Joint inclusion bounds for random greedy matroid bases

Scott Schubr
Luminary Oddities

20 November 2026
Archive edition 2

Abstract

Assign independent priorities from one common atomless law to the elements of a finite matroid and construct its minimum-weight basis greedily. We bound every joint inclusion probability in terms of the corresponding individual marginals, with a factorial coefficient and an explicit marginal-dependent improvement. The pair case gives a universal factor two for every finite matroid. For random minimum spanning trees this establishes the global factor two left open in Gupta's Problem 7.1, improving his factor eight and placing the optimal constant in $[13938405/13872419, 2]$, where $13938405/13872419 \approx 1.00476$. The proof first conditions on all unmarked priorities. Forcing a marked element into the basis and contracting it yields a conditional cylinder inequality; conditional Jensen and Hölder inequalities then return to the original one-element closure thresholds. Bernoulli sprinkling bounds these thresholds in convex order by a beta distribution with the same mean. We obtain a sharp moment hierarchy, characterize equality by rank-one direct summands, and derive exclusion and occupancy bounds. Finite-cylinder limits transfer the estimates to countable finitary-matroid restriction limits and to both free and wired minimal spanning forests. On transitive unimodular graphs this gives a uniform exponential bound for the wired-forest degree distribution.

Keywords. Matroid; random greedy basis; minimum spanning forest; joint inclusion; convex order.
2020 Mathematics Subject Classification. 05B35, 60C05, 60E15.

1 Introduction and main result

Let M be a finite matroid on a ground set E. Assign independent priorities $(U_e)_{e\in E}$ from one common atomless probability law. Scan the elements in increasing priority and accept an element when adding it preserves independence. The resulting basis B minimizes the sum of priorities. Only their relative order matters, so its law is also obtained from independent uniform priorities on $[0,1]$. We call this the *random greedy basis law*. Write

$$p_e = \mathbb{P}(e \in B), \qquad p_S = \mathbb{P}(S \subseteq B).$$

For $\ell \geq 1$ and $p = (p_1, \dots, p_\ell) \in [0,1]^\ell$, put

$$F_\ell(p) = \frac{\ell! \prod_{i=1}^{\ell} p_i}{\left(\prod_{i=1}^{\ell} \prod_{j=1}^{\ell-1} (1 + j p_i)\right)^{1/\ell}}. \tag{1}$$

Empty products equal one, and $F_0 = 1$. Each F_ℓ is continuous on its closed cube.

Theorem 1.1 (Joint inclusion)**.** *For every finite matroid and every set of distinct elements* $S = \{e_1, \dots, e_\ell\}$*,*

$$\mathbb{P}(S \subseteq B) \le F_\ell(p_{e_1}, \dots, p_{e_\ell}) \le \ell! \prod_{i=1}^{\ell} p_{e_i}. \tag{2}$$

The assertion includes loops, coloops and nonrepresentable matroids. For $\ell \ge 2$*, equality in the first inequality holds exactly when* S *contains a loop or every member of* S *is a coloop.*

Corollary 1.2 (Pairs)**.** *For distinct* e, f*,*

$$\mathbb{P}(e, f \in B) \le \frac{2p_e p_f}{\sqrt{(1+p_e)(1+p_f)}} \le 2p_e p_f. \tag{3}$$

Equality in the first inequality holds precisely when one element is a loop or both elements are coloops.

The graphical specialization is a minimum spanning tree on a connected finite graph, with loops and parallel edges allowed. Gupta [1] proves the universal pair bound $8p_e p_f$ in that setting, answering a question of R. Lyons, and shows that the optimal constant is at least $13938405/13872419 \approx 1.00476$. His Problem 7.1 asks for the least such constant; he notes that his small-marginal estimate gives $2 + o(1)$ but not a global factor two. Corollary 1.2 supplies that global factor, so the optimal constant lies in $[13938405/13872419, 2]$, and the same bound holds for every finite matroid. The same corollary places the simple-graph supremum of his Problem 7.2 in $[13938405/13872419, 2]$. The moment hierarchy below refines Gupta's sharp inequality $\mathbb{E}T^2 \le 2(\mathbb{E}T)^2$ to $\mathbb{E}T^2 \le 2(\mathbb{E}T)^2/(1+\mathbb{E}T)$, which is attained exactly by the parallel classes of his Remark 3.2, where his bound is approached only as $\mathbb{E}T \to 0$. His proof combines conditional Harris correlation with a moment estimate for connection thresholds, then compares thresholds in the graph with both marked edges deleted. Here we keep the original single-element thresholds. Conditional expectation removes the two deletion-comparison losses, and a stronger sprinkling identity retains the denominator in (3). The conditional argument also extends to arbitrarily many marked elements and to every finite matroid.

Random greedy bases have classical antecedents. Mayhew [2] studies their exact distribution through linear extensions of fundamental-circuit partial orders and identifies matroids for which the distribution is uniform. Huh, Schröter and Wang [3] prove correlation bounds for the law assigning a basis b mass proportional to $\prod_{e \in b} w_e$, with prescribed positive weights. That product-weighted law is a different model. For example, the graphical product-weighted law has negatively correlated distinct edges, whereas positive correlations occur for random minimum spanning trees [4, 1]. The present estimates concern the latter random-priority law.

Makowiec [5] obtains exact asymptotics for inclusion of each fixed forest in the minimum spanning tree of K_n. Those asymptotics identify much finer complete-graph constants. Our theorem gives finite, marginal-based control uniformly over matroids and marked sets. The shared tools include threshold representations and Hölder's inequality; the conditional forcing-and-contraction step supplies the all-orders return used here.

The threshold comparison has a useful independent formulation. If T_e is the first priority level at which the other elements span e, then $\mathbb{E}T_e = p_e$. Its stop-loss function is bounded by that of $\mathrm{Beta}(1, p_e^{-1} - 1)$. Mean residual life, convex order and the linear-residual-life beta family belong to classical reliability theory; see [6]; for a recent use of mean residual life in variance comparisons see [7]. We give direct proofs of the specific threshold admission, all moments and equality structure.

The argument proceeds from closure thresholds and sprinkling to a conditional cylinder inequality. It then derives (2), dual exclusion bounds and occupancy estimates. Finally, continuous

finite-cylinder passage gives countable restriction limits and the classical free and wired minimal spanning forests of Lyons, Peres and Schramm [4].

2 Greedy bases and conditional monotonicity

We use standard matroid definitions [8]. The rank of $A \subseteq E$ is the maximum size of an independent subset of A, and

$$\mathrm{cl}(A) = \{e \in E : r(A \cup \{e\}) = r(A)\}$$

is its closure. A circuit is a minimal dependent set. A loop is in no basis and a coloop is in every basis.

Lemma 2.1 (Greedy characterization)**.** *For distinct priorities the greedy algorithm produces the unique minimum-weight basis. An element e is accepted exactly when*

$$e \notin \mathrm{cl}\{g \neq e : U_g < U_e\}. \tag{4}$$

Lowering the priority of an accepted element leaves the entire minimum basis unchanged.

Proof. After any prefix of the priority order, the accepted set is a maximal independent subset of that prefix and hence a basis of it. It has the same closure as the prefix, proving (4). The usual exchange argument proves optimality: if a basis omits the next greedy element, its fundamental circuit after adding that element contains an element outside the already accepted prefix. Exchanging such an element preserves a basis and replaces it by an element of smaller priority. Repeating compares the greedy basis with any other basis and proves strict optimality when they differ.

If a minimum basis contains e, lowering U_e subtracts the same amount from the cost of every basis containing e, and subtracts nothing from costs of bases omitting e. Its unique minimum therefore remains unchanged. □

It follows from (4) that acceptance of e is nonincreasing in U_e and nondecreasing in every competitor priority. Increasing a competitor's priority removes it from some lighter sets, and closure is monotone under inclusion.

We also record the exact product-measure inequality needed below.

Lemma 2.2 (Harris inequality)**.** *For a finite family of independent real coordinates, bounded coordinatewise increasing functions f, g satisfy $\mathbb{E}(fg) \geq \mathbb{E}f\,\mathbb{E}g$. If one function is increasing and the other decreasing, the inequality reverses.*

Proof. For one coordinate X and an independent copy X', twice the covariance is

$$\mathbb{E}\big[(f(X) - f(X'))(g(X) - g(X'))\big] \geq 0.$$

Induct on the number of coordinates. Conditional on the last coordinate, apply the induction hypothesis to the others. Their conditional means are increasing functions of the last coordinate, so the one-coordinate result completes the induction. Apply the increasing case to $f, -g$ for the reversed inequality. This is the classical product-measure inequality of Harris [9]. □

3 Closure thresholds and beta comparison

For uniform priorities define

$$T_e = \inf\{t \in [0,1] : e \in \mathrm{cl}\{g \neq e : U_g < t\}\}. \tag{5}$$

The infimum of an empty family is set to 1; a loop has threshold 0. This threshold depends only on priorities other than U_e. Equation (4) gives, almost surely,

$$\mathbf{1}_{\{e \in B\}} = \mathbf{1}_{\{U_e < T_e\}}, \qquad \mathbb{E}T_e = p_e. \tag{6}$$

Equality of the two independent variables has probability zero, including the endpoint conventions.

Lemma 3.1 (Multiplicative sprinkling)**.** *Let $q(t) = \mathbb{P}(T_e > t)$ for $0 \le t < 1$, and $p = \mathbb{E}T_e$. Then*

$$q(s + (1-s)t) \le q(s)q(t), \qquad 0 \le s, t < 1, \tag{7}$$

and

$$R(s) := \int_s^1 q(u)\,du \le p(1-s)q(s). \tag{8}$$

Proof. The function $q(t)$ is the probability that a Bernoulli(t) subset of $E \setminus \{e\}$ fails to span e. Take independent Bernoulli subsets of densities s, t. Their union has density $s+(1-s)t$, independently at each element. Failure of the union to span e implies failure of both subsets, by closure monotonicity. Independence gives (7). Integrating it in t and substituting $u = s + (1-s)t$ gives (8), because $\int_0^1 q(t)\,dt = p$. □

The argument applies to any increasing spanning property of finite subsets. The matroid enters through the greedy admission (6).

Theorem 3.2 (Threshold distribution and moments)**.** *For $0 < p_e < 1$, let $Y_{p_e} \sim \mathrm{Beta}(1, p_e^{-1} - 1)$. Then T_e is smaller than Y_{p_e} in convex order. More explicitly, for $0 \le s \le 1$,*

$$\mathbb{E}(T_e - s)_+ \le p_e(1-s)^{1/p_e}. \tag{9}$$

For every integer $m \ge 1$,

$$\mathbb{E}T_e^m \le \frac{m! p_e^m}{\prod_{j=1}^{m-1}(1 + jp_e)}. \tag{10}$$

The cases $p_e = 0, 1$ have thresholds $0, 1$ almost surely and satisfy (10)*.*

Proof. On the interval where $R > 0$, equation (8) and $R' = -q$ give

$$\frac{R'(s)}{R(s)} \le -\frac{1}{p_e(1-s)}$$

almost everywhere. Integrate from 0, using $R(0) = p_e$, to obtain (9); after the first zero of R the bound is immediate. The beta variable has survival function $(1-s)^{p_e^{-1}-1}$ and stop-loss function $p_e(1-s)^{1/p_e}$. Both variables have mean p_e.

For completeness, stop-loss domination with equal means implies convex order directly on $[0,1]$: every piecewise-linear convex function is an affine function plus a finite nonnegative linear combination of $(x-s)_+$, and every continuous convex function on this compact interval is uniformly approximated by such functions. Thus all convex expectations have the stated order.

There is also a direct moment proof. Put $a_m = \mathbb{E}T_e^m = m\int_0^1 t^{m-1}q(t)\,dt$. For $m \ge 2$, multiply (8) by s^{m-2} and integrate. Tonelli's theorem gives

$$\frac{a_m}{m(m-1)} \le p_e\left(\frac{a_{m-1}}{m-1} - \frac{a_m}{m}\right), \qquad [1 + (m-1)p_e]a_m \le mp_e a_{m-1}.$$

Starting from $a_1 = p_e$ proves (10). The mean identity (6) handles the deterministic endpoints. □

4 A conditional cylinder inequality

Fix a marked set $S \subseteq E$ and condition on all outside priorities $W = (U_g)_{g \notin S}$. Write

$$a_e(W) = \mathbb{P}(e \in B \mid W), \qquad e \in S.$$

The following statement allows arbitrary fixed distinct outside priorities in $[0, 1]$, while marked priorities remain independent uniforms.

Proposition 4.1 (Conditional marked inclusion)**.** *For every such outside realization,*

$$\mathbb{P}(S \subseteq B \mid W) \leq \prod_{e \in S} a_e(W). \tag{11}$$

Proof. Induct on $|S|$. The singleton case is equality. If S is dependent, its joint inclusion probability is zero. Otherwise choose $e \in S$ and force its priority to -1. It is a nonloop, so it is accepted first, and the remaining algorithm is greedy on the contracted matroid M/e, with all other priorities unchanged. Let A_0 be the event that this contracted run accepts every element of $S \setminus \{e\}$.

On the event that the original basis contains e, Lemma 2.1 says that lowering its priority to -1 leaves the whole basis unchanged. Consequently

$$\mathbf{1}_{\{S \subseteq B\}} = \mathbf{1}_{\{U_e < T_e\}} \mathbf{1}_{A_0}$$

almost surely. Integrating U_e first yields

$$\mathbb{P}(S \subseteq B \mid W) = \mathbb{E}[T_e \mathbf{1}_{A_0} \mid W]. \tag{12}$$

As a function of the other marked priorities, T_e is increasing. The event A_0 is decreasing in each of them: whenever all remaining marked elements belong to the contracted minimum basis, lowering any one's priority preserves that entire basis by Lemma 2.1. Harris's inequality therefore gives

$$\mathbb{E}[T_e \mathbf{1}_{A_0} \mid W] \leq a_e(W) \mathbb{P}(A_0 \mid W),$$

where $a_e(W) = \mathbb{E}[T_e \mid W]$ follows by integrating its own independent priority.

Apply the induction hypothesis in M/e to the remaining marked set. If $b_f(W)$ denotes its contracted marginal, then $\mathbb{P}(A_0 \mid W) \leq \prod_{f \in S \setminus \{e\}} b_f(W)$. Finally, $b_f(W) \leq a_f(W)$: coupling by lowering U_e to -1 adds a competitor before f, and the closure characterization shows pointwise that this can only decrease acceptance of f. Combining these inequalities proves (11). □

Proof of Theorem 1.1. The cases $\ell = 0, 1$ are immediate; assume $\ell \geq 2$. Average (11) over W and use Hölder's inequality with $\ell = |S|$ identical exponents:

$$\mathbb{P}(S \subseteq B) \leq \mathbb{E} \prod_{e \in S} a_e(W) \leq \prod_{e \in S} \left(\mathbb{E} a_e(W)^\ell\right)^{1/\ell}.$$

Each $a_e(W)$ is the conditional expectation of the *original* threshold T_e. Conditional Jensen gives $\mathbb{E} a_e(W)^\ell \leq \mathbb{E} T_e^\ell$. Substitute (10) and take the product to obtain (2). □

5 Sharp threshold equality and pair strictness

A parallel class of nonloops is a set whose distinct members form two-element circuits. It is a rank-one direct summand when no circuit meets both the class and its complement.

Theorem 5.1 (Equality in threshold moments)**.** *For a nonloop e in a finite matroid, equality in* (10) *for any integer $m \geq 2$ holds exactly when e lies in a rank-one direct summand. If that summand has $k+1$ elements, then*

$$q(t) = (1-t)^k, \qquad p_e = \frac{1}{k+1}.$$

For $k \geq 1$ its threshold is the minimum of k independent uniforms; $k = 0$ is a coloop. Loops have threshold zero.

Proof. First take $m = 2$. The proof of its moment inequality integrates the nonnegative continuous deficit

$$D(s,t) = q(s)q(t) - q(s + (1-s)t)$$

against the strictly positive weight $1-s$ on $(0,1)^2$. Equality forces $D = 0$ there and hence everywhere on $[0,1)^2$. The nonloop reliability polynomial q is positive on $[0,1)$, since choosing no other element has positive probability. Therefore $h(x) = q(1-e^{-x})$ is continuous, positive and satisfies $h(x+y) = h(x)h(y)$, with $h(0) = 1$. It follows that $q(t) = (1-t)^c$ for some real $c \geq 0$; since q is a polynomial, $c = k$ is a nonnegative integer.

Differentiation at zero shows that $k = -q'(0)$ is the number of elements parallel to e: precisely those singleton subsets span e. Call this set of peers P. Nonspanning implies that none of P is chosen, giving $q(t) \leq (1-t)^{|P|}$. If a circuit C of size at least three contained e, it would contain no peer of e, by minimality of circuits. For every $0 < t < 1$, selecting $C \setminus \{e\}$ while selecting no peer has positive probability and spans e, making this inequality strict. Thus every circuit containing e is a parallel pair. Replacing a parallel element by e in a circuit preserves the circuit property, since parallel elements give the same closure after either is adjoined. A circuit meeting this class and its complement would therefore give a longer circuit containing e. The class is a direct summand.

Conversely an isolated class of size $k+1$ has the stated minimum threshold and attains every moment bound. For general $m \geq 2$, equality in the iterated positive moment recurrence of Theorem 3.2 forces equality in the second-moment step, completing the characterization. □

Equality in Theorem 1.1 and Corollary 1.2. Fix $\ell \geq 2$. A marked loop or an all-coloop set gives equality. If every marked marginal is positive, equality in the proof of Theorem 1.1 forces equality in every threshold ℓth-moment bound. Theorem 5.1 therefore puts each marked element in a rank-one direct summand. Two marked elements in one parallel class give joint probability zero. Otherwise the classes are distinct, their choices are independent and the joint probability is the product of their marginals. Since

$$\prod_{j=1}^{\ell-1} (1 + jp_i) \leq \ell!,$$

with equality exactly when $p_i = 1$, the prefactor in (1) equals one exactly when every marked marginal is one. In a finite matroid every nonloop has positive greedy probability and every noncoloop has probability less than one: placing it first, or last after a basis omitting it, respectively, proves these assertions. This establishes the equality classification in Theorem 1.1.

The pair inequality and its equality statement are the case $\ell = 2$. □

6 Exclusion and occupancy

Corollary 6.1 (Joint exclusion)**.** *For every marked set $S = \{e_1, \dots, e_\ell\}$,*

$$\mathbb{P}(S \cap B = \varnothing) \le F_\ell(1 - p_{e_1}, \dots, 1 - p_{e_\ell}). \tag{13}$$

Proof. Complements of bases of M are bases of its dual M^*. Subtracting basis weight from the total weight over E shows that $E \setminus B$ is the maximum-weight dual basis. Reversing priorities by $U'_e = 1 - U_e$ turns this into the minimum-weight dual basis, since every dual basis has the same cardinality. Reversed uniform priorities are still iid, so Theorem 1.1 applies with marginals $1 - p_e$. □

Corollary 6.2 (Occupancy bounds)**.** *For $A \subseteq E$, set $Z = |B \cap A|$ and $\mu = \sum_{e \in A} p_e$. Then, for each integer $\ell \ge 0$,*

$$\mathbb{E}\binom{Z}{\ell} \le \mu^\ell. \tag{14}$$

If $\mu > 0$ and $0 \le z < 1/\mu$, then

$$\mathbb{E}(1+z)^Z \le \frac{1}{1 - z\mu}, \qquad \mathbb{P}(Z \ge r) \le \inf_{0 < z < 1/\mu} \frac{(1+z)^{-r}}{1 - z\mu}. \tag{15}$$

In particular $\mathbb{P}(Z \ge r) \le 2(1 + 1/(2\mu))^{-r}$. If $\mu = 0$, then $Z = 0$ almost surely.

Proof. Sum the weaker bound in (2) over ℓ-element subsets of A. The sum of their marginal products is the elementary symmetric polynomial $e_\ell((p_e)_{e \in A})$, and $\ell! e_\ell \le \mu^\ell$ by expanding μ^ℓ. This proves (14). Expand $(1+z)^Z$ and bound its nonnegative finite sum by the geometric series. Markov's inequality gives (15); choose $z = 1/(2\mu)$ for the stated example. □

7 Countable limits and minimal spanning forests

Proposition 7.1 (Finite-cylinder passage)**.** *Suppose finite random greedy basis laws converge in finite-dimensional distributions on a countable labeled set E, with every finite marked set eventually present in their ground sets. Let $X \subseteq E$ have the limit law and $p_e = \mathbb{P}(e \in X)$. Then* (2) *and* (13) *hold for X and every finite marked set, with these limit marginals. If $A \subseteq E$ is countable and $\mu = \sum_{e \in A} p_e < \infty$, the occupancy conclusions* (14)–(15) *also hold.*

Proof. Cylinder probabilities and individual marginals converge. Continuity of F_ℓ on the closed cube passes both inequalities to the limit. Exhaust a countable A by finite sets. Its counts increase to Z, and Tonelli gives $\mathbb{E}Z = \mu$, so Z is finite almost surely. Monotone convergence passes each binomial moment and each exponential bound, using the common finite bound $(1 - z\mu)^{-1}$. □

7.1 Finitary matroid restrictions

A matroid on a countable ground set is *finitary* when all its circuits are finite.

Proposition 7.2. *Let M be finitary on countable E, with iid atomless priorities. Define*

$$X = \{e : e \text{ is not the largest-priority member of any circuit}\}. \tag{16}$$

For any finite exhaustion $E_n \uparrow E$, the greedy bases of $M|E_n$ converge almost surely on every finite cylinder to X. The set X is independent and satisfies the conclusions of Proposition 7.1.

Proof. Once $e \in E_n$, its acceptance indicator decreases with n by closure monotonicity. Rejection is witnessed by a circuit containing e and only lighter elements. Every such circuit is finite and eventually belongs to E_n, so the indicator limit is exactly (16). Countability gives simultaneous convergence. A circuit contained in X would have a largest-priority member and contradict the definition; hence X is independent. Apply Proposition 7.1. □

In a countably infinite rank-one parallel class, every element has a lighter peer almost surely, so this restriction limit is empty. This example describes the possible loss of maximality in the restriction construction.

For the threshold statements choose independent uniform priorities, as permitted by order invariance. For each e, the finite thresholds decrease to

$$T_e = \inf_{C \ni e} \max_{f \in C \setminus \{e\}} U_f,$$

where C runs over finite circuits, the empty-family infimum is 1, and a loop's empty maximum is 0. The limit is independent of U_e, so $\mathbf{1}_{\{e \in X\}} = \mathbf{1}_{\{U_e < T_e\}}$ almost surely and $\mathbb{E}T_e = p_e$. Bounded convergence passes every moment and stop-loss inequality of Theorem 3.2 to this threshold.

7.2 Free and wired minimal spanning forests

Let $G = (V, E)$ be connected and locally finite, with iid uniform labels. The *free minimal spanning forest* F_f consists of edges that are not the largest-labeled member of a finite cycle. The *wired minimal spanning forest* F_w consists of edges that are the least-labeled member of some cut $\partial_E W$ with $W \subseteq V$ finite. These are the classical free-cycle and wired finite-cut formulations [4, Section 3].

Theorem 7.3 (Minimal spanning forests)**.** *Both F_f and F_w satisfy the joint inclusion, joint exclusion and finite-mean occupancy conclusions above, using their respective actual edge marginals. Their original one-edge path or cut thresholds satisfy the beta comparison and moment bounds of Theorem 3.2.*

Proof. Choose finite connected induced vertex sets $V_n \uparrow V$. The free graph G_n contains all edges between vertices in V_n. Form G_n^w by identifying all outside vertices to one boundary vertex and discarding resulting loops at that vertex. Retain every original edge incident to V_n, with its separate label, so parallel boundary edges remain distinct. Local finiteness makes these finite graphs.

Their greedy bases are minimum spanning trees. For each eventually internal edge e, free acceptance decreases to the finite-cycle condition defining F_f. Wired acceptance has the equivalent finite cut criterion

$$e \in B_n^w \quad \Longleftrightarrow \quad e \text{ is least on some } \partial_E W, \text{ with } W \subseteq V_n.$$

Indeed, a finite tree edge belongs to a minimum spanning tree exactly when it is the unique least edge across some cut: its lighter-edge component gives the cut in one direction, and the exchange argument proves the other. The side avoiding the boundary vertex is a subset of V_n, and its cut is exactly the original cut because all boundary edges were retained. Thus wired indicators increase to the finite-cut condition defining F_w. This also reproduces the finite-exhaustion identification of [4, Proposition 3.1]. Proposition 7.1 now applies to both laws.

For threshold comparison, cap all thresholds at 1. Free thresholds are the bottleneck minima over paths connecting the endpoints in $G_n \setminus \{e\}$ and decrease. Wired thresholds are

$$T_e^{w,n} = \max_{\substack{W \subseteq V_n \\ e \in \partial_E W}} \min_{f \in \partial_E W \setminus \{e\}} U_f,$$

with empty minimum 1, and increase. The finite path-cut identity follows from the same minimum-tree criterion. Both limits depend only on labels other than U_e and describe acceptance by $U_e < T_e$ almost surely. Their means equal their respective edge marginals. Bounded convergence passes the finite moments and stop-loss bounds. Loops have threshold zero. These threshold conventions agree with [4, Lemma 3.2]. □

Corollary 7.4 (Uniform wired-forest degree tail)**.** *Let G be an infinite transitive unimodular locally finite graph and D_v the WMSF degree at a vertex. Then*

$$\mathbb{E}(1+z)^{D_v} \le \frac{1}{1-2z}, \qquad 0 \le z < \tfrac{1}{2}, \qquad \mathbb{P}(D_v \ge r) \le 2(5/4)^{-r}.$$

For $r > 2$ the optimized bound is

$$\mathbb{P}(D_v \ge r) \le \frac{r+1}{3}\left(\frac{2(r+1)}{3r}\right)^r.$$

Proof. The classical expected-degree theorem gives $\mathbb{E}D_v = 2$ [4, Theorem 1.1]. Apply the occupancy bound to incident edges with $\mu = 2$. The explicit exponential estimate uses $z = 1/4$; optimizing at $z = (r-2)/(2(r+1))$ gives the last formula. □

Declarations

AI assistance. OpenAI ChatGPT/Codex assisted with literature discovery, mathematical exploration and proof development, manuscript drafting, verification-code development and internal critical checking. The author assumes responsibility for all content.

Competing interests. Scott Schubr owns Luminary Oddities, which operates the Research Archive hosting the preprint of this article.

Funding. This work received no specific grant from any funding agency, commercial or not-for-profit sector.

Data and source availability. Complete proofs are included in the article. Editable LaTeX source and reproducible build instructions are available with the Research Archive preprint.

References

[1] A. Gupta, *Pairwise edge correlations in random minimum spanning trees: a universal bound and complete-graph negative correlation*, arXiv:2608.06816v1, 2026. `https://arxiv.org/abs/2608.06816`.

[2] D. Mayhew, *Equitable matroids*, Electron. J. Combin. **13** (2006), no. 1, Research Paper 41, 8 pp. `https://doi.org/10.37236/1067`.

[3] J. Huh, B. Schröter and B. Wang, *Correlation bounds for fields and matroids*, J. Eur. Math. Soc. **24** (2022), no. 4, 1335–1351. `https://doi.org/10.4171/JEMS/1119`.

[4] R. Lyons, Y. Peres and O. Schramm, *Minimal spanning forests*, Ann. Probab. **34** (2006), no. 5, 1665–1692. `https://doi.org/10.1214/009117906000000269`.

[5] L. Makowiec, *Fixed forests in the minimum spanning tree and cubic volume growth*, arXiv:2608.10654v1, 2026. https://arxiv.org/abs/2608.10654.

[6] M. Shaked and J. G. Shanthikumar, *Stochastic Orders*, Springer Series in Statistics, Springer, New York, 2007. https://doi.org/10.1007/978-0-387-34675-5.

[7] A. Di Crescenzo, G. Pisano and G. Psarrakos, *Comparisons of variances through the probabilistic mean value theorem and applications*, Adv. Appl. Probab. **58** (2026), no. 1, 269–300. https://doi.org/10.1017/apr.2025.10019.

[8] J. Oxley, *Matroid Theory*, 2nd ed., Oxford Graduate Texts in Mathematics 21, Oxford University Press, Oxford, 2011.

[9] T. E. Harris, *A lower bound for the critical probability in a certain percolation process*, Proc. Cambridge Philos. Soc. **56** (1960), no. 1, 13–20. https://doi.org/10.1017/S0305004100034241.