AIRR.SCIENCE · ARR-2026-44PEE4YMY96EBHH v1 [v1] 29 Nov 2026

Bhattacharyya–Volatility Laplacian Decompositions, Harmonic-Risk Duality, Universal Factor-of-2 Le Cam Sandwiches, and Exact Chernoff Tensorization Exponents for χ^2 Strong Data-Processing Inequalities

Lluis Eriksson
Independent Researcher, Östersund, Sweden
editor@airr.science

November 29, 2026

Abstract

Let $K \in \mathbb{R}^{m\times n}_{\geq 0}$ be a finite discrete Markov channel ($m, n \geq 2$) and let $\eta_{\chi^2}(q, K)$ and $\eta_{\chi^2}(K) := \sup_{q\in\Delta_m} \eta_{\chi^2}(q, K)$ denote its prior-dependent and global Pearson χ^2 strong data-processing inequality (SDPI) contraction coefficients. We develop a unified spectral, convex-analytic, and information-geometric framework governing both the full singular spectrum of K at fixed priors and the global multi-input and tensor-power contraction behavior of K.

First, for every full-support prior $q \in \Delta_m^+$ with output marginal $r = qK$ (restricted to $\mathrm{supp}(r)$) and symmetrized forward–backward Gram operator $G(q, K) := D_q^{1/2} K D_r^{-1} K^\top D_q^{1/2}$, we prove an exact positive-semidefinite Laplacian splitting

$$I_m - G(q, K) \;=\; \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) \;+\; \widetilde{\mathcal{L}}_{\mathrm{Var}}(q) \;\succeq\; \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) \;\succeq\; \mathcal{L}_{\mathrm{BC}}(q) \;\succeq\; 0,$$

where $\widetilde{\mathcal{L}}_{\mathrm{BC}}(q)$ carries support-normalized Bhattacharyya edge affinities $\widetilde{B}_{ik}(q) := \mathrm{BC}(K_i, K_k)^2/r(S_{ik})$ ($S_{ik} := \mathrm{supp}(K_i) \cap \mathrm{supp}(K_k)$) and $\widetilde{\mathcal{L}}_{\mathrm{Var}}(q)$ carries conditional posterior-volatility edge weights $V_{ik}(q) := r(S_{ik})\,\mathrm{Var}_{Y\sim r(\cdot|S_{ik})}(\sqrt{K_{iY}K_{kY}}/r_Y) \geq 0$. This yields all-eigenvalue Weyl contraction bounds $\sigma_r^2(q, K) \leq 1 - \lambda_r(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q))$, exact zero-volatility rigidity for every single-overlap channel ($|S_{ik}| \leq 1$ for all $i \neq k$, including simple weighted graph-incidence channels and nearest-neighbor lazy cycle walks W_m where Dobrushin's coefficient $\eta_{\mathrm{TV}}(W_m) = 1$ is trivial while $\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)$ recovers the exact gap $\sin^2(\pi/m)$ and improves unnormalized $\mathcal{L}_{\mathrm{BC}}(u)$ by a factor of m), and a support-normalized Fourier–Bhattacharyya inequality on arbitrary finite abelian groups Γ.

Second, for any binary-input channel $K = \binom{p}{q}$, we establish the exact Bayesian harmonic-risk identity $1 - \eta_{\chi^2}(y; p, q) = H_{p,q}(y) := \sum_{j\in S} p_j q_j/((1-y)p_j + yq_j)$, express its derivatives in closed form via posterior-risk covariance $H'_{p,q}(y) = -\mathrm{Cov}_{r_y}(\Pi, \Pi(1-\Pi))/(y^2(1-y)^2)$, classify the minimizer into an exact four-regime (flat BEC plus three non-degenerate phase) trichotomy governed by the support-restricted Pearson discrepancies (Δ_0, Δ_1), and prove that global Bhattacharyya saturation $\eta_{\chi^2}(p, q) = 1 - \mathrm{BC}(p, q)^2$ holds if and only if (p, q) belongs to exactly three canonical classes: trivial/disjoint pairs, generalized Z-channels, or full-support two-ratio channels isomorphic to 2×2 binary asymmetric channels.

Third, combining a self-contained five-line polytope-vertex proof of binary-input reduction $\eta_{\chi^2}(K) = \max_{i<k} \eta_{\chi^2}(K_i, K_k)$ (Cohen–Kemperman–Zbăganu; Ordentlich–Polyanskiy) with harmonic-risk duality, we establish the universal two-sided sandwich

$$\begin{aligned}\max_{i<k} H^2(K_i, K_k) \;&\leq\; \max_{i<k} \mathrm{LC}(K_i, K_k) \;\leq\; \eta_{\chi^2}(K) \\ &\leq\; \min\{\eta_{\mathrm{TV}}(K),\, 1 - \mathrm{BC}_{\min}(K)^2\} \;\leq\; 2 \max_{i<k} H^2(K_i, K_k),\end{aligned}$$

proving that the closed-form symmetric Le Cam divergence $\mathrm{LC}_{\max}(K)$ approximates $\eta_{\chi^2}(K)$ universally within a sharp factor of 2 (attained asymptotically by Z-channels Z_ϵ as $\epsilon \to 0^+$ with ratio $2 - \epsilon$, and equal to 1 on every involution-symmetric binary channel).

Fourth, we resolve the dichotomy on d-fold tensor powers $K^{\otimes d}$: while fixed product priors remain dimension-free ($\eta_{\chi^2}(q^{\otimes d}, K^{\otimes d}) = \eta_{\chi^2}(q, K)$), an elementary Chebyshev-tilted proof establishes the uniform finite-d two-sided sandwich $\frac{1}{2}\exp(-dC_{\max}(K) - \sqrt{2dV_{\max}(K)}) \le 1 - \eta_{\chi^2}(K^{\otimes d}) \le \exp(-dC_{\max}(K))$ governed by the exact pairwise Chernoff information $C_{\max}(K)$, and for $\mathrm{BSC}(p)^{\otimes d}$ we derive the exact parity-resolved hyperbolic-secant lattice prefactor law $1 - \eta_{\chi^2}(\mathrm{BSC}(p)^{\otimes d}) \sim \sqrt{2/(\pi d)}\, S_\rho(\alpha)\, e^{-dC(p)}$ ($d \equiv \rho \pmod 2$). All theorems are verified computationally to machine precision via an open-source reproducibility suite.

Keywords: strong data-processing inequalities, χ^2-contraction coefficient, Bhattacharyya affinity, graph Laplacian, Le Cam divergence, Chernoff information, Markov channels.

1 Introduction

Strong data-processing inequalities (SDPIs) quantify the strict loss of statistical distinguishability induced by a noisy Markov channel $K \in \mathbb{R}_{\ge 0}^{m\times n}$ ($\sum_{j=1}^n K_{ij} = 1$ for all $i \in [m] := \{1, \dots, m\}$). Given a convex function $f : (0,\infty) \to \mathbb{R}$ with $f(1) = 0$ and the associated Csiszár f-divergence $D_f(p \,\|\, q)$, the prior-dependent and global contraction coefficients are

$$\eta_f(q, K) := \sup_{\substack{p \in \Delta_m \setminus \{q\} \\ \mathrm{supp}(p) \subseteq \mathrm{supp}(q)}} \frac{D_f(pK \,\|\, qK)}{D_f(p \,\|\, q)}, \qquad \eta_f(K) := \sup_{q \in \Delta_m} \eta_f(q, K),$$

where $\Delta_m := \{q \in \mathbb{R}_{\ge 0}^m : \sum_{i=1}^m q_i = 1\}$ is the input probability simplex (with the convention $\eta_f(e_i, K) := 0$ for point-mass priors where the feasible set is empty) and $\Delta_m^+ := \{q \in \Delta_m : q_i > 0\ \forall i\}$. Among all f-divergences, the Pearson χ^2-divergence $\chi^2(p \,\|\, q) := \sum_i (p_i - q_i)^2/q_i$ occupies a uniquely central position [1, 4, 9, 8]:

1. It equals the squared Hirschfeld–Gebelein–Rényi (HGR) maximal correlation $\rho_m(X;Y)^2 = \sigma_2(D_q^{1/2} K D_{qK}^{-1/2})^2$ at every prior q (with $D_{qK}^{-1/2}$ restricted to $\mathrm{supp}(qK)$) [2, 6], bridging information theory with spectral graph theory and reversible Markov chains.
2. It universalizes lower and upper bounds across f-divergences: for every twice-differentiable strictly convex f with $f''(1) > 0$, $\eta_{\chi^2}(q, K) \le \eta_f(q, K) \le \eta_{\mathrm{TV}}(K)$, where $\eta_{\mathrm{TV}}(K) := \max_{i<k} \mathrm{TV}(K_i, K_k)$ is Dobrushin's ergodicity coefficient [5, 8].
3. As conjectured by Cohen, Kemperman, and Zbăganu [4] and resolved in full generality by Ordentlich and Polyanskiy [7], the global contraction coefficient $\eta_{\chi^2}(K)$ reduces to binary-input subchannels, whereas the local spectral contraction $\eta_{\chi^2}(q, K)$ at a fixed prior q is genuinely multi-input and sensitive to the global connectivity of the channel's overlap graph.

Despite this mature foundation, four fundamental structural questions have remained fragmented across spectral Markov theory, binary hypothesis testing, and f-divergence geometry:

1. **Exact spectral decomposition beyond Dobrushin.** When a multi-input channel K is sparse or local—for example, a nearest-neighbor channel on a cycle where distant input rows have disjoint supports—Dobrushin's coefficient satisfies $\eta_{\mathrm{TV}}(K) = 1$ and pairwise minimum-overlap bounds vanish identically. Can $I_m - G(q, K)$ be decomposed *exactly* into an explicit Bhattacharyya graph Laplacian plus a positive-semidefinite remainder that captures full-spectrum contraction even when $\eta_{\mathrm{TV}}(K) = 1$?
2. **Exact geometry and saturation rigidity of binary-input SDPIs.** While the closed-form identity $\eta_{\chi^2}(\mathrm{BAC}(a,b)) = (\sqrt{a(1-b)} - \sqrt{b(1-a)})^2 = 1 - \mathrm{BC}^2$ is classical for 2×2 channels $\mathrm{BAC}(a,b) := \left(\begin{smallmatrix} a & 1-a \\ b & 1-b \end{smallmatrix}\right)$ [1, 8], what governs the optimal prior $y^* \in [0,1]$ for arbitrary $2 \times n$ channels, when does y^* collapse to a boundary 0^+ or 1^-, and which channels *exact-saturate* the Bhattacharyya ceiling $1 - \mathrm{BC}(p,q)^2$?

3. **Closed-form universal approximation of $\eta_{\chi^2}(K)$.** Computing $\eta_{\chi^2}(K)$ requires optimizing over priors $q \in \Delta_m$. Does a single elementary, closed-form symmetric divergence approximate $\eta_{\chi^2}(K)$ within a small universal constant across *all* discrete Markov channels, and what is the sharp worst-case constant?

4. **Tensor-power amplification vs. product-prior invariance.** While $\eta_{\chi^2}(q^{\otimes d}, K^{\otimes d}) = \eta_{\chi^2}(q, K)$ is dimension-free for product priors, allowing correlated priors on $K^{\otimes d}$ drives $\eta_{\chi^2}(K^{\otimes d}) \to 1$. What is the exact finite-d exponential rate and sub-exponential prefactor governing $1 - \eta_{\chi^2}(K^{\otimes d})$?

Summary of Main Results

This paper resolves all four questions through four tightly coupled theorems:

- **Theorem 2.2 (Exact Support-Normalized Bhattacharyya–Volatility Laplacian Decomposition):** For every $K \in \mathbb{R}_{\geq 0}^{m \times n}$ and $q \in \Delta_m^+$, $I_m - G(q, K) = \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) + \widetilde{\mathcal{L}}_{\mathrm{Var}}(q) \succeq \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) \succeq \mathcal{L}_{\mathrm{BC}}(q) \succeq 0$. This yields full-spectrum Weyl bounds $\sigma_r^2(q, K) \leq 1 - \lambda_r(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q))$, exact trace-contraction bounds, zero-volatility rigidity for every single-overlap channel ($|S_{ik}| \leq 1$, recovering the exact spectrum $\sin^2(\pi t/m)$ of the lazy cycle W_m with a factor-m gain over $\mathcal{L}_{\mathrm{BC}}(u)$), and a support-normalized Fourier–Bhattacharyya inequality on every finite abelian group Γ.

- **Theorem 3.1 (Binary-Input Harmonic-Risk Duality, Four-Regime Trichotomy, and Saturation Rigidity):** For every binary-input channel $K = \binom{p}{q}$, $1 - \eta_{\chi^2}(y; p, q) = H_{p,q}(y) = \mathbb{E}_{r_y}[\Pi(1-\Pi)]/(y(1-y))$ is convex with $H'_{p,q}(y) = -\operatorname{Cov}_{r_y}(\Pi, \Pi(1-\Pi))/(y^2(1-y)^2)$, admits an exact classification into the flat BEC regime plus Phases I–III via (Δ_0, Δ_1), and saturates $\eta_{\chi^2}(p, q) = 1 - \mathrm{BC}(p, q)^2$ if and only if (p, q) is trivial/disjoint, a generalized Z-channel, or a full-support two-ratio channel isomorphic to $\mathrm{BAC}(a, b)$.

- **Theorem 4.1 (Polytope-Vertex Binary Reduction and Universal Factor-of-2 Le Cam Sandwich):** Pairing a self-contained five-line polytope-vertex proof of $\eta_{\chi^2}(K) = \max_{i<k} \eta_{\chi^2}(K_i, K_k)$ with Theorem 3.1, we prove $H^2_{\max}(K) \leq \mathrm{LC}_{\max}(K) \leq \eta_{\chi^2}(K) \leq \min\{\eta_{\mathrm{TV}}(K), 1 - \mathrm{BC}_{\min}(K)^2\} \leq 2H^2_{\max}(K)$, yielding the universally sharp factor-of-2 approximation $1 \leq \eta_{\chi^2}(K)/\mathrm{LC}_{\max}(K) \leq 2$ (saturated by Z_ϵ as $\epsilon \to 0^+$, and exact equality $\eta_{\chi^2} = \mathrm{LC}$ for every involution-symmetric binary channel).

- **Theorem 5.1 (Chernoff Tensorization Exponents and Parity-Resolved sech-Lattice Law):** While product priors satisfy $\eta_{\chi^2}(q^{\otimes d}, K^{\otimes d}) = \eta_{\chi^2}(q, K)$, an elementary Chebyshev-tilted proof yields $\frac{1}{2} e^{-dC_{\max}(K) - \sqrt{2dV_{\max}(K)}} \leq 1 - \eta_{\chi^2}(K^{\otimes d}) \leq e^{-dC_{\max}(K)}$ for all $d \geq 1$ when $S_{ik} \neq \emptyset$, establishing $\lim_{d\to\infty} -\frac{1}{d}\log(1 - \eta_{\chi^2}(K^{\otimes d})) = C_{\max}(K)$. For $\mathrm{BSC}(p)^{\otimes d}$, we prove the exact even/odd parity prefactor law $1 - \eta_{\chi^2}(\mathrm{BSC}(p)^{\otimes d}) \sim \sqrt{2/(\pi d)}\, S_\rho(\alpha)\, e^{-dC(p)}$ ($d \equiv \rho \pmod 2$) governed by the Jacobi hyperbolic-secant lattice sums $S_0(\alpha)$ and $S_1(\alpha)$.

All analytical formulas, spectral bounds, saturation classes, and asymptotic laws are verified numerically in an open-source Python suite (https://github.com/lluiseriksson/sdpi-bhattacharyya-laplacian).

2 Exact Bhattacharyya–Volatility Laplacian Decomposition

Throughout, $K \in \mathbb{R}_{\geq 0}^{m \times n}$ ($m, n \geq 2$) is a row-stochastic matrix with rows $K_1, \ldots, K_m \in \Delta_n$. For any full-support input prior $q \in \Delta_m^+$, let $r := qK \in \Delta_n$ be the induced output marginal ($r_j = \sum_{i=1}^m q_i K_{ij}$); restricting outputs to $\mathrm{supp}(r) = \{j \in [n] : r_j > 0\}$ leaves all divergences invariant, so without loss of generality $r_j > 0$ for all $j \in [n]$ in local statements.

Define the normalized channel operator $A(q,K) := D_q^{1/2} K D_r^{-1/2} \in \mathbb{R}_{\geq 0}^{m\times n}$ and the symmetrized forward–backward Gram matrix

$$G(q,K) := A(q,K)A(q,K)^\top = D_q^{1/2} K D_r^{-1} K^\top D_q^{1/2} \in \mathbb{R}_{\geq 0}^{m\times m},$$

whose entries are $G(q,K)_{ik} = \sqrt{q_i q_k}\, W_{ik}(q)$ with

$$W_{ik}(q) := \sum_{j\in \mathrm{supp}(r)} \frac{K_{ij}K_{kj}}{r_j}. \tag{1}$$

Since $G(q,K)\sqrt{q} = \sqrt{q}$ and $G(q,K) \succeq 0$ has eigenvalues $1 = \sigma_1^2 \geq \sigma_2^2 \geq \cdots \geq \sigma_m^2 \geq 0$, the variational characterization of Pearson χ^2 contraction [1, 6] gives $\eta_{\chi^2}(q,K) = \sigma_2^2(q,K) = \lambda_2(G(q,K))$ and, more generally, $I_m - G(q,K)$ has eigenvalues $0 = 1-\sigma_1^2 \leq 1-\sigma_2^2 \leq \cdots \leq 1-\sigma_m^2 \leq 1$.

Definition 2.1 (Prior-Weighted Graph Laplacian and Support-Normalized Affinities)**.** For any symmetric nonnegative matrix $M \in \mathbb{R}_{\geq 0}^{m\times m}$ and prior $q \in \Delta_m^+$, define the symmetric prior-weighted normalized graph Laplacian

$$\mathcal{L}_M(q) := \mathrm{diag}(Mq) - D_q^{1/2} M D_q^{1/2} \in \mathbb{R}^{m\times m}.$$

For each pair $(i,k) \in [m]^2$, let $S_{ik} := \{j \in [n] : K_{ij}K_{kj} > 0\}$ be the joint support, $r(S_{ik}) := \sum_{j\in S_{ik}} r_j \in [0,1]$ its marginal output mass, and $\mathrm{BC}(K_i,K_k) := \sum_{j=1}^n \sqrt{K_{ij}K_{kj}}$ the Bhattacharyya coefficient. Define the three symmetric nonnegative affinity matrices $B, \widetilde{B}(q), V(q) \in \mathbb{R}_{\geq 0}^{m\times m}$ by $B_{ik} := \mathrm{BC}(K_i,K_k)^2$,

$$\widetilde{B}_{ik}(q) := \begin{cases} \dfrac{\mathrm{BC}(K_i,K_k)^2}{r(S_{ik})}, & S_{ik} \neq \emptyset, \\ 0, & S_{ik} = \emptyset, \end{cases}$$

$$V_{ik}(q) := \begin{cases} r(S_{ik})\ \mathrm{Var}_{Y\sim r(\cdot|S_{ik})}\left(\dfrac{\sqrt{K_{iY}K_{kY}}}{r_Y}\right), & S_{ik} \neq \emptyset, \\ 0, & S_{ik} = \emptyset, \end{cases}$$

where $r(\cdot \mid S_{ik})$ is the conditional probability distribution $r_j/r(S_{ik})$ on S_{ik}. Set $\mathcal{L}_{\mathrm{BC}}(q) := \mathcal{L}_B(q)$, $\widetilde{\mathcal{L}}_{\mathrm{BC}}(q) := \mathcal{L}_{\widetilde{B}(q)}(q)$, and $\widetilde{\mathcal{L}}_{\mathrm{Var}}(q) := \mathcal{L}_{V(q)}(q)$.

Note that for any $v = D_q^{1/2} f \in \mathbb{R}^m$ and any symmetric $M \in \mathbb{R}_{\geq 0}^{m\times m}$,

$$v^\top \mathcal{L}_M(q) v = \sum_{i,k=1}^m q_i q_k M_{ik} f_i^2 - \sum_{i,k=1}^m q_i q_k M_{ik} f_i f_k = \frac{1}{2}\sum_{i,k=1}^m q_i q_k M_{ik}(f_i - f_k)^2 \geq 0, \tag{2}$$

with $\mathcal{L}_M(q)\sqrt{q} = 0$, and the diagonal entries M_{ii} cancel identically in $\mathcal{L}_M(q)$. We label the eigenvalues of any PSD Laplacian $\mathcal{L}_M(q)$ in non-decreasing order as $0 = \lambda_1(\mathcal{L}_M(q)) \leq \lambda_2(\mathcal{L}_M(q)) \leq \cdots \leq \lambda_m(\mathcal{L}_M(q))$.

Theorem 2.2 (Exact Bhattacharyya–Volatility Laplacian Decomposition)**.** *Let $K \in \mathbb{R}_{\geq 0}^{m\times n}$ be a Markov channel and $q \in \Delta_m^+$ any full-support prior.*

(a) ***Exact PSD Operator Splitting and Loewner Hierarchy.*** *The forward–backward contraction Laplacian decomposes identically as*

$$I_m - G(q,K) \;=\; \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) \;+\; \widetilde{\mathcal{L}}_{\mathrm{Var}}(q), \tag{3}$$

and satisfies the three-tier Loewner ordering

$$I_m - G(q,K) \;=\; \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) + \widetilde{\mathcal{L}}_{\mathrm{Var}}(q) \;\succeq\; \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) \;\succeq\; \mathcal{L}_{\mathrm{BC}}(q) \;\succeq\; 0. \tag{4}$$

(b) ***Full-Spectrum Weyl Bounds and Overlap-Graph Connectivity.*** *For every index* $r \in \{2, \dots, m\}$*, the squared singular values* $\sigma_r^2(q, K)$ *of* $D_q^{1/2} K D_{qK}^{-1/2}$ *satisfy*

$$\sigma_r^2(q,K) \;\le\; 1 - \lambda_r(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q)) \;\le\; 1 - \lambda_r(\mathcal{L}_{\mathrm{BC}}(q)) \;\le\; 1 - \min_{i \neq k} \mathrm{BC}(K_i, K_k)^2. \tag{5}$$

Summing over $r = 2, \dots, m$ *yields the exact Hilbert–Schmidt trace contraction bound*

$$\sum_{r=2}^{m} \sigma_r^2(q,K) \;\le\; (m-1) - \sum_{\substack{i,k=1 \\ i \neq k}}^{m} q_k \widetilde{B}_{ik}(q) \;\le\; (m-1) - \sum_{\substack{i,k=1 \\ i \neq k}}^{m} q_k \mathrm{BC}(K_i, K_k)^2. \tag{6}$$

In particular, $\eta_{\chi^2}(q, K) = \sigma_2^2(q, K) < 1$ *whenever the Bhattacharyya overlap graph* $G_{\mathrm{BC}} := ([m], \{\{i, k\} : \mathrm{BC}(K_i, K_k) > 0\})$ *is connected, even when* $\eta_{\mathrm{TV}}(K) = 1$ *and* $\min_{i \neq k} \mathrm{BC}(K_i, K_k) = 0$.

(c) ***Zero-Volatility Rigidity.*** *Equality* $I_m - G(q, K) = \widetilde{\mathcal{L}}_{\mathrm{BC}}(q)$ *(equivalently* $\widetilde{\mathcal{L}}_{\mathrm{Var}}(q) = 0$*) holds if and only if for every pair* $i \neq k$ *with* $S_{ik} \neq \emptyset$*, the ratio* $j \mapsto \sqrt{K_{ij} K_{kj}}/r_j$ *is constant on* S_{ik}*. In particular, every* single-overlap channel *(*$|S_{ik}| \le 1$ *for all* $i \neq k$*, including all simple weighted graph-incidence channels,* $\mathrm{BEC}(\epsilon)$*, Z-channels, and nearest-neighbor lazy cycle walks) satisfies* $I_m - G(q, K) = \widetilde{\mathcal{L}}_{\mathrm{BC}}(q)$ *identically for* every *prior* $q \in \Delta_m^+$.

(d) ***Support-Normalized Finite Abelian Group Fourier–Bhattacharyya Inequality.*** *Let* Γ *be a finite abelian group,* $c \in \Delta(\Gamma)$ *a probability distribution on* Γ *with shifted distributions* $(\tau_d c)(x) := c(x + d)$ *and overlap supports* $S_d := \{x \in \Gamma : c(x)c(x + d) > 0\}$*, and* $\hat{c}(\chi) := \sum_{x \in \Gamma} c(x)\chi(x)$ *its Fourier transform for any character* $\chi \in \widehat{\Gamma}$*. Then*

$$1 - |\hat{c}(\chi)|^2 \;\ge\; \sum_{\substack{d \in \Gamma \\ S_d \neq \emptyset}} \frac{\mathrm{BC}(c, \tau_d c)^2}{|S_d|} (1 - \mathrm{Re}\, \chi(d)) \;\ge\; \frac{1}{|\Gamma|} \sum_{d \in \Gamma} \mathrm{BC}(c, \tau_d c)^2 (1 - \mathrm{Re}\, \chi(d)), \tag{7}$$

with exact equality in the first inequality for all $\chi \in \widehat{\Gamma}$ *whenever* $|S_d| \le 1$ *for all* $d \neq 0$.

Proof. **(a)** First compute $(Wq)_i$ for each $i \in [m]$: using $r_j = \sum_{k=1}^m q_k K_{kj}$ and row stochasticity $\sum_j K_{ij} = 1$,

$$(Wq)_i = \sum_{k=1}^{m} q_k \sum_{j \in \mathrm{supp}(r)} \frac{K_{ij} K_{kj}}{r_j} = \sum_{j \in \mathrm{supp}(r)} \frac{K_{ij}}{r_j} \left(\sum_{k=1}^{m} q_k K_{kj} \right) = \sum_{j \in \mathrm{supp}(r)} K_{ij} = 1.$$

Hence $\mathrm{diag}(Wq) = I_m$, and therefore $I_m - G(q, K) = \mathrm{diag}(Wq) - D_q^{1/2} W D_q^{1/2} = \mathcal{L}_W(q)$.

Now fix any pair $(i, k) \in [m]^2$. If $S_{ik} = \emptyset$, then $W_{ik}(q) = 0 = \widetilde{B}_{ik}(q) + V_{ik}(q)$. If $S_{ik} \neq \emptyset$, then since $q \in \Delta_m^+$ we have $r_j \ge q_i K_{ij} + q_k K_{kj} > 0$ for every $j \in S_{ik}$, so $r(S_{ik}) > 0$. Define the random variable $Z_{ik}(Y) := \sqrt{K_{iY} K_{kY}}/r_Y$ under the conditional probability law $Y \sim r(\cdot \mid S_{ik})$. Direct evaluation gives

$$\mathbb{E}_{r(\cdot|S_{ik})}[Z_{ik}] = \sum_{j \in S_{ik}} \frac{r_j}{r(S_{ik})} \frac{\sqrt{K_{ij} K_{kj}}}{r_j} = \frac{\mathrm{BC}(K_i, K_k)}{r(S_{ik})},$$
$$\mathbb{E}_{r(\cdot|S_{ik})}[Z_{ik}^2] = \sum_{j \in S_{ik}} \frac{r_j}{r(S_{ik})} \frac{K_{ij} K_{kj}}{r_j^2} = \frac{W_{ik}(q)}{r(S_{ik})}.$$

By the identity $\mathbb{E}[Z_{ik}^2] = (\mathbb{E}[Z_{ik}])^2 + \mathrm{Var}(Z_{ik})$, multiplying both sides by $r(S_{ik})$ yields

$$W_{ik}(q) \;=\; \frac{\mathrm{BC}(K_i, K_k)^2}{r(S_{ik})} \;+\; r(S_{ik})\, \mathrm{Var}_{r(\cdot|S_{ik})}(Z_{ik}) \;=\; \widetilde{B}_{ik}(q) + V_{ik}(q). \tag{8}$$

Since the map $M \mapsto \mathcal{L}_M(q)$ is linear in M, substituting (8) into $I_m - G(q,K) = \mathcal{L}_W(q)$ establishes (3). Because $V_{ik}(q) \geq 0$ and $\widetilde{B}_{ik}(q) \geq B_{ik} \geq 0$ (as $r(S_{ik}) \leq 1$ whenever $S_{ik} \neq \emptyset$), (2) implies $\widetilde{\mathcal{L}}_{\mathrm{Var}}(q) \succeq 0$ and $\widetilde{\mathcal{L}}_{\mathrm{BC}}(q) - \mathcal{L}_{\mathrm{BC}}(q) = \mathcal{L}_{\widetilde{B}(q)-B}(q) \succeq 0$, proving (4).

(b) By Weyl's monotonicity theorem applied to (4), we have $1 - \sigma_r^2(q,K) = \lambda_r(I_m - G(q,K)) \geq \lambda_r(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q)) \geq \lambda_r(\mathcal{L}_{\mathrm{BC}}(q))$ for every $r \in \{2,\dots,m\}$. Moreover, writing $b_{\min} := \min_{i\neq k} \mathrm{BC}(K_i,K_k)^2$, for any $v = D_q^{1/2} f \perp \sqrt{q}$ with $\|v\|_2 = 1$ (so $\sum_i q_i f_i = 0$ and $\sum_i q_i f_i^2 = 1$), (2) gives

$$\begin{aligned} v^\top \mathcal{L}_{\mathrm{BC}}(q) v &= \frac{1}{2}\sum_{i\neq k} q_i q_k B_{ik}(f_i - f_k)^2 \\ &\geq \frac{b_{\min}}{2}\sum_{i,k} q_i q_k (f_i - f_k)^2 = b_{\min}\left(\sum_i q_i f_i^2 - \left(\sum_i q_i f_i\right)^2\right) = b_{\min}, \end{aligned}$$

so $\lambda_r(\mathcal{L}_{\mathrm{BC}}(q)) \geq \lambda_2(\mathcal{L}_{\mathrm{BC}}(q)) \geq b_{\min}$, proving (5). For (6), note that $\mathrm{Tr}(\mathcal{L}_M(q)) = \sum_i (Mq)_i - \sum_i q_i M_{ii} = \sum_{i\neq k} q_k M_{ik}$, and taking traces in (4) gives $\sum_{r=2}^m (1-\sigma_r^2) = \mathrm{Tr}(I_m - G(q,K)) \geq \mathrm{Tr}(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q)) \geq \mathrm{Tr}(\mathcal{L}_{\mathrm{BC}}(q))$. Finally, by (2), $v^\top \widetilde{\mathcal{L}}_{\mathrm{BC}}(q) v = 0$ iff $f_i = f_k$ whenever $\mathrm{BC}(K_i,K_k) > 0$; hence if G_{BC} is connected, $\ker(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q)) = \mathrm{span}\{\sqrt{q}\}$, so $\lambda_2(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q)) > 0$ and $\eta_{\chi^2}(q,K) \leq 1 - \lambda_2(\widetilde{\mathcal{L}}_{\mathrm{BC}}(q)) < 1$.

(c) By (2), $\widetilde{\mathcal{L}}_{\mathrm{Var}}(q) = 0$ iff $q_i q_k V_{ik}(q) = 0$ for all $i \neq k$. Since $q_i, q_k > 0$, this holds iff $\mathrm{Var}_{r(\cdot|S_{ik})}(\sqrt{K_{iY}K_{kY}}/r_Y) = 0$ whenever $S_{ik} \neq \emptyset$, i.e., $j \mapsto \sqrt{K_{ij}K_{kj}}/r_j$ is constant on S_{ik}. When $|S_{ik}| \leq 1$, any function on S_{ik} is trivially constant, so $V_{ik}(q) = 0$ for every $q \in \Delta_m^+$.

(d) Consider the convolution channel $K_{xy} := c(y-x)$ on Γ with uniform prior $u_x = 1/|\Gamma|$. Then $r_y = 1/|\Gamma|$ for all $y \in \Gamma$, and for any $x, z \in \Gamma$ with $d := z - x$, the joint support $S_{x,x+d} = \{y \in \Gamma : c(y-x)c(y-x-d) > 0\}$ has cardinality $|S_{x,x+d}| = |S_{-d}| = |S_d|$ (via $w = y - x - d$) and marginal mass $r(S_{x,x+d}) = |S_d|/|\Gamma|$, while $\mathrm{BC}(K_x, K_{x+d}) = \mathrm{BC}(c, \tau_d c)$. Thus $G(u,K)$, $\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)$, and $\mathcal{L}_{\mathrm{BC}}(u)$ are all Γ-invariant convolution operators, simultaneously diagonalized by the orthonormal character basis $\{|\Gamma|^{-1/2}\chi : \chi \in \widehat{\Gamma}\}$. On the eigenvector χ, $G(u,K)$ has eigenvalue $|\hat{c}(\chi)|^2$, while $\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)$ has eigenvalue

$$\frac{1}{|\Gamma|}\sum_{\substack{d\in\Gamma \\ S_d\neq\emptyset}} \frac{\mathrm{BC}(c,\tau_d c)^2}{|S_d|/|\Gamma|}(1-\chi(d)) = \sum_{\substack{d\in\Gamma \\ S_d\neq\emptyset}} \frac{\mathrm{BC}(c,\tau_d c)^2}{|S_d|}(1-\mathrm{Re}\,\chi(d)),$$

using symmetry $|S_{-d}| = |S_d|$ and $\mathrm{BC}(c,\tau_{-d}c) = \mathrm{BC}(c,\tau_d c)$ to take real parts. Evaluating (4) on the character χ yields (7), and part (c) gives exact equality when $|S_d| \leq 1$ for all $d \neq 0$. □

Example 2.3 (Lazy Cycle Walk W_m on $\mathbb{Z}_m$: Exactness and Factor-m Improvement)**.** Let $\Gamma = \mathbb{Z}_m$ ($m \geq 4$) and $c = \frac{1}{2}(\delta_0 + \delta_1)$, defining the lazy cycle walk channel $(W_m)_{ij} = \frac{1}{2}(\delta_{j,i} + \delta_{j,i+1 \bmod m})$. Because rows 0 and $\lfloor m/2 \rfloor$ have disjoint supports, Dobrushin's coefficient is trivial ($\eta_{\mathrm{TV}}(W_m) = 1$) and $\min_{i\neq k} \mathrm{BC} = 0$. Yet $S_{\pm 1}$ has cardinality $|S_{\pm 1}| = 1$ with $\mathrm{BC}(c, \tau_{\pm 1} c) = 1/2$, while $S_d = \emptyset$ for all $d \notin \{0, \pm 1\}$. By Theorem 2.2(c)–(d), $\widetilde{\mathcal{L}}_{\mathrm{Var}}(u) \equiv 0$, and for every non-trivial Fourier character $\chi_t(x) = e^{2\pi i t x/m}$ ($t \in \{1,\dots,m-1\}$), the corresponding mode eigenvalue $\tilde{\sigma}_t^2(u, W_m) = |\hat{c}(\chi_t)|^2$ of $G(u, W_m)$ and $\tilde{\lambda}_t(\widetilde{\mathcal{L}}_{\mathrm{BC}}(u))$ of $\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)$ satisfy

$$1 - \tilde{\sigma}_t^2(u, W_m) = \tilde{\lambda}_t(\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)) = 2 \cdot \frac{(1/2)^2}{1}(1 - \cos(2\pi t/m)) = \sin^2\left(\frac{\pi t}{m}\right),$$

matching the exact spectrum identically (with spectral gap $\lambda_2(\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)) = \sin^2(\pi/m)$ at $t \in \{1, m-1\}$), whereas unnormalized $\mathcal{L}_{\mathrm{BC}}(u)$ yields $\frac{1}{m}\sin^2(\pi t/m)$—a strict factor-$m$ improvement.

3 Binary-Input Harmonic-Risk Duality and Saturation Rigidity

We now specialize to an arbitrary binary-input discrete channel $K = \binom{p}{q}$ $(p, q \in \Delta_n)$ with joint support $S := \{j \in [n] : p_j q_j > 0\}$. Under input prior Bern$(y)$ $(y \in (0,1))$, the output marginal is $r_y := (1-y)p + yq$, and the Bayesian posterior probability of $X = 1$ given $Y = j$ is $\Pi(j) := yq_j/r_y(j)$ for $j \in \mathrm{supp}(r_y)$.

For $m = 2$, the unique unit vector $v \perp \sqrt{(1-y, y)}$ corresponds to $f_1 = \sqrt{y/(1-y)}$, $f_2 = -\sqrt{(1-y)/y}$, so $q_1 q_2 (f_1 - f_2)^2 = 1$. Substituting into $I_2 - G(y, K) = \mathcal{L}_W(y)$ from the proof of Theorem 2.2(a) yields the exact *harmonic overlap risk*

$$H_{p,q}(y) \;:=\; 1 - \eta_{\chi^2}(y; p, q) \;=\; W_{12}(y) \;=\; \sum_{j \in S} \frac{p_j q_j}{(1-y)p_j + yq_j}, \qquad y \in (0,1). \tag{9}$$

Notice that $H_{p,q}(y)$ extends continuously to $[0,1]$ with boundary limits $H_{p,q}(0^+) = \sum_{j \in S} q_j = Q(\mathrm{supp}(p))$ and $H_{p,q}(1^-) = \sum_{j \in S} p_j = P(\mathrm{supp}(q))$, where $P(A) := \sum_{j \in A} p_j$ and $Q(A) := \sum_{j \in A} q_j$. Define the support-restricted Pearson discrepancies

$$\Delta_0(p, q) := \sum_{j \in S} q_j \left(\frac{q_j}{p_j} - 1 \right), \qquad \Delta_1(p, q) := \sum_{j \in S} p_j \left(\frac{p_j}{q_j} - 1 \right). \tag{10}$$

Theorem 3.1 (Harmonic-Risk Calculus, Four-Regime Trichotomy, and Exact Saturation Rigidity)**.** *Let $K = \binom{p}{q}$ $(p, q \in \Delta_n)$ be any binary-input channel with joint support S.*

(a) ***Exact Posterior-Covariance Calculus and Strict Convexity.*** *On $(0,1)$, $H_{p,q} \in C^\infty(0,1)$ satisfies*

$$H_{p,q}(y) = \frac{\mathbb{E}_{r_y}[\Pi(1-\Pi)]}{y(1-y)}, \tag{11}$$

$$H'_{p,q}(y) = \sum_{j \in S} \frac{p_j q_j (p_j - q_j)}{r_y(j)^2} = -\frac{\mathrm{Cov}_{r_y}(\Pi,\, \Pi(1-\Pi))}{y^2(1-y)^2}, \tag{12}$$

$$H''_{p,q}(y) = 2 \sum_{j \in S} \frac{p_j q_j (q_j - p_j)^2}{r_y(j)^3} = \frac{2}{y^3(1-y)^3}\, \mathbb{E}_{r_y}\left[\Pi(1-\Pi)(\Pi - y)^2\right] \geq 0. \tag{13}$$

Consequently, $H_{p,q}$ is convex on $[0,1]$, and strictly convex unless $p_j = q_j$ for all $j \in S$. Moreover, for every $y \in (0,1)$, the exact Bhattacharyya slack equals the posterior-volatility variance:

$$(1 - \mathrm{BC}(p, q)^2) - \eta_{\chi^2}(y; p, q) \;=\; \mathrm{Var}_{r_y}\left(\frac{\sqrt{p(Y)q(Y)}}{r_y(Y)} \right) \;=\; \frac{\mathrm{Var}_{r_y}(\sqrt{\Pi(1-\Pi)})}{y(1-y)} \;\geq\; 0. \tag{14}$$

(b) ***Exact Minimizer Phase Classification.*** *Assume $S \neq \emptyset$ (otherwise $H_{p,q} \equiv 0$ and $\eta_{\chi^2}(p, q) = 1$). Then $H'_{p,q}(0^+) = -\Delta_0(p, q)$ and $H'_{p,q}(1^-) = +\Delta_1(p, q)$, and (p, q) falls into exactly one of four mutually exclusive regimes:*

- Flat Generalized-BEC Regime ($p_j = q_j$ $\forall j \in S$, equivalently $\Delta_0 = \Delta_1 = 0$): *$H_{p,q}(y) \equiv P(S) = Q(S)$ is constant on $[0,1]$ and $\eta_{\chi^2}(p, q) = 1 - P(S) = \mathrm{TV}(p, q)$.*
- Phase I (Left-Boundary Z-Regime: $\Delta_0 \leq 0 < \Delta_1$): *$H_{p,q}$ is strictly increasing on $[0,1]$, uniquely minimized at $y^* = 0^+$, and $\eta_{\chi^2}(p, q) = 1 - Q(\mathrm{supp}(p)) = Q(\mathrm{supp}(p)^c)$.*
- Phase II (Right-Boundary Reverse-Z Regime: $\Delta_1 \leq 0 < \Delta_0$): *$H_{p,q}$ is strictly decreasing on $[0,1]$, uniquely minimized at $y^* = 1^-$, and $\eta_{\chi^2}(p, q) = 1 - P(\mathrm{supp}(q)) = P(\mathrm{supp}(q)^c)$.*
- Phase III (Strict Interior Harmonic Regime: $\Delta_0 > 0$ and $\Delta_1 > 0$): *$H_{p,q}$ attains a unique strict interior minimum at $y^* \in (0,1)$ solving $\mathrm{Cov}_{r_{y^*}}(\Pi, \Pi(1-\Pi)) = 0$. In particular, every full-support pair ($\mathrm{supp}(p) = \mathrm{supp}(q)$, $p \neq q$) lies in Phase III.*

(c) ***Exact Three-Class Global Bhattacharyya Saturation Rigidity.*** *The upper bound* $\eta_{\chi^2}(p,q) \le 1 - \mathrm{BC}(p,q)^2$ *holds with equality if and only if* (p,q) *belongs to one of three canonical classes:*

(i) Trivial or disjoint: $p = q$ *(*$\eta_{\chi^2} = 0$*) or* $S = \emptyset$ *(*$\eta_{\chi^2} = 1$*).*

(ii) Generalized Z-channel (boundary saturation): $S \neq \emptyset$ *and either* $\mathrm{supp}(p) \subseteq \mathrm{supp}(q)$ *with* $q_j = cp_j$ *for all* $j \in \mathrm{supp}(p)$ *(*$c \in (0,1)$*, Phase I), or* $\mathrm{supp}(q) \subseteq \mathrm{supp}(p)$ *with* $p_j = cq_j$ *for all* $j \in \mathrm{supp}(q)$ *(*$c \in (0,1)$*, Phase II); in both cases* $\eta_{\chi^2}(p,q) = 1 - c = 1 - \mathrm{BC}(p,q)^2$.

(iii) Full-support two-ratio channel (interior saturation): $\mathrm{supp}(p) = \mathrm{supp}(q)$ *and the likelihood-ratio spectrum* $\{p_j/q_j : j \in S\}$ *takes exactly two distinct values* $\{r_1, r_2\}$ *(so merging outputs with equal likelihood ratio identifies* K *with a* 2×2 *binary asymmetric channel* $\mathrm{BAC}(a,b)$*), with unique optimal prior* $y^* = \sqrt{r_1 r_2}/(1 + \sqrt{r_1 r_2})$ *and* $\eta_{\chi^2}(p,q) = (\sqrt{a(1-b)} - \sqrt{b(1-a)})^2 = 1 - \mathrm{BC}(p,q)^2$.

Proof. **(a)** For every $j \in \mathrm{supp}(r_y)$, $\Pi(j) = yq_j/r_y(j)$ and $1 - \Pi(j) = (1-y)p_j/r_y(j)$, so $r_y(j)\Pi(j)(1-\Pi(j)) = y(1-y)p_jq_j/r_y(j)$ (which vanishes automatically outside S), giving (11). Differentiating (9) termwise yields the first equalities in (12) and (13). Moreover,

$$y - \Pi(j) = y(1-\Pi(j)) - (1-y)\Pi(j) = y(1-y)\,\frac{p_j - q_j}{r_y(j)},$$

so $\frac{p_jq_j(p_j-q_j)}{r_y(j)^2} = \frac{1}{y^2(1-y)^2} r_y(j)\Pi(j)(1-\Pi(j))(y-\Pi(j))$. Summing over $j \in \mathrm{supp}(r_y)$ and using $\mathbb{E}_{r_y}[\Pi] = y$ gives $\mathbb{E}_{r_y}[\Pi(1-\Pi)(y-\Pi)] = -\mathrm{Cov}_{r_y}(\Pi, \Pi(1-\Pi))$, proving (12). Similarly,

$$\frac{p_jq_j(q_j-p_j)^2}{r_y(j)^3} = \frac{1}{y^3(1-y)^3}\, r_y(j)\,\Pi(j)(1-\Pi(j))\,(\Pi(j)-y)^2,$$

proving (13). Clearly $H''_{p,q}(y) = 0$ iff $p_j = q_j$ for all $j \in S$. Finally, setting $Z(Y) := \sqrt{p(Y)q(Y)}/r_y(Y) = \sqrt{\Pi(Y)(1-\Pi(Y))}/\sqrt{y(1-y)}$ under $Y \sim r_y$, we have $\mathbb{E}_{r_y}[Z] = \sum_j \sqrt{p_jq_j} = \mathrm{BC}(p,q)$ and $\mathbb{E}_{r_y}[Z^2] = \sum_{j\in S} p_jq_j/r_y(j) = H_{p,q}(y)$, so $H_{p,q}(y) - \mathrm{BC}(p,q)^2 = \mathrm{Var}_{r_y}(Z)$, proving (14).

(b) Letting $y \to 0^+$ and $y \to 1^-$ in the finite sum (12) gives $H'_{p,q}(0^+) = \sum_{j\in S} q_j(1 - q_j/p_j) = -\Delta_0(p,q)$ and $H'_{p,q}(1^-) = \sum_{j\in S} p_j(p_j/q_j - 1) = \Delta_1(p,q)$. Since $H''_{p,q} \ge 0$, $H'_{p,q}$ is non-decreasing on $[0,1]$, so $-\Delta_0 \le \Delta_1$ always holds, with equality $-\Delta_0 = \Delta_1$ iff $H''_{p,q} \equiv 0$, i.e., $p_j = q_j$ for all $j \in S$ (in which case $\Delta_0 = \Delta_1 = 0$ and $H_{p,q} \equiv P(S) = 1 - \mathrm{TV}(p,q)$). Otherwise $H_{p,q}$ is strictly convex and $-\Delta_0 < \Delta_1$, partitioning cleanly into: $0 \le -\Delta_0 < \Delta_1$ (Phase I, $H' > 0$ on $(0,1)$, minimum at 0^+), $-\Delta_0 < \Delta_1 \le 0$ (Phase II, $H' < 0$ on $(0,1)$, minimum at 1^-), and $-\Delta_0 < 0 < \Delta_1$ (Phase III, unique interior root of $H'(y^*) = 0$). When $\mathrm{supp}(p) = \mathrm{supp}(q)$ and $p \neq q$, $\Delta_0 = \chi^2(q \,\|\, p) > 0$ and $\Delta_1 = \chi^2(p \,\|\, q) > 0$, placing (p,q) automatically in Phase III.

(c) Recall $\eta_{\chi^2}(p,q) = 1 - \min_{y\in[0,1]} H_{p,q}(y)$. Assume $p \neq q$ and $S \neq \emptyset$. We analyze each regime from part (b):

- *Flat* BEC *regime:* $p_j = q_j$ on S and $p \neq q$ imply $0 < P(S) < 1$ and $\mathrm{BC}(p,q) = P(S)$, so $H_{p,q} \equiv P(S) > P(S)^2 = \mathrm{BC}(p,q)^2$; saturation never occurs.
- *Phase I (*$\Delta_0 \le 0 < \Delta_1$*):* Here $\min H_{p,q} = H_{p,q}(0^+) = Q(S)$. By Cauchy–Schwarz on S,

$$\mathrm{BC}(p,q)^2 = \left(\sum_{j\in S} \sqrt{p_jq_j}\right)^2 \le P(S)\,Q(S) \le Q(S) = H_{p,q}(0^+),$$

with equality iff $P(S) = 1$ (so $\mathrm{supp}(p) \subseteq \mathrm{supp}(q)$) *and* there exists $c > 0$ such that $q_j = cp_j$ for all $j \in S$. Summing over $j \in S$ forces $c = Q(S) \in (0,1)$ (since $p \neq q$). Conversely, any such generalized Z-channel has $\Delta_0 = c(c-1) < 0 < c^{-1} - 1 = \Delta_1$, hence lies in Phase I with $\eta_{\chi^2} = 1 - c = 1 - \mathrm{BC}^2$. Phase II is identical by exchanging p and q.

- *Phase III ($\Delta_0, \Delta_1 > 0$):* Here the minimum is achieved at a unique interior $y^* \in (0,1)$. By (14), $H_{p,q}(y^*) = \mathrm{BC}(p,q)^2$ iff $Z(j) = \sqrt{p_j q_j}/r_{y^*}(j)$ is constant across *all* $j \in \mathrm{supp}(r_{y^*}) = \mathrm{supp}(p) \cup \mathrm{supp}(q)$. Since $Z(j) > 0$ on $S \neq \emptyset$ and $Z(j) = 0$ outside S, constancy forces $\mathrm{supp}(p) = \mathrm{supp}(q) = S$ and the existence of $\kappa > 0$ such that for every $j \in S$,

$$\frac{\sqrt{p_j q_j}}{(1-y^*)p_j + y^* q_j} = \kappa \quad \Longleftrightarrow \quad (1-y^*)\sqrt{\frac{p_j}{q_j}} + y^* \sqrt{\frac{q_j}{p_j}} = \frac{1}{\kappa}.$$

Setting $\Lambda_j := \sqrt{p_j/q_j} > 0$, this is the quadratic equation $(1-y^*)\Lambda_j^2 - \kappa^{-1}\Lambda_j + y^* = 0$, which has at most two positive roots Λ_+, Λ_- satisfying $\Lambda_+\Lambda_- = y^*/(1-y^*)$. Since $p \neq q$, the likelihood ratio $r_j := p_j/q_j = \Lambda_j^2$ must take *exactly two* distinct values $\{r_1, r_2\}$ on S (one > 1, one < 1), and $y^*/(1-y^*) = \sqrt{r_1 r_2} \iff y^* = \sqrt{r_1 r_2}/(1+\sqrt{r_1 r_2})$. Merging outputs with $p_j/q_j = r_1$ into symbol 1 (total mass a under p, b under q) and outputs with $p_j/q_j = r_2$ into symbol 2 yields a $\mathrm{BAC}(a,b)$ with identical $H_{p,q}(y)$ and $\mathrm{BC}(p,q) = \sqrt{ab} + \sqrt{(1-a)(1-b)}$, giving $1 - \mathrm{BC}(p,q)^2 = (\sqrt{a(1-b)} - \sqrt{b(1-a)})^2$. Conversely, any full-support two-ratio channel satisfies this quadratic equation at $y^* = \sqrt{r_1 r_2}/(1+\sqrt{r_1 r_2})$ with constant $Z(j)$, so $\mathrm{Var}_{r_{y^*}}(Z) = 0$ and saturation holds. □

4 Polytope-Vertex Binary Reduction and the Universal Factor-of-2 Le Cam Sandwich

We next lift Theorem 3.1 to arbitrary $m \times n$ Markov channels K. Recall that the squared Hellinger distance is $H^2(p,q) := 1 - \mathrm{BC}(p,q) = \frac{1}{2}\sum_{j=1}^n (\sqrt{p_j} - \sqrt{q_j})^2$ and the symmetric Le Cam (triangular) discrimination is

$$\mathrm{LC}(p,q) := \frac{1}{2} \sum_{j : p_j + q_j > 0} \frac{(p_j - q_j)^2}{p_j + q_j} = 1 - \sum_{j \in S} \frac{2 p_j q_j}{p_j + q_j} = 1 - H_{p,q}\left(\tfrac{1}{2}\right).$$

For an $m \times n$ channel K, define the pairwise maxima $H^2_{\max}(K) := \max_{i<k} H^2(K_i, K_k)$, $\mathrm{LC}_{\max}(K) := \max_{i<k} \mathrm{LC}(K_i, K_k)$, and $\mathrm{BC}_{\min}(K) := \min_{i<k} \mathrm{BC}(K_i, K_k)$.

Theorem 4.1 (Polytope-Vertex Binary Reduction and Universal Factor-of-2 Le Cam Sandwich)**.** *Let $K \in \mathbb{R}_{\geq 0}^{m \times n}$ ($m, n \geq 2$) be any Markov channel.*

(a) ***Binary-Input Reduction*** *([4, 7]; self-contained polytope proof). The global χ^2-contraction coefficient reduces identically to binary-input subchannels:*

$$\eta_{\chi^2}(K) \;=\; \max_{1 \leq i < k \leq m} \eta_{\chi^2}(K_i, K_k) \;=\; 1 - \min_{1 \leq i < k \leq m} \inf_{y \in (0,1)} H_{K_i, K_k}(y). \tag{15}$$

(b) ***Universal Two-Sided Sandwich.*** *Every Markov channel K satisfies*

$$H^2_{\max}(K) \;\leq\; \mathrm{LC}_{\max}(K) \;\leq\; \eta_{\chi^2}(K) \;\leq\; \min\Big\{\eta_{\mathrm{TV}}(K),\; 1 - \mathrm{BC}_{\min}(K)^2\Big\} \;\leq\; 2H^2_{\max}(K). \tag{16}$$

In particular, whenever $\eta_{\chi^2}(K) > 0$,

$$1 \;\leq\; \frac{\eta_{\chi^2}(K)}{\mathrm{LC}_{\max}(K)} \;\leq\; \frac{1 - \mathrm{BC}_{\min}(K)^2}{\mathrm{LC}_{\max}(K)} \;\leq\; 2. \tag{17}$$

(c) ***Universal Sharpness and Complementary Extremals.***

(i) Exact lower bound for involution-symmetric channels: *If* (p, q) *is involution-symmetric—meaning there exists an involution* π *of* $[n]$ *(*$\pi^2 = \mathrm{id}$*) such that* $q_j = p_{\pi(j)}$ *(and hence simultaneously* $p_j = q_{\pi(j)}$*) for all* $j \in [n]$*, which covers every binary memoryless symmetric (BMS) channel—then* $H_{p,q}(1-y) = H_{p,q}(y)$*,* $y^* = 1/2$*, and* $\eta_{\chi^2}(p, q) = \mathrm{LC}(p, q)$ *identically. (Without* $\pi^2 = \mathrm{id}$*, cyclic permutations on* $n \geq 3$ *outputs can fail this symmetry.) Moreover,* $\eta_{\chi^2}(p, q)/\mathrm{LC}(p, q) \to 1$ *both in the* interior high-noise limit *(*$p \neq q$ *with fixed support* $\mathrm{supp}(q) = \mathrm{supp}(p)$ *and* $\max_{j \in \mathrm{supp}(p)} |q_j/p_j - 1| \to 0$*, in sharp contrast to singular support perturbations such as* Z_ϵ *below) and in the disjoint limit* $\mathrm{BC}(p, q) \to 0$*.*

(ii) Sharpness of the constant 2: *For the Z-channel family* $Z_\epsilon = \left(\begin{smallmatrix} 1 & 0 \\ 1-\epsilon & \epsilon \end{smallmatrix}\right)$ *(*$\epsilon \in (0, 1)$*),* $\eta_{\chi^2}(Z_\epsilon) = \epsilon$ *while* $\mathrm{LC}(Z_\epsilon) = \epsilon/(2-\epsilon)$*, so* $\eta_{\chi^2}(Z_\epsilon)/\mathrm{LC}(Z_\epsilon) = 2 - \epsilon \to 2$ *as* $\epsilon \to 0^+$*.*

(iii) Incomparability of the Dobrushin and Bhattacharyya ceilings: *On* $\mathrm{BSC}(\frac{1}{2} - \varepsilon)$*,* $(1 - \mathrm{BC}^2)/\eta_{\mathrm{TV}} = 2\varepsilon \to 0$ *(quadratic improvement), whereas on* $\mathrm{BEC}(1-\delta)$*,* $\eta_{\mathrm{TV}}/(1-\mathrm{BC}^2) = 1/(2-\delta) \to 1/2$*.*

Proof. **(a)** Obviously $\eta_{\chi^2}(K) \geq \max_{i<k} \eta_{\chi^2}(K_i, K_k)$ by restricting priors to two-point supports. For the reverse inequality, fix any $q \in \Delta_m$ and any output function $g \in \mathbb{R}^n$ with $\mathrm{Var}_{qK}(g) > 0$. By the dual variational characterization $\eta_{\chi^2}(q, K) = \sup_{\mathrm{Var}_{qK}(g)>0} \mathrm{Var}_q(Kg)/\mathrm{Var}_{qK}(g)$, shift $g \mapsto g - \mathbb{E}_{qK}[g]\mathbf{1}$ so that $\mathbb{E}_q[Kg] = \mathbb{E}_{qK}[g] = 0$. Set $h_i := (Kg)_i$ and $s_i := (Kg^2)_i \geq h_i^2 \geq 0$ (by Jensen's inequality), and let $\lambda := (\sum_i q_i h_i^2)/(\sum_i q_i s_i) \in [0, 1]$ (assuming $\lambda > 0$, otherwise trivial). Then q lies in the compact convex polytope $P_h := \{w \in \Delta_m : \sum_{i=1}^m w_i h_i = 0\}$ and satisfies $\sum_{i=1}^m q_i(h_i^2 - \lambda s_i) = 0$. Writing $q = \sum_\alpha t_\alpha w^{(\alpha)}$ as a convex combination of extreme points (vertices) of P_h, every vertex of the intersection of the simplex Δ_m with a single hyperplane has $|\mathrm{supp}(w^{(\alpha)})| \leq 2$: either a single-coordinate vertex e_i with $h_i = 0$ (where $h_i^2 - \lambda s_i = -\lambda s_i \leq 0$) or a two-point vertex $w = (1-y)e_i + ye_k$ with $h_i > 0 > h_k$ and $y = h_i/(h_i - h_k) \in (0, 1)$. Because $\sum_\alpha t_\alpha \sum_r w_r^{(\alpha)}(h_r^2 - \lambda s_r) = 0$ and $\sum_i q_i h_i^2 > 0$, at least one two-point vertex $w^* = (1-y)e_i + ye_k \in P_h$ satisfies $\sum_r w_r^*(h_r^2 - \lambda s_r) \geq 0$. Since $(1-y)s_i + ys_k \geq (1-y)h_i^2 + yh_k^2 > 0$, we obtain

$$\eta_{\chi^2}(K_i, K_k) \;\geq\; \eta_{\chi^2}(w^*, K) \;\geq\; \frac{\mathrm{Var}_{w^*}(Kg)}{\mathrm{Var}_{w^*K}(g)} = \frac{(1-y)h_i^2 + yh_k^2}{(1-y)s_i + ys_k} \;\geq\; \lambda.$$

Taking the supremum over g and q establishes (15).

(b) For any pair (p, q) of distributions on $[n]$:

- Comparing term-by-term for $p_j + q_j > 0$, $(\sqrt{p_j} + \sqrt{q_j})^2 = p_j + q_j + 2\sqrt{p_j q_j} \geq p_j + q_j$, so

$$\mathrm{LC}(p, q) = \frac{1}{2}\sum_j \frac{(\sqrt{p_j} - \sqrt{q_j})^2(\sqrt{p_j} + \sqrt{q_j})^2}{p_j + q_j} \geq \frac{1}{2}\sum_j (\sqrt{p_j} - \sqrt{q_j})^2 = H^2(p, q).$$

- By (9), $\eta_{\chi^2}(p, q) = 1 - \inf_{y \in (0,1)} H_{p,q}(y) \geq 1 - H_{p,q}(1/2) = \mathrm{LC}(p, q)$.
- For every $y \in (0, 1)$ and $j \in S$, $(1-y)p_j + yq_j \leq \max(p_j, q_j)$, so $H_{p,q}(y) \geq \sum_j \min(p_j, q_j) = 1 - \mathrm{TV}(p, q)$, giving $\eta_{\chi^2}(p, q) \leq \mathrm{TV}(p, q)$; simultaneously Theorem 3.1(a) gives $\eta_{\chi^2}(p, q) \leq 1 - \mathrm{BC}(p, q)^2$.
- Finally, $1 - \mathrm{BC}(p, q)^2 = (1 + \mathrm{BC}(p, q))(1 - \mathrm{BC}(p, q)) \leq 2H^2(p, q)$.

Taking $\max_{1 \leq i < k \leq m}$ across all inequalities and invoking (15) yields (16) and (17).

(c) For (i), if π is an involution ($\pi^2 = \mathrm{id}$) with $q_j = p_{\pi(j)}$ and $p_j = q_{\pi(j)}$ for all j, then re-indexing the sum (9) via $j \mapsto \pi(j)$ swaps p_j and q_j simultaneously and yields $H_{p,q}(1-y) = H_{p,q}(y)$; by convexity, $H_{p,q}(1/2) \leq \frac{1}{2}(H_{p,q}(y) + H_{p,q}(1-y)) = H_{p,q}(y)$ for all $y \in (0, 1)$, so $y^* = 1/2$ and $\eta_{\chi^2}(p, q) = \mathrm{LC}(p, q)$. For the interior high-noise limit, write $q_j = p_j(1 + \delta_j)$ on the common support $S = \mathrm{supp}(p) = \mathrm{supp}(q)$ with $\sum_{j \in S} p_j \delta_j = 0$ and $\varepsilon := \max_{j \in S} |\delta_j| \to 0^+$:

then $\mathrm{LC}(p,q) = \frac{1}{4}\sum_{j\in S} p_j\delta_j^2(1+O(\varepsilon))$ and $1-\mathrm{BC}(p,q)^2 = \frac{1}{4}\sum_{j\in S} p_j\delta_j^2(1+O(\varepsilon))$, forcing $1 \le \eta_{\chi^2}/\mathrm{LC} \le (1-\mathrm{BC}^2)/\mathrm{LC} \to 1$; as $\mathrm{BC}\to 0$, $H^2\to 1$ and $1-\mathrm{BC}^2\to 1$, again forcing $\eta_{\chi^2}/\mathrm{LC}\to 1$. Parts (ii) and (iii) follow immediately by direct substitution into $H_{p,q}$, LC, TV, and BC. □

5 Tensor-Power Chernoff Amplification and Parity-Resolved sech-Lattice Laws

Consider the d-fold product channel $K^{\otimes d} \in \mathbb{R}_{\ge 0}^{m^d\times n^d}$ acting independently on d parallel uses of K. For any two distributions $p,q\in\Delta_n$, define the Chernoff function $\psi_{p,q}(t) := \sum_{j\in S} p_j^t q_j^{1-t}$ ($t\in[0,1]$, extended by $\psi_{p,q}(0)=Q(S)$ and $\psi_{p,q}(1)=P(S)$) and the Chernoff information [3]

$$C(p,q) := -\log \inf_{t\in[0,1]} \psi_{p,q}(t) \in [0,\infty], \qquad C_{\max}(K) := \max_{1\le i<k\le m} C(K_i,K_k).$$

Whenever $S_{ik}\neq\emptyset$ for all $i,k\in[m]$, define the finite channel log-likelihood spread

$$V_{\max}(K) := \max_{1\le i<k\le m}\ \max_{j\in S_{ik}} \left(\log\frac{K_{ij}}{K_{kj}}\right)^2 < \infty.$$

Theorem 5.1 (Product-Prior Invariance, Exact Chernoff Tensorization Exponent, and Parity sech-Lattice Law)**.** *Let $K\in\mathbb{R}_{\ge 0}^{m\times n}$ be a discrete Markov channel.*

(a) ***Product-Prior Dimension-Free Invariance.*** *For every full-support prior $q\in\Delta_m^+$ and every $d\ge 1$, $\eta_{\chi^2}(q^{\otimes d},K^{\otimes d}) = \eta_{\chi^2}(q,K)$.*

(b) ***Uniform Finite-d Chebyshev-Tilted Sandwich and Exact Chernoff Rate.*** *Assume $S_{ik}\neq\emptyset$ for all $i,k\in[m]$ (otherwise $\eta_{\chi^2}(K)=\eta_{\chi^2}(K^{\otimes d})=1$ for all $d\ge 1$). Then for every integer $d\ge 1$,*

$$\begin{aligned}\frac{1}{2}\exp\Big(-d\,C_{\max}(K) - \sqrt{2d\,V_{\max}(K)}\Big) \;&\le\; 1-\eta_{\mathrm{TV}}(K^{\otimes d})\\ &\le\; 1-\eta_{\chi^2}(K^{\otimes d}) \;\le\; \exp\big(-d\,C_{\max}(K)\big).\end{aligned} \tag{18}$$

Consequently, the global χ^2- and TV-contraction deficits share the exact asymptotic Chernoff exponent:

$$\lim_{d\to\infty} -\frac{1}{d}\log(1-\eta_{\chi^2}(K^{\otimes d})) \;=\; \lim_{d\to\infty} -\frac{1}{d}\log(1-\eta_{\mathrm{TV}}(K^{\otimes d})) \;=\; C_{\max}(K). \tag{19}$$

(c) ***Exact Parity-Resolved* sech*-Lattice Law for* $\mathrm{BSC}(p)^{\otimes d}$.** *For the binary symmetric channel $\mathrm{BSC}(p)$ ($p\in(0,1/2)$) with Chernoff information $C(p) := -\log(2\sqrt{p(1-p)})$ and log-odds $\alpha := \log((1-p)/p) > 0$,*

$$1-\eta_{\chi^2}(\mathrm{BSC}(p)^{\otimes d}) \;=\; 2^{-d}\,e^{-dC(p)}\sum_{k=0}^{d}\binom{d}{k}\operatorname{sech}\left(\left(k-\frac{d}{2}\right)\alpha\right), \tag{20}$$

and as $d\to\infty$ along even or odd d separately,

$$\lim_{\substack{d\to\infty\\ d\equiv\rho \pmod 2}} \sqrt{\frac{\pi d}{2}}\,e^{dC(p)}\,(1-\eta_{\chi^2}(\mathrm{BSC}(p)^{\otimes d})) \;=\; S_\rho(\alpha) \qquad (\rho\in\{0,1\}), \tag{21}$$

where $S_0(\alpha) := \sum_{m\in\mathbb{Z}}\operatorname{sech}(m\alpha) = \frac{\pi}{\alpha}\sum_{n\in\mathbb{Z}}\operatorname{sech}(\pi^2 n/\alpha)$ and $S_1(\alpha) := \sum_{m\in\mathbb{Z}}\operatorname{sech}((m+\frac{1}{2})\alpha) = \frac{\pi}{\alpha}\sum_{n\in\mathbb{Z}}(-1)^n\operatorname{sech}(\pi^2 n/\alpha)$ satisfy $S_1(\alpha) < \frac{\pi}{\alpha} < S_0(\alpha)$ strictly for every $\alpha>0$.

Proof. **(a)** Because $G(q^{\otimes d}, K^{\otimes d}) = G(q,K)^{\otimes d}$ has eigenvalues $\{\prod_{r=1}^{d}\sigma_{j_r}^2 : j_1,\dots,j_d \in [m]\}$ with $\sigma_1^2 = 1 \geq \sigma_2^2 \geq \cdots \geq \sigma_m^2$, the second-largest eigenvalue of the tensor product is $1^{d-1}\cdot\sigma_2^2 = \eta_{\chi^2}(q,K)$.

(b) *Upper bound:* Choose (i^*,k^*) attaining $C(K_{i^*},K_{k^*}) = C_{\max}(K)$ and let $p := K_{i^*}$, $q := K_{k^*}$. For any $t \in (0,1)$ and any output tuple $\mathbf{j} \in S^d$, weighted AM–GM gives $(1-t)p_{\mathbf{j}} + tq_{\mathbf{j}} \geq p_{\mathbf{j}}^{1-t}q_{\mathbf{j}}^{t}$, so

$$1-\eta_{\chi^2}(K^{\otimes d}) \;\leq\; 1-\eta_{\chi^2}(p^{\otimes d},q^{\otimes d}) \;\leq\; H_{p^{\otimes d},q^{\otimes d}}(t) = \sum_{\mathbf{j}\in S^d}\frac{p_{\mathbf{j}}q_{\mathbf{j}}}{(1-t)p_{\mathbf{j}}+tq_{\mathbf{j}}} \;\leq\; \sum_{\mathbf{j}\in S^d} p_{\mathbf{j}}^{t}q_{\mathbf{j}}^{1-t} = \psi_{p,q}(t)^d.$$

Taking the infimum over $t \in (0,1)$ (and noting $\psi_{p,q}$ is continuous on $[0,1]$) gives $1-\eta_{\chi^2}(K^{\otimes d}) \leq e^{-dC_{\max}(K)}$.

Lower bound: By Theorem 4.1(b), $1-\eta_{\chi^2}(K^{\otimes d}) \geq 1-\eta_{\mathrm{TV}}(K^{\otimes d}) = \min_{\mathbf{i},\mathbf{k}\in[m]^d}(1-\mathrm{TV}(K_{\mathbf{i}},K_{\mathbf{k}}))$. Fix *any* input pair $(\mathbf{i},\mathbf{k}) \in [m]^d\times[m]^d$, set $P := K_{\mathbf{i}} = \bigotimes_{r=1}^{d}K_{i_r}$, $Q := K_{\mathbf{k}} = \bigotimes_{r=1}^{d}K_{k_r}$, and $S_{\mathbf{i},\mathbf{k}} := \prod_{r=1}^{d}S_{i_r,k_r} \neq \emptyset$. Let $\psi(t) := \prod_{r=1}^{d}\psi_{K_{i_r},K_{k_r}}(t)$ attain its minimum on $[0,1]$ at $t^* \in [0,1]$. Crucially,

$$\psi(t^*) = \prod_{r=1}^{d}\psi_{K_{i_r},K_{k_r}}(t^*) \;\geq\; \prod_{r=1}^{d}\inf_{s\in[0,1]}\psi_{K_{i_r},K_{k_r}}(s) = \prod_{r=1}^{d}e^{-C(K_{i_r},K_{k_r})} \;\geq\; e^{-d\,C_{\max}(K)}.$$

Define the tilted product probability measure $R(\mathbf{j}) := P(\mathbf{j})^{t^*}Q(\mathbf{j})^{1-t^*}/\psi(t^*)$ on $S_{\mathbf{i},\mathbf{k}}$, and let $L := \sum_{r=1}^{d}L_r$ with independent summands $L_r := \log(K_{i_r,Y_r}/K_{k_r,Y_r})$ under $Y \sim R$. Since each $|L_r| \leq \sqrt{V_{\max}(K)}$ almost surely, $\mathrm{Var}_R(L) = \sum_{r=1}^{d}\mathrm{Var}_{R_r}(L_r) \leq d\,V_{\max}(K)$. Setting $\mu := \mathbb{E}_R[L] = \psi'(t^*)/\psi(t^*)$ and $\tau := \sqrt{2dV_{\max}(K)}$, we observe that if $V_{\max}(K) = 0$ then $L \equiv \mu = 0$ R-almost surely so $R(|L-\mu| \leq \tau) = 1 \geq 1/2$ trivially, while if $V_{\max}(K) > 0$ then Chebyshev's inequality ensures $R(|L-\mu| \leq \tau) \geq 1/2$.

- If $t^* \in (0,1)$, then $\mu = 0$. On the event $E := \{|L| \leq \tau\}$, we have $\min(P(\mathbf{j}),Q(\mathbf{j})) = P(\mathbf{j})^{t^*}Q(\mathbf{j})^{1-t^*}\min(e^{(1-t^*)L},e^{-t^*L}) \geq \psi(t^*)R(\mathbf{j})e^{-\tau}$.
- If $t^* = 0$, then $\psi'(0^+) \geq 0$, so $\mu \geq 0$. On $E := \{L \geq \mu-\tau\}$ (which has $R(E) \geq 1/2$), $L \geq -\tau$, so $P(\mathbf{j}) = Q(\mathbf{j})e^{L} \geq Q(\mathbf{j})e^{-\tau}$ and hence $\min(P(\mathbf{j}),Q(\mathbf{j})) \geq Q(\mathbf{j})e^{-\tau} = \psi(0)R(\mathbf{j})e^{-\tau}$.
- If $t^* = 1$, then $\psi'(1^-) \leq 0$, so $\mu \leq 0$; symmetrically on $E := \{L \leq \mu+\tau\}$, $\min(P(\mathbf{j}),Q(\mathbf{j})) \geq \psi(1)R(\mathbf{j})e^{-\tau}$.

In all three cases, summing over $\mathbf{j} \in E$ yields

$$1-\mathrm{TV}(K_{\mathbf{i}},K_{\mathbf{k}}) = \sum_{\mathbf{j}\in S_{\mathbf{i},\mathbf{k}}}\min(P(\mathbf{j}),Q(\mathbf{j})) \;\geq\; \psi(t^*)\,e^{-\tau}\,R(E) \;\geq\; \frac{1}{2}\,e^{-d\,C_{\max}(K)-\sqrt{2dV_{\max}(K)}}.$$

Since this lower bound is uniform over all $(\mathbf{i},\mathbf{k}) \in [m]^d\times[m]^d$, taking the minimum yields (18), and taking $-\frac{1}{d}\log(\cdot)$ as $d\to\infty$ gives (19).

(c) For $\mathrm{BSC}(p)^{\otimes d}$, by binary reduction (15), any pair $(\mathbf{i},\mathbf{k}) \in \{0,1\}^d$ differing in $h \leq d$ coordinates gives a subchannel isomorphic to $\mathrm{BSC}(p)^{\otimes h}$ whose η_{χ^2} is non-decreasing in h (since appending independent noisy coordinates can only increase distinguishability by data processing), so the maximum is attained at Hamming distance $h = d$, i.e., $P = (1-p,p)^{\otimes d}$ and $Q = (p,1-p)^{\otimes d}$. Because (P,Q) is involution-symmetric (under bit-flip), Theorem 4.1(c)(i) gives $1-\eta_{\chi^2}(\mathrm{BSC}(p)^{\otimes d}) = H_{P,Q}(1/2)$. Grouping outputs by Hamming weight $k \in \{0,\dots,d\}$:

$$\begin{aligned} H_{P,Q}\left(\tfrac{1}{2}\right) &= \sum_{k=0}^{d}\binom{d}{k}\frac{2(p(1-p))^d}{p^k(1-p)^{d-k}+p^{d-k}(1-p)^k} \\ &= 2^{-d}\left(2\sqrt{p(1-p)}\right)^d\sum_{k=0}^{d}\binom{d}{k}\frac{2}{e^{(k-d/2)\alpha}+e^{-(k-d/2)\alpha}}, \end{aligned}$$

which is (20). Setting $m := k - d/2 \in \mathbb{Z} + \rho/2$ (for $d \equiv \rho \pmod 2$), Stirling's formula gives $\sqrt{\pi d/2}\, 2^{-d}\binom{d}{d/2+m} \to 1$ pointwise for every fixed m while remaining uniformly bounded above by a constant C_0 (specifically $2^{-d}\binom{d}{\lfloor d/2 \rfloor} \le \sqrt{2/(\pi d)} \cdot \sqrt{\pi/2}$ for all $d \ge 1$). Since $\sum_{m \in \mathbb{Z}+\rho/2} \operatorname{sech}(m\alpha) < \infty$, the dominated convergence theorem yields (21). Finally, the Fourier transform of $x \mapsto \operatorname{sech}(\alpha x)$ is $\int_{-\infty}^{\infty} \operatorname{sech}(\alpha x) e^{-2\pi i \xi x}\, dx = \frac{\pi}{\alpha} \operatorname{sech}(\pi^2 \xi/\alpha)$; Poisson summation on $\mathbb{Z}$ and $\mathbb{Z} + \frac{1}{2}$ gives the dual series $S_0(\alpha) = \frac{\pi}{\alpha} \sum_{n \in \mathbb{Z}} \operatorname{sech}(\pi^2 n/\alpha) > \frac{\pi}{\alpha}$ and $S_1(\alpha) = \frac{\pi}{\alpha} \sum_{n \in \mathbb{Z}} (-1)^n \operatorname{sech}(\pi^2 n/\alpha) < \frac{\pi}{\alpha}$ (an alternating strictly decreasing series in $|n|$). □

6 Computational Verification and Benchmark Tables

All four theorems were verified programmatically using `scripts/verify_all.py` and the automated test suite `tests/test_sdpi_bhattacharyya.py` (4/4 test modules passing). Key numerical benchmarks are summarized below and visualized in Figures 1 and 2.

- **Theorem 2.2 (480 random channels across $m \in \{2,3,4,6\}$, $n \in \{2,3,5,7\}$):** Exact splitting error $\|(I_m - G) - (\widetilde{\mathcal{L}}_{\mathrm{BC}} + \widetilde{\mathcal{L}}_{\mathrm{Var}})\|_\infty \le 4.72 \times 10^{-16}$; zero Weyl spectral violations ($\le 1.33 \times 10^{-15}$); zero trace-bound violations (0.0). On the lazy cycle W_{20}, $\|\widetilde{\mathcal{L}}_{\mathrm{Var}}(u)\|_\infty = 0.0$ and $\lambda_2(\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)) = 0.02447174 = \sin^2(\pi/20)$ identically (a 20.0× improvement over unnormalized $\mathcal{L}_{\mathrm{BC}}(u) = 0.00122359$). On $\mathbb{Z}_{12}$ with asymmetric four-point support $c = (0.4, 0.3, 0.2, 0.1)$ on $\{0,1,2,5\}$, the Fourier–Bhattacharyya inequality (7) holds across all 11 non-trivial characters with minimum strict margin 1.38×10^{-3}.
- **Theorem 3.1 (200 random binary channels + 7 canonical phase/saturation benchmarks):** Posterior-covariance identities (11)–(14) matched direct formulas to max errors 7.77×10^{-16} (H), 8.88×10^{-16} (H'), 2.13×10^{-14} (H''), and 3.87×10^{-15} (variance slack). Table 1 tests the following explicit distributions: BEC(0.4) ($p = (0.6, 0.4, 0), q = (0, 0.4, 0.6)$), $Z_{0.3}$ ($p = (1, 0), q = (0.7, 0.3)$), Gen-$Z$ ($p = (0.6, 0.4, 0), q = (0.42, 0.28, 0.30)$), non-proportional Rev-$Z$ ($p = (0.5, 0.2, 0.3), q = (0.75, 0.25, 0)$), BAC(0.85, 0.25) ($p = (0.85, 0.15), q = (0.25, 0.75)$), Two-Ratio ($p = (0.45, 0.40, 0.10, 0.05), q = (0.15, \frac{2}{15}, \frac{43}{90}, \frac{43}{180})$ with exact likelihood ratios 3 and $\frac{9}{43}$), and Three-Ratio ($p = (0.60, 0.30, 0.10), q = (0.15, 0.35, 0.50)$), confirming exact saturation on the three canonical classes and strict slack on the remaining three.
- **Theorem 4.1 (150 random multi-input channels + canonical benchmarks):** Zero violations of the sandwich (16) (0.0) and zero instances of multi-input priors exceeding binary reduction (0.0). On Z_ϵ at $\epsilon = 10^{-4}$, $\eta_{\chi^2}/\mathrm{LC} = 1.99990000 = 2 - \epsilon$. Table 2 compares all local and global contraction functionals across six canonical channels, where $\mathrm{SC}_4(0.30) := 0.70\, I_4 + 0.10(J_4 - I_4)$ denotes the 4-ary symmetric channel with total off-diagonal error probability 0.30.
- **Theorem 5.1 (Tensor powers and parity sech-lattice law):** On the asymmetric 3×3 channel K_{asym} with rows $(0.65, 0.25, 0.10)$, $(0.20, 0.55, 0.25)$, $(0.10, 0.30, 0.60)$ ($C_{\max} = 0.256543$, achieved by rows 1 and 3), product-prior invariance at $q = (0.2, 0.5, 0.3)$ held to 1.11×10^{-16} ($d = 2$) and 2.78×10^{-16} ($d = 3$), and the finite-d Chebyshev sandwich (18) held strictly for all $d = 1, \ldots, 7$. For $\mathrm{BSC}(0.10)^{\otimes d}$ at $d = 500$ (even) and $d = 501$ (odd), the scaled deficits 1.491056 and 1.362758 matched the exact parity limits $S_0(\alpha) = 1.494572$ and $S_1(\alpha) = 1.366463$ within 0.24% and 0.27%, respectively (Figure 2(b)).

7 Discussion and Open Directions

We have established a unified geometric and spectral theory of Pearson χ^2 strong data-processing inequalities linking local multi-input contraction spectra, binary harmonic-risk calculus, closed-

Table 1: Verification of the four-regime phase classification and three-class Bhattacharyya saturation rigidity (Theorem 3.1).

Channel (p, q)	Regime	y^*	$\eta_{\chi^2}(p, q)$	$1 - \mathrm{BC}^2$	Slack	$\Delta_0,\ \Delta_1$
BEC(0.4)	Flat BEC	any (0.5)	0.600000	0.840000	2.40×10^{-1}	0.00, 0.00
$Z_{0.3}$	Phase I (Left Z)	0^+	0.300000	0.300000	0 ($< 10^{-15}$)	-0.21, $+0.43$
Gen-Z (3 outputs)	Phase I (Left Z)	0^+	0.300000	0.300000	0 ($< 10^{-15}$)	-0.21, $+0.43$
Non-prop. Rev-Z	Phase II (Right Z)	1^-	0.300000	0.301139	1.14×10^{-3}	$+0.44$, -0.21
BAC$(0.85, 0.25)$	Phase III (Interior)	0.451941	0.365767	0.365767	0 ($< 10^{-15}$)	$+2.82$, $+1.92$
Two-Ratio (4 outputs)	Phase III (Interior)	0.442091	0.329862	0.329862	0 (0.0)	$+2.52$, $+1.58$
Three-Ratio (3 outputs)	Phase III (Interior)	0.487626	0.270356	0.281500	1.11×10^{-2}	$+1.95$, $+1.68$

Table 2: Comparison of uniform-prior spectral bounds (Theorem 2.2) and global contraction bounds (Theorem 4.1) across canonical channels.

Channel K	$\eta_{\chi^2}(u, K)$	$1 - \lambda_2(\widetilde{\mathcal{L}}_{\mathrm{BC}}(u))$	$\mathrm{LC}_{\max}(K)$	$\eta_{\chi^2}(K)$	$1 - \mathrm{BC}^2_{\min}$	$\eta_{\mathrm{TV}}(K)$
BSC(0.45)	0.010000	0.010000	0.010000	0.010000	0.010000	0.100000
BSC(0.20)	0.360000	0.360000	0.360000	0.360000	0.360000	0.600000
BEC(0.80)	0.200000	0.200000	0.200000	0.200000	0.360000	0.200000
$Z(0.05)$	0.025641	0.025641	0.025641	0.050000	0.050000	0.050000
$\mathrm{SC}_4(0.30)$	0.360000	0.468340	0.450000	0.450000	0.468340	0.600000
W_8 (lazy cycle)	0.853553	0.853553	1.000000	1.000000	1.000000	1.000000

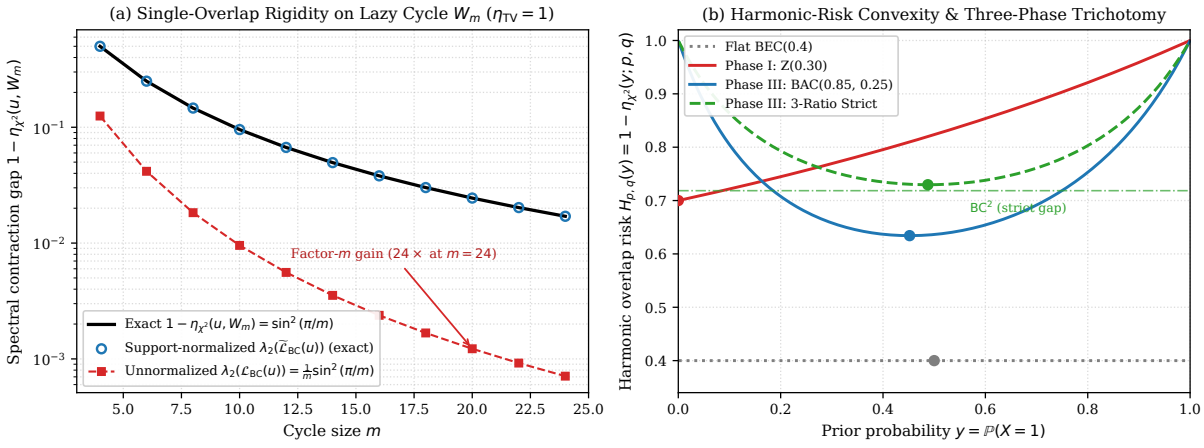


Figure 1: **(a)** Spectral contraction gap $1 - \eta_{\chi^2}(u, W_m) = \sin^2(\pi/m)$ on the lazy cycle walk W_m ($\eta_{\mathrm{TV}} = 1$): support-normalized $\widetilde{\mathcal{L}}_{\mathrm{BC}}(u)$ is exact for every m (zero volatility), whereas unnormalized $\mathcal{L}_{\mathrm{BC}}(u)$ loses a factor of m. **(b)** Harmonic overlap risk $H_{p,q}(y)$ illustrating the flat BEC regime, left-boundary Z-regime (Phase I), interior two-ratio saturation (BAC), and strict Bhattacharyya slack on a 3-ratio channel.

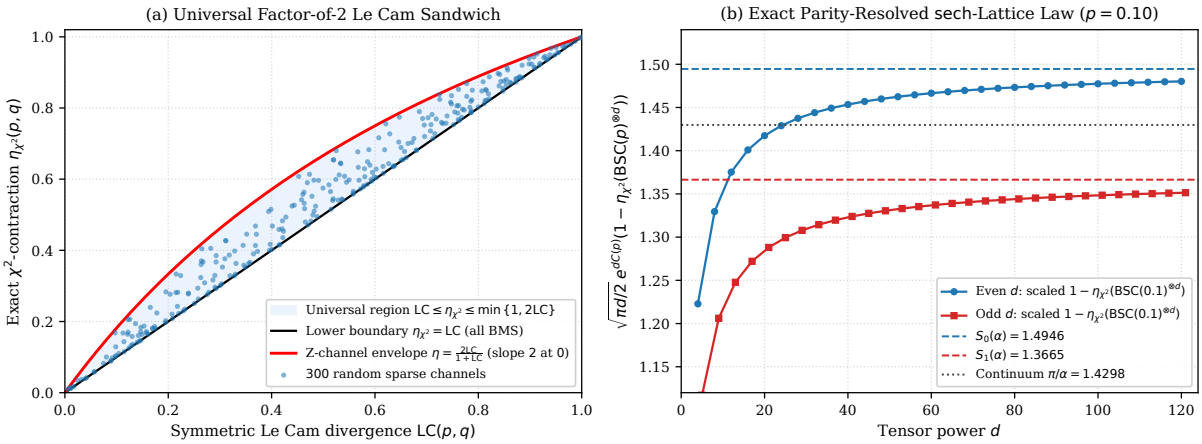


Figure 2: **(a)** Universal factor-of-2 Le Cam sandwich across 300 random sparse channels: every channel lies between $\eta_{\chi^2} = \text{LC}$ (exact on all involution-symmetric binary channels) and the sharp Z-channel envelope $\eta = 2\text{LC}/(1 + \text{LC})$ with slope 2 at the origin. **(b)** Exact parity-resolved sech-lattice prefactor convergence for $\text{BSC}(0.10)^{\otimes d}$: even and odd tensor powers converge to distinct limits $S_0(\alpha) = 1.4946 > \pi/\alpha > S_1(\alpha) = 1.3665$.

form Le Cam approximations, and tensor-power Chernoff amplification. Two natural open directions emerge directly from our framework:

1. **Quantum Petz–Bhattacharyya Laplacian splitting.** Does the operator splitting $I_m - G(q, K) = \widetilde{\mathcal{L}}_{\text{BC}}(q) + \widetilde{\mathcal{L}}_{\text{Var}}(q)$ admit a non-commutative generalization for quantum CPTP maps Φ using Uhlmann fidelity $F(\Phi(\rho_i), \Phi(\rho_k))^2$ and the Petz recovery channel?
2. **Parity-resolved prefactors for asymmetric channels.** Theorem 5.1(c) reveals sharp even/odd sech-lattice oscillations on $\text{BSC}(p)^{\otimes d}$; extending the Bahadur–Rao prefactor analysis to non-lattice and asymmetric channels in Phase III via the tilted harmonic-risk saddle point remains an appealing problem.

Reproducibility and Open-Source Code

All numerical experiments, tables, figures, and automated verification tests (4/4 pytest modules) are openly archived under the MIT license at https://github.com/lluiseriksson/sdpi-bhattacharyya-laplacian.

References

[1] R. Ahlswede and P. Gács, Spreading of sets in product spaces and hypercontraction of the Markov operator, *Ann. Probab.* **4** (1976), 925–939.

[2] V. Anantharam, A. Gohari, S. Kamath, and C. Nair, On maximal correlation, hypercontractivity, and the data processing inequality studied by Erkip and Cover, arXiv:1304.6133 (2013).

[3] H. Chernoff, A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations, *Ann. Math. Statist.* **23** (1952), 493–507.

[4] J. E. Cohen, J. H. B. Kemperman, and G. Zbăganu, *Comparisons of Stochastic Matrices with Applications in Information Theory, Statistics, Economics and Population Sciences*, Birkhäuser, Boston, 1998.

[5] R. L. Dobrushin, Central limit theorem for nonstationary Markov chains. I, *Theory Probab. Appl.* **1** (1956), 65–80.

[6] A. Makur and Y. Polyanskiy, Comparison of channels: Criteria for domination by a symmetric channel, *IEEE Trans. Inform. Theory* **64** (2018), 5704–5725.

[7] O. Ordentlich and Y. Polyanskiy, Strong data processing inequality for f-divergences: Reduction to binary inputs, *IEEE Trans. Inform. Theory* **68** (2022), 7726–7737.

[8] Y. Polyanskiy and Y. Wu, Strong data-processing inequalities for channels and Bayesian networks, in *Convexity and Concentration*, IMA Vol. Math. Appl. **161**, Springer, New York, 2017, pp. 211–249.

[9] M. Raginsky, Strong data processing inequalities and Φ-Sobolev inequalities for discrete channels, *IEEE Trans. Inform. Theory* **62** (2016), 3355–3389.