Foregger's Conjecture on Minimum Permanents of Nearly Decomposable Doubly Stochastic Matrices

Yair Lavi

4 November 2026

Abstract

We prove Foregger's conjecture: every nearly decomposable doubly stochastic matrix A of order $n \geq 2$ satisfies

$$\operatorname{per} A \geq 2^{1-n}.$$

Equality holds exactly for matrices obtained from $\frac{1}{2}(I_n + P_n)$ by independent row and column permutations, where P_n is a single-cycle permutation matrix. In fact, for every fixed nearly decomposable support whose bipartite graph is not a single cycle, the minimum permanent on its closed doubly stochastic face is strictly larger than 2^{1-n}.

Keywords. Permanent; doubly stochastic matrix; nearly decomposable matrix; tight cut; capacity.
2020 Mathematics Subject Classification. 15A15, 05C70, 05A20, 15B51.

1 Introduction

For an $n \times n$ matrix $A = (a_{ij})$, write

$$\operatorname{per} A = \sum_{\pi \in S_n} \prod_{i=1}^{n} a_{i,\pi(i)}.$$

A matrix is **doubly stochastic** if it is nonnegative and every row and column sums to one. Its support is the set of its positive positions. We regard a square support as a balanced bipartite graph with rows and columns as vertices and positive positions as edges. A square support D of order $n \geq 2$ is **fully indecomposable** (FI) if it has no $s \times (n-s)$ zero submatrix for $1 \leq s < n$. It is **nearly decomposable** (ND) if it is FI and deleting any one support edge destroys full indecomposability. These properties depend only on the support. In order one, by convention, a support is FI exactly when it is nonempty; the singleton support is ND. We write $\Omega(D)$ for the closed face of doubly stochastic matrices supported within D; thus a matrix with support exactly D lies in its relative interior.

Let P_n be the permutation matrix of the n-cycle $(1\,2\,\ldots\,n)$. A **cycle support** is a connected 2-regular bipartite support. For $n = 2$, $\frac{1}{2}(I_n + P_n)$ is the all-$\frac{1}{2}$ matrix.

Theorem 1.1 (Foregger's conjecture)**.** *Let $n \geq 2$ and let A be an ND doubly stochastic $n \times n$ matrix. Then*

$$\operatorname{per} A \geq 2^{1-n},$$

with equality if and only if A is obtained from $\frac{1}{2}(I_n + P_n)$ by independent row and column permutations. Moreover, if an ND support D is not a cycle support, then

$$\min_{A \in \Omega(D)} \operatorname{per} A > 2^{1-n}.$$

Nearly decomposable matrices were introduced by Sinkhorn and Knopp [15]. Their supports are the minimal elementary bipartite graphs studied by Lovász and Plummer [10]; related structural results for nearly decomposable matrices were developed by Hartfiel [8] and by Brualdi and Hedrick [1]. The bound 2^{1-n} and its equality classification are Foregger's conjecture, stated under his name as Conjecture 15 by Minc [13] and by Cheon and Wanless [3]. Foregger's own paper proved the bound for $2 \le n \le 9$ [5, Theorem 6] and credited R. A. Brualdi with suggesting the problem [5, p. 85]. His induction uses interior face minimizers and degree-two contractions [5, Lemmas 1–2], but the order-12 support obtained by replacing every edge of $K_{3,3}$ with a path of length three has no ND-preserving contraction [5, pp. 84–85]. Foregger later described minimizers for faces built from internally disjoint odd paths and proved the bound for fully indecomposable minimizers on certain more general graph-defined faces [6, Observation 2 and Theorem 3]. Brualdi and Shader determined the minimum on generalized Hessenberg faces [2].

Foregger's degree-two contraction identity reduces the order whenever the contracted support remains ND. We show that a support for which all such contractions fail is obtained from a smaller balanced bipartite graph, possibly with parallel edges, by replacing selected edges with paths of length three. An identity involving the entropy of perfect matchings then lets us combine bounds across a cut crossed exactly once by every perfect matching. If no such cut separates two parts of at least three vertices, a direct capacity estimate gives the strict bound. Otherwise, a smallest part supplies a strict estimate and induction controls the rest. These steps prove Theorem 1.1 and strictness on every fixed noncycle ND face.

Section 2 gives the minimizer, contraction, and support reductions. Sections 3–6 prove a strict bound for the remaining supports using matching entropy and tight cuts. Section 7 completes the induction and equality classification.

All logarithms are natural, and $0 \log 0 = 0$.

2 Support reductions

An **occurrence** is one labelled edge of a bipartite multigraph; parallel occurrences have the same endpoints but remain distinct when taking matchings or deleting edges. For a bipartite graph D, $N_D(S)$ is the set of neighbors of S in the opposite vertex class, $\deg_D(v)$ counts edges incident to v, and $D[S,T]$ is the set of edges joining S to T; parallel occurrences are counted separately. We write $N(S)$ when D is clear, and call a row or column a **line**.

A graph is **matching-covered** (or, in matrix terminology, has **total support**) if every edge lies in a perfect matching. For order at least two, it is FI precisely when it is connected and matching-covered [11, §4.1]; equivalently, its row sets satisfy strict Hall inequalities

$$|N(S)| \ge |S| + 1 \qquad (\varnothing \ne S \subsetneq U),$$

where U is the row set. The same assertions hold for occurrence-labelled bipartite multigraphs; parallel edges count separately as occurrences but not as distinct neighbors. For an order-one multigraph, FI means at least one occurrence and ND means exactly one.

2.1 Face minimizers

Lemma 2.1 (Foregger [5, Lemma 1])**.** *If D is FI of order at least two and A minimizes the permanent on $\Omega(D)$, then the support of A is FI and*

$$\operatorname{per} A(i \mid j) = \operatorname{per} A \quad \text{at every positive position } (i,j). \tag{2.1}$$

Consequently, if D is ND, every face minimizer has support exactly D.

Foregger's cofactor assertion uses a result of Hedrick [9], as credited in [5, Lemma 1]. Since $\Omega(D)$ is compact and its minimizers lie in its relative interior when D is ND, every point on its relative boundary has permanent strictly above the face minimum.

2.2 Contraction

For a degree-two column c meeting rows r, s, **contract at** c by deleting c and merging r, s into one row, adding their entries in every surviving column. The contracted matrix $C(A)$ is doubly stochastic of order $n-1$. A degree-two row has the transposed contraction. When A is positive on an FI support, the support of $C(A)$ has total support because $C(A)$ is doubly stochastic; contraction also preserves connectedness, so that support is FI.

Lemma 2.2 (Foregger's contraction identity [5, Lemma 2])**.** *At a permanent minimizer A on an FI face, contracting any row or column of degree two in* $\operatorname{supp} A$ *satisfies*

$$\operatorname{per} C(A) = 2 \operatorname{per} A. \tag{2.2}$$

Indeed, the two cofactors at a degree-two line equal $\operatorname{per} A$ by (2.1), and multilinearity gives (2.2).

Lemma 2.3 (no four-cycles)**.** *An ND support of order at least three contains no four-cycle. Hence, if a degree-two column c meets rows r, s, those rows have no common neighbor other than c.*

Proof. Suppose rows r, s both meet columns c, d. For $x \in \{c, d\}$, deletion of (r, x) violates strict Hall. Since D itself satisfies strict Hall, there is a nonempty proper row set S_x with $r \in S_x$, $|N_D(S_x)| = |S_x| + 1$, and r the only row of S_x adjacent to x; since s is also adjacent to x, we have $s \notin S_x$. Put $\sigma(S) = |N_D(S)| - |S|$. Both $S_c \cap S_d$ and $S_c \cup S_d$ are nonempty proper: r belongs to both S_c, S_d, whereas s belongs to neither. Submodularity gives $\sigma(S_c \cap S_d) + \sigma(S_c \cup S_d) \le \sigma(S_c) + \sigma(S_d) = 2$; strict Hall therefore forces $\sigma(T) = 1$ for $T = S_c \cap S_d$. Removing r from T removes both c, d from its neighborhood, so

$$|N_D(T \setminus \{r\})| \le |T| - 1 = |T \setminus \{r\}|.$$

Strict Hall forces $T = \{r\}$, and then $N_D(r) = \{c, d\}$. Symmetrically $N_D(s) = \{c, d\}$. The proper row set $\{r, s\}$ now violates strict Hall, a contradiction. □

Foregger [5, p. 76] states this fact, citing Hartfiel [8] and Minc [12, Theorem 1] for the equivalent matrix formulation and Lovász and Plummer [10] for the four-cycle formulation.

Minc proved that an ND support of order $n \ge 3$ has at most $3n-3$ edges [12, Theorem 2]. Every FI line has degree at least two. Thus every ND support of order at least three has a degree-two row and a degree-two column. The following structural lemma describes the case in which none of their contractions remains ND.

2.3 Terminal supports

An ND support of order at least three is **terminal** if no contraction at a degree-two row or column is ND. In a simple support, a vertex of degree at least three is a **hub**. An **occurrence core** is a balanced FI bipartite multigraph. Its **occurrence degree** counts incident labels with multiplicity, while the number of distinct neighbors counts each adjacent vertex once. A **bundle** is the set of occurrences with the same endpoints.

Lemma 2.4 (the two contraction mechanisms)**.** *Let c be a degree-two column of an ND support of order at least three. If both its neighboring rows have degree two, or both have degree at least three, then contracting c gives an ND support. The transposed assertion holds for a degree-two row.*

Proof. Write D' for the contraction and m for its merged row. Lemma 2.3 makes the surviving neighborhoods of r, s disjoint, so each edge f of D' has a unique original edge $\widetilde{f}$ in D. Averaging the perfect matchings of the FI support D gives a doubly stochastic matrix positive on D; its contraction shows that D' is FI, as observed above.

First suppose f is not incident to m, and that $D' - f$ is FI. Choose a doubly stochastic matrix positive exactly on $D' - f$. Let $\alpha, \beta > 0$ be the masses of row m on the columns inherited from r, s, respectively; $\alpha + \beta = 1$. Restore the two rows with their respective surviving entries, and put β at (r, c) and α at (s, c). The resulting matrix is doubly stochastic and positive exactly on $D - \widetilde{f}$. This support is connected because $D' - f$ is connected and the restored path $r - c - s$ joins the two parts of the merged row. Hence $D - \widetilde{f}$ is FI, contrary to ND.

Now let $f = (m, j)$ come from (r, j). Since $D - (r, j)$ is not FI, choose a nonempty proper row set S with $|N_{D-(r,j)}(S)| \leq |S|$. Strict Hall for D, together with deletion of only one edge, forces $|N_{D-(r,j)}(S)| = |S|$, $|N_D(S)| = |S| + 1$, $r \in S$, and r to be the only row of S adjacent to j. Taking T to be the complement of $N_{D-(r,j)}(S)$ among the columns gives $j \in T$, $|S| + |T| = n$, and $D[S, T] = \{(r, j)\}$. Strict Hall for the columns of D then gives $N_D(T) = S^c \cup \{r\}$. Also $c \notin T$, since $r \in S$ and (r, c) remains after deleting (r, j). If $s \in S$, put $S' = (S \setminus \{r, s\}) \cup \{m\}$. This is a nonempty proper row set of $D' - f$, with $|S'| = |S| - 1$. Its neighborhood is contained in $N_D(S) \setminus \{c, j\}$: contraction deletes c, and deleting f removes j, since r is the only row of S adjacent to j. Hence

$$|N_{D'-f}(S')| \leq (|S| + 1) - 2 = |S| - 1 = |S'|,$$

contrary to strict Hall. If $s \notin S$ and $|S| \geq 2$, then T is a nonempty proper column set of $D' - f$ with at most $|T|$ neighbors, again contradicting strict Hall. The remaining case $S = \{r\}$ forces $\deg_D(r) = 2$. It cannot occur when both neighbors of c are hubs; when both have degree two, m has degree two and deleting f directly destroys FI. Thus every edge deletion from D' destroys FI. The row case follows by transposition. □

Foregger [5, p. 76] attributes the hub case to Roberts.

For an FI occurrence core H, write $E(H)$ for its labelled occurrence set. An occurrence is **removable** if deleting it leaves an FI multigraph, and **essential** otherwise. Given $L \subseteq E(H)$, let $D(H, L)$ replace an occurrence $e = (u, v) \in L$, from core row u to core column v, by the path $u - c_e - r_e - v$, where c_e is a private column and r_e a private row. We call an occurrence in L **long**, and its three-edge replacement a **long path**. Occurrences outside L stay **direct**. Parallel direct occurrences form one edge of the simple support, while their labels remain distinct in core matchings and flows. Let $L(H)$ be the removable occurrences and $D_{\mathrm{can}}(H) := D(H, L(H))$ its **canonical expansion**. Parallel occurrences are removable: deleting one leaves the simple support unchanged, and another occurrence with the same endpoints can replace it in any perfect matching. Hence they give separate paths in $D_{\mathrm{can}}(H)$. If h is the order of H, this expansion has order $h + |L(H)|$.

For $H = K_{3,3}$, all nine occurrences are removable, so $D_{\mathrm{can}}(H)$ replaces every edge by a three-edge path and has order $3 + 9 = 12$. For a six-cycle core with one perfect matching doubled, the six occurrences in doubled bundles are removable and the other three are essential. Its canonical expansion has order $3 + 6 = 9$, with the former occurrences long and the latter direct.

Lemma 2.5 (terminal classification)**.** *Every noncycle terminal ND support is $D_{\mathrm{can}}(H)$ for a balanced FI occurrence core H whose minimum occurrence degree is at least three. Conversely $D_{\mathrm{can}}(H)$*

is ND for every FI core H, including lower-degree cores. If H has minimum occurrence degree at least three, every private-line contraction in $D_{\mathrm{can}}(H)$ is FI but not ND.

Proof. We first recover a core from a terminal support. We then show that canonical expansions of FI cores are ND. Finally, we identify which recovered occurrences are long and analyze their contractions.

Step 1: Recovering the core. Because D is connected, has minimum degree at least two, and is not a cycle, it has a hub. Lemma 2.4 forces every degree-two vertex to have exactly one degree-two neighbor and one hub neighbor. Thus the degree-two vertices form disjoint adjacent row-column pairs, each of which is the interior of a three-edge path between two hubs. Suppressing these pairs leaves exactly the hubs as vertices of a balanced core H: each path becomes one labelled occurrence, while each remaining hub-to-hub edge becomes a direct occurrence. A perfect matching chooses either both outer edges of a private path or its middle edge; this gives a bijection with labelled core matchings. Suppression also preserves connectedness, so H is FI. At each hub, its incident edges in D correspond bijectively to incident occurrences in H; hence its occurrence degree is at least three.

Step 2: Canonical expansion and deletion. Every occurrence of an FI core belongs to a perfect matching. Unless H is the one-occurrence order-one core, each occurrence can also be avoided: choose a different occurrence at its row and extend it. Thus replacing occurrences by private paths preserves connectedness and matching coverage, so $D(H, L)$ is FI except when that one-occurrence core has its sole occurrence long. For the exceptional core, $L(H) = \varnothing$ and $D_{\mathrm{can}}(H)$ is its FI singleton. Deleting any path edge of an FI expansion leaves a private degree-one vertex and destroys FI. If a direct support edge represents a single essential occurrence e of H, deleting it also destroys FI after expansion: an FI expanded deletion would suppress to an FI $H - e$. Conversely, if $H - e$ is FI, its expansion after deleting the direct edge e is FI except when $H - e$ is an order-one core with one long occurrence; that expansion is a path. In $D_{\mathrm{can}}(H)$ every direct occurrence is essential and represents a single support edge, since parallel occurrences are removable and long. Hence $D_{\mathrm{can}}(H)$ is ND for every FI core H.

Step 3: Terminal identification and private contractions. For the terminal support D, the core has minimum occurrence degree at least three, so deleting any one occurrence cannot leave the exceptional one-occurrence core. By construction of H, every direct edge of D represents exactly one core occurrence. Deleting such an edge preserves FI exactly when deleting its core occurrence does, by the expansion argument above; thus every direct occurrence of D is essential in H. Contracting a private degree-two line of a long occurrence e produces $D(H, L \setminus \{e\})$. Lemma 2.3 ensures that its new direct edge does not merge with an existing direct edge. The remaining path edges stay essential, and the old direct edges stay essential by the deletion test. Consequently the contracted support is ND exactly when e is essential in H. Since D is terminal, all its long occurrences are removable, while its direct occurrences are essential. Thus it is $D_{\mathrm{can}}(H)$. Conversely, when the core has minimum occurrence degree at least three, every private-line contraction of $D_{\mathrm{can}}(H)$ is FI but not ND because its newly direct occurrence is removable. □

Remark 2.6 (a qualification to Foregger's subdivision claim). Foregger's general assertion about full subdivisions [5, p. 85] requires an additional hypothesis: a fully subdivided FI core of minimum occurrence degree at least three has no ND-preserving private-line contraction exactly when every occurrence is removable. For example, take two disjoint copies of $K_{3,3}$, delete r_1c_1 from the first and r_2c_2 from the second, and add the cross edges r_1c_2 and r_2c_1. The resulting core is connected and matching-covered: internal edges occur in matchings within the copies, and the cross edges occur together in a matching. It has degree three at every vertex, but its two cross edges are essential: deleting one leaves the other uncovered by any matching. Its full subdivision is ND, and contracting the private line of either cross edge still gives an ND support.

In the theorem's induction, a support with an ND-preserving contraction is handled by Lemma 2.2 and a smaller-order instance of Theorem 1.1. The remaining noncycle supports are terminal and are classified by Lemma 2.5.

3 Matching entropy on expanded cores

First we turn a permanent bound on an entire face into an entropy bound at each allowed choice of edge weights. This lets us combine estimates when we split a support in Section 4.

Let H be an FI balanced bipartite multigraph of order h. A **feasible occurrence flow** $x = (x_e)_{e\in E(H)}$ is a nonnegative vector satisfying

$$\sum_{e\ni u} x_e = 1 \quad (u \text{ a row or column vertex}). \tag{3.1}$$

Let $\mathcal{M}(H)$ denote the set of labelled perfect matchings of H; a law on $\mathcal{M}(H)$ is a probability distribution on that set. Aggregate the flows of parallel occurrences to obtain a doubly stochastic matrix, and decompose it into permutation matrices by the Birkhoff–von Neumann theorem. For each matrix edge selected by a permutation, choose a labelled occurrence with probability proportional to its flow; bundles of zero total flow are never selected. This gives a law on labelled perfect matchings with marginals x. Conversely, the marginals of any such law satisfy (3.1). If $x_e > 0$ for all occurrences, we call the flow positive. Write $\mathrm{Ent}(p) = -\sum p_i \log p_i$ for Shannon entropy, $h_2(t) = -t\log t - (1-t)\log(1-t)$ for binary entropy, and

$$\mathrm{Ent}^*_H(x) = \max_{\substack{\mu \text{ a law on } \mathcal{M}(H)\\ \mu(e\in M)=x_e}} \mathrm{Ent}(\mu), \qquad F_H(x) = \mathrm{Ent}^*_H(x) - \sum_{u \text{ row}} \mathrm{Ent}((x_e)_{e\ni u}). \tag{3.2}$$

The feasible laws form a nonempty compact polytope, so continuity of entropy gives the maximum. These definitions apply equally to simple supports.

Lemma 3.1 (face variational identity)**.** *For a square total-support graph D and $X \in \Omega(D)$,*

$$\log \mathrm{per}\, X \ge F_D(X), \qquad \min_{X\in\Omega(D)} F_D(X) = \log \min_{X\in\Omega(D)} \mathrm{per}\, X. \tag{3.3}$$

Consequently, for every $X \in \Omega(D)$, including boundary points,

$$F_D(X) \ge \log \min_{Y\in\Omega(D)} \mathrm{per}\, Y.$$

Proof. Restrict to the support S of X. Every law μ with edge marginals X is supported on $\mathcal{M}(S)$. Put $w(M) = \prod_{e\in M} x_e > 0$, $Z = \sum_{M\in\mathcal{M}(S)} w(M) = \mathrm{per}\, X$, and $q(M) = w(M)/Z$. The finite Gibbs inequality is the nonnegativity of relative entropy:

$$\log Z = \mathrm{Ent}(\mu) + \sum_M \mu(M)\log w(M) + D_{\mathrm{KL}}(\mu\|q) \ge \mathrm{Ent}(\mu) + \sum_M \mu(M)\log w(M),$$

where $D_{\mathrm{KL}}(\mu\|q) = \sum_M \mu(M)\log(\mu(M)/q(M)) \ge 0$. The expectation of $\log w(M)$ is $\sum_e x_e \log x_e = -\sum_u \mathrm{Ent}(x_u)$. Maximizing over μ proves the first inequality of (3.3).

For the converse inequality, every edge of S lies in a perfect matching because X is doubly stochastic. Let B be the average of the matching matrices of S. Since X is positive on S, $(X - \varepsilon B)/(1-\varepsilon)$ is doubly stochastic and supported within S for sufficiently small $\varepsilon > 0$. Express it as a convex combination of matchings and mix the corresponding law with ε times the uniform law on

$\mathcal{M}(S)$. The result has marginal X and assigns positive probability to every matching of S. Mixing any feasible law having a zero coordinate with this positive law increases entropy by order $-\varepsilon \log \varepsilon$, so the entropy maximizer μ^* is positive. Because μ^* is positive, Lagrange stationarity applies with one constraint for normalization and one for each edge marginal, giving

$$\log \mu^*(M) = \lambda_0 + \sum_{e \in M} \lambda_e \qquad (M \in \mathcal{M}(S)).$$

Thus μ^* has product weights $w_e = e^{\lambda_e} > 0$. These weights are positive on S, and every edge of S lies in a perfect matching; hence $W = (w_e)$ has total support. The Sinkhorn–Knopp scaling theorem [14] gives a doubly stochastic scaling Y of W on S. Since scaling multiplies every perfect-matching weight by the same constant, Y induces μ^*, so

$$\mathrm{Ent}(\mu^*) = \log \operatorname{per} Y - \sum_{ij} x_{ij} \log y_{ij}.$$

Subtracting the row entropies gives

$$F_D(X) = \log \operatorname{per} Y + \sum_{ij} x_{ij} \log \frac{x_{ij}}{y_{ij}} \geq \log \operatorname{per} Y \geq \log \min_{Q \in \Omega(D)} \operatorname{per} Q. \tag{3.4}$$

The first inequality in (3.4) is the sum of the rowwise relative-entropy inequalities. The restriction to S handles all boundary points directly. Combining (3.4) with the first inequality proves (3.3). □

Lemma 3.2 (aggregation of parallel occurrences)**.** *Let H be an FI occurrence core with feasible flow x. Let $\bar{H}$ be obtained by merging one parallel bundle into a single occurrence, and let $\bar{x}$ assign that occurrence the sum of the bundle flows, leaving all other flows unchanged. Then $F_H(x) = F_{\bar{H}}(\bar{x})$. Repeated application merges any collection of bundles.*

Proof. If the bundle has total flow $s > 0$ and labelled flows (x_e), independent choice of its label adds $s \,\mathrm{Ent}((x_e/s)_e)$ to maximum matching entropy. Conversely, for any labelled law μ projecting to an aggregated law ν, the entropy chain rule and concavity give $\mathrm{Ent}(\mu) - \mathrm{Ent}(\nu) \leq s \,\mathrm{Ent}((x_e/s)_e)$. Grouping the bundle at its row endpoint adds exactly the same quantity to the row-entropy sum, so the contributions cancel. At $s = 0$ both contributions vanish. □

Every doubly stochastic matrix supported on $D(H, L)$ admits a feasible flow lift x: a long path has entries $(x_e, 1 - x_e, x_e)$ in path order, and a direct matrix entry is the sum of its parallel direct-occurrence flows. Write X_x for the doubly stochastic matrix induced by x. When $D(H, L)$ is FI, the matrix is interior exactly when it admits a positive lift. In $D_{\mathrm{can}}(H)$ all parallel occurrences are long, so the lift is unique. Labelled matchings of H and $D(H, L)$ correspond as in Lemma 2.5. Each long path adds one private row of entropy $h_2(x_e)$, so

$$F_{D(H,L)}(X_x) = F_H(x) - \sum_{e \in L} h_2(x_e). \tag{3.5}$$

If a direct matrix entry has several labelled flow lifts, Lemma 3.2 shows that changing its split leaves $F_H(x)$ unchanged; those direct occurrences do not appear in the sum over L. Thus (3.5) is independent of the lift.

Put $\delta(t) = \log 2 - h_2(t) \geq 0$ and define the normalized gap

$$G_{H,L}(x) = F_H(x) + (h - 1) \log 2 + \sum_{e \in L} \delta(x_e). \tag{3.6}$$

Since $D(H, L)$ has order $h + |L|$, (3.5) says

$$G_{H,L}(x) = F_{D(H,L)}(X_x) + (h + |L| - 1)\log 2. \tag{3.7}$$

For the canonical expansion the gap is $G_{H,L(H)}$. By the first inequality of (3.3), positivity of this gap at every positive flow excludes an interior face minimum at or below 2^{1-n}.

4 Tight cuts and expansion accounting

To use induction across a cut, we split the core while preserving matching marginals and record exactly which occurrences remain long. Tight cuts and their contractions are standard in matching-covered graph theory [4]. A **tight cut** of an FI core H is a partition of its vertices into two shores such that every perfect matching uses exactly one edge across it. For a fixed cut, an occurrence is **internal** if both endpoints lie in the same shore, and **crossing** otherwise. It is **nontrivial** if each shore contains at least three vertices. Let one shore contain a rows and b columns. In any perfect matching, the number of edges leaving its rows minus the number entering its columns is $a - b$. Tightness forces $a - b = +1$ or -1, with every crossing edge oriented the corresponding way: each edge is used by some matching. In particular, both shores have odd size, so the condition that each has at least three vertices is equivalent to requiring that neither shore be a single vertex. Let S_1 be the shore with one extra row and S_2 the shore with one extra column. Close S_1 by a marker column and S_2 by a marker row, retaining the labels of all crossing occurrences, as in Figure 1. This gives balanced FI factors H_1, H_2, respectively, of orders m_1, m_2: each factor edge appears in a projected parent matching, and every component of a shore attaches to the marker because the parent is connected. Thus the factors are matching-covered and connected. Each crossing occurrence $e = (u, v)$, with $u \in S_1$ and $v \in S_2$, becomes $(u, \text{marker column})$ in H_1 and $(\text{marker row}, v)$ in H_2, retaining label e in both. Their orders satisfy

$$h = m_1 + m_2 - 1. \tag{4.1}$$

Their matchings glue exactly when they choose the same crossing label. Every factor matching extends, because its label occurs in a parent matching and hence in a matching of the opposite factor. For a feasible parent flow x, project any matching law with these marginals to the two factors. The inherited factor flows keep x_e on internal occurrences and assign x_e to the marker edge with crossing label e. They are feasible; in particular, the marker sum is $\sum_{e \text{ crossing}} x_e = 1$ because the cut is tight.

A **leaf** is a closed tight-cut factor with no nontrivial tight cut of its own.

Lemma 4.1 (entropy and removability across a tight cut)**.** *At factor marginals inherited from a parent flow,*

$$F_H = F_{H_1} + F_{H_2}. \tag{4.2}$$

An internal occurrence is removable in H exactly when it is removable in its factor. For a crossing occurrence e, whose label appears in both factors,

$$H - e \text{ is fully indecomposable} \iff H_1 - e \text{ and } H_2 - e \text{ are fully indecomposable.} \tag{4.3}$$

Proof. Let q be the crossing-label distribution. Projecting any parent matching law to the factors and using conditional entropy gives $\mathrm{Ent}^*_H \le \mathrm{Ent}^*_{H_1} + \mathrm{Ent}^*_{H_2} - \mathrm{Ent}(q)$. Conversely, glue entropy-maximizing factor laws independently conditional on their common label. This gives the prescribed parent edge marginals and attains equality. The two factor row-entropy sums exceed the parent

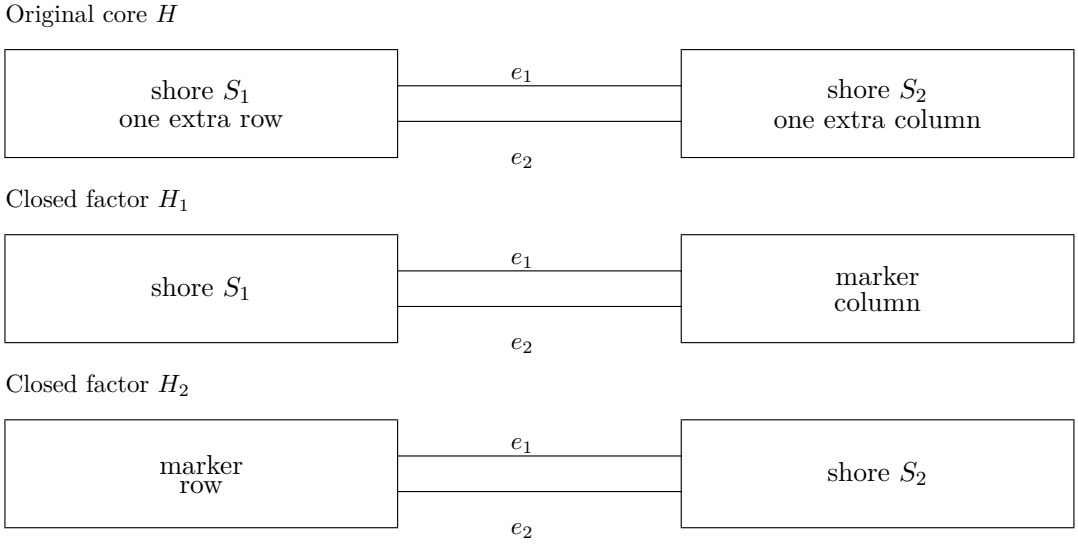


Figure 1: A tight cut with representative crossing labels and the two closed factors. The same labels appear at both markers; factor matchings glue when they choose the same label.

row-entropy sum by precisely $\mathrm{Ent}(q)$, from the added marker row. This proves (4.2), also when some labels have zero probability.

Every FI factor of order at least two is 2-connected [11, §4.1], as also follows from strict Hall, so deleting its marker leaves a connected shore.

If both factors remain FI after deleting a crossing occurrence e, each marker-deleted shore is connected by 2-connectedness. Every surviving factor edge belongs to a factor matching. That matching chooses a surviving crossing label, which also lies in a matching of the opposite factor because the latter is matching-covered. Gluing along this common label covers the chosen edge in a parent matching; the connected shores and a surviving crossing edge make the parent connected. Thus the parent deletion is FI. Conversely, parent matchings after deletion project to matchings covering every surviving factor edge, and contracting the opposite shore preserves factor connectedness. This proves (4.3). The same argument with one factor unchanged proves the internal assertion. □

For the asymmetric accounting below, choose either factor and call it H_1, relabelling the factors of Figure 1 if necessary; call the other H_2. Thus the marker of H_1 may be a row or a column. Let L_B consist of its removable internal occurrences, and set $B = D(H_1, L_B)$, leaving all marker-incident occurrences direct. Call $R = D_{\mathrm{can}}(H_2)$ the **remainder**, and let N be the crossing occurrences essential in H_1 but removable in H_2. Write $G_B = G_{H_1,L_B}$ and $G_R = G_{H_2,L(H_2)}$. This asymmetry is used in Section 7: H_1 is chosen as a leaf, with its marker-incident occurrences left direct for Lemma 5.1 (or Lemma 6.3 at order two), while the canonical expansion R of H_2 is bounded by induction. The internal and crossing rules in Lemma 4.1 give the exact occurrence identity

$$L(H) = L_B \sqcup \big(L(H_2) \setminus N\big). \tag{4.4}$$

Equations (3.6), (4.1)–(4.4) therefore give

$$G_{H,L(H)} = G_B + G_R - \sum_{e \in N} \delta(x_e). \tag{4.5}$$

The three terms use the same parent marginal. The subtraction removes exactly the terms counted in R for labels in N but absent from the parent expansion.

5 A permanent bound for one shortened star

This is the quantitative input for every leaf of order at least three. We state it on the entire closed face, so Lemma 3.1 can transfer it to arbitrary inherited marginals.

The **star** of a core vertex v is its set of incident occurrences. To **shorten** that star, leave those occurrences direct and replace every other occurrence by a three-edge path.

Lemma 5.1 (one-star bound)**.** *Let K be an FI balanced occurrence core of order $k \geq 3$, with at least three distinct neighbors at every vertex. Choose one row v, called the* **root***. Expand every occurrence not incident to v to a three-edge path, leave the occurrences incident to v direct, and aggregate any parallel occurrences at v into single matrix entries. Let ℓ be the number of non-v occurrences. Each of the $k-1$ nonroot rows has at least three occurrences, so $\ell \geq 3(k-1)$. Every doubly stochastic matrix A on this expanded face satisfies*

$$\frac{\operatorname{per} A}{2^{1-k-\ell}} \geq \frac{8192}{6561}\left(\frac{177147}{131072}\right)^{\ell-3(k-1)} \geq \frac{8192}{6561} > 1. \tag{5.1}$$

The same holds with a distinguished column.

Proof. It suffices first to take positive occurrence flows x. For every nonroot row i let S_i be its set of occurrences and $r_i = |S_i| \geq 3$. Thus $0 < x_e < 1$ for $e \in S_i$, as required by the divisions below. Eliminating the private paths yields the $k \times k$ nonnegative matrix

$$M_{vj} = \sum_{e:vj} x_e, \qquad M_{ij} = \sum_{\substack{e \in S_i \\ \operatorname{col}(e)=j}} x_e^2 \prod_{f \in S_i \setminus \{e\}} (1 - x_f) \quad (i \neq v). \tag{5.2}$$

The permanent expansion selects one occurrence in every core row and column. A selected long occurrence contributes x_e^2 and an unselected one $1 - x_e$. Hence $\operatorname{per} A = \operatorname{per} M$ exactly, including parallel occurrences.

For $y_j > 0$, the feasible-flow condition $\sum_{e \in S_i} x_e = 1$ lets us apply weighted arithmetic-geometric mean with weights x_e in each nonroot row:

$$\begin{aligned} (My)_i &= \prod_{f \in S_i} (1 - x_f) \sum_{e \in S_i} \frac{x_e^2 y_{\operatorname{col}(e)}}{1 - x_e} \\ &\geq \prod_{e \in S_i} x_e^{x_e} (1 - x_e)^{1 - x_e} \prod_{e \in S_i} y_{\operatorname{col}(e)}^{x_e}. \end{aligned} \tag{5.3}$$

The root row contributes at least $\prod_j y_j^{M_{vj}}$. Multiplying (5.3) over the rows and using the column flow sums cancels all powers of y_j in the capacity denominator. Therefore

$$\operatorname{Cap}(M) := \inf_{y_j > 0} \frac{\prod_i (My)_i}{\prod_j y_j} \geq \exp\left(- \sum_{i \neq v} \sum_{e \in S_i} h_2(x_e) \right). \tag{5.4}$$

Put $g(d) = ((d-1)/d)^{d-1}$ for $d \geq 2$, with $g(1) = 1$. Gurvits's degree-sensitive capacity inequality [7, Theorem 2.4 and Corollary 2.7(a)] applies to the homogeneous polynomial $\prod_i (My)_i$:

every factor has positive real part when all variables do, so the polynomial is half-plane stable. Let Δ_j be the number of nonzero entries in column j of M. Define c_j as the number of nonroot occurrences ending at j, plus one if v meets j. Then $\Delta_j \le c_j$, even with parallel occurrences. Gurvits's product depends on the current column order through $\min\{\Delta_j, j\}$. We use that order, and since g decreases [7, Definition 2.1(6)],

$$g(\min\{\Delta_j, j\}) \ge g(\min\{c_j, j\}).$$

The inequality therefore gives

$$\operatorname{per} M \ge \operatorname{Cap}(M) \prod_{j=2}^{k} g(\min\{c_j, j\}). \tag{5.5}$$

Every $c_j \ge 3$, because K has at least three distinct neighbors at each column. If $\rho_v = |N_K(v)|$ counts distinct root neighbors after aggregation, then

$$\sum_{j=1}^{k} c_j = \ell + \rho_v \le \ell + k. \tag{5.6}$$

Concavity of binary entropy gives $\sum_{e\in S_i} h_2(x_e) \le r_i h_2(1/r_i)$. Define the row factor A_d and column factor C_d by

$$A_d = \frac{2^d (d-1)^{d-1}}{d^d}, \qquad C_d = 2\left(\frac{d-1}{d}\right)^{d-1}. \tag{5.7}$$

Equivalently, $A_d = 2^d \exp[-d\, h_2(1/d)]$ and $C_d = 2g(d)$. Since $\sum_{i\ne v} r_i = \ell$, normalization of (5.4)–(5.5) by $2^{1-k-\ell}$ gives

$$\frac{\operatorname{per} A}{2^{1-k-\ell}} \ge \prod_{i\ne v} A_{r_i} \prod_{j=3}^{k} C_{\min\{c_j, j\}}. \tag{5.8}$$

The omitted $j = 2$ factor is exactly $C_2 = 1$.

For real $t > 1$, direct differentiation gives

$$(\log A_t)'' = \frac{1}{t(t-1)}, \qquad (\log C_t)'' = \frac{1}{t^2(t-1)}. \tag{5.9}$$

For each convex function in (5.9), the secant line through $t = 3, 4$ lies below the graph at every integer $d \ge 3$, with equality at $d = 3, 4$:

$$A_d \ge \frac{32}{27}\left(\frac{729}{512}\right)^{d-3}, \qquad C_d \ge \frac{8}{9}\left(\frac{243}{256}\right)^{d-3}. \tag{5.10}$$

Also C_d decreases and $C_d \le 8/9$ for $d \ge 3$. Since $\min\{c_j, j\} \le c_j$,

$$\prod_{j=3}^{k} C_{\min\{c_j, j\}} \ge \left(\frac{9}{8}\right)^2 \prod_{j=1}^{k} C_{c_j}. \tag{5.11}$$

The row and column products in (5.8) have the separate bounds

$$\prod_{i\ne v} A_{r_i} \ge \left(\frac{32}{27}\right)^{k-1} \left(\frac{729}{512}\right)^{\ell-3(k-1)}, \qquad \prod_{j=1}^{k} C_{c_j} \ge \left(\frac{8}{9}\right)^{k} \left(\frac{243}{256}\right)^{\ell-2k}.$$

The second uses $\sum_j c_j \le \ell + k$ from (5.6) and $243/256 < 1$. Combining these bounds with (5.8) and (5.11) gives

$$\begin{aligned}\frac{\operatorname{per} A}{2^{1-k-\ell}} &\ge \left(\frac{9}{8}\right)^2 \left(\frac{32}{27}\right)^{k-1} \left(\frac{729}{512}\right)^{\ell-3k+3} \left(\frac{8}{9}\right)^k \left(\frac{243}{256}\right)^{\ell-2k} \\ &= \frac{8192}{6561} \left(\frac{177147}{131072}\right)^{\ell-3(k-1)}.\end{aligned} \tag{5.12}$$

The last equality uses $(32/27)(8/9)(243/256) = 1$ to cancel the $k-3$ extra factors. It uses $(729/512)(243/256) = 177147/131072$ for the power $\ell - 3(k-1)$, leaving the constant $8192/6561$. Continuity of the permanent extends the bound to the closed face. Transposition proves the column version. □

The leading constant is sharp: for $K = K_{3,3}$ with any root row and $x_e = 1/3$ on every occurrence, $k = 3$, $\ell = 6$, and $\operatorname{per} A = 32/6561$, giving equality in (5.1).

Corollary 5.2. *Under the hypotheses of Lemma 5.1, let v be the distinguished row or column and L the occurrences of K not incident to v. Every feasible occurrence flow x on K satisfies*

$$G_{K,L}(x) \ge \log \frac{8192}{6561} > 0.$$

Proof. Apply Lemma 3.1 to the whole-face bound (5.1), then use (3.7). The column case follows by transposition. □

6 Tight-cut leaves and a strict gap

We find a leaf by choosing a smallest shore, then bound its contribution to (4.5).

Lemma 6.1 (a cut-free core)**.** *Let K be FI of order $k \ge 3$ with no nontrivial tight cut. Then every vertex has at least three distinct neighbors, and every occurrence is removable.*

Proof. A vertex with exactly two distinct neighbors forms, together with them, a nontrivial tight shore. Strict Hall excludes degree one. Thus all distinct degrees are at least three. If a nonempty row set S has $|S| \le k-2$, strict Hall gives $|N(S)| \ge |S| + 1$. Equality would make $S \cup N(S)$ a nontrivial tight shore: every perfect matching sends the $|S|$ rows of S into $N(S)$, leaving exactly one of its $|S|+1$ columns to be matched across the cut. Thus $|N(S)| \ge |S| + 2$. Deleting one occurrence can remove at most one neighbor, and hence strict Hall remains true. If $|S| = k-1$, each column has at least two neighbors in S, so one deletion cannot remove that column from $N(S)$. Thus the deletion remains FI. The column version is the same. □

Lemma 6.2 (minimum-shore leaf)**.** *Let H be an FI core with a nontrivial tight cut. Choose a shore X of minimum cardinality among all nontrivial tight-cut shores of H, and close it to a factor K with one marker. Then K is a leaf, and each nonmarker vertex has the same occurrence degree as in H.*

Proof. If K had a nontrivial tight cut, take its shore Z avoiding the marker. Then $3 \le |Z| \le |X|-2$, since the other shore in K contains the marker and at least two vertices of X; the complement of Z in H also has at least three vertices. Every parent matching projects to a factor matching, preserving the number of edges crossing Z. Hence Z would be a smaller nontrivial tight shore of H, a contradiction. Thus K is a leaf. The marker replaces the opposite shore's crossing endpoints without changing any incidence count at a vertex of X, proving the degree assertion. □

When H comes from Lemma 2.5, each nonmarker vertex of K therefore has occurrence degree at least three.

If K has order $k \geq 3$, Lemma 6.1 makes all its occurrences removable. Thus the correction set N in (4.5) is empty, every nonmarker occurrence is long in B, and Corollary 5.2 gives

$$G_B \geq \log(8192/6561) > 0. \tag{6.1}$$

The case $k = 2$ requires a separate argument because the marker may have only two distinct neighbors.

Lemma 6.3 (four-vertex leaf)**.** *Let H be an FI core of minimum occurrence degree at least three, and let K be an order-two leaf factor of a nontrivial tight cut of H. Its marker may be a row or a column. Form B, N as in Section 4 with $H_1 = K$, and let $\ell_B = |L_B|$. At every positive parent flow,*

$$G_B - \sum_{e \in N} \delta(x_e) > 0. \tag{6.2}$$

Moreover, $\ell_B \geq 2$ and $|N| \leq 1$.

Proof. The underlying simple graph is $K_{2,2}$. Let the marker bundles have multiplicities $a_1, a_2 \geq 1$, and the internal bundles have multiplicities $b_1, b_2 \geq 1$. By Lemma 6.2, the three nonmarker vertices retain their occurrence degrees from H; the minimum-degree hypothesis therefore gives

$$a_1 + b_1 \geq 3, \qquad a_2 + b_2 \geq 3, \qquad b_1 + b_2 \geq 3. \tag{6.3}$$

An occurrence is removable exactly when its bundle has multiplicity at least two. Let the aggregate flow of the first marker bundle be t, so the second has flow $1-t$. If the marker is a column, the two rows have aggregated distributions $(t, 1-t)$ and $(1-t, t)$. If the marker is a row, the marker row and the other row have those same distributions. The aggregated graph has two perfect matchings with probabilities t and $1-t$, so its matching entropy is $h_2(t)$, while its row entropies sum to $2h_2(t)$. The label terms cancel by Lemma 3.2, giving $F_K = -h_2(t)$, and

$$G_B = \delta(t) + \sum_{e \text{ internal long}} \delta(x_e). \tag{6.4}$$

Only singleton marker bundles can contribute to N. If both are singletons, the opposite factor H_2 has order at least two by nontriviality of the cut, and its marker has occurrence degree two. Deleting either label leaves that marker with degree one, so neither label is removable in H_2 and N is empty. Otherwise there is at most one singleton marker bundle. Therefore $|N| \leq 1$, and its possible term is $\delta(t)$, since $\delta(1-t) = \delta(t)$.

Finally, $b_1 + b_2 \geq 3$ implies an internal bundle of multiplicity at least two, giving $\ell_B \geq 2$. Its aggregate flow lies strictly between zero and one at a positive flow. Its occurrence flows cannot all equal $1/2$, so at least one internal $\delta(x_e)$ is strictly positive. By (6.4),

$$G_B - \sum_{e \in N} \delta(x_e) \geq \sum_{e \text{ internal long}} \delta(x_e) > 0,$$

which proves (6.2). □

Let n, n_R be the orders of the parent and remainder expansions. The core-order identity (4.1) and the occurrence identity (4.4) give the exact count

$$n = (k-1) + \ell_B + n_R - |N|. \tag{6.5}$$

For $k \geq 3$, $N = \varnothing$; for $k = 2$, Lemma 6.3 gives $\ell_B \geq 2$ and $|N| \leq 1$. Thus in both cases $n_R < n$. For example, take a six-cycle core and duplicate the three edges of one perfect matching. Choosing one row and its two neighbors as the tight shore gives $k = 2$, $\ell_B = 2$, $n_R = 6$, and $N = \varnothing$, so (6.5) reads $9 = (2-1) + 2 + 6$. The remainder is ND by the converse direction of Lemma 2.5, which applies to the canonical expansion of any FI core, including one whose marker has degree two.

7 Proof of Theorem 1.1

We prove the theorem simultaneously for every ND support, by induction on matrix order. Although Theorem 1.1 starts at $n = 2$, we include order one in the induction: by convention its sole ND support is the singleton, with permanent $1 = 2^{1-1}$. In order two, the only ND support is $K_{2,2}$, a cycle support. More generally, on a cycle support every doubly stochastic matrix is, after independent row and column permutations, $tI_n + (1-t)P_n$. Its permanent is $t^n + (1-t)^n$, uniquely minimized at $t = 1/2$ for $n \geq 2$. This proves the theorem on all cycle faces. For a noncycle support, the induction has five branches:

1. An ND-preserving degree-two contraction, handled by Lemma 2.2 and induction.
2. A terminal core of order $h = 1$, handled by (7.2).
3. A terminal core of order $h = 2$, handled by the $K_{2,2}$ calculation below.
4. A terminal core of order $h \geq 3$ without a nontrivial tight cut, handled by Lemma 6.1 and Corollary 5.2.
5. A terminal core of order $h \geq 3$ with a nontrivial tight cut, handled by Lemmas 6.2–6.3, (4.5), and induction.

We use induction in two ways. An ND-preserving degree-two contraction gives a direct bound. For a smaller canonical ND remainder R, Lemma 2.1 places the minimum of $\Omega(R)$ at full support, so the inductive bound holds on the whole face. Lemma 3.1 and (3.7) then give $G_R \geq 0$ at every inherited marginal.

Assume the claim below order $n \geq 3$, and fix an ND support D of order n. The compact face $\Omega(D)$ has a permanent minimizer A, whose support is D by Lemma 2.1.

Minc's edge bound supplies degree-two rows and columns. If some degree-two line has an ND contraction, choose one. Lemma 2.2 and the induction hypothesis give

$$2 \operatorname{per} A = \operatorname{per} C(A) \geq 2^{2-n}, \quad \text{so} \quad \operatorname{per} A \geq 2^{1-n}. \tag{7.1}$$

If no degree-two contraction is ND, D is terminal and Lemma 2.5 writes it as $D_{\mathrm{can}}(H)$ with minimum core occurrence degree at least three. We show $G_{H,L(H)}(x) > 0$ at every positive core flow x.

If $h = 1$, the core consists of a bundle of $q \geq 3$ occurrences, all removable and hence long in $D_{\mathrm{can}}(H)$. Here $F_H = 0$, and entropy concavity yields

$$G_{H,L(H)} = \sum_e \delta(x_e) \geq q\delta(1/q) = \log \frac{2^q (q-1)^{q-1}}{q^q} > 0. \tag{7.2}$$

By (3.3) and (3.7), (7.2) gives $\operatorname{per} X \geq (q-1)^{q-1}/q^q$ throughout this face. At the uniform flow $x_e = 1/q$, its q matchings each have weight $q^{-2}(1-1/q)^{q-1}$, so equality holds and the face minimum is $(q-1)^{q-1}/q^q = 2^{-q} A_q$.

If $h = 2$, the underlying core is $K_{2,2}$. After aggregating parallel occurrences and labelling the columns, its bundle-flow matrix is $\left(\begin{smallmatrix} t & 1-t \\ 1-t & t \end{smallmatrix}\right)$. The entropy calculation (6.4) without a marker gives $G_{H,L(H)} = \delta(t) + \sum_{e \text{ long}} \delta(x_e) > 0$: the degree-three premise forces a parallel bundle, and at a positive flow some occurrence in that bundle has $\delta(x_e) > 0$.

Let $h \geq 3$. If H has no nontrivial tight cut, Lemma 6.1 makes every occurrence removable. Shorten any one row star. Corollary 5.2 gives a gap at least $\log(8192/6561)$ for the shortened expansion. Making the star paths long again adds only nonnegative δ-terms, so $G_{H,L(H)} > 0$.

Otherwise choose a minimum tight shore. Lemma 6.2 gives a leaf factor K, whose expansion is $B = D(K, L_B)$ as in Section 4. Equation (6.5) gives a smaller canonical ND remainder R. Applying the induction hypothesis to the ND support R of order $n_R < n$ bounds its full-support minimizer, hence all of $\Omega(R)$ by Lemma 2.1. Lemma 3.1 and (3.7) then give $G_R \geq 0$ at the inherited positive flow. For a leaf of order at least three, (6.1) and $N = \varnothing$ make (4.5) strictly positive. For a four-vertex leaf, Lemma 6.3 makes (4.5) strictly positive. Thus $G_{H,L(H)}(x) > 0$ in every terminal case. By (3.7) and the first inequality of (3.3), the permanent at an interior terminal point is strictly above 2^{1-n}. This completes the lower-bound induction.

For equality, let A be ND with support D and $\operatorname{per} A = 2^{1-n}$. The lower bound just proved on $\Omega(D)$ makes A a face minimizer, so the preceding contraction/terminal analysis applies. The terminal branch is strict. In the contraction branch, equality in (7.1), together with the induction hypothesis, makes $C(A)$ equal to $\frac{1}{2}(I_{n-1} + P_{n-1})$ after independent row and column permutations. Its support has $2(n-1)$ edges. Lemma 2.3 says the two merged rows have disjoint surviving neighborhoods, so contraction removes exactly two support edges and D has $2n$ edges. Every FI line has degree at least two, and D is connected; hence every line has degree exactly two and D is a cycle support. Its unique minimizing weights are $1/2$.

Finally, a fixed noncycle ND face $\Omega(D)$ is compact and has a minimizer. By Lemma 2.1 that minimizer has support exactly D; equality at its minimum would then contradict the classification above. Thus its minimum is strictly greater than 2^{1-n}. □

Acknowledgments

The proof of this conjecture was carried out by OpenAI's GPT models (Sol and Astra) and Anthropic's Claude Fable models, under the guidance of the author. The author has reviewed the resulting proof arguments. Responsibility for the final text rests with the author.

References

[1] R. A. Brualdi and M. B. Hedrick, "A unified treatment of nearly reducible and nearly decomposable matrices," *Linear Algebra and its Applications* **24** (1979), 51–73. doi:10.1016/0024-3795(79)90146-0.

[2] R. A. Brualdi and B. L. Shader, "Minimum permanents on special faces of the polytope of doubly stochastic matrices," *Linear Algebra and its Applications* **201** (1994), 103–111. doi:10.1016/0024-3795(94)90109-0.

[3] G.-S. Cheon and I. M. Wanless, "An update on Minc's survey of open problems involving permanents," *Linear Algebra and its Applications* **403** (2005), 314–342. doi:10.1016/j.laa.2005.02.030.

[4] J. Edmonds, L. Lovász, and W. R. Pulleyblank, "Brick decompositions and the matching rank of graphs," *Combinatorica* **2** (1982), 247–274. doi:10.1007/BF02579233.

[5] T. H. Foregger, "On the minimum value of the permanent of a nearly decomposable doubly stochastic matrix," *Linear Algebra and its Applications* **32** (1980), 75–85. doi:10.1016/0024-3795(80)90008-7.

[6] T. H. Foregger, "Minimum permanents of multiplexes," *Linear Algebra and its Applications* **87** (1987), 197–211. doi:10.1016/0024-3795(87)90167-4.

[7] L. Gurvits, "Van der Waerden/Schrijver–Valiant like conjectures and stable (aka hyperbolic) homogeneous polynomials: One theorem for all," *Electronic Journal of Combinatorics* **15** (2008), Research Paper 66. doi:10.37236/790; arXiv:0711.3496.

[8] D. J. Hartfiel, "On constructing nearly decomposable matrices," *Proceedings of the American Mathematical Society* **27** (1971), 222–228. doi:10.1090/S0002-9939-1971-0268062-4.

[9] M. B. Hedrick, "The permanent at a minimum on certain classes of doubly stochastic matrices," *Bulletin of the American Mathematical Society* **80** (1974), 836–838. doi:10.1090/S0002-9904-1974-13532-9.

[10] L. Lovász and M. D. Plummer, "On minimal elementary bipartite graphs," *Journal of Combinatorial Theory, Series B* **23** (1977), 127–138. doi:10.1016/0095-8956(77)90062-4.

[11] L. Lovász and M. D. Plummer, *Matching Theory*, Annals of Discrete Mathematics **29**, North-Holland, Amsterdam, 1986.

[12] H. Minc, "Nearly decomposable matrices," *Linear Algebra and its Applications* **5** (1972), 181–187. doi:10.1016/0024-3795(72)90027-4.

[13] H. Minc, "Theory of permanents 1978–1981," *Linear and Multilinear Algebra* **12** (1983), 227–263. doi:10.1080/03081088308817488.

[14] R. Sinkhorn and P. Knopp, "Concerning nonnegative matrices and doubly stochastic matrices," *Pacific Journal of Mathematics* **21** (1967), 343–348. doi:10.2140/pjm.1967.21.343.

[15] R. Sinkhorn and P. Knopp, "Problems involving diagonal products in nonnegative matrices," *Transactions of the American Mathematical Society* **136** (1969), 67–75. doi:10.1090/S0002-9947-1969-0233830-7.