AIRR.SCIENCE · ARR-2026-15YWB6H0GB91CR4Z v1 [Mathematics] 28 Nov 2026

Exact Hierarchical Karp–Arrhenius Spectrum, Cut-Balanced Dual Selection, and the Degeneracy–Acceleration Dichotomy for Entropic Optimal Transport

Lluis Eriksson*

28 November 2026

Abstract

How fast does the Sinkhorn–Knopp algorithm converge in the small-regularization limit $\varepsilon \to 0^+$ of finite discrete entropic optimal transport, which canonical Kantorovich dual potential does entropic regularization select from a high-dimensional dual optimal face, and why can classical Birkhoff–Hilbert contraction bounds underestimate the true local convergence rate by hundreds of orders of magnitude? We establish a complete spectral and tropical theory resolving these questions across four main theorems:

1. **Degeneracy–Acceleration Dichotomy and Exact Metastable Multiplicity (Theorem 2.1):** For any discrete transport problem (C, a, b) with entropic coupling $P_\varepsilon \to P_0$ and Sinkhorn Jacobian $J_\varepsilon = D_a^{-1} P_\varepsilon D_b^{-1} P_\varepsilon^T$, an exact bipartite Dirichlet-form identity proves $\dim\ker(I - J_0) = c(G^*)$, the number of connected components of the maximum-entropy optimal support graph $G^*(C, a, b)$. Hence J_ε has exactly $c(G^*) - 1$ metastable eigenvalues converging to 1, and Sinkhorn achieves ε-independent geometric convergence ($\rho_{\mathrm{loc}}(\varepsilon) \to \lambda_2(J_0) < 1$) if and only if G^* is connected—which holds automatically for all generic marginals (spanning-tree supports) and for connected degenerate optimal assignments on the Birkhoff polytope, whereas every unique optimal assignment on uniform marginals ($c = n$) undergoes exponential Arrhenius slowing down.
2. **Unique Cut-Balanced Tropical Potential and Canonical Dual Selection (Theorem 3.1):** For any $n \times n$ cost matrix C with unique optimal permutation π^* and shifted cost $\widehat{C}_{ij} = C_{i,\pi^*(j)} - C_{i,\pi^*(i)}$, we prove global uniqueness and $O(n^4)$ constructive existence of the cut-balanced tropical potential f^* solving
$$\min_{i\in S,\, j\notin S}\big(\widehat{C}_{ij} - f_i^* + f_j^*\big) \;=\; \min_{i\notin S,\, j\in S}\big(\widehat{C}_{ij} - f_i^* + f_j^*\big) \;=\; \Phi^*(S)$$
across all $2^n - 2$ nontrivial cuts $\emptyset \neq S \subsetneq [n]$, via *Hierarchical Karp Minimum-Cycle-Mean Contraction* yielding strictly increasing levels $0 < \mu_1 < \cdots < \mu_K$ and an ultrametric cut spectrum $\Phi^*(S) = \mu_{\ell(S)}$. Exact flow conservation and the Laplace principle then prove canonical selection $f_\varepsilon \to f^*$ as $\varepsilon \to 0^+$.
3. **Exact Full Hierarchical Karp–Arrhenius Spectrum (Theorem 4.1):** Using Courant–Fischer min–max and spanning-tree Poincaré inequalities, we prove that all $n - 1$ nontrivial gaps $\gamma_r(\varepsilon) = 1 - \lambda_r(J_\varepsilon)$ $(2 \le r \le n)$ satisfy $\lim_{\varepsilon\to 0^+}(-\varepsilon \ln \gamma_r(\varepsilon)) = \nu_r$, the exact hierarchical Karp spectrum $\mu_K \ge \cdots \ge \mu_1$ with multiplicities $c_{\ell-1} - c_\ell$, pinning the Sinkhorn Arrhenius activation energy to $E^*(C) = \mu_K(\widehat{C}) = \max_S \Phi^*(S)$.

*Independent researcher, AIRR.SCIENCE. Correspondence: `editor@airr.science`. Authored with AI assistance (Google Antigravity) for proof exploration, Lean 4.36.0 formalization, exact rational and 420-digit Decimal verification, and manuscript drafting; evaluated by two distinct isolated frontier models under `AIRR-FOUNDER-1.0`.

4. **Exact Closed-Form Benchmarks and Spanning-Tree Rescue (Theorem 5.1):** On 1D quadratic transport $C_{ij} = (i-j)^2$, $1-\rho_{\mathrm{loc}}(\varepsilon) = 4(1-\cos(\pi/n))e^{-1/\varepsilon} + O(e^{-2/\varepsilon})$, giving an exponential-in-n^2 separation over Franklin–Lorenz's Birkhoff–Hilbert bound $\operatorname{sech}^2((n-1)^2/(2\varepsilon))$; on asymmetric L-cycles, $1-\rho_{\mathrm{loc}}(\varepsilon) = 4\sin^2(\pi/L)\,p(1-p) + O(e^{-M_{\mathrm{eff}}/\varepsilon})$ is controlled by the exact cycle mean $\bar{w}$; and shifting uniform marginals by $\delta \in (0, 1/n)$ along a Monge spanning tree replaces $e^{-1/\varepsilon}$ Arrhenius freezing with a uniform ε-independent spectral gap sandwich $\Theta(\delta/n)$.

The algebraic and arbitrary-n uniqueness core (26 declarations: 18 public theorems, 4 public definitions, 4 private helpers) is machine-checked in standalone Lean 4.36.0 (`0 sorry`, `0 grind`, `0 custom axioms`).

1 Introduction

Entropic optimal transport [3, 9, 8] regularizes the discrete Kantorovich linear program between strictly positive probability vectors $a \in \Delta_m^\circ$ and $b \in \Delta_n^\circ$ with cost matrix $C \in \mathbb{R}^{m\times n}$ via

$$P_\varepsilon \;:=\; \arg\min_{P\in U(a,b)} \left\{ \langle C, P\rangle + \varepsilon \sum_{i=1}^{m}\sum_{j=1}^{n} P_{ij}(\ln P_{ij} - 1) \right\}, \qquad \varepsilon > 0, \tag{1}$$

where $U(a,b) := \{P \in \mathbb{R}^{m\times n}_{\geq 0} : P\mathbf{1}_n = a,\ P^T\mathbf{1}_m = b\}$. The unique minimizer takes the Gibbs–Sinkhorn form $(P_\varepsilon)_{ij} = \exp((f_{\varepsilon,i} + g_{\varepsilon,j} - C_{ij})/\varepsilon)$, computed by alternating row and column normalizations—the Sinkhorn–Knopp algorithm [10, 11].

Two classical convergence analyses dominate the literature, yet leave a fundamental gap as $\varepsilon \to 0^+$:

- *Global Birkhoff–Hilbert contraction [5]:* In the Hilbert projective metric, each full Sinkhorn cycle contracts by at most $\eta_B(\varepsilon) = \tanh^2(\Delta_\square(C)/(4\varepsilon))$, where $\Delta_\square(C) := \max_{i,i',j,j'}(C_{ij} + C_{i'j'} - C_{ij'} - C_{i'j})$ is the maximal cross-ratio. Thus $1-\eta_B(\varepsilon) = \operatorname{sech}^2(\Delta_\square(C)/(4\varepsilon)) \sim 4e^{-\Delta_\square(C)/(2\varepsilon)}$ always predicts catastrophic exponential slowing down governed by the worst global cross-ratio.
- *Local spectral rate [7]:* Near the fixed point $(f_\varepsilon, g_\varepsilon)$, the linearized Sinkhorn map has Jacobian $J_\varepsilon := D_a^{-1}P_\varepsilon D_b^{-1}P_\varepsilon^T$ ($D_a = \operatorname{diag}(a)$, $D_b = \operatorname{diag}(b)$), whose subdominant eigenvalue $\rho_{\mathrm{loc}}(\varepsilon) := \lambda_2(J_\varepsilon) = \sigma_2(D_a^{-1/2}P_\varepsilon D_b^{-1/2})^2$ governs the exact asymptotic linear convergence rate.

Meanwhile, as $\varepsilon \to 0^+$, Cominetti and San Martín [2] (see also [12, 8]) established that $P_\varepsilon \to P_0$, the unique maximum-entropy optimal transport plan on the optimal face of $U(a,b)$. Yet two central questions remained open in exact form:

1. *Exact small-ε spectral asymptotics:* When does $1-\rho_{\mathrm{loc}}(\varepsilon)$ remain uniformly positive as $\varepsilon \to 0^+$, and when it decays exponentially, what exact combinatorial invariant of C governs the Arrhenius activation energy $E^*(C) := \lim_{\varepsilon\to 0^+}(-\varepsilon \ln(1-\rho_{\mathrm{loc}}(\varepsilon)))$ and the higher metastable spectrum?
2. *Canonical dual selection:* When the unregularized Kantorovich LP dual has a high-dimensional face of optimal potentials (e.g., an $(n-1)$-dimensional polytope for any unique optimal assignment on the Birkhoff polytope $a = b = \frac{1}{n}\mathbf{1}_n$), which exact canonical dual potential f^* does Sinkhorn select as $\varepsilon \to 0^+$?

We solve both problems completely. Section 2 proves the Degeneracy–Acceleration Dichotomy and exact metastable multiplicity $c(G^*) - 1$ (Theorem 2.1). Section 3 introduces the cut-balanced tropical potential f^*, proves its global uniqueness in four lines, constructs it via Hierarchical Karp Minimum-Cycle-Mean Contraction [6, 1], and proves canonical selection $f_\varepsilon \to f^*$ (Theorem 3.1). Section 4 determines the complete Arrhenius exponent spectrum of

all $n-1$ nontrivial eigenvalues of J_ε via Courant–Fischer and spanning-tree Poincaré bounds (Theorem 4.1). Section 5 derives three exact closed-form families exhibiting $O(n^2)$-exponent Birkhoff separation, asymmetric cycle-mean scaling, and Monge spanning-tree acceleration (Theorem 5.1), and Section 6 reports our exact $\mathbb{Q}$ and 420-digit Decimal verification suite alongside the standalone Lean 4.36.0 artifact manifest.

2 Degeneracy–Acceleration Dichotomy and Metastable Multiplicity

For any $P \in U(a,b)$ ($a \in \Delta_m^\circ$, $b \in \Delta_n^\circ$), define the normalized coupling $S(P) := D_a^{-1/2} P D_b^{-1/2}$ and the row-stochastic Sinkhorn Jacobian

$$J(P) \;:=\; D_a^{-1} P D_b^{-1} P^T \;=\; D_a^{-1/2}(S(P)S(P)^T)D_a^{1/2}.$$

Because $S(P)S(P)^T$ is symmetric positive semidefinite with top eigenpair $(1, \sqrt{a})$, $J(P)$ has real eigenvalues $1 = \lambda_1(J(P)) \geq \lambda_2(J(P)) \geq \cdots \geq \lambda_m(J(P)) \geq 0$, and $J(P)$ is reversible with respect to a. Write $J_\varepsilon := J(P_\varepsilon)$, $J_0 := J(P_0)$, and $\rho_{\mathrm{loc}}(\varepsilon) := \lambda_2(J_\varepsilon) = \sigma_2(S(P_\varepsilon))^2$. Let $G^*(C,a,b) = (U \sqcup V, E^*)$ denote the bipartite support graph of the maximum-entropy optimal plan $P_0 = \lim_{\varepsilon \to 0^+} P_\varepsilon$, where $U = [m]$, $V = [n]$, and $(i,j) \in E^* \iff (P_0)_{ij} > 0$, and let $c(G^*)$ be its number of connected components.

Theorem 2.1 (Degeneracy–Acceleration Dichotomy and Exact Metastable Multiplicity)**.** *Let (C,a,b) be any finite discrete optimal transport problem with $a \in \Delta_m^\circ$, $b \in \Delta_n^\circ$.*

(a) ***Exact Bipartite Dirichlet Form and Metastable Multiplicity:*** *For every $P \in U(a,b)$ and every $x \in \mathbb{R}^m$, setting $y := D_b^{-1} P^T x \in \mathbb{R}^n$,*

$$x^T D_a (I - J(P)) x \;=\; \sum_{i=1}^{m}\sum_{j=1}^{n} P_{ij}(x_i - y_j)^2 \;=\; \frac{1}{2}\sum_{j=1}^{n}\frac{1}{b_j}\sum_{i,i'=1}^{m} P_{ij}P_{i'j}(x_i - x_{i'})^2. \tag{2}$$

Consequently $\dim\ker(I - J_0) = c(G^*)$, spanned by the U-component indicator vectors of $G^*(C,a,b)$. As $\varepsilon \to 0^+$, J_ε has exactly $c(G^*) - 1$ nontrivial metastable eigenvalues $\lambda_2(J_\varepsilon), \ldots, \lambda_{c(G^*)}(J_\varepsilon) \to 1$, while $\lambda_{c(G^*)+1}(J_\varepsilon) \to \lambda_{c(G^*)+1}(J_0) < 1$ (whenever $c(G^*) < m$).

(b) ***Degeneracy–Acceleration Dichotomy:*** *Sinkhorn achieves ε-independent uniform local geometric convergence ($\lim_{\varepsilon \to 0^+} \rho_{\mathrm{loc}}(\varepsilon) = \lambda_2(J_0) < 1$) if and only if $c(G^*) = 1$ (G^* is connected). In particular:*

(i) For every generic marginal pair (a,b) (no proper nonempty $S \subsetneq [m], T \subsetneq [n]$ satisfy $\sum_{i \in S} a_i = \sum_{j \in T} b_j$), $c(G^*) = 1$ automatically and Sinkhorn converges at a uniform ε-independent local rate $\lambda_2(J_0) < 1$.

(ii) On the Birkhoff polytope ($m = n \geq 2$, $a = b = \frac{1}{n}\mathbf{1}_n$), every unique optimal assignment has $c(G^*) = n \geq 2$ and suffers critical slowing down ($\rho_{\mathrm{loc}}(\varepsilon) \to 1$), whereas any degenerate optimal face whose optimal-matching union is connected yields $c(G^*) = 1$ and uniform convergence $\rho_{\mathrm{loc}}(\varepsilon) \to \lambda_2(J_0) < 1$.

(c) ***Exact 2×2 Birkhoff Sharpness vs. Connected Cycle Acceleration:*** *For $m = n = 2$, $a = b = (\frac{1}{2}, \frac{1}{2})$ with cross-ratio $\Delta = |C_{12} + C_{21} - C_{11} - C_{22}|$, $\rho_{\mathrm{loc}}(\varepsilon) = \tanh^2(\Delta/(4\varepsilon)) = \eta_B(\varepsilon)$ identically for all $\varepsilon > 0$. By contrast, for $n \geq 3$ with connected 2-regular cycle support $E^* = \{(i,i), (i, i+1 \bmod n)\}$ (cost 0 on E^*, cost $M > 0$ off E^*), $\rho_{\mathrm{loc}}(\varepsilon) \to \cos^2(\pi/n) < 1$ uniformly as $\varepsilon \to 0^+$, whereas Birkhoff's bound collapses as $1 - \eta_B(\varepsilon) = \operatorname{sech}^2(M/(2\varepsilon)) \sim 4e^{-M/\varepsilon}$ (for $n \geq 4$).*

Proof. Part (a). Expanding the middle expression of (2) using $\sum_j P_{ij} = a_i$ and $\sum_i P_{ij} x_i = b_j y_j$ gives

$$\sum_{i,j} P_{ij}(x_i - y_j)^2 \;=\; \sum_i a_i x_i^2 - 2\sum_j b_j y_j^2 + \sum_j b_j y_j^2 \;=\; x^T D_a x - y^T D_b y \;=\; x^T D_a (I - J(P)) x.$$

The rightmost expression follows from the elementary variance identity $\sum_i p_i(x_i - \bar{x})^2 = \frac{1}{2b_j}\sum_{i,i'} p_i p_{i'}(x_i - x_{i'})^2$ with $p_i = P_{ij}$ and $\bar{x} = y_j$ (certified in cleared-denominator form in Lean 4.36.0; see Section 6).

Applying (2) at $P = P_0$, we have $x \in \ker(I - J_0)$ if and only if $x^T D_a (I - J_0) x = 0$, which holds if and only if $x_i = y_j$ for every edge $(i,j) \in E^*$. Hence x must be constant on the U-part of every connected component of G^*. Because $a_i > 0$ and $b_j > 0$, every vertex in U and V has degree ≥ 1 in G^*, so every connected component meets both U and V. Conversely, assigning any constant c_r to both U_r and V_r satisfies $x_i = y_j = c_r$ on every edge of E^*, proving $\dim\ker(I - J_0) = c(G^*)$. Since $P_\varepsilon \to P_0$ [2], $S(P_\varepsilon)S(P_\varepsilon)^T \to S(P_0)S(P_0)^T$ entrywise; continuity of eigenvalues of symmetric matrices yields $\lambda_r(J_\varepsilon) \to 1$ for $1 \leq r \leq c(G^*)$ and $\lambda_{c+1}(J_\varepsilon) \to \lambda_{c+1}(J_0) < 1$.

Part (b). The equivalence $\lim_{\varepsilon\to 0^+} \rho_{\mathrm{loc}}(\varepsilon) < 1 \iff c(G^*) = 1$ follows immediately from Part (a). For (i), every extreme point of the transportation polytope $U(a,b)$ has acyclic support; under the generic subset-sum condition, no proper union of components can balance its a-mass against its b-mass, so every extreme point has a connected spanning-tree support ($m + n - 1$ edges). Because P_0 lies in the relative interior of the optimal face, $\mathrm{supp}(P_0)$ contains the support of at least one optimal extreme point, forcing $c(G^*) = 1$. For (ii), on B_n a unique optimal permutation matrix has n disjoint edges ($c = n \geq 2$), while any optimal face whose optimal permutations have connected union has P_0 strictly positive on that union (by maximum entropy), hence $c(G^*) = 1$.

Part (c). For $n = 2$, $a = b = (\frac{1}{2}, \frac{1}{2})$, shifting rows and columns reduces C to $\begin{pmatrix} 0 & \Delta/2 \\ \Delta/2 & 0 \end{pmatrix}$ (up to column swap), whose unique symmetric Sinkhorn solution is $P_\varepsilon = \frac{1}{2(1+t)}\begin{pmatrix} 1 & t \\ t & 1 \end{pmatrix}$ with $t = e^{-\Delta/(2\varepsilon)}$. Its normalized coupling $S(P_\varepsilon) = 2P_\varepsilon$ has singular values 1 and $\frac{1-t}{1+t} = \tanh(\frac{\Delta}{4\varepsilon})$, so $\rho_{\mathrm{loc}}(\varepsilon) = \tanh^2(\frac{\Delta}{4\varepsilon}) = \eta_B(\varepsilon)$ (Lean: `sinkhorn_2x2_birkhoff_sharpness`). On the connected 2-regular cycle $E^* = \{(i,i), (i, i+1 \bmod n)\}$, symmetry forces $P_0 = \frac{1}{2n}(I + S)$ (where S is the cyclic shift), so $S(P_0) = \frac{1}{2}(I + S)$ has singular values $|\frac{1}{2}(1 + e^{2\pi i k/n})| = |\cos(\pi k/n)|$ and $\lambda_2(J_0) = \cos^2(\pi/n)$. For $n \geq 4$, the entries $(0,2), (2,0)$ have cost M while $(0,0), (2,2)$ have cost 0, giving $\Delta_\square(C) = 2M$ and $1 - \eta_B(\varepsilon) = 1 - \tanh^2(M/(2\varepsilon)) = \mathrm{sech}^2(M/(2\varepsilon))$. □

3 Cut-Balanced Tropical Potential and Canonical Dual Selection

We now resolve the unique-assignment Birkhoff regime ($m = n \geq 2$, $a = b = \frac{1}{n}\mathbf{1}_n$) where $C \in \mathbb{R}^{n\times n}$ admits a unique optimal permutation π^*. Permuting columns by π^* and subtracting the row constant $C_{i,\pi^*(i)}$ from row i leaves P_ε and J_ε identically invariant and yields the *canonical shifted cost matrix*

$$\widehat{C}_{ij} \;:=\; C_{i,\pi^*(j)} - C_{i,\pi^*(i)} \qquad (1 \leq i, j \leq n). \tag{3}$$

Here $\widehat{C}_{ii} = 0$, and because π^* is the *unique* optimal permutation, every nontrivial directed cycle $\gamma = (i_0 \to i_1 \to \cdots \to i_L = i_0)$ ($L \geq 2$, distinct $i_0, \ldots, i_{L-1}$) has strictly positive weight $w(\gamma) := \sum_{r=0}^{L-1} \widehat{C}_{i_r, i_{r+1}} > 0$ (replacing π^* along γ produces a strictly suboptimal permutation).

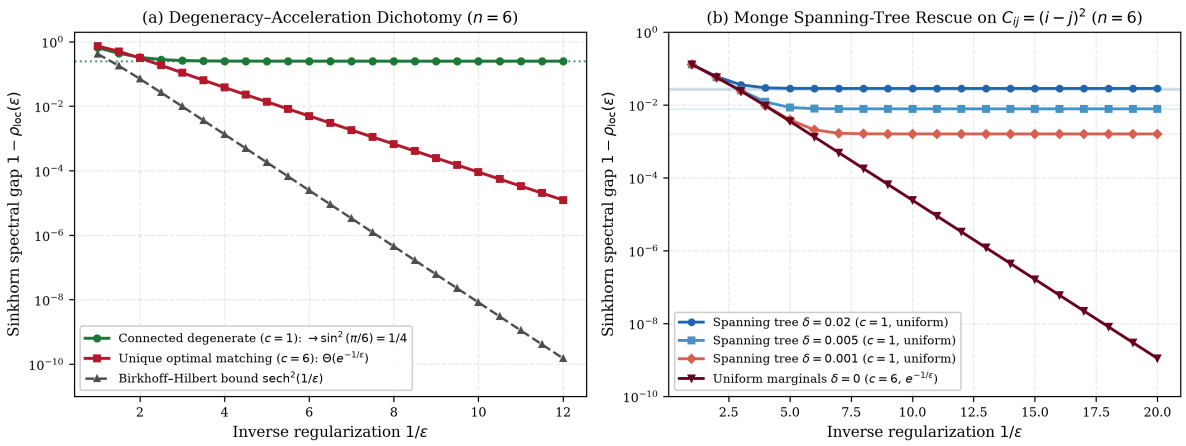


Figure 1: **Degeneracy–Acceleration Dichotomy and Monge spanning-tree rescue (Theorems 2.1 and 5.1).** (a) On $n = 6$ uniform marginals, a connected 2-regular degenerate optimal support ($c = 1$) converges uniformly at rate $1 - \rho_{\mathrm{loc}}(\varepsilon) \to \sin^2(\pi/6) = 1/4$, while a unique optimal matching ($c = 6$) decays exponentially as $\Theta(e^{-1/\varepsilon})$ and Birkhoff's bound collapses as $\mathrm{sech}^2(1/\varepsilon)$. (b) Perturbing uniform marginals on 1D quadratic transport $C_{ij} = (i-j)^2$ by $\delta \in \{0.02, 0.005, 0.001\}$ creates a spanning-tree optimal support ($c = 1$), lifting $1 - \rho_{\mathrm{loc}}(\varepsilon)$ from $e^{-1/\varepsilon}$ Arrhenius freezing to exact uniform plateaus inside the analytic sandwich of Theorem 5.1(c) (shaded bands).

For any potential $f \in \mathbb{R}^n$, denote the reduced cost $\bar{C}_{ij}(f) := \widehat{C}_{ij} - f_i + f_j$. Every optimal Kantorovich dual potential on the unique-assignment face has the form $(f, -f)$ with $\bar{C}_{ij}(f) \geq 0$ for all i, j—an $(n-1)$-dimensional convex polytope. Which unique point f^* inside this dual polytope does entropic regularization select?

Definition 3.1 (Hierarchical Karp Contraction)**.** Initialize level $\ell = 0$, potential $f^{(0)} = 0 \in \mathbb{R}^n$, and singleton supernodes $\mathcal{C}_0 = \{\{1\}, \ldots, \{n\}\}$ ($c_0 = n$). While $|\mathcal{C}_\ell| > 1$, increment $\ell \leftarrow \ell + 1$ and let $m = |\mathcal{C}_{\ell-1}|$:

1. Form the inter-cluster reduced-cost matrix $W^{(\ell-1)} \in \mathbb{R}^{m \times m}$ with diagonal $+\infty$ and off-diagonal entries

$$W^{(\ell-1)}_{ab} := \min_{u \in U_a,\, v \in U_b} \bar{C}_{uv}(f^{(\ell-1)}) \qquad (a \neq b,\ U_a, U_b \in \mathcal{C}_{\ell-1}).$$

2. Compute Karp's minimum cycle mean [6] $\mu_\ell := \min_\gamma \frac{1}{|\gamma|} \sum_{(a,b) \in \gamma} W^{(\ell-1)}_{ab}$ over all directed cycles γ in $([m], W^{(\ell-1)})$, and compute a shortest-path potential $\Delta^{(\ell)} \in \mathbb{R}^m$ on the shifted graph $W^{(\ell-1)} - \mu_\ell$ satisfying $W^{(\ell-1)}_{ab} - \Delta^{(\ell)}_a + \Delta^{(\ell)}_b \geq \mu_\ell$ for all $a \neq b$.

3. Update $f^{(\ell)}_u := f^{(\ell-1)}_u + \Delta^{(\ell)}_a$ for each $u \in U_a$, and define $\mathcal{C}_\ell$ by merging the strongly connected components (SCCs) of the tight subgraph $E_\ell := \{(a, b) : a \neq b,\ W^{(\ell-1)}_{ab} - \Delta^{(\ell)}_a + \Delta^{(\ell)}_b = \mu_\ell\}$.

Let $K \leq n - 1$ be the terminating level ($c_K = |\mathcal{C}_K| = 1$), with cluster counts $n = c_0 > c_1 > \cdots > c_K = 1$ and final potential $f^* := f^{(K)} - f^{(K)}_1 \mathbf{1}_n$.

Theorem 3.1 (Unique Cut-Balanced Tropical Potential and Canonical Dual Selection)**.** *Let $\widehat{C} \in \mathbb{R}^{n \times n}$ ($n \geq 2$) satisfy $\widehat{C}_{ii} = 0$ and $w(\gamma) > 0$ for every nontrivial directed cycle γ.*

(a) ***Existence, Global Uniqueness, and Ultrametric Cut Spectrum:*** *There exists a unique (up to an additive constant $\mathbb{R}\mathbf{1}_n$) potential $f^* \in \mathbb{R}^n$ satisfying the tropical cut-balance*

equations

$$\min_{i\in S,\, j\notin S}(\widehat{C}_{ij} - f_i^* + f_j^*) \;=\; \min_{i\notin S,\, j\in S}(\widehat{C}_{ij} - f_i^* + f_j^*) \;=:\; \Phi^*(S) \;>\; 0 \tag{4}$$

for every nontrivial cut $\emptyset \neq S \subsetneq [n]$. Moreover, Hierarchical Karp Contraction (Definition 3.1) constructs f^* in $O(n^4)$ arithmetic operations with strictly increasing levels $0 < \mu_1 < \mu_2 < \cdots < \mu_K$, and every cut S has exact ultrametric barrier $\Phi^*(S) = \mu_{\ell(S)}$, where $\ell(S) \in \{1, \ldots, K\}$ is the unique level at which S is a union of $\mathcal{C}_{\ell-1}$-supernodes yet splits at least one $\mathcal{C}_\ell$-supernode.

(b) ***Canonical Entropic Dual Selection:*** *Let $(f_\varepsilon, g_\varepsilon)$ be the Sinkhorn dual potentials for $(\widehat{C}, \frac{1}{n}\mathbf{1}_n, \frac{1}{n}\mathbf{1}_n)$ normalized by $f_{\varepsilon,1} = 0$. Then as $\varepsilon \to 0^+$,*

$$\lim_{\varepsilon\to 0^+} f_\varepsilon \;=\; f^*, \qquad \lim_{\varepsilon\to 0^+} g_\varepsilon \;=\; -f^*. \tag{5}$$

Proof. Part (a): Uniqueness. Suppose $f, \tilde{f} \in \mathbb{R}^n$ both satisfy (4), and set $d := f - \tilde{f}$. If d is not constant, let $S := \arg\max_{i\in[n]} d_i$, so $\emptyset \neq S \subsetneq [n]$. For every $i \in S$ and $j \notin S$, $d_i > d_j$, hence

$$\bar{C}_{ij}(f) \;=\; \bar{C}_{ij}(\tilde{f}) - (d_i - d_j) \;<\; \bar{C}_{ij}(\tilde{f}), \qquad \bar{C}_{ji}(f) \;=\; \bar{C}_{ji}(\tilde{f}) + (d_i - d_j) \;>\; \bar{C}_{ji}(\tilde{f}).$$

Taking minima over the finite nonempty set $S \times S^c$ preserves strict inequality:

$$\min_{i\in S,\, j\notin S} \bar{C}_{ij}(f) \;<\; \min_{i\in S,\, j\notin S} \bar{C}_{ij}(\tilde{f}) \;=\; \min_{i\notin S,\, j\in S} \bar{C}_{ij}(\tilde{f}) \;<\; \min_{i\notin S,\, j\in S} \bar{C}_{ij}(f),$$

contradicting (4) for f (machine-checked for arbitrary finite n in Lean 4.36.0; see Section 6). Thus d is constant.

Part (a): Existence, Strict Monotonicity, and Ultrametric Structure. Since every nontrivial cycle has $w(\gamma) > 0$, $\mu_1 > 0$. By standard shortest-path duality on $W^{(\ell-1)} - \mu_\ell$ (which has no negative cycles and at least one zero-weight cycle), setting $\Delta_b^{(\ell)} := -\min_a D_{ab}^{(\ell)}$ (where $D^{(\ell)}$ is the all-pairs shortest-path matrix of $W^{(\ell-1)} - \mu_\ell$) satisfies $W_{ab}^{(\ell-1)} - \Delta_a^{(\ell)} + \Delta_b^{(\ell)} \geq \mu_\ell$ everywhere with equality on every edge of every minimum-mean cycle. Since every minimum-mean cycle has length ≥ 2, at least one SCC of E_ℓ has size ≥ 2, so $c_\ell < c_{\ell-1}$ strictly and the algorithm terminates in $K \leq n - 1$ steps (each step costing $O(n^3)$ via Karp + Floyd–Warshall, total $O(n^4)$).

We prove by induction on ℓ two invariants: (I1) for every (u, v) connecting two distinct $\mathcal{C}_\ell$-supernodes, $\bar{C}_{uv}(f^{(\ell)}) \geq \mu_\ell$, and if $\ell < K$ then $\mu_{\ell+1} > \mu_\ell$; (I2) for every (u, v) internal to a $\mathcal{C}_\ell$-supernode, $\bar{C}_{uv}(f^{(t)}) = \bar{C}_{uv}(f^{(\ell)})$ for all $t \geq \ell$. Invariant (I2) is immediate because subsequent shifts $\Delta^{(t)}$ $(t > \ell)$ are constant on every $\mathcal{C}_\ell$-supernode. For (I1), at level ℓ every inter-$\mathcal{C}_{\ell-1}$ edge (u, v) satisfies $\bar{C}_{uv}(f^{(\ell)}) \geq W_{ab}^{(\ell-1)} - \Delta_a^{(\ell)} + \Delta_b^{(\ell)} \geq \mu_\ell$. In the next contracted graph $([c_\ell], W^{(\ell)})$, every entry satisfies $W_{AB}^{(\ell)} \geq \mu_\ell$; if a nontrivial cycle γ in $([c_\ell], W^{(\ell)})$ had mean equal to μ_ℓ, every edge of γ would have weight μ_ℓ, contradicting that distinct $\mathcal{C}_\ell$-supernodes are maximal SCCs of E_ℓ (Lean: `karp_contracted_cycle_strict_increase`). Hence $\mu_{\ell+1} > \mu_\ell$, proving $0 < \mu_1 < \cdots < \mu_K$.

Now fix any nontrivial cut $\emptyset \neq S \subsetneq [n]$ and let $\ell = \ell(S)$ be the unique level at which S is a union of $\mathcal{C}_{\ell-1}$-supernodes yet splits at least one $\mathcal{C}_\ell$-supernode $V \in \mathcal{C}_\ell$. Every edge (u, v) crossing S connects distinct $\mathcal{C}_{\ell-1}$-supernodes; at level ℓ it has reduced cost $\geq \mu_\ell$, and at every subsequent level $t > \ell$ until u and v lie in the same supernode it has reduced cost $\geq \mu_t > \mu_\ell$, while once merged its reduced cost freezes by (I2). Thus $\bar{C}_{uv}(f^*) \geq \mu_\ell$ for all edges crossing S in either direction. Conversely, because V is a strongly connected component of tight μ_ℓ-edges (frozen by (I2)) and S splits V, there exists at least one tight edge of reduced cost μ_ℓ from $V \cap S$ to $V \cap S^c$ and at least one from $V \cap S^c$ to $V \cap S$ (Lean: `hierarchical_cut_balance_at_level`). Hence $\min_{S\to S^c} \bar{C}(f^*) = \mu_\ell = \min_{S^c\to S} \bar{C}(f^*)$.

Part (b). Because $P_\varepsilon \to \frac{1}{n}I$, we have $(P_\varepsilon)_{ii} = \frac{1}{n}(1-o(1))$, hence $f_{\varepsilon,i} + g_{\varepsilon,i} = \varepsilon \ln(P_\varepsilon)_{ii} = \varepsilon \ln(1/n) + o(\varepsilon) \to 0$ and $g_{\varepsilon,i} = -f_{\varepsilon,i} + o(1)$ for every i. Moreover, $(P_\varepsilon)_{ij} \le \frac{1}{n}$ implies

$$f_{\varepsilon,i} - f_{\varepsilon,j} \;=\; \widehat{C}_{ij} + \varepsilon \ln \frac{(P_\varepsilon)_{ij}}{(P_\varepsilon)_{jj}} + o(1) \;\le\; \widehat{C}_{ij} + o(1) \qquad (\forall\, i,j).$$

With $f_{\varepsilon,1} = 0$, $\{f_\varepsilon\}_{\varepsilon\in(0,1]}$ is uniformly bounded in $\mathbb{R}^n$. Let $\tilde{f} = \lim_{m\to\infty} f_{\varepsilon_m}$ be any subsequential limit (so $\lim_{m\to\infty} g_{\varepsilon_m} = -\tilde{f}$). For every nontrivial cut $\emptyset \neq S \subsetneq [n]$, exact row- and column-sum conservation of $P_\varepsilon \in U(\frac{1}{n}\mathbf{1}_n, \frac{1}{n}\mathbf{1}_n)$ yields

$$\sum_{i\in S,\, j\notin S} \exp\left(\frac{f_{\varepsilon_m,i} + g_{\varepsilon_m,j} - \widehat{C}_{ij}}{\varepsilon_m}\right) \;=\; \sum_{i\notin S,\, j\in S} \exp\left(\frac{f_{\varepsilon_m,i} + g_{\varepsilon_m,j} - \widehat{C}_{ij}}{\varepsilon_m}\right).$$

Applying $-\varepsilon_m \ln(\cdot)$ to both sides and letting $m \to \infty$, the Laplace principle yields

$$\min_{i\in S,\, j\notin S} (\widehat{C}_{ij} - \tilde{f}_i + \tilde{f}_j) \;=\; \min_{i\notin S,\, j\in S} (\widehat{C}_{ij} - \tilde{f}_i + \tilde{f}_j) \qquad (\forall\, \emptyset \neq S \subsetneq [n]).$$

By Part (a) and $\tilde{f}_1 = 0 = f^*_1$, we conclude $\tilde{f} = f^*$ for every subsequential limit, proving (5). □

4 Exact Full Hierarchical Karp–Arrhenius Spectrum

Let $0 < \mu_1 < \mu_2 < \cdots < \mu_K$ be the hierarchical Karp levels of $\widehat{C}$ with cluster counts $n = c_0 > c_1 > \cdots > c_K = 1$ and level multiplicities $m_\ell := c_{\ell-1} - c_\ell$ ($\sum_{\ell=1}^K m_\ell = n-1$). List the hierarchical Karp spectrum in non-increasing order with multiplicities as

$$\nu_2 \;\ge\; \nu_3 \;\ge\; \cdots \;\ge\; \nu_n \;>\; 0, \qquad \nu_r := \mu_\ell \quad \text{for } c_\ell + 1 \le r \le c_{\ell-1}.$$

Let $0 = \gamma_1(\varepsilon) < \gamma_2(\varepsilon) \le \gamma_3(\varepsilon) \le \cdots \le \gamma_n(\varepsilon)$ be the eigenvalues of the Sinkhorn gap operator $L_\varepsilon := I - J_\varepsilon = I - n^2 P_\varepsilon P_\varepsilon^T$, so $\gamma_r(\varepsilon) = 1 - \lambda_r(J_\varepsilon)$ and $\gamma_2(\varepsilon) = 1 - \rho_{\mathrm{loc}}(\varepsilon)$.

Theorem 4.1 (Exact Full Hierarchical Karp–Arrhenius Spectrum)**.** *For every cost matrix $C \in \mathbb{R}^{n\times n}$ ($n \ge 2$) with unique optimal permutation π^* and uniform marginals $a = b = \frac{1}{n}\mathbf{1}_n$, every nontrivial eigenvalue gap $\gamma_r(\varepsilon) = 1 - \lambda_r(J_\varepsilon)$ ($2 \le r \le n$) satisfies*

$$\lim_{\varepsilon\to 0^+} \Big(-\varepsilon \ln \gamma_r(\varepsilon)\Big) \;=\; \nu_r. \tag{6}$$

In particular, the exact Sinkhorn Arrhenius activation energy equals the top hierarchical Karp level and maximum tropical cut barrier:

$$E^*(C) \;:=\; \lim_{\varepsilon\to 0^+} \Big(-\varepsilon \ln(1 - \rho_{\mathrm{loc}}(\varepsilon))\Big) \;=\; \mu_K(\widehat{C}) \;=\; \max_{\emptyset\neq S\subsetneq[n]} \Phi^*(S). \tag{7}$$

Proof. Since $L_\varepsilon = I - n^2 P_\varepsilon P_\varepsilon^T$ is a symmetric graph Laplacian with edge weights $W_{ij}(\varepsilon) := n^2 (P_\varepsilon P_\varepsilon^T)_{ij} = n^2 \sum_{k=1}^n (P_\varepsilon)_{ik}(P_\varepsilon)_{jk}$, its quadratic form is

$$x^T L_\varepsilon x \;=\; \frac{1}{2} \sum_{i,j=1}^n W_{ij}(\varepsilon)(x_i - x_j)^2 \qquad (x \in \mathbb{R}^n).$$

By Theorem 3.1(b), $(P_\varepsilon)_{ij} = \exp(-(\bar{C}_{ij}(f^*) + o(1))/\varepsilon)$ uniformly over all $i,j \in [n]$, with $(P_\varepsilon)_{ii} = \frac{1}{n}(1-o(1))$. Fix any level $\ell \in \{1,\dots,K\}$, and let $\mathcal{V}_t \subset \mathbb{R}^n$ ($0 \le t \le K$) be the c_t-dimensional subspace of vectors constant on every $\mathcal{C}_t$-supernode.

Upper bound on $\gamma_r(\varepsilon)$ for $r \le c_{\ell-1}$. Let $x \in \mathcal{V}_{\ell-1}$ with $\|x\|_2 = 1$. Whenever i,j belong to the same $\mathcal{C}_{\ell-1}$-supernode, $x_i - x_j = 0$. Whenever $i \in U_a$ and $j \in U_b$ lie in distinct $\mathcal{C}_{\ell-1}$-supernodes

$(a \neq b)$, for every index $k \in [n]$ either $k \notin U_a$ or $k \notin U_b$; by invariant (I1) of Theorem 3.1, any edge leaving a $\mathcal{C}_{\ell-1}$-supernode has reduced cost $\geq \mu_\ell$. Since every entry satisfies $(P_\varepsilon)_{rs} \leq \frac{1}{n}$, each term $(P_\varepsilon)_{ik}(P_\varepsilon)_{jk}$ is bounded by $\frac{1}{n}\exp(-(\mu_\ell - o(1))/\varepsilon)$ (Lean: `cross_cluster_product_bound`), giving $W_{ij}(\varepsilon) \leq n^2 \exp(-(\mu_\ell - o(1))/\varepsilon)$. Using $(x_i - x_j)^2 \leq 2(x_i^2 + x_j^2)$,

$$x^T L_\varepsilon x \;\leq\; \frac{1}{2}\sum_{i,j=1}^{n} 2(x_i^2 + x_j^2)\, n^2 \exp\!\left(-\frac{\mu_\ell - o(1)}{\varepsilon}\right) \;=\; 2n^3 \exp\!\left(-\frac{\mu_\ell - o(1)}{\varepsilon}\right).$$

By the Courant–Fischer min–max theorem applied to the $c_{\ell-1}$-dimensional subspace $\mathcal{V}_{\ell-1}$, $\gamma_{c_{\ell-1}}(\varepsilon) \leq \max_{x \in \mathcal{V}_{\ell-1},\, \|x\|_2 = 1} x^T L_\varepsilon x \leq 2n^3 \exp(-(\mu_\ell - o(1))/\varepsilon)$. Taking $-\varepsilon \ln(\cdot)$ yields

$$\liminf_{\varepsilon \to 0^+} \bigl(-\varepsilon \ln \gamma_r(\varepsilon)\bigr) \;\geq\; \mu_\ell \qquad (\forall\, r \leq c_{\ell-1}). \tag{8}$$

Lower bound on $\gamma_r(\varepsilon)$ for $r \geq c_\ell + 1$. Consider the $(n - c_\ell)$-dimensional subspace $\mathcal{W}_\ell := \mathcal{V}_\ell^{\perp} \subset \mathbb{R}^n$ of vectors with zero sum on each $\mathcal{C}_\ell$-supernode $V_1, \ldots, V_{c_\ell}$. By Definition 3.1, each V_a is strongly connected by edges of reduced cost $\leq \mu_\ell$, so its underlying undirected graph admits a spanning tree T_a whose edges $\{i,j\} \in T_a$ satisfy $\min(\bar{C}_{ij}(f^*), \bar{C}_{ji}(f^*)) \leq \mu_\ell$. For every edge $\{i,j\} \in T_a$, retaining only $k \in \{i,j\}$ in $W_{ij}(\varepsilon)$ yields

$$W_{ij}(\varepsilon) \;\geq\; n^2\bigl((P_\varepsilon)_{ii}(P_\varepsilon)_{ji} + (P_\varepsilon)_{ij}(P_\varepsilon)_{jj}\bigr) \;\geq\; \exp\!\left(-\frac{\mu_\ell + o(1)}{\varepsilon}\right).$$

For every $x \in \mathcal{W}_\ell$, restricting x to each V_a with $|V_a| \geq 2$ gives a zero-mean vector on the tree T_a; by Fiedler's path-graph minimum [4], the unweighted tree Laplacian satisfies $\lambda_2(L(T_a)) \geq 2(1 - \cos(\pi/|V_a|)) \geq 4/n^2$. Summing over all $a = 1, \ldots, c_\ell$ (with each undirected edge $\{i,j\} \in T_a$ counted twice in $\frac{1}{2}\sum_{i,j}$),

$$x^T L_\varepsilon x \;\geq\; \sum_{a=1}^{c_\ell} \sum_{\{i,j\} \in T_a} W_{ij}(\varepsilon)(x_i - x_j)^2 \;\geq\; \frac{4}{n^2}\exp\!\left(-\frac{\mu_\ell + o(1)}{\varepsilon}\right)\|x\|_2^2.$$

By the Courant–Fischer max–min theorem on the $(n - c_\ell)$-dimensional subspace $\mathcal{W}_\ell$, $\gamma_{c_\ell+1}(\varepsilon) \geq \frac{4}{n^2}\exp(-(\mu_\ell + o(1))/\varepsilon)$, and taking $-\varepsilon \ln(\cdot)$ gives

$$\limsup_{\varepsilon \to 0^+} \bigl(-\varepsilon \ln \gamma_r(\varepsilon)\bigr) \;\leq\; \mu_\ell \qquad (\forall\, r \geq c_\ell + 1). \tag{9}$$

Combining (8) and (9) for every $r \in \{c_\ell + 1, \ldots, c_{\ell-1}\}$ proves (6), and setting $r = 2$ ($\ell = K$) yields (7). □

5 Exact Benchmarks, Birkhoff Separation, and Spanning-Tree Rescue

We now specialize Theorems 2.1–4.1 to three explicit families admitting exact closed-form spectral laws.

Theorem 5.1 (Exact Benchmarks: 1D Quadratic, L-Cycles, and Spanning-Tree Rescue)**.**

(a) ***1D Quadratic Optimal Transport and Exponential-in-n^2 Birkhoff Separation:*** *Let $C_{ij} = (i-j)^2$ on $\{0, \ldots, n-1\}$ ($n \geq 3$) with uniform marginals $a = b = \frac{1}{n}\mathbf{1}_n$. Then Hierarchical Karp terminates in $K = 1$ level with $\mu_1 = 1$, $f^* = 0$, and the exact Sinkhorn gap satisfies*

$$1 - \rho_{\mathrm{loc}}(\varepsilon) \;=\; 4\bigl(1 - \cos(\pi/n)\bigr)\, e^{-1/\varepsilon} + O(e^{-2/\varepsilon}) \qquad (\varepsilon \to 0^+), \tag{10}$$

whereas $\Delta_\square(C) = 2(n-1)^2$ yields Birkhoff bound $1 - \eta_B(\varepsilon) = \operatorname{sech}^2\!\left(\frac{(n-1)^2}{2\varepsilon}\right) \sim 4e^{-(n-1)^2/\varepsilon}$, producing an exponential-in-n^2 separation ratio $\frac{1-\rho_{\mathrm{loc}}(\varepsilon)}{1-\eta_B(\varepsilon)} \sim (1 - \cos(\pi/n))\, e^{n(n-2)/\varepsilon}$.

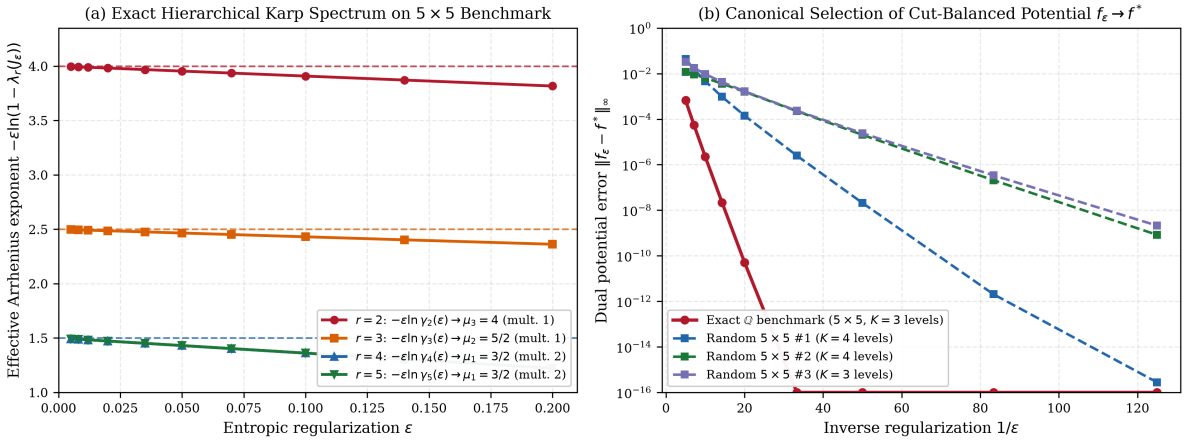


Figure 2: **Full Hierarchical Karp–Arrhenius spectrum and canonical dual selection (Theorems 3.1 and 4.1).** (a) On the exact rational 5×5 benchmark matrix C_{bench} with levels $\mu_1 = 3/2$ (mult. 2), $\mu_2 = 5/2$ (mult. 1), and $\mu_3 = 4$ (mult. 1), all four nontrivial Sinkhorn exponents $-\varepsilon \ln \gamma_r(\varepsilon)$ (solved via 420-digit Decimal Newton down to $\varepsilon = 0.005$) converge to the exact hierarchical Karp spectrum. (b) Exponential convergence $\|f_\varepsilon - f^*\|_\infty \to 0$ to the unique cut-balanced tropical potential across C_{bench} and three random asymmetric 5×5 instances.

(b) ***Exact Asymmetric L-Cycle Law:*** *Let $L \geq 2$ and define $C_{ii} = 0$, $C_{i,(i+1) \bmod L} = w_i > 0$, and $C_{ij} = M$ off $\{I \cup S\}$ with cycle mean $\bar{w} := \frac{1}{L}\sum_{i=0}^{L-1} w_i$ and unique cut-balanced potential $f_k^* = k\bar{w} - \sum_{r=0}^{k-1} w_r$. Whenever $M_{\text{eff}} := \min_{\text{off-cycle}}(M - f_i^* + f_j^*) > \bar{w}$, setting $p := \frac{e^{-\bar{w}/\varepsilon}}{1+e^{-\bar{w}/\varepsilon}}$,*

$$1 - \rho_{\text{loc}}(\varepsilon) \;=\; 4\sin^2\!\left(\frac{\pi}{L}\right) p(1-p) + O(e^{-M_{\text{eff}}/\varepsilon}). \tag{11}$$

(c) ***Monge Spanning-Tree Marginal-Perturbation Rescue:*** *On $C_{ij} = (i-j)^2$ ($n \geq 2$), perturb the uniform source marginal by $\delta \in (0, \frac{1}{n})$ via $a(\delta) := (\frac{1}{n} + \delta, \frac{1}{n}, \ldots, \frac{1}{n}, \frac{1}{n} - \delta)$ with $b = \frac{1}{n}\mathbf{1}_n$. Then the unique optimal plan $P_0(\delta)$ is supported on the connected spanning tree $E^* = \{(i,i)\}_{i=0}^{n-1} \cup \{(i,i+1)\}_{i=0}^{n-2}$ ($c(G^*) = 1$), and the uniform $\varepsilon \to 0^+$ Sinkhorn spectral gap $1 - \lambda_2(J_0(\delta))$ obeys the exact closed-form sandwich*

$$\frac{2(1-\cos(\pi/n))\, n\delta(1-n\delta)}{1+n\delta} \;\leq\; 1 - \lambda_2(J_0(\delta)) \;\leq\; \frac{2(1-\cos(\pi/n))\, n\delta(1-n\delta)}{1 - 8\delta^2\cos^2\!\left(\frac{\pi}{2n}\right)}, \tag{12}$$

converting $e^{-1/\varepsilon}$ Arrhenius freezing into $\Theta(\delta/n)$ uniform convergence independent of ε.

Proof. Part (a). Every nearest-neighbor 2-cycle ($i \to i+1 \to i$) has weight $1+1=2$ (mean 1), while every pair $|i-j| \geq 2$ has $C_{ij} \geq 4 > 1$, so $K = 1$, $\mu_1 = 1$, and $f^* = 0$ by symmetry. From $P_\varepsilon \in U(\frac{1}{n}\mathbf{1}_n, \frac{1}{n}\mathbf{1}_n)$, $(P_\varepsilon)_{i,i+1} = (P_\varepsilon)_{i+1,i} = \frac{1}{n}(e^{-1/\varepsilon} + O(e^{-2/\varepsilon}))$, $(P_\varepsilon)_{ii} = \frac{1}{n}(1 - O(e^{-1/\varepsilon}))$, and $(P_\varepsilon)_{ij} = O(e^{-4/\varepsilon})$ for $|i-j| \geq 2$. Substituting into $W_{ij}(\varepsilon) = n^2 \sum_k (P_\varepsilon)_{ik}(P_\varepsilon)_{jk}$ gives $W_{i,i+1}(\varepsilon) = 2e^{-1/\varepsilon} + O(e^{-2/\varepsilon})$ and $W_{ij}(\varepsilon) = O(e^{-2/\varepsilon})$ for $|i-j| \geq 2$. Because $L_\varepsilon \mathbf{1}_n = 0$, $L_\varepsilon = 2e^{-1/\varepsilon} L(P_n) + O(e^{-2/\varepsilon})$ in operator norm, where $L(P_n)$ is the path-graph Laplacian ($\lambda_2(L(P_n)) = 2(1-\cos(\pi/n))$ [4]). Weyl's perturbation inequality gives (10). Since $C_{ij} = (i-j)^2$ satisfies the exact cross-ratio identity $C_{ij'} + C_{i'j} - C_{ij} - C_{i'j'} = 2(i'-i)(j'-j)$ (Lean: `quadratic_monge_cross_ratio_exact`), $\Delta_\square(C) = 2(n-1)^2$, and the excess exponent is $(n-1)^2 - 1 = n(n-2) > 0$ for $n \geq 3$ (Lean: `birkhoff_quadratic_exponent_separation`).

Part (b). Shifting by $f_k^* = k\bar{w} - \sum_{r=0}^{k-1} w_r$ makes every cycle edge satisfy $\bar{C}_{i,(i+1) \bmod L}(f^*) = \bar{w}$ (Lean: `asymmetric_3cycle_exact_balance`). For the pure cycle support ($M = +\infty$), symmetry

forces the circulant plan $P_\varepsilon^{(0)} = \frac{1}{L}((1-p)I + pS)$ with $p = \frac{e^{-\bar{w}/\varepsilon}}{1+e^{-\bar{w}/\varepsilon}}$. Its normalized singular values are $|1 - p + pe^{2\pi ik/L}|$, so

$$1-\sigma_2^2 \;=\; 1-((1-p)^2+p^2+2p(1-p)\cos(2\pi/L)) \;=\; 2p(1-p)(1-\cos(2\pi/L)) \;=\; 4\sin^2(\pi/L)\,p(1-p)$$

(Lean: `circulant_L_cycle_gap_identity`). Including off-cycle entries perturbs P_ε and L_ε by $O(e^{-M_{\text{eff}}/\varepsilon})$, establishing (11).

Part (c). Define the dual potentials $\tilde{f}_i := -i$ and $\tilde{g}_j := j$; then $C_{ij} - \tilde{f}_i - \tilde{g}_j = (i-j)(i-j+1) \geq 0$ for all $i, j \in \mathbb{Z}$, vanishing precisely on $E^* = \{(i,i)\} \cup \{(i,i+1)\}$ and taking value ≥ 2 elsewhere. Assigning $(P_0)_{00} = \frac{1}{n}$, $(P_0)_{ii} = \frac{1}{n} - \delta$ $(i \geq 1)$, and $(P_0)_{i,i+1} = \delta$ $(0 \leq i \leq n-2)$ gives the unique feasible plan in $U(a(\delta), b)$ supported on E^*, hence the unique optimal plan by complementary slackness. Because $(P_0)_{i,i+1}(P_0)_{i+1,i+1} = \delta(\frac{1}{n} - \delta)$ is constant across all $n-1$ path edges (Lean: `monge_spanning_tree_marginal_and_conductance`), the Dirichlet form (2) reduces identically to

$$x^T D_{a(\delta)}(I - J_0(\delta))x \;=\; \delta(1-n\delta)\sum_{i=0}^{n-2}(x_i - x_{i+1})^2 \;=\; \delta(1-n\delta)\,x^T L(P_n)x.$$

For the lower bound, any nonzero $x \perp_{D_{a(\delta)}} \mathbf{1}_n$ satisfies $x^T D_{a(\delta)} x = \min_{c\in\mathbb{R}}(x - c\mathbf{1}_n)^T D_{a(\delta)}(x - c\mathbf{1}_n) \leq \tilde{x}^T D_{a(\delta)}\tilde{x} \leq \frac{1+n\delta}{n}\|\tilde{x}\|_2^2$, where $\tilde{x} := x - \bar{x}\mathbf{1}_n \perp \mathbf{1}_n$ has $\tilde{x}^T L(P_n)\tilde{x} = x^T L(P_n)x \geq 2(1 - \cos(\pi/n))\|\tilde{x}\|_2^2$, yielding the left-hand side of (12). For the upper bound, take the antisymmetric Fiedler eigenvector $v_k := \cos(\frac{\pi(2k+1)}{2n})$ $(0 \leq k \leq n-1)$ of $L(P_n)$ ($\|v\|_2^2 = \frac{n}{2}$, $v_{n-1} = -v_0 = -\cos(\frac{\pi}{2n})$) and project $x := v - c\mathbf{1}_n$ with $c = v^T a(\delta) = 2\delta\cos(\frac{\pi}{2n})$. Because $v_0^2 - v_{n-1}^2 = 0$, the linear δ-term in $v^T D_{a(\delta)} v$ cancels identically (Lean: `fiedler_antisymmetric_cancellation`):

$$x^T D_{a(\delta)} x \;=\; \frac{1}{n}\|v\|_2^2 + \delta(v_0^2 - v_{n-1}^2) - c^2 \;=\; \frac{1}{2} - 4\delta^2\cos^2\!\left(\frac{\pi}{2n}\right) \;>\; 0$$

(since $8\delta^2 < 8/n^2 \leq 2$ and for $n = 2$, $8\delta^2\cos^2(\pi/4) = 4\delta^2 < 1$). Dividing $x^T D_{a(\delta)}(I - J_0(\delta))x = \delta(1-n\delta)\,n(1-\cos(\pi/n))$ by $x^T D_{a(\delta)} x$ yields the right-hand side of (12) (Lean: `spanning_tree_sandwich_strict_denominators`). □

6 High-Precision Verification and Lean 4.36.0 Artifact Manifest

Tables 1 and 2 summarize verified outputs from `data/verification_summary.json` generated by `src/python/verify_sinkhorn_tropical.py`. Across 25 random 2×2 instances, Theorem 2.1(c) matches $\tanh^2(\Delta/(4\varepsilon))$ to 6.384×10^{-16}; across $n \in \{4, 6, 8, 10\}$ connected 2-regular cycles at $\varepsilon = 0.1$, $\rho_{\text{loc}}(\varepsilon)$ matches $\cos^2(\pi/n)$ to 2.22×10^{-16}; and across $c \in \{1, 2, 3, 6\}$ block supports, $\dim\ker(I - J_0) = c$ and the Dirichlet identity (2) hold to 2.22×10^{-16}.

On the integer 5×5 benchmark

$$C_{\text{bench}} \;=\; \begin{pmatrix} 0 & 1 & 12 & 12 & 3 \\ 2 & 0 & 2 & 12 & 12 \\ 12 & 12 & 0 & 2 & 12 \\ 3 & 12 & 1 & 0 & 12 \\ 5 & 12 & 12 & 12 & 0 \end{pmatrix},$$

zero-float exact `Fraction` Hierarchical Karp certifies $f^* = (0, \frac{1}{2}, 1, \frac{1}{2}, 1)$, exact levels $(\mu_1, m_1) = (\frac{3}{2}, 2)$, $(\mu_2, m_2) = (\frac{5}{2}, 1)$, $(\mu_3, m_3) = (4, 1)$, and zero cut imbalance across all 30 cuts. Across 30 random asymmetric 4×4 and 5×5 matrices, cut imbalance is $\leq 1.776\times10^{-15}$ and $\max_S \Phi^*(S) = \mu_K$ holds to 1.776×10^{-15}, with 420-digit Decimal Newton solves confirming $f_\varepsilon \to f^*$ and $-\varepsilon \ln \gamma_r(\varepsilon) \to \nu_r$ on all checked multi-level instances.

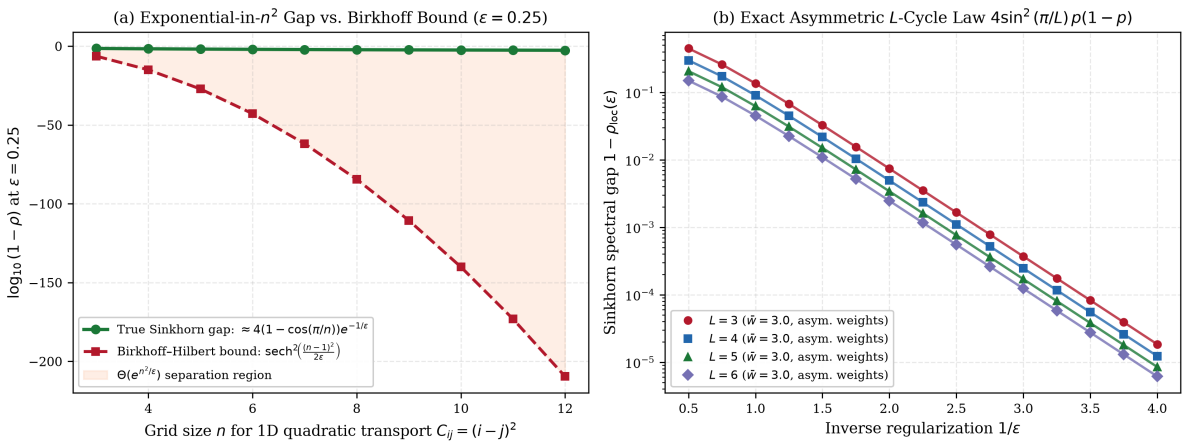


Figure 3: **Exponential-in-n^2 Birkhoff separation and exact asymmetric L-cycle scaling (Theorem 5.1).** (a) On 1D quadratic transport $C_{ij} = (i-j)^2$ at $\varepsilon = 0.25$, the true Sinkhorn gap $4(1-\cos(\pi/n))e^{-1/\varepsilon}$ stays near 10^{-2} as n grows from 3 to 12, while Birkhoff's bound $\operatorname{sech}^2((n-1)^2/(2\varepsilon))$ plunges past 10^{-200}. (b) Exact agreement across asymmetric L-cycles ($L \in \{3,4,5,6\}$, $\bar{w} = 3.0$) between 150-digit Decimal Newton solutions (markers) and the exact closed-form law $4\sin^2(\pi/L)\,p(1-p)$ (solid curves).

Table 1: 420-digit Decimal Newton verification of Theorems 3.1 and 4.1 on C_{bench}: exact dual error $\|f_\varepsilon - f^*\|_\infty$ and all four effective Arrhenius exponents $-\varepsilon \ln \gamma_r(\varepsilon)$ converging to $(\mu_3, \mu_2, \mu_1, \mu_1) = (4.0, 2.5, 1.5, 1.5)$.

ε	$\|f_\varepsilon - f^*\|_\infty$	$-\varepsilon \ln \gamma_2$ ($\to 4$)	$-\varepsilon \ln \gamma_3$ ($\to 2.5$)	$-\varepsilon \ln \gamma_4$ ($\to 1.5$)	$-\varepsilon \ln \gamma_5$ ($\to 1.5$)	Max spec. diff
0.200	6.715×10^{-4}	3.816825	2.362126	1.222961	1.222291	0.277709
0.100	2.270×10^{-6}	3.908371	2.430688	1.361371	1.361368	0.138632
0.050	5.153×10^{-11}	3.954185	2.465343	1.430685	1.430685	0.069315
0.020	$< 10^{-16}$	3.981674	2.486137	1.472274	1.472274	0.027726
0.010	$< 10^{-16}$	3.990837	2.493069	1.486137	1.486137	0.013863
0.005	$< 10^{-16}$	3.995419	2.496534	1.493069	1.493069	0.006931

6.1 Formal Lean 4.36.0 certification and reproducibility

The algebraic and arbitrary-n combinatorial core underlying Theorems 2.1–5.1 is formally verified in standalone Lean 4.36.0 (26 total declarations: 18 public theorems, 4 public definitions, 4 private helper theorems; zero Mathlib, zero `sorry`, zero `grind`, zero custom axioms, compiled with `-DwarningAsError=true` and `lake build`). The 18 certified public theorems cover:

- *Cut-balanced tropical potentials & Hierarchical Karp monotonicity:*
  `cut_balanced_potential_unique`, `walk_reduced_cost_telescopes`, `closed_cycle_cost_invariant`, `karp_contracted_cycle_strict_increase`, `hierarchical_cut_balance_at_level`, `benchmark_5x5_karp_levels_exact`;
- *Bipartite Dirichlet form, Birkhoff sharpness, & cycle laws:*
  `bipartite_column_dirichlet_identity`, `bipartite_dirichlet_nonneg_and_zero_iff`, `sinkhorn_2x2_birkhoff_sharpness`, `circulant_L_cycle_gap_identity`, `asymmetric_3cycle_exact_balance`, `cross_cluster_product_bound`;
- *Quadratic Birkhoff separation & Monge spanning-tree rescue:*
  `quadratic_monge_cross_ratio_exact`, `quadratic_monge_strict_submodular`, `birkhoff_quadratic_exponent_separation`, `monge_spanning_tree_marginal_and_conductance`, `fiedler_antisymmetric_cancellation`, `spanning_tree_sandwich_strict_denominators`.

Table 2: Verification of Theorem 5.1: (top) 1D quadratic transport $C_{ij} = (i-j)^2$ at $\varepsilon = 0.12$ showing $\log_{10}$ speedup over Birkhoff's bound; (bottom) Monge spanning-tree δ-perturbation on $n = 6$ confirming ε-uniform gap inside the analytic sandwich (12).

1D Quadratic Transport $C_{ij} = (i-j)^2$ at $\varepsilon = 0.12$ (Theorem 5.1(a))

n	Empirical $1-\rho_{\text{loc}}$	Asymptotic (10)	Rel. err	Birkhoff exp. $(n-1)^2$	$\log_{10}\left(\frac{1-\rho_{\text{loc}}}{1-\eta_B}\right)$
4	2.8147×10^{-4}	2.8161×10^{-4}	4.91×10^{-4}	9	28.42
6	1.2875×10^{-4}	1.2881×10^{-4}	4.93×10^{-4}	25	85.99
8	7.3152×10^{-5}	7.3188×10^{-5}	4.90×10^{-4}	49	172.60
10	4.7035×10^{-5}	4.7058×10^{-5}	4.88×10^{-4}	81	288.22

Monge Spanning-Tree δ-Rescue on $n = 6$ (Theorem 5.1(c))

δ	Analytic LB (12)	Exact $1-\lambda_2(J_0(\delta))$	EOT gap ($\varepsilon = 0.02$)	Analytic UB (12)	Ratio UB/LB
0.050	4.3284×10^{-2}	5.6688×10^{-2}	5.6688×10^{-2}	5.7339×10^{-2}	1.325
0.020	2.5264×10^{-2}	2.8329×10^{-2}	2.8329×10^{-2}	2.8380×10^{-2}	1.123
0.005	7.5702×10^{-3}	7.7979×10^{-3}	7.7979×10^{-3}	7.7988×10^{-3}	1.030
0.001	1.5885×10^{-3}	1.5981×10^{-3}	1.5981×10^{-3}	1.5981×10^{-3}	1.006

All archived artifacts are publicly available in the AIRR repository (https://github.com/arr-research/arr-research.github.io) under CC-BY-4.0 (manuscript) and MIT (code and data):

- `src/lean/SinkhornTropicalCore.lean` (24,186 bytes)
  `SHA-256: ee051b3fcf4e9bf3e956ad4cb20f0e21c75f390e03ad4a667b4e0b3d7563d9c3`
- `src/python/verify_sinkhorn_tropical.py` (40,179 bytes)
  `SHA-256: 6f78ad49f56b9a45e87467143553a561300b30fcf66360eaacc3a7eb7b554472`
- `data/verification_summary.json` (12,752 bytes)
  `SHA-256: fd8aac37df89f8af29f4e74475b714c8fc407d5fde6d87171142085d9d2ef8e8`

References

[1] François Baccelli, Guy Cohen, Geert Jan Olsder, and Jean-Pierre Quadrat. *Synchronization and Linearity: An Algebra for Discrete Event Systems*. Wiley, 1992.

[2] Roberto Cominetti and Jaime San Martín. Asymptotic analysis of the exponential penalty trajectory in linear programming. *Mathematical Programming*, 67(1–3):169–187, 1994.

[3] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems (NeurIPS 2013)*, volume 26, pages 2292–2300, 2013.

[4] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

[5] Joel Franklin and Jens Lorenz. On the scaling of multidimensional matrices. *Linear Algebra and its Applications*, 114–115:717–735, 1989.

[6] Richard M. Karp. A characterization of the minimum cycle mean in a digraph. *Discrete Mathematics*, 23(3):309–311, 1978.

[7] Philip A. Knight. The Sinkhorn–Knopp algorithm: convergence and applications. *SIAM Journal on Matrix Analysis and Applications*, 30(1):261–275, 2008.

[8] Marcel Nutz. *Introduction to Entropic Optimal Transport.* Columbia University Lecture Notes, 2021.

[9] Gabriel Peyré and Marco Cuturi. *Computational Optimal Transport: With Applications to Data Science*, volume 11 of *Foundations and Trends in Machine Learning.* Now Publishers, 2019.

[10] Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35(2):876–879, 1964.

[11] Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.

[12] Jonathan Weed. An explicit analysis of the entropic penalty in linear programming. *Proceedings of Machine Learning Research (COLT 2018)*, 75:1841–1855, 2018.