjoint and their periods divide 2^N, so this is a Boolean function of mean θ_N.

Every interval of N consecutive addresses has

$$\sum_{j=1}^{N} f_N(x-j) = N\theta_N + O(\log(N+2)) \tag{7.3}$$

uniformly in x, including cyclic intervals. To prove it, put $m = \lfloor \log_4 N \rfloor$. For each $e \le m$, one residue class has a count differing from its mean by at most one. Every flagged address from $e > m$ is a multiple of 4^{m+1}, so the whole remaining family contributes at most one address to such an interval. Its expected count is at most $N/(3 \cdot 4^{m+1})$. Summing proves (7.3).

Set

$$k_N = 1 + \max_x \sum_{j=1}^{N} f_N(x-j) \tag{7.4}$$

and add k_N optional qubits. Let R_0 rotate their contents from slot i to slot $i+1$ cyclically. Define

$$V|x,y\rangle = |x+1, R_0^{f_N(x)} y\rangle, \qquad F = f_N(x), \quad Q = FP_{\text{slot }0}, \quad Y = (I-F) + FX_{\text{slot }0}. \tag{7.5}$$

Define the translates F_j, Q_j by (5.16) and set

$$B = \sum_{j\geq 0} 3\cdot 4^{-j-1} Q_j, \qquad C = \sum_{j\geq 0} 3\cdot 4^{-j-1} F_j. \tag{7.6}$$

These are explicit finite diagonal matrices. Indeed V is a permutation of a finite basis, so the series are periodic geometric sums. They are positive contractions. Their terms commute, so (5.12) is exact. Shifting the sums proves (5.13) exactly; the controlled flip gives (5.14) exactly.

At an active Q_j and a fixed current address x, its occupied slot is

$$\sum_{r=x-j}^{x-1} f_N(r) \pmod{k_N}.$$

If $1 \leq j \leq N$, this sum is positive because $f_N(x-j) = 1$, and is at most $k_N - 1$ by (7.4). Thus Y preserves Q_j for those j. It always commutes with C, and the uncontrolled remainder satisfies

$$\|[Y, VBV^*]\| \leq 2\sum_{j>N} 3\cdot 4^{-j} = 2\cdot 4^{-N}. \tag{7.7}$$

This verifies (5.15).

To couple V to the address clock, let T_0, X_0 be its ordinary cyclic rotation and first-bit flip and put

$$h_N(x) = \sum_{0\leq r<x} f_N(r), \quad D = \bigoplus_x R_0^{h_N(x)}, \quad T_a = D(T_0\otimes I)D^*, \quad X_a = D(X_0\otimes I)D^*. \tag{7.8}$$

The coordinate A_a and digit P_a commute with D, so all address clock defect bounds from Theorem 4.4 are unchanged. In the D^* gauge, V is ordinary increment except at the all-ones address, where it carries a fixed optional-register permutation. The recursive increment agrees away from that address. At all ones its two output addresses are distinct, with the same optional permutation. The guard is 4^{-N} there. Hence its operator defect is $\sqrt{2}\,4^{-N}$ and its flat defect is $\sqrt{2}\,4^{-N}2^{-N/2}$.

For $H(x)$ use 4^{-k} when the address has $k < N$ initial ones and then a zero, and use zero at all ones. Rotation removes one initial one, so (5.1) is exact, including its all-ones case. All commutations in (5.2) are exact, and off all ones

$$A_a(x+1) - A_a(x) = -1 + \tfrac{7}{4}\,4^{-k}.$$

At all ones the defect of the last equation is 4^{-N}. Its flat norm is $4^{-N}2^{-N/2}$. Let U be ordinary cyclic time increment and use ordinary clocks on fuel and time. Their defects have the same bounds as before.

It remains to verify the whole program on these same registers. For a nonzero address let e be the position of its first nonzero base-four digit. Set $R(x) = 2^{-e}$ and $L(x) = 1$ precisely when that digit is one. At the all-zero address set $R = L = 0$. Since N is even, rotation by

two bits removes one whole base-four digit. Thus (5.3)–(5.4) hold exactly, including at the all-zero address.

For fuel with k initial ones then a zero, put

$$Z(q,t) = 1_{\{t \equiv 0 \pmod{2^k}\}}.$$

On all-ones fuel use $k = N$. When the first fuel bit is zero, $Z = 1$. Otherwise the factor $I - P_t$ in (5.6) requires the first time bit to be zero and the rotated equation tests the remaining $k-1$ bits. On all-ones fuel, the recursive test is true precisely at the all-zero time word. This proves (5.6) in every case.

At a basis point (x, q, t) assign every workspace counter and control the deterministic configuration

$$\mathrm{Step}^{t \bmod 2^k}(\mathrm{Init}(x)). \tag{7.9}$$

These are diagonal entries on existing registers; they add no quantum tensor factor. For ordinary counter entries set $z_i = 1, s_i = 0$ when $r_i = 1$ and $z_i = 0, s_i = 2r_i$ otherwise; the idle $r_i = 0$ case uses $z_i = s_i = 0$. These are contractions and give (5.8) exactly. The next time point is initialized if its residue modulo 2^k is zero, and is otherwise one step of the old configuration. This is exactly (5.9)–(5.10), also at the cyclic time boundary. A displayed halt occurs in fewer than 2^N steps on an input $e < N/2$, so its address is flagged in (7.2). Hence (5.11) is exact.

All workspace and reset variables are diagonal in (x, q, t), and the optional register carries no workspace labels. Every specified remaining commutation therefore holds. The only nonzero defining defects are the clock affine and tail-flip defects, the guarded increments, the carry step and (7.7). Their operator norms are bounded by the theorem's constant. The three compulsory registers and the optional register have dimension 2^{3N+k_N}, and (7.3) gives (7.1). Finally each halting input eventually meets both cutoffs in K_N, while the address weights have a summable geometric tail. This proves $\theta_N \uparrow \theta_K$. □

8 Counting measurements on the actual state

The count must tolerate measurements that commute only approximately and a density matrix that is not flat. Its proof keeps the measurement branches together. Its last step is the usual random-access-code argument [57], using Holevo's information bound [37].

Theorem 8.1 (Triangular measurement count)**.** *Let m PVMs on $\mathbb{C}^d$ have two or three outcomes. Let ρ be a density matrix and $\xi = \sqrt{\rho}$. Assume each outcome has state commutator at most η, and outcomes of different measurements have weighted commutator at most η. For each measurement i, assume a unitary U_i has state commutator at most η, exchanges its own two binary outcomes, or its two active ternary outcomes while fixing the inactive one, and preserves every later measurement. Each of these intertwining equations is required in weighted norm with error at most η.*

If b measurements are binary and F_i is the active projection of a ternary measurement, then for an absolute constant C, with $e = Cm^4\eta^2 \le 1/2$,

$$\log d \ge b + \sum_{i \ \mathrm{ternary}} \mathrm{Tr}(\rho F_i) - Cm^3\eta - m\big[h_2(e) + e\big]. \tag{8.1}$$

Here h_2 is binary entropy, with $h_2(0) = 0$.

Proof. For a PVM (P_a) define left and right Hilbert–Schmidt isometries

$$L_P(Z) = (P_a Z)_a, \qquad R_P(Z) = (Z P_a)_a.$$

They differ on ξ by at most $\sqrt{3}\eta$. Every left multiplication commutes with every right multiplication. Moving a product of s left recording maps through ξ to right maps in reverse order therefore changes the recording vector by at most $s\sqrt{3}\eta$, by telescoping isometries.

Suppose a column of left operators $D = (D_c)_c$ has column operator norm at most L and $\|D\xi\|_2 \le \alpha$. It commutes with all right recording maps. After any prefix of s measurements, the preceding comparison gives

$$\left(\sum_{\mathbf{a},c} \|D_c P_{s,a_s} \cdots P_{1,a_1}\xi\|_2^2\right)^{1/2} \le \alpha + Ls\sqrt{3}\eta. \tag{8.2}$$

This estimates all branches together. In particular, no inverse of a small branch probability appears.

Measure in reverse index order. For the suffix starting at i, move U_i through its recorded measurements. It exchanges its own two labels and preserves later ones. A single intertwining column is a difference of two isometries, has column operator norm at most two and weighted norm $O(\eta)$ on ξ. Equation (8.2) bounds it after any preceding measurements by $O(m\eta)$. At most m intertwiners, followed by the initial state-invariance error, give total vector error $O(m^2\eta)$. Distances of purifications bound total variation of the classical records. Thus the suffix distribution differs from its own i-th outcome swap by at most $Cm^2\eta$.

For a binary outcome with conditional law $(p, 1-p)$, concavity gives $h_2(p) \ge 1 - |2p-1|$. The suffix swap variation is exactly the average of $|2p-1|$ over the later record. Hence its conditional entropy is at least one minus that variation. For ternary probabilities,

$$H(p_0, p_1, p_2) \ge (p_1 + p_2) - |p_1 - p_2|.$$

Indeed split off the active event and apply the binary inequality within it. The ternary swap variation is the corresponding average of $|p_1 - p_2|$. The chain rule now bounds the full record entropy below by b plus recorded active probabilities minus $Cm^3\eta$.

Those recorded active probabilities can be replaced by their values on ρ. For one PVM, $(P_a\xi)_a$ and $(\xi P_a)_a$ purify, respectively, its nonselectively measured state and ρ on the original system. Their distance is at most $\sqrt{3}\eta$, so those states differ by trace norm at most $C\eta$. If Φ_j are the nonselective measurement channels, expand

$$\Phi_s \cdots \Phi_1(\rho) - \rho = \sum_{j=1}^{s} \Phi_s \cdots \Phi_{j+1}\big(\Phi_j(\rho) - \rho\big).$$

Trace-norm contractivity gives an $O(m\eta)$ bound for every prefix. Thus all active expectations change by at most $Cm^2\eta$ in total, already covered by the preceding error term.

We next bound record entropy above by dimension. Fix one recorded variable. Move its measurement to the end of the actual sequence by adjacent swaps. An elementary swap column has components $[P_a, Q_b]$, weighted norm $O(\eta)$ and column operator norm at most two: it is the difference of the two recording isometries, with record factors reordered. Equation (8.2) bounds its norm after a prefix by $O(m\eta)$. After at most m swaps the total recording vectors differ by $O(m^2\eta)$.

In the reordered process this PVM is last. Remeasuring it on the final system recovers its recorded value with zero error, since equal projective measurements are adjacent. In the original process, apply the same failure projection to the difference of recording vectors. Its squared norm bounds the decoding error by $e = Cm^4\eta^2$.

Let Z be the classical record and Q_{fin} the final system, of dimension d. The decoder just described acts only on this system, not on the purification reference. Holevo's bound,

conditional entropy subadditivity and the error-indicator bound for at most three outcomes give

$$\begin{aligned} H(Z) &= I(Z : Q_{\mathrm{fin}}) + H(Z \mid Q_{\mathrm{fin}}) \\ &\le \log d + \sum_i H(Z_i \mid Q_{\mathrm{fin}}) \\ &\le \log d + m[h_2(e) + e]. \end{aligned}$$

For the last line, measuring the proposed decoder can only increase conditional entropy; the classical conditional entropy is at most $h_2(e) + e \log(3-1)$. Combining the upper and lower record bounds proves (8.1). □

9 The matched lower bound

Theorem 9.1. *Every weighted precision-δ model of $\mathcal{P}_K$ on $\mathbb{C}^d$ satisfies*

$$\log d \ge (3 + \theta_K) \log_4(1/\delta) - o(\log(1/\delta)). \tag{9.1}$$

The error is uniform over all states and dimensions.

9.1 The carry estimate

Identify $\mathcal{K}$ with binary sequences viewed as 2-adic integers, and let g_i add the nonnegative integer i. If two sequences first differ at digit j, their Cantor coordinates differ by between $\frac{1}{2}4^{-j}$ and 4^{-j}. Addition preserves that first differing digit. Consequently every g_i is 2-Lipschitz on $\mathcal{K}$, uniformly in i.

Let $h(a) = 4^{-k}$ if the sequence starts with exactly k ones followed by a zero, and put $h(1) = 0$. It is also 2-Lipschitz: equal first-zero positions give equal values; otherwise the smaller first-zero position j is the first differing digit, and the value difference is at most 4^{-j}. Its identities are

$$h = 1 - p + \tfrac{1}{4} p(h \circ s), \qquad g_1(a) = a - 1 + \tfrac{7}{4} h(a). \tag{9.2}$$

The latter follows by summing the coordinate change when k initial ones become zero and the next zero becomes one. The all-ones case follows by the limit $k \to \infty$. Lipschitz functions on $\mathcal{K}$ can be extended linearly across its gaps with the same constant.

Lemma 9.2 (Translations with polynomial cost)**.** *For an absolute constant C, every model of the address relations satisfies*

$$|H - h(\widehat{A}_a)|_\rho \le C\delta, \tag{9.3}$$

$$|V^{-i} \widehat{A}_a V^i - g_i(\widehat{A}_a)|_\rho \le C(i+1)\delta, \tag{9.4}$$

$$|[V^i J V^{-i}, P_{a,j}]|_\rho \le C(i+1)4^j\delta \quad (J \in \{F, Q, Y\}). \tag{9.5}$$

Proof. Let $p = p(\widehat{A}_a)$ and $E = H - h(\widehat{A}_a)$. Subtract the first identity of (9.2) from (5.1), retaining its error r. The exact subtraction is

$$\begin{aligned} E = &-(P_a - p) + \tfrac{1}{4} P_a T_a E T_a^* \\ &+ \tfrac{1}{4} P_a (T_a h(\widehat{A}_a) T_a^* - h(s(\widehat{A}_a))) \\ &+ \tfrac{1}{4} (P_a - p) h(s(\widehat{A}_a)) + r. \end{aligned}$$

The first, third, fourth and last terms have weighted size $O(\delta)$ by Lemmas 2.3 and 4.2. In the fourth term, moving the bounded right factor past ξ costs $O(\delta)$, since $h \circ s$ is Lipschitz with

fixed constant. The second term is bounded by $\frac{1}{4}|E|_\rho + O(\delta)$ using $\|E\| \le 1$. Absorb it to prove (9.3). The last equation of (5.2), together with spectral repair and unitary transport, gives

$$|V^* \widehat{A}_a V - g_1(\widehat{A}_a)|_\rho \le C_1 \delta.$$

Put $e_i = |V^{-i} \widehat{A}_a V^i - g_i(\widehat{A}_a)|_\rho$, so $e_0 = 0$. Since $g_i = g_{i-1} \circ g_1$ on $\mathcal{K}$ and functional calculus commutes with unitary conjugation, for $i \ge 1$

$$\begin{aligned} V^{-i} \widehat{A}_a V^i - g_i(\widehat{A}_a) = V^*(V^{-(i-1)} \widehat{A}_a V^{i-1} - g_{i-1}(\widehat{A}_a))V \\ + (g_{i-1}(V^* \widehat{A}_a V) - g_{i-1}(g_1(\widehat{A}_a))). \end{aligned}$$

The bracket in the first term is a difference of positive contractions, so it has norm at most one, and Lemma 2.2 bounds the first term by $e_{i-1} + \|[V, \xi]\|_2 \le e_{i-1} + \delta$. Both $V^* \widehat{A}_a V$ and $g_1(\widehat{A}_a)$ have spectrum in $\mathcal{K}$, where g_{i-1} is 2-Lipschitz. Hence (2.3), the bound $\|[g_1(\widehat{A}_a), \xi]\|_2 \le 2\|[\widehat{A}_a, \xi]\|_2$ and (4.4) bound the second term by $2(C_1 + 4C)\delta$. Thus $e_i \le e_{i-1} + C_2\delta$ with C_2 independent of i, and induction gives $e_i \le C_2 i\delta$. This proves (9.4). Each step pays one transport of the accumulated error and one functional-calculus loss with the uniform constant 2; no iterated factor four occurs.

Put $f_j = p \circ s^j$. Replace $P_{a,j}$ by $f_j(\widehat{A}_a)$, at cost $C4^j\delta$ by (4.5), and conjugate by V^{-i}, at transport cost $i\delta$. Then apply (2.3) to (9.4), with the Lipschitz constant $2 \cdot 4^j$ of f_j. The result is

$$|V^{-i} P_{a,j} V^i - f_j(g_i(\widehat{A}_a))|_\rho \le C(i+1)4^j\delta.$$

The composition $f_j \circ g_i$ has Lipschitz constant at most $4 \cdot 4^j$ on $\mathcal{K}$. Each $J \in \{F, Q, Y\}$ commutes with A_a to precision δ, and hence with $\widehat{A}_a$ to precision $C\delta$. Equation (2.4) bounds its commutator with this composition by $C4^j\delta$. Replace back at cost $C(i+1)4^j\delta$: for the self-adjoint difference Δ just estimated, $|[J, \Delta]|_\rho \le 2|\Delta|_\rho + \|\Delta\|\|[J, \xi]\|_2$. Conjugate by V^i, at transport cost $2i\delta$. This gives (9.5). □

9.2 Joint digits and conditional measurements

Lemma 9.3. *Fix $M \ge 1$. The $3M$ compulsory binary measurements are the raw digit PVMs $(P_{c,j}, I - P_{c,j})$, where $P_{c,j} = T_c^j P_c T_c^{-j}$ for $c \in \{a, f, t\}$ and $j < M$, with flips $X_{c,j} = T_c^j X_c T_c^{-j}$. The M conditional ternary measurements are $(I - F_i, F_i - Q_i', Q_i')$ for $i < M$, with flips Y_i, where Q_i' is the repaired data projection constructed in the proof. Index the conditional measurements first, then the address, fuel and time digits, each family in increasing digit order. All commutations between outcomes of different measurements, all flip equations required in this order, and all state commutations have error at most*

$$\eta = C(M+1)^6 4^M \delta. \tag{9.6}$$

Each Q_i' differs from Q_i by weighted size at most $C\eta$, and the conditional measurements are exact PVMs on the same state.

Proof. Two elementary rules are used throughout. First, if D is self-adjoint and Y is a contraction, then $DY\xi = D\xi Y + D[Y, \xi]$ gives

$$|[D, Y]|_\rho \le 2|D|_\rho + \|D\| \, \|[Y, \xi]\|_2. \tag{9.7}$$

Hence, if $X = X_0 + D_X$ and $Y = Y_0 + D_Y$ with self-adjoint corrections, the identity

$$[X, Y] = [D_X, Y] + [X_0, D_Y] + [X_0, Y_0]$$

shows that replacing both sides of a commutator by nearby operators adds the two replacement errors; it never multiplies them. Second, suppose every letter of a word w commutes to precision

δ with every letter of a word w', adjoint letters included. Moving letters one at a time and applying Lemma 2.2 to each interchange gives $|[w, w']|_\rho \le C|w|\,|w'|(|w| + |w'| + 1)\delta$. We call this the word-interchange bound.

The required estimates are the following.

(i) *One compulsory clock.* Lemma 4.2 gives the digit commutations, own-bit flips, later-digit preservation and state commutators within each clock, with error $C(M+1)^3 4^M \delta$. Its raw digit commutators already replace both digits by commuting spectral digits, at additive cost.

(ii) *Different compulsory clocks.* For $c \neq c'$, $P_{c,j}$ and $P_{c',k}$ are words of lengths $2j+1$ and $2k+1$ in two complete clock sets, which commute by hypothesis. The word-interchange bound gives $|[P_{c,j}, P_{c',k}]|_\rho \le C(M+1)^3\delta$. The same bound shows that $X_{c,j}$ preserves every $P_{c',k}$. These estimates must go through the raw conjugated digits. Comparing the spectral digits $p(s^j(\widehat{A}_c))$ and $p(s^k(\widehat{A}_{c'}))$ by two applications of (2.4) would multiply their Lipschitz constants. That costs order $4^{j+k}\delta$, which for j, k near M exceeds (9.6) by a factor of order $4^M/(M+1)^6$ and would defeat the choice (9.8) below. If spectral digits were measured instead, put $D_{c,j} = p(s^j(\widehat{A}_c)) - P_{c,j}$. It is self-adjoint, with norm at most one and $|D_{c,j}|_\rho \le C4^j\delta$ by (4.5). Then

$$[p(s^j(\widehat{A}_c)), p(s^k(\widehat{A}_{c'}))] = [D_{c,j}, p(s^k(\widehat{A}_{c'}))] + [P_{c,j}, D_{c',k}] + [P_{c,j}, P_{c',k}].$$

By (9.7), with $\|[p(s^k(\widehat{A}_{c'})), \xi]\|_2 \le C4^k\delta$ from Lemma 2.3 and $\|[P_{c,j}, \xi]\|_2 \le (2j+1)\delta$, its weighted norm is at most $C4^M\delta + C(M+1)^3\delta$. Thus either choice satisfies (9.6); we fix the raw digits.

(iii) *Conditional against conditional.* Use the affine-only part of the spectral repair on (B, Q, V) and (C, F, V). If $i \le j$,

$$[Q_i, F_j] = V^i [Q, F_{j-i}] V^{-i}.$$

Recover F_{j-i} from the rounded flag coordinate and apply (2.4) with $[Q, C] = 0$. The error is $O(4^{j-i}\delta)$ plus polynomial transport. If $i > j$, use $[F, B] = 0$ instead. The same argument with $[Q, B] = 0$ controls $[Q_i, Q_j]$, and with $[F, C] = 0$ controls $[F_i, F_j]$. Only one side is replaced by a coordinate function here, so there is a single depth loss.

(iv) *Conditional against address digits.* By (9.5) with $J \in \{F, Q\}$, the commutators of F_i and Q_i with $P_{a,j}$ are at most $C(M+1)4^M\delta$. With $J = Y$, the same bound shows that Y_i preserves every address digit.

(v) *Conditional against fuel and time digits.* The operators F_i, Q_i and Y_i are words of length $2i+1$ in V, V^* and one of F, Q, Y. Each of these letters commutes with every fuel and time clock variable by hypothesis. The digits $P_{f,k}$ and $P_{t,k}$ are words of length $2k+1$ in the fuel and time sets. The word-interchange bound gives commutators at most $C(M+1)^3\delta$, and shows that Y_i preserves every fuel and time digit. These bounds are again taken through the raw conjugated words. Writing F_i or Q_i as a digit of a conditional coordinate, and a fuel or time digit as a spectral digit, would stack two depth losses.

(vi) *Flips.* The own conditional flip is the V^i-translate of (5.14). For later flags, $[Y, C] = 0$ and digit recovery imply preservation by Y_i. For later data, use $[Y, VBV^*] = 0$ and digit recovery on the tail coordinate, with depth $j - i - 1$ when $j > i$. This again costs at most one factor 4^{j-i}. Conditional flips preserve all compulsory digits by (iv) and (v). A compulsory flip $X_{c,j}$ exchanges $P_{c,j}$ and preserves later digits of its own clock by (i), and every digit of a later clock family by (ii).

(vii) *State commutators.* State commutators of translated raw variables are at most $C(M+1)\delta$.

All the estimates just listed fit (9.6): cubic word-interchange costs and linear address costs are dominated by its sixth power. No factor 4^{i+j} is used.

We finally repair inclusion. For projections F, Q with $|(I-F)Q|_\rho \le \eta$ and state commutators at most η, put

$$A_0 = FQF, \quad L_0 = (I-F)QF, \quad Q' = 1_{[1/2,1]}(A_0) \le F.$$

The weighted error rule and adjoints give $|Q - A_0|_\rho \le 4\eta$ and $|L_0|_\rho \le 2\eta$. Since $A_0 - A_0^2 = L_0^* L_0$ and

$$|1_{[1/2,1]}(t) - t|^2 \le t(1-t) \quad (0 \le t \le 1),$$

we get $|Q' - A_0|_\rho \le 2\eta$ and thus $|Q' - Q|_\rho \le 6\eta$. This difference is self-adjoint, so its state commutator is at most twice its weighted norm. Replacing Q by Q' in a commutator or a flip equation therefore costs only $C\eta$, by Lemma 2.2. Apply this independently to the raw pairs F_i, Q_i, after absorbing their polynomial transport constants into (9.6). The outcomes $(I-F_i, F_i - Q_i', Q_i')$ are now exact PVMs. All previous estimates survive with a fixed factor. □

Proof of Theorem 9.1. Index the optional measurements first, then the three compulsory clock families, in increasing digit order within each family. Optional flips preserve all compulsory measurements. A compulsory flip needs to preserve only later compulsory measurements, not the earlier optional family. Thus this ordering meets the hypotheses of Theorem 8.1, with $m = 4M$ and $b = 3M$.

Put $t = \log_4(1/\delta)$ and, for sufficiently small δ, choose

$$M = \lfloor t - 30 \log_4(2+t) \rfloor . \tag{9.8}$$

Then $4^M \delta \le (2+t)^{-30}$, so (9.6) gives $\eta \le C(1+t)^{-24}$. The error terms of (8.1) are $o(1)$. Consequently

$$\log d \ge 3M + \sum_{i<M} \operatorname{Tr}(\rho F_i) - o(1).$$

Unitary state transport gives

$$|\operatorname{Tr}(\rho F_i) - \operatorname{Tr}(\rho F)| \le \|V^{-i} \rho V^i - \rho\|_1 \le 2i\delta.$$

Apply Theorem 6.2 to the same ρ. The result is

$$\log d \ge (3 + \theta_K - \omega(\delta))M - o(M).$$

Since $M = t - O(\log(t+2))$ and $\omega(\delta) \to 0$, this is (9.1). Every estimate used the actual model state. □

Remark 9.4 (Where the guards are used). Two structural facts about the proof follow from tracing where each relation enters.

(a) The guarded relations (4.7) are used only in the exact tracial limit. In the proof of Theorem 1.2 they enter only through Proposition 4.5, in the exact part of Theorem 6.2, and the approximate part of that theorem reaches the exact case by compactness, not by a quantitative estimate. The quantitative lower bound of this section uses instead the unguarded carry equation, the last relation of (5.2), whose defect in the finite models is 4^{-N} (Theorem 7.1). Consequently Proposition 4.6 is not used in the proof of Theorem 1.2, and its guard-removal loss, of order $\delta^{1/6}$ up to logarithms, is never paid.

(b) The guarded increment (4.7) for the address increment V is used by no proof. Section 6 applies Proposition 4.5 only to U, and V acts on the address digits only through the carry relations (5.1)–(5.2), in Lemma 9.2. That relation is therefore redundant. It is harmless: the finite models satisfy it with operator defect $\sqrt{2}\,4^{-N}$ (Theorem 7.1).

10 The game and its exponent

Corollary 10.1 (Matched source profiles)**.** *The presentation $\mathcal{P}_K$ satisfies*

$$\log W(4^{-N}) = (3+\theta_K)N + o(N) = \log H(4^{-N}). \tag{10.1}$$

One computable comparison sequence is $A_N = 3N + \lfloor N\theta_N \rfloor$, and both logarithmic profiles equal $A_N + o(N)$.

Proof. Theorem 9.1 is the lower bound for W, and $W \le H$. Theorem 7.1 is the matching upper bound for H at a fixed multiple of the requested precision. Replacing N by an even integer at a fixed greater distance absorbs that multiple. Such a change costs $O(1)$ depth, and the upper and lower limiting slopes are both $3+\theta_K$ since $\theta_N \to \theta_K$. Monotonicity interpolates intermediate depths. Finally $N\theta_N = N\theta_K + o(N)$, giving the comparison sequence. □

The Boolean correspondence theorem [26, Theorem 4.9] supplies both game directions with the actual marginal state kept. We do not reprove that general correspondence or its return theorem. To make the precision accounting for this construction transparent, we give the elementary matrix estimate responsible for its square-root extraction and explain how it applies here.

Lemma 10.2 (Consistency and the square root of the state)**.** *Let A, B be reflections and Ξ a rectangular amplitude matrix with $\rho_B = \Xi\Xi^*$. If $\|B\Xi - \Xi A\|_2 \le \eta$, then*

$$\|[B, \sqrt{\rho_B}]\|_2 \le \sqrt{2}\,\eta. \tag{10.2}$$

Proof. Set $Y = (\Xi + B\Xi A)/2$. Then $BY = YA$, so B commutes with YY^* and its square root. Also $\|\Xi - Y\|_2 \le \eta/2$. For any rectangular matrices M, N,

$$\|\sqrt{MM^*} - \sqrt{NN^*}\|_2 \le \sqrt{2}\,\|M - N\|_2. \tag{10.3}$$

Indeed apply the eigenbasis proof of (2.2) to $f(t) = |t|$ and the self-adjoint block matrices $\left(\begin{smallmatrix} 0 & M \\ M^* & 0 \end{smallmatrix}\right)$ and its N analogue, with $Z = I$. Their absolute values are block diagonal with square-root blocks. The right-hand squared norm is $2\|M - N\|_2^2$. Discard the second left-hand block to obtain (10.3). Since B commutes with $\sqrt{YY^*}$, its commutator with $\sqrt{\rho_B}$ is bounded by twice their difference, proving (10.2). □

Proposition 10.3 (The two transfers needed here)**.** *Let G be the fixed context-variable game supplied by Theorem 3.1 for $\mathcal{P}_K$. There are constants $c, C > 0$ such that a loss-ε strategy with local dimensions at most d gives a weighted source model of precision $C\sqrt{\varepsilon}$ in dimension at most cd. A flat source model of precision δ in dimension d gives a strategy of dimension at most cd and loss at most $C\delta^2$.*

Proof. The upper direction is precisely the explicit flat strategy in Theorem 3.1. For the lower direction apply the established Boolean correspondence, followed by that theorem's weighted soundness. The correspondence is applicable because every compiled context is an ordinary Boolean predicate with a nonempty set of valid answers and every tested incidence has fixed positive probability.

Here is also a direct check of the precision ingredients in its extraction, without invoking its more general return theorem. A mixed strategy may first be replaced by a pure component of at least the same winning probability. Compressing to Schmidt supports gives local dimension at most d. Alice's context POVMs have simultaneous fixed-size ancilla dilations, and Bob's binary contractions have reflection dilations on a common extra qubit. The dimensions grow by a fixed factor; the initial ancilla state is included in the marginal. Invalid Alice answers

can be reassigned to a fixed valid answer, since an invalid answer was already a certain loss. After these reductions, consistency at a tested incidence gives

$$\|(A \otimes I - I \otimes B)\psi\| \le 2\sqrt{\varepsilon/\pi_{\min}}.$$

Pull a fixed Bob context word across to Alice in reverse order. Every replacement costs this bound, and Alice's context observables commute and satisfy their truth table exactly. Thus each Bob context commutator and forbidden assignment polynomial has error $C\sqrt{\varepsilon}$ on the marginal. Lemma 10.2 gives the same order for its state commutators, with no additional square root. These are exactly the Boolean-model conditions used by the compiler. Hence the composed source precision is $C\sqrt{\varepsilon}$ on that state. □

Proof of Theorem 1.2. Apply Proposition 10.3 and Corollary 10.1. For $\varepsilon = 2^{-n}$, precision $C\sqrt{\varepsilon}$ has

$$\log_4(1/(C\sqrt{\varepsilon})) = n/4 + O(1).$$

The weighted lower profile therefore gives

$$\log D_G(2^{-n}) \ge \tfrac{3+\theta_K}{4} n - o(n).$$

The flat models at precision 4^{-N} give loss at most $C16^{-N}$ and logarithmic dimension $(3+\theta_K)N + o(N)$. Choosing an even $N = n/4 + O(1)$ at or above the required precision gives the matching upper bound. Fixed dimension multipliers add only $O(1)$ to the logarithm. Monotonicity also gives (1.4) for arbitrary $\varepsilon \downarrow 0$. These strategies have losses tending to zero, so $\omega_q(G) = 1$. A finite-dimensional perfect strategy would extract a source model at every positive precision in one fixed dimension, contradicting Theorem 9.1.

An oracle for K computes θ_K by finite summation and a geometric tail bound. Conversely, suppose membership below e is known and subtract those contributions. With $w_e = 4^{-(e+1)}$, the residual is at most $w_e/3$ when $e \notin K$, and at least w_e when $e \in K$. An approximation of error less than $w_e/6$, compared with threshold $2w_e/3$, distinguishes these cases. Induction computes K from a real oracle for θ_K. Thus $\theta_K \equiv_T K$, and the same holds for $\beta = (3+\theta_K)/4$. Lemma 2.4 now excludes a computable certified convergence rate. □

11 Every positive left-c.e. exponent

The halting game uses separated address cylinders so that its coefficient has an explicit Turing degree. This realization step treats arbitrary left-c.e. reals by reading the whole address. The clocks, reset, carry and conditional measurement estimates stay the same. Section 16 will replace monotone acceptance by stop-or-continue verification to reach all positive Δ^0_2 exponents.

The finite reader, its forcing lemma and its simultaneous finite models are given in Appendix C. We apply Lemmas 11.1 and 11.2 below.

11.1 Adding clocks and powering equations

Proposition 11.3. *For every $r \in \mathbb{N}$, integer $q \ge 1$ and left-c.e. $\theta \in [0,1]$, there is a fixed finite rational BCS game $G_{\theta,r,q}$ of quantum value one, with no finite-dimensional perfect strategy, and*

$$\log D_{G_{\theta,r,q}}(\varepsilon) = \frac{3+r+\theta}{4q} \log(1/\varepsilon) + o(\log(1/\varepsilon)). \tag{11.8}$$

Proof. Append r copies of the compulsory clock, requiring each complete new generator set to commute with all generators outside that clock. Tensoring the upper models adds rN to logarithmic dimension. For the lower bound, append the new compulsory digit families to the order in Lemma 9.3. Its word-interchange argument bounds all new cross terms by a polynomial in M times δ. The same error form $C_r(M+1)^6 4^M \delta$ holds, with $m = (4+r)M$ and $b = (3+r)M$. Neither the carry argument nor any other quantitative step uses the removed decoder. Applying Theorem 8.1 and Lemma 11.1, at depth (9.8), gives

$$\log W(\delta) \geq (3 + r + \theta) \log_4(1/\delta) - o(\log(1/\delta)). \tag{11.9}$$

Indeed the intermediate count is $(3+r)M + \sum_{i<M} \operatorname{Tr}(\rho F_i) - o(1)$, and state transport bounds the sum below by $M(\theta - \omega_\theta(\delta)) - O(M^2\delta)$.

Replace every defining polynomial equation $p = 0$, including specified cross-commutators, by $h_q(p) = 0$, where

$$h_q(p) = \begin{cases} (p^*p)^{q/2}, & q \text{ even}, \\ p(p^*p)^{(q-1)/2}, & q \text{ odd}. \end{cases} \tag{11.10}$$

Keep the exact variable domains and the weighted state-commutator conditions. Since $h_q(p)^* h_q(p) = (p^*p)^q$, scalar Jensen applied to the probability measure $B \mapsto \operatorname{Tr}(\rho 1_B(p^*p))$ gives

$$\|p\sqrt{\rho}\|_2 \leq \|h_q(p)\sqrt{\rho}\|_2^{1/q}. \tag{11.11}$$

No commutation between ρ and p^*p is needed. A powered precision-η model is consequently an original precision-$\eta^{1/q}$ model, also for its state commutators when $\eta \leq 1$. Conversely the operator-norm bounds of Lemma 11.2 give

$$\|h_q(p)\| = \|p\|^q \leq 3^q 4^{-qN},$$

and the flat state commutators vanish. All powered equations are fixed rational $*$-polynomials. Theorem 3.1 and the game transfer give original precision $(C\sqrt{\varepsilon})^{1/q}$ for the lower bound, and upper loss $C16^{-qN}$. Taking an even $N = \log(1/\varepsilon)/(4q) + O(1)$ above the required precision proves (11.8). All compiler dimension multipliers are constant for each fixed presentation. The upper losses tend to zero; the lower bound diverges, excluding a finite-dimensional perfect strategy. □

The left-c.e. dimension case of Theorem 1.1. Given a positive left-c.e. α, choose integers $q \geq 1$ and $m \geq 3$ with $m \leq 4q\alpha < m + 1$. Put $r = m - 3$ and $\theta = 4q\alpha - m \in [0, 1)$. This is left-c.e.; an increasing rational name is $\max(0, 4q\alpha_s - m)$. Proposition 11.3 gives exponent α. The integers q, m are chosen once for the given α. The usable-entanglement assertion is proved in Section 12. □

Corollary 11.4. *Uncomputable dimension exponents of fixed finite BCS games accumulate at zero. In fact powering the original halting presentation gives exponents $(3 + \theta_K)/(4q)$, each Turing-equivalent to K.*

Proof. Apply (11.11) and the operator upper estimate of Theorem 7.1 to $\mathcal{P}_K$. The invertible rational affine map preserves Turing degree, and these positive exponents tend to zero. □

12 Usable entanglement of the original state

Dimension alone does not give a distillation protocol. The separable state $d^{-1} \sum_x |xx\rangle\langle xx|$ has marginal I/d and no distillable entanglement. The measurement count of Section 8 must therefore be supplemented by agreement between the players' actual operators. This

section extends the binary triangular extraction of [26, Theorem 6.2] to measurements with an inactive outcome. Keeping the random targets preserves the mean number of active Bell pairs.

12.1 Triangular extraction with an inactive outcome

For operators D_ℓ and a vector ψ, write

$$\mathcal{R}_\psi((D_\ell)_\ell) = \left(\sum_\ell \|D_\ell\psi\|^2\right)^{1/2}.$$

Opposite-party tensor factors are suppressed below.

Lemma 12.1 (Extraction on a mixed state)**.** *Let ψ_{ABE} purify a finite-dimensional state ρ_{AB}. For $1 \le j \le m$, each party $D = A, B$ has an exact three-outcome PVM $(P^D_{j,\perp}, P^D_{j,0}, P^D_{j,1})$. Bob has a reflection U^B_j and Alice a unitary U^A_j. Put $F^D_j = P^D_{j,0} + P^D_{j,1}$. Suppose each of the following quantities is at most η:*

(i) For every j, $\mathcal{R}_\psi((P^A_{j,a} - P^B_{j,a})_a)$ and $\|(U^A_j - U^B_j)\psi\|$.

(ii) For each $s < j$, with $a, b \in \{\perp, 0, 1\}$, the two separate quantities $\mathcal{R}_\psi(([P^B_{j,a}, P^B_{s,b}])_{a,b})$ and $\mathcal{R}_\psi(([U^B_j, P^B_{s,b}])_b)$.

(iii) For every j, with $b \in \{0, 1\}$, $\mathcal{R}_\psi((U^B_j P^B_{j,b} U^B_j - P^B_{j,1-b})_b)$.

Set $\kappa = 7m^3\eta$ and $g(t) = (1+t)h_2(t/(1+t))$. If $\kappa \le 1/2$, then

$$E^{A\to B}_D(\rho_{AB}) \ge \sum_{j=1}^m \mathrm{Tr}(\rho_B F^B_j) - \tfrac{1}{2}m(m-1)\eta - 6m\kappa - 2g(\kappa). \tag{12.1}$$

A compulsory binary stage is included by $P^D_{j,\perp} = 0$. Different controls need not commute.

Proof. One stage. Alice samples a uniform bit a and sends it to Bob. Each party applies the local isometry

$$W^a_D = |0,0\rangle P^D_\perp + \sum_{b=0,1} |1,b\rangle (U_D)^{a\oplus b} P^D_b. \tag{12.2}$$

The first recorded bit is activity and the second is a coherent output qubit. Orthogonality of the three labels gives $(W^a_D)^* W^a_D = I$. A purification of the actual random operation is

$$\mathcal{A}\zeta = 2^{-1/2} \sum_a |a,a,a\rangle (W^a_A \otimes W^a_B)\zeta.$$

The third target copy records classical randomness in a new environment. It does not act on E. Compare this with the isometry

$$\mathcal{I}\zeta = \sum_a |a,a,a\rangle \left[2^{-1/2}|00\rangle_{\text{flags}}|00\rangle_{\text{qubits}} P^B_\perp \zeta + |11\rangle_{\text{flags}}|\Phi_2\rangle P^B_a \zeta\right]. \tag{12.3}$$

Its squared branch norms sum to $\|\zeta\|^2$.

For an arbitrary input ζ, let p, u, c be its projection-mirror aggregate, control-mirror norm and active-swap aggregate. Then

$$\|(\mathcal{A} - \mathcal{I})\zeta\| \le p + (u+c)/\sqrt{2}. \tag{12.4}$$

To check this, put $V_{a,\perp} = I$, $V_{a,b} = U^{a\oplus b}$ and use the exact identity

$$V^A_{a,r}P^A_r = P^B_r V^B_{a,r} + V^A_{a,r}(P^A_r - P^B_r) + P^B_r(V^A_{a,r} - V^B_{a,r}).$$

Multiplication by Bob's $V^B_{a,o}P^B_o$ kills the first term unless $o = r$. Equal inactive outcomes give $P^B_\perp$; equal active outcomes give P^B_a, apart from one swap defect when $r = 1 - a$. Summing squares over all records bounds the three error families by p, $u/\sqrt{2}$, and $c/\sqrt{2}$. For example $\sum_o \|V^B_{a,o}P^B_o z\|^2 = \|z\|^2$, and the control error survives only at $r = 1 - a$. Minkowski proves (12.4). No swap condition on Alice is used.

All histories. The ideal residual Kraus operators on either party are

$$R_{j,a,0} = P_{j,\perp}/\sqrt{2}, \qquad R_{j,a,1} = P_{j,a}.$$

They are self-adjoint and $\sum_{a,f} R^2_{j,a,f} = I$. For a history of j stages let K^B_h be the chronological product on Bob and Q^A_h the reversed product on Alice. Summing the innermost outcome repeatedly, without interchanging measurements, gives

$$\sum_h K^{B*}_h K^B_h = \sum_h K^B_h K^{B*}_h = I, \tag{12.5}$$

and both analogous identities for Q^A_h. The recursion

$$R^B K^B_h - Q^A_h R^A = R^B(K^B_h - Q^A_h) + Q^A_h(R^B - R^A)$$

therefore implies

$$\Delta_j := \mathcal{R}_\psi((K^B_h - Q^A_h)_{|h|=j}) \le j\eta.$$

For any family of Bob operators D_ℓ, opposite-party commutation and completeness give the useful estimate

$$\left(\sum_{h,\ell} \|D_\ell K^B_h \psi\|^2\right)^{1/2} \le \mathcal{R}_\psi((D_\ell)_\ell) + \left(\sum_\ell \|D_\ell\|^2\right)^{1/2} \Delta_j. \tag{12.6}$$

Indeed write $D_\ell K^B_h \psi = Q^A_h D_\ell \psi + D_\ell(K^B_h - Q^A_h)\psi$ and sum squares.

Expand a current operator's commutator with a history one factor at a time. Later outcomes disappear by completeness. At a prefix of length s, equation (12.6) bounds each commutator column by $(1 + Ls)\eta$. For the three current projections one can take $L = 3/2$, since commutators of two projections have norm at most $1/2$. For the current control take $L = \sqrt{3}$, since a unitary's commutator with a projection has norm at most one. The two scaled inactive copies count as one in the squared sum. Thus, at a stage preceded by j stages $(0 \le j < m)$, their history commutator aggregates are at most

$$(j + \tfrac{3}{4}j(j-1))\eta \le 3j^2\eta, \qquad (j + \tfrac{\sqrt{3}}{2}j(j-1))\eta \le 2j^2\eta.$$

For a mirror use $(G_A - G_B)K^B_h = K^B_h(G_A - G_B) - [G_B, K^B_h]$, rather than inserting a bipartite defect into (12.6). Its two aggregates obey

$$p_j \le (3j^2 + 1)\eta, \qquad u_j \le (2j^2 + 1)\eta.$$

The two Bob swap defects have operator root norm at most $2\sqrt{2}$, so (12.6) gives $c_j \le (1 + 2\sqrt{2}j)\eta$.

The ideal prefix is an orthogonal sum of histories, with normalized Bell or inactive output factors. Apply (12.4) to these histories together. The error of the stage preceded by j stages is at most

$$[(3 + \sqrt{2})j^2 + 2j + 1 + \sqrt{2}]\eta \le 7(j+1)^2\eta.$$

Telescoping the compositions, with actual isometries on the left and ideal prefixes on the right, bounds the final vector difference by $\sum_{j=0}^{m-1} 7(j+1)^2\eta \le 7m^3\eta = \kappa$. Completeness, rather than the number of histories, controls every sum.

Mean activation. Discarding the ideal stage's records gives Bob the channel $\Lambda_j(X) = \sum_r P^B_{j,r} X P^B_{j,r}$. The two recording vectors $(P^B_{j,r}\psi)_r$ and $(P^A_{j,r}\psi)_r$ are unit vectors at distance at most η. Their Bob marginals are $\Lambda_j(\rho_B)$ and ρ_B, respectively. Hence $\|\Lambda_j(\rho_B) - \rho_B\|_1 \le 2\eta$. Contractivity and telescoping bound a prefix disturbance by $2(j-1)\eta$. For a traceless self-adjoint X and $0 \le F \le I$, $|\operatorname{Tr}(FX)| \le \|X\|_1/2$. The ideal number N_{act} of active outputs therefore has

$$\mathbb{E}N_{\mathrm{act}} \ge \sum_j \operatorname{Tr}(\rho_B F^B_j) - \tfrac{1}{2}m(m-1)\eta. \tag{12.7}$$

The targets remain random throughout this argument. Selecting one fixed target could select a different activation density.

Hashing the actual output. Discard the original systems and targets and locally dephase the activity flags. Keep the output qubits coherent. The ideal output is

$$\sigma = \sum_f p_f |f,f\rangle\langle f,f| \otimes \bigotimes_{j=1}^{m} \begin{cases} \Phi_2, & f_j = 1, \\ |00\rangle\langle 00|, & f_j = 0. \end{cases}$$

Its coherent information is $I(A\rangle B)_\sigma = \mathbb{E}N_{\mathrm{act}}$: $S(B) = H(p) + \mathbb{E}N_{\mathrm{act}}$ and $S(AB) = H(p)$. The actual output ω has half trace distance at most κ from σ. Each local output space has dimension 4^m.

For states on dimension D at half trace distance t,

$$|S(\alpha) - S(\beta)| \le t \log D + g(t). \tag{12.8}$$

For a direct proof write $\alpha - \beta = \Delta_+ - \Delta_-$ with $\operatorname{Tr}\Delta_\pm = t$ and compare the two binary-mixture decompositions of $(\alpha + \Delta_-)/(1+t) = (\beta + \Delta_+)/(1+t)$. Concavity and the binary-mixture upper bound give (12.8) in both directions. Also g is increasing. Apply the bound to B and AB, of dimensions 4^m and 16^m. Then

$$I(A\rangle B)_\omega \ge \mathbb{E}N_{\mathrm{act}} - 6m\kappa - 2g(\kappa).$$

The hashing inequality [23, Theorem 10] makes this an achievable Alice-to-Bob rate. All preprocessing was one-way LOCC on ρ_{AB}, so composing it with hashing lower-bounds $E_D^{A\to B}(\rho_{AB})$. Together with (12.7) this proves the lemma. □

12.2 The bridge from an actual mixed strategy

Proposition 12.2. *For a loss-ε strategy of the compiled game of any presentation containing the clock, carry and conditional-clock relations of Section 5 and r extra compulsory clocks, powered by q (in particular G_K, the games of Proposition 11.3 and those of Theorem 15.7), the hypotheses of Lemma 12.1 hold on its original mixed state, after adjoining pure local ancillas, with $m = (4+r)M$ and*

$$\eta = C_{\mathcal{P},r,q}(M+1)^6 4^M \delta, \qquad \delta = (C_{\mathcal{P},r,q}\sqrt{\varepsilon})^{1/q}. \tag{12.9}$$

The active expectations are $(3+r)M + \sum_{i<M} \operatorname{Tr}(\rho_B F_i)$ for the reconstructed source model on the actual Bob marginal.

Proof. The actual-state interface is [26, Proposition 6.1]; the relevant precision ingredients are also visible in the compiler and Lemma 10.2. Here they must be applied before any

pure-component selection. Relabel illegal Alice answers to a valid answer and dilate the finitely many POVMs with pure local ancillas. On a purification ψ_{ABE} of the original state, fix for each variable one incident Alice context. Its reflection A_v and Bob's reflection B_v satisfy

$$\|(A_v - B_v)\psi\| \le 2\sqrt{\varepsilon/\pi_{\min}} =: \mu.$$

This statement is valid with the inaccessible E retained. Moving a Bob word to Alice reverses its order and costs at most its length times μ. For a fixed polynomial certificate, transfer its right factors, then its context relation, which vanishes exactly on Alice. This gives all Boolean residuals $O(\mu)$. Lemma 10.2, with the rectangular amplitude matrix for AE against B, gives the square-root-state commutators at the same scale. Theorem 3.1(a) reconstructs the powered source on this actual marginal. Equation (11.11) gives the unpowered precision δ in (12.9). Pure local ancillas do not change distillable entanglement.

The raw measurements and controls in Lemma 9.3 are word polynomials of length $O(M)$ in Bob's reflections. Alice uses their opposite word polynomials. These are exact projections and reflections for the raw translated observables: for example, the opposite of $T_B^j P_B T_B^{-j}$ is $T_A^{-j} P_A T_A^j$, with T_A the opposite word for T_B. Their mirrors cost $O(M\mu) \le O(M\delta)$. Applying Lemma 10.2 with parties exchanged gives their Alice square-root-state commutators as well.

Repair each raw pair F_i, Q_i locally on both sides by $Q_i' = 1_{[1/2,1]}(F_i Q_i F_i)$. Bob's error is $O(\eta)$ by Lemma 9.3. On Alice, transferring $(I - F_i)Q_i$ reverses its order, giving the adjoint of her inclusion defect. For any operator R,

$$\|R^*\sqrt{\rho}\|_2 \le \|R\sqrt{\rho}\|_2 + \|[R, \sqrt{\rho}]\|_2.$$

The raw mirrors and state commutators therefore give Alice's inclusion defect $O(\eta)$. The repair proof in Lemma 9.3, using $FQF - (FQF)^2 = ((I-F)QF)^*((I-F)QF)$, gives $\|(Q_i' - Q_i)\sqrt{\rho}\|_2 = O(\eta)$ on each party. The triangle inequality now yields $\|(Q_i'^A - Q_i'^B)\psi\| = O(\eta)$. Thus the ternary PVMs have actual mirrors at the same scale.

Use the reverse of the count order: the last compulsory family first, each digit family reversed, then the optional indices reversed. A flip's preservation of later count measurements is exactly preservation of earlier extraction measurements. Bob's reflection converts the conjugation-form estimates into commutators at cost $O(\eta)$ by right multiplication and its state commutator. Lemma 9.3, with the extra clocks as in Proposition 11.3, supplies every swap and triangular relation. A fixed enlargement of $C_{\mathcal{P},r,q}$ changes individual errors into the finite root aggregates. Only Q_i was repaired, so the activation flags are the original F_i. This proves every assertion. □

Theorem 12.3. *The games $G_{\theta,r,q}$ satisfy*

$$d^{\rightarrow}_{G_{\theta,r,q}}(\varepsilon) = \frac{3 + r + \theta}{4q} \log(1/\varepsilon) + o(\log(1/\varepsilon)). \tag{12.10}$$

The original halting game satisfies the same formula with coefficient $c_K = (3 + \theta_K)/4$.

Proof. Put $t = \log_4(1/\delta)$ and choose M by (9.8). Then $\eta = O(t^{-24})$, $\kappa = O(t^{-21})$, $m\kappa = O(t^{-20})$, and every extraction error in (12.1) tends to zero. On every original strategy state,

$$E_D^{A\to B}(\rho_{AB}) \ge (3 + r)M + \sum_{i<M} \mathrm{Tr}(\rho_B F_i) - o(1).$$

Lemma 11.1 and $|\mathrm{Tr}(\rho_B F_i) - \mathrm{Tr}(\rho_B F)| \le 2i\delta$ bound this below by $(3 + r + \theta - \omega_\theta(\delta))M - O(M^2\delta) - o(1)$. Since

$$M = \frac{1}{4q} \log(1/\varepsilon) - O(\log\log(1/\varepsilon)),$$

taking the infimum proves the lower bound. For the upper bound, $E_D^{A\to B}(\rho) \le \log d_B$ for any state on Bob dimension d_B, so $d_G^{\to}(\varepsilon) \le \log D_G(\varepsilon)$. Proposition 11.3 matches the lower coefficient. The explicit upper strategies are in fact maximally entangled, with only constant compiler overhead. For G_K, use Theorems 6.2 and 7.1 in the same argument; its retained raw measurements are identical. This also proves the left-c.e. usable-entanglement assertion of Theorem 1.1. □

These are equalities for the infimum demand. A particular successful strategy may carry arbitrary unused extra entanglement. The non-effective remainder comes from flag-density convergence; the extraction error above tends to zero absolutely.

Corollary 12.4. *Every positive left-c.e. real belongs to both coefficient sets $\mathcal{A}_{\log}$ and $\mathcal{A}_{\dim}$ of [26, Section 11.6].*

Proof. In that notation $\mathcal{A}_{\log}$ consists of the finite limits $d_G^{\to}(\varepsilon)/\log_2(1/\varepsilon)$ for finite value-one games, and $\mathcal{A}_{\dim}$ uses $\log_2 D_G(\varepsilon)$ instead. Theorem 1.1 uses these exact conventions. The present section proves the left-c.e. usable case; Section 16 strengthens the inclusion to every positive Δ_2^0 real, including the positive-computable-real inclusion conjectured in the discussion of [26, Problem 18.3(i)]. It is consistent with that paper's Theorem 11.14, which bounds $\mathcal{A}_{\log}$ by the reals left-c.e. relative to $0''$. □

13 Heat depends on the order of reset and communication

Use the thermodynamic model of *The heat price of precision* [24]: degenerate energy endpoints, asymptotic reversible preparation and erasure, and all working systems included in the reset. Work and bath heat coincide in its dimensionless units, with one unit equal to $k_B T \ln 2$. Write $K_{\text{ans}} = |\mathcal{A}||\mathcal{B}|$ for the product of the two answer-alphabet sizes; it is unrelated to the halting set K.

The local schedule allows no communication until both working devices are blank. The joint schedule permits quantum reunion after answer release and before reset. Theorem 2.1 of [24] gives, for every attainable loss,

$$2d_G^{\to}(\varepsilon) \le C_G(\varepsilon) \le w_{\text{local},G}(\varepsilon) \le C_G(\varepsilon) + 2\log K_{\text{ans}} \le 2\log D_G(\varepsilon) + 2\log K_{\text{ans}}, \tag{13.1}$$

where C_G is the infimum mutual information of a successful state. The same theorem gives $0 \le w_{\text{joint},G}(\varepsilon) \le 2\log K_{\text{ans}}$. These are cited results, not new erasure claims. The matching-sandwich argument of its Corollary 2.2 applies here with $o(\log(1/\varepsilon))$ remainders rather than its sharper hypothesis for the earlier coefficient family.

Corollary 13.1 (Order of operations)**.** *For every positive Δ_2^0 real α, the same game G_α of Theorem 1.1 satisfies, with $L = \log_2(1/\varepsilon)$,*

$$w_{\text{local},G_\alpha}(\varepsilon) = 2\alpha L + o(L), \qquad 0 \le w_{\text{joint},G_\alpha}(\varepsilon) \le 2\log_2 K_{\text{ans}}. \tag{13.2}$$

For the original halting game the local coefficient is $(3+\theta_K)/2$, Turing-equivalent to K, while joint-reset heat stays bounded.

Proof. Substitute $d_{G_\alpha}^{\to}(\varepsilon) = \alpha L + o(L)$ and $\log D_{G_\alpha}(\varepsilon) = \alpha L + o(L)$ into (13.1). The fixed answer term is $O(1)$. The joint bound is the cited theorem. The halting coefficient is $2c_K$, an invertible rational affine image of θ_K. □

Thus one fixed device realizes any prescribed positive Δ_2^0 local-heat coefficient (replace α by half that coefficient), while its joint-reset cost stays bounded. The order of communication and reset alone selects between these regimes. These are limiting coefficients; whether the finite-precision heat function is computable is not decided here.

13.1 Recovery before the blank deadline

Use exactly the block model of [24, Section 4, Theorem 4.1 and Lemma 4.2]. Fix a finite strategy and loss first, then process independent copies after answer release, allowing coherent local compression. Per play, recovery transfers at most q qubits and c physically moved fully dephased bits, up to sublinear block corrections, with an Alice-to-Bob route available. The blank deadline follows recovery. Degenerate endpoints, independent local baths, thermal answer ports and reversible preparation/erasure are retained; private workspace, communication hardware, copied messages and failed branches all enter reset. There is no uncharged catalyst. Here q is a communication budget, unrelated to a residual power.

Let $E_G(\varepsilon)$ be the infimum of entanglement entropy of successful *pure* strategies, not entanglement cost, and set $B_G = \log_2(|\mathcal{X}||\mathcal{Y}|K_{\mathrm{ans}})$. Write $x_+ = \max\{x, 0\}$ and

$$\mathcal{F}(e; q, c) = (e - q)_+ + (e - q - c)_+.$$

The cited theorem gives the general bound

$$\mathcal{F}(d_G^{\rightarrow}(\varepsilon); q, c) \le w_{q,c,G}(\varepsilon) \le \mathcal{F}(E_G(\varepsilon); q, c) + 4B_G. \tag{13.3}$$

Its proof supplies an arbitrary-original-state lower bound; no new recovery inequality is assumed here.

Corollary 13.2 (Programmable quantum-recovery threshold)**.** *For each matched game of Theorem 1.1, in the stated block model,*

$$w_{q,c,G_\alpha}(\varepsilon) = \mathcal{F}(\alpha L; q, c) + o(L), \qquad L = \log_2(1/\varepsilon), \tag{13.4}$$

uniformly over $q, c \ge 0$. In particular, for $q/L \to \kappa$ and $c/L \to \chi$ with $\kappa, \chi \ge 0$, the heat coefficient is

$$(\alpha - \kappa)_+ + (\alpha - \kappa - \chi)_+. \tag{13.5}$$

The critical quantum-recovery rate that removes the leading heat term is α, even with unlimited classical recovery. It can be any positive Δ_2^0 real.

Proof. For fixed measurements the score is affine in the state, so a successful mixed strategy has a successful pure component in the same local spaces. Pure-state distillable entanglement equals its marginal entropy. Consequently

$$d_G^{\rightarrow}(\varepsilon) \le E_G(\varepsilon) \le \log_2 D_G(\varepsilon).$$

Theorem 1.1 squeezes both endpoints to $\alpha L + o(L)$. The function $\mathcal{F}$ is nondecreasing and 2-Lipschitz in its first argument, with constant independent of q, c. Apply (13.3) to obtain (13.4); homogeneity and continuity give (13.5). Since $\mathcal{F}(e; q, c) \ge (e - q)_+$, a rate $\kappa < \alpha$ cannot remove the leading cost for any classical budget, whereas $\kappa \ge \alpha$ does so even with $c = 0$. □

The imported inequality applies to these games; with the present $o(L)$ demand laws it gives $o(L)$ heat at the critical rate, which need not be bounded.

Proposition 13.3 (Undecidable critical rates under a two-value promise)**.** *There is a computably generated family of finite rational matched games whose critical recovery rates lie in $\{15/16, 1\}$, on which no uniform estimator has a guaranteed fixed error less than $1/32$. On this promised family neither strict side of $31/32$ has a uniform semidecision procedure.*

Proof. Given a machine index e, the left-c.e. reader of Section 11, with no extra clock or residual power, uniformly realizes

$$\theta_e = \tfrac{3}{4} + \tfrac{1}{4}1_{\{e\in K\}}, \qquad \alpha_e = \frac{3+\theta_e}{4} = \tfrac{15}{16} + \tfrac{1}{16}1_{\{e\in K\}}.$$

Its increasing rational approximation simply simulates machine e for longer finite times. An estimate within any fixed error less than $1/32$, compared with $31/32$, decides halting.

Include also a tagged complementary family from Section 15: let $A_e = \{0\}$ if $e \in K$ and $A_e = \varnothing$ otherwise. Its uniform c.e. program enumerates zero exactly on halting. Theorem 15.7 gives the matched rate

$$\alpha'_e = \frac{4-\theta_{A_e}}{4} = 1 - \tfrac{1}{16}1_{\{e\in K\}}.$$

Corollary 13.2's proof uses only matched demands, so these are the respective critical rates. A semidecision for the low side on the union, restricted to the reader family, would enumerate nonhalting indices. A semidecision for the high side, restricted to the complementary family, would do the same. Both are impossible. Each individual rate is a rational; the obstruction is choosing it uniformly from a finite game description. □

13.2 An arithmetic upper bound for heat coefficients

Proposition 13.4. *For finite rational games of quantum value one, $C_G(\varepsilon)$ is uniformly right-c.e. at positive rational ε. Every existing finite local-reset heat coefficient in the stated thermodynamic model is left-c.e. relative to $0'$.*

Proof. At a fixed local dimension d, density matrices and POVMs with loss at most ε form a rational semialgebraic feasible set, after splitting complex entries into real and imaginary parts as in Lemma 2.4. Its algebraic points are dense even when it has empty interior. Indeed, intersect it with any rational open ball meeting it. This is a nonempty rational semialgebraic set, so real-closed-field transfer gives a point with real algebraic coordinates. Quantifier elimination decides nonemptiness and can supply such a point. Enumerating dimensions and rational balls therefore enumerates a dense algebraic feasible family in each dimension.

Spectra of algebraic density matrices, and hence their entropies and mutual informations, are computable. Continuity at fixed dimension, including $-t\log_2 t = 0$ at zero, shows that the infimum over this family is exactly $C_G(\varepsilon)$. Running minima of certified rational upper approximations to the enumerated mutual informations decrease to that infimum. A first feasible algebraic strategy can be found by dimension search under the value-one promise. Thus the approximation is uniform and total on that promise; no bound on the required dimension is assumed.

Choose total computable rationals $c_{n,s} \downarrow C_G(2^{-n})$ for $n \geq 1$. If the finite local-heat coefficient is h, the sandwich $C_G \leq w_{\mathrm{local},G} \leq C_G + 2\log K_{\mathrm{ans}}$ implies $C_G(2^{-n})/n \to h$. For rational a,

$$a < h \quad \Longleftrightarrow \quad \exists r \in \mathbb{Q}\ \exists N \quad r > a \ \text{ and } \ \forall n \geq N\ \forall s \quad c_{n,s} \geq rn. \tag{13.6}$$

If $a < h$, choose $a < r < h$ and use convergence. Conversely the right side gives $C_G(2^{-n}) \geq rn$ on that tail, hence $h \geq r > a$. The rational comparisons are decidable, so the lower cut is Σ^0_2, equivalently c.e. in $0'$. □

If $\mathcal{H}^+_{\mathrm{local}}$ denotes the positive finite local-heat coefficients of rational value-one games in this model, then

$$\{\alpha > 0 : \alpha \in \Delta^0_2\} \ \subseteq \ \mathcal{H}^+_{\mathrm{local}} \ \subseteq \ \{h > 0 : h \text{ is left-c.e. in } 0'\}. \tag{13.7}$$

For the first inclusion choose the matched dimension coefficient to be half the desired heat coefficient and use Corollary 13.1. The second is Proposition 13.4. Equality for the full heat spectrum remains open. The right-c.e. assertion concerns finite-precision mutual information, not distillable entanglement.

14 The exact-trace barrier and matrix approximation

Let $\mathcal{T}(\mathcal{P})$ be the exact tracial states of a finite bounded rational presentation, and let $\mathcal{T}_{\mathrm{mat}}(\mathcal{P})$ be those obtained as moment limits of flat finite-dimensional models with defects tending to zero. Both spaces are viewed in the compact space of bounded word moments. A diagonal choice of approximants shows that $\mathcal{T}_{\mathrm{mat}}$ is closed. In this section the relevant space is assumed nonempty and F is a fixed rational self-adjoint polynomial. These results explain why the construction of Sections 11–12 reaches only left-c.e. targets; they do not classify all games.

14.1 The exact-trace barrier

Proposition 14.1. *Suppose every exact trace has $\tau(F) \geq \theta$, and some sequence of flat approximate models with defects tending to zero has F expectations tending to θ. Then*

$$\theta = \min_{\tau \in \mathcal{T}(\mathcal{P})} \tau(F) \tag{14.1}$$

is left-c.e. The matched fixed-observable mechanism used here realizes exactly the left-c.e. targets in $[0, 1]$.

Proof. A moment limit of the flat models is tracial and has $\tau(p^*p) = 0$ for every defining relation p. The GNS argument in Theorem 6.2 makes it an exact model. Its F expectation is θ, proving (14.1).

Here is an effective lower hierarchy for that minimum. At degree $2j$ use complex word moments, impose normalization, adjoints, cyclic invariance, a positive moment matrix, the domain localizing constraints, and every relation insertion visible at that degree. Add $|y_w| \leq 1$ for contraction words. These bounded constraints express all declared domains. Alternatively use (3.1) for positive contractions and $P = (I + x)/2$, $x = x^* = x^{-1}$, for projections. The exact encodings are finite and rational and preserve tracial models and the substituted observable's trace. The finite feasible sets are compact rational semialgebraic sets in real and imaginary coordinates. Their lower minima ℓ_j increase. A diagonal limit of minimizers is a positive tracial functional with all relations and bounded domains, so its GNS tuple is an exact model. Consequently $\ell_j \uparrow \theta$.

Real quantifier elimination decides whether $\ell_j > r$ for a rational r. Enumerate these certified lower rationals and take running maxima, starting with a rational norm lower bound for F. They converge to θ. This proves left semicomputability without an assumption about numerical SDP duality. This is the bounded tracial moment method [12, 42]; it optimizes over all exact traces. It does not identify matrix-approximable traces, just as general NPA convergence concerns commuting models [56].

Lemma 11.1 and Lemma 11.2 give the converse realization for every left-c.e. $\theta \in [0, 1]$. Fixed positive rational affine changes, extra clocks and residual powering preserve this class. □

Complementing F changes a minimum to one minus a maximum. It therefore does not reverse this one-sided forcing argument. If all exact traces gave one common value, the lower hierarchy for F and for $-F$ would make that value computable. Similarly, an irreversible acceptance program presented with a nonmonotone approximation keeps earlier accepted cylinders: their union is effectively open and its measure remains left-c.e.

14.2 Weighted limits are matrix-approximable

The next compression uses Connes' spectral-cut inequality [20], in the form Vidick uses for almost synchronous correlations [72]. Its role here is to identify possible limiting traces. Its dimension cost is never charged to the actual-state measurement bound.

Lemma 14.2 (Moment compression)**.** *Let contraction generators $X_i \in M_d(\mathbb{C})$ lie in the declared domains of Definition 2.1, and suppose $\max_i \|[X_i, \sqrt{\rho}]\|_2 \le \eta$. For each integer $l \ge 1$ and $\zeta > 0$ there is a finite-dimensional tuple Y in the same domains such that every $*$-word w of length at most l satisfies*

$$|\mathrm{tr}_D(w(Y)) - \mathrm{Tr}(\rho w(X))| \le 2l\sqrt{2\eta} + \zeta. \tag{14.2}$$

No bound on D is asserted.

Proof. Put $P_t = 1_{(t,\infty)}(\rho)$, $0 \le t \le 1$. Spectral integration gives $\rho = \int_0^1 P_t\,dt$ and $\int_0^1 \mathrm{rank}\,P_t\,dt = 1$. In an eigenbasis of ρ,

$$\begin{aligned}\int_0^1 \|[X_i, P_t]\|_2^2\,dt &= \sum_{a,b} |(X_i)_{ab}|^2 |\lambda_a - \lambda_b| \\ &\le \|[X_i, \sqrt{\rho}]\|_2 \left(2\,\mathrm{Tr}(\rho X_i X_i^*) + 2\,\mathrm{Tr}(\rho X_i^* X_i)\right)^{1/2} \le 2\eta.\end{aligned}$$

The middle inequality factors $|\lambda_a - \lambda_b|$ into the difference and sum of their square roots and applies Cauchy–Schwarz.

On $\mathrm{ran}\,P_t$ compress to $B_i = P_t X_i P_t$. Insert P_t between consecutive letters of a word of length k. Each of the $k-1$ trace errors is at most $\sqrt{\mathrm{rank}\,P_t}\,\|[X_i, P_t]\|_2$. Integrating and applying Cauchy–Schwarz bounds the total by $(k-1)\sqrt{2\eta}$.

Restore the domains inside each compressed space. Positive contractions need no repair. Round a projection at $1/2$, a reflection by spectral sign, and a unitary by polar decomposition, completing the partial isometry on its kernel. Each resulting Z_i satisfies

$$\|Z_i - B_i\|_2^2 \le \|[X_i, P_t]\|_2^2.$$

For projections use $\min(s^2, (1-s)^2) \le s(1-s)$ and $\mathrm{Tr}(B_i - B_i^2) = \|(I - P_t)X_i P_t\|_2^2$. For unitaries use $(1-s)^2 \le 1 - s^2$ and $\mathrm{Tr}(I - B_i^* B_i) = \|(I - P_t)X_i P_t\|_2^2$; spectral sign is the self-adjoint case. Telescoping the k repaired letters adds at most $k\sqrt{2\eta}$ after integration.

There are finitely many spectral cuts. The integrated repaired moments form a convex combination of normalized traces of the compressed tuples, with weights $a_h = (\mathrm{rank}\,P_h)\,\Delta t_h$ summing to one. If their dimensions are r_h, choose integer multiplicities n_h with $n_h r_h / \sum_i n_i r_i$ approximating a_h in ℓ^1 within ζ. The direct sum realizes those weights and changes every contraction-word moment by at most ζ. Combining both errors proves (14.2). □

Corollary 14.3. *Moment limits of weighted approximate models are exactly $\mathcal{T}_{\mathrm{mat}}(\mathcal{P})$. In particular, if $\lambda = \min_{\mathcal{T}_{\mathrm{mat}}} \tau(F)$, then uniformly over weighted precision-δ models,*

$$\mathrm{Tr}(\rho F) \ge \lambda - \omega(\delta), \qquad \omega(\delta) \to 0. \tag{14.3}$$

Proof. In a convergent weighted sequence pass to a subsequence with commutator errors $\eta_j \le j^{-4}$. Apply Lemma 14.2 to length j with $\zeta_j \to 0$. The resulting flat tuples have the same limiting moments. Each defining p^*p moment tends to zero, so their flat relation defects tend to zero. The reverse inclusion takes $\rho = I/d$. Compactness and the closedness of $\mathcal{T}_{\mathrm{mat}}$ give attainment of λ. A sequence violating (14.3) would have a weighted moment limit below λ, a contradiction. □

14.3 Arithmetic and the necessary trace gap

Proposition 14.4. *Set*

$$m_n = \inf\{\operatorname{tr}_d F(X) : d \geq 1,\ X \text{ a flat precision-}2^{-n} \text{ model}\}.$$

Then $\lambda = \sup_n m_n$ *is left-c.e. relative to* $0'$*. If* $a = \min_{\mathcal{T}} \tau(F)$ *and* λ *is not left-c.e., then* $a < \lambda$*.*

Proof. The feasible sets shrink, so m_n increases. Approximants to any trace in $\mathcal{T}_{\mathrm{mat}}$ show $m_n \leq \lambda$. Conversely choose models with costs at most $m_n + 1/n$. A moment subsequence gives a trace in $\mathcal{T}_{\mathrm{mat}}$ of cost $\lim_n m_n$, proving equality.

For fixed n, d, all domains, squared flat residuals and the cost are rational semialgebraic data; real quantifier elimination decides whether there is a feasible tuple of cost below a rational r. Enumerating d makes m_n uniformly right-c.e. More explicitly,

$$\begin{aligned} s < \lambda \iff \exists n\ \exists r \in \mathbb{Q} \quad & s < r \text{ and} \\ & \text{every dimension lacks a precision-}2^{-n} \text{ model of cost } < r. \end{aligned} \tag{14.4}$$

The last finite-dimensional predicate is decidable. The displayed lower cut is therefore Σ^0_2, or c.e. in $0'$. Proposition 14.1's hierarchy proves that a is left-c.e., and $a \leq \lambda$. A non-left-c.e. λ must consequently satisfy $a < \lambda$.

For any rational r between them there is an exact trace of cost below r, but some fixed positive precision admits no flat finite-dimensional model of cost at most r. Otherwise a sequence of such models would limit to an approximable trace below λ. This is the necessary trace gap. □

If all approximable traces have one common value, apply the same arithmetic argument to $-F$. Both cuts are then c.e. in $0'$, so the value is Δ^0_2. An arbitrary minimum above is only asserted to be left-c.e. in $0'$; that statement alone does not make it Δ^0_2.

14.4 Sharpness of the approximable-minimum range

Theorem 14.5 (All approximable projection minima)**.** *The values* $\min_{\tau \in \mathcal{T}_{\mathrm{mat}}(\mathcal{P})} \tau(F)$*, over finite bounded rational presentations with nonempty* $\mathcal{T}_{\mathrm{mat}}(\mathcal{P})$ *and a declared projection* F*, are exactly the reals in* $[0,1]$ *left-c.e. relative to* $0'$*. More precisely, for every total computable sequence* $a_s \in \mathbb{Q} \cap [0,1]$ *there is such a presentation* $\mathcal{Q}[a]$ *with*

$$\min_{\mathcal{T}_{\mathrm{mat}}(\mathcal{Q}[a])} \tau(F) = 1 - \limsup_s a_s. \tag{14.5}$$

Proof. Use the dynamic presentation of Section 16.1, but remove R_*, the stop counter with its zero test and witness, the stop branch with its scalar verifier, the REJECT control and their equations. Section 16.1 and Lemma 16.2 do not use Theorem 14.5 or Corollary 14.6, so this forward reference is not circular. At every depth the program continues: compute the packed list for $C(j,w)$ in (16.3), execute exactly one LOAD, test the entire list, READ the next bit and increment j. Keep every other scalar, clock, register and data relation, deleting R_* only from the commutation lists. No decoded-region restriction is imposed. The list computation terminates without searching for a perfect witness. This is one finite rational instruction table.

Put $u = \limsup_s a_s$. If $v(x) > u$, it has a prefix w of some length J such that $\ell(w) > u$ and $a_s < \ell(w)$ for all $s \geq J$. The gate $C(J,w)$ is false. Its complete formal execution has a

finite runtime independent of its data. If $\tau(B_w(I-F))>0$, choose 2^k above that runtime. The retained fuel/time fairness gives

$$\tau(B_wC_kD_k(I-F))=2^{-(2k+1)}\tau(B_w(I-F))>0.$$

Lemma 16.2 applies: its proof uses the finite initialized interval and one LOAD, not convergence or a stop counter. Every shifted variable reduces this corner; the same tuple is tested throughout; its positive packed projection corner would give flat approximants of the false gate. This is impossible. Fuel/time flips establish only the displayed scalar mass identity, not transport of dynamic data. Every such bad cylinder has zero unflagged mass. Cover the compact set $\{v\ge u+\zeta\}$ by finitely many bad cylinders and refine to a common depth. Fair address digits then bound the remaining unflagged mass by $u+\zeta$ plus a vanishing dyadic endpoint error. Letting the error and ζ decrease proves $\tau(F)\ge 1-u$.

For the upper bound fix a dyadic $0<b<u$ and an even N beyond its denominator depth. Set $I-F=1_{\{v_N<b\}}$. Every prefix of each unflagged finite address has $\ell(w)\le v_N<b$, including after all N bits have been read and the reader supplies zeros. Arbitrarily late a_s exceed b, so every continuation gate on that address is true. In the finitely many address/fuel/time histories at this fixed N, only finitely many gates reach LOAD. Choose an exact witness for each whole packed list, even if reset truncates its test, and tensor these finitely many factors. There is no bound on this data dimension.

Assign the actual initialized scalar and register histories blockwise at time $t \bmod 2^k$, with $k=N$ for all-ones fuel. Load the assigned tuple on each outgoing LOAD edge, keep it until the next LOAD or reset, use $(0,I,I)$ from a reset until the first LOAD after it, and use constant default data $(0,I,I)$ on flagged addresses. The contraction registers have their blockwise two-unitary representations without a new factor. Thus all new equations and declared coordinate/workspace commutations are exact: ordinary steps hold data, LOAD permits replacement, reset removes the holding guard, and cyclic wrap is a reset because 2^k divides 2^N. Every completed unflagged CHECK uses a true witness. The address dressing commutes with the new assignments, so the old residuals remain at most $3\cdot 4^{-N}$. The flag mean is exactly $1-b$.

Let $N\to\infty$ at fixed b and take a moment subsequence. It gives a trace in $\mathcal{T}_{\mathrm{mat}}$ with cost $1-b$. Let $b\uparrow u$ and use compactness and closedness to attain $1-u$. For $u=0$, fully flagged histories with constant default data give cost one directly; the argument with $b<1$ includes $u=1$. This proves (14.5) and nonemptiness.

Downey–Jockusch–Schupp [27, Theorem 5.1(ii)–(iii)] characterize reals in $[0,1]$ with Σ^0_2 lower cuts as the liminf of computable rational sequences. These are precisely the reals left-c.e. in $0'$. Clipping the sequence to $[0,1]$ preserves its liminf. For $\lambda=\liminf_s c_s$, apply (14.5) to $a_s=1-c_s$. Proposition 14.4 supplies the reverse inclusion. □

Corollary 14.6 (An unavoidable witness-budget obstruction)**.** *Let $\mathcal{P}$ be a finite bounded rational presentation containing the relations used in Lemma 9.3, with flag F and $\lambda=\min_{\mathcal{T}_{\mathrm{mat}}(\mathcal{P})}\tau(F)$, for example $\mathcal{P}=\mathcal{Q}[a]$. If λ is not Δ^0_2, there are no flat models at every sufficiently large even N with operator residuals $O(4^{-N})$ and*

$$\log_2 d_N=(3+\lambda)N+o(N). \tag{14.6}$$

Proof. Corollary 14.3 gives $\mathrm{Tr}(\rho F)\ge\lambda-\omega(\delta)$ on the actual precision-δ state, with $\omega(\delta)\to 0$. Set $t=\log_4(1/\delta)$ and $M=\lfloor t-30\log_4(2+t)\rfloor$. Lemma 9.3 bounds the retained measurements' joint error by $C(M+1)^6 4^M\delta=O(t^{-24})$. The raw clock count of Theorem 8.1 and flag transport give

$$\log_2 d\ge 3M+\sum_{i<M}\mathrm{Tr}(\rho F_i)-o(M)\ge(3+\lambda-\omega(\delta))M-O(M^2\delta)-o(M).$$

The new data do not enter these measurements. The compiler gives $\delta = C\sqrt{\varepsilon}$, hence lower game exponent $(3+\lambda)/4$. Hypothetical models (14.6) compile to loss $O(16^{-N})$ with constant dimension overhead. Neighboring even depths have bounded logarithmic-precision gaps, so they give the matching upper exponent at every small loss. The rational finite game has value one. Lemma 2.4 and the limit argument of Corollary 16.7 would make $(3+\lambda)/4$, hence λ, Δ^0_2, a contradiction. □

For example, $\lambda = \sum_{e \in K^{0'}} 4^{-(e+1)}$ is left-c.e. in $0'$ and has degree $0''$: the tail after each digit is at most one third of its weight, so successive accurate approximations recover membership in $K^{0'}$, and that set computes the sum. It is not Δ^0_2. The theorem realizes this minimum, but the corollary forbids its matched budget. This obstruction concerns all sufficiently large even depths; it asserts neither a fixed positive overhead gap nor a bound on isolated subsequences.

Projection charges in $[0,1]$	Exact arithmetic range
Exact-trace minima	Left-c.e.
Matrix-approximable minima	Left-c.e. relative to $0'$
Charges with a matched three-clock budget	Δ^0_2

The first row is Proposition 14.1 and the reader realization; the second is Theorem 14.5. In the third, "matched" means the retained clock lower count and the all-scale bound (14.6). Corollary 14.6 gives necessity, and Theorem 16.1, with $\theta = 1 - \lambda$, realizes every Δ^0_2 charge. Feasibility and feasibility within that budget are distinct.

Context: why a uniform lower hierarchy is different. Use the finite-dimensional *synchronous* value in Mousavi–Nezhadi–Yuen's hardness reduction [54, Lemma 6.7], which directly makes synchronous value one Π^0_2-hard. Place the game projection relations under a central projection p and set $F = I - p$. The scalar branch $p = 0$ always exists. The synchronous/tracial correspondence [60, 41] shows that synchronous value one gives approximable traces with $p = I$, hence minimum zero. Otherwise a positive-trace p corner would give finite synchronous strategies approaching one: compress to the representing projections, normalize by their positive limiting trace, and repair each finite PVM. For that last step, successively compress and round its first outcomes into the complement of earlier repaired projections, then use the remaining complement as the last outcome. All errors tend to zero for a fixed finite family. Thus a uniform ordinary lower hierarchy would enumerate the Σ^0_2-complete complement of synchronous value one. This contextual reduction produces only the individual minima zero and one. It uses the synchronous-value hardness itself, without a conversion from unrestricted strategies that would require extra assumptions on diagonal question weights.

Towards the next construction. The exact-trace barrier leaves open a different matched mechanism: a non-left-c.e. minimum over matrix-approximable traces, together with precision-4^{-N} flat models of dimension $2^{(3+\lambda)N+o(N)}$. The trace gap alone supplies no such upper models. Section 15 constructs them using finite perfect witnesses and an elementary slow-admission schedule. It thereby proves that the dimension spectrum strictly extends beyond the left-c.e. reals. Section 16 completes the positive classification; the remaining questions are in Section 17.

15 Escaping the barrier: a complementary halting flag

The barrier in Section 14 concerns the minimum over *all* exact traces. We now construct a presentation whose minimum over matrix-approximable traces is

$$\lambda_A = 1 - \theta_A, \qquad \theta_A = \sum_{e \in A} 4^{-(e+1)},$$

for any c.e. set A. For $A = K$ this minimum is not left-c.e. The clocks retain their matched dimension budget. Every statement in this section is proved here, except the MIP* = RE input of Section 15.1 and, in Theorem 15.7, the cited Boolean correspondence, actual-state interface and heat sandwich.

15.1 Finite perfect witnesses as an external input

We use Ji–Natarajan–Vidick–Wright–Yuen's MIP* = RE theorem (arXiv version 4, accepted to Annals of Mathematics) in its finite perfect-completeness form [39, Theorem 12.6, pp. 193–194]. The theorem gives a computable map from machines to finite rational games: a halting machine yields a finite-dimensional perfect PCC strategy, and a nonhalting machine yields quantum value at most 1/2. Its proof starts with the trivial perfect PCC strategy of Lemma 12.3 and applies Claim 12.7 in finitely many downward steps, using the compression completeness of Theorem 12.1(1), p. 184. Only the finite-dimensional perfect PCC strategy of Theorem 12.6(1) is used: a state ψ on $H \otimes H$ and projective measurements, each player's measurements consistent on ψ in the sense of [39, Definition 5.10]. PCC means projective, consistent and commuting in that paper's sense; no symmetry between the players is assumed. The informal synchronous wording of Section 1.2 of [39] is not needed; mere quantum value one would not suffice.

Apply the map to the machine which halts precisely when $e \in A$, and denote the resulting game by $H_{A,e}$, with question sets $X_{\mathrm{A}}, X_{\mathrm{B}}$. Introduce projections p_{xa} ($x \in X_{\mathrm{A}}$) and q_{yb} ($y \in X_{\mathrm{B}}$), impose the PVM equations of both families, and impose $p_{xa}q_{yb} = 0$ for each tested losing pair. Untested pairs impose no zero-product condition. Call this finite rational projection presentation $\mathcal{A}_{A,e}$.

Lemma 15.1. *If $e \in A$, then $\mathcal{A}_{A,e}$ has an exact unital matrix model. If $e \notin A$, it has no sequence of flat matrix models whose defining defects tend to zero. Exact matrix certificates, without repeated indices, can be enumerated effectively, and their indices are exactly A.*

Proof. Let (A, B, ψ) be a finite-dimensional perfect PCC strategy and write $\psi = \mathrm{vec}(S)$. Consistency of each player's measurements reads $A_a^x S = S(A_a^x)^{\mathsf{T}}$ and $S(B_b^y)^{\mathsf{T}} = B_b^y S$; taking adjoints shows that every A_a^x and B_b^y commutes with Alice's reduced state $\rho = SS^*$. Write $\rho = \sum_c cP_c$. Its eigenspaces reduce both projection families. For a tested losing pair, consistency of Bob's measurement and perfect play give

$$0 = \langle \psi, A_a^x \otimes B_b^y \, \psi \rangle = \mathrm{Tr}(\rho A_a^x B_b^y) = \sum_{c>0} c \, \|(P_c B_b^y P_c)(P_c A_a^x P_c)\|_2^2.$$

Every summand is nonnegative. On each nonzero positive-eigenvalue eigenspace the two compressed families are PVMs and every tested losing product vanishes, including its adjoint. Restrict to any such eigenspace to obtain the unital matrix model. The commuting clause of PCC is not needed for this argument.

Conversely, approximate projection partitions can be repaired to PVMs in the same dimension. For projections $P_1, \dots, P_k$ with $S_0 = \sum_i P_i$ and $\|S_0 - I\|_{2,d} \to 0$, where $\|X\|_{2,d}^2 = \mathrm{tr}_d(X^*X)$,

$$\sum_{i \neq j} \mathrm{tr}_d(P_i P_j) = \mathrm{tr}_d(S_0^2) - \mathrm{tr}_d(S_0) \longrightarrow 0.$$

Each term is nonnegative and equals $\|P_iP_j\|_{2,d}^2$. Keep P_1, successively compress P_i to the complement R of the previously repaired outcomes and round at $1/2$, and use the last complement as the final outcome. If $\alpha = \|(I-R)P_i\|_{2,d}$, then

$$\|P_i - RP_iR\|_{2,d} \le 2\alpha, \qquad \|1_{[1/2,1]}(RP_iR) - RP_iR\|_{2,d}^2 \le \mathrm{tr}_d(RP_iR - (RP_iR)^2) \le \alpha^2.$$

Finite induction and the approximate sum equation give corrections tending to zero. Bounded-product telescoping preserves every tested losing-product defect. Using the repaired PVMs on a maximally entangled state, with Alice using the transposed p_{xa}^{T} and Bob q_{yb}, gives losing probabilities $\mathrm{tr}_d(p_{xa}q_{yb}) = \|q_{yb}p_{xa}\|_{2,d}^2 \to 0$. This contradicts the soundness bound when $e \notin A$.

For fixed e, d, real quantifier elimination decides the rational semialgebraic matrix equations, after splitting complex entries into real and imaginary parts. A nonempty set has an effectively obtainable algebraic point. Dovetail over e, d and suppress repeated indices. The first two assertions make its output exactly A. No dimension bound computable from e is asserted. □

15.2 Packing and a fixed polynomial-testing program

For a projection presentation $\mathcal{A}$ with $m \ge 1$ generators, use two projections D, E and a unitary W. Set $G_j = W^{-j}DW^j$, $0 \le j < m$, and impose

$$\sum_{j<m} W^jEW^{-j} = I, \qquad [E, G_j] = 0. \tag{15.1}$$

Substitute EG_jE for its generators in every equation of $\mathcal{A}$, replacing the constant identity by E. Clear denominators, and call the finite integer-coefficient list consisting of (15.1) and the substituted equations $\mathcal{R}(\mathcal{A})$. In $\widehat{\mathcal{P}}_A$ these are not defining relations; the program tests them.

Lemma 15.2 (Three data generators)**.** *A d-dimensional model of $\mathcal{A}$ extends to an md-dimensional model of $\mathcal{R}(\mathcal{A})$. Conversely the E corner of every tracial model of the latter carries $\mathcal{A}$, and $\tau(E) = 1/m$.*

Proof. List the projection generators as $g_0, \dots, g_{m-1}$, including both players' families for $\mathcal{A}_{A,e}$. Take $D = \mathrm{diag}(g_0, \dots, g_{m-1})$, $E = |0\rangle\langle 0| \otimes I_d$, and let W cyclically shift the m blocks forward. Then $W^{-j}DW^j$ restricts to g_j on block zero, proving the extension, also when $m = 1$. Conversely EG_jE is a projection in the unital E corner and satisfies every substituted relation. Tracing (15.1) gives $m\tau(E) = 1$. Arbitrary models need not satisfy $W^m = I$. □

Write $m_e = \#\{p_{xa}\} + \#\{q_{yb}\}$ for the total number of projections in both families of $\mathcal{A}_{A,e}$, adding an unused projection if needed, and $\mathcal{R}_e = \mathcal{R}(\mathcal{A}_{A,e})$. The map $e \mapsto \mathcal{R}_e$ is total computable. Retain the decoder L, R, reset Z, clocks, carry and conditional-clock relations of Section 5. Replace the halting recognizer by the following fixed counter transducer, delete (5.11), and add

$$(I-L)(I-F) = 0. \tag{15.2}$$

Thus an unflagged address must be in the decoded region. Keep the scalar counter domains, zero tests, controls, initialization and transition rules; their start and input are still determined by L, R.

Add the data generators D, E, W. Require them to commute with A_a, L, R, F, Z, every fuel and time generator including U, and every instantaneous scalar workspace variable. Do not require them to commute with the address shift or flip, or with one another.

The transducer has two operator registers, a monomial M and an accumulator Σ. Represent each $S \in \{M, \Sigma\}$ by two unitary generators:

$$S = (u_S + v_S)/2. \tag{15.3}$$

Every such S is a contraction, and every finite square contraction has this form. Indeed extend its polar partial isometry to a unitary V_0, put $J = |S| + i(I - |S|^2)^{1/2}$, and take $u_S = V_0 J$, $v_S = V_0 J^*$. This model assignment does not add irrational coefficients to (15.3).

On input e, compute the finite list $\mathcal{R}_e$. For each polynomial r, express its integer coefficients as repeated signed monomials, and compute an integer $b \geq 0$ with 2^b at least their total number. Clear Σ, set $M = I$ for each monomial, build it by successive right multiplications, and add its signed value times 2^{-b} to Σ. Enter CHECK, then process the next polynomial and finally idle. The extra instruction alphabet is

$$\Sigma := 0, \quad M := I, \quad M := Md \ (d \in \{D, E, W, W^*\}), \quad \Sigma := \Sigma \pm r_b M,$$

and holding a register fixed. The ordinary counter encoding supplies $r_b = 2^{-b}$. The final idle branch clears the registers to their initialized values. Every partial monomial and accumulator is a contraction, including on flagged addresses with arbitrary data.

Counter-machine universality compiles this algorithm into one fixed finite instruction table. Its finite branch selectors q_α and updates define rational polynomials Φ_S. The polynomial lists and word lengths are program outputs, rather than additional defining equations. Put $\mathrm{Init}(M) = I$, $\mathrm{Init}(\Sigma) = 0$, and impose

$$ZS = Z\,\mathrm{Init}(S), \quad U^*SU = Z_+\,\mathrm{Init}(S) + (I - Z_+)\Phi_S, \quad Z_+ = U^*ZU, \tag{15.4}$$

$$q_{\mathrm{CHECK}}\Sigma(I - F) = 0. \tag{15.5}$$

For each of the four representing unitaries in (15.3), impose commutation with the instantaneous scalar workspace and with $A_a, A_f, A_t, L, R, F, Z$. Impose no commutation with other operator registers, data, time shifts, or fuel/time flips. Scalar workspace here excludes M, Σ. This completes a finite bounded rational presentation, denoted $\widehat{\mathcal{P}}_A$.

Lemma 15.3 (Programmed testing)**.** *Every faithful exact tracial model satisfies, for every e and $r \in \mathcal{R}_e$,*

$$B_e(I - F)r(D, E, W) = 0. \tag{15.6}$$

Proof. Fix the finite CHECK time h for this relation and choose $2^k > h$. Put $Q_0 = B_e C_k D_k (I - F)$, using (5.7). The scalar initialization lemma gives $ZQ_0 = Q_0$, with no reset before h, and forces the scalar controls and counters throughout this finite history. Non-self-adjoint registers require an extra justification: $[S, Z] = 0$ gives

$$SQ_0 = SZQ_0 = ZSQ_0 = Z\,\mathrm{Init}(S)Q_0.$$

Consequently $MQ_0 = Q_0$ and $\Sigma Q_0 = 0$ initially.

Every data word commutes with Q_0; the address digits are spectral functions of A_a in an exact clock. The data also commute with U. If $S_l = U^{-l}SU^l$ has right action $S_l Q_0 = wQ_0$, then $S_l dQ_0 = wdQ_0$ for a data letter d. A conjugated scalar selector commutes with every data word, since the data commute with that selector and with U. Its forced scalar right action therefore holds on wQ_0 as well. Induction in (15.4) proves the initialized operator history. Products do not require commuting computed registers: if $AQ_0 = aQ_0$, $BQ_0 = bQ_0$, and $[b, Q_0] = 0$, then $ABQ_0 = AQ_0 bQ_0 = abQ_0$. Before h the next-reset factor vanishes on Q_0, whose range also contains every computed right action. At CHECK,

$$\Sigma_h Q_0 = 2^{-b} r(D, E, W)Q_0, \qquad (q_{\mathrm{CHECK}})_h Q_0 = Q_0.$$

Conjugate (15.5) by U^{-h} and apply it to Q_0. The flag is time invariant, and the conjugated selector commutes with data words, so $r(D, E, W)Q_0 = 0$.

The operator $r(D, E, W)B_e(I - F)$ commutes with all fuel/time flips. Conjugate its zero product with $C_k D_k$ by the earliest specified fuel flip and add: that flip exchanges the two outcomes of its digit and preserves later digits. Repeat through the final zero of C_k, then through D_k. This removes both cylinders and proves (15.6). No conditional trace independence is needed. Since $\mathcal{R}_e$ includes the packed relations for both families, the E corner of every nonzero unflagged cylinder therefore carries both the p_{xa} and q_{yb} PVMs and all tested losing products. □

The approximable flag minimum and all complementary finite-model and resource estimates are proved in Appendix D: Proposition 15.4, Lemmas 15.5 and 15.6, Theorem 15.7, Corollary 15.8 and Proposition 15.9.

15.3 From a c.e. gate to a convergent approximation

The predicate certified by a finite perfect witness is c.e.: finite exact matrix feasibility can be searched by real quantifier elimination. This limits the predicate at one gate, not the exponent of a whole construction. Section 16 keeps these c.e. continuation gates and adds a different scalar tail test after stopping. A compactly encoded stopping parameter is safe only because a bad address is ruled out both at every finite stop and on an infinite continuation. Thus the right-c.e., d.c.e. and general Δ^0_2 logarithmic realization questions are all answered by one construction. No subtraction of games or disjunction of all witness candidates is needed. Appendix B and the ancillary README describe the finite checks of the complementary-construction algebra and of the complete construction.

16 Every positive limit-computable exponent

The reader of Section 11 keeps every accepted address. That monotone rule realizes left-c.e. flag densities. We now allow a model to choose when to stop reading, while checking every later rational approximation. Finite quantum witnesses certify each decision to continue. A scalar computation verifies the tail after a finite stop. Together these two tests handle an arbitrary computable convergent approximation, including one of unbounded total variation.

Theorem 16.1 (A flag for a computable limit)**.** *Let $(a_s)_{s\ge 0}$ be a total computable sequence of rationals in $[0,1]$, with limit θ. There is a finite bounded rational presentation $\mathcal{P}[a]$, retaining the three compulsory clocks and the conditional clock, such that*

$$\min_{\tau\in\mathcal{T}_{\mathrm{mat}}(\mathcal{P}[a])} \tau(F) = 1-\theta. \tag{16.1}$$

For every sufficiently large even N it has flat models whose defining operator-norm residuals are at most $3\cdot 4^{-N}$ and whose dimensions satisfy

$$\log_2 d_N = (4-\theta)N + o(N). \tag{16.2}$$

Every new defining equation is exact in these models.

The presentation is specified by the finite program computing (a_s); its real limit is not a coefficient of any relation. We give the presentation and both sides of (16.1) before deriving the game laws.

16.1 A finite program with two different tests

For a binary word w of length j, with integer binary value $p(w)$ in reader order, put

$$\ell(w) = 2^{-j}p(w), \qquad u(w) = 2^{-j}(p(w)+1).$$

The radical-inverse value of every extending address lies between these endpoints. At the empty word they are zero and one. The continuation predicate is the c.e. statement

$$C(j,w) \iff \exists s\ge j \quad a_s > \ell(w). \tag{16.3}$$

Apply Lemma 15.1 to the search machine for this predicate. Its projection presentation $\mathcal{A}_{j,w}$ has separate Alice and Bob families. A true predicate has an exact finite matrix witness; a false one admits no sequence of flat matrix models with defects tending to zero. Lemma 15.2 packs its complete relation list into D, E, W. Computing that list always terminates: it applies the finite game reduction to a machine description, rather than waiting for the search machine to halt.

A stop at (j,w) instead starts an ordinary counter computation of $a_j, a_{j+1}, \ldots$. It enters an absorbing REJECT control when it finds

$$a_s \le u(w), \qquad s\ge j. \tag{16.4}$$

Impose $q_{\mathrm{REJECT}} \le F$. Absence of rejection is an infinite tail condition, not a finite quantum certificate.

Start with Section 11's reader presentation: delete the old recognition program and its halt-flag equation, retain the reader domains and READ updates, and retain all clock, carry, reset, flag and conditional-clock relations. Thus the decoder L, R of Section 5 is absent, and no decoded-region restriction (15.2) is imposed.

For clarity, the inherited scalar workspace consists of all counter coordinates, their zero-test projections and positive-contraction witnesses, the mutually orthogonal control projections summing to I, and $A_{\mathrm{r}}, C_{\mathrm{r}}, P_{\mathrm{r}}$. They commute pairwise and with Z, A_a, A_f, A_t, F. Retain their domains, zero tests (5.8) and reader equation (11.2). For every ordinary counter, control and the reader impose the initialization and transition forms (5.9)–(5.10), using the new finite instruction table; in particular $ZA_{\mathrm{r}} = ZA_a$. Zero-test and reader witnesses satisfy their domain equations at every time and need no separate updates. The flag still commutes with A_a and the complete fuel/time sets, including U. Add two projection data generators D, E and a unitary W, and represent M, Σ by the four unitaries of (15.3). LOAD maps M to I and Σ to zero without reading data on its outgoing edge; other scalar instructions hold them unless an operator instruction updates them.

Add a positive contraction R_* and a counter c initialized by

$$Zc = ZR_*, \qquad U^*cU = Z_+R_* + (I - Z_+)\,\mathrm{Step}(c), \qquad Z_+ = U^*ZU. \tag{16.5}$$

All ordinary counters are initialized at I, the start control at I, the other controls at zero, and the reader at A_a. Counters retain the zero-test relation (5.8). Counter value zero has encoding $c = 1$; a nonzero decrement doubles c. The choice $R_* = 2^{-m}$ therefore means stop at depth m, while $R_* = 0$ continues at every finite depth.

At depth j the program first zero-tests c. Its zero branch starts the scalar verifier (16.4). Its nonzero branch computes the packed list for (16.3), executes one LOAD, tests that entire list, reads the next address bit, increments j and decrements c. A CHECK never branches on a matrix value. Every gate has *exactly one LOAD, before any data-using instruction*; there is no LOAD between its monomials or relations. After LOAD, initialize the monomial and accumulator registers, then use the signed right-multiplication program of Section 15. All classical arithmetic is ordinary counter computation. This algorithm has a fixed finite instruction table, although its computed relation lists and runtimes are unbounded.

The precise new commutations. Require R_* to commute with A_a, F, Z, every fuel and time clock generator including U, every instantaneous variable of the commutative workspace, the data D, E, W, and the four register-representing unitaries. Each data generator commutes with

$$A_a, A_f, A_t, F, Z, R_* \quad \text{and every instantaneous workspace variable.}$$

The four unitaries representing M, Σ by (15.3) have that same coordinate/workspace commutation list. They need not commute with the data or with one another. These are the complete additional lists: in particular, the new data do not inherit Section 15's commutations with U, fuel/time rotations or fuel/time flips. No address-shift or address-flip commutation is imposed on the new data or registers. Their required square-root-state commutators in approximate models are those of Definition 2.1, as for every declared generator.

Instead of time invariance, for $X \in \{D, E, W\}$ impose

$$(U^*XU - X)(I - Z_+)(I - q_{\mathrm{LOAD}}) = 0. \tag{16.6}$$

The tuple at the next time can change only after a current LOAD or at a next-time reset. Use the register domains, initialization and transition formulas (15.3) and (15.4), with the new finite branch selectors, and impose

$$q_{\mathrm{CHECK}}\Sigma(I - F) = 0, \qquad q_{\mathrm{REJECT}} \le F. \tag{16.7}$$

The normalization of signed monomials bounds every partial accumulator by one, even on flagged addresses. All displayed abbreviations have finite rational polynomial definitions. These equations and commutations complete $\mathcal{P}[a]$.

16.2 Why every completed gate tests one tuple

Write $Y_h = U^{-h} Y U^h$. For a finite address word w, choose k and the initialized unflagged projection

$$Q_0 = B_w C_k D_k (I - F). \tag{16.8}$$

The reader follows the specified address bits while it has read at most $|w|$ of them. We consider a prescribed finite computation shorter than 2^k; Lemma 6.1 gives reset at time zero, and Proposition 4.5 gives no later reset during this interval, as in (6.3).

Lemma 16.2 (A fixed tuple between loads)**.** *Suppose the scalar control flow on Q_0 reaches one continuation gate. From immediately after that gate's LOAD through all its CHECKs, every data generator has the same restriction to Q_0. Every CHECK for a packed polynomial π forces $\pi(d_{\mathrm{load}})Q_0 = 0$. If $\tau \in \mathcal{T}_{\mathrm{mat}}(\mathcal{P}[a])$ and $\tau(Q_0) > 0$, the corresponding continuation predicate is true.*

Proof. Shifted variables reduce the initialized corner. Let Y be any data generator, register unitary or instantaneous workspace variable. It commutes with A_a, A_f, A_t, F. The time increment commutes with the address and fuel sets and with F. Proposition 4.5 says that $U^h D_k U^{-h}$ is a sum of finite time-prefix cylinders, hence a spectral function of A_t. The other factors of $U^h Q_0 U^{-h}$ are unchanged. Consequently

$$[Y_h, Q_0] = 0. \tag{16.9}$$

This proves reduction for every shifted variable needed below without assuming $[Y, U] = 0$.

Conjugate (16.6). Off a programmed LOAD and before the next reset, its right guard acts as I on Q_0, giving

$$X_{h+1} Q_0 = X_h Q_0. \tag{16.10}$$

If LOAD occurs at h_L, put $d_{\mathrm{load}} = (D_{h_L+1}, E_{h_L+1}, W_{h_L+1})$. Its restrictions stay fixed throughout the gate. They have their exact domains on the reducing Q_0 corner. In particular the restricted unitary has the same fixed adjoint.

Register induction. Each register is cleared after LOAD. Suppose $S_h Q_0 = P(d_{\mathrm{load}}) Q_0$. Every polynomial in shifted data commutes with Q_0 by (16.9). For a current data letter d_h and a shifted scalar selector $q_{\sigma,h}$ with scalar action σQ_0, the required identities are

$$\begin{aligned} S_h d_h Q_0 &= S_h Q_0 d_{\mathrm{load}} = P(d_{\mathrm{load}}) d_{\mathrm{load}} Q_0, \\ q_{\sigma,h} P' Q_0 &= q_{\sigma,h} Q_0 P' = \sigma P' Q_0, \\ Z_{h+1} P' Q_0 &= Z_{h+1} Q_0 P' = 0, \qquad [P', Q_0] = 0. \end{aligned}$$

They handle multiplication, branch selection and the absent reset term in the register transition. The scalar dyadic coefficient acts by its programmed value on the same corner. We have not commuted a selector through a data word at a different time. Non-self-adjoint initialization is justified by $[S, Z] = 0$: $SQ_0 = SZQ_0 = ZSQ_0 = Z\,\mathrm{Init}(S)Q_0$. Induction therefore gives at each CHECK

$$\Sigma_h Q_0 = 2^{-b} \pi(d_{\mathrm{load}}) Q_0. \tag{16.11}$$

The CHECK selector equals I there, and F is time invariant. Equation (16.7) forces this polynomial to vanish.

Positive corners. If the packed presentation has m projection generators, the packing relation and traciality give

$$\tau(Q_0 E_{h_L+1}) = \tau(Q_0)/m.$$

The projection $Q_0 E_{h_L+1}$ is a fixed polynomial in the full generators, including the finite time and fuel cylinders. If its trace is positive, it carries the recovered Alice and Bob projections of $\mathcal{A}_{j,w}$.

It also inherits matrix approximation. In flat approximants, self-adjointize its polynomial and round at $1/2$. The pointwise inequality $|1_{[1/2,\infty)}(t) - t| \le 2|t^2 - t|$ makes the normalized Hilbert–Schmidt correction vanish; its rank fraction tends to the positive trace above. Compress and round the recovered projection generators, not the packed unitary itself, into that corner. Bounded-product telescoping preserves all their fixed relations. Renormalizing trace costs a factor tending to the inverse square root of that positive trace. This yields flat approximate models of $\mathcal{A}_{j,w}$. Lemma 15.1 excludes them when the predicate is false. We keep $C_k D_k$ in this corner; no commutation of the new data with fuel/time flips is used. □

16.3 The stopping parameter cannot escape to infinity

Lemma 16.3 (Commutative right actions and stop sectors)**.** *On a fixed finite address itinerary in Q_0, the counter and control actions used by the program are Borel functions of R_* times Q_0. If $Q_0\, 1_{\{2^{-i}\}}(R_*) = 0$ for every $i < j$, so that all of Q_0 has continued through the depths below j, then the depth-j zero test at time h_j satisfies*

$$c_{h_j} Q_0 = 2^j R_* Q_0, \qquad Q_0 z_{c,h_j} = Q_0\, 1_{\{2^{-j}\}}(R_*). \tag{16.12}$$

The right-hand projection is the polynomial projection on the left. Every shifted program or data variable reduces it.

Proof. The reader action is $s^i(A_a)Q_0$ after i reads, as in (11.5); its next-bit action is scalar on the specified itinerary. Ordinary counter initialization and updates are unchanged. The exception is c: using $[c, Z] = 0$, $ZQ_0 = Q_0$ and $[R_*, Q_0] = 0$, its initialization gives $cQ_0 = R_* Q_0$.

The scalar induction must now take place in the commutative algebra of functions of R_* on Q_0. If $AQ_0 = f(R_*)Q_0$ and $BQ_0 = g(R_*)Q_0$, then

$$ABQ_0 = AQ_0 g(R_*) = f(R_*)g(R_*)Q_0. \tag{16.13}$$

Thus the polynomial transition maps preserve these right actions. For a counter with such a self-adjoint right action, taking adjoints shows that Q_0 reduces it. The zero-test relation forces $z = 1_{\{1\}}(c)$, since $c = 1$ on z and $0 \le c \le 1/2$ on $I - z$. Functional calculus on Q_0 supplies the next zero-test projection as a Borel function of R_*. Only $(I - z)s$ of the auxiliary witness is used; the unconstrained block zs is never read by the program. All right actions entering the induction are therefore functions of R_*. Control projections may be such functions while different stop choices coexist; they are not assumed scalar at that point.

Conjugating the transition of c in (16.5) kills its reset term on the no-reset interval. On a surviving nonzero branch, c is held fixed until its decrement and then doubled. After j such decrements it acts as $2^j R_*$. The spectral zero-test formula gives (16.12).

Every shifted variable reduces Q_0 by (16.9). It commutes with R_* because its instantaneous variable does and $[U, R_*] = 0$. It therefore reduces each stop sector. On the depth-j sector the stop counter has value one, and the subsequent verifier is the ordinary scalar computation prescribed by that prefix.

No discrete spectrum was assumed for R_*. On a continuing sector the zero-test domain itself forces R_* into $[0, 2^{-(j+1)}]$ after the depth-j atom has been removed. The value zero always continues. These conclusions follow at every finite depth and cover any continuous spectral part as well. □

Proposition 16.4 (The approximable lower bound)**.** *Every $\tau \in \mathcal{T}_{\mathrm{mat}}(\mathcal{P}[a])$ satisfies $\tau(F) \geq 1 - \theta$.*

Proof. Call a word w of length J bad when

$$\ell(w) > \theta, \qquad a_s < \ell(w) \quad (s \geq J). \tag{16.14}$$

Every address with $v(x) > \theta$ has a bad prefix: choose a gap between its value and the limit, then take J large enough both for the prefix lower endpoint to lie inside that gap and for the sequence's entire remaining tail to lie below it.

Fix a bad word. Each possible stop at an ancestor of depth $j \leq J$ has $u(w|j) \geq \ell(w) > \theta$, so its verifier eventually finds a rejection. There are finitely many such stop depths. Their first rejecting runtimes, including preceding continuation work, are finite. Computing each gate list and its instruction stream through depth J also has a fixed finite runtime independent of the data and of R_*. Choose one k with 2^k above all these runtimes.

If $\tau(B_w(I - F)) > 0$, the retained fair fuel/time cylinders give

$$\tau(Q_0) = 2^{-(2k+1)}\tau(B_w(I - F)) > 0. \tag{16.15}$$

Only $B_w(I - F)$, not the new data, must commute with the fuel/time flips for this equality.

Eliminate the stop sectors in increasing order $j = 0, \ldots, J$. Once the sectors of smaller depth are known to vanish, Lemma 16.3 identifies the next one as $Q_j = Q_0 z_{c,h_j}$. On that reducing projection the scalar verifier reaches REJECT before reset. Since $q_{\mathrm{REJECT}} \leq F$ and $Q_j \leq I - F$, it follows that $Q_j = 0$. The continue branch is then forced on the entire remaining Q_0. At the next depth the same argument applies. Intermediate continuation CHECKs do not branch, so their finite control schedule is fixed. After all these stop sectors vanish, the computation must execute the gate for $C(J, w)$.

Condition (16.14) makes that predicate false. Lemma 16.2 contradicts (16.15). Thus $\tau(B_w(I - F)) = 0$ for every bad word, including on a never-stopping sector $R_* = 0$.

For a finite measure argument, fix $\zeta > 0$. The compact set $\{v \geq \theta + \zeta\}$ is covered by finitely many bad cylinders. Refine them to a common address depth. Their orthogonal sum has zero unflagged trace, and fairness bounds the trace of the remaining cylinders by $\theta + \zeta$ plus an arbitrarily small dyadic endpoint error. Let that error and then ζ decrease to zero. We obtain $\tau(I - F) \leq \theta$. The possibly enormous runtime chosen above is used only for exact tracial forcing, not in the quantitative measurement depth. □

16.4 Matched finite models with a fixed witness budget

Lemma 16.5 (Finite threshold models)**.** *For each dyadic $0 < b < \theta$ there are integers m and $\ell_b \geq 1$, independent of N, such that for every even $N \geq m$ the presentation has flat models with all new equations exact, old residuals at most $3 \cdot 4^{-N}$ in operator norm, and*

$$\mathrm{Tr}(F)/d_{N,b} = 1 - b, \qquad d_{N,b} = 2^{3N+k_N}\ell_b, \qquad k_N = (1 - b)N + O(\log(N + 2)). \tag{16.16}$$

Proof. Choose m with $2^m b$ integral and $a_s > b$ for every $s \geq m$. Convergence guarantees this fixed choice. On the N-bit address cycle set

$$I - F = 1_{\{v_N(x) < b\}}, \qquad R_* = 2^{-m} I.$$

Every unflagged address stops after m reads. Its final prefix has $u(w_m) \leq b < a_s$ for all $s \geq m$, so its scalar verifier never rejects. At every preceding depth $j < m$, $\ell(w_j) \leq v_N(x) < b < \theta$, so the continuation predicate is true.

Only finitely many prefix nodes occur on these paths. Choose an exact matrix witness at each, pack both players' families, and tensor their spaces. The resulting common data factor has some finite dimension ℓ_b depending on b, m, not on N. Use the three compulsory N-qubit registers and the optional-slot construction of Theorem 7.1 for this Boolean mask.

For fuel with k initial ones use the actual initialized history at time t mod 2^k; all-ones fuel uses $k = N$. On each address/fuel/time block assign the commutative workspace its scalar history. On unflagged blocks, after a gate's single LOAD assign that gate's packed tuple on its tensor factor and identities on the others, holding it fixed until the next LOAD or reset. From a reset until the first LOAD after it, use the default tuple $(0, I, I)$. On flagged addresses take one default tuple $(D, E, W) = (0, I, I)$ *constant in time*, even at LOAD. CHECK is vacuous there and scalar rejection is permitted.

Assign M, Σ their actual bounded operator histories. Within each block choose their average-of-two-unitary representations as in Section 15; this adds no tensor factor. All declared coordinate/workspace commutations hold because these assignments are block diagonal in the existing address, fuel and time registers. The scalar R_* has every required commutation. We impose no new commutation with a fuel/time rotation or flip.

At an ordinary non-LOAD transition the data are unchanged, so (16.6) is exact. At LOAD its final guard vanishes. At the next-time reset its reset guard vanishes. Since every fuel period divides 2^N, the cyclic boundary is also a reset. Register transitions follow the assigned history, including resets that cut a gate before completion. Each completed unflagged CHECK uses its one exact witness tuple; every unflagged tail comparison passes. Hence all new equations hold exactly at every block and boundary. Address dressing acts only on address and optional slots and commutes with these added assignments. The old residual bounds are therefore unchanged.

In an aligned block of 2^h consecutive address integers, the radical-inverse grid count below b differs from $2^h b$ by at most one. Decompose a length-N interval into at most two aligned blocks of each relevant size. Uniformly over starts, including cyclic intervals,

$$\sum_{i=1}^{N} F(x - i) = (1 - b)N + O(\log(N + 2)).$$

Taking one more optional slot than the largest such count gives (16.16). The finite grid has exactly a fraction b of unflagged addresses because $N \geq m$. □

Proof of Theorem 16.1. Choose dyadics $b_i \uparrow \theta$ and associated m_i, ℓ_i from Lemma 16.5. Choose strictly increasing even thresholds N_i with

$$N_i \geq m_i, \qquad N_i \geq i^2 \max\{1, \log_2 \ell_i\}.$$

For $N_i \leq N < N_{i+1}$ use the b_i model. Its witness overhead divided by N is at most i^{-2}; the uniform discrepancy bound and $b_i \to \theta$ give (16.2). Its flag trace tends to $1 - \theta$. A moment limit attains that value in $\mathcal{T}_{\mathrm{mat}}$, and Proposition 16.4 gives the reverse inequality. When $\theta = 0$, use $F = I$, $R_* = I$ and constant default data directly. The case $\theta = 1$ uses the same thresholds below one. The diagonal choices are existential: the finite presentation is still specified by the program for (a_s), and no effective convergence modulus or uniform witness bound is claimed. □

16.5 The game laws and the positive classification

Theorem 16.6 (Matching dimension and usable entanglement)**.** *For every total computable rational sequence $a_s \to \theta \in [0, 1]$, integers $r \geq 0$ and $q \geq 1$, there is a fixed finite rational*

BCS game $G^{\mathrm{stop}}_{a,r,q}$ of quantum value one, without a finite-dimensional perfect strategy, such that, with $L_\varepsilon = \log_2(1/\varepsilon)$ and $\beta = (4 + r - \theta)/(4q)$,

$$\begin{aligned} \log_2 D_{G^{\mathrm{stop}}_{a,r,q}}(\varepsilon) &= \beta L_\varepsilon + o(L_\varepsilon), \\ d^{\rightarrow}_{G^{\mathrm{stop}}_{a,r,q}}(\varepsilon) &= \beta L_\varepsilon + o(L_\varepsilon). \end{aligned} \tag{16.17}$$

Proof. By Corollary 14.3 and Theorem 16.1, uniformly on every original weighted precision-δ model,

$$\mathrm{Tr}(\rho F) \geq 1 - \theta - \omega(\delta), \qquad \omega(\delta) \longrightarrow 0. \tag{16.18}$$

Otherwise a moment limit with a fixed positive shortfall would contradict (16.1). Append r compulsory clocks and power every defining residual, including all new equations and cross-commutators, as in Proposition 11.3. Retain the exact domains and square-root-state requirements. Jensen gives unpowered precision $\delta = (C_{\mathcal{P},r,q}\sqrt{\varepsilon})^{1/q}$ from a loss-ε strategy.

Set $t = \log_4(1/\delta)$ and $M = \lfloor t - 30\log_4(2 + t)\rfloor$. The retained raw clock measurements have error

$$\eta = C_{\mathcal{P},r,q}(M + 1)^6 4^M \delta = O(t^{-24}).$$

The new data, the stopping parameter R_*, the stop counter and the registers occur in none of those measurements or controls. The same count as in Theorem 8.1, with the extra clocks, gives

$$\begin{aligned} \log_2 d &\geq (3 + r)M + \sum_{i<M} \mathrm{Tr}(\rho F_i) - o(M) \\ &\geq (4 + r - \theta - \omega(\delta))M - O(M^2\delta) - o(M). \end{aligned}$$

The second line uses (16.18) on this same marginal and $|\,\mathrm{Tr}(\rho F_i) - \mathrm{Tr}(\rho F)| \leq 2i\delta$. Since $M = L_\varepsilon/(4q) - O(\log L_\varepsilon)$, this proves the dimension lower coefficient.

For the other direction, the flat models of Theorem 16.1 gain exactly rN logarithmic dimension from the extra clocks. Every powered operator residual is at most $3^q 4^{-qN}$. The compiler gives loss at most $C'16^{-qN}$ and constant dimension overhead. Neighboring even depths change logarithmic precision by a bounded amount. Their dimension law is $(4 + r - \theta)N + o(N)$, proving the matching upper coefficient at every sufficiently small loss. The upper losses tend to zero and the positive lower bound diverges, so the game has value one and no finite perfect strategy.

For usable entanglement, Proposition 12.2 applies on the original mixed strategy state and its original purifying environment. Its Boolean reconstruction is generic in the finite presentation; time invariance of the additional data is not a hypothesis. Opposite words, independent local inclusion repairs and reverse count order supply Lemma 12.1 with $m = (4 + r)M$ and the same η. Only Q_i is repaired, leaving the actual activation flags. The vector, mean and entropy-continuity errors vanish at this depth. The one-way protocol and hashing yield on every original state

$$E_D^{A\to B}(\rho_{AB}) \geq (3 + r)M + \sum_{i<M} \mathrm{Tr}(\rho_B F_i) - o(1) \geq (4 + r - \theta)M - o(M).$$

Take the infimum, and use $d^{\rightarrow}_G(\varepsilon) \leq \log_2 D_G(\varepsilon)$ for the upper bound. Compression in (16.18) identifies limiting traces; it is not a replacement of the strategy used for extraction. □

Corollary 16.7 (Complete positive dimension spectrum)**.**

$$\mathcal{B}_{\mathrm{BCS}} = \{\alpha > 0 : \alpha \text{ is } \Delta^0_2\}. \tag{16.19}$$

Every real on the right is also the original-state usable-entanglement coefficient of a game in this class. In particular all positive right-c.e. and all positive d.c.e. reals occur.

Proof. For $\alpha > 0$ computable in $0'$, choose a fixed q with $4q\alpha \geq 4$ and put

$$m = \lceil 4q\alpha \rceil, \qquad r = m - 4, \qquad \theta = m - 4q\alpha \in [0, 1).$$

Shoenfield's limit lemma [64] supplies a total computable rational sequence converging to θ; clipping it to $[0, 1]$ preserves convergence. Theorem 16.6 realizes $(4 + r - \theta)/(4q) = \alpha$, for both resources. These fixed integer choices need not be uniform in a limit name for α.

Conversely, Lemma 2.4 makes $D_G(2^{-n})$ computable for a rational finite value-one game. Approximate $\log_2 D_G(2^{-n})/n$ by computable rationals with error tending to zero. If its finite exponent exists, that sequence converges. A $0'$ oracle can find a term with no later deviation larger than a given rational tolerance; convergence guarantees one. Its distance from the limit has that bound. Thus the limit is Δ^0_2. This proves Theorem 1.1. □

Real strategies and question weights. The distributed real simulation of McKague–Mosca–Gisin [47, Section II.C] adds one rebit per party and preserves all measurement probabilities. Real strategies are also complex strategies, so

$$D_{\mathbb{C},G}(\varepsilon) \leq D_{\mathbb{R},G}(\varepsilon) \leq 2D_{\mathbb{C},G}(\varepsilon).$$

The two logarithms differ by at most one. Their exponents therefore coincide whenever either exists, and the dimension classification holds verbatim for real quantum strategies. By the reweighting lemma [26, Lemma 11.12, Section 11.6], exponents and logarithmic resource coefficients are unchanged by any strictly positive reweighting on the same finite support; the dimension classification thus also covers real, even noncomputable, positive question weights.

16.6 Why this passes the earlier barriers

Each continuation gate still tests a c.e. predicate. What changes is the choice of a finite stopping depth followed by a universal scalar tail condition:

$$\exists j \; \forall s \geq j \quad a_s > u(w_j).$$

The arbitrary-length finite histories verify that condition without a finite quantum certificate of the whole tail. An infinite continuation cannot evade a false gate above the limit, by Proposition 16.4. This is the additional mechanism missing from the fixed finite-witness rule of Section 15.

The exact-trace barrier remains intact. If $1 - \theta$ is not left-c.e., Propositions 14.1 and 14.4 give

$$\min_{\mathcal{T}(\mathcal{P}[a])} \tau(F) < \min_{\mathcal{T}_{\mathrm{mat}}(\mathcal{P}[a])} \tau(F) = 1 - \theta.$$

Exact traces can contain tracial models of false continuation gates; matrix approximation excludes precisely those positive corners. Thus even when the quantum and commuting values of the compiled game both equal one, the failure of Connes' embedding conjecture changes the attainable computability class in this construction: its non-left-c.e. coefficients require different exact and matrix-approximable flag minima. The carry-checker proof of Corollary Q4.4 reaches the same class with no exact-versus-approximable gap, so this gap belongs to the flag mechanism and is not necessary for the classification.

Convergence is the matching hypothesis. For a bounded nonconvergent sequence, the same bad-prefix argument gives a flag lower bound $1 - \limsup_s a_s$, whereas safe finite-stop models approach $1 - \liminf_s a_s$. These bounds need not agree, and no equality or exponent realization for an arbitrary oscillating sequence follows from this argument. The unequal-limit problem is stated separately below. Appendix B and the ancillary README describe finite checks of the new matrix algebra and histories on known-answer toy games.

Power-scale exponents

P1 Power spectra and the realization theorem

The logarithmic laws above concern $\log D_G(\varepsilon)$ and $d_G^{\rightarrow}(\varepsilon)$ of order $\log(1/\varepsilon)$. We now classify a second scale. Write

$$p_G(\varepsilon) = \inf_{\substack{\text{pure original strategies} \\ \text{loss}\le\varepsilon}} S(\rho_B), \qquad f_G(\varepsilon) = \inf_{\text{loss}\le\varepsilon} E_F(\rho_{AB}).$$

Here E_F is the single-copy convex roof. A resource r_G has power exponent β when $r_G(\varepsilon) = \varepsilon^{-\beta+o(1)}$, equivalently $\log(1 + r_G(\varepsilon))/\log(1/\varepsilon) \to \beta > 0$. For the dimension resource take $r_G = \log D_G$, rather than D_G. All strategies have finite local spaces; the lower bounds concern their original states.

Theorem P1.1 (Power-scale realization and classification)**.** *Given a program for a total computable rational sequence $a_s \to \beta \in (0, 1]$, one can effectively construct one fixed finite rational BCS game G of quantum value one, with no finite-dimensional perfect strategy, such that*

$$d_G^{\rightarrow}(\varepsilon) = \varepsilon^{-\beta+o(1)}, \qquad \log D_G(\varepsilon) = \varepsilon^{-\beta+o(1)}.$$

The game uses neither β as a coefficient nor a modulus of convergence. For each of the resources $\log D_G$, p_G and f_G, the set of existing power exponents in $(0, 1]$ is exactly

$$\Delta_2^0 \cap (0, 1].$$

The usable-entanglement spectrum contains this class. Its reverse inclusion remains open. If a fixed game has a finite positive usable exponent outside Δ_2^0, its efficient original states have a polynomial formation/distillation gap along a sequence, as quantified in Proposition P6.2.

The proof occupies Sections P2 to P6. First, a finite-moment source has logarithmic dimension $3N + O(\sqrt{N})$ and every fixed residual moment of order $N^{-\nu}$. Second, a polynomial-cost clock localizes a growing interval of physical instructions on the same state. Third, bounded scaled counters run a fixed program while an old rational continues to generate pulses. The density of one pulses tends to β with polynomial delay in pulse count. A stopped normalized product converts that density into a residual calibration exponent. Moment models give simultaneous flat upper assignments, and the original-state lamp extraction gives the lower bound. No witness Hilbert space is added by the computation.

We use the original-state and flat-completeness clauses of Theorem 3.1. For source extraction we use the written source-level proof of [26, Theorem 12.4], based on its signed regular extraction theorem [26, Theorem 6.4]. At source error e, including actual mirrors, inverse mirrors, state commutators and fixed marker visibility, it gives

$$E_D^{A\to B}(\rho_{AB}) \ge e^{-2+o(1)}. \tag{P1.1}$$

Indeed, with $M = 2^q$ lamps and $\Lambda_q \le 2^{(\log q+23)^2}$, its budget is $M\Lambda_q^4 e^2 \le c$. Taking $q = 2\log(1/e) - O((\log\log(1/e))^2)$ gives a rate at least $c'2^q$. The constant-probability Alice filter is a physical one-way operation and its probability is charged to that rate. Thus (P1.1) concerns the original mixed state, not its entropy or a selected pure component.

Remark P1.2 (The exact-trace barrier has a different hypothesis). Proposition 14.1 concerns the matched minimum expectation of one fixed rational self-adjoint polynomial in exact traces. The exponent here is $\lim_J k(J)/J$, where $k(J)$ counts one pulses among the first J outputs

of a program. The initialized cylinder, history columns and normalized products change with J, and J grows with precision. No fixed polynomial is asserted to have minimum trace β; in every exact trace the pulse projection has the computable value c_{one} of Lemma Q8.2 (the unguarded clock has the same exact digit action), which is unrelated to β. Consequently the fixed-observable barrier does not apply. Early rational blocks have finite lengths and disappear in the limiting density. Slow or oscillating convergence changes an ineffective threshold in the estimates, not the finite coefficients of the game.

P2 Finite-moment models for the lamp source

We use the six source unitaries p, q, t, d, x, z and the 27 words of [50, Definition 7.1]. Set $s = [p, q]$, $B = psp^{-1}$, $t_i = p^{i-1}tp^{1-i}$, $J = (xz)^2$, and normalize the residuals by $R_w = (w - I)/2$, together with $R_J = (J + I)/2$. Every residual has norm at most one on unitary assignments.

Theorem P2.1 (All fixed finite moments)**.** *For every fixed real $\nu > 0$ and every sufficiently large integer N there are flat unitary assignments of this same source with $x^2 = z^2 = I$, $J = -I$,*

$$\max_w \tau_{D_N}(|R_w|^{2\nu}) \le C_\nu N^{-\nu}, \qquad \log D_N = 3N + O_\nu(\sqrt{N}), \qquad R_J = 0.$$

The presentation does not depend on ν. No uniform estimate as $\nu \to \infty$ is claimed.

P2.1 Character moments and Clifford lamps

Fix $\nu \ge 1$, $k = \lceil \nu \rceil$, $m = \lceil \log N \rceil$ and $F = \mathbb{F}_{2^m}$. Choose distinct $\alpha_1, \dots, \alpha_N \in F$ and a binary basis, and put

$$\zeta_j = (1, \alpha_j, \dots, \alpha_j^{2k-1}) \in \mathbb{F}_2 \oplus F^{2k-1} \cong \mathbb{F}_2^{n_0}, \quad n_0 = 1 + (2k-1)m.$$

Pad into $W = \mathbb{F}_2^n$ for any $n \ge n_0$. Define

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

Then $\langle v_u, v_w\rangle = c(u + w)$ and, with $b_k = (2k-1)!!$,

$$\frac{1}{2^n - 1}\sum_{u \ne 0} |c(u)|^{2\nu} \le b_k^{\nu/k} N^{-\nu}. \tag{P2.1}$$

To prove this, any binary dependence of at most $2k$ distinct columns has even length, by its first coordinate. For the dependent column set S, its field power sums of orders $0, \dots, |S| - 1$, with $|S| \le 2k$, then vanish. The Vandermonde matrix on the distinct α_j would annihilate the all-ones vector, a contradiction. Thus the signs are $2k$-wise independent and unbiased for uniform u. In their $2k$-th moment only even multiplicities survive; pairing positions bounds their number by $b_k N^k$. Divide by N^{2k}. Removing $u = 0$ decreases the average since $c(0) = 1$ and $|c| \le 1$, and concavity of the power ν/k proves (P2.1).

Use full exterior Fock space $\mathcal{H}_0 = \Lambda(\mathbb{C}^{3N})$, of dimension 2^{3N}, with Majoranas satisfying $\{\gamma(f), \gamma(g)\} = 2\langle f, g\rangle I$ on real vectors. In three orthogonal copies set

$$A_u = \gamma(v_u \otimes e_1), \quad B_u = \gamma(v_u \otimes e_2), \quad C_u = \gamma(v_u \otimes e_3), \qquad X_u = iA_uB_u, \quad Z_u = iA_uC_u.$$

These are reflections, $X_uZ_u = B_uC_u = -Z_uX_u$ and $(X_uZ_u)^2 = -I$. The Clifford construction of overlapping qubits is due to Chao, Reichardt, Sutherland and Vidick [14, Theorem 3.1]. We

use its three-family Clifford construction; the fixed finite moments and simultaneous affine covariance needed here are established below. For $c = c(u + w)$, CAR expansion gives

$$X_u Z_w - Z_w X_u = 2cB_u C_w, \qquad X_u X_w - X_w X_u = -2c(A_w A_u + B_w B_u - 2cI).$$

Since $B_u C_w$ is unitary, the absolute normalized mixed group residual is exactly $|c|I$. For the Majorana reflections A_u, A_w (and likewise B_u, B_w), the identity $UV + VU = 2cI$ gives $(UV - cI)^*(UV - cI) = (1 - c^2)I$. Put $P = A_u A_w$ and $Q = B_u B_w$. They commute, $X_u X_w = PQ$, and the same-type group commutator is $(PQ)^2$. Since $P + P^* = Q + Q^* = 2cI$, write $P = cI + i\sqrt{1 - c^2}R$ and $Q = cI + i\sqrt{1 - c^2}S$ for commuting reflections when $|c| < 1$. The scalar joint-spectrum values give $\|(PQ)^2 - I\|/2 \le 2|c|\sqrt{1 - c^2} \le 2|c|$. The endpoint cases follow directly, and the identical calculation applies to Z.

Let $D_u = \mathrm{diag}((-1)^{\langle \zeta_j, u \rangle})$ and $U_u = \Lambda(D_u \otimes I_3)$. Full second quantization is a genuine representation, so $U_u U_w = U_{u+w}$ and $U_u X_w U_u^* = X_{u+w}$, likewise for Z; there is no projective sign ambiguity. On $\ell^2(G) \otimes \mathcal{H}_0$, $G = \mathrm{GL}(n, 2)$, put

$$L_h(\delta_g \otimes \psi) = \delta_{hg} \otimes \psi, \qquad T_u(\delta_g \otimes \psi) = \delta_g \otimes U_{g^{-1}u}\psi, \qquad \widetilde{X}_w|_g = X_{g^{-1}w}, \quad \widetilde{Z}_w|_g = Z_{g^{-1}w}.$$

Direct substitution proves $L_h T_u L_h^* = T_{hu}$, $T_u T_w = T_{u+w}$ and $L_h \widetilde{X}_w L_h^* = \widetilde{X}_{hw}$, $T_u \widetilde{X}_w T_u^* = \widetilde{X}_{u+w}$, with the same identities for Z. Thus this is an exact affine representation. Every linear action commutes with the zero lamps. Its logarithmic dimension is at most $n^2 + 3N$.

Finite-field independence and small-bias spaces have established precedents in Naor–Naor [55] and Alon–Goldreich–Håstad–Peralta [1]. The Vandermonde argument and even-moment pairing are standard tools; the additional requirements here are simultaneous affine covariance, the literal source certificate and the collective clock. The finite exhaustive search in the small-bias side bound below is not claimed to be a new efficient derandomization.

P2.2 The collective clock in the affine representation

Here we give the geometric lift explicitly. Its permutation regions and spreading method are those of [25, Propositions 6.1 and 6.5, Lemma 6.4] and [26, Proposition 13.9]. Take $M \ge 48$, $n = 4M + 4$, $\ell = 2M + 2$ and right-coordinate cycles $c_r = (1\ r\ r - 1\ \cdots\ 2)$. At $0 \le i, j < M$ use

$$A_{i,j} = c_{\ell - i - j}, \quad B_{i,j} = c_{2\ell - j}\ (j < M - 1), \quad B_{i,M-1} = D_i,$$
$$D_0 = c_{2\ell - M + 1}, \qquad D_{i+1} = c_{\ell - i - (M-1)}^{-1}(1\ 2) D_i c_{\ell - i}.$$

Make the preliminary exact moves $(1\ M + 4)$ at the a-link $(M - 1, M - 1)$ and $(5\ M + 5)$ at the b-link $(0, M - 2)$, multiplying on the right-action left. The only exceptional plaquettes have labels

$$P_B(q_B) = \prod_{r=5}^{M+3} (r\ r + M), \quad P_T(q_T) = (5\ M + 5), \quad q_B = (M - 1, M - 1),\ q_T = (0, M - 2).$$

The bulk region contains q_B, the column $(M - 1, j)$ for $1 \le j \le M - 2$, and $0 \le i \le M - 3$, $0 \le j \le M - 2$, $i + j \le M - 3$. The top region contains q_T and $i, j \ge 0$, $i + j \le M - 3$. Their areas are

$$A_B = M(M - 1)/2, \qquad A_T = 1 + (M - 1)(M - 2)/2.$$

The rth transported bulk transposition at a vertex (i, j) has endpoints

$$\{r + M - 1 - i - j, r + 2M - 1 - j\}\ (i \le M - 2), \qquad \{r + M - 1 - j, r + 2M - 1 - j\}\ (i = M - 1).$$

Take their product for $P_B(i, j)$; the top label is $\{M + 3 - i - j, 2M + 3 - j\}$. The cited covariance identities restrict to this bulk region: labels are disjoint, base labels avoid $1, 2$,

and initial labels of active interior edges avoid $2,3,4$. Only those edges carry potentials; the source is attached by its bottom edge. These are finite permutation identities, so they can be checked directly from the displayed cycles and labels.

Use the inverse coordinate map to pass from right coordinates to left linear actions, and represent through L above, denoting the representation by Π. Set $Q_R(v) = (I - \Pi(P_R(v)))/2$. Then Q_R is a projection, $e^{i\pi Q_R} = \Pi(P_R)$, covariance survives and commuting labels give commuting projections. A label fixing $2,3,4$ commutes with translations on those coordinates and the transvection $E_{3\leftarrow 2}$; its represented projection and all its exponentials do too. Every linear action fixes configuration zero, so all these corrections also commute with both zero lamps. This verifies the protected-affine operations rather than assuming that an arbitrary representation acts on named qubit factors.

Solve $\operatorname{curl}\varphi_R = \pi/A_R - \pi 1_{q_R}$ with zero boundary potentials. The required curls sum to zero; eliminating leaves of the connected dual spanning tree assigns each parent edge to meet its leaf curl, and leaves the root equation automatically. Use the corrected links $F_a(v) = \prod_R e^{i\varphi_R(a\text{-edge})Q_R(v)}\Pi(A_v)$ and similarly F_b. Define

$$p|c,\psi\rangle = |c - e_1, F_a(c-e_1)\psi\rangle, \qquad q|c,\psi\rangle = |c - e_2, F_b(c-e_2)\psi\rangle.$$

Covariance brings each correction in a plaquette to its base. Commuting projections add their angles; at a source $e^{i\pi Q_R}$ cancels the original defect. The remaining holonomy is

$$e^{i1_B\pi Q_B/A_B}e^{i1_T\pi Q_T/A_T}, \qquad \|[p,q] - s_0\| \le h_M := \pi/A_B + \pi/A_T \le 8\pi/M^2,$$

where $s_0 = \Pi((1\ 2))$. Put $t = T_{e_1}$, $d = L_{E_{2\leftarrow 1}}$, $x = \widetilde{X}_0$, $z = \widetilde{Z}_0$. The protected identities give

$$ptp^* = qtq^* = T_{e_2}, \quad pT_{e_2}p^* = T_{e_3}, \quad pdp^* = qdq^* = L_{E_{3\leftarrow 2}}.$$

Also $ps_0p^* = qs_0q^* = B_0 = \Pi((2\ 3))$ and $p^2s_0p^{-2} = \Pi((3\ 4))$, so the independent-$s_0$ Houghton words are exact. The point swaps give $B_0s_0(pdp^*)s_0B_0 = d$, and similarly with q, while s_0p, s_0q commute with t. The genuine affine identities give $t^2 = I$, commuting coordinate translations, $dT_{e_1}d^{-1} = T_{e_1}T_{e_2}$ and $[d,T_{e_2}] = [d,T_{e_3}] = I$. Every base/correction and d commute with x,z; their onsite relations and $J = -I$ are exact.

For completeness, the entire 27-word certificate has the following normalized bounds, with $h = h_M$, $K = \max_{u\neq 0}|c(u)|$. Rows containing s use its actual value $[p,q]$.

row	word	bound	row	word	bound
1	s^2	h	2	$(spsp^{-1})^3$	$3h$
3	$[s,p^2sp^{-2}]$	$2h$	4	$psp^{-1}qs^{-1}q^{-1}$	h
5	t^2	0	6	$[t,sp]$	h
7	$[t,sq]$	h	8	$[d,Bsp]$	$2h$
9	$[d,Bsq]$	$2h$	10	$[t_1,t_2]$	0
11	$dt_1d^{-1}(t_1t_2)^{-1}$	0	12	$[d,t_2]$	0
13	$[d,t_3]$	0	14–19	$[p,x],[p,z],[q,x],[q,z],[d,x],[d,z]$	0
20–22	$x^2,z^2,(xz)^4$	0	23	$[x,txt^{-1}]$	$\le \min(1,2K)$
24	$[x,tzt^{-1}]$	K	25	$[z,txt^{-1}]$	K
26	$[z,tzt^{-1}]$	$\le \min(1,2K)$	27	$[J,t]$	0

The first four words have $2,6,4,2$ occurrences of $s^{\pm 1}$; rows $6,7,8,9$ have $2,2,4,4$, after expanding B. Replace each by $s_0^{\pm 1}$ and telescope, then divide by two. This proves every host bound. The zero rows follow from the displayed exact affine and onsite identities. For the four lamp rows, t sends zero to e_1 and the g-block labels are $0, g^{-1}e_1$. Uniform g makes the latter uniform on $W\setminus\{0\}$, with equal stabilizer multiplicities. Thus rows $24,25$ have exactly the moment in (P2.1); rows $23,26$ have at most 4^ν times it. Their clock factor is an identical multiplicity, as these words contain only t,x,z.

Proof of Theorem P2.1. Choose

$$M = \max\{48, \lceil N^{1/4} \rceil, \lceil (n_0 - 4)/4 \rceil\}, \qquad n = 4M + 4.$$

The complete space has dimension $D_N = M^2 |\mathrm{GL}(n,2)| 2^{3N}$, hence $\log D_N = 3N + O_\nu(\sqrt{N})$, since $n_0 = O_\nu(\log N)$ and $\log |\mathrm{GL}(n,2)| \le n^2$. Every nonzero host residual is at most $3h_M \le 24\pi M^{-2}$, whose 2ν-moment is $O_\nu(N^{-\nu})$. The character calculation gives the same order for all lamp rows; the marker residual is zero. This proves the theorem for $\nu \ge 1$. For $0 < \nu < 1$, use the $\nu = 1$ assignments and scalar concavity: $\tau((|R|^2)^\nu) \le \tau(|R|^2)^\nu$. □

Corollary P2.2 (Explicit rational powers)**.** *For integers $1 \le a \le b$, one fixed rational BCS game has*

$$c\varepsilon^{-a/b} 2^{-C(\log\log(1/\varepsilon))^2} \le \vec{d}_G(\varepsilon) \le \log D_G(\varepsilon) \le C'\varepsilon^{-a/b}.$$

Proof. For each normalized source residual replace its zero equation by $R = VS^a$, $S^b = 0$, with V unitary and S a positive contraction. At weighted precision δ, the spectral measure $\mathrm{Tr}(\rho 1_B(S))$ gives $|S^a|_\rho \le |S^b|_\rho^{a/b}$, without requiring $[S, \rho] = 0$. Thus $|R|_\rho \le 2\delta^{a/b}$. The compiler's original decoded words have actual mirrors, inverse mirrors and state commutators $O(\sqrt{\varepsilon})$, within $e = O(\varepsilon^{a/(2b)})$. The marker gives fixed visibility, so the explicit budget behind (P1.1) proves the lower bound on every original mixed state. For the upper bound take moment order $\nu = b/a$, extend each square-matrix polar partial isometry to a unitary $R = V|R|$, and set $S = |R|^{1/a}$. The first equation is exact and $\|S^b\|_{2,D_N}^2 \le CN^{-b/a}$. Flat commutators vanish. Compiler completeness gives loss $O(N^{-b/a})$ with only a constant dimension multiplier and a maximally entangled original state. An integer $N \asymp \varepsilon^{-a/b}$ covers all small losses. Upper losses tend to zero, and the divergent lower bound excludes finite perfection. □

These are finite-moment models, not an operator-norm endpoint theorem: for $u_* = (1, 0, \ldots, 0)$ all character signs are -1, so $X_{u_*} = X_0$, $Z_{u_*} = Z_0$ and the mixed residual has norm one on an exceptional affine block. In contrast its higher finite moments decay. Choosing a uniformly small-bias list gives a separate operator-norm upper estimate $\log D_\eta = O(\eta^{-5/2})$: independent characters obey $\Pr(|c(u)| > \kappa) \le 2e^{-N\kappa^2/2}$ by $\cosh t \le e^{t^2/2}$. A union bound gives a list for $N \ge 2\kappa^{-2} \ln(4(2^n - 1))$, selectable by finite exact search. Take $M \asymp \eta^{-1/2}$, $n = 4M + 4$, $\kappa = \eta/2$ in the table. This neither proves nor refutes a quadratic operator-norm cost.

The distinction between fixed-word moments and operator-norm control matters in the surrounding approximation literature. Ozawa–Rørdam–Sato [58] prove qualitative quasidiagonality for elementary amenable groups, without the dimension rate needed here. Gowers–Hatami [32] study normalized Schatten approximation, and uniform pair-defect stability, while Fournier-Facio–Rangarajan [29] establish uniform Ulam stability for lamplighters and related groups. Becker–Lubotzky [7] and Arzhantseva–Păunescu [2] address stability in their specified Hilbert–Schmidt or Hamming metrics. None of these hypotheses turns a small exceptional fixed-relator block into a small operator norm. No quadratic general filling estimate is assumed from the Houghton literature [44]; the proof uses the explicit clock regions and their checked affine lift.

P3 Polynomial localization and stopped products

Throughout the power part $\xi = \sqrt{\rho}$, $|R|_\rho = \|R\xi\|_2$. A weighted precision-δ model has exact bounded generator domains, defining residuals at most δ, and $\|[Y, \xi]\|_2 \le \delta$ for every generator. Constants depend on the fixed finite presentation, not dimension, state or horizon. For a contraction word v of length l, Lemma 2.2 gives

$$|Rv|_\rho \le |R|_\rho + \|R\| \|[v, \xi]\|_2, \qquad \|[v, \xi]\|_2 \le l\delta. \tag{P3.1}$$

Adjoint and translated fixed-polynomial residuals obey the same rule with fixed constants.

P3.1 The complete clock interface

Use two commuting complete Cantor-clock copies, fuel and time, each with positive contraction A, projection P, unitary T and reflection X satisfying

$$[A,P]=0, \qquad TAT^*=4A-3P, \qquad XPX=I-P, \qquad [X,TAT^*]=0.$$

Require U unitary, commuting with the entire fuel copy, with the *unguarded* increment

$$U=X_t\big[(I-P_t)+(T_tUT_t^*)P_t\big].$$

The reset projection Z commutes with A_f, A_t and, for $W=T_fT_t$, obeys

$$Z=(I-P_f)+P_f(I-P_t)WZW^*. \tag{P3.2}$$

This is the reset relation (5.6); the increment used here is the unguarded one of Proposition 4.5, not cancellation of a small guard. Write $Y_h=U^{-h}YU^h$ and $Y_+=U^*YU$. Every instantaneous working contraction commutes with A_f, A_t. Static inputs T_0, S_0 also commute with the full clock copies and U; they need not commute with ρ.

Proposition P3.1 (Polynomial clock localization)**.** *For $H\ge 2$, choose $H<2^k\le 2(H+1)$, put $P_{f,i}=T_f^iP_fT_f^{-i}$ and similarly for time, and take the ordered contraction*

$$Q=\left(\prod_{i<k}P_{f,i}\right)(I-P_{f,k})\left(\prod_{i<k}(I-P_{t,i})\right).$$

There are fixed C,a such that, for $r=C(H+1)^a\delta$,

$$\begin{aligned}
&|(I-Z)Q|_\rho\le r, && |Z_hQ|_\rho\le r \quad (1\le h\le H),\\
&|[Y_h,Q]|_\rho\le r \quad (0\le h\le H), && |T_0|_\rho\le 2^{k+1/2}\|T_0Q\xi\|_2+r,\\
&\|S_0Q\xi\|_2\le |S_0|_\rho+r.
\end{aligned}$$

Here Y ranges over any fixed finite list of instantaneous working contractions. No inter-time commutation is imposed, and Q need not be a projection.

Proof. Linear spectral repair in Lemma 4.2, with (4.4)–(4.5), gives raw digit, digit commutator and triangular flip errors to depth k at most $C(k+1)^3 4^{k+1}\delta$. Replacing a raw digit by $p(s^i(\widehat{A}))$, using the Hilbert–Schmidt Lipschitz estimate of Lemma 2.3 with constant at most $2\cdot 4^i$, and replacing it back gives the same bound for $[Y,P_i]$. Indeed $[Y,\widehat{A}]$ and $[\widehat{A},\xi]$ are $O(\delta)$ by linear repair and (P3.1). No commutation of Y with a clock shift or flip is used.

Set $V=T_tUT_t^*$, $B=(I-P_t)+VP_t$, $W_0=X_tU$. The increment says that $B-W_0$ and its adjoint have weighted norm $O(\delta)$; W_0 is unitary and $\|B\|\le 2$. The literal identities

$$BB^*=I-P_t+VP_tV^*, \qquad BB^*-I=(B-W_0)B^*+W_0(B^*-W_0^*)$$

and right transport give $|VP_tV^*-P_t|_\rho=O(\delta)$ and then $|[V,P_t]|_\rho=O(\delta)$. This is a linear vector estimate, not a square root of an expectation. Expanding B^*P_jB and using triangular flips yields, to the preceding weighted error,

$$U^*P_0U=I-P_0, \quad U^*P_jU=(I-P_0)P_j+P_0T_t(U^*P_{j-1}U)T_t^* \quad (j\ge 1).$$

The canonical output bit is $F_j = P_j + (I - 2P_j)\prod_{i<j} P_i$. Substituting the preceding canonical expression and reducing digit squares and commutators gives $E_j \le cE_{j-1} + c(j+1)^b 4^j\delta$, hence $E_j \le C_0^j(j+1)^{b+1}\delta$ for a fixed $C_0 > \max(c, 4)$. Each F_j has bounded coefficient mass, and a prefix or complementary prefix has mass at most 5^k and word length polynomial in k. Reduction in the commuting Boolean algebra gives a single translated raw binary prefix with error $C_*^k\,\mathrm{poly}(k)\delta$.

At the next physical translation use this raw prefix again. Induction through $h \le H$ adds $hC_*^k\,\mathrm{poly}(k)\delta + \mathrm{poly}(h,k)\delta$ by (P3.1); it does not multiply C_*^k at every time. Since $2^k \le 2(H+1)$, every cost is a fixed polynomial in H. Unwind (P3.2) through the first k fuel ones and the next fuel zero. On Q, the terms before depth k vanish, the depth-k term equals Q, and its remainder is killed by that zero. On a translated time residue $1 \le h \le H < 2^k$, one of the first k time bits is one, so the reset term vanishes. The expansion has depth $k+1$ and coefficient mass only exponential in k. This proves the first two estimates. The conjugate U^hQU^{-h} is the unchanged fuel prefix times a translated time prefix, up to this polynomial error. It commutes with instantaneous Y by the digit estimate. Conjugate back and pay $O(H\delta)$ to obtain $[Y_h, Q]$. Moving ξ across Q proves the S_0 estimate.

For the input charge, read the $2k+1$ specified raw bits in order as successive projective measurements, with outcome columns Q_b. Eliminating the leftmost projection gives exactly $\sum_b Q_b^*Q_b = I$, irrespective of digit commutation. Thus $\sum_b \|Q_bT_0\xi\|_2^2 = |T_0|_\rho^2$. A triangular flip exchanges its own bit and preserves later specified bits to the digit error. For a fixed later suffix, move this flip from left to right; its commutators with that suffix and T_0, its own-bit error and its commutator with ξ bound the difference of the two outcome *amplitudes* by $e = \mathrm{poly}(H)\delta$. For $|u - v| \le e$, $\sqrt{u^2+v^2} \le \sqrt{2}u + e$. Eliminate bits successively, using digit commutators when reordering. The main factor is $2^{(2k+1)/2}$; a crude error factor $2^{2k+1} \le 8(H+1)^2$ remains polynomial. Finally commute static T_0 through the selected column. This proves the input charge and every assertion. □

For $\phi = Q\xi$, any fixed local residual translated to a time at most H and inserted in an ordered history has bound

$$\eta_H = C(H+1)^a\delta. \tag{P3.3}$$

Increase a once to cover the finite list. A shifted history has expanded word length $O(H^2)$; use (P3.1). Reset terms are removed by commuting Q across the current polynomial and history, applying Z_hQ, and paying for the right suffix. We never normalize a retained branch or sum over an exponential tree of machine histories.

P3.2 The norm estimate behind stopping

Lemma P3.2 (Stopped products)**.** *Suppose exact columns L_h, K_H, B_j satisfy $\sum_h L_h^*L_h + K_H^*K_H + \sum_j B_j^*B_j = I$, with $\|\phi\| \le 1$. Let C_h be contractions and positive numbers M_h, N_h satisfy*

$$\|L_hT_0\phi - M_h^{-1}C_hL_h\phi\| \le r, \quad \|L_hS_0\phi - (2N_h)^{-1}L_h\phi\| \le r,$$

$$\|K_HT_0\phi - M_H^{-1}C_HK_H\phi\| \le r, \qquad \left(\sum_j \|B_jT_0\phi\|^2\right)^{1/2} \le r.$$

If $M_h \ge 1$, then with $R_H = \max_{h<H} 2N_h/M_h$,

$$\|T_0\phi\| \le M_H^{-1} + R_H(\|S_0\phi\| + \sqrt{H}r) + C\sqrt{H+1}\,r.$$

Proof. Put $a_h = \|L_h\phi\|$, $b_h = \|L_hS_0\phi\|$. The stop estimate gives $a_h \le 2N_h(b_h + r)$, and the other gives $\|L_hT_0\phi\| \le M_h^{-1}a_h + r$. Euclidean summation and $\sum_h b_h^2 \le \|S_0\phi\|^2$ bound all

stopped columns by $R_H(\|S_0\phi\| + \sqrt{Hr}) + \sqrt{Hr}$. The terminal column is at most $M_H^{-1} + r$, and the bad columns have aggregate at most r. Apply exact completeness to $T_0\phi$ and the Euclidean triangle inequality. □

For constant products $M_h = m^h$, $N_h = n^h$, the same proof with stop coefficient n^{-h-1} gives $\|T_0\phi\| \le C_{m,n}[m^{-H} + (n/m)^H(\|S_0\phi\| + \sqrt{Hr}) + \sqrt{Hr}]$ when $m \le n$. Dividing each update by its actual product makes errors additive. With $H = \log_n(1/\delta) - O(\log\log(1/\delta))$, the resulting calibrated exponent is $\log m/\log n$, including $m = n$. This is the constant-multiplier specialization of the stopped estimate; the variable products below use the same complete-column argument.

P4 Scaled counters and a stream with polynomial delay

P4.1 A finite rational presentation for a counter program

Take a finite deterministic register program with nonnegative counters $n_1, \dots, n_r$, initially zero, and finite controls Σ, initially σ_0. Instructions are $\mathrm{INC}(i,\tau)$, $\mathrm{DECJZ}(i,\tau_0,\tau_1)$ and $\mathrm{NOP}(\tau)$. A zero test decrements only on its nonzero branch. Two designated output NOP controls emit zero and one; all other instructions emit no pulse. Each instruction changes at most one counter by one, so $n_i(h) \le h$ at physical time h.

Introduce positive contractions A, Y_i, W_i and projections E_i, Q_σ, F. They represent reciprocal age, scaled counters, nonzero witnesses, zero tests, controls and pulse occurrence. Abbreviate $G = Q_{\mathrm{one}}$. Impose

$$\begin{gathered} Q_\sigma Q_\tau = 0 \quad (\sigma \ne \tau), \qquad \sum_\sigma Q_\sigma = I, \qquad F = Q_{\mathrm{zero}} + Q_{\mathrm{one}}, \\ Y_i E_i = 0, \qquad (Y_i - A)(I - E_i) = W_i(I - E_i). \end{gathered} \tag{P4.1}$$

The instantaneous list is exactly

$$Z, A, Y_i, W_i, E_i, Q_\sigma, F, T_0, S_0, C, D, H_{\mathrm{run}}, P_{\mathrm{stop}}.$$

Impose pairwise commutation in this list and commutation of every listed variable with A_f, A_t. Only static T_0, S_0 additionally commute with the full fuel/time copies and U. The unitary polar variable and the six original source unitaries are excluded; there are no other inter-time or cross-checker commutations.

For each instruction use its selector Q_σ, or $Q_\sigma E_i, Q_\sigma(I - E_i)$ for a test. Let $B_\tau(Q,E)$ be the sum of selectors entering τ, and $\Delta_i(Q,E)$ the sum incrementing i minus the sum decrementing it. These are fixed rational polynomials. Impose

$$\begin{gathered} Z(A - I) = 0, \quad ZY_i = 0, \quad Z(Q_\sigma - 1_{\sigma=\sigma_0} I) = 0, \\ A_+(I + A) = A + Z_+, \\ (Y_i)_+(I + A) = (I - Z_+)(Y_i + A\Delta_i), \\ (Q_\sigma)_+ = Z_+ 1_{\sigma=\sigma_0} I + (I - Z_+)B_\sigma. \end{gathered} \tag{P4.2}$$

All coefficients are rational and all domains are accepted by Theorem 3.1. The order "next variable times $(I+A)$" is part of the relation. Zero tests, witnesses and F need no independent updates.

Proposition P4.1 (Exact scalar counter histories)**.** *These relations have the prescribed histories on every finite fuel/time block, including a reset interrupting a computation, with no additional tensor factor for counter size or running time.*

Proof. At age h since reset assign

$$A_h = \frac{1}{h+1}, \quad Y_{i,h} = \frac{n_i(h)}{h+1}, \quad E_{i,h} = 1_{n_i(h)=0}, \quad W_{i,h} = \frac{\max(n_i(h)-1,0)}{h+1},$$

and $Q_{\sigma,h} = 1_{\sigma(h)=\sigma}$, with F_h, G_h their indicated outputs. All domains hold because $0 \le n_i(h) \le h$. Both zero-test relations are exact. Off reset, multiplying $n_i(h+1)/(h+2)$ by $1 + 1/(h+1)$ gives $(n_i(h) + \Delta_i)/(h+1)$; the age equation is the same identity. At a next reset the values are $A_+ = 1$, $Y_+ = 0$ and initial control, exactly as (P4.2) requires even for large current counters. On an N_c-bit clock block use age modulo its fuel-dependent period: 2^k for k initial fuel ones followed by zero, and 2^{N_c} for all-ones fuel, reset only at all-zero time. These periods divide the full time cycle. The reset recursion (P3.2) holds by stripping the first fuel/time bits until a fuel zero or a time one is met. Every transition remains exact across the cyclic boundary. All workspaces are diagonal scalar functions on the existing clock/source blocks, so their commutations are exact and they add no matrix factor. □

P4.2 Noncommuting same-state stability

Proposition P4.2 (Polynomial counter stability)**.** *Fix a weighted precision-δ model and horizon H. Put $\phi = Q\xi$ as in Proposition P3.1 and $\eta_H = C(H+1)^a\delta$ as in (P3.3). Let $\sigma(h), n_i(h)$ be the unique classical run. The prescribed good history has reciprocal-age error $O((h+1)\eta_H)$ and scaled-counter error $O((h+1)^2\eta_H)$. Wrong control and zero-test columns have total amplitude on $T_0\phi$ bounded by a fixed polynomial in H times δ, even with an arbitrary current-time stopping projection. No inter-time commutation or positive lower bound on a retained branch's norm is required.*

Proof. Write $q_h = (Q_{\sigma(h)})_h$, start $K_0 = I$, and split successively

$$B_h^c = (I - q_h)K_h, \quad V_h = q_h K_h, \quad B_h^z = (I - e_h)V_h, \quad V_h' = e_h V_h,$$

where $e_h = (E_i)_h$ for a truly zero test and $I - (E_i)_h$ for a truly nonzero test. Without a test take $V_h' = V_h$, $B_h^z = 0$. Finally set

$$L_h = (P_{\text{stop}})_h V_h', \qquad K_{h+1} = (I - (P_{\text{stop}})_h)V_h'. \tag{P4.3}$$

Every factor is an exact projection. The two-outcome identity $(pK)^*pK + ((I-p)K)^*(I-p)K = K^*K$ telescopes to

$$\sum_{h<H} \left((B_h^c)^* B_h^c + (B_h^z)^* B_h^z + L_h^* L_h\right) + K_H^* K_H = I. \tag{P4.4}$$

There is one retained path and at most three discarded columns per time, irrespective of commutation.

Put $a_h = 1/(h+1)$, $b_{i,h} = n_i(h)/(h+1)$. We prove

$$\begin{aligned} \|(A_h - a_h I)K_h\phi\| &\le C(h+1)\eta_H, \\ \|((Y_i)_h - b_{i,h}I)K_h\phi\| &\le C(h+1)^2\eta_H, \\ \|(I - q_h)K_h\phi\| &\le C\eta_H. \end{aligned}$$

Initialization follows from (P4.2), the current $[Z, Y]$ relations and localization. The reason errors do not amplify is the pair of literal ordered identities, for positive contractions A, A', Y, Y', scalars a, b, $s \in \{-1, 0, 1\}$ and $a' = a/(1+a)$, $b' = (b + sa)/(1+a)$:

$$(1+a)(A' - a'I) = [A'(I+A) - A] + (I - A')(A - aI), \tag{P4.5}$$

$$(1+a)(Y' - b'I) = [Y'(I+A) - Y - sA] + (Y - bI) + (sI - Y')(A - aI). \tag{P4.6}$$

Expansion in the displayed order verifies them; no operators at different times have been interchanged. Their multiplier norms are at most one and two respectively, and $1 + a \geq 1$.

Restrict the current invariants to $K_{h+1}\phi$ by the current $q_h, e_h, P_{\text{stop}}$ commutations, paying $O(\eta_H)$. Selection fixes the current control, kills the other controls by orthogonality, and fixes the selected zero test. Thus Δ_i acts as its classical increment $s_i(h)$ to $O(\eta_H)$. Eliminate resets by (P3.3); the age and counter update brackets in (P4.5)–(P4.6) are then $O(\eta_H)$ on that vector. If e_h, d_h are their errors, these identities give

$$e_{h+1} \leq e_h + C\eta_H, \qquad d_{h+1} \leq d_h + 2e_h + C\eta_H.$$

Hence $e_h = O((h+1)\eta_H)$, $d_h = O((h+1)^2\eta_H)$. The factor two multiplies the separately additive age error, not the old counter error. Selection has already fixed the current control, so the successor polynomial gives the next prescribed control with $O(\eta_H)$ error, independently of the preceding control bound. This closes the induction.

The bad control column has norm $O(\eta_H)$. For a bad zero-test vector $v = B_h^z\phi$, restrict the current invariants again using only current commutators. If $n_i(h) > 0$, the wrong branch has $(E_i)_h v = v$ and its zero relation gives $\|(Y_i)_h v\| \leq C\eta_H$. Since $b_{i,h} \geq 1/(h+1)$, the counter invariant yields $\|v\| \leq C(h+1)^3\eta_H$. If $n_i(h) = 0$, the wrong branch has $(I - E_i)_h v = v$; its nonzero-witness relation gives

$$\|((W_i)_h + a_h I)v\| \leq C(h+1)^2\eta_H.$$

Exact positivity of W_i makes the left side at least $a_h\|v\|$, giving the same bound. This is linear vector coercivity. Its inverse gap is charged only to the discarded branch and never inserted into a later good recurrence. Static T_0 commutes with U and the instantaneous control, test and stop projections. Expand its commutator through the bad-column word, whose shifted letters have total expanded length $O(H^2)$, and use (P3.1). The same polynomial bound holds for $B_h^c T_0\phi, B_h^z T_0\phi$. Euclidean summation over $O(H)$ columns gives the asserted aggregate. The vectors have never been renormalized. Every constant is uniform in dimension and ρ. □

P4.3 A pulse stream for any computable limit

Lemma P4.3 (Polynomial physical delay)**.** *From a total computable rational sequence $a_s \to \beta \in [0,1]$ one can effectively construct a finite counter program as above emitting infinitely many bits $w_1, w_2, \ldots$ with*

$$k(j) := \sum_{l \leq j} w_l = \beta j + o(j), \qquad t_j \leq C(j+1)^2,$$

where t_j is physical time just after the jth pulse. Neither a convergence modulus nor a computation-time bound for a_s is needed.

Proof. Clamp each rational to $[0,1]$. Use foreground and background coroutines with disjoint finite counter sets. The active rational starts at $1/2$. For active p/q, $q \geq 1$, the foreground emits repeated periods of p ones followed by $q - p$ zeros. Stored p, q are preserved by scratch-counter copy/countdown loops. Copying or restoring unary counters of size at most q uses $O(q+1)$ elementary INC/DECJZ instructions, so each gap between adjacent pulses costs that order, apart from a stage switch.

After every pulse give the background exactly one elementary instruction, with fixed dispatch overhead. It computes the next $a_s = p_s/q_s$, then performs padding with at least q_s^2 elementary instructions before signalling READY. Nested loops building and clearing q_s^2 in scratch suffice. Only at a completed foreground period boundary is a READY rational installed. Clearing/copying the old and new foreground registers costs $O(q_{\text{old}} + q_{\text{new}} + 1)$; all

arbitrary background scratch cleanup occurs through pulse-triggered background instructions. The next stage begins only after installation.

This is one finite program: take the product of the coroutine control sets and fixed mode flags. If there are several emitting locations, a bounded return-address counter routes them through the two designated output NOP controls and back with fixed overhead. It changes by elementary instructions and has a fixed maximum. Each step of a long background computation is an actual INC/DECJZ/NOP, not a unit-cost unbounded call. For completeness, finite counters compute any total computable rational name: encode the two tape halves as integer stacks in a fixed base. Push multiplies by the base and adds a digit; pop takes quotient and remainder. Elementary increment and zero-test/decrement loops with finitely many scratch counters implement both; machine state and current symbol are finite control. Empty stacks supply blanks. Integer input and rational numerator/denominator output use the same loops. This is the classical register-machine universality of Shepherdson–Sturgis and Minsky [63, 51]; no time efficiency is assumed.

If a stage were never installed, the current finite q would generate infinitely many pulses, so the background would receive infinitely many instructions. Its total computation and finite padding would finish, and a period boundary would follow. Thus every stage is eventually installed. Before installing denominator q, at least q^2 pulse-triggered instructions have occurred; at least q^2 pulses have already been emitted. Hence the current denominator at pulse count j is at most $\max(2, \sqrt{j})$. There is at most one installation between consecutive pulses: after an installation another READY requires a new pulse-triggered background step. Each gap therefore costs $O(\sqrt{j+1}+1)$ actual instructions, including copying, clearing and dispatch, and $t_j = O((j+1)^{3/2})$, in particular the claimed quadratic bound.

Every complete period has exactly its prescribed rational average; switches occur only between completed periods. At pulse prefix j the unfinished-period discrepancy is at most the current $q = O(\sqrt{j}+1)$. Given $\zeta > 0$, all sufficiently late installed rationals lie within ζ of β; earlier stages contribute finitely many pulses. Therefore

$$|k(j) - \beta j| \le C_\zeta + \zeta j + O(\sqrt{j}+1).$$

Divide by j and let $\zeta \downarrow 0$. This uses convergence alone, with no variation or monotonicity hypothesis. □

P5 Programmed calibration and matched game laws

Take the program of Lemma P4.3 for $\beta \in (0,1]$. For each normalized source residual R use a fresh fuel/time clock pair and program workspace from Section P4. Add positive contractions T_0, S_0, C, D, projections $H_{\mathrm{run}}, P_{\mathrm{stop}}$, and a unitary V. Use exactly the instantaneous and static commutation lists stated there, and impose

$$\begin{gathered}
R = VT_0, \qquad (I - H_{\mathrm{run}})P_{\mathrm{stop}} = 0, \\
Z(C - T_0) = 0, \quad Z(D - S_0) = 0, \quad Z(H_{\mathrm{run}} - I) = 0, \\
(H_{\mathrm{run}})_+ = Z_+ + (I - Z_+)(H_{\mathrm{run}} - P_{\mathrm{stop}}), \\
C_+ = Z_+T_0 + (I - Z_+)C[I + G(H_{\mathrm{run}} - P_{\mathrm{stop}})], \\
D_+ = Z_+S_0 + (I - Z_+)D[I + F(H_{\mathrm{run}} - P_{\mathrm{stop}})], \\
(D - I/2)P_{\mathrm{stop}} = 0, \qquad S_0 = 0.
\end{gathered} \tag{P5.1}$$

This list, the two complete clock lists, (P3.2), (P4.1)–(P4.2), exact domains and the declared commutations are the complete finite rational presentation. The old equations $R = 0$ are *replaced*, not also tested. Different residual checkers have no additional mutual workspace

or input commutations. The program does not depend on running or stopping and keeps computing after calibration stops. For an inhomogeneous estimate omit only the last test $S_0 = 0$.

Proposition P5.1 (Programmed retained-state calibration)**.** *Put $\eta = \delta + |S_0|_\rho$. In every weighted precision-δ model of the inhomogeneous presentation,*

$$|T_0|_\rho \le \eta^{\beta - o(1)}.$$

Precisely, for each $0 < \gamma < \beta$ there are constants $C_\gamma, \eta_\gamma > 0$, depending only on the fixed program, such that $|T_0|_\rho \le C_\gamma \eta^{\beta-\gamma}$ for $\eta < \eta_\gamma$. No effective modulus for these thresholds is asserted.

Proof. Let f_h, g_h be the classical pulse and one indicators, $g_h \le f_h$, and set

$$j(h) = \sum_{u<h} f_u, \quad k(h) = \sum_{u<h} g_u, \quad N_h = 2^{j(h)}, \quad M_h = 2^{k(h)}.$$

Use the columns (P4.3)–(P4.4). Selection of the prescribed control makes $F_h = f_h$ and $G_h = g_h$ on current selected histories to $O(\eta_H)$. The program invariants and bad columns are bounded by Proposition P4.2. Initialization and the running transition give $(H_{\mathrm{run}})_h K_h \phi = K_h \phi + O(\mathrm{poly}(H)\delta)$. To restrict it use current commutations and $P_h K_{h+1} = 0$ exactly. The multiplier $H_h - P_h$ is a difference of projections and has norm at most one even before repairing approximate order. Thus running errors add.

The normalized counter inductions give, uniformly through H,

$$K_h T_0 \phi = M_h^{-1} C_h K_h \phi + O(r_H), \qquad K_h S_0 \phi = N_h^{-1} D_h K_h \phi + O(r_H), \quad r_H = C(H+1)^b \delta. \tag{P5.2}$$

For the C step, left-restrict the old identity by the next current control, test and continuation projections. Commute only current C_h past these current projections. On this retained continuation, $G_h = g_h$ and $H_h - P_h = 1$ to the polynomial error. The update therefore says $C_{h+1} = 2^{g_h} C_h$ there. Divide by $M_{h+1} = 2^{g_h} M_h$. Every new residual has coefficient at most $M_h^{-1} \le 1$; the old error is merely left-restricted by a contraction. The D proof is identical with f_h, N_h. Idle instructions have factor one and add only a local error. The running, program, static-input and history-transport costs are fixed polynomials and there are H steps; enlarging b covers all of them. In particular no factor 2^H or product of zero-test denominators appears.

Multiply these identities by the current stop projection. Its current commutations and $(D_h - I/2)P_h = 0$ give

$$\|L_h T_0 \phi - M_h^{-1} C_h L_h \phi\| \le r_H, \qquad \|L_h S_0 \phi - 2^{-j(h)-1} L_h \phi\| \le r_H.$$

The terminal estimate follows from (P5.2); enlarge r_H to bound the bad-column aggregate as well. Lemma P3.2 now gives

$$\|T_0 \phi\| \le M_H^{-1} + R_H(\|S_0\phi\| + \sqrt{H} r_H) + C\sqrt{H+1} r_H, \qquad R_H = \max_{h<H} 2^{j(h)+1-k(h)}. \tag{P5.3}$$

This invocation is just Euclidean summation of exact measurement columns, so it introduces no inter-time or inter-column commutation. The integer $j(h) - k(h)$ is nondecreasing: a zero pulse increases it, a one pulse and idle work leave it unchanged. Consequently $R_H \le 2^{j(H)-k(H)+1}$ even for arbitrarily oscillating rational approximations.

For J pulses choose $H = t_J$. Then $H \le C(J+1)^2$, $j(H) = J$ and $k(H) = \beta J + o(J)$. Every error and initialized-branch factor is therefore a fixed polynomial in J. With $L = 1 + \log(1/\eta)$ take

$$J = \lfloor \log(1/\eta) - K \log L \rfloor,$$

where fixed K exceeds all required polynomial degrees by at least five. Localization gives $\|S_0\phi\| \le |S_0|_\rho + \mathrm{poly}(H)\delta$ and the input charge $|T_0|_\rho \le C(H+1)\|T_0\phi\| + \mathrm{poly}(H)\delta$. In (P5.3), $2^J\eta$ times every required polynomial is bounded. The result is at most a polynomial in J times $2^{-k(H)}$, plus $\eta\,\mathrm{poly}(J)$. Since $J = \log(1/\eta) - O(\log L)$ and $\beta \le 1$, both have the asserted power bound. The density limit yields its precise fixed-margin formulation, uniformly over every original state. Its convergence threshold need not be computable. □

Proposition P5.2 (Exact upper histories and their envelope)**.** *There is one fixed scalar function $s : [0,1] \to [0,1/2]$, with $s(0) = 0$, such that on each $T_0 = t$ eigenspace the non-penalty calibrator and program relations have simultaneous exact histories on every finite clock block, and*

$$s(t) = t^{1/\beta + o(1)} \quad (t \downarrow 0), \qquad s(t) \le C_v t^v \quad \textit{for every fixed } 0 < v < 1/\beta.$$

No continuity or endpoint bound with $v = 1/\beta$ is required.

Proof. For $t > 0$ set $\ell(t) = \lfloor \log(1/t) \rfloor$. Since $\beta > 0$, there are infinitely many ones and their cumulative count assumes every nonnegative integer. Choose the first physical age $h(t)$ with $k(h(t)) = \ell(t)$, let $J(t) = j(h(t))$ and put $s(t) = 2^{-J(t)-1}$. For $\ell(t) = 0$ use age zero. At age h since reset assign the program as in Proposition P4.1, and

$$C_h = 2^{k(\min(h,h(t)))}t, \quad D_h = 2^{j(\min(h,h(t)))}s(t), \quad H_h = 1_{h \le h(t)}, \quad P_h = 1_{h=h(t)}.$$

Before and at the stop, $C_h \le 2^{\ell(t)}t \le 1$, $D_h \le 1/2$. A continuing one pulse doubles both counters, a zero pulse only D, and an idle instruction neither. At the stop $D = 1/2$ exactly; afterward counters freeze and running turns off, while the program continues. At each reset they initialize. For $t = 0$, take $C = D = 0$, $H = 1, P = 0$ forever. On every finite block use age modulo its reset period. If that period ends before $h(t)$, no stop is asserted; all transitions still hold through the reinitialization. This covers arbitrarily delayed stops, interrupted runs, zero, one and exact dyadic boundaries. For large ℓ, $\ell = \beta J(t) + o(J(t))$ and $\ell = \log(1/t) + O(1)$, proving the asymptotic envelope. It gives the v-bound below some threshold; on the rest of $[0,1]$, enlarge C_v using $s \le 1/2$. Thresholds may be ineffective, but no such constant is a game coefficient. For a matrix residual take a square-matrix unitary polar extension $R = V|R|$, $T_0 = |R|$ and $S_0 = s(T_0)$. Use these scalar histories on each spectral/clock block. All current-time and static commutations hold, with no cross-residual commutations. The program workspace depends only on age; calibrator variables depend on age and this one T_0. They are operators on existing source/clock spaces, not additional tensor registers for values or time. Spectral calculus on a finite spectrum requires no continuity of s. □

Realization clause of Theorem P1.1. Combine the full checker for each of the 28 normalized source residuals, replacing their old zero tests, and apply Theorem 3.1. This is a single finite rational presentation and game, determined by the finite program for the promised rational name. Fix $0 < v < 1/\beta$ and take the source assignments of Theorem P2.1 at this fixed order. On each residual use Proposition P5.2. The only nonzero new errors are the penalties $S_0 = 0$ and inherited clock baselines, since

$$\|S_0\|^2_{2,D_N} \le C_v^2 \tau(|R|^{2v}) \le C_v' N^{-v}.$$

On an N_c-bit clock the Cantor defects are $O(4^{-N_c})$ and the deliberately unguarded increment has rank-one defect

$$(|0^{N_c}\rangle - |10^{N_c-1}\rangle)\langle 1^{N_c}|,$$

of flat norm $\sqrt{2}\,2^{-N_c/2}$. Choose $N_c = \lceil v \log N \rceil + O_v(1)$. There are only finitely many clock pairs, adding $O_v(\log N)$ to logarithmic dimension. All assignments are flat, hence all state commutators vanish. Thus

$$\text{flat precision} \le C_v N^{-v/2}, \qquad \log D_{\text{total}} \le 3N + O_v(\sqrt{N}) + O_v(\log N).$$

Compiler completeness gives a maximally entangled original strategy, a fixed dimension multiplier and loss at most $C_v N^{-v}$. The models exist for every sufficiently large integer N, so choosing an integer N at least a suitable constant times $\varepsilon^{-1/v}$ covers every small loss. Hence the limsup exponent of $\log D_G$, and of each entanglement resource, is at most $1/v$. The game is unchanged as v increases to $1/\beta$; take that limit only after the small-loss limsup to obtain at most β. The $v < 1$ concavity clause of Theorem P2.1 covers $\beta = 1$.

Conversely, start from any original finite mixed strategy of loss ε. Same-state soundness gives exact domains and precision $\delta = C\sqrt{\varepsilon}$ on its original marginal. The tested penalties and Proposition P5.1 give

$$|R|_\rho \le \delta + |T_0|_\rho \le \varepsilon^{\beta/2 - o(1)}.$$

The original six source unitaries are fixed decoded words, with actual opposite-party mirrors, inverse mirrors and square-root marginal commutators $O(\sqrt{\varepsilon})$. Since $\beta \le 1$, these fit in the source error $e = \varepsilon^{\beta/2-o(1)}$. The marker residual gives fixed visibility for small loss. The explicitly probability-charged extraction (P1.1) therefore yields

$$E_D^{A\to B}(\rho_{\text{original}}) \ge e^{-2+o(1)} = \varepsilon^{-\beta+o(1)}.$$

This holds for every original state, without pure-component selection. Take infima; $E_D \le \log D$ also gives the dimension lower bound. The maximally entangled upper models match it. Their losses tend to zero, so the value is one; a finite-dimensional perfect strategy would contradict the divergent all-state lower bound. Every positive Δ_2^0 real has a total computable rational name, completing realization. □

P6 Ceilings, explicit exponents and the remaining gap

P6.1 The three exact spectra

Proposition P6.1 (Effective ceilings)**.** *Every existing finite power exponent of* $\log D_G$*, pure original-state demand* p_G*, or formation demand* f_G *is* Δ_2^0*, for a fixed finite rational game of quantum value one. At every feasible error,*

$$f_G(\varepsilon) \le p_G(\varepsilon) \le f_G(\varepsilon) + 1.$$

The ceiling also holds for arbitrary positive question probabilities on a fixed finite support, by constant reweighting.

Proof. Fixed-dimensional strategy feasibility at rational loss is compact rational semialgebraic and decidable by real quantifier elimination, as in Lemma 2.4. Under the value-one promise, dimension search terminates, so $D_G(2^{-n})$ is computable. Computable rational approximations to $\log(1 + \log D_G(2^{-n}))/n$ converge to an existing power exponent; Shoenfield's limit lemma gives the asserted class.

For pure demand set $e_n = 2^{-n}$ and compute $D_G(e_n/4)$. Let U_n be its integer bit length plus two. A pure purification on Alice of any strategy at that dimension has Bob entropy at most its logarithm, so $p_G(e_n/2) + 1 \le U_n$. Set $t_n = (e_n/8)^2$ and $R_n = 2^{\lceil U_n/t_n \rceil}$. For decreasing Schmidt probabilities, $\lambda_i \le 1/i$. If entropy $S \le U_n$, its mass beyond rank R_n is at most $S/\log R_n \le t_n$. Truncation and normalization change half trace distance by at most $\sqrt{t_n}$ and give entropy at

most $S/(1-t_n)$: discard the positive tail entropy in the normalization formula. Compressing measurements to the Schmidt supports preserves valid POVMs and increases game loss by at most $\sqrt{t_n}$. Let q_n be the minimum entropy over pure rank-at-most-R_n strategies of loss at most e_n. It is computable on a compact rational semialgebraic domain: quantifier elimination decides box intersection and entropy is computable uniformly continuously in this fixed dimension (equivalently approximate $-x\log x$ by rational polynomials to dimension-adjusted error). The logarithm itself is not assumed semialgebraic. The feasible set is nonempty by the truncation argument, and near-minimizers at loss $e_n/2$ give

$$p_G(e_n) \le q_n \le p_G(e_n/2)/(1-t_n).$$

Thus computable rational approximations to $\log(1+q_n)/n$ converge to an existing pure power exponent. There is no assumption about an effectively bounded rank before truncation or an effective convergence modulus.

For formation, take a finite pure decomposition of any low-loss mixed strategy and record each component's loss and pure entropy. Minimize its average entropy over the finite simplex with mean loss at most ε. An optimal extreme point has at most two components, because there are only total-weight and mean-loss constraints. Coherently superpose those two states with common local orthogonal labels, and use block-diagonal original measurements. This is a pure strategy of the same mean loss and entropy $h_2(w)+wS_1+(1-w)S_2$, at most one plus the selected average. Taking infima over decompositions and original strategies proves $f_G \le p_G \le f_G+1$, hence the same positive power exponents. This uses the single-copy convex roof [9], not regularized entanglement cost; the pure-state rate is its entropy [8].

Finally replace any positive finite question law by uniform weights on its support. The two losses bound each other at fixed positive multiplicative constants. Their resource demands bracket at constant rescalings of ε, which preserve an existing exponent. This proves the extension. □

Classification clause of Theorem P1.1. The ceilings above supply necessity for each of the three resources. For each game constructed in Section P5, the all-original-state directional-rate lower bound also lower-bounds formation and pure entropy, while the maximally entangled upper strategies match it. Thus every Δ_2^0 exponent in $(0,1]$ occurs for all three resources. For $d_G^{\rightarrow}$ this proves only inclusion. Neither existence of limits for all games nor exclusion of exponents above one is asserted. □

Proposition P6.2 (A non-limit-computable usable exponent forces irreversibility)**.** *If a fixed game has $d_G^{\rightarrow}(\varepsilon)=\varepsilon^{-\beta+o(1)}$, $0<\beta<\infty$, with $\beta\notin\Delta_2^0$, there are $\gamma>0$ and $\varepsilon_i\downarrow 0$ such that every efficient original strategy sequence, meaning $E_D^{A\to B}(\rho_i)\le\varepsilon_i^{-\beta+o(1)}$ at losses at most ε_i, satisfies*

$$E_F(\rho_i)\ge\varepsilon_i^{-\beta-\gamma+o(1)},\qquad \frac{E_F(\rho_i)}{E_D^{A\to B}(\rho_i)}\ge\varepsilon_i^{-\gamma+o(1)},$$
$$S(\rho_i)\ge\varepsilon_i^{-\beta-\gamma+o(1)},\qquad \operatorname{rank}\rho_i\ge 2^{\varepsilon_i^{-\beta-\gamma+o(1)}}.$$

The justified general arithmetic ceiling for β is left-c.e. relative to $0''$; this is weaker than Δ_2^0.

Proof. Suppose for every fixed $a<1$ the limsup of

$$\frac{\log(1+p_G(\varepsilon^a))-\log(1+d_G^{\rightarrow}(\varepsilon))}{\log(1/\varepsilon)}$$

were nonpositive. Rescaling $\delta=\varepsilon^a$ gives pure limsup at most β/a, and $d_G^{\rightarrow}\le p_G$ gives pure liminf at least β. Let $a\uparrow 1$. Its exponent would exist and equal β, contradicting

Proposition P6.1. Therefore some fixed slack has a positive gap; choose γ below that gap and a subsequence with $1 + p_G(\varepsilon_i^a) \geq (1 + d_G^{\rightarrow}(\varepsilon_i))\varepsilon_i^{-\gamma}$. For every low-loss mixed state, $E_F \geq f_G(\varepsilon_i) \geq p_G(\varepsilon_i) - 1 \geq p_G(\varepsilon_i^a) - 1$. This proves the formation lower bound and its ratio for efficient states. Entropy concavity gives $E_F \leq S(\rho_B)$; hashing [23] gives $E_D^{A\to B} \geq S(\rho_B) - S(\rho_{AB})$. Hence $S(\rho_{AB}) \geq E_F - E_D$, of the same larger polynomial order, and $S(\rho) \leq \log \operatorname{rank} \rho$ proves the last bound.

For the general ceiling use the regularized-instrument formula and compact finite-dimensional optimization [26, Theorem 11.14]: $d_G^{\rightarrow}(2^{-n})$ is uniformly right-c.e. relative to $0'$. The increasing computable transformation $x_n = \log(1 + d_G^{\rightarrow}(2^{-n}))/n$ has the same property. If $x_n \to \beta$, then $\beta = \sup_m \inf_{n\geq m} x_n$. Each inner infimum is right-c.e. in $0'$. For rational r, the assertion $r < \beta$ is equivalent to existence of m and rational $r' > r$ with $\inf_{n\geq m} x_n \geq r'$, a Σ_3^0 condition. Its lower cut is therefore c.e. in $0''$, as claimed. This does not compute the mixed-state infimum or rule out the gap. □

P6.2 Explicit examples, closures and strict extension

Corollary P6.3 (Computable families)**.** *Every algebraic $\beta \in (0, 1]$, every $\log m/\log n \in (0, 1]$ for integers $m, n \geq 2$, and every admissible finite algebraic mixing*

$$\beta = \frac{\sum_i \alpha_i \log m_i}{\log n} \in (0, 1], \qquad \alpha_i \in [0, 1] \text{ real algebraic}, \quad m_i, n \geq 2,$$

has matching usable and log-dimension exponents. Finite products of such exponents are also realized.

Proof. Root isolation, logarithms of positive integers and the indicated finite arithmetic are uniformly computable. Each expression has a total computable convergent rational name, so Theorem P1.1 applies. Only the averaged displayed admissibility condition is required; neither $\prod_i m_i \leq n$ nor $\sum_i \alpha_i \leq 1$ is assumed. For example use rational root-isolating approximations to the root $\sqrt{2}/2$ of $2x^2 - 1$ in $(7/10, 3/4)$. These approximations, not an irrational coefficient, are computed by its finite program. □

The rotation description of the algebraic examples uses $b_h = \lfloor (h + 1)\beta \rfloor - \lfloor h\beta \rfloor$, whose prefix sum is $\lfloor j\beta \rfloor$; this is the classical irrational-rotation coding of Morse–Hedlund [52]. For an algebraic irrational, the elementary estimate $\operatorname{dist}(r\beta, \mathbb{Z}) \geq cr^{1-d}$ follows from a degree-d integer polynomial: its nonzero value at p/r is at least r^{-d}, and a bounded derivative gives the separation. A rational isolated-root register has $|B - \beta I|_\rho \leq C|f(B)|_\rho$ by the bounded inverse of $f(x)/(x - \beta)$. Such a rotation can select the multipliers with polynomially charged discarded carries. The general program proof above replaces the need for this particular root and gap interface; it handles computable rotations with gaps too small for that estimate as well.

For instance set $n_1 = 2$, $n_{k+1} = 2^{2n_k}$ and $\beta_* = \sum_{k\geq 1} 2^{-n_k}$. Its effective tail makes it computable; with $q_k = 2^{n_k}$, $0 < \operatorname{dist}(q_k\beta_*, \mathbb{Z}) \leq 2^{n_k+1-q_k^2}$. It is irrational, since its difference from a dyadic partial sum would otherwise be at least $1/(qq_k)$ for a fixed rational denominator q. Thus a uniform subexponential inverse-gap estimate fails at $J = q_k$, while Theorem P1.1 still realizes β_*. This limits that carry estimate, not realizability.

The scalar interfaces themselves admit composition, more specifically than the existence corollary. If checker i has inhomogeneous exponent β_i, remove the first penalty test and identify $S_1 = T_2$. Then $|S_1|_\rho \leq |T_2|_\rho + \delta \leq \delta^{\beta_2 - o(1)}$, and applying the first inhomogeneous estimate at $\eta = \delta + |S_1|_\rho$ gives exponent $\beta_1\beta_2$. The upper assignment uses the composed spectral penalty $s_2(s_1(T_1))$. For fixed $v_i < 1/\beta_i$ it is at most $CT_1^{v_1v_2}$. The moment theorem applies at that order; size *both* clock pairs at the final precision, $N_c \geq v_1v_2 \log N + O(1)$.

The extra logarithmic dimension remains $O(\log N)$. Taking v_i to their endpoints proves the product law for these specified checker interfaces, not an unrestricted assertion about products of arbitrary games.

No algebraic irrational equals $\log m/\log n$: otherwise $n^\beta = m$ contradicts Gelfond–Schneider [31]. The qualification "irrational" is essential. Baker's algebraic-coefficient linear independence for rationally independent logarithms [4] also makes $(\sqrt{2}/2)\log 2/\log 3$ transcendental. It does not show that this particular number is outside every log-ratio class: cross-multiplication would introduce products of logarithms, beyond that linear theorem. The following strict extension needs no such unproved independence.

Proposition P6.4 (Strictly beyond finite explicit expression families)**.** *The realized exponents include a noncomputable real and a computable real outside all finite algebraic mixing and product expressions just described.*

Proof. For a fixed noncomputable c.e. halting set K, take its finite stage approximations K_s and

$$a_s = \frac{1}{4} + \frac{1}{4}\sum_{\substack{e\in K_s\\ e\le s}} 4^{-e-1}, \qquad \beta_K = \lim_s a_s \in [1/4, 1/3].$$

This is a total computable rational name. If β_K were computable, so would be $x = 4\beta_K - 1$, whose base-four digits are zero or one according to K. The first-digit cases lie in the separated intervals $[0, 1/12]$ and $[1/4, 1/3]$, since the later tail is at most one third of the current place. A sufficiently accurate approximation decides the digit; replace x by $4x$ minus that digit and repeat. This would decide K. All the finite algebraic/logarithmic expressions are computable, so this already strictly extends their class. Replacing the plus sign by a minus sign and starting at $1/2$ gives a noncomputable right-c.e., hence non-left-c.e., example too, consistently with Remark P1.2.

There is also a computable strict extension. Enumerate all candidate finite expressions from effectively isolated algebraic numbers, logarithms of integers at least two, finite sums/products and division only by known nonzero rational constants or such logarithms. This is a uniformly computable sequence $x_0, x_1, \ldots$ containing the indicated families and closures. No test of nonzeroness for arbitrary expressions or membership in $(0, 1]$ is needed. Start at $I_0 = [1/4, 1/2]$. If I_e has length l and midpoint m, approximate x_e by a rational v to error below $l/16$. For $v \le m$ choose the rightmost closed quarter, otherwise the leftmost. The chosen child excludes x_e by at least $3l/16$. The intervals are nested and have lengths $|I_0|4^{-e}$, so their unique common point β_d is computable, differs from every x_e, and has the rational midpoints as a total computable name. Apply Theorem P1.1 to both names. No halting oracle or noncomputable coefficient enters either game. □

P6.3 What the architecture does not imply

Binary pulses force $0 \le k(J) \le J$, so this construction has $\beta \le 1$. A dimension-uniform weighted calibration $|T_0|_\rho \le C|S_0|_\rho^\alpha$ with $\alpha > 1$ cannot hold for this interface. Put rational weight w on the binary history $T_0 = 1$, $S_0 = 1/2$ stopped at age zero, and the remaining weight on $T_0 = S_0 = 0$, never stopped. The two norms are $\sqrt{w}$ and $\sqrt{w}/2$, so their ratio is $2^\alpha w^{(1-\alpha)/2}$, which diverges as w decreases to zero. Choose the finite clock large enough that its boundary error is smaller than the penalty. This restricts the calibration mechanism, not all games.

Fixed positive mixing of finitely many resource laws $\varepsilon^{-\beta_i+o(1)}$ has exponent $\max_i \beta_i$: a maximal term bounds the sum below, and finitely many upper bounds at any fixed exponent margin bound it above. Thus fixed density mixing does not select a new power exponent. Likewise a qualitative error $e(\delta) \le C\delta^\beta + \omega(\delta)$, $\omega \to 0$, does not imply the power estimate:

$\omega(\delta) = 1/\log(\mathrm{e}/\delta)$ has exponent zero. These observations explain the need for the polynomial same-state bounds above rather than a qualitative compactness modulus. The full usable arithmetic ceiling and existing finite exponents beyond one remain open. The quantitative local ceiling and the contrasting unbounded upper exponents are proved below.

P6.4 A ceiling for quantitatively local lamp extensions

The upper endpoint one has a model-theoretic explanation for the lamp constructions, stronger than the binary-pulse restriction. It is a statement about a specified extension class. A finite list of equations containing the lamp letters need not belong to that class.

Definition P6.5 (Block-sharp models and locality)**.** Let Σ be the six source letters of Section P2, J its marker, and W a finite set of source-trivial words. A (W,ν)-block-sharp model is a unitary assignment on $H = \bigoplus_{b\in B} H_b$, with equal-dimensional blocks, whose generators permute the blocks, whose marker is $-I$, and for which every $w \in W$ fixes every block outside one common set B_{bad} pointwise, with $|B_{\mathrm{bad}}| \le \nu|B|$. Write Π for its bad-block projection.

Let W_L comprise the source-trivial words of length at most a fixed L. A finite bounded rational extension $\mathcal{P}$ is *L-local* if there are fixed $K, c \ge 1$ such that every (W_L,ν) model extends on $H \otimes H_{\mathrm{aux}}$, retaining the source letters and exact variable domains, with

$$\log_2 \dim H_{\mathrm{aux}} \le c(1+\log_2(1/\nu)), \quad Q_j = Q_j(\Pi\otimes I) + E_j, \quad \|Q_j\| \le K, \quad \|E_j\|_{2,\tau} \le K\sqrt{\nu}$$

for every defining residual Q_j, including commutators. Here τ is the flat state. The constants and L are fixed before precision tends to zero. This is a quantitative extension hypothesis.

Lemma P6.6 (Simultaneous sharp source models)**.** *For each finite W there are (W,ν)-block-sharp models with*

$$\log_2 \dim H \le C_W \nu^{-1} \log_2(2/\nu).$$

No connectivity or irreducibility is required.

Proof. Use the block construction in [50, Proposition 9.4], whose hash-register mechanism builds on the overlapping-qubit construction [14]. Let ℓ be the largest word length, increased to one if necessary. At an interior clock position take the union of the lit sites of all words in W; its size is at most $\ell|W|$. Set

$$M = \lceil 4\ell/\nu \rceil, \qquad k = \left\lceil \log_2 \frac{4\ell^2 \max(1,|W|)^2}{\nu} \right\rceil.$$

The clock-boundary fraction is at most $2\ell/M \le \nu/2$, and pairwise independence bounds collision on that common lit set by $(\ell|W|)^2 2^{-k-1} \le \nu/8$. Outside their union each word fixes the entire block, simultaneously. The source's displayed dimension formula gives

$$\log_2 \dim H = 1 + 2\log_2(2M+1) + k(6M+10) + 2^k = O_W(\nu^{-1}\log(2/\nu)).$$

Generators permute equal-sized blocks and $J = -I$. Their action may preserve distinct rank sectors of the hash label; the proof uses only pointwise identity on good blocks. □

Theorem P6.7 (Local lamp ceiling)**.** *For the rational BCS game compiled from an L-local extension,*

$$\log_2 D_G(\varepsilon) \le C_{\mathcal{P}} \varepsilon^{-1} \log(1/\varepsilon)$$

at all sufficiently small losses. The same upper bound holds for $d_G^{\rightarrow}, p_G, f_G$. The game has quantum value one, and all of these upper power exponents are at most one.

Proof. In the flat state, $\tau(\Pi \otimes I) \le \nu$, so every residual in Definition P6.5 has norm at most $2K\sqrt{\nu}$. Square-root-state commutators vanish. The compiler Theorem 3.1 gives loss at most $C_0\nu$ and local dimension at most a fixed multiple of $\dim H \dim H_{\mathrm{aux}}$. Take $\nu = \varepsilon/C_0$ and apply Lemma P6.6. The upper state is maximally entangled, so its pure entropy, formation and directional distillable rate are the same logarithm. Its losses tend to zero. □

Lemma P6.8 (Membership of the displayed calibrators)**.** *The stopped and programmed lamp calibrators used in this part, their admitted static phase variants, and their finite mixing and chaining satisfy Definition P6.5. This includes the oscillating power-pair construction of Theorem Q6.1.*

Proof. For every normalized source residual $R_i = (\varphi(r_i) - I)/2$, its relator unitary fixes the good space pointwise and hence preserves its orthogonal complement. Thus $R_i = \Pi R_i \Pi$ and $T_i = |R_i| \le \Pi$; choose a unitary polar extension. Use separate finite fuel/time clocks, and assign the controller, reciprocal age, counters, zero tests and positive witnesses their exact scalar histories through cyclic resets, as in Proposition P4.1.

In each calibrator set $C = D = S = 0$, $H_{\mathrm{run}} = I$ and $P_{\mathrm{stop}} = 0$. Running, stopping, multiplier updates and domains are then exact; only the initialization of C and the reset term of its update can fail. Those failures are $-ZT_i$ and $-Z_+T_i$, up to sign and order, and are supported on Π. The controller need not reach a nonzero-input stopping time. The declared current commutations hold: the source polar factor and original source letters are absent from that list, and different residual workspaces are not required to commute.

For a finite chain set every intermediate input and penalty to zero. Only its first link to the original modulus has a bad-supported failure. Mixing changes bounded scalar coefficients or controller histories. Static algebraic phase parameters act as scalar roots in the admitted phase interface. The unguarded increment defect is $\sqrt{2}\,2^{-N_c/2}$ and the Cantor defects are $O(4^{-N_c})$; $N_c = \lceil \log_2(1/\nu) \rceil + O(1)$ makes them $O(\sqrt{\nu})$. Finitely many clock factors have the required logarithmic auxiliary cost. Powering a bounded residual preserves this estimate, since $\|h_q(Q)\|_{2,\tau} \le \|Q\|^{q-1}\|Q\|_{2,\tau}$.

The hold stream changes none of these assignments. The separate logarithmic profile construction has Cantor sources and a carry increment, rather than this lamp source; its $O(n)$ log dimension at loss 2^{-n} already has power exponent zero. It is not being identified with the unguarded lamp construction. □

Lemma P6.9 (A local invariance majorant)**.** *Adjoining positive contractions K, Y with $K - T - Y = 0$, tests $[K, s] = 0$ for source generators, and penalty $K = 0$, to an L-local extension remains $\max(L, 34)$-local for a source-relator modulus T.*

Proof. Let A be the blocks on which an actual source-certificate relator fails, and $N_j(A)$ their radius-j neighborhood in the undirected generator graph. Set $K = \Pi_{N_1(A)}$ and $Y = K - T$. Since $0 \le T \le \Pi_A \le K$, their domains and $K - T - Y = 0$ hold exactly. The commutator $[K, \varphi(s)]$ is supported on $N_2(A)$. For a block there, a path p of length at most two reaches an actual failure of some relator r of length at most thirty. Consequently $p^{-1}rp$ fails on that block and has length at most thirty-four. Thus $N_2(A)$ lies in the common declared bad set for $W_{\max(L,34)}$, as does the support of K. All residuals are bounded and bad-supported, with no auxiliary factor. A padded declared bad set cannot replace the actual failure set in this argument. □

Corollary P6.10 (Saturation and the class spectrum)**.** *For the stopped lamp checker with fixed multipliers $a > b \ge 2$, all four resource power exponents equal one. Among games compiled from the local class of Definition P6.5, the positive existing spectra for $\log D_G, p_G, f_G$ are exactly $\Delta_2^0 \cap (0, 1]$. The directional usable spectrum contains this set and has upper exponents at most one; its general arithmetic ceiling is not established here.*

Proof. The complete stopping columns L_j and terminal column K_J of the stopped-product proof remain valid when $a \geq b$. With ϕ the initialized column and r its error allowance, they give

$$\|L_j T\phi\| \leq b(b/a)^j(\|L_j S\phi\| + r) + r \leq b\|L_j S\phi\| + (b+1)r, \qquad \|K_J T\phi\| \leq a^{-J} + r.$$

Exact column completeness and the Euclidean triangle inequality imply

$$\|T\phi\| \leq a^{-J} + b\|S\phi\| + ((b+1)\sqrt{J} + 1)r.$$

Choose $J = \lfloor \log_b(1/\delta) - B\log_b(1+\log(1/\delta)) \rfloor$ with fixed large B. Since $a^{-J} \leq b^{-J}$, localization, the tested penalty and charge back give $|T|_\rho \leq C\delta(1+\log(1/\delta))^c$. At loss ε, the original-state lamp extraction used in Theorem P1.1 therefore has scale $e = O(\sqrt{\varepsilon}\,\mathrm{polylog}(1/\varepsilon))$, including mirrors and state commutators, and yields $E_D^{A\to B} \geq e^{-2+o(1)} = \varepsilon^{-1+o(1)}$. The local upper bound matches it. This lower bound also bounds pure entropy, formation and log dimension. The nominal smooth-model exponent $\log a/\log b > 1$ is not a minimum over strategies.

For the class spectrum use the effective ceilings of Proposition P6.1, the local upper bound and the realizations of Theorem P1.1, whose membership was proved above. This proves equality for the three stated existing spectra, and only realization inclusion for the usable one. □

P6.5 Unbounded upper exponents for fixed games

The ceiling in Theorem P6.7 is specific to its quantitative locality hypothesis. The arithmetic hierarchy gives a different conclusion for general games, and even for Boolean constraint-system games.

Theorem P6.11 (No computable eventual envelope)**.** *For each total computable $F : \mathbb{N}^2 \to \mathbb{N}$ there is a finite rational two-player value-one game G such that, for every integer C,*

$$D_G(2^{-n}) > F(C,n) \quad \textit{for infinitely many } n.$$

The game can be chosen to be a BCS game. In particular, one such fixed game has infinite upper power exponent for $\log D_G, p_G, f_G$.

Proof. For fixed local dimension d, the assertion $R(G,d,n)$ that a strategy wins with probability at least $1-2^{-n}$ is a decidable real semialgebraic feasibility problem, uniformly in the rational game, d, n. Mixed states and POVMs are included; compactness makes the weak threshold exact. Suppose the first assertion fails for F. Every value-one game would then have C, n_0 such that $R(G, F(C,n), n)$ holds for all $n \geq n_0$. Conversely, this condition implies value one. Thus the set of value-one games would have the definition

$$\exists C\, \exists n_0\, \forall n \geq n_0\ R(G, F(C,n), n),$$

placing it in Σ_2^0. Mousavi–Nezhadi–Yuen prove that this set is Π_2^0-complete [54, Theorem 1.1]. Closure under computable preimages would collapse these two levels, contrary to the strict arithmetic hierarchy. The Boolean version follows from Lemma P6.12 below, which proves the same hardness within that class.

Take $F(C,n) = 2^{2^{Cn}}$. The conclusion for every C says exactly that the upper power exponent of $\log D_G$ is infinite. To transfer it to pure entropy, let a loss-ε pure strategy have entropy S. Its Schmidt tail past rank R is at most $S/\log_2 R$: each ordered Schmidt probability beyond R is at most $1/R$. Set $R = 2^{\lceil (S+1)/\varepsilon^2 \rceil}$. Normalized Schmidt truncation changes the state by half trace distance at most ε, so compression to its Schmidt supports gives

$$\log_2 D_G(2\varepsilon) \leq (p_G(\varepsilon)+1)\varepsilon^{-2} + 1.$$

Infima are handled by arbitrarily near minimizing strategies. A finite upper power exponent for p_G would give a finite one for $\log D_G$, a contradiction. The two-support equal-loss ensemble and coherent-flag argument used in Proposition P6.1 gives $f_G \le p_G \le f_G + 1$, proving the formation assertion too. □

Lemma P6.12 (Synchronous value one and Boolean value one)**.** *There is an effective map from finite rational synchronous games G whose question distribution gives every diagonal pair (x, x) positive probability to finite rational BCS games H_G such that $\omega_q(G) = 1$ if and only if $\omega_q(H_G) = 1$. Completeness has loss $O_G(\zeta)$ and soundness loss $O_G(\sqrt{\eta})$, where ζ, η are the respective input losses.*

Proof. This is the constraint-variable (SynchBCS) construction of Paddock [59, Section 4.3 and Proposition 4.22], whose BCS algebra is $*$-isomorphic to the synchronous algebra. We record the value-one transfer with the explicit loss exponents used here. For each question x and answer a, introduce a projection $P_{x,a}$. Impose the sum-to-one equations, pairwise orthogonality within a question, and $P_{x,a}P_{y,b} = 0$ for each losing pair on a supported question pair. This is a bounded rational presentation of the allowed kind in Theorem 3.1; compile it to H_G.

For games with positive weight on every diagonal question pair, Mousavi–Nezhadi–Yuen show that value one admits finite-dimensional synchronous strategies approaching one [54, Theorem 1.6]. Such strategies are tracial PVM strategies. Split the finite tracial algebra into matrix blocks and choose a block whose loss is at most the average. For losing projections P, Q in that block, $\|PQ\|_{2,\tau}^2 = \tau(PQ)$. The minimum positive question probability is fixed. A loss-ζ strategy therefore supplies a flat model of precision $O_G(\sqrt{\zeta})$, with the PVM equations exact. Flat completeness of the compiler gives loss $O_G(\zeta)$ for H_G.

Conversely, a loss-η strategy for H_G gives exact projection variables in a weighted model with $\xi = \sqrt{\rho}$, defining errors at most $\delta = O_G(\sqrt{\eta})$ and $\|[P_{x,a}, \xi]\|_2 \le \delta$. We construct valid POVMs directly. For one question suppress x, order its k answers, put $V_0 = I$, $V_j = (I - P_j)V_{j-1}$, and set

$$E_a = V_{a-1}^* P_a V_{a-1} \ (a < k), \qquad E_k = V_{k-1}^* V_{k-1}.$$

These effects are positive, and their sum is exactly I by telescoping $V_{j-1}^* V_{j-1} - V_j^* V_j$.

Here are weighted estimates requiring no inverse and no tracial approximation. Telescoping gives $V_j - I = -\sum_{l \le j} P_l V_{l-1}$. For $a > l$, the residual $P_a P_l$ has weighted norm at most δ. Moving ξ past the following contraction word costs at most $k\delta$, by its factor commutators. Hence $\|P_a(V_{a-1} - I)\xi\|_2 = O_k(\delta)$. For any fixed polynomial A,

$$\|A^*\xi\|_2 \le \|A\xi\|_2 + \|[A^*, \xi]\|_2,$$

since $\|\xi A^*\|_2 = \|A\xi\|_2$. Applying this to the preceding polynomial gives the adjoint bound. Expand $V_{a-1}^* P_a V_{a-1} - P_a$ into its two differences and transport the remaining factors; the result is

$$\|(E_a - P_a)\xi\|_2 = O_k(\delta), \qquad \|[E_a, \xi]\|_2 = O_k(\delta).$$

The sum equation gives the first bound for the last answer as well, and the second follows from the product commutator rule.

Share the canonical purification $\operatorname{vec}\xi$ and measure $E_{x,a}^{\mathsf{T}}$ on Alice and $E_{y,b}$ on Bob, with the matching vectorization convention. Their correlation is $\operatorname{Tr}(\xi E_{x,a} \xi E_{y,b})$. Cauchy–Schwarz and the two estimates replace this by $\operatorname{Tr}(\rho P_{x,a} P_{y,b})$ with error $O_G(\delta)$. On a losing pair its absolute value is at most $\|P_{x,a}P_{y,b}\xi\|_2 \le \delta$. The actual correlations are nonnegative and there are finitely many pairs, so total loss is $O_G(\delta)$. This is a strategy in the same local dimension. Value one is therefore preserved in both directions. The Π_2^0-hard synchronous instances in

the proof of [54, Lemma 6.7] use the uniform distribution on all question pairs. They satisfy the diagonal-support hypothesis, so the lemma transfers that hardness to BCS games and proves the Boolean clause of Theorem P6.11. □

Why diagonal support is needed. Take questions and answers $\{0,1\}$, uniform weight on $(0,1),(1,0)$ and zero weight on the diagonal. On $(0,1)$ require $a = b = 0$; on $(1,0)$ require $a = b = 1$, with the synchronous rule on the unasked diagonal. Alice's answers $f(0) = 0, f(1) = 1$ and Bob's answers $g(0) = 1, g(1) = 0$ win perfectly. But the losing-pair relations force $P_{0,1} = 0$ and hence $P_{0,0} = I$; they then force both $P_{1,0} = P_{1,1} = 0$, contradicting their sum equation. The same finite sum/product identities rule out a sequence of vanishing-defect normalized models, so compiler soundness makes the Boolean game's value strictly below one.

The game in Theorem P6.11 depends on F. For each fixed value-one game its finite-precision dimension is computable by the same feasibility search; no one game defeats every computable function. Nor does this theorem assert an infinite usable exponent or any existing power limit above one. It establishes an infinite upper exponent.

P6.6 What a global pairwise certificate can establish

The following obstruction concerns a certificate using only global pair errors, mirrors and square-root-state commutators. It places no upper bound on the actual resource demand of a game with additional relations.

Proposition P6.13 (Pairwise hash models)**.** *For N labelled qubits and $0 < \kappa \le 1$, there are reflections with exact squares, exact on-site anticommutation, exact mirrors and zero square-root-state commutators, whose distinct-site signed pair errors are at most κ, with a pure flat model satisfying*

$$\log_2 D \le 8\kappa^{-2} + 2\log_2 \max(N, 8\kappa^{-2}) + 2.$$

For $N \ge 2$ another such model satisfies

$$\log_2 D \le 12\kappa^{-2} + 3\log_2(C\kappa^{-2}\log_2 N)$$

for an absolute constant C. There are mixed models with the first model's pair errors and both directional distillable entanglements and formation exactly $K = \lfloor 8\kappa^{-2} \rfloor$, independently of N.

Proof. Let $K \ge 1$ and take a prime $p \ge \max(N, K)$. Label items by distinct elements of $\mathbb{F}_p$, and average uniformly over $h_{a,b}(u) = ((au + b) \bmod p) \bmod K$, $a \ne 0$. On the block for h, use Pauli X, Z on qubit $h(u)$ in a K-qubit register. Squares and on-site anticommutation are exact; equal-type Paulis always commute. For distinct items a cross commutator is zero off collisions and has norm two on a collision.

Write $p = qK + r$, $0 \le r < K$. A distinct ordered pair of field elements is uniform under (a, b), and the residue classes have sizes q or $q + 1$. Thus

$$\Pr[h(u) = h(v)] = \frac{Kq(q-1) + 2rq}{p(p-1)} \le \frac{1}{K}.$$

Indeed, after multiplying by K, the difference between the denominator and numerator is $Kq(K-1) + r(r-1) \ge 0$. The squared cross error is four times this probability. Take $K = \lfloor 8\kappa^{-2} \rfloor \ge 4\kappa^{-2}$ and, by Bertrand's theorem, $p < 2\max(N, K)$; the endpoint two is harmless. The direct-sum Hilbert space has dimension $2^K p(p-1)$, proving the first bound.

Its maximally entangled state supplies exact transposed mirrors and a scalar square-root marginal.

For the refinement put $d = \lceil \log_2 N \rceil$, $K = \lfloor 12\kappa^{-2} \rfloor$, and take a prime $dK \le q < 2dK$. There are at least N polynomials over $\mathbb{F}_q$ of degree at most d; assign distinct ones P_u. Choose a uniform evaluation point z and an independent affine hash of $\mathbb{F}_q$ into K residues. Distinct polynomials agree at z with probability at most $d/q \le 1/K$, and, conditional on distinct values, the preceding exact count bounds their hash collision by $1/K$. Thus collision is at most $2/K$ and squared pair error at most $8/K \le \kappa^2$. The family has fewer than q^3 members, giving the second dimension bound. For one item one qubit suffices.

Finally replace the coherent flat hash register by shared classical flags:

$$\rho = \frac{1}{|\mathcal{H}|} \sum_h |h,h\rangle\langle h,h| \otimes \Phi_{2^K}.$$

The controlled reflections have the same weighted errors, flat local marginals and exact mirrors. Discarding the flags distils K ebits in either direction, while the displayed ensemble has formation at most K. Since distillation is bounded by formation, all three equal K. □

Proposition P6.14 (Junk sectors and the restricted barrier)**.** *Suppose a lower-bound certificate applies a universal pairwise lamp theorem to global weighted errors valid for every loss-ε strategy, with word decoders respecting direct sums and giving scalar reflections on a deterministic classical sector. Its common pair tolerance must allow $\kappa = \Omega(\sqrt{\varepsilon})$. Such a certificate can conclude at most $O(\varepsilon^{-1})$ for original-state distillation or formation, and at most $O(\varepsilon^{-1} + \log(L+1))$ for a pure or dimension bound when all lamp words have length at most L over a fixed finite alphabet.*

Proof. Take a strategy S of loss at most $\varepsilon/2$ and a deterministic classical strategy S_0. The flagged direct sum $(1-w)S \oplus wS_0$, with $w = \varepsilon/2$, has loss at most ε, dimension $\dim S + 1$, and

$$|q|_\rho^2 = (1-w)|q|_S^2 + w|q|_{S_0}^2$$

for every decoded polynomial residual. In the scalar sector the on-site anticommutator of two reflected lamps has modulus two. Its weighted error is therefore at least $2\sqrt{w}$. This forces the common tolerance. Squares, mirrors and cross commutators can be exactly zero there; the lower floor is not a statement about every residual.

A universal bound using only the pairwise data must also hold for Proposition P6.13. Its mixed flags bound the distillation and formation conclusion by $\lfloor 8\kappa^{-2} \rfloor$. Its polynomial hash model bounds a pure or dimension conclusion by $O(\kappa^{-2} + \log(\kappa^{-1} \log(N+1)))$. Word counting gives $\log\log(N+1) \le \log(L+1) + O(1)$, not equality. Substitute the forced tolerance. Selecting sectors, using additional higher-order relations or exploiting more of the game is outside this certificate hypothesis. □

P6.7 A conditional route above one

Definition P6.15 (A stronger source interface)**.** A finite marked source has strength $\gamma > 0$ if the following hold. Its lower interface, with fixed marker visibility, bounds the directional distillable entanglement of the *original* mixed strategy state by $e^{-2\gamma + o(1)}$ whenever its decoded residuals, mirrors, inverse mirrors and square-root-state commutators are at most e. The marker residual is among the tested residuals. Its upper interface supplies flat models with marker $-I$,

$$\log_2 D_N \le N^{\gamma + o(1)}, \qquad \tau(|R_i|^{2v}) \le C_v N^{-v} \quad (v \ge 1)$$

for every fixed v and every large N, simultaneously for its finitely many normalized residuals. The finite upper models must admit the same clock and calibrator extensions used in Theorem P1.1, at logarithmic auxiliary cost.

Theorem P6.16 (Transfer from a stronger source)**.** *If Definition P6.15 holds with positive $\gamma \in \Delta_2^0$, every $\beta \in \Delta_2^0 \cap (0, \gamma]$ has one fixed rational value-one BCS game with matched existing power exponent β for $\log D_G, d_G^{\rightarrow}, p_G, f_G$.*

Proof. Use the stream and calibrator of Theorem P1.1 with pulse density $\beta' = \beta/\gamma \in \Delta_2^0 \cap (0, 1]$. Limit-computable reals form a field. Include a calibrator and residual test for the marker, retaining its fixed lower-interface visibility.

The same-state calibration gives source error $e = \varepsilon^{\beta'/2 - o(1)}$ at loss ε. Mirror and state commutator errors are $O(\sqrt{\varepsilon})$, hence no larger asymptotically because $\beta' \le 1$. The lower interface gives $E_D^{A \to B} \ge \varepsilon^{-\beta + o(1)}$ on the original state.

For upper models fix $0 < v < 1/\beta'$. If $v < 1$, concavity gives $\tau(|R_i|^{2v}) \le \tau(|R_i|^2)^v \le C_1^v N^{-v}$; otherwise use the stated moment bound. The scalar calibrator envelope, its spectral integration and the finite residual sum give penalty at most $C_v |\mathcal{R}|^v N^{-v}$. Its fixed constants do not depend on N. The complete model has loss $O_v(N^{-v})$ and log dimension $N^{\gamma + o(1)} + O_v(\log N) = N^{\gamma + o(1)}$. Thus the upper power exponent is at most γ/v. Let v increase to $1/\beta'$. The upper states are maximally entangled, so the same upper bound holds for all four resources, matching the lower bound. □

Proposition P6.17 (Necessary costs of such a source)**.** *Under Definition P6.15, sharp models admitting the required auxiliary extension of negligible logarithmic cost must have $\log_2 \dim H \ge \nu^{-\gamma + o(1)}$ at bad fraction ν. A universal lower interface derived only from the pairwise data of Proposition P6.13 cannot have $\gamma > 1$.*

Proof. A sharp model has normalized residual errors at most $O(\sqrt{\nu})$, exact marker, and exact mirrors in its flat purification. Compile its admissible extension. The lower interface gives original-state distillation at least $\nu^{-\gamma + o(1)}$, bounded by its log dimension. Removing a negligible auxiliary logarithm proves the first claim. For the second, the mixed hash models have all the stipulated errors at most e with distillation and formation $O(e^{-2})$, irrespective of the number of labels. A universal conclusion $e^{-2\gamma + o(1)}$ would contradict them when $\gamma > 1$. □

The lamp source has $\gamma = 1$. No stronger source is supplied here. These necessary conditions concern the stated interface and universal pairwise conclusions; they do not exclude sector selection or richer relations. Existing finite power limits above one remain open, alongside the infinite upper exponents of Theorem P6.11.

Precision profiles and exponent pairs

Q1 A carry checker for logarithmic profiles

The finite presentation in this part converts a programmed pulse count into a precision profile. It uses the compiler, clocks, weighted transport, triangular count and original-state extraction of Sections 3, 4, 8 and 12, and the counter and stopped-product arguments of Sections P4 and P5. Its controller has no game-witness register. Throughout, $\xi = \sqrt{\rho}$ and $|R|_\rho = \|R\xi\|_2$; logarithms are base two. All constants depend only on a fixed finite presentation. The precision convention is that of Section P3, including square-root-state commutators.

Q1.1 The cost of a checker increment

Remark Q1.1 (Standard-model overflow). On the standard N_c-bit cyclic time register, the unguarded increment in Section P3 has residual

$$(|0^{N_c}\rangle - |10^{N_c-1}\rangle)\langle 1^{N_c}|, \qquad \|R\|_f = \sqrt{2}\,2^{-N_c/2}.$$

At unpowered precision $\delta = 2^{-n/2}$ it therefore needs $N_c \geq n+1$, whereas the clock count resolves $\log_4(1/\delta) = n/4$ digits per checker. Two such registers leave a linear difference between these upper-model costs and that count. The carry model below has overflow amplitude 4^{-N_c}, so $N_c = n/4 + O(1)$ suffices. This comparison concerns the displayed models and makes no lower-bound claim about all models of the unguarded relation.

The rank-one residual follows by evaluating the increment on the all-ones input; all other input columns agree. Its two output basis vectors are orthogonal, giving the stated flat norm. The carry recursion of Sections 5 and 7 instead assigns $H(x) = 4^{-r(x)}$ when $r(x) < N_c$ initial bits are one, and zero on the all-ones point. Its recursion is exact and its increment residual is 4^{-N_c} on that point. The clock's affine and tail-flip residuals are $3 \cdot 4^{-N_c}$ times a projection or commutator, and set the common precision scale. This is the vanishing-amplitude alternative to the fixed-rejection-gap obstruction of Appendix A.

Q1.2 The complete presentation

Fix a finite deterministic INC/DECJZ/NOP program π, initially in its zero-counter start configuration, with two output NOP controls. Let $w \in \{0,1\}^{\mathbb{N}}$ be its pulse word, $k(J)$ its one-count among the first J pulses, and t_J the physical time just after pulse J. Set $k(0) = 0$, $t_0 = 0$ and $\kappa(l) = \min\{J : k(J) = l\}$, with $\kappa(0) = 0$. Take fixed integers $m, q \geq 1$.

The *checker* has two Cantor-clock copies, fuel and time, satisfying (4.1). Add a positive time carry H_t and a unitary time increment U with

$$\begin{gathered} H_t = (I - P_t) + \tfrac{1}{4} P_t T_t H_t T_t^*, \qquad [P_t, T_t H_t T_t^*] = [H_t, A_t] = 0, \\ U^* A_t U = A_t - I + \tfrac{7}{4} H_t. \end{gathered} \tag{Q1.1}$$

The complete fuel set commutes with the complete time set, including H_t, and with U. The reset projection Z commutes with A_f, A_t and satisfies (P3.2) for $W = T_f T_t$. The controller uses exactly the domains and rational equations (P4.1)–(P4.2), with $Y_+ = U^* Y U$. In particular $F = Q_{\text{zero}} + Q_{\text{one}}$ and $G = Q_{\text{one}}$.

For each $i \leq m$ introduce a source clock $\mathcal{S}_i = (A_i, P_i, T_i, X_i)$. Its relations $[A_i, P_i] = 0$ and $X_i P_i X_i = I - P_i$ are retained. The two other clock relations are replaced by

$$T_i A_i T_i^* - 4A_i + 3P_i = V_{i1}\Theta_{i1}, \qquad [X_i, T_i A_i T_i^*] = V_{i2}\Theta_{i2}, \tag{Q1.2}$$

where the $V_{i\nu}$ are unitaries and the $\Theta_{i\nu}$ positive contractions. Complete source-clock sets commute with one another and with both complete checker-clock sets. There are no other source commutations.

For every (i,ν) use a calibrator (P5.1), with input $T_0 = \Theta_{i\nu}$, its own $S, C, D, H_{\mathrm{run}}, P_{\mathrm{stop}}$, and the shared controller and reset. Thus its equations are

$$\begin{aligned} (I - H_{\mathrm{run}})P_{\mathrm{stop}} = 0, \quad Z(C-\Theta) = Z(D-S) = Z(H_{\mathrm{run}} - I) = 0, \\ (H_{\mathrm{run}})_+ = Z_+ + (I - Z_+)(H_{\mathrm{run}} - P_{\mathrm{stop}}), \\ C_+ = Z_+\Theta + (I - Z_+)C[I + G(H_{\mathrm{run}} - P_{\mathrm{stop}})], \\ D_+ = Z_+ S + (I - Z_+)D[I + F(H_{\mathrm{run}} - P_{\mathrm{stop}})], \\ (D - I/2)P_{\mathrm{stop}} = 0, \qquad S = 0. \end{aligned} \tag{Q1.3}$$

For a calibrator its instantaneous list consists of Z, all controller variables and Θ, S, C, D, together with $H_{\mathrm{run}}, P_{\mathrm{stop}}$. Require pairwise commutation in that list and commutation with A_f, A_t. Static Θ, S commute with every checker generator and U. Distinct calibrator workspaces need not commute; no additional commutation is imposed on $V_{i\nu}$. The domains are those in the counter and calibrator constructions. This is a finite rational bounded $*$-polynomial presentation. Power every defining residual, commutators included, by the polynomial construction (11.10), retaining domains and state-commutator tests, and apply Theorem 3.1. Denote the resulting rational BCS game by $G[\pi, m, q]$.

Q1.3 Polynomial localization with the carry coordinate

Theorem Q1.2 (Carry-checker interface)**.** *For every weighted precision-δ model of this presentation and horizon $H \geq 2$, choose $H < 2^k \leq 2(H+1)$ and put*

$$\Gamma_k = \left(\prod_{j<k} P_{f,j}\right)(I - P_{f,k}), \quad D_k = \prod_{j<k}(I - P_{t,j}), \quad Q = \Gamma_k D_k,$$

with raw digits $P_{c,j} = T_c^j P_c T_c^{-j}$ and products in the displayed order. For Y an instantaneous self-adjoint contraction and Θ a static input there are fixed C, a such that, writing $Y_h = U^{-h} Y U^h$ and $r = C(H+1)^a\delta$,

$$\begin{aligned} &|(I-Z)Q|_\rho \leq r, && |Z_h Q|_\rho \leq r \quad (1 \leq h \leq H), \\ &|[Y_h, Q]|_\rho \leq r \quad (0 \leq h \leq H), && |\Theta|_\rho \leq C(H+1)\|\Theta Q\xi\|_2 + r, \\ &\|\Theta Q\xi\|_2 \leq |\Theta|_\rho + r. \end{aligned}$$

Proof. Use Lemmas 2.2, 2.3, 4.2 and 9.2 and the word-interchange argument in Lemma 9.3. For self-adjoint B and a contraction Y, moving the right factor across ξ gives

$$|[B,Y]|_\rho \leq 2|B|_\rho + \|B\|\|[Y,\xi]\|_2.$$

The clock repair $\widehat{A}_c$ has spectrum in the Cantor set $\mathcal{K}$, differs from A_c by $O(\delta)$ in weighted norm, and has square-root-state commutator $O(\delta)$. Its raw digits differ from $f_j(\widehat{A}_c)$ by $C4^j\delta$. Own-bit flips and preservation of later digits have errors at most $C(M+1)^3 4^M\delta$ through depth M.

Identify $\mathcal{K}$ with $\mathbb{Z}_2$ by $x = \sum_{j\geq 0} 3b_j 4^{-j-1}$ and let g_h be addition by h. Addition preserves the first differing binary digit, so every g_h is 2-Lipschitz on $\mathcal{K}$. The functions f_j have Lipschitz constant $2\cdot 4^j$, while $f_j \circ g_{-h}$ has constant at most 4^{j+1} on $\mathcal{K}$. Such functions extend linearly across Cantor gaps without increasing their constants. The carry translation estimate gives

$$|U^{-h}\widehat{A}_t U^h - g_h(\widehat{A}_t)|_\rho \leq C(h+1)\delta.$$

Conjugating by U^h costs $O(h\delta)$ by right transport. Apply the 2-Lipschitz map g_{-h}, using functional calculus on operators with Cantor spectrum, to obtain also

$$|U^h \widehat{A}_t U^{-h} - g_{-h}(\widehat{A}_t)|_\rho \le C(h+1)\delta. \tag{Q1.4}$$

The needed state-commutator term is $O(\delta)$. Thus neither direction incurs an iterated Lipschitz factor.

Time cylinders. Let $B_c \subset \mathcal{K}$ be the residue class c mod 2^k and $E_c = 1_{B_c}(\widehat{A}_t)$. Its indicator is $2 \cdot 4^k$-Lipschitz on $\mathcal{K}$. Replace the raw factors of D_k from the right. Each replacement costs its digit error plus the state commutator of the remaining spectral product, at most $C4^k\delta$. Consequently

$$|D_k - E_0|_\rho \le Ck4^k\delta, \qquad |U^h D_k U^{-h} - E_h|_\rho \le C(h+1)k4^k\delta. \tag{Q1.5}$$

The second bound uses (Q1.4) and $1_{B_0} \circ g_{-h} = 1_{B_h}$. Also $E_0 E_h = 0$ for $1 \le h < 2^k$.

Fuel and reset. The weighted commutators of Z, Θ, E_c with Γ_k are bounded by $C(k+2)^4 4^k\delta$. For Z, pass from its commutator with A_f to $\widehat{A}_f$, then to its spectral digits, and back to raw digits using the preceding commutator inequality. For E_c use commutation of fuel generators with A_t and functional calculus. For static Θ use its generator commutations. Expand product commutators with their right-factor state terms. Word interchange similarly gives

$$|U^h \Gamma_k U^{-h} - \Gamma_k|_\rho \le C(h+1)^2(k+2)^6\delta.$$

Translate the reset equation by W^j. Interchanging fuel/time words has polynomial cost in j, giving

$$Z^{(j)} = (I - P_{f,j}) + P_{f,j}(I - P_{t,j})Z^{(j+1)} + e_j, \quad |e_j|_\rho \le C(j+1)^3\delta, \quad \|e_j\| \le 8.$$

For $G_j = \prod_{i<j} P_{f,i}(I - P_{t,i})$, ordered substitution through k yields

$$Z = \sum_{j \le k} G_j(I - P_{f,j}) + G_{k+1}Z^{(k+1)} + \sum_{j \le k} G_j e_j.$$

On left multiplication by Γ_k, the $j < k$ terms vanish up to digit/interchange errors by $P_{f,j}(I - P_{f,j}) = 0$. In the $j = k$ term, move $P_{f,k}$ against $I - P_{f,k}$ and drop the repeated earlier fuel digits. The remainder is $\Gamma_k D_k$. The final term has the same contradictory $P_{f,k}$ factor; transport of its remaining right factor costs $O(k)\delta$. Hence

$$|\Gamma_k Z - \Gamma_k D_k|_\rho \le C(k+2)^6 4^k\delta. \tag{Q1.6}$$

Move Z past Γ_k and use $D_k^2 \approx E_0^2 = E_0 \approx D_k$ to get $ZQ \approx Q$. For $h \ge 1$, right transport and (Q1.5) give $ZU^h Q U^{-h} \approx Z\Gamma_k E_h \approx \Gamma_k D_k E_h \approx \Gamma_k E_0 E_h = 0$. This proves the first two interface bounds, including all right-factor state commutators.

Transported workspaces. Conjugate $[Y, \widehat{A}_t]$ and replace the translated clock by $g_h(\widehat{A}_t)$. The commutator inequality gives $|[Y_h, g_h(\widehat{A}_t)]|_\rho \le C(h+1)\delta$. Apply functional calculus to $f_j \circ g_{-h}$, then return to raw digits, obtaining $|[Y_h, P_{t,j}]|_\rho \le C4^j(h+1)\delta$. Here $\|[Y_h, \xi]\|_2 \le (2h+1)\delta$. Fuel digits use commutation of U with the fuel generators and the same spectral argument for Y. Product expansion over the $2k+1$ factors proves the third interface bound.

Amplitude charge. Take the ordered measurement columns of these $2k+1$ projection factors. Exact completeness holds even when different factors fail to commute. Static Θ moves through their words with the displayed polynomial error. The flip at a factor exchanges its two outcomes and preserves the suffix: use later-digit preservation within a clock and generator commutation between clocks. It also approximately commutes with Θ and ξ. Put

$\epsilon_k = C(k+2)^4 4^k \delta$. Moving Θ to the right, applying exact two-column completeness, and moving it back gives $|\sqrt{a_0^2 + a_1^2} - a| \le C\epsilon_k$ for the two outcome amplitudes and their preceding amplitude. The flip gives $|a_0 - a_1| \le C\epsilon_k$. These are amplitude estimates, so no square root of an error is taken. Writing α_i for the successive chosen-branch amplitudes, they imply

$$\alpha_i \ge 2^{-1/2}\alpha_{i+1} - C\epsilon_k.$$

Iterate over all $2k+1$ bits. The additive errors form a geometric sum, so $\|\Theta Q\xi\|_2 \ge 2^{-(2k+1)/2}|\Theta|_\rho - C\epsilon_k$. Since $2^{(2k+1)/2} = O(H+1)$, this proves the amplitude charge. Finally $\Theta Q\xi = Q\Theta\xi + [\Theta, Q]\xi$ gives the last estimate. Every remaining error is a fixed polynomial in H times δ, because $k = O(\log H)$. □

Q2 Hold streams and profile calibration

Lemma Q2.1 (A stream with prescribed density endpoints)**.** *Let (a_s) be a total computable sequence of rationals in $[c_0, 1]$, where $c_0 > 0$ is rational. A finite counter program can emit an infinite binary word whose pulse times and one-count satisfy*

$$t_J \le C(J+1)^2, \qquad \liminf_J \frac{k(J)}{J} = \liminf_s a_s, \quad \limsup_J \frac{k(J)}{J} = \limsup_s a_s.$$

Every window of length ℓ in the first j pulses contains at least

$$c_0\ell - C(\sqrt{j}+2)\log_2(j+4)$$

ones.

Proof. Use the generator/background coroutine in Lemma P4.3: the current rational p/q emits periods of p ones and $q-p$ zeros, one elementary background instruction is executed per pulse, the next denominator is padded by q_{new}^2 background instructions before READY, and switching occurs only at a period boundary. Start with rational 1. Add a pulse counter P, stage counter σ and hold countdown D. Each pulse increments P and decrements positive D. Install a READY rational only at a period boundary with $D = 0$; increment σ and set $D = (\sigma+2)P$ by $\sigma+2$ non-destructive copies using a scratch counter. These operations have finite INC/DECJZ implementations. Copying the active generator uses $O(q_{\text{old}} + q_{\text{new}} + 1)$ instructions; arbitrary background scratch remains in the background.

Let j_s be the installation position of stage s, carrying $b_s = a_{s-1}$ for $s \ge 1$. The hold implies $j_{s+1} \ge (s+3)j_s$, and there are $O(\log J)$ installations before pulse J. The added work per pulse is bounded. Completed installation costs are bounded by $C\sum_{s<s^*}(s+3)j_s \le C\sum j_{s+1} = O(J)$; an installation occurring between adjacent pulses costs at most $O(J\log(J+2))$. The generator's denominator padding gives its existing $O(J^{3/2})$ instruction bound. Together these give $t_J \le C(J+1)^2$.

Every background stage finishes by totality, every finite hold expires, and the continuing generator reaches another period boundary. Thus every stage is installed. Complete periods have exact density b_s, with at most one incomplete period. Padding gives $q \le \sqrt{j} + O(1)$, so for $j_s \le j \le j_{s+1}$,

$$\frac{k(j)}{j} = \frac{j_s}{j}\frac{k(j_s)}{j_s} + \left(1 - \frac{j_s}{j}\right) b_s + O(j^{-1/2}). \tag{Q2.1}$$

At stage end the first coefficient is at most $1/(s+3)$, and $|k(j_{s+1})/j_{s+1} - b_s| \le 1/(s+3) + O(j_{s+1}^{-1/2})$. These subsequences attain both endpoint limits. Every intermediate ratio lies within vanishing error of the segment between the starting ratio and b_s, proving the reverse bounds. A window meets at most $\log_2 j + 4$ stages and loses at most one partial period of length $\sqrt{j}+2$ per stage. Since all active rationals are at least c_0, the window estimate follows. □

Lemma Q2.2 (Calibration at a prescribed pulse depth)**.** *For the presentation of Section Q1 there are fixed C, a_1 such that, in every weighted precision-δ model,*

$$2^J(t_J+2)^{a_1}\delta \le 1 \quad\Longrightarrow\quad |\Theta_{i\nu}|_\rho \le C(t_J+2)^{a_1}2^{-k(J)}. \tag{Q2.2}$$

This assertion does not require convergence of the pulse density.

Proof. Fix one calibrator and put $H = t_J$, $\phi = Q\xi$. The five conclusions of Theorem Q1.2 supply initialization, absence of resets through H, current-workspace localization and the two static-input estimates. The proof of Proposition P4.2, the normalized identities (P5.2) and the complete-column stopped-product argument therefore apply with precisely this interface. With $r_H = C(H+1)^b\delta$ they give

$$\|\Theta\phi\|_2 \le 2^{-k(J)} + 2^{J-k(J)+1}\big(\|S\phi\|_2 + \sqrt{H}\, r_H\big) + C\sqrt{H+1}\, r_H.$$

The integer pulse count minus one-count is nondecreasing, so the maximum stopped multiplier is $2^{J-k(J)+1}$ even for oscillating streams. The test $S=0$ and the last interface estimate bound $\|S\phi\|_2$ by a fixed polynomial in H times δ. Charge back by $|\Theta|_\rho \le C(H+1)\|\Theta\phi\|_2 + C(H+1)^a\delta$. Choose $a_1 \ge a+b+3$, enlarging it to absorb all fixed degrees. Under the displayed condition every term is bounded by $C(H+2)^{a_1}2^{-k(J)}$, since $\delta \le 2^{-J} \le 2^{-k(J)}$.

Each application uses only that calibrator's equations and the common controller. No step reorders different calibrator workspaces. In upper models the controller is scalar on each checker-age block and each calibrator is a spectral function of its own input and age. Thus the shared checker introduces no additional tensor factor or commutation assumption. □

Q3 The matched logarithmic profile

Theorem Q3.1 (Precision profile on the original state)**.** *Let π emit an infinite pulse word with positive lower one-density. Suppose $t_J \le T(J)$, where T is nondecreasing and $\log T(J) = o(J)$. For every fixed $m, q \ge 1$, $G = G[\pi, m, q]$ has quantum value one and no finite-dimensional perfect strategy. Set $\Lambda = \lfloor n/(2q) \rfloor$. Then*

$$\log_2 D_G(2^{-n}) = \Lambda + \tfrac{m}{2}k(\Lambda) + O\big(\log T(\Lambda) + \log n\big), \tag{Q3.1}$$

$$d_G^{\rightarrow}(2^{-n}) = \Lambda + \tfrac{m}{2}k(\Lambda) + O\big(\log T(\Lambda) + \log n\big). \tag{Q3.2}$$

The common upper bound has additive error $O(1)$. For polynomial delay the two-sided error is $O(\log n)$, and the lower bound holds on every original mixed strategy state of the indicated loss.

Proof. Simultaneous upper histories. Use two checker registers of N_c bits and m source registers of N_s bits. Each carries the finite clock of Theorem 4.4. Set U to cyclic addition on the time register, use the carry H_t of Remark Q1.1, and let $Z(f,t) = 1_{t\equiv 0 \pmod{2^{r(f)}}}$, with $r(f) = N_c$ on all-ones fuel. The carry recursion and reset hold exactly; clock and increment defects are bounded by $3\cdot 4^{-N_c}$.

For a source register put $\tau = 3\cdot 4^{-N_s}$. Its affine and tail-flip residuals are respectively $\tau P_{i,0}$ and $\tau(I-2P_{i,0})X_i$. The latter polar factor is unitary. Thus set $\Theta_{i1} = \tau P_{i,0}$, $V_{i1} = I$, $\Theta_{i2} = \tau I$ and $V_{i2} = (I-2P_{i,0})X_i$; (Q1.2) is exact.

On a static input eigenvalue $t > 0$ put $l = \lfloor \log_2(1/t) \rfloor$ and stop at the first pulse with one-count l, using age zero when $l = 0$. Set

$$s(t) = 2^{-\kappa(l)-1}, \qquad s(0) = 0. \tag{Q3.3}$$

At physical age h after reset, use the exact scalar counter history of Proposition P4.1. If h_* is the stop age, set

$$C_h = 2^{k(j(\min(h,h_*)))}t, \qquad D_h = 2^{j(\min(h,h_*))}s(t), \qquad H_h = 1_{h\le h_*}, \quad P_h = 1_{h=h_*},$$

where $j(h)$ counts pulses before physical age h. These histories are bounded, initialize correctly, and satisfy every non-penalty equation exactly, including reset before a programmed stop and the cyclic boundary. This is the direct substitution used in Proposition P5.2; it uses the finite count k, not a limiting density. All workspaces are diagonal in the existing checker and source-digit basis, and impose no new matrix factor. Every declared commutation is exact.

After powering, residual norms are at most $\max(3\cdot 4^{-N_c}, s(\tau))^q$. The compiler gives loss at most C_0 times its square and local dimension at most $c\,2^{2N_c+mN_s}$. Write $x = (n+\log_2 C_0)/(2q)$ and choose

$$N_c = \left\lceil \frac{x+\log_2 3}{2} \right\rceil, \qquad N_s = \left\lceil \frac{k(\lceil x\rceil - 2)+3}{2} \right\rceil.$$

Finite initial indices can be absorbed by increasing these sizes. Indeed $\lfloor \log_2(1/\tau)\rfloor = 2N_s - 2$, and for integers $y \ge 1$,

$$\kappa(l) \ge y \quad\Longleftrightarrow\quad l \ge k(y-1)+1.$$

This identity includes all plateaux of k. The choices give loss at most 2^{-n} and, since k has unit increments, $2N_c + mN_s = \Lambda + (m/2)k(\Lambda) + O(1)$. The upper states are maximally entangled, giving the same usable upper bound and quantum value one.

One joint measurement family. On an arbitrary original loss-ε strategy in local dimension d, the compiler transfer and residual powering give an unpowered weighted model in dimension at most cd, at precision

$$\delta = (C\sqrt{\varepsilon})^{1/q}, \qquad \lambda = \log_2(1/\delta) = n/(2q) - O(1).$$

Enlarge the fixed compiler constant so that $\lambda \le \Lambda$. Take an admissible J in (Q2.2) with

$$J \ge \lambda - a_1 \log_2(T(\lambda)+2) - O(1).$$

Monotonicity of T makes this choice admissible, after a fixed shift. The calibrated source clocks have a safe precision $\delta_s \ge \delta$, chosen as the maximum of δ and their explicit calibration bound. Since k has unit increments,

$$\log_2(1/\delta_s) \ge k(\Lambda) - O\big(\log T(\Lambda) + \log n\big). \tag{Q3.4}$$

Positive lower density and subexponential delay imply $t_s = \log_4(1/\delta_s) \ge c\lambda$ eventually. Use checker depth

$$M = \lfloor \lambda/2 - 30\log_4(2+\lambda/2)\rfloor$$

and source depth $M_s = \lfloor t_s - 30\log_4(2+t_s)\rfloor$. The within-clock errors are

$$C(M+1)^3 4^M \delta = O(\lambda^{-27}), \qquad C(M_s+1)^3 4^{M_s}\delta_s = O(\lambda^{-27}).$$

Every cross-clock commutator and flip-preservation error is a fixed polynomial in the depths times the original δ, because the complete clock-generator sets commute at that precision. These errors are exponentially small in λ. Thus a single $\eta = O(\lambda^{-26})$ controls the joint family of $b = 2M + mM_s = O(\lambda)$ binary measurements. Order the families fuel, time, then sources, each with increasing digits. Each flip preserves later measurements. By Theorem 8.1, $\log_2(cd) \ge b - o(1)$: its terms $Cb^3\eta + b[h_2(e)+e]$ tend to zero for $e = Cb^4\eta^2$. Combining this with (Q3.4) proves the dimension lower bound.

Extraction from the original mixed state. Use the proof of Proposition 12.2, rather than restricting its literal statement to the earlier source family. Every raw digit here is a projection conjugated by a unitary word; Alice's opposite word is again an exact projection, and her opposite flip is a unitary. On a purification of the actual mixed input, including its inaccessible environment, their mirror errors obey

$$C\ell\sqrt{\varepsilon} \le C\ell\delta^q \le C\ell\delta, \qquad \ell = O(\max(M, M_s)) = O(\lambda). \tag{Q3.5}$$

Alice's square-root-state commutators have the same polynomial word-length overhead. All these quantities are exponentially small and hence bounded by the common η above, also when $q = 1$. Bob's swap, triangular and preservation estimates are exactly those already proved for the joint count. Reverse its measurement order to put the preserved measurements before the current stage, as required by Lemma 12.1. All b stages are compulsory binary stages, with no inactive outcomes or inclusion operation. That lemma gives

$$E_D^{A\to B}(\rho_{AB}) \ge b - o(1) \ge \Lambda + \tfrac{m}{2}k(\Lambda) - O(\log T(\Lambda) + \log n).$$

Pure local dilation ancillas do not change the rate. This argument neither selects a pure component nor normalizes a favored clock branch. Taking the infimum over every original strategy proves (Q3.2). The compulsory contribution diverges, excluding a finite-dimensional perfect strategy. □

Q4 Exponent pairs and the witness-free classification proof

Use the real-number hierarchy convention of Zheng and Weihrauch [73]: finite liminfs of computable rational sequences form Σ_2^0, finite limsups form Π_2^0, and their common limit class is Δ_2^0. The following elementary naming fact will be useful.

Lemma Q4.1 (Simultaneous endpoint names)**.** *If $0 < L \le U < \infty$, $L \in \Sigma_2^0$ and $U \in \Pi_2^0$, there is a total computable bounded rational sequence with lower limit L and upper limit U. If both endpoints lie in $(0, 1]$, the sequence may be chosen in $[c_0, 1]$ for any fixed rational $0 < c_0 < L$.*

Proof. For $L < U$, choose a rational separator r strictly between them, a computable liminf name (u_s) for L, and a computable limsup name (v_s) for U. Fix rational bounds $a < L < U < b$. Interleave

$$\min(r, \max(a, u_s)), \qquad \max(r, \min(b, v_s)).$$

The first subsequence has liminf L and is at most r; the second has limsup U and is at least r. Their interleaving has the required endpoints. If $L = U$, membership in both classes gives a convergent computable name. For endpoints in $(0, 1]$, apply $x \mapsto \min(1, \max(c_0, x))$. Continuity and monotonicity preserve both endpoints. Upper clipping may change infinitely many terms when $U = 1$, which does not affect this conclusion. □

Theorem Q4.2 (The logarithmic exponent pairs)**.** *The positive finite lower/upper logarithmic dimension-exponent pairs of finite rational value-one BCS games are exactly*

$$\{(L, U) : 0 < L \le U < \infty,\ L \in \Sigma_2^0,\ U \in \Pi_2^0\}. \tag{Q4.1}$$

Every such pair is realized by one fixed game with no finite-dimensional perfect strategy and with the same pair for $d_G^{\to}(2^{-n})/n$. Separately, the positive finite lower dimension exponents are exactly the positive Σ_2^0 reals, and the positive finite upper ones exactly the positive Π_2^0 reals.

Proof. By Lemma 2.4, $D_G(2^{-n})$ is a total computable integer sequence under the value-one promise. Approximate its logarithm by rationals to error at most $1/n$. The normalized sequence has the same finite lower and upper limits, giving the two necessary hierarchy classes. This necessity concerns dimension, not an algorithm for distillable entanglement.

For sufficiency choose integers

$$q = \lfloor 1/(2L) \rfloor + 1, \qquad m = \lceil 4qU - 2 \rceil, \quad L' = \frac{4qL - 2}{m}, \quad U' = \frac{4qU - 2}{m}.$$

Then $m \geq 1$ and $0 < L' \leq U' \leq 1$. Positive rational affine maps preserve the respective arithmetic classes. Use Lemma Q4.1 and Lemma Q2.1 for a rational $c_0 < L'$. By Theorem Q3.1,

$$\frac{\log_2 D_G(2^{-n})}{n} = \frac{1}{2q} + \frac{m}{4q}\frac{k(\Lambda)}{\Lambda} + O(\log n/n),$$

and the same formula holds for usable demand. Every integer Λ occurs at $2q$ consecutive precision indices, giving lower limit L and upper limit U. These parameter and separator choices are existential; no uniform procedure extracting them from arbitrary endpoint names is asserted. To realize an arbitrary positive Σ^0_2 lower endpoint alone, pair it with a larger rational upper endpoint. For a positive Π^0_2 upper endpoint choose a smaller positive rational lower endpoint. The separate necessity arguments also apply when the other endpoint is infinite. □

Corollary Q4.3 (Local-heat pairs for the constructed games)**.** *In the reset model of Section 13, the polynomial-delay games above satisfy*

$$w_{\mathrm{local},G}(2^{-n}) = 2\Lambda + mk(\Lambda) + O(\log n).$$

Their local-heat coefficient pairs are $(2L, 2U)$, and joint-reset heat is bounded. This is an inclusion in the full heat-pair spectrum.

Proof. Apply the cited thermodynamic sandwich (13.1), namely $2d^{\rightarrow}_G \leq w_{\mathrm{local},G} \leq 2\log_2 D_G + 2\log_2 K_{\mathrm{ans}}$, to the matched profiles. The answer alphabet is fixed. The joint-reset bound is the same input from [24, Theorem 2.1]; no new thermodynamic inequality is used. □

Corollary Q4.4 (Classification without a game witness)**.** *Every positive Δ^0_2 logarithmic dimension coefficient has a matched original-state usable realization using no* MIP* = RE *witness. Together with Lemma 2.4, this proves Theorem 1.1. For each polynomial-delay profile construction, the explicit function $\Lambda + (m/2)k(\Lambda)$ is computable in polynomial time in n and approximates both demands to additive error $O(\log n)$.*

Proof. Take $L = U = \alpha$ in Theorem Q4.2; a convergent pulse name suffices. Its presentation uses only the polynomial compiler, Cantor clocks, carry, finite counter simulation, count and extraction already specified, with no call to the witness construction of Section 16. The converse is the computable-dimension limit argument. Dyadic losses imply the full $\varepsilon \downarrow 0$ law: monotonicity sandwiches a loss between adjacent dyadic losses, whose linear targets differ by $O(1)$. For the algorithmic assertion, run the fixed counter program through Λ pulses and evaluate its one-count. Polynomial delay makes this polynomial time. The error statement follows from Theorem Q3.1; it is an approximation to optimized demands, not an exact polynomial-time evaluation of either one. □

The exact-trace obstruction of Section 14 is compatible with these results. It concerns a fixed rational polynomial whose exact-trace minimum is the target coefficient. Here the initialized cylinder, history columns, multipliers and counted measurement family change

with precision. In an exact trace (Q2.2) forces every calibrated input to vanish, since positive lower density dominates the subexponential horizon; that zero is independent of the target endpoints. No stationary expectation of the output control is identified with $k(J)/J$. The witness-based proof in Section 16 retains its separate projection-minimum and exact-trace conclusions.

Q5 Profiles computable in small space

Corollary Q5.1 (Space-bounded target profiles)**.** *Let $A : \mathbb{N} \to \mathbb{N}$ be computable on input n in work space $O(\log n)$, respectively $o(n)$. Suppose fixed rationals $0 < \alpha \le \beta$ and $c \ge 0$ satisfy, for every $n' \ge n$,*

$$\alpha(n' - n) - c \le A(n') - A(n) \le \beta(n' - n) + c. \tag{Q5.1}$$

Some fixed finite rational value-one BCS game has both $\log_2 D_G(2^{-n})$ and $d_G^{\rightarrow}(2^{-n})$ equal to $A(n) + O(\log n)$, respectively $A(n) + o(n)$.

Proof. Choose integer $q \ge 1/\alpha$ and integer $m \ge \max(1, 2(2q\beta - 1))$, and put $\tau(J) = (2/m)(A(2qJ) - J)$. For $J' \ge J$ its increments lie between $(2/m)(J' - J) - 2c/m$ and $(J' - J) + 2c/m$. Set $k(0) = 0$ and

$$k(J + 1) = k(J) + 1_{k(J)<\tau(J+1)}.$$

The discrepancy is bounded by $1 + 2c/m + |\tau(0)|$. For its upper side, take the last increment before J and use the lower interval-slope bound after that point. For its lower side, take the last non-increment and use the upper interval-slope bound; if either point is absent, the initial value gives the same bound. Thus the word of unit increments has positive lower one-density and $\Lambda + (m/2)k(\Lambda) = A(2q\Lambda) + O(1)$.

Let $s(n)$ be the target computation's space use and set

$$S(n) = \max_{j \le n} s(j) + O(\log(n + 2)). \tag{Q5.2}$$

This includes input-head positions, loop bookkeeping and the $O(\log n)$ output bits, since the slope bound gives $A(n) = O(n)$. A halting deterministic computation cannot revisit a complete configuration. Its configuration count, including these data, is $2^{O(S(n))}$. Store tapes as counter stacks and simulate each elementary step by bounded copying and counter operations; stack sizes and the total sequential time through all inputs $j \le n$ remain $2^{O(S(n))}$. Computing the greedy word sequentially is therefore a finite counter program with that delay bound.

If $s(n) = O(\log n)$ the delay is polynomial. If $s(n) = o(n)$, its prefix maximum is also $o(n)$ and the delay is subexponential. Apply Theorem Q3.1, with fixed rescaling $n = 2qJ$, to obtain respectively $O(\log n)$ and $o(n)$ error. Finally (Q5.1) makes $A(n) - A(2q\Lambda) = O(1)$. □

Conversely a fixed counter program run for $T(J)$ steps has all counters at most $T(J)$ and its pulse word is computable in space $O(\log T(J))$. This describes small-space representatives produced by this architecture. It does not characterize the exact optimized finite-scale values of every game.

Q6 Power-scale exponent pairs

Theorem Q6.1 (Power pairs in the unit interval)**.** *For every $0 < L \le U \le 1$ with $L \in \Sigma_2^0$ and $U \in \Pi_2^0$, one fixed finite rational value-one BCS game has lower/upper pair (L, U) for both*

$$\frac{\log_2 \log_2 D_G(2^{-n})}{n}, \qquad \frac{\log_2 d_G^{\rightarrow}(2^{-n})}{n}.$$

The same pair is realized for pure entropy and formation demands. More precisely, the constructed game's two logarithmic profiles equal

$$2k(\lfloor n/2 \rfloor) + O(\sqrt{n}\log n). \tag{Q6.1}$$

Conversely finite lower and upper power exponents of $\log D_G$ *belong respectively to* Σ_2^0 *and* Π_2^0*. Thus dimension pairs in* $(0,1]$ *are classified by these endpoint conditions. No upper ceiling of one for arbitrary games, or arithmetic ceiling for arbitrary usable pairs, is asserted.*

Proof. Use the finite-moment source and full programmed source construction of Sections P2 and P5. Replace its convergent stream by the hold stream from Lemmas Q2.1 and Q4.1, with endpoints L, U and a fixed rational $c_0 < L$. This retains the original unguarded checker interface: at this scale its $O(n)$ logarithmic dimension is negligible compared with the source cost. No source-moment or extraction input is replaced by a logarithmic clock count.

The complete-column calibration proof gives

$$e \le \mathrm{poly}(n) 2^{-k(\lfloor n/2 \rfloor - O(\log n))} + O(2^{-n/2}).$$

The source mirrors have error $O(2^{-n/2})$, which fits this envelope because $k(J) \le J$. The original-state source extraction in Theorem P1.1 then yields

$$\log_2 d_G^{\rightarrow}(2^{-n}) \ge 2k(\lfloor n/2 \rfloor) - O(\log^2 n).$$

This is the source-level operational estimate already proved in the power part, on every original mixed state.

For the upper model fix one integer moment order $v > 1/c_0 + 1$. The source model of size N has $\tau(|R_i|^{2v}) \le C_v N^{-v}$. Put $l_0 = \lfloor \frac{1}{2}\log_2 N \rfloor - 1$. For the spectral measure of $|R_i|$, the dyadic band $(2^{-l-1}, 2^{-l}]$ has mass at most $\min(1, C_v 2^{-2v(l_0-l)})$. Its calibrated penalty is $s = 2^{-\kappa(l)-1}$; zero eigenvalues contribute zero. Apply the window lower-density bound between the l-th and l_0-th ones to obtain

$$\kappa(l_0) - \kappa(l) \le \frac{l_0 - l}{c_0} + O\big(\sqrt{\kappa(l_0)}\log(\kappa(l_0)+2)\big).$$

The fixed factor $1/c_0$ multiplying the window error is absorbed in the constant. For $l < l_0$, multiply the band mass by its squared penalty and sum the convergent geometric series with ratio $2^{2/c_0 - 2v} < 1$. For $l \ge l_0$, monotonicity of κ and the total mass bound suffice. Thus

$$\|S_i\|_2^2 \le C 2^{-2\kappa(l_0) + O(\sqrt{\kappa(l_0)}\log(\kappa(l_0)+2))}. \tag{Q6.2}$$

Choose a sufficiently large fixed B and set explicitly

$$J_* = \left\lceil n/2 + B\sqrt{n+2}\log_2(n+2) \right\rceil, \qquad \log_2 N = 2k(J_*) + 4.$$

Then $l_0 = k(J_*) + 1$ and $\kappa(l_0) > J_*$ exactly, including plateaux. Positive lower density gives $\kappa(l_0) = O(n)$. Increasing B or the fixed additive constant absorbs compiler constants in (Q6.2), so all penalties reach the required loss. Size the unguarded checker registers at $n + O(1)$ bits to control their overflow defects. The simultaneous source upper bound is

$$\log_2 D_G(2^{-n}) \le 3N + O(\sqrt{N}) + O(n),$$

hence $\log_2 \log_2 D_G(2^{-n}) \le 2k(\lfloor n/2 \rfloor) + O(\sqrt{n}\log n)$. Unit increments of k control the shift from J_*, and the maximally entangled upper states match the usable lower bound. Pure entropy and formation are squeezed by the same two bounds. Dividing by n proves the endpoint realization. Finally $D_G(2^{-n})$ is computable, so rational approximations to its double logarithm give the stated necessary liminf/limsup classes. □

The window condition is needed in the power upper bound because the source spectrum is controlled by moments. The logarithmic construction has only one nonzero calibrated source eigenvalue and needs no such spectral summation.

Q7 A time bound for whole profiles

Write $A_G(n) = \log_2 D_G(2^{-n})$. Finite-dimensional feasibility imposes a restriction on the entire function, stronger than the endpoint hierarchy alone.

Proposition Q7.1 (Doubly exponential computation)**.** *For every fixed finite rational value-one game with finite upper logarithmic dimension exponent, there is a constant c_G such that $D_G(2^{-n})$ is computable exactly in time $2^{2^{c_G n}}$. A rational approximation to $A_G(n)$ with error at most 2^{-n} is computable within the same form of time bound. More generally a requested logarithm precision has the additional standard numerical overhead; no uniform procedure finding c_G from a game description is asserted.*

Proof. For fixed local dimension d, pure states suffice for feasibility: with measurements fixed, a mixed state's winning score is an average of pure-state scores in the same spaces. This reduction concerns dimension feasibility only. Represent a pure state by a complex vector in $\mathbb{C}^{d^2}$ and each POVM effect by YY^*. Separate real and imaginary coordinates. Normalization, positivity-by-factorization, completeness and the winning inequality give an existential real sentence with $O_G(d^2)$ variables and constraints, degree at most six, and coefficient bit lengths $O_G(n)$. Its dense description has polynomial size in d, n.

Singly exponential existential-real algorithms, rather than unrestricted iterated quantifier elimination, give bit complexity $2^{O_G(d^2 \log(d+2))}\,\mathrm{poly}_G(n)$; see [5, Sections 2.2 and 2.5]. Search $d = 1, 2, \ldots$ for the first feasible dimension. The value-one promise ensures termination. A finite upper exponent gives $D_G(2^{-n}) \le 2^{C_G n}$ eventually, so the total search cost is at most $2^{2^{c_G n}}$ after enlarging the constant. Finitely many early indices are absorbed by a finite table. The integer answer is exact. Its logarithm is computed to the stated rational accuracy by standard arithmetic with $O_G(n)$ integer/output bits and precision n, negligible in this bound. A bounded-error logarithm, or its integer part, is enough for the diagonal arguments below. □

Theorem Q7.2 (Computable profiles outside every game)**.** *There is a total computable $A : \mathbb{N} \to \mathbb{N}$, with $A(0) = 0$ and $A(n+1) - A(n) \in \{0, 1\}$, such that*

$$\tfrac{1}{4} \le \liminf A(n)/n \le \limsup A(n)/n \le \tfrac{3}{4},$$

and no finite rational value-one game satisfies $A_G(n) = A(n) + o(n)$. More generally, for every $0 < L < U < \infty$ with $L \in \Sigma^0_2$, $U \in \Pi^0_2$, there is such an unrealizable computable profile with these exact endpoints, $A(0) = 0$, and integer increments between zero and a fixed integer C. The unit-increment condition in the first example is not asserted for arbitrary endpoints above one.

Proof. A direct repeated-index diagonal. Let $e(i)$ be a computable enumeration in which every machine index occurs infinitely often, and set $T(i) = 2^{2^{8^i}}$. Simulate machine $e(i)$ on input i for $T(i)$ of its own steps. If it halts with a bit, define b_i to be the opposite bit; otherwise define $b_i = 1$. This is a total computable sequence. Any machine eventually computing it within time $T(i)$ disagrees at arbitrarily large indices assigned to that machine. The computation defining b itself has no claimed bound T.

Let a continuous piecewise-linear function have slope $1/4 + b_i/2$ on $[4^i, 4^{i+1})$, and take its integer part, with a harmless nondecreasing finite prefix from $A(0) = 0$. Its integer increments are zero or one and its normalized asymptotic values lie in $[1/4, 3/4]$. Across block i,

$$A(4^{i+1}) - A(4^i) = \tfrac{3}{4}\,4^i + \tfrac{3}{2} b_i 4^i + O(1).$$

An $o(n)$ approximation at both endpoints recovers b_i eventually by comparing this difference with $(3/2)4^i$. If that approximation were A_G, Proposition Q7.1 would compute it in time

$2^{2^{O_G(4^i)}} < T(i)$ eventually. Bounded numerical error in the two logarithms is harmless. This contradicts the repeated-index diagonal without requiring an effective modulus of the assumed $o(n)$ error.

Every prescribed off-diagonal pair. Fix rationals

$$0 < a < L < c_1 < c_2 < U < b < C,$$

with C an integer. By Lemma Q4.1, take a total computable rational sequence $q_r \in [a, b]$ with endpoints L, U. Put $i_r = 2^{2^r}$ and $N_r = 4^{i_r}$. Define hard bits b_r by the same repeated-index diagonal, now using time $T(r) = 2^{2^{8^{i_r}}}$ on input r. On the bit interval $[N_r, 2N_r]$ set

$$A(n) = A(N_r) + \lfloor c_{1+b_r}(n - N_r) \rfloor. \tag{Q7.1}$$

On the gap from $2N_r$ to N_{r+1} use the fixed next target q_{r+1} and the greedy update

$$A(n+1) = \min\big(A(n) + C,\ \max(A(n), \lfloor q_{r+1}(n+1) \rfloor)\big). \tag{Q7.2}$$

Before N_0 choose any nonnegative prefix starting at zero with increments at most C. This defines a total computable integer function. Computing a target or a hard bit may take arbitrarily long; no target-runtime bound is imposed in this theorem. Both rules have increments in $[0, C]$, and $0 \le A(n) \le Cn$.

To prove that a gap catches its target, suppose it starts at m. If above the line, the constant value is caught by time at most $(C/a)m + O(1)$. If below, increasing by C catches a line of slope at most $b < C$ by time at most $(C/(C-b))m + O(1)$. Once caught, (Q7.2) tracks its integer part. There is therefore one fixed R such that every gap tracks by $Rm + O(1)$, independently of r and q_r. Since $N_{r+1}/(2N_r) \to \infty$, every sufficiently late gap ends on its target:

$$A(N_r)/N_r = q_r + O(1/N_r).$$

During a gap the normalized ratio remains in the hull of its starting ratio and the target, up to $O(1/n)$: above it waits, below it increases until caught, then tracks. During a bit interval the ratio is in the hull of $A(N_r)/N_r$ and c_1 or c_2, with the same floor error. As both slopes are strictly inside (L, U), all late ratios lie in $[L - o(1), U + o(1)]$; the values at N_r attain both endpoint limits. Thus the exact endpoints are L, U.

If $A_G = A + o(n)$, compare its values at $2N_r$ and N_r with threshold $((c_1 + c_2)/2)N_r$. The two possible increments in (Q7.1) differ by $(c_2 - c_1)N_r + O(1)$, so this decodes b_r eventually. By Proposition Q7.1, the decoder costs $2^{2^{O_G(4^{i_r})}} < T(r)$ for all sufficiently large r, again contradicting the diagonal. This proves the complete prescribed-endpoint assertion. □

When $L = U = \alpha$, every profile with those endpoints is $\alpha n + o(n)$ and Corollary Q4.4 realizes that class. The whole-profile obstruction is therefore off the diagonal, even though every admissible endpoint pair is realized by some other profile.

Q7.1 The remaining frontiers

Problem Q7.3 (The space/time profile gap)**.** Which classes of logarithmic dimension profiles modulo $o(n)$ are realized by finite rational value-one games? Under (Q5.1), Corollary Q5.1 realizes small-space representatives, while Proposition Q7.1 and Theorem Q7.2 impose a doubly exponential time restriction and exclude some computable bounded-increment targets.

The counter architecture has counters bounded by elapsed physical time, and polynomial-horizon stability permits only subexponential physical delay at logarithmic precision. It therefore supplies space-$o(n)$ representatives. Extending the realized range would require a

different effective precision-to-computation charge, or a stronger general feasibility algorithm narrowing the excluded range. The witness forcing of Section 16, which uses compactness, does not supply such a scale-pinned modulus.

Conjecture Q7.4 (The full usable endpoint ceiling)**.** For a finite rational value-one game, positive finite lower and upper logarithmic usable exponents belong respectively to Σ_2^0 and Π_2^0. The analogous endpoint ceiling holds at the power scale.

The matched games prove the realization inclusion, not this ceiling. Finite-dimensional feasibility computes dimension but does not compute the regularized directional rate. A counterexample may have operational endpoints outside these classes without contradicting the dimension results.

Sapir–Birget–Rips [62] provide the resource-function template for the time restriction: their stated Dehn-law digit conditions have double-exponential sufficiency and triple-exponential necessity, and their broader time-function characterization has additional hypotheses. No Dehn-to-game reduction is used here. Meyerovitch [49, Theorem 5.1] classifies lower, upper and existing entropy dimensions with classes $\Sigma_2, \Pi_3, \Delta_2$; the upper class differs because global pattern extension is not the computable finite-dimensional feasibility problem here. Hochman–Meyerovitch [36] give the finite-local-rule arithmetic-spectrum precedent, Zheng–Weihrauch [73] the liminf/limsup normal forms, and Brady–Bridson [10] the dense algebraic-scaling predecessor. These are comparisons of constructions and invariants, not priority or reduction claims.

Q8 Group profiles and the trace obstruction

There is an effective upper classification for group profiles, but transferring the realizations requires additional structure. A trace obstruction explains why the particular counter presentations here cannot be transferred by fixed polynomial decoders into algebraic group rings.

Q8.1 Effective bounds for group profiles

Fix a finite presentation $\Gamma = \langle S \mid R \rangle$ with distinguished central involution J. Let $W_\Gamma(e)$ be the minimum d admitting unitaries U_s on $\mathbb{C}^d$ and a positive ξ with $\operatorname{Tr} \xi^2 = 1$ such that

$$\|(U_r - I)\xi\|_2 \le e, \quad \|[U_s, \xi]\|_2 \le e, \quad \|(U_J + I)\xi\|_2 \le e.$$

Let $h_\Gamma(e)$ instead require $U_J = -I$ and normalized Frobenius relator errors at most e. These definitions use words in the free group, not a word-problem oracle. Clearly $W_\Gamma(e) \le h_\Gamma(e)$.

Theorem Q8.1 (Group arithmetic and time ceilings)**.** *Under finiteness at every positive precision, W_Γ, h_Γ are computable on positive rationals. Their finite lower logarithmic exponents belong to Σ_2^0, their finite upper exponents to Π_2^0, and existing finite exponents to Δ_2^0. If either profile has finite upper exponent, both integer dimensions at precision 2^{-n} are computable in time $2^{2^{O_\Gamma(n)}}$, and their logarithms to additive error one in the same time.*

For every admissible distinct pair $0 < L < U < \infty$, $L \in \Sigma_2^0$, $U \in \Pi_2^0$, there is a total computable bounded-increment profile A with these endpoints such that neither $\log_2 W_\Gamma(2^{-n})$ nor $\log_2 h_\Gamma(2^{-n})$ can equal $A(n) + o(n)$ for any finitely presented group.

Proof. In dimension d, use real and imaginary matrix coordinates for the unitaries and write $\xi = BB^*$, $\operatorname{Tr} \xi^2 = 1$. Squared defects and unitarity give rational polynomial equalities and inequalities of fixed degree in $O_\Gamma(d^2)$ variables. In the flat case impose $\|U_r - I\|_F^2 \le de^2$

and $U_J = -I$, avoiding an irrational scalar $1/\sqrt{d}$. Existential real decision and increasing search in d compute the minima under the finiteness promise. Rational approximations to $\log_2 W(2^{-n})/n$, or to the flat analogue, form a computable sequence with vanishing error. The normal forms used in Lemma Q4.1 and Theorem Q4.2 give the arithmetic inclusions. Monotonicity between adjacent dyadics gives the same finite endpoints over all real precisions.

We justify the dimension cap needed for the time assertion even when only the weighted upper exponent is known. For a weighted η-model put $P_t = 1_{[\sqrt{t},\infty)}(\xi)$. Spectral integration gives

$$\int_0^\infty \operatorname{Tr} P_t \, dt = 1, \qquad \int_0^\infty \|(U_r - I)P_t\|_F^2 dt \le \eta^2.$$

The same identity holds for the negative-J defect. Spectral coarea [68, Lemma 5.5] gives

$$\int_0^\infty \|[U_s, P_t]\|_F^2 dt \le 2\|[U_s, \xi]\|_F \le 2\eta.$$

Sum the finitely many errors. Some nonzero P_t has total squared error at most $C\eta \operatorname{rank} P_t$. For completeness, the coarea inequality follows by diagonalizing two positive matrices: the spectral overlap weights give a measure on pairs (x, y), and the integral of the squared difference of their threshold indicators is $|x^2 - y^2|$. Cauchy–Schwarz bounds its integral by the product of the Frobenius norms of the difference and sum. Apply this to ξ and $U_s \xi U_s^*$, whose sum norm is at most two.

Compress each generator to $P_t H$ and round its singular values to one. If $s_i \le 1$ are those values, $\sum(1 - s_i)^2 \le \sum(1 - s_i^2)$, the latter being the off-block squared mass. Thus normalized rounding costs $O_\Gamma(\sqrt{\eta})$. Telescoping every fixed relator, using the off-block bounds at its intermediate factors, has the same order. The negative defect permits replacing the rounded J by $-I$. Consequently

$$h_\Gamma(K\sqrt{\eta}) \le W_\Gamma(\eta)$$

for a fixed K. A finite upper exponent for either profile therefore bounds both dimensions at 2^{-n} by 2^{Cn}, after changing C. For the fixed-degree real system, singly exponential existential decision costs at most $2^{O_\Gamma(d^2 \log(d+1))} \operatorname{poly}(n)$ [5]. Sum this over $d \le 2^{Cn}$ to obtain the doubly exponential bound. The constants depend on the presentation and promise; no uniform test of that promise is claimed.

Here is the whole-profile diagonal explicitly. Choose rationals $0 < a < L < c_1 < c_2 < U < b < C$, with integer C, and a computable rational sequence $q_r \in [a, b]$ with endpoints L, U. Put $i_r = 2^{2^r}$, $N_r = 4^{i_r}$ and $T(r) = 2^{2^{8^{i_r}}}$. Enumerate machine indices with every index repeated infinitely often. On input r simulate the assigned machine for $T(r)$ steps and choose the opposite bit if it returns a bit, and one otherwise. This total computable bit sequence b_r has no eventual $T(r)$-time algorithm, by its repeated-index diagonal.

On $[N_r, 2N_r]$ define $A(n) = A(N_r) + \lfloor c_{1+b_r}(n - N_r) \rfloor$. In each following gap use

$$A(n+1) = \min\{A(n) + C,\ \max\{A(n), \lfloor q_{r+1}(n+1) \rfloor\}\}.$$

Start from a bounded-increment nonnegative prefix. From above a target line of slope at least a catches a constant value after a fixed multiple of the gap's initial index; from below speed $C > b$ catches it in a fixed multiple as well. Since $N_{r+1}/N_r \to \infty$, every late gap reaches and then tracks its target. Thus $A(N_r)/N_r = q_r + O(1/N_r)$. Within a gap the normalized ratio stays in the hull of its initial ratio and its target, and within a bit block in the hull of its initial ratio and c_1 or c_2, up to vanishing floor error. The endpoints are exactly L, U, and increments lie in $[0, C]$.

An $o(n)$ approximation by either group profile would recover b_r eventually by comparing its increment from N_r to $2N_r$ with $((c_1 + c_2)/2)N_r$. The time bound already proved makes

this a decoder of time $2^{2^{O_\Gamma(4^{i_r})}} < T(r)$ eventually, a contradiction. This treats both profiles separately without a convergence modulus. □

These are necessary group classes, not group realizations. Constant precision rescaling shifts a dyadic index by $O(1)$; a square-root game interface divides the corresponding exponent by two. The linear-system comparison [50, Corollary 5.5] respects these conventions but does not supply the missing realization inclusion.

Q8.2 Traces forced by the counter presentation

Let $\mathcal{A}(\mathcal{P}_\pi)$ be the algebraic unitary-domain presentation of the counter construction for a program π. Encode a projection as $(1-s)/2$ for a self-adjoint involution s, and a positive contraction as $(2+u+u^*)/4$ for a unitary u. We use the carry-clock presentation of Section Q1; the unguarded presentation has the same exact digit action. Powered residuals are understood in positive representations. In particular, $(p^*p)^q = 0$ implies $p = 0$ in the positive GNS representation, not as an abstract rule about nilpotents in an arbitrary algebra.

Lemma Q8.2 (Trace rigidity of initialized controls)**.** *Every tracial state of this bounded presentation assigns to the control projection Q_σ the value*

$$c_\sigma = \sum_{r\ge 0} 2^{-2r-1} N_\sigma(2^r),$$

where $N_\sigma(M)$ counts the visits to σ at physical times $0 \le h < M$ of the initialized classical program. A single visit at time h contributes $(2/3)4^{-\operatorname{bitlength}(h)}$, with $\operatorname{bitlength}(0) = 0$.

Proof. Pass to the finite von Neumann GNS algebra, with faithful normal trace. The positive clock relation separates the low and high intervals, so its digit projection is the high spectral digit of its Cantor coordinate. Iteration gives all later digits as spectral functions. The flip fixes the tail and interchanges the first digit; trace invariance therefore makes that digit fair independently of the tail. Iterating the shift proves the fair Bernoulli law. Commutation of the two clock copies, together with their separate flips and shifts, proves joint independence of all fuel and time digits.

Let Γ_r be the fuel cylinder with exactly r initial ones. Its trace is 2^{-r-1}. The carry recursion has remainder norm at most 4^{-n} and implements addition by one on time digits. On Γ_r, reset is the indicator of time residue zero modulo 2^r. The commuting cylinder projections

$$C_h = \Gamma_r 1_{\{\text{time}\equiv h \bmod 2^r\}}, \quad 0 \le h < 2^r,$$

have traces 2^{-2r-1} and are cyclically permuted by U. Every instantaneous variable commutes with them. Initialization fixes age one, zero counters and initial control on C_0.

Induct on h. The age and counter equations force the scalar values $AC_h = (h+1)^{-1}C_h$, $Y_iC_h = n_i(h)(h+1)^{-1}C_h$. If $n_i(h) > 0$, $Y_iE_i = 0$ forces $E_iC_h = 0$. If $n_i(h) = 0$, put $R = (I - E_i)C_h$. The positive witness equation gives $RW_iR = -(h+1)^{-1}R$; positivity forces $R = 0$. Thus the zero test, the next control and the next counter are exactly the deterministic ones, including cyclic reset. In particular $Q_\sigma C_h = 1_{\sigma(h)=\sigma}C_h$; multiplicity spaces carry no alternative history.

The fuel-all-ones projection has trace at most 2^{-R} for every R, hence zero. Normality now sums the cylinder identities and gives the formula. Its total over all controls is $\sum_r 2^{-r-1} = 1$. A visit at h occurs in precisely the terms $2^r > h$, whose geometric sum is $(2/3)4^{-\operatorname{bitlength}(h)}$, including $h = 0$ and powers of two. □

Lemma Q8.3 (An elementary irrational marker)**.** *Any of the polynomial-delay pulse programs can be combined with an elementary counter program so that a finite sum of its orthogonal control projections has irrational trace in every tracial state. The pulse word and all profile and exponent conclusions are preserved.*

Proof. Use counters L, O, T, I, A, initially zero, and four increments to initialize $L = 4$. The following loop uses seventeen controls. Write $+X; v$ for increment and jump, $-X; v_0, v_1$ for decrement-or-zero-test, and $0; v$ for a NOP.

Control	Instruction	Control	Instruction
MARK	0; CO	CO	$-L$; RO, CO1
CO1	$+O$; CO2	CO2	$+T$; CO
RO	$-T$; OUTER, RL	RL	$+L$; RO
OUTER	$-O$; CLEAR, CI	CI	$-L$; RI, CI1
CI1	$+I$; CI2	CI2	$+T$; CI
RI	$-T$; INNER, RL2	RL2	$+L$; RI
INNER	$-I$; OUTER, AI	AI	$+A$; INNER
CLEAR	$-L$; MOV, CLEAR	MOV	$-A$; MARK, ML
ML	$+L$; MOV		

The copy through T preserves L while putting it in O. Each of the L outer iterations copies L into I and transfers I into A. Clearing L and moving A back squares L and leaves all scratch counters zero. Initial copy and restoration cost $5L + 2$. Each outer iteration costs $7L + 3$ internally, and the outer tests cost $L + 1$, giving $7L^2 + 4L + 1$. Clearing costs $L + 1$, moving costs $2L^2 + 1$, and MARK costs one. Thus there are exactly $9L^2 + 10L + 6$ instructions between MARK visits. All counters are at most elapsed elementary time.

The first visit times are $v_i = 4, 194, 2664, 595054$, with $L_i = 4, 16, 256, 65536$. In general $L_i = 4^{2^i}$ and $v_i = 9L_i + O(L_{i-1})$. Alternate one main instruction with one marker instruction, using finite product controls. Marker visits then occur at $t_i = 2v_i + 1$; their bit lengths satisfy $k_i = 2^{i+1} + O(1)$, initially $4, 9, 13, 21$. They are eventually distinct with unbounded gaps. By Lemma Q8.2, the projection for marker visits, a finite orthogonal sum of product controls, has trace $(2/3)\sum_i 4^{-k_i}$. A finite prefix only changes a rational number. After multiplication by $3/2$, its base-four digits have infinitely many isolated ones with unbounded zero gaps, and are not eventually periodic. Its trace is irrational. At least one constituent control has irrational trace; no effective identification of that constituent is required.

Allowing finite sets of output controls gives exactly doubled main delays. To retain the literal convention of two output NOP controls, store the finite return tag in extra counters, pass through the unique appropriate output control and dispatch through a finite decision tree. This takes a bounded number of primitive instructions; pad to a fixed constant slowdown. The marker gap argument survives this change. Pulse labels and their order are unchanged, polynomial pulse delay is preserved, and constant physical slowdown changes the logarithmic profile accounting by only $O(\log n)$. Thus every admissible matched pair still has a realizing presentation with this obstruction. □

Q8.3 Canonical group-ring traces

Lemma Q8.4 (Rational matrix idempotent traces)**.** *For every group Γ and idempotent $E \in M_N(\mathbb{C}[\Gamma])$, $(\mathrm{Tr} \otimes \tau_0)(E)$ is rational, where τ_0 is coefficient at the identity. Every nonzero projection has strictly positive canonical matrix trace.*

Proof. Corollary 1.2 of Burger–Valette already gives rationality on K_0, hence for matrix idempotents. The following optional normalization also derives it directly from Zalesskii's scalar theorem for arbitrary idempotents in $\mathbb{C}[\Gamma]$; see Burger–Valette [13, Theorem 1.3 and

Corollary 1.2]. To see the matrix normalization directly, take the order-N^3 finite Heisenberg group. Its primitive central-character corner is $M_N(\mathbb{C})$: the shift and clock matrices give its irreducible N-dimensional representation, and their N^2 products span the corner. The central-character idempotent has identity coefficient $1/N$. On this simple block the canonical trace is therefore $N^{-2}\,\mathrm{Tr}$. Embed E as an idempotent in $\mathbb{C}[H_N \times \Gamma]$. Scalar rationality there gives $N^{-2}(\mathrm{Tr}\otimes\tau_0)(E) \in \mathbb{Q}$, proving the claim. For projections, positivity and faithfulness follow in the regular representation; equivalently the canonical trace of X^*X is the sum of squared coefficients, also for matrices. Kaplansky's positivity and endpoint result is a separate classical input [13, Theorem 2.1]. □

Theorem Q8.5 (The counter presentation has no algebraic group-ring decoder)**.** *For the marker-augmented counter presentation above, there is no nonzero $*$-homomorphism into $M_N(\mathbb{C}[\Gamma])$, for any N and any group Γ, including nonunital homomorphisms. The same holds for its compiled Boolean algebra and for targets that are negative-J corners of solution-group rings, or finite central-character corners with rational trace normalization.*

Proof. Suppose α is such a map, and put $p = \alpha(1)$. This is a nonzero projection. With $t = \mathrm{Tr}\otimes\tau_0$, Lemma Q8.4 gives rational $t(p) > 0$. Hence $t\alpha/t(p)$ is a tracial state of the source presentation. Domain encodings are positive in this representation, so Lemmas Q8.2 and Q8.3 apply. The marker projection and its image are projections; the latter has rational trace, contradicting its forced irrational normalized trace.

The compiler proof gives a unital map $\iota : \mathcal{A}(\mathcal{P}) \to \mathcal{A}(\mathcal{B})$, not merely a statement about approximate models: each reconstructed source generator is a fixed polynomial in Boolean generators, and every source equation is a finite ideal consequence of Boolean relations. The explicit reconstruction identities in Section 3, including its orthogonality identity and gadget converses, give precisely these consequences. Thus a hypothetical nonzero map β from the Boolean algebra composes with ι to a nonzero map, since $\beta\iota(1) = \beta(1) \neq 0$. Injectivity of ι is unnecessary.

The negative-J quotient of a solution-group ring is the corner by $(1-J)/2$, with normalized canonical trace $2\tau_0$. Finite central-character corners likewise have rational normalization, so the same argument applies. This does not conflict with Slofstra's group embedding theorem [67], which is a theorem about groups and does not embed every Boolean algebra into a solution-group ring. □

Corollary Q8.6 (An approximate version with a canonical limit)**.** *Suppose a group with central J has a hyperlinear trace τ with $\tau(J) = -1$. Put $H = \{g : |\tau(g)| = 1\}$ and $\chi = \tau|_H$, and assume χ has finite image. There cannot be fixed unital polynomial decoders for the preceding counter presentation that are $O(e)$-sound in* every *flat e-model of that group.*

Proof. Equality in GNS Cauchy–Schwarz makes every $h \in H$ act on the cyclic vector by $\chi(h)$; the trace property gives a normal subgroup and conjugation-invariant character. The tensor traces $\tau|\tau|^{2m}$ are hyperlinear and negative on J, and converge to the trace τ_∞ equal to χ on H and zero elsewhere. A diagonal choice of their microstates proves hyperlinearity of the limit. Their negative-J defect tends to zero, so replacing J by $-I$ gives flat negative models with vanishing fixed-relator defects. Universal soundness forces each decoded relation-square trace to vanish in this limit.

The quotient is the twisted group ring of Γ/H. Its different coset basis vectors are orthogonal in the canonical GNS trace, which is faithful on finite sums. Consequently the decoded relations hold as algebraic identities there. Let $K = \ker\chi$. Then K is normal and H/K is finite central in Γ/K. The twisted ring is the corresponding character corner of $\mathbb{C}[\Gamma/K]$, with rationally normalized canonical trace. This supplies the forbidden unital

homomorphism in Theorem Q8.5. Soundness on a selected microstate sequence would not justify this argument, and finite image is essential to the stated reduction. □

The obstruction uses rationality of idempotent traces in algebraic group rings. Projections in group von Neumann algebras behave differently: Grabowski embeds Turing machines into integral group rings and obtains kernel projections of irrational and transcendental trace [33], following Austin's irrational example [3]; see also Pichot–Schick–Żuk [61]. These kernel projections are analytic spectral functions of group-ring elements, rather than polynomial functions, which is why Theorem Q8.5 excludes only fixed polynomial decoders.

This obstruction concerns these presentations and algebraic polynomial decoders. Precision-dependent analytic truncations are one possible route around it. A fixed infinite analytic decoder lies outside the algebraic group ring, and another presentation realizing the same profile can have different control traces. No conclusion that all group realizations require precision dependence follows from this theorem; the group realization problem remains open.

Q9 Weak pure rounding of original strategies

The remaining usable-spectrum question is an optimized fixed-game problem. All distillation below is Alice-to-Bob on the original mixed state, whose purifying environment is inaccessible. Pure comparison strategies may use different measurements and dimensions.

Paddock [59, Corollary 1.3] rounds near-perfect BCS and synchronous strategies to near-perfect strategies with maximally entangled states. Weak pure rounding asks instead for pure strategies whose entropy is bounded by the original state's one-way distillable rate.

Proposition Q9.1 (A sufficient one-way purification criterion)**.** *Let a finite one-way local instrument on an original strategy state ρ produce Bob-flagged actual game-strategy branches τ_j, with probabilities p_j, losses ℓ_j, mean loss $b = \sum_j p_j \ell_j$ and mean joint entropy $s = \sum_j p_j S(\tau_j)$. Arbitrary local garbage erasure and nonflat branches are allowed. Then*

$$p_G(b) \le E_D^{A\to B}(\rho) + s + 1.$$

Consequently, for every fixed $0 < a < 1$, selecting at each small ε an $o(L)$-near-minimizer for $d_G^{\rightarrow}(\varepsilon)$ and such an instrument with $b \le \varepsilon^a$, $s = o(L)$, $L = \log_2(1/\varepsilon)$, implies weak pure rounding

$$p_G(\varepsilon^a) \le d_G^{\rightarrow}(\varepsilon) + o(L).$$

This criterion is sufficient; its necessity is not asserted.

Proof. The output state is $\omega = \sum_j p_j \tau_j \otimes |j\rangle\langle j|_B$. One-way monotonicity and hashing [23] give

$$E_D^{A\to B}(\rho) \ge S(\omega_B) - S(\omega) = \sum_j p_j (S(\tau_{j,B}) - S(\tau_j)).$$

The Shannon flag entropies cancel. Bob's flag may include his private outcome; Alice need not learn it.

For each branch fix its game measurements and losing effect Λ_j. Put $\ell_j = \operatorname{Tr} \tau_j \Lambda_j$ and restrict to its support. The Hermitian matrix $T = \sqrt{\tau_j}(\Lambda_j - \ell_j I)\sqrt{\tau_j}$ is traceless. Fourier-transform an eigenbasis of T to obtain an orthonormal basis u_i with $\langle u_i, T u_i\rangle = 0$. The vectors $\sqrt{\tau_j} u_i$ decompose τ_j; every nonzero normalized vector has loss exactly ℓ_j. Entropy concavity on Bob's marginal selects one of entropy at most $S(\tau_{j,B})$. Denote this pure strategy's entropy by e_j. The preceding inequality gives $\sum_j p_j e_j \le E_D^{A\to B}(\rho) + s$.

Minimize the mean of the fixed numbers e_j over weights $q_j \ge 0$ with $\sum q_j = 1$ and $\sum q_j \ell_j \le b$. The weights p_j are feasible. A vertex optimum has at most two positive entries,

since there are just normalization and at most one active loss constraint. Coherently superpose those at most two pure strategies with orthogonal local flags and block-diagonal measurements. The resulting pure strategy has loss at most b and entropy $H(q) + \sum q_j e_j \le 1 + E_D^{A\to B}(\rho) + s$. This proves the bound. These branchwise pure witnesses are an existence argument; they need not form a protocol implemented by Alice while ignorant of Bob's private flag. Finally apply the bound to the selected near-minimizer and use monotonicity of p_G. □

After refining all local discarded registers into fine Kraus outcomes, let J_A be Alice's communicated outcome and J_B Bob's private outcome. Each conditional retained-AB plus environment-E state is pure. The exact identity

$$s = \sum_{j_A, j_B} p_j S(E|j) = S(E) - \chi(J_A, J_B : E)$$

expresses the residual joint entropy. All erased registers must be included in this refinement. The instrument criterion concerns a selected near-minimizing family, not all states or all inefficient strategies.

Corollary Q9.2 (What weak rounding would classify)**.** *For a fixed game satisfying weak rounding for every $a \in (0,1)$, if either $d_G^{\rightarrow}(\varepsilon)/L$ or $p_G(\varepsilon)/L$ has a finite limit, both do and their limits agree. A positive common coefficient belongs to Δ_2^0. In the reset model the corresponding local-heat coefficient is twice it.*

Proof. Always $d_G^{\rightarrow} \le p_G$. If $d_G^{\rightarrow}/L \to \alpha$, substitute $\delta = \varepsilon^a$ into weak rounding to bound the upper pure coefficient by α/a; its lower coefficient is at least α. Let a increase to one. If instead $p_G/L \to \alpha$, weak rounding gives lower usable coefficient at least $a\alpha$, while $d_G^{\rightarrow} \le p_G$ gives upper coefficient at most α. Again let a increase to one. For the arithmetic claim apply the effective truncation argument of Proposition P6.1: its computable rank-restricted entropy minima q_n satisfy $p_G(2^{-n}) \le q_n \le p_G(2^{-n-1})/(1-t_n)$, with $t_n \to 0$. Thus $q_n/n \to \alpha$, so the coefficient is limit-computable. The heat sandwich in Section 13 gives twice this coefficient in the reset model. □

Q9.1 One obstruction to effect-only purification

The next example fixes an effect and the underlying space. It explains why a general proof must use structure specific to games; it is not a separation for optimized strategies of one fixed game.

Proposition Q9.3 (A high-entanglement subspace obstruction)**.** *For arbitrarily large d there are a projection P on $\mathbb{C}^d \otimes \mathbb{C}^d$, effect $\Lambda = I - P$ and zero-loss state $\rho = P/r$, $r = \operatorname{rank} P$, such that*

$$E_D^{A\to B}(\rho) \le k_d = \log_2(d^2/r) = O(\log\log d),$$

while every pure vector of effect loss ℓ has entropy at least

$$(1 - \sqrt{\ell}) \log_2 d - c_0 - 1, \qquad c_0 = 1 + 1/\ln 2.$$

For any one-way channel on ρ and comparison pure vectors in this same space with this same effect, let their mean loss be b, the mean joint branch entropy s, and their mean entropy allowance above branch Bob entropy c. Then

$$s + c \ge (1 - \sqrt{b}) \log_2 d - c_0 - 1 - k_d.$$

Proof. Hayden–Leung–Winter [35, Theorem IV.1], with $d_A = d_B = d$, $\alpha = 1$ and $\Gamma = 1/1753$, provide a subspace $\mathcal{S}$ of dimension

$$r \geq \frac{\Gamma d^2}{2(\log_2 d)^{5/2}}$$

for sufficiently large d, in which every pure vector has entropy at least $\log_2 d - c_0$. The factor two absorbs the floor in their formula. Take P to be its projection. The separable state I/d^2 and the relative-entropy bound on distillation [71] give

$$E_D^{A\to B}(\rho) \leq E_R(\rho) \leq D(\rho\|I/d^2) = k_d \leq \tfrac{5}{2}\log_2\log_2 d + \log_2(2/\Gamma).$$

If a pure vector has loss $\ell < 1$, its normalized projection into $\mathcal{S}$ is at half trace distance $\sqrt{\ell}$. Partial trace contracts that distance, and entropy continuity loses at most $\sqrt{\ell}\log_2 d + 1$. This proves the pure lower bound. For $\ell = 1$ the displayed lower bound is negative and automatic.

For channel branches τ_j and same-effect pure comparison vectors of loss ℓ_j and entropy at most $S(\tau_{j,B}) + c_j$, flagged hashing as in Proposition Q9.1 gives

$$E_D^{A\to B}(\rho) + s + c \geq \sum_j p_j E(\psi_j) \geq (1 - \sum_j p_j\sqrt{\ell_j})\log_2 d - c_0 - 1.$$

Use concavity of the square root and $E_D^{A\to B}(\rho) \leq k_d$. No enlargement by coherent flags is used in this fixed-effect comparison. □

Thus $b = o(1)$ forces residual entropy plus allowance $(1 - o(1))\log d - O(\log\log d)$. A universal effect-only purification theorem cannot supply the required small residual. Actual game effects have finitely many positive product summands tied to common POVM constraints, and the game optimization permits changing those measurements. The dimension-dependent generic subspace supplies neither a fixed game nor a universal lower bound under that optimization. It therefore does not refute weak pure rounding.

17 Open spectra and sharper resource laws

The dimension classification concerns a positive finite limit and a fixed rational finite game. It leaves the following questions, including the distinction between dimension and operational entanglement.

Problem 17.1 (Zero without perfection)**.** Can a finite rational value-one BCS game with no finite-dimensional perfect strategy have dimension exponent zero? Can its logarithmic usable coefficient be zero with diverging sublogarithmic usable demand?

Finite-dimensional perfection forces a zero dimension exponent, but this does not settle the converse. The positive lower bounds in our constructions exclude that case without addressing subpolynomial dimension growth in other games. The witness-based architecture forces coefficient at least $3/(4q)$ for every fixed residual power q, from its three compulsory clocks alone. The carry-checker architecture of Section Q1 forces at least $1/(2q)$ from its two compulsory checker clocks. A zero example needs a verifier cheaper than its compulsory clocks or a different mechanism; changing flags or taking one larger fixed power cannot produce zero.

Problem 17.2 (Weak pure rounding and the usable spectrum)**. Weak pure rounding conjecture.** For every fixed finite rational value-one game and every fixed $0 < a < 1$, does

$$p_G(\varepsilon^a) \le d_G^{\rightarrow}(\varepsilon) + o(\log_2(1/\varepsilon))$$

hold for the original mixed-state demand, with the purifying environment inaccessible? The finite-logarithmic regime $d_G^{\rightarrow}(\varepsilon) = O(\log(1/\varepsilon))$ is the remaining case relevant to positive usable coefficients.

Paddock's maximally entangled rounding [59, Corollary 1.3] is a close precedent, but does not bound the rounded entropy by the original state's distillable rate. Every positive Δ_2^0 coefficient is already realized by Theorem 1.1; the opposite inclusion remains open because finite-dimensional feasibility does not compute a regularized distillation rate. A sufficient theorem is the one-way instrument criterion of Proposition Q9.1: on a selected $o(L)$-near-minimizer, Bob-flagged actual game-strategy branches should have mean loss at most ε^a and mean joint entropy $o(L)$, allowing arbitrary garbage erasure and nonflat states. After fine Kraus refinement the latter is exactly $S(E) - \chi(J_A, J_B : E) = o(L)$, including Bob's private outcomes and all discarded registers. This operational condition is sufficient, not claimed necessary. By Corollary Q9.2, weak rounding would identify the usable and pure coefficients and conditionally the doubled heat coefficient. The single effect-only obstruction in Proposition Q9.3 shows why arbitrary bipartite effects cannot establish it.

A refutation must give one fixed G, fixed $a \in (0,1)$, $\gamma > 0$ and $\varepsilon_i \to 0$ with $p_G(\varepsilon_i^a) \ge d_G^{\rightarrow}(\varepsilon_i) + \gamma \log_2(1/\varepsilon_i)$, using actual original-state rate upper bounds and universal pure-strategy lower bounds. Failure of a restricted instrument or a varying-effect example does not meet these quantifiers.

Problem 17.3 (Linear systems and groups)**.** Does the positive Δ_2^0 realization theorem hold for linear-system games, with matched original-state usable laws? Which admissible group-profile endpoint pairs and whole profiles are realizable by finitely presented groups?

Theorem Q8.1 gives the necessary arithmetic classes, the doubly exponential time restriction and whole-profile exclusions, not group realizations. Theorem Q8.5 forbids fixed polynomial matrix group-ring decoders for the specified irrational-control-trace counter presentations. Precision-dependent analytic truncations are one possible way around that

obstruction. Fixed infinite analytic decoders and different realizing presentations remain outside its scope; a blanket necessity of precision dependence is not established. Grabowski's Turing-machine encodings in group rings [33], following Austin [3] and with related constructions of Pichot–Schick–Żuk [61], are natural starting points for analytic decoders. Any proposed transfer must preserve simultaneous finite upper models and the original-state lower bound.

Problem 17.4 (Power-law usable demand)**.** Every Δ^0_2 exponent in $(0,1]$ has matched original-state usable, log-dimension, pure and formation realizations. Is every existing positive usable exponent in this interval Δ^0_2? Which existing finite power limits above one occur for a fixed finite rational game?

The local lamp class has ceiling one and positive existing log-dimension, pure and formation spectra exactly $\Delta^0_2 \cap (0,1]$ by Corollary P6.10. Its usable spectrum has the realization inclusion and the same upper bound; the arithmetic reverse inclusion remains open. The stopped checker with multipliers $a > b \geq 2$ saturates one. Beyond this class, Theorem P6.11 gives fixed rational BCS games with infinite upper exponents for log dimension, pure entropy and formation. These are upper exponents, not existing limits. Existing finite limits above one remain open. Theorem P6.16 and Proposition P6.17 state a conditional stronger-source route and necessary costs; Proposition P6.14 restricts universal global pairwise certificates. Theorem Q6.1 realizes admissible endpoint pairs within $(0,1]$, while Conjecture Q7.4 states the remaining general usable-pair ceiling.

The endpoint classes are settled for dimension: by Theorem Q4.2 the positive finite lower/upper logarithmic dimension pairs are exactly $0 < L \leq U$ with $L \in \Sigma^0_2$ and $U \in \Pi^0_2$, each with a matched original-state usable realization. Computability of finite-precision dimension separates the upper class from Meyerovitch's upper entropy-dimension class Π_3 [49, Theorem 5.1].

Problem 17.6 (The full local-heat spectrum)**.** Which positive finite local-reset heat coefficients occur in the block thermodynamic model of Section 13?

Equation (13.7) places this spectrum between the positive Δ^0_2 reals and the positive reals left-c.e. in $0'$. Neither the dimension classification nor the arithmetic upper bound proves equality for heat.

Corollary Q4.3 realizes every doubled admissible logarithmic endpoint pair in the local-reset model. It supplies the pair inclusion, while the full heat spectrum and a general heat-pair ceiling remain open. Problem Q7.3 concerns whole profiles rather than their two endpoints.

How this work was done

The authors selected the questions and proof standards, directed the research, and take responsibility for every claim. AI systems from more than one developer assisted with exploration, derivation, programming and drafting.

The mathematical evidence consists of the proofs in this paper and the explicitly cited external theorems. The ancillary programs provide reproducible finite checks. They test the compiler gadgets, finite clock and reader models, extraction protocols, gate and stop-or-continue histories, finite-moment lamps, the affine clock lift, stopped calibrators and scaled counters. The independent engine in `verify/` checks the finite models at $N = 4, 6, 8$. The ancillary README gives the commands, outputs, dependencies and limitations of each suite.

The power checks include finite CAR and affine calculations, exact scalar histories, ordered noncommutative identities and small weighted perturbations. They do not implement a complete universal counter compiler, instantiate the full compiled game, establish a uniform

stability theorem by sampling, or certify a noncomputable density. The asymptotic conclusions rest on the written proofs.

The profile suites check exact carry-clock and reset equations, initialized cylinders, the hold scheduler, finite calibration histories and profile accounting. An independent suite tests overflow-state deletion, noncommuting shared calibration inputs, greedy tracking, endpoint rescaling, hold accounting and the diagonal decoder. These finite tests do not establish the operator estimates, asymptotic endpoints or non-realizability theorem. The proofs in the profile part give those conclusions.

The ceiling and group suites additionally check simultaneous bad-block bounds, exact scalar calibrator histories, stopped-column saturation, affine hash collisions, weighted sequential POVMs, finite exact counter histories and the literal elementary irrational-marker program. They test the matrix normalization of the canonical group trace. The hierarchy argument, universal pairwise restriction, positive-GNS trace rigidity and optimized weak-rounding criterion are written proofs, not consequences of numerical sampling. The effect-only example is a cited subspace theorem with its explicit entropy comparison; it is not presented as a fixed-game counterexample.

A Lean formalization of Sections 2–10 is in progress and incomplete. Sections 11–17, the power-scale part and the profile part are not yet formalized. This preprint has not been peer reviewed.

A Why exact periodic verification is insufficient

This appendix records two obstructions that explain the use of approximate clocks and vanishing guard amplitudes. Neither is a general impossibility theorem for finite presentations. For the distinction between locally invariant finite marginals and globally extendible measures, see Chazottes, Gambaudo, Hochman and Ugalde [16]; the compactness argument below uses all box sizes, not extension of one finite marginal.

Proposition A.1 (Periodic extrema)**.** *Let X be a nonempty shift of finite type on $\mathbb{Z}^r$ and f a rational local observable. Put $\alpha = \min_\mu \int f \, d\mu$, where μ ranges over invariant probability measures. If averages of periodic configurations approach α, then α is computable.*

Proof. On each finite box, consider probability distributions on locally admissible patterns and impose every translation-invariance equation whose two supports stay in the box. This is a finite rational polytope. Minimize the expected value of f at the origin, obtaining a computable rational lower bound ℓ_N. Restriction from larger boxes makes these bounds monotone. A convergent subsequence of minimizers has consistent marginals on every finite support, gives an invariant measure on X, and has expectation $\lim \ell_N$. Thus $\ell_N \uparrow \alpha$. Periodic configurations and their finite fundamental domains are enumerable; local validity and their rational averages are decidable. Their successive minimum averages give upper bounds converging to their infimum. By hypothesis that infimum is α. Search until the upper and lower bounds differ by the desired error. □

Hochman–Meyerovitch's noncomputable extremal frequencies [36] therefore cannot be combined with arbitrarily good exact periodic models of those minima. In the present upper construction the workspace histories close as labels, but the complete clocks remain approximate. Proposition A.1 does not apply to those clocks.

Theorem A.2 (Extraction under a fixed rejection gap)**.** *Suppose $\mathbb{C}^d = \bigoplus_{v\in\mathcal{V}} \mathcal{H}_v$, with $M = |\mathcal{V}|$ nonzero sector projections P_v. Unitaries $U_1, \dots, U_s$ permute the sectors exactly and*

satisfy $\|[U_i, \xi]\|_2 \le \delta$. *For forbidden sets* $B_j \subseteq \mathcal{V}$, $1 \le j \le r$, *suppose*

$$D_j^* D_j \ge \kappa P_{B_j}, \qquad |D_j|_\rho \le \delta, \qquad P_{B_j} = \sum_{v \in B_j} P_v,$$

where $\kappa > 0$. *Let* $F = \sum_v c_v P_v$, $0 \le c_v \le 1$, *and* $a = M\delta$. *If* $2(s + r/\kappa)a^2 < 1$, *a transition orbit avoiding all forbidden sets has average cost at most*

$$\mathrm{Tr}(\rho F) + \sqrt{2s}\, a + 2(s + r/\kappa)a^2. \tag{A.1}$$

Proof. Let $p_v = \mathrm{Tr}(\rho P_v)$ and $q_v = \sqrt{p_v}$. If π_i is the sector permutation of U_i, the reverse triangle inequality applied to the row blocks $P_v\xi$ and $P_v U_i \xi U_i^*$ gives

$$\sum_v (q_v - q_{\pi_i(v)})^2 \le \delta^2.$$

The union B of forbidden sets has $p(B) \le (r/\kappa)\delta^2$. For an orbit O, let $E_O = \sum_i \sum_{v \in O} (q_v - q_{\pi_i(v)})^2$, so $\sum_O E_O \le s\delta^2$. A simple path in the connected orbit graph has at most $|O| - 1$ edges; Cauchy–Schwarz gives

$$(q_v - q_w)^2 \le (|O| - 1)E_O.$$

Choose a bad vertex w_O in each invalid orbit. Then

$$\sum_{v \in O} q_v^2 \le 2M q_{w_O}^2 + 2M^2 E_O.$$

Summing shows that the mass b of invalid orbits is at most $2(s + r/\kappa)a^2$.

Average p uniformly within each orbit, obtaining $\bar{p}$, and put $u_v = \sqrt{\bar{p}_v}$. The variance identity and the path bound give

$$\sum_{v \in O} (q_v - \bar{q}_O)^2 \le M^2 E_O,$$

where $\bar{q}_O$ is the ordinary average of q on O. If the orbit has mass $p_O > 0$, set $t_O = (\sum_{v \in O} q_v)/\sqrt{|O| p_O}$. Then

$$\sum_{v \in O} (q_v - u_v)^2 = 2p_O(1 - t_O) \le 2p_O(1 - t_O^2) = 2 \sum_{v \in O} (q_v - \bar{q}_O)^2.$$

Consequently $\|q - u\|_2 \le \sqrt{2s}\, M\delta$. Since $\mathrm{TV}(p, \bar{p}) \le \frac{1}{2}\|q - u\|_2 \|q + u\|_2 \le \|q - u\|_2$, orbit averaging changes expected cost by at most this amount. Conditioning $\bar{p}$ on valid orbits changes total variation by $b < 1$. The result is a mixture of uniform valid-orbit laws, one of which has cost at most (A.1). □

Corollary A.3. *Assume valid finite orbits are effectively enumerable, have computable rational costs, and every such cost is at least a left-c.e. real* θ. *If near-minimizing models have flag cost tending to* θ *and* $M\delta \to 0$, *then* θ *is computable. If* θ *is uncomputable and flat models have flag cost tending to* θ, *their dimensions eventually satisfy* $d \ge \kappa\delta^{-2}$.

Proof. Theorem A.2 makes the infimum of valid finite-orbit costs equal to θ under the first hypothesis. Enumerating those costs gives computable upper bounds tending to θ; left semicomputability gives lower bounds. This computes θ.

More generally, for uncomputable θ the infimum θ_{fin} of valid finite-orbit costs must be strictly greater than θ, or there are no valid finite orbits: equality would already make θ computable by the same enumeration. If a flat model had no nonzero forbidden projection,

every orbit would be valid. Exact sector permutation makes ranks constant on an orbit, so its flat cost would be a mixture of uniform orbit costs, at least θ_{fin}. Thus near-minimizing flat models eventually have a nonzero P_{B_j}. The rejection gap gives

$$\delta^2 \geq d^{-1} \operatorname{Tr}(D_j^* D_j) \geq \kappa/d.$$

This proves the claimed dimension bound. □

The semantic assumption about all valid finite-orbit costs is essential. A program orbit may fail to extend to the full clock system, or may have a cheap flag cost that the extra clock relations exclude. The corollary does not apply automatically to every operator presentation. Its flat bound explains the amplitude issue: the target dimensions here are below order δ^{-2}, so a construction satisfying these hypotheses could not use fixed-gap bad sectors. Our increment defect instead has amplitude 4^{-N} at overflow, as Proposition 4.5 proves. The initialization proof uses exact tracial cancellation and compactness, not a fixed-gap finite-sector extraction.

B Finite computations and reproducibility

The code in `anc/audit/check_models.py` represents an operator as a finite sum of weighted basis permutations. For each defining relation it enumerates all basis columns of the tensor factors that occur in that relation. The absent factors are identities, so averaging squared column norms computes exactly the full-space normalized Hilbert–Schmidt expression, before floating-point rounding. Equal output rows are combined before squaring, avoiding cancellation of large squared terms. The periodic coordinates in (7.6) are summed with integer numerators and denominator $4^{\ell} - 1$ for the permutation period ℓ, then rounded once to floating point.

The test program has one counter and five controls. From the starting control, zero halts, while nonzero decrements and moves to an odd-parity control. There zero enters an absorbing nonhalting control, while nonzero decrements and returns to start. Idle and halt are absorbing. Hence input e halts exactly when it is even, at step $e + 1$. This is a test of the finite-model mechanism, not a computation of the universal halting set.

Table 1: Finite-model checks for the disclosed test program. Each ratio is a flat Hilbert–Schmidt defect divided by 4^{-N}. The first two rows apply separately to all three compulsory clocks.

Quantity	$N = 4$	$N = 6$	$N = 8$
Clock affine ratio	$3/\sqrt{2}$	$3/\sqrt{2}$	$3/\sqrt{2}$
Clock tail-flip ratio	3	3	3
Guarded increment ratio	0.353553	0.176777	0.088388
Carry-step ratio	0.25	0.125	0.0625
Conditional tail ratio	0.00585946	0.000160228	0.00256366
K_N	$\{0\}$	$\{0, 2\}$	$\{0, 2\}$
k_N	2	4	4
Dimension	16 384	4 194 304	268 435 456
Recorded checks	390	394	398

The remaining defining defects are zero to within 10^{-14} in these tests. Complete per-relation values and their reference scales are archived, not just the maxima in Table 1. The two guarded increments have identical ratios in this example. The theorem's operator-norm bounds are proved in Section 7; a script's maximum column norm is not generally an operator norm. The displayed dimensions are represented without dense matrices. The original $N = 8$ run had an observed peak working set about 1.16 GB.

The gadget script `anc/gadgets/gadget_check.py` tests self-adjointness, squares, all pairwise commutations within each context, context products, and the auxiliary conjugacies on random finite instances. It is a numerical check of the explicit extension, not a proof of the converse; that converse is proved in Section 3.

The ancillary README gives commands, dependencies, recorded outputs and limitations. All computations are run with one numerical-library thread. The finite-model script caps its exhaustive workspace implementation at $N = 8$ to stay within the stated memory budget. The independent engine described in the methods section is in the repository directory `verify/`, with its recorded output, and the two scripts of the independent proof check are in `anc/crosscheck/`.

Reader and extraction sections. The spectrum program in `anc/spectrum/` builds 24 finite reader models at $N = 2, 4, 6$ and eight rational thresholds, including both endpoints. It tests reset boundaries, reads beyond N, halt inclusion, slot collisions, dressing and defining operator residuals on the nontrivial kernels. Its other tests cover radical-inverse windows, compression with domain repairs, finite word moments and noncommuting-state Jensen. The protocol programs in `anc/extraction/` evaluate the stage inequality, full exact and perturbed protocols with mixed inputs and noncommuting controls, small perturbations with positive lower bounds, bounded adversarial searches, and one-sided inclusion repair. Their recorded outputs are preserved in the respective `recorded/` directories. The ancillary README states each file's scope and distinguishes the original allowances from the constants of Lemma 12.1.

Complementary-construction checks. The script `anc/gate/toy_gate.py` retains its original recorded results in `gate/recorded/` and writes a new result beside the script. It executes the signed packed polynomial instructions for three-colouring of K_3, a pair of noncommuting binary bases, and a flagged bad three-colouring of K_4. It checks contraction-register decompositions, initialized finite histories through every distinct reset period, cyclic boundaries, simultaneous data factors, decoder and reset recursions, complementary window counts and address dressing. This script uses a single projection family for both players; its history, reset and window tests are independent of that choice and cover 16,893 history transitions, 35 cyclic reset boundaries and 2,396,744 mask/start cases across 254 masks. The separate two-family suite in `anc/gate/crosscheck/` tests the presentation of Section 15.1, with separate families for the two players: K_3 colouring has 18 generators, and the magic-square witness has 42 generators and packed dimension 168. It also checks the reduced-state eigenspaces, an anti-synchronous PCC control, 53,248 history points, sparse address models and windows through depth 22. Recorded logs and rerun commands are specified in the ancillary README. These are tests of the new matrix algebra and finite histories on known-answer toy games, not the MIP* games themselves. They do not certify the halting reduction, arbitrary-state lower bound or compiled game.

Stop-or-continue classification. The four scripts in `anc/classification/` test the new stop-or-continue algebra on finite known-answer instances. They cover oscillating rational approximations, finite stopping parameters, tail rejection, per-gate LOAD boundaries, matrix-valued registers, noncommuting witness tuples, reset periods, flagged default data and finite threshold/window budgets. The operator-level check evaluates the new equations as sparse operators built from a toy history, with negative controls for an omitted LOAD factor and imposed time invariance. The matrix-interface check displays the shifted-corner and stop-sector identities on a hand-built instance, in which the stop-sector identity holds by construction, with a negative control for dropped coordinate commutation. The toy scripts add negative controls that change data between the checks of a false gate (`toy_classification.py`) or one step

before a LOAD (`toy_attack.py`). The two register checks initialize $M = I$ and $\Sigma = 0$ on the outgoing LOAD edge. The operator-level script also checks a synthetic LOAD with incoming $M = 0$, $\Sigma = I$, testing both the instruction history and the polynomial selectors against that transition; holding the incoming registers gives nonzero negative controls. Their recorded outputs are archived and hashed. They do not instantiate the MIP* reduction, synthesize a universal counter compiler, certify a noncomputable limit, or test the full compiled game. The lower bound and the nonuniform slow-diagonal upper bound are written proofs in Section 16, not conclusions of these finite tests.

Power-scale sections. The seven programs in `anc/power/` are run, with fixed seeds, one process at a time and one numerical-library thread. The two moment scripts enumerate finite-field code moments, check full-Fock CAR and small affine actions, and rebuild the collective clock in natural and sparse coordinate representations, including $M = 48$. They verify finite examples of the 27-row certificate's geometric ingredients; they do not allocate the full large affine representation. The stopped and rotation scripts test finite cyclic histories, zero and dyadic inputs, reset interruptions and the inherited unguarded increment defect. The independent stack tests use small noncommuting matrices and diagonal histories to stress ordered updates, column completeness, polynomial carry/clock charges and scalar product envelopes. The scaled-machine script checks all new counter, control, witness and calibrator equations on finite clock blocks, literal noncommutative polynomial identities, weighted bad-zero-test perturbations and a finite scheduler arithmetic simulation. Its scheduler test is not a compiled universal register program; compilation and elementary-step bounds are proved in Lemma P4.3. The independently written counter tests use exact rational scalar histories, small weighted noncommuting finite-window perturbations and arbitrary ordered measurement columns; they do not include a complete vanishing-defect clock/compiler strategy. Their scope and every output are retained under the corresponding `recorded/` directories, with SHA-256 hashes. These checks support, and do not replace, Theorem P2.1 and Propositions P3.1, P4.2 and P5.1.

Precision-profile checks. The profile checks in `anc/profiles/` and their individual scopes are described here. The runner `anc/profiles/run_checks.py` executes three suites. The carry suite uses exact rational diagonal/permutation calculations for the carry and unguarded overflow, reset and initialized/no-reset cylinders, Cantor Lipschitz constants and the checker-cost comparison. The hold suite simulates twelve million pulses, checks denominator padding and holds, exact calibrator histories for thirteen inputs, logarithmic profile accounting and power moment bands. Its scheduler is an abstract instruction-cost simulation, not a complete emitted universal machine table. The independent suite imports neither supplied engine; it tests overflow-state deletion through twelve bits, noncommuting shared inputs, exact greedy tracking, affine endpoint parameters, separate hold accounting and the diagonal bit decoder. Recorded JSON/logs are archived with SHA-256 hashes. Finite outcomes support these interfaces; no numerical run certifies a noncomputable endpoint, a uniform operator theorem or exclusion by every game.

Ceiling, pairwise and group checks. Four suites in `anc/frontiers/` test the additional finite interfaces. The supplied ceiling program checks finite clock defects, exact scalar calibrator domains and residuals, sharp-versus-smooth penalties and stopped multiplier saturation. The supplied hash program exhaustively checks affine collisions and the pairwise dimension bound. Independently specified tests cover weighted sequential POVM positivity, telescoping and effect-approximation ratios, noncommuting complete columns, actual-failure neighborhoods, mixed classical hash flags and exact rational stopped histories. The group suite checks

sixteen finite counter programs on joint fuel/time grids, the single-visit trace kernel including zero and powers of two, the seventeen-state squaring loop through four marker visits, and finite-Heisenberg matrix trace normalization. Its elementary marker executes 595,050 instructions after initialization. These finite checks do not prove the arithmetic-hierarchy theorem, asymptotic irrationality, arbitrary-state trace rigidity, or an algebraic group-transfer theorem. The written arguments give those conclusions. JSON and console outputs are archived and hashed; the runner uses one numerical-library thread.

C The finite reader construction

This appendix supplies the reader used in Section 11. Its lemmas and equations use the prefix 11, consistently with their applications in that section.

C.1 A finite reader

Fix a computable increasing sequence of rationals $0 \le q_s \le 1$ tending to $\theta \in [0,1]$. For a binary address $x = (b_0, b_1, \dots)$ put

$$v(x) = \sum_{j \ge 0} b_j 2^{-j-1}, \qquad p_s = \sum_{j<s} b_j 2^{s-1-j}.$$

At stage $s \ge 1$, a fixed program computes q_s, reads the next bit, updates p_s, and halts if

$$(p_s + 1)/2^s < q_s. \tag{11.1}$$

Counter instructions implement the integer arithmetic and strict rational comparison. Each stage takes at least one instruction. Acceptance implies $v(x) \le (p_s+1)/2^s < \theta$; conversely, if $v(x) < \theta$, the two sequences $(p_s+1)/2^s \to v(x)$ and $q_s \to \theta$ eventually satisfy (11.1). Thus the program halts exactly on $\{v < \theta\}$. Only the finite program for (q_s) occurs in the presentation; θ is not an operator coefficient.

Here is the full change to Section 5. Delete L, R, equations (5.3)–(5.4), and their occurrences in initialization and commutation lists. Retain all other clock, reset, carry and conditional relations. Add positive contractions $A_{\mathrm{r}}, C_{\mathrm{r}}$ and a projection P_{r} to the instantaneous commuting workspace, with

$$4A_{\mathrm{r}} = 3P_{\mathrm{r}} + C_{\mathrm{r}}, \qquad [P_{\mathrm{r}}, C_{\mathrm{r}}] = 0. \tag{11.2}$$

The workspace commutes with Z, A_a, A_f, A_t, F as before. The flag still commutes with A_a and the complete fuel and time sets, including U; all the conditional-clock commutations are retained. Initialize every ordinary counter to I, the start control to I, the other controls to zero, and A_{r} to A_a. A READ control q_σ has selectors $q_\sigma P_{\mathrm{r}}$ and $q_\sigma(I - P_{\mathrm{r}})$. If Q_{r} is the sum of READ controls, set

$$\mathrm{Step}(A_{\mathrm{r}}) = Q_{\mathrm{r}} C_{\mathrm{r}} + (I - Q_{\mathrm{r}}) A_{\mathrm{r}}. \tag{11.3}$$

Include A_{r}, together with every ordinary counter and control, in (5.9)–(5.10). In particular,

$$Z A_{\mathrm{r}} = Z A_a, \qquad U^* A_{\mathrm{r}} U = Z_+ A_a + (I - Z_+)\,\mathrm{Step}(A_{\mathrm{r}}).$$

The bit and tail are witnesses, like the zero-test witnesses; they need no separate update equations. Call this finite rational presentation $\mathcal{P}_\theta$.

Lemma 11.1 (Reader forcing)**.** *Every exact tracial model of $\mathcal{P}_\theta$ has $\tau(F) \ge \theta$. Uniformly over its weighted precision-δ models,*

$$\mathrm{Tr}(\rho F) \ge \theta - \omega_\theta(\delta), \qquad \omega_\theta(\delta) \longrightarrow 0. \tag{11.4}$$

Proof. On the kernel and range of P_{r}, equation (11.2) places the spectrum of A_{r} in $[0, 1/4]$ and $[3/4, 1]$, respectively. Therefore

$$P_{\mathrm{r}} = p(A_{\mathrm{r}}), \qquad C_{\mathrm{r}} = s(A_{\mathrm{r}}).$$

This conclusion does not require Cantor spectrum in advance.

Let B_w specify an s-bit address word w satisfying (11.1). Every extension follows the same computation to a halt, at some time h. It may halt at an earlier stage; in either case it halts by stage s after at most s reads. Choose $2^k > h$ and put $Q_0 = B_w C_k D_k (I - F)$. These factors commute. The reset part of Lemma 6.1 is independent of the deleted decoder and initializes this subspace.

The reader requires an operator induction in addition to the scalar counter induction of Theorem 6.2. After j reads, track

$$(A_{\mathrm{r}})_\ell Q_0 = s^j(A_a) Q_0. \tag{11.5}$$

The right side commutes with Q_0. Taking adjoints shows that Q_0 reduces the self-adjoint left side. Functional calculus on that reducing subspace gives

$$(P_{\mathrm{r}})_\ell Q_0 = b_j Q_0, \qquad (C_{\mathrm{r}})_\ell Q_0 = s^{j+1}(A_a) Q_0$$

whenever the next bit is read. All bits read before the halt belong to w. On other instructions the reader remains fixed. The ordinary counters and controls have their prescribed scalar right actions, with zero tests forced by Lemma 6.1. There is no reset before h. Thus (5.10) proves the initialized history through its halt. As in Theorem 6.2, $q_{\mathrm{halt}} \le F$ and $[F, U] = 0$ force $Q_0 = 0$. Conjugating by successive fuel and time flips and adding removes C_k, D_k, giving $B_w \le F$.

The length-s address cylinders have trace 2^{-s}. The number accepted at stage s is

$$a_s = \max\{0, \lceil 2^s q_s \rceil - 1\}, \qquad 2^{-s} a_s \ge q_s - 2^{-s}.$$

Summing accepted cylinders gives $\tau(F) \ge q_s - 2^{-s}$, hence $\tau(F) \ge \theta$. This includes $\theta = 0$ and $\theta = 1$: the possibly unaccepted all-ones address has trace zero. Finally the bounded moment-limit and tracial GNS argument in Theorem 6.2 applies to all the added bounded variables and proves (11.4). □

Lemma 11.2 (Finite reader models)**.** *For even $N \ge 2$, $\mathcal{P}_\theta$ has flat models of dimension*

$$d_N = 2^{3N + k_N}, \qquad k_N = \theta N + o(N), \tag{11.6}$$

with every defining residual at most $3 \cdot 4^{-N}$ in operator norm. Every new program and reader equation is exact.

Proof. Use the three N-bit registers of Theorem 7.1, with bit zero least significant for increment. Write

$$v_N(x) = \sum_{j<N} b_j(x) 2^{-j-1}, \quad \vartheta_N = q_{2^N}, \quad f_N(x) = 1_{\{v_N(x) < \vartheta_N\}}.$$

Read the N bits followed by zeros. The Cantor reader and its witnesses remain exact after any number of reads. Assign all workspace entries the history (7.9), with this initialization. The next history point is reset precisely when its time residue is zero, including the cyclic boundary and all-ones fuel. Thus all transitions are exact. A displayed halt has accepting stage $s \le 2^N$, so $v_N(x) < q_s \le q_{2^N}$; hence $q_{\mathrm{halt}} \le F$ is exact. Specifying q_{2^N} may take arbitrarily long classical computation. It does not lengthen a displayed history or add a tensor factor.

The moving-window estimate holds for this new mask. On an aligned block of 2^r consecutive address integers, the radical-inverse values v_N are $(j+u)/2^r$, $0 \le j < 2^r$, for a fixed $u \in [0,1)$. Their count below ϑ_N differs from $2^r\vartheta_N$ by at most one. Decompose an ordinary interval of length L into at most two aligned dyadic blocks of each size, and a cyclic interval into at most two ordinary intervals. Uniformly in the starting point,

$$\sum_{j=1}^{N} f_N(x-j) = N\vartheta_N + O(\log(N+2)). \tag{11.7}$$

This is the elementary radical-inverse discrepancy argument associated with van der Corput [70]. It is proved here for the moving windows used by the construction.

Use this f_N and $k_N = 1 + \max_x \sum_{j=1}^{N} f_N(x-j)$ in (7.5)–(7.8). The slot noncollision, geometric series and dressing arguments in Theorem 7.1 use only a Boolean mask and this slot bound. Their residual estimates therefore still hold. Since $\vartheta_N \to \theta$, equation (11.7) gives (11.6). The full address mean differs from ϑ_N by at most 2^{-N}, so these flat models also have $\mathrm{Tr}(\rho F) \to \theta$. □

D Complementary flag estimates and resource laws

The presentation and testing program are defined in Section 15. Here we prove their approximable flag minimum, finite models, resource laws and trace gap. Results and equations use the prefix 15, consistently with their applications in that section.

D.1 The approximable flag minimum

Proposition 15.4. *Every $\tau \in \mathcal{T}_{\mathrm{mat}}(\widehat{\mathcal{P}}_A)$ satisfies $\tau(F) \ge \lambda_A$.*

Proof. Let $Q_e = B_e(I-F)$. It commutes with the data and, by Lemma 15.3, its corner satisfies $\mathcal{R}_e$. Tracing the first packed equation gives $\tau(Q_eE) = \tau(Q_e)/m_e$. If $\tau(Q_e) > 0$, the Q_eE corner carries $\mathcal{A}_{A,e}$, with generators given by fixed bounded polynomials $Q_eEW^{-j}DW^jE$.

This positive-trace corner inherits flat approximation. In the flat approximants, self-adjointize the polynomial for Q_eE and round at $1/2$. Its projection defect tends to zero, and for every real t,

$$|1_{[1/2,\infty)}(t) - t| \le 2|t^2 - t|.$$

The correction tends to zero in normalized Hilbert–Schmidt norm, and the rounded projection's rank fraction tends to $\tau(Q_eE) > 0$. Compress and self-adjointize the finitely many corner generators, and round them by the same estimate. Fixed polynomial relations survive by bounded-product telescoping. Normalizing the corner trace multiplies norms by a bounded factor tending to $\tau(Q_eE)^{-1/2}$. These are flat approximate models of $\mathcal{A}_{A,e}$. Lemma 15.1 therefore gives $\tau(B_e(I-F)) = 0$ for $e \notin A$.

The decoder exhausts L in trace. Iterating (5.3) yields

$$L = \sum_{e<s} B_e + R_s, \qquad 0 \le R_s \le \prod_{j<s} T_a^{2j} E_{00} T_a^{-2j}, \qquad \tau(R_s) \le 4^{-s}.$$

Indeed $L - \sum_{e<s} B_e$ is a projection whose range lies in the all-00 prefix; no commutation between rotated L and all earlier digits is assumed. Since $I - F \le L$, the disjoint cylinders of trace $4^{-(e+1)}$ give

$$\tau(I-F) = \sum_e \tau(B_e(I-F)) \le \sum_{e\in A} 4^{-(e+1)} = \theta_A.$$

Only the passage excluding nonmember corners uses matrix approximability. The testing identity holds for every exact trace. □

D.2 Slow admission and simultaneous finite models

Lemma 15.5 (Elementary witness scheduling)**.** *Suppose finite witnesses with positive integer costs a_j arrive in a computable stream without repeated indices. At even depth N admit its longest prefix arrived by stage N, with every index $e_j < N/2$ and total cost at most $\lfloor\sqrt{N}\rfloor$. These prefixes are nested and eventually contain every fixed witness. Their total cost is $o(N)$.*

Proof. For a fixed prefix its arrival time, largest index and total cost are finite constants. All three cutoffs eventually exceed them, and none decreases with N. The cost bound is $\sqrt{N} = o(N)$. An enormous early witness may cause an enormous finite delay; no convergence rate is claimed. More generally any nondecreasing computable unbounded budget $o(N)$ has the same property. □

Apply this lemma to the certificates $(e_j, d_j, \mathrm{model}_j)$ of Lemma 15.1, with $a_j = \lceil \log_2 \max(2, m_{e_j} d_j) \rceil$. Use a dovetailing with finitely many outputs per stage. Let S_N be the admitted indices. Then $S_N \uparrow A$, and the common data factor

$$\mathcal{H}_{\mathrm{data},N} = \bigotimes_{e_j \in S_N} (\mathbb{C}^{m_{e_j}} \otimes \mathbb{C}^{d_j}), \qquad \ell_N = \dim \mathcal{H}_{\mathrm{data},N}$$

has $\log \ell_N \le \sqrt{N}$; the empty product has dimension one.

Lemma 15.6 (Complementary finite models)**.** *For even $N \ge 2$, $\widehat{\mathcal{P}}_A$ has flat models with all defining residuals at most $3 \cdot 4^{-N}$ in operator norm and*

$$d_N = 2^{3N+k_N} \ell_N, \qquad k_N = \lambda_N N + O(\log(N+2)), \qquad \lambda_N = 1 - \sum_{e \in S_N} 4^{-(e+1)} \downarrow \lambda_A. \tag{15.7}$$

Every new defining equation is exact, and $\mathrm{tr}_{d_N}(F) = \lambda_N$.

Proof. On the N-bit address cycle set

$$f_N(x) = 1 - \sum_{e \in S_N} 1_{\{x \equiv 4^e \pmod{4^{e+1}}\}}, \qquad k_N = 1 + \max_x \sum_{j=1}^{N} f_N(x-j). \tag{15.8}$$

The classes are disjoint, and their periods divide 2^N. Bit zero is least significant for increment, so these are precisely the cylinders B_e. The mask is Boolean, has mean λ_N, and satisfies (15.2). For $s = \lfloor \log_4 N \rfloor$, each selected class with $e \le s$ has interval-count error at most one. The disjoint union of the remaining classes lies in the multiples of 4^{s+1} and has at most one point in an interval of length $N < 4^{s+1}$; its expected count is at most $N/(3 \cdot 4^{s+1})$. Hence, uniformly in the subset and starting point,

$$\left| \sum_{j=1}^{N} f_N(x-j) - N\lambda_N \right| \le s + 2 + \frac{N}{3 \cdot 4^{s+1}}. \tag{15.9}$$

Cyclic intervals lift to integer intervals because all periods divide the cycle.

Use the three compulsory N-qubit registers, the k_N optional qubits, and the common data factor. Theorem 7.1's conditional permutation, series and dressing require only a Boolean mask and that k_N exceed every N-window count. They therefore apply to (15.8). On address $x \in B_{e_j}$ for an admitted index, assign the packed witness to factor j and tensor identities on the others. On other addresses take $(D, E, W) = (0, I, I)$. These address-block data are independent of fuel, time and optional slots. They commute with A_a, L, R, F, Z and all scalar workspace, and with the complete fuel/time sets. Address dressing acts on address and

optional slots and leaves these commutations intact. The data need not commute with the dressed address shift or flip.

For fuel with k initial ones, use the actual initialized scalar and operator history at time $t \bmod 2^k$; all-ones fuel uses $k = N$. On each (x, q, t) block the scalar workspace is a scalar multiple of the data identity. Every M, Σ is a contraction by the signed-term bound. Choose their two representing unitaries within each block and take direct sums. All required representing-unitary commutators hold; naming these matrices adds no tensor factor.

When the next time residue is zero, the transition prescribes initialization; otherwise it prescribes the next instruction. Since $2^k \mid 2^N$, the cyclic transition from $2^N - 1$ to zero is also a reset. Thus (15.4) holds exactly at every ordinary and reset boundary. On an unflagged block every completed CHECK tests the packed witness belonging to that address, so it is zero. An unfinished CHECK imposes no test there; on flagged blocks $I - F = 0$. Therefore (15.5) is exact regardless of program speed.

The address wraparound is handled as in Theorem 7.1. In the gauge of the dressing (7.8), V is ordinary increment except at the all-ones address, where it also applies the fixed optional permutation $R_0^{\sum_x f_N(x)}$. There its output address differs from that of the recursive increment, so the guarded defect is $\sqrt{2}\,4^{-N}$ in operator norm, as before.

All new equations are exact simultaneously. The remaining clock, guarded-increment, carry and conditional-tail residuals have their old bounds, at most $3 \cdot 4^{-N}$. Counting the registers gives (15.7), and

$$\log d_N = (3 + \lambda_N)N + O(\log(N + 2)) + O(\sqrt{N}) = (3 + \lambda_A)N + o(N).$$

A fixed increase of even depth absorbs the constant three. Bounded gaps between even depths suffice for every small precision. Computing a certificate or its admission stage is classical model-description work; it consumes neither simulated fuel nor an extra quantum register. □

D.3 Dimension, usable entanglement and heat

Theorem 15.7. *For every c.e. set A, integer $r \geq 0$ and integer $q \geq 1$, there is a fixed finite rational BCS game $\widehat{G}_{A,r,q}$ of quantum value one and with no finite-dimensional perfect strategy. With*

$$\beta_{A,r,q} = \frac{4 + r - \theta_A}{4q}, \qquad L_\varepsilon = \log_2(1/\varepsilon),$$

its resource laws are

$$\log_2 D_{\widehat{G}_{A,r,q}}(\varepsilon) = \beta_{A,r,q} L_\varepsilon + o(L_\varepsilon), \tag{15.10}$$

$$d^{\rightarrow}_{\widehat{G}_{A,r,q}}(\varepsilon) = \beta_{A,r,q} L_\varepsilon + o(L_\varepsilon), \tag{15.11}$$

$$w_{\mathrm{local},\widehat{G}_{A,r,q}}(\varepsilon) = 2\beta_{A,r,q} L_\varepsilon + o(L_\varepsilon), \qquad 0 \leq w_{\mathrm{joint},\widehat{G}_{A,r,q}}(\varepsilon) \leq 2\log_2 K_{\mathrm{ans}}. \tag{15.12}$$

The heat statement uses exactly the model and units of Section 13.

Proof. The actual-state flag bound. Moment compactness of Lemma 15.6's flat models supplies a matrix-approximable trace of flag expectation λ_A. Proposition 15.4 gives the reverse inequality, so

$$\min_{\mathcal{T}_{\mathrm{mat}}(\widehat{\mathcal{P}}_A)} \tau(F) = \lambda_A. \tag{15.13}$$

Corollary 14.3 now proves, uniformly on every original weighted precision-δ model,

$$\mathrm{Tr}(\rho F) \geq \lambda_A - \omega_A(\delta), \qquad \omega_A(\delta) \longrightarrow 0. \tag{15.14}$$

Indeed a fixed positive shortfall along $\delta \to 0$ would have a weighted moment limit in $\mathcal{T}_{\mathrm{mat}}$ below (15.13). The bound holds on the original state ρ; compression only identifies the limit class.

Dimension. Every relation used by Lemma 9.3 remains in the presentation. For $t = \log_4(1/\delta)$ and $M = \lfloor t - 30\log_4(2+t) \rfloor$, its joint errors are at most $C(M+1)^6 4^M \delta$. Theorem 8.1 and state transport give

$$\log d \geq 3M + \sum_{i<M} \mathrm{Tr}(\rho F_i) - o(M) \geq (3 + \lambda_A - \omega_A(\delta))M - O(M^2\delta) - o(M),$$

since $|\mathrm{Tr}(\rho F_i) - \mathrm{Tr}(\rho F)| \leq 2i\delta$. This is $(3+\lambda_A)\log_4(1/\delta) - o(\log(1/\delta))$. The compiler sends game loss ε to source precision $C\sqrt{\varepsilon}$, with constant dimension overhead. Lemma 15.6 gives the converse upper loss $C'16^{-N}$. Choose an even depth $N \geq \log_{16}(C'/\varepsilon)$ within a fixed additive constant to match the bound through all small losses. This proves (15.10) for $r = 0, q = 1$.

Append r compulsory clocks with the cross-commutations of Proposition 11.3. Their upper cost is exactly rN, and their digit families add rM to the lower count. Power *every* defining residual, including the new tests and cross-commutators, by (11.10); keep domains and state-commutator requirements. Equation (11.11) turns powered precision η into unpowered precision $\eta^{1/q}$, whereas the explicit operator-norm residuals become at most $3^q 4^{-qN}$. Thus the game lower precision is $\delta = (C\sqrt{\varepsilon})^{1/q}$, upper loss is $C'16^{-qN}$, and the coefficient is $(3 + r + \lambda_A)/(4q)$. The upper losses tend to zero and this positive lower bound diverges, excluding a finite-dimensional perfect strategy.

Usable entanglement on the original mixed state. Proposition 12.2 applies. On a purification of an arbitrary mixed loss-ε strategy, keep the original environment, relabel illegal answers, and dilate locally with pure ancillas. The finite Boolean certificates reconstruct this presentation on the actual Bob marginal with powered precision $C\sqrt{\varepsilon}$; their constants depend on the fixed transducer but not on M or ε. Jensen then gives the same unpowered δ. New data/register variables enlarge a fixed certificate list only. None occurs in the raw translated clock measurements or their controls.

Consequently the same opposite-word construction on Alice, independent local inclusion repairs on both sides, and reverse count order supply the hypotheses of Lemma 12.1, with

$$m = (4 + r)M, \qquad \eta = C_{A,r,q}(M+1)^6 4^M \delta.$$

As in Proposition 12.2, only Q_i is repaired, so the activation flags are the raw F_i on the actual marginal.

At the chosen depth $\eta = O(t^{-24})$, and the vector, activation and entropy-continuity errors of Lemma 12.1 tend to zero. Its one-way protocol and hashing therefore give on *every original strategy state*

$$\begin{aligned} E_D^{A\to B}(\rho_{AB}) &\geq (3+r)M + \sum_{i<M} \mathrm{Tr}(\rho_B F_i) - o(1) \\ &\geq (3 + r + \lambda_A - \omega_A(\delta))M - O(M^2\delta) - o(1). \end{aligned}$$

Here (15.14), rather than an all-exact-trace bound, supplies the activation estimate. Since $M = L_\varepsilon/(4q) - O(\log L_\varepsilon)$, the infimum gives the lower side of (15.11). The general inequality $d_{\overrightarrow{G}}(\varepsilon) \leq \log D_G(\varepsilon)$ and the proved dimension upper bound give the other side.

Heat. Substitute the two matching laws into (13.1). Its fixed answer term is $O(1)$, so both endpoints have coefficient $2\beta_{A,r,q}$. The joint bound is the same cited heat theorem. This proves (15.12). □

Corollary 15.8 (Beyond the left-c.e. reals)**.** *The game $\widehat{G}_K := \widehat{G}_{K,0,1}$ has dimension and usable-entanglement coefficient $(4 - \theta_K)/4$, and local-heat coefficient $(4 - \theta_K)/2$. These are*

noncomputable right-c.e. reals, not left-c.e., and have the Turing degree of K. Recall the positive dimension spectrum $\mathcal{B}_{\mathrm{BCS}}$ from Section 1. This escape establishes

$$\{\textit{positive left-c.e. reals}\} \subsetneq \mathcal{B}_{\mathrm{BCS}} \subseteq \{\textit{positive } \Delta^0_2 \textit{ reals}\}. \tag{15.15}$$

Both $\mathcal{A}_{\log}$ and $\mathcal{A}_{\dim}$ also strictly contain the positive left-c.e. reals. The new non-left-c.e. examples accumulate at zero by taking $q \to \infty$.

Proof. The weights $w_e = 4^{-(e+1)}$ have tail at most $w_e/3$ after position e. Given earlier membership bits, absence leaves residual at most $w_e/3$, whereas presence leaves at least w_e. Approximating θ_K within $w_e/6$ distinguishes them at threshold $2w_e/3$. Thus $K \leq_T \theta_K$. Conversely K computes the summable series, so $\theta_K \equiv_T K$. Its complement and the displayed positive rational affine images are noncomputable right-c.e.; if also left-c.e., the two rational approximations would compute them. Theorem 15.7 supplies these coefficients, giving an intermediate case of Theorem 1.1. The left-c.e. reader gives the first inclusion, and Lemma 2.4, with the limit argument of Corollary 16.7, the dimension outer class. No Δ^0_2 upper bound for the entire usable-entanglement class is inferred here. □

D.4 Why the trace gap is consistent

Proposition 15.9. *For $\widehat{\mathcal{P}}_K$ the exact and approximable minima differ strictly:*

$$\min_{\mathcal{T}} \tau(F) < \min_{\mathcal{T}_{\mathrm{mat}}} \tau(F) = 1 - \theta_K.$$

An exact minimizing trace has a positive-trace nonhalting unflagged corner carrying $\mathcal{A}_{K,e}$. Such a trace is not matrix approximable.

Proof. The bounded exact-trace hierarchy of Proposition 14.1 makes the exact minimum left-c.e.; compactness supplies a minimizer. It is at most the approximable minimum and cannot equal its non-left-c.e. value. At the minimizer, $\tau(I - F) > \theta_K$. Decoder exhaustion and the total capacity of halting cylinders force $\tau(B_e(I - F)) > 0$ for some $e \notin K$. Testing and packing give its unital tracial model. Left and right multiplication in the tracial GNS representation yield a perfect commuting-operator strategy for $H_{K,e}$, using its two projection families; this is the usual tracial representation construction [60, 41]. Soundness concerns finite-dimensional value, so there is no contradiction. □

This is the matrix-versus-commuting distinction behind the negative solution of Connes' embedding problem in [39], made quantitative in the resource coefficient. Slofstra's nonclosure and undecidability results [66] are essential context for distinguishing finite exact perfection from limiting perfection; the finite exact witnesses used above are the stronger specified input, not a consequence of nonclosure.

There is also an independent consistency argument. The index set $C_{\mathrm{tr}} = \{e : \mathcal{A}_{K,e} \text{ has a unital tracial model}\}$ is Π^0_1. Each bounded tracial moment level, including domain localizers and all visible relation insertions, is a rational semialgebraic feasibility problem decidable by real quantifier elimination. Infeasibility of some level enumerates the complement; feasibility at every level gives a tracial state by bounded-moment compactness and GNS. Since $K \subseteq C_{\mathrm{tr}}$, equality would make K co-c.e. Hence some nonhalting indices, in fact infinitely many, already admit perfect tracial strategies independently of the clock construction. This argument uses the tracial presentation, rather than assuming arbitrary perfect commuting strategies are tracial. Lin's $\mathrm{MIP}^{\mathrm{co}} = \mathrm{coRE}$ [45] supplies the commuting-complexity context; no equality of commuting and finite-dimensional values is used.

References

[1] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992. doi:10.1002/rsa.3240030308.

[2] Goulnara Arzhantseva and Liviu Păunescu. Almost commuting permutations are near commuting permutations. *Journal of Functional Analysis*, 269(3):745–757, 2015. arXiv:1410.2626, doi:10.1016/j.jfa.2015.02.013.

[3] Tim Austin. Rational group ring elements with kernels having irrational dimension. *Proceedings of the London Mathematical Society*, 107(6):1424–1448, 2013. doi:10.1112/plms/pdt029.

[4] A. Baker. Linear forms in the logarithms of algebraic numbers. *Mathematika*, 13(2):204–216, 1966. doi:10.1112/S0025579300003971.

[5] Saugata Basu. Algorithms in real algebraic geometry: A survey, 2014. Sections 2.2 and 2.5 discuss singly exponential algorithms and bit complexity. URL: https://arxiv.org/abs/1409.1534v1, arXiv:1409.1534v1.

[6] Johannes Bausch, Toby S. Cubitt, and James D. Watson. Uncomputability of phase diagrams. *Nature Communications*, 12:452, 2021. arXiv:1910.01631, doi:10.1038/s41467-020-20504-6.

[7] Oren Becker and Alexander Lubotzky. Group stability and Property (T). *Journal of Functional Analysis*, 278(1):108298, 2020. arXiv:1809.00632, doi:10.1016/j.jfa.2019.108298.

[8] Charles H. Bennett, Herbert J. Bernstein, Sandu Popescu, and Benjamin Schumacher. Concentrating partial entanglement by local operations. *Phys. Rev. A*, 53:2046–2052, 1996. arXiv:quant-ph/9511030, doi:10.1103/PhysRevA.53.2046.

[9] Charles H. Bennett, David P. DiVincenzo, John A. Smolin, and William K. Wootters. Mixed-state entanglement and quantum error correction. *Phys. Rev. A*, 54(5):3824–3851, 1996. arXiv:quant-ph/9604024, doi:10.1103/PhysRevA.54.3824.

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

[11] Noel Brady, Martin R Bridson, Max Forester, and Krishnan Shankar. Snowflake groups, Perron–Frobenius eigenvalues and isoperimetric spectra. *Geometry & Topology*, 13(1):141–187, 2009. doi:10.2140/gt.2009.13.141.

[12] Sabine Burgdorf and Igor Klep. The truncated tracial moment problem. *Journal of Operator Theory*, 68(1):141–163, 2012. arXiv:1001.3679.

[13] Marc Burger and Alain Valette. Idempotents in complex group rings: theorems of Zalesskii and Bass revisited. *Journal of Lie Theory*, 8:219–228, 1998. Author scan: https://people.math.ethz.ch/~burger/pub/1998_Idempotents.pdf. URL: https://ftp.gwdg.de/pub/misc/EMIS/journals/JLT/vol.8_no.2/3.html.

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

[15] Rui Chao, Ben W. Reichardt, Chris Sutherland, and Thomas Vidick. Test for a large amount of entanglement, using few measurements. *Quantum*, 2:92, 2018. arXiv:1610.00771, doi:10.22331/q-2018-09-03-92.

[16] Jean-René Chazottes, Jean-Marc Gambaudo, Michael Hochman, and Edgardo Ugalde. On the finite-dimensional marginals of shift-invariant measures. *Ergodic Theory and Dynamical Systems*, 32(5):1485–1500, 2012. arXiv:1011.2442, doi:10.1017/S0143385711000526.

[17] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *J. Math. Phys.*, 58(1):012202, 2017. arXiv:1606.02278, doi:10.1063/1.4973422.

[18] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *Automata, Languages, and Programming (ICALP 2014)*, volume 8572 of *Lecture Notes in Comput. Sci.*, pages 320–331. Springer, 2014. arXiv:1209.2729, doi:10.1007/978-3-662-43948-7_27.

[19] 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. arXiv:1904.02350, doi:10.22331/q-2020-06-18-282.

[20] Alain Connes. Classification of injective factors. Cases II_1, II_∞, III_λ, $\lambda \neq 1$. *Ann. of Math. (2)*, 104(1):73–115, 1976. doi:10.2307/1971057.

[21] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. Undecidability of the spectral gap. *Forum of Mathematics, Pi*, 10:e14, 2022. arXiv:1502.04573, doi:10.1017/fmp.2021.15.

[22] Mikael de la Salle. Spectral gap and stability for groups and non-local games. *Journal de lÉcole polytechnique — Mathématiques*, 12:1417–1444, 2025. Result numbers refer to arXiv version 2. arXiv:2204.07084v2, doi:10.5802/jep.314.

[23] Igor Devetak and Andreas Winter. Distillation of secret key and entanglement from quantum states. *Proceedings of the Royal Society A*, 461(2053):207–235, 2005. The hashing inequality is Theorem 10 of the arXiv version. arXiv:quant-ph/0306078, doi:10.1098/rspa.2004.1372.

[24] Seth Douglas and Nidhal Mghirbi. The heat price of precision, 2026. Version 1.0.0. Unrefereed preprint; theorem numbers refer to this version. doi:10.5281/zenodo.23701732.

[25] Seth Douglas and Nidhal Mghirbi. Houghton's group H_3 has superpolynomial sofic profile, 2026. Version 1.1.2, Zenodo. Result numbers refer to that version. doi:10.5281/zenodo.23070869.

[26] 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; theorem numbers refer to this version. doi:10.5281/zenodo.23697346.

[27] Rodney G. Downey, Carl G. Jockusch, Jr., and Paul E. Schupp. Asymptotic density and computably enumerable sets. *Journal of Mathematical Logic*, 13(2):1350005, 2013. Theorem 5.1 refers to the arXiv version. arXiv:1307.0040, doi:10.1142/S0219061313500050.

[28] Marco Fanizza, Larissa Kroell, Arthur Mehta, Connor Paddock, Denis Rochette, William Slofstra, and Yuming Zhao. The NPA hierarchy does not always attain the commuting operator value, 2025. Result numbers refer to arXiv version 4. arXiv:2510.04943.

[29] Francesco Fournier-Facio and Bharatram Rangarajan. Ulam stability of lamplighters and Thompson groups. *Mathematische Annalen*, 389(3):2469–2497, 2024. arXiv:2301.03970, doi:10.1007/s00208-023-02708-5.

[30] Honghao Fu, Carl A. Miller, and William Slofstra. The membership problem for constant-sized quantum correlations is undecidable. *Communications in Mathematical Physics*, 406(5):96, 2025. arXiv:2101.11087v5, doi:10.1007/s00220-024-05229-7.

[31] A. Gelfond. Sur le septième problème de Hilbert. *Bull. Acad. Sci. URSS, Cl. Sci. Math. Nat. (VII sér.), no. 4*, pages 623–634, 1934.

[32] W. T. Gowers and O. Hatami. Inverse and stability theorems for approximate representations of finite groups. *Sb. Math.*, 208(12):1784–1817, 2017. arXiv:1510.04085, doi:10.1070/SM8872.

[33] Łukasz Grabowski. On Turing dynamical systems and the Atiyah problem. *Invent. Math.*, 198(1):27–69, 2014. doi:10.1007/s00222-013-0497-5.

[34] Andrew Hawkins, Adam Skalski, Stuart White, and Joachim Zacharias. On spectral triples on crossed products arising from equicontinuous actions. *Mathematica Scandinavica*, 113(2):262–291, 2013. arXiv:1103.6199, doi:10.7146/math.scand.a-15572.

[35] Patrick Hayden, Debbie W. Leung, and Andreas Winter. Aspects of generic entanglement. *Communications in Mathematical Physics*, 265(1):95–117, 2006. Theorem numbering refers to arXiv version 2. arXiv:quant-ph/0407049, doi:10.1007/s00220-006-1535-6.

[36] Michael Hochman and Tom Meyerovitch. A characterization of the entropies of multidimensional shifts of finite type. *Ann. of Math. (2)*, 171(3):2011–2038, 2010. doi:10.4007/annals.2010.171.2011.

[37] Alexander S. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Probl. Inf. Transm.*, 9(3):177–183, 1973. Translated from Problemy Peredachi Informatsii 9(3), 3–11.

[38] Zhengfeng Ji, Debbie Leung, and Thomas Vidick. A three-player coherent state embezzlement game. *Quantum*, 4:349, 2020. arXiv:1802.04926, doi:10.22331/q-2020-10-26-349.

[39] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*=RE, 2026. arXiv version 4, accepted to Annals of Mathematics. Result and page numbers refer to this version. URL: https://arxiv.org/abs/2001.04383v4, arXiv:2001.04383.

[40] Carl G. Jockusch, Jr. and Paul E. Schupp. Generic computability, Turing degrees, and asymptotic density. *Journal of the London Mathematical Society*, 85(2):472–490, 2012. Theorem 2.21 refers to the arXiv version. arXiv:1010.5212, doi:10.1112/jlms/jdr051.

[41] Se-Jin Kim, Vern I. Paulsen, and Christopher Schafhauser. A synchronous game for binary constraint systems. *Journal of Mathematical Physics*, 59(3):032201, 2018. arXiv:1707.01016, doi:10.1063/1.4996867.

[42] Igor Klep, Victor Magron, and Jurij Volčič. Optimization over trace polynomials. *Annales Henri Poincaré*, 23:67–100, 2022. arXiv:2006.12510, doi:10.1007/s00023-021-01095-4.

[43] Artemy Kolchinsky and David H. Wolpert. Thermodynamic costs of Turing machines. *Physical Review Research*, 2:033312, 2020. arXiv:1912.04685, doi:10.1103/PhysRevResearch.2.033312.

[44] Sang Rae Lee. Geometry of Houghton's groups, 2012. arXiv:1212.0257.

[45] Junqiao Lin. MIP^{co}=coRE, 2025. Commuting-operator interactive proofs; arXiv version 1. URL: https://arxiv.org/abs/2510.07162, arXiv:2510.07162.

[46] Laura Mančinska and Thomas Vidick. Unbounded entanglement in nonlocal games. *Quantum Inf. Comput.*, 15(15–16):1317–1332, 2015. arXiv:1402.4145, doi:10.26421/QIC15.15-16-4.

[47] Matthew McKague, Michele Mosca, and Nicolas Gisin. Simulating quantum systems using real Hilbert spaces. *Physical Review Letters*, 102(2):020505, 2009. arXiv:0810.1923, doi:10.1103/PhysRevLett.102.020505.

[48] Arthur Mehta, William Slofstra, and Yuming Zhao. Positivity is undecidable in tensor products of free algebras. *J. Eur. Math. Soc. (JEMS)*, 2026. Published online. Result numbers refer to arXiv version 2. arXiv:2312.05617, doi:10.4171/JEMS/1786.

[49] 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. Theorem 5.1 refers to the arXiv version. arXiv:0902.0223, doi:10.1007/s00222-010-0296-1.

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

[51] Marvin L. Minsky. Recursive unsolvability of Post's problem of "tag" and other topics in theory of Turing machines. *The Annals of Mathematics*, 74(3):437–455, 1961. doi:10.2307/1970290.

[52] Marston Morse and Gustav A. Hedlund. Symbolic dynamics II. Sturmian trajectories. *American Journal of Mathematics*, 62(1):1–42, 1940. doi:10.2307/2371431.

[53] Hamoon Mousavi, Seyed Sajjad Nezhadi, and Henry Yuen. On the complexity of zero gap MIP*. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, volume 168 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 87:1–87:12, 2020. URL: https://arxiv.org/abs/2002.10490, arXiv:2002.10490, doi:10.4230/LIPIcs.ICALP.2020.87.

[54] 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 (STOC 2022)*, pages 1–11. ACM, 2022. Result numbers refer to arXiv version 2. arXiv:2110.04651, doi:10.1145/3519935.3519949.

[55] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993. doi:10.1137/0222053.

[56] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10:073013, 2008. arXiv:0803.4290, doi:10.1088/1367-2630/10/7/073013.

[57] Ashwin Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, pages 369–376. IEEE, 1999. arXiv:quant-ph/9904093, doi:10.1109/SFFCS.1999.814608.

[58] Narutaka Ozawa, Mikael Rørdam, and Yasuhiko Sato. Elementary amenable groups are quasidiagonal. *Geometric and Functional Analysis*, 25(1):307–316, 2015. arXiv:1404.3462, doi:10.1007/s00039-015-0315-x.

[59] Connor Paddock. Rounding near-optimal quantum strategies for nonlocal games to strategies using maximally entangled states, 2025. Result numbers refer to arXiv version 5. URL: https://arxiv.org/abs/2203.02525v5, arXiv:2203.02525v5.

[60] Vern I. Paulsen, Simone Severini, Daniel Stahlke, Ivan G. Todorov, and Andreas Winter. Estimating quantum chromatic numbers. *Journal of Functional Analysis*, 270(6):2188–2222, 2016. arXiv:1407.6918, doi:10.1016/j.jfa.2016.01.010.

[61] Mikaël Pichot, Thomas Schick, and Andrzej Żuk. Closed manifolds with transcendental L^2-betti numbers. *Journal of the London Mathematical Society*, 92(2):371–392, 2015. arXiv:1005.1147, doi:10.1112/jlms/jdv026.

[62] Mark V. Sapir, Jean-Camille Birget, and Eliyahu Rips. Isoperimetric and isodiametric functions of groups. *Annals of Mathematics*, 156(2):345–466, 2002. arXiv:math/9811105, doi:10.2307/3597195.

[63] J. C. Shepherdson and H. E. Sturgis. Computability of recursive functions. *Journal of the ACM*, 10(2):217–255, 1963. doi:10.1145/321160.321170.

[64] Joseph R. Shoenfield. On degrees of unsolvability. *Annals of Mathematics*, 69(3):644–653, 1959. doi:10.2307/1970028.

[65] William Slofstra. A group with at least subexponential hyperlinear profile, 2018. arXiv:1806.05267.

[66] William Slofstra. The set of quantum correlations is not closed. *Forum Math. Pi*, 7:e1, 2019. Result numbers refer to arXiv version 2. arXiv:1703.08618, doi:10.1017/fmp.2018.3.

[67] William Slofstra. Tsirelson's problem and an embedding theorem for groups arising from non-local games. *J. Amer. Math. Soc.*, 33(1):1–56, 2020. arXiv:1606.03140, doi:10.1090/jams/929.

[68] William Slofstra and Thomas Vidick. Entanglement in non-local games and the hyperlinear profile of groups. *Ann. Henri Poincaré*, 19(10):2979–3005, 2018. Result numbers refer to arXiv version 2. arXiv:1711.10676, doi:10.1007/s00023-018-0718-y.

[69] Ernst Specker. Nicht konstruktiv beweisbare Sätze der Analysis. *Journal of Symbolic Logic*, 14(3):145–158, 1949. doi:10.2307/2267043.

[70] Johannes G. van der Corput. Verteilungsfunktionen. I. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen*, 38:813–821, 1935. Binary radical-inverse construction; the window estimate used here is proved in the text.

[71] Vlatko Vedral and Martin B. Plenio. Entanglement measures and purification procedures. *Physical Review A*, 57:1619–1633, 1998. arXiv:quant-ph/9707035, doi:10.1103/PhysRevA.57.1619.

[72] Thomas Vidick. Almost synchronous quantum correlations. *Journal of Mathematical Physics*, 63(2):022201, 2022. arXiv:2103.02468, doi:10.1063/5.0056512.

[73] Xizhong Zheng and Klaus Weihrauch. The arithmetical hierarchy of real numbers. *Mathematical Logic Quarterly*, 47(1):51–65, 2001. doi:10.1002/1521-3870(200101)47:1<51::AID-MALQ51>3.0.CO;2-W.