

Descriptor-Bounded Ordered Positive Propagation implies Extremal Projective Rigidity modulo Scaling

Abstract

We prove a closed structural rigidity principle for finite-type order-preserving positive dynamics.

A *descriptor-bounded ordered positive propagation* model is presented as a finite-template positive matrix cocycle over a shift of finite type, together with an extremal support set Σ_* and an invariant face $(\mathbb{R}_+^S)^\circ$. If there exists a single *Doebelin block* B on S that recurs on Σ_* with bounded gaps (syndetically), then Hilbert-projective contraction occurs with uniform frequency along every orbit in Σ_* . Consequently, forward products admit a *unique extremal projective direction modulo scaling* and converge to it exponentially.

The quantitative layer is explicit and computable from finite data. Writing $m_B := \min_{i,j \in S} B_{ij}$ and $M_B := \max_{i,j \in S} B_{ij}$, one may take the Hilbert contraction modulus $\tau(B) \leq (r-1)/(r+1)$ with $r := M_B/m_B$, hence an explicit exponential rate $\kappa = -\log \tau(B)/(L_{\text{syn}} + |w|)$ and an explicit projective “spectral gap” parameter $\delta = -\log \tau(B)$. When a uniform positivity window is available on S , these bounds depend only on finite template parameters such as $(|S|, |w|, M_*/m_*)$.

Our contribution is a synthesis in a finite, referee-checkable form: a single recurrent Doebelin block plus a bounded-gap return certificate forces a unique extremal projective direction, and in a canonical transfer-matrix archetype we also obtain a closed extremal classification with an explicit contraction modulus.

We record two short universality embeddings (positive cocycles over primitive shifts of finite type; primitive weighted directed graphs/transfer matrices) and an optional selection add-on: a gap certificate that forces extremals onto a critical subsystem on which the rigidity theorem applies.

1 Introduction, main results, and referee map

This paper isolates one flagship implication for finite-type positive dynamics:

descriptor-bounded ordered positive propagation $\Rightarrow$ *unique extremal projective rigidity modulo scaling*.

Why this matters. The same finite-type positive geometry underlies transfer-matrix methods and Perron–Frobenius growth, locally constant positive cocycles over shifts of finite type, and deterministic avatars of random matrix product theory. Our goal is to state the rigidity mechanism so that *every hypothesis is a finite object* and *every quantitative constant is computable* from finite data.

Main results and contributions at a glance.

- **Flagship rigidity theorem (finite certificate $\Rightarrow$ unique extremal direction).** For a descriptor-bounded order-preserving propagation cocycle together with an extremal support

set Σ_* , a *single* certified Doeblin word w plus a finite bounded-gap return certificate on Σ_* forces a unique σ -equivariant projective section and exponential projective contraction along every orbit in Σ_* (Theorem 1).

- **Closed extremal classification in the canonical finite-template archetype.** In the worked finite-template weighted-graph propagation model, extremal behavior admits a finite, exhaustive taxonomy (dominant primitive component, ties/degeneracy, reducible obstructions), yielding a classification theorem that closes the logical world of possibilities (Theorem 3).
- **Explicit quantitative modulus from finite block data.** If the Doeblin block $B = B(w)$ has entry window (m_B, M_B) and returns syndetically with gap L_{syn} , then the contraction modulus and exponential rate are explicit and computable (Corollary 4); in uniform-window regimes one may bound the same quantities directly in terms of (n, M_*, m_*) .
- **Classical reach.** The contract holds in standard classes (Appendix D): locally constant positive cocycles over primitive SFTs (transfer-matrix/thermodynamic formalism models) and primitive transfer matrices / weighted directed graphs.

Referee map (finite hypotheses $\Rightarrow$ exact discharge point).

Finite hypothesis / datum	Used for	Discharged at
Finite-template cocycle $(G, \mathcal{T}, A)$ and (optional) window (m_*, M_*)	finite descriptor closure, local constancy, template bounds	Def. 1 (p. 5)
Invariant face S and extremal support Σ_*	domain where rigidity is asserted	Def. 2 (p. 5); Sec. 3
Doeblin block B on S (with m_B, M_B)	strict Hilbert contraction input	Def. 3 (p. 6); Lem. 2 (p. 6); Lem. 1 (p. 6)
Syndetic return bound for the Doeblin word w	uniform frequency of contraction blocks	Lem. 4 (p. 7); finite certificate Lem. 3 (p. 7)
Flagship rigidity conclusion	unique projective direction modulo scaling	Thm. 1 (p. 8)
Explicit quantitative modulus	computable τ and κ from finite block data	Cor. 1 (p. 9); Prop. 1 (p. 15)
Minimal coupling caveat (no overclaim)	transparency of (n, M, ε) -dependence	Rem. 1 (p. 5)

Constants summary (how to compute the contraction bound). Fix the certified word w of length $\ell := |w|$ and its block product $B := A(w)$ on S . Compute

$$m_B := \min_{i,j \in S} B_{ij}, \quad M_B := \max_{i,j \in S} B_{ij}, \quad r := \frac{M_B}{m_B}.$$

Then $\Delta(B) \leq 2 \log r$ and one may take

$$\tau(B) \leq \hat{\tau} := \frac{r-1}{r+1}, \quad \kappa := \frac{-\log \hat{\tau}}{L_{\text{syn}} + \ell}, \quad \delta := -\log \hat{\tau}.$$

If a uniform positivity window is available on S (parameter $\varepsilon := m_* > 0$ with maximal entry $M_* < \infty$), then r admits the crude bound of Corollary 4:

$$r \leq |S|^{\ell-1} \left(\frac{M_*}{\varepsilon} \right)^\ell.$$

In the constant primitive transfer-matrix case, one may take $\ell \leq (d-1)^2 + 1$ (Wielandt bound), and ε, M_* may be chosen as the minimal and maximal positive entries of the matrix. **How to apply (recipe).**

- Identify the extremal support set Σ_* (or the face S in the worked archetype) from the finite template data.
- Find a Doeblin word w on Σ_* and compute the block product $B = B(w)$; record (m_B, M_B) .
- Certify bounded-gap return of w on Σ_* (syndeticity) and record L_{syn} .
- Plug $(m_B, M_B, |w|, L_{\text{syn}})$ into Corollary 4 to obtain an explicit contraction modulus and rate.

Notation cheat sheet.

Σ / Σ_*	full symbolic space / extremal support (where rigidity is asserted)
$w / w $	Doeblin word / its length
$B = B(w)$	block product along w (strictly positive on the core)
m_B, M_B	minimal / maximal entry of B on the core indices
L_{syn}	maximal gap between consecutive occurrences of w on Σ_*
τ / κ	contraction modulus / exponential rate ($\kappa = -\log \tau$)

Guide for the reader. The flagship rigidity implication is Theorem 1 in Section 6. Sections 3–5 set up the finite-type cocycle model, the Hilbert-metric contraction mechanism for a Doeblin block, and the syndetic-frequency lemma that upgrades local contraction to global projective rigidity on Σ_* . Section 7 records an optional selection add-on (a gap certificate that forces extremals onto a critical subsystem). Appendix A records a concrete reduction archetype and the reduction pipeline. Appendix B records finite, checkable return certificates (avoidance automata and explicit gap bounds). Appendix C records standard consequences (equivariant projective section and additive reduction). Appendix D records two short universality embeddings into classical classes.

2 Positioning and reduction interface

Where this sits in the literature

The mechanism proved here is the projective-regularity backbone that often lies beneath transfer-matrix methods and Perron–Frobenius growth phenomena, random products of positive matrices and Lyapunov-exponent theory, and symbolic-dynamics formulations of cocycles. We use the Hilbert projective metric (Birkhoff/Bushell) as the contraction geometry [1, 2], and keep the symbolic base as a finite directed graph (a shift of finite type) [5], so hypotheses and certificates remain finite and checkable. For the “random products” and growth-rate viewpoint, see [3, 4].

Relation to standard Perron–Frobenius / cocycle contraction results (one-table view).

Classical viewpoint	Certificate-contract viewpoint (this paper)	Payoff
Uniform positivity or primitivity imposed “globally” (e.g. $A(\omega) > 0$ each step, or a fixed primitive matrix)	A <i>single</i> recurrent Doeblin block B plus a finite bounded-gap return certificate on the relevant core Σ_*	Rigidity and contraction from <i>finite data</i> even when positivity is not uniform step-by-step
Ergodic/measure-theoretic statements common (Lyapunov exponents, Oseledets)	Deterministic, extremal-support statements: contraction and uniqueness hold on Σ_* , the support of maximizers	“Extremal” rigidity tuned to the maximizing set, not an ambient typical set
Constants often implicit or asymptotic	Explicit modulus from $(m_B, M_B, w , L_{\text{syn}})$; windowed bounds from (n, m_*, M_*) in finite-template regimes	Quantitative control that is checkable from a finite certificate
Examples treated class-by-class	Two universality embeddings (primitive SFT cocycles; primitive transfer matrices) show classical reach without changing the mechanism	Immediate connection to transfer-matrix methods and symbolic dynamics

Reduction interface (recorded separately)

A formal finite-data reduction interface statement is recorded in Appendix A. It is not used in the proof of the flagship theorem; it exists only to document how a descriptor-bounded propagation model can be presented as a finite-template cocycle on an SFT together with an extremal support set Σ_* .

3 Finite-type positive cocycles and extremal rigidity

Finite-template cocycles

Definition 1 (Finite-template positive cocycle). *Fix $d \geq 2$. Let $\mathcal{T} \subset \{0, 1\}^{d \times d}$ be a finite family of support templates and fix an upper bound $0 < M_* < \infty$. For $T \in \mathcal{T}$ define*

$$\mathcal{M}(T; M_*) := \{A \in \mathbb{R}_+^{d \times d} : A_{ij} = 0 \text{ if } T_{ij} = 0, \quad 0 < A_{ij} \leq M_* \text{ if } T_{ij} = 1\}.$$

If, in addition, a uniform lower bound $m_ > 0$ is part of the data, define the uniform-window subclass*

$$\mathcal{M}(T; m_*, M_*) := \{A \in \mathcal{M}(T; M_*) : m_* \leq A_{ij} \leq M_* \text{ whenever } T_{ij} = 1\}.$$

Let $G = (V, E)$ be a finite directed graph and let $\Sigma \subset E^{\mathbb{Z}}$ be the associated edge shift (an SFT). A finite-template cocycle is a map $A : \Sigma \rightarrow \mathbb{R}_+^{d \times d}$ that is locally constant on edges: for each $e \in E$, $A(\omega)$ depends only on $\omega_0 = e$, and satisfies $A(e) \in \mathcal{M}(T(e); M_)$ for a template map $T : E \rightarrow \mathcal{T}$. If additionally $A(e) \in \mathcal{M}(T(e); m_*, M_*)$ for all admissible e , we say A has a uniform positivity window (m_*, M_*) on the relevant invariant face.*

Remark 1 (On minimal coupling and quantitative constants). *Descriptor-boundedness is the finiteness of the symbolic base and template family. By itself it does not force a uniform lower bound on positive entries. Accordingly, quantitative moduli are stated in one of two transparent ways: either directly from the certified Doeblin block B (via m_B and M_B), or from an additional uniform-window parameter $m_* > 0$ on the invariant face. Whenever a bound is expressed in terms of (m_*, M_*) , this dependence is explicit, and we make no claim that descriptor boundedness alone controls m_* .*

For $\omega \in \Sigma$ define the forward product

$$A^{(n)}(\omega) := A(\sigma^{n-1}\omega) \cdots A(\sigma\omega)A(\omega), \quad n \geq 1.$$

Faces and extremal rigidity

For $S \subset \{1, \dots, d\}$ define the coordinate face

$$\mathbb{R}_+^S := \{x \in \mathbb{R}_+^d : x_i = 0 \text{ for } i \notin S\}, \quad (\mathbb{R}_+^S)^\circ := \{x \in \mathbb{R}_+^S : x_i > 0 \text{ for } i \in S\}.$$

Definition 2 (Extremal projective rigidity). *Let $\Sigma_* \subset \Sigma$ be shift-invariant (the extremal support set). We say the cocycle has extremal projective rigidity on Σ_* along $(\mathbb{R}_+^S)^\circ$ if there exists a map $\xi : \Sigma_* \rightarrow \mathbb{P}((\mathbb{R}_+^S)^\circ)$ such that for every $\omega \in \Sigma_*$ and every $x \in (\mathbb{R}_+^S)^\circ$,*

$$\lim_{n \rightarrow \infty} [A^{(n)}(\omega)x] = \xi(\sigma^n\omega),$$

and the limit is independent of x (equivalently, all forward images become asymptotically collinear).

4 Hilbert metric and Doeblin contraction

Hilbert projective metric

On $(\mathbb{R}_+^S)^\circ$ define the Hilbert metric

$$d_H(x, y) := \log \left(\frac{\max_{i \in S} x_i / y_i}{\min_{i \in S} x_i / y_i} \right), \quad x, y \in (\mathbb{R}_+^S)^\circ.$$

For a positive linear map B preserving $(\mathbb{R}_+^S)^\circ$, define its projective diameter

$$\Delta(B) := \sup\{d_H(Bx, By) : x, y \in (\mathbb{R}_+^S)^\circ\} \in [0, \infty].$$

Lemma 1 (Birkhoff–Bushell contraction coefficient). *If $B : (\mathbb{R}_+^S)^\circ \rightarrow (\mathbb{R}_+^S)^\circ$ is linear and $\Delta(B) < \infty$, then*

$$d_H(Bx, By) \leq \tau(B) d_H(x, y) \quad \text{for all } x, y \in (\mathbb{R}_+^S)^\circ,$$

where one may take

$$\tau(B) = \tanh(\Delta(B)/4) \in [0, 1).$$

Remark 2. *Even when $\Delta(B) = \infty$, the Hilbert metric is nonexpansive in the sense that the best Lipschitz constant is ≤ 1 . The strict contraction input is therefore exactly the finiteness of $\Delta(B)$ for a recurrent block.*

Doeblin blocks

We use a concrete finitary sufficient condition for $\Delta(B) < \infty$.

Definition 3 (Doeblin block on a face). *A matrix $B \in \mathbb{R}_+^{d \times d}$ is a Doeblin block on S if*

- $B(\mathbb{R}_+^S) \subset \mathbb{R}_+^S$,
- there exist constants $0 < m_B \leq M_B < \infty$ such that for all $i, j \in S$,

$$m_B \leq B_{ij} \leq M_B.$$

Lemma 2 (Doeblin block gives uniform finite diameter). *If B is a Doeblin block on S with bounds (m_B, M_B) , then $\Delta(B) < \infty$ and*

$$\Delta(B) \leq 2 \log\left(\frac{M_B}{m_B}\right), \quad \tau(B) \leq \tanh\left(\frac{1}{2} \log\left(\frac{M_B}{m_B}\right)\right) < 1.$$

Proof. For $x \in (\mathbb{R}_+^S)^\circ$ and $i \in S$ we have

$$m_B \sum_{j \in S} x_j \leq (Bx)_i \leq M_B \sum_{j \in S} x_j.$$

Thus for any $x, y \in (\mathbb{R}_+^S)^\circ$ and any $i \in S$,

$$\frac{(Bx)_i}{(By)_i} \in \left[\frac{m_B}{M_B} \cdot \frac{\sum_{j \in S} x_j}{\sum_{j \in S} y_j}, \frac{M_B}{m_B} \cdot \frac{\sum_{j \in S} x_j}{\sum_{j \in S} y_j} \right].$$

Taking the ratio of the maximal to minimal coordinate ratios yields $d_H(Bx, By) \leq 2 \log(M_B/m_B)$, so $\Delta(B)$ is finite with the stated bound. The contraction bound follows from the previous lemma. $\square$

Remark 3. *In applications, B is produced as the product of a word w on the extremal core. Under finite-template bounds (m_*, M_*) , once the product is primitive on S (all entries on $S \times S$ positive), one can take $(m_B, M_B) = (m_*^{|w|}, M_*^{|w|})$ after restricting to the $S \times S$ block. This is the finite certificate used by the Engineer’s Kit.*

5 Syndetic return and exponential contraction

Fix a shift-invariant extremal set $\Sigma_* \subset \Sigma$. Let w be an admissible word in the underlying directed graph of Σ_* and write $[w] \subset \Sigma$ for the cylinder set.

Definition 4 (Syndeticity of a word on an invariant set). *The word w is syndetic on Σ_* if there exists $L_{\text{syn}} \in \mathbb{N}$ such that every $\omega \in \Sigma_*$ contains an occurrence of w starting in each interval of indices of length L_{syn} . Equivalently, along every $\omega \in \Sigma_*$ the gaps between successive occurrences of w are uniformly bounded by L_{syn} .*

Lemma 3 (Finite syndeticity certificate on an SFT). *Let $\Sigma_* \subset \Sigma$ be an SFT presented by a finite directed graph. Fix an admissible word w of length $|w|$ on that graph. Consider the finite automaton whose states are the length- $(|w| - 1)$ suffixes of admissible words and whose transitions track extension by one symbol along the graph while avoiding the appearance of w . Then:*

- *if the avoidance automaton contains a directed cycle, there exists $\omega \in \Sigma_*$ with no occurrences of w (so w is not syndetic);*
- *if the avoidance automaton is acyclic, there is a finite bound L_{syn} on the length of any w -free path, hence w is syndetic on Σ_* .*

In the acyclic case one may take a uniform lower return density

$$\rho_0 \geq \frac{1}{L_{\text{syn}} + |w|}.$$

Proof. A directed cycle in the avoidance automaton generates an arbitrarily long admissible path that avoids w , hence an infinite bi-infinite path in Σ_* with no w occurrence. If the automaton is acyclic, every w -free path has length at most the length of the longest directed path, call it L_{syn} ; thus every orbit must hit $[w]$ at least once in each window of length $L_{\text{syn}} + |w|$. This implies the stated uniform density bound. $\square$

Lemma 4 (Syndetic Doeblin return forces exponential contraction). *Assume there exists a nonempty face S such that:*

- (Face invariance) *for every $\omega \in \Sigma_*$ and every $n \geq 1$, $A^{(n)}(\omega)(\mathbb{R}_+^S) \subset \mathbb{R}_+^S$, and $A^{(n)}(\omega)$ maps $(\mathbb{R}_+^S)^\circ$ into itself;*
- (Doeblin block) *there exists an admissible word w on Σ_* such that the block product B associated to w is a Doeblin block on S , hence has $\tau(B) \in (0, 1)$;*
- (Syndetic return) *the word w is syndetic on Σ_* with bound L_{syn} .*

Then there exist constants $C \geq 1$ and $\kappa > 0$ such that for every $\omega \in \Sigma_$, every $n \geq 1$, and all $x, y \in (\mathbb{R}_+^S)^\circ$,*

$$d_H(A^{(n)}(\omega)x, A^{(n)}(\omega)y) \leq Ce^{-\kappa n} d_H(x, y),$$

with κ depending only on $\tau(B)$ and $L_{\text{syn}}, |w|$.

Proof. Along any $\omega \in \Sigma_*$, the syndetic bound implies that in the first n steps there are at least

$$r(n) \geq \left\lfloor \frac{n}{L_{\text{syn}} + |w|} \right\rfloor$$

occurrences of w . Decompose the product into alternating segments

$$A^{(n)}(\omega) = U_r B U_{r-1} B \cdots U_1 B U_0,$$

where each U_j is the inter-block product. By face invariance, all factors map $(\mathbb{R}_+^S)^\circ$ into itself, so d_H is defined throughout. Each inter-block product is nonexpansive in Hilbert metric (Lipschitz constant ≤ 1), while each B contracts by $\tau(B)$. Thus

$$d_H(A^{(n)}(\omega)x, A^{(n)}(\omega)y) \leq \tau(B)^{r(n)} d_H(x, y).$$

Since $r(n) \geq n/(L_{\text{syn}} + |w|) - 1$, we get

$$\tau(B)^{r(n)} \leq \tau(B)^{-1} \exp\left(-\frac{|\log \tau(B)|}{L_{\text{syn}} + |w|} n\right),$$

which yields the claim with $C = \tau(B)^{-1}$ and $\kappa = |\log \tau(B)|/(L_{\text{syn}} + |w|)$. $\square$

6 Flagship rigidity theorem

Theorem 1 (Descriptor-bounded ordered positive propagation $\Rightarrow$ extremal rigidity). *Let (Σ, A) be a finite-template positive cocycle. Let $\Sigma_* \subset \Sigma$ be shift-invariant and suppose there exists a nonempty face S and an admissible word w such that:*

- $A^{(n)}(\omega)$ preserves $(\mathbb{R}_+^S)^\circ$ for all $\omega \in \Sigma_*$ and all $n \geq 1$;
- the block product B associated to w is a Doeblin block on S (hence $\tau(B) \in (0, 1)$);
- w is syndetic on Σ_* .

Then the cocycle has extremal projective rigidity on Σ_ along $(\mathbb{R}_+^S)^\circ$. More precisely, there exists a unique map $\xi : \Sigma_* \rightarrow \mathbb{P}((\mathbb{R}_+^S)^\circ)$ such that for every $\omega \in \Sigma_*$ and every $x \in (\mathbb{R}_+^S)^\circ$,*

$$\lim_{n \rightarrow \infty} [A^{(n)}(\omega)x] = \xi(\sigma^n \omega),$$

and the convergence is exponential in the Hilbert metric. Moreover, if L_{syn} is a syndetic gap bound for w on Σ_ and $\ell := |w|$, then for all $\omega \in \Sigma_*$, all $x, y \in (\mathbb{R}_+^S)^\circ$, and all $n \geq 0$,*

$$d_H(A^{(n)}(\omega)x, A^{(n)}(\omega)y) \leq \tau(B)^{\lfloor \frac{n}{L_{\text{syn}} + \ell} \rfloor} d_H(x, y) \leq e^{-\kappa n} d_H(x, y),$$

where one may take $\kappa := \frac{-\log \tau(B)}{L_{\text{syn}} + \ell}$. Quantitative addendum. In particular, for numerical finite-template data, the contraction modulus is explicitly computable: set $m_B := \min_{i,j \in S} B_{ij}$, $M_B := \max_{i,j \in S} B_{ij}$, and $r := M_B/m_B$. Then $\Delta(B) \leq 2 \log r$ and one may take

$$\tau(B) \leq \tanh\left(\frac{1}{2} \log r\right) = \frac{r-1}{r+1} \in [0, 1),$$

hence $\hat{\tau}$ and κ are explicit. If uniform template bounds m_, M_* on S are part of the contract, then r admits the crude descriptor bound of Corollary 4.*

Proof. Fix $\omega \in \Sigma_*$. The previous lemma gives exponential contraction of d_H along ω . Hence for any $x \in (\mathbb{R}_+^S)^\circ$, the sequence $[A^{(n)}(\omega)x]$ is Cauchy in the projective topology induced by d_H , so it converges to a limit projective point. The same contraction implies independence of the initial vector, giving a well-defined projective limit class.

Define $\xi(\omega)$ to be the limit of $[A^{(n)}(\omega)x]$ as $n \rightarrow \infty$ for any $x \in (\mathbb{R}_+^S)^\circ$. Cocycle consistency yields the covariance relation $\xi(\sigma\omega) = [A(\omega)v]$ for any representative v of $\xi(\omega)$, and uniqueness follows because projective space quotients out positive scaling. $\square$

Corollary 1 (Explicit quantitative contraction modulus from finite certificates). *Assume the hypotheses of Theorem 1 and let $B := A(w)$ be the certified Doeblin block on S . Then the numbers*

$$m_B := \min_{i,j \in S} B_{ij}, \quad M_B := \max_{i,j \in S} B_{ij}, \quad r := \frac{M_B}{m_B}$$

are finite and effectively computable from the finite template data. Consequently one may take the Hilbert contraction modulus

$$\hat{\tau} := \frac{r-1}{r+1} \in [0, 1),$$

and the exponential contraction rate

$$\kappa := \frac{-\log \hat{\tau}}{L_{\text{syn}} + |w|}.$$

Equivalently, the projective spectral gap $\delta := -\log \hat{\tau}$ is explicit once B is computed.

Corollary 2 (Periodic-core closure). *If Σ_* is supported on a directed cycle and w is the full cycle word, then w is syndetic. If the corresponding cycle product is Doeblin on S , the theorem yields rigidity on the periodic core.*

Corollary 3 (Finite-step inevitable contraction under strict dominance (transfer-matrix class)). *Work in the primitive weighted-graph / transfer-matrix archetype. Assume the maximal cycle-mean value is attained on a single aperiodic strongly connected component C_* of size d_* . Let T_* be the restricted transfer matrix on C_* . Then there exists an exponent $N \leq (d_* - 1)^2 + 1$ such that $B := T_*^N$ is strictly positive entrywise, hence a Doeblin block on the dominant component. Consequently, Hilbert-metric contraction is forced once every N steps: for all $x, y \in (\mathbb{R}_+^{C_*})^\circ$ and all $n \geq 0$,*

$$d_H(T_*^n x, T_*^n y) \leq \tau(B)^{\lfloor n/N \rfloor} d_H(x, y).$$

In particular, extremal projective rigidity holds on the cycle-mean extremal support set Σ_^{cm} supported on C_* .*

Proof. By Proposition 9, there is $N \leq (d_* - 1)^2 + 1$ such that $B = T_*^N$ is strictly positive, hence a Doeblin block with $\tau(B) \in (0, 1)$. Because the cocycle is constant, the block B occurs deterministically at times $N, 2N, 3N, \dots$ along every orbit, so the syndetic-return hypothesis is automatic with $L_{\text{syn}} = N$. Applying Theorem 1 (with w the length- N time block) yields the stated contraction bound and rigidity. $\square$

Corollary 4 (Quantitative contraction modulus from contract bounds). *Assume the hypotheses of Theorem 1. Assume further that the finite-template data come with uniform entry bounds on the invariant face S : for every admissible symbol e occurring on Σ_* and every $i, j \in S$, either $(A(e))_{ij} = 0$ or $m_* \leq (A(e))_{ij} \leq M_*$. Let $s := |S|$ and $\ell := |w|$. Then with*

$$r := s^{\ell-1} \left(\frac{M_*}{m_*} \right)^\ell, \quad \hat{\tau} := \frac{r-1}{r+1} \in [0, 1),$$

one has $\tau(B) \leq \hat{\tau}$ and may take the exponential rate in Theorem 1 as

$$\kappa \geq \frac{-\log \hat{\tau}}{L_{\text{syn}} + \ell}.$$

In the constant primitive transfer-matrix case, one may take $\ell \leq (d-1)^2 + 1$ and m_*, M_* as the minimal and maximal positive entries of the matrix, yielding an explicit modulus in terms of $(d, M_*/m_*)$.

Proof. Immediate from Proposition 1 (Appendix B) together with Theorem 1. $\square$

7 Optional add-on: spectral separation and selection

The flagship theorem above is purely projective. To connect it to a concrete extremal selection mechanism, introduce an edge cost $c : E \rightarrow \mathbb{R}$ and the minimum cycle mean λ_* . Let Σ_{crit} be the SFT induced by SCCs attaining λ_* .

Lemma 5 (Selection under a gap certificate). *Assume there exists $\delta > 0$ such that any SCC with minimum mean strictly larger than λ_* has mean at least $\lambda_* + \delta$. Then every invariant minimizing measure for c is supported on Σ_{crit} .*

Proof. Shift costs by $c'(e) = c(e) - \lambda_*$. On critical SCCs every cycle has nonnegative c' -weight and some have zero. On every other SCC, every directed cycle has total c' -weight at least δ times its length. Any invariant measure decomposes into an average of cycle means along its visited SCCs; if it assigns positive mass outside Σ_{crit} , the average c' -cost is bounded below by a positive constant proportional to that mass, hence cannot be minimal. $\square$

Remark 4. *Once the gap certificate forces extremals onto Σ_{crit} , the rigidity theorem applies provided one verifies a Doeblin block and syndetic return on that critical support. This is exactly the bridge targeted by the Engineer’s Kit certificates.*

8 Discussion: scope, stability, and what is not assumed

The flagship theorem proves a rigidity mechanism under the weakest hypotheses that make uniform projective contraction possible. It is important to record what is *not* assumed.

No global primitivity requirement

Many classical Perron–Frobenius arguments begin by assuming a single matrix is primitive or that every generator is strictly positive. Here neither is required. The cocycle may have many symbols with sparse templates or even reducible behavior outside the invariant face. All that matters for extremal rigidity is that along the extremal support set there exists one Doeblin block that returns with bounded gaps.

Weak quantitative inevitability (finite-step contraction)

The flagship theorem is qualitative in the sense that it does not attempt to express the optimal contraction coefficient purely from descriptor data. Nevertheless, once a Doeblin word w and a syndetic gap bound L_{syn} are available, there is an immediate “finite-step inevitability” statement.

Remark 5. *If w occurs on Σ_* with gaps at most L_{syn} , then with $L := L_{\text{syn}} + |w|$ every length- L window along every $\omega \in \Sigma_*$ contains an occurrence of w . Consequently, along every $\omega \in \Sigma_*$ the cocycle map contracts by at least the Doeblin factor $\tau(B)$ once every L steps, and the exponential rate in Theorem 1 is obtained by iterating this finite-step contraction.*

Remark 6 (Automatic Doeblin block in the constant primitive transfer-matrix class). *In the primitive weighted-graph / transfer-matrix archetype, once the maximizing support is a single aperiodic strongly connected component, primitivity forces a finite power of the restricted transfer matrix to be strictly positive (hence a Doeblin block) with an exponent bounded only by the component size. Thus in that classical class, the existence of a contraction block is structurally forced by the finite combinatorics of the maximizing component. See Corollary 3 (and Proposition 9 for the finite-step positivity input).*

Stability under contract-preserving perturbations

All hypotheses in the flagship mechanism are formulated with uniform margins: a finite descriptor model, an invariant face, a strictly positive Doeblin block on that face, and a bounded-gap return property on the extremal support. Within any perturbation regime that preserves these margins (for example, a uniform positivity window and the positivity of the chosen Doeblin block on $S \times S$), the same conclusion persists, with constants that vary continuously with the underlying finite data. This provides a natural notion of structural stability for the rigidity mechanism.

No probabilistic assumptions

Although the literature on random matrix products provides a natural context [3, 4], the present statement is deterministic. It applies to every orbit in the shift-invariant extremal set Σ_* . This is essential for a variational principle: extremals are selected deterministically, not sampled.

Spectral gap as optional glue

A spectral separation or selection gap is often the hardest part of a full program. The present paper isolates the projective-regularity consequence once such a gap has been established, but it does not require it. If a gap theorem later forces Σ_* onto a canonical critical subsystem, the flagship theorem applies immediately.

What “breaks” without syndetic return

If the Doeblin word occurs only occasionally, one may still have contraction along some trajectories but not uniformly along all extremals. In that regime one expects a family of projective limits indexed by past histories or by components of Σ_* . Thus syndeticity is precisely the hypothesis that turns local projective contraction into global uniqueness on the extremal set.

A Reduction archetype and pipeline

Theorem 2 (Reduction interface theorem: descriptor bound $\Rightarrow$ finite-template cocycle). *Assume a propagation model satisfies the DBOPP interface data above. Then there exists a finite-template cocycle $A : \Sigma \rightarrow \mathbb{R}_+^{d \times d}$ over the SFT $\Sigma \subset E^{\mathbb{Z}}$ (the edge shift of G) such that:*

- *the propagation dynamics along an admissible descriptor trajectory is represented by the corresponding cocycle product $A^{(n)}(\omega)$,*

- the underlying variational/extremal principle selects a shift-invariant support set $\Sigma_* \subset \Sigma$ on which extremal growth is realized,
- all hypotheses and certificates in the flagship rigidity theorem can be formulated as finite checks on $(G, \mathcal{T}, M_*, A, \Sigma_*)$ (and, in the uniform-window regime, on (m_*, M_*) as well).

In particular, once a coordinate face $(\mathbb{R}_+^S)^\circ$ is identified as invariant on Σ_* , the rigidity theorem below applies on that face whenever a recurrent Doeblin block and a return-density certificate are available on Σ_* .

Purpose. This appendix records one concrete reduction archetype and the finite reduction pipeline that connects a general extremal-selection problem to the hypotheses of the flagship theorem. None of this material is used in the proof of Theorem 1; it is included only to clarify how the finite certificates typically arise.

A concrete reduction archetype (local ordered propagation with bounded descriptors)

To make the interface tangible while keeping the flagship proof independent of any particular model, we record a standard archetype. Assume an order-preserving, positively homogeneous recursion on a cone, where a *descriptor* $v_n \in V$ (finite) selects a local update rule at time n :

$$x_{n+1} = F_{v_n}(x_n), \quad x_n \in \mathbb{R}_+^d,$$

with each F_v monotone and positively homogeneous. In many propagation models (maximizing or minimizing linear-fractional updates, constrained growth, transfer updates), one may represent each F_v on an invariant face as a nonnegative matrix action

$$x_{n+1} = A(v_n) x_n,$$

where only finitely many templates occur; all nonzero entries are bounded above by M_* , and in the uniform-window regime one has $m_* \leq A_{ij} \leq M_*$ whenever the template allows (i, j) on the invariant face. Admissible descriptor transitions are encoded by a finite directed graph $G = (V, E)$, so (v_n) is a path in G and the base dynamics is a vertex shift (conjugate to an edge shift by a standard higher-block presentation). An extremal principle then selects a shift-invariant support set Σ_* of descriptor paths.

Worked archetype instance. Let $V = \{A, B, C\}$ and let admissible transitions be given by the aperiodic strongly connected graph

$$A \rightarrow A, \quad A \rightarrow B, \quad A \rightarrow C, \quad B \rightarrow C, \quad C \rightarrow A.$$

Define a vertex-locally-constant cocycle on $\mathbb{R}_+^2$ by

$$A(A) = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \quad A(B) = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \quad A(C) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

This system is order-preserving on $\mathbb{R}_+^2$ and positively homogeneous. The face is $S = \{1, 2\}$ (full), and the word $w = A$ is a Doeblin block because $A(A)$ has all entries strictly positive. Moreover, in the underlying SFT every admissible path returns to A with gaps at most 3 (the longest A -free excursion is $B \rightarrow C$), so w is syndetic. The flagship theorem therefore yields unique projective rigidity on the extremal support set Σ_* (for this toy archetype one may take $\Sigma_* = \Sigma$). This is the exact pattern used abstractly in the main proof: a single recurrent Doeblin block plus a finite return certificate forces projective regularity.

Remark 7 (Explicit constants in the worked archetype). *In the worked archetype instance, the Doeblin block is the single symbol A , with $B(w) = A(A)$. On the full face $S = \{1, 2\}$ one may take $m_B = 1$ and $M_B = 2$, hence*

$$\Delta(B) \leq 2 \log\left(\frac{M_B}{m_B}\right) = 2 \log 2, \quad \tau(B) \leq \tanh(\log 2/2) < 1.$$

Since the return-gap bound is $L_{\text{syn}} = 2$, one may take $\rho_0 \geq 1/3$. Lemma 4 then produces an explicit $\kappa > 0$ depending only on these finite constants. The numerical values are not the point; the point is that they are determined from finite data.

Remark 8 (Mid-size constants illustration). *This short computation illustrates how the explicit modulus behaves once a certificate is in hand. Suppose a certified Doeblin block B on a core of size $|S| = 12$ has entry window $(m_B, M_B) = (0.08, 0.32)$, so that $r := M_B/m_B = 4$ and $\tau(B) \leq (r-1)/(r+1) = 3/5$. If the word returns with gap $L_{\text{syn}} = 9$ and length $|w| = 7$, then Corollary 4 gives a global contraction rate*

$$d_H(A^{(n)}(\omega)x, A^{(n)}(\omega)y) \leq e^{-\kappa n} d_H(x, y), \quad \kappa \geq \frac{-\log(3/5)}{L_{\text{syn}} + |w|} = \frac{\log(5/3)}{16} \approx 0.0319.$$

The point is not sharpness (the certificate could be better), but that the bound is explicit, checkable, and already nontrivial for moderate-sized cores when the Doeblin block has a reasonable entry ratio.

Remark 9 (Mid-size numeric illustration (constant primitive transfer matrix)). *Take $d = 12$ and consider the constant primitive transfer matrix*

$$T := \mathbf{1}\mathbf{1}^\top + I_d,$$

so every entry of T is either 1 (off-diagonal) or 2 (diagonal). Then T is already strictly positive, so one may take the Doeblin word w of length $\ell = 1$ with block $B = T$ and syndetic gap bound $L_{\text{syn}} = 1$. Here $m_B = 1$ and $M_B = 2$, hence $r = 2$ and

$$\hat{\tau} = \frac{r-1}{r+1} = \frac{1}{3}, \quad \delta = -\log \hat{\tau} = \log 3 \approx 1.0986, \quad \kappa = \frac{\log 3}{2} \approx 0.5493.$$

This example is intentionally elementary: it shows that even at moderate template size ($d \approx 10-20$), the explicit modulus can be nontrivial when the dominant primitive block has uniformly bounded positive entries.

Reduction pipeline (why finite-type is the right interface)

The reduction interface theorem (Theorem 2) is intentionally minimal: it asserts only that the original propagation model can be represented by a finite-template cocycle on a finite SFT, together with an extremal support set Σ_* . In practice, the reduction typically follows a common pipeline, which we record here as a guide to what is being abstracted away.

- **Descriptor closure.** Identify a finite set of local “contexts” or “types” that determine the update rule. These may be literal states, local boundary configurations, or equivalence classes of local neighborhoods. The descriptor bound is the assertion that only finitely many such types occur along extremal trajectories.

- **Symbolic base.** Encode admissible transitions of descriptors by a finite directed graph G . Depending on the model, the cocycle may be vertex-locally-constant (matrices attached to descriptors) or edge-locally-constant (matrices attached to transitions). These are conjugate via higher-block presentations.
- **Template family.** Prove that each local update uses one of finitely many support patterns (templates) and that all nonzero entries are uniformly bounded above on the relevant invariant face. For quantitative moduli, one additionally assumes a uniform lower bound $m_* > 0$ (a “positivity window”) on that face. This is what makes the quantitative certificate layer finite.
- **Extremal support.** Use the variational principle defining the extremals to isolate a shift-invariant subset Σ_* of admissible paths. This may be the full shift, a critical subshift, or a proper sofic/SFT subsystem.

The present paper starts *after* this pipeline: once $(G, \mathcal{T}, M_*, A, \Sigma_*)$ are in hand (and, when quantitative constants are desired, a uniform window parameter $m_* > 0$ is fixed), the projective-regularity mechanism reduces to two finite dynamical certificates (a Doeblin word and a return-density certificate). This is what allows the proof to be short and the assumptions to be checkable.

Remark 10 (Why “modulo scaling” is the natural rigidity level). *Because the underlying propagation is positively homogeneous, the natural state space is projective: only directions in the cone matter. The Hilbert metric is precisely the projective geometry compatible with positivity. Thus the correct notion of uniqueness is a unique projective class (a unique ray), not a unique vector.*

Finite certificate interface (what must be checked)

The flagship theorem requires only two dynamical inputs on Σ_* , both of which admit finite certificates.

- **Doeblin block.** A finite word w in the edge shift whose product matrix $B(w)$ is strictly positive on the invariant face S and admits an explicit lower bound $m_B > 0$ on its $S \times S$ entries.
- **Return density.** A certificate that w recurs on Σ_* with bounded gaps (syndetically), yielding an explicit gap bound L_{syn} and an explicit lower bound $\rho_0 > 0$ on the frequency of contraction blocks along every orbit in Σ_* .

Aperiodic cores are handled by an avoidance-automaton certificate: on a finite SFT, syndeticity of a word is equivalent to the nonexistence of directed cycles in the word-avoidance automaton. The Engineer’s Kit implements these checks, but the checks themselves are finite and independent of any computation layer.

Object	Finite check	Output constant
Invariant face S	closure of supports on Σ_* (coordinates outside S never created)	face S
Doeblin word w	$B(w)$ has all $S \times S$ entries positive and bounded below	$m_B, \Delta(B), \tau(B)$
Syndetic return	avoidance automaton has no directed cycle	L_{syn}, ρ_0

B Return certificates and finite checks

Purpose. The flagship theorem is stated so that each hypothesis can be discharged on finite data. This appendix records the finite certificates and their explicit constants, including the avoidance-automaton certificate for syndetic return.

The flagship implication is deliberately organized so that every hypothesis can be verified on finite data. This section records the finite certificates and the constants they produce, without relying on any computation layer.

Certificate A: invariant face

Given $(G, \mathcal{T}, m_*, M_*, A, \Sigma_*)$, a coordinate face $S \subset \{1, \dots, d\}$ is *invariant on Σ_** if for every admissible edge (or vertex) symbol used on Σ_* , the update never creates mass outside S from mass inside S . In the finite-template setting this reduces to support closure:

$$i \notin S, j \in S \implies A_{ij}(e) = 0 \text{ for every symbol } e \text{ occurring on } \Sigma_*.$$

Equivalently, templates restrict to a closed sub-cone $(\mathbb{R}_+^S)^\circ$.

Certificate B: Doeblin block and contraction factor

A word w (in the chosen SFT presentation of Σ_*) is a *Doeblin word on S* if its block product $B(w)$ has all $S \times S$ entries strictly positive. In finite-template form, one can compute explicit bounds

$$m_B \leq (B(w))_{ij} \leq M_B, \quad i, j \in S,$$

either directly from the realized matrices or from the template bounds. The Hilbert diameter satisfies $\Delta(B) \leq 2 \log(M_B/m_B)$, and the Birkhoff contraction coefficient satisfies

$$\tau(B) \leq \tanh(\Delta(B)/4) \in (0, 1).$$

This yields an explicit “one-hit” contraction factor whenever w occurs along an orbit.

Proposition 1 (Crude explicit contraction modulus from template bounds). *Assume the finite-template data come with uniform entry bounds on the invariant face S : for every admissible symbol e occurring on Σ_* and every $i, j \in S$, either $(A(e))_{ij} = 0$ or $m_* \leq (A(e))_{ij} \leq M_*$. Let $w = e_0 \cdots e_{\ell-1}$ be a Doeblin word on S of length ℓ and let*

$$B := B(w) = A(e_{\ell-1}) \cdots A(e_0).$$

Write $s := |S|$ and set $m_B := \min_{i,j \in S} B_{ij}$ and $M_B := \max_{i,j \in S} B_{ij}$. Then

$$m_B \geq m_*^\ell, \quad M_B \leq s^{\ell-1} M_*^\ell, \quad r := \frac{M_B}{m_B} \leq s^{\ell-1} \left(\frac{M_*}{m_*} \right)^\ell.$$

Consequently the Hilbert contraction coefficient of B admits the explicit bound

$$\tau(B) \leq \tanh\left(\frac{1}{2} \log r\right) = \frac{r-1}{r+1} \in (0, 1).$$

If, in addition, w returns syndetically on Σ_* with gap bound L_{syn} , then the extremal contraction rate in Theorem 1 can be taken explicitly as

$$\kappa \geq \frac{-\log \tau(B)}{L_{\text{syn}} + \ell}.$$

In particular, if one only knows the existence of some Doeblin word of length $\ell \leq L_{\text{Dob}}$ (e.g. Proposition 2), the same bound holds with ℓ replaced by L_{Dob} .

Proof. Fix $i, j \in S$. By matrix multiplication, B_{ij} is a sum over index strings $(k_1, \dots, k_{\ell-1}) \in S^{\ell-1}$ of products

$$(A(e_{\ell-1}))_{ik_{\ell-1}} \cdots (A(e_0))_{k_1 j}.$$

Each product is at most M_*^ℓ and there are at most $s^{\ell-1}$ such strings, so $B_{ij} \leq s^{\ell-1} M_*^\ell$. Since w is Doeblin, $B_{ij} > 0$, hence at least one summand is nonzero; every nonzero factor is $\geq m_*$, so $B_{ij} \geq m_*^\ell$. Thus $r \leq s^{\ell-1} (M_*/m_*)^\ell$. The bound on $\tau(B)$ follows from $\Delta(B) \leq 2 \log r$ and $\tau(B) \leq \tanh(\Delta(B)/4)$. Finally, the exponential rate bound is the one used in the proof of Theorem 1. $\square$

Certificate C: return-density (syndeticity) via avoidance automata

On a finite SFT, the statement “ w returns with bounded gaps on Σ_* ” admits a *finite* certificate: build the word-avoidance automaton $\mathcal{A}_{\text{avoid}}(w)$ and test whether it contains a directed cycle. In the acyclic case, the longest directed path provides an explicit gap bound L_{syn} and hence an explicit uniform return-density lower bound. A formal statement with explicit bounds appears as Proposition 3 and Proposition 4 below.

Optional: bounded automatic search

Because all data are finite, one may search for a short Doeblin–syndetic certificate by enumerating words w up to a length cutoff, testing Doeblin positivity of $B(w)$ on S , and testing syndeticity by the avoidance-automaton cycle test. If a word exists within the cutoff, the procedure finds one and outputs $(m_B, M_B, \tau(B), L_{\text{syn}}, \rho_0)$.

Certificate D: a descriptor-only bound on the length of Doeblin words

The Doeblin hypothesis is existential: there exists an admissible word w whose block product $B(w)$ is strictly positive on $S \times S$. Even without any quantitative entry bounds, the *existence* of such a word admits a purely finite, descriptor-only completeness bound.

Let $s = |S|$ and for each admissible symbol e on Σ_* define its *support pattern* $P(e) \in \{0, 1\}^{s \times s}$ by

$$P(e)_{ij} = 1 \iff (A(e))_{ij} > 0, \quad i, j \in S.$$

For a word $w = e_0 e_1 \cdots e_{n-1}$ define the Boolean product

$$P(w) := P(e_{n-1}) \odot \cdots \odot P(e_0),$$

where $\odot$ is matrix multiplication over the Boolean semiring. Since all matrices are nonnegative, $(B(w))_{ij} > 0$ holds if and only if there exists a Boolean-positive path of indices, so $B(w)$ is strictly positive on $S \times S$ if and only if $P(w)$ is the all-ones pattern J .

Construct the *pattern-augmented graph* with vertex set

$$\mathbf{H} := V_* \times \mathcal{P}, \quad \mathcal{P} := \{0, 1\}^{s \times s},$$

and edges

$$(v, Q) \rightarrow (v', Q') \quad \text{whenever } (v, v') \in E_* \text{ is an admissible step with symbol } e, \text{ and } Q' = P(e) \odot Q.$$

Start states are (v, I) for $v \in V_*$, where I is the Boolean identity pattern on S . A Doeblin word exists on S if and only if some state (v, J) is reachable in $\mathbf{H}$.

Proposition 2 (Descriptor-only Doeblin length bound). *Assume there exists a Doeblin word on S along Σ_* . Then there exists such a word of length at most*

$$L_{\text{Dob}} \leq |V_*| \cdot 2^{s^2} - 1.$$

Equivalently, enumerating admissible words w up to length $|V_| \cdot 2^{s^2} - 1$ is complete for finding a Doeblin word, whenever one exists.*

Proof. Consider a shortest path in the pattern-augmented graph H from some start state (v_0, I) to a target state (v_n, J) . Along this path the sequence of visited states has length $n + 1$ and all states are distinct (otherwise a repeated state would yield a shortcut, contradicting minimality). Since $|H| = |V_*| \cdot |\mathcal{P}| = |V_*| \cdot 2^{s^2}$, we must have $n + 1 \leq |V_*| \cdot 2^{s^2}$, hence $n \leq |V_*| \cdot 2^{s^2} - 1$. The corresponding word is Doeblin and has the claimed length. $\square$

This bound is intentionally crude, but it supplies a fully finite “weak inevitability” statement: if strict positivity on S is attainable at all, then it is attainable within a length bounded by the descriptor data $(|V_*|, |S|)$. It also justifies bounded automatic search as a complete certificate procedure once the cutoff exceeds $|V_*| \cdot 2^{|S|^2} - 1$.

Engineer’s Kit as a computation layer

The Engineer’s Kit should be viewed as a convenience layer that performs the finite checks above and reports explicit constants. The mathematics of the flagship theorem does not require any software: the proof uses only the finite certificates and the constants they produce.

The flagship theorem is stated so that every assumption is *finite and checkable*. In particular, the syndetic return hypothesis is not a qualitative recurrence claim: on an SFT it admits a deterministic certificate built from a finite automaton. This section records that certificate in a form suitable for refereeing.

Avoidance automaton and explicit gap bounds

Fix an SFT Σ_* presented by a finite directed graph $G_* = (V_*, E_*)$. View points $\omega \in \Sigma_*$ as bi-infinite admissible vertex paths $(v_n)_{n \in \mathbb{Z}}$ with $(v_n, v_{n+1}) \in E_*$. Let $w = (w_0, \dots, w_{\ell-1})$ be an admissible vertex-word of length ℓ in G_* . Define the *avoidance automaton* $\mathcal{A}_{\text{avoid}}(w)$ as follows:

- States are admissible length- $(\ell - 1)$ vertex blocks $(u_0, \dots, u_{\ell-2})$ that occur in Σ_* .
- A transition $(u_0, \dots, u_{\ell-2}) \rightarrow (u_1, \dots, u_{\ell-2}, a)$ exists whenever $(u_{\ell-2}, a) \in E_*$ and the resulting length- ℓ block $(u_0, \dots, u_{\ell-2}, a)$ is *not* equal to w .

This automaton recognizes exactly the finite admissible words in G_* that avoid w .

Proposition 3 (Algorithmic syndeticity certificate on an SFT). *Let Σ_* be an SFT presented by G_* and let w be an admissible word of length ℓ . Then w is syndetic on Σ_* if and only if the avoidance automaton $\mathcal{A}_{\text{avoid}}(w)$ is acyclic. In the acyclic case, the maximum length of a w -free path equals the length of the longest directed path in $\mathcal{A}_{\text{avoid}}(w)$; denote this bound by L_{syn} . Consequently, every orbit in Σ_* hits the cylinder $[w]$ at least once in each window of length $L_{\text{syn}} + \ell$, and one may take*

$$\rho_0 \geq \frac{1}{L_{\text{syn}} + \ell}.$$

Moreover, the decision and the bound L_{syn} are computable by a topological sort and dynamic programming in time linear in the size of $\mathcal{A}_{\text{avoid}}(w)$.

Proposition 4 (Crude descriptor-only gap bound). *Assume the avoidance automaton $\mathcal{A}_{\text{avoid}}(w)$ is acyclic and has N states. Then its longest directed path has length at most $N-1$, hence $L_{\text{syn}} \leq N-1$. Moreover, since states are admissible $(\ell-1)$ -blocks on the vertex set V_* , one always has the crude bound*

$$N \leq |V_*|^{\ell-1}.$$

Consequently, whenever w is syndetic one may take the explicit (but typically nonsharp) bounds

$$L_{\text{syn}} \leq |V_*|^{\ell-1} - 1, \quad \rho_0 \geq \frac{1}{|V_*|^{\ell-1} - 1 + \ell}.$$

Corollary 5 (Finite-step contraction window from descriptor data). *Under the hypotheses of the flagship theorem, along every $\omega \in \Sigma_*$ there exists a time n in each window of length $L_{\text{syn}} + \ell$ at which the Doeblin block $B = A(w)$ appears, and at that time the Hilbert distance contracts by at least $\tau(B)$. In particular, using Proposition 4 one obtains the descriptor-only window bound*

$$n \leq (|V_*|^{\ell-1} - 1) + \ell.$$

This is a weak quantitative inevitability statement: it provides a finite step scale on which strict contraction must occur, without attempting to bound $\tau(B)$ purely from descriptor size.

Cycle test and longest-path bound (finite-time computation)

For completeness we spell out the finite computation underlying the certificate.

- **Cycle detection.** A directed cycle in $\mathcal{A}_{\text{avoid}}(w)$ can be detected by a standard depth-first search or by attempting a topological sort. Failure of the topological sort is equivalent to the existence of a cycle.
- **Longest path in a DAG.** If the automaton is acyclic, compute a topological ordering and then dynamic-program the longest path length L_{syn} by a single forward pass.

A simple description is:

Initialize $d(s) = 0$ for each state s of the automaton.

Process states in topological order. For each edge $s \rightarrow t$, set $d(t) = \max(d(t), d(s) + 1)$.

Then $L_{\text{syn}} = \max_s d(s)$.

The point is not efficiency but finiteness: the certificate is a graph-theoretic computation on an explicitly constructed finite automaton, so the return hypothesis is as checkable as primitivity of a finite matrix.

Remark 11. *This proposition is the formal version of Lemma 3 and is exactly what is implemented in `engineers.kit/code/syndetic.py`: construct $\mathcal{A}_{\text{avoid}}(w)$, check for directed cycles, and in the acyclic case compute the longest path length.*


Worked toy example (aperiodic forcing)

To make the certificate concrete, consider the vertex shift on the directed graph with vertices $\{A, B, C\}$ and edges

$$A \rightarrow A, \quad A \rightarrow B, \quad A \rightarrow C, \quad B \rightarrow C, \quad C \rightarrow A.$$

This graph is strongly connected and aperiodic (the self-loop at A forces period 1). Every admissible bi-infinite path visits the vertex A with gaps at most 3: any excursion away from A must start with $A \rightarrow B$ and then follow $B \rightarrow C \rightarrow A$, so the longest A -free block is the length-2 path $B \rightarrow C$. Thus the length-1 vertex-word $w = A$ is syndetic with $L_{\text{syn}} = 2$ and one may take

$$\rho_0 \geq \frac{1}{L_{\text{syn}} + |w|} = \frac{1}{3}.$$



The diagram above is the avoidance automaton for the vertex-word $w = A$: it is the directed subgraph on $\{B, C\}$ obtained by deleting every step that lands in A . It is a DAG, and its longest directed path has length 2 (the word $B \rightarrow C$), yielding $L_{\text{syn}} = 2$.

The avoidance automaton for $w = A$ is acyclic: the only A -free paths are B and $B \rightarrow C$. The “no-cycle” outcome is the finite, referee-checkable reason that A returns with bounded gaps. (If one prefers an edge-shift presentation, pass to the standard higher-block conjugate where vertex cylinders become edge cylinders.)

Automatic search for a Doeblin word (bounded)

In some applications one does not want to hand-select the Doeblin word w . Since all data are finite, one may search for a short certificate.

Proposition 5 (Bounded search for a Doeblin–syndetic certificate). *Fix a face S and bounds (m_*, M_*) for the templates on $S \times S$. Given a finite SFT presentation of Σ_* and a length cutoff L_{max} , enumerate admissible words w in Σ_* with $|w| \leq L_{\text{max}}$. For each candidate, compute the block product $B(w)$ on S and test:*

- Doeblin test: $B(w)$ has strictly positive entries on $S \times S$ (hence is Doeblin with $(m_B, M_B) = (m_*^{|w|}, M_*^{|w|})$);
- Return test: w is syndetic on Σ_* via the avoidance-automaton certificate.

If the procedure returns a word, it outputs explicit constants $\tau(B)$, L_{syn} , and ρ_0 and therefore an explicit contraction rate κ from Lemma 4. If there exists a word w with $|w| \leq L_{\text{max}}$ satisfying both tests, the procedure finds one.

Remark 12. This bounded-search idea is implemented in `engineers_kit/code/wd_search.py` and can be used either as an automatic certificate generator or as a diagnostic tool to suggest human-readable w .

C Projective section and additive reduction

Purpose. The flagship theorem yields uniform projective contraction and therefore a canonical σ -equivariant projective section on the extremal support. This appendix records standard consequences that are often useful when connecting rigidity back to variational principles.

The contraction conclusion can be packaged in a form that is often useful when connecting the flagship theorem back to an extremal principle: a unique projective section on Σ_* that is equivariant under the cocycle.

Fix a norm $\|\cdot\|$ on $\mathbb{R}_+^S$ and let $\mathbb{P}((\mathbb{R}_+^S)^\circ)$ denote the projective space of positive rays.

Corollary 6 (Unique equivariant projective section). *Assume the hypotheses of Theorem 1. Then there exists a unique map*

$$u : \Sigma_* \rightarrow \mathbb{P}((\mathbb{R}_+^S)^\circ)$$

with the equivariance property

$$u(\sigma\omega) = [A(\omega) \hat{u}(\omega)],$$

where $\hat{u}(\omega)$ is any representative vector of the ray $u(\omega)$ and $[\cdot]$ denotes projective class. Moreover, for every $\omega \in \Sigma_$ and every $x \in (\mathbb{R}_+^S)^\circ$,*

$$\lim_{n \rightarrow \infty} \frac{A^{(n)}(\omega)x}{\|A^{(n)}(\omega)x\|} = \hat{u}(\sigma^n \omega)$$

with convergence that is uniform in ω and exponential in n in the Hilbert metric.

Proof sketch. Choose $u(\omega)$ to be the common projective limit of normalized products along ω , whose existence and uniqueness follow from exponential Hilbert contraction. Equivariance is the identity obtained by shifting the product by one symbol. Uniqueness follows because any other equivariant section would have to coincide with the projective limit produced by the contraction. $\square$

Lemma 6 (Cylinder regularity of the section). *Assume the cocycle is locally constant on symbols and the hypotheses of Theorem 1 hold. Then there exist constants $C > 0$ and $\theta \in (0, 1)$ such that if two points $\omega, \omega' \in \Sigma_*$ agree on their first n symbols, then*

$$d_H(\hat{u}(\omega), \hat{u}(\omega')) \leq C \theta^n.$$

In particular, the projective section is Hölder continuous with respect to the standard symbolic metric.

Proof sketch. If ω and ω' agree on the first n symbols, then the first n factors in the products $A^{(n)}(\omega)$ and $A^{(n)}(\omega')$ coincide. Apply the exponential contraction estimate from Lemma 4 to compare the images of two arbitrary initial vectors, and then pass to the projective limit that defines $\hat{u}$. The constants depend only on the finite certificate data. $\square$

Corollary 7 (Additive reduction along the section). *Under the same hypotheses, there exists a real-valued function $\alpha : \Sigma_* \rightarrow \mathbb{R}$ such that for representatives normalized by $\|\hat{u}(\omega)\| = 1$ one has*

$$A(\omega) \hat{u}(\omega) = e^{\alpha(\omega)} \hat{u}(\sigma\omega).$$

Consequently, for each $n \geq 1$,

$$A^{(n)}(\omega) \hat{u}(\omega) = \exp\left(\sum_{k=0}^{n-1} \alpha(\sigma^k \omega)\right) \hat{u}(\sigma^n \omega).$$

Remark 13. In many extremal problems, α is the object that enters the variational principle as an “additive cost” or “pressure” along extremals. The flagship theorem supplies $\hat{u}$ and α canonically on the extremal support set once a Doeblin–return certificate is known.

Corollary 8 (Periodic orbits pick a unique positive eigenvector). *If $\omega \in \Sigma_*$ is periodic with period p , then $A^{(p)}(\omega)$ has a unique positive eigenvector on $\mathbb{R}_+^S$ up to scaling, and its projective class equals $u(\omega)$.*

Proof. Along a periodic orbit, the contraction applies to the p -step product map, forcing a unique attracting projective fixed point, which is exactly a positive eigenvector class. $\square$

D Universality and classical reach (corollaries)

Role in the paper. Theorem 1 is the conceptual center of gravity. The short results in this appendix are included only to record classical classes where the flagship hypotheses hold, and may be skipped without loss of the main proof.

D.1 Positive matrix cocycles over primitive subshifts of finite type

We record a clean universality embedding for *positive, locally constant matrix cocycles* over a *primitive* subshift of finite type.

Classical system definition. Let $\mathcal{A} = \{1, \dots, q\}$ be a finite alphabet and let $P \in \{0, 1\}^{q \times q}$ be a *primitive* adjacency matrix (some power P^N is strictly positive). The associated two-sided SFT is

$$\Sigma_P = \{ \omega = (\omega_k)_{k \in \mathbb{Z}} \in \mathcal{A}^{\mathbb{Z}} : P_{\omega_k, \omega_{k+1}} = 1 \text{ for all } k \}, \quad \sigma(\omega)_k = \omega_{k+1}.$$

A *locally constant positive cocycle* of dimension d assigns to each allowed transition (i, j) with $P_{ij} = 1$ a matrix

$$M_{ij} \in \mathbb{R}_{>0}^{d \times d}, \quad \text{and sets} \quad A(\omega) := M_{\omega_0, \omega_1}.$$

The n -step product is $A^{(n)}(\omega) = A(\sigma^{n-1}\omega) \cdots A(\omega)$.

We focus on the standard “finite descriptor” regime in which the local matrices come from a *finite template family* with a *uniform positivity window*:

$$0 < m \leq (M_{ij})_{ab} \leq M < \infty \quad \text{for all allowed } (i, j) \text{ and all } a, b \in \{1, \dots, d\}. \quad (1)$$

This is the natural hypothesis for transfer-matrix models and for positive cocycles in thermodynamic formalism: the system is finite-type, order-preserving, and uniformly Doeblin at the one-step level.

Verification of contract conditions. We verify the flagship hypotheses in lemma form. Each check is *finite*.

Lemma 7 (Finite descriptor and ordered positivity). *A positive locally constant cocycle over Σ_P is specified by finite data: the directed graph of allowed transitions $E = \{(i, j) : P_{ij} = 1\}$ and the finite matrix family $\{M_{ij}\}_{(i, j) \in E}$. In particular the descriptor is bounded by $|E|$ together with (d, m, M) . Moreover for each $\omega \in \Sigma_P$, the map $x \mapsto A(\omega)x$ sends $\mathbb{R}_{>0}^d$ into itself and is order-preserving: $x \leq y$ implies $A(\omega)x \leq A(\omega)y$ coordinatewise.*

Lemma 8 (H3': uniform Doeblin and uniform projective contraction). *Assume the uniform positivity window (1). Then each one-step matrix $A(\omega)$ has finite Hilbert diameter on $\mathbb{R}_{>0}^d$ and a uniform Birkhoff contraction coefficient $\tau \in (0, 1)$ depending only on m, M :*

$$\Delta(A(\omega)) \leq 2 \log(M/m), \quad \tau(A(\omega)) \leq \tanh\left(\frac{1}{2} \log(M/m)\right) =: \tau < 1.$$

Consequently, for all $\omega \in \Sigma_P$, all $n \geq 1$, and all $x, y \in \mathbb{R}_{>0}^d$,

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

Proof. The diameter bound is a standard estimate for strictly positive matrices with entry ratio bounded by M/m , and the contraction coefficient bound is the Birkhoff–Bushell formula $\tau = \tanh(\Delta/4)$; see the references cited in the Flagship theorem section. The n -step bound follows by iterating the one-step contraction. $\square$

Remark 14 (How this fits the flagship return hypothesis). *Lemma 8 is strictly stronger than the “Doeblin word + syndetic return” hypothesis used in the flagship theorem. Here contraction happens at every step, so the return-density requirement is automatic (density 1). This is exactly why positive cocycles over finite-type bases are the safest universality target.*

Universality statement.

Corollary 9 (Universality: positive cocycles over primitive SFTs). *Let (Σ_P, σ) be a primitive SFT and let $A : \Sigma_P \rightarrow \mathbb{R}_{>0}^{d \times d}$ be a locally constant cocycle satisfying the uniform positivity window (1). Then the system satisfies the descriptor-bounded ordered positive propagation contract and therefore admits the rigidity conclusions of the flagship theorem: there exists a unique σ -equivariant projective section $\xi : \Sigma_P \rightarrow \mathbb{P}(\mathbb{R}_{>0}^d)$ and exponential projective contraction along every orbit. In particular, projective dynamics is unique modulo scaling and the contraction rate is explicit:*

$$d_H(A^{(n)}(\omega)x, A^{(n)}(\omega)y) \leq e^{-\kappa n} d_H(x, y), \quad \kappa := -\log \tau > 0,$$

where $\tau = \tanh(\frac{1}{2} \log(M/m))$.

Proof. Lemmas 7 and 8 verify the contract hypotheses (in fact in a stronger “uniform Doeblin” form). The rigidity conclusions follow by the same Hilbert-metric contraction mechanism used in the flagship proof, with return density equal to 1. $\square$

Remark 15 (Why record universality). *Corollary 9 shows that the hypotheses of the Flagship theorem arise in a standard symbolic-dynamics class. No additional mechanism is introduced here; the role of the section is only to exhibit classical reach from finite data.*

D.2 Primitive weighted directed graphs and transfer matrices

A second classical target is the *transfer-matrix* setting: propagation by a fixed weighted directed graph, equivalently by a fixed primitive nonnegative matrix (ubiquitous in combinatorics and statistical mechanics) [6, 7].

Classical system definition. Let $G = (V, E)$ be a finite directed graph with $V = \{1, \dots, d\}$. Assign to each directed edge $i \rightarrow j$ a strictly positive weight $w_{ij} > 0$, and set $w_{ij} = 0$ if $i \rightarrow j \notin E$. The associated *transfer matrix* is the nonnegative matrix

$$T = (w_{ij})_{1 \leq i, j \leq d} \in \mathbb{R}_{\geq 0}^{d \times d}.$$

Propagation is the linear update $x_{n+1} = Tx_n$ on $\mathbb{R}_{\geq 0}^d$, so that $x_n = T^n x_0$. We assume T is *primitive*: there exists $N_0 \geq 1$ with T^{N_0} strictly positive.

Verification of contract conditions.

Lemma 9 (Descriptor boundedness and ordered positivity). *The transfer-matrix system is finite-type: it is specified by finite data $(V, E, \{w_{ij}\})$, and the descriptor is bounded by $(d, |E|)$ together with the finite set of weights. Moreover T is order-preserving on the positive cone: $x \leq y$ implies $Tx \leq Ty$ coordinatewise.*

Lemma 10 (Doebelin block from primitivity). *If T is primitive, then there exists $N_0 \geq 1$ such that $B := T^{N_0}$ is strictly positive. Let*

$$m_B := \min_{i,j} (B)_{ij} > 0, \quad M_B := \max_{i,j} (B)_{ij} < \infty.$$

Then B has finite Hilbert diameter $\Delta(B) \leq 2 \log(M_B/m_B)$ and Birkhoff contraction coefficient $\tau(B) \leq \tanh(\Delta(B)/4) \in (0, 1)$.

Proposition 6 (Weak quantitative inevitability in the primitive transfer-matrix class). *Let $T \in \mathbb{R}_{\geq 0}^{d \times d}$ be primitive. Then there exists an exponent $N_0 \geq 1$ with T^{N_0} strictly positive. Moreover, one may choose N_0 bounded purely in terms of the dimension d ; for instance a classical bound gives*

$$N_0 \leq (d-1)^2 + 1.$$

Consequently, strict projective contraction occurs at a finite step bounded by a function of the descriptor (here d): there exists $L(d)$ such that $T^{L(d)}$ is a Doeblin block and hence contracts the Hilbert metric on $\mathbb{R}_{>0}^d$.

Proof. Primitivity implies finite exponent: $T^n > 0$ for some n . A universal bound $n \leq (d-1)^2 + 1$ is standard in nonnegative matrix theory (often attributed to Wielandt); see, e.g., [6]. Once $B := T^{N_0}$ is strictly positive, Lemma 10 yields strict Hilbert-metric contraction for B . $\square$

Proposition 7 (Crude explicit contraction bound from descriptor data). *Let $T \in \mathbb{R}_{\geq 0}^{d \times d}$ be primitive and set*

$$m := \min\{T_{ij} : T_{ij} > 0\} > 0, \quad M := \max_{i,j} T_{ij} < \infty.$$

Let $N_0 \geq 1$ be such that $B := T^{N_0}$ is strictly positive. Then

$$m_B := \min_{i,j} B_{ij} \geq m^{N_0}, \quad M_B := \max_{i,j} B_{ij} \leq d^{N_0-1} M^{N_0}.$$

Consequently the entry ratio $r := M_B/m_B$ satisfies

$$r \leq d^{N_0-1} \left(\frac{M}{m} \right)^{N_0},$$

and the Hilbert contraction coefficient obeys the explicit bound

$$\tau(B) \leq \frac{r-1}{r+1} \in (0, 1).$$

In particular, using Proposition 6 one obtains a fully explicit (typically nonsharp) bound depending only on $(d, M/m)$.

Proof. Each entry B_{ij} is a sum over length- N_0 paths $i \rightarrow \cdots \rightarrow j$ of products of the corresponding edge weights. Each product is at most M^{N_0} and the number of length- N_0 index sequences is at most d^{N_0-1} , hence $B_{ij} \leq d^{N_0-1} M^{N_0}$. Since B is strictly positive, for each pair (i, j) there exists at least one length- N_0 path contributing to B_{ij} , and every factor in such a product is at least m , so $B_{ij} \geq m^{N_0}$. Thus $r \leq d^{N_0-1} (M/m)^{N_0}$. The displayed bound on $\tau(B)$ follows from the Hilbert diameter estimate $\Delta(B) \leq 2 \log r$ and $\tau(B) \leq \tanh(\Delta(B)/4) = \tanh((1/2) \log r) = (r-1)/(r+1)$. $\square$

Remark 16 (Explicit contraction in the constant primitive case). *In the transfer-matrix setting one may take $B = T^{N_0}$. If B is strictly positive with entry ratio $r := M_B/m_B$, then*

$$\Delta(B) \leq 2 \log r, \quad \tau(B) \leq \tanh(\Delta(B)/4) \leq \tanh((1/2) \log r) = \frac{r-1}{r+1}.$$

Thus the quantitative contraction coefficient in this constant primitive case admits an explicit bound depending only on an entry-ratio parameter.

Lemma 11 (Syndetic return is automatic). *For the constant cocycle $A(\omega) \equiv T$ over the one-symbol shift, the Doeblin block $B = T^{N_0}$ occurs with bounded gaps along every orbit: every window of length N_0 contains one occurrence. Equivalently, one may take $L_{\text{syn}} = N_0 - 1$ and a return-density bound $\rho_0 \geq 1/N_0$.*

Universality statement.

Corollary 10 (Universality: primitive weighted directed graphs / transfer matrices). *Let $T \in \mathbb{R}_{\geq 0}^{d \times d}$ be the transfer matrix of a finite weighted directed graph, and assume T is primitive. Then the system satisfies the descriptor-bounded ordered positive propagation contract. Consequently it inherits the flagship rigidity mechanism: there exists a unique projective direction in $\mathbb{P}(\mathbb{R}_{>0}^d)$ attracting every orbit, and the projective contraction rate is explicit. In particular, for $B = T^{N_0}$ as in Lemma 10 and all $x, y \in \mathbb{R}_{>0}^d$,*

$$d_H(T^n x, T^n y) \leq \tau(B)^{\lfloor n/N_0 \rfloor} d_H(x, y) \leq e^{-\kappa n} d_H(x, y), \quad \kappa := \frac{-\log \tau(B)}{N_0} > 0.$$

Proof. Lemma 9 gives the finite descriptor and ordered positivity. Lemma 10 provides a Doeblin block with a strict contraction factor, and Lemma 11 supplies syndetic return. The flagship proof then yields global contraction and uniqueness modulo scaling. $\square$

D.3 Classical reach in one paragraph

The two universality corollaries together place the flagship mechanism inside the standard landscape: transfer matrices and Perron–Frobenius growth systems (constant primitive case), symbolic dynamics and locally constant cocycles (finite-type time-inhomogeneous case), and the broader setting of products of positive matrices and Lyapunov exponents studied in random matrix product theory (but used here deterministically). In particular, the paper isolates the *structural* ingredient behind many “unique eigenvector/equilibrium” conclusions: a finite Doeblin block together with a finite return-density certificate, which upgrades local primitivity into global projective rigidity.

E Structural classification in the weighted-graph archetype

This appendix records the finite classification skeleton that underlies many descriptor-bounded ordered propagation reductions. The purpose is not to introduce new machinery, but to make explicit that the “extremal core” is governed by finite graph theory and Perron–Frobenius structure on finitely many components.

Weighted directed graphs and maximal cycle mean

Let $G = (V, E)$ be a finite directed graph and let $a : E \rightarrow \mathbb{R}$ be an edge weight. For a directed cycle $\gamma = e_1 \cdots e_k$ define its *mean weight*

$$\mu(\gamma) := \frac{1}{k} \sum_{j=1}^k a(e_j), \quad \mu_* := \max_{\gamma \text{ cycle}} \mu(\gamma).$$

Let $G_* \subseteq G$ be the (possibly disconnected) subgraph obtained by taking the union of all strongly connected components that contain at least one cycle attaining the maximal mean μ_* .

Theorem 3 (Finite classification of extremal supports in the cycle-mean archetype). *Consider the shift space Σ_G of bi-infinite directed paths in G . The set of paths whose asymptotic average weight equals the maximal cycle mean,*

$$\Sigma_*^{\text{cm}} := \left\{ \omega \in \Sigma_G : \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=0}^{n-1} a(\omega_t \rightarrow \omega_{t+1}) = \mu_* \right\},$$

is a shift-invariant subshift supported on G_ . Moreover, G_* has finitely many strongly connected components, and exactly one of the following holds:*

- Σ_*^{cm} *is supported on a single primitive component (an aperiodic core);*
- Σ_*^{cm} *is supported on a single periodic component (a maximizing periodic orbit);*
- Σ_*^{cm} *is a finite union of transitive components whose maximal means tie at the same value μ_* .*

In particular, there are no hidden infinite-dimensional corner cases: the extremal possibilities are exhausted by a finite component comparison.

Proposition 8 (Stability of the classification under small perturbations). *Let $a, a' : E \rightarrow \mathbb{R}$ be two edge-weight functions on the same finite directed graph $G = (V, E)$ and assume*

$$\|a - a'\|_\infty := \max_{e \in E} |a(e) - a'(e)| \leq \varepsilon.$$

Then every cycle mean shifts by at most ε , hence $|\mu_(a) - \mu_*(a')| \leq \varepsilon$. If, for a , the maximizing cycle-mean value $\mu_*(a)$ is separated from the runner-up by a gap $\delta > 0$ (no ties at the top), then for all $\varepsilon < \delta/2$ the maximizing subgraph G_* is unchanged and Σ_*^{cm} is supported on the same maximizing component(s). In particular, in the nondegenerate (gap) regime, the extremal-support classification is structurally stable under contract-preserving perturbations.*

Proof. For a cycle $\gamma = e_1 \cdots e_k$,

$$|\mu_{a'}(\gamma) - \mu_a(\gamma)| = \left| \frac{1}{k} \sum_{j=1}^k (a'(e_j) - a(e_j)) \right| \leq \varepsilon.$$

Taking maxima gives $|\mu_*(a) - \mu_*(a')| \leq \varepsilon$. If the top mean is separated by δ , then no cycle whose mean is $\leq \mu_*(a) - \delta$ can overtake the maximizer under a perturbation of size $\varepsilon < \delta/2$, so the set of maximizing cycles (and hence the union of maximizing components) is unchanged. $\square$

Remark 17. *In the ordered-propagation reductions that motivate this paper, the extremal support set Σ_* is produced by a finite maximization criterion of the above type (often a cycle-mean or subadditive-pressure maximization on a finite template graph), and then refined by the Doeblin-return certificate on the selected component. Theorem 3 explains why the “world of extremals” is closed once the descriptor reduction is in place: the only remaining question is whether a unique maximizing component is selected, and whether it carries a recurrent Doeblin block.*

Proposition 9 (Strict dominance $\Rightarrow$ finite-step positivity and inevitable contraction). *Work in the weighted-graph / transfer-matrix archetype. Assume the maximal cycle-mean value μ_* is attained on a single strongly connected component $C_* \subseteq V$ and that C_* is primitive (aperiodic). Let $d_* := |C_*|$ and let T_* be the $d_* \times d_*$ transfer matrix obtained by restricting T to C_* (equivalently, the adjacency of C_* with the induced positive weights). Then there exists an exponent $N_*(d_*) \geq 1$ depending only on d_* such that*

$$(T_*^{N_*})_{ij} > 0 \quad \text{for all } i, j \in C_*.$$

In particular, taking $B := T_^{N_*}$ yields a Doeblin block on the dominant primitive component, hence a strict Hilbert-metric contraction for B . Consequently, on the extremal support Σ_*^{cm} (which is supported on C_* in this strict-dominance regime), projective contraction is inevitable at a finite step bounded by a function of the descriptor: there is a finite window length $L(d_*) := N_*(d_*)$ such that every length- $L(d_*)$ block produces strict projective contraction.*

Proof. Since C_* is aperiodic and strongly connected, its adjacency matrix is primitive. A classical primitivity-exponent bound (often attributed to Wielandt) gives a number

$$N_*(d_*) \leq (d_* - 1)^2 + 1$$

such that the 0–1 adjacency matrix of C_* has all entries positive at power $N_*(d_*)$. Because all weights on edges are strictly positive, every admissible length- N_* path contributes a strictly positive term to the corresponding entry of $T_*^{N_*}$, hence $T_*^{N_*}$ is strictly positive entrywise. Strict positivity implies $B := T_*^{N_*}$ maps the positive cone into its interior, hence has finite Hilbert diameter and a strict Birkhoff contraction coefficient $\tau(B) \in (0, 1)$. Finally, in the constant transfer-matrix setting, the block B occurs deterministically at times $N_*, 2N_*, 3N_*, \dots$, so contraction occurs within every length- N_* window along every orbit supported on C_* . $\square$

Proposition 10 (Generic absence of top ties). *Fix a finite directed graph $G = (V, E)$ and regard edge-weights as a parameter $a \in \mathbb{R}^E$. Then the set of weights a for which the maximal cycle-mean value $\mu_*(a)$ is attained on two distinct strongly connected components (i.e. a genuine “tie at the top” in Theorem 3) is contained in a finite union of proper affine hyperplanes in $\mathbb{R}^E$. In particular it has Lebesgue measure zero and is unstable under generic small perturbations.*

Proof. Every directed cycle has the same mean as its underlying simple cycle, so it suffices to take the maximum over the finite set of simple cycles. For each simple cycle γ the mean weight $\mu(\gamma)$ is a linear functional of a . Hence $\mu_*(a) = \max_\gamma \mu(\gamma)$ is the maximum of finitely many linear functionals. If two different components both contain maximizing cycles, then there exist two distinct simple cycles $\gamma_1 \neq \gamma_2$ with $\mu(\gamma_1) = \mu(\gamma_2) = \mu_*(a)$. This forces a to lie in the hyperplane $\{\mu(\gamma_1) = \mu(\gamma_2)\}$. Taking the finite union over all pairs (γ_1, γ_2) yields the claim. $\square$

Remark 18 (When ties persist). *The previous proposition formalizes the informal principle that top ties are nongeneric. Ties can persist only when the parameter space is constrained, for example by an imposed symmetry: if the weight assignment is forced to be invariant under a nontrivial automorphism of G , then symmetric cycles are forced to have equal means and a top tie may be structurally enforced. This is the natural setting in which “equal maximal means” should be expected to occur.*

F References

References

- [1] G. Birkhoff. Extensions of Jentzsch’s theorem. *Transactions of the American Mathematical Society* **85** (1957), 219–227.
- [2] P. J. Bushell. Hilbert’s metric and positive contraction mappings in a Banach space. *Archive for Rational Mechanics and Analysis* **52** (1973), 330–338.
- [3] H. Furstenberg and H. Kesten. Products of random matrices. *Annals of Mathematical Statistics* **31** (1960), 457–469.
- [4] P. Bougerol and J. Lacroix. *Products of Random Matrices with Applications to Schrödinger Operators*. Progress in Probability and Statistics, Birkhäuser, 1985.
- [5] D. Lind and B. Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, 1995.
- [6] E. Seneta. *Non-negative Matrices and Markov Chains*. Springer Series in Statistics, Springer, 2006.
- [7] D. Ruelle. *Thermodynamic Formalism: The Mathematical Structures of Equilibrium Statistical Mechanics*. Addison–Wesley, 1978.