Proper-Delay Spectra Majorize Subspace Rotation: Exact Finsler Resource Laws for Passive Networks

Lluis Eriksson

10 August 2026

Abstract

A subspace quantum speed limit usually records only the largest canonical angle. Passive scattering theory usually records either the largest proper delay or their total trace. Both compressions discard how a rank-k routing event is distributed across modes. We derive the missing vector law. If $S(t) \in U(N)$ is absolutely continuous, $Q(t) = -\mathrm{i}S(t)^*\dot{S}(t)$ is its Hermitian generator, and $\beta_1 \geq \cdots \geq \beta_k$ are the canonical angles between a fixed rank-k subspace and its transported image, then, for every $r \leq \min\{k, N-k\}$,

$$2\sum_{j=1}^{r} \beta_j \leq \int \sum_{j=1}^{r} \big(\lambda_j(Q) - \lambda_{N-j+1}(Q)\big)\, \mathrm{d}t.$$

Thus the integrated spectral-spread vector weakly majorizes twice the canonical-angle vector. More strongly, for every symmetric gauge Φ and every pair of k-subspaces in $N \geq 2k$ dimensions, we solve the variational problem exactly:

$$\inf_S \int \Phi(\mathrm{Spr}^+(Q(t)))\, \mathrm{d}t = 2\Phi(\beta).$$

The same value holds for the leading-proper-delay action when $Q \succeq 0$, and one constant coupled-mode path minimizes all gauges simultaneously. The Ky Fan members recover the bandwidth-type limit at $r = 1$, strengthen the trace-action law at $r = k$, and expose every intermediate modal budget. An exact nonnegative slack decomposition separates three causes of inefficiency: common-mode delay, failure to use the available spectral spread, and a nongeodesic subspace path. Robust corollaries turn heterogeneous pass/stop leakage spectra and operator-norm tomography errors into certified Lorenz curves for delay and, for rational inner networks, McMillan-degree lower bounds. The matrix spectral-spread inequality and Grassmann metrics used in the proof are established results; the contribution is their dynamical composition with Wigner–Smith positivity, robust routing data and realization degree. A deterministic artifact tests 8,192 random generators, 1,536 random time-dependent paths, sharp equality families and 2,048 noisy tomography instances. Source, frozen data and independent replay are available in the immutable companion release.

Keywords: quantum speed limit; Wigner–Smith time delay; canonical angles; weak majorization; Finsler metric; symmetric gauge; passive quantum network; McMillan degree; robust spectral routing.

1 One routing event, an entire Lorenz curve

Suppose that a device moves a k-dimensional family of inputs through output space. A scalar description asks only how far the worst direction moved, or sums all motion and all delay. Neither distinguishes a coherent rotation of all k directions from one violently delayed direction accompanied by $k-1$ nearly stationary ones. The distinction matters in multiport filters, frequency-selective reservoirs and passive quantum networks.

The correct statement is vector-valued. Write the ordered endpoint angles as $\beta = (\beta_1, \ldots, \beta_k)$. Write the ordered eigenvalues of the instantaneous generator as $\lambda_1(Q) \geq \cdots \geq \lambda_N(Q)$, and set

$$\mathrm{Spr}_j^+(Q) = \lambda_j(Q) - \lambda_{N-j+1}(Q), \qquad 1 \leq j \leq \lfloor N/2 \rfloor. \tag{1}$$

Our central law is

$$2\beta \ \prec_w \ \int \mathrm{Spr}^+(Q(t))\, \mathrm{d}t, \tag{2}$$

with both vectors truncated to $m = \min\{k, N-k\}$. Here $x \prec_w y$ means $\sum_{j=1}^r x_j \leq \sum_{j=1}^r y_j$ for every prefix r. The hierarchy is invariant under $Q(t) \mapsto Q(t) + c(t)I$: common phase does not rotate a subspace.

For a passive monotone scattering path, $Q \succeq 0$ and its eigenvalues are the proper delays. Since every bottom partial sum is nonnegative,

$$2\sum_{j=1}^r \beta_j \leq \int \sum_{j=1}^r \mathrm{Spr}_j^+(Q)\, \mathrm{d}t \leq \int \sum_{j=1}^r \lambda_j(Q)\, \mathrm{d}t. \tag{3}$$

The first inequality says that no collection of fewer than r large proper delays can hide a collective rotation of the first r canonical directions. The full list of prefix inequalities is a Lorenz-curve certificate.

The result refines the preceding robust routing law [6]. That work proved only the endpoints $r = 1$ and a trace bound at $r = k$. Here every intermediate r is controlled, the generator need not be positive for the spectral-spread theorem, heterogeneous directional leakage is retained rather than collapsed to one tolerance, and an exact slack identity diagnoses near saturation.

Our claims are deliberately separated from the ingredients. Unitarily invariant Grassmann metrics are due to Qiu, Zhang and Li [14]. The sharp block inequality $2s(B) \prec_w \mathrm{Spr}^+(Q)$ and its constant-generator subspace consequence were proved by Massey, Stojanoff and Zárate [11]. Maximal-angle subspace speed limits are known from Albeverio and Motovilov [2, 1]; geometric operator speed limits and bandwidth-constrained transformations provide neighboring frameworks [3, 5, 10]; bi-invariant action functionals have also been used to derive exact gate-speed limits [15]. The new result here is the complete time-dependent subspace majorization law, the exact symmetric-gauge variational value with a universal common minimizer, and its operational consequences for positive Wigner–Smith spectra, robust tomography, exact memory accounting and rational realization degree.

2 Geometry and generators

Let $S : [a, b] \to U(N)$ be absolutely continuous and define

$$Q(t) = -\mathrm{i}S(t)^* \dot{S}(t). \tag{4}$$

Then $Q(t)$ is Hermitian almost everywhere and $\dot{S}(t) = \mathrm{i}S(t)Q(t)$. Fix a rank-$k$ projection P_0 and let

$$\mathcal{X}(t) = \mathrm{ran}(S(t)P_0), \qquad \Pi(t) = S(t)P_0S(t)^*. \tag{5}$$

Let $\beta_1 \geq \cdots \geq \beta_k \in [0, \pi/2]$ be the canonical angles between $\mathcal{X}(a)$ and $\mathcal{X}(b)$. Only $m = \min\{k, N-k\}$ can be nonzero.

Relative to $\mathrm{ran}\, P_0 \oplus \mathrm{ran}(I - P_0)$, write

$$Q(t) = \begin{bmatrix} A(t) & B(t)^* \\ B(t) & C(t) \end{bmatrix}, \qquad B(t) = (I - P_0)Q(t)P_0. \tag{6}$$

The diagonal blocks change phases inside the two instantaneous sectors. The horizontal block B alone moves the point of the Grassmannian. Let $s_1(B) \geq \cdots \geq s_m(B)$ be its singular values and define the Ky Fan velocity

$$\nu_r(B) = \sum_{j=1}^r s_j(B), \qquad 1 \leq r \leq m. \tag{7}$$

Lemma 2.1 (Grassmann path length)**.** *For every* $1 \le r \le m$,

$$\sum_{j=1}^{r} \beta_j \le \int_a^b \nu_r(B(t))\,\mathrm{d}t. \tag{8}$$

Proof. For the symmetric gauge $\Phi_r(x) = \sum_{j=1}^{r} x_j^{\downarrow}$, $d_r(\mathcal{U},\mathcal{V}) = \Phi_r(\Theta(\mathcal{U},\mathcal{V}))$ is a unitarily invariant Grassmann metric [14]. The metric derivative of $\mathcal{X}(t)$ is $\Phi_r(s(B(t)))$: conjugating $\dot{\Pi} = \mathrm{i}S[Q,P_0]S^*$ back by S leaves the horizontal tangent block B. The distance between endpoints does not exceed the length of an absolutely continuous path, which is exactly (8). □

3 The Ky Fan speed-limit hierarchy

The algebraic bridge is a sharp result from [11, Theorem 2.7]. We restate it in the notation needed for the dynamical argument and include a short Ky Fan proof.

Lemma 3.1 (Spectral-spread block inequality)**.** *Let* $Q = Q^*$ *have the block form* (6)*. For every* $1 \le r \le m$,

$$2\sum_{j=1}^{r} s_j(B) \le \sum_{j=1}^{r} \big(\lambda_j(Q) - \lambda_{N-j+1}(Q)\big). \tag{9}$$

The constant 2 *is optimal.*

Proof. Put $J = 2P_0 - I$. Then

$$Q - JQJ = 2\begin{bmatrix} 0 & B^* \\ B & 0 \end{bmatrix},$$

whose r largest eigenvalues are $2s_1(B),\dots,2s_r(B)$. Ky Fan's eigenvalue inequality [8] applied to $Q + (-JQJ)$ gives

$$\sum_{j=1}^{r} \lambda_j(Q - JQJ) \le \sum_{j=1}^{r} \lambda_j(Q) + \sum_{j=1}^{r} \lambda_j(-JQJ).$$

Unitary conjugation preserves the spectrum and $\lambda_j(-Q) = -\lambda_{N-j+1}(Q)$, proving (9). Sharpness follows already for $Q = \left[\begin{smallmatrix} 0 & B^* \\ B & 0 \end{smallmatrix}\right]$. □

Theorem 3.2 (Time-dependent spectral-spread QSL)**.** *Let* S*,* Q*,* P_0 *and* β *be as above. Then for every* $1 \le r \le m$,

$$\boxed{2\sum_{j=1}^{r} \beta_j \le \int_a^b \sum_{j=1}^{r} \big(\lambda_j(Q(t)) - \lambda_{N-j+1}(Q(t))\big)\,\mathrm{d}t.} \tag{10}$$

Equivalently, $2\beta \prec_w \int_a^b \mathrm{Spr}^+(Q(t))\mathrm{d}t$*, after truncation to* m *entries. The theorem is invariant under an arbitrary time-dependent scalar shift of* Q.

Proof. Apply Theorem 3.1 almost everywhere, integrate, and combine with Theorem 2.1. Both sides of (10) are unchanged by $Q(t) \mapsto Q(t) + c(t)I$. □

Corollary 3.3 (Positive proper-delay hierarchy)**.** *If* $Q(t) \succeq 0$ *almost everywhere, then for every* $r \le m$,

$$2\sum_{j=1}^{r} \beta_j \le \int_a^b \sum_{j=1}^{r} \mathrm{Spr}_j^+(Q(t))\,\mathrm{d}t \le \int_a^b \sum_{j=1}^{r} \lambda_j(Q(t))\,\mathrm{d}t. \tag{11}$$

In particular, for $k \le N/2$,

$$2\beta_1 \le \int \lambda_1(Q)\,\mathrm{d}t, \qquad 2\sum_{j=1}^{k} \beta_j \le \int \sum_{j=1}^{k} \lambda_j(Q)\,\mathrm{d}t \le \int \mathrm{tr}\,Q\,\mathrm{d}t. \tag{12}$$

The first endpoint in (12) is the familiar maximal-angle, bandwidth-type subspace speed limit. The second is stronger than the trace law because only the k leading proper delays can pay for a rank-k rotation. The intermediate prefixes prevent delay concentration from being mistaken for collective transport.

Corollary 3.4 (Peak constraints)**.** *If $I = [a, b]$ has width $\Delta > 0$, then*

$$\operatorname*{ess\,sup}_{t \in I} \sum_{j=1}^{r} \mathrm{Spr}_j^+(Q(t)) \geq \frac{2}{\Delta} \sum_{j=1}^{r} \beta_j. \tag{13}$$

For $Q \succeq 0$, the same lower bound holds with $\sum_{j=1}^{r} \lambda_j(Q(t))$ on the left.

4 An exact variational principle for every symmetric gauge

The prefix hierarchy is equivalent, by Ky Fan dominance, to a simultaneous family of norm inequalities. It therefore solves a larger optimization problem. Let Φ be any symmetric gauge on $\mathbb{R}^k$: a norm invariant under coordinate permutations and sign changes. It induces a unitarily invariant matrix norm and a Finsler length on the Grassmannian [14]. For an absolutely continuous unitary path define

$$\mathcal{A}_\Phi^{\mathrm{spr}}[S] = \int_a^b \Phi\big(\mathrm{Spr}_1^+(Q(t)), \ldots, \mathrm{Spr}_k^+(Q(t))\big)\, \mathrm{d}t, \tag{14}$$

$$\mathcal{A}_\Phi^{+}[S] = \int_a^b \Phi(\lambda_1(Q(t)), \ldots, \lambda_k(Q(t)))\, \mathrm{d}t, \tag{15}$$

where the second action is restricted to $Q \succeq 0$.

Theorem 4.1 (Universal gauge-optimal action)**.** *Let $N \geq 2k$, and let $\mathcal{X}, \mathcal{Y} \subset \mathbb{C}^N$ be rank-k subspaces with canonical-angle vector β. Among all absolutely continuous unitary paths transporting $\mathcal{X}$ to $\mathcal{Y}$,*

$$\boxed{\inf_S \mathcal{A}_\Phi^{\mathrm{spr}}[S] = 2\Phi(\beta).} \tag{16}$$

Restricting to positive generators gives the same exact value:

$$\boxed{\inf_{S:\, Q \succeq 0} \mathcal{A}_\Phi^{+}[S] = 2\Phi(\beta).} \tag{17}$$

One constant positive path attains both infima simultaneously for every symmetric gauge Φ.

Proof. Theorem 3.2 gives $2\beta \prec_w \int_a^b \mathrm{Spr}^+(Q(t))\mathrm{d}t$. Ky Fan dominance and the triangle inequality for Φ therefore imply

$$2\Phi(\beta) \leq \Phi\left(\int_a^b \mathrm{Spr}^+(Q(t))\, \mathrm{d}t\right) \leq \int_a^b \Phi(\mathrm{Spr}^+(Q(t)))\, \mathrm{d}t.$$

For $Q \succeq 0$, use Theorem 3.3 in the same way with the leading k eigenvalues. To attain both bounds, choose principal-vector pairs for $\mathcal{X}, \mathcal{Y}$ and use the constant coupled-plane generator in (19) below, conjugated into those pairs. Its spread and leading-delay vectors are both $2\beta/(b-a)$, so positive homogeneity gives action $2\Phi(\beta)$ for every Φ. □

This theorem identifies an exact resource value, rather than only a lower bound for a chosen norm. The same physical path is universally optimal:

Gauge $\Phi(x)$	Exact cost	Interpretation
$\lVert x\rVert_\infty$	$2\beta_1$	maximal-angle / bandwidth
$\sum_{j=1}^{r} x_j^{\downarrow}$	$2\sum_{j=1}^{r} \beta_j$	Ky Fan r modal budget
$\lVert x\rVert_2$	$2(\sum_j \beta_j^2)^{1/2}$	Euclidean / Hilbert–Schmidt budget
$\lVert x\rVert_1$	$2\sum_j \beta_j$	total angular action

The intermediate prefixes contain strictly more information than the two scalar endpoints. For example,

$$\beta = (1, 0.8, 0.2), \qquad \widetilde{\beta} = (1, 0.5, 0.5) \tag{18}$$

have the same largest angle and the same total angle. Hence the $r = 1$ and $r = 3$ laws assign identical costs. Yet the sharp $r = 2$ costs are 3.6 and 3.0, respectively—a 20% separation invisible to both scalar summaries. The construction below realizes both vectors, so this is an operational separation, not merely a comparison of unattained lower bounds.

5 Sharpness and the exact Pareto frontier

The hierarchy loses no constant, even when all prefixes are requested simultaneously.

Theorem 5.1 (Simultaneous saturation). *Fix $k \geq 1$, $\Delta > 0$ and any ordered vector $\beta_1 \geq \cdots \geq \beta_k \in [0, \pi/2]$. There exists a constant positive generator on $\mathbb{C}^{2k}$ whose path over an interval of width Δ has exactly these canonical angles and satisfies equality in* (10) *and* (11) *for every $r = 1, \ldots, k$.*

Proof. Let $D = \Delta^{-1}\operatorname{diag}(\beta_1, \ldots, \beta_k)$ and set

$$Q_\beta = \begin{bmatrix} D & -D \\ -D & D \end{bmatrix} \succeq 0, \qquad S(t) = \exp(\mathrm{i}tQ_\beta), \quad 0 \leq t \leq \Delta. \tag{19}$$

The eigenvalues of Q_β are $2\beta_1/\Delta, \ldots, 2\beta_k/\Delta$ and k zeros. Starting from the first coordinate sector, each paired plane evolves independently: its desired and complementary amplitudes have moduli $\cos(t\beta_j/\Delta)$ and $\sin(t\beta_j/\Delta)$. Hence the endpoint angles are β_j, while the first r integrated leading delays sum to $2\sum_{j=1}^{r}\beta_j$. The bottom-delay, block-inequality and path-length slacks all vanish. $\square$

Thus 2β is the exact lower Lorenz curve for the proper-delay action. This statement is stronger than showing that each scalar inequality is separately sharp: one physical device attains every prefix at once. Adding $c(t)I$ produces signed/gauge-related equality families without changing the routing.

6 Where excess memory goes: an exact slack identity

Near equality should have a diagnostic meaning rather than a vague "approximately optimal" label. Define, for each r,

$$\delta_r(t) = \sum_{j=1}^{r} \operatorname{Spr}_j^{+}(Q(t)) - 2\nu_r(B(t)) \geq 0, \tag{20}$$

$$g_r = \int_a^b \nu_r(B(t))\,\mathrm{d}t - \sum_{j=1}^{r}\beta_j \geq 0. \tag{21}$$

The first is algebraic coupling slack; the second is Grassmann detour slack.

Theorem 6.1 (Exact slack decomposition). *For every Hermitian flow and $r \leq m$,*

$$\int_a^b \sum_{j=1}^{r} \operatorname{Spr}_j^{+}(Q)\,\mathrm{d}t - 2\sum_{j=1}^{r}\beta_j = \int_a^b \delta_r(t)\,\mathrm{d}t + 2g_r. \tag{22}$$

If $Q \succeq 0$, then the leading-proper-delay surplus is

$$\int_a^b \sum_{j=1}^{r} \lambda_j(Q)\,\mathrm{d}t - 2\sum_{j=1}^{r}\beta_j \tag{23}$$

$$= \underbrace{\int_a^b \sum_{j=1}^{r} \lambda_{N-j+1}(Q)\,\mathrm{d}t}_{\textit{common/background delay}} + \underbrace{\int_a^b \delta_r(t)\,\mathrm{d}t}_{\textit{inefficient spectral coupling}} + \underbrace{2g_r}_{\textit{Grassmann detour}}.$$

All three terms are nonnegative.

Proof. Insert (20)–(21) and cancel $2\int \nu_r$. For (23), split the leading-delay sum into the spectral spread sum and the corresponding bottom-eigenvalue sum. $\square$

Corollary 6.2 (Quantitative near-rigidity). *Assume $Q \succeq 0$ and that the leading-delay surplus in* (23) *is at most ε. Then*

$$\int_a^b \sum_{j=1}^{r} \lambda_{N-j+1}(Q)\,\mathrm{d}t \le \varepsilon, \qquad \int_a^b \delta_r(t)\,\mathrm{d}t \le \varepsilon, \qquad g_r \le \frac{\varepsilon}{2}. \tag{24}$$

Moreover, for every $\tau > 0$, the sets on which the bottom-delay density or $\delta_r(t)$ exceeds τ each have measure at most ε/τ.

Proof. Each term in (23) is nonnegative, so each is bounded by their sum. The measure statements are Markov's inequality applied to the two nonnegative densities. $\square$

Consequently, a sequence of devices can approach the Pareto frontier only if all three nonnegative quantities vanish. This is a quantitative rigidity statement without asserting a fragile classification of all equality cases. In particular, adding large common delay cannot improve routing speed; it is exposed by the first term in (23).

7 Robust routing from heterogeneous leakage spectra

Let the output space split orthogonally as $A \oplus B$, with both sectors of dimension at least k. At the pass endpoint let $X_p \in \mathbb{C}^{N\times k}$ be an isometric output frame and define the ordered leakage singular values

$$p_1 \ge \cdots \ge p_k, \qquad p = s(R_B X_p). \tag{25}$$

At the stop endpoint define

$$q_1 \ge \cdots \ge q_k, \qquad q = s(R_A X_s). \tag{26}$$

Their leakage angles are $a_j = \arcsin p_j$ and $b_j = \arcsin q_j$.

Proposition 7.1 (Directional robust separation). *For $r = 1, \ldots, k$, put*

$$\Gamma_r = \left[\frac{r\pi}{2} - \sum_{j=1}^{r} a_j - \sum_{j=1}^{r} b_j\right]_+. \tag{27}$$

Then the endpoint canonical angles obey $\sum_{j=1}^{r} \beta_j \ge \Gamma_r$. Therefore

$$2\Gamma_r \le \int_I \sum_{j=1}^{r} \mathrm{Spr}_j^+(Q)\,\mathrm{d}t, \tag{28}$$

and, if $Q \succeq 0$, also $2\Gamma_r \le \int_I \sum_{j=1}^{r} \lambda_j(Q)\mathrm{d}t$.

Proof. The Ky Fan r distance from $\operatorname{ran} X_p$ to the Grassmannian of k-subspaces contained in A is $\sum_{j=1}^{r} a_j$; the corresponding distance at the stop endpoint to a subspace contained in B is $\sum_{j=1}^{r} b_j$. Every such pair of sector subspaces is orthogonal and has distance $r\pi/2$. The triangle inequality for d_r gives $d_r(\operatorname{ran} X_p, \operatorname{ran} X_s) \ge \Gamma_r$. Apply Theorems 3.2 and 3.3. $\square$

This retains heterogeneity. A single weak direction need not erase the certificate carried by the stronger ones. If only operator leakages $p_1 \le \varepsilon_p$ and $q_1 \le \varepsilon_s$ are known, then

$$\Gamma_r \ge r\alpha, \qquad \alpha = \left[\frac{\pi}{2} - \arcsin \varepsilon_p - \arcsin \varepsilon_s\right]_+. \tag{29}$$

This recovers the preceding scalar robust law at $r = 1$ and $r = k$, while the new hierarchy supplies every intermediate budget. The heterogeneous bound (27) is a triangle-inequality certificate; unlike the uniform family in Theorem 5.1, we do not claim it is sharp for every pair of heterogeneous leakage spectra.

7.1 Finite-error tomography

The bound can be made fail-closed. Suppose the measured pass signal block is $\widehat{G}_p = R_A X_p + E_p$ with $\|E_p\|_{\mathrm{op}} \le \eta_p$, and let $\widehat{g}_1 \ge \cdots \ge \widehat{g}_k$ be its singular values. Suppose the measured stop leakage block is $\widehat{G}_s = R_A X_s + E_s$ with $\|E_s\|_{\mathrm{op}} \le \eta_s$ and singular values $\widehat{h}_1 \ge \cdots \ge \widehat{h}_k$. Weyl's singular-value bound and $X_p^* R_A X_p + X_p^* R_B X_p = I$ give valid upper bounds

$$p_j^{\mathrm{cert}} = \sqrt{1 - [\widehat{g}_{k-j+1} - \eta_p]_+^2}, \tag{30}$$

$$q_j^{\mathrm{cert}} = \min\{1, \widehat{h}_j + \eta_s\}. \tag{31}$$

Clipping all displayed quantities to $[0, 1]$ is understood. Replacing p_j, q_j by these upper bounds in (27) can only decrease Γ_r, hence remains rigorous. The output is not one number but a certified lower Lorenz curve $r \mapsto 2\Gamma_r$ for the required delay spectrum.

8 Additivity and rational memory

Consider disjoint arcs $I_1, \ldots, I_L$ with endpoint leakage data and certificates $\Gamma_{\ell,r}$. Positivity of the integrands gives, for every fixed r,

$$\int_{\cup_\ell I_\ell} \sum_{j=1}^{r} \mathrm{Spr}_j^+(Q)\,\mathrm{d}t \ge 2 \sum_{\ell=1}^{L} \Gamma_{\ell,r}. \tag{32}$$

There is no cancellation between separated switches.

For a causal rational inner transfer matrix on the disk, the boundary Wigner–Smith matrix

$$Q(\theta) = -\mathrm{i} S(\mathrm{e}^{\mathrm{i}\theta})^* \partial_\theta S(\mathrm{e}^{\mathrm{i}\theta})$$

is positive, by its Blaschke–Potapov factorization [13, 4], and its trace satisfies the topological identity

$$\int_0^{2\pi} \mathrm{tr}\, Q(\theta)\,\mathrm{d}\theta = 2\pi\ \deg_{\mathrm{McM}}(S). \tag{33}$$

This is the stored-memory law developed in [7].

Corollary 8.1 (Hierarchy of degree obstructions)**.** *Let S be rational inner of McMillan degree n, and suppose L disjoint routing arcs have certificates $\Gamma_{\ell,r}$. Then*

$$\boxed{n \ge \left\lceil \frac{1}{\pi} \max_{1 \le r \le k} \sum_{\ell=1}^{L} \Gamma_{\ell,r} \right\rceil.} \tag{34}$$

For $2M$ alternating transitions with uniform leakages,

$$n \ge \left\lceil \frac{2Mk}{\pi} \left[\frac{\pi}{2} - \arcsin \varepsilon_p - \arcsin \varepsilon_s\right]_+ \right\rceil, \tag{35}$$

and at zero error $n \ge Mk$.

Proof. For every r, combine (32) with $\sum_{j=1}^r \mathrm{Spr}_j^+(Q) \le \sum_{j=1}^r \lambda_j(Q) \le \mathrm{tr}\, Q$, integrate over the circle, and use (33). Maximize over r and use that n is an integer. The uniform specialization follows from (29) with $r = k$ and $L = 2M$. □

The scalar degree lower bound is numerically the same as in the fifth paper when all directions have one common tolerance. The advance is structural: heterogeneous data choose the strongest prefix automatically, and (32) simultaneously dictates how the memory must appear in the leading proper-delay spectrum on every arc. Realization degree alone cannot express that allocation.

9 Physical interpretation and scope

Wigner and Smith introduced scattering delay through energy/frequency derivatives of phase and the lifetime matrix [18, 16]. Its eigenvalues—proper delays—are directly meaningful in multiport wave systems [12]. In passive quantum linear systems, internal modes implement a rational inner scattering law, while feedback and propagation delay may produce more general inner responses [9, 17].

The theorem applies at the scattering layer. If selected ports couple a system to reservoirs with a weak-coupling Markovian reduction, the squared transfer amplitudes enter Kossakowski rates. A routing Lorenz curve therefore limits how rapidly a passive architecture can redistribute decoherence among orthogonal reservoir sectors. This interpretation requires the stated system–bath model; the geometric theorem itself does not assume Markovianity and does not claim that Wigner–Smith delay is literally decoherence time.

Three distinctions from standard QSLs are useful.

1. The evolution parameter may be frequency rather than laboratory time. The generator is Wigner–Smith delay, not necessarily a system Hamiltonian.

2. The resource is the full spectral-spread Lorenz curve. The maximal-angle/bandwidth result is only its first prefix.

3. For rational passive devices, integrating the trace closes the loop to an integer realization resource through (33).

The signed theorem is also a genuine time-dependent subspace QSL. It is compatible with the maximal-angle theory of [2, 1] and with geometric operator-flow bounds [5, 10], but answers a different question: every partial sum of canonical angles is charged to the matching partial sum of instantaneous spectral spreads. Hilbert–Schmidt bounds instead aggregate quadratically, and the operator norm sees only $r = 1$.

10 Reproducible adversarial audit

The results are analytic; computation is used to attack implementation mistakes and hidden ordering assumptions. The immutable artifact runs six deterministic suites:

1. four coupled-mode equality families of ranks $1, 2, 4, 6$, checking all prefixes and the $\ell_\infty, \ell_2, \ell_1$ variational optima to machine precision;

2. 4,096 random signed Hermitian generators, testing 11,697 pointwise partial-sum inequalities;

3. 4,096 random positive generators, testing 11,905 inequalities;

4. 768 signed and 768 positive piecewise-constant paths, testing 4,653 endpoint hierarchy inequalities after exact matrix exponentials;

5. 2,048 heterogeneous noisy tomography instances, verifying that the fail-closed certificate never exceeds the exact leakage certificate; and

6. 4,096 nonnegative slack decompositions, checking the identity independently of the matrix calculations, together with the strict scalar-endpoint separation in (18).

The largest observed pointwise ratios were 0.99435 (signed) and 0.98134 (positive); the largest path ratios were 0.98068 and 0.88671. These random ratios are not estimates of optimal constants: the explicit family (19) attains ratio one for every prefix. Their role is to probe generic noncommuting, nongeodesic cases. An independent verifier recomputes the equality arithmetic, launches a separately coded matrix probe, checks minimum campaign sizes and refuses any negative slack. Continuous integration is not certified by gridding: paths are piecewise constant, so their action integrals are exact.

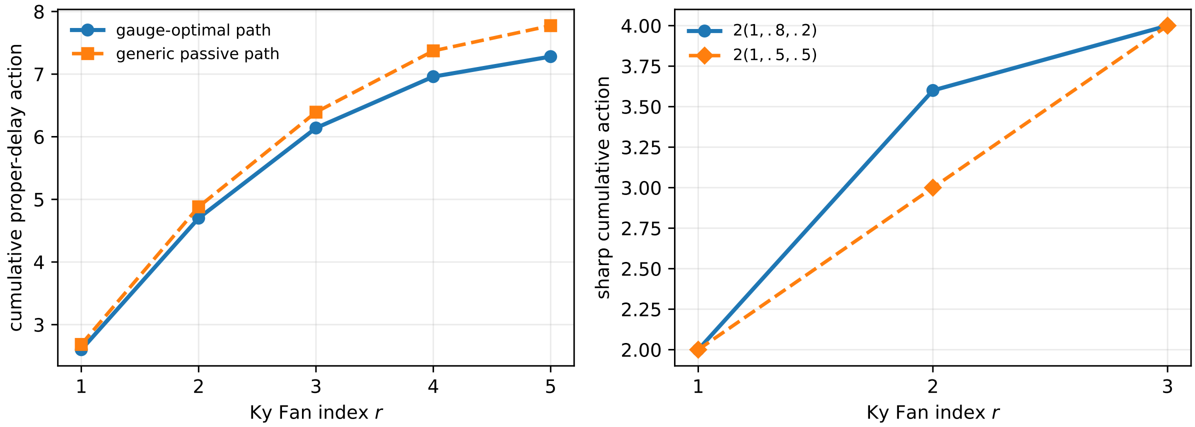

Figure 1: Left: the required cumulative action $2\sum_{j\le r}\beta_j$ is an exact lower Lorenz curve; generic passive motion carries nonnegative surplus resolved by Theorem 6.1. Right: two angle spectra with equal maximum and total cost have different sharp intermediate prefixes, proving strict information gain over the two scalar endpoint laws.

11 Limitations, falsifiers and open directions

Limitation 11.1 (No positivity from unitarity alone). A general unitary path has Hermitian Q but need not have $Q \succeq 0$. The spectral-spread theorem still holds; the proper-delay and degree corollaries require positivity. A measured negative proper delay is a falsifier for applying those passive-monotone corollaries, not for Theorem 3.2.

Limitation 11.2 (Boundary model). The theorem constrains a unitary boundary scattering path. Loss, gain, unobserved ports and calibration drift must be incorporated through a larger unitary dilation or an explicit error model. Treating a visibly contractive subblock as if it were the full unitary would invalidate the generator identity.

Limitation 11.3 (Lower bound, not synthesis classification). The coupled-mode construction proves all constants optimal but does not classify every saturating path or solve degree-minimal interpolation for arbitrary heterogeneous samples. Equation (23) provides the quantitative componentwise stability bounds in Theorem 6.2, but does not yet convert small slack into operator-norm distance from the complete set of optimizers.

Promising next questions are: a stability theorem converting small total slack into distance from the coupled-mode geodesic family; dilation-aware certificates for measured lossy scattering matrices; and a synthesis theorem characterizing which proper-delay Lorenz curves above the sharp lower frontier are realizable by rational inner networks of fixed degree.

12 Conclusion

Subspace rotation consumes a spectrum, not a scalar. The complete law is $2\beta \prec_w \int \mathrm{Spr}^+(Q)$: every prefix of canonical-angle motion must be paid for by the matching prefix of generator spectral spread. Positive Wigner–Smith flows turn this into a hierarchy of leading proper-delay budgets. Equivalently, every symmetric-gauge action has the exact minimum $2\Phi(\beta)$, and one positive coupled-mode geodesic is optimal for all gauges at once. The hierarchy admits an exact three-term slack decomposition and quantitative near-rigidity bounds, survives heterogeneous leakage and finite tomography error, adds over separated switches and closes to an integer McMillan-degree obstruction for rational passive networks. It therefore upgrades a trace memory law into a mode-resolved, experimentally certifiable resource theory.

Data and code availability

The manuscript source, frozen JSON certificate, generated figure, independent verifier, release manifest and continuous-integration workflow are archived at the immutable GitHub release. The principal claims are proved analytically in the text; numerical tests are adversarial audits, not substitutes for the proofs.

References

[1] Sergio Albeverio and Alexander K. Motovilov. Optimal bounds on the speed of subspace evolution. *Journal of Physics A: Mathematical and Theoretical*, 55:215301, 2022.

[2] Sergio Albeverio and Alexander K. Motovilov. Quantum speed limits for time evolution of a system subspace. *Physics of Particles and Nuclei*, 53(2):287–291, 2022.

[3] David Allan, Niklas Hörnedal, and Ole Andersson. Time-optimal quantum transformations with bounded bandwidth. *Quantum*, 5:462, 2021.

[4] Daniel Alpay, Palle Jorgensen, and Izchak Lewkowicz. Characterizations of rectangular (para)-unitary rational functions. *Opuscula Mathematica*, 36(6):695–716, 2016.

[5] Nicoletta Carabba, Niklas Hörnedal, and Adolfo del Campo. Quantum speed limits on operator flows and correlation functions. *Quantum*, 6:884, 2022.

[6] Lluis Eriksson. Every spectral switch costs memory: Sharp robust wigner–smith speed limits for passive quantum networks. GitHub release, 2026. Immutable artifact.

[7] Lluis Eriksson. Lossless calibration is stored memory: A topological mcmillan-degree and wigner–smith law for passive quantum networks. GitHub release, 2026. Immutable artifact.

[8] Ky Fan. Maximum properties and inequalities for the eigenvalues of completely continuous operators. *Proceedings of the National Academy of Sciences USA*, 37:760–766, 1951.

[9] John E. Gough and Guofeng Zhang. On realization theory of quantum linear systems. *Automatica*, 59:139–151, 2015.

[10] Niklas Hörnedal, Nicoletta Carabba, Kazutaka Takahashi, and Adolfo del Campo. Geometric operator quantum speed limit, wegner hamiltonian flow and operator growth. *Quantum*, 7:1055, 2023.

[11] Pedro Massey, Demetrio Stojanoff, and Sebastián Zárate. The spectral spread of hermitian matrices. *Linear Algebra and its Applications*, 616:19–44, 2021.

[12] Utkarsh R. Patel and Eric Michielssen. Wigner–smith time-delay matrix for electromagnetics: Theory and phenomenology. *IEEE Transactions on Antennas and Propagation*, 69(2):902–917, 2021.

[13] V. P. Potapov. The multiplicative structure of J-contractive matrix functions. *Trudy Moskovskogo Matematicheskogo Obshchestva*, 4:125–236, 1955. In Russian; English translation in AMS Translations, Series 2, Vol. 15.

[14] Li Qiu, Yanxia Zhang, and Chi-Kwong Li. Unitarily invariant metrics on the Grassmann space. *SIAM Journal on Matrix Analysis and Applications*, 27(2):507–531, 2005.

[15] Benjamin Russell and Susan Stepney. The geometry of speed limiting resources in physical models of computation. *International Journal of Foundations of Computer Science*, 28(4):321–333, 2017.

[16] Felix T. Smith. Lifetime matrix in collision theory. *Physical Review*, 118:349–356, 1960.

[17] Gil Tabak and Hideo Mabuchi. Trapped modes in linear quantum stochastic networks with delays. *EPJ Quantum Technology*, 3:3, 2016.

[18] Eugene P. Wigner. Lower limit for the energy derivative of the scattering phase shift. *Physical Review*, 98:145–147, 1955.