A RICE THEOREM FOR STATIONARY STATES UNDER FIXED NOISE

SETH DOUGLAS AND NIDHAL MGHIRBI

Abstract. Which macroscopic questions about a noisy lattice can an algorithm decide from its local rule? We show that already on the line, for deterministic finite-range rules followed by independent uniform replacement at one fixed positive rational rate, every nontrivial property determined by the number of extreme translation-invariant stationary laws, all-law ergodicity, and the existence of a translation-invariant stationary law with nondecaying local covariance is undecidable. A refined theorem retains the numbers of extreme translation-invariant stationary states with nondecaying covariance and with a Bragg atom at some nonzero wavevector. The same holds for circuits of local quantum channels on the line with independent single-qubit depolarization after every layer, and, with uniqueness among all Gibbs measures in place of ergodicity, for Gibbs measures on the square lattice given by positive rational Boltzmann weights. On the line, ergodicity excludes nondecaying covariance, and two or more translation-invariant stationary states force nondecaying covariance and exclude ergodicity, so only four kinds of coarse data can occur. In both line categories, for every $r \geq 1$ the cut "at most r extremes" is Π_2^0 and halting-hard, and for every $k \geq 2$ the cut "at least k" is Σ_2^0 and nonhalting-hard. In dimension seven the same Rice theorem holds for quantum circuits followed by whole-site depolarization, and under a stated published-source condition C_T, at another fixed rational rate, these cuts are Π_2^0-complete and Σ_2^0-complete in both categories. The scope has a decidable boundary: a classical translation-invariant stationary law attains the noise entropy floor exactly when the deterministic rule is constant.

Claude Opus 5.5, GPT-6.1 Sol, GPT-6.2 Sol and GPT-6.2 Astra assisted with this work; the authors take responsibility for every claim.

Archival record: doi:10.5281/zenodo.24241163.

1. Introduction

Positive noise gives every local transition a positive probability. It does not give an algorithm for deciding which stationary states survive in infinite volume. We study a finite deterministic lattice rule followed, at every site and time, by independent uniform replacement. The local rule is the input; the dimension and replacement rate are fixed throughout the decision problem.

Our main assertion has the Rice form: every nontrivial property of a specified stationary datum is undecidable. Here nontrivial means that the property holds for at least one valid input and fails for another. A translation-invariant stationary law is extreme when it is an extreme point of the convex set of such laws. We count these extreme points with their full cardinality. All-law (quantumly, all-state) ergodicity means local convergence to one stationary law or state from every initial law or state, including spatially noninvariant ones. We use the term all-law ergodicity in both categories.

The headline results concern the infinite line. At one fixed positive rational rate, Theorems 3.1 and 3.2 give the Rice assertion for classical replacement and for independent single-qubit depolarization after every circuit layer, respectively. Theorem 3.3 retains the ordered and Bragg-ordered extreme counts. The quantum gates are arbitrary exact finite

Date: Version 2.0.0, November 2026.

2020 *Mathematics Subject Classification.* 60K35, 37B15, 03D35, 81P45.

Key words and phrases. Probabilistic cellular automata, stationary states, undecidability, Rice theorem, Bragg order, local quantum channels.

local channels; their finite arity may grow with the input. Theorem 7.1 gives a square-lattice Gibbs analogue for strictly positive rational local weights, with temperature absorbed into the weights. The whole-site, once-per-round quantum category is not covered on the line, nor is the class of binary energy tables at one fixed temperature.

For comparison, the following seven-dimensional assertion remains part of the result.

Theorem A. *For deterministic finite-range rules on $\mathbb{Z}^7$ followed by uniform replacement at one fixed positive rational rate, every nontrivial property determined by the number of extreme translation-invariant stationary laws, all-law ergodicity, and existence of a translation-invariant stationary law with nondecaying local covariance is undecidable.*

For this seven-dimensional assertion the rate is explicit in (2.3). The alphabet and finite range may vary. The conclusion includes uniqueness of the translation-invariant stationary law, ergodicity, and every nontrivial cardinality condition on the extreme set. The same parameters work for all properties at once.

One might expect every property of the stationary laws to be undecidable. That is false, and the counterexample fixes the scope. We prove that a stationary spatial Shannon entropy density attains the minimum forced by replacement noise exactly when the deterministic local rule is constant. A finite table scan decides this property. Binary reset and noisy identity have the same datum in Theorem A, but opposite entropy-floor verdicts. Thus the datum, rather than the phrase "stationary property", identifies the Rice class.

The covariance flag in Theorem A allows mixtures. Two distinct translation-invariant stationary laws already give a mixture with nondecaying covariance, by the variance of spatial averages. Order in an individual extreme state is a more discriminating question. Our stronger theorem retains both the number of ordered extreme states and the number having a Bragg atom at some nonzero wavevector. On the line it holds in the two categories just specified; in dimension seven it holds also for finite-depth local quantum circuits with whole-site depolarization. The construction and its two outcomes appear in Figure 1.

On the line, the controller consists of two copies of the unchanged hardwired hierarchy of [37, Theorem C]. If the program halts, a common graphical map forgets every initial controller row and the coupled system has exactly the datum of the arbitrary input, called the host. If it never halts, high-rank spatial addresses reveal the relative origin of the two hierarchies with summable error. Their compatible relative origins label continuum many extreme TI stationary states. Each selected state has a single-row Bragg atom at π. One universal datum therefore absorbs every host, giving the Rice reduction. Figure 1 shows the two alternatives.

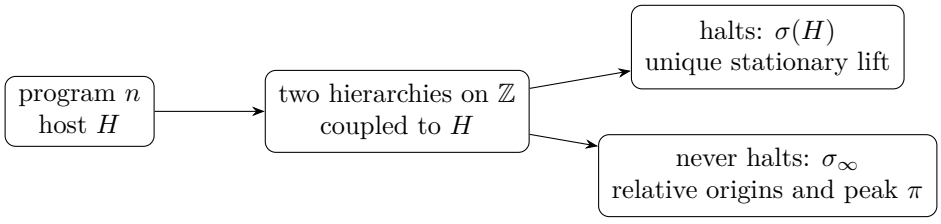

FIGURE 1. The line construction preserves every host datum in the halting case and has continuum many Bragg-ordered extreme TI stationary states in the nonhalting case. Relative spatial origins survive the averaging needed for translation invariance and stationarity.

The separate seven-dimensional proof uses a Wang controller on six coordinates and independent slices in the seventh. Its crystalline amplification and whole-site quantum implementation are retained in Sections 9–11.

On the line the finite upper total-count cuts belong to Π_2^0 and are halting-hard; finite lower cuts belong to Σ_2^0 and are nonhalting-hard (Theorem 8.3). Matching second-level hardness on the line is not established here.

In dimension seven the total phase count also has a conditional precise arithmetical degree. Here a phase means an extreme TI stationary law or state. Under the published-source condition C_T in Section 12, Theorem C makes "at most r phases" Π_2^0-complete for every $r \geq 1$ and "at least k phases" Σ_2^0-complete for every $k \geq 2$, in both input categories at another single fixed rational rate. A stationary version of host preservation gives one-sided second-level hardness for every nontrivial predicate of $(N, N_{\mathrm{ord}}, N_{\mathrm{Br}}, L)$ and excludes nontrivial Δ_2^0 members. The total-count cuts are the concrete complete instances; the ordered counts and covariance flag have no arithmetic upper bounds established here.

1.1. **Relation to earlier work.** Rice's theorem concerns nontrivial semantic properties of computably enumerable sets [1]. Kari established the corresponding theorem for limit sets of one-dimensional cellular automata [5]; Guillon and Richard refined its fixed-alphabet aspects [7]. Delacourt proved Rice theorems for μ-limit sets and generic limit sets [8, 9]. Hellouin de Menibus and Sablik, and Delacourt and Hellouin de Menibus, obtained Rice theorems for sets of limit measures from random initial configurations [11, 10]. Their dynamics is deterministic. Table 1 places ongoing noise beside these limit constructions.

Sources	Object and randomness	Undecidable question
Kari; Delacourt	Deterministic CA limit sets; random initial data for μ-limits	Nontrivial properties of the limit object
Hellouin de Menibus–Sablik; Delacourt–Hellouin de Menibus	Limit-measure sets from a Bernoulli initial law	Nontrivial properties of these sets
Kurdyumov; Toom; Bušić–Mairesse–Marcovici	PCA allowing zero transition probabilities	Ergodicity
Marsan–Sablik	Uniform perturbation as noise tends to zero	Stochastic stability
Marcovici–Taati	Positive-rate PCA in dimension at least two, with fixed uniform noise	Stationary Bernoulli existence versus nonexistence
Douglas–Mghirbi preprints	Fixed-noise controllers and noise diagrams	Ergodicity and stationary multiplicity
This paper	Stationary classical and quantum states at fixed positive noise	Every nontrivial predicate of the stated datum

TABLE 1. Undecidability results for cellular automata and PCA, by object and source of randomness.

Undecidability of PCA ergodicity goes back to Kurdyumov and Toom [15, 16, 17, 19, 20]. Their local rules may have zero transition probabilities. Bušić, Mairesse and Marcovici prove that ergodicity of a deterministic one-dimensional CA is equivalent to nilpotency [22], whose undecidability is due to Kari [6]. Unique ergodicity for deterministic CA is also undecidable [23]; that temporal notion concerns invariant measures of the deterministic map. It is distinct from uniqueness within the TI stationary set used for our count.

Marsan and Sablik prove undecidability of stochastic stability under uniform perturbations in the zero-noise limit [26]. In every dimension at least two, Marcovici and Taati prove computable inseparability between positive-rate PCA admitting the uniform stationary Bernoulli law and those admitting no stationary Bernoulli law [25, Theorem 3]. Their construction permits any chosen rational replacement rate in $(0, 1)$. The additional assertion here is the Rice closure

over our stationary datum. Work of Toom and Gács explains how persistent memory can coexist with positive rates [18, 40, 21]. The hardwired hierarchy imported from [37] adapts the reliable-computation and self-repair organization of Gács [40]; the new steps here are the relative-origin construction and the host correspondence. Marsan, Sablik and Törmä exhibit a perturbed CA with two ergodicity transitions [28].

Gu, Weedbrook, Perales and Nielsen formulated a Rice programme for macroscopic ground-state properties of periodic Ising systems and expected it to extend to positive temperature [29]. At a fixed positive temperature, Sandhu proved that uniqueness of the Gibbs state is undecidable for a computable family of translation-invariant finite-range interactions on $\mathbb{Z}^2$ [30]; see also [36]. The space-time Gibbs representation of PCA goes back to Domany and Kinzel [41] and was developed by Goldstein, Kuik, Lebowitz and Maes [42] and by Lebowitz, Maes and Speer [43]. Miękisz and Taati use noise-resilient automata to construct nonperiodic low-temperature Gibbs states for finite-range interactions on $\mathbb{Z}^4$ [44]. The all-DLR uniqueness implication used in our square-lattice reduction is [38, Theorem 3.1]. For quantum ground states, spectral-gap and phase-diagram undecidability were established by Cubitt, Pérez-García and Wolf and by Bausch, Cubitt and Watson [32, 33]. Our quantum dynamics is dissipative and locally implemented. Its hardness is inherited from a classical measure-and-prepare controller. Carrasco-Vargas proved that every nontrivial dynamical property of sofic subshifts on groups with undecidable domino problem is undecidable, and that computable real invariants of SFTs monotone under disjoint unions and products are constant [12]; our bath invariance plays a role similar to his product monotonicity.

Effective symbolic systems and entropy invariants are treated more broadly by Hochman and by Hochman and Meyerovitch [13, 14]. Zero-noise limit-measure sets are characterized by Marsan and Sablik [27].

For isolated quantum dynamics, Shiraishi and Matsumoto prove undecidability of thermalization even for one-dimensional shift-invariant nearest-neighbor systems with a fixed product initial state [31].

Our earlier preprints prove that ergodicity is undecidable at one fixed rate on $\mathbb{Z}^6$, $\mathbb{Z}^4$ and $\mathbb{Z}^3$, and that a stacked controller on $\mathbb{Z}^7$ has either exactly one invariant measure or a continuum of extreme ones, with no algorithm deciding which [34, Corollary 3.11]. They also give one noisy automaton with an uncomputable noise diagram [35], and they prove spontaneous magnetization and Gibbs uniqueness undecidable at a fixed positive temperature [36]. The classical uniqueness and ergodicity instances of Theorem A, and every total-count predicate that distinguishes one extreme from continuum many, therefore already follow from [34], whose controller theorem we import. The full Rice closure over realized total counts uses the host correspondence. The hardwired line construction of [37, Theorem C] already proves fixed-rate ergodicity undecidability on $\mathbb{Z}$; its one-site temporal witness alone does not count stationary row laws. Here the high-rank decoder and two-hierarchy relative origins establish continuum many ordered extreme stationary laws on the line. The host correspondence yields Rice closure in the classical and specified quantum categories. The Gibbs result uses [38, Theorem 3.1], with uniqueness among all DLR measures. The seven-dimensional arguments retain the ordered counts, host correspondence and conditional completeness theorem described above. Changes from version 1.0.0. The present paper adds the theorems on the line: the classical, quantum and refined Rice theorems (Theorems 3.1–3.3), the square-lattice Gibbs version for positive rational weights (Theorem 7.1), the necessary data on the line (Theorem 8.1) and partial degree results (Theorem 8.3). Theorems A, B and C and the other seven-dimensional results are unchanged in statement from version 1.0.0 (10.5281/zenodo.24241164).

1.2. **Reading guide.** Section 2 fixes the datum and the seven-dimensional input conventions. Section 3 specifies the line categories and statements. Sections 4–6 prove the hierarchy estimates, relative-origin construction, quantum implementation and full host correspondence.

Section 7 proves the Gibbs theorem; Section 8 gives obstructions and partial degrees. Sections 9–11 give the seven-dimensional proofs. The conditional totality reduction is in Section 12, the decidable entropy boundary in Section 13, and effective convex geometry in Appendix A.

2. Inputs, stationary data and the refined theorem

2.1. **The seven-dimensional input categories.** A classical input consists of a nonempty finite alphabet A, a finite stencil $\mathcal{N} \subset \mathbb{Z}^7$, and a deterministic table $F : \mathsf{A}^{\mathcal{N}} \to \mathsf{A}$. One round applies F at every site and then replaces each whole output symbol independently by a uniform symbol with probability ε_0. The flags and replacement values are independent across sites and times. The resulting local transition probabilities are

$$p(a \mid u) = (1-\varepsilon_0)\mathbf{1}_{\{a=F(u)\}} + \varepsilon_0/|\mathsf{A}|.$$

All entries are positive. Finite alphabets and stencils are unbounded over the input category.

A quantum input has a finite-dimensional site space $\mathbb{C}^d$ and a finite list of layers of finite-range local completely positive trace-preserving (CPTP) gates. Gates have exact algebraic Choi or Kraus entries and use fresh independent product ancillas. A layer acts on disjoint registers; a finite periodic schedule is permitted when the overall channel Φ is translation invariant. CPTP validity and overall translation invariance are promises on the presentation. After Φ, every whole site independently undergoes

$$\mathcal{D}_{\varepsilon_0}(X) = (1-\varepsilon_0)X + \varepsilon_0 \operatorname{Tr}(X) I_d/d.$$

Arbitrary initial states, including entangled ones, are allowed. Disjoint sufficiently distant output cones use disjoint circuits and independent environments. A causal channel supplied with globally shared randomness need not have this property and is outside this category.

2.2. **The datum.** Write $\mathcal{K}(M)$ for the compact convex set of all translation-invariant stationary laws or states of M. Translation invariance always means invariance under every spatial translation. In both categories, local convergence means convergence of every finite-region marginal. A state is extreme here precisely when it belongs to $\operatorname{Ext}\mathcal{K}(M)$.

For real classical or Hermitian quantum local observables f, g, write

$$\operatorname{Cov}_\omega(f, \tau_x g) = \omega(f\tau_x g) - \omega(f)\omega(\tau_x g).$$

Let $\ell(\omega)$ mean that some such f, g satisfy $\limsup_{|x|\to\infty} |\operatorname{Cov}_\omega(f, \tau_x g)| > 0$. For $B_n = \{0, \dots, n-1\}^7$ and a nonzero wavevector $q \in (\mathbb{R}/2\pi\mathbb{Z})^7$, define

$$\operatorname{Av}_n^q(f) = \frac{1}{|B_n|} \sum_{x \in B_n} e^{-iq\cdot x} \tau_x f, \qquad b_q(\omega) \iff \exists f : \limsup_n \omega\big(\operatorname{Av}_n^q(f)^* \operatorname{Av}_n^q(f)\big) > 0. \tag{2.1}$$

When $b_q(\omega)$ holds we say that ω has a Bragg atom at q. To see the spectral meaning, implement translations in the classical L^2 space or the quantum GNS space. For the commuting unitaries $V_i = e^{-iq_i} U_i$, cube averages vanish on every coboundary $v - V_i v$ by telescoping. The closed span of these coboundaries has orthogonal complement the joint fixed space. Contraction and density show that the cube averages converge to its orthogonal projection. The second moment in (2.1) therefore has a limit equal to the squared norm of the projection of the observable vector.

The refined datum is

$$\sigma(M) = (N, N_{\mathrm{ord}}, N_{\mathrm{Br}}, E, L), \tag{2.2}$$

where

$$\begin{aligned} N &= |\operatorname{Ext}\mathcal{K}(M)|, \\ N_{\mathrm{ord}} &= |\{\omega \in \operatorname{Ext}\mathcal{K}(M) : \ell(\omega)\}|, \\ N_{\mathrm{Br}} &= |\{\omega \in \operatorname{Ext}\mathcal{K}(M) : \exists q \neq 0 \ b_q(\omega)\}|. \end{aligned}$$

Here E means that a unique stationary state among all states attracts every initial state locally; L means $\exists\omega \in \mathcal{K}(M) : \ell(\omega)$. The cardinalities are full cardinalities. In particular N_{Br} counts states with a peak at *some* nonzero wavevector.

A predicate is a function of σ if equal data force equal verdicts. We allow arbitrary coupled dependence on its five coordinates; only realized values matter. An independent constant bath is a second on-site track whose pre-noise rule is constant (quantumly, a fixed on-site preparation), followed by its own independent replacement or depolarization at the host rate. To discuss independent constant baths and arbitrary on-site unitaries, define the same datum on the larger class of locally implemented media with unrestricted channel entries, including independent bath products. Such operations can leave the single-replacement or exact-algebraic input category. Decision problems and nontriviality below are restricted to the finite exact inputs just specified.

The full result strengthens Theorem A.

Theorem B (Stationary-state Rice theorem)**.** *Set*

$$K = 16^4, \quad \Delta = 18 \cdot 5K, \quad B = 2\Delta^2, \quad D = 7, \quad \varepsilon_0 = B^{-200}. \tag{2.3}$$

In each input category separately, every nontrivial predicate of σ is undecidable. Let

$$\sigma_\infty = (\mathfrak{c}, \mathfrak{c}, \mathfrak{c}, \mathrm{false}, \mathrm{true}), \qquad \mathfrak{c} = 2^{\aleph_0}.$$

This value is realized. If the predicate is false at σ_∞, halting many-one reduces to it; if true there, nonhalting many-one reduces to it. Every predicate of the datum is invariant under on-site relabeling or unitary conjugation, as appropriate, and adjoining an independent constant bath.

Exactly-k conditions are subject to nontriviality. For each finite $k \geq 2$, realizability of $N = k$ at these parameters remains open; if realized, deciding $N = k$ is halting-hard. The universal condition $N \geq 1$ is trivial. The theorem concerns this stationary datum, with varying finite alphabet or site dimension and range. Full stationary-set isomorphism, entanglement, storage, topological order and equilibrium Gibbs properties require other arguments. The fixed-alphabet question remains open. The line results use the separate input conventions in Section 3.

2.3. **Compactness and extreme lifts.** We give the elementary facts needed in the proof.

Lemma 2.1. *The set $\mathcal{K}(M)$ is nonempty compact and metrizable, and has at most $\mathfrak{c}$ points. Every nonempty compact face contains an extreme point. An autonomous on-site factor induces an onto continuous affine map of TI stationary sets, and every extreme factor state has an extreme joint lift.*

Proof. A classical law is specified by consistent finite-cylinder probabilities; a quantum state by consistent positive finite-region density matrices of trace one. Each coordinate lies in a compact finite-dimensional set. A diagonal subsequence and consistency give compactness and metrizability. TI equations are closed. Stationarity is closed because the one-step Heisenberg image of a local observable is local. Cesàro averages of iterates of a TI initial state have stationarity defect at most $2\|f\|/T$ on f and hence yield a stationary TI limit. There are countably many determining real local coordinates, which also bound the cardinality by $\mathfrak{c}$. The same argument without TI gives stationary existence among all states.

On a compact face, maximize the first coordinate, then the second on that maximizing face, and continue. Nested compactness gives an intersection in which all determining coordinates are fixed, hence one point. Any convex decomposition stays in every maximizing face, so this point is extreme in the original set. If a compact convex set has two points, a separating coordinate has disjoint minimizing and maximizing faces, each with an extreme point. Consequently $N = 1$ is equivalent to uniqueness within $\mathcal{K}(M)$.

Given a stationary factor state, start the joint process with that marginal and a fixed TI auxiliary state. The factor marginal remains fixed, and a Cesàro limit is a stationary joint lift. Projection is continuous and affine. The fiber over an extreme factor state is a compact face: a convex decomposition projects to a decomposition of that extreme state. Applying the preceding construction in the fiber produces an extreme joint lift. □

The statistic is coarse. In a compact metrizable convex set the nonextreme points form a countable union of compact sets: use midpoints of pairs whose distance is at least $1/j$. Thus the extreme set is G_δ. The perfect-set theorem for Polish spaces gives $N \in \{1, 2, \dots, \aleph_0, \mathfrak{c}\}$ without a continuum-hypothesis assumption. Moreover $N \geq 2$ implies $E =$ false and $L =$ true.

Lemma 2.2. *Two distinct TI stationary states imply L. A Bragg atom at a nonzero wavevector implies ℓ for that state.*

Proof. Choose a local real or Hermitian f with distinct means a_0, a_1 in two states and take their equal mixture ω. For the ordinary cube average $\overline{f}_n$, variance decomposition gives $\mathrm{Var}_\omega(\overline{f}_n) \geq (a_1 - a_0)^2/4$. If the covariance tended to zero, expansion as a double sum would make this variance tend to zero: only $O(|B_n|)$ pairs have any fixed bounded displacement, and far covariances are arbitrarily small. This contradiction proves L.

For $q \neq 0$, the weighted mean $\omega(\mathrm{Av}_n^q(f))$ tends to zero by a geometric-sum factor in a coordinate with $q_i \neq 0$. If connected covariance decayed, its weighted double sum would vanish by the same near-pair/far-pair argument. Thus a positive Bragg moment forces nondecay. □

3. Rice theorems on the infinite line

For the line categories, use the definitions of Section 2 with spatial lattice $\mathbb{Z}$ and averaging intervals $[0, l) \cap \mathbb{Z}$. In particular E requires a unique stationary state among *all* states, attracting every initial state locally; extreme points are taken in the jointly TI and stationary set. Use the same five-coordinate notation σ and the same maximal value σ_∞. The wavevector is in $\mathbb{R}/2\pi\mathbb{Z}$ and is nonzero for a Bragg label.

A classical line input has a nonempty finite alphabet, an arbitrary finite stencil in $\mathbb{Z}$, a deterministic finite table, and independent uniform whole-symbol replacement at rate e after each update. A quantum line input has finitely many qubits per site and a finite repeating list of finite-range local CPTP gates with exact algebraic entries and fresh product environments. Gates within a layer act on disjoint internal tensor factors. After *every* layer, independently on every qubit, apply

$$\mathcal{D}_e^{(1)}(X) = (1 - e)X + e\,\mathrm{Tr}(X)I_2/2.$$

This includes preparation, reset and idle layers. The full noisy round is promised TI. The finite presentation, gate arity, number of wires, range and period may vary with the input. Arbitrary entangled initial states are allowed. Gate arity in the constructed family grows with the program; no bounded-arity primitive-library assertion is made. The general-site, once-per-round whole-site quantum category of Theorem B is not treated on the line.

Fix the native interpreter, track order and tape template of [37, Section 3 and Appendix A]. Choose w to be the least admissible integer for the finite requirements of that construction, including $w \geq 10^{10}$. Put $P = 5 \cdot 2^w$ and let T be its fixed native clock period. Both exceed $2w$. These choices are independent of the program. Define

$$\begin{aligned} F_* &= \max(P, T), & M_* &= 1000w(P + T + 1)^3, \\ D_* &= 8(P + 2)^2(T + 2)^2, & A_* &= (4M_*(P + 1) + 8)(4M_*(T + 1) + 8), \\ R_* &= 4M_*F_*, & V_* &= 10^9 R_*^4 A_* F_*^4, \\ e &= \frac{1}{4D_*w^2V_*^2}. \end{aligned} \tag{3.1}$$

These are definite computable integers and a definite rational once the fixed interpreter is chosen. Only the existence of one such rate is asserted; no decimal value is claimed. In particular $0 < e < 1/32$.

Theorem 3.1 (Classical Rice theorem on the line)**.** *At the fixed rational rate* (3.1)*, every predicate of* (N, E, L) *nonconstant on the realized classical line data is undecidable.*

Theorem 3.2 (Quantum Rice theorem on the line)**.** *For the line quantum inputs with independent single-qubit depolarization after every layer specified above, at the same rate* (3.1)*, every predicate of* (N, E, L) *nonconstant on the realized data is undecidable.*

Theorem 3.3 (Refined line theorem)**.** *In each line category separately, every predicate of* $\sigma = (N, N_{\mathrm{ord}}, N_{\mathrm{Br}}, E, L)$ *nonconstant on the realized values is undecidable at* (3.1)*. The value* σ_∞ *is realized. If a predicate is false there, halting many-one reduces to it; if true there, nonhalting many-one reduces to it. These conclusions hold also when every infinite count is reported as a single infinity symbol.*

The proofs occupy Sections 4–6. Compactness, extreme lifts and the implication from a nonzero Bragg atom to covariance order follow by the full proofs of Lemmas 2.1 and 2.2, with intervals in place of cubes. A line circuit has a finite backward cone and independent environments in distant cones, exactly as those proofs require. For quantum changes of coordinates, the depolarizing convention is covariant under products of single-qubit unitaries and wire permutations. An entangling on-site conjugation need not remain in this input category.

4. Two hierarchies and their relative spatial origin

4.1. **The imported deterministic construction.** Let f_n be exactly the radius-one hardwired rule of [37, Theorem C and Section 3.5], with its unique empty symbol. Its literal is the true initial configuration of $U(n)$, doubled at successive promotions. A promotion executes one machine step. No external parameter reader is used. Geometry, decoding, interpreter and native transition are unchanged. Section 3 of that reference makes the finite table computable from n, at fixed P, T, w. A logical step below means one native update.

We use the following numbered statements of [37]. Definition 4.1 defines i-sparsity by deleting one translate of $3w(P^{i-1}, T^{i-1})$ in every $w(P^i, T^i)$ window; 0-sparsity means emptiness. Lemma 4.2 gives disjoint child witnesses. Definition 5.1 supplies nominal and served tile observations, whose spatial support is inside the tile and whose lookahead is at most $T^j - 1$. Proposition 5.3 gives organization for every translated grid, every history with sparse deviations, true read-only operands and ordinary unhalted lower ranks; in particular nominally live tiles with distinct geographical origins cannot overlap. Lemma 5.24, with six consecutive served positive inputs of one tribe and two prescribed positive central transitions, gives service of both central homes for two periods, the correct lower clock and address residues, and the two served prescribed upper successors. It assumes a fully protected $(j-1)$-sparse rank-j test. Lemmas 6.2 and 6.3 give the halted all-start cut and ascent from any live physical site to a failed centered test after a common burn-in. Lemma 6.4 gives the clean initialized negative-time hierarchy in the hardwired family. Lemma 6.5 constructs its deterministic nested tests. The clock witness in Theorem C alone would not imply a count of stationary row laws; we need the spatial-residue conclusion of Lemma 5.24 and the decoder below.

4.2. **A quantitative mark estimate.**

Lemma 4.1 (Sparse tests and fixed-rank paths)**.** *Suppose deviations of any native history are contained in iid marks of rate at most* e*. For every prescribed centered test with half-sides* $M_*(P^{i+1}, T^{i+1})$*, failure of* i*-sparsity has probability at most*

$$a_i = A_* V_*^{-2\cdot 2^i}, \qquad r_i = R_* F_*^i$$

and is supported in a max-norm ball of radius r_i. Moreover

$$\sum_i a_i < \tfrac{1}{4}, \qquad \sum_i (i+1)a_i < \infty, \qquad \sum_i 10^6 r_i^4 \sqrt{a_i} < \tfrac{1}{2}.$$

For each fixed K, a nearest- or diagonal-neighbor path of vertices with a failed centered test of some rank $0 \le i \le K$, starting at a prescribed point and reaching max-distance h, has probability at most $C_K \exp(-c_K h)$.

Proof. All tests needed here fit the stated half-sides. One-rank work is in a w-window and its collar has $\lfloor w/5 \rfloor$ cells. The nested tests of [37, Lemma 6.5] have fewer than w cells at their own scale. Cap ancestry at scale $i+1$ costs at most $(4T/P + 4, 6)$ cells before its work collar and rounding [37, Lemma 6.3]. Covering every translate of a tile through one point costs at most two further cells per side. All these finite multipliers are below M_*.

Let b_i bound failure on a full doubled box with side vector $2w(P^i, T^i)$. Its base bound is $b_0 \le 4w^2 e$. The disjoint witnesses of [37, Lemma 4.2] have at most $(2P+3)(2T+3)$ possible child positions, hence at most D_* pairs. Independence on disjoint child boxes gives $b_{i+1} \le D_* b_i^2$. Induction yields

$$b_i \le D_*^{-1}(4D_* w^2 e)^{2^i} = D_*^{-1} V_*^{-2\cdot 2^i}.$$

Cover a larger prescribed box by a mesh of spacing $w(P^i, T^i)$, including one mesh step past each boundary. Every moving w-window has its intersection covered by one doubled box. There are at most A_* such boxes. If all pass, restriction gives the deletion required in every moving window. At rank zero the same cover bounds every marked site. The union bound proves the asserted a_i; the enclosing box has radius at most r_i. The estimate applies to arbitrary error values, since sparsity is inherited by subsets of the marks.

Using $2^i \ge i+1$, the first two sums are bounded respectively by

$$\frac{A_*}{V_*^2(1 - V_*^{-2})}, \qquad \frac{A_*}{V_*^2(1 - V_*^{-2})^2}.$$

The third is at most

$$\frac{10^6 R_*^4 \sqrt{A_*}}{V_*(1 - F_*^4/V_*)} < \tfrac{1}{2}.$$

The definitions in (3.1) give all three inequalities.

To prove the path assertion, cover a finite such path by its failing test balls. Greedily retain disjoint balls in decreasing radius. Every removed ball is contained in a retained triple ball. A simple chain of adjacent retained triple balls connects the endpoints. The sum of its radii is at least $h/8$: the diameters are at most $6r$ and a unit jump costs at most one further radius. The first center has at most $(6r+1)^2 \le 49r^2$ choices. Successive centers of radii r, s have at most $(6r + 6s + 3)^2 \le 225 r^2 s^2$ choices. For an ordered type list the total placement count is therefore at most the product of $10^6 r_i^4$. Disjoint retained supports make the failure probability at most the product of their a_i. Reserve half of each exponent. Since the radius sum is at least $h/8$, this half contributes at most

$$\exp\left[-\frac{h}{8} \min_{0 \le i \le K} \frac{-\log a_i}{2r_i}\right].$$

Sum the remaining half over type lists and placements. The resulting geometric series converges because $\sum_i 10^6 r_i^4 \sqrt{a_i} < 1/2$. This proves the claimed exponential estimate. Only K enters its exponent. The entire event is a function of the common dominating marks, independently of which histories might realize discrepancies. □

4.3. **Forgetting on halting.**

Lemma 4.2 (All-input common-map forgetting)**.** *If $U(n)$ halts, a native history whose deviations are contained in the marks of Lemma 4.1 forgets all starting rows on a common event with one-site failure probability $C\exp(-ct)$, provided that with innovations fixed every update, even a marked update, is a deterministic function of its three old symbols. The constants may depend on n. The bound is for the event that the entire graphical map at the observed site is nonconstant.*

Proof. Let K be the first halted rank. Its true literal forces the protected strong issue to be empty on every triple, including forged input and work symbols; all lower ranks are ordinary. For $K \geq 1$, [37, Lemmas 6.2–6.3] therefore apply literally. Their premises use the native rule off deviations, the unchanged geometry and true hardwired operands, all of which hold here. Choose the common burn-in greater than twice the largest radius of the finitely many centered tests below K, plus two. After this burn-in any live vertex has a failed centered sparsity test below K. This inclusion is simultaneous for every starting row on the same marks.

At a discrepancy of two trajectories at least one symbol is live, since the empty symbol is unique. If the three parents were identical, the common deterministic update with fixed innovations would give identical children. Trace a discrepant parent backwards to half the elapsed time. For t beyond twice the burn-in plus four, this gives a nearest- or diagonal-neighbor path of failed tests with time span at least $\lfloor t/2 \rfloor - 1$. Lemma 4.1 makes its probability exponentially small. The same noise-only path is necessary if *some* pair of starting rows disagrees, so no union over rows is needed. Finite exceptional times enter the prefactor.

At $K = 0$ an unmarked output is always empty. A discrepancy must follow a backward path all of whose vertices are marked. There are at most 3^t length-t paths and each has probability at most e^t, giving $(3e)^t$. For exact classical replacement, the rank-zero constant rule in fact coalesces in one update. A marked encoded update need not erase a discrepancy, so the path bound is used for that implementation.

If an ascent test reads beyond the observation date, independently extend the innovations to those dates. They can witness the necessary path without changing the nonconstancy event, which depends only on the actual backward cone. Finally a union bound over an observed finite set gives the bound for its whole common map. □

4.4. **Decoding a spatial origin.**

Lemma 4.3 (Stopped descent and finite decoders)**.** *If $U(n)$ never halts, initialize its clean aligned hierarchy from* [37, Lemma 6.4]*. For every $j \geq 1$, a finite space-time decoder at (x, t), $t \geq T^j$, returns the left edge of its true rank-j tile modulo P^j, except with probability at most*

$$d_j = \sum_{i \geq j-1} a_i, \qquad \sum_{j \geq 1} d_j < \infty.$$

The bound is uniform in x, t and integer spatial shifts of the initialization. At rank zero, with probability at least $1 - \sum_i a_i > 3/4$, the physical root is live with both initialized structural residues, in particular its spatial address modulo P.

Proof. Every rank is ordinary on a nonhalting input. Join the fully occupied clean negative-time history to any forward history obeying the native rule off the marks. For $j \geq 1$, examine all integer tiles of width P^j and height T^j containing (x, t), using half-open tiles. Compute their finite nominal observations from [37, Definition 5.1]. Return the common left edge modulo P^j if at least one is nominally live and all nominally live candidates have that edge; otherwise fail. The decoder is a cylinder function, possibly reading dates on either side of (x, t).

Start the geometrical ascent of tests with the entire true tile. Its dimensions are less than P^{j+1} and $T^{j+1}/2$, as required by [37, Lemma 6.5]. For completeness, if a smaller time span

is $[a, b)$ of length less than $T^k/2$, put $h = \lfloor a/T^k - 1/2 \rfloor$; it lies in $[(h + 1/2)T^k, (h + 2)T^k)$. If its spatial span is $[c, e')$ of length less than P^k, choose $k' = \lfloor c/P^k \rfloor$ and the two homes $k', k' + 1$. Add the six bottom reads, their work, covering homes and the prescribed collar. The resulting box again fits the next rank, by the strict inequalities of that lemma. Every choice is determined by the true grid and the observation, independently of the noise.

At a sufficiently large finite top rank r choose first test date $-T^r$. Enlarge r so $4(t+2) < T^r$ and the bottom tile and collar fit. All six top reads, including their lookahead at most $T^r - 1$, are strictly before zero. The initialized past supplies their common tribe and both prescribed positive successors. Impose $(k - 1)$-sparsity for these tests at $k = j + 1, \dots, r$. Apply [37, Lemma 5.24] at the top. It serves both central homes throughout the covering two periods with their correct child residues. Its two upper successors supply the actual positive transitions needed at the next descent. Repeating the same argument, with those actual served inputs and successors at each rank, and stopping at j, gives a served live observation on the true tile. A served live observation is nominally live. No recovery of arbitrary mutable root work data is asserted.

For uniqueness, put the true tile, all overlapping candidates and their collars in one prescribed test of scale j. The size bound was checked in Lemma 4.1. Its $(j - 1)$-sparsity implies j-sparsity. The classification of [37, Proposition 5.3(i)] excludes nominally live overlapping candidates with different spatial origins. In that proposition geographical origin is the full space-time offset modulo (P^j, T^j); its shifted-child-grid argument uses both offsets. The decoder's common-edge convention also works using only the spatial-origin consequence. The true tile supplies existence, and all live candidates therefore have its spatial residue. The failed-test union has probability at most $\sum_{i \geq j-1} a_i$. Negative-time deviations are empty and may be dominated by an independent extension of the marks. This proves the uniform bound, and the weighted sum in Lemma 4.1 proves summability. Continuing the descent through rank one instead of stopping gives the physical root and both residues by the explicit conclusion of [37, Lemma 5.24], with error at most $\sum_i a_i$. □

4.5. **Stationary laws with distinct relative origins.**

Proposition 4.4 (Continuum many ordered extreme states)**.** *Use two autonomous copies of the nonhalting hierarchy on the same line. Their noises may be dependent, provided each satisfies Lemmas 4.1–4.3. More generally there may be a fixed positive number of logical steps per observed round. The joint controller has continuum many extreme TI stationary states with a single-row Bragg atom at π.*

Proof. Let $G = \varprojlim_j \mathbb{Z}/P^j\mathbb{Z}$. It has cardinality $\mathfrak{c}$: each digit may independently be zero or one, while it is a countable product of finite sets. Fix $d \in G$. For $h \geq 1$, let $\hat{d}_h \in [0, P^h)$ be its representative. Initialize the first copy with spatial origin zero and the second with origin $\hat{d}_h$, on the same time grid. Neither rule nor noise reads d. In an encoded implementation start at the boundary preceding fanout with encoded centers and default auxiliaries. Run the actual dynamics. Average its trajectory over spatial shifts $0, \dots, h - 1$ and time shifts by $0, \dots, h - 1$ complete rounds, extending negative dates arbitrarily for compactness. Take a common weakly convergent subsequence on every finite window.

Boundary defects in either average are at most $2/h$, so the limit is TI and invariant under complete-round time translation. The conditional transition identities against any finite past cylinder are the true kernels: each initialized experiment uses fresh innovations, complete-round shifts retain the same schedule phase, and the proportion of shifts whose cylinder reaches before the initialization tends to zero. These identities pass to the limit on cylinders and then by approximation. Thus the limit is the canonical two-sided trajectory of its stationary boundary row, with the intermediate kernels inserted in their specified order. For any stationary row law this trajectory map is affine on finite cylinders and Borel; use a

past boundary and stationarity for two-sided cylinders. In the quantum implementation all boundary controller states are diagonal, as proved in Lemma 5.2, so the same construction applies to the full basis-valued register and its local logical decode.

Apply the two rank-j decoders at the space-time origin and subtract the first returned residue from the second modulo P^j, failing if either fails. Common spatial shifts cancel; time shifts do not move a spatial origin. For $h \geq j$ the correct difference is $d \bmod P^j$. At fixed j the proportion of starts too early is $O(T^j/h)$, and on the others the error is at most $2d_j$, with d_j from Lemma 4.3. Each decoder is a finite cylinder of the intermediate history, so the same bound holds in the limit. Borel–Cantelli gives probability one to the Borel event B_d that every sufficiently large output is the residue of d. For distinct d these events are disjoint, since a differing residue remains different at all higher ranks.

Let f on the first lane be $(-1)^a$ on a live symbol with spatial address a, and zero on the empty symbol; for encoded rows use the locally decoded observable. Since P is even, Lemma 4.3 gives in the origin-zero experiment, for every site and date, including round boundaries,

$$\mathbb{E}[(-1)^x f(x)] \geq 1 - 2\sum_i a_i > \tfrac{1}{2}.$$

A spatial shift changes the prescribed sign by one common parity. For every interval length l, Jensen's inequality before the shift mixture gives

$$\mathbb{E}\left|\frac{1}{l}\sum_{x=0}^{l-1}(-1)^x f(x)\right|^2 \geq \tfrac{1}{4}.$$

This is a finite row observable. The bound passes through both averages and the weak limit. The phase-twisted mean ergodic argument in Section 2 shows that its limit is the Bragg second moment at π.

The TI stationary controller set is compact convex metrizable in the locally convex space of its countable real cylinder coordinates. The metrizable representation theorem [39, Theorem 4.2, pp. 316–317] therefore represents our row state by a probability measure carried by its extreme points. The mixture identity for their canonical trajectories holds on cylinders by linearity of kernels and stationarity, hence on all Borel sets by the monotone class theorem. Since B_d has probability one, almost every representing component also gives it probability one. The finite Bragg second moments are bounded affine functions of a row state and converge in every TI state. Bounded convergence shows that the integral of the limiting moments is at least $1/4$. A set of positive representing measure has a positive moment, so some extreme component has both label d and a peak at π. A stationary row state has a unique canonical trajectory and cannot carry two disjoint probability-one labels. Varying d gives $\mathfrak{c}$ distinct such extremes; countable coordinates give the matching upper bound. No uniqueness of the representation, spatial ergodicity of its components, or conditioning of the dynamics on a noncausal event is used. □

5. The quantum controller with noise after every layer

Lemma 5.1 (Encoded native updates and independent marks)**.** *At rate e of (3.1), for every n and host period $J \geq 2$ one can effectively construct a quantum controller with exactly J noisy layers per round. It performs $s = \lfloor J/2 \rfloor \geq 1$ native steps per round. Its effective classical update equals f_n off a common iid dominating set of marks of probability at most e, and with innovations fixed it is a deterministic function of its three previous effective symbols on every update, including marked ones.*

Proof. Let $b = \max(1, \lceil \log_2 |\mathsf{A}_n| \rceil)$ and choose an injective b-bit code of the native alphabet. Decode unused words as the unique empty letter. Repeat each bit on m wires, where m is the least odd positive integer with $4b2^{-m} \leq e$, and decode by majority. A finite basis

measurement and preparation is an exact local CPTP gate: enumerate its finite input-output table to obtain algebraic Kraus matrices. This is legal in the arbitrary finite-arity channel category, without a noisy bounded-arity compilation claim.

Use two layers per native step. In the fanout layer, one on-site trace/measure-and-prepare gate decodes the center and prepares three repetition-encoded copies for destinations $x-1, x, x+1$, discarding all old contents. The zero-offset port is the center bundle itself: each lane has exactly three bm-wire words per site, center, left port and right port. In the lookup layer, the gate at destination y consumes its three dedicated incoming words, decodes them, applies f_n, writes the encoded result into the center at y, and prepares zeros on the other two consumed words. Every port is used once. The lookup gates have spatial diameter two and act on disjoint internal registers, even where their sites overlap. Use s pairs of layers, with the possible one remaining idle layer immediately after the final lookup. Noise acts on every wire after all these layers, including fanout, lookup and idle. All auxiliaries are included in the physical site space.

At a pre-fanout boundary the effective letter is the decoded center. Resolve a depolarization of a basis wire into identity with probability $1-e$ and replacement by a fresh fair bit with probability e. For each logical destination and step, mark if at least half the wires were touched in any of its $3b$ received bit bundles after fanout, or in any of its b output-center bundles after lookup or the following idle. Add independent dummy idle flags on shorter steps only to the bounding event. Each wire then has touch probability at most $2e$. Different destination/step pairs use disjoint wire-time samples: incoming ports have dedicated wires, and a center's post-lookup noise and next post-fanout noise are different samples. Thus, after dummy flags, the bounding events are independent across destination/step pairs and identically distributed, each of probability at most e; they contain every deviation of the effective update. Indeed the per-bundle binomial bound is

$$\Pr\{\text{at least } m/2 \text{ touched wires}\} \le 2^m (2e)^{m/2} \le 2^{-m},$$

since $e \le 1/32$. The union over $4b$ bundles gives the claimed rate. Independent extra marks can enlarge it to exactly e if desired. On an unmarked update all majorities are correct, hence the output is exactly f_n on the three old effective symbols. On a marked update the fanout has already discarded all raw input except those three decoded symbols. With all innovations fixed the output is still a deterministic function of those symbols. This checks both the arbitrary-deviation hypothesis in Lemma 4.3 and the parent-dependence hypothesis in Lemma 4.2. The estimates charge majority failures, not the occurrence of any reset somewhere in a growing native word. □

Lemma 5.2 (The entire physical controller)**.** *Every stationary quantum controller state is diagonal at the observed round boundary. In the halting case the common-map forgetting of Lemma 4.2 holds for the full physical register, with a finite enlargement of the observed support. It applies to arbitrary entangled initial states. In the nonhalting case the full controller has the stationary states of Proposition 4.4.*

Proof. The first fanout measures or discards every controller wire. All later preparations, lookups and depolarizations keep outputs basis diagonal, including at whole-round boundaries. A stationary state is consequently a classical boundary law. For fixed innovations, equality of the three effective parents at the last step gives equality of its complete raw center output and subsequent corrupted bits. Both noncenter words have been overwritten by constants in lookup and subsequently see only fresh noise. Thus an enlargement of a finite support by the radius-one parents at the preceding step suffices to determine the full raw output there. There are at least $s \ge 1$ logical steps per round and at most a one-step initial delay. Lemma 4.2 and a finite union bound give full-register forgetting on halting.

At channel level, on every input including entangled ones,

$$\mathcal{D}_e^{(1)} = (1-e)\,\mathrm{id} + \frac{e}{2}|0\rangle\langle 0|\,\mathrm{Tr} + \frac{e}{2}|1\rangle\langle 1|\,\mathrm{Tr}\,.$$

The controller channel over finitely many rounds is therefore a convex combination, over innovations, of a deterministic classical map after basis measurement of centers and tracing out ports. On a good common-map branch its reduced map on the observed output is $X \mapsto \mathrm{Tr}(X)|c\rangle\langle c|$. This identity, rather than a claim only about basis inputs, removes input entanglement. Dummy marks and additional future innovations used for bounds change no actual channel.

On nonhalting, initialize encoded native centers and default ports. Lemma 5.1 supplies the uniform native estimates at every logical step. The repeating schedule has s steps per round and identical reference time grids for both lanes, so every premise of Proposition 4.4 holds. Its canonical histories include the intermediate basis-valued registers and locally decoded logical rows; its Bragg observable is on the boundary row itself. □

Lemma 5.3 (Hosts with one layer)**.** *Every line quantum input whose repeating period has one layer has datum* $(1, 0, 0, \mathrm{true}, \mathrm{false})$ *at* $e > 0$*. It may be replaced, when choosing a finite witness datum, by a fixed depth-two constant-preparation input of that datum.*

Proof. One repeated layer has a fixed partition into disjoint finite gate blocks, with idle wires regarded as singleton blocks. The union of blocks meeting a finite observed support is finite and closed under repeated ancestry. If it contains k wires, all output wires reset at a tick with probability e^k, giving a trace-and-prepare branch on that union. Its marginal contracts trace distance by at most $1 - e^k$ per tick, proving uniqueness and all-state local attraction. Begin with a product over gate blocks. Observations separated beyond twice the maximum gate diameter plus their widths have disjoint block sets and factor at every date. Their unique stationary limit has finite-range independence. It has no covariance order or nonzero Bragg atom, hence the stated datum. A fixed product preparation followed by noisy layers realizes the same datum at depth two. This replacement does not raise a general host channel to a power. For $J \geq 2$ the preceding implementation instead tensors host and controller on disjoint wires layer by layer, leaving the host's actual round and observation boundary intact. Independent qubit noise factors across them, so the joint noisy round is a tensor product of TI round channels, even if individual host layers use only a periodic schedule. □

6. Exact host correspondence and the line reductions

For a classical host H, take its deterministic product with two copies of f_n, then use one uniform replacement on the entire product alphabet with probability e. Every coordinate has exactly its original replacement marginal. In particular the controller lanes may share a flag; their individual estimates are unchanged. For a quantum host of period $J \geq 2$, tensor two encoded controllers with it in each of its J layers. Independent qubit noise makes the host and controller round channels literal tensor factors. Denote either joint construction by $J(H, n)$, using the equivalent depth-two witness of Lemma 5.3 when necessary.

Proposition 6.1 (Full correspondence on halting)**.** *If* $U(n)$ *halts, projection induces an affine homeomorphism between the TI stationary sets of* $J(H, n)$ *and* H*. It preserves* E*, and for each corresponding stationary state it preserves covariance order and every fixed nonzero-wavevector Bragg label. Consequently* $\sigma(J(H, n)) = \sigma(H)$*, without a finiteness or spatial-ergodicity assumption on the host stationary set.*

Proof. Let $\alpha_m \to 0$ bound the probability that the recent m-round graphical controller map to one physical site depends on any initial controller input. For the two lanes, add the bounds in Lemma 4.2, including the finite last-parent enlargement of Lemma 5.2. They hold under the true common innovations. For a host state μ, let $W_m(\mu)$ be the output after m rounds

starting from μ and a constant TI controller row. For any joint input ρ and an observation containing k controller sites, the local distance between its evolved output and $W_m(\rho_H)$ is at most $k\alpha_m$; distance means total variation or half trace norm.

Classically use the same initial host row and the same joint flags and replacement values in the two experiments. On controller constancy the complete observed pair agrees. This comparison retains all correlations induced by the shared host flags. Quantumly, condition the finite backward circuit on controller reset flags and bits. The channel identity in Lemma 5.2 makes each good reduced controller branch trace-and-prepare. Tensoring it with the independent host channel yields the same output from an entangled input as from the product comparison using its host marginal. Bad branches cost at most one in half trace norm. In fact, for quantum inputs the exact tensor product gives

$$W_m(\mu) = H^m(\mu) \otimes C^m(c),$$

where C is the whole controller and c the fixed constant start. On halting its unique stationary law is ν_C, and every stationary joint state is $\mu \otimes \nu_C$ for a stationary host state μ.

For existence of a stationary lift, start with a stationary host marginal times a constant controller and average iterates. Compactness of finite cylinder probabilities or density matrices gives a limit; stationarity and TI are closed, since finite observations have finite backward cones. The factor remains the required host marginal. For two stationary lifts of μ, their restrictions are within $2k\alpha_m$ for every m, hence identical. Projection is thus an affine continuous bijection of compact metrizable stationary sets, so its inverse is continuous. The identical argument without TI gives uniqueness of lifts among all stationary states.

Suppose the host has E. To prove attraction of the joint process, apply the comparison only to its last m rounds at time t. For fixed m, W_m depends continuously on finitely many host input coordinates. The host marginal at $t - m$ converges locally to μ, so the comparison converges to $W_m(\mu)$. The stationary joint lift is also within $k\alpha_m$ of that state. First let $t \to \infty$, then $m \to \infty$. This proves all-state attraction without a uniform host convergence rate. Conversely projection carries attraction to the host: every initial host state has an extension and every stationary host state a stationary lift. Thus E is equivalent.

For local Hermitian joint f, evaluating its comparison channel against the constant controller defines a local host observable $G_m(f)$ with $\|G_m(f)\| \le \|f\|$. The comparison is TI. For translates with disjoint m-round backward cones,

$$W_m(\mu)(f\tau_x g) = \mu\big(G_m(f)\tau_x G_m(g)\big). \tag{6.1}$$

Classically the two cones' innovations are independent, even though flags *within* each cone are shared across tracks. Quantumly fresh product environments give the same operator identity before any host state is evaluated. First moments satisfy the analogous identity by definition. If each controller support has at most r sites, the finite-observation bound on the product and the first moments gives a stationary/comparison covariance error at most

$$8r\|f\|\|g\|\alpha_m.$$

If the host state has no covariance order, let $|x| \to \infty$ at fixed m in (6.1), and then $m \to \infty$. The joint covariances decay. Host-only observables give the converse.

For any fixed local f and $q \neq 0$, expand the normalized weighted second moment over the l^2 pairs of translates in $[0, l)$. Each stationary/comparison pair error is at most $4r\|f\|^2\alpha_m$. At fixed m only $O_m(l)$ pairs have intersecting backward cones. Every other pair obeys (6.1). After division by l^2 the comparison and the host Bragg moment of $G_m(f)$ have the same limit. First let $l \to \infty$, then $m \to \infty$. If the host has no atom at q for any local observable, neither has the joint state. Host-only observables give the reverse implication for every q. The affine homeomorphism preserves extreme points, so it preserves $N, N_{\mathrm{ord}}, N_{\mathrm{Br}}$ and L in addition to E. □

Proposition 6.2 (The two line branch values)**.** *The construction is total computable from the finite presentations, and at the single rate* (3.1) *satisfies, for every host,*

$$\sigma(J(H,n)) = \begin{cases} \sigma(H), & U(n) \text{ halts}, \\ \sigma_\infty, & U(n) \text{ never halts}. \end{cases}$$

Proof. The first branch is Proposition 6.1 and, if needed, Lemma 5.3. In the second, Proposition 4.4 gives continuum many extreme TI stationary controller states with peak π. The factor projection onto the controller is onto: start the joint process with that stationary factor marginal and any TI host row, and take Cesàro limits. This works for shared classical flags because the projected replacement is exactly uniform on the controller alphabet. For each extreme controller state its fiber is a nonempty compact face. Successively maximize a countable separating coordinate family there to obtain an extreme joint lift, as in Lemma 2.1. Distinct controller marginals distinguish the lifts, and controller-only observables retain their peak. Hence

$$\mathfrak{c} \le N_{\mathrm{Br}} \le N_{\mathrm{ord}} \le N \le \mathfrak{c}.$$

Multiple stationary states preclude E and imply L by Lemma 2.2. The full datum is therefore σ_∞ for every host.

The native table is computable from n, with finite alphabet, fixed geometry and no word-width dependence in e. Classical product rules and their finite tables are enumerable. Quantum code length, repetition length, fanout and lookup tables are finite computable data; their Kraus entries are algebraic. For host period $J \ge 2$ the construction uses exactly J layers and preserves its observed round. The one-layer case is detected from the finite presentation and replaced as already proved. Thus the whole map is effective in the specified input categories. □

Proof of Theorem 3.3. A fixed program that never halts realizes σ_∞ by Proposition 6.2. Let Q be a predicate nonconstant on the realized data of the chosen category. There is a finite host H with $Q(\sigma(H)) \neq Q(\sigma_\infty)$, for otherwise the predicate would be constant. Fix that finite host as a constant of the reduction. No procedure finding witnesses from predicates is needed. Given a program n, output $J(H,n)$. Its datum is $\sigma(H)$ on halting and σ_∞ on nonhalting. If $Q(\sigma_\infty)$ is false, the output satisfies Q exactly when $U(n)$ halts; if true, exactly when $U(n)$ never halts. These are the asserted many-one reductions. Reporting all infinite counts by one symbol preserves both branches and the same finite-witness reasoning. □

Proofs of Theorems 3.1 and 3.2. A predicate of (N,E,L) is a predicate of the refined datum. Nonconstancy on its realized coarse values is exactly nonconstancy of the lifted predicate on the realized refined values. Theorem 3.3 applies in each category separately. □

7. A square-lattice Gibbs Rice theorem

A Gibbs input is a finite-alphabet, finite-range, translation-invariant interaction on $\mathbb{Z}^2$ given by finite tables of strictly positive rational local Boltzmann weights. Products are exact; at inverse temperature 1 the potentials are minus their logarithms. Temperature is absorbed into the weights. This is a computable class of rational weight tables; the class of $\{0,1\}$ energy tables at one fixed temperature is not covered. Let $\mathcal{K}_G$ be the set of all $\mathbb{Z}^2$-TI DLR measures of the specification. Define $N, N_{\mathrm{ord}}, N_{\mathrm{Br}}, L$ as in Section 2, using plane translations, square averages and wavevectors in $(\mathbb{R}/2\pi\mathbb{Z})^2$. Replace E by U, uniqueness among *all* DLR measures, and write $\sigma_G = (N, N_{\mathrm{ord}}, N_{\mathrm{Br}}, U, L)$.

Theorem 7.1 (Gibbs Rice theorem for rational weights)**.** *Every predicate of σ_G nonconstant on its realized values is undecidable for the rational-weight inputs just defined. This includes every nonconstant predicate of (N,U,L) and remains true when all infinite counts are reported by one infinity symbol.*

Lemma 7.2 (The controller Gibbs alternatives)**.** *There is a computable rational-weight controller interaction whose halting case has a unique DLR measure with exponential covariance decay for every pair of local plane observables, and whose nonhalting case has continuum many extreme TI DLR measures with a Bragg atom at* $(\pi, 0)$*.*

Proof. Use the two-track classical controller with one joint replacement of probability e. If its alphabet has size a, put the factor

$$p(c \mid \xi) = (1 - e)\mathbf{1}_{\{c = f_n^{\times 2}(\xi)\}} + e/a$$

on each child and its three parents on the preceding row. This is a finite-alphabet, radius-one, strictly positive rational PCA table normalized by $\sum_c p(c \mid \xi) = 1$. The interaction is TI, finite range and finite energy. Every finite-volume kernel contains all factors whose child-parent footprint meets the volume, including exterior children with an interior parent.

If $U(n)$ halts, add the two per-lane common-map bounds of Lemma 4.2. For every deterministic row η, finite I and time t, comparison to a stationary start gives

$$d_{\mathrm{TV}}((\eta P^t)_I, \mu_I) \le \sum_{x \in I} \sup_{\zeta} \Pr\{F_t(\eta)(x) \ne F_t(\zeta)(x)\} \le 2C|I| \exp(-ct).$$

More precisely the probability of nonconstancy bounds each supremum, so this inequality is uniform over the initial row, not merely an estimate for each fixed pair. A stationary row law exists by compactness and Cesàro averaging. This is uniform exponential ergodicity in the exact sense of [38, Theorem 3.1]. That theorem, for the normalized space-time specification of a finite-alphabet finite-radius strictly positive PCA on the line, gives a unique DLR measure among all boundaries and exponential covariance decay between every pair of fixed local space-time observations. All its hypotheses have therefore been checked for this joint table itself. No independence between the two lane estimates is needed. Denote this unique measure by γ_C; its constants may depend on the finite table.

If $U(n)$ never halts, the stationary trajectories in Proposition 4.4 are $\mathbb{Z}^2$-TI. We verify their DLR property directly. For a finite interior volume Λ and an exterior cylinder test, include every factor touching Λ, all sites of that test, and their parent closure down to a common earlier row. This is a finite cone. Given the bottom row its forward weight is the product of the PCA factors. All children of sites of Λ are included, including exterior children. When summing over the interior assignment, precisely the factors touching Λ vary; all other factors and the bottom law are independent of that assignment. Dividing by the sum of these weights is exactly the Gibbs kernel. This proves the integrated specification identity against every exterior cylinder. Exhaustion by cylinders and the monotone class theorem give the DLR conditional identity. This argument uses stationary kernels and all touching factors; it does not assume a converse identifying every DLR law with a trajectory.

The events B_d constructed in Proposition 4.4 are now ordinary Borel events of full plane configurations and have probability one. Average $(-1)^x f(x, t)$ over an l by l square. Before mixing translations, its expected value has one common sign and absolute value at least $1/2$ throughout the square: the spatial grid origin is constant in time. Jensen therefore gives the square second moment at least $1/4$ for every l. Its limit exists by the mean ergodic theorem for the two commuting translations of $\mathbb{Z}^2$ (space and time), applied to phase-twisted square averages. The coboundary proof in Section 2 applies verbatim to these two commuting unitaries.

The TI DLR set is compact convex metrizable. Indeed, positivity and finite range make each local kernel continuous, and its cylinder identities and TI equations are closed. Apply [39, Theorem 4.2] directly in this set. Probability one on B_d passes to almost every extreme component by the ordinary mixture identity on Borel sets. Bounded convergence of the finite square moments shows that some such component has a positive limit at $(\pi, 0)$. Distinct d give distinct components. There are at most $\mathfrak{c}$ measures by countable cylinder coordinates, proving exactly the asserted continuum of ordered extremes. □

Lemma 7.3 (Arbitrary Gibbs hosts)**.** *Take the product interaction of any rational-weight host with the controller of Lemma 7.2. On halting it preserves every coordinate of σ_G; on nonhalting its datum is* $(\mathfrak{c}, \mathfrak{c}, \mathfrak{c}, \text{false}, \text{true})$ *independently of the host.*

Proof. Multiply the local weights on separate host and controller tracks. The finite-volume specification kernels factor as $\Gamma_\Lambda = \gamma_\Lambda^H \otimes \gamma_\Lambda^C$ for each fixed joint boundary. Projecting the joint DLR identity, with exterior tests on one track, proves that every marginal is DLR. Products of separate DLR measures are joint DLR measures by the factorized kernels. Every host has a TI DLR measure: take finite-volume measures, use compactness with receding boundaries, pass the continuous local DLR identities to a limit, then average translations. The same reasoning gives compactness of all the relevant TI DLR sets.

On halting, the controller's unique all-DLR measure is γ_C. We must exclude correlated joint Gibbs laws. Let ρ be any joint DLR law, h a bounded host cylinder, k a controller-exterior cylinder for Λ, and f a controller cylinder. The joint kernel applied to $h\,k\,(f - \gamma_\Lambda^C f)$ is zero: its host factor is $\gamma_\Lambda^H h$, while its controller factor is $k\,\gamma_\Lambda^C(f - \gamma_\Lambda^C f) = 0$. This uses idempotence and properness of the specification, and holds even if h reads host sites inside Λ. The DLR identity gives

$$\rho[h\,k\,(f - \gamma_\Lambda^C f)] = 0.$$

Extend the host test by the monotone class theorem to every bounded function of the entire host configuration. Disintegrating over that configuration, and using a countable determining collection of cylinder identities, shows that the conditional controller law is DLR almost surely. It must equal γ_C. Thus every joint DLR measure is $\nu \otimes \gamma_C$ with ν host DLR. Projection is an affine homeomorphism also on TI DLR sets, and it preserves uniqueness among all DLR measures.

For order, expand any local joint observable as a finite sum $\sum_i a_i \otimes b_i$ in the finite tensor-product basis. Evaluate b_i in γ_C to define the local host observable $G(f) = \sum_i \gamma_C(b_i) a_i$. For two distant translates, the difference between their joint two-point moment and the host moment of their G-observables is a finite sum of terms bounded by host norms times controller covariances. Those tend exponentially to zero in every plane direction, by Lemma 7.2. First moments agree exactly. Thus covariance order is preserved. For weighted square second moments, separate pairs at bounded displacement from the remaining pairs. The first fraction tends to zero as the square size grows; the second has uniformly small covariance error when the displacement cutoff is large. This proves equality of the limiting Bragg moments at each nonzero wavevector. Host-only observables give both converse implications. The affine bijection preserves extremality and all three counts, and therefore L as well.

On nonhalting, product Gibbs measures give a TI joint lift of every extreme controller state with any fixed TI host Gibbs measure. Each fiber over such an extreme is a nonempty compact face. Its extreme points are extreme in the joint set; the successive coordinate maximization argument of Lemma 2.1 proves existence. Distinct controller marginals give distinct joint extremes with the inherited peak at $(\pi, 0)$. All three counts are bounded above by $\mathfrak{c}$, so they equal $\mathfrak{c}$. Uniqueness among all DLR measures fails, and the ordered states give L. This proves the second branch. □

Proof of Theorem 7.1. The maximal datum is realized by a fixed nonhalting controller. For a nonconstant predicate choose a finite rational-weight host with the opposite verdict and hold it fixed. Given n, output its product with the controller interaction. The factors are positive rationals and effectively enumerable. Lemma 7.3 gives the host datum on halting and the maximal datum on nonhalting, so the same two many-one reductions as in Theorem 3.3 apply. The argument depends only on realized values and survives collapsing infinite counts. □

8. Obstructions, finite rings and partial degrees

Theorem 8.1 (Necessary data on the line)**.** *In both line categories, E forces L false, without any uniform attraction assumption. Multiple TI stationary states force E false and L true. Always* $N_{\mathrm{Br}} \le N_{\mathrm{ord}} \le N$*; if* $N = 1$*, then* $N_{\mathrm{ord}} = 1$ *exactly when L holds. The total count belongs to* $\{1, 2, \dots, \aleph_0, \mathfrak{c}\}$*. Consequently the only possible coarse data* (N, E, L) *are*

$$(1, \text{true}, \text{false}), \quad (1, \text{false}, \text{false}), \quad (1, \text{false}, \text{true}), \quad (k, \text{false}, \text{true}),$$

where the last case has $k \in \{2, 3, \dots, \aleph_0, \mathfrak{c}\}$*. At (3.1), the first case and the last with* $k = \mathfrak{c}$ *are realized in both categories. No realization or exclusion of the other cases is established here.*

Proof. Nonemptiness and compactness follow from finite marginal coordinates, closed TI and stationarity equations, and Cesàro averaging. For distinct TI stationary states choose a real local, or Hermitian, observable with means $a \neq b$. In their half-mixture the variance of its normalized length-l spatial average is at least $(a-b)^2/4$ for every l. If all distant covariances decayed, a near/far split of its double sum would make that variance tend to zero: only $O(l)$ pairs have any fixed bounded displacement. Thus multiplicity gives L and precludes E. The Bragg/covariance implication proved in Lemma 2.2 gives the count inequalities. If $N = 1$, the compact convex TI stationary set is a singleton, so its only point is ordered exactly when L holds.

We prove that E forbids nondecaying covariance. In the classical case, let μ be the unique attracting law and P the transition operator on observables. For each deterministic initial row η, $P^t f(\eta) \to \mu(f)$. Bounded convergence gives convergence to this constant in $L^2(\mu)$. For distant enough translations the two t-step backward cones are disjoint. Independent innovations and stationarity imply

$$\mathrm{Cov}_\mu(f, \tau_z g) = \mathrm{Cov}_\mu(P^t f, \tau_z P^t g).$$

Cauchy–Schwarz bounds the right side by the product of the centered L^2 norms, uniformly in sufficiently large z. Letting $t \to \infty$ proves covariance decay without uniform convergence over starts.

For the quantum case write P for the Heisenberg round and μ for its unique stationary state, which is TI by uniqueness. Suppose a local covariance fails to decay and choose $|z_j| \to \infty$ along which it converges to a value different from the product of means. In the cyclic GNS representation (π, Ω) of μ, the Hilbert space is separable: the chain algebra has a countable norm-dense subset. By a diagonal subsequence the bounded operators $\pi(\tau_{z_j} b)$ converge weakly simultaneously for a countable dense set of b, hence for all b by norm approximation. Call the limit $Z(b)$. It belongs to the generated von Neumann algebra and commutes with every local operator, by eventual spatial disjointness and local approximation. Thus Z is positive, unital and linear into the center. Define

$$\Lambda(a, b) = \langle \Omega, \pi(a) Z(b) \Omega \rangle.$$

For sufficiently distant local a, b, independent backward circuits give $P(a\tau_z b) = P(a)\tau_z P(b)$. Stationarity and translation covariance, then the weak limit, imply $\Lambda(a, b) = \Lambda(Pa, Pb)$, and by iteration the same identity for every t. The chosen covariance means that $\Lambda(f, g)$ is not $\mu(f)\mu(g)$ for some local pair.

We now express Λ as an integral of products of two states. For finitely many Hermitian b, partition the spectra of their commuting central $Z(b)$ into intervals of width η, and let Q_i be the joint central spectral projections. For $p_i = \langle \Omega, Q_i \Omega \rangle > 0$ put

$$\alpha_i(a) = \frac{\langle \Omega, Q_i \pi(a) \Omega \rangle}{p_i}, \qquad \beta_i(b) = \frac{\langle \Omega, Q_i Z(b) \Omega \rangle}{p_i}.$$

Centrality and positivity make these positive unital functionals on the chain algebra. On each projection, $Z(b)$ differs from its conditional scalar $\beta_i(b)$ by norm at most η for the chosen b.

Consequently

$$\left|\Lambda(a,b) - \sum_i p_i \alpha_i(a)\beta_i(b)\right| \le \eta\|a\|.$$

Exhaust a countable dense Hermitian set and let $\eta \downarrow 0$. The product of the two compact metrizable chain state spaces has compact probability-measure space, so a subsequence of the laws $\sum_i p_i \delta_{(\alpha_i,\beta_i)}$ converges to a probability law θ. It gives $\Lambda(a,b) = \int \alpha(a)\beta(b)\, d\theta$ for every a, b. For each pair of states, pointwise attraction gives $\alpha(P^t a) \to \mu(a)$ and $\beta(P^t b) \to \mu(b)$. Their product is bounded by $\|a\|\|b\|$. Dominated convergence, together with invariance of Λ, forces $\Lambda(a,b) = \mu(a)\mu(b)$, a contradiction. Thus E implies L false in both categories.

The nonextreme points of a compact metrizable convex set form a countable union of compact sets of midpoints of pairs at distance at least $1/j$. Hence its extreme set is G_δ and is Polish. The perfect-set theorem gives the asserted possibilities for N, without a continuum-hypothesis assumption. The implications already proved give the displayed necessary coarse list. A constant classical table gives its one-step iid replacement law. A fixed on-site quantum preparation followed by qubit noise gives a product stationary state. Both have $(1, 0, 0, \text{true}, \text{false})$; a one-symbol classical alphabet is also allowed. A fixed never-halting controller gives σ_∞ by Proposition 4.4. These are the two claimed realizations. The Rice reductions range over the actual image and require no classification of the remaining possibilities. □

Remark 8.2 (Fixed finite rings). A classical ring of l sites has a complete reset with probability e^l each step. Its whole-state total-variation contraction coefficient is at most $1 - e^l$. If a finite quantum ring has k physical wires, simultaneous reset of every wire in the final layer has probability e^k and makes the full output trace-and-prepare. Its trace-distance contraction coefficient is at most $1 - e^k$. Iteration gives a unique stationary state attracting every initial state in either case. An unbounded-distance covariance label concerns the infinite line and has no such meaning on one fixed ring. These contraction bounds need not be uniform as the ring grows; growing storage horizons are consistent with fixed-ring uniqueness. At replacement probability one every input is reset; the theorems use the strictly smaller rate (3.1). A noisy identity also has the reset datum, since a fixed finite observation forgets its starting values after all its sites or wires have been reset.

Theorem 8.3 (Partial degree results on the line)**.** *Uniformly on valid presentations in either line category, for each $r \ge 1$ the cut $N \le r$ belongs to Π^0_2 and is halting-hard. For each $k \ge 2$ the cut $N \ge k$ belongs to Σ^0_2 and is nonhalting-hard. Matching second-level hardness on the line is not asserted.*

Proof. We give the effective finite-coordinate argument, including its hypotheses for the per-layer channel. Enumerate real local coordinates $x_1, x_2, \ldots$ bounded by one: cylinder probabilities classically, and real and imaginary density-matrix entries with Hermiticity built in quantumly. Use all finite intervals and their translates. Normalization, consistency, TI and stationarity are countably many effective linear constraints. Classical positivity is nonnegativity. Quantum positivity is the family $v^* \rho_I v \ge 0$ for every rational complex vector v on every finite interval I; density makes this equivalent to positive semidefiniteness. The backward circuit computing a local Heisenberg image includes every depolarizing layer. Rational e and algebraic gates give exact algebraic coefficients with decidable arithmetic and order. Thus all these linear constraints are effectively enumerable in both categories.

Let $b(m) \ge m$ increase to cover the coordinates used by the first m constraints. Let P_m be their nonempty bounded algebraic polytope in those coordinates with coordinate bounds $[-1, 1]$. All earlier constraints are retained. Every genuine stationary state projects into every P_m. Conversely a coordinatewise limit of outer points with $m \to \infty$ satisfies every constraint and is genuine. Product compactness implies that for each fixed n the first-n projections of P_m converge in Hausdorff distance to those of $\mathcal{K}$: otherwise distant outer points have a convergent subsequence contradicting that characterization.

Every vertex of P_m is computable: enumerate full-rank active constraint systems, solve, and check all inequalities. This also handles lower-dimensional polytopes, since at a vertex active normals must span the ambient space or a feasible segment remains. A bounded polytope is the convex hull of its vertices, by induction on face dimension using boundary-to-boundary segments. Exact algebraic linear feasibility is therefore decidable. Put $\eta = 2^{-n}$. If V is the number of vertices, choose an integer $L_{m,n} \geq 8V/\eta$ and let $\mathcal{Q}_{m,n}$ be all their convex combinations whose coefficients are multiples of $1/L_{m,n}$. This finite set is an $\eta/4$-net in the first n coordinates: round the first $V-1$ weights down and put the remainder on the last vertex; total weight error is at most $2(V-1)/L_{m,n}$.

For $m \geq n$, let $\mathsf{T}_r(n,m)$ mean that some r points of $\mathcal{Q}_{m,n}$, with repetitions, have a convex hull within η of every vertex in first-n maximum norm. For each candidate and vertex this is finite algebraic feasibility in simplex coefficients λ_i, with inequalities $|v_j - \sum_i \lambda_i y_{i,j}| \leq \eta$, $1 \leq j \leq n$. So T_r is decidable. Vertex coverage covers the polytope because the sum of a convex hull and a norm ball is convex. We claim

$$N \leq r \quad \Longleftrightarrow \quad \forall n \geq 1 \; \exists m \geq n : \mathsf{T}_r(n,m).$$

A compact convex set with finitely many extremes is their hull: if a point lay outside, a finite-coordinate separating functional would have a maximizing compact face disjoint from that hull; successive coordinate maximization there gives a further extreme, a contradiction. Hence if $N \leq r$, choose r generators, repeating as needed. For fixed n, choose m so the outer projection is within $\eta/4$ of $\mathcal{K}$, then approximate those generators by $\mathcal{Q}_{m,n}$ within $\eta/4$. Their hull covers the outer projection within $\eta/2$, proving the test.

Conversely choose successful stages $m_n \geq n$ and their r centers. A common coordinatewise convergent subsequence has limits $y_1, \ldots, y_r \in \mathcal{K}$, since every constraint eventually holds. For any $x \in \mathcal{K}$, vertex coverage provides simplex coefficients approximating its first n coordinates within 2^{-n}. A further subsequence of these coefficients converges, giving $x = \sum_i \lambda_i y_i$ in each coordinate. The limiting centers were chosen before x, so these same r points generate all of $\mathcal{K}$. There are at most r extremes. This proves the equivalence and the Π^0_2 membership; complementation gives the Σ^0_2 membership of $N \geq k$.

Use a reset host in Proposition 6.2. Its two counts are 1 on halting and $\mathfrak{c}$ on nonhalting. This reduces halting to every listed finite upper cut and nonhalting to every listed finite lower cut. Such a two-branch reduction gives no matching hardness at the second arithmetical level. The totality construction in Section 12 uses transverse slices and a different controller and is not a line construction. □

9. From Wang tilings to crystalline extreme states

9.1. **The imported controller.** The reduction starts from the domino problem [2, 3]. We use its effective unseeded form: from a machine index e, one computes a finite oriented Wang set R_e such that

$$R_e \text{ tiles } \mathbb{Z}^2 \quad \Longleftrightarrow \quad e \text{ never halts on blank input.} \tag{9.1}$$

One implementation is the forced variable-zoom macrotile construction of Durand, Romashchenko and Shen [4, Sections 2.2–2.3 and 5]. Its local coordinate, wire and verifier rules force successively longer computations in every plane tiling. A halt obstructs a sufficiently high level. Without a halt, compatible finite levels exist and compactness gives a plane tiling. This is an unseeded plane construction; its finite tables are computable from e.

Here is the exact dynamical theorem imported from our forgetting preprint. We record the rule to fix its noise representation. For a finite Wang set R, let $\mathcal{A}_R = R \sqcup \{\perp\}$. On $\mathbb{Z}^2 \times \mathbb{Z}^4$, first let a cardinal checker retain the central tile when the five letters at $\Gamma = \{0, \pm e_1, \pm e_2\}$ are

live and the four central shared edges match; otherwise it returns blank. Put

$$\mathcal{V}_i = v_i + \{0,\dots,15\}^4, \quad v_0 = 0, \quad v_i = 160e_i \ (1 \le i \le 4), \qquad \mathcal{V} = \bigcup_{i=0}^{4} \mathcal{V}_i.$$

The rule G_R at (u, P) returns a tile r if at least $4K$ of the $5K$ checked values at $(u, P + a)$, $a \in \mathcal{V}$, equal r, and returns blank otherwise. At most one tile qualifies. Its stencil is $\Gamma \times \mathcal{V}$, of size $25K$.

Theorem 9.1 (Imported Wang controller)**.** *Use K, B, ε_0 from (2.3), and let*

$$\kappa = 4B^{-99} < 1, \qquad c = \log(1/\kappa)/50.$$

For every finite R and every $0 < \varepsilon \le \varepsilon_0$, follow G_R by independent uniform whole-symbol replacement at rate ε. At each site and time, a replacement flag and a replacement value, together called the innovations, determine a map from parent words to output letters. Compose these maps using the same innovations for all initial rows, and put

$$a_t = \Pr\{x \mapsto X_t^x(0) \text{ is nonconstant over all initial rows}\}.$$

If R has no legal $n \times n$ pattern, then

$$a_t \le \frac{\exp\{c(762n + 1 - t)\}}{1 - \kappa} \quad (t \ge 0). \tag{9.2}$$

The bound also holds for replacement-value laws independent of the configuration and shared across starts. If R tiles, there are fully $\mathbb{Z}^6$-TI stationary laws μ_1, μ_0 with nonblank probabilities at least $1 - 2\varepsilon$ and at most 2ε, respectively.

This is the natural Wang classification theorem, its all-input coalescence lemma, and its persistent-separation argument with planar averaging, in [34, Theorem 2.1, Lemmas 2.10 and 2.13, and Section 2.20]. Non-tiling has a forbidden finite square by compactness. The source proves the common-map estimate by a flag-only contour bound and deterministic coherence and mortality. The estimate covers all starts on one event, including a pair selected after seeing the innovations. A bound for each separately fixed pair would not supply the argument below. A union bound gives common-map failure at most $|O|a_t$ on a finite set O.

9.2. **Pairing, slicing and a uniform parity conjugacy.** From now on the rate is ε_0. Pair two copies of G_R, followed by one uniform whole-pair replacement, and denote this six-dimensional medium by $G_R^{(2)}$. Its deterministic rule and noise commute with the on-site track swap S. The two marginal kernels are exactly the original controller kernels. If R tiles, start the pair with marginals μ_1, μ_0; both remain stationary under evolution. A TI Cesàro limit gives a stationary joining ξ_{10} with these marginals. Write $\xi_{01} = \mathsf{S}_*\xi_{10}$ and

$$\xi_p = (1 - p)\xi_{10} + p\xi_{01}, \qquad g = \mathbf{1}_{\{\text{first live}\}} - \mathbf{1}_{\{\text{second live}\}}, \qquad \delta_p = (1 - 2p)\delta, \tag{9.3}$$

where $0 < p < 1/2$ and $\delta = \xi_{10}(g) \ge 1 - 4\varepsilon_0 > 1/2$.

Let the rule ignore a seventh coordinate, and place independent identically distributed ξ_p whole slices along that coordinate. The law $\nu_p = \bigotimes_{j\in\mathbb{Z}} \xi_p$ is fully TI and stationary. Write $\mathrm{par}(x) = \sum_i x_i \pmod 2$, and let T apply S at odd-parity sites and the identity elsewhere. Write $F_{\mathrm{loc}} = G_R \times G_R$ for the local rule of the paired medium. Conjugating by T gives the *uniform* deterministic local rule

$$F'(y)(x) = F_{\mathrm{loc}}((\mathsf{S}^{\mathrm{par}(a)} y(x + a))_{a\in\mathcal{N}}). \tag{9.4}$$

Indeed $\mathrm{par}(x + a) = \mathrm{par}(x) + \mathrm{par}(a)$ and global swap equivariance cancels both factors depending on x. Uniform pair replacement is swap invariant. Thus the conjugated medium, denoted C_R, is a legal TI input with the same whole-symbol noise. The parity map is a way of deriving its finite table, rather than an inhomogeneity in the input.

If R does not tile, each original track has coefficient $a_m \to 0$. A union over tracks and conjugacy show that C_R has one-site all-input failure at most

$$\alpha_m = 2a_m \longrightarrow 0. \tag{9.5}$$

The slice restriction leaves the estimate unchanged even for initially correlated slices.

Proposition 9.2 (Crystalline amplification)**.** *If R tiles, C_R has continuum many distinct TI stationary laws that are extreme among all TI laws and have a Bragg atom at $q_\pi = (\pi, \dots, \pi)$. Each is an equal mixture of two distinct non-TI periodic stationary laws.*

Proof. Set

$$\rho_{p,0} = T_*\nu_p, \qquad \rho_{p,1} = T_*\mathsf{S}_*\nu_p, \qquad \rho_p = \tfrac{1}{2}(\rho_{p,0} + \rho_{p,1}).$$

Both components are stationary by conjugacy. Even-parity translations preserve each and odd-parity translations exchange them, because ν_p is TI and $\tau_x T = \mathsf{S}^{\mathrm{par}(x)} T \tau_x$. Their average is therefore fully TI.

Fix a bounded local real f and its ordinary cube average $\overline{f}_n$. Each component is a product over whole seventh-coordinate slices, with alternating slice laws. Two translates of f are independent if their finite sets of slices are disjoint. Among the $|B_n|^2$ pairs, only $O_f(|B_n|^2/n)$ can share a slice. Boundedness therefore gives $\mathrm{Var}_{\rho_{p,i}}(\overline{f}_n) = O_f(1/n)$. The component means of translates of f take two parity values, interchanged between $i = 0, 1$. The fractions of even and odd sites in a cube differ by at most $1/|B_n|$. Thus the two component cube means have the same limit. Variance decomposition yields

$$\mathrm{Var}_{\rho_p}(\overline{f}_n) \longrightarrow 0. \tag{9.6}$$

If $\rho_p = t\omega_0 + (1-t)\omega_1$, $0 < t < 1$, for TI laws, the variance is at least $t(1-t)(\omega_0(f) - \omega_1(f))^2$. Equation (9.6) forces equality on every local observable, so the decomposition is trivial. This proves extremality among all TI laws, hence among TI stationary laws.

With $\chi(x) = (-1)^{\sum_i x_i}$, the one-site means of g_x in the two components are $\chi(x)\delta_p$ and $-\chi(x)\delta_p$. Consequently the staggered average $\mathrm{Av}_n^{q_\pi}(g)$ has component means δ_p and $-\delta_p$, for every n, and

$$\rho_p\big((\mathrm{Av}_n^{q_\pi}(g))^2\big) \geq \delta_p^2 > 0.$$

Independence of distinct slices also gives

$$\rho_p(g_0) = 0, \qquad \rho_p(g_0 g_{je_7}) = (-1)^j \delta_p^2 \quad (j \neq 0).$$

In particular $\rho_p(g_0 g_{2e_7}) = \delta_p^2$ distinguishes all $p \in (0, 1/2)$. There are continuum many such laws, and Lemma 2.1 gives the matching upper bound. The two components have opposite nonzero g_0 means, so they are distinct and non-TI, as claimed. □

Classically, extremality among all TI laws is spatial ergodicity. In the quantum realization of Section 10, the same variance proof applies to every local Hermitian observable: the state is diagonal and a product over slices in each component. It is therefore extreme among all TI quantum states as well. The witnesses are layered and anisotropic. Their extremality is relative to spatial translation invariance; they are explicitly not extreme among all stationary states.

10. Preserving an arbitrary host

The controller must forget jointly with its graphical maps. Forgetting of its marginal law alone would leave open the possibility of stationary correlations between the host and controller. We prove the required joint statement for classical hosts and for local quantum circuits.

10.1. **Quantum presentations and a local implementation.** The classical controller has an exact implementation in the quantum category of Section 2. Represent its alphabet by an orthonormal basis. First dephase the old registers and introduce fresh scratch and output registers in fixed product states. For every read offset a, copy the basis value of main_{x+a} to its designated scratch register at x. For fixed a, these gates act on pairwise disjoint registers as x varies: every source and every destination scratch register is used once, and main and scratch registers are distinct. Thus one translation-invariant layer suffices for each offset. A controlled modular-addition gate performs the copy; initial dephasing makes this a copy of classical data. At each destination, compute the finite lookup table into its output register and discard the old and scratch registers. The resulting channel has exact algebraic entries and finite depth and range. It prepares basis-diagonal outputs, so its stationary quantum states are exactly the diagonal embeddings of the classical stationary laws.

10.2. **Shared replacement and finite-past comparison.** Let C denote the crystalline controller C_R, and let H be an arbitrary host in either input category. Form J by applying the two pre-noise updates separately and then one uniform replacement of the whole pair, at rate ε_0. Quantumly the reset branch is $\mathcal{R}_H \otimes \mathcal{R}_C$, where $\mathcal{R}_H(X) = \mathrm{Tr}(X)I_H/d_H$, and similarly for C. Partial trace shows that the two marginal updates are exactly those of H and C. Their replacement flags are shared.

On the non-tiling branch, the imported controller bound [34] and the two-track construction give a one-site all-input common-map nonconstancy bound $\alpha_m = 2a_m \to 0$. The event concerns all initial controller configurations simultaneously and depends only on the recent flags and replacement symbols. A union bound gives error $|O_C|\alpha_m$ on a finite controller observation region O_C.

For a host law or state μ, let $W_m(\mu)$ be the joint output after m rounds when the host starts from μ and the controller starts from a fixed constant basis configuration. The following bound holds for every joint initial law or state ρ, with host marginal ρ_H:

$$d_O(\rho P_J^m, W_m(\rho_H)) \le |O_C|\alpha_m. \tag{10.1}$$

Here d_O is total variation on a finite classical observation or trace distance $\frac{1}{2}\|\cdot\|_1$ on a finite quantum observation, and O_C is its controller support. Classically, retain the same host start and the same recent innovations, and replace only the controller start. The observed outputs agree on the all-input-good event, which proves the bound.

For the quantum proof, resolve the controller reset as

$$\mathcal{R}_C = \frac{1}{d_C}\sum_z \mathcal{R}_{C,z}, \qquad \mathcal{R}_{C,z}(X) = \mathrm{Tr}(X)|z\rangle\langle z|.$$

Condition on the finite graphical data ω in the backward observation cone: the shared replacement flags and the uniform controller replacement symbols. Branch by branch, the channel on the two complete tracks factors as $\Phi_{H,\omega} \otimes \Phi_{C,\omega}$. The host factor includes the specified host reset locations. The controller factor starts with complete dephasing and thereafter applies classical graphical maps. On the good event every input basis word gives the same observed controller output c_ω. Dephasing kills off-diagonal matrix units, so the reduced controller channel is the linear map

$$X \longmapsto \mathrm{Tr}(X)P_\omega, \qquad P_\omega = |c_\omega\rangle\langle c_\omega|.$$

For every joint input, including an entangled one,

$$\big(\Phi_{H,\omega} \otimes [\mathrm{Tr}(\cdot)P_\omega]\big)(\rho_{HC}) = \Phi_{H,\omega}(\rho_H) \otimes P_\omega.$$

Thus the actual output and the comparison output agree on each good branch. Their trace distance on a bad branch is at most one. Averaging proves (10.1). This is a channel identity on finite cones, so it applies to arbitrary lattice states.

Proposition 10.1 (Stationary preservation)**.** *On the non-tiling branch, host projection is an affine homeomorphism $\mathcal{K}(J) \to \mathcal{K}(H)$. It preserves E, and for each corresponding pair of states it preserves the local covariance-order label and b_q for every fixed nonzero wavevector q. Consequently $\sigma(J) = \sigma(H)$.*

Proof. Stationary correspondence. Every stationary host state has a stationary joint lift: start with that host marginal and a constant controller, and take a Cesàro limit. The host marginal remains fixed at every time. This also works without translation invariance; with a translation-invariant host and constant controller, the lift is translation invariant. If γ is any stationary lift of μ, (10.1) gives

$$d_O(\gamma, W_m(\mu)) \le |O_C|\alpha_m.$$

Any two lifts therefore agree on every finite observation as $m \to \infty$. Projection is a continuous affine bijection between compact stationary sets, hence a homeomorphism; its inverse is affine. In particular, it bijects their extreme points, including when the extreme sets are infinite.

Attraction of all initial states. Suppose first that the host has a unique stationary state μ attracting every initial state locally. The preceding argument, without translation invariance, gives its unique stationary joint lift γ. For any joint initial state ρ, apply (10.1) to the final m rounds before time $t \ge m$. The comparison output is $W_m(\rho_H P_H^{t-m})$. At fixed m, each finite restriction of W_m depends continuously on finitely many host coordinates, so this converges locally to $W_m(\mu)$ as $t \to \infty$. The latter lies within $|O_C|\alpha_m$ of γ. First take $t \to \infty$, then $m \to \infty$. This proves attraction to γ. Conversely, every host initial state has a joint extension, and every stationary host state has a stationary lift. Joint attraction and uniqueness therefore imply host attraction and uniqueness by projection. Hence $E(J) = E(H)$.

Covariance order for each state. Fix corresponding translation-invariant stationary states μ, γ, and real or Hermitian joint local observables f, g. Let r bound the number of controller sites in either support. The two-observation version of (10.1), together with the first- and second-moment bounds, gives, uniformly in x,

$$\left|\mathrm{Cov}_\gamma(f, \tau_x g) - \mathrm{Cov}_{W_m(\mu)}(f, \tau_x g)\right| \le 8\|f\|\|g\| r \alpha_m. \tag{10.2}$$

At fixed m, evaluating the recent innovations and the fixed initial controller gives local host observables $G_m^f = W_m^*(f)$ and $G_m^g = W_m^*(g)$. Classically these are conditional expectations given the initial host configuration. Distant backward cones have independent innovations, so

$$W_m(\mu)(f\tau_x g) = \mu(G_m^f \tau_x G_m^g).$$

Quantumly the two host observables are Hermitian. Because gates use independent product ancillas, disjoint backward circuits give the operator identity

$$W_m^*(f\tau_x g) = G_m^f \tau_x G_m^g.$$

This holds before evaluating the host state, including spatially entangled inputs. Equality of first moments therefore makes the comparison covariance $\mathrm{Cov}_\mu(G_m^f, \tau_x G_m^g)$. If μ lacks ℓ, this tends to zero for every fixed m. Taking separation to infinity in (10.2), then $m \to \infty$, excludes ℓ in γ. Conversely, witnessing host observables retain their moments in the joint state. Thus ℓ is preserved for every corresponding pair, and the affine bijection preserves N_{ord} and L.

Bragg order at a fixed wavevector. Fix $q \ne 0$ and a real or Hermitian f with controller-support size r; put $G_m = W_m^*(f)$. Expanding the weighted second moment gives products $(\tau_x f)(\tau_y f)$, whose coefficients all have modulus one. For fixed m, only $O_m(|B_n|)$ of the pairs $(x, y) \in B_n^2$ have intersecting backward cones. The factorization above is exact on every other pair. Boundedness on the remaining pairs gives

$$\left|W_m(\mu)(\mathrm{Av}_n^q(f)^* \,\mathrm{Av}_n^q(f)) - \mu(\mathrm{Av}_n^q(G_m)^* \,\mathrm{Av}_n^q(G_m))\right| \le \frac{c_{m,f}^{\mathrm{Br}}}{|B_n|}.$$

Applying (10.1) to the union of the two translated supports bounds the additional stationary/comparison moment error by $8\|f\|^2 r\alpha_m$, independently of n. If μ lacks b_q, the host weighted moment tends to zero for each fixed G_m. First send $n \to \infty$, then $m \to \infty$, to exclude b_q in γ. Presence passes to the joint state by using the witnessing host observable. Thus b_q is preserved for every fixed q, and the corresponding states have identical sets of nonzero Bragg wavevectors. The bijection of extreme points therefore also preserves N_{Br}. □

11. Rice separation and invariance

We now combine the two controller branches. Couple an arbitrary host H to C_R by taking their deterministic product rule, or tensor product of pre-noise circuits, followed by one whole-symbol or whole-site replacement layer. Denote the result by $J(H, R)$. Each marginal has its exact original kernel. In particular this is a legal input in the corresponding category.

Proposition 11.1 (The two branch values)**.** *For every host H,*

$$\sigma(J(H,R)) = \begin{cases} \sigma(H), & R \text{ does not tile}, \\ \sigma_\infty, & R \text{ tiles}. \end{cases}$$

The construction is total computable from the finite presentations.

Proof. For a non-tiling R, (9.5) and Proposition 10.1 give the first branch. If R tiles, Proposition 9.2 gives continuum many crystalline extreme controller states. Each has an extreme joint lift by Lemma 2.1, and its controller observables retain their Bragg moments in the lift. Distinct marginals give distinct lifts. Thus $N_{\mathrm{Br}} \geq \mathfrak{c}$. Since $N_{\mathrm{Br}} \leq N_{\mathrm{ord}} \leq N \leq \mathfrak{c}$, by Lemma 2.2 all three counts are exactly $\mathfrak{c}$. Multiple stationary states make E false and L true. Arbitrary-host extreme lifts need not be spatially ergodic; extremality within the TI stationary set is sufficient here.

Classically, the controller alphabet has $(|R|+1)^2$ letters. The product has $Q = |A_H|(|R|+1)^2$ letters and a stencil of size $n \leq |\mathcal{N}_H|+25K$, so its table has Q^n entries. Enumerating this finite table is an algorithm. Quantumly, dephasing, basis copying, finite lookup and preparation give an exact finite controller circuit; tensoring with the host and appending depolarization are effective finite operations. The rate and dimension remain (2.3). □

Proof of Theorem B. A singleton Wang tile with matching opposite edge colors tiles the plane, so its controller realizes σ_∞. Let P be a nontrivial predicate of the datum on the chosen exact input category. There is a finite valid host H whose verdict differs from $P(\sigma_\infty)$; otherwise the predicate would be constant. Fix this witness as a finite constant of the reduction.

Given e, compute R_e from (9.1) and output $J(H, R_e)$. By Proposition 11.1, halting gives datum $\sigma(H)$ and nonhalting gives σ_∞. If $P(\sigma_\infty)$ is false, e halts exactly when the output satisfies P. If $P(\sigma_\infty)$ is true, e never halts exactly when it satisfies P. These are the asserted many-one reductions. The reduction depends on P only through the fixed witness H. Invariance is proved below. □

Proof of Theorem A. Every predicate of (N, E, L) is also a predicate of the refined datum. Nontriviality on classical inputs is unchanged. The classical part of Theorem B applies. □

Proposition 11.2 (Bath and coordinate invariance)**.** *On-site relabeling or unitary conjugation, and adjoining an independent constant bath, preserve every coordinate of σ.*

Proof. A uniform on-site relabeling or unitary commutes with translations. It transports stationary states, extreme points and local observables bijectively, and preserves local convergence, covariance and weighted second moments. Depolarization is unitarily covariant.

For an independent constant bath let ω_{B} be its fixed iid output after one round, including its own noise. On arbitrary correlated or entangled inputs, the one-round product dynamics has output

$$(\rho_H P_H) \otimes \omega_{\mathrm{B}}.$$

Indeed the bath channel is trace-and-prepare. Every TI stationary joint state is therefore $\mu \otimes \omega_{\mathrm{B}}$, uniquely for $\mu \in \mathcal{K}(H)$. This is an affine homeomorphism; it preserves N and, by the same one-round identity, E.

For local joint observables f, g, evaluate the bath in ω_{B} to obtain local host observables G^f, G^g. When two bath supports are disjoint,

$$(\mu \otimes \omega_{\mathrm{B}})(f\tau_x g) = \mu(G^f \tau_x G^g), \qquad (\mu \otimes \omega_{\mathrm{B}})(f) = \mu(G^f).$$

Thus ℓ is preserved for each corresponding state. Expanding a Bragg moment over all pairs gives the same identity except for $O(|B_n|)$ pairs with overlapping supports. After division by $|B_n|^2$ their bounded contribution tends to zero. Every fixed-q Bragg label is therefore preserved. Host observables give both converse implications. This preserves $N_{\mathrm{ord}}, N_{\mathrm{Br}}$ and L. □

A shared-flag constant bath also preserves the datum inside the original input category: its all-input coefficient is zero after one round, so Proposition 10.1 applies. An independent bath product generally has two flags and belongs to the comparison universe specified in Section 2.

Corollary 11.3. *In both input categories, TI stationary uniqueness, all-law ergodicity, L, existence of an ordered or Bragg-ordered extreme TI stationary state, and every positive finite lower threshold on N_{ord} or N_{Br} are undecidable. For total N, every lower threshold $N \geq k$, $k \geq 2$, and the infinite and uncountable count conditions are undecidable. Infinite and uncountable ordered or Bragg-ordered extreme counts are undecidable as well.*

Proof. Constant reset has datum $(1, 0, 0, \mathrm{true}, \mathrm{false})$; its unique iid stationary state has decaying covariance and no nonzero Bragg atom. The self-matching-tile controller has σ_∞. These witnesses establish nontriviality of every listed predicate. Theorem B gives the reductions. With reset as host, uniqueness, E and finite upper total-count cuts are halting-hard; order existence and the lower, infinite and uncountable counts are nonhalting-hard. Complements reverse these orientations. □

For classical inputs, the statements about uniqueness, E and the total count N already follow from [34, Corollary 3.11].

For each $k \geq 1$, existence of at least k spatially ergodic stationary states with a Bragg atom at some nonzero wavevector is also nonhalting-hard by using C_{R_e} alone: its non-tiling branch has $N_{\mathrm{Br}} = 0$ by Proposition 10.1 with a one-point host, and its tiling branch has the witnesses of Proposition 9.2. Quantumly, spatial ergodicity here means extremality among all TI states. This separate reduction needs no spatially ergodic lifting through an arbitrary host.

12. Phase-count completeness from totality

The halting reduction preserves all-law ergodicity because its controller forgets every start. For total-count completeness, stationary common-noise couplings suffice. This section uses the input categories of Section 2 with the rate replaced by

$$\varepsilon_1 = 160000^{-101}. \tag{12.1}$$

12.1. **The source condition.** *Condition C_T* is the validity of the repaired Törmä-based totality source with the following interface, using [23, Lemma 25, Corollary 26 and Lemmas 27–29] as published. There is a total computable family of one-dimensional radius-one CA U_e, with finite alphabet D_e, blank 0, and local table $u_e : \mathsf{D}_e^3 \to \mathsf{D}_e$, such that:

(i) $u_e(0, 0, 0) = 0$. If the partial computable function φ_e is total, every fully $\mathbb{Z}^2$-TI law on the two-sided input-erasure trajectory SFT is all blank.

(ii) If φ_e is not total, the ordinary two-sided trajectory SFT has a fully $\mathbb{Z}^2$-TI law with positive nonblank activity.

Here input erasure means replacing selected parent symbols by 0 from one common preceding row. The repaired source reserves zero for each bookkeeping field's blank value, normalizes invalid own bookkeeping to blank, leaves simulated data unrestricted, and increments age only while live; field widths are fixed before self-reference. Quiescence of the universal layer is part of this source premise. The interface is stronger than deterministic unique ergodicity. The cited published lemmas supply the macrocell, overlap, rigidity, connecting and live-coverage ingredients; their internal proofs are used as published.

12.2. **A controller with diagonal stationary couplings.** A canonical common-noise pair uses the same replacement flag and replacement letter in both replicas at each site and time. Say a controller has the *diagonal-coupling property* if every TI stationary law of this pair kernel is supported on equal replicas.

Lemma 12.1. *Under C_T, one can compute from e a deterministic rule G_e on $\mathbb{Z}^6$, of maximum-norm radius at most* 11*, followed by uniform replacement at ε_1, such that totality gives the diagonal-coupling property and nontotality gives two distinct TI stationary laws.*

Proof. Certificates and quorum. Put

$$\mathsf{L}_e = \{0\} \sqcup \{(a,w) : a \neq 0,\ w \in \mathsf{D}_e^3,\ u_e(w) = a\}, \qquad \mathrm{sym}(0) = 0, \quad \mathrm{sym}(a,w) = a.$$

On $u = (i,s) \in \mathbb{Z}^2$, let $H_{0,e}$ retain a nonblank label (a,w) exactly when every nonzero w_j agrees with the symbol at $(i+j, s-1)$, $j = -1, 0, 1$; otherwise delete it. Blank centers remain blank. A fixed field projects to an input-erasure trajectory, using its stored witness at a nonblank center and erasing all parents at a blank one. Conversely, an ordinary trajectory lifts by storing its actual predecessor word.

Physical sites are $(u,P) \in \mathbb{Z}^2 \times \mathbb{Z}^4$. Let H_e require, in addition, equality of the full label with all eight transverse unit neighbors. Its stencil is

$$\mathcal{N}_* = \big(\{0, (-1,-1), (0,-1), (1,-1)\} \times \{0\}\big) \cup \{(0, \pm e_j) : 1 \le j \le 4\}.$$

Use providers at $c_0 = 0$, $c_j = 10e_j$, $1 \le j \le 4$. The rule G_e returns a nonblank label if four of the five provider values of H_e equal it, and returns 0 otherwise. The 60 read offsets are distinct and have radius at most 11. Since $|\mathsf{L}_e| \le 1 + |\mathsf{D}_e|^3$, its finite table is computable from the source table.

Noise marks. Write $\mathsf{Flag}_t(v)$ for the iid replacement flags. Mark a site if it is flagged or if two distinct providers each have a marked $\mathcal{N}_*$-parent. Empty starts receding to the past give the least two-sided mark field E^*, a stationary causal factor of the flags. Every mark has a finite support graph ending at distinct physical flags.

Here are the flag estimates used below. The backward graph has undirected degree at most 120; use $d_* = 200$. The five potentials

$$\psi_j(P,t) = -P_j - 2t \quad (1 \le j \le 4), \qquad \psi_0(P,t) = \sum_{j=1}^{4} P_j + 8t$$

sum to zero. Provider j is excluded for ψ_j, $0 \le j \le 4$. Along a backward provider step the gain is at least 1 except at the excluded provider, and every gain is at least -9. Two distinct providers therefore offer a good step for each potential.

To count a finite support graph while reusing terminals correctly, group frontier vertices by overlap of their descendant-terminal sets. Exposing the two supports of one vertex can split its component into at most two and cannot merge different components. Start with one pole at the root, and keep one terminal endpoint of each potential in each component. At a split add a pole, sending its arm of each potential to the daughter lacking that endpoint. Only that first edge may be bad. In the original root's arm component, n_{f} distinct terminal flags give n_{f} poles and at most $5(n_{\mathrm{f}} - 1)$ bad edges. For each potential the arms biject these poles with the terminal flags. The sum of all arm gains is zero, since the potentials balance at each pole

and terminal. Telescoping bounds the total number of arm edges by $50(n_{\mathrm{f}}-1)$: replacing a gain of at least 1 by one of at least -9 loses at most 10. Physical projection cannot increase edge count or time span.

A rooted connected graph with a edges has at most d_*^{2a} doubled-walk encodings and 2^{a+1} choices of terminal subset. Distinct flags cost $\varepsilon_1^{n_{\mathrm{f}}}$, even when several arms reach the same flag. Thus, with

$$\lambda = 2d_*^2\varepsilon_1^{1/50} = \tfrac{1}{2}\,160000^{-51/50} < \tfrac{1}{2},$$

the mark density satisfies

$$p_* := \Pr(E^*(0)=1) \le \varepsilon_1 + 2\varepsilon_1 \sum_{a\ge 1} \lambda^a < 3\varepsilon_1 < \tfrac{1}{2}. \tag{12.2}$$

For a marked nonflagged backward path of length n, take complete support graphs in successive disjoint time slabs, restarting below the earliest terminal of each and keeping the last graph complete. Their distinct n_{f} flags cover $n \le 50n_{\mathrm{f}}$ path steps. The connected union has $n \le a \le 100n_{\mathrm{f}}$ edges. The same counting gives

$$\Pr(\text{such a path from a fixed root}) \le \frac{2\gamma_*^n}{1-\gamma_*}, \qquad \gamma_* = 2d_*^2\varepsilon_1^{1/100} = \tfrac{1}{2}\,160000^{-1/100} < \tfrac{1}{2}. \tag{12.3}$$

Consequently infinite marked nonflagged backward paths have probability zero, simultaneously at all countably many roots.

Clean extraction. Call an unmarked site clean. Two clean nonblank labels at the same u, t agree, even at different transverse positions. Indeed each has four supporters, so their supporter sets intersect in at least three indices. Each clean root has at most one dirty provider; after excluding both dirty indices a common wholly clean supporter remains. Its transverse equality checks move one label one unit toward the other at the preceding physical time. Iterating decreases their transverse ℓ^1 distance until they occupy the same site, proving equality of the full labels.

Augment any TI stationary noisy law by its two-sided iid innovations, using consistent finite time windows, and append E^*. Let $y_t(u)$ be the unique nonblank label found at a clean transverse position, or 0 if none exists. A clean live witness has a wholly clean successful provider, which supplies its same old center and every positive parent prescribed by its witness. Thus $y_t \le y_{t-1}$ in deletion order. Time stationarity makes their nonblank probabilities equal, so they agree almost surely. Countability makes the extracted field time constant and $H_{0,e}$-fixed, with a fully $\mathbb{Z}^2$-TI law. On totality, source clause (i) makes it blank. Every nonblank vertex in every stationary TI noisy margin is therefore marked.

Now take an arbitrary TI stationary canonical common-noise pair and augment both replicas by the same innovations and marks. A discrepancy is unflagged, since shared replacement gives equal letters. Distinct deterministic outputs require a discrepant parent. Each discrepancy has a nonblank margin, hence is marked. Tracing parents into the two-sided past would give an infinite marked nonflagged path, excluded by (12.3). The pair is diagonal. This is the arbitrary-pair quantifier in the usual stationary uniqueness argument made explicit.

Separation on nontotality. Source clause (ii) gives a positive-activity $H_{0,e}$-fixed reference law and also the blank reference law. Clone either reference constantly in P. It is fixed by G_e. Start the noisy field at its reference and the mark field in its stationary least law, independently of the reference. The closed relation

$$X_t(v) \ne \mathrm{Ref}(v) \implies E_t(v) = 1$$

is preserved: at a clean site four wholly clean providers have the correct reference inputs, including when the reference is blank. Fixed-marginal Cesàro limits give stationary TI noisy lifts with this error domination.

For fixed u, the transverse mark field is mixing: approximate its indicator in L^1 by a finite-past function of finitely many flags, whose distant translates are independent. Its transverse cube averages have variance tending to zero and mean $p_* < 1/2$. Let δ_R be the probability that at least half the sites of the radius-R transverse cube are marked; then $\delta_R \to 0$. Strict majority decodes $\mathrm{Ref}(u)$ except on this event. If the active reference has nonblank probability $a_* > 0$, the two noisy lifts have local decoded nonblank probabilities at least $a_* - \delta_R$ and at most δ_R. Some finite R separates them. This observable is needed only to prove distinctness; the compiler uses the finite source table. □

12.3. **The simultaneous count reduction.** One reduction family establishes hardness for every finite upper and lower threshold.

Theorem C (Phase-count completeness under C_T)**.** *Assume C_T. At $D = 7$ and the single rate $\varepsilon_1 = 160000^{-101}$, in each input category separately, $\{M : N(M) \le r\}$ is Π^0_2-complete for every finite $r \ge 1$, and $\{M : N(M) \ge k\}$ is Σ^0_2-complete for every finite $k \ge 2$, under computable many-one reductions. One computable family has $N = 1$ on total programs and $N = \mathfrak{c}$ on nontotal programs.*

Proof. Run G_e independently in each seventh-coordinate slice, with the same rate ε_1, and call the result C_e. On totality, a TI stationary common-noise pair restricts on each slice to such a pair for G_e, hence to equal replicas by Lemma 12.1. Countably many slices make the full pair diagonal. Given any two TI stationary laws, evolve their product using common noise and take a Cesàro limit. The marginals stay fixed and the limit is a stationary TI pair. Diagonality equates its marginals. Existence follows from Lemma 2.1, so $N(C_e) = 1$.

On nontotality choose two distinct stationary TI slice laws μ_0, μ_1, and a bounded real local h separating their means. For $0 \le p \le 1$, put

$$\mu_p = (1-p)\mu_0 + p\mu_1, \qquad \nu_p = \bigotimes_{j \in \mathbb{Z}} \mu_p.$$

These laws are stationary and fully TI. Translates of any local observable far enough apart along e_7 depend on disjoint sets of independent slices. Their Cesàro average has variance tending to zero. A decomposition of ν_p into two TI laws with distinct local means would give a positive lower bound on that variance. Thus ν_p is extreme among all TI laws. The h-means distinguish continuum many p, and Lemma 2.1 bounds their number by $\mathfrak{c}$. Hence

$$\varphi_e \text{ total} \implies N(C_e) = 1, \qquad \varphi_e \text{ nontotal} \implies N(C_e) = \mathfrak{c}. \tag{12.4}$$

The classical range stays at most 11; the finite alphabet may depend on e.

The set $\mathrm{TOT} = \{e : \forall n\ \exists s\ [\varphi_e(n) \text{ halts by } s]\}$ is Π^0_2-complete. For a decidable relation $R(a, n, t)$, compile a program on input n that searches for a t satisfying $R(a, n, t)$. It is total exactly when $\forall n\, \exists t\, R(a, n, t)$. Complementation gives Σ^0_2-completeness of $\overline{\mathrm{TOT}}$. Equation (12.4) reduces TOT to every finite upper cut $N \le r$, and its complement to every lower cut $N \ge k$, $k \ge 2$. The unconditional upper bounds of Theorem A.1 give the matching membership statements.

For quantum inputs, implement C_e by the dephase/copy/lookup/prepare circuit of Section 10, followed by whole-site depolarization at ε_1. Every finite output block is basis diagonal after one round. Thus every stationary quantum state is diagonal, and its stationary TI set is affinely exactly the classical one. Equation (12.4) holds for the entire TI stationary quantum set. Theorem A.1 also gives quantum membership. On the quantum presentation promise, membership means a total arithmetic formula agrees on valid inputs; every reduction here outputs a valid exact circuit. □

12.4. **Stationary preservation from diagonal couplings.** For the general stationary Rice clause, discard the E coordinate:

$$\sigma_{\mathrm{stat}} = (N, N_{\mathrm{ord}}, N_{\mathrm{Br}}, L), \qquad \sigma_{\mathrm{stat},\infty} = (\mathfrak{c}, \mathfrak{c}, \mathfrak{c}, \mathrm{true}).$$

Proposition 12.2. *At any common rate, coupling a host to a deterministic TI controller followed by uniform replacement and having the diagonal-coupling property, using one whole-pair replacement, preserves σ_{stat}, in both categories. Host projection is an affine homeomorphism of TI stationary sets and preserves ℓ and every b_q per state.*

Proof. The controller has a unique TI stationary law ϖ, by the fixed-marginal Cesàro argument above. Start its common-noise pair from $(\varpi, \mathbf{0})$, and let d_m be the one-site disagreement after m rounds. Every subsequential Cesàro pair limit is TI stationary and diagonal. Since disagreement is a cylinder event,

$$\frac{1}{T}\sum_{m<T} d_m \longrightarrow 0.$$

Otherwise a compactness limit would have positive disagreement. Choose an unbounded subsequence m_j with $d_{m_j} \to 0$.

Every TI stationary host state μ has a stationary joint lift. In any such lift ζ, the controller marginal is ϖ. Retain its host start and future innovations, but reset the comparison controller to $\mathbf{0}$. The controller pair has exactly the law started from $(\varpi, \mathbf{0})$, regardless of initial host correlations. For the finite-past comparison $W_m(\mu)$ this gives

$$d_O(\zeta, W_m(\mu)) \le |O_C| d_m. \tag{12.5}$$

Along m_j, all stationary lifts of μ agree. The continuous affine onto projection is therefore an affine homeomorphism and bijects the full extreme sets.

Quantumly, a stationary joint state is diagonal on the measure-and-prepare controller. On a finite input cone write it as $\sum_c \rho_H^c \otimes |c\rangle\langle c|$, with positive subnormalized host matrices and traces given by the ϖ-marginal. At fixed innovations the channel factors between tracks. The actual and reset outputs have the same host images of ρ_H^c; their normalized trace distance is at most the total weight of strings whose controller outputs differ on O_C. Averaging gives exactly the classical mismatch bound $|O_C| d_m$. This proves (12.5) also with initially entangled host data.

For joint observables f, g, let r bound the number of controller sites in either support. The covariance error in (10.2) is now at most $8\|f\|\|g\| r d_m$. At fixed m, distant finite cones factor through local host observables G_m^f, G_m^g; quantumly the factorization follows from the independent product environments. If the host state lacks ℓ, send separation to infinity and then $m = m_j \to \infty$ to exclude it in the lift. Host observables give the converse. For each $q \neq 0$, expand the weighted second moment: only $O_m(|B_n|)$ pairs have intersecting cones, and the remaining products are host products. The stationary/reset error is bounded uniformly in n by $8\|f\|^2 r d_m$. First $n \to \infty$, then $m = m_j \to \infty$, preserves each b_q, exactly as in Proposition 10.1. Thus $N_{\mathrm{ord}}, N_{\mathrm{Br}}$ and L are preserved as well. □

This proposition controls stationary controller marginals; its conclusion concerns σ_{stat}, excluding all-law ergodicity. A constant bath sharing the replacement flags is a trivial diagonal-coupling controller: replicas agree after one step. Its invariance follows from this proposition. The product lift $\mu \otimes \omega_{\mathrm{B}}$ in Proposition 11.2 refers to a bath with independent noise in the comparison universe.

12.5. **The stationary Rice degrees.**

Theorem 12.3 (Stationary Rice degrees under C_T)**.** *Assume C_T, $D = 7$ and rate ε_1. In each input category, every nontrivial predicate P of σ_{stat} is Π_2^0-hard if it is false at $\sigma_{\mathrm{stat},\infty}$, and Σ_2^0-hard if it is true there. There is no nontrivial Δ_2^0 predicate of this datum. If such a predicate separately belongs to Π_2^0, or to Σ_2^0, it is complete for that class.*

Proof. Pair two copies of G_e, slice in the seventh direction, and apply the parity conjugacy (9.4), all at ε_1. Denote the resulting controller by Q_e. On totality the diagonal-coupling property survives slicing and pairing, by restriction to slices and to corresponding registers of the two

replicas. For parity conjugacy, pull back a fully TI stationary pair through T in both replicas. The pulled-back law is stationary, invariant under even translations, and each odd translation acts on it by the global swap in both replicas. Averaging it with its swap image makes it fully TI, hence diagonal by the ungauged property. Each component of this positive average is diagonal, so the original pair is diagonal. Proposition 12.2 therefore preserves every host's σ_{stat}.

On nontotality, let h be any real local separator of the two stationary slice laws, with mean gap $\delta > 0$. Fixed-marginal Cesàro averaging gives a paired joining ξ_{10}; let ξ_{01} be its swap. Set

$$\xi_p = (1-p)\xi_{10} + p\xi_{01}, \quad g = h(\text{first}) - h(\text{second}), \quad \delta_p = (1-2p)\delta, \quad 0 < p < \tfrac{1}{2}.$$

Use iid ξ_p slices and average the two parity-conjugate components as in Proposition 9.2. For every local observable the component cube-average variance is $O(1/n)$, because only $O(|B_n|^2/n)$ pairs share a slice. Their parity means average to the same limit. Mixture variance therefore proves that the TI average ρ_p is extreme among all TI laws, and among all TI quantum states in its diagonal realization.

Here g may read several sites inside one slice. Use the local observable $f = T^*g$, anchored at the origin. Swap oddness gives

$$T^*\tau_x(T^*g) = (-1)^{\sum_i x_i}\tau_x g.$$

The staggered component means are δ_p and $-\delta_p$, so the q_π Bragg moment is at least δ_p^2. Moreover $\rho_p(f\tau_{2e_7}f) = \delta_p^2$; independence of the two whole slices proves this identity and distinguishes continuum many p. Thus Q_e has continuum many Bragg-ordered extreme TI stationary states.

Write $H \boxdot Q_e$ for the product of the host and controller updates followed by one whole-pair replacement. For every host H, projection of its stationary set onto the controller is onto. The fiber over each extreme ρ_p is a compact face and has an extreme joint lift by Lemma 2.1. The lift retains the controller Bragg moment. By Lemma 2.2 and the continuum upper bound, every nontotal product has datum $\sigma_{\mathrm{stat},\infty}$. Total products have datum $\sigma_{\mathrm{stat}}(H)$. The quantum assertions follow from the same diagonal controller and the trace-distance argument just proved.

A known nontotal program realizes $\sigma_{\mathrm{stat},\infty}$. Nontriviality supplies a finite valid host H with the opposite verdict; fix it. The computable map $e \mapsto H \boxdot Q_e$ reduces TOT to P when its universal verdict is false, and $\overline{\mathrm{TOT}}$ to P when it is true. Only finite tables and circuits enter this map.

For completeness, hierarchy strictness forces the matching orientations. The relation

$$V(e,x) \iff \forall n\, \exists s\, [\text{program } e \text{ on } (x,n) \text{ halts by } s]$$

is universal for Π_2^0. Its diagonal complement $\neg V(x,x)$ is Σ_2^0 and cannot be Π_2^0: an index d for such a definition would give $\neg V(d,d) = V(d,d)$. Hence $\mathrm{TOT} \notin \Sigma_2^0$ and $\overline{\mathrm{TOT}} \notin \Pi_2^0$. A Π_2^0 predicate with true universal verdict would place $\overline{\mathrm{TOT}}$ in Π_2^0, a contradiction. It must have false verdict and is Π_2^0-complete. The dual statement holds for Σ_2^0; belonging to both classes is impossible for a nontrivial predicate. On the quantum promise domain these preimages are ordinary arithmetic predicates because every output presentation is valid. □

The concrete complete predicates established here are the total-count cuts of Theorem C. For $N_{\mathrm{ord}}, N_{\mathrm{Br}}$ and L, the general clause supplies one-sided hardness and the absence of nontrivial Δ_2^0 members; an arithmetic upper bound would be needed to invoke its matching-completeness clause. Exact $N = k$, $k \geq 2$, is Π_2^0-hard if realized and has the unconditional $D_2(\Sigma_2^0)$ upper bound, while realizability remains open. Finite and infinite total count retain their Σ_3^0 and Π_3^0 upper bounds.

13. A decidable stationary entropy boundary

The next criterion explains why a Rice theorem cannot encompass every stationary property. Entropy methods for noisy CA have a broader history in [24, 25].

Theorem 13.1 (Classical entropy floor)**.** *In any dimension $D \geq 1$, for a Q-symbol deterministic finite-range TI rule followed by uniform replacement at $0 < \varepsilon < 1$, put*

$$\beta = \varepsilon/Q, \qquad \alpha = 1 - \varepsilon + \beta, \qquad h_{\mathrm{bath}} = -\alpha \log \alpha - (Q-1)\beta \log \beta. \tag{13.1}$$

Some TI stationary law has spatial Shannon entropy density h_{bath} if and only if the deterministic rule is constant. For a nonconstant rule every such law has an explicit strictly positive entropy-density surplus. Consequently entropy-floor attainment is decidable by a finite table scan.

Proof. For a TI law μ, subadditivity of finite Shannon entropy proves existence of

$$h(\mu) = \lim_n \frac{H_\mu(X_{B_n})}{|B_n|} = \inf_n \frac{H_\mu(X_{B_n})}{|B_n|}.$$

To verify this, tile a large cube by fixed k-cubes plus a remainder of vanishing density. Subadditivity and TI bound the limsup by the k-cube ratio; the liminf is at least the infimum of all ratios. The remainder costs at most $\log Q$ per site.

Every replacement row $r_a = \beta \mathbf{1} + (1-\varepsilon)e_a$ has entropy h_{bath}. Conditional on the pre-noise field, outputs are independent rows of this form. Removing conditioning gives $H(X_{B_n}) \geq |B_n| h_{\mathrm{bath}}$ after one round, so every stationary entropy density is at least the floor. A constant rule sends every initial law in one round to an iid row law and attains it. The one-symbol and empty-stencil cases are constant and immediate.

Suppose the rule on the nonempty stencil $\mathcal{N}$ is nonconstant. Choose input words giving different output symbols a, b, and set

$$\upsilon = \beta^{|\mathcal{N}|}, \qquad \Delta_h = 2\upsilon \left[H\left(\frac{r_a + r_b}{2} \right) - h_{\mathrm{bath}} \right] > 0.$$

Strict concavity gives positivity. Any output-row mixture assigning weights at least υ to r_a and r_b has entropy at least $h_{\mathrm{bath}} + \Delta_h$: split off weight 2υ with the displayed half-mixture, and apply concavity to the remaining row mixture. This also covers $2\upsilon = 1$.

Start from any stationary TI law and run two rounds $X^0 \to X^1 \to Y$. Condition on the entire field $Z = F(X^0)$ before the first replacement layer. Given Z, the X^1 sites are independent and each symbol has probability at least β. Every chosen $\mathcal{N}$-word therefore has conditional probability at least υ. The preceding strict-concavity estimate gives

$$H(Y_v \mid Z) \geq h_{\mathrm{bath}} + \Delta_h.$$

Choose an integer s larger than every coordinate span of $\mathcal{N}$. For $\Lambda_{\mathrm{sparse}} = \Lambda \cap s\mathbb{Z}^D$, the input neighborhoods are disjoint, so the Y_v, $v \in \Lambda_{\mathrm{sparse}}$, are conditionally independent given Z. Use these sites first in the entropy chain rule. For the other sites condition further on X^1; their fresh final noises remain independent of $Y_{\Lambda_{\mathrm{sparse}}}$ and give exactly h_{bath} per site. Thus

$$\begin{aligned} H(Y_\Lambda) &\geq H(Y_\Lambda \mid Z) \\ &\geq |\Lambda_{\mathrm{sparse}}|(h_{\mathrm{bath}} + \Delta_h) + |\Lambda \setminus \Lambda_{\mathrm{sparse}}| h_{\mathrm{bath}}. \end{aligned}$$

Stationarity and cube limits imply

$$h(\mu) \geq h_{\mathrm{bath}} + \Delta_h / s^D > h_{\mathrm{bath}}.$$

This holds uniformly over all TI stationary laws of the given nonconstant rule. The finite criterion follows without computing logarithms or stationary laws. □

Theorem 13.2 (Quantum one-site entropy floor)**.** *For a valid exact local quantum circuit with final whole-site depolarization at $0 < \varepsilon < 1$, some TI stationary state attains the one-site von Neumann entropy floor (13.1), with $Q = d$, if and only if the pre-noise one-site output map is constant pure preparation. Finite exact matrix equality and rank tests decide this property.*

Proof. Let Ψ be the CPTP map from its finite input cone $\mathcal{N}$ to the one-site output before final noise, with ancillas evaluated. In a stationary state the preceding final noise gives

$$\rho_{\mathcal{N}} \geq (\varepsilon/d)^{|\mathcal{N}|} I_{\mathcal{N}},$$

since one positive branch resets all sites in that region. If λ_i are the eigenvalues of $\Psi(\rho_{\mathcal{N}})$, the final eigenvalues are $\varepsilon/d + (1-\varepsilon)\lambda_i$. Strict concavity on the probability simplex shows that their entropy is minimized exactly at a pure pre-noise state.

Suppose $\Psi(\rho_{\mathcal{N}}) = P$ is rank one and $\rho_{\mathcal{N}} \geq cI$. Positivity gives $0 \leq c\Psi(I) \leq P$. For every positive A, the inequality $\Psi(A) \leq \|A\|\Psi(I)$ puts its range in the same rank-one space. Trace preservation forces $\Psi(A) = \mathrm{Tr}(A)P$, hence by linearity this identity holds on all matrices. Conversely, constant pure one-site output maps force every finite pre-noise output to be the product pure state: a state with each marginal supported on a rank-one projector is supported on their tensor product. Final noise therefore gives an iid bath state, stationary and at the floor.

The finite circuit computes the algebraic matrix of Ψ. With input dimension $d_{\mathcal{N}} = d^{|\mathcal{N}|}$ and $\theta = \Psi(I/d_{\mathcal{N}})$, check

$$\Psi(|i\rangle\langle j|) = \delta_{ij}\theta \quad \text{for all matrix units,} \qquad \mathrm{rank}(\theta) = 1.$$

These are exact algebraic equality and rank tests. They decide the property on every valid input. □

Proposition 13.3 (The boundary lies outside the Rice class)**.** *At the parameters of Theorem B, entropy-floor attainment is nontrivial and decidable in each category but is not a function of σ.*

Proof. Use binary constant reset and binary noisy identity. Reset attains the appropriate floor. Under noisy identity, the probability that some site of a finite observed region O has not yet been refreshed by time t is at most $|O|(1-\varepsilon_0)^t$. On the other branches the observed law or density matrix is maximally mixed. Thus identity converges locally to the unique uniform iid state, even from spatially correlated or entangled starts. Reset and identity both have $\sigma = (1, 0, 0, \mathrm{true}, \mathrm{false})$, whereas identity's entropy is $\log 2 > h_{\mathrm{bath}}$. Their opposite floor verdicts prove the assertion. □

The classical floor property is invariant under relabeling and under adding a shared-flag constant bath: the product deterministic rule is constant exactly when the host rule is constant. For independent baths the same holds if the entropy floor is taken per track. An ambient whole-alphabet floor uses a different noise convention. Thus bath neutrality by itself does not substitute for dependence on the stationary datum. The boundary example identifies a limitation of this Rice class, rather than a maximality assertion about all possible classes. The floor predicate also escapes the reduction for a transparent reason: a host coupled to a nonconstant controller has a nonconstant product rule, so the forgetting branch does not preserve floor attainment.

14. Consequences and open questions

The line theorems apply at the single rational rate in (3.1). Their nonhalting construction produces continuum many extreme TI stationary states with a Bragg peak at π, distinguished by compatible relative spatial origins. Their halting correspondence preserves every fixed nonzero-wavevector Bragg label of each corresponding host state. The Rice datum counts states

with a peak somewhere: counts at another specified wavevector can retain host information and are outside the assertion.

On the line, finite upper total-count cuts are Π_2^0 and halting-hard; finite lower cuts are Σ_2^0 and nonhalting-hard. Matching second-level hardness is unresolved here. In dimension seven, the published-source condition C_T gives the completeness conclusions of Theorem C at ε_1, and, under the same condition, the stationary coupling argument of Theorem 12.3 excludes nontrivial Δ_2^0 predicates of σ_{stat}. The transverse-slice construction used there does not supply a line totality reduction. Arithmetic upper bounds for the ordered and Bragg subset counts and L, and matching degrees for all-law ergodicity, remain open.

Only the reset and maximal coarse data are witnessed on the line in this work. Realizability of the other possibilities in Theorem 8.1, including exactly $k \geq 2$ extreme TI stationary states, remains unresolved. So do fixed-alphabet restrictions, the line quantum category with once-per-round whole-site depolarization, a bounded-arity primitive-gate implementation of the present quantum controller, and the square-lattice Gibbs class with $\{0, 1\}$ energies at one fixed temperature. The rational-weight Gibbs theorem absorbs temperature into its weights. Fixed finite rings have a unique attracting stationary state, as explained in Remark 8.2; storage times growing with ring length are compatible with that conclusion.

Appendix A. Effective bounds on the total extreme count

Finite local stationarity constraints give upper bounds on the arithmetical complexity of counting extreme translation-invariant stationary states. The argument applies to rational classical PCA and algebraically presented local quantum circuits.

Theorem A.1. *Uniformly on valid finite presentations, for every fixed integer $r \geq 1$, the condition $N \leq r$ is Π_2^0. Consequently $N \geq k$ is Σ_2^0 for $k \geq 2$, and $N = 1$ is Π_2^0. For $k \geq 2$, the condition $N = k$ belongs to $D_2(\Sigma_2^0) \subseteq \Delta_3^0$. Finite total extreme count is Σ_3^0, and infinite total extreme count is Π_3^0.*

Here $D_2(\Sigma_2^0)$ denotes differences of two Σ_2^0 sets. All the assertions are upper bounds on the total count N.

Proof. Linear descriptions of the stationary set. Choose an effective countable determining family of real local coordinates $x_1, x_2, \ldots$, each bounded in absolute value by one. Classically use cylinder probabilities. Quantumly use real and imaginary entries of finite-region density matrices, with Hermiticity built into the coordinates. A state is specified by countably many linear equalities and inequalities: normalization, consistency under restriction, translation invariance, stationarity and positivity. Classical positivity is entrywise nonnegativity. For a finite quantum region Λ, impose

$$v^* \rho_\Lambda v \geq 0 \qquad \text{for every vector } v \text{ with rational real and imaginary parts.}$$

Each is a linear inequality in the chosen real coordinates. Density of these vectors makes the countable family equivalent to positive semidefiniteness of ρ_Λ. Stationarity on a local Hermitian basis has exact algebraic coefficients, computed from the finite backward circuit. Thus all constraints can be enumerated effectively and their coefficients admit exact arithmetic and order tests.

Let $b(m) \geq m$ increase so that the first m constraints use only the first $b(m)$ coordinates. In those coordinates let P_m be the bounded algebraic polytope defined by the first m constraints and the coordinate bounds $[-1, 1]$. Choose the enumeration so every earlier constraint is retained. The restriction of every genuine stationary state belongs to every P_m, so these polytopes are nonempty. Conversely, if $m_j \to \infty$ and points of P_{m_j}, extended arbitrarily in the unused bounded coordinates, converge coordinatewise, their limit satisfies every constraint. It is therefore a genuine translation-invariant stationary law or state.

Write $\mathcal{X} = \mathcal{K}(M)$. For every fixed n, the first-n coordinate projections of P_m converge in Hausdorff distance to the corresponding projection of $\mathcal{X}$. Indeed, a sequence of outer points staying a positive distance away would have a coordinatewise convergent subsequence, whose limit would belong to $\mathcal{X}$. Genuine points are contained at every stage.

Finite decidable approximation tests. The vertices of each P_m are computable with exact algebraic arithmetic. Enumerate full-rank systems of active constraints, solve their linear equations and check all inequalities. This also covers lower-dimensional polytopes: if the active normals at a vertex failed to span the ambient coordinate space, a small segment in a common null direction would remain feasible. Every bounded polytope is the convex hull of its vertices, as follows by expressing an interior point on a boundary-to-boundary segment and inducting on face dimension. Finite algebraic linear feasibility is consequently decidable by the same finite linear-algebra operations.

For $m \geq n$, put $\eta = 2^{-n}$, and let $V \geq 1$ be the number of vertices of P_m. Choose an integer $L_{m,n} \geq 8V/\eta$. Let $\mathcal{Q}_{m,n}$ contain all convex combinations of these vertices with nonnegative weights in multiples of $1/L_{m,n}$. It is a finite subset of P_m. It is an $\eta/4$-net in the first n coordinates: round the first $V-1$ weights down to multiples of $1/L_{m,n}$, putting the remainder on the last vertex. The sum of absolute weight errors is at most $2(V-1)/L_{m,n}$, so each coordinate error is at most $\eta/4$. The case $V = 1$ is immediate.

Define the finite test $\mathsf{Test}_r(n, m)$ to hold if some r-tuple of points in $\mathcal{Q}_{m,n}$, with repetitions allowed, has the following property: every vertex of P_m is within η, in first-n coordinate maximum norm, of that tuple's convex hull. For a fixed tuple and vertex v, the test is feasibility of

$$\lambda_i \geq 0, \qquad \sum_{i=1}^{r} \lambda_i = 1, \qquad \left| v_j - \sum_{i=1}^{r} \lambda_i y_{i,j} \right| \leq \eta \quad (1 \leq j \leq n).$$

This is a finite algebraic linear-feasibility problem. Define the test to be false on malformed or empty stage data. Enumerating the finitely many tuples makes $\mathsf{Test}_r(n, m)$ decidable. Covering the vertices also covers P_m, since the sum of a convex hull and a maximum-norm ball is convex.

The quantifier characterization. We claim that

$$N \leq r \quad \Longleftrightarrow \quad \forall n \geq 1 \; \exists m \geq n : \; \mathsf{Test}_r(n, m). \tag{A.1}$$

First recall why a compact convex set with finitely many extreme points is their convex hull in this setting. Suppose a point of $\mathcal{X}$ lies outside the hull of the alleged finite list. Compactness of the coefficient simplex gives a finite coordinate projection in which it still lies outside that hull. Finite-dimensional separation then supplies a local linear functional whose maximum on $\mathcal{X}$ exceeds its maximum on the list. Its maximizing compact face is disjoint from the hull. Successively maximizing a countable determining coordinate family inside that face leaves a singleton intersection. This point is extreme in $\mathcal{X}$, contradicting completeness of the list.

If $N \leq r$, choose r generators of $\mathcal{X}$, repeating points when necessary. At fixed n, choose $m \geq n$ so large that the first-n projection of P_m lies within $\eta/4$ of that of $\mathcal{X}$. Approximate the r generators by points of $\mathcal{Q}_{m,n}$ within $\eta/4$ in those coordinates. Their hull then covers the projected P_m within $\eta/2$, so $\mathsf{Test}_r(n, m)$ holds.

Conversely, choose stages $m_n \geq n$ and successful centers $y_1^{(n)}, \ldots, y_r^{(n)} \in P_{m_n}$. After extending unused coordinates in $[-1, 1]$, simultaneous product compactness gives a subsequence along which every center converges coordinatewise to a fixed point $y_i \in \mathcal{X}$. Membership in $\mathcal{X}$ follows because $m_n \geq n \to \infty$. Now fix any $x \in \mathcal{X}$. Its restriction lies in each relevant P_{m_n}, so vertex coverage supplies simplex coefficients $\lambda_i^{(n)}$ approximating its first n coordinates within 2^{-n}. A further subsequence makes these coefficients converge. Passing to each fixed

coordinate yields

$$x = \sum_{i=1}^{r} \lambda_i y_i.$$

The limiting centers were chosen before x; only the coefficients depend on it. Hence $\mathcal{X}$ is the convex hull of these same r points and has at most r extreme points. This proves (A.1).

The right side is a universal-existential formula with decidable matrix, proving the Π_2^0 bound. Its complement gives $N \geq k \in \Sigma_2^0$ for $k \geq 2$. Since $\mathcal{X}$ is nonempty compact convex, it has an extreme point, so $N = 1$ is already the case $N \leq 1$. For $k \geq 2$, write

$$\{N = k\} = \{N \geq k\} \setminus \{N \geq k+1\}.$$

This gives $D_2(\Sigma_2^0)$, contained in Δ_3^0. Finally, prefixing (A.1) with an existential integer r makes finite total count Σ_3^0; its complement is Π_3^0. □

References

[1] H. G. Rice, Classes of recursively enumerable sets and their decision problems, *Trans. Amer. Math. Soc.* **74** (1953), no. 2, 358–366. doi:10.1090/S0002-9947-1953-0053041-6.

[2] R. Berger, *The Undecidability of the Domino Problem*, Memoirs of the American Mathematical Society, no. 66, 1966. doi:10.1090/memo/0066.

[3] R. M. Robinson, Undecidability and nonperiodicity for tilings of the plane, *Invent. Math.* **12** (1971), 177–209. doi:10.1007/BF01418780.

[4] B. Durand, A. Romashchenko and A. Shen, Fixed-point tile sets and their applications, *J. Comput. Syst. Sci.* **78** (2012), no. 3, 731–764. doi:10.1016/j.jcss.2011.11.001.

[5] J. Kari, Rice's theorem for the limit sets of cellular automata, *Theoret. Comput. Sci.* **127** (1994), no. 2, 229–254. doi:10.1016/0304-3975(94)90041-8.

[6] J. Kari, The nilpotency problem of one-dimensional cellular automata, *SIAM J. Comput.* **21** (1992), no. 3, 571–586.

[7] P. Guillon and G. Richard, Revisiting the Rice theorem of cellular automata, in *STACS 2010*, LIPIcs **5** (2010), 441–452. doi:10.4230/LIPIcs.STACS.2010.2474.

[8] M. Delacourt, Rice's theorem for μ-limit sets of cellular automata, in *ICALP 2011, Part II*, Lecture Notes in Computer Science **6756** (2011), 89–100. doi:10.1007/978-3-642-22012-8_6.

[9] M. Delacourt, Rice's theorem for generic limit sets of cellular automata, in *AUTOMATA 2021*, OASIcs **90** (2021), 6:1–6:12. arXiv:2106.07907.

[10] M. Delacourt and B. Hellouin de Menibus, Characterisation of limit measures of higher-dimensional cellular automata, *Theory Comput. Syst.* **61** (2017), no. 4, 1178–1213. doi:10.1007/s00224-017-9753-1.

[11] B. Hellouin de Menibus and M. Sablik, Characterization of sets of limit measures of a cellular automaton iterated on a random configuration, *Ergodic Theory Dynam. Systems* **38** (2018), no. 2, 601–650. doi:10.1017/etds.2016.46.

[12] N. Carrasco-Vargas, On a Rice theorem for dynamical properties of SFTs on groups, preprint, arXiv:2401.10347v3, 2025.

[13] M. Hochman, On the dynamics and recursive properties of multidimensional symbolic systems, *Invent. Math.* **176** (2009), no. 1, 131–167. doi:10.1007/s00222-008-0161-7.

[14] M. Hochman and T. Meyerovitch, A characterization of the entropies of multidimensional shifts of finite type, *Ann. of Math.* (2) **171** (2010), no. 3, 2011–2038. doi:10.4007/annals.2010.171.2011.

[15] G. L. Kurdyumov, An algorithm-theoretic method for the study of uniform random networks, in *Multicomponent Random Systems*, R. L. Dobrushin and Ya. G. Sinai (eds.), Advances in Probability and Related Topics **6**, Marcel Dekker, New York, 1980, 471–503.

[16] A. Toom, Algorithmical unsolvability of the ergodicity problem for binary cellular automata, *Markov Process. Related Fields* **6** (2000), no. 4, 569–577. https://math-mprf.org/journal/articles/id882/.

[17] A. Toom, Algorithmical unsolvability of the ergodicity problem for locally interacting processes with continuous time, *J. Stat. Phys.* **98** (2000), 495–501. doi:10.1023/A:1018699527637.

[18] A. L. Toom, Stable and attractive trajectories in multicomponent systems, in *Multicomponent Random Systems*, R. L. Dobrushin and Ya. G. Sinai (eds.), Advances in Probability and Related Topics **6**, Marcel Dekker, New York, 1980, 549–575.

[19] A. L. Toom, N. B. Vasiliev, O. N. Stavskaya, L. G. Mityushin, G. L. Kurdyumov and S. A. Pirogov, Discrete local Markov systems, in *Stochastic Cellular Systems: Ergodicity, Memory, Morphogenesis*, R. L. Dobrushin, V. I. Kryukov and A. L. Toom (eds.), Manchester University Press, Manchester, 1990, 1–182.

[20] P. Gács, Probabilistic cellular automata with Andrei Toom, preprint, arXiv:2306.13226v3, 2024.

[21] P. Gács, Reliable cellular automata with self-organization, *J. Stat. Phys.* **103** (2001), 45–267. doi:10.1023/A:1004823720305.
[22] A. Bušić, J. Mairesse and I. Marcovici, Probabilistic cellular automata, invariant measures, and perfect sampling, *Adv. Appl. Probab.* **45** (2013), no. 4, 960–980. doi:10.1239/aap/1386857853.
[23] I. Törmä, A uniquely ergodic cellular automaton, *J. Comput. Syst. Sci.* **81** (2015), no. 2, 415–442. doi:10.1016/j.jcss.2014.10.001.
[24] I. Marcovici, M. Sablik and S. Taati, Ergodicity of some classes of cellular automata subject to noise, *Electron. J. Probab.* **24** (2019), paper no. 41, 1–44. doi:10.1214/19-EJP297.
[25] I. Marcovici and S. Taati, Positive-rate PCA and IPS with stationary Bernoulli measures are rapidly forgetful, preprint, arXiv:2605.16904, 2026.
[26] H. Marsan and M. Sablik, Strong stochastic stability of cellular automata, *Electron. J. Probab.* **30** (2025), 1–32. doi:10.1214/25-EJP1396.
[27] H. Marsan and M. Sablik, Characterization of the set of zero-noise limits measures of perturbed cellular automata, preprint, arXiv:2412.04672, 2024.
[28] H. Marsan, M. Sablik and I. Törmä, A perturbed cellular automaton with two phase transitions for the ergodicity, *J. Stat. Phys.* **193** (2026), no. 1, article 6. doi:10.1007/s10955-025-03564-0.
[29] M. Gu, C. Weedbrook, A. Perales and M. A. Nielsen, More really is different, *Physica D* **238** (2009), 835–839. doi:10.1016/j.physd.2008.12.016.
[30] H. Sandhu, Finite certificates and undecidable Gibbs uniqueness, preprint, version 1.0, Zenodo, 6 November 2026. doi:10.5281/zenodo.23983514.
[31] N. Shiraishi and K. Matsumoto, Undecidability in quantum thermalization, *Nature Commun.* **12** (2021), article 5084. doi:10.1038/s41467-021-25053-0.
[32] T. S. Cubitt, D. Pérez-García and M. M. Wolf, Undecidability of the spectral gap, *Nature* **528** (2015), 207–211. doi:10.1038/nature16059.
[33] J. Bausch, T. S. Cubitt and J. D. Watson, Uncomputability of phase diagrams, *Nature Commun.* **12** (2021), article 452. doi:10.1038/s41467-020-20504-6.
[34] S. Douglas and N. Mghirbi, *Forgetting under noise is undecidable*, preprint, version 2.2.0, November 2026. doi:10.5281/zenodo.24207965.
[35] S. Douglas and N. Mghirbi, *One noisy cellular automaton with an uncomputable noise diagram*, preprint, version 1.4.0, November 2026. doi:10.5281/zenodo.24204195.
[36] S. Douglas and N. Mghirbi, *Spontaneous magnetization and thermal decay are undecidable at positive temperature*, preprint, version 1.2.0, 2026. doi:10.5281/zenodo.23962706.
[37] S. Douglas and N. Mghirbi, *An uncomputable noise diagram on the line and a finite-energy Gibbs diagram on the square lattice*, preprint, version 2.0.0, 2026. doi:10.5281/zenodo.24352353.
[38] S. Douglas and N. Mghirbi, *Exponential ergodicity and the space-time Gibbs states of one-dimensional probabilistic cellular automata*, preprint, version 1.2.0, 2026. doi:10.5281/zenodo.24122307.
[39] E. Bishop and K. de Leeuw, The representations of linear functionals by measures on sets of extreme points, *Ann. Inst. Fourier* **9** (1959), 305–331. doi:10.5802/aif.95.
[40] P. Gács, Reliable computation with cellular automata, *J. Comput. Syst. Sci.* **32** (1986), no. 1, 15–78. doi:10.1016/0022-0000(86)90002-4.
[41] E. Domany and W. Kinzel, Equivalence of cellular automata to Ising models and directed percolation, *Phys. Rev. Lett.* **53** (1984), no. 4, 311–314. doi:10.1103/PhysRevLett.53.311.
[42] S. Goldstein, R. Kuik, J. L. Lebowitz and C. Maes, From PCA's to equilibrium systems and back, *Commun. Math. Phys.* **125** (1989), 71–79. `https://cmsr.rutgers.edu/images/people/lebowitz_joel/publications-archived/jll.pub_313.pdf`.
[43] J. L. Lebowitz, C. Maes and E. R. Speer, Statistical mechanics of probabilistic cellular automata, *J. Stat. Phys.* **59** (1990), nos. 1–2, 117–170. doi:10.1007/BF01015566.
[44] J. Miękisz and S. Taati, Quasicrystalline Gibbs states in 4-dimensional lattice-gas models with finite-range interactions, *J. Stat. Phys.* **193** (2026), article 166. doi:10.1007/s10955-026-03727-7.