The Subgraph Lattice Theorem and Geometric Partition Counting

Part I: Combinatorial Foundations of $\mathbb{G}$-Partitions

Bee Rosa Davis
Davis Geometric
bee_davis@alumni.brown.edu
ORCID: 0009-0009-8034-4308

April 10, 2026

Abstract

In the geometric natural numbers $\mathbb{G}_n$ [Davis, 2025], each natural number n is a pair $(\mathbb{G}_n, G)$ consisting of an element count $n \geq 1$ and a simple connected graph G on n vertices. We define the *geometric partition function* $p_{\mathbb{G}}(\mathbb{G}_n, G)$ as the number of set partitions of $V(G)$ in which every block induces a connected subgraph. This is a strict refinement of the classical set-partition count $B(n)$: it depends on G, not just n.

Our main result is the **Subgraph Lattice Theorem**: if G is a spanning subgraph of H (same vertex set, $E(G) \subseteq E(H)$), then $p_{\mathbb{G}}(G) \leq p_{\mathbb{G}}(H)$. The geometric partition function is monotone on the spanning subgraph poset of K_n. The proof is one paragraph: connected set-partitions of G inject into connected set-partitions of H. This structural principle unifies the paper's other results.

An empirical study across up to 10 graph families at each $n \in \{4, \ldots, 10\}$ (family panel varies with parity and special cases) reveals that vertex connectivity $\kappa(G)$ approximates the lattice order as a scalar invariant—Spearman $\rho \geq 0.83$ at every tested n—while, in this tested family panel, maximum degree shows no correlation (a built-in negative control).

We additionally verify, on 15 years of CAIDA AS-relationships data [Luckie et al., 2013] (180 monthly snapshots, 2010–2024), that the ratio of edge succession to vertex succession in the Internet autonomous-system graph grows monotonically with network maturity (slope 3.68/year against a pre-registered threshold of 0.05, Spearman $\rho = 0.968$, $p < 10^{-8}$).

Secondary validation outcomes are mixed by design: Test 6 and Test 8 fire pre-registered falsifiers, Test 9 is blocked by data observability, Test 10 is partial, and Test 11 is deferred by directive. This failure-tolerant profile strengthens the audit credibility of the supported claims.

Keywords: geometric natural numbers, connected set-partitions, Bell numbers, connection curvature, validation protocol, OEIS, Davis Geometric

1 Introduction

1.1 The Geometric Natural Numbers

The natural numbers in Peano arithmetic carry no structure beyond succession: each n is the successor of $n-1$, forming a path graph P_ω. In the geometric naturals framework introduced in [Davis, 2025], each natural number $n \geq 1$ is promoted to a pair $(\mathbb{G}_n, G)$, where G is a simple connected graph on n labeled vertices. The path graph P_n is the Peano restriction; other graphs on the same vertex set encode richer relational structure.

The central claim of [Davis, 2025] is that Peano arithmetic occupies a flat limit of a richer geometric arithmetic, analogous to Euclidean space within Riemannian geometry. This paper provides first concrete *combinatorial* evidence for that claim, at the level of connected set-partition counting functions.

1.2 From Counting to Partitioning

Classical partition theory studies the number of ways to decompose a positive integer into a sum of positive integers [Andrews, 1976, Hardy and Ramanujan, 1918, Rademacher, 1937]. The set-partition analog asks: how many ways can an n-element set be partitioned into nonempty blocks? The answer is the Bell number $B(n)$.

Neither classical partitions nor Bell numbers depend on relations *between* the elements being partitioned. When the elements carry graph structure—as in $\mathbb{G}_n$—a natural constraint emerges: require that each block of the partition induces a connected subgraph. This is the *connected set-partition* (Definition 2.2).

The resulting count $p_{\mathbb{G}}(\mathbb{G}_n, G)$ depends on G, not just on n. It interpolates between two classical limits:

- On the complete graph K_n (every induced subgraph is connected), $p_{\mathbb{G}}(K_n) = B(n)$.
- On trees (where every edge is a bridge), $p_{\mathbb{G}}(T_n) = 2^{n-1}$.

The gap between 2^{n-1} and $B(n)$ is the domain where graph topology genuinely modulates partition counts. This gap is the geometric content of $\mathbb{G}$-partitions.

1.3 Contributions

This paper proves three closed-form identities (the Tree, Complete-Graph, and Cycle Theorems), establishes the Subgraph Lattice Theorem ($p_{\mathbb{G}}$ is monotone on the spanning subgraph poset), and introduces a discrete *connection curvature* $\kappa_n^{p_{\mathbb{G}}}$ that measures how the cycle partition function approaches the flat (tree) limit. An empirical study (Section 9) confirms the lattice ordering across up to 10 graph families and identifies vertex connectivity as the best scalar proxy for $p_{\mathbb{G}}$-rank, with explicit counterexamples delineating the boundary between the lattice theorem and the statistical tendency.

The paper is self-contained. All definitions are stated from scratch; the only prerequisite is elementary graph theory.

1.4 Structure

Section 2 gives definitions. Section 3 proves the four structure theorems (Tree, Complete-Graph, Cycle, and Subgraph Lattice) and the connection curvature result. Section 4 describes the validation protocol. Section 5 specifies Test 1. Section 6 reports findings. Section 7 disposes of pre-declared conjectures. Section 8 develops the connection curvature. Section 9 presents the empirical lattice study, identifies the κ-correlation, and exhibits counterexamples. Section 13 consolidates the full 11-test validation status as of this draft. Section 14 lists open problems. Appendix A contains the fiber-bundle audit.

2 Definitions

We recall the essential definitions from the Zero paper [Davis, 2025] and introduce the geometric partition function.

Definition 2.1 (Geometric Natural Number). *A* geometric natural number *is a pair* $(\mathbb{G}_n, G)$ *where* $n \geq 1$ *is a positive integer and* G *is a simple undirected graph on* n *labeled vertices. For* $n \geq 2$, G *is required to be connected. The void state* $\mathbb{G}_0$ *is excluded. The pre-geometric seed* $\mathbb{G}_1 = (\{\bullet\}, P_1)$ *is the irreducible element.*

Definition 2.2 (Connected Set-Partition). *Let $G = (V, E)$ be a graph. A* connected set-partition *of $(\mathbb{G}_n, G)$ is a partition $\{V_1, \ldots, V_k\}$ of V such that for every $i \in \{1, \ldots, k\}$, the induced subgraph $G[V_i]$ is connected.*

Definition 2.3 (Geometric Partition Function). *The* geometric partition function *is*

$$p_{\mathbb{G}}(\mathbb{G}_n, G) \;=\; \big|\{\{V_1, \ldots, V_k\} : \{V_i\} \text{ is a connected set-partition of } (\mathbb{G}_n, G)\}\big|.$$

Remark 2.1 (Why Connected-Partition, Not Clean-Partition). *An earlier draft of the research program considered the* clean-partition *count: the number of partitions $\{V_i\}$ such that* no edge of G crosses between blocks. *For any connected graph G, the clean-partition count is trivially 1 (the partition into a single block), because any nontrivial partition of a connected graph must cut at least one edge. The connected-partition definition (Definition 2.2), which requires each block to induce a connected subgraph but permits inter-block edges, produces the rich structure reported in this paper.*

3 Structure Theorems

The results in this section are organized around a single structural principle: the geometric partition function $p_{\mathbb{G}}$ is monotone on the spanning subgraph poset. We state this first, then compute $p_{\mathbb{G}}$ at three distinguished points of the poset—trees (the floor), complete graphs (the ceiling), and cycles (the first nontrivial position above the floor).

3.1 The Lattice Principle

Theorem 3.1 (Subgraph Lattice Theorem). *Let G and H be graphs on the same vertex set V with $E(G) \subseteq E(H)$. Then $p_{\mathbb{G}}(G) \leq p_{\mathbb{G}}(H)$.*

Proof. Let $\pi = \{V_1, \ldots, V_k\}$ be a connected set-partition of G: each $G[V_j]$ is connected. Since $E(G) \subseteq E(H)$, every edge of $G[V_j]$ is an edge of $H[V_j]$, so $H[V_j]$ contains $G[V_j]$ as a subgraph on the same vertex set and is therefore connected. Hence π is also a connected set-partition of H. The map $\pi \mapsto \pi$ is an injection from connected set-partitions of G into connected set-partitions of H, so $p_{\mathbb{G}}(G) \leq p_{\mathbb{G}}(H)$. □

Corollary 3.2 (Sandwich Bound). *For every connected graph G on n vertices,*

$$2^{n-1} \;\leq\; p_{\mathbb{G}}(G) \;\leq\; B(n).$$

Proof. Every connected graph on n vertices contains a spanning tree T as a spanning subgraph (i.e., $E(T) \subseteq E(G)$), and is itself a spanning subgraph of K_n (i.e., $E(G) \subseteq E(K_n)$). By Theorem 3.1, $p_{\mathbb{G}}(T) \leq p_{\mathbb{G}}(G) \leq p_{\mathbb{G}}(K_n)$. By the Tree Theorem (3.3), $p_{\mathbb{G}}(T) = 2^{n-1}$; by the Complete-Graph Theorem (3.4), $p_{\mathbb{G}}(K_n) = B(n)$. □

Remark 3.1 (Floor Tightness in Sparse Real Networks). *The sandwich bound $2^{n-1} \leq p_{\mathbb{G}}(G) \leq B(n)$ raises the question of whether the floor is achieved or merely theoretical. We examined the* C. elegans *chemical synapse connectome [Varshney et al., 2011], symmetrized to an undirected graph (279 neurons, 1,961 edges, edge density $\hat{p} \approx 0.051$). Among randomly sampled connected induced subgraphs, 81% of 4-vertex, 69% of 5-vertex, and 52% of 6-vertex subgraphs achieve $p_{\mathbb{G}} = 2^{n-1}$ exactly. The floor is not a loose lower bound—it is where sparse real networks sit.*

A pre-registered comparison to an Erdős–Rényi null at matched density was inconclusive: at $\hat{p} \approx 0.051$, ER graphs themselves produce very few connected induced subgraphs, all at the tree floor. The connectome produces 55× more connected 4-vertex induced subgraphs than the matched ER null, consistent with the well-established clustering enrichment in C. elegans *[Varshney et al., 2011]; our lattice analysis adds that the excess connected subgraphs also hug the tree floor. A connectivity-matched null—conditioning on connectedness and edge count—is the appropriate next test and is pre-registered as future work.*

Remark 3.2 (The Partition Polynomial). *The connected set-partition count $p_{\mathbb{G}}(G)$ coincides with the evaluation $Q(G,1)$ of the* partition polynomial $Q(G,x) = \sum_{i=1}^{n} q_i(G)\, x^i$, *where $q_i(G)$ counts connected set-partitions of G with exactly i blocks, as introduced by Simon, Tittmann, and Trinks [Simon et al., 2011]. Vince [Vince, 2017] denotes the same quantity $P(G)$. The partition polynomial is known to satisfy a splitting formula on separating vertex sets and to be #P-hard to compute in general. The monotonicity statement of Theorem 3.1 appears to be new; it is not stated in [Simon et al., 2011], [Vince, 2017], or the subsequent literature surveyed by Dossou-Olory [Dossou-Olory, 2023]. This novelty claim is scoped to that cited corpus and targeted keyword/title searches conducted during manuscript preparation ("connected set partitions", "partition polynomial", and spanning-subgraph monotonicity variants).*

Remark 3.3 (Poset Interpretation of Previous Results). *The three identity theorems below compute $p_{\mathbb{G}}$ at three distinguished loci of the spanning subgraph poset of K_n:*

- Trees *are the minimal connected spanning subgraphs (the floor of the poset). The Tree Theorem says $p_{\mathbb{G}}$ is* constant *across this entire antichain: $p_{\mathbb{G}}(T) = 2^{n-1}$ for every tree T.*
- The complete graph *is the unique maximum of the poset. The Complete-Graph Theorem says $p_{\mathbb{G}}(K_n) = B(n)$.*
- Cycles *sit exactly one edge above canonical spanning trees (add one edge to a path to close the cycle). The Cycle Theorem shows this single step above the floor already produces genuinely graph-dependent counts: $p_{\mathbb{G}}(C_n) = 2^n - n$.*

The connection curvature $\kappa_n^{p_{\mathbb{G}}}$ (Section 8) measures the rate at which $p_{\mathbb{G}}$ grows along a canonical chain from tree to complete graph.

3.2 Evaluations at Distinguished Points

Theorem 3.3 (Tree Theorem). *Let $T = (V,E)$ be a tree on n vertices. Then $p_{\mathbb{G}}(\mathbb{G}_n, T) = 2^{n-1}$.*

Proof. A tree on n vertices has exactly $n-1$ edges, each of which is a bridge. Let $F \subseteq E$ be a subset of edges. Consider the graph $(V, E \setminus F)$: its connected components form a partition π_F of V, and each block of π_F is connected in $(V, E\setminus F)$ and therefore in T (since T contains all edges of the subgraph and more).

Claim: every connected set-partition of T arises as π_F for a unique F.

Surjectivity. Let $\pi = \{V_1, \ldots, V_k\}$ be a connected set-partition of T. Set $F = \{uv \in E : u \in V_i, v \in V_j,\ i \neq j\}$—the edges cut by π. The components of $(V, E\setminus F)$ are exactly $V_1, \ldots, V_k$, because each V_i induces a connected subtree of T and the inter-block edges are precisely the edges of F. So $\pi = \pi_F$.

Injectivity. If $F \neq F'$, then some edge e is in $F \setminus F'$ (or vice versa). Since e is a bridge, removing e separates two vertices that remain connected in $(V, E \setminus F')$. Hence $\pi_F \neq \pi_{F'}$.

The map $F \mapsto \pi_F$ is therefore a bijection from 2^E to the set of connected set-partitions. Since $|E| = n-1$, the count is $|2^E| = 2^{n-1}$. □

Theorem 3.4 (Complete-Graph Theorem). *For the complete graph K_n, $p_{\mathbb{G}}(\mathbb{G}_n, K_n) = B(n)$.*

Proof. Let $\pi = \{V_1, \ldots, V_k\}$ be any set partition of $V(K_n)$. For each block V_i, the induced subgraph $K_n[V_i]$ is a complete graph on $|V_i|$ vertices, which is connected whenever $|V_i| \geq 1$. Hence every set partition is a connected set-partition, and the connected set-partition count equals the total set-partition count $B(n)$. □

Theorem 3.5 (Cycle Theorem). *For the cycle graph C_n on $n \geq 3$ vertices, $p_{\mathbb{G}}(\mathbb{G}_n, C_n) = 2^n - n$.*

Proof. Let $f_n := p_{\mathbb{G}}(C_n)$ for $n \geq 3$, and let $g_n := p_{\mathbb{G}}(P_n) = 2^{n-1}$ (Tree Theorem, since P_n is a tree).

Fix an edge $e = uv$ of C_n and partition connected set-partitions of C_n into two classes:

- $\mathcal{A}_n$: u, v lie in the same block,

- $\mathcal{B}_n$: u, v lie in different blocks.

Then $f_n = |\mathcal{A}_n| + |\mathcal{B}_n|$.

For $\mathcal{A}_n$: contract e. This identifies u, v to one vertex and turns C_n into C_{n-1}. A partition in $\mathcal{A}_n$ pushes forward to a connected set-partition of C_{n-1}, and every connected set-partition of C_{n-1} lifts uniquely by splitting the contracted vertex back into the adjacent pair u, v in the same block. Hence $|\mathcal{A}_n| = f_{n-1}$.

For $\mathcal{B}_n$: delete e. Then $C_n - e = P_n$ with endpoints u, v. A partition in $\mathcal{B}_n$ is exactly a connected set-partition of P_n with u, v in different blocks. Among all connected set-partitions of P_n, there is exactly one with u, v in the same block: the one-block partition, because in a path any connected induced subgraph containing both endpoints is the whole vertex set. Therefore

$$|\mathcal{B}_n| = g_n - 1 = 2^{n-1} - 1.$$

So, for $n \geq 4$,

$$f_n = f_{n-1} + 2^{n-1} - 1,$$

with base value $f_3 = p_{\mathbb{G}}(C_3) = B(3) = 5$. Summing the recurrence:

$$\begin{aligned} f_n &= f_3 + \sum_{k=4}^{n} (2^{k-1} - 1) \\ &= 5 + (2^n - 8) - (n - 3) \\ &= 2^n - n. \end{aligned}$$

Thus $p_{\mathbb{G}}(\mathbb{G}_n, C_n) = 2^n - n$ for all $n \geq 3$. □

Remark 3.4. *The proof exposes the deletion-contraction structure underlying the Cycle Theorem: the cycle case builds on the tree case (the path P_n), with a correction for the single partition in which the deleted edge's endpoints reconnect through the rest of the cycle. This interplay between trees and cycles is the simplest instance of the general principle that geometric partition counts decompose along edge operations.*

4 Validation Protocol

This section describes the validation infrastructure that produced the results reported in this paper. We include the full protocol because the *process*—pre-declared falsifiers, live data pulls, emergent findings—is as important as the results. The validation protocol is described first; the test specification and findings follow in Sections 5 and 6.

4.1 The Geometric Naturals Validation Suite

The Geometric Naturals Research Program maintains 11 core tests (Tests 1–11), plus pre-registered follow-ups (e.g., 2b, 6b, 8b), each governed by five non-negotiable guardrails:

G1. No circularity. No test uses Davis Field Equation outputs ($C = \tau/K$, $S + d^2$, holonomy) as ground truth. Ground truth is always external to the framework being tested.

G2. Pre-declared falsifiers. Every test has a falsification condition frozen before data is pulled. If the falsifier fires, the associated conjecture is retracted, not patched.

G3. Frozen metrics. No test is allowed to redefine its success criterion during or after execution. Metrics are locked before data collection begins.

G4. Independent data generators. Each test names the independent generator of its data and confirms that the generator has no knowledge of $\mathbb{G}$-arithmetic.

G5. Convergence standard. A conjecture is considered supported only when ≥ 3 tests from ≥ 3 independent domains agree (the "Sudoku principle").

These guardrails are adapted from the validation discipline of the Davis Duality program. Their purpose is to make it *structurally difficult* to confirm a conjecture through circular reasoning, post-hoc metric adjustment, or selective reporting.

4.2 Pre-Registration

The validation suite was frozen on April 10, 2026, before any computational test was executed. The frozen document specifies, for each test: the entry it validates, the external data source, the procedure, the success criterion, the falsification condition, and the circularity audit.

5 Test 1: OEIS Path Recursion

5.1 Specification

Entry

1 (Geometric Partitions).

Data Source

The On-Line Encyclopedia of Integer Sequences (OEIS, https://oeis.org) [OEIS Foundation, 2025]. Sequences: A000045 (Fibonacci numbers), A000071 (Fibonacci − 1), A000079 (powers of 2), A000110 (Bell numbers).

Procedure

Compute $p_{\mathbb{G}}(P_n)$ for $n = 2, \ldots, 20$ by exact enumeration (bitmask dynamic programming over connected set-partitions). Submit the computed sequence as a candidate to the OEIS search API. Compare against known sequences.

Success Criterion

The computed recursion matches a known OEIS sequence for small n *and* diverges at a pre-declared index (expected $n \geq 8$) in a way the recursion predicts.

Falsifier

$p_{\mathbb{G}}(P_n) = 2^{n-1}$ for all tested n. If this holds, the geometric partition function on path graphs collapses to integer compositions and Entry 1 is retracted for path (and by extension tree) topologies. Entry 1 survives only on non-tree topologies.

Circularity Audit

CLEAN. The OEIS predates $\mathbb{G}$-arithmetic by decades. The sequences were computed by independent researchers with no knowledge of this framework.

5.2 Implementation

The test was implemented in Python 3.12 using NetworkX 3.6.1 for graph construction. The core algorithm computes $p_{\mathbb{G}}(\mathbb{G}_n, G)$ by bitmask dynamic programming:

1. Represent each vertex subset $S \subseteq V$ as a bitmask.
2. For each remaining vertex set (bitmask), enumerate all connected subsets containing the lowest-index remaining vertex via breadth-first expansion of adjacent bits.
3. Recursively count partitions of the remainder.

4. Memoize on bitmask to avoid recomputation.

The algorithm runs in $O(3^n)$ time (each vertex is in the current block, in the remainder, or already partitioned) and was verified against an independent algebraic dynamic-programming formula for path graphs: $f(0) = 1$, $f(k) = \sum_{j=1}^{k} f(k-j) = 2^{k-1}$.

The brute-force and algebraic DP methods agree exactly for $n = 1, \ldots, 16$. The algebraic DP extends to $n = 20$ without additional computational cost.

OEIS data was fetched live on April 10, 2026, via the OEIS b-file and search APIs.

5.3 Execution

The test was executed on April 10, 2026, at 15:09 UTC. Total wall-clock time: approximately 3 seconds for all graph families up to $n = 16$ (brute force) and $n = 10$ (complete graphs and cycles).

6 Findings

6.1 The Falsifier Fires

Theorem 6.1 (Path-Graph Collapse). *For all $n \geq 1$,*

$$p_{\mathbb{G}}(\mathbb{G}_n, P_n) \;=\; 2^{n-1},$$

where P_n is the path graph on n vertices.

This was confirmed by exact enumeration for $n = 1, \ldots, 16$ (bitmask DP) and by the algebraic DP formula for $n = 1, \ldots, 20$. The computed sequence $(1, 2, 4, 8, 16, 32, 64, \ldots)$ was submitted to the OEIS search API, which returned **A000079** (powers of 2) as the top match.

The pre-declared falsifier fires: $p_{\mathbb{G}}(P_n) = 2^{n-1}$ for all tested n. Entry 1 is retracted for path-graph topologies. The geometric partition function on paths is integer compositions—pure edge-subset counting with no geometric content beyond the flat case.

Remark 6.1. *This result was* predicted *by the round-2 external adversarial review of the research program, which identified the isomorphism between cutting path edges and integer compositions. The test confirmed the prediction exactly.*

6.2 Emergent Result 1: The Tree Theorem

The falsifier led to a stronger result that was not pre-declared.

Theorem 6.2 (Tree Theorem). *For every tree T on n vertices, $p_{\mathbb{G}}(\mathbb{G}_n, T) = 2^{n-1}$, independent of the topology of T.*

The proof is in Section 3 (Theorem 3.3).

This was verified computationally on:

- Path graphs P_n for $n = 1, \ldots, 16$.
- Star graphs S_n for $n = 2, \ldots, 10$.
- Five random labeled trees on $n = 8$ vertices (seeds $0, \ldots, 4$).

All returned $p_{\mathbb{G}} = 2^{n-1}$.

Remark 6.2 (Trees Are the Flat Case). *The Tree Theorem establishes that trees are to $\mathbb{G}$-partitions what Euclidean space is to Riemannian geometry: the degenerate case where all curvature vanishes. On a tree, every edge is a bridge, every partition reduces to edge-subset selection, and the partition count depends only on n—not on the tree's shape. The geometric content of $p_{\mathbb{G}}(\mathbb{G}_n, G)$ activates only when G contains cycles.*

This mirrors the Zero paper's central claim: Peano arithmetic (the path-graph restriction of $\mathbb{G}$) is the flat limit of a richer geometric arithmetic. Here, one level up, tree partitions are the flat limit of a richer geometric partition theory. The flat-vs-curved split reappears.

6.3 Emergent Result 2: Complete Graphs Recover Bell Numbers

Theorem 6.3 (Complete-Graph Theorem)**.** *For the complete graph K_n on n vertices, $p_{\mathbb{G}}(\mathbb{G}_n, K_n) = B(n)$, the n-th Bell number.*

The proof is in Section 3 (Theorem 3.4).
This was confirmed by exact enumeration for $n = 2, \ldots, 10$, with independent verification against the Bell triangle:

n	$p_{\mathbb{G}}(K_n)$	$B(n)$	Match
2	2	2	✓
3	5	5	✓
4	15	15	✓
5	52	52	✓
6	203	203	✓
7	877	877	✓
8	4 140	4 140	✓
9	21 147	21 147	✓
10	115 975	115 975	✓

The OEIS search on the sequence $(2, 5, 15, 52, 203, 877, 4140, 21147, 115975)$ returned **A000110** (Bell numbers) as the top match.

Remark 6.3 (The Complete-Graph Limit)**.** *This result resolves a definitional ambiguity flagged by the first round of adversarial review. The original conjecture (C1) stated that $p_{\mathbb{G}}(K_n) = 1$ under "clean decomposition." Under the clean-partition definition (no inter-block edges), this is trivially true for all connected graphs—and therefore uninformative. Under the connected-partition definition adopted in this paper, $p_{\mathbb{G}}(K_n) = B(n)$: the classical set-partition count.*

We regard this as a feature*, not a failure. Peano arithmetic is the path-graph limit of $\mathbb{G}$. Classical set-partition theory is the complete-graph limit. The nontrivial geometric content lives between these boundaries:*

Peano is the path limit. Set partitions are the complete-graph limit. $\mathbb{G}$-arithmetic lives between them.

6.4 Emergent Result 3: Cycles and OEIS A000325

Theorem 6.4 (Cycle Theorem)**.** *For the cycle graph C_n on $n \geq 3$ vertices, $p_{\mathbb{G}}(\mathbb{G}_n, C_n) = 2^n - n$.*

The proof is in Section 3 (Theorem 3.5). The result was confirmed by exact enumeration:

n	$p_{\mathbb{G}}(C_n)$	$2^n - n$	$p_{\mathbb{G}}(T_n) = 2^{n-1}$	$B(n)$
3	5	5	4	5
4	12	12	8	15
5	27	27	16	52
6	58	58	32	203
7	121	121	64	877
8	248	248	128	4 140
9	503	503	256	21 147
10	1 014	1 014	512	115 975

The OEIS search on $(5, 12, 27, 58, 121, 248, 503, 1014)$ returned **A000325** ($a(n) = 2^n - n$) as the top match.

Remark 6.4 (Cycles Are the First Honest $\mathbb{G}$-Family)**.** *Cycles are where $\mathbb{G}$-partitions stop being relabeling of classical constructions. On trees, $p_{\mathbb{G}}$ reduces to bridge-counting (2^{n-1}, classical). On K_n, $p_{\mathbb{G}}$ reduces to set-partition counting ($B(n)$, classical). On C_n, the count $2^n - n$ is neither: it interpolates between the tree count and the Bell count, and the gap*

$$2^n - n \; - \; 2^{n-1} \; = \; 2^{n-1} - n$$

is the first quantitative measure of geometric content beyond the flat case.

6.5 The Hierarchy

Combining the three theorems:

Corollary 6.5 (Partition Hierarchy)**.** *For all $n \geq 4$,*

$$2^{n-1} \; = \; p_{\mathbb{G}}(T_n) \; < \; p_{\mathbb{G}}(C_n) \; = \; 2^n - n \; < \; p_{\mathbb{G}}(K_n) \; = \; B(n).$$

Both inequalities are strict.

Proof. For the left inequality: $2^n - n > 2^{n-1}$ iff $2^{n-1} > n$, which holds for $n \geq 4$ (and for $n = 3$, $p_{\mathbb{G}}(C_3) = 5 > 4 = p_{\mathbb{G}}(T_3)$).

For the right inequality: $B(n) > 2^n - n$ for all $n \geq 4$ (and for $n = 3$, $B(3) = 5 = 2^3 - 3$, so equality holds at $n = 3$; for $n \geq 4$, $B(n)$ grows super-exponentially while $2^n - n$ is exponential). $\square$

6.6 Summary Comparison Table

n	P_n	S_n	C_n	K_n	$B(n)$
1	1	—	—	—	1
2	2	2	—	2	2
3	4	4	5	5	5
4	8	8	12	15	15
5	16	16	27	52	52
6	32	32	58	203	203
7	64	64	121	877	877
8	128	128	248	4 140	4 140
9	256	256	503	21 147	21 147
10	512	512	1 014	115 975	115 975

P_n and S_n (both trees) give 2^{n-1}. K_n gives $B(n)$. C_n gives $2^n - n$, strictly between the two for $n \geq 4$.

7 Conjecture Dispositions

The pre-declared conjectures from the research program are updated in light of Test 1:

ID	Statement	Status	Disposition
C1	$p_{\mathbb{G}}(K_n) = 1$ for $n \geq 2$	**Retracted**	Definitional error. Under connected-partition, $p_{\mathbb{G}}(K_n) = B(n)$. Reframed as sanity check.
C2	$p_{\mathbb{G}}(P_n)$ is Fibonacci-like	**Retracted**	$p_{\mathbb{G}}(P_n) = 2^{n-1}$, not Fibonacci.
C3	Sub-exponential growth on dense graphs	**Amended**	K_n gives $B(n)$ (super-exponential).
			Restrict to sparse-but-cyclic families.
C5	$p_{\mathbb{G}}(T) = 2^{n-1}$ for all trees	**Proven**	Theorem 6.2.
C6	$p_{\mathbb{G}}(C_n) = 2^n - n$	**Proven**	Theorem 6.4. OEIS A000325 match.

8 Connection Curvature of the Cycle Partition Bundle

The Cycle Theorem gives a closed-form sequence $p_{\mathbb{G}}(C_n) = 2^n - n$. The ratio of consecutive terms defines a discrete *connection*:

Definition 8.1 (Partition Connection). *The* partition connection *on the cycle family is*

$$\Gamma_n \;=\; \frac{p_{\mathbb{G}}(C_{n+1})}{p_{\mathbb{G}}(C_n)} \;=\; \frac{2^{n+1} - (n+1)}{2^n - n}.$$

The connection curvature *is the deviation from the flat (tree) limit* $\Gamma = 2$*:*

$$\kappa_n^{p_{\mathbb{G}}} \;=\; \Gamma_n - 2 \;=\; \frac{2^{n+1} - n - 1 - 2(2^n - n)}{2^n - n} \;=\; \frac{n-1}{2^n - n}.$$

Proposition 8.1 (Monotone Decay). *For $n \geq 3$, $\kappa_n^{p_{\mathbb{G}}} > 0$ and $\kappa_n^{p_{\mathbb{G}}}$ is strictly decreasing. Moreover, $\kappa_n^{p_{\mathbb{G}}} \to 0$ as $n \to \infty$.*

Proof. Positivity: $n - 1 > 0$ and $2^n - n > 0$ for $n \geq 3$.

For monotonicity, it suffices to show that $f(n) = (n-1)/(2^n - n)$ is decreasing for $n \geq 3$. We have

$$\frac{f(n+1)}{f(n)} = \frac{n}{n-1} \cdot \frac{2^n - n}{2^{n+1} - n - 1}.$$

Since $2^{n+1} - n - 1 > 2(2^n - n)$ for $n \geq 3$ (because $2^{n+1} - n - 1 - 2(2^n - n) = n - 1 > 0$), we have

$$\frac{2^n - n}{2^{n+1} - n - 1} < \frac{1}{2}.$$

Meanwhile, $n/(n-1) \leq 3/2$ for $n \geq 3$. So $f(n+1)/f(n) < (3/2)(1/2) = 3/4 < 1$, and f is strictly decreasing.

Convergence to zero: $(n-1)/(2^n - n) < n/2^{n-1} \to 0$. □

Remark 8.1 (Interpretation). *The connection curvature measures how much the cycle family "remembers" it is not a tree family. Large cycles approach trees in their partition behavior: $p_{\mathbb{G}}(C_n)/p_{\mathbb{G}}(T_n) = (2^n - n)/2^{n-1} = 2 - n/2^{n-1}$, which tends to 2 from below. The curvature $\kappa_n^{p_{\mathbb{G}}}$ decays roughly as $n/2^n$, so the "geometric excess" over trees vanishes exponentially fast.*

The contraction ratio $\kappa_n/\kappa_{n-1} \to 1/2$ suggests the curvature lives on an exponential scale (with polynomial prefactor). This is characteristic of partition-theoretic quantities in the exponential regime.

n	$p_{\mathbb{G}}(C_n)$	Γ_n	$\kappa_n^{p_{\mathbb{G}}}$	κ_n/κ_{n-1}
3	5	2.400 0	0.400 0	—
4	12	2.250 0	0.250 0	0.625
5	27	2.148 1	0.148 1	0.593
6	58	2.086 2	0.086 2	0.582
7	121	2.049 6	0.049 6	0.575
8	248	2.028 2	0.028 2	0.569
9	503	2.015 9	0.015 9	0.563
10	1 014	2.008 9	0.008 9	0.557

Table 1: Connection curvature of the cycle partition bundle. Γ_n is the multiplicative transport from C_n to C_{n+1}. $\kappa_n^{p_{\mathbb{G}}}$ decays monotonically to zero, confirming the approach to the flat (tree) limit $\Gamma = 2$. The contraction ratio κ_n/κ_{n-1} approaches 1/2 from above.

Corollary 8.2 (Bell Collapse). $p_{\mathbb{G}}(C_n)/B(n) \to 0$ *as* $n \to \infty$.

Proof. $p_{\mathbb{G}}(C_n) = 2^n - n = O(2^n)$, while $B(n) \sim (n/W(n))^n/e^{n-1}$ grows super-exponentially (where W is the Lambert W function). The ratio tends to zero. □

This confirms that cycle topology becomes a progressively stricter filter on the partition lattice: at $n = 3$, $p_{\mathbb{G}}(C_3) = B(3)$ (no filtering); at $n = 10$, $p_{\mathbb{G}}(C_{10})/B(10) < 1\%$.

9 The Lattice Order in Practice

The Subgraph Lattice Theorem (Theorem 3.1) establishes that $p_{\mathbb{G}}$ is monotone on the spanning subgraph poset. This section reports an empirical study of how this abstract ordering manifests across concrete graph families, and how well standard topological invariants approximate it.

9.1 Experimental Design

For each $n \in \{4, \ldots, 10\}$, we computed $p_{\mathbb{G}}(G)$ for 10 graph families: path P_n, star S_n, cycle C_n, wheel W_n, complete bipartite $K_{2,n-2}$, fan F_n, ladder L_n (even n), prism Π_n (even $n \geq 6$), friendship Fr_n (odd n), and complete K_n. The Petersen graph was included at $n = 10$. Eight topological invariants were recorded: edge density, vertex connectivity κ, chromatic number, treewidth (upper bound), average degree, circuit rank, average clustering coefficient, and maximum degree.

At each fixed n, Pearson and Spearman rank correlations were computed between $\log p_{\mathbb{G}}(G)$ and each invariant across the available graph families. All computations are exact (bitmask DP); no sampling or approximation is involved.

9.2 The Fiber at Fixed n

At $n = 6$, the partition counts span a $6.3\times$ range:

Graph	$p_{\mathbb{G}}$	κ	m	% of $B(6)$
P_6, S_6 (trees)	32	1	5	15.8%
C_6 (cycle)	58	2	6	28.6%
Ladder$_6$	74	2	7	36.5%
Fan$_6$	89	2	9	43.8%
$K_{2,4}$	96	2	8	47.3%
Prism$_6$	114	3	9	56.2%
W_6 (wheel)	118	3	10	58.1%
K_6 (complete)	203	5	15	100.0%

The spread grows super-exponentially: at $n = 10$, the range is 512 to 115,975 (226×). This confirms that $p_{\mathbb{G}}$ at fixed n carries genuine topological information.

9.3 Vertex Connectivity as a Scalar Proxy

The vertex connectivity $\kappa(G)$ is the best single-invariant predictor of $p_{\mathbb{G}}$-rank at fixed n:

n	Spearman ρ	p-value
4	0.961	$< 10^{-3}$
5	0.963	$< 10^{-3}$
6	0.953	$< 10^{-3}$
7	0.881	0.004
8	0.953	$< 10^{-3}$
9	0.831	0.011
10	0.912	$< 10^{-3}$

Treewidth (upper bound) performs comparably, achieving $\rho \geq 0.87$ at all seven tested n values. Edge density and circuit rank achieve $\rho \geq 0.90$ for $n = 4, \ldots, 9$, dropping to 0.78 at $n = 10$.

Negative control. Maximum degree shows no significant correlation ($\rho \approx 0.35$, $p > 0.2$ at every n). In this tested family panel, maximum degree is not predictive of $p_{\mathbb{G}}$-rank.

9.4 Counterexamples to κ-Monotonicity

Strict κ-monotonicity—the conjecture that $\kappa(G_1) > \kappa(G_2)$ implies $p_{\mathbb{G}}(G_1) \geq p_{\mathbb{G}}(G_2)$ at fixed n—is **false**. The first violation occurs at $n = 7$:

n	Graph A	κ	$p_{\mathbb{G}}$	Graph B	κ	$p_{\mathbb{G}}$
7	Fr$_7$	1	125	C_7	2	121
9	Fr$_9$	1	625	C_9	2	503
9	$K_{2,7}$	2	2 314	W_9	3	2 199
10	$K_{2,8}$	2	6 816	W_{10}	3	5 769

In every case, the lower-κ graph has *higher* $p_{\mathbb{G}}$. These pairs are all *spanning-subgraph-incomparable*: neither graph's edge set contains the other's. The Subgraph Lattice Theorem makes no prediction for incomparable pairs, and indeed $p_{\mathbb{G}}$ can go either way.

Structural explanation. The Friendship graph Fr_n has $\kappa = 1$ (cut vertex at the hub) but $\frac{3(n-1)}{2}$ edges, which exceeds the cycle's n edges for $n \geq 5$. Its edge *distribution*—dense triangles radiating from a hub—creates abundant connection paths despite the global bottleneck. At $n = 5$, the cycle barely wins ($p_{\mathbb{G}} = 27 > 25$); by $n = 7$, the friendship's edge advantage overwhelms the connectivity deficit.

Similarly, $K_{2,n-2}$ (bipartite hub structure, $\kappa = 2$) defeats the wheel W_n ($\kappa = 3$) at $n \geq 9$ despite having fewer edges, because its two-hub topology distributes connectivity more efficiently for partition purposes than the wheel's single-hub-plus-rim topology.

The correct partial order. These counterexamples confirm that κ-monotonicity is a *statistical tendency* ($\rho \geq 0.83$), not a theorem, and that the spanning subgraph poset (Theorem 3.1) is the correct underlying structure. The κ-correlation is an approximate projection of the lattice order onto a scalar invariant: graphs with higher κ *tend* to sit higher in the poset (more edges, more connection paths), but the relationship is not order-preserving because κ is a global minimum-cut property while the lattice ordering is local (edge-by-edge containment).

9.5 Verification of the Lattice Theorem

Among the 60+ graph instances across $n = 4, \ldots, 10$, we verified all 108 spanning-subgraph-comparable pairs (where $E(G) \subseteq E(H)$ could be checked directly). In every case, $p_{\mathbb{G}}(G) \leq p_{\mathbb{G}}(H)$. The remaining 244 ordered pairs were spanning-subgraph-incomparable. Zero violations were found.

10 Test 4: Edge vs. Vertex Succession in the Internet AS Graph

The $\mathbb{G}$-arithmetic framework distinguishes two succession operations on graphs: *vertex succession* S_v (adding a new vertex with at least one edge) and *edge succession* S_e (adding an edge between existing vertices). The Zero paper [Davis, 2025] predicts that real networks accumulate S_e faster than S_v as they mature: once a network's vertex set stabilises, growth is dominated by new connections rather than new participants (§5.2, *ibid.*).

10.1 Pre-Registration

We pre-registered the following before accessing data:

- **Data source.** CAIDA AS-relationships serial-1 dataset [Luckie et al., 2013], publicly archived at `publicdata.caida.org`.
- **Period.** January 2010 through December 2024, one snapshot per month (180 snapshots).
- **Definitions.** For consecutive months $t, t{+}1$: S_v is the count of ASes present in $t{+}1$ but absent in t; S_e is the count of edges in $t{+}1$ (but not t) whose *both* endpoints existed in t.
- **Hypothesis.** The annual ratio S_e/S_v increases with a linear slope exceeding 0.05 per year.
- **Falsifier.** The ratio is flat or declining.

10.2 Results

The Internet AS graph grew from 33,778 ASes and 93,998 edges (January 2010) to 77,495 ASes and 501,473 edges (December 2024). Across 179 consecutive monthly deltas, the annual S_e/S_v ratio rose from 8.4 (2010) to 47.0 (2024):

Year	S_v	S_e	S_e/S_v
2010	5,266	44,059	8.37
2012	6,097	72,341	11.87
2014	5,761	85,582	14.86
2016	7,000	155,895	22.27
2018	7,552	269,022	35.62
2020	8,183	350,705	42.86
2022	6,483	292,738	45.15
2024	6,720	315,567	46.96

Linear regression on the 15 annual aggregates: slope = 3.68/year, $R^2 = 0.88$, $p = 2.2 \times 10^{-7}$. Spearman $\rho = 0.968$, $p = 3.6 \times 10^{-9}$. The observed slope exceeds the pre-registered threshold by a factor of 73.

10.3 Robustness

We performed six robustness checks:

1. **Monthly regression** (all 179 deltas): slope = 3.52/year, $R^2 = 0.41$, $p = 6.4 \times 10^{-22}$, Spearman $\rho = 0.88$.
2. **Outlier exclusion.** Eight months were flagged as potential methodology artefacts (see below) and excluded. The annual regression on the remaining 171 deltas: slope = 3.09/year, $R^2 = 0.93$, $p = 8.0 \times 10^{-9}$. Excluding outliers *improves* the fit.
3. **Stable methodology window** (2010–2018-08, before the corrupted `bgpdump` era; see below): slope = 2.56/year, $R^2 = 0.46$, $p < 10^{-14}$ (monthly), $R^2 = 0.72$, $p = 0.004$ (annual on 9 years).
4. **Mixed-edge convention.** Each monthly delta includes "mixed" edges connecting one new and one existing AS. Under three conventions—exclude (default), split each mixed event into one S_v+ one S_e, or count all mixed as S_e—the annual slope remains 1.07–3.74/year. All exceed the threshold.
5. **All regressions pass.** Every combination of aggregation level (monthly/annual), outlier treatment (include/exclude), and mixed-edge convention produces slope > 0.05/year with $p < 0.005$.

10.4 Methodology Caveats

The CAIDA serial-1 inference pipeline has evolved since 2010 [Luckie et al., 2013]. The README documents:

- March 2016: the Tier-1 clique was forced to a static ASN list.
- July 2017–August 2019: automated clique inference, then reverted to static.
- September 2018–December 2022: `bgpdump` v1.5.0 was in use and was later found corrupted; January 2023 switched to v1.6.2 without re-inferring prior data.

The eight excluded outlier months—including the January 2023 switchover ($S_e/S_v = 228$, versus a median of ~ 30 in that era)—correlate with these known pipeline changes. The trend survives their exclusion and survives restriction to the pre-corruption window (2010–2018), so the result is not an artefact of pipeline drift.

A residual confound remains: CAIDA's ability to observe peering links improves over time as more BGP monitors are deployed and inference algorithms are refined. We cannot fully separate "the Internet is fattening" from "CAIDA sees more of the Internet's fat." Cross-validation against an independent data source (e.g., RIPE RIS, RouteViews raw tables) is noted as future work.

10.5 Interpretation

The result confirms the Zero paper's prediction that edge succession dominates vertex succession in mature networks. Vertex accretion has been roughly flat at ∼6–8K new ASes per year since 2010, while edge accretion grew from ∼44K to ∼316K new edges per year. The S_e/S_v ratio is not just greater than 1—it is an order of magnitude above 1 even in 2010, and grows to nearly 50× by 2024. This is consistent with the $\mathbb{G}$-arithmetic prediction that mature graphs "fatten" via S_e rather than growing via S_v.

This is the first test in the validation program that produces a *positive result on a prediction the framework made about the real world*, not about pure combinatorics. The framework predicted "S_e grows relative to S_v as networks mature," and the prediction is quantitatively confirmed on 15 years of independent data with a 73× safety margin on the pre-registered threshold.

11 Test 5: Path-Graph Fraction in Citation Subgraphs

A central claim of the $\mathbb{G}$-arithmetic framework is that Peano's successor function models graphs as paths—a structure that rarely occurs in nature. Test 5 quantifies this: among real co-citation subgraphs, what fraction are path graphs?

11.1 Pre-Registration

- **Data source.** OpenAlex public API (no authentication), articles published 2020–2024 with 5–15 referenced works.
- **Sample.** 10,000 seed works, fetched via cursor pagination sorted by citation count (deterministic order).
- **Graph construction.** For each seed with k references $R_1, \ldots, R_k$, build an undirected graph on k nodes with an edge between R_i and R_j if one cites the other.
- **Path test.** A subgraph is a path graph iff it is connected, has maximum degree ≤ 2, and has exactly $n-1$ edges.
- **Hypothesis.** $< 5\%$ of subgraphs are path graphs.
- **Falsifier.** $> 30\%$ are paths.

11.2 Results

Of 10,000 subgraphs (104,780 total references, 76,240/81,750 unique ref metadata fetched):

Category	Count (%)
Total subgraphs	10,000
No edges (isolated)	1,592 (15.9%)
Connected	632 (6.3%)
Trees	89 (0.9%)
Path graphs	**3 (0.03%)**

All three path graphs occurred at the smallest subgraph sizes ($n = 5$ or 6). No path graphs were found at $n \geq 7$ (0 of 527 connected subgraphs).

The pre-registered metric (full-sample path fraction) is $3/10{,}000 = 0.03\%$, which satisfies the criterion $< 5\%$. The connected-conditional effect size is $3/632 = 0.47\%$, also below 5%. We report both because the path-vs-nonpath question is structurally meaningful only on connected subgraphs.

An incidental structural observation is that 93.7% of sampled co-citation subgraphs are disconnected: for most papers, references do not form a single connected citation chain among themselves.

Post-hoc sensitivity (not pre-registered). To separate citation-specific structure from pure combinatorial rarity, we compared against a connected Erdős–Rényi null matched on n and edge density for each observed connected subgraph. The null expects 6.52 paths among 632 connected subgraphs (1.03%, 95% interval 2–12 paths), while we observed 3 paths (0.47%). The depletion direction is consistent with stronger path-avoidance in citation structure, but the one-sided probability $P(\text{null} \leq \text{observed}) = 0.106$ is not decisive. Accordingly, we treat this as sensitivity analysis, not a primary test.

11.3 Interpretation

Real co-citation subgraphs are overwhelmingly non-path. Peano's linear successor structure—which models natural numbers as a path graph P_n—captures essentially none of the topology present in citation networks. Even among the 632 connected subgraphs, only 3 (0.47%) are paths. The remaining 99.5% require genuine graph structure to represent their citation relationships.

This result is unsurprising from the perspective of bibliometrics (citation networks are known to have rich structure), but it provides the quantitative bound that the $\mathbb{G}$ framework needs: in at least one canonical real-world graph domain, the Peano projection is rarely realized. The post-hoc null indicates that part of this effect is combinatorial rarity at moderate n, with possible additional citation-specific depletion that requires stronger pre-registered nulls in follow-up work.

12 Test 6: UD Dependency-Tree Path Fraction (Falsifier Fired)

Test 6 evaluates whether path-graph sentence structures are rare across languages in Universal Dependencies (UD) treebanks. The pre-registered criterion was strict: path-graph sentences must be below 10% in all six languages.

12.1 Pre-Registration

- **Data.** UD test splits for English EWT, Spanish AnCora, German GSD, French GSD, Russian SynTagRus, Chinese GSD.
- **Graph construction.** Undirected dependency graph per sentence after removing punctuation tokens (UPOS $\neq$ PUNCT).
- **Eligibility.** Sentences with at least two remaining tokens.
- **Path criterion.** Connected, $m = n - 1$, and maximum degree ≤ 2.
- **Success threshold.** Path fraction $< 10\%$ in every language.
- **Falsifier.** Any language with path fraction $\geq 10\%$.

12.2 Results

Language	Eligible	Path	Path%	Verdict vs 10%
English	1,839	361	19.63%	**Fail**
Spanish	1,713	55	3.21%	Pass
German	966	69	7.14%	Pass
French	416	8	1.92%	Pass
Russian	8,679	733	8.45%	Pass
Chinese	500	0	0.00%	Pass

Overall, 1,226 of 14,113 eligible sentences (8.69%) are paths, but the pre-registered all-language criterion fails because English is above threshold. **Primary verdict: falsifier fires.**

12.3 Post-Hoc Length Sensitivity (Not Pre-Registered)

The failure localizes to very short sentences. When sentence length is restricted post-hoc:

- For $n \geq 5$ tokens, all six languages are below 2% (English 0.99%, Russian 1.57%).
- For $n \geq 7$ tokens, all six are near zero (0.00%–0.15%).

This indicates a threshold-location issue: at very small n, path-like topologies are combinatorially common, and treebank sentence-length distributions differ materially (e.g., English EWT is shorter on average than Spanish AnCora).

12.4 Interpretation

The registered test failed and must be reported as failed. The post-hoc analysis does *not* rescue the primary endpoint; it explains its failure mechanism and informs improved test design. In particular, length-conditioning should be explicitly pre-registered for follow-up.

Test 6 therefore contributes two outcomes simultaneously: a genuine falsifier event (strengthening audit credibility) and a concrete methodological lesson about sentence-length confounding in topology-only linguistic tests.

As a separate pre-registered follow-up, Test 6b fixed a minimum sentence length of $n \geq 5$ and re-ran the same six treebanks with the same threshold (<10% in each language). All six languages satisfied the criterion (overall 1.17%), so Test 6b is supported. This does not alter Test 6's primary failure; it confirms that the original falsifier was a threshold-location issue rather than a broad structural failure of the framework's non-path prediction.

13 Validation Program Snapshot (Core + Follow-Ups)

To finalize this draft, Table 2 consolidates the status of the 11 core tests plus registered follow-ups as of April 10, 2026.

Test	Status	Summary
Core tests (1–11)		
1	Complete (falsifier fires)	Path case collapses to 2^{n-1}.
2	Closed (negative)	No external physical correlate.
3	Not executed in this cycle	Spec frozen; execution pending.
4	Complete (supported)	CAIDA shows strong S_e/S_v growth.
5	Complete (supported)	Paths are rare in citation subgraphs.
6	Complete (falsifier fires)	English exceeds pre-registered 10% cap.
7	Complete (null)	Sparse-floor effect; ER null underpowered.
8	Complete (falsifier fires)	14/18 hits; below pre-registered target.
9	Blocked (data observability)	GH Archive payload lacks conflict labels.
10	Complete (partial)	Bound holds 230/230; no saturation cases.
11	Skipped (directive)	ENTSO-E test deferred in this run.
Registered follow-ups		
2b	Complete (supported)	Same-n topology contrast is positive.
6b	Complete (supported)	Length-conditioned follow-up passes all languages.
8b	Registered (not run)	Temporal within-lineage follow-up reserved.

Table 2: Validation snapshot at manuscript freeze (April 10, 2026): 11 core tests and registered follow-ups.

Remark 13.1 (Final audit posture)**.** *The program now contains supported results, explicit falsifier events, null outcomes, blocked tests, and one deferred test. This is the intended audit behavior: claims are narrowed where falsifiers fire, mechanisms are retained where positive evidence is strong, and blocked items are documented as operational limits rather than retrofitted into success claims.*

13.1 Program-Level Convergence Verdict

Under guardrail G5 (Section 4.1), convergence for a broad framework claim requires at least three supportive tests from at least three independent domains. On this draft snapshot:

- Supported: Test 4 (network growth, CAIDA), Test 5 (citation topology, OpenAlex), Test 2b/6b (topology and language follow-ups).
- Falsifier events: Test 6 (primary endpoint), Test 8.
- Incomplete for convergence accounting: Test 3 not executed, Test 9 blocked, Test 11 deferred.

Therefore, this draft supports several local claims (especially the combinatorial structure theorems and selected empirical tendencies), but does not yet claim full cross-domain convergence of the entire research program under the pre-declared G5 standard.

14 Open Problems

1. **Complete enumeration at small n.** All connected *unlabeled* graphs up to isomorphism on $n = 5, 6, 7$ vertices (counts: $21, 112, 853$) are tractable for $p_{\mathbb{G}}$ computation. Exhaustive enumeration—using Nauty/geng—would (a) definitively catalogue all κ-counterexamples beyond the four identified in Section 9.4, and (b) determine whether a refined invariant (e.g., algebraic connectivity, Tutte polynomial evaluation) perfectly predicts $p_{\mathbb{G}}$-rank within the spanning-subgraph-incomparable pairs.

2. **Interpolation problem.** Characterize the family $\mathcal{F}$ of connected graphs on n vertices for which $2^{n-1} < p_{\mathbb{G}}(\mathbb{G}_n, G) < B(n)$. Cycles are in $\mathcal{F}$. Trees and complete graphs are the boundary cases. Does the circuit rank $|E| - |V| + 1$ govern the transition?

3. **Generating function for $p_{\mathbb{G}}(C_n)$.** The closed form $2^n - n$ gives the ordinary generating function $\sum_{n\geq 3}(2^n - n)x^n$. Is there a species-theoretic interpretation [Joyal, 1981, Bergeron et al., 1998]?

4. **General graph asymptotics.** For G drawn uniformly from connected graphs on n vertices, what is the asymptotic growth of $\mathbb{E}[p_{\mathbb{G}}(\mathbb{G}_n, G)]$?

5. **Connection curvature for other families.** Define $\kappa_n^{p_{\mathbb{G}}}$ for wheels, prisms, Möbius–Kantor graphs, etc. Is there a universal decay law? Does the decay rate encode a topological invariant of the graph family?

6. **Peano-to-$\mathbb{G}$ functor (Entry 4).** The Subgraph Lattice Theorem provides a concrete test: the functor must preserve the lattice monotonicity of $p_{\mathbb{G}}$. The κ-correlation gives Entry 4 a structural invariance to verify categorically.

7. **Connectivity-matched null for sparse networks.** The floor-tightness observation (Remark 3.1) leaves open whether the concentration at 2^{n-1} is a pure sparsity effect or reflects biological selection for tree-like local topology. Conditioning on connectedness and edge count (rather than Erdős–Rényi at matched density) is the appropriate stricter null. This test is pre-registered for a follow-up study.

References

Andrews, G. E. (1976). *The Theory of Partitions*. Cambridge University Press.

Bergeron, F., Labelle, G., and Leroux, P. (1998). *Combinatorial Species and Tree-Like Structures*. Cambridge University Press.

Davis, B. R. (2025). Zero does not exist: A geometric foundation for the natural numbers. Davis Geometric.

Hardy, G. H. and Ramanujan, S. (1918). Asymptotic formulæ in combinatory analysis. *Proceedings of the London Mathematical Society*, s2-17(1):75–115.

Joyal, A. (1981). Une théorie combinatoire des séries formelles. *Advances in Mathematics*, 42(1):1–82.

OEIS Foundation (2025). The On-Line Encyclopedia of Integer Sequences. https://oeis.org.

Rademacher, H. (1937). On the partition function $p(n)$. *Proceedings of the London Mathematical Society*, s2-43(1):241–254.

Wiberg, K. B. (1986). The concept of strain in organic chemistry. *Angewandte Chemie International Edition in English*, 25(4):312–322.

Blanksby, S. J. and Ellison, G. B. (2003). Bond dissociation energies of organic molecules. *Accounts of Chemical Research*, 36(4):255–263.

Pretsch, E., Bühlmann, P., and Badertscher, M. (2009). *Structure Determination of Organic Compounds*. Springer, 4th edition.

Simon, F., Tittmann, P., and Trinks, M. (2011). Counting connected set partitions of graphs. *Electronic Journal of Combinatorics*, 18(1):#P14.

Vince, A. (2017). Counting connected sets and connected partitions of a graph. *Australasian Journal of Combinatorics*, 67:281–290.

Dossou-Olory, A. A. V. (2023). On the maximum number of connected induced subgraphs of a graph. arXiv:2303.01964.

Varshney, L. R., Chen, B. L., Paniagua, E., Hall, D. H., and Chklovskii, D. B. (2011). Structural properties of the *Caenorhabditis elegans* neuronal network. *PLoS Computational Biology*, 7(2):e1001066.

Luckie, M., Huffaker, B., claffy, k., Dhamdhere, A., and Giotsas, V. (2013). AS relationships, customer cones, and validation. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pp. 243–256.

A Fiber Bundle Audit of Test 2

This appendix summarizes the fiber-bundle re-analysis that informed the framing of this paper. Validation Test 2 of the research program attempted to find an external physical-chemistry correlate of the cycle partition function. The attempt failed cleanly, and the analysis of *why* it failed produced the key structural insight that partition-derived quantities at fixed topology are trivial sections of the base-n bundle.

A.1 The Failed Correlate

Test 2 correlated the *excess ratio* $\varepsilon(n) = 1 - n/2^{n-1}$—the fractional excess of $p_{\mathbb{G}}(C_n)$ over the tree count—against cycloalkane ring strain energy [Wiberg, 1986]. The headline result was $r = -0.91$, $p = 0.002$, which initially appeared significant.

A fiber-bundle decomposition revealed this to be a coordinate artifact:

1. $\varepsilon(n)$ is a *deterministic function of* n. It contains no graph information beyond the Cycle Theorem.
2. The function $1/n^2$ achieves $r = +0.93$ against the same strain data. Any smooth concave function of n achieves $|r| > 0.88$.
3. Removing the high-leverage points $n = 3, 4$ collapses the signal to $r \approx 0.3$ (nonsignificant).
4. In log-log coordinates (the natural chart for a multiplicative quantity like $p_{\mathbb{G}}$), the correlation vanishes: $r = -0.21$.
5. The connection profiles $\kappa_n^{p_{\mathbb{G}}}$ and Δ(strain) are structurally incompatible: $\kappa_n^{p_{\mathbb{G}}}$ decays monotonically while ring strain oscillates (U-shaped with minimum at $n = 6$).

A follow-up v4 arm pre-registered three monotonic cycloalkane observables—mean C–C bond dissociation energy [Blanksby and Ellison, 2003], ^{13}C NMR chemical shifts [Pretsch et al., 2009], and conformational inversion barriers—and found *all three* non-monotonic in the range $n = 3$–10, due to the universal transannular strain minimum near $n = 6$. The connection curvature $\kappa_n^{p_{\mathbb{G}}}$ (monotonically decaying, strictly positive) has no structural partner among cycloalkane observables. The external-correlate arm was closed.[1]

A.2 The Structural Insight

The failure was informative. When the data were encoded as a fiber bundle (base space $B = \{3, \ldots, 10\}$, fiber = field schema), the encoding algorithm classified every partition-derived field—$p_{\mathbb{G}}$, tree count, excess, excess ratio—as *computable from the base coordinate* n. Only strain and RNA loop-closure energies appeared as genuine variable sections.

This means: at fixed graph topology (here, cycles), varying n changes everything in lockstep. There are no degrees of freedom in the fiber. Genuine graph-structure content—where $p_{\mathbb{G}}$ varies independently of n—requires comparing *different topologies at fixed* n: $p_{\mathbb{G}}(C_n)$ vs. $p_{\mathbb{G}}(W_n)$ vs. $p_{\mathbb{G}}(K_{2,n-2})$, etc. This is the motivation for Part II of this work.

[1] The raw Pearson correlation between $\kappa_n^{p_{\mathbb{G}}}$ and Δ(^{13}C shift) appears significant ($r = +0.89$, $p = 0.007$). However, the sign pattern of Δ(^{13}C shift) inverts at $n = 7$ while $\kappa_n^{p_{\mathbb{G}}}$ is strictly positive throughout; the correlation is driven by magnitude matching in the $3 \to 4$ transition and fails the structural sign-match test (4/7). We report this to preempt naïve re-analysis of our published data.