Discrete Jensen inequalities for alcoved lattice counts, planar flows and transportation tables

Scott Schubr
Luminary Oddities

November 2026

Abstract

Leake and Mohammadi Yekta asked whether lattice-point counts of unimodular fibers satisfy a discrete Jensen inequality, $\log N(\mathbb{E}a) \geq \mathbb{E}\log N(a)$, for every distribution of right-hand sides with integral mean. Two-point midpoint inequalities give log-concavity along lattice lines, while a general distribution with integral mean may have no pair of support points with an integral midpoint. We prove the full Jensen inequality for type-A alcoved polytopes, $x_u - x_v \leq h_{uv}$, in every dimension, with every directed difference bound allowed to vary simultaneously and for every integrable distribution of bounds. This settles the type-A_4 branch of the dimension-four reduction of the question, and covers all type-A systems in any dimension. The proof balances finitely many quantiles with the entropy midpoint inequality of Gozlan, Roberto, Samson and Tetali, lifts the balance to several coordinates by a conditional quantile coupling, and rounds all coordinates with one common offset, which preserves every integral directed difference. The same argument proves Jensen inequalities for partition functions of integer height functions and of planar integer flows with convex edge energies, with conservation data and edge capacities varying together. For the second branch, 3×3 transportation tables, we prove a fixed-order obstruction for coordinatewise rounding and then compare whole fiber counts. We prove the arbitrary-mixture Jensen inequality on explicit families of that branch: for jointly varying uniform cell caps up to six; for one, two or three fixed cells with independently unbounded caps when the remaining common cap is at most three, two or one; and for a fixed 2×2 rectangle of unbounded widths with binary widths elsewhere. The proofs combine affine upper supports, clipping of unbounded intervals and exact integer certificates. For three colors with integral mean cell caps at most two we prove a coefficientwise comparison at every fixed aggregate, which gives the equal-weight three-input product inequality $N(b_1)N(b_2)N(b_3) \leq N(b)^3$ for margins and caps b_i with mean b. Finally, for fixed numbers of rows and colors with unit caps, at a fixed aggregate the balanced margins uniquely maximize the count once every row cut is crossed by a pair of rows that are both flexible in sufficiently many columns; for three rows and three colors, sufficiently many flexible columns of any types already suffice.

Keywords. Alcoved polytopes, unimodular fibers, discrete Jensen inequality, entropy, integer flows.
Mathematics Subject Classification (2020). 52B20, 60E15, 05C21.

1 Introduction

Let M be a unimodular integer matrix. Leake and Mohammadi Yekta [5, Question 1.7] ask whether the number of nonnegative integral solutions to $My = a$ satisfies a discrete Jensen inequality for every distribution of a with integral mean. This question extends an exact log-concavity question of Barvinok for contingency tables [1, (2.3)]. Leake and Mohammadi Yekta prove the integral

two-point midpoint inequality for unimodular systems by comparing centered fibers of boxes [5, Theorems 1.4 and 3.1]. That proof gives log-concavity along lattice lines. An arbitrary distribution with integral mean may involve no pair of distinct support points with integral midpoint.

The earlier paper [7] proves a stronger power inequality when the central fiber has dimension at most three. It also reduces the four-dimensional case to two normal systems. The first is type A_4, with arbitrary directed differences among five vertex coordinates. The second comes from cell-bounded 3×3 transportation tables. We answer the logarithmic Jensen question for the first system. The same proof works in every dimension and for all the bond bounds at once. A second consequence is a partition-function inequality for planar integer flows, where the edge capacities and divergences may vary together.

Our transport proof starts with the entropy midpoint theorem of Gozlan, Roberto, Samson and Tetali [2, Theorem 8]. Slomka [8, Theorem 1.4 and §3.1] proved the two-law lattice entropy inequality for more general complementing operations using conditional monotone coupling. Halikias, Klartag and Slomka [3, Theorems 1.2 and 1.3] likewise give two-input functional inequalities for complementary operations, including weighted midpoint operations. We use a repeated scalar balance to treat an arbitrary finite list of input laws and then use the same conditional coupling strategy. This gives an entropy inequality averaged over one common rounding offset for arbitrary real weights.

There is also a distinct marginal approach. Speyer [9, Theorem 5.1] proves that summing over coordinates preserves L-log-concavity for L-convex sets. This applies to alcoved counts with moving bounds adjacent to the anchor and fixed internal bounds [7, Theorem 1.8]. In the problem here an internal bound h_{uv} is itself a parameter in $x_u - x_v \le h_{uv}$. The corresponding joint system in the parameters and the lattice point need not be L-convex. Our proof rounds the lattice points after taking their weighted average at the integral mean of the parameters.

Sections 6–9 return to the transportation branch. In view of the fixed-order obstruction of Section 5, these sections compare whole fiber counts. Section 6 constructs affine upper supports to the logarithm of the count and certifies them exactly, for jointly varying uniform caps up to six and for up to three distinguished cells whose caps are unbounded. Section 7 lets an entire 2×2 rectangle have unbounded widths. Sections 6 and 7 thus prove the arbitrary-mixture inequality of Leake and Mohammadi Yekta on these families, and Section 6.7 extends it to nonplanar networks obtained by gluing planar and covered transportation blocks at cut vertices. Section 8 proves the equal-weight three-input product inequality for three colors with integral mean caps at most two, by a coefficientwise comparison at every fixed aggregate. Section 9 proves fixed-aggregate balancing results for multicolor counts with unit caps: exact results for two-row arrays and small cubic cores, and, for fixed numbers of rows and colors, unique maximization at balance once every row cut is crossed by a pair of rows that are both flexible in sufficiently many columns.

This is the second edition of this article. The second edition adds Sections 6–9 and retitles the article accordingly.

2 Alcoved lattice counts

Fix vertices $0, \dots, d$ and a finite labeled collection $\mathcal{A}$ of ordered pairs $e = (u, v)$. Labels allow multiple constraints on the same directed difference. For $h \in \mathbb{Z}^{\mathcal{A}}$ write

$$P(h) = \{x \in \mathbb{R}^{d+1} : x_0 = 0,\ x_u - x_v \le h_e \text{ for } e = (u, v) \in \mathcal{A}\}, \qquad N(h) = |P(h) \cap \mathbb{Z}^{d+1}|.$$

With all directed differences present this is the anchored type-A_d description of an alcoved polytope [4].

Theorem 1 (Discrete Jensen for directed difference bounds)**.** *Let A be an integrable $\mathbb{Z}^{\mathcal{A}}$-valued random vector and let $b = \mathbb{E}A \in \mathbb{Z}^{\mathcal{A}}$. If $N(A) > 0$ almost surely and $N(b)$ is finite, then the central count is positive, all counts at support atoms are finite, and*

$$\log N(b) \;\geq\; \mathbb{E}\log N(A).$$

The proof first develops the entropy inequality for a finite family. We then use it on uniform distributions over the endpoint lattice points. An exact-mean approximation gives the countable-support case.

3 Entropy under simultaneous rounding

For a finitely supported random variable X, write $\mathsf{H}(X) = -\sum_x \Pr(X = x)\log\Pr(X = x)$ for its Shannon entropy. For a probability law μ on $\mathbb{Z}$, its quantile at $u \in (0,1)$ is the least integer t with $\mu((-\infty, t]) \geq u$.

Lemma 2 (Balancing scalar quantiles)**.** *Let U be uniform on $(0,1)$, let $f_1, \ldots, f_m$ be finite-valued nondecreasing functions from $(0,1)$ to $\mathbb{Z}$, and put $S = \sum_i f_i(U)$. Then*

$$\sum_{r=0}^{m-1} \mathsf{H}\left(\left\lfloor \frac{S+r}{m} \right\rfloor\right) \;\geq\; \sum_{i=1}^{m} \mathsf{H}(f_i(U)).$$

Proof. For two such functions, [2, Theorem 8, equation (9)] says that replacing their values by the floor and ceiling of their mean can only increase the sum of their Shannon entropies. Indeed, the common uniform variable gives the monotone coupling in that theorem, and its relative entropy with respect to counting measure has the opposite sign.

Ignore the endpoints of the finite common interval partition on which all f_i are constant. Whenever a pair differs by at least two on an interval, replace the pair everywhere by the floor and ceiling of its mean. The functions remain nondecreasing and constant on the same partition. Each replacement preserves the pointwise sum and does not decrease total entropy. The sum of the squared values over the functions and the positive-length intervals is a nonnegative integer. Balancing a pair never increases that sum and strictly decreases it on every interval where the pair differs by at least two. Thus the procedure ends with all values at each point differing by at most one.

Sort the resulting functions pointwise using a fixed finite sequence of pair comparisons, such as a sorting network of adjacent comparisons. Each comparison is another permitted entropy operation: for two integers at distance at most one the floor and ceiling of the mean are their minimum and maximum. The ordering at every point is therefore attained without decreasing total entropy. If $S = mq + k$ with $0 \leq k < m$, the sorted list has $m - k$ entries q and k entries $q + 1$. Its r-th entry, indexed from zero, is $\lfloor (S + r)/m \rfloor$. □

For finitely supported laws $\mu_1, \ldots, \mu_m$ on $\mathbb{Z}^d$, construct a joint law of $X_1, \ldots, X_m$ by using a common independent uniform variable U_j for each coordinate j. At coordinate j, take $X_{i,j}$ to be the μ_i-conditional quantile at U_j given its own preceding coordinates. This is the coordinatewise Knothe construction used for two laws in [8, §3.1]. Each X_i has law μ_i.

Theorem 3 (Entropy of rounded weighted averages)**.** *Use this coupling for finitely supported laws $\mu_1, \ldots, \mu_m$ on $\mathbb{Z}^d$. Suppose $\lambda_i \geq 0$ and $\sum_i \lambda_i = 1$. Write $W = \sum_i \lambda_i X_i$, and for $0 \leq s < 1$ set $Y(s) = \lfloor W + s\mathbf{1} \rfloor$, using the same offset in each coordinate. Then*

$$\int_0^1 \mathsf{H}(Y(s))\, ds \;\geq\; \sum_i \lambda_i \mathsf{H}(X_i).$$

Here the entropy in the integrand is for a fixed offset s.

Proof. For equal weights set $Y_r = \lfloor W + (r/m)\mathbf{1} \rfloor$ for $0 \le r < m$. At coordinate j, condition on the full input prefix $B_j = (X_{i,\ell})_{i,\,\ell<j}$. Since this prefix depends only on the earlier uniforms, U_j remains independent of B_j. The conditional coordinates $X_{i,j}$ are nondecreasing functions of U_j, so Lemma 2 gives, after averaging over the prefix,

$$\sum_r \mathsf{H}(Y_{r,j} \mid B_j) \;\ge\; \sum_i \mathsf{H}(X_{i,j} \mid B_j).$$

The output prefix $Y_r^{<j}$ is determined by B_j. Conditioning on $Y_r^{<j}$ therefore yields entropy at least $\mathsf{H}(Y_{r,j} \mid B_j)$. Each input coordinate uses only its own prefix and U_j, so $\mathsf{H}(X_{i,j} \mid B_j) = \mathsf{H}(X_{i,j} \mid X_i^{<j})$. The entropy chain rule gives $\sum_r \mathsf{H}(Y_r) \ge \sum_i \mathsf{H}(X_i)$. Since mW is integral, $Y(s) = Y_r$ for $r/m \le s < (r+1)/m$. Dividing the last inequality by m proves the equal-weight case.

For rational weights, repeat each law according to its numerator over a common denominator. Repeated copies of any one law use the same uniforms and the same conditional quantiles, so they agree almost surely in every coordinate. Thus the equal-weight result applies to the weighted average in the stated coupling. For real weights, approximate them by rational weights in the simplex and retain the same input random variables. Their joint array has finite support. For almost every fixed offset, none of its finitely many limiting weighted coordinates lands at a rounding boundary. At these offsets the rounded laws eventually agree with the limiting rounded law. The rounded variables lie in one fixed finite box, so dominated convergence proves the stated entropy inequality. □

Proof of Theorem 1. When $P(h)$ is feasible, the directed graph of its constraints, with the arc $v \to u$ of length h_e, has no negative directed cycle. Each simple cycle length at b is the mean of nonnegative simple cycle lengths at the support atoms. A new source with zero-length arcs to all the vertices then gives integral shortest-path potentials at b. Subtracting the potential of vertex 0 gives a central lattice point. If any nonempty endpoint were unbounded, the common rational recession cone would contain a nonzero integral ray. Adding its nonnegative integer multiples to the central lattice point would contradict the finiteness of $N(b)$.

First let the bound distribution be finite, with positive weights λ_i at h_i. Use uniform laws on the endpoint lattice fibers in Theorem 3. For each realization, $W = \sum_i \lambda_i X_i$ belongs to $P(b)$ by linearity, with its anchor coordinate zero. For $e = (u, v)$, the integral upper bound gives

$$W_u \le W_v + b_e \quad\Longrightarrow\quad \lfloor W_u + s \rfloor \le \lfloor W_v + s \rfloor + b_e.$$

Thus every $Y(s)$ belongs to the central fiber. Since a law supported on at most $N(b)$ points has entropy at most $\log N(b)$, Theorem 3 gives the result.

For countable support, write the distinct atoms as h_i with positive masses p_i. Choose an affine basis $c_0, \dots, c_q$ of their affine span from among the atoms. Truncate the law after n atoms for n large enough to retain the basis. The omitted joint mass and moment $\sum_{i>n} p_i(1, h_i)$ tend to zero by integrability. Their unique coordinates $\delta_{n,k}$ in the basis $(1, c_k)$ also tend to zero. Add $\delta_{n,k}$ to the retained mass of c_k. These masses stay positive for large n and the finite corrected law has mass one and mean b. Finite Jensen applies. The correction to the expectation of $\log N$ tends to zero, while the uncorrected partial sums increase because every positive lattice count is at least one. This establishes the inequality and the finiteness of its right side. □

Corollary 4 (Graphic kernel normals)**.** *Suppose a family of unimodular fibers has a lattice parameterization in which each fiber is given by fixed signed type-A roots with integer-affine right-hand*

sides in the demand. Its counts satisfy the discrete Jensen inequality for every integrable demand law with an integral mean and finite central count, provided the endpoint counts are positive almost surely. This includes the type-A_4 branch of [7, Proposition 9.1].

Proof. After adjoining the anchored vertex, every signed root is a directed difference. Keep the raw inequalities as labeled constraints, including parallel ones that share a root. Their right-hand sides commute with averaging. Any zero-normal constraints also hold at the integral mean by affinity. Apply Theorem 1 in the lattice coordinates. The fixed kernel parameterization in the active face used in [7, §5 and Proposition 9.1] has these properties in the type-A_4 branch. □

4 Partition functions of heights and planar flows

Let E_{en} be a fixed finite list of directed bonds on the vertex set. For each bond $e = (u, v)$ let $\phi_e : \mathbb{Z} \to \mathbb{R}$ have nonnegative second differences, so its continuous piecewise-linear interpolation $\widetilde{\phi}_e$ is convex. Define

$$\mathcal{E}(x) = \sum_{e=(u,v)\in E_{\mathrm{en}}} \phi_e(x_u - x_v), \qquad Z(h) = \sum_{x\in P(h)\cap\mathbb{Z}^{d+1}} \exp[-\mathcal{E}(x)].$$

Theorem 5 (Partition functions for integer heights)**.** *Let h_i be a finite list of bound vectors with positive finite counts. Suppose their weighted mean $b = \sum_i \lambda_i h_i$ is integral, with positive weights summing to one. Then $Z(b) \geq \prod_i Z(h_i)^{\lambda_i}$.*

Proof. The recession cone is fixed across fibers and the endpoint fibers are nonempty and bounded, so the central fiber is finite. Use the endpoint Gibbs laws $e^{-\mathcal{E}}/Z(h_i)$ in Theorem 3. The proof of Theorem 1 shows that every $Y(s)$ is central. For a fixed input realization, the rounded bond $Y_u(s) - Y_v(s)$ is the floor or ceiling of $W_u - W_v$. For $t = W_u - W_v$, subtract $\lfloor W_v + s \rfloor$ inside $\lfloor W_u + s \rfloor$. The difference is then $\lfloor t + \{W_v + s\} \rfloor$, where $\{\,\cdot\,\}$ denotes the fractional part. Also $\int_0^1 \lfloor t + s \rfloor \, ds = t$ for each real t. It follows that the rounded bond has mean $W_u - W_v$ over offsets. The interpolation $\widetilde{\phi}_e$ is affine between those two neighboring integers, so

$$\int_0^1 \phi_e(Y_u(s) - Y_v(s))\, ds = \widetilde{\phi}_e(W_u - W_v) \leq \sum_i \lambda_i \phi_e(X_{i,u} - X_{i,v}).$$

Summing over the bonds and taking expectations gives the energy comparison. On any finite fiber the relative entropy from its Gibbs law gives $\mathsf{H}(X) - \mathbb{E}\mathcal{E}(X) \leq \log Z$, with equality for that Gibbs law. Apply this inequality to each $Y(s)$ at the central fiber and integrate. Theorem 3 and the energy comparison give the conclusion. □

For example, impose $|x_u - x_v| \leq 1$ along every edge of a finite connected graph and take $\phi_e(k) = -|k| \log t$ with $0 < t \leq 1$. This is the integer Lipschitz height weight $t^{\#\{e : x_u \neq x_v\}}$. Prescribed boundary values are imposed by two opposite constraints to the anchor. Theorem 5 applies to any finite mixture of boundary values and bond limits whose averages are integral. OpenAI [6, Theorem 3.3 and Lemma 3.4] use the two-input midpoint energy comparison for this height model.

Corollary 6 (Planar integer flows)**.** *Let a fixed connected plane graph have oriented edges. For an integer divergence vector a of total sum zero and integer edge bounds $\ell \leq u$, define*

$$Z(a, \ell, u) = \sum_{\substack{x\in\mathbb{Z}^E \\ \operatorname{div} x = a,\ \ell\leq x\leq u}} \exp\left[-\sum_{e\in E} \phi_e(x_e)\right],$$

with fixed discrete convex edge energies ϕ_e. For every finite mixture of feasible data (a_i, ℓ_i, u_i) with an integral mean $(\bar{a}, \bar{\ell}, \bar{u})$,

$$Z(\bar{a}, \bar{\ell}, \bar{u}) \;\geq\; \prod_i Z(a_i, \ell_i, u_i)^{\lambda_i}.$$

In particular, the lattice counts satisfy this inequality by taking zero energy. The count statement holds also for countable integrable laws on bounded data fibers. For nonnegative flows without upper capacities, it holds for finite mixtures of finite fibers.

Proof. Choose a spanning tree. There is an integer-linear particular flow $p(a)$ supported on the tree, since each tree-edge flow is the signed sum of the divergences on one side of its cut. Every circulation has the form

$$x_e - p_e(a) = z_{f_+(e)} - z_{f_-(e)},$$

where $f_+(e)$ and $f_-(e)$ are the faces on the two sides of the oriented edge, in the order determined by the signed face boundaries. Anchor the outer face at zero. The face-boundary map takes values in the circulation space and is injective after anchoring because the dual graph is connected. Euler's formula gives equal dimensions for the two real spaces. Integral circulations correspond to integral potentials by propagating along a dual spanning tree. A bridge has the same face on both sides and contributes a zero normal.

In these coordinates the lower and upper edge bounds give directed differences with integer-affine right-hand sides. This gives the count results by Theorem 1. For the weighted result, use the conditional quantile coupling of endpoint Gibbs laws in the face coordinates. Because $p(\bar{a}) = \sum_i \lambda_i p(a_i)$ is integral, common-offset rounding of the face potentials makes each central edge flow a floor or ceiling of the mean input edge flow and preserves that mean over offsets. Therefore the energy and entropy argument of Theorem 5 applies with $\phi_e(x_e)$ on each full edge flow. If upper capacities are absent and the finitely many nonnegative endpoint and central flow fibers are finite, add a common cap above all their feasible lattice coordinates. This leaves their counts unchanged. □

5 A fixed-order obstruction for 3×3 tables

The second normal system in [7, Proposition 9.1] consists of the nine cell functionals for a 3×3 table in the four free corner coordinates $x_{11}, x_{12}, x_{21}, x_{22}$. The following example concerns rounding schemes that choose a floor or ceiling in each coordinate while preserving the two inputs' coordinate sum. Call such a scheme *causal in a fixed coordinate order* if a rounding decision can use the endpoint data and the already observed input coordinates, including the current one.

Proposition 7. *There is a single pair of 3×3 endpoint data with integral midpoint such that, for every fixed coordinate order, no causal rounding scheme can produce two feasible midpoint tables with probability one for each feasible input pair.*

Proof. At the first endpoint all margins are zero, so the only table is the zero table. At the second endpoint every margin is two. At both endpoints take lower bound zero and upper bound two on all cells. In particular the second endpoint contains

$$T^{(0)} = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}, \qquad T^{(1)} = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}.$$

At the midpoint all row and column margins are one. First suppose the last observed coordinate is x_{22}. The two input pairs agree at each preceding coordinate. For $T^{(0)}$, both output last coordinates are zero. Each output has $x_{33} = \sum_{i,j\le 2} x_{ij} - 1 \ge 0$, so the two output corner sums, which add to two, both equal one. The up-rounds of x_{12} and x_{21} must go to different outputs. For $T^{(1)}$, the combined second corner-row sum is two and so is the combined second corner-column sum. Each such sum is at most one in each output, hence both equal one. The up-rounds of x_{12} and x_{21} must now go to the same output. The earlier prefix cannot distinguish these input pairs. The same proof applies to randomized decisions required to succeed almost surely because the events demanded at that prefix are disjoint. Finally, swapping the first two rows or columns preserves both endpoint data and places any prescribed last corner at x_{22}. □

6 Discrete Jensen for capped three-by-three transportation tables

This section proves the arbitrary-mixture discrete Jensen inequality for counts of 3×3 integer transportation tables in four cap families, with caps and margins varying jointly: a common cap at most six; or a common cap at most three, two or one on the ordinary cells together with one, two or three fixed distinguished cells of independently unbounded caps. Clipping and signed affine upper supports transfer exact finite certificates to the unbounded caps, and the comparison extends to networks glued from planar and covered transportation blocks at cut vertices.

6.1 The comparisons

For a nonnegative integer 3×3 cap array u and nonnegative integer triples r, s of equal total, let $N(u, r, s)$ count integer tables in the box $[0, u]$ with row sums r and column sums s.

Theorem 8. *Let $X = (u, r, s)$ be a random transportation parameter with integrable data and positive counts almost surely. Suppose its mean is integral. Assume that its cap arrays belong to one of the following families.*

1. *Every cell has the same cap $c \in \{0, \dots, 6\}$.*
2. *Eight cells have the same cap $c \in \{0, \dots, 3\}$ and the other cell has a finite nonnegative integer cap d of arbitrary size. The distinguished location is fixed across the mixture.*
3. *Seven cells have the same cap $c \in \{0, 1, 2\}$. The other two cells are fixed across the mixture and each has a finite nonnegative integer cap of arbitrary size.*
4. *Six cells share a cap $c \in \{0, 1\}$. The other three cells are fixed across the mixture and each has a finite nonnegative integer cap of arbitrary size.*

Then $\log N(\mathbb{E}X) \ge \mathbb{E} \log N(X)$.

The central parameter is feasible. Average a real table chosen from each input fiber and use the integrality of the bipartite flow constraints at the integral mean. The family in part 1 is finite. In parts 2–4 the bounded entries determine the distinguished entries at fixed margins, so counts are bounded on each cap family.

Leake and Mohammadi Yekta prove the integral two-point midpoint comparison for unimodular fibers and ask whether this extends to arbitrary mixtures [5, Theorem 3.1 and Question 1.7]. Introducing slack variables y_{ij} with $x_{ij} + y_{ij} = u_{ij}$ gives a totally unimodular system. Expand minors along slack unit columns and the remaining cell-bound rows to reduce to minors of the

bipartite incidence matrix. The cographic reduction in dimension four identifies cell-bounded 3×3 transportation as the branch alongside the proven alcoved case [7, Proposition 9.1]. The alcoved comparison of Sections 2–4 covers planar networks. Sections 8 and 9 below give exact three-input and fixed-aggregate comparisons on the transportation branch. Theorem 8 allows arbitrary mixtures and compares entire fibers in its four cap families.

6.2 Uniform caps and affine upper supports

Use six independent coordinates

$$p = (c, r_1, r_2, r_3, s_1, s_2), \qquad s_3 = r_1 + r_2 + r_3 - s_1 - s_2.$$

For every feasible center b, we will give a vector $\theta \in \mathbb{R}^6$ satisfying

$$(p - b) \cdot \theta \geq \log\bigl(N(p)/N(b)\bigr) \qquad \text{for every feasible } p. \tag{1}$$

Taking expectations at mean b proves the theorem. Conversely, by finite-dimensional linear programming duality, an inequality of this form exists whenever the arbitrary-mixture comparison holds at b. The primal problem maximizes $\sum_p \alpha_p \log N(p)$ under $\alpha_p \geq 0$, $\sum_p \alpha_p = 1$, and $\sum_p \alpha_p p = b$. An affine dual majorant has the form in (1) if its optimum equals $\log N(b)$. Only sufficiency is needed for the proof.

Row and column permutations, transposition, and simultaneous complementation $x_{ij} \mapsto c - x_{ij}$ preserve the counts and act linearly on p. Complementation sends r_i, s_j to $3c - r_i, 3c - s_j$, which is linear even when c varies. After dropping s_3 these maps are still invertible, because the dropped coordinate is recovered by the displayed total constraint. We sort each of the two margin triples, then retain one center under transposition and simultaneous complementation followed by sorting. An upper support transports to every center in its orbit.

At a sorted center, impose decreasing row coefficients and decreasing column coefficients of θ, with the coefficient of s_3 set to zero. The four corresponding homogeneous inequalities are included in the certificate. For an arbitrary input p, sorting its margins in increasing order minimizes the dot product with these coefficients by rearrangement. This operation leaves the count invariant. Therefore it suffices to verify (1) against the sorted inputs for all admissible caps. The verifier generates these inputs and their counts by enumerating all 3×3 tables at each cap and retaining the tables with the two already sorted margins.

6.3 Exact verification of the uniform family

For each center the certificate lists six numbered linear constraints. Each one has integer direction d_i and positive rational right-hand ratio q_i, indicating $d_i \cdot \theta = \log q_i$. The possible rows are the inequalities (1), the four order constraints, and artificial coordinate bounds with right-hand ratio 10^{-8} and directions $\pm e_j$. The coordinate bounds give a bounded search polytope when finding an intersection. Replay uses any bounds listed among the six rows to define the candidate and checks only the count and order inequalities afterward.

Let M have the six selected directions as rows. The verifier checks invertibility and computes M^{-1} over the rationals. The implicit slope is $\theta = M^{-1}(\log q_1, \ldots, \log q_6)^{\mathsf{T}}$. Take a positive integer D clearing all denominators of M^{-1}. For any inequality with direction d and right-hand ratio q, verification reduces exactly to

$$\prod_{i=1}^{6} q_i^{D(d^{\mathsf{T}} M^{-1})_i} \;\geq\; q^D.$$

Every exponent is integral. Cross multiplication therefore reduces each comparison to two nonnegative integers, with no floating-point logarithms in the verifier. All labeled supports are thereby derived from the representative centers.

Cap	Sorted input profiles	Center representatives
0	1	1
1	34	12
2	325	100
3	1556	426
4	5 135	1 373
5	13 466	3 503
6	30 219	7 803
Total	50 736	13 218

Every representative is tested against the same combined input set. The reported certificate is the run of this integer-power verifier, written with the Python standard library only, on all 13,218 representatives; the program and its list of six-row witnesses are included in the source package. An independently implemented NumPy verifier replays the same witnesses: it enumerates by convolution of row counts and uses provable rational log enclosures, with exact integer powers for comparisons close to equality. All 13,218 supports also passed this independent replay. An optimizer was used to find candidate active rows. The identity of the rows and the complete integer verification are the only computational inputs to the proof.

6.4 An arbitrary distinguished cap

Put the distinguished cell at $(3, 3)$ and write $w = x_{31} + x_{32}$ for the sum of the two ordinary entries in the third row. For the eight ordinary entries, keep the four margins r_1, r_2, s_1, s_2. The remaining entry is determined by $x_{33} = r_3 - w$, and the last column margin then follows from the equality of totals. Since $0 \le w \le 2c$, the count equals $F(p)$, the number of ordinary configurations with the retained margins and $A \le w \le B$, for

$$p = (c, A, B, r_1, r_2, s_1, s_2), \qquad A = \max(r_3 - d, 0), \qquad B = \min(r_3, 2c).$$

A positive count has $0 \le A \le B \le 2c$. Conversely, each feasible clipped profile p embeds in a transportation profile by setting $d = B - A$, $r_3 = B$, and $s_3 = r_1 + r_2 + B - s_1 - s_2$. The map is integer-linear, and all the counts agree.

At the original mean data, the clipped endpoints obey $A_* \le \mathbb{E}A$ and $B_* \ge \mathbb{E}B$, by convexity of the maximum and concavity of the minimum. The other coordinates commute with taking means. Suppose we have an affine upper support at the clipped mean p_*, with slope $\theta_A \le 0$ on the lower bound and $\theta_B \ge 0$ on the upper bound. Then

$$\mathbb{E}\log F(p) - \log F(p_*) \le \theta \cdot (\mathbb{E}p - p_*) \le 0.$$

Thus such supports prove Jensen for the original parameters, even when clipping and averaging give different endpoints.

The sign requirement also follows abstractly from the finite Jensen comparison. Use the trim coordinates $U = A$ and $V = 2c - B$. Increasing either trim decreases the count at fixed other data. If a support has a positive slope in one trim, set it to zero. For inputs whose trim is at most the

center's, the upper bound increases. For an input whose trim exceeds the center's, reduce that trim to the center value and apply the original upper support at the relaxed input. The count is at least the original count and the resulting right side is the revised support. Repeat for the other trim. The finite signed-support comparison is therefore equivalent to arbitrary-mixture Jensen for the distinguished cap family, using the integer-linear embedding for the converse.

6.4.1 Exact verification of the clipped family

We enumerate feasible clipped profiles from the eight ordinary entries, reducing only the first two row margins and the first two column margins to sorted order. Decreasing slope pairs suffice to test sorted inputs by the rearrangement argument of Section 6.2. Complementing the eight entries and restoring these sorted pairs sends

$$(c, A, B, r_1, r_2, s_1, s_2) \mapsto (c, 2c - B, 2c - A, 3c - r_2, 3c - r_1, 3c - s_2, 3c - s_1).$$

This is a linear involution preserving counts and exchanging the two endpoint sign requirements. We choose the least element of each resulting orbit.

Each representative supplies seven independent active rows, including possible order, endpoint-sign and coordinate-box rows. The rational inverse determines the support in terms of logarithms of positive rational ratios. The replay encloses the logarithms by rationals before using fixed-point integer comparisons, and falls back to exact powered products for comparisons too close to resolve in the enclosure. It verifies both sign inequalities exactly. The certificate passes every representative against all sorted profiles.

Ordinary cap	Sorted input profiles	Center representatives
0	1	1
1	214	110
2	4 101	2 070
3	29 698	14 887
Total	34 014	17 068

An independent additive interval enumeration reproduced every count. Twenty-seven supports also passed a separate replay by direct integer powers against all sorted profiles. The signed certificate and the transfer above prove part 2 of Theorem 8.

6.5 Two independent distinguished caps

There are two geometries up to row and column permutations and transpose. In either one, clipping gives two interval constraints on ordinary sums in $[0, 2c]$. The two lower endpoints are convex functions of the original parameters and the two upper endpoints are concave. Therefore the signed-support argument of Section 6.4 applies to both constraints together.

For cells on a shared line choose $(3, 1), (3, 2)$ and retain r_1, r_2, s_3. Let $w_i = x_{1i} + x_{2i}$ for $i = 1, 2$. The eliminated entries are $x_{3i} = s_i - w_i$, so

$$A_i = \max(s_i - d_i, 0), \qquad B_i = \min(s_i, 2c)$$

are the two endpoints on w_i. The remaining row margin follows from the equality of totals. Conversely, a feasible clipped profile lifts by taking $s_i = B_i$, $d_i = B_i - A_i$, and $r_3 = B_1 + B_2 + s_3 - r_1 - r_2$.

For cells in distinct rows and columns choose $(1,1),(2,2)$. Let w_i and v_i be the ordinary row and column sums at index i for $i=1,2$. Retain r_3,δ_1,δ_2 with $\delta_i=r_i-s_i=w_i-v_i$. The eliminated entries are $x_{ii}=r_i-w_i$, so the endpoints are

$$A_i=\max(r_i-d_i,0),\qquad B_i=\min(r_i,2c).$$

The clipped profile lifts by taking $r_i=B_i$, $s_i=B_i-\delta_i$, $d_i=B_i-A_i$, and $s_3=r_3+\delta_1+\delta_2$. In both cases these lifts are integer-linear and the counts agree.

The joint finite certificates use eight coordinates, consisting of c, the four endpoints, and the three retained margins or differences. They require a nonpositive slope on every A_i and a nonnegative slope on every B_i.

In the shared-line geometry sort r_1,r_2 and impose the corresponding order on their slopes. There is no ordinary margin pair to sort in the disjoint geometry. Both use an exchange of the two distinguished positions and complementation of all ordinary entries, with the ordinary pair resorted when used. Exchange swaps the two intervals and, in the disjoint geometry, also swaps δ_1,δ_2. Complementation sends (A_i,B_i) to $(2c-B_i,2c-A_i)$. On the retained coordinates its images are $(3c-r_2,3c-r_1,3c-s_3)$ in the shared-line case and $(3c-r_3,-\delta_1,-\delta_2)$ in the disjoint case. These commuting involutions preserve counts and all endpoint sign requirements. Each center is represented by the least of its four images.

The exact replay lists eight active rows for each representative and verifies the upper support by the same rational and integer-power comparisons. Each ordinary table is counted for every interval pair containing its two ordinary sums. With c varying from zero to two, the full coverage is

Geometry	Sorted input profiles	Center representatives
Shared line	17 389	4 676
Distinct rows and columns	30 796	7 806

An independently implemented direct replay checked 23 centers against all inputs for each geometry. Thus all signed supports hold, and clipping and averaging prove part 3 of Theorem 8.

6.6 Three distinguished caps on a forest

Three chosen cells are three edges of the bipartite graph from rows to columns. They cannot form a cycle. The four possible forests up to the row, column and transpose symmetries are a three-star, a path of length three, a wedge with a disjoint edge, and a three-edge matching. The four placement orbits have sizes $6,36,36,6$, respectively.

Root each tree, including the singleton components. For a distinguished edge e let S_e be its child subtree and $\omega_e=1$ if the row-to-column orientation points out of it, or $\omega_e=-1$ otherwise. If τ_e is the signed ordinary outflow across this cut, the prescribed divergence sum is

$$a(S_e)=\tau_e+\omega_e x_e.$$

On the ordinary box, $\alpha_e(c)\le\tau_e\le\alpha_e(c)+m_e c$, where $\alpha_e(c)$ is minus the sum of incoming ordinary capacities and m_e is the number of ordinary cut edges. Write $w_e=\tau_e-\alpha_e(c)$. The constraint $0\le x_e\le d_e$ is an interval on w_e. Before clipping its endpoints are

$$\begin{cases}[\,a(S_e)-\alpha_e(c)-d_e,\ a(S_e)-\alpha_e(c)\,], & \omega_e=1,\\ [\,a(S_e)-\alpha_e(c),\ a(S_e)-\alpha_e(c)+d_e\,], & \omega_e=-1.\end{cases}$$

Clamp that interval to $[0, m_e c]$. Retain the aggregate prescribed divergence of two of the three tree components. The third follows from total divergence zero. These totals constrain only the ordinary entries. The subtree equations then reconstruct all distinguished entries uniquely, since subtracting the equations for immediate child subtrees recovers the original vertex divergences.

Conversely, from a feasible clipped interval $[A_e, B_e]$ choose $d_e = B_e - A_e$. Choose $a(S_e) = \alpha_e(c) + B_e$ for $\omega_e = 1$, or $a(S_e) = \alpha_e(c) + A_e$ for $\omega_e = -1$. The resulting distinguished entry is $B_e - w_e$ or $w_e - A_e$, respectively, and the vertex divergences are integer-linear in these data and in the component totals. This reverse lift preserves the counts.

As before, the lower clipped endpoints are convex and the upper endpoints are concave. At an integral original mean, signed supports at the clipped mean prove the desired inequality. For each placement the nine-coordinate certificate uses c, the six endpoints and two component divergences. Complementing the ordinary entries sends w_e to $m_e c - w_e$ and complements each component divergence about its all-c value. The map is a linear involution preserving counts and the endpoint slope constraints. We choose one center per orbit.

The verifier enumerates the six ordinary entries and counts each assignment in all interval triples containing its shifted cut values. Exact replay of the nine-row supports, including all six endpoint signs, gives

Forest	Input profiles	Center representatives
Three-star	1 513	757
Path of length three	5 173	2 587
Wedge and disjoint edge	1 513	757
Matching	1 513	761

An independent check verified the four placement orbits and directly counted the original tables for sixteen reverse lifts per placement. This completes part 4 of Theorem 8.

6.7 Flow counts on graphs joined at cut vertices

The comparison extends to nonplanar networks that mix different types of blocks. Fix a finite graph with oriented edges. Assume every block is planar or is a copy of $K_{3,3}$. At each input prescribe an integral divergence a and finite integral edge intervals $[\ell_e, u_e]$. In every nonplanar block the widths must belong to one of the four families of Theorem 8, fixed for that block across the mixture. The widths may vary between inputs and between blocks. Planar blocks allow arbitrary widths. If the parameter law is integrable, the input counts are positive almost surely and the mean data are integral, then the count at the mean has logarithm at least the expected log count.

For a block B, remove its edges. Each connected component of the resulting graph in its original connected component contains exactly one vertex v of B. By the cut identity, the divergence of the restriction to B at v is the sum of the prescribed divergences on that component. Conversely, these block divergences sum to $a(v)$ at each v because the incident block branches partition the other vertices and the component's total divergence is zero. The flow count therefore factors over the blocks with an integer-linear map on the divergence data.

On a planar block, subtract an integer-linear particular flow and express the remaining circulation in anchored face potentials. The cell bounds become directed difference bounds, to which the countable-mixture alcoved theorem (Theorem 1) applies. On $K_{3,3}$, orient the edges from the row side to the column side. Change variables to $y_e = x_e - \ell_e$ for edges already oriented this way and $y_e = u_e - x_e$ for reversed edges. These variables have precisely the given interval widths as caps. Their margins are integer-linear in the original data, so each of the four transportation

comparisons applies to the corresponding widths. The logarithm of each positive count is at most $\sum_e \log(1 + u_e - \ell_e)$, which is integrable. Adding the block inequalities gives the assertion.

The factorization uses that the nonplanar blocks meet the rest of the network at cut vertices.

7 Discrete Jensen for three-by-three transportation tables with an unbounded four-cycle

This section proves the arbitrary-mixture discrete Jensen inequality for 3×3 transportation tables in which the four entries of a fixed rectangle have independently varying interval bounds of unbounded width and the other five have binary widths. An analytic affine support treats every center with normalized slack at least nine, and exact logarithmic certificates treat the remaining 27,130 representative centers.

7.1 Statement

Let $N(\xi)$ count 3×3 integer tables with given integral row and column margins and finite entry intervals with integral endpoints. The vector ξ consists of all margins and interval endpoints. The expectation of these parameters is taken in the original affine coordinates. Use $\log 0 = -\infty$.

Theorem 9. *Fix a 2×2 rectangle. Let ξ be integrable random transportation data with integral mean $\bar{\xi}$. Suppose the widths outside the rectangle take values in $\{0, 1\}$. The four widths on the rectangle may vary independently over the nonnegative integers. Then*

$$\log N(\bar{\xi}) \;\geq\; \mathbb{E} \log N(\xi).$$

The arbitrary-mixture conclusion is stronger than linewise midpoint log concavity. Leake and Mohammadi Yekta's midpoint theorem applies after the bounds are encoded by nonnegative slack variables adjoining identity rows and columns to the network incidence matrix [5, Theorem 3.1]. Their Question 1.7 asks for the arbitrary-mixture conclusion for general unimodular fibers. The present proof treats one nonplanar family with binary exterior widths. Planar networks have a separate arbitrary-mixture theorem (Corollary 6). Bounded general 3×3 transportation families are also discussed in [7].

7.2 One circulation coordinate

If an input has zero count on an event of positive probability, the assertion is immediate. Otherwise average feasible real flows and use total unimodularity of incidence constraints to obtain an integer flow at the integral mean. Each binary exterior width equals its mean almost surely, because the mean endpoints and hence the mean width are integral. Subtract the exterior lower bounds, shifting the divergences affinely. If an exterior width is zero, delete that edge and apply the planar theorem to the resulting subgraph of $K_{3,3}$ with one edge removed. It remains to treat all five exterior intervals equal to $[0, 1]$.

Orient the rectangle consistently and negate the two variables traversed backward, including reversal of their intervals. Order its vertices as R_1, C_1, R_2, C_2. Let h be the flow on the fourth cycle edge. The other three edge values are $h + D_i - A_i$, where D_i, A_i are the prescribed and exterior divergence sums on the first i vertices. The final divergence equation is the component total. After fixed translations the exterior fluxes take the form

$$f = (a,\ a - c' + 1,\ a - c' + b + 1,\ 0), \qquad w = (1, 2, 3, 0), \qquad 0 \leq f_i \leq w_i.$$

Here a, b, c', d, e are the exterior bits in positions $13, 23, 31, 32, 33$. We index exterior fibers by the contracted margins $r = c' + d + e$, $c = a + b + e$. The four lower and upper bounds on h are $A_i + f_i$ and $B_i + f_i$, respectively, where A, B are affine in the original endpoints and divergences. In the sequel A, B always denote these effective offsets.

Put $L = \max_i A_i$, $U = \min_i(B_i + w_i)$, and $s = U - L + 1$. Define trims

$$t_i = (A_i + w_i - L)_+, \qquad v_i = (U - B_i)_+, \qquad 0 \le t_i, v_i \le w_i.$$

The number of lifts of a given exterior assignment to an integer circulation is $(s - D(f))_+$, where

$$D(f) = \max_i(f_i + t_i - w_i) + \max_i(v_i - f_i).$$

Both maxima are nonnegative since the fourth terms vanish. More generally the raw maxima attain nonnegative values at an index realizing L or U. Replacing a negative raw trim by zero adds at most a nonpositive term, so the maxima and every lift count are unchanged. The third coordinate is largest both in f and in $w - f$, by the binary formulas. Thus $D(f) \le f_3 + (3 - f_3) = 3$.

Write $p = (r, c)$, and let Y_p denote the exterior fiber. Its cardinality M_p, in order of increasing r, c from zero to three, is

$$M = \begin{pmatrix} 1 & 2 & 1 & 0 \\ 2 & 5 & 4 & 1 \\ 1 & 4 & 5 & 2 \\ 0 & 1 & 2 & 1 \end{pmatrix}.$$

This follows by summing over e, since each pair of binary entries has sum counts $(1, 2, 1)$. The count is $C_p(s) = \sum_{f \in Y_p}(s - D(f))_+$. It is at most $M_p s$ when positive, and is exactly affine in s for $s \ge 3$.

7.3 Large slack with independently varying endpoints

Every positive contracted center $p_0 = (r_0, c_0)$ has a strict rational multiplicative support. The table gives the slope bases and the minimum ratio $M_{p_0} u^{r - r_0} v^{c - c_0} / M_p$ over other positive profiles. Transposition and the complement $(r, c) \mapsto (3 - r, 3 - c)$ cover the remaining centers.

p_0	u	v	minimum ratio
$(0, 0)$	4	4	2
$(0, 1)$	5	1	2
$(0, 2)$	32	1/4	2
$(1, 1)$	7/5	7/5	7/4
$(1, 2)$	9/4	4/9	16/9

Each entry follows by substituting into the 4×4 count matrix.

Fix a center with $s_0 \ge 9$. For input offsets define the mean untruncated length

$$\Phi(A, B) = \frac{1}{M_{p_0}} \sum_{f \in Y_{p_0}} [\min_i(B_i + f_i) - \max_i(A_i + f_i) + 1].$$

Selecting one minimizing upper index and one maximizing lower index for each f at the center gives an affine upper support $H \ge \Phi$ with $H_0 = \Phi_0 \ge 6$. Uniformly, $H \ge s - 3$, because every deficit is at most three. An affine support of the full log count is

$$\log(M_{p_0} H_0) + (r - r_0) \log u + (c - c_0) \log v + (H - H_0)/H_0.$$

For $p \neq p_0$, use $C_p \leq M_p(H+3)$ and the tangent to log at H_0. Its additional term $3/H_0 \leq 1/2 < \log(7/4)$ is absorbed by the strict contracted support. For $p = p_0$ and $s \geq 3$, use $C_{p_0} = M_{p_0}\Phi \leq M_{p_0}H$.

In the final case $0 < s < 3$, the hinge bound $(s-D)_+ \leq (s/3)(3-D)$ gives $C_{p_0}/M_{p_0} \leq s(H+3-s)/3$, and also $C_{p_0}/M_{p_0} \leq s$. For $H < s$, the logarithm of the first bound minus H/H_0 increases up to $H = s$, since $H_0 \geq 3$. For $H \geq s$, use the second bound. Thus $\log(C_{p_0}/M_{p_0}) - H/H_0 \leq \log s - s/H_0 \leq \log H_0 - 1$, as needed.

7.4 Small slack and the exact certificate

Translate the circulation so $L = 0$, then clip endpoints to the canonical representative

$$A_i' = t_i - w_i, \qquad B_i' = s - 1 - v_i.$$

This changes neither count nor normalized slack. Seek a support with slopes $-\alpha_i$ on A_i and β_i on B_i, where $\alpha_i, \beta_i \geq 0$ and both sums equal μ. Such a support is invariant under the common translation and only rises when a clipped input is relaxed to its raw offsets. To ensure contact at the raw center, set a slope to zero if the corresponding endpoint could have been strictly clipped there. For $i \leq 3$ this means $t_i = 0$ or $v_i = 0$ for a lower or upper endpoint, respectively. For $i = 4$ the lower slope is zero if some other $t_i = w_i$, and the upper slope is zero if some other $v_i = w_i$. If no other index realizes the associated extreme, the fourth endpoint is forced to realize it and was not clipped.

After removing the fourth trim coordinate, the support has the two margin slopes, the six trim slopes, and μ. Require each sum of three trim slopes to be at most μ, with equality if the fourth slope is required to vanish. For each positive canonical center of slack $s_0 \leq 8$ it suffices to compare against all positive canonical data with slack at most 12. Comparison with the same center profile at $s_0 + 1$ implies

$$\mu \geq \log((s_0+1)/s_0) \geq \log(9/8) > 1/9,$$

since each hinge grows by at least the displayed ratio. At every comparator with fixed trims and profile, the derivative of its affine-tail log count on $[12, \infty)$ is at most $1/9$. The finite support thus lifts to all integer slacks. Positive inputs require $s \geq 1$.

The six trim coordinates have 24^2 possibilities. Over the fourteen positive margin profiles there are 52,132 positive canonical centers with slack at most eight. Complementing the exterior bits and reversing the circulation sends f to $w - f$, A to $-B - w$, B to $-A - w$, and (r, c) to $(3 - r, 3 - c)$. This affine involution swaps the two trim vectors and preserves counts and allowable support slopes. Keeping only lexicographic $t \leq v$ leaves 27,130 centers. The comparison table through slack 12 was also verified by counting integer h directly.

For each representative a numerical LP proposes a full-rank set of tight linear equations for the support slopes. It does not certify feasibility. We solve the proposed basis in the rational span of logarithms of positive integers, then check every required slope sign and comparison. If a candidate has a zero trim slope, the corresponding input trim can be set to zero, since its count can only rise. Thus the comparison grid can be restricted to trims supported on the active mask. All log signs are checked using rational series bounds with an integer product fallback. Bounded integer arrays are used only after checking an overflow bound. Six runs verify respectively $4522, 4522, 4522, 4522, 4522, 4520$ centers. The source package includes the exact comparison code and replay instructions.

The analytic argument or the finite certificate now gives an affine upper support of $\log N$ in the original data at every integral positive mean. Integrating this support proves Theorem 9.

7.5 Scope and transfer

The one-circulation reduction applies whenever the distinguished edges form a simple cycle. Suppose a finite collection of exterior profiles has nonempty finite sets of exterior flows, the circulation deficits admit a uniform bound $D \le D_{\max}$, and the contracted log count at the center has a strict affine support with gap $\Delta > 0$ against all other profiles. At a large enough central slack, the proof of Section 7.3 still applies, replacing three by $D_{\max}$. It suffices that the mean untruncated center length H_0 satisfies $H_0 \ge D_{\max}$ and $D_{\max}/H_0 \le \Delta$. This gives an analytic transfer to other single-cycle networks with bounded exterior profiles. The small-slack certificate of Section 7.4 is tailored to the binary exterior and rectangle geometry of Theorem 9.

8 Three-color transportation with mean cell caps at most two

This section proves a coefficientwise comparison for three 3×3 transportation fibers with possibly different cell caps whose cellwise means are integers at most two: at every fixed aggregate, the allocation count is largest when each color receives one third of every line sum and the common mean caps. Summing over aggregates gives the equal-weight three-input product inequality. The proof combines an interval translation with an exact integer comparison of 5,578,824 symmetry classes.

8.1 Result and relation to other comparisons

For nonnegative integer arrays $u^{(c)}$ of shape 3×3, let $N(r^{(c)}, s^{(c)}; u^{(c)})$ denote the number of integer arrays between zero and $u^{(c)}$ with row sums $r^{(c)}$ and column sums $s^{(c)}$, for $c = 1, 2, 3$. A triple of such arrays has an aggregate X obtained by adding them cellwise. Let M_X count triples with a specified aggregate as well as the given profiles and caps. Bars denote arithmetic means over the three colors.

Theorem 10. *Suppose every entry of $\bar{u}$ is an integer in $\{0, 1, 2\}$, and $\bar{r}, \bar{s}$ are integral. For every fixed aggregate X,*

$$M_X((r^{(c)}, s^{(c)}, u^{(c)})_{c=1}^3) \le M_X((\bar{r}, \bar{s}, \bar{u})_{c=1}^3).$$

Corollary 11. *Under these assumptions,*

$$\prod_{c=1}^{3} N(r^{(c)}, s^{(c)}; u^{(c)}) \le N(\bar{r}, \bar{s}; \bar{u})^3.$$

The corollary follows by summing Theorem 10 over X. Leake and Mohammadi Yekta prove a coefficientwise inequality at an integral midpoint for two profiles under shared box bounds [5, Theorem 3.1], and ask for arbitrary integral means in their Question 1.7. The dimension-four reduction of [7, Proposition 9.1] isolates cell-bounded 3×3 transportation as a cographic branch. Our theorem is a three-input result on that branch for integral mean caps at most two. The binary three-color 3×3 comparison of Section 9.4 and the product inequality for cellwise mean caps 0 or 1 in Section 9.5 are the special cases of Theorem 10 and Corollary 11 with every entry of $\bar{u}$ in $\{0, 1\}$. Section 9 extends the unit-cap comparison to any number of colors, to larger arrays and to growing cuts. The proof here is finite and uses no growing-cut hypothesis.

8.2 Local support reduction

It suffices to consider X with positive old count. Then every line sum of X is divisible by three. At a cell write $a = \bar{u}$, $u = (u^{(1)}, u^{(2)}, u^{(3)})$ and $x = X$ for its aggregate. For a signed integer

$k \in \{-2,-1,0,1,2\}$, let S_k be the signed compositions of $|k|$ into three nonnegative parts, with $S_0 = \{0\}$. In the target, an endpoint cell is forced and can be translated to S_0. When $x \le a$, old splits are a subset of S_x, and the target support equals S_x. When $3a - x \le a$, subtract the respective full cap vectors from the old and target splits. The old support is now a subset of S_{x-3a}, which is the full target support. These possibilities cover every cell except $a = 2, x = 3$.

In the middle case the target support after subtracting $\mathbf{1}$ is

$$R = \{0\} \cup \{e_i - e_j : 1 \le i, j \le 3,\ i \ne j\}.$$

The coordinate range of the old support has lower and upper endpoints $L_c = \max\{0, u^{(c)} - 3\}$ and $U_c = \min\{u^{(c)}, 3\}$. If no $u^{(c)}$ equals three then $U_c - L_c \le 2$ for each c. Choose integers $\ell_c \in [U_c - 2, L_c]$ with sum zero. These intervals are nonempty, their lower endpoints sum to at most zero since $\sum_c U_c \le 6$, and their upper endpoints sum to at least zero. Subtracting $\ell + \mathbf{1}$ embeds the old support into R. The only exceptions, up to color permutation, are $u = (3,3,0)$ and $u = (3,2,1)$. In these cells retain the exact residual support

$$H_u = \{z - \mathbf{1} : 0 \le z \le u,\ \textstyle\sum_c z_c = 3,\ z \in \mathbb{Z}^3\}.$$

Every translation changes the old profiles by fixed constants. It is therefore enough to bound the maximum of the old count over all residual profiles after enlarging the indicated supports. In the target each removed vector was equal across colors. Thus the target profile in the residual cells still gives every color one third of each residual line total.

8.3 Symmetries and their independent count

The underlying pattern records the common target support in each cell. Residue zero modulo three allows S_0 and R. Residue one allows S_1, S_{-2}, and residue minus one allows S_{-1}, S_2. The line divisibility conditions give 81 residue patterns, because the top-left 2×2 residue block determines the other five entries. There are 2^9 local type choices per residue array. Each R can be left ordinary or marked as exceptional. A marked cell receives one of nine ordered cap vectors, namely the distinct permutations of the two exceptions.

The action consists of row and column permutations, transposition, simultaneous reflection of the signed supports and one common permutation of the three colors. Reflection sends S_k to S_{-k}, preserves R, and sends H_u to $H_u + 2\mathbf{1} - u$ using $z \mapsto u - z$. This last fixed shift just relabels the old profiles. The comparison at issue is invariant under all these actions, even with multiple exceptional cells.

The generator first takes underlying target representatives under the geometric and sign action. It selects marked subsets under the stabilizer of each representative. The verifier assigns the nine cap labels to marks and retains the lexicographic minimum under the product of the subset's induced stabilizer and the global color action. Its prefix check rejects an assignment only when the first differing coordinates of an action and the original word are both known.

For an independent coverage calculation use Burnside's lemma on all $72 \cdot 2 \cdot 6 = 864$ geometric, sign and color actions. For each action loop over the 81 residue arrays and keep those invariant under its geometry and sign. For a geometric cycle of length ℓ with residue zero, its fixed label choices have generating polynomial

$$2 + \#\,\mathrm{Fix}(\tau^\ell)t^\ell,$$

where τ is the color permutation on the nine exceptional variants and the exponent counts marked cells. A cycle with nonzero residues has two ordinary choices, since invariance of its residues already handles the sign reversal. Each cycle is independent. Multiplying the factors, summing over invariant residues and the 864 actions, and dividing by 864 gives the orbit counts in Table 1.

Exceptional cells	Orbits	Exceptional cells	Orbits
0	470	5	301,416
1	1,116	6	835,148
2	5,456	7	1,608,580
3	21,876	8	1,804,940
4	86,456	9	913,366

Table 1: Orbits of normalized capped split systems. All 5,578,824 comparisons pass.

8.4 Exact comparison of allocation counts

In each cell add a fixed nonnegative offset when needed so every encoded entry is in $[0, 3]$. In a middle cell the added vector is $\mathbf{1}$, including for H_u. In an S_k cell the added vector is $\max(0, -k)\mathbf{1}$. The underlying total in each cell remains fixed. Group row allocations by the first two color totals. The third color total follows from the row aggregate. Within a group encode the first two colors of the first two cells by four base-ten digits. The full row is uniquely determined by this code and the row profile. Across three rows each column digit is at most nine, so adding codes has no carry. Together with the three row profiles, the sum code determines every column-color total.

For each row-profile triple, convolve the first two row code sets, keeping every multiplicity, and then add each element of the third set. The largest multiplicity over the resulting codes is the largest old allocation count for that row-profile triple. For fixed column code, any two row codes determine the third, and within its fixed profile the third code can occur at most once. Thus the product of sizes of any two groups bounds every multiplicity. The verifier can skip a row-profile triple using this bound. It can also use the maxima of the relevant group sizes to skip multiple triples at once. The target is computed directly by taking the ordinary supports everywhere and counting the single prescribed balanced profile. All operations are on integer codes and counts.

The generator and verifier included in the source package compute every orbit in Table 1; every old maximum is bounded by its target. In particular the nine-exception case is included. The largest possible count within any profile is at most 7^9, so signed 32-bit counts suffice. With the local translations of Section 8.2, the comparisons prove Theorem 10.

8.5 Consequences and scope

At a fixed aggregate on a larger bipartite grid with integral mean caps, delete the endpoint cells, whose splits are forced under both sets of caps, and adjust the prescribed divergences. At an articulation vertex v, delete v and take the component containing the rest of a particular incident block. The block's required divergence at v is the negative sum of the prescribed divergences over that component, since those edges of the block are precisely the edges crossing the cut. Its value is determined, so the allocation count factors over blocks. If the global aggregate line sums are divisible by three, so are the aggregate block divergences. Deleted cell totals are zero or three times a mean cap, and the cut sums preserve divisibility. Theorem 10 therefore applies blockwise when each cyclic block has at most three vertices on either side and every mean cap therein is at most two. Bridge splits are determined by the color divergences, so no bound on their integral mean caps is required.

The certificate concerns three colors and mean caps at most two, which is the range in which every middle cell has the supports listed in Section 8.2.

9 Balanced margins for multicolor transportation with unit cell caps

This section counts ordered decompositions of a fixed integral array into binary color tables with prescribed margins. For fixed numbers of rows and colors, balanced margins uniquely maximize the count once every row cut is crossed by a pair of rows that are both flexible in sufficiently many columns; for three rows and three colors, many flexible columns suffice. Exact finite balancing holds for two-row arrays, binary 3×3 tables and arrays with small reduced cubic components, and a marking reduction extends these comparisons to unequal caps at cells within one unit of empty or full.

9.1 Setting

Question 1.7 of Leake and Mohammadi Yekta concerns discrete log-concavity beyond integral midpoints for unimodular fiber counts [5]. Their Theorem 3.1 gives a coefficientwise two-input midpoint inequality. A shared rounding offset yields an arbitrary-mixture result for planar flow and type-A alcoved fibers (Sections 2–4). The low-dimensional unimodular reduction also singles out cell-bounded three-by-three transportation as a residual case [7, Section 9]. Its full row-column support is nonplanar.

We address a fixed-aggregate version with unit caps. For an $r \times n$ aggregate X and q colors, the count is the number of ordered binary tables with sum X that realize prescribed row and column margins in each color. The sum of the prescribed margins over the colors equals the corresponding aggregate margin. Balanced margins give every color one q-th of each aggregate margin, assumed integral. Exact maximization at balance is a coefficientwise statement, since X is held fixed.

The Fourier argument uses an anisotropic local limit theorem with two derivatives (Lemma 13) and exact color symmetry to locate the lattice maximum. The small-core argument compares counts directly and gives a finite three-input product inequality after summing over aggregates. A last reduction applies the fixed-aggregate results to unequal caps at cells near empty or full saturation.

9.2 Growing interactions across every row cut

Fix $r, q \geq 2$. Let X be an $r \times n$ integer array with entries in $\{0, \dots, q\}$, and suppose all row and column sums are divisible by q. Consider ordered decompositions of X into q binary arrays with prescribed margins for each color. In an aggregate column, call a row flexible if its entry lies strictly between 0 and q. For distinct rows i, j, define

$$w_{ij} = \#\{\ell : \text{rows } i, j \text{ are both flexible in column } \ell\}.$$

Theorem 12. *There is $M_{r,q} < \infty$, depending only on r and q, with the following property. If every nontrivial cut of the row set has some crossing pair with $w_{ij} \geq M_{r,q}$, then among all prescribed color margins the count is uniquely maximized by giving each color one q-th of every aggregate margin.*

The number of columns, the proportions of the column patterns and the number of forced columns are arbitrary.

9.2.1 One-column comparison and its exact lattice

For a fixed aggregate column with row sizes $k_1, \dots, k_r$, a color assignment is a binary $r \times q$ matrix with row sums k_i. Each color's prescribed total in the aggregate column is a column sum of this

matrix. The number of such assignments is uniquely maximized at the balanced color totals, when these are integral. To see strictness, suppose two color totals differ by at least two. First condition on the complete assignments of the other colors. In each row the two remaining colors have residual capacity zero, one or two. Rows of residual capacity two contribute both colors, and the rows of residual capacity one choose one. Conditional on the other colors, the number with prescribed two-color totals is one binomial coefficient. Moving one unit from the larger total to the smaller increases it strictly whenever the old number is positive. It never decreases the total after summing the conditionings. Iterating reaches the balanced vector. An infeasible old profile contributes zero.

A balanced assignment exists: list the colors periodically from 1 to q and allocate consecutive blocks of lengths k_i to the rows. No row gets a repeated color, and the total used length is a multiple of q. For any two flexible rows, put them consecutively in this construction. Their two sets of colors are nonnested. If their lengths sum to at most q they are disjoint. Otherwise each has an exclusive portion of size q minus the other's length. Choose one exclusive color from each, and switch those colors between the two rows. This gives two balanced assignments whose difference is a 2×2 switch. Relabeling colors gives any chosen color pair for the switch.

Subtract the means from the uniform balanced assignments. Every forced row becomes zero. Let U be the flexible row set for the column. The centered step lies in the subspace of matrices supported in U with zero row and column sums. Its differences generate the full integer lattice of that subspace, by the switches just constructed. In particular the covariance is positive definite on it. For fixed r, q there are finitely many column patterns. There is a common $\theta < 1$ bounding the ratio of any unbalanced column count to its balanced count, by the preceding strictness.

9.2.2 Choosing the scales and the Fourier estimate

Take a maximum-weight spanning tree T on the r rows with edge weights w_{ij}. For an edge e of T, write w_e for its weight. If removing e separates i and j, then $w_{ij} \le w_e$, since replacing e by ij would otherwise improve the tree. Thus the cut hypothesis forces every w_e above the threshold. Orient the edges arbitrarily and write $r_e = e_i - e_j$ for the root of the edge $i \to j$. These roots form an integral basis of the zero-sum row lattice, as can be seen by successively removing leaves. Let V_{col} be the real color-sum-zero space and $L_{\mathrm{col}} = V_{\mathrm{col}} \cap \mathbb{Z}^q$. In the tree coordinates the ambient lattice L of integer matrices with zero row and column sums is a product of $r-1$ copies of L_{col}.

Write a matrix $Y \in L$ as $Y = \sum_{e \in T} r_e \otimes y_e$ with $y_e \in L_{\mathrm{col}}$, and call $y = (y_e)_{e \in T}$ its tree coordinates; real matrices with zero row and column sums have real tree coordinates $y_e \in V_{\mathrm{col}}$ in the same way. Removing e from T splits the rows into two sides, and y_e is, up to a sign, the sum of the rows of Y on one side. Pair frequencies $t = (t_e)_{e \in T}$, $t_e \in V_{\mathrm{col}}$, with tree coordinates by $\langle t, y \rangle = \sum_e \langle t_e, y_e \rangle$. Let $C^*_{\mathrm{col}} \subset V_{\mathrm{col}}$ be the closed Voronoi cell about 0 of the dual lattice $L^*_{\mathrm{col}} = \{\tau \in V_{\mathrm{col}} : \langle \tau, v \rangle \in 2\pi\mathbb{Z} \text{ for all } v \in L_{\mathrm{col}}\}$, and put $C^* = \prod_{e \in T} C^*_{\mathrm{col}}$, a fundamental domain for the dual lattice of L. Put $m = (r-1)(q-1)$ and let v_L be the covolume of L in tree coordinates, so that $\operatorname{vol} C^* = (2\pi)^m / v_L$. If a random vector S takes values in one coset $x_0 + L$ and has characteristic function $\phi_S(t) = \mathbb{E} e^{i\langle t, S \rangle}$, Fourier inversion on L gives

$$\Pr(S = x) = F_S(x) := \frac{1}{\operatorname{vol} C^*} \int_{C^*} \phi_S(t) e^{-i\langle t, x \rangle} \, dt, \qquad x \in x_0 + L.$$

We call the smooth function F_S on the real tree-coordinate space the Fourier interpolant of S. It is real-valued, because $C^* = -C^*$ and $\phi_S(-t) = \overline{\phi_S(t)}$.

For column ℓ, let U_ℓ be its flexible set and let W_ℓ be its centered balanced step: a uniformly random balanced color assignment of the column minus its mean, written in tree coordinates. Let E_ℓ be the set of tree edges e for which U_ℓ meets both sides of the cut made by removing e. If

$e \notin E_\ell$, the e-th coordinate of W_ℓ vanishes, because the forced rows are zero after centering and the centered rows of U_ℓ sum to zero. Let $K_{r,q}$ bound the Euclidean norm of every tree coordinate of every centered balanced step, and let $n_{r,q} \le 2^{rq}$ bound the number of balanced assignments of a column. Let g_Σ denote the centered Gaussian density with covariance Σ on the tree-coordinate space.

Lemma 13 (Uniform anisotropic local limit theorem)**.** *Fix $r, q \ge 2$, $A \ge 1$ and $K \ge 0$. Let $W_1, \dots, W_n$ be independent centered balanced steps of columns with arbitrary patterns, and let B be an independent random vector with $\|B\| \le K$ such that $S = B + \sum_\ell W_\ell$ takes values in a single coset of L. Let $(w_e)_{e \in T}$ be positive numbers such that for every $e \in T$*

$$\textit{(i)}\ \ \#\{\ell : e \in E_\ell\} \le A w_e, \qquad \textit{(ii)}\ \ w'_e := \#\{\ell : \textit{both endpoints of } e \textit{ lie in } U_\ell\} \ge w_e/A.$$

Let D multiply the e-th tree coordinate by $\sqrt{w_e}$, put $J = \det D = \prod_e w_e^{(q-1)/2}$, and let $\Sigma = D^{-1} \operatorname{Cov}(\sum_\ell W_\ell) D^{-1}$. Then:

(a) $\lambda \|s\|^2 \le s^{\mathsf T} \Sigma s \le \Lambda \|s\|^2$ with $0 < \lambda \le \Lambda$ depending only on r, q, A;

(b) $\sup_x \Pr(S = x) \le C_{r,q} \prod_e (1 + w'_e)^{-(q-1)/2}$, and this part uses neither (i) nor (ii);

(c) for $k = 0, 1, 2$,

$$\sup_y \big\| D_y^k \big(J F_S(Dy) \big) - v_L D_y^k g_\Sigma(y) \big\| \le \delta\big(\min_e w_e\big),$$

where δ depends only on r, q, A, K and $\delta(w) \to 0$ as $w \to \infty$.

Proof. Envelope. Let W'_ℓ be an independent copy of W_ℓ and ϕ_ℓ the characteristic function of W_ℓ. Then $1 - |\phi_\ell(t)|^2 = \mathbb{E}[1 - \cos\langle t, W_\ell - W'_\ell \rangle]$. If both endpoints of e lie in U_ℓ, then for every pair of colors $a \ne b$ the switch $r_e \otimes (e_a - e_b)$ is a difference of two balanced assignments, by the one-column construction above, so it is a value of $W_\ell - W'_\ell$ of probability at least $n_{r,q}^{-2}$. Its tree coordinates are $e_a - e_b$ at e and zero elsewhere. Hence

$$1 - |\phi_\ell(t)|^2 \ \ge\ n_{r,q}^{-2} \sum_{\substack{e \in T \\ \text{both ends in } U_\ell}} \ \sum_{a \ne b} \big(1 - \cos\langle t_e, e_a - e_b \rangle\big).$$

The inner sum is continuous on C^*_{col} and vanishes exactly on L^*_{col}, since the vectors $e_a - e_b$ generate L_{col}; the only point of L^*_{col} in the closed Voronoi cell is 0. Near 0 the inner sum equals $\frac{1}{2} \sum_{a \ne b} \langle t_e, e_a - e_b \rangle^2 + O(\|t_e\|^4)$, a positive-definite quadratic form on V_{col} plus a higher-order term. By compactness the inner sum is at least $c_q \|t_e\|^2$ on C^*_{col} for some $c_q > 0$. With $c = c_q n_{r,q}^{-2}$ we get

$$1 - |\phi_\ell(t)|^2 \ \ge\ c \sum_{\substack{e \in T \\ \text{both ends in } U_\ell}} \|t_e\|^2, \qquad t \in C^*.$$

Since $|\phi_\ell| \le (1 - (1 - |\phi_\ell|^2))^{1/2} \le \exp(-\frac{1}{2}(1 - |\phi_\ell|^2))$ and $|\phi_B| \le 1$, multiplying over the columns gives the Gaussian envelope

$$|\phi_S(t)| \ \le\ \exp\Big(-\frac{c}{2} \sum_{e \in T} w'_e \|t_e\|^2 \Big), \qquad t \in C^*. \tag{2}$$

Part (b). Integrate (2) over C^*. Each factor $\int_{C^*_{\mathrm{col}}} \exp(-\frac{c}{2} w'_e \|\tau\|^2)\, d\tau$ is at most a constant times $\min(1, w_e'^{-(q-1)/2})$, hence at most a constant times $(1 + w'_e)^{-(q-1)/2}$.

Part (a). Since $1-\cos x \le x^2/2$, we have $1-|\phi_\ell(t)|^2 \le \frac{1}{2}\mathbb{E}\langle t, W_\ell - W_\ell'\rangle^2 = \mathrm{Var}\langle t, W_\ell\rangle$. Summing the envelope inequality over ℓ gives $\sum_\ell \mathrm{Var}\langle t, W_\ell\rangle \ge c\sum_e w_e'\|t_e\|^2$ for $t \in C^*$, and then for all t by homogeneity, since C^* contains a neighborhood of 0. With $t = D^{-1}s$ and (ii), $s^\mathsf{T}\Sigma s \ge c\sum_e (w_e'/w_e)\|s_e\|^2 \ge (c/A)\|s\|^2$. Conversely $|\langle D^{-1}s, W_\ell\rangle| \le K_{r,q}\sum_{e\in E_\ell}\|s_e\| w_e^{-1/2}$, so by Cauchy–Schwarz $\langle D^{-1}s, W_\ell\rangle^2 \le (r-1)K_{r,q}^2\sum_{e\in E_\ell}\|s_e\|^2/w_e$. Summing over ℓ and using (i) gives $s^\mathsf{T}\Sigma s \le (r-1)K_{r,q}^2 A\|s\|^2$.

Part (c). Write $\psi(s) = \phi_S(D^{-1}s)$ and $G(s) = \exp(-\frac{1}{2}s^\mathsf{T}\Sigma s)$. Since $\langle t, Dy\rangle = \langle Dt, y\rangle$, the substitution $s = Dt$ gives

$$JF_S(Dy) = \frac{1}{\mathrm{vol}\, C^*}\int_{DC^*}\psi(s)e^{-i\langle s,y\rangle}\,ds, \qquad v_{Lg_\Sigma}(y) = \frac{1}{\mathrm{vol}\, C^*}\int_{\mathbb{R}^m} G(s)e^{-i\langle s,y\rangle}\,ds.$$

Differentiating k times under the integrals multiplies the integrands by $(-is)^{\otimes k}$, so the quantity in (c) is at most

$$\frac{1}{\mathrm{vol}\, C^*}\Big(\int_{DC^*}\|s\|^k|\psi(s)-G(s)|\,ds + \int_{\mathbb{R}^m\setminus DC^*}\|s\|^k G(s)\,ds\Big).$$

Let $\rho > 0$ be the inradius of C^*_{col} and $w = \min_e w_e$. Then DC^* contains the ball of radius $\rho\sqrt{w}$. On DC^*, (2) and (ii) give $|\psi(s)| \le \exp(-\frac{c}{2A}\|s\|^2)$, and (a) gives $G(s) \le \exp(-\frac{\lambda}{2}\|s\|^2)$. Fix $R \ge 1$ with $R \le \rho\sqrt{w}$. The integral of $\|s\|^k(e^{-c\|s\|^2/(2A)} + e^{-\lambda\|s\|^2/2})$ over $\|s\| > R$ is at most some $\tau(R)$, depending only on r, q, A, with $\tau(R) \to 0$ as $R \to \infty$. This bounds the second integral and the part of the first integral over $\|s\| > R$.

On $\|s\| \le R$ let $X_\ell = \langle D^{-1}s, W_\ell\rangle$. Then $|X_\ell| \le \eta := (r-1)K_{r,q}Rw^{-1/2}$, and we assume $\eta \le \frac{1}{2}$. For a centered real random variable X with $|X| \le \eta$, the number $z = \mathbb{E}e^{iX} - 1$ satisfies $|z + \frac{1}{2}\mathbb{E}X^2| \le \frac{1}{6}\mathbb{E}|X|^3$ and $|z| \le \frac{1}{2}\mathbb{E}X^2 \le \frac{1}{8}$. By Cauchy–Schwarz, $(\mathbb{E}X^2)^2 \le \mathbb{E}|X|\,\mathbb{E}|X|^3 \le \eta\,\mathbb{E}|X|^3$, since $\mathbb{E}|X| \le \eta$. Hence $|\log(1+z) - z| \le |z|^2 \le \frac{1}{4}(\mathbb{E}X^2)^2 \le \frac{1}{4}\eta\,\mathbb{E}|X|^3$, and so $|\log \mathbb{E}e^{iX} + \frac{1}{2}\mathbb{E}X^2| \le \mathbb{E}|X|^3$. Since $|X_\ell|^3 \le (r-1)^2K_{r,q}^3R^3\sum_{e\in E_\ell} w_e^{-3/2}$, summing over ℓ with (i) gives

$$\Big|\log\prod_\ell \phi_\ell(D^{-1}s) + \tfrac{1}{2}s^\mathsf{T}\Sigma s\Big| \le (r-1)^2K_{r,q}^3AR^3\sum_e w_e^{-1/2} \le (r-1)^3K_{r,q}^3AR^3w^{-1/2} =: \varepsilon.$$

Also $|\phi_B(D^{-1}s) - 1| \le \mathbb{E}|\langle D^{-1}s, B\rangle| \le KRw^{-1/2}$. Therefore $|\psi(s) - G(s)| \le G(s)(\varepsilon e^\varepsilon + KRw^{-1/2})$ on $\|s\| \le R$, and the integral of $\|s\|^k$ times this is at most $C_{r,q,A}(\varepsilon e^\varepsilon + KRw^{-1/2})$. Altogether the quantity in (c) is at most

$$\frac{1}{\mathrm{vol}\, C^*}\Big(2\tau(R) + C_{r,q,A}\big(\varepsilon e^\varepsilon + KRw^{-1/2}\big)\Big).$$

Choose $R = w^{1/12}$. For large w the constraints $\eta \le \frac{1}{2}$ and $R \le \rho\sqrt{w}$ hold, $\varepsilon \to 0$ and $\tau(R) \to 0$. This bound is the required $\delta(w)$. □

9.2.3 Exact balancing and the defective-column comparison

Proof of Theorem 12. For prescribed color totals in every column and every row, the count equals $\prod_\ell n_\ell \cdot \Pr(\sum_\ell Z_\ell = R)$. Here n_ℓ counts the assignments of column ℓ with its prescribed color totals, the Z_ℓ are independent and uniform on these assignments, and R is the prescribed row-by-color matrix. Call a column balanced if its prescribed color totals are balanced and defective otherwise. Infeasible prescriptions give count zero.

The family of all columns satisfies the hypotheses of Lemma 13 with $B = 0$, the tree weights w_e and $A = 2r^2$. Indeed, (ii) holds with $w_e' = w_e$. For (i), if $e \in E_\ell$ then some pair i, j crossing

the cut of e is flexible in column ℓ; there are fewer than r^2 such pairs and each has $w_{ij} \le w_e$. Every subfamily obtained by deleting at most $w_e/2$ columns still satisfies (i) and (ii) with the same A. The cut hypothesis gives $\min_e w_e \ge M_{r,q}$, so we may take $\min_e w_e$ as large as needed, with constants depending only on r, q.

With every column margin balanced, the full means are integral by the row-sum hypothesis. The affine lattice coset is the integer zero-row-sum and zero-column-sum lattice, and contains zero. Let $S = \sum_\ell W_\ell$, so the balanced row profile corresponds to $S = 0$. Each step is invariant under color permutations, and so is C^*. Hence F_S is color invariant, and its gradient at zero is a color-invariant vector of $\bigoplus_e V_{\mathrm{col}}$. It vanishes because the color-sum-zero space has no nonzero fixed vector. By Lemma 13(a) the normalized covariance Σ lies in the compact set $\{\lambda \le \Sigma \le \Lambda\}$. Choose $\rho_0 > 0$ with $\rho_0^2 \le \lambda^2/(2\Lambda)$. For $\|y\| \le \rho_0$,

$$D^2 g_\Sigma(y) = g_\Sigma(y)\big(\Sigma^{-1} y y^{\mathsf{T}} \Sigma^{-1} - \Sigma^{-1}\big) \le -\tfrac{1}{2} g_\Sigma(y) \Sigma^{-1} \le -\kappa I,$$

with $\kappa > 0$ depending only on r, q. By Lemma 13(c) with $k = 2$, for large $\min_e w_e$ the function $y \mapsto JF_S(Dy)$ is strictly concave on the ball $\|y\| \le \rho_0$, and its gradient vanishes at 0. Hence $F_S(Dy) < F_S(0)$ for $0 < \|y\| \le \rho_0$. For $\|y\| > \rho_0$, part (c) with $k = 0$ gives

$$JF_S(Dy) \le v_L g_\Sigma(0) e^{-\rho_0^2/(2\Lambda)} + \delta < v_L g_\Sigma(0) - \delta \le JF_S(0)$$

once δ is small, since $g_\Sigma(0) \ge (2\pi\Lambda)^{-m/2}$. Thus $\Pr(S = x) < \Pr(S = 0)$ for every $x \ne 0$, and among profiles with all columns balanced the balanced profile is the unique maximum. Moreover $\Pr(S = 0) = J^{-1}(v_L g_\Sigma(0) + o(1))$.

Now let $d \ge 1$ columns be defective. Each has $n_\ell \le \theta\, n_\ell^{\mathrm{bal}}$. Condition on the assignments in the defective columns. This only shifts the target for the sum S_{rem} of the centered steps of the remaining balanced columns, so the ratio of the count to the balanced count is at most

$$\theta^d \, \frac{\sup_x \Pr(S_{\mathrm{rem}} = x)}{\Pr(S = 0)}.$$

The remaining columns satisfy $w'_e \ge (w_e - d)_+$. By Lemma 13(b), the ratio is at most

$$C'_{r,q}\, \theta^d \prod_e \Big(\frac{w_e}{1 + (w_e - d)_+} \Big)^{(q-1)/2} \le C'_{r,q} (d+1)^{m/2} \theta^d,$$

since $w/(1 + (w-d)_+) \le d+1$. Choose d_0 with $C'_{r,q}(d+1)^{m/2}\theta^d < 1$ for all $d > d_0$. For $1 \le d \le d_0$ and $\min_e w_e \ge 2d_0$, apply Lemma 13(c) with $k = 0$ to the remaining columns in the same scaling D. This gives $\sup_x \Pr(S_{\mathrm{rem}} = x) \le J^{-1}(v_L g_{\Sigma_{\mathrm{rem}}}(0) + \delta)$. By the upper bound in the proof of part (a), applied to the d deleted columns, $\|\Sigma - \Sigma_{\mathrm{rem}}\| \le (r-1) K_{r,q}^2 d_0 / \min_e w_e$, so $g_{\Sigma_{\mathrm{rem}}}(0) = g_\Sigma(0)(1 + o(1))$ uniformly. The ratio is therefore at most $\theta(1 + o(1)) < 1$ for large $\min_e w_e$. Taking $M_{r,q}$ large enough for all of these estimates proves the theorem. □

9.3 Three rows and three colors

Let X be a $3 \times n$ array with entries in $\{0, 1, 2, 3\}$, and assume all its row and column sums are divisible by three. Color decompositions are ordered triples of zero-one arrays with sum X. A column is flexible if it has an entry equal to 1 or 2.

Theorem 14. *There is an absolute N_0 such that, if X has at least N_0 flexible columns, the count for prescribed color margins is uniquely maximized when each color has one third of the aggregate row and column margins.*

The number of forced columns and the relative counts of the flexible types are arbitrary.

For this proof use the present-color mark at a cell of aggregate one and the absent-color mark at a cell of aggregate two. The latter contributes a negative mark and an additive constant of one in each color. A flexible column is either a complete column of one sign with three flexible rows, or a pair with opposite signs. A balanced complete column has six assignments and a balanced pair has three.

If the assertion failed, take a sequence of counterexamples with the number of flexible columns tending to infinity. After taking a subsequence, each of the two complete sign counts and the six ordered pair counts is fixed or tends to infinity. If the pairs of rows that share unboundedly many flexible columns connect all three rows, Theorem 12 with $r = q = 3$ applies. Otherwise there is exactly one unbounded unordered pair and only finitely many flexible columns outside it. We now handle this case.

Suppose the only unbounded overall count comes from one unordered pair, say rows 1 and 2, with $M \to \infty$ such columns. All complete and other pair counts are fixed along the chosen subsequence. Call these remaining flexible columns the finite set W. Forced columns contribute fixed constants and can be omitted. The large pair columns leave row 3 unchanged. For a balanced column in W, every possible row-3 marked color has exactly two completions if the column is complete, and exactly one if the column is a pair. The sign on that row is fixed by the aggregate. Thus the balanced marginal count for row 3, up to the common product of the two-completion factors, is the law of a sum of independent signed categorical steps. The balanced row-3 count is an integer vector because the aggregate margin is divisible by three. It is feasible, by pairing plus and minus positions and distributing the remaining positions, a multiple of three, evenly among the colors.

The signed categorical marginal has its unique maximum at balance. If two color totals x_i, x_j differ by at least two, fix the allocation of the third color in both the plus and minus positions. Among the remaining positions, let P, N count the two signs. The number of assignments giving the chosen two-color totals is $\binom{P+N}{x_i+N}$, since $x_i + N$ counts plus positions allocated to i and minus positions allocated to j. Its imbalance from the midpoint is $(x_i - x_j)/2$. Moving one unit from the larger total to the smaller strictly increases the count in every conditioning with positive count. The new count is at least zero on the other conditionings. Repeating reaches the balanced totals, establishing strictness for every feasible noncentral target.

An unbalanced column in W has, for each fixed row-3 marked color, at most as many completions as the balanced column. For a complete column with multiset of type $2 + 1$, those numbers across the three colors are $2, 1, 0$, compared with $2, 2, 2$; for a monochromatic column they are $1, 0, 0$. An unbalanced pair pins both marks, instead of allowing three common choices. The pointwise domination is strict for some marked color at any defective column. At the central row-3 target, every color at every column occurs in some balanced assignment, by feasibility and color symmetry. Consequently any unbalanced margin in W, or any noncentral row-3 target, gives a strictly smaller finite marginal count than the central balanced value $W_0 > 0$.

For the large pair, apply Lemma 13 with $r = 2$ to rows 1 and 2, whose tree has the single edge 12. The large-pair columns leave row 3 unchanged, so their centered steps lie in $r_{12} \otimes V_{\mathrm{col}}$, and each step is $\pm(e_a - \frac{1}{3}\mathbf{1})$ with a uniform. Take $w = M$ and $A = 2$. Then (i) holds since only the M large-pair columns enter, and (ii) holds for every subfamily of at least $M/2$ of these columns. The normalized covariance of the full family is a fixed matrix Σ_0, the covariance of e_a for uniform a. With $J = M$ and $g_0 = v_L g_{\Sigma_0}(0) > 0$, Lemma 13(c) shows that the maximum mass is $(g_0 + o(1))/M$, that the mass at every bounded offset in the correct coset has the same asymptotic, and that the two-derivative Gaussian estimate holds with scale $\sqrt{M}$. With all balanced columns, condition on the finitely many assignments in W giving the central row-3 target. The remaining bounded

centered offsets lie in the strong two-dimensional plane and form a color-invariant distribution. All required offsets for the central total belong to the large-pair affine lattice: every realized total and the central target are integer vectors shifted by the same column means, so each required offset is a difference of two integer vectors lying in the plane of the large-pair steps, and hence lies in the large-pair lattice. Thus the balanced count is $3^M W_0(g_0 + o(1))/M$. No profile with a strict finite marginal gap can attain it for large M.

If d large-pair columns have defective margins, each fixes its marks and loses a factor three. For bounded $d \geq 1$ the sharp maximum mass in the remaining $M - d$ columns is asymptotic to g_0/M, by Lemma 13(c) applied to these columns, whose normalized covariance is $\frac{M-d}{M}\Sigma_0$. For $d \to \infty$ the ratio of the upper bound to the central count is at most $C(d+1)3^{-d}$, using the bound $C/(1 + (M - d)_+)$ of Lemma 13(b). Both give strict loss, even at the maximal weak marginal. It remains to compare strong targets when there are no defects and row 3 is centered. The conditioned finite offsets form an independent bounded vector B with a color-invariant law, whose bound depends only on the fixed set W. Lemma 13(c) with this B gives the two-derivative approximation with the same leading Gaussian. The gradient at zero is exactly zero, by color invariance of the law of the sum. Strict concavity on a fixed rescaled ball and the Gaussian gap outside it give the unique exact maximum at the central strong target.

This excludes counterexamples on every subsequence and proves Theorem 14.

9.4 Exact finite cases

9.4.1 Two rows with any number of colors

For two rows the fixed-aggregate comparison holds at every width and with any number q of colors. A nonforced column has row sizes $k, q-k$ with $0 < k < q$, by divisibility of its aggregate total. At its balanced color margins the first row can choose any size-k subset of colors and the second row uses its complement. At any other color margins each first-row subset still determines at most one second-row choice, and some subsets are forbidden. Indeed an unbalanced feasible column has a color requested either zero times or twice.

For balanced columns the first-row count is therefore the coefficient of the product of the corresponding elementary symmetric polynomials. Equivalently it counts a sequence of subsets of prescribed sizes with prescribed color totals. These counts have a unique maximum at balanced first-row totals: fixing the placements of the other colors leaves only forced pairs and independent single choices between any two selected colors, so the binomial balancing argument applies. The center is feasible by choosing consecutive blocks in a periodic word of colors, since the sum of the subset sizes is divisible by q. At the center every subset of the requisite size is realized at each particular column, by color symmetry. Hence an unbalanced column strictly reduces the count even at the central row profile. The two row totals across all colors and the column color totals determine the second-row profile once the first is fixed. This proves exact uniqueness.

9.4.2 Three colors on small cubic cores

Let X be an integral 3×3 array with entries in $\{0, 1, 2, 3\}$. Assume each row and column sum of X is divisible by three. A decomposition is an ordered triple of zero-one tables $Y^{(1)}, Y^{(2)}, Y^{(3)}$ with sum X. For each color prescribe its row and column sums, with the prescribed margins adding to those of X. Then the number of decompositions is maximized when all three colors have the same margins, each equal to one third of the corresponding margin of X.

This statement is the case $u^{(1)} = u^{(2)} = u^{(3)} = \bar{u} \equiv 1$ of Theorem 10. We give a direct combinatorial proof, which extends to the larger arrays below.

9.4.3 Mark the flexible cells

Cells with $X_{ij} = 0$ or 3 are forced. At a cell with $X_{ij} = 1$ mark its unique present color, and at a cell with $X_{ij} = 2$ mark its unique absent color. Write + and −, respectively, for these cell types. The contribution of a minus cell to the color-margin vector is $\mathbf{1} - e_a$ if its marked color is a, while a plus cell contributes e_a.

Treat the rows and columns as the vertices of a bipartite graph, with an edge for each flexible cell. Direct plus edges from row to column and minus edges from column to row. At a vertex, the number of plus signs minus the number of minus signs equals the aggregate margin modulo three. Because this margin is divisible by three and the degree is at most three, the possible patterns are no edges, two opposite signs, or three identical signs. A degree-two vertex therefore has one incoming and one outgoing edge, and a degree-three vertex is a source or a sink.

The prescribed margin vector at a degree-two vertex, after subtracting the forced contributions, either requires the two marks to coincide or fixes both marks individually. Indeed $e_a - e_b = 0$ is precisely equality and each nonzero difference determines the ordered pair. On a directed path of degree-two vertices between degree-three endpoints, either all marks are a single common variable or all are fixed. To see this, if one degree-two vertex fixes both adjacent marks, propagate through the equality vertices in both directions. If there are more fixing vertices, each interval between them has forced marks or the data are inconsistent. The same applies to a directed cycle containing only degree-two vertices. Such a cycle has at most three assignments and has exactly three at the balanced margins.

Suppress the degree-two vertices in each path connecting degree-three endpoints, keeping the path as a labeled edge. Its endpoints are a source and a sink. The resulting bipartite multigraph has the same number p of sources and sinks, each of degree three, with $0 \le p \le 3$ because the original graph has just six vertices. An edge in this graph either has a single color shared at both ends or has fixed colors at its ends, which can differ. At each endpoint the multiset of its three incident marks is prescribed, since the three original signs there were identical. At the balanced margins every edge has a variable common mark and every endpoint requires one mark of each color. Thus the balanced count is the number of proper three-edge-colorings of this labeled multigraph, times the degree-two cycle counts.

9.4.4 Count the degree-three part

For $p = 0$ there is no degree-three part. If $p = 1$, the multigraph is three parallel edges. The number of choices is at most $3! = 6$ from the source multiset, and six is attained at the balanced margins.

For $p = 2$, its multiplicity array is

$$\begin{pmatrix} a & 3-a \\ 3-a & a \end{pmatrix}, \qquad a \in \{0, 1, 2, 3\}.$$

If $a = 0$ or 3, there are two independent triples of parallel edges, already handled. Otherwise label the two single edges as the diagonal and the two double edges as the off-diagonal. Each off-diagonal pair can be assigned in at most two ways once the marks on the single edges are given, since its multiset at its source is then fixed. If a single edge is fixed, it contributes no choice. If both are variable, say with mark colors x, y, compare the prescribed mark counts at the source and sink incident to one off-diagonal pair. The variable edges in that pair contribute the same counts at both ends, and the fixed edges contribute known vectors at each end. Consequently $e_x - e_y$ is prescribed. If nonzero it fixes the two colors, and if zero it gives at most three choices for their

common color. With one or both single edges fixed there are also at most three choices. The total is at most $3 \cdot 2 \cdot 2 = 12$. At balance, the two single edges must share their color, and either order of the other two colors is allowed on each parallel pair. Hence there are exactly twelve choices there.

For $p = 3$ all six vertices have degree three, so the graph is the original $K_{3,3}$ with one edge in every cell and all signs identical. If all margins prescribe one of each mark per line, the count is the twelve Latin squares of order three. This follows by choosing a first row in six ways, a second row in two ways as a derangement of the first, and a forced third row.

If any prescribed line is unbalanced, make it the first row by permuting or transposing. That row has at most three assignments, since a color repeats. Fix one such assignment. The remaining multiset of two marks in each column is now known. Where the two marks differ, the second row chooses between them. Order the three colors strictly, and represent an assignment by the set of columns in which the second row takes the larger option. For any two strictly nested sets, the sums of the numerical color labels in the second row differ strictly. A fixed second-row multiset therefore selects an antichain of subsets of at most three columns. Its size is at most three, by the antichain bound for the three-element Boolean lattice. The third row is forced. Thus the number is at most nine in this case.

Multiplying by the independent cycle counts proves the theorem for every flexible pattern. Infeasible input margin data contribute zero.

As a consequence, let $N(b)$ count binary 3×3 tables with prescribed row and column margins b. For three integral margin vectors b_1, b_2, b_3 with integral arithmetic mean b,

$$N(b_1)N(b_2)N(b_3) \leq N(b)^3.$$

Group both ordered triple counts by their fixed aggregate and apply the theorem in each group. The same proof allows a common pattern of forbidden cells, since such cells have aggregate zero on both sides.

9.4.5 Extension by suppressing parallel pairs

The proof extends to rectangular arrays of arbitrary size with unit cell caps, aggregate margins divisible by three, and at most three flexible cells per line, under the following graph condition. Suppress the degree-two paths, then suppress parallel pairs by the operation below. It suffices that some such reduction ends with every simple cubic component having at most five sources and five sinks.

After suppressing paths, a triple of parallel edges is an isolated component with at most six assignments and exactly six at balance. Remove it. If two edges are parallel from source u to sink v, each of u, v has one other incident edge. Their other endpoints are a sink v' and a source u', respectively. Let x be the mark of the edge $u \to v'$ at u, and let y be the mark of $u' \to v$ at v. The difference of the prescribed mark counts at u, v, after subtracting the known endpoint marks on any fixed parallel edges, prescribes $e_x - e_y$. If both external edges are variable, this either requires the same color or fixes both colors. If one external edge is fixed, the other is fixed or the instance is inconsistent.

Replace u, v and all their incident edges by one edge $u' \to v'$. When the two external colors must agree and both are variable, the replacement edge has this variable common color. Otherwise it carries the fixed colors at the two outer endpoints, which may differ. For each assignment of the replacement edge, there are at most two assignments of the deleted parallel pair by its prescribed multiset at u. At balance every assignment has exactly two extensions, because the external colors agree and the two parallel edges must use the other two colors. Thus the reduction preserves the claimed comparison whenever it holds for the resulting graph. Fixed colors that differ at the two

ends may also be changed to a single pinned common color by changing the prescribed counts at those ends accordingly.

The only simple cubic bipartite graph with three vertices on each side is $K_{3,3}$, handled above even when edges are pinned. In that case if an edge is pinned and every margin is balanced, the colorings form a subset of the balanced colorings without pins. The only one with four vertices on each side is $K_{4,4}$ minus a perfect matching. It has 24 balanced assignments. The first-color matching is a derangement, and the three derangements of type $(2,2)$ leave two alternating four-cycles, while the six of type (4) leave one alternating eight-cycle. The remaining two colors therefore give $3\cdot 4+6\cdot 2 = 24$. For any unbalanced prescribed margin, put an unbalanced line first, transposing if necessary. It has at most three assignments. If all margins are balanced, the upper bound holds immediately by dropping any pins. For each fixed first row, its three adjacent columns each have a prescribed two-mark multiset left, with at most two ways to assign it. The remaining column is missing the first-row edge, and its three marks are determined individually by their remaining row margins. Thus at most $3 \cdot 2^3 = 24$ assignments are possible.

With five vertices on each side, the bipartite complement is a simple two-regular graph, hence C_{10} or $C_4 \sqcup C_6$. Pins can be dropped after adjusting their fixed endpoint marks to agree, so it suffices to bound the count with ordinary prescribed multisets.

Complement C_{10}. Index rows and columns by $\mathbb{Z}/5\mathbb{Z}$, with row i missing columns $i, i+1$. The number of first-color matchings is 13 by inclusion–exclusion, since the numbers of sets of j disjoint forbidden edges on C_{10} are $1, 10, 35, 50, 25, 2$. The five four-cycles occur on adjacent row pairs. For a given four-cycle the forced first-color edges out of its two rows and into its two columns are disjoint and determine the remaining first-color edge. This follows, for rows $0, 1$ and columns $3, 4$, by taking the pairs $(0,2), (1,0), (4,3), (2,4), (3,1)$, and rotation gives the others. A complement of a matching cannot have two four-cycles, as it has ten vertices and its cycles have even length at least four. Thus five matchings leave a four-cycle and a six-cycle, and the other eight leave one ten-cycle. The balanced count is $5 \cdot 4 + 8 \cdot 2 = 36$.

For an unbalanced profile, by symmetry and transpose put an unbalanced line at row 0, with at most three assignments. Fix one of them. The residual row pairs of columns $2, 3, 4$ form the path $3 - 4 - 1 - 2$, with at most two choices in each column. After choosing these columns, the remaining edges for rows $1, 4$ are forced. The only four unforced entries are the 2×2 square in rows $2, 3$ and columns $0, 1$. It has at most two assignments, and two require the prescribed two-mark row multisets of rows $2, 3$ to be equal and mixed. If some column on the path has only one choice, there are at most four path choices and at most eight completions. Otherwise, flipping either end bit of the path changes the residual multiset of its corresponding square row injectively. Thus at most two pairs of end choices give equal square-row multisets, and at most four full path choices give two square completions. There are at most $8 + 4 = 12$ completions per first row, giving the upper bound 36.

Complement $C_4 \sqcup C_6$. The original graph comprises two copies of $K_{2,3}$, with the three degree-two vertices of one copy matched to the three degree-two vertices of the other. The join multiset is fixed by subtracting the sum of the two degree-three-vertex multisets from the sum of the three remaining prescribed multisets on one side of a gadget. At balance each join color is distinct. Each gadget has two assignments after the join colors are fixed, giving $6 \cdot 2^2 = 24$.

For arbitrary margins and a join word x, let $L(x)$ and $R(x)$ count the completions of the two gadgets. Each has at most three completions by the same antichain argument used for $K_{3,3}$. Three completions require the three residual pairs all to be the same mixed pair P. Indeed, selecting among differing marks with a fixed two-row multiset gives an antichain in the three-element Boolean lattice. An antichain of size three consists of the three singletons or the three doubletons. Equality of the selected color sums as vectors forces the three difference vectors to agree, so the residual

pairs coincide.

Consider the left gadget. If two join words x, x' have three completions, with common pair multiplicity vectors p, p', then $e_{x_j} - e_{x'_j} = p' - p$ for each of their three positions j. If $x \neq x'$, the nonzero difference determines the same ordered color pair everywhere, so both words are monochromatic. For a fixed nonmonochromatic join multiset, at most one word therefore has three completions. If that multiset has two colors, there are three words, and $\sum_x L(x)^2 \leq 9 + 2 \cdot 4$. For a monochromatic multiset there is just one word and the sum is at most 9. For three distinct colors and no word with three completions, the sum is at most $6 \cdot 4$. Otherwise let x be the word with three completions. The color outside its common mixed pair occurs in the prescribed three-mark multiset only at its one position in x. Any feasible word must assign that color at that position, leaving at most one other word. The sum is at most $9 + 4$ in this case. The same bounds hold for R. By Cauchy–Schwarz, $\sum_x L(x)R(x) \leq 24$. This completes the five-source comparison and hence the extension.

9.4.6 Comparison and scope

The planar alcoved-Jensen argument (Sections 2–4) rounds shared potential differences on planar networks. A full 3×3 transportation support has the nonplanar graph $K_{3,3}$. The binary three-color proof handles that case by suppressing directed paths and counting at most three branching vertices on either side. Parallel-pair reductions extend the argument to larger arrays. Mean caps up to two, whose cells admit color splits beyond a single present or absent mark, are treated for three colors in Section 8.

The proof compares aggregates over external marks and retains the full color multiset, and both features are essential, as two small cores show. First, pointwise transfer domination is not available on a four-cycle of a cubic core. Give its two external source marks two distinct colors and give its two external sink marks the same ordered pair. With all four margins balanced, there is exactly one internal square assignment. Under the unbalanced margins that prescribe the external mark plus the same two distinct residual marks at all four vertices, there are two alternating internal assignments. Thus a square reduction needs an aggregate comparison over external marks.

Second, the comparison uses the full color multiset, not only the divergence over $\mathbb{F}_2^2$. Identify the three colors with its nonzero elements. On the all-plus $K_{3,3}$ core the zero-divergence count is the proper-coloring count 12. If the prescribed divergence is the same nonzero color at all vertices, the marks of the other two colors form disjoint monochrome cycles. The graph has nine four-cycles and six six-cycles, and cannot contain two vertex-disjoint cycles. The count is therefore $1+2(9+6) = 31$. A flow reduction needs to retain multiplicity information beyond this modular divergence.

A separate exhaustive check of the signed 3×3 patterns verified all 81 possibilities. The balanced counts had frequencies $1, 30, 48, 2$ at counts $1, 3, 6, 12$, respectively. The proofs above are analytic.

9.5 Unequal caps through a binary translation

Let the q nonnegative integer caps in a cell be $u_1, \ldots, u_q$, with integral average $a = q^{-1} \sum_c u_c$. Assume its aggregate value x is in

$$\{0, 1, qa - 1, qa\} \cap [0, qa].$$

At $x = 0$ the allocation is the zero vector, and at $x = qa$ it is the cap vector u. At $x = 1$, it is e_c, with mark c allowed exactly when $u_c > 0$. At $x = qa - 1$, it is $u - e_c$ with the same restriction. If the two middle values coincide, use the first representation. For cells of zero average use only the zero allocation.

Define an effective binary aggregate by replacing the four cases by $0, 1, q-1, q$, respectively. Its line totals are congruent modulo q to those of the original aggregate. A full or minus-mark cell now has the fixed contribution $\mathbf{1}$ in place of u, and keeps the same variable mark when it has one. Thus changing the constants translates each prescribed color margin by a fixed amount. Every original decomposition injects, by its marks, into the binary decompositions with the translated profile. The binary count allows all marks, including those whose capacity was zero.

For the averaged caps $a\mathbf{1}$, all marks are allowed in the flexible cells, since there $a \geq 1$. The correspondence with binary decompositions is bijective. Its changes of constants are uniform across the colors, so the balanced margins correspond exactly to the balanced margins for the effective binary aggregate. Consequently each fixed-aggregate comparison proved above extends to these unequal caps, whenever the effective aggregate meets its hypotheses. This includes arbitrary unequal caps with cellwise means zero or one for three colors in the exact 3×3 result, a case also contained in Theorem 10, and Theorem 14. The statement transferred to the original unequal-cap profiles is the maximization at balance; uniqueness is asserted for the effective binary aggregate.

If three upper-cap arrays on 3×3 have cellwise integral averages all equal to zero or one, the fixed-aggregate result can also be summed. For three prescribed margin vectors with an integral mean, the product of the three counts under their respective caps is at most the cube of the count at the mean margins under the mean caps. Every feasible input aggregate automatically meets the cell restrictions in this case. Lower bounds with integral cellwise averages can first be subtracted, and the same conclusion applies when the resulting cap widths satisfy the hypotheses. This product inequality is also the case of Corollary 11 with every entry of $\bar{u}$ in $\{0, 1\}$.

9.5.1 Translations with linewise integral averages

The marking reduction is a special case of a criterion that allows intermediate aggregate values. In each cell choose an integer vector ℓ such that every feasible allocation y for the original caps and aggregate lies in $\ell + \{0, 1\}^q$. Write $R = \sum_c \ell_c$ in that cell and $X' = X - R$ for the resulting binary aggregate. Suppose the row and column sums of R are divisible by q. Denote the averaged cap by a, and define a base interval $[l, u]$ for each cell by

$$[l, u] = \begin{cases} [0, a], & X' = 0, \\ [-1, a-1], & X' = q, \\ [0, a-1], & 0 < X' < q. \end{cases}$$

Assume $R \in [ql, qu]$ cellwise. Then the binary comparison for X' implies the original unequal-cap comparison for X.

To see this, the map $y \mapsto y - \ell$ first injects the original decompositions into the effective binary ones at a translated margin profile. Its balanced count can be lifted to the averaged caps by finding q integer base tables with sum R, each with the same row and column margins, equal to the corresponding line sums of R divided by q, and with entries in $[l, u]$. For this construction, suppose s base tables remain and let R_s be their remaining aggregate. The entry constraints for the next table are the interval

$$[\, \max(l, R_s - (s-1)u),\ \min(u, R_s - (s-1)l)\,].$$

The fractional table R_s/s is feasible in this box with the required margins. The capacitated bipartite flow polytope with integer bounds and margins has an integer point, by total unimodularity of the vertex-edge incidence matrix after reversing one side's signs. Choose such a table, subtract it, and continue. Each balanced binary allocation z now lifts by adding these fixed base tables. The defining intervals ensure the lifted values lie in $[0, a]$ in every color and cell. The equal base margins

and the balanced binary margins sum to the balanced original margins. This is an injection into the averaged-cap count.

The cell test for an available unit cube uses the coordinate minima $L_c = \max(0, u_c-(qa-X))$ and maxima $U_c = \min(u_c, X)$, where u_c are the original caps. A cube exists precisely when $U_c - L_c \le 1$ for every color. For a varying coordinate $\ell_c = L_c$, and a fixed coordinate allows $\ell_c = L_c$ or $L_c - 1$. Thus the possible values of R form the integer interval $[\sum_c L_c - \#\{c : U_c = L_c\}, \sum_c L_c]$, to which the line divisibility and lifting conditions apply. For example, if each cell of a three-by-three array has caps $(5, 1, 0)$ and aggregate 3, one may take $\ell = (2, 0, 0)$. This gives $X' = 1$ and $R = 2$ in every cell, satisfying the lift despite R being indivisible by three in each cell. The criterion thus applies exactly to the cells in which every color's feasible coordinate interval has length at most one.

AI assistance. OpenAI Codex and Anthropic Claude assisted with developing and checking the arguments, numerical verification and manuscript preparation. The author is responsible for the work.

References

[1] A. Barvinok, *Brunn–Minkowski inequalities for contingency tables and integer flows*, Adv. Math. 211 (2007), 105–122. doi:10.1016/j.aim.2006.07.012.

[2] N. Gozlan, C. Roberto, P.-M. Samson and P. Tetali, *Transport proofs of some discrete variants of the Prékopa–Leindler inequality*, Ann. Sc. Norm. Super. Pisa Cl. Sci. (5) 22 (2021), 1207–1232. doi:10.2422/2036-2145.201906_002. Theorem numbering cited from arXiv:1905.04038v1.

[3] D. Halikias, B. Klartag and B. A. Slomka, *Discrete variants of Brunn–Minkowski type inequalities*, Ann. Fac. Sci. Toulouse Math. (6) 30 (2021), 267–279. doi:10.5802/afst.1674.

[4] T. Lam and A. Postnikov, *Alcoved polytopes I*, Discrete Comput. Geom. 38 (2007), 453–478. arXiv:math/0501246.

[5] J. Leake and M. Mohammadi Yekta, *Log-concavity and approximate counting for totally unimodular polytopes*, 2026. arXiv:2609.39917v1.

[6] OpenAI, *Gaussian free-field limits of weighted integer Lipschitz heights*, preprint, 6 October 2026. OpenAI Mathematics research release.

[7] S. Schubr, *Discrete Jensen inequalities for unimodular fiber counts of dimension at most three*, Luminary Research Archive, edition 6, 20 November 2026. doi:10.5281/zenodo.24406132. Archive record.

[8] B. A. Slomka, *A remark on discrete Brunn–Minkowski type inequalities via transportation of measure*, Israel J. Math. 261 (2024), 791–807. doi:10.1007/s11856-023-2596-3. Theorem numbering cited from arXiv:2008.00738v2.

[9] D. E. Speyer, *L-log-concavity and a proof of the conjecture of Lam, Postnikov and Pylyavskyy*, 2026. arXiv:2601.05007v2.