The stack-sorting degree grows faster than 1.80^n: a refutation of the Defant–Propp interval conjecture

DRAFT for root review, 3 November 2026. Prepared with AI assistance: the discovery, review and manuscript lanes were run by Claude Opus 5.5 agents and the Lean formalisation rounds by GPT-6.1 Sol agents, under a human-directed research campaign. No human authorship or affiliation is asserted. Zero accepted credit; acceptance is decided by the campaign root alone.

Abstract. For a map f from a finite set X to itself, Defant and Propp defined the degree of noninvertibility $\deg(f) = |X|^{-1} \Sigma_{x\in X} |f^{-1}(f(x))|$. For West's stack-sorting map s they proved that the limit of $\deg(s : S_n \to S_n)^{1/n}$ exists and lies in [1.12462, 4], and they conjectured, from random sampling, that it lies in (1.68, 1.73) (arXiv:2002.07144v3/v4, Conjecture 4.4; Conjecture 23 in the journal version). We show that the conjecture is false. Writing P_n for the number of pairs $(\sigma, \tau) \in S_n \times S_n$ with $s(\sigma) = s(\tau)$, so that $\deg(s : S_n \to S_n) = P_n / n!$, we give an exact recursion for P_n on pairs of tuples of words, polynomial in n, and use it with the source's own superadditivity inequality. An elementary root-split inequality and twenty exact values give $d_n^{1/n} > 1.73$ along $n = 131\cdot 2^k - 1$, so every limit is at least 1.73. The exact values through n = 82 give the stronger bound $\lim \deg(s : S_n \to S_n)^{1/n} \geq 1.80468$, and the numerical value of the limit appears to be about 1.95. The refutation with the bound 1.73, together with a non-vacuity companion, has been formalised in Lean 4 with Mathlib; the bound 1.80468 and the estimate are ordinary computations, not part of the formal packet. The MathDB paraphrase of the problem (record 345545), which reads deg as the maximum fibre size, is refuted trivially and earns no credit.

mathdb:2611.00040v1 04 Nov 2026

1. Introduction and results

Let s be West's stack-sorting map [W, DP]. For a word w with distinct letters, $s(\emptyset) = \emptyset$ and, writing w = L m R with m the largest letter, s(w) = s(L) s(R) m. Defant and Propp [DP] introduced the degree of noninvertibility of a map and asked for the exponential growth of the degree of s. Throughout,

$$d_n = \deg(s : S_n \to S_n) = P_n / n!, \quad P_n = \#\{(\sigma, \tau) \in S_n \times S_n : s(\sigma) = s(\tau)\}.$$

The second equality is the identity $\deg(f) = |X|^{-1} \Sigma_y |f^{-1}(y)|^2$ recorded in [DP, Section 1]: both sides count pairs with equal image.

Theorem 1.1 (formalised). For every $k \geq 0$, with $a_k = 131\cdot 2^k - 1$, we have $d_{a_k} > 1.73^{a_k}$. Consequently $d_n^{1/n} > 1.73$ for infinitely many n; if $d_n^{1/n} \to L$ then $L \geq 1.73$; and no L in the open interval (1.68, 1.73) is the limit of $d_n^{1/n}$. In particular Conjecture 4.4 of [DP] is false.

Theorem 1.2 (exact computation, not formalised). Let P_{82} be the 147-digit integer given in the dossier and produced by the recursion of Theorem 4.5. Then $P_{82} \cdot 100000^{83} > 4 \cdot 83 \cdot 82! \cdot 180468^{83}$, and hence $d_a^{1/a} > 1.80468$ along $a = 83\cdot 2^k - 1$. By the existence part of [DP, Theorem 2.5], $\lim d_n^{1/n} \geq 1.80468$. Numerically, $d_{82}^{1/82} = 1.95107$ and $d_{82} / d_{81} = 1.95411$, so the limit appears to be about 1.95.

The proof of Theorem 1.1 has four ingredients: the source's inequality (1), used only in the doubling case; a series-free iteration of it; a new elementary root-split inequality that propagates lower bounds from n = 20 to n = 130; and exact values $P_1, \ldots, P_{20}$ obtained from a new exact recursion (Theorem 4.5). Section 6 explains why the random-sampling evidence in [DP] pointed to the wrong interval. Section 8 describes the Lean formalisation and its trust limits.

2. The source statement

We quote the source [DP] from cached copies of arXiv v1–v4 and of the journal PDF [DPj], fetched by the independent review lane on 3 November 2026 and preserved with receipts in the dossier; this lane made no new download. The journal version has the same mathematics with different numbering (Definition 1, Theorem 8, Conjecture 23).

Definition 1.1 [DP]. "Let X and Y be finite sets, and let $f : X \to Y$ be a function. The degree of noninjectivity of f is the quantity $\deg(f) = (1/|X|) \Sigma_{x\in X} |f^{-1}(f(x))|$. When X = Y, we call deg(f) the degree of noninvertibility of f." The authors note that $\deg(f) = |X|^{-1} \Sigma_{y\in Y} |f^{-1}(y)|^2$ and that $|X| \deg(f)$ is the number of pairs (x, x') with $f(x) = f(x')$.

Stack-sorting map [DP, Section 2.3]. "If π is a permutation with largest entry m, then we can write π = LmR for some permutations L and R. Then s(π) = s(L)s(R)m." The example given is s(416352) = 143256.

Theorem 2.5 [DP]. "The limit $\lim_{n\to\infty} \deg(s : S_n \to S_n)^{1/n}$ exists and satisfies $1.12462 \leq \lim_{n\to\infty} \deg(s : S_n \to S_n)^{1/n} \leq 4$."

In its proof the authors relabel two colliding pairs $(\pi, \pi') \in S_{m-1}^2$ and $(\sigma, \sigma') \in S_{n-1}^2$ increasingly onto an $(m-1)$-set $A \subseteq [m+n-2]$ and onto its complement, and insert the letter $m+n-1$ between them. This map $(\pi, \pi', \sigma, \sigma', A) \mapsto (\pi_A\,(m+n-1)\,\sigma_A,\ \pi'_A\,(m+n-1)\,\sigma'_A)$ is injective and lands in colliding pairs, which gives $C(m+n-2, m-1)\,((m-1)!\,d_{m-1})\,((n-1)!\,d_{n-1}) \le (m+n-1)!\,d_{m+n-1}$, that is,

$$(1)\quad d_{m-1}\,d_{n-1} \le (m+n-1)\,d_{m+n-1}.$$

With $a_n = d_{n-1}/n^2$ for $n \ge 8$ (and $a_n = 0$ for $n \le 7$), (1) gives $a_m\,a_n \le a_{m+n}$ when $1/2 \le n/m \le 2$, and a theorem of de Bruijn and Erdős [dBE] yields the **sup formula** $\lim_n d_n^{1/n} = \sup_{n\ge1} a_n^{1/n}$. The lower bound 1.12462 is a computed value of $a_{10}^{1/10}$; our table reproduces it ($a_{10}^{1/10} = 1.124627$ from $P_9 = 117448188$). The upper bound comes from $|s^{-1}(\pi)| \le C_n$, the Catalan number.

Conjecture 4.4 [DP, v3 and v4]. "The value of the limit $\lim_{n\to\infty} \deg(s : S_n \to S_n)^{1/n}$ lies in the interval (1.68, 1.73)." The evidence immediately before it reads: "Using the 'decomposition lemma' described in [6], we have computed $|s^{-1}(s(\pi))|^{1/n}$ for several random permutations in S_n in order to gauge the size $\lim_{n\to\infty} \deg(s : S_n \to S_n)^{1/n}$. We first computed this quantity for 1000 random permutations with $n = 100$; the mean and standard deviation were 1.69 and 0.09. We then tried 100 random permutations with $n = 300$; the mean and standard deviation were 1.70 and 0.06." Problem 4.3 asks to improve the estimates of Theorem 2.5.

Version history. arXiv:2002.07144 has four versions: v1 (17 February 2020), v2 (18 February 2020), v3 (29 March 2020) and v4 (27 September 2020). Versions v1 and v2 contain no Conjecture 4.4. At the same place they state only: "This suggests that lim $\deg_n(s)^{1/n}$ should lie in the interval (1.6, 1.8)." The numbered conjecture with (1.68, 1.73) first appears in v3, is kept in v4, and is Conjecture 23 of the journal version [DPj] (Electron. J. Combin. 27(3) (2020), #P3.51; submitted 27 March 2020, published 4 September 2020). Theorem 1.2 refutes the earlier (1.6, 1.8) remark as well. Defant later restated the conjecture as open in "Troupes, Cumulants, and Stack-Sorting" [D22, Section 7.6], where Theorem 7.29 proves $\lim d_n^{1/n} > 1.62924$ and Remark 7.30 observes that brute force cannot go much beyond $n \approx 14$.

The MathDB record. MathDB record 345545 ("The growth constant for the stack-sorting map") copies Conjecture 4.4 but writes "let $\deg(s : S_n \to S_n)$ denote its maximum fiber size". This differs from Definition 1.1. Under that paraphrase the statement is false for a trivial reason. Let $F_1 = \{1\}$ and $F_n = \{n\ p : p \in F_{n-1}\} \cup \{p\ n : p \in F_{n-1}\}$. Since $s(n\ p) = s(p)\ n = s(p\ n)$, every element of F_n sorts to the identity. The two halves are disjoint, so $|F_n| = 2^{n-1}$. Hence the maximum fibre has size at least 2^{n-1}, its n-th root is at least $2^{4/5} > 1.741$ for $n \ge 5$, and any limit is at least 2. (In fact the maximum fibre is $s^{-1}(\mathrm{id})$, of Catalan size, and the limit is 4.) This reading is formalised for completeness and earns no credit. The substantive problem is the source reading, Definition 1.1.

3. Preliminaries

All words have distinct letters from the positive integers. A word on a finite set A is a word whose letters are exactly A. For $n \ge 0$ we write $[n] = \{1, \ldots, n\}$.

Lemma 3.1 (standardisation). If φ is strictly increasing on the letters of w, then $s(\varphi(w)) = \varphi(s(w))$. Consequently, for any k-element set A, the number of pairs of words (u, v) on A with $s(u) = s(v)$ is P_k.

Proof. φ maps the largest letter of w to the largest letter of $\varphi(w)$ and preserves the split $w = L\ m\ R$. Induct on the length of w. The count follows by transporting along the increasing bijection $[k] \to A$. □

Lemma 3.2 (decreasing tree and postorder readout). For a word w, let T(w) be the binary plane tree that is empty if w is empty and otherwise has root m, the largest letter, left subtree T(L) and right subtree T(R), where $w = L\ m\ R$. The in-order reading of T(w) is w, and the postorder reading of T(w) is s(w).

Proof. In-order reading gives $\mathrm{in}(T(L))\ m\ \mathrm{in}(T(R)) = L\ m\ R$. Postorder gives $\mathrm{post}(T(L))\ \mathrm{post}(T(R))\ m = s(L)\ s(R)\ m = s(w)$ by induction. □

We use three consequences. (a) s(w) is a rearrangement of w, so $|s(w)| = |w|$ and s(w) has the same letters as w. (b) The largest letter of w is the last letter of s(w). (c) The root of T(w) has 0, 1 or 2 children, and in the one-child case the child may be a left or a right child. These are the multiplicities 1, 2, 1 of Lemma 4.3. In tree language, P_n counts pairs of decreasing binary trees on [n] with the same postorder word, which is the viewpoint of the original discovery; the proof below needs only (a)–(c).

mathdb:2611.00040v1 04 Nov 2026

mathdb:2611.00040v1 04 Nov 2026

4. An exact recursion for P_n

Tuples of words. For $M \geq 0$ and $r \geq 0$, an r-tuple on $[M]$ is a sequence $F = (F_1, \ldots, F_r)$ of nonempty words whose concatenation $F_1 F_2 \cdots F_r$ is a permutation of $[M]$. Its output is $\mathrm{out}(F) = s(F_1)\, s(F_2) \cdots s(F_r)$, a permutation of $[M]$ by Lemma 3.2(a). Its set of internal boundaries is $\beta(F) = \{|F_1|, |F_1| + |F_2|, \ldots, |F_1| + \cdots + |F_{r-1}|\} \subseteq [M-1]$. Because $|s(F_j)| = |F_j|$, these are also the positions in $\mathrm{out}(F)$ where the blocks $s(F_j)$ end. A pair (F, G) of an r-tuple and an s-tuple on $[M]$ is a collision if $\mathrm{out}(F) = \mathrm{out}(G)$. It is irreducible if moreover $\beta(F) \cap \beta(G) = \emptyset$, and it is maximum-final if $\mathrm{out}(F)$ ends with M. Let

$$X(M, r, s),\ I(M, r, s),\ \Phi(M, r, s)$$

be the numbers of collisions, irreducible collisions and maximum-final collisions with arities r and s. Thus $X(0, 0, 0) = 1$ (the pair of empty tuples), $X(0, r, s) = 0$ otherwise, and $I(0, r, s) = 0$. A 1-tuple has no internal boundary, so $X(M, 1, 1) = I(M, 1, 1) = P_M$. All counts vanish if $r > M$ or $s > M$. By Lemma 3.1, the same counts are obtained for tuples on any M-element set of letters.

Lemma 4.1 (T1, common cut). For $M \geq 1$ and $r, s \geq 1$,

$$X(M, r, s) = I(M, r, s) + \sum_{b=1}^{M-1} C(M, b) \sum_{a=1}^{r-1} \sum_{c=1}^{s-1} X(M-b, r-a, s-c)\, I(b, a, c).$$

Proof. Let (F, G) be a reducible collision and let p be the largest element of $\beta(F) \cap \beta(G)$. Cut F after the word ending at output position p, and G likewise, giving $F = F'F''$ and $G = G'G''$. Then $\mathrm{out}(F') = \mathrm{out}(G')$ is the first p letters of the common output, and $\mathrm{out}(F'') = \mathrm{out}(G'')$ is the rest. By Lemma 3.2(a), F'' and G'' have the same letter set A, namely the letters of the last $M - p$ output positions; put $b = |A| = M - p$. The pair (F'', G'') has no common internal boundary, since such a boundary would give a common boundary of (F, G) larger than p. Let a and c be the arities of F'' and G''. Then $1 \leq a \leq r-1$, $1 \leq c \leq s-1$, and $1 \leq b \leq M-1$. Conversely, take b and $A \subseteq [M]$ with $|A| = b$, a collision (F', G') on $[M] \setminus A$ with arities $r - a$ and $s - c$, and an irreducible collision (F'', G'') on A with arities a and c. The concatenation is a collision with common boundary $M - b$, and no larger common boundary, because such a boundary would be internal to (F'', G''). The two constructions are inverse, and Lemma 3.1 counts each family. □

Lemma 4.2 (T2, the maximum is last). If (F, G) is an irreducible collision on $[M]$ with $M \geq 1$, then $\mathrm{out}(F)$ ends with M, and M lies in the last word of F and in the last word of G.

Proof. Let F_j contain M. By Lemma 3.2(b), $s(F_j)$ ends with M, so the output position p of M is the end of block j. Hence $p \in \beta(F)$ or $p = M$. The same holds for G at the same position p, because the outputs agree. If $p < M$, then $p \in \beta(F) \cap \beta(G)$, which contradicts irreducibility. So $p = M$, and since the words are nonempty, the block ending at M is the last one. □

The same argument shows that a collision is maximum-final exactly when M lies in the last word of both tuples.

Lemma 4.3 (T3, deletion of the maximum). For $M \geq 1$ and $r, s \geq 1$, with $\mu(0) = \mu(2) = 1$ and $\mu(1) = 2$,

$$\Phi(M, r, s) = \sum_{t=0}^{2} \sum_{u=0}^{2} \mu(t)\, \mu(u)\, X(M-1, r-1+t, s-1+u).$$

Proof. Let F be an r-tuple on $[M]$ with M in its last word, $F_r = A\, M\, B$. Then $s(F_r) = s(A)\, s(B)\, M$. Delete M and replace F_r by the nonempty words among A and B, in order. This gives an $(r - 1 + t)$-tuple F^- on $[M-1]$, where $t \in \{0, 1, 2\}$ is the number of nonempty words among A and B, and $\mathrm{out}(F) = \mathrm{out}(F^-)\, M$. The type of F_r is one of: leaf ($A = B = \emptyset$), left child ($B = \emptyset$), right child ($A = \emptyset$), or two children. Conversely, from an $(r - 1 + t)$-tuple H on $[M-1]$ and a type with t children, take the last t words of H as the children (possible since $r \geq 1$) and rebuild F_r as M, $C\, M$, $M\, C$ or $C_1\, M\, C_2$. These maps are inverse. For two tuples with M in their last words, $\mathrm{out}(F) = \mathrm{out}(G)$ if and only if $\mathrm{out}(F^-) = \mathrm{out}(G^-)$. The number of types with t children is $\mu(t)$. Summing gives the formula. For $M = 1$ only $t = u = 0$ and $r = s = 1$ contribute, through $X(0, 0, 0) = 1$. □

Lemma 4.4 (T4, fixed maximum). For $M \geq 1$ and $r, s \geq 1$,

$$\Phi(M, r, s) = I(M, r, s) + \sum_{b=1}^{M-1} C(M-1, b-1) \sum_{a=1}^{r-1} \sum_{c=1}^{s-1} X(M-b, r-a, s-c)\, I(b, a, c).$$

Proof. By Lemma 4.2 every irreducible collision is maximum-final. Decompose a reducible collision as in Lemma 4.1, with suffix letter set A. The suffix is an irreducible collision on A, so by Lemma 4.2 (applied after standardisation) its output ends with $\max A$, and so does the whole output. Thus the collision is maximum-final if and only if $M \in A$. The number of b-element sets $A \subseteq [M]$ containing M is $C(M-1, b-1)$. □

Theorem 4.5 (top-max recursion). For $M \geq 1$ and $r, s \geq 1$,

$$I(M, r, s) = \Sigma_{t,u}\, \mu(t)\, \mu(u)\, X(M - 1, r - 1 + t, s - 1 + u) - \Sigma_{b=1}^{M-1} C(M - 1, b - 1)\, K_b,$$

$$X(M, r, s) = I(M, r, s) + \Sigma_{b=1}^{M-1} C(M, b)\, K_b, \quad \text{where } K_b = \Sigma_{a=1}^{r-1} \Sigma_{c=1}^{s-1} X(M - b, r - a, s - c)\, I(b, a, c).$$

Every term on the right involves a size smaller than M. With $X(0, 0, 0) = 1$, this determines all X and I, and $P_n = X(n, 1, 1)$.

Proof. Combine Lemmas 4.3 and 4.4 to obtain I, and Lemma 4.1 to obtain X. The subtraction is exact because both expressions count $\Phi(M, r, s)$. □

To compute P_N it suffices to use arities $r, s \le \min(M, N + 1 - M)$ at size M. Deletion lowers M by one and raises an arity by at most one, while the convolution lowers both. So the table has $O(N^3)$ entries, each a sum of $O(N^3)$ products of integers with $O(N \log N)$ digits. This is polynomial in N. Exact values through $N = 82$ take about 80 seconds in plain Python. Scope of the formal version: the common-cut identity is formalised without clipping for all $M \ge 1$ and all arities, and the deletion step as an equivalence of finite families for all sizes. The recursion actually used, however, is proved only on the clipped domain $r, s \le \min(M, 21 - M)$, which is exactly the dependency cone of $P_1, \ldots, P_{20}$; the executable table is not a formal theorem about every P_n. Theorem 4.5 in full generality, and therefore the values $P_{21}, \ldots, P_{82}$, rest on the ordinary proof above. In both versions the identities are written additively, as $\Phi = I +$ (correction) and $\Phi =$ (deletion), so no truncated subtraction hides a counting deficit.

5. Lower bounds and the proof of Theorem 1.1

Lemma 5.1 (doubling, from (1)). For every $a \ge 0$, $C(2a, a)\, P_a^{\,2} \le P_{2a+1}$, equivalently $d_a^{\,2} \le (2a + 1)\, d_{2a+1}$.

Proof. Take $m = n = a + 1$ in the injection behind (1). □

Lemma 5.2 (series-free iteration). Let (d_n) be nonnegative with $d_a^{\,2} \le (2a + 1)\, d_{2a+1}$ for all a. Fix $a \ge 0$ and $c \ge 0$ with $4(a + 1)\, c \le d_a$. Put $a_0 = a$ and $a_{k+1} = 2a_k + 1$, so that $a_k + 1 = 2^k (a + 1)$. Then $d_{a_k} \ge 4\, (a_k + 1)\, c^{2^k}$ for all k.

Proof. The case $k = 0$ is the hypothesis. If it holds for k, then $d_{a_{k+1}} \ge d_{a_k}^{\,2} / (2a_k + 1) \ge 16\, (a_k + 1)^2\, c^{2^{k+1}} / (2a_k + 1) \ge 8\, (a_k + 1)\, c^{2^{k+1}} = 4\, (a_{k+1} + 1)\, c^{2^{k+1}}$, using $2(a_k + 1) \ge 2a_k + 1$. □

Corollary 5.3. Let $a \ge 1$ and $\beta = (d_a / (4(a + 1)))^{1/(a+1)} > 1$. Then $d_{a_k}^{\,1/a_k} > \beta$ for every k. Hence $\limsup d_n^{\,1/n} \ge \beta$, and $\lim d_n^{\,1/n} \ge \beta$ whenever the limit exists. In particular, if $4(a + 1) \cdot 1.73^{a+1} \le d_a$, then $d_{a_k} > 1.73^{a_k}$ for all k.

Proof. Apply Lemma 5.2 with $c = \beta^{a+1}$: $d_{a_k} \ge 4\, (a_k + 1)\, \beta^{a_k + 1} > \beta^{a_k}$. No logarithms, series or existence of the limit are used. □

Lemma 5.4 (root split). For $n \ge 2$,

$$P_n \ge 4\, P_{n-1} + \Sigma_{i=1}^{n-2} C(n - 1, i)\, P_i\, P_{n-1-i}.$$

Proof. For $0 \le i \le n - 1$, choose an i-set $A \subseteq [n - 1]$, two words L, L′ on A with $s(L) = s(L')$, and two words R, R′ on $[n - 1] \setminus A$ with $s(R) = s(R')$. Then $s(L\ n\ R) = s(L)\ s(R)\ n = s(L'\ n\ R')$. These are $C(n - 1, i)\, P_i\, P_{n-1-i}$ pairs, with $P_0 = 1$. The cases $i = 0$ and $i = n - 1$ give $2\, P_{n-1}$. In addition, take R, L′ on $[n - 1]$ with $s(R) = s(L')$. Then $s(n\ R) = s(R)\ n = s(L'\ n)$, and symmetrically $s(L\ n) = s(n\ R')$. This gives $2\, P_{n-1}$ further pairs. All these pairs are distinct: the positions of n in the two words separate the families, and within a family the words L, R, L′, R′ are recovered by splitting at n. □

Lemma 5.5 (base-130 certificate). Define $G_n = P_n$ for $n \le 20$ and $G_n = 4\, G_{n-1} + \Sigma_{i=1}^{n-2} C(n - 1, i)\, G_i\, G_{n-1-i}$ for $n \ge 21$. Then $G_n \le P_n$ for all n, and

$$4 \cdot 131 \cdot 130! \cdot 173^{131} < G_{130} \cdot 100^{131}.$$

Proof. The first claim follows by strong induction from Lemma 5.4, since the right-hand side is monotone in the earlier values. The integer G_{130} has 254 digits; it is the literal `lower130` of the formal packet, and the inequality is a closed integer comparison. □

Proof of Theorem 1.1. Theorem 4.5 gives $P_1, \ldots, P_{20}$ (Table 1). Lemma 5.5 gives $d_{130} \ge G_{130} / 130! > 4 \cdot 131 \cdot 1.73^{131}$. Corollary 5.3 with $a = 130$ gives $d_{a_k} > 1.73^{a_k}$ for $a_k = 131 \cdot 2^k - 1$. Hence $d_n^{\,1/n} > 1.73$ for infinitely many n. If $d_n^{\,1/n} \to L$, then $L \ge 1.73$ along this subsequence, so $L \notin (1.68, 1.73)$. □

The proof never assumes that the limit exists. The non-vacuity statement $(d_n^{\,1/n} > 1.73$ frequently, with explicit witnesses 130, 261, 523, 1047, …) shows that the refutation is not the vacuous one available if the limit failed to exist. Together with [DP, Theorem 2.5] the same argument gives $\lim d_n^{\,1/n} \ge (G_{130} / (4 \cdot 131 \cdot 130!))^{1/131} > 1.73006$. This is an ordinary consequence; the formal packet proves only that every possible limit is at least 1.73, and it does not

mathdb:2611.00040v1 04 Nov 2026

formalise existence of the limit.

Proof of Theorem 1.2. Corollary 5.3 with a = 82 and the exact P_{82} gives $\beta_{82} = 1.8046832\ldots > 1.80468$. The integer form $P_{82} \cdot 100000^{83} > 4 \cdot 83 \cdot 82! \cdot 180468^{83}$ was checked with exact integer arithmetic. The comparison with 1.8047 fails: the true value rounds to 1.80468, so a bound of 1.8047 would be too strong. The same test passes at threshold 1.8 for a = 79, …, 82, which refutes the v1/v2 remark (1.6, 1.8). Independently of Corollary 5.3, the source's sup formula with the exact P_{78} gives $(d_{78} / 79^2)^{1/79} = 1.73190 > 1.73$. That route needs no new inequality, but it does need the exact value of P_{78}. □

Remark 5.6 (the reviewer's observation: the depth of the bare refutation). The independent review of this result pointed out that the bare refutation needs much less than Theorem 4.5. With only the exact values P_1, …, P_{13} as seeds, the recursion G of Lemma 5.5 first passes the 1.73 test at a = 910 (seeds through 14, 16 and 20 give a = 446, 230 and 130). The value $P_{13} = 30533141865700$ is a sum over 13! ≈ 6.2×10^9 permutations, within the brute-force range that Defant estimated in [D22, Remark 7.30]. With seeds through 12 the review found no passing a ≤ 3000; the asymptotic ratio of G is then about 1.73008. So the conjecture was refutable with brute-force data and the five-line Lemma 5.4. This lowers the depth of the bare refutation. The exact polynomial-time recursion, the values through n = 82, the bound 1.80468 and the estimate of about 1.95 go beyond the bare refutation. We do not claim that the new recursion is essential for refuting the conjecture.

6. Why the random-sampling heuristic failed

For a uniform random $\pi \in S_n$ put $X_n = |s^{-1}(s(\pi))|$. Then $E[X_n] = d_n$. The experiment of [DP] estimated $E[X_n^{1/n}]$ by sample means. Since $t \mapsto t^{1/n}$ is concave, Jensen's inequality gives $E[X_n^{1/n}] \le (E\, X_n)^{1/n} = d_n^{1/n}$. So the experiment measured a smaller quantity, and the data never supported the upper end 1.73. Note that $d_n = \Sigma_y |s^{-1}(y)|^2 / n!$ weights each fibre by the square of its size. It is therefore dominated by the rare outputs with exponentially large fibres. A sample mean of $X_n^{1/n}$ instead tracks the typical fibre, roughly $\exp(E[\log X_n]/n)$. Exponential growth rates of a mean and of a typical value can differ when large deviations dominate, and the exact data show a large gap: $d_{82}^{1/82} = 1.951$ while the sampled mean at n = 100 was 1.69. We have not proved anything about the typical growth rate. The point is only that the conjectured interval estimated the wrong statistic.

7. Data

Table 1 lists exact values of P_n for n ≤ 20; these are the twenty seeds proved in Lean. Table 2 lists selected n up to 82. There $\beta_n = (d_n / (4(n+1)))^{1/(n+1)}$ is the doubling bound of Corollary 5.3, and $\sigma_{n+1} = (d_n / (n+1)^2)^{1/(n+1)}$ is the source's sup-formula bound. Decimals are rounded to five places; every inequality claimed in the text was checked with exact integers. The exact P_1, …, P_{82} are in the dossier. They were computed by the discovery lane, by an independent reviewer (a top-max subtraction recursion to n = 82 in two modes, a label-insertion automaton to n = 22 and brute force to n = 10), and again in this lane by a fresh implementation of Theorem 4.5, with full agreement. In the computed range, $d_n^{1/n}$ increases strictly for $2 \le n \le 82$ and d_n / d_{n-1} decreases strictly for $4 \le n \le 82$. A three-point fit of the ratios at n = 80, 81, 82 to $L + A/n + B/n^2$ gives 1.9532. These trends are heuristic, which is why we state the limit only as "about 1.95".

mathdb:2611.00040v1 04 Nov 2026

mathdb:2611.00040v1 04 Nov 2026

Table 1. Exact $P_n = n!\, d_n$ for $n \le 20$.

n	P_n	$d_n^{1/n}$	d_n / d_{n-1}
1	1	1.00000	–
2	4	1.41421	2.00000
3	26	1.63032	2.16667
4	224	1.74787	2.15385
5	2356	1.81384	2.10357
6	29064	1.85213	2.05603
7	411172	1.87536	2.02102
8	6573624	1.89032	1.99844
9	117448188	1.90063	1.98518
10	2322885008	1.90821	1.97780
11	50429847572	1.91407	1.97364
12	1192782788192	1.91875	1.97103
13	30533141865700	1.92258	1.96910
14	841025569061736	1.92575	1.96748
15	24802500974885020	1.92841	1.96605
16	779708932637648392	1.93067	1.96479
17	26028895710444355516	1.93259	1.96370
18	919586611535156918680	1.93426	1.96275
19	34279121480711537327388	1.93570	1.96193
20	1344580237342589147760320	1.93697	1.96122

Table 2. Selected n. The 1.73 test first passes at $n = 50$ for β_n and at $n + 1 = 79$ for σ_{n+1}; $\beta_{49} = 1.72959$.

n	P_n (leading digits)	$d_n^{1/n}$	d_n / d_{n-1}	β_n	σ_{n+1}
13	3.0533×10^{13}	1.92258	1.96910	1.37637	1.25856
20	1.3445×10^{24}	1.93697	1.96122	1.51991	1.40451
25	2.4774×10^{32}	1.94149	1.95879	1.58297	1.47302
30	1.2188×10^{41}	1.94421	1.95739	1.62892	1.52480
40	3.0807×10^{59}	1.94727	1.95587	1.69179	1.59843
49	9.6011×10^{76}	1.94877	1.95516	1.72959	1.64439
50	9.3855×10^{78}	1.94889	1.95510	1.73308	1.64870
60	2.0924×10^{99}	1.94988	1.95464	1.76245	1.68546
70	2.4500×10^{120}	1.95054	1.95435	1.78449	1.71364
78	4.9273×10^{137}	1.95092	1.95418	1.79855	1.73190
79	7.6067×10^{139}	1.95096	1.95416	1.80014	1.73397
80	1.1891×10^{142}	1.95100	1.95414	1.80168	1.73600
81	1.8822×10^{144}	1.95103	1.95413	1.80320	1.73799
82	3.0160×10^{146}	1.95107	1.95411	1.80468	1.73994

8. Formal verification

Packet. The Lean 4 development (Lean 4.34.1 with a pinned Mathlib) is the source-only archive `stacksort-source6.zip` (SHA-256 9b734e2f…cf5c2f): 92 modules and 8 scripts, built over six rounds. The final endpoints are in `StackSortFinal6.lean`, namespace `StackSort.Round6`:

```
theorem mathdb_345545 : ¬ Conj
theorem limit_ge : ∀ L, Tendsto (fun n => deg n ^ (1 / (n : ℝ))) atTop (𝓝 L) → (1.73 : ℝ) ≤ L
theorem nonvacuity : ∃ᶠ n in atTop, (1.73 : ℝ) < deg n ^ (1 / (n : ℝ))
theorem subsequence_companion : ∀ k, (1.73 : ℝ) < roots (index 130 k)
theorem mathdb_345545_max : ¬ ConjMax
theorem max_companion : ∀ n, 5 ≤ n → (1.73 : ℝ) < maxRoots n
```

Here `Conj := ∃ L : ℝ, Tendsto roots atTop (𝓝 L) ∧ L ∈ Set.Ioo 1.68 1.73`, `roots n = deg n ^ (1/(n:ℝ))`, `deg n = (P n : ℝ) / n!`, and `index 130 k = 131 · 2^k − 1` (index 130 0 = 130, index a (k+1) = 2 · index a k + 1). Both readings are stated literally: `ConjMax` is the same proposition for `maxRoots n`, the n-th root of `maxFib n`, the largest fibre size.

Table 3. Correspondence between the paper and Lean.

Paper	Lean (main names)
s(LmR) = s(L)s(R)m	`StackSort.s` (structural fuel `sortFuel`), `s_split`; example `s [4,1,6,3,5,2] = [1,4,3,2,5,6]`
S_n, with n! elements	`words n` = permutations of `List.range n` (letters 0,…,n−1); `words_card`
P_n, d_n (Def. 1.1)	`P`, `deg = P/n!`; `P_eq_sum`, `deg_eq_average`, `fibre_square_identity`
Lemma 3.1	`s_map_orderEmbedding`, `standardisation_proved`
Lemma 5.1 (from (1))	`DoublingInjection`, `Round3.doubling_proved`
Lemma 5.2, Cor. 5.3	`iteration`, `subsequence_above`, `limit_ge_of_subsequence`, `frequently_of_subsequence`
Lemma 5.4	`RootSplit`, `Round3.rootSplit_proved`
Lemma 5.5	`G`, `G_le_P`, `lower130`, `certificate130`, `Round6.lower130`
Tuples, X, I	`tuples`, `output`, `XI`, `II`, clipped `literalX`, `literalI`
Lemma 4.1 (T1)	`Cut5.commonCutCount_proved`
Lemma 4.2 (T2)	`irreducible_maximum_last`, `maximum_in_last_word`
Lemma 4.3 (T3)	`Del5.deletionEquiv`, `Del6.taggedDeletionCount_proved`
Lemma 4.4 (T4)	`Fixed6.fixedMaximumDecomposition_proved`
Theorem 4.5	`Seeds2.AdditiveRecursion`, `Round6.additiveRecursion_proved`
$P_1,\ldots,P_{20}$	`ExactSeeds20`, `Round6.exactSeeds20_proved` (20 kernel-checked table rows)
Theorem 1.1	`Round6.mathdb_345545`, `limit_ge`, `nonvacuity`, `subsequence_companion`
MathDB paraphrase	`maxFib`, `ConjMax`, `family`; `Round6.mathdb_345545_max`, `max_companion`

Two fidelity points. First, `¬ Conj` alone would be vacuously true if the limit did not exist. `nonvacuity` and `subsequence_companion` exclude that reading, and `limit_ge` is the stronger universal form. No existence theorem is added. Second, letters 0, …, n − 1 replace 1, …, n, which is justified by the proved order-equivariance. Third, at n = 0 real division gives 1/0 = 0, so `roots 0 = 1`; a single initial term does not affect limits or frequency at infinity. The maximum reading takes the supremum over attained outputs, which suffices because unattained outputs have empty fibres.

Table 4. Formalisation rounds. Hypotheses are ordinary Lean propositions passed as explicit arguments, never axioms.

Round	Modules	Archive SHA-256	New unconditional content	Remaining hypotheses
1	11	67a3893f…	literal s, P, deg, fibre identities, standardisation, iteration, wrap-up, integer certificate, maximum-fibre endpoints	Doubling, Lower130
2A	27	a7aa7c4b…	factorial conversion, fixed and end-family injections, structural G with 110 kernel rows, G_{130} = lower130	DoublingInjection, RootSplit, ExactSeeds20
2B	41	ae124e54…	tuple domain, X(M,1,1) = P_M, word-level deletion with weights 1, 2, 1, T2, pruned recursion with uniqueness, table to 20	AdditiveRecursion
3	66	f55b6933…	labelled shuffles; RootSplit and DoublingInjection proved	AdditiveRecursion
4	73	46c01f20…	last common cut and uniqueness, labelled tuples, deletion lemmas	CommonCutCount, MaximumFinalCount
5A	79	31dfd87c…	CommonCutCount (T1) proved	MaximumFinalCount
5B	80	3ca49673…	deletionEquiv (T3 as an equivalence)	FixedMaximumDecomposition, TaggedDeletionCount
6	92	9b734e2f…	merge of 5A and 5B (86 inherited modules, byte-identical); tagged deletion count and fixed-maximum decomposition proved	none

Audits and controls. The following are reported by the round-6 lane and were not rerun here. They are listed separately, as the fidelity review requested.

- *Compilation and checker scope.* A fresh extraction of the source-only archive, a strict serial rebuild of all 92 modules (warnings as errors, −M3500 memory cap) and an explicit-module leanchecker run on each module passed.
- *Endpoint axioms.* `#print axioms` for each of the twelve round-6 theorems prints exactly `propext`, `Classical.choice` and `Quot.sound`, the only allowed axioms.

mathdb:2611.00040v1 04 Nov 2026

- *Source scan and closure audit.* There is no `sorry`, `admit` or `native_decide`. The endpoint import closure has 349,566 declarations, of which 1,110 are local and 90 new in round 6. All local `unsafe`, `partial`, `implemented_by`, `extern` and `axiom` flags are zero; upstream metadata is reported separately in the round-6 results.
- *Negative controls,* checked by `decide`: $P_3 \neq 27$; the wrong convention $s(416352) \neq 641532$ is rejected; the base-130 certificate fails at threshold 1.95; the analogous certificate at base 129 fails; a perturbed lower-recursion row and seed are rejected; the seed table rejects $P_5 = 2357$; and all-unit deletion weights are rejected.
- *Independent diagnostics.* A Python regression in the round-6 lane enumerated literal collisions through size 7.

The independent fidelity review separately rebuilt, with strict settings, the 84-module import closure of the endpoint module (a resumed, bounded build), ran leanchecker on `StackSortFinal6`, printed the same three axioms for all six endpoints, and rescanned all 92 modules for forbidden tokens, finding none. Its own Python enumeration through $n = 7$, its recomputation of G_{130} and of the certificate, and its own Lean examples (kernel reductions of $P_0, \ldots, P_4$) agree. It found no surrogate definition and no vacuity, and required no Lean source fix.

Checks in this lane (no Lean run). The archive copied into the dossier is byte-identical to the round-6 archive. A safe-extraction test found 100 regular files, no absolute or parent paths, no links and no compiled artifacts. All 92 module hashes equal the round-6 record, all six endpoints are declared, and no forbidden token occurs outside comments and strings. A fresh Python implementation in this lane brute-forced $P_1, \ldots, P_8$. It verified Lemmas 4.1–4.4 literally on all tuple pairs with $M \leq 6$, recomputed $P_1, \ldots, P_{82}$ by Theorem 4.5, and checked Lemma 5.4 for $n \leq 82$ and Lemma 5.1 for $2a + 1 \leq 81$. It confirmed that G_{130} and G_{129} equal the literals `lower130` and `lower129`, and it reproduced the certificate and both numerical controls.

Trust limits. (i) Theorem 1.1 is the formal content. The bound 1.80468, the values $P_{21}, \ldots, P_{82}$ and the estimate 1.95 are ordinary computations (exact integer arithmetic in Python, three independent implementations); they are not kernel-checked. (ii) The direct kernel evaluations of the literal counts $P_5 = 2356$ and $P_6 = 29064$ by enumeration did not complete within the resource caps. P_5 hit Lean's −M3500 memory limit with `decide` and `decide +kernel`; P_6 was stopped by the 6 GiB watchdog. These values are nonetheless proved, through the proved recursion (`exactSeeds20_proved`); only the redundant direct regression is missing. (iii) The formal recursion is clipped to the dependency cone of the twenty seeds (Section 4). (iv) Limit existence is the source's Theorem 2.5 and is not formalised; the formal statements are that every possible limit is at least 1.73 and that 1.73 is exceeded frequently. (v) The root's decisive fresh replay (lane stacksort-parent-final-r1) is pending: no replay result existed when this draft was finished, and no claim is made about it. (vi) The independent formal-fidelity review (lane stacksort-fidelity-r1, finished 3 November 2026) reports gates A (definition fidelity), B (proof structure), C (independent semantic sanity) and D (bounded formal spot check) as PASS, and E (manuscript readiness) as qualified pass with seven integration items, all applied here (`gate-E-map.json`). Its gate D is bounded: a strict build of the 84-module import closure of the endpoint module, resumed once after an unexplained interruption, with leanchecker run on the endpoint module only. It is not the root's replay. (vii) This lane ran no Lean, lake or leanchecker.

9. Related work, prior art and access limits

The stack-sorting map goes back to Knuth's stack [K1] and West's thesis [W]; see Bóna [B1, B2]. Defant developed methods for counting preimages under s [D17, D20], including the decomposition lemma used for the sampling in [DP]. The lower bound 1.62924 of [D22, Theorem 7.29] comes from free cumulants and the sup formula of [DP]. It is consistent with our bounds and lies below the conjectured interval, so it did not decide the conjecture. A bounded search by the independent review (3 November 2026) covered the OpenAlex cited-by lists of [DP] (5 works) and [D22] (22 works, screened by title, the plausible ones read), keyword searches, the arXiv and journal pages (no erratum or withdrawal), and OEIS. It found no refutation of the conjecture, no growth-constant bound above 1.62924, and no OEIS entry for 1, 4, 26, 224, 2356, 29064, 411172, 6573624, nor for 2322885008. Google Scholar, MathSciNet and zbMATH were not consulted. The root's registry refresh at 2026-11-03T19:52Z found the OpenAI math repository head unchanged (4eaee875…). A path screen of the complete current Omega trees (head 813fe492…) found only an unrelated vincular stack-sorting item, with no degree or noninvertibility content. That screen is a path screen, not a full-body proof search. This lane made no fresh external access and used frozen local copies only. We make no priority claim. The result is offered as a correction to a published conjecture, and the recursion and data are offered for checking.

mathdb:2611.00040v1 04 Nov 2026

10. Status

This is a draft. Theorem 1.1 rests on the Lean packet described in Section 8 and on the ordinary proof above. The independent fidelity review passes gates A–D and gives qualified pass for manuscript readiness; its seven items are applied. The root's decisive fresh replay was pending when this draft was finished. Acceptance is outside this manuscript. AI assistance is described in the header. The campaign records zero accepted credit until the root reads this paper and decides.

References

[B1] M. Bóna, Combinatorics of Permutations, CRC Press, 2012 (as cited in [DP]).

[B2] M. Bóna, A survey of stack-sorting disciplines, Electron. J. Combin. 9 (2003) (as cited in [DP]).

[dBE] N. G. de Bruijn and P. Erdős, Some linear and some quadratic recursion formulas II, Proc. Kon. Ned. Akad. Wetensch. Ser. A 14 (1952), 152–163 (as cited in [DP]).

[D17] C. Defant, Preimages under the stack-sorting algorithm, Graphs Combin. 33 (2017), 103–122 (as cited in [DP]).

[D20] C. Defant, Counting 3-stack-sortable permutations, J. Combin. Theory Ser. A 172 (2020) (as cited in [DP]).

[D22] C. Defant, Troupes, cumulants, and stack-sorting, arXiv:2004.11367v2 (10 January 2022); Adv. Math. 399 (2022), article 108270, according to the OpenAlex record W3018498331 (publisher page not accessed).

[DP] C. Defant and J. Propp, Quantifying noninvertibility in discrete dynamical systems, arXiv:2002.07144, v1–v4 (2020).

[DPj] C. Defant and J. Propp, Quantifying noninvertibility in discrete dynamical systems, Electron. J. Combin. 27(3) (2020), #P3.51. DOI 10.37236/9475 (resolved to the journal landing page on 3 November 2026).

[K1] D. E. Knuth, The Art of Computer Programming, Vol. 1: Fundamental Algorithms, Addison–Wesley, 1973 (as cited in [DP]).

[MDB] MathDB record 345545, "The growth constant for the stack-sorting map", https://mathdb.com/p/345545/the-growth-constant-for-the-stack-sorting-map (retrieved through the 2026-09-29 snapshot; raw line hash in the dossier).

[W] J. West, Permutations with restricted subsequences and stack-sortable permutations, Ph.D. thesis, MIT, 1990.

mathdb:2611.00040v1 04 Nov 2026