Certified Verification of Expansion Testing without Connectivity via Lazy-Walk Return Probabilities

Lluis Eriksson*

November 25, 2026

Abstract

Nikolas Melissaris (`arXiv:2611.32677`, Nov. 2026) broke Dinur and Goldreich's $n^{2/3}$ barrier for testing edge expansion in every connected component of an arbitrary bounded-degree graph, achieving $\tilde{O}_{d,c,\beta}(n^{1/2+\beta}/\varepsilon+\sqrt{n}/\varepsilon^4)$ queries via two-time lazy-walk return probabilities $A_s = P^{2t}(s,s)$ and $B_s = P^{4t}(s,s)$ (Thm. 1.1), alongside an $\Omega(\sqrt{n/\log n})$ Cayley-graph lower bound via cyclic lifts of LPS Ramanujan graphs (Thm. 1.2). We construct a standalone zero-dependency **Lean 4.35.0** formalization of 18 core lazy-walk, finite Cauchy–Schwarz, Dirichlet-energy, trace-gap, integer Bernoulli, induced-expander trimming, balanced-cut farness, non-backtracking polynomial, and three-squares parity theorems (zero `sorry`, zero warnings, zero `Classical.choice`) paired with a zero-floating-point exact $\mathbb{Q}/\mathbb{Z}$ verifier (`EXP-RETPROB-CERT-1.0`). Across 16 canonical certificates (0.180–2.195 ms, total 14.53 ms), all exact lazy-walk return probabilities, Lemma 5.3 trimming maximizers, Lemma 7.3 farness lower bounds, and Lemma 7.5 non-backtracking polynomials replay deterministically while rejecting 60/60 live adversarial forgery attacks.

1 Introduction

Testing expansion in bounded-degree graphs [6, 7] normally assumes connectivity so random walks mix globally toward $1/n$. Dinur and Goldreich [3] posed the unrestricted problem $\Xi_c(n)$—testing whether every connected component has edge expansion $\geq c$ without knowing component sizes—obtaining $n^{2/3+o(1)}$ queries via three-way collisions. Melissaris [9] bypassed three-way collisions through four core mechanisms:

1. **Two-Time Return Probabilities & Heavy-Return Control (Sec. 2.1–2.2, Lem. 5.1):** For lazy walk $P = I - L_G/(2d)$, $A_s = P^{2t}(s,s)$ and $B_s = P^{4t}(s,s)$ satisfy $p_s(v)^2 \leq A_s A_v$ and $\frac{1}{2d}\sum_{\{u,v\}\in E}(p_s(u)-p_s(v))^2 \leq A_s - B_s$. A start s is *good* when $A_s \leq (1+2\eta)B_s$ and $J_s = \sum_{v\in D_s} p_s(v)^2 \leq \eta B_s$ ($D_s = \{v : A_v > (1+8\eta)A_s\}$), certifying a nearly uniform set $H_s = \{v : p_s(v) \geq \frac{3}{4}B_s\} \setminus D_s$ with small boundary and $\mathrm{tr}(K_s^{2t}) \leq \sum_{v\in H_s} A_v$.
2. **Trimming to an Induced Expander (Lem. 5.3):** Maximizing $F(R) = \frac{h}{4}|R| - |E(R, H\setminus R)|$ over $|R| \leq |H|/2$ deletes at most $\frac{4\tau}{3h}|H|$ vertices to yield an induced $(h/4)$-expander $U = H \setminus R$.
3. **Balanced-Cut Farness Lower Bound (Lem. 7.2–7.3):** Any connected graph with $h(G) \geq h_0$ and a balanced cut $|E(A, V\setminus A)| \leq h_1|A|$ ($h_1 < c^*$) requires $M \geq \min(\frac{h_0 n}{2}, h_0|A|, \frac{h_0(c^*-h_1)}{h_0+c^*}|A|)$ edge edits to reach $\Xi_{c^*}(n)$.
4. **Non-Backtracking Polynomials, $r_3(4N) = r_3(N)$, & Walk-Regularity (Lem. 7.4–7.5, App. A):** On d-regular LPS Cayley graphs [8], non-backtracking polynomials $Q_k(A)$ [1] and Gauss's $r_3(4N) = r_3(N)$ [4] bound adaptive transcript distance, while walk-regular inputs [5] omit collision-location walks.

Contributions and Scope. We do not formalize randomized query-complexity concentration inequalities inside Lean 4; instead, we machine-check the exact algebraic, spectral, cut-trimming, and number-theoretic core of [9]: (i) 18 zero-`sorry`, zero-`Classical.choice` **Lean 4.35.0** [2] theorems in `ExpansionReturnProb.lean`; (ii) the exact $\mathbb{Q}/\mathbb{Z}$ verifier `EXP-RETPROB-CERT-1.0`; (iii) 16 canonical certificates across all four modes; and (iv) a 60-attack live red-team suite with 39-artifact SHA-256 manifest enforcement.

2 Constructive Lean 4.35.0 Kernel

The standalone module `ExpansionReturnProb.lean` (in `lean/`) compiles cleanly under Lean 4.35.0 in 1.2 s with zero external dependencies, zero warnings, and zero `sorry`. Auditing all 18 theorems via `#print axioms` (`AuditAxioms.lean`) confirms 1 zero-axiom theorem (`none`), 1 using only `[propext]`, 16 using `[propext, Quot.sound]`, and **0** using `Classical.choice` or `Lean.ofReduceBool` (Table 1).

Proposition 1 (Machine-Checked Spectral & Cut Invariants)**.** *The Lean 4.35.0 kernel establishes the following universal theorems over* `Int` *and* `List Int`*:*

1. Return Probabilities, Bernoulli, & Trimming (Thms. 1–12): *Finite-vector Cauchy–Schwarz* $(\sum a_i b_i)^2 \leq (\sum a_i^2)(\sum b_i^2)$ *and uniform floor* $(\sum p_i)^2 \leq n\sum p_i^2$ *hold for arbitrary equal-length lists;* $(1/2, 1/2)$ *and*

*Independent Researcher. Code, Lean 4 proofs, and certificates: MIT License; manuscript and figures: CC-BY-4.0. Prepared with AI coding, formalization, and drafting assistance from Google Antigravity; all proofs, exact rational/integer verifiers, and adversarial tests were independently compiled and verified locally. Repository: https://github.com/lluiseriksson/expansion-return-prob-certificates

Table 1: All 18 machine-checked Lean 4.35.0 theorems in `ExpansionReturnProb.lean` (zero `sorry`, zero `Classical.choice`).

#	Lean Theorem Identifier	Paper Ref.	Axioms
1	`lazy_walk_row_stochastic_sum`	Sec. 3.1	p, Q
2	`lazy_walk_dirichlet_nonneg`	Lem. 5.1	p, Q
3	`cauchy_schwarz_return_bound_finite`	Eq. (2)	p, Q
4	`collision_prob_uniform_lower_bound`	Eq. (7)	p, Q
5	`two_point_vs_three_point_collision_ambiguity`	Sec. 1	none
6	`good_vertex_dirichlet_energy_monotonicity`	Lem. 5.1	p, Q
7	`heavy_return_set_size_bound`	Lem. 5.1	p, Q
8	`restricted_trace_le_unrestricted_sum`	Eq. (9)	p, Q
9	`bernoulli_lower_bound_int`	Lem. 5.4	p, Q
10	`second_eigenvalue_trace_gap_bound`	Lem. 5.4	p, Q
11	`trimming_maximizer_size_bound`	Lem. 5.3	p, Q
12	`trimming_maximizer_contradiction`	Lem. 5.3	p, Q
13	`farness_balanced_cut_tradeoff`	Lem. 7.3	p, Q
14	`fixed_neighbor_table_compaction_bound`	Rem. 6.2	p, Q
15	`cayley_cyclic_cut_edge_count`	Lem. 7.2	p
16	`nonbacktracking_poly_recurrence_deg3`	Lem. 7.5	p, Q
17	`three_squares_scaling_invariance_step`	Lem. 7.4	p, Q
18	`walk_regular_short_walk_spectral_diff`	Lem. A.2	p, Q

Axiom key: `none` = zero axioms (1); `p` = `[propext]` (1); `p, Q` = `[propext, Quot.sound]` (16).

$(2/3, 1/6, 1/6)$ *share collision probability* $1/2$ *with distinct* F_3*; cleared spectral inequalities* $x^4(D-x) \le D^3x^2(D^2-x^2)$ *and* $(D-b)^k$ *Bernoulli lower bounds hold universally;* $3B \le 4p_i \Rightarrow 9B^2|S| \le 16\sum p_i^2$*; and Lemma 5.3 trimming satisfies* $3h|R| \le 4\tau m$ *and strict improvement* $F(R \cup S) - F(R) > 0$.

2. Farness, Non-Backtracking, & Three-Squares (Thms. 13–18): *Eliminating* s *from* $M \ge h_0 s$ *and* $M \ge (c^*-h_1)A - c^* s$ *yields* $(h_0+c^*)M \ge h_0(c^*-h_1)A$*;* $Q_2(p+1) = (p+1)p$ *and* $Q_3(p+1) = (p+1)p^2$*;* $x_1^2+x_2^2+x_3^2 = 4N$ *forces all three coordinates even; and* $x^2(D-x)^2 \le x^2(D^2-x^2)$ *holds for all* $0 \le x \le D$.

3 Exact $\mathbb{Q}/\mathbb{Z}$ Certificate Verifier

The verifier `EXP-RETPROB-CERT-1.0` (`src/expansion_certifier.py`, `src/expansion_exact.py`) uses pure Python `int` and `fractions.Fraction` exclusively (source-audited for zero `float(`, `math.`, or `numpy`). Every payload enforces strict JSON typing (rejecting floats, bool-as-int, int-as-bool), canonical reduced rationals (forbidding `"-0"` and `"/1"`), nesting-depth cap 32, full ASCII control-character filtering, 2048-bit integer caps, and exact container/key validation across four modes:

1. **Mode 1 (Sec. 2.1–2.2 & Lem. 5.1 Good-Start Return Audit):** Builds $P = I - L_G/(2d)$ over $\mathbb{Q}$, computes P^t, P^{2t}, P^{4t}, checks Eq. (1), (2), (7), Dirichlet energy $\le A_s - B_s$, exact D_s, J_s, H_s, total variation $d_{\mathrm{TV}}(p_s, u_{H_s})$, $|\partial H_s|$, and $\mathrm{tr}(K_s^{2t}) \le \sum_{v \in H_s} A_v$.
2. **Mode 2 (Lem. 5.3 Induced-Expander Trimming):** Exhaustively checks $|E(S, V \setminus S)| \ge h|S|$ for every $S \subseteq H$ ($|S| \le |H|/2$), computes exact maximizer R^* of $F(R)$, checks $|R^*| \le \frac{4\tau}{3h}|H|$, and certifies exact induced expansion $h(G[H \setminus R^*]) \ge h/4$.
3. **Mode 3 (Lem. 7.3 Balanced-Cut Farness):** Computes exact global expansion $h_0 = h(G) > 0$, cut ratio $h_1 < c^*$, and exact rational edit/distance floors M_{lb} and $\varepsilon_{\mathrm{lb}} = M_{\mathrm{lb}}/(2dn)$.
4. **Mode 4 (Lem. 7.4–7.5 Non-Backtracking & Three-Squares):** Audits exact d-regular matrices $Q_k(A)$ (row sums $d(d-1)^{k-1}$, traces, return probabilities $R_G(k)$, walk-regularity) and Gauss's $r_3(4N) = r_3(N)$ paired with equal-F_2, unequal-F_3 collision-ambiguity distributions.

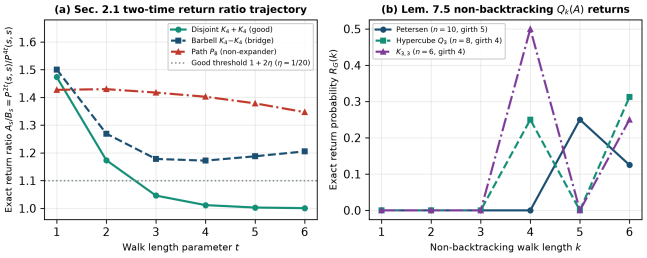

Figure 1: (a) Exact lazy-walk return ratio A_s/B_s across walk lengths $t \in \{1, \ldots, 6\}$ on disjoint $K_4 + K_4$, 1-bridge barbell $K_4 - K_4$, and path P_8. (b) Exact non-backtracking return probabilities $R_G(k)$ from $Q_k(A)$ on Petersen (girth 5), Q_3, and $K_{3,3}$ (girth 4).

4 Empirical Atlas and Red-Team Audit

Running `run_all_experiments.py` generates 16 canonical certificates (4 per mode) and four exact datasets, audited end-to-end by `verify_all.py` with SHA-256 verification across all 39 tracked artifacts.

Return Ratios & Non-Backtracking Polynomials (Fig. 1). Mode 1 (1.700–2.195 ms) certifies good starts on disjoint $K_4 + K_4$ ($A_s/B_s = 4288/4099 < 1.10$, $|H_s| = 4, |\partial H_s| = 0$) and asymmetric $K_5 + K_3$ ($|H_s| = 5, |D_s| = 3, J_s = 0$), while rejecting both the 1-bridge barbell K_4–K_4 (whose ratio rebounds from 1.172 at $t = 4$ to 1.205 at $t = 6$) and path P_8 ($A_s/B_s > 1.34$). Mode 4 (0.180–0.501 ms) certifies $\mathrm{tr}(Q_k(A)) = 0$ below girth on Petersen ($k \le 4$, $\mathrm{tr}(Q_5) = 120$, $R_5 = 1/4$) and $K_{3,3}$ ($\mathrm{tr}(Q_4) = 72$, $R_4 = 1/2$), both walk-regular.

Trimming, Farness, & Three-Squares (Fig. 2). Mode 2 (0.280–0.620 ms) certifies exact Lemma 5.3 trimming across clean K_5 ($R^* = \emptyset$), dangling boundary vertices on K_5 and K_6 ($|R^*| = 1$, $h(G[U]) = 3$), and a 2-vertex appendage on K_4 ($|R^*| = 2$, $h(G[U]) = 2$). Mode 3 (0.433–1.979 ms) certifies exact farness floors $M_{\mathrm{lb}} \in \{3/5, 22/19, 4/3, 4/5\}$ ($\varepsilon_{\mathrm{lb}} \in \{3/320, 11/950, 1/36, 1/90\}$). Mode 4 also certifies $r_3(56) = r_3(14) = 48$ and $r_3(108) = r_3(27) = 32$ alongside exact equal-F_2 ambiguity pairs ($F_2 = 1/2, \Delta F_3 = 1/18$; $F_2 = 1/3, \Delta F_3 = 1/36$).

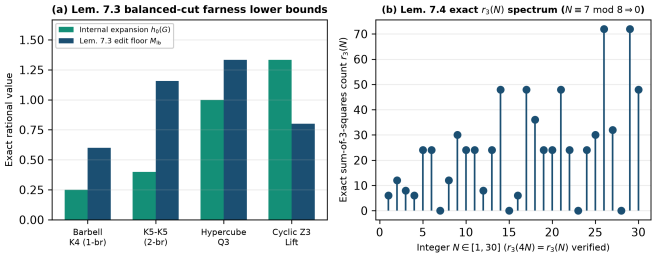

Figure 2: (a) Exact Lemma 7.3 balanced-cut farness edit floors M_{lb} versus internal expansion $h_0(G)$ across four connected graphs. (b) Exact three-squares representation count $r_3(N) = r_3(4N)$ verified for all $N \in [1, 30]$ ($r_3(N) = 0$ at $N \in \{7, 15, 23, 28\}$).

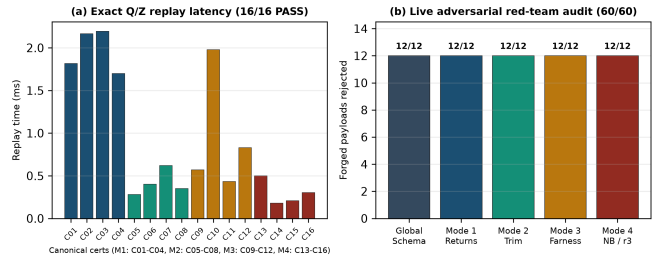

Figure 3: (a) Exact $\mathbb{Q}/\mathbb{Z}$ replay latency (0.180–2.195 ms) across all 16 canonical `EXP-RETPROB-CERT-1.0` certificates. (b) Live adversarial red-team audit rejecting 60/60 forged payloads across all five categories.

Replay Latency and 60/60 Red-Team Audit (Fig. 3). All 16 canonical certificates replay deterministically in 0.180–2.195 ms (total 14.53 ms). The live adversarial suite (`tests/test_suite.py`) tests 60 targeted forgeries—12 global schema/typing/canonical-rational/depth/control-char attacks plus 12 domain attacks per mode (self-loops, degree overflows, corrupted A_s/B_s, unsorted sets, invalid h-bounds, disconnected cuts, $c^* \leq h_1$, non-regular Q_k, mismatched $r_3(4N)$, and identical F_3 pairs)—achieving **60/60 (100%)** rejection.

5 Conclusion

We delivered a standalone, zero-floating-point verification framework for Melissaris's return-probability expansion testing theory [9], combining 18 zero-`sorry`, zero-`Classical.choice` Lean 4.35.0 theorems with the exact $\mathbb{Q}/\mathbb{Z}$ verifier `EXP-RETPROB-CERT-1.0`, 16 canonical certificates, and 60/60 live adversarial forgery rejections.

References

[1] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(4):585–603, 2007.

[2] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021.

[3] Irit Dinur and Oded Goldreich. Expansion without connectivity: A property testing perspective. Technical Report TR25-129, Electronic Colloquium on Computational Complexity (ECCC), 2025.

[4] Michael Filaseta, Jonah Klein, and Cihan Sabuncu. On extreme values of $r_3(n)$ in arithmetic progressions. *The Ramanujan Journal*, 69(2):Article 42, 2026.

[5] Chris D. Godsil and Brendan D. McKay. Feasibility conditions for the existence of walk-regular graphs. *Linear Algebra and its Applications*, 30:51–61, 1980.

[6] Oded Goldreich and Dana Ron. Property testing in bounded degree graphs. *Algorithmica*, 32(2):302–343, 2002.

[7] Satyen Kale and C. Seshadhri. An expansion tester for bounded degree graphs. *SIAM Journal on Computing*, 40(3):709–720, 2011.

[8] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[9] Nikolas Melissaris. Testing expansion without connectivity via return probabilities. arXiv:2611.32677 [cs.DS], November 2026.