Strict Generic Constructive Cycle-Core Universality for Reduced Ordered Positive Propagation Systems

Drew Higgins[1]

[1](affiliation omitted)

Abstract

We present a closed, obstruction-theoretic universality theory for a broad class of ordered positive propagation systems after a canonical locality reduction to a finite template with bounds. The reduced output is a finite directed graph with a finite-dimensional parameterization and an admissible constant ledger. The paper is centered around four master results: (i) a Master Reduction Theorem identifying canonical, representation-invariant reduced objects; (ii) a Cycle-Core Universality Theorem giving a sharp boundary equivalence between uniform minimizer-facing projective regularity and a finite set of certifiable contract gates; (iii) a Strict Genericity Theorem showing that unique minimizing cycle structure is open–dense and full measure in the fixed reduced parameter space; and (iv) a Constructive Witness Theorem producing a canonical witness family covering all minimizing components. Projective contraction rates and the minimizer lift are presented as corollaries. Implementation and certification details are deferred to an appendix.

Contents

1 Introduction

A recurring theme in positive propagation and transfer-matrix models is that quantitative regularity of the induced projective dynamics is governed not only by positivity but by a finite collection of structural and recurrence features. This paper formalizes that principle in a reduced setting where "locality" produces a finite template-with-bounds interface and an associated finite variational object.

The goal is not an absolute "locality forces spectral gap" theorem. Counterexamples exist without additional mixing/positivity hypotheses. Instead, we prove the strongest honest endpoint available in this architecture: a *complete obstruction-theoretic boundary characterization* of uniform minimizer-facing projective regularity (PR), together with a *strict genericity theorem* showing that a rigid, periodic minimizing core is prevalent in the fixed reduced parameterization.

What is meant by PR on minimizers. Given a nonnegative matrix cocycle acting on the positive cone, PR refers to uniform projective contraction (e.g. in the Hilbert metric) along the minimizing set of trajectories/controls selected by the reduced variational problem. The main equivalence theorem identifies a finite set of certifiable "contract gates" whose conjunction is necessary and sufficient for such uniform contraction.

How cycle-core enters. The locality reduction yields a finite directed graph $G = (V, E)$ and an additive cost c on edges. The *cycle-core* is the union of minimum-mean cycles. Minimizing invariant measures are supported on this set, so minimizer-facing statements reduce to core-facing statements.

Genericity. Within the fixed reduced parameter space (edge-cost coordinates $c \in \mathbb{R}^E$), ties between distinct cycle means lie on proper affine hyperplanes. Hence unique minimizing cycle structure is open–dense and full measure. This prevalence result collapses recurrence on the minimizing set to a periodic phenomenon on a full-measure chamber set.

Constructiveness. A canonical witness family is synthesized directly from the cycle-core SCC decomposition, removing hand-chosen witness dependence at the covering stage. The remaining nontrivial boundary burden is a certified projective contraction event (Doeblin/primitive block) and excursion-safety.

Related literature. The projective-metric approach to positive operators goes back to Birkhoff's contraction coefficient and subsequent refinements (see, e.g., [1, 2]). Transfer-matrix and subadditive-ergodic perspectives connect to random matrix products and Lyapunov exponents [3, 4]. The symbolic dynamics viewpoint (subshifts, cocycles, minimizing measures) underlies the cycle-core formalism [5].

Organization. Section 2 fixes the reduced objects and PR notion. Section 3 states the Master Reduction Theorem. Section 4 states the Cycle-Core Universality (boundary equivalence) theorem and derives the minimizer lift and PR rate bounds as corollaries. Section 5 gives strict genericity. Section 6 gives constructive witness synthesis. Appendix material contains the certification interface, constants, obstruction certificates, and proof skeletons.

2 Reduced setting and projective regularity

2.1 Reduced finite data

We assume a locality mechanism (left abstract) produces a *reduced output* consisting of:

- a finite directed graph $G = (V, E)$,
- a finite-dimensional parameter $c \in \mathbb{R}^E$ varying in an open admissible region $\mathcal{U} \subset \mathbb{R}^E$,
- an additive edge cost $c(e)$ and the minimum cycle mean value $\lambda_*(c)$,
- a positive propagation cocycle (or family of cocycles) satisfying a finite-template-with-bounds (FTG) interface on a cone $K = \mathbb{R}^d_+$.

The FTG interface is the only place where uniform constants enter; the full ledger is relegated to Appendix C. Intuitively: supports come from a finite template set and all nonzero coefficients lie in a fixed interval $[m_*, M_*]$.

2.2 Cycle means and cycle-core

For a directed cycle γ in G define the mean cost

$$\overline{c}(\gamma) = \frac{1}{|\gamma|} \sum_{e \in \gamma} c(e), \qquad \lambda_*(c) = \min_{\gamma \text{ cycle}} \overline{c}(\gamma).$$

The *cycle-core* $\mathcal{C}(c) \subseteq V$ is the union of all $\lambda_*(c)$-critical cycles. We also use the SCC decomposition

$$\mathcal{C}(c) = \bigsqcup_{j=1}^{J} \mathcal{C}_j(c).$$

2.3 Projective regularity (PR) on minimizers

Let $K^\circ = (0, \infty)^d$ and d_H be the Hilbert projective metric on K°. A positive N-step block product B has Birkhoff contraction coefficient $\tau(B) < 1$ when it has finite projective diameter.

Definition 1 (PR on minimizers)**.** *We say* PR on minimizers holds *if there exists $\kappa > 0$ such that, for every minimizing invariant measure μ (for the reduced variational problem) and μ-a.e. trajectory ω, the cocycle products satisfy*

$$d_H\big(A^{(n)}(\omega)x,\ A^{(n)}(\omega)y\big) \le C(\omega)\, e^{-\kappa n}\, d_H(x, y) \quad \textit{for all } x, y \in K^\circ,$$

with κ uniform over all minimizers.

The central boundary theorem identifies a finite conjunction of certifiable gates equivalent to this property and gives an explicit constant-closed lower bound κ_{safe} whenever PR holds.

3 Master Reduction Theorem

This paper treats locality reduction as a black box, but we require that it outputs a finite arena in which all later statements are representation-invariant.

Theorem 2 (Master Reduction Theorem: finite canonical arena)**.** *Assume a locality reduction produces, on a reduction stratum, a reduced output*

$$\mathsf{R} = (G = (V, E),\ \mathcal{U} \subset \mathbb{R}^E,\ c \in \mathcal{U},\ \textit{FTG ledger})$$

with G finite and $\mathcal{U}$ open.

Then the following objects are canonical (independent of the chosen representative within the standard cost-coboundary / projective gauge transformations) and depend only on R*:*

- *the minimum cycle mean value $\lambda_*(c)$,*
- *the cycle-core $\mathcal{C}(c)$ and its SCC decomposition $\mathcal{C}(c) = \bigsqcup_j \mathcal{C}_j(c)$,*
- *the minimizing set of invariant measures for the reduced variational problem, and the restriction of these measures to the cycle-core dynamics,*
- *the validity of contract gates (defined in Section 4) and the resulting certified exponent output when the gates hold.*

If the reduction output varies with parameters, the statements above are interpreted stratumwise: on each stratum where R *is constant, the canonical objects above are well-defined and stable under small perturbations inside the stratum.*

Remark 3 (Representation invariance)**.** *Two invariances are used repeatedly: (i) cost cohomology invariance $c'(u \to v) = c(u \to v) + \phi(u) - \phi(v)$ preserves all cycle totals and cycle means; (ii) projective gauge invariance preserves Hilbert-metric contraction properties of the cocycle. Precise statements are recorded in Appendix F.*

The reduction theorem fixes the finite arena. The remaining results identify what is universally determined inside this arena (cycle-core), what is boundary-equivalent (PR), what is prevalent (genericity), and what is canonical and constructive (witness synthesis).

4 Cycle-Core Universality and the PR boundary

This section states the central boundary equivalence: uniform PR on minimizers is true if and only if a finite conjunction of certifiable contract gates holds. The gates are designed to be non-overlapping: each failure mode blocks PR for a distinct reason and admits a finite obstruction certificate (Appendix E).

4.1 Contract gates

Fix a reduced output R on a stratum, with graph $G = (V, E)$, parameter $c \in \mathcal{U}$, and cycle-core $\mathcal{C}(c) = \bigsqcup_j \mathcal{C}_j(c)$.

Definition 4 (Contract gates)**.** *The contract gates are the conjunction of:*

- ***Gate B1 (Admissible contraction / constant closure).*** *The FTG ledger closes and certifies a projective contraction factor $\tau \in (0,1)$ at some block length $N \geq 1$ (e.g. via a Doeblin minorization event for an N-step block).*
- ***Gate B2 (Uniformity over minimizers).*** *A witness family covers every cycle-core SCC $\mathcal{C}_j$ so that the statement is uniform over all minimizing components.*
- ***Gate B3 (Recurrence on minimizing support).*** *A positive return density $\rho > 0$ is certified for the witness family along minimizing trajectories/measures (via a symbolic density witness, or via forced periodic recurrence in the single-cycle chamber when the witness lies on the core).*
- ***Gate B4 (Excursion-safety).*** *The reduction's visibility/safety conditions prevent witness frequency from vanishing along minimizers (no "escaping" minimizers that avoid the contracting blocks at asymptotic scale).*

4.2 Boundary equivalence

Theorem 5 (Cycle-Core Universality Theorem: PR on minimizers $\iff$ contract gates)**.** *Fix a reduced output on a stratum. Then PR on minimizers holds if and only if the contract gates in Definition 4 hold.*

Moreover, when the gates hold the framework returns a constant-closed exponent bound

$$\kappa_{\mathrm{safe}} \;\geq\; \rho\,|\log \tau|,$$

where τ is the certified projective contraction factor and ρ is the certified return density from Gate B3.

When PR fails, exactly one gate blocks the proof and a finite obstruction certificate can be produced consistent with that gate.

Remark 6 (Where the cycle-core is universal)**.** *The cycle-core is the unique minimizing support object: minimizer-facing recurrence and contraction statements reduce to core-facing statements, and the only place where parameter dependence enters is through the finite reduced costs and the cycle-core geometry.*

4.3 Corollaries: minimizer lift and rate extraction

Corollary 7 (Minimizer lift)**.** *Every minimizing invariant measure for the reduced variational problem is supported on the cycle-core $\mathcal{C}(c)$. In particular, the set of minimizers decomposes over the SCC components $\mathcal{C}_j(c)$.*

Corollary 8 (PR rate extraction from a witnessed contraction event)**.** *Assume the contract gates hold with contraction factor $\tau \in (0,1)$ and return density $\rho > 0$. Then along minimizing trajectories one has an exponential projective contraction rate at least $\kappa_{\mathrm{safe}} \geq \rho|\log \tau|$ (after the usual block-length normalization).*

Remark 9 (Quantitative limits)**.** *The theory is obstruction-complete but does not imply a uniform quantitative lower bound on κ_{safe} over all admissible parameters: near-ties in cycle means allow arbitrarily small margins and therefore arbitrarily small return-density or contraction-frequency constants. The correct quantitative statement is local (stratumwise) and certificate-driven.*

5 Strict genericity in the reduced parameter space

This section proves prevalence of unique minimizing cycle structure in a fixed reduced graph parameterization. The theorem is purely finite-dimensional: it does not replace the boundary equivalence, but it identifies a large chamber set where recurrence simplifies to periodicity.

Theorem 10 (Strict Genericity Theorem: unique minimizing cycle is open–dense and full measure)**.** *Fix a finite directed graph $G = (V, E)$ and edge-cost parameterization $c \in \mathbb{R}^E$ with Euclidean topology. For each directed cycle γ, the mean cost $\mu_\gamma(c) = \bar{c}(\gamma)$ is affine-linear in c.*

Let

$$\mathcal{G} \;:=\; \{c \in \mathbb{R}^E : \text{ the minimizing cycle is unique}\}.$$

Then $\mathcal{G}$ is the complement of a finite union of proper affine hyperplanes and hence is:

- *open,*
- *dense,*
- *full Lebesgue measure.*

On $\mathcal{G}$, the cycle-core $\mathcal{C}(c)$ is a single directed cycle $\gamma_*(c)$ and is locally constant on a neighborhood whose size is controlled by the cycle margin

$$\Delta(c) \;:=\; \min_{\gamma \neq \gamma_*} \big(\mu_\gamma(c) - \mu_{\gamma_*}(c)\big) \;>\; 0.$$

Remark 11 (Stratified interpretation)**.** *If the reduction output $(G, \mathcal{U})$ varies with parameters, apply Theorem 10 stratumwise on each region where the reduced graph is constant, and take the union of the resulting open–dense/full-measure chamber sets.*

Remark 12 (What genericity does and does not provide)**.** *Genericity collapses the minimizing support to a single periodic orbit (hence explicit return density), but it does not itself create projective contraction. Contraction remains a boundary property supplied by Gate B1 and the FTG constants.*

6 Constructive witness synthesis

The boundary theorem requires a witness family that (i) covers each minimizing SCC and (ii) recurs with positive density along minimizing trajectories. This section makes the covering step canonical and provides a generic recurrence collapse on the unique-cycle chamber.

Theorem 13 (Constructive Witness Theorem: canonical SCC covering and generic recurrence collapse)**.** *Fix a reduced output on a stratum with cycle-core SCC decomposition $\mathcal{C}(c) = \bigsqcup_{j=1}^{J} \mathcal{C}_j(c)$.*

- ***Canonical covering.*** *There exists a canonical procedure which selects, for each SCC $\mathcal{C}_j$, a short core cycle word w_j supported inside $\mathcal{C}_j$. The resulting witness family $\mathcal{W} = \{w_j\}_{j=1}^{J}$ covers the cycle-core (Gate B2).*
- ***Generic periodic recurrence.*** *On the open–dense/full-measure chamber set $\mathcal{G}$ of Theorem 10, the cycle-core is a single directed cycle γ_* of period $p = |\gamma_*|$. If a valid witness word*

w lies on γ_* (e.g. w is a period word or a bounded iterate of the period), then minimizing measures are periodic on γ_* and the return density is forced:

$$\rho \;\geq\; \frac{1}{p}.$$

Thus, on $\mathcal{G}$ the only nontrivial boundary burden in Theorem 5 is admissible contraction (Gate B1) together with excursion-safety (Gate B4).

Remark 14 (Autosynthesis and finiteness)**.** *The selection in Theorem 13 can be made purely combinatorial from the finite reduced data (choose a representative directed cycle in each SCC, then take a bounded iterate as needed for a primitivity/Doeblin event). Details and an explicit autosynthesis routine are provided in Appendix B.*

7 Discussion and submission posture

Claim level. The results give *strict generic constructive cycle-core universality* inside the admissible reduced class. The architecture is obstruction-complete: failures of PR on minimizers are explained by a finite menu of blocking gates with explicit certificates.

What has been maximized. Within this model class, no stronger universality claim is available without strengthening hypotheses. In particular, "locality alone forces a spectral gap" is false in general; the correct universal statement is the boundary equivalence (Theorem 5) plus strict genericity in the fixed reduced parameter space (Theorem 10).

How to present. For a clean submission package, we recommend centering the exposition around the four master results: Theorem 2, Theorem 5, Theorem 10, and Theorem 13, with Corollaries 7–8 as immediate consequences. Implementation and certification mechanics belong in the appendix.

Referee attack surface reduction. The only points where the reader must trust a "black box" are (i) the locality reduction producing a finite reduced output and (ii) the FTG constant ledger. Both are isolated: (i) is encapsulated in the hypotheses of Theorem 2; (ii) is an appendix contract interface from which all quantitative constants follow.

Open directions. Two natural extensions are: (a) widening the admissible class by relaxing the FTG bounds while preserving enough positivity/mixing to certify Gate B1; (b) obtaining quantitative lower bounds on PR exponents on restricted parameter regions away from hyperplane tie sets.

A Sharpness and realizability of boundary gates

We show that each obstruction class appearing in the Cycle-Core Universality Theorem (Theorem 5) is *realizable* inside the admissible reduced class. In particular, none of the contract gates is redundant: removing any gate admits examples where PR on minimizers fails for a mathematically intrinsic reason.

Throughout this appendix we work on a fixed finite reduced directed graph $G = (V, E)$ and the FTG interface of Section 2. The constructions are elementary and purely finite: we choose costs $c \in \mathbb{R}^E$ and (when relevant) a family of cone-preserving one-step matrices respecting the FTG template bounds.

A.1 Failure of unique minimizing core (tie realizability)

Proposition 15 (Tie realizability)**.** *There exist admissible costs $c \in \mathbb{R}^E$ for which multiple simple directed cycles attain the minimum mean, hence the cycle-core has more than one component.*

Proof. Assume G contains two vertex-disjoint simple cycles γ_1, γ_2 of the same length L (or, more generally, two cycles whose means can be tuned independently). Set $c(e) = 0$ for $e \in \gamma_1 \cup \gamma_2$ and $c(e) = 1$ for all other edges. Then

$$\mu_{\gamma_1}(c) = \mu_{\gamma_2}(c) = 0 = \lambda_*(c),$$

so both cycles are minimizing and the cycle-core has at least two SCC components.

More generally, for any two distinct cycles $\gamma \neq \gamma'$ in a fixed graph, the tie set

$$\mu_\gamma(c) - \mu_{\gamma'}(c) = 0$$

is a proper affine hyperplane in $\mathbb{R}^E$. Hence non-uniqueness is nongeneric (codimension one) but not empty. □

A.2 Failure of Doeblin witness (minorization failure)

Proposition 16 (Doeblin witness is not automatic)**.** *There exist admissible FTG propagation systems for which the minimizing cycle-core is unique, but no admissible Doeblin minorization event occurs on that core. Consequently no strict projective contraction occurs on minimizers.*

Proof. Fix a graph G whose minimizing cycle-core is a single simple cycle γ. Choose costs so that γ is the unique minimizing cycle (e.g., set $c(e) = 0$ for $e \in \gamma$ and $c(e) = 1$ for $e \notin \gamma$).

Take $d \geq 2$ and define the one-step matrices along γ to be diagonal and bounded, for instance

$$A(e) = \mathrm{diag}(1, \dots, 1) \quad \text{for every edge } e \in \gamma,$$

with the FTG template allowing diagonal support and all nonzero entries in $[m_*, M_*]$.

Then every block product along the unique minimizing trajectory is diagonal. In particular no block product is primitive on $K = \mathbb{R}^d_+$, so no Doeblin inequality of the form recorded in `appendix/DEFINITIONS.tex` can hold at any block length. Equivalently, every block product has infinite Hilbert diameter on K°, so its Birkhoff coefficient equals $\tau = 1$ and strict projective contraction is impossible. □

A.3 Failure of recurrence on minimizing support (zero-frequency realizability)

Proposition 17 (Recurrence gate is necessary)**.** *There exist admissible reduced systems in which a Doeblin witness block exists on the minimizing core SCC, but some minimizing ergodic measure assigns zero frequency to every such witness. Hence PR on minimizers fails.*

Proof. Work inside a single strongly connected subgraph $H \subseteq G$ and set $c(e) = 0$ for $e \in H$, $c(e) = 1$ for $e \notin H$. Then $\lambda_*(c) = 0$ and every invariant ergodic measure supported on H is minimizing.

Choose an alphabet realization of H so that there is a finite word w whose block product A_w is an S-Doeblin event on the relevant face (so $\tau_S(A_w) < 1$), but such that H also contains a directed cycle η whose periodic orbit avoids w entirely.

The periodic orbit measure μ_η is minimizing (it is supported on H) and satisfies $\mu_\eta([w]) = 0$. Along μ_η-typical trajectories there are no Doeblin blocks, so no uniform positive projective exponent can hold. This realizes the "recurrence on minimizing support" obstruction. □

A.4 Failure of excursion-safety (vanishing-frequency realizability)

Proposition 18 (Vanishing-frequency obstruction is realizable)**.** *There exist admissible reduced systems for which every minimizing measure sees some Doeblin event, but the infimum of Doeblin-event frequencies over minimizing ergodic measures is* 0*. Consequently no uniform exponent $\kappa > 0$ can hold in the PR definition.*

Proof. As in the previous proof, let $H \subseteq G$ be a strongly connected minimizing core with $c \equiv 0$ on H.

Arrange that H contains a short cycle segment that realizes a Doeblin word w (so $\tau_S(A_w) = \tau < 1$), and also contains arbitrarily long directed detours that avoid w but return to its basepoint. For each n choose a periodic orbit $\omega^{(n)}$ that traverses w exactly once and then follows a w-avoiding detour of length n before closing up. The associated periodic orbit measure μ_n is minimizing and satisfies

$$\mu_n([w]) \asymp \frac{1}{n + |w|} \longrightarrow 0.$$

Along μ_n-typical trajectories, the disjoint Doeblin-block density is $\rho_n \to 0$, and the best available exponent bound scales like $\rho_n |\log \tau| \to 0$ (compare Theorem 5). Since PR requires a *single* $\kappa > 0$ to work uniformly over *all* minimizing ergodic measures, the existence of the family $\{\mu_n\}$ forces the optimal uniform exponent to be 0, so PR fails.

Remark (gap-dilution mechanism). If one perturbs the above example so that the long detour has mean $\lambda_*(c) + \varepsilon$ instead of $\lambda_*(c)$, then the same construction produces near-minimizers with average cost $\lambda_*(c) + o(1)$ and Doeblin frequency $o(1)$. This is the "cheap avoidance" mechanism isolated by the excursion-safe SG certificate in the hybrid contract statement. □

A.5 Summary: each gate is sharp

Each obstruction corresponds to an explicit, finite configuration inside the admissible reduced class.

Gate / obstruction	Realizable?	Genericity (within fixed G)
Unique core failure (ties)	Yes	Nongeneric (affine hyperplanes)
Doeblin failure (no minorization)	Yes	Admissibility-dependent
Recurrence failure (zero frequency)	Yes	Possible in multi-cycle cores
Excursion / vanishing-frequency failure	Yes	Typical unless forced recurrence

Therefore the boundary partition in Theorem 5 is obstruction-complete and non-redundant: no gate can be dropped without admitting genuine counterexamples.

B Appendix overview: certification and engineers-kit material

This appendix contains the engineering-facing certification layer supporting the main theorem suite. The main text is intentionally free of implementation narrative; here we record the constant ledger, the Doeblin/Birkhoff contraction lemmas with explicit constants, representative obstruction certificates, and proof skeletons.

What is certified. The boundary equivalence (Theorem 5) is implemented as a finite gate system:

- Gate B1: a certified contraction event (typically an N-step Doeblin block), together with explicit $\tau < 1$;
- Gate B2: a witness family covering each cycle-core SCC;
- Gate B3: a certified return density $\rho > 0$ for the witness family along minimizers;
- Gate B4: excursion-safety preventing witness frequency from vanishing along minimizers.

What is produced when a gate fails. If PR fails, the framework returns (i) the unique blocking gate and (ii) a finite obstruction certificate witnessing that failure mode (Appendix E). This is the "obstruction-complete" content.

What is constructive. The witness-covering step is made canonical by selecting one short core cycle word per SCC and applying a bounded iterate when a primitivity/Doeblin event is required. In the unique-cycle chamber, recurrence collapses to periodicity and the return density is forced.

The remaining appendix items are organized as follows:

- Appendix C: FTG ledger, Hilbert metric, and explicit contraction constants.
- Appendix E: obstruction certificate templates for each gate.
- Appendix F: proof skeletons and invariance lemmas.

C Constant ledger and contract lemmas

We include here the constant ledger and contract-level lemma suite used to certify Gate B1 and to convert a primitive/Doeblin event into an explicit projective contraction factor.

Parameter and Notation Ledger (Contract Level)

Throughout, let $d \in \mathbb{N}$ and $K = \mathbb{R}_+^d$ with interior $K^\circ = (0,\infty)^d$.

Finite-Template-with-Bounds Interface (FTG)

A matrix cocycle $\{A(\omega)\}$ satisfies **FTG** with parameters $(d, \mathcal{T}, m_*, M_*)$ if:

- $\mathcal{T} \subset \{0,1\}^{d\times d}$ is a finite set of *support templates*.

- For each realized matrix $A(\omega)$, $\mathrm{supp}(A(\omega)) \in \mathcal{T}$.
- For every (i, j) with $\mathrm{supp}(A(\omega))_{ij} = 1$, the coefficient is bounded:

$$m_* \le A_{ij}(\omega) \le M_*.$$

- Cone preservation: $A(\omega)K \subset K$.

No finite-valued coefficient assumption: coefficients may vary continuously inside $[m_*, M_*]$.

Hilbert metric and operator diameter

For $x, y \in K^\circ$, the Hilbert projective metric is

$$d_H(x, y) = \log \frac{\max_i x_i/y_i}{\min_i x_i/y_i}.$$

For an operator $B : K^\circ \to K^\circ$ (e.g. a positive matrix),

$$\Delta(B) = \sup_{x,y \in K^\circ} d_H(Bx, By)$$

is its projective diameter, and the Birkhoff contraction coefficient is

$$\tau(B) = \tanh\Big(\frac{\Delta(B)}{4}\Big).$$

Primitive block certificate

A nonnegative matrix Q is *primitive* if $Q^k > 0$ for some k. A *primitivity exponent* is an integer k_0 such that $Q^{k_0} > 0$. (One may always take $k_0 \le W(d)$ where $W(d)$ is the Wielandt bound.)

In applications, Q is typically a *core block product*:

$$Q = A^{(B)}(\omega) := A(\sigma^{B-1}\omega) \cdots A(\omega),$$

with block length B fixed by the recurrent core word/return block.

We will write

$$N := Bk_0$$

for the total number of one-step matrices in the derived Doeblin block.

Doeblin minorization form

Given $\varepsilon > 0$, the Doeblin (rank-one floor) inequality is:

$$Bv \;\ge\; \varepsilon\langle \mathbf{1}, v\rangle\, \mathbf{1} \quad \forall v \in K,$$

where $\mathbf{1} = (1, \ldots, 1)^\top$ and $\langle \mathbf{1}, v\rangle = \sum_i v_i$.

Strict gap (SG) as variational certificate

SG is treated as a *finite certification* condition on the reduced finite system: a computable separation $\delta > 0$ between the minimal cycle-mean growth and all competitors. This is external to FTG and is not derived from locality in this contract layer.

Contract Lemma Suite (Pass 1)

Lemma A (Primitive block $\Rightarrow$ explicit Doeblin minorization)

Let $\{A(\omega)\}$ satisfy $\mathrm{FTG}(d, \mathcal{T}, m_*, M_*)$. Fix a block length $B \in \mathbb{N}$ and define the block product

$$Q(\omega) = A(\sigma^{B-1}\omega) \cdots A(\omega).$$

Assume there exists a primitivity exponent $k_0 \in \mathbb{N}$ such that, for the relevant core block,

$$Q(\omega)^{k_0} > 0.$$

Set $N := Bk_0$ and $B_* := Q(\omega)^{k_0}$ (so B_* is an N-step product of one-step matrices).

Then every entry of B_* admits the conservative bounds

$$(B_*)_{ij} \geq m_*^N, \qquad (B_*)_{ij} \leq (dM_*)^N.$$

In particular, for every $v \in K$ and every coordinate i,

$$(B_* v)_i \;\geq\; m_*^N \langle \mathbf{1}, v \rangle.$$

Equivalently,

$$B_* v \;\geq\; \varepsilon \langle \mathbf{1}, v \rangle \, \mathbf{1} \quad \text{with} \quad \varepsilon := m_*^N.$$

Notes for tightening. The lower bound uses only: (i) nonnegativity, (ii) $Q^{k_0} > 0$, and (iii) per-edge floor m_*. Sharper ε can be obtained by exploiting known path multiplicities in the core templates.

Lemma B (Doeblin floor $\Rightarrow$ finite diameter and explicit contraction)

Let $B : K \to K$ be a nonnegative linear operator satisfying:

$$Bv \geq \varepsilon \langle \mathbf{1}, v \rangle \, \mathbf{1} \quad \forall v \in K$$

for some $\varepsilon > 0$. Assume also the coordinate-wise upper bound:

$$(Bv)_i \leq U \langle \mathbf{1}, v \rangle \quad \forall v \in K, \ \forall i,$$

for some finite $U > 0$.

Then for all $x, y \in K^\circ$,

$$d_H(Bx, By) \leq 2 \log\Big(\frac{U}{\varepsilon}\Big),$$

hence

$$\Delta(B) \leq 2 \log\Big(\frac{U}{\varepsilon}\Big) < \infty, \qquad \tau(B) \leq \tanh\Big(\frac{1}{2} \log\Big(\frac{U}{\varepsilon}\Big)\Big) = \frac{\frac{U}{\varepsilon} - 1}{\frac{U}{\varepsilon} + 1} < 1.$$

Contract choice of U under FTG. For an N-step product B of one-step FTG matrices, a safe bound is

$$U := (dM_*)^N,$$

so Lemma A + Lemma B yield

$$\Delta(B) \leq 2N \log\Big(\frac{dM_*}{m_*}\Big), \qquad \tau(B) \leq \tanh\Big(\frac{N}{2} \log\Big(\frac{dM_*}{m_*}\Big)\Big).$$

Corollary (Primitive core block $\Rightarrow$ derived Birkhoff contraction event)

Under the hypotheses of Lemma A, the N-step block operator B_* satisfies:

$$\Delta(B_*) < \infty, \qquad \tau(B_*) < 1,$$

with constants depending only on (d, m_*, M_*, N) through the bounds above.

D Return Density Certification and Derived Contraction on Minimizers

D.1 SG as a finite certification and periodic-core output

Let $G = (V, E)$ be a finite directed graph with an additive edge-cost $c : E \to \mathbb{R}$. For a directed cycle γ define the mean cost

$$\bar{c}(\gamma) = \frac{1}{|\gamma|} \sum_{e \in \gamma} c(e), \qquad \lambda_* := \min_{\gamma \text{ cycle}} \bar{c}(\gamma).$$

Let G_* denote the union of all λ_*-cycles (the critical subgraph).

Definition 19 (SG–PC certificate)**.** *We say that (G, c) satisfies* SG–PC *if:*

- *(strict gap) there exists $\delta > 0$ such that every directed cycle γ not contained in G_* satisfies $\bar{c}(\gamma) \geq \lambda_* + \delta$;*
- *(periodic-core structure) every strongly connected component of G_* is a simple directed cycle.*

Lemma 20 (Minimizing measures are periodic under SG–PC)**.** *Assume SG–PC. Then every ergodic invariant measure minimizing $\int c\, d\mu$ is supported on one of the periodic orbits corresponding to the critical cycles. In particular, if a critical cycle has period L, then the associated ergodic minimizing measure μ satisfies*

$$\mu([w]) = \frac{1}{L}$$

for the period word w (so w occurs with asymptotic frequency $1/L$ along μ-typical trajectories).

D.2 Doeblin block and contraction along minimizing trajectories

Let $A : \Sigma \to \mathbb{R}^{d \times d}_{\geq 0}$ be a cocycle over a subshift Σ satisfying FTG: there is a finite template family $\mathcal{T}$ and bounds $m_*, M_* > 0$ such that each realized matrix has support in $\mathcal{T}$ and all nonzero entries lie in $[m_*, M_*]$.

Fix a support set $S \subset \{1, \ldots, d\}$ and write $K_S = \{x \geq 0 : \mathrm{supp}(x) \subseteq S\}$.

Definition 21 (S-Doeblin block)**.** *A finite allowed word w (length L) is an S-*Doeblin word *if its block product $B(w)$ is strictly positive on $S \times S$. Under FTG bounds this yields a minorization:*

$$B(w)v \geq \varepsilon \langle \mathbf{1}_S, v \rangle \mathbf{1}_S,$$

for some $\varepsilon = \varepsilon(|S|, m_*, M_*, L) > 0$, and hence a Birkhoff contraction coefficient $\tau = \tau(\varepsilon) < 1$ in the Hilbert metric on K_S°.

Proposition 22 (Repeated Doeblin blocks imply exponential contraction)**.** *Let w be an S-Doeblin word with Birkhoff coefficient $\tau < 1$. For any trajectory ω and any n, let $N_w(n;\omega)$ be the number of disjoint occurrences of w among the first n symbols of ω. Then for all $x, y \in K_S^\circ$,*

$$d_H(A^{(n)}(\omega)x,\ A^{(n)}(\omega)y) \le \tau^{N_w(n;\omega)}\, d_H(x,y).$$

If $N_w(n;\omega) \ge \rho n - o(n)$ along a set of trajectories (e.g. under a minimizing ergodic measure), then one obtains an exponential rate $d_H(\cdot) \le C(\omega)e^{-\kappa n} d_H(x,y)$ with $\kappa = \rho|\log \tau|$ (up to the usual block-length normalization).

Remark 23 (How the contract supplies ρ)**.** *Under SG–PC (Lemma 20), minimizing measures are periodic and yield an explicit return density $\rho = 1/L$ for the period word. Choosing an S-Doeblin word as a suitable iterate of the period (e.g. k_0 periods so that the product is primitive) makes Proposition 22 fully quantitative with constants depending only on (d, m_*, M_*, k_0, L).*

E Obstruction certificates

This appendix describes the finite obstruction outputs returned when PR on minimizers fails, organized by the blocking gate in Definition 4. The philosophy is: each gate corresponds to a logically distinct failure mode; the certificate is a finite witness of that mode in the reduced finite arena.

Gate B1 obstruction: no certified contraction event

Typical certificate content:

- the alleged contracting block length N,
- the candidate witness block product template(s),
- verification failure: either the block does not become positive on the intended support, or the FTG ledger cannot close (missing bounds).

Gate B2 obstruction: incomplete SCC covering

Typical certificate content:

- the SCC decomposition $\mathcal{C} = \bigsqcup_j \mathcal{C}_j$,
- the chosen witness family $\mathcal{W}$,
- an uncovered SCC index j with a short cycle word inside $\mathcal{C}_j$.

Gate B3 obstruction: zero return density

Typical certificate content:

- a minimizing component (SCC or orbit class) and an invariant measure supported there,
- a witness family $\mathcal{W}$,
- a verified upper bound $N_w(n;\omega) = o(n)$ along minimizing trajectories, i.e. witness blocks occur sublinearly.

Gate B4 obstruction: excursion/visibility failure

Typical certificate content:

- a minimizing sequence/trajectory that spends asymptotically vanishing time in the visible region where witnesses can occur,
- the reduction visibility condition that fails, and the finite reduced data witnessing the escape route.

Note. The obstruction outputs are finite and stratum-respecting: they live in the reduced graph/cost arena and do not require returning to the unreduced model.

F Proof skeletons and invariance lemmas

We record the proof dependencies used in the main text.

Invariance under cost coboundaries

If $c'(u \to v) = c(u \to v) + \phi(u) - \phi(v)$, then every cycle total and hence every cycle mean is unchanged. Therefore $\lambda_*(c)$, cycle-core $\mathcal{C}(c)$, and the unique-cycle chamber hyperplane arrangement are invariant under coboundary changes.

Projective gauge invariance

Projective contraction statements in the Hilbert metric are invariant under multiplication by positive scalar cocycles and under coordinate scalings that preserve the cone.

Boundary theorem skeleton

Theorem 5 follows a standard structure:

- Gate B1 produces a block map with $\tau < 1$ (Appendix C).
- Gate B3 produces a linear-in-n count of disjoint contracting blocks along minimizers (Appendix C).
- Gate B2 ensures this holds on every minimizing SCC, giving uniformity.
- Gate B4 prevents minimizers from avoiding the witness region and collapsing the frequency estimate.

Conversely, if PR holds, then each gate must hold: PR forces the existence of contraction events at some scale (B1), forces nonzero contraction frequency (B3), forces coverage across all minimizing components (B2), and prohibits escape routes that suppress the frequency (B4). Each failure yields a finite obstruction certificate (Appendix E).

Genericity theorem skeleton

Theorem 10 is finite-dimensional: for distinct cycles γ
neq
gamma$'$, the tie set $\mu_{gamma}(c) =$
cycmeangamma$'c$ is a proper affine hyperplane in
mathbbR^{E}. Taking the finite union over all pairs yields a closed measure-zero nowhere-dense exceptional set.

Engineers-kit note. The attached markdown file `appendix/PROOF_SKELETONS.md` contains additional proof-outline material used during development; it is included for completeness but is not required for the main paper.

G Boundary partition of PR

Theorem 5 induces a clean partition of the PR boundary into disjoint failure types. The purpose of this appendix is to make the logical dependency flow explicit and to reduce redundancy in discussion.

Gate	Status type	Interpretation on the minimizing set
B1	contraction failure	no certified primitive/Doeblin block, hence no projective contraction event
B2	coverage failure	some minimizing SCC is not controlled by any witness family
B3	frequency failure	witness blocks occur sublinearly along minimizers (zero return density)
B4	escape failure	minimizers avoid the visible/witness region at asymptotic scale (excursion/visibility leak)

Table 1: Disjoint PR failure modes by contract gate.

Each failure mode admits a finite obstruction certificate (Appendix E). When all gates hold, the framework certifies PR and returns an explicit lower bound on the contraction exponent (Corollary 8).

H Strict genericity: full proof in the finite reduced parameter space

This appendix records a complete proof of Theorem 10 with explicit parameter-space hypotheses and the structural stability clause.

H.1 Finite reduced parameter space

Fix a finite directed graph $G = (V, E)$ produced by the reduction layer. Work in the edge-cost coordinate space

$$\mathcal{P} := \mathbb{R}^{E},$$

with Euclidean topology and Lebesgue measure. If the reduction restricts to an admissible region (positivity floors, bounded ratios, strict inequalities), denote it by $\mathcal{P}_{\text{adm}} \subset \mathcal{P}$ and assume it is open.

H.2 Cycle means are continuous linear functionals

For any directed cycle γ in G with length $|\gamma|$, define its mean cost

$$\mu_\gamma(c) := \frac{1}{|\gamma|} \sum_{e \in \gamma} c_e.$$

This is a linear functional in c and hence continuous on $\mathcal{P}$.

A standard finite reduction observation is that the minimum cycle mean is attained by at least one *simple* directed cycle. Since G is finite, there are finitely many simple directed cycles; let $\mathcal{S}$ denote this finite set.

H.3 Generic uniqueness: open, dense, full measure

For distinct simple cycles $\gamma \neq \gamma'$, define the tie set

$$H_{\gamma,\gamma'} := \{c \in \mathcal{P} : \ \mu_\gamma(c) = \mu_{\gamma'}(c)\}.$$

Because $\mu_\gamma(\cdot) - \mu_{\gamma'}(\cdot)$ is a nonzero linear functional, $H_{\gamma,\gamma'}$ is a proper affine hyperplane in $\mathcal{P}$. Hence it is closed, has empty interior, and has Lebesgue measure zero.

Define the non-unique minimizer set

$$\mathcal{N} := \bigcup_{\gamma \neq \gamma' \in \mathcal{S}} H_{\gamma,\gamma'}.$$

This is a finite union of hyperplanes, so $\mathcal{N}$ is closed, nowhere dense, and measure zero. Therefore its complement

$$\mathcal{U} := \mathcal{P} \setminus \mathcal{N}$$

is open, dense, and full measure. Intersecting with $\mathcal{P}_{\text{adm}}$ yields an open–dense/full-measure chamber set $\mathcal{U}_{\text{adm}}$ in the admissible region.

H.4 Structural stability: local constancy of the minimizing cycle-core

Fix $c \in \mathcal{U}$. Then there is a unique simple cycle $\gamma_*(c) \in \mathcal{S}$ minimizing $\mu_\gamma(c)$. Define the strict margin

$$\Delta(c) := \min_{\gamma \in \mathcal{S} \setminus \{\gamma_*(c)\}} \big(\mu_\gamma(c) - \mu_{\gamma_*}(c)\big) \ > \ 0.$$

By continuity of the finitely many $\mu_\gamma(\cdot)$, there exists $\varepsilon > 0$ such that whenever $\|c' - c\| < \varepsilon$ we have

$$\mu_{\gamma_*}(c') < \mu_\gamma(c') \quad \text{for all } \gamma \neq \gamma_*.$$

Thus $\gamma_*(c') = \gamma_*(c)$ on that neighborhood. Consequently, the cycle-core equals the single cycle γ_* throughout the neighborhood; uniqueness is structurally stable on $\mathcal{U}$.

H.5 Recurrence consequence

If γ_* has length p, the unique minimizing ergodic measure on the cycle-core is the periodic orbit measure. Any cylinder block that appears at least once per period has frequency at least $1/p$. This is the regime in which Gate B3 becomes structural once a witness block is known to lie on the core.

H.6 Lower-dimensional parametrizations

If costs are parametrized by $\theta \in \mathbb{R}^k$ through a continuous map $\theta \mapsto c(\theta)$, the same conclusion holds provided the family is nondegenerate: for each pair $\gamma \neq \gamma'$, the function $\theta \mapsto \mu_\gamma(c(\theta)) - \mu_{\gamma'}(c(\theta))$ is not identically zero on the admissible region. In the affine case $c(\theta) = A\theta + b$, this reduces to excluding parameterizations that collapse distinct cycle means into identical functionals.

I Witness synthesis details

This appendix expands the constructive witness procedure used in Theorem 13.

I.1 Cycle-core SCC extraction

Given the reduced graph $G = (V, E)$ and a parameter c, the cycle-core $\mathcal{C}(c)$ is determined by minimum mean cycle computation (e.g. Karp's algorithm on c). Restricting to $\mathcal{C}(c)$, compute the SCC decomposition

$$\mathcal{C}(c) = \bigsqcup_{j=1}^{J} \mathcal{C}_j(c).$$

I.2 Canonical witness selection

For each SCC $\mathcal{C}_j$ choose a directed cycle γ_j entirely contained in $\mathcal{C}_j$ (e.g. select the lexicographically first cycle discovered by a DFS on the SCC). Define the associated cycle word w_j as the edge list of γ_j.

Doeblin iterate. If the FTG ledger requires a fixed length N for a primitive/Doeblin witness block, replace w_j by an iterate $(w_j)^{k_j}$ so that $|(w_j)^{k_j}| \geq N$ and the corresponding block product contains the certified minorization event. The iterate is bounded in terms of N and the SCC size.

I.3 How the witness family enforces Gate B2

By construction, each SCC $\mathcal{C}_j$ contains its witness word w_j supported entirely inside $\mathcal{C}_j$. Therefore any minimizing invariant measure supported on $\mathcal{C}_j$ is a candidate for witnessing through w_j, and the family $\{w_j\}$ provides SCC-level uniformity.

I.4 Return density in the unique-cycle chamber

On the unique-cycle chamber set, $\mathcal{C}(c)$ is a single cycle γ_* of period p. If the chosen witness is a subword of γ_* (or an iterate of γ_*), then every minimizer is the periodic orbit measure and occurrences of the witness are forced with frequency at least $1/p$ (Gate B3 with $\rho \geq 1/p$).

Outside the unique-cycle chamber, Gate B3 requires an explicit symbolic density witness. A standard approach is to produce a finite-state density witness set (SDW) based on the induced subshift constraints; see Appendix C for how this density enters the quantitative exponent bound.

J Certification report schema (for reproducibility)

For reproducibility and referee clarity, it is helpful to standardize the certification output of Theorem 5. The following report schema is suggested; it is purely descriptive and does not add mathematical content.

Minimal report fields

- Reduced graph: (V, E) sizes and an explicit edge list.
- Parameter point: the edge-cost vector $c \in \mathcal{U}$ (or a parameter θ together with the map $\theta \mapsto c$).
- Cycle-core: list of critical cycles and SCC decomposition.
- Gate status: Boolean status for each gate B1–B4.
- If PR holds: $(N, \tau, \rho, \kappa_{\text{safe}})$ and the witness family used.
- If PR fails: blocking gate and a finite obstruction certificate.

Example JSON-like shape

```
{
  "graph": {"|V|": ..., "|E|": ..., "edges": [...]},
  "parameter": {"c": [...]},
  "cycle_core": {"cycles": [...], "sccs": [...]},
  "gates": {"B1": true/false, "B2": ..., "B3": ..., "B4": ...},
  "pr_if_true": {"N": ..., "tau": ..., "rho": ..., "kappa_safe": ...},
  "obstruction_if_false": {"blocking_gate": "B3", "certificate": {...}}
}
```

This schema makes the dependency flow explicit and keeps the universality claim honest: PR is certified if and only if the gates pass, and failures are accompanied by explicit finite witnesses.

K IV. Universality and Sharp Boundary (Breakthrough Upgrade)

This module makes the universality story referee-safe by proving two things that are actually true.

- Generic universality can be claimed only inside natural model classes where the critical minimizing core is structurally simple or conservativity is forced.
- Structural optimality can be claimed at the level of the paper's geometric output: uniform projective regularity on minimizers holds exactly when the contract gates that enable it hold.

The blocked direction (deriving a strict gap from minimal locality alone) is replaced by an obstruction lemma that explains why it is structurally impossible without extra hypotheses.

K.1 IV.A The dilution obstruction and what a strict gap really demands

The hybrid uses an excursion-safe SG notion because simple-cycle separation is not enough.

K.1.1 Lemma IV.A1 (Dilution kills any all-cycle gap inside a nonconservative critical SCC)

Let $G = (V, E)$ be a finite directed graph with edge costs $c : E \to \mathbb{R}$. Let λ_* be the minimum cycle mean.

Let $C \subseteq G$ be a strongly connected component that contains a cycle γ_* with mean λ_*. Assume C also contains a cycle γ with mean $\lambda_* + \alpha$ for some $\alpha > 0$.

Then for every $\varepsilon > 0$ there exists a directed cycle $\Gamma \subseteq C$ with

$$\lambda_* < \mathrm{mean}(\Gamma) < \lambda_* + \varepsilon.$$

Consequently, any "gap over all cycles" satisfies $\delta = 0$.

Proof. Because C is strongly connected, there exist directed paths $p_\to$ from a vertex of γ_* to a vertex of γ and $p_\leftarrow$ back, with fixed finite lengths and fixed finite total costs. For $k \geq 1$, form a cycle Γ_k that traverses k copies of γ_*, then follows $p_\to$, then traverses γ, then follows $p_\leftarrow$. Write

- L_* for the length of γ_*,
- L for the length of γ,
- B for the total length of $p_\to \cup p_\leftarrow$,
- S_* for the total cost of γ_*,
- S for the total cost of γ,
- R for the total cost of $p_\to \cup p_\leftarrow$.

Then

$$\mathrm{mean}(\Gamma_k) = \frac{kS_* + S + R}{kL_* + L + B} = \lambda_* + \frac{(S - \lambda_* L) + (R - \lambda_* B)}{kL_* + L + B}.$$

Since $S - \lambda_* L = \alpha L > 0$ and $R - \lambda_* B$ is fixed, the excess term is positive and tends to 0 as $k \to \infty$. Choose k large enough so the excess is smaller than ε. $\square$

K.1.2 Corollary IV.A2 (Closure is not optional if $\delta > 0$ is demanded)

Fix a critical SCC C. Shift costs by $c'(e) = c(e) - \lambda_*$.

- If C contains any cycle with strictly positive total c'-weight, then the all-cycle gap collapses to $\delta = 0$ by Lemma IV.A1.
- Therefore, any excursion-safe SG must force: every cycle in C has total c'-weight 0.

This is the conservative / coboundary test used in `SG_CERTIFICATION.md`.

K.1.3 Corollary IV.A3 (Why raw "SG is generic" is false in the ambient cost space)

On a fixed strongly connected graph that contains at least two distinct directed cycles, the property

- "the shifted cost is conservative on the critical SCC"

is rigid and fails under generic perturbations of c. By Lemma IV.A1, that immediately forces $\delta = 0$.

Any generic universality statement must therefore be scoped to a model class that prevents this obstruction, either structurally or by an optimality forcing mechanism.

K.2 IV.B Generic universality inside natural model classes

The right generic theorem is not "strong SG is generic." The right generic theorem is:

- once closure is built in, the remaining contract failures are essentially tie sets, hence meagre / measure zero.

K.2.1 Definition IV.B0 (Closure-built parameter classes)

A parameter class $\mathcal{P} \subseteq \mathbb{R}^E$ is *closure-built* if whenever an SCC is minimizing in $\mathcal{P}$, the shifted cost on that SCC is automatically conservative.

Standard closure-built classes include:

- PC-by-structure: the minimizing SCC is a directed cycle, so every cycle is a multiple of it and conservativity is automatic.
- Coboundary families: costs are parametrized as $c(e) = \psi(e) + \varphi(\mathrm{head}(e)) - \varphi(\mathrm{tail}(e))$ with ψ fixed and φ free.
- Reduction-output families: the reduction mechanism outputs costs already cohomologous to a constant on the selected minimizing core.

K.2.2 Theorem IV.B1 (Generic UC and positive SCC separation in closure-built classes)

Fix a finite directed graph G and a closure-built parameter class $\mathcal{P} \subseteq \mathbb{R}^E$ whose parametrization is finite-dimensional and piecewise affine.

Then, on a residual subset of $\mathcal{P}$ and on a full-measure subset with respect to Lebesgue measure in parameter coordinates, the following hold.

- UC: there is a unique minimizing SCC.
- Positive SCC separation: the minimum mean in the next-best SCC exceeds λ_* by a strict positive amount.

A fully explicit open–dense / full-measure version of the "tie sets are hyperplanes" argument in the edge-cost coordinate space is recorded in:

- `IV_GENERICITY_STRICT_PARAMETER_SPACE.md`

Mechanism. Equality of the minimum SCC means is a finite collection of affine equalities on each region where the identity of the minimizing cycle in each SCC is fixed. Hence ties lie in a finite union of codimension-one sets, which are meagre and measure zero.

K.2.3 Theorem IV.B2 (Generic periodic-core recurrence in standard symbolic classes)

In standard symbolic settings where the reduced problem is a continuous potential on a transitive SFT, generic potentials in natural function spaces select periodic minimizing measures.

When that happens:

- PC holds on the selected minimizing support, hence the recurrence bound $\mu([w]) \geq 1/L$ is automatic for any witness word w contained in the periodic orbit.
- The hybrid contract reduces to: Doeblin witness existence plus the finite SG/UC checks.

This is the clean "probabilistic universality" route: contract violations become nongeneric in the function-space sense, once the reduction places the variational layer in a classical ergodic-optimization setting.

Finite-parameter version (the regime of this paper). In the reduced interface here, the variational object is an edge-cost function on a finite directed graph, i.e. a locally constant potential on an edge shift. In that finite-dimensional space, periodic minimizing measures are automatic and uniqueness is generic (ties are codimension-one). The remaining universal ingredient is therefore a *cycle-core* (PC-by-structure) condition on the closure-built class so that the minimizing support is a single directed cycle.

A proof-ready statement of this "ρ_0 becomes explicit + contract is generic in cycle-core classes" upgrade is isolated in:

- `engineers_kit/docs/IV_GENERICITY_RHO0_AND_CONTRACT.md`

K.2.4 Theorem IV.B3 (Cycle-core can be extracted from optimality in the finite reduced graph)

In the finite reduced variational layer, assume a calibrated Bellman potential u exists (equivalently: an additive eigenvector for the min-plus operator at level λ_*). Then there exists a deterministic calibrated selector σ whose orbit is eventually periodic. Its periodic tail is a minimizing directed cycle C_*.

Consequences:

- A periodic minimizing measure exists without any "PC-by-structure" hypothesis.
- Along that periodic minimizer, the recurrence density is explicit: $ho_0 \geq 1/|C_*|$ for any witness word that appears on the cycle (or by using the one-lap product as the witness block).
- Therefore, any existence-level PR claim can be reduced to checking a single extracted minimizing cycle.

This extraction step is formalized in:

- `engineers_kit/docs/IV_CYCLE_CORE_FROM_REDUCTION.md`

K.2.5 Remark IV.B4 (What still needs an extra gate)

The hybrid contract's PR conclusion is uniform over minimizing measures. To upgrade the extracted-cycle statement to a uniform claim, you still need either:

- PC-by-structure (all minimizers collapse to that cycle),
- SD / Assumption R (a uniform return-density bound), or
- a uniqueness/selection theorem showing contraction forces the minimizing set to shadow the extracted cycle.

Literature anchor (optional, for context). In classical ergodic optimization, generic periodicity of ground states is proved in expanding-map settings (e.g. Contreras, *Invent. Math.* 205 (2016), 383–412) and related genericity results are surveyed by Jenkinson.

K.3 IV.C Structural optimality: the contract is the boundary of projective regularity

The strongest upgrade is an if-and-only-if theorem stated at the level of the geometric conclusion.

K.3.1 Definition IV.C0 (PR: uniform projective regularity on minimizers)

PR holds if there exist constants $\kappa > 0$ and $\tau \in (0, 1)$ such that for every minimizing ergodic measure μ on the selected core and for μ-almost every ω, the projective action contracts at a uniform exponential rate

$$d_H(A^{(n)}(\omega)x, A^{(n)}(\omega)y) \le C(\omega)\, e^{-\kappa n}\, d_H(x, y) \quad \forall x, y \in K_S^\circ,\ \forall n \ge 1,$$

where d_H is Hilbert's projective metric on the invariant face.

K.3.2 Theorem IV.C1 (Sufficiency: contract gates imply PR)

This is the bridge theorem pipeline already implemented in the kit.

- FTG plus a certified primitive / Doeblin word w yields a block contraction coefficient $\tau < 1$.
- SG localizes minimizers to the critical SFT Σ_*.
- A recurrence gate yields a uniform lower bound $\mu([w]) \ge \rho_0$ for every minimizing ergodic μ.
- PR follows with κ proportional to $\rho_0|\log \tau|$, with overlap-safe normalization.

K.3.3 IV.C necessity layer (where the "boundary" comes from)

The sufficiency direction is the bridge pipeline already implemented in the kit.

The necessity direction is where "structural optimality" lives: uniform projective regularity cannot hold unless the gates that generate contraction events exist and recur with a uniform positive density across the minimizing measures.

A proof-ready necessity package with explicit failure-mode constructions is provided in:

- `engineers_kit/docs/IV_BOUNDARY_NECESSITY_FAILURE_MODES.md`

The key implications used in the paper draft are the following.

K.3.4 Lemma IV.C2 (Doeblin existence is necessary)

If uniform projective regularity holds with exponent $\kappa > 0$ on the minimizing core, then there exists a finite word w whose block product has finite projective diameter on the face, equivalently w is an S-Doeblin witness with $\tau_S(A_w) < 1$.

K.3.5 Lemma IV.C3 (Recurrence is necessary for a uniform exponent)

Fix an S-Doeblin witness w with $\tau_S(A_w) < 1$. If there exists a minimizing ergodic measure μ with $\mu([w]) = 0$, then no uniform positive exponent $\kappa > 0$ can hold on minimizers.

This is the necessity side of the kit's recurrence gate, implemented either as PC or as SD.

K.3.6 Lemma IV.C4 (Excursion-safe SG is necessary to prevent vanishing Doeblin density)

If the reduced variational layer admits near-minimizers whose itineraries avoid w at arbitrarily small cost penalty, then one can build minimizing limits with arbitrarily small Doeblin frequency, forcing the best uniform exponent to drop to 0.

Excursion-safe SG is the finite certificate that blocks this "cheap avoidance" mechanism.

K.3.7 Theorem IV.C5 (Sharp boundary statement)

A single boxed equivalence statement with explicit constants (ready to insert verbatim) is provided in:

- `engineers_kit/docs/IV_BOXED_THEOREM_PR_IFF_CONTRACT.md`

A constant-closed specialization of the exponent (eliminating τ in favor of (d, m_*, M_*, N) via the conservative FTG bounds) is recorded in:

- `engineers_kit/docs/IV_CONSTANT_CLOSURE_ADDENDUM.md`

In the finite reduced setting of the kit, PR holds on the minimizing core if and only if the following explicit gates hold.

- FTG on the cocycle templates and bounds.
- Existence of an S-Doeblin witness block.
- A recurrence lower bound for that witness, supplied either by PC, by SD, or by an externally proved ρ_0.
- Excursion-safe SG, so near-minimizing measures cannot avoid the witness at vanishing penalty.

This is the structural optimality upgrade: every failure of PR corresponds to a named contract violation, and every contract violation has an explicit failure mode construction.

K.4 IV.D Optimality forces closure on the minimizing core

The closure gate in SG is not a random coincidence. It is forced by calibrated optimality on the minimizing subsystem.

K.4.1 Lemma IV.D1 (Calibrated equality implies a coboundary identity on tight edges)

Consider the min-plus dynamic programming operator on vertices

$$(Tu)(v) = \min_{e:v\to w}(c(e) + u(w)).$$

Suppose there exist $u : V \to \mathbb{R}$ and $\lambda_* \in \mathbb{R}$ such that

- $Tu = u + \lambda_*$ on the minimizing subsystem,
- $Tu \geq u + \lambda_*$ on all vertices.

If an edge $e : v \to w$ attains the minimum in the definition of $Tu(v)$, then

$$c(e) - \lambda_* = u(v) - u(w).$$

Consequently, on the subgraph of tight edges, $c - \lambda_*$ is a vertex coboundary.

K.4.2 Corollary IV.D2 (Closure on the calibrated minimizing subgraph)

Any directed cycle composed of tight edges has total shifted cost 0. In particular, on the calibrated minimizing core, every cycle has mean exactly λ_*.

This is the conceptual reason the SG closure gate belongs in the engine: it is the finite certificate that the minimizing core is calibrated and immune to dilution.

K.5 IV.E Literature anchors for referees

These are standard places to point a referee so the bridge story reads as a natural synthesis rather than an isolated trick.

- P. J. Bushell, "Hilbert's metric and positive contraction mappings in a Banach space," Archive for Rational Mechanics and Analysis 52 (1973), 330–338.
- H. Furstenberg and H. Kesten, "Products of Random Matrices," Annals of Mathematical Statistics 31 (1960), 457–469.
- P. Bougerol and J. Lacroix, Products of Random Matrices with Applications to Schrödinger Operators, Progress in Probability, Birkhäuser, 1985.
- P. Walters, "A generalized Ruelle–Perron–Frobenius theorem and some applications," Astérisque 40 (1976), 183–192.
- O. Jenkinson, "Ergodic Optimization," Discrete and Continuous Dynamical Systems 15 (2006), 197–224.
- G. Contreras, "Ground States are generically a periodic orbit," preprint 2013.
- W. Parry and M. Pollicott, Zeta functions and the periodic orbit structure of hyperbolic dynamics, Astérisque 187–188 (1990).

K.6 Pass 23: cycle-core PR rate is the default "universality" statement

The strongest genuinely universal statement available without extra mixing assumptions is now the cycle-core statement:

- the *global* SG gate can fail by dilution without affecting minimizers
- the *cycle-core* is invariant under dilution and captures minimum-mean structure
- if a Doeblin witness occurs on a core cycle, the kit produces an explicit minimizer-facing exponent $\kappa_{\text{core,pc}}$
- if SD/SD-W holds on the cycle-core, the kit produces a uniform exponent $\kappa_{\text{core,sd}}$ for all minimizing measures supported on the core

This is exactly the "optimality induces projective regularity" mechanism in finite form: reduction identifies the recurrent aperiodic core; a Doeblin block gives uniform projective contraction on that core; recurrence density upgrades contraction to an exponential rate along minimizers.

See `IV_BOXED_THEOREM_CORE_PR_IFF_CORE_CONTRACT.md` for the boxed statement and the report fields.

K.7 Pass 25: lifting cycle-core rates to *all* minimizing measures

Pass 23 already gives minimizer-facing rates on the cycle-core. Pass 25 closes the final logical gap:

- any minimizing invariant measure on a finite graph is supported on the union of λ_*-cycles,
- therefore any cycle-core contraction certificate is automatically a statement about minimizers,
- if the cycle-core has multiple SCCs, a uniform statement over **all** minimizers requires SD/SD-W (or the chosen recurrence witness) on **each** SCC.

The selection/return lift is isolated in:

- `IV_LIFT_FROM_CORE_TO_MINIMIZERS.md`

and implemented in the report as a component-wise minimizer block that aggregates the worst-case rate when possible.

K.8 IV.F Generic Universality Corollary (what Pass 27 makes possible)

Pass 27 upgrades the program from conditional classification to prevalence.

- In the reduced finite parameter space, strict tie conditions live on lower-dimensional boundaries.
- Away from those boundaries, the minimizing cycle-core is typically a single simple directed cycle.

In that single-cycle regime, the remaining PR gate is purely about whether the Doeblin witness intersects the minimizing cycle. If it does, the periodic minimizing measure yields an explicit density $\rho_{pc} \geq 1/p$ and therefore an explicit exponent.

The strongest closed universality posture in the kit is obtained by also certifying a recurrence gate that removes witness-avoidance across the full reduced graph.

- The report now includes an optional `syndetic_doeblin_global` certificate (SD or SD-W on the full reduced graph).
- When it holds, the witness set is cycle-covering in particular, so every minimizing cycle in every strict-genericity chamber contains a witness.
- Hence PR holds on a full-measure open dense subset of the admissible region, and failure occurs only on the certified contract-violation boundaries.

See `IV_GENERIC_UNIVERSALITY_COROLLARY.md` for the submission-ready formulation.

K.9 Pass 29 addendum: the covering gate

After strict genericity, the only remaining universality obstruction is *witness placement*: if minimizers split across multiple cycle-core SCCs, uniform rates require the witness set to be cycle-core covering (one component-compatible witness per core SCC). The kit now certifies this gate and emits explicit uncovered components as finite obstructions.

L IV.C′ Boxed Boundary Theorem (PR $\iff$ Contract Gates)

This file gives a single, paper-ready theorem statement that can be inserted as the "boxed" culmination of Section IV. It consolidates the sufficiency pipeline (bridge theorem) and the necessity layer into one equivalence with explicit constants.

L.1 Theorem IV.C′ (Uniform projective regularity on minimizers is equivalent to the hybrid contract gates)

Work in the finite hybrid setting produced by the reduction:

- a finite reduced additive graph $G = (V, E, c)$ and its induced critical subshift Σ_*,
- an FTG template layer providing a face K_S and overlap-safe normalization data,
- a family of cone-preserving matrices $A(\omega_0)$ indexed by symbols $\omega_0 \in \mathcal{A}$ (the boundary alphabet),
- the Hilbert metric d_H on K_S°.

Define **PR (uniform projective regularity on minimizers)** as the existence of $\kappa > 0$ such that for every minimizing ergodic measure μ on Σ_*, for μ-almost every ω, and for all $x, y \in K_S^\circ$,

$$d_H\big(A^{(n)}(\omega)x,\ A^{(n)}(\omega)y\big) \leq C(\omega)\, e^{-\kappa n}\, d_H(x, y) \quad \forall n \geq 1.$$

Then the following are equivalent.

- **(A) PR holds on minimizers** (uniformly over all minimizing ergodic measures).
- **(B) The hybrid contract gates hold**, meaning there exist explicit finite data $(S, \mathcal{W})$ and a uniform density lower bound such that
- **FTG is in force on the face** K_S (fixed $d, \mathcal{T}, m_*, M_*$, overlap-safe normalization).
- **Doeblin witness on** K_S**:** there exists a finite Doeblin event. Concretely, either
- a **single word** w such that the block product A_w preserves K_S and has finite projective diameter, equivalently $\tau_S(A_w) = \tau < 1$, or
- a **finite set of words** $\mathcal{W}$ (all of a common edge-length N) such that each $w \in \mathcal{W}$ is a Doeblin event on K_S with contraction coefficient $\tau < 1$.
- **Recurrence gate (uniform density of Doeblin events):** every minimizing ergodic measure μ sees Doeblin events with a uniform positive lower frequency.

In the engineer's kit this is supplied either by

- a **PC** certificate (minimizers are periodic on a cycle-core, giving an explicit return density), or
- an **SD** certificate (the witness word is syndetic in the minimizing language), or
- an **SD-W** certificate (the witness set is syndetic: every long enough path hits at least one word in $\mathcal{W}$), or
- an external argument establishing the stated density.
- **Excursion-safe SG gate:** the reduced variational layer is certified so that near-minimizers cannot avoid Doeblin events at vanishing cost penalty; concretely, the SG certificate blocks sequences of invariant measures μ_n with $\int c\, d\mu_n \downarrow \lambda_*$ and Doeblin-event frequency $\downarrow 0$.

Moreover, when (B) holds, PR holds with an explicit exponent.

L.2 Explicit exponent (constant-closed statement)

Let $\tau \in (0,1)$ be the certified Birkhoff contraction coefficient for one Doeblin event on K_S°. Let $\rho_{\text{dis}} > 0$ be a uniform lower bound on the **density of disjoint Doeblin events per unit time** along every minimizing ergodic measure.

Then one can take

$$\kappa \;=\; \rho_{\text{dis}} \left|\log \tau\right| .$$

Interpretation.

- $|\log \tau|$ is the contraction strength of one Doeblin event.
- ρ_{dis} is the guaranteed asymptotic number of disjoint Doeblin events per step.

In the kit:

- for **PC**, ρ_{dis} is the certified return density (typically 1/period);
- for **SD/SD-W**, the kit provides a conservative ρ_{dis} computed from the minimal bounded-gap constant N_{gap} and the shortest witness length.

L.3 FTG-only closure of τ (using only (d, m_*, M_*, N))

When the Doeblin event has edge-length N and the overlap-safe normalization constants from FTG are in place, one can bound τ explicitly using only (d, m_*, M_*, N).

From the standard Doeblin floor and uniform upper bounds,

$$\varepsilon := m_*^N, \qquad U := (dM_*)^N, \qquad \frac{U}{\varepsilon} = \Big(\frac{dM_*}{m_*}\Big)^N =: r > 1.$$

This yields a Hilbert diameter bound and a Birkhoff contraction bound

$$\Delta \le 2\log(r) = 2N\,\log\Big(\frac{dM_*}{m_*}\Big), \qquad \tau \le \tanh(\Delta/4) = \frac{r-1}{r+1}.$$

Therefore a fully explicit safe exponent is

$$\kappa_{\mathrm{safe}} \;:=\; \rho_{\mathrm{dis}}\,\log\Big(\frac{r+1}{r-1}\Big) \;=\; \rho_{\mathrm{dis}}\,\big|\log\big(\tfrac{r-1}{r+1}\big)\big|.$$

- This uses only conservative FTG bounds.
- If the transcript provides a sharper measured τ for the actual witness block, replacing $(r-1)/(r+1)$ by that τ improves κ immediately.

L.4 Proof spine (one paragraph per direction)

L.4.1 (B) $\Rightarrow$ (A) (sufficiency)

- FTG + the Doeblin witness yields a block contraction coefficient $\tau < 1$ on K_S° and overlap-safe normalization.
- The recurrence gate gives a uniform lower bound $\rho_{\mathrm{dis}} > 0$ on the density of disjoint Doeblin events under every minimizing ergodic measure.
- Counting disjoint occurrences yields the exponential rate $\kappa = \rho_{\mathrm{dis}}|\log\tau|$.

This is the implemented bridge pipeline (see `BRIDGE_THEOREM_OPTIMALITY_FORCES_REGULARITY.md` and `engineers_kit/PROOF_SKELETONS.md`).

L.4.2 (A) $\Rightarrow$ (B) (necessity)

- If PR holds with $\kappa > 0$, then along every minimizing ergodic measure there must exist strict contraction events on K_S°. By finiteness of words of fixed length, at least one finite word must realize such an event on a positive-measure set; this forces the existence of a Doeblin witness (a word with $\tau_S(A_w) < 1$).
- If some minimizing ergodic measure μ had zero Doeblin-event density, then no uniform positive exponent could be sustained. Hence a uniform positive lower bound on Doeblin-event frequency is necessary, and the kit isolates robust ways to certify it (PC or SD/SD-W or an external ρ_{dis}).
- Finally, if near-minimizers could avoid Doeblin events at arbitrarily small variational penalty, one can build minimizing limits with arbitrarily small Doeblin frequency, forcing the best uniform exponent to drop to 0. The excursion-safe SG gate is exactly the finite certificate blocking this mechanism.

A proof-ready necessity suite with explicit failure-mode constructions is in `engineers_kit/docs/IV_BOUNDARY_`

M IV.C Necessity Layer — Boundary Equivalence and Explicit Constructions

This module upgrades the "structural optimality" route into a referee-safe necessity story.

The goal is to justify the sharp boundary claim:

- uniform projective regularity on the minimizing core holds **only if** the hybrid gates that enable it hold,
- and each gate failure yields an explicit failure-mode construction that forces the best uniform exponent to drop to zero.

This module is designed to be cited by `engineers_kit/docs/IV_UNIVERSALITY_BREAKTHROUGH.md` and cross-linked from `BRIDGE_THEOREM_OPTIMALITY_FORCES_REGULARITY.md`.

M.1 Setup: projective contraction data on a face

Fix a face K_S of the nonnegative cone and its relative interior K_S°. Let d_H denote Hilbert's projective metric on K_S°.

For a linear map A that preserves K_S, define its projective diameter on the face:

$$\Delta_S(A) = \sup_{x,y \in K_S^\circ} d_H(Ax, Ay) \in [0, \infty].$$

Define the corresponding Birkhoff contraction coefficient:

$$\tau_S(A) = \begin{cases} \tanh(\Delta_S(A)/4)\&\Delta_S(A) < \infty, \\ 1\&\Delta_S(A) = \infty. \end{cases}$$

When $\Delta_S(A) < \infty$, the classical Birkhoff inequality gives

$$d_H(Ax,Ay) \le \tau_S(A)\, d_H(x,y) \quad \text{for all } x,y \in K_S^\circ.$$

The Engineer's Kit uses the term **S-Doeblin** for "$\Delta_S(A) < \infty$" because, on a face, finite projective diameter is equivalent to a uniform Doeblin minorization (positivity on S plus uniform comparability).

M.2 Lemma C.N1 (Finite diameter is the Doeblin event on a face)

Let A preserve K_S.

- If $\Delta_S(A) < \infty$, then A maps K_S° into a subset of K_S° with uniformly bounded coordinate ratios on S, hence A satisfies the S-Doeblin condition used by the certification layer.
- If A fails the S-Doeblin condition, then $\Delta_S(A) = \infty$ and $\tau_S(A) = 1$.

This equivalence is the reason the kit treats "Doeblin witness exists" as the primitive projective contraction gate.

M.3 Lemma C.N2 (Counting Doeblin blocks controls the exponent)

Fix a finite word w with A_w preserving K_S and $\tau_S(A_w) = \tau < 1$.

For any admissible concatenation that contains k **disjoint** copies of w, the associated product $A^{(n)}$ satisfies the bound

$$d_H(A^{(n)}x, A^{(n)}y) \;\le\; C\,\tau^k\, d_H(x,y),$$

where C depends only on the overlap-safe renormalization constants already tracked by FTG.

Consequently, along a trajectory ω, the best achievable exponential rate is bounded in terms of the asymptotic frequency of w:

$$\limsup_{n\to\infty} \frac{1}{n} \log \frac{d_H(A^{(n)}x, A^{(n)}y)}{d_H(x,y)} \;\le\; \Big(\limsup_{n\to\infty} \frac{k_n(\omega)}{n}\Big) \log \tau,$$

where $k_n(\omega)$ is the maximum number of disjoint w-occurrences inside $\omega_{0:n-1}$.

M.4 Proposition C.N3 (Recurrence lower bounds are necessary for a uniform positive exponent)

Assume there is a Doeblin witness w with $\tau_S(A_w) = \tau < 1$.

If there exists an invariant ergodic measure μ on the minimizing subshift with $\mu([w]) = 0$, then uniform projective regularity on minimizers is impossible.

- Along μ-typical trajectories, w never occurs.
- By Lemma C.N2, no exponential contraction contribution can be extracted from w.
- Any uniform lower bound $\kappa > 0$ on the exponent would force infinitely many strict contraction events, contradicting the absence of w.

This is the necessity side of the recurrence gate in the kit.

M.5 Proposition C.N4 (Two ways recurrence can be forced, and why the kit isolates them)

Uniform projective regularity is a statement that must hold for **every** minimizing ergodic measure μ.

There are only two robust mechanisms that can force a positive lower bound on Doeblin frequency across all such μ.

- **Periodic core (PC).** The minimizing simplex collapses to a single periodic orbit. Every minimizing measure is that orbit measure. Any witness word w occurring on the orbit satisfies a uniform frequency bound automatically.

- **Syndetic Doeblin (SD).** The witness word w is unavoidable in the minimizing language: every admissible path of a fixed length contains w. This implies a uniform cylinder lower bound $\mu([w]) \geq 1/N$ for every invariant μ.

Without PC or SD, a transitive SFT typically supports ergodic measures that avoid a given fixed word, so $\inf_\mu \mu([w])$ drops to 0, and the best uniform exponent drops with it.

This is why the kit treats SD as a separate, decidable gate and offers PC as an alternative recurrence certifier.

M.6 Lemma C.N5 (Doeblin existence is necessary for any positive uniform exponent)

Assume uniform projective regularity holds on the minimizing subsystem with exponent $\kappa > 0$.

Then there exists a finite word w that occurs with positive frequency on every minimizing ergodic measure and whose block product satisfies $\tau_S(A_w) < 1$.

Mechanism sketch that can be converted to a full proof in the paper draft.

- Uniform exponent $\kappa > 0$ implies that for sufficiently large n, a positive fraction of length-n blocks along μ-typical sequences must be strictly contracting in Hilbert metric.

- On a finite alphabet shift, there are only finitely many length-n words. Therefore at least one word w of length n must carry a strict contraction event on a set of positive μ-measure.

- By Lemma C.N1, strict projective contraction on K_S° forces $\Delta_S(A_w) < \infty$, which is exactly the S-Doeblin witness property.

This is the necessity side of the “Doeblin witness exists” gate.

M.7 Lemma C.N6 (Excursion-safe SG is necessary to prevent vanishing Doeblin frequency)

Assume there is a Doeblin witness w with $\tau_S(A_w) < 1$.

If the reduced variational layer admits a sequence of invariant measures μ_n with

- $\int c \, d\mu_n \downarrow \lambda_*$,

- and $\mu_n([w]) \downarrow 0$,

then no uniform positive projective exponent can hold on minimizers.
This is the exact obstruction SG is designed to block in an excursion-safe way.

- If near-minimizers can "buy" avoidance of the Doeblin witness at arbitrarily small variational penalty, then one can realize minimizing limits with arbitrarily rare contraction events.
- By Lemma C.N2, the best achievable exponent scales with the Doeblin frequency, forcing the uniform lower exponent bound to drop to 0.

In the Engineer's Kit, excursion-safe SG supplies the missing implication:

- near-minimizers that avoid the witness pay a strict penalty,
- hence minimizing measures cannot have vanishing Doeblin frequency.

M.8 The boundary statement in the finite hybrid setting

On the finite reduced graph and induced critical SFT Σ_* used by the kit, uniform projective regularity on minimizers is equivalent to the conjunction of explicit gates.

- FTG, to fix a face and ensure normalization / overlap safety.
- Doeblin witness existence, to supply at least one strictly contracting block on the face.
- A recurrence certifier (PC or SD or an external lower bound), to prevent measures with $\mu([w]) = 0$ or $\mu([w])$ arbitrarily small.
- Excursion-safe SG, to prevent "cheap avoidance" of the witness by near-minimizers.

Every failure of projective regularity can be traced to a violation of one of these gates, and each gate violation has a constructive failure mode documented in `engineers_kit/docs/FAILURE_MODES.md`.
For the *engineering* version of these failure modes (finite, copy/paste witnesses emitted by the kit), see `engineers_kit/docs/IV_OBSTRUCTION_CERTIFICATES.md` and the implementation in `engineers_kit/code/obstruction_certificates.py`.

N IV. Automatic Cycle-Core Witness Set Synthesis (Pass 30)

This note makes the "witness covering" requirement constructive.
Earlier passes isolated the correct universality gate:

- minimizers live on the **cycle-core SCCs**
- to certify **uniform-over-minimizers** contraction, the witness family must intersect **every** cycle-core SCC

Pass 29 turned that requirement into a checkable certificate.
Pass 30 goes one step further:

- whenever the cycle-core is nonempty, the kit *constructs* a canonical finite witness set that covers it.

This is a graph-theoretic lemma, independent of any mixing or positivity assumption.

N.1 Cycle-Core Witness Synthesis Lemma

Let $G = (V, E)$ be a finite directed graph. Let $\mathcal{C} \subseteq V$ be the cycle-core vertex set at λ_*, and let

- $\mathcal{C} = \bigsqcup_{j=1}^{J} \mathcal{C}_j$

  be its decomposition into strongly connected components (SCCs).
  Then for each j, the induced subgraph on $\mathcal{C}_j$ contains a directed cycle.
  Therefore, there exists a finite family of cycle words

- $W_{\mathrm{cyc}} = \{w_j\}_{j=1}^{J}$

  such that each w_j is a directed cycle word lying entirely inside $\mathcal{C}_j$.
  Consequently, W_{cyc} is automatically **cycle-core covering**.

N.2 What the kit actually constructs

For each cycle-core SCC $\mathcal{C}_j$, the kit synthesizes a *short* directed cycle as a vertex-word.

- It searches over directed edges $u \to v$ inside $\mathcal{C}_j$
- It closes the edge by a shortest directed path $v \Rightarrow u$ inside $\mathcal{C}_j$ (BFS)
- It returns the cycle word
- $[u, v, \ldots, u]$

  Tie-breaking is deterministic:

- shorter cycles are preferred
- remaining ties break lexicographically

  The output is stable under reruns and suitable for use as a reproducibility artifact.

N.3 How this is used in the universality pipeline

These synthesized cycle words are **symbolic** recurrence witnesses.
They are used by the bounded-gap recurrence gates:

- SD (single word)
- SD-W (word family)

Once your reduction layer supplies an actual Doeblin block length N or a certified Doeblin word family W, these recurrence gates become fully checkable on the minimizing subsystem.

The purpose of Pass 30 is to remove the last "hand-chosen witness" dependency from the **cycle-core covering** step.

N.4 What this does and does not prove

What it gives:

- a canonical, finite family of words W_{cyc} that intersects every cycle-core SCC
- an automatic covering certificate (Pass 29 gate becomes constructive)
- optional SD-W checks on the cycle-core for W_{cyc}
- componentwise minimizer recurrence certificates driven by W_{cyc}

What it does not give:

- it does not certify that these words are Doeblin blocks at the matrix level

That matrix-level property remains a reduction/positivity input.
Pass 30 isolates the remaining burden cleanly:

- the only non-structural step is validating that each minimizing SCC contains at least one **true** Doeblin witness.

N.5 Report outputs

The JSON report now includes:

- `cycle_core.synthesized_cyclecore_witness`
- synthesized words and per-component synthesis summaries
- `cycle_core.synthesized_covering_cycle_core`
- covering certificate for the synthesized set
- `cycle_core.synthesized_syndetic_doeblin_cycle_core`
- SD-W certificate on the cycle-core for the synthesized set (when computable)
- `minimizers.autosynth`
- componentwise minimizer recurrence certificates using the synthesized words
- a uniform ρ lower bound when all components certify

References

[1] G. Birkhoff. Extensions of Jentzsch's theorem. *Transactions of the American Mathematical Society*, 85(1):219–227, 1957.

[2] P. J. Bushell. Hilbert's metric and positive contraction mappings in a Banach space. *Archive for Rational Mechanics and Analysis*, 52(4):330–338, 1973.

[3] H. Furstenberg and H. Kesten. Products of random matrices. *Annals of Mathematical Statistics*, 31(2):457–469, 1960.

[4] P. Bougerol and J. Lacroix. *Products of Random Matrices with Applications to Schrödinger Operators*. Birkhäuser, 1985.

[5] P. Walters. *An Introduction to Ergodic Theory*. Springer, 1982.

[6] W. Parry and M. Pollicott. *Zeta Functions and the Periodic Orbit Structure of Hyperbolic Dynamics*. Astérisque 187–188, 1990.

[7] M. Viana. *Lectures on Lyapunov Exponents*. Cambridge Studies in Advanced Mathematics, Cambridge University Press, 2014.