DIFFERENTIALLY PRIVATE CYCLE COUNTING: LOCAL LOWER BOUNDS AND SHUFFLE ESTIMATORS

Anonymous authors
Paper under double-blind review

ABSTRACT

The shuffle model of differential privacy can improve aggregate statistics, but its benefits for graph pattern counting remain unclear. We identify an algebraic parity split for cycles. Let A be the adjacency matrix of an n-vertex graph and $W = A^2$ its matrix of length-two walk counts. Even closed-walk counts satisfy $\mathrm{tr}(A^{2m}) = \mathrm{tr}(W^m)$, whereas odd counts satisfy $\mathrm{tr}(A^{2m+1}) = \mathrm{tr}(W^m A)$ and retain one factor of A. This distinction motivates our estimators, but does not by itself characterize shuffle amplification. For every fixed odd k and $n^{-1/2} \le \varepsilon \le \varepsilon_0$, with $\varepsilon_0 > 0$ a small constant (and $\varepsilon \ge \sqrt{\log n/n}$ if $k = 3$), we give an unbiased shuffle estimator of $\mathrm{tr}(A^k)$ with dense-graph root mean squared error (RMSE) $\Theta_k(n^{k-1}/\varepsilon)$. Correcting walks that repeat vertices gives a simple-cycle estimator with $O_k(n^{k-1}/\varepsilon)$ RMSE, and the same Θ_k rate when $\varepsilon\sqrt{n} \to \infty$; for C_5, the rate is $\Theta(n^4/\varepsilon)$ throughout this range. We also prove an $\Omega(n^{k-1}/\varepsilon)$ lower bound for every fixed $k \ge 3$ in the non-interactive per-edge local model, where each edge has its own pure ε-DP channel. For C_5, on the dense random graphs that witness this bound and for $\varepsilon \ge n^{-1/4}$, suppose the analyzer also receives W with Gaussian noise. Noise of order $\widetilde{\Theta}(n/\varepsilon)$ per entry is then a threshold: above it the bound still holds, and well below it spectral deflation, which removes the top eigencomponent, beats the bound. With a provably private multi-message release of W, deflation reaches RMSE $\widetilde{O}(n^{3.5}/\varepsilon)$ on these graphs. This (ε, δ) shuffle protocol improves on the per-edge lower bound by $\widetilde{\Theta}(\sqrt{n})$ on that family, which is not a worst-case separation from the full local model.

1 INTRODUCTION

Counting small subgraphs such as triangles, wedges, and short cycles is a basic task in network analysis. When edges represent sensitive relationships, accuracy under differential privacy (DP) depends on the trust model. In the *central* model a trusted curator adds noise (Dwork et al., 2006); in the *local* model each user randomizes their adjacency row (Kasiviswanathan et al., 2011; Imola et al., 2021). The *shuffle* model adds an anonymizing shuffler that permutes messages before the analyst sees them. Amplification by shuffling can strengthen privacy (Erlingsson et al., 2019; Feldman et al., 2021). Imola et al. (2022) show that shuffling improves triangle and 4-cycle estimation over the local baseline. Fang & Yi (2025) sharpen these results with a negative-binomial mechanism for subgraphs on up to four nodes, yet call a provable shuffle-vs-local separation "elusive," citing missing sparse-graph local lower bounds.

Suppakitpaisarn et al. (2025) already give non-interactive local protocols for any k-node graphlet at rate $O(n^{k-1})$, with a matching lower bound for cliques and only $\Omega(n^{k-1.5})$ for cycles. Can shuffling improve the dependence on n for cycles, and in which protocol classes? On dense graphs (all degrees $\Theta(n)$) and for $n^{-1/2} \le \varepsilon \le \varepsilon_0$, with ε_0 a small constant (and $\varepsilon \ge \sqrt{\log n/n}$ if $k = 3$), our odd-cycle shuffle estimators match the local rate in n. We prove a matching worst-case lower bound for non-interactive per-edge local protocols, which randomize each edge independently under pure ε-DP. But on the random graphs behind our lower bound, a multi-message shuffle protocol improves on it by $\widetilde{\Theta}(\sqrt{n})$ for 5-cycles when $\varepsilon \ge n^{-1/4}$.

Our starting point is the *wedge matrix* $W = A^2 = \sum_v a_v a_v^\top$ of common-neighbor counts, where a_v is row v of the adjacency matrix A. The trace $\mathrm{tr}(A^k)$ counts closed walks of length k, which

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

may revisit vertices, whereas a simple k-cycle (C_k) visits k distinct vertices. Even traces $\mathrm{tr}(A^{2m})$ are polynomials in W, while odd ones have the form $\mathrm{tr}(W^m A)$ and are not functions of W. For instance, C_6 and two disjoint triangles share W up to relabeling, but only the latter has triangles. We call the unpaired factor A in $\mathrm{tr}(W^m A)$ the *seam* and its entries A_{qr} *bare edges*. Only q and r hold A_{qr}, but neither can compute its weight $(W^m)_{rq}$, a global count. Whether parity also splits shuffle accuracy is open. For even cycles, a shuffle gain is known only for C_4.

Contributions.

- **Parity split.** We show that even traces factor through $W = A^2$ and odd ones do not (Proposition 1). Theorem 1 records the C_4 shuffle rate $\widetilde{O}(n^{2.75})$ of Fang & Yi (2025).
- **Odd cycles.** For odd k we give shuffle estimators with a *fixed privacy ledger* of $(k+1)/2$ independent releases. On dense graphs (all degrees $\Theta(n)$), for $n^{-1/2} \le \varepsilon \le \varepsilon_0$ ($\varepsilon \ge \sqrt{\log n/n}$ if $k = 3$), their trace RMSE $\Theta_k(n^{k-1}/\varepsilon)$ matches the local rate in n. The simple-cycle RMSE is $O_k(n^{k-1}/\varepsilon)$, and $\Theta_k(n^{k-1}/\varepsilon)$ when $\varepsilon\sqrt{n} \to \infty$, and for C_5 throughout.
- **Per-edge lower bound.** We prove that if each edge bit passes non-interactively through its own pure ε-DP randomizer, every C_k-count estimator has worst-case dense-graph RMSE $\Omega(n^{k-1}/\varepsilon)$ for all $k \ge 3$. The model excludes general row-local and shuffle protocols.
- **Wedge-noise threshold.** On the hard family $\mu_\pm = G(n, 1/2 \pm \Theta(1/(\varepsilon n)))$, the per-edge C_5 bound holds even given W plus Gaussian noise above a threshold $\widetilde{\Theta}(n/\varepsilon)$ per entry. A *certified* multi-message wedge release (Fang & Yi, 2025), private on every graph, has noise $\widetilde{O}(n^{1/2}/\varepsilon)$, far below it. With it, for $\varepsilon \ge n^{-1/4}$, we reach C_5 RMSE $\widetilde{O}(n^{3.5}/\varepsilon)$ on $\mu_\pm$ by *spectral deflation*, which reads the top-eigenvalue term of $\mathrm{tr}(A^5)$ from the noisy W.

Results and scope. Our shuffle protocols are multi-message and pair-indexed. Each user sends one report per vertex pair in each wedge release, so one edge changes $\Theta(n)$ scalar reports and the protocols lie outside $\mathcal{C}$ (Definition 3). They are private on every graph and leave open the seam barrier, our conjecture that $\mathcal{C}$ cannot beat the per-edge rate by a $\mathrm{poly}(n)$ factor for odd k. Simulations with the specified noise laws check only the error exponents (§6).

Result	Model	RMSE on	Status
Odd-cycle estimators	shuffle, pair-indexed	every dense graph	proved, Thms. 2–4
Per-edge lower bound	per-edge local, non-interactive, pure ε	worst dense graph	proved, Thm. 5
Wedge-noise threshold	per-edge local plus wedge oracle	$\mu_\pm$	proved, Thms. 6–7
C_5 gain $\widetilde{\Theta}(\sqrt{n})$ over the per-edge bound	shuffle, pair-indexed	$\mu_\pm$ only	proved, Cor. 1
Seam barrier	class $\mathcal{C}$, $O(1)$ scalar reports per edge	worst dense graph	conjectured, Conj. 1–2

2 PRELIMINARIES

Differential privacy. We study *edge privacy*: two graphs are neighboring if one is obtained from the other by adding or removing a single edge.

Definition 1 (Edge differential privacy)**.** *A randomized mechanism $\mathcal{M}$ satisfies (ε, δ)-edge-DP if for every pair of neighboring graphs G, G' (differing in one edge) and every measurable $S \subseteq \mathrm{range}(\mathcal{M})$, $\Pr[\mathcal{M}(G) \in S] \le e^\varepsilon \Pr[\mathcal{M}(G') \in S] + \delta$. When $\delta = 0$ we write ε-DP.*

Edge-DP is closed under post-processing: any data-independent map applied to $\mathcal{M}(G)$ remains (ε, δ)-edge-DP. Composing L mechanisms that are each ε'-DP gives $L\varepsilon'$-DP or, for small ε', $\big(O(\varepsilon'\sqrt{L\log(1/\delta)}), \delta\big)$-DP by advanced composition (Dwork & Roth, 2014). A real-valued query f with ℓ_1-sensitivity $\Delta f = \max_{G \sim G'} |f(G) - f(G')|$ is released ε-DP by the Laplace mechanism (Dwork et al., 2006), adding noise drawn from $\mathrm{Lap}(\Delta f/\varepsilon)$.

Trust models. Besides the central, local (Kasiviswanathan et al., 2011) and shuffle (Cheu et al., 2019) models of §1, in §5 we use the narrower *non-interactive per-edge ε-local model*. There the analyzer sees only a transcript $T = (Y_e)_{e \in \binom{[n]}{2}} \sim \bigotimes_e Q_e$, each Q_e being ε-DP in the single edge bit A_e. Estimators are arbitrary, and every derived quantity, including $W = A^2$, is a function of T.

Definition 2 (Shuffle amplification (informal)). *Let each of n users apply an $\varepsilon_{\mathrm{loc}}$-local randomizer, and let the analyst observe the uniformly shuffled multiset of their messages. Then the composed mechanism is (ε, δ)-edge-DP with $\varepsilon = O\big(\varepsilon_{\mathrm{loc}}\sqrt{\log(1/\delta)/n}\big)$ for $\varepsilon_{\mathrm{loc}} = O(1)$, a $\Theta(\sqrt{n})$ gain over the local guarantee (Erlingsson et al., 2019; Feldman et al., 2021), provided each private bit influences only $O(1)$ of the shuffled messages.*

The cited theorems treat one record per user; the proviso is our edge-level adaptation, and the class $\mathcal{C}$ (Definition 3, §3) formalizes it. Our proved shuffle protocols lie outside $\mathcal{C}$, and their privacy proofs use direct noise summation and composition but no amplification-by-shuffling theorem (Theorem 2, Lemma 10). The per-edge randomizers Q_e are more restricted than those of $\mathcal{C}$ and exclude the wedge releases of §4, so Theorem 5 does not lower-bound $\mathcal{C}$ (Remark 3 discusses what would).

Primitives. Our base randomizer, randomized response (Warner, 1965), reports each private bit truthfully with probability $e^{\varepsilon_{\mathrm{loc}}}/(1+e^{\varepsilon_{\mathrm{loc}}})$ and flips it otherwise. For wedge/degree counts we use geometric (discrete-Laplace) noise $\Pr[z] \propto e^{-|z|/b}$ on $\mathbb{Z}$, optimal for a single count query (Ghosh et al., 2009). The negative-binomial mechanism of Fang & Yi (2025) splits it into per-user shares.

Graph notation. $G = (V, E)$ is a simple graph on $n = |V|$ vertices with symmetric adjacency matrix $A \in \{0,1\}^{n\times n}$, $A_{uu} = 0$. Each vertex v is a user holding row $a_v \in \{0,1\}^n$. The degree of v is $d_v = \sum_u A_{vu}$ and $d_{\max} = \max_v d_v$. Our Θ rates need the *dense regime*, in which the minimum degree is $\Theta(n)$. The trace $\mathrm{tr}(A^k)$ counts closed k-walks and $N_k(G)$ counts simple (non-induced) k-cycles. Each k-cycle is traversed by $2k$ closed walks, so $\mathrm{tr}(A^k) = 2kN_k + D_k$, where D_k counts closed walks that revisit a vertex and $D_3 = 0$ (Lemma 6). We call estimators of $\mathrm{tr}(A^k)$ and N_k *trace* and *simple-count* estimators. Induced counts, which also forbid chords, can differ from N_k at leading order on dense graphs and are out of scope. $\mathrm{RMSE} = \mathbb{E}[(\widehat{N} - N)^2]^{1/2}$ is taken over the mechanism's randomness on a fixed graph. Rates 'on dense graphs' hold on every dense graph, Theorem 5 bounds the worst one, and Corollary 1 averages over a random-graph family.

3 The Algebraic Parity Split

3.1 Setup and the amplification target

Let $W := A^2$ be the *wedge matrix*, $W_{ij} = \#\{u : iu, uj \in E\}$ the number of 2-paths between i and j. The structural fact we exploit is that W decomposes over users,

$$W = \sum_{u=1}^{n} a_u a_u^\top, \tag{1}$$

each summand $a_u a_u^\top$ computable by user u from its own row alone. An edge bit A_{iu} occurs in exactly two rows, so it affects only two summands of (1) but $\Theta(n)$ entries of each. Definition 3 therefore bounds how many scalar reports one edge can change.

Definition 3 (One-round sum-based shuffle protocol). *A protocol in the class $\mathcal{C}$ consists of: (i) a local randomizer per user that reads only that user's own row and emits messages made of scalar reports (one index and one scalar value each; a vector message counts as one report per coordinate), with each edge bit changing $O(1)$ of these reports; (ii) a trusted shuffler that releases the messages as an unordered multiset; and (iii) an analyzer that applies a fixed $O(1)$-degree polynomial map g to each shuffled message, forms the symmetric sum $\sum_y g(y)$ over the messages y, and returns a fixed $O(1)$-degree polynomial of that sum (for a row message, g may form its outer product). The class excludes multi-round and interactive protocols and cross-user-correlated randomizers. (In: the blind per-edge randomized-response release of §4, and the one-message-per-user (row-report) variant, whose coordinates move $O(1)$ per edge. Out: the pair-indexed release of Lemma 10, which moves $\Theta(n)$.)*

Because differential privacy is closed under post-processing, every polynomial in A computed from a single noisy release of A is trivially (ε, δ)-edge-DP; privacy alone cannot separate patterns. The operative question is whether the $\sqrt{n}$ accuracy gain of shuffling survives the polynomial readout, i.e., whether each edge keeps a bounded influence on the final estimate.

Definition 4 (Shuffle-amplifiability). *A cycle count ($\mathrm{tr}(A^k)$ or N_k) is $\mathcal{C}$-amplifiable if some estimator in $\mathcal{C}$ whose shuffled multiset is (ε, δ)-edge-DP, with $\delta \le n^{-2}$ and ε in the range of Theorem 5, has*

worst-case dense-graph RMSE smaller by a $\mathrm{poly}(n)$ *factor than the per-edge local rate* $\Theta(n^{k-1}/\varepsilon)$ *of Theorem 5.*

3.2 EVEN TRACES

Theorem 1 (Even-trace factorization; C_4 instantiation)**.** *For every* $m \geq 2$*, the closed-walk count* $\mathrm{tr}(A^{2m}) = \mathrm{tr}(W^m)$ *is a fixed degree-*m *polynomial in the per-user-decomposable aggregate* W*, so the count has no seam, the structural obstruction of the odd case. For* $m = 2$ *the utility is proven: the closed-4-walk trace* $\mathrm{tr}(W^2) = \langle W, W\rangle_F$ *(the 4-cycle count up to the corrections of Lemma 6) is releasable in the shuffle model at RMSE* $\widetilde{O}(n^{2.75})$ *(Fang & Yi, 2025), below the best known local-DP upper bound* $O(n^3)$ *(Suppakitpaisarn et al., 2025).*

Proof sketch. Since $A^{2m} = W^m$, the count has no seam, and privacy of $\mathrm{tr}(\widehat{W}^m)$ is post-processing. The utility claim is the $m = 2$ instance, the rate of Fang & Yi (2025), whose pair-indexed release lies outside $\mathcal{C}$ (Remark 1). For $m \geq 3$ we bound only the first-order noise term, so we make no utility claim there (full sketch in Appendix A.11). □

3.3 ODD TRACES AND THE SEAM

An odd power of A cannot be written using only even powers:

$$\mathrm{tr}(A^{2m+1}) = \mathrm{tr}(W^m A) = \sum_{q,r} (W^m)_{rq}\, A_{qr}, \tag{2}$$

leaving exactly one factor of A, the seam, whose entries A_{qr} are bare edges. For C_5 this is $\mathrm{tr}(A^5) = \mathrm{tr}(WAW) = \langle A, W^2\rangle$, which weights each bare edge A_{qr} by the global coefficient $(W^2)_{rq}$, the number of 4-walks from r to q.

Proposition 1 (Odd counts are not functions of W)**.** *The map* $A \mapsto A^2$ *is not injective on graphs: there exist* G, G'*, dense after a blow-up, with* $A(G)^2 = A(G')^2$ *up to relabeling but* $\mathrm{tr}(A(G)^{2m+1}) \neq \mathrm{tr}(A(G')^{2m+1})$ *for every* $m \geq 1$*. Hence* $\mathrm{tr}(A^{2m+1})$ *is not a polynomial in* W *for any* m*, and the bare edge in (2) is unavoidable.*

Proof sketch. $G = C_6$ and $G' = 2K_3$ share the square A^2 yet differ in every odd trace: 0 for the bipartite C_6 versus $2(2^{2m+1} - 2) > 0$ (60 at $m = 2$) for $2K_3$, whose spectrum is two copies of $(2, -1, -1)$. The two graphs have the same eigenvalue moduli but different signs. A blow-up $A_0 \otimes J_t$ (t copies per vertex, J_t all-ones) squares to $A_0^2 \otimes tJ_t$, so equal squares stay equal while the pair becomes dense. Full proof in Appendix A.2. □

The bare edge A_{qr} is owned by its endpoints, yet must be multiplied by the global coefficient $(W^m)_{rq}$, which neither endpoint can compute. Heuristically, shuffling hides a sender only among look-alike messages, but applying the weight $(W^m)_{rq}$ needs each report of A_{qr} to carry the edge label $\{q, r\}$. Only q or r can send such a report. We call this obstruction the *seam barrier*.

Conjecture 1 (Seam barrier)**.** *No protocol in* $\mathcal{C}$ *amplifies the seam:* $\mathrm{tr}(A^{2m+1})$ *is not* $\mathcal{C}$*-amplifiable (Definition 4). For* C_5*, more sharply: worst-case dense RMSE* $\Omega(n^4/\varepsilon)$ *for every* $\mathcal{C}$*-estimator private as in Definition 4.*

We have not proved Conjecture 1. Theorem 5 proves the corresponding bound only for non-interactive per-edge randomizers, which are more restricted than those of $\mathcal{C}$ (§2).

3.4 FROM CLOSED WALKS TO SIMPLE CYCLE COUNTS

The degenerate part $D_k = \mathrm{tr}(A^k) - 2kN_k$ is a fixed combination of counts of the smaller graphs that a closed k-walk traces when it revisits vertices (its *images*). Using Lemma 6, we estimate every image count from the debiased randomized-response release of §4. For fixed k and $\varepsilon \geq n^{-1/2}$, the correction RMSE is $O_k(n^{k-1}/\varepsilon)$, so conversion preserves the target upper rate without a new privacy charge. The correction is lower order when $\varepsilon\sqrt{n} \to \infty$, but not at $\varepsilon = \Theta(n^{-1/2})$ for odd $k \geq 7$: for the image $C_3 \vee C_{k-3}$ (a triangle and a $(k-3)$-cycle sharing one vertex), the error term

of degree k in the noise alone has order $n^{(k-1)/2}/\varepsilon^k$. For C_4 the corrections fold into the proven wedge release, and for C_5 a direct image enumeration gives a lower-order correction throughout $\varepsilon \geq n^{-1/2}$.

4 A SHUFFLE-MODEL 5-CYCLE PROTOCOL

The natural C_5 extension of prior wedge shuffling, the *pivot baseline*, anchors an edge and sums products of two wedge counts over pivots. Its per-pivot sub-mechanisms cannot be composed in parallel, so its ledger grows with n and caps how many pivots it can sample. For $n^{-1/2} \leq \varepsilon = O(1)$, its RMSE on dense graphs is $\widetilde{O}(n^{4.5}/\varepsilon)$, a $\sqrt{n}$ factor above ours (Theorem 3), and randomized response on its subsampled anchor edges dominates it (Lemma 8, Corollary 2, Appendix A.3). We instead factor $\mathrm{tr}(A^5) = \mathrm{tr}(WAW)$, whose bare edge is unavoidable (Proposition 1, §3), and give each factor its own independent release.

4.1 CONSTRUCTION

The estimator is the trilinear form

$$\widehat{T} = \mathrm{tr}(\widehat{W}_1 \widehat{A} \widehat{W}_2), \tag{3}$$

built from three independent releases under a 3-way split $\varepsilon = \varepsilon_1 + \varepsilon_2 + \varepsilon_3$ of the edge-DP budget (Algorithm 1). Two of them, $\widehat{W}_1$ and $\widehat{W}_2$, release the wedge matrix $W = A^2$ and satisfy the wedge-release interface: each is entrywise unbiased, including the diagonal $W_{qq} = d_q$, has a public noise law, and is $(\varepsilon_i, \delta_i)$-edge-DP on all graphs. We certify both with the pair-indexed negative-binomial mechanism of Fang & Yi (2025) off the diagonal plus an independent degree release on it, composed inside each budget (Lemma 10). The row-report variant in $\mathcal{C}$ is conditional (Remark 1). The third, $\widehat{A}$, is a *blind* randomized-response release of the adjacency (Warner, 1965), covering every vertex pair with no sampling. Each pair's bit is reported once, by its lower-indexed endpoint $\min(u, w)$, then mirrored and debiased so $\mathbb{E}[\widehat{A}] = A$ with $\{\widehat{A}_{uw}\}_{u<w}$ independent.

Algorithm 1 Shuffle-model C_5 estimator

Require: rows $\{a_u\}_{u\in[n]}$; budget split $\varepsilon = \varepsilon_1 + \varepsilon_2 + \varepsilon_3$

1: **Wedge release 1:** each user sends, for every pair $\{q, r\}$, its 2-path indicator through the negative-binomial randomizer; the analyzer sums the shuffled messages into $(\widehat{W}_1)_{qr}$ and fills the diagonal from an independent geometric degree release (Lemma 10(e)), so $\mathbb{E}[\widehat{W}_1] = W$ entrywise. {budget ε_1}

2: **Wedge release 2:** independently repeat Step 1 with fresh randomness to obtain $\widehat{W}_2$, $\mathbb{E}[\widehat{W}_2] = W$. {budget ε_2}

3: **Blind edge release:** each user applies randomized response (Warner, 1965) to the edges it owns (min-endpoint convention); the analyzer mirrors and debiases to $\widehat{A}$, $\mathbb{E}[\widehat{A}] = A$. {budget ε_3}

4: **return** $\widehat{T} = \mathrm{tr}(\widehat{W}_1 \widehat{A} \widehat{W}_2)$

4.2 UNBIASEDNESS

Lemma 1 (Unbiasedness)**.** *Let $\widehat{W}_1 = W + N_1$, $\widehat{W}_2 = W + N_2$, $\widehat{A} = A + M$ be the releases of Algorithm 1, mutually independent with $\mathbb{E}N_1 = \mathbb{E}N_2 = \mathbb{E}M = 0$ entrywise. Then $\mathbb{E}[\widehat{T}] = \mathrm{tr}(WAW) = \mathrm{tr}(A^5)$.*

Proof. Each term of $\widehat{T} = \sum_{u,v,w} (\widehat{W}_1)_{uv} (\widehat{A})_{vw} (\widehat{W}_2)_{wu}$ factors over the three independent releases, so every noise cross-term contains an independent mean-zero factor and vanishes. Sharing a noisy wedge ($N_1 = N_2$) or a degree diagonal would break independence. At $k = 5$ sharing keeps the estimator unbiased, as $\widehat{A}$ has zero diagonal, but its cross terms become correlated (Lemma 7). For odd $k \geq 7$, sharing adjacent wedge factors biases it through middle factors with nonzero diagonal. □

4.3 PRIVACY

Theorem 2 (Privacy of the certified instantiation)**.** *Algorithm 1, with the wedge releases instantiated per Lemma 10(e) (pair-indexed off the diagonal, degree release on it) at budgets $(\varepsilon_i, \delta_i)$, each $\varepsilon_i < 2$, $\delta_i \in (0, \frac{1}{2})$, is (ε, δ)-edge-DP on all graphs with $\varepsilon = \varepsilon_1 + \varepsilon_2 + \varepsilon_3$ and $\delta = \delta_1 + \delta_2$.*

Proof sketch. The ledger has a fixed length $L = 3$: its three privacy-consuming mechanisms $\widehat{W}_1, \widehat{W}_2, \widehat{A}$ are combined by basic composition (Dwork & Roth, 2014), and the contraction (3) is post-processing. The wedge releases are $(\varepsilon_i, \delta_i)$-edge-DP on all graphs by Lemma 10. Per-invocation privacy comes from direct noise summation. The $\leq 2(n-1)$ pair invocations that one edge influences are composed inside the cited mechanism by advanced composition, and basic composition adds the two degree reports one edge touches. No amplification-by-shuffling theorem is used. The blind release is pure ε_3-DP, with one randomized bit per owned edge. □

Remark 1 (Certified vs. conditional)**.** The certified instantiation is pair-indexed: one edge influences $\Theta(n)$ scalar pair reports, so it exits $\mathcal{C}$ (Remark 3). The in-$\mathcal{C}$ row-report variant is conditional on the unproven within-row-composition ledger of Remark 3. All rate claims hold for the certified instantiation, and class membership matters only for Conjecture 1 (§3) and Conjecture 2 (§5).

4.4 ACCURACY

Theorem 3 (Accuracy)**.** *On dense graphs (all degrees $\Theta(n)$), for $n^{-1/2} \leq \varepsilon = O(1)$, and under a balanced split $\varepsilon_i = \Theta(\varepsilon)$, $\delta_i = \Theta(\delta)$ with $\delta = n^{-\Theta(1)}$, the estimator $\widehat{T}$ has $\mathrm{RMSE} = \Theta(n^4/\varepsilon)$. On families with $\mathrm{tr}(A^5) = \Theta(n^5)$ the relative error is $\Theta\big(1/(\varepsilon n)\big)$; density alone does not guarantee this—dense bipartite graphs have all odd traces zero.*

Almost all of this error comes from the bare-edge channel, the error term linear in the noise of $\widehat{A}$.

Proof sketch. By Lemma 1 the error is $X_{N_1} + X_{N_2} + X_M$ plus terms of degree at least two in the noise. The bare-edge channel $X_M = \sum_{v,w} M_{vw}(A^4)_{wv}$ has variance $\Theta(v_A \|A^4\|_F^2) = \Theta(n^8/\varepsilon^2)$ on dense graphs, where $v_A = \Theta(1/\varepsilon^2)$ is the per-entry variance of blind randomized response. Distinct noise monomials are uncorrelated, so this lower-bounds the total. The wedge channels $X_{N_i} = \mathrm{tr}(N_i A^3)$ and the higher-degree terms total $\widetilde{O}(n^7/\varepsilon^2)$ for $\varepsilon \geq n^{-1/2}$ (Lemmas 7 and 10), a factor $\widetilde{O}(n)$ below (full sketch in Appendix A.11). □

Remark 2. The simple count $\widehat{N}_5 = (\widehat{T} - \widehat{D}_5)/10$ has the same dense-regime $\Theta(n^4/\varepsilon)$ rate and ledger: Lemma 6 gives correction RMSE $O(n^3/\varepsilon + n^{2.5}/\varepsilon^3 + n^2/\varepsilon^4) = o(n^4/\varepsilon)$ throughout $\varepsilon \geq n^{-1/2}$.

4.5 EXTENSION TO ALL ODD CYCLES

Nothing in the construction is specific to $k = 5$. For odd $k = 2m + 1 \geq 3$ we factor $\mathrm{tr}(A^k) = \mathrm{tr}(W^m A)$ and give each factor its own independent release:

$$\widehat{T}_k := \mathrm{tr}\big(\widehat{W}_1 \widehat{W}_2 \cdots \widehat{W}_m \widehat{A}\big),$$

with m independent certified wedge releases (Lemma 10, Remark 1) and one blind randomized-response release under an $(m+1)$-way budget split.

Theorem 4 (All odd cycles; trace and simple counts)**.** *Fix odd $k = 2m + 1 \geq 3$ and use a balanced constant-factor split $\varepsilon_i = \Theta_k(\varepsilon)$, $\delta_i = \Theta_k(\delta)$, with $\varepsilon = \sum_{i=1}^{m+1} \varepsilon_i$, $\delta = \sum_{i=1}^{m} \delta_i$ and $\delta = n^{-\Theta(1)}$. The protocol is (ε, δ)-edge-DP, has fixed ledger length $L = m + 1$, and is unbiased for $\mathrm{tr}(A^k)$. On dense graphs, for $n^{-1/2} \leq \varepsilon \leq \varepsilon_0$ ($\varepsilon \geq \sqrt{\log n/n}$ if $k = 3$), where $\varepsilon_0 = \varepsilon_0(k) > 0$ is a sufficiently small constant, its trace RMSE is $\Theta_k(n^{k-1}/\varepsilon)$. Let $\widehat{D}_k$ be the correction plug-in of Lemma 6, computed from the already released $\widehat{A}$, and set $\widehat{N}_k = (\widehat{T}_k - \widehat{D}_k)/(2k)$. Then $\widehat{N}_k$ is unbiased and has $O_k(n^{k-1}/\varepsilon)$ RMSE with no extra ledger entry. If $\varepsilon\sqrt{n} \to \infty$, this particular simple-count estimator has $\Theta_k(n^{k-1}/\varepsilon)$ RMSE; the same Θ conclusion holds for $k = 5$ throughout $\varepsilon \geq n^{-1/2}$. For odd $k \geq 7$ at $\varepsilon = \Theta(n^{-1/2})$, only the displayed O_k simple-count guarantee is claimed.*

Appendix A.11 sketches the proof; Appendix A.2 gives the full computation.

5 Per-Edge Local Lower Bound and Wedge-Noise Transition

The $\Theta(n^{k-1}/\varepsilon)$ rate of §4 matches a lower bound that holds for every estimator in the non-interactive per-edge local model and every cycle length $k \geq 3$. We bound N_k directly, because the closed-walk proxy $\mathrm{tr}(A^k) = 2kN_k + D_k$ carries a degenerate part D_k as large as $\Theta(n^{k-1})$ for $k \geq 4$ (Lemma 6). This is the order of the minimax risk, so transfer through the proxy is lossy.

Theorem 5 (Per-edge local lower bound, every cycle length)**.** *Fix $k \geq 3$. In the dense regime (all degrees $\Theta(n)$) there is a constant $\varepsilon_0(k) \in (0, 1]$ such that for every $\varepsilon \in (0, \varepsilon_0(k)]$ with $\varepsilon \gg 1/n$, every estimator $\widehat{N}(T)$ in the non-interactive per-edge ε-local model obeys*

$$\sup_G \mathbb{E}\big[(\widehat{N}(T) - N_k(G))^2\big]^{1/2} \geq c_k \frac{n^{k-1}}{\varepsilon}, \qquad c_k > 0.$$

The rate is tight for every k: for $\varepsilon \geq n^{-(k-2)/(2k-2)}$ (any fixed ε in particular) the debiased randomized-response plug-in attains $O_k(n^{k-1}/\varepsilon)$ RMSE.

This bound exceeds the $\Omega(n^{k-1.5})$ cycle bound of Suppakitpaisarn et al. (2025) by $\sqrt{n}$ because our model is narrower and randomizes every edge separately.

Proof sketch. Take biased product priors $\mu_\pm = \mathrm{Bern}(\frac{1}{2} \pm \gamma^*)^{\otimes \binom{n}{2}}$ at bias $\gamma^* = \Theta(1/(\varepsilon n))$, where no ε-local transcript reliably tells them apart (Lemma 4). A Margulis–Russo computation gives a mean separation $\Theta(n^{k-1}/\varepsilon)$ of the realized count, because all $\Theta(n^2)$ per-edge gradients $\Delta_e N_k = \#\{k\text{-cycles through } e\} = \Theta(n^{k-2})$ are strictly positive, hence sign-coherent. Their mean vector $c := (\mathbb{E}\Delta_e N_k)_e$ has $\|c\|_2 = \Theta(n^{k-1})$. Efron–Stein caps the count's standard deviation at $\Theta(n^{k-1})$. An estimator-agnostic test reduction (Le Cam) then converts the transcript's bounded total variation into RMSE $\Omega(n^{k-1}/\varepsilon)$. Full proof in Appendix A.1, which also records why a two-point argument cannot reach the rate (Remark 5). □

Theorem 5's hard family relies on the analyzer not seeing W exactly (Appendix A.4).

Theorem 6 (Exact wedges defeat the hard family)**.** *On the dense product family $\mu_\pm$ of Theorem 5, assume $\varepsilon \leq \varepsilon_0$ and $\varepsilon n \to \infty$. If the analyzer is additionally given the exact, noiseless wedge matrix $W = A^2$, then a non-interactive per-edge ε-local estimator attains simple-C_5 RMSE $O(n^{2.5}/\varepsilon)$.*

Thus the $\Omega(n^4/\varepsilon)$ construction of Theorem 5 fails on its own family under a free exact-W oracle. We leave open whether another dense family witnesses n^4 hardness with free W.

Conjecture 2. *For odd k, the $\Omega(n^{k-1}/\varepsilon)$ bound extends to the private estimators (Definition 4) in the one-round sum-based shuffle class $\mathcal{C}$ of Definition 3, in which the wedge messages are shuffle-amplified rather than raw ε-local, each edge influencing $O(1)$ scalar reports.*

We now locate, on $\mu_\pm$, the wedge-noise level at which the per-edge hardness disappears.

Theorem 7 (Wedge-noise phase transition)**.** *In the setting of Theorem 5 with $k = 5$, give the analyzer the per-edge local transcript T and a graph-independent wedge oracle $\widehat{W} = W + E$. There are constants $0 < c < C$ such that: (i) if the upper-triangular entries of E are i.i.d. Gaussian with standard deviation $\sigma \geq Cn/\varepsilon$ and E is independent of T, every estimator from $(T, \widehat{W})$ incurs RMSE $\Omega(n^4/\varepsilon)$ on the dense family; (ii) suppose instead that the off-diagonal entries are independent centered sub-exponential variables of standard deviation σ, while the diagonal entries are independent centered sub-exponential variables of public standard deviations at most $\sigma_d \leq C_d\sigma$, and E is independent of T (assumption (N*), Appendix A.7). If $\sigma \leq cn/\varepsilon$ and $\varepsilon \geq n^{-1/4}$, a spectral-deflation estimator attains*

$$\widetilde{O}\left(n^3\sigma + n^{2.5}\sigma_d + \frac{n^{2.5}}{\varepsilon}\right).$$

For the Gaussian oracle $\sigma_d = \sigma$, so this reduces to $\widetilde{O}(n^3\sigma + n^{2.5}/\varepsilon)$. Hence the Gaussian threshold is $\sigma^* = \widetilde{\Theta}(n/\varepsilon)$ up to the upper-side polylogarithms, with Theorem 6 as the exact-W endpoint.

Together, (i) and (ii) say that on $\mu_\pm$, Gaussian wedge noise with $\sigma \geq Cn/\varepsilon$ keeps the per-edge bound, while for $\varepsilon \geq n^{-1/4}$ Gaussian wedge noise with σ polynomially below n/ε gives deflation error polynomially below n^4/ε.

Proof idea. For (i), by the KL chain rule the oracle's coherent mean shift adds $O(n^2/(\varepsilon^2\sigma^2))$ to the divergence. For (ii), $\|E\|_{\mathrm{op}} = \widetilde{O}(\sigma\sqrt{n})$ stays below the $\Theta(n^2)$ spectral gap of W; the wedge noise then costs $\widetilde{O}(n^3\sigma + n^{2.5}\sigma_d)$ and the deflated readouts of $\widehat{A}$ cost $\widetilde{O}(n^{2.5}/\varepsilon)$ (Appendix A.7). □

The spectral-deflation estimator. With high probability under $\mu_\pm$, A has a simple top eigenvalue $\lambda_1 = \Theta(n)$ with unit eigenvector v_1, and all other eigenvalues are $O(\sqrt{n})$ (Lemma 14). With $\mu_1 = \lambda_1^2$ and $\Pi_1 = v_1 v_1^\top$,

$$\mathrm{tr}(A^5) = \mu_1^{5/2} + \langle W^2 - \mu_1^2\Pi_1, A\rangle, \qquad \|W^2 - \mu_1^2\Pi_1\|_F = O(n^{2.5}) \ll \|W^2\|_F = \Theta(n^4).$$

The Perron term $\mu_1^{5/2}$ depends on W alone and is read from $\widehat{W}$. Only the deflated coefficient multiplies $\widehat{A} = A + M$, so the bare-edge channel costs $O(n^{2.5}/\varepsilon)$ rather than $\Theta(n^4/\varepsilon)$ (Theorem 3). With $(\hat{\mu}_1, \hat{v}_1)$ the top eigenpair of $\widehat{W}$ and $\hat{\mu}_1^c := \hat{\mu}_1 - (n-1)\sigma^2/\hat{\mu}_1$, the estimate of $\mathrm{tr}(A^5)$ is

$$\widehat{F} := (\hat{\mu}_1^c)^{5/2} - (n-1)\sigma^2(\hat{\mu}_1^c)^{1/2} + \big\langle \widehat{W}^2 - (\hat{\mu}_1^c)^2 \hat{v}_1\hat{v}_1^\top,\, \widehat{A}\big\rangle,$$

where $\hat{\mu}_1^c$ and the middle term remove the second-order biases of the top eigenvalue and eigenvector under wedge noise, both using public noise variances. The diagonal of $\widehat{W}^2$, debiased in step (S3) of Appendix A.7, never enters, as $\widehat{A}$ has zero diagonal. We read out D_5 the same way and clip $\widehat{N}_5 = (\widehat{F} - \widehat{D}_5)/10$ to $[0, n^5]$, as in steps (S1)–(S7) there.

Remark 3 (Where protocols sit)**.** The transition suggests reading Conjecture 2 as a question about the effective wedge noise σ_{prot} a one-round shuffle protocol can afford. First, wedge statistics computed from the ε-local transcript itself carry per-user-correlated noise whose Perron component alone costs $\Theta(n/\varepsilon)$, already at the threshold, so no post-processing of T helps (Theorem 5). Second, within $\mathcal{C}$, where row reports touch $O(1)$ coordinates per edge, a heuristic, unproven within-row-composition ledger caps the per-bit budget at $\widetilde{O}(n^{-1/2})$. This puts $\sigma_{\mathrm{prot}} = \widetilde{O}(n)$ at the threshold, so within $\mathcal{C}$ the conjecture would hinge on polylogs we cannot resolve. Third, outside $\mathcal{C}$ no heuristic is needed. If each wedge pair is released as its own shuffled aggregate, so that one edge influences $\Theta(n)$ scalar pair reports, the negative-binomial release of Fang & Yi (2025) with sampling rate $q = 1$ and degree cap $d = n - 1$ certifies this transcript (Lemma 10). Its i.i.d. graph-independent discrete-Laplace wedge noise has scale $\sigma_{\mathrm{FY}} = \widetilde{O}(\sqrt{n}/\varepsilon) \ll \sigma^*$. Theorem 7(ii) then gives Corollary 1.

Corollary 1 (Shuffle beats per-edge local on the hard family)**.** *There are constants $\varepsilon_0 \in (0, 1]$ and $c, C > 0$ such that for all $\varepsilon \in [n^{-1/4}, \varepsilon_0]$ and $\delta = n^{-\Theta(1)}$, on the product family $\mu_\pm$ of Theorem 5 at $k = 5$: the one-round shuffle protocol that releases (a) the off-diagonal wedge matrix by the pair-indexed negative-binomial mechanism of Fang & Yi (2025) ($q = 1$, $d = n - 1$; interface derived in Lemma 10), (b) all degrees—supplying the diagonal $W_{qq} = d_q$—by their degree release (their Lem. 3.1), and (c) the adjacency by randomized response, under a constant 3-way budget split, is (ε, δ)-edge-DP on all graphs, uses $\widetilde{O}(n^2)$ bits of communication per user, and admits a spectral-deflation post-processing estimator (Theorem 7(ii), applied to the unequal off-diagonal and degree noise scales of (a)–(b)) with*

$$\max_{\pm} \mathbb{E}_{\mu_\pm}\big[(\widehat{N}_5 - N_5)^2\big]^{1/2} \leq C\, n^{3.5}\, \mathrm{polylog}(n)/\varepsilon,$$

while by Theorem 5 (whose proof bounds exactly this prior risk) every estimator over a product of pure ε-DP per-edge channels pays $\max_\pm \mathbb{E}_{\mu_\pm}[(\widehat{N} - N_5)^2]^{1/2} \geq c\, n^4/\varepsilon$ on the same family. On this family the shuffle model thus outperforms the non-interactive per-edge local model by $\widetilde{\Theta}(\sqrt{n})$.

Remark 4 (Scope of the separation)**.** (a) Corollary 1 bounds only the Bayes risk on $\mu_\pm$. Near-bipartite dense inputs defeat deflation, and the worst-case dense shuffle complexity of C_5 is open. (b) Corollary 1 compares an (ε, δ)-edge-DP shuffle protocol with pure-ε per-edge randomizers, a strict subclass of row-local ones. Against the full local model, it improves on the best known upper bound $O(n^4)$ (Suppakitpaisarn et al., 2025), but no row-local lower bound above $\Omega(n^{3.5})$ is known. (c) The protocol lies outside $\mathcal{C}$, so Conjectures 1 and 2 are untouched. Our gain, like the certified C_4 rate of Theorem 1, costs $\Theta(n)$ messages per edge, and whether fewer suffice is open.

6 EXPERIMENTS

We evaluate the C_5 protocol of §4 against local randomized response (RR) (Warner, 1965; Imola et al., 2021) and central Laplace (Dwork et al., 2006) on $G(n, \frac{1}{2})$ and on the hard family $\mu_\pm$, both dense, with $n \le 1600$. We simulate each shuffle mechanism through its asymptotic noise law, with normalized constants (polylog and budget-split constants set to 1), and do not run the protocols end to end at the message level. The plots therefore validate scaling exponents. They do not measure privacy-calibrated practical performance, and a gap between a shuffle curve and a baseline curve is not a measured constant-factor advantage. The estimand is the simple N_5 (exact ground truth via Lemma 6). Each point averages 16 independent repetitions, and $\pm$ is the least-squares standard error of the four-point log–log fit. Figure 1 and full details are in Appendix A.10.

Rates on $G(n, \frac{1}{2})$. Our estimator and local RR show the same n^4 rate. Their RMSEs fit slopes 4.02 ± 0.06 and 3.97 ± 0.07, against 3.20 for central Laplace, and the theoretical curve of the pivot baseline (Lemma 8) lies a factor $\sqrt{n}$ above both. Our ε-slope is -0.97 ± 0.04 (predicted -1). As Theorem 3 predicts, the bare-edge channel carries the error. Alone it fits slope 4.06 and matches the full RMSE at $n = 1600$, while the wedge channels fit ≈ 3.6 and lie $18\times$ below. At $n = 1600$ and $\varepsilon = 1$, RMSE/$\mathbb{E}N_5$ is 0.59% for ours, 0.78% for local RR and 0.015% for central Laplace.

Separation on $\mu_\pm$. The slopes are consistent with Corollary 1. With the full estimator of Theorem 7(ii), deflation fits slope 3.35 ± 0.08 against 3.98 ± 0.12 for per-edge RR. The difference, 0.63, has paired-bootstrap 95% CI $[0.27, 1.01]$. The gap grows from $2.3\times$ at $n = 200$ to $7.8\times$ at $n = 1600$, where RMSE/$\mathbb{E}N_5$ at $\varepsilon = 1$ is 0.10% for deflation and 0.76% for per-edge RR. In both shuffle protocols, each user sends $\widetilde{O}(n^2)$ bits, one report per pair per wedge release. The analyzer sums $\Theta(n^3)$ messages and computes $O(1)$ matrix products, plus a top eigenpair for deflation.

7 RELATED WORK

Shuffle-model graph counting. Imola et al. (2022) first showed that shuffling wedge messages improves triangle and 4-cycle counting over the pure local model; Fang & Yi (2025) generalized this via a negative-binomial mechanism. Both handle patterns on at most four nodes, and to our knowledge no shuffle-model 5-cycle protocol precedes ours (Corollary 2, Appendix A.3).

Local-DP graph counting. In the pure local model Imola et al. (2021) gave triangle and k-star estimators; the sharp one-round triangle analysis is due to Imola et al. (2022) and Eden et al. (2023). For general patterns, Suppakitpaisarn et al. (2025) give a non-interactive graphlet upper bound of $O(n^{k-1})$ with a matching lower bound only for cliques. Their cycle lower bound $\Omega(n^{k-1.5})$ holds even for interactive protocols. For non-interactive per-edge randomizers, a strict subclass of their row-level randomizers, Theorem 5 closes this clique-versus-cycle gap with explicit ε-dependence. Remark 4(b) and Appendix A.9 discuss the gap that remains for row-level and interactive protocols. Hillebrand et al. (2023) debias polynomials of Laplace-perturbed local reports and Betzer et al. (2024) publish clipped per-vertex walk counts (not traces or cycle counts) over several local rounds. We instead get unbiasedness from independent one-round releases of the factors (Lemma 1).

8 CONCLUSION

For closed walks, parity decides whether the count factors through $W = A^2$. Even traces do, and odd traces keep one bare edge. The 4-cycle utility improvement is known (Fang & Yi, 2025). For odd k our fixed-ledger trace estimator has dense rate $\Theta_k(n^{k-1}/\varepsilon)$ and the simple-count estimator preserves the corresponding upper rate. These rates match the pure-ε per-edge local lower bound (Theorem 5). For C_5 on the hard family, a certified protocol outside $\mathcal{C}$ has wedge noise far below the threshold of Theorem 7 and improves on the lower bound of Theorem 5 by $\widetilde{\Theta}(\sqrt{n})$. We leave open the seam barrier inside $\mathcal{C}$ (Conjectures 1 and 2), whether interaction recovers the same factor in the full local model, and the trade-off between messages per edge and attainable wedge noise.

AI USE STATEMENT

In this work, we used generative AI tools to provide ingredients for proofs, to assist in writing and checking them, to formulate and correct some of the mathematical claims, and to implement the simulation code. We have not used them to generate data sets; translation, data cleaning and qualitative analysis are not applicable. Additionally, we used them to edit and rewrite parts of the paper for readability, to find and summarize related literature, and to format references. We have reviewed all AI-assisted work: the authors re-derived and checked every proof, reviewed and tested all code (Appendix A.10), and verified every AI-suggested reference. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work is a theoretical and simulation study of differentially private graph statistics. It uses only synthetically generated random graphs, involves no human subjects, personal data, or real-world networks, and releases no data set or model that could be misused. Its purpose is to clarify which graph statistics can be released under stronger privacy protection; we are not aware of negative societal consequences specific to this work.

REPRODUCIBILITY STATEMENT

All main results are stated with their assumptions in the main text, where Lemma 1 is proved and Proposition 1 and Theorems 1, 2, 3, 5, and 7 carry proof sketches (fuller sketches of Theorems 1 and 3, the sketch of Theorem 4, and a remark on Theorem 2, in Appendix A.11), with the tier computation behind Theorem 3 in Lemma 7 (Appendix A.2). Appendix A gives complete proofs of the per-edge local lower bound (Theorem 5) with its supporting lemmas (Appendix A.1); the walk-to-cycle conversion (Lemma 6), Proposition 1, and the full tier computation for Theorem 4 (Appendix A.2); the corrected-ledger pivot baseline (Lemma 8, Appendix A.3); Theorem 6 and Theorem 7(i) (Appendix A.4); Theorem 7(ii) (Appendix A.7); and the Fang–Yi interface (Lemma 10, Appendix A.5). Corollary 1 is assembled in Appendix A.8 from these results and the discrete-Laplace transfer lemmas of Appendix A.6. The experiments of §6 are noise-law simulations that need no external data: Appendix A.10 lists every graph family, parameter, noise scale, estimator, and the statistical procedure behind each reported slope. The complete simulation script, its dependency list (NumPy and Matplotlib; developed with Python 3.12 and NumPy 2.2), the raw results file with all fitted slopes, their standard errors, and the paired-bootstrap confidence interval for the separation slope difference, and the four figures are provided as anonymous supplementary material. The script is driven by fixed random seeds, runs on one commodity CPU core in 10–20 minutes, and verifies its own exact-count identities, multilinearity, and estimator unbiasedness at start-up.

REFERENCES

Louis Betzer, Vorapong Suppakitpaisarn, and Quentin Hillebrand. Publishing number of walks and Katz centrality under local differential privacy. In *Proceedings of the 40th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 244 of *PMLR*, pp. 377–393, 2024.

Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *Advances in Cryptology – EUROCRYPT 2019*, 2019.

Cynthia Dwork and Aaron Roth. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science, 2014.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference (TCC)*, 2006.

Talya Eden, Quanquan C. Liu, Sofya Raskhodnikova, and Adam Smith. Triangle counting with local edge differential privacy. In *50th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 261 of LIPIcs, pp. 52:1–52:21, 2023. doi: 10.4230/LIPIcs.ICALP.2023.52. Full version: arXiv:2305.02263.

Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2019.

Juanru Fang and Ke Yi. Counting subgraphs under shuffle differential privacy. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2025. doi: 10.1145/3719027.3765047.

Vitaly Feldman, Audra McMillan, and Kunal Talwar. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling. In *62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2021.

Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, 2009.

Quentin Hillebrand, Vorapong Suppakitpaisarn, and Tetsuo Shibuya. Unbiased locally private estimator for polynomials of Laplacian variables. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 741–751, 2023. doi: 10.1145/3580305.3599537.

Jacob Imola, Takao Murakami, and Kamalika Chaudhuri. Locally differentially private analysis of graph statistics. In *30th USENIX Security Symposium (USENIX Security)*, 2021.

Jacob Imola, Takao Murakami, and Kamalika Chaudhuri. Differentially private triangle and 4-cycle counting in the shuffle model. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2022.

Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

Vorapong Suppakitpaisarn, Donlapark Ponnoprat, Nicha Hirankarn, and Quentin Hillebrand. Counting graphlets of size k under local differential privacy. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2025. arXiv:2505.12954.

Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

A Technical Appendix

This appendix proves Theorems 5, 6, and 7 (part (ii) under the unified off-diagonal/diagonal noise model), Proposition 1, and Lemmas 6 and 10; gives the full computation for Theorem 4 and the corrected-ledger pivot baseline (Lemma 8); derives Corollary 1 via the discrete-Laplace transfer lemmas; records the experimental details; and collects the proof sketches for Theorems 1–4 (Appendix A.11). Numbering of theorems, lemmas, equations, and figures continues that of the main paper.

A.1 The per-edge local lower bound: Theorem 5

Setup and notation. Fix a cycle length $k \geq 3$; $k = 5$ is the paper's running instance, but nothing below uses $k = 5$. Edges are indexed by $e \in \binom{[n]}{2}$, $m = \binom{n}{2} = \Theta(n^2)$; $A \in \{0,1\}^{n\times n}$ is the symmetric adjacency matrix with $A_e \in \{0,1\}$ the free bits. The target functional is the simple-cycle count $N = N_k(A)$ itself. We deliberately do not route through the closed-walk proxy $F = \mathrm{tr}(A^k)$: the two differ by the degeneracy corrections D_k of Lemma 6, and D_k can be as large as $\Theta(n^{k-1})$ for $k \geq 4$ in the dense regime—for $k = 5$ the tadpole walks (triangle plus pendant edge) alone contribute $\Theta\big(\sum_i (d_i - 2)(A^3)_{ii}\big) = \Theta(n^4)$—the same order as the minimax risk we prove, so a transfer through F would be lossy. (At $k = 3$, $D_3 = 0$ and the transfer is lossless; running on N directly is simply uniform in k.) Running the argument on N directly costs nothing, because N has the two properties the lemmas need: it is monotone in the edge bits, and its discrete derivatives are

order-stable. For a product parameter $p \in (0,1)$ let $\mu_p = \mathrm{Bern}(p)^{\otimes m}$ and $\psi(p) := \mathbb{E}_{\mu_p}[N]$; the two hypotheses are $\mu_\pm := \mu_{1/2\pm\gamma}$ with $\gamma \in (0, \frac{1}{4}]$ to be fixed. Each degree stays $\Theta(n)$ w.h.p. across the prior bulk, since $\gamma = o(1)$. For a channel product $Q = \bigotimes_e Q_e$ write $P_\pm := Q \circ \mu_\pm$ for the transcript laws. Discrete derivatives: $\Delta_e N := N|_{A_e=1} - N|_{A_e=0}$, equal to the number of simple k-cycles through $e = \{q, r\}$ when e is present, so $\Delta_e N \geq 0$; on the bulk $\Delta_e N = \Theta(n^{k-2})$, being a sum over the $k-2$ free vertices of the completing path.

Lemma 2 (Mean separation via Margulis–Russo)**.** *For all $p \in [\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]$, $\psi'(p) = \sum_e \mathbb{E}_{\mu_p}[\Delta_e N] = \Theta(n^k)$, uniformly. Consequently $2s := \psi(\frac{1}{2} + \gamma) - \psi(\frac{1}{2} - \gamma) = \Theta(n^k \gamma)$.*

Proof. Margulis–Russo for a product measure gives $\frac{d}{dp}\mathbb{E}_{\mu_p}[N] = \sum_e \mathbb{E}_{\mu_p}[\Delta_e N]$. Every summand is ≥ 0 (adding an edge never destroys a simple cycle) and equals the expected number of k-cycles through e: cycles through e correspond bijectively to ordered tuples of $k-2$ distinct vertices in $[n] \setminus \{q, r\}$, each completing the cycle with probability $p^{k-1} = \Theta(1)$, so $\mathbb{E}_{\mu_p}[\Delta_e N] = \Theta(n^{k-2})$, strictly positive and of matched order, uniformly for $p = \frac{1}{2} + o(1)$. Summing $m = \Theta(n^2)$ sign-coherent terms gives $\psi'(p) = \Theta(n^k)$. Integrate over $[\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]$. □

Lemma 3 (Efron–Stein fluctuation cap)**.** *For every $p = \frac{1}{2} + o(1)$, $\mathrm{Var}_{\mu_p}(N) \leq \frac{1}{4}\sum_e \mathbb{E}_{\mu_p}[(\Delta_e N)^2] = \Theta(n^{2k-2})$, hence $\mathrm{sd}_{\mu_p}(N) \leq \sigma := C_\sigma n^{k-1}$.*

Proof. Efron–Stein for independent coordinates:

$$\mathrm{Var}(N) \leq \sum_e \mathbb{E}[\mathrm{Var}_e(N)] = \sum_e p(1-p)\mathbb{E}[(\Delta_e N)^2] \leq \frac{1}{4}\sum_e \mathbb{E}[(\Delta_e N)^2].$$

Each $\Delta_e N$ is a sum of $\Theta(n^{k-2})$ completing-path indicators (products of $k-1$ edge bits over the $k-2$ free vertices), and $\mathbb{E}[(\Delta_e N)^2] = (\mathbb{E}\Delta_e N)^2 + \mathrm{Var}(\Delta_e N)$: indicator pairs with disjoint edge sets are independent under μ_p; pairs sharing an edge share an internal vertex and hence number $O(n^{2k-5})$, with covariance at most 1 each; and the diagonal is $\Theta(n^{k-2})$. So $\mathrm{Var}(\Delta_e N) = O(n^{2k-5})$ and $\mathbb{E}[(\Delta_e N)^2] = \Theta(n^{2k-4})$. Summing, $\sum_e \mathbb{E}[(\Delta_e N)^2] = \Theta(n^2) \cdot \Theta(n^{2k-4}) = \Theta(n^{2k-2})$, whence $\mathrm{sd}(N) \leq C_\sigma n^{k-1}$. □

Lemma 4 (ε-local indistinguishability; DJW contraction)**.** *Let $Q = \bigotimes_e Q_e$ be any product of ε-local channels, $\varepsilon \leq 1$. Then $\mathrm{TV}(P_+, P_-)^2 \leq \frac{1}{2}\mathrm{KL}(P_+ \| P_-) \leq C\varepsilon^2 m\gamma^2$, and, for $\varepsilon n \geq c$, there is $\gamma^* = \Theta(1/(\varepsilon n))$ with $\mathrm{TV}(P_+, P_-) \leq \frac{1}{4}$ for all $\gamma \leq \gamma^*$. The bound is achieved (up to constants) by randomized response with bias $\Theta(\varepsilon)$ (flip probability $\frac{1}{2} - \Theta(\varepsilon)$), whose debiased plug-in attains RMSE $O_k(n^{k-1}/\varepsilon)$ for $\varepsilon \geq n^{-(k-2)/(2k-2)}$.*

Proof. The Duchi–Jordan–Wainwright strong data-processing / Fisher-information contraction for non-interactive ε-local channels gives $\mathrm{KL}(P_+ \| P_-) \leq 4(e^\varepsilon - 1)^2 \sum_e (2\gamma)^2 = 16(e^\varepsilon - 1)^2 m\gamma^2$, each edge-marginal TV being 2γ. For $\varepsilon \leq 1$, $(e^\varepsilon - 1)^2 \leq (e-1)^2\varepsilon^2$, so by Pinsker $\mathrm{TV}^2 \leq C\varepsilon^2 m\gamma^2$ with $m = \Theta(n^2)$; setting $C\varepsilon^2 m(\gamma^*)^2 = \frac{1}{16}$ yields $\gamma^* = \Theta(1/(\varepsilon\sqrt{m})) = \Theta(1/(\varepsilon n))$. Tightness: RR with bias $\Theta(\varepsilon)$ observes each edge with variance $v_A = \Theta(1/\varepsilon^2)$; the debiased multilinear plug-in for N_k (itself multilinear, as simple cycles have distinct edges) has gradient ℓ_2-norm $\big(\sum_e (\Delta_e N)^2\big)^{1/2} = O(n^{k-1})$ on dense graphs, so its single-noise tier has RMSE $O(n^{k-1}/\varepsilon)$. For the higher tiers, expand over noise-edge sets S: the tier with $|S| = q$ has variance $v_A^q \sum_S c_S^2$, c_S the completion count. A proper subset of a cycle is a union of paths, giving $\sum_S c_S^2 = O(n^{2k-q-1})$ and tier-to-gradient ratio $O\big((n\varepsilon^2)^{-(q-1)}\big)$, dominated once $\varepsilon \geq n^{-1/2}$; the one closed set ($S$ the whole cycle, $c_S = 1$) contributes $\Theta(n^k)v_A^k = \Theta(n^k/\varepsilon^{2k})$, dominated by n^{2k-2}/ε^2 exactly when $\varepsilon \geq n^{-(k-2)/(2k-2)}$—the binding constraint ($n^{-1/4}$ at $k=3$, $n^{-3/8}$ at $k=5$, approaching $n^{-1/2}$ as $k \to \infty$). □

Lemma 5 (Estimand-to-test reduction: loss against the realized random count)**.** *Let $\mu_\pm$ have transcript laws $P_\pm$, let $S := \frac{1}{2}(\psi_+ + \psi_-)$ with $\psi_\pm := \mathbb{E}_{\mu_\pm}[N]$ and $2s := \psi_+ - \psi_- > 0$, and suppose $\mathrm{sd}_{\mu_\pm}(N) \leq \sigma$ and $\mathrm{TV}(P_+, P_-) \leq \frac{1}{4}$. If $s \geq 2\sqrt{8}\,\sigma$ then every transcript estimator $\widehat{N}(T)$ satisfies $\max_\pm \mathbb{E}_{\mu_\pm}[(\widehat{N} - N)^2] \geq \frac{1}{16}s^2$.*

Proof. Set threshold $c_0 := S$ and back-off $\tau := \sqrt{8}\,\sigma \le s/2$, and define the good events $G_+ = \{N \ge c_0 + s - \tau\}$, $G_- = \{N \le c_0 - s + \tau\}$. Since $\mathbb{E}_{\mu_+}[N] = c_0 + s$ with $\mathrm{sd} \le \sigma$, Chebyshev gives $\mu_+(G_+^c) \le \sigma^2/\tau^2 \le \frac{1}{8} =: \beta$, and symmetrically for μ_-; note $s - \tau \ge s/2$. For any $\widehat{N}$ form the test $\phi := \mathbf{1}[\widehat{N} \ge c_0]$, a function of T only. On G_+, $\phi = 0 \Rightarrow |\widehat{N} - N| \ge s - \tau \ge s/2$, so $\mathbb{E}_{\mu_+}[(\widehat{N} - N)^2] \ge (s/2)^2(P_+(\phi = 0) - \beta)$; symmetrically for μ_-. Averaging and applying Le Cam $P_+(\phi = 0) + P_-(\phi = 1) \ge 1 - \mathrm{TV}(P_+, P_-)$,

$$\begin{aligned}\max_{\pm} \mathbb{E}_{\mu_\pm}[(\widehat{N} - N)^2] &\ge \frac{1}{2}\left(\frac{s}{2}\right)^2 (1 - \mathrm{TV} - 2\beta) \\ &\ge \frac{1}{2}\left(\frac{s}{2}\right)^2 \left(1 - \frac{1}{4} - \frac{1}{4}\right) = \frac{1}{16} s^2.\end{aligned}$$

No structure on $\widehat{N}$ was used, so the bound holds for all estimators—linear, spectral, U-statistic, arbitrary. □

Proof of Theorem 5. Fix $k \ge 3$. Take $\mu_\pm = \mathrm{Bern}(\frac{1}{2} \pm \gamma^*)^{\otimes m}$ with $\gamma^* = \Theta(1/(\varepsilon n))$ from Lemma 4, so $\mathrm{TV}(P_+, P_-) \le \frac{1}{4}$ for every product of ε-local channels (Lemma 4 is k-free). By Lemma 2, $2s = \Theta(n^k \gamma^*) = \Theta(n^{k-1}/\varepsilon)$, i.e. $s \ge C_s n^{k-1}/\varepsilon$ with $C_s = C_s(k)$; by Lemma 3, $\sigma \le C_\sigma n^{k-1}$, $C_\sigma = C_\sigma(k)$. The hypothesis $s \ge 2\sqrt{8}\,\sigma$ of Lemma 5 reads $C_s/\varepsilon \ge 2\sqrt{8}\,C_\sigma$, i.e. $\varepsilon \le \varepsilon_0(k) := C_s/(2\sqrt{8}\,C_\sigma)$. For such ε, Lemma 5 gives $\max_\pm \mathbb{E}_{\mu_\pm}[(\widehat{N} - N)^2] \ge \frac{1}{16} s^2 = \Omega(n^{2k-2}/\varepsilon^2)$, whence RMSE $\ge \Omega(n^{k-1}/\varepsilon)$ against the realized random N. To respect the dense-regime scoping, intersect the good events of Lemma 5 with the bulk event $\mathcal{B} := \{\text{all degrees } \Theta(n)\}$: $\mu_\pm(\mathcal{B}^c) = e^{-\Omega(n)}$ folds into β and the witness graphs are dense; the supremum over G then dominates the prior average. Tightness is Lemma 4. The $\sqrt{n}$ between $\|c\|_2 = \Theta(n^{k-1})$ and the worst per-vertex block $\Theta(n^{k-1.5})$ is the honest per-edge tensorization gain $\sqrt{m} = \Theta(n)$: the gradient is sign-coherent (Lemma 2) and the channels are per-edge, so detection anchors at the m edge channels, not a single vertex block. This edge-versus-vertex anchoring is exactly the $\sqrt{n}$ separating the bound from the $\Omega(n^{k-1.5})$ cycle bound of Suppakitpaisarn et al. (2025). □

Remark 5 (Why two-point arguments cannot reach n^4)**.** The bound above needs the full $\Theta(n^2)$-coordinate product family; a two-point/Le Cam argument over a single same-square pair cannot substitute. If A, A' are connected non-bipartite dense graphs with $A^2 = A'^2$, they share the eigenvalue-modulus multiset, so $|\operatorname{tr}(A^5) - \operatorname{tr}(A'^5)| \le 2\sum_{i \ge 2} |\lambda_i|^5$, the Perron term $|\lambda_1| = \Theta(n)$ being sign-locked by nonnegativity. On a quasirandom base this is $O(n^{3.5})$—a factor $\sqrt{n}$ short of $\|c\|_2$—which is why Theorem 5 routes through the product family rather than a fixed-W fiber; the same ceiling extends to N_5, the corrections D_5 being a shared-W linear read-out of A (Lemma 6) obeying the same modulus bound.

A.2 Image types, noise tiers, and all odd cycles

Lemma 6 (Image-type and noise-tier conversion)**.** *Fix $k \ge 3$. For a simple graph H, write*

$$\mathrm{inj}(H, G) := \sum_{\phi : V(H) \hookrightarrow [n]} \prod_{uv \in E(H)} A_{\phi(u)\phi(v)}.$$

There is a finite collection $\mathcal{D}_k$ of connected simple graphs and positive integer coefficients $c_{k,H}$ such that

$$\operatorname{tr}(A^k) = 2kN_k + D_k, \qquad D_k = \sum_{H \in \mathcal{D}_k} c_{k,H}\, \mathrm{inj}(H, G),$$

where every $H \in \mathcal{D}_k$ is the edge-support image of a non-injective closed k-walk and satisfies $v(H) \le k - 1$, $e(H) \le k$. Moreover $D_3 = 0$, $D_k = O_k(n^{k-1})$, and every image count has per-edge influence $O_k(n^{k-3})$.

Let $\widehat{A} = A + M$ be the hollow (zero-diagonal) debiased randomized-response release and assume $\max_{i<j} \|M_{ij}\|_2 \le \tau$. Define $\widehat{\mathrm{inj}}(H)$ by replacing each distinct edge bit in the preceding display by the corresponding entry of $\widehat{A}$. This plug-in is unbiased. If $v_H(F)$ is the number of vertices incident

to a nonempty edge set $F \subseteq E(H)$, then

$$\|\widehat{\mathrm{inj}}(H) - \mathrm{inj}(H, G)\|_2 \le C_H \sum_{\varnothing \neq F \subseteq E(H)} n^{v(H)-v_H(F)/2} \tau^{|F|}.$$

Consequently, for debiased randomized response and $n^{-1/2} \le \varepsilon \le \varepsilon_0$,

$$\|\widehat{D}_k - D_k\|_2 = O_k(n^{k-1}/\varepsilon).$$

The correction is $o_k(n^{k-1}/\varepsilon)$ whenever $\varepsilon\sqrt{n} \to \infty$, but not uniformly at $\varepsilon = n^{-1/2}$ for odd $k \ge 7$. For $k = 5$,

$$\|\widehat{D}_5 - D_5\|_2 = O\left(\frac{n^3}{\varepsilon} + \frac{n^{2.5}}{\varepsilon^3} + \frac{n^2}{\varepsilon^4}\right),$$

so the correction is $o(n^4/\varepsilon)$ throughout $\varepsilon \ge n^{-1/2}$ and is $O(n^3/\varepsilon)$ once $\varepsilon \ge n^{-1/4}$.

Proof. Group non-injective closed k-walks by their simple edge-support image H. For a fixed nonempty $F \subseteq E(H)$, let $f = |F|$, $s = v_H(F)$, and let X_F be the expansion tier using noise exactly on F. Group its summands by the image edge set $S = \phi(F)$. There are $O_H(n^s)$ possible sets S, and a fixed S admits at most $O_H(n^{v(H)-s})$ extensions of the remaining vertices. Distinct noise monomials are orthogonal because some independent centered edge variable occurs exactly once. Hence

$$\|X_F\|_2^2 \le C_H n^s n^{2(v(H)-s)} \tau^{2f},$$

which proves the tier bound and unbiasedness.

Write $v = v(H)$ and $e = e(H)$. Since H is connected, the cycle rank of the subgraph spanned by F is at most $e - v + 1$, so $f - s \le e - v$. With $\varepsilon^{-1} \le \sqrt{n}$,

$$n^{v-s/2}\varepsilon^{-f} \le \frac{1}{\varepsilon} n^{v+(f-s-1)/2} \le \frac{1}{\varepsilon} n^{(v+e-1)/2} \le \frac{n^{k-1}}{\varepsilon}.$$

Summing over the finitely many image types proves the uniform rate. If $r_n = \varepsilon\sqrt{n} \to \infty$, every tier with $f \ge 2$ gains a factor $r_n^{-(f-1)}$, while the $f = 1$ tier is $O_k(n^{k-2}/\varepsilon)$, proving the little-o statement. For $k = 5$, the only proper supports are a triangle and a triangle with one pendant edge; applying the displayed tier bound to their edge subsets gives the stated refinement. □

Remark 6 (Sharpness at the endpoint)**.** For odd $k \ge 7$, let $H = C_3 \vee C_{k-3}$. Then $v(H) = k - 1$, $e(H) = k$, and H is Eulerian, so it occurs with positive coefficient in D_k. Its all-noise tier has RMSE $\Theta_k(n^{(k-1)/2}/\varepsilon^k)$, equal to the target order at $\varepsilon = n^{-1/2}$. For $k = 7$ and $\varepsilon = n^{-3/8}$ this tier is $n^{45/8}$, whereas $n^{k-2}/\varepsilon = n^{43/8}$; thus no uniform $O_k(n^{k-2}/\varepsilon)$ threshold of the former form is valid.

Proof of Proposition 1. Take $G = C_6$, $G' = 2K_3$. Labeling C_6 so its parity classes are $\{1, 3, 5\}$ and $\{2, 4, 6\}$, the square $A(C_6)^2$ has diagonal 2 and off-diagonal support exactly on the distance-two pairs—two disjoint triangles on the parity classes; $(J_3 - I_3)^2 = J_3 + I_3$ gives the identical matrix for $2K_3$. Yet C_6 is bipartite, so $\mathrm{tr}(A^5) = 0$, while $2K_3$ has spectrum two copies of $(2, -1, -1)$ and $\mathrm{tr}(A^{2m+1}) = 2\big(2^{2m+1} + 2(-1)^{2m+1}\big) = 2(2^{2m+1} - 2) > 0$ for every $m \ge 1$ (60 at $m = 2$): the pair shares W—hence all eigenvalue moduli—but differs in eigenvalue signs (so the pair is not cospectral; no same-W pair with different odd traces can be, since each $\mathrm{tr}(A^{2m+1})$ is a spectral invariant). Blowing each vertex up to an independent set of size t (edges to complete bipartite blocks) replaces A_0 by $A_0 \otimes J_t$, whose square is $A_0^2 \otimes tJ_t$: equal squares stay equal, all degrees become $\Theta(n)$, and the $(2m + 1)$-st traces scale by t^{2m+1}, preserving the separation—at every odd power—in the dense regime. □

Full computation for Theorem 4. Unbiasedness: the $m + 1$ factors come from mutually independent, entrywise mean-correct releases, so every noise cross-term vanishes as in Lemma 1. Privacy: the ledger length $L = m + 1 = (k + 1)/2$ is a constant depending only on k—never on n or on any pivot count—so composition is exactly as in Theorem 2. Accuracy: the seam channel is

$X_M = \mathrm{tr}(W^m M) = \langle M, A^{k-1}\rangle$; with M mirrored and independent above the diagonal ($M_{uu} = 0$: the blind release reports no diagonal),

$$\mathrm{Var}[X_M] = 4v_A \sum_{u<v} (A^{k-1})^2_{uv} = \Theta\big(v_A \|A^{k-1}\|_F^2\big) = \Theta(n^{2k-2}/\varepsilon^2),$$

using that $\sum_u (A^{k-1})^2_{uu} = O(n^{2k-3}) = o(\|A^{k-1}\|_F^2)$ in the dense regime, so the off-diagonal sum is $(\frac{1}{2} - o(1))\|A^{k-1}\|_F^2$ ($\|A^{k-1}\|_F = \Theta(n^{k-1})$ by the spectral bounds $\lambda_1^{k-1} \le \|A^{k-1}\|_F \le \|A\|_2^{k-2}\|A\|_F$ with $\lambda_1 = \Omega(n)$ from the average degree), i.e. RMSE $\Theta(n^{k-1}/\varepsilon)$, again pinned by the un-amplified bare edge. The i-th wedge channel is $\mathrm{tr}(N_i A^{k-2})$ (note $W^{m-i}AW^{i-1} = A^{k-2}$ for every i). Under the certified instantiation the off-diagonal entries of N_i are i.i.d. with per-entry variance $\widetilde{O}(n/\varepsilon^2)$ (Lemma 10), giving $\widetilde{O}(n/\varepsilon^2)\|A^{k-2}\|_F^2 = \widetilde{O}(n^{2k-3}/\varepsilon^2)$, a factor $\widetilde{O}(n)$ below the seam; the degree-noise diagonal (per-entry variance $O(1/\varepsilon^2)$, coefficients $(A^{k-2})_{qq} \le n^{k-3}$) contributes $O(n^{2k-5}/\varepsilon^2)$, lower order. (In the conditional per-user variant the entries of N_i are correlated through per-user randomness, so the variance is $\Theta\big(v_0 \,\mathrm{tr}(A^{2k-2})\big) = \Theta\big(v_0\|A^{k-1}\|_F^2\big)$ with v_0 the per-coordinate report variance—the same coefficient mass as the seam; un-amplified ($v_0 = \Theta(1/\varepsilon^2)$) it would be co-dominant, and the amplified $v_0 = \widetilde{O}(1/(n\varepsilon^2))$—conditional on the within-row ledger of Remark 3—restores $\widetilde{O}(n^{2k-3}/\varepsilon^2)$, exactly what amplification would buy.) Higher-order noise products are L^2-orthogonal to the linear tier and, by the norm counting of Lemma 7 (k fixed), fit within the seam budget for $\varepsilon \ge n^{-1/2}$ ($\varepsilon \ge \sqrt{\log n/n}$ if $k = 3$). For the simple count, use $\widehat{N}_k = (\widehat{T}_k - \widehat{D}_k)/(2k)$ with the correction estimator of Lemma 6. It is unbiased and costs no extra privacy. Minkowski's inequality gives the $O_k(n^{k-1}/\varepsilon)$ rate. When $\varepsilon\sqrt{n} \to \infty$, or for $k = 5$ throughout the stated range, the correction is little-o of the trace error, so the reverse triangle inequality gives the claimed Θ_k rate. At $k = 3$, $D_3 = 0$ and the estimator reduces to $\mathrm{tr}(\widehat{W}\widehat{A})/6$, matching the non-interactive triangle frontier.

Lemma 7 (Orthogonal noise decomposition; the nonlinear tiers)**.** *Write the error of the trilinear estimator as* $\widehat{T} - \mathrm{tr}(A^5) = \sum_{\varnothing \ne S \subseteq \{N_1, M, N_2\}} T_S$*, where* T_S *collects the monomials of* $\mathrm{tr}\big((W + N_1)(A + M)(W + N_2)\big)$ *containing exactly the noise blocks in* S *(so* $T_{\{M\}} = X_M$*, etc.). Then: (i) the seven terms are pairwise uncorrelated, so* $\mathrm{Var}[\widehat{T}] = \sum_S \mathrm{Var}[T_S] \ge \mathrm{Var}[X_M]$*—the seam channel lower-bounds the total variance; (ii) in the dense regime, under either wedge instantiation of Theorem 3 and for* $\varepsilon \ge n^{-1/2}$*, the nonlinear tiers total* $\sum_{|S|\ge 2} \mathrm{Var}[T_S] = \widetilde{O}(n^7/\varepsilon^2)$*—the same order as the wedge channels, a factor* $\widetilde{O}(n)$ *below the seam. The analogous statements hold for the odd-*k *estimator* $\widehat{T}_k$ *with blocks* $\{N_1, \dots, N_m, M\}$*: the nonlinear tiers total* $\widetilde{O}_k(n^{2k-3}/\varepsilon^2 + n^k/\varepsilon^{k+1})$*, which for* $k \ge 5$ *and* $\varepsilon \ge n^{-1/2}$ *is* $O(n^{2k-2}/\varepsilon^2)$ *(so Theorem 4 is unaffected) and a factor* $\widetilde{O}(n)$ *below the seam; at* $k = 3$ *(certified instantiation) the single tier* $\mathrm{tr}(N_1 M)$ *is a* $\Theta(\log n)$ *factor above the seam at* $\varepsilon = n^{-1/2}$*, co-dominant at* $\varepsilon = \Theta(\sqrt{\log n/n})$*, and falls strictly below it for* $\varepsilon = \widetilde{\Omega}(n^{-1/2+c})$*,* $c > 0$*.*

Proof. (i) Fix $S \ne S'$ and pick a block $b \in S' \setminus S$ (swap S, S' if needed). Each block appears at most once in each monomial, so $T_{S'}$ is multilinear in the entries of b with coefficients not depending on b, while T_S does not involve b. Since b is mean-zero entrywise and independent of the remaining blocks and of A, conditioning on everything but b gives $\mathbb{E}[T_{S'} \mid \text{rest}] = 0$, whence $\mathbb{E}[T_S T_{S'}] = 0$; every T_S is centered for the same reason, so covariances vanish and the variance is additive. The bound $\mathrm{Var}[\widehat{T}] \ge \mathrm{Var}[X_M]$ is immediate. (ii) Independence across blocks makes $\mathrm{Var}[T_S]$ a sum, over index tuples, of products of per-entry variances times squared A-measurable coefficients. For $k = 5$, with $v_A = \Theta(1/\varepsilon^2)$ for M and per-entry wedge variance $v_W = \widetilde{O}(n/\varepsilon^2)$ (certified instantiation; the conditional variant is covered by the per-user computation in the proof of Theorem 4, with the same conclusions), the quadratic tiers are

$$\mathrm{Var}[\mathrm{tr}(N_1 MW)],\ \mathrm{Var}[\mathrm{tr}(WMN_2)] = \widetilde{O}\big(v_W v_A \cdot n\|W\|_F^2\big) = \widetilde{O}(n^6/\varepsilon^4),$$
$$\mathrm{Var}[\mathrm{tr}(N_1 AN_2)] = \widetilde{O}\big(v_W^2 \cdot n\|A\|_F^2\big) = \widetilde{O}(n^5/\varepsilon^4),$$

and the cubic tier is $\mathrm{Var}[\mathrm{tr}(N_1 MN_2)] = \widetilde{O}(v_W^2 v_A \cdot n^3) = \widetilde{O}(n^5/\varepsilon^6)$. For $\varepsilon \ge n^{-1/2}$ every extra factor $1/\varepsilon^2 \le n$ is absorbed, so each tier is $\widetilde{O}(n^7/\varepsilon^2)$; the counts run over index tuples up

to the $O(1)$ multiplicity from mirrored index pairs and lower-order diagonal collisions, and the degree-noise diagonals have per-entry variance $O(1/\varepsilon^2)$ and only shrink them. (For the conditional per-user variant, condition on all blocks but one: each tier is linear in the remaining correlated block, so the per-user variance computation of Theorem 4 applies with a random coefficient matrix, and taking expectations gives bounds no larger than the certified counts.) The odd-k case is the identical norm counting with coefficients of the form $W^a A W^b$ (k fixed); the largest tier is the all-noise one, $v_W^m v_A \cdot n^{m+1} = \widetilde{O}(n^k/\varepsilon^{k+1})$, absorbed below n^{2k-3}/ε^2 exactly when $k \geq 5$. At $k = 3$ the single nonlinear tier gets its own display: the error of $\widehat{T}_3$ has nonlinear part $\mathrm{tr}(N_1 M) = \sum_{i,j}(N_1)_{ij} M_{ji}$, and by block independence

$$\mathrm{Var}[\mathrm{tr}(N_1 M)] = \widetilde{O}\big(v_W v_A \cdot n^2\big) = \widetilde{O}\left(\frac{n}{\varepsilon^2} \cdot \frac{1}{\varepsilon^2} \cdot n^2\right) = \widetilde{O}\left(\frac{n^3}{\varepsilon^4}\right),$$

with $\widetilde{O}$ hiding exactly the factor $\log(1/\delta) = \Theta(\log n)$ of $v_W = \Theta(b^2)$ (Lemma 10(b)); as $\mathrm{tr}(N_1 M) = 2\sum_{i<j}(N_1)_{ij} M_{ij}$ is a sum of independent centered terms, this variance is in fact $\Theta(n^3 \log n/\varepsilon^4)$, a $\Theta(\log n)$ factor above the seam variance at $\varepsilon = n^{-1/2}$. It is O of the seam variance $\Theta(n^4/\varepsilon^2)$ for $\varepsilon \geq \sqrt{\log n/n}$—co-dominant at that boundary, strictly below it for $\varepsilon = \widetilde{\Omega}(n^{-1/2+c})$, $c > 0$—so Theorem 4's $\Theta_3(n^2/\varepsilon)$ rate holds at $k = 3$ for $\varepsilon \geq \sqrt{\log n/n}$. □

Remark 7 (Wedge-channel variance of the conditional per-user variant)**.** In the proof sketch of Theorem 3 (Appendix A.11), the certified pair-indexed instantiation gives $\mathrm{Var}[X_{N_i}] = \widetilde{O}(n^7/\varepsilon^2)$ through $\|A^3\|_F^2$. The conditional per-user variant (Remark 1) lands at the same bound by a different route: its entries are correlated through per-user randomness, so the variance is governed not by $\|A^3\|_F^2$ but by the per-user computation $\Theta\big(v_0 \,\mathrm{tr}(A^8)\big) = \Theta\big(v_0 \|A^4\|_F^2\big)$—the same coefficient mass as the seam—which un-amplified ($v_0 = \Theta(1/\varepsilon^2)$) would be co-dominant, and amplified ($v_0 = \widetilde{O}(1/(n\varepsilon^2))$) restores $\widetilde{O}(n^7/\varepsilon^2)$.

A.3 The pivot-based baseline

The natural extension of the wedge-shuffling machinery of Imola et al. (2022); Fang & Yi (2025) to 5-cycles anchors an edge and sums, over pivots p, products of two wedge counts (entries of $W = A^2$ restricted to random sides): sampling $t \leq \lfloor n/2 \rfloor$ disjoint anchor pairs and k pivots per pair, it runs one shuffle binary-sum sub-mechanism per (pair, pivot, side), $L = 2tk$ in all, plus randomized response on each anchor edge. Parallel composition across pivots is invalid: one edge (v_j, u) feeds the wedge report for every pivot of its pair, and a wedge edge (u, p) recurs across pairs, so the per-pivot sub-accounts must be charged on a single sequential ledger of length $L = \Theta(tk)$—in contrast to our fixed $L = 3$.

Lemma 8 (Corrected-ledger rate of the pivot baseline)**.** *Calibrate each wedge sub-mechanism by advanced composition, $\varepsilon_0 = \Theta\big(\varepsilon/\sqrt{L\log(1/\delta)}\big)$, $\delta_0 = \Theta(\delta/L)$, and the anchor-edge reports at $\Theta(\varepsilon)$. The baseline is (ε, δ)-edge-DP and unbiased, and on dense graphs (all degrees $\Theta(n)$) its RMSE is $\widetilde{O}(n^{4.5}/\varepsilon)$ at the optimal choice $t = \Theta(n)$, $k = \Theta(\varepsilon^2 n)$; no choice of t, k improves the rate within this design.*

Proof. Each shuffle-sum estimate carries variance $\tau^2 = \widetilde{O}(1/\varepsilon_0^2) = \widetilde{O}(tk/\varepsilon^2)$: the ledger's composition tax. Writing $d = d_{\max}$, a product of two independent debiased wedge estimates (true counts $\leq d$) has variance $O(\tau^4 + d^2\tau^2)$, and pivot subsampling adds $O(n^2 d^4/k)$; with the pivot upscaling $(n-2)/k$, the edge factor ($\mathbb{E}[\hat{e}_s^2] = O(1/\varepsilon^2)$), and the true per-edge count $\leq d^3$, the per-pair estimator has

$$\mathrm{Var}[\widehat{T}_s] = \widetilde{O}\big(\varepsilon^{-2}\{n^2 d^4/k + n^2 t d^2/\varepsilon^2 + n^2 t^2 k/\varepsilon^4\} + d^6/\varepsilon^2\big).$$

Upscaling by $\binom{n}{2}/(5t)$ gives

$$\mathrm{RMSE} \leq \frac{n^2}{\sqrt{t}}\,\mathrm{Var}[\widehat{T}_s]^{1/2} = \widetilde{O}\left(\frac{n^3 d^2}{\varepsilon\sqrt{tk}} + \frac{n^3 d}{\varepsilon^2} + \frac{n^3\sqrt{tk}}{\varepsilon^3} + \frac{n^2 d^3}{\varepsilon\sqrt{t}}\right).$$

The first and third terms balance at $tk \asymp \varepsilon^2 d^2$; the composition tax $n^3\sqrt{tk}/\varepsilon^3$ forbids a larger pivot budget. At $d = \Theta(n)$ with $t = n$, $k \asymp \varepsilon^2 n$, every wedge term is $\widetilde{O}(n^4/\varepsilon^2)$, while the bare-edge

channel $n^2d^3/(\varepsilon\sqrt{t})$—randomized response on the anchor edge times the intrinsic per-edge count d^3, untouched by L but subsampled, with disjointness capping $t \le \lfloor n/2 \rfloor$—contributes $\widetilde{O}(n^{4.5}/\varepsilon)$ and dominates for $\varepsilon = \Omega(n^{-1/2})$. No choice of t, k does better: the last term alone is $\widetilde{\Omega}(n^{4.5}/\varepsilon)$ once $t \le n$. This sits a $\sqrt{n}$ factor above Theorem 3; the gap is the anchor-edge subsampling that our blind full-matrix release $\widehat{A}$ avoids. □

Corollary 2 (Comparison)**.** *To our knowledge no shuffle-model 5-cycle protocol precedes this work (Fang & Yi, 2025 and Imola et al., 2022 cap at 4-node patterns). The natural pivot-based extension of their wedge-shuffling machinery incurs $\widetilde{O}(n^{4.5}/\varepsilon)$ under a corrected composition ledger (per-pivot sub-accounts force $L = \Theta(tk)$-fold composition; Lemma 8); the trilinear design improves this baseline by a $\sqrt{n}$ factor, matching the local-DP rate $O(n^4)$ in n—optimal per-edge-locally by Theorem 5. Within $\mathcal{C}$ we claim no shuffle-vs-local separation (that is the open Conjecture 2); outside it, Corollary 1 beats the per-edge floor by $\widetilde{\Theta}(\sqrt{n})$ on the hard family.*

A.4 Exact wedges and the wedge-noise transition: Theorems 6 and 7(i)

The mechanism behind Theorem 6: the free W fixes the dominant term $\lambda_1^5 = \Theta(n^5)$ at zero variance (sign by Perron–Frobenius), leaving the subdominant eigenvalues, of aggregate coefficient norm $\Theta(n^{2.5})$, to estimate. Remark 10 (Appendix A.10) records why the free W already defeats the Theorem 5 construction: the scalar $\mathrm{tr}(W)$ alone separates μ_+ from μ_-.

Proof of Theorem 6. Since A commutes with $W = A^2$, choose a common eigenbasis and write $A = \sum_i s_i\sqrt{\mu_i}v_iv_i^\top$, $W = \sum_i \mu_i v_i v_i^\top$. On the bulk event, Perron–Frobenius fixes $s_1 = +1$. With $\Pi_1 = v_1v_1^\top$, $P = I - \Pi_1$, and $\Delta = W^2 - \mu_1^2\Pi_1$,

$$\mathrm{tr}(A^5) = \mu_1^{5/2} + \langle \Delta, A\rangle, \qquad \|\Delta\|_F^2 = O(n^5).$$

Thus the dominant Perron term is known from W and only a deflated linear readout is estimated from the debiased randomized-response matrix $\widehat{A} = A + M$.

For the simple-cycle correction, write $D_5 = 5\langle W+B, A\rangle$ with $B_{ij} = \frac{1}{2}(d_i+d_j-4)W_{ij}$. The exact diagonal of W gives every d_i and the exact average degree $\bar{d} = n^{-1}\,\mathrm{tr}(W)$. Set $C = \mathrm{diag}(d_i - \bar{d})$. Then

$$B = (\bar{d}-2)W + \tfrac{1}{2}(CW + WC),$$
$$\|PBP\|_F \le |\bar{d}-2|\|PWP\|_F + \|CW\|_F.$$

The first term has squared norm $O(n^5)$. Also

$$\|CW\|_F^2 \le n^3 \sum_i (d_i - \bar{d})^2.$$

If $m_d = \mathbb{E}d_i$ under either product prior, the sample average $\bar{d}$ minimizes squared deviations pointwise, so

$$\mathbb{E}\sum_i (d_i - \bar{d})^2 \le \mathbb{E}\sum_i (d_i - m_d)^2 = \sum_i \mathrm{Var}(d_i) = O(n^2).$$

Hence $\mathbb{E}\|PBP\|_F^2 = O(n^5)$ without any condition $\varepsilon \ge n^{-1/2}$; centering at $\bar{d}$ removes the common density shift.

For every symmetric X, $\langle X, A\rangle = \sqrt{\mu_1}v_1^\top X v_1 + \langle PXP, A\rangle$, so all Perron parts are known from W. The final estimator is a known W-measurable scalar plus $\langle K(W), \widehat{A}\rangle$, where $K(W) = \Delta - 5P(W+B)P$ and $\mathbb{E}\|K(W)\|_F^2 = O(n^5)$. Conditional on the graph, randomized response gives variance at most $O(\|K(W)\|_F^2/\varepsilon^2)$. Averaging and dividing by 10 yields $O(n^{2.5}/\varepsilon)$ RMSE; clipping handles the exponentially small complement of the bulk event. □

Lemma 9 (Posterior stability under an ε-local observation)**.** *Fix $\varepsilon \le 1$ and an ε-DP channel Q_e for the edge bit A_e. For priors $\pi_\pm = \frac{1}{2} \pm \gamma$, $\gamma \le \frac{1}{4}$, let $p_\pm(y) := \Pr_{\pi_\pm}[A_e = 1 \mid Y_e = y]$ be the posteriors after observing $Y_e = y$. Then $\sup_y |p_+(y) - p_-(y)| \le 2e^3\gamma$.*

Proof. Write $r(y) := q(y \mid 1)/q(y \mid 0) \in [e^{-\varepsilon}, e^{\varepsilon}]$ (ε-DP in the bit) and $p_\pi(y) = \pi r/(\pi r + 1 - \pi)$. Then $\partial_\pi p_\pi = r/(\pi r + 1 - \pi)^2$, and since $\pi r + 1 - \pi$ is a convex combination of r and 1, it is at least $\min(r, 1) \geq e^{-\varepsilon}$; hence $\partial_\pi p_\pi \leq e^{\varepsilon} \cdot e^{2\varepsilon} = e^{3\varepsilon} \leq e^3$ for $\varepsilon \leq 1$. The mean value theorem over $\pi \in [\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]$ concludes. □

Proof of Theorem 7(i). Run the product family of Theorem 5 and account for the oracle by the chain rule: $\mathrm{KL}(P_+ \| P_-) \leq \mathrm{KL}_T + \mathbb{E}_T \mathrm{KL}_{\widehat{W}|T}$. The transcript part is Lemma 4. For the oracle part, the family's signal in $\widehat{W}$ is a coherent mean shift, not a fluctuation-level one: under $\mu_\pm$, $\mathbb{E}W_{qr} = (n-2)(\frac{1}{2} \pm \gamma)^2$, so the per-entry gap is exactly $2(n-2)\gamma$—all n wedge paths shift together. Conditionally on T the edges are independent with posterior parameters $p_\pm(y_e)$; Lemma 9 moves each parameter by $O(\gamma)$ uniformly in y_e, and since $p_\pi(y)$ is increasing in π, an edgewise monotone coupling of the two posterior edge configurations flips each edge with probability $p_+(y_e) - p_-(y_e) = O(\gamma)$. Writing $\Delta_{qr} := W^+_{qr} - W^-_{qr}$ for the coupled difference, joint convexity of KL over the coupling—each conditional law of $\widehat{W}_\sigma$ is a Gaussian mixture, and the coupling bounds the mixture KL by the averaged mean-shift KL—gives $\mathrm{KL}_{\widehat{W}|T} \leq \sum_{q \leq r} \mathbb{E}[\Delta^2_{qr}]/(2\sigma^2)$, with $\mathbb{E}[\Delta^2_{qr}] = O(n^2\gamma^2 + n\gamma)$: the n^2 two-flip cross terms contribute $O(n^2\gamma^2)$ and the n single-flip terms $O(n\gamma)$ (the n diagonal entries of W at lower order); hence $\mathrm{KL}_{\widehat{W}|T} = O\big((n^4\gamma^2 + n^3\gamma)/\sigma^2\big)$. At $\gamma^* = \Theta(1/(\varepsilon n))$ this is $O\big(n^2/(\varepsilon^2\sigma^2)\big)$ (the second term is lower order for $\varepsilon \leq 1$), which is $O(1)$ exactly when $\sigma = \Omega(n/\varepsilon)$; then Lemmas 2–5 conclude as in Theorem 5 (their σ is the count's fluctuation cap $C_\sigma n^{k-1}$ of Lemma 3, not the oracle scale), with Lemma 5 applied to the augmented transcript $T' = (T, \widehat{W}_\sigma)$ (its proof uses only $\mathrm{TV}(P_+, P_-) \leq \frac{1}{4}$ of the observation law). Part (ii)—the spectral-deflation upper bound—is proved in full in Appendix A.7, after the discrete-Laplace transfer lemmas it invokes (Appendix A.6); it also supplies the noise-law generality behind Corollary 1. □

A.5 The Fang–Yi wedge-release interface

The separation consumes prior work through a precise interface, which we now derive from the components of Fang & Yi (2025) rather than cite wholesale.

Lemma 10 (Fang–Yi wedge-release interface)**.** *Let $\varepsilon < 2$ and $\delta \in (0, 1/2)$. Instantiate the Fang–Yi 2-star randomizer R_{star} (Fang & Yi, 2025, Alg. 6) with sampling probability $q = 1$ and degree parameter $d = n - 1$, invoked once for every unordered pair $\{q, r\} \subseteq [n]$ (all $m = n$ groups in the notation of Fang & Yi, 2025, §4.1.3), with per-pair budget $\delta'' = \delta/(4(n-1))$ and*

$$\varepsilon'' = \max\left\{\hat{\varepsilon} : \sqrt{4(n-1)\ln(2/\delta)}\,\hat{\varepsilon} + 2(n-1)\hat{\varepsilon}(e^{\hat{\varepsilon}} - 1) \leq \varepsilon\right\}. \tag{4}$$

Then the analyzer obtains $\widehat{W}_{qr} = W_{qr} + \eta_{qr}$ for every pair, where $W = A^2$, and: (a) $\mathbb{E}[\widehat{W}_{qr}] = W_{qr}$ exactly (no debiasing needed); (b) the η_{qr} are i.i.d. across pairs and independent of the graph, with $\eta_{qr} \sim \mathrm{DLap}(b)$ for $b = 4/(3\varepsilon'') = \Theta\big(\sqrt{n\ln(1/\delta)}/\varepsilon\big) = \widetilde{O}(\sqrt{n}/\varepsilon)$, a public function of (ε, δ, n) alone; (c) the shuffled transcript—hence the released table $\{\widehat{W}_{qr}\}$, by post-processing—is (ε, δ)-edge-DP on all n-node graphs; (d) each user sends $\widetilde{O}(n^2 + n^{3/2}/\varepsilon)$ bits in expectation, which is $\widetilde{O}(n^2)$ whenever $\varepsilon = \Omega(n^{-1/2})$; (e) (full matrix) running (a)–(d) at budget $\varepsilon/2$ (replace ε by $\varepsilon/2$ in (4)) and, with fresh coins, having each user report its degree through the geometric degree release (Fang & Yi, 2025, Lem. 3.1) at budget $\varepsilon/4$, the analyzer sets $\widehat{W}_{qq} = d_q + \eta_q$ with $\eta_q \sim \mathrm{DLap}(\Theta(1/\varepsilon))$ i.i.d., independent of the off-diagonal noise and of the graph; the combined table $\{\widehat{W}_{qr}\}_{q \leq r}$ is entrywise unbiased for W, (ε, δ)-edge-DP on all graphs, and retains the Θ-scales of (b) and (d).

Proof. For $T = \{q, r\}$, $k = 2$, the FY summand is $f_{T,2}(x_i) = x_{iq}x_{ir}\binom{d_i - 2}{0} = A_{qi}A_{ir}$ for $i \notin T$ and 0 otherwise, so $\sum_i f_{T,2}(x_i) = (A^2)_{qr} = W_{qr}$, with sensitivity parameter $\Delta = \binom{d-2}{0} = 1$ (Fang & Yi, 2025, Alg. 6, line 1) for any d. At $q = 1$, Alg. 6 passes $(\varepsilon'', \delta'')$ unchanged to the negative-binomial randomizer (Fang & Yi, 2025, Alg. 3): per pair, user i sends (i) the value $f_{T,2}(x_i)$ if nonzero, (ii) $z^{(i)}_{+1}$ copies of $+1$ and $z^{(i)}_{-1}$ copies of -1 with $z^{(i)}_{\pm 1} \sim \mathrm{NB}_{/n}(1, 1 - e^{-3\varepsilon''/4})$ i.i.d., and (iii) count-masking $(+1, -1)$ pairs that cancel in any sum. Hence $\widehat{W}_{qr} = W_{qr} + \sum_i (z^{(i)}_{+1} -$

$z_{-1}^{(i)}$), and by $\sum_{i\in[n]} \mathrm{NB}_{/n}(r,p) = \mathrm{NB}(r,p)$ together with $\mathrm{NB}(1, 1-e^{-\varepsilon_1}) - \mathrm{NB}(1, 1-e^{-\varepsilon_1}) = \mathrm{DLap}(1/\varepsilon_1)$ (Fang & Yi, 2025, §2.3), we get $\eta_{qr} \sim \mathrm{DLap}(4/(3\varepsilon''))$ exactly. The noise distributions are computed from $(\varepsilon'', \delta'', \Delta, n)$ only and sampled before the data enters (Alg. 3, lines 3–4), so η_{qr} is graph-independent and mean-zero, giving (a). Each pair is a separate invocation of R_{star} with fresh coins and $O(\log n)$-bit pair-tagged messages (Fang & Yi, 2025, proofs of Thms. 4.2, 4.3), so the η_{qr} are independent across pairs, and identically distributed since all pairs share $(\varepsilon'', \delta'', \Delta)$; this is (b). For (c), let $G \sim G'$ differ in edge (a, c). The invocation for $\{q, r\}$ is affected only if $\{q, r\} = \{c, w\}$ with $w \sim a$ or $\{q, r\} = \{a, w\}$ with $w \sim c$: at most $d_a + d_c \le 2(n-1)$ invocations, unconditionally, since $d = n - 1$ upper-bounds every degree of every graph. Within each affected invocation exactly one user's value changes ($f_{T,2} \equiv 0$ on T itself), so it is $(\varepsilon'', \delta'')$-edge shuffle DP (Fang & Yi, 2025, Lem. 2.10; Lem. 3.5, case $k = |T|$), using $\varepsilon'' < 4$. Advanced composition over the $\le 2(n-1)$ affected mechanisms with slack $\delta/2$ costs at most $\sqrt{4(n-1)\ln(2/\delta)}\,\varepsilon'' + 2(n-1)\varepsilon''(e^{\varepsilon''} - 1) \le \varepsilon$ in ε and $2(n-1)\delta'' + \delta/2 = \delta$ in δ; unaffected invocations are identically distributed under G and G'. This is precisely the R_{star} composition step in the proof of Fang & Yi (2025, Thm. 4.1), with their free parameter ε' set to ε. Solving (4) at $\varepsilon < 2$ gives $\varepsilon'' = \Theta\big(\varepsilon/\sqrt{n\ln(1/\delta)}\big)$ (Fang & Yi, 2025, §4.1.2), hence $b = 4/(3\varepsilon'') = \Theta(\sqrt{n\ln(1/\delta)}/\varepsilon)$. For (d), user i sends $\binom{d_i}{2} = O(n^2)$ value messages and $\widetilde{O}(1/(n\varepsilon''))$ expected noise messages per pair (Fang & Yi, 2025, Lem. 2.10); over $\binom{n}{2}$ pairs, with $O(\log n)$-bit tags, this totals $\widetilde{O}(n^2 + n/\varepsilon'') = \widetilde{O}(n^2 + n^{3/2}/\varepsilon)$ bits, matching Fang & Yi (2025, Thm. 4.3) at $q = 1$, $d = n-1$. For (e): the degree query has sensitivity 1 per report, and one edge (a, c) changes exactly the two reports d_a and d_c, so the degree release is $2 \cdot (\varepsilon/4) = \varepsilon/2$-edge-DP by basic composition (pure DP, no δ); basic composition with the $(\varepsilon/2, \delta)$ off-diagonal release gives (ε, δ)-edge-DP. Fresh coins give independence, the geometric release is exactly unbiased, and the extra communication is one $O(\log n)$-bit report per user; the constants in (b) shift under $\varepsilon \mapsto \varepsilon/2$, the Θ-forms do not. □

Remark 8. Run verbatim, Fang & Yi (2025, Thm. 4.1) splits its budget as $\varepsilon' = \varepsilon/2$ between an adjacency-matrix release (R_{adj}, unused by our interface) and the wedge invocations; the wedge sub-release alone then satisfies $(\varepsilon/2, \delta)$-edge-DP. Lemma 10 reallocates the full budget by setting $\varepsilon' := \varepsilon$ in (4); the composition step quoted above never uses $\varepsilon' = \varepsilon/2$, so this is a re-parameterization, not a new proof. Constants in b shift by the corresponding factor; the Θ-form is unchanged. Fresh independent coins per invocation, though not spelled out in their algorithm box, are asserted in their variance proofs ("the estimate for any two sets of nodes is independent") and are forced by their own Theorem 4.5, whose unbiasedness for 4-cycles fails without them.

A.6 From Gaussian to discrete Laplace: transfer lemmas for Corollary 1

By Lemma 10, the analyzer holds $W_{qr} + \eta_{qr}$ with i.i.d., graph-independent $\mathrm{DLap}(b)$ noise of public scale $b = \widetilde{O}(\sqrt{n}/\varepsilon)$; the degree release (Fang & Yi, 2025, Lem. 3.1) supplies the diagonal $W_{qq} = d_q$ up to $\mathrm{DLap}(O(1/\varepsilon))$ noise. The proof of Theorem 7(ii) in Appendix A.7 is carried out directly under the sub-exponential model (N*), which contains the Gaussian oracle ($\sigma_d = \sigma$) as a special case; the three routine lemmas below record the facts it uses for this release.

Lemma 11 (Operator norm for heterogeneous wedge noise)**.** *Let $E = H + Z$, where H is hollow symmetric with independent centered sub-exponential off-diagonal entries of variance σ^2, and $Z = \mathrm{diag}(\zeta_i)$ has independent centered sub-exponential entries of standard deviation at most $\sigma_d \le C\sigma$. For every fixed $c' > 0$,*

$$\Pr\big(\|E\|_{\mathrm{op}} > C_{c'}(\sigma\sqrt{n} + \sigma_d)\log n\big) \le n^{-c'}.$$

This includes the Gaussian oracle and the Fang–Yi off-diagonal DLap release with its independent degree-noise diagonal.

Proof sketch. Truncate and recenter the off-diagonal entries at $O(\sigma \log n)$ and apply the Füredi–Komlós trace-moment bound to H. A union bound gives $\|Z\|_{\mathrm{op}} = \max_i |\zeta_i| = O(\sigma_d \log n)$. No Frobenius-norm shortcut is used for the spectral gap. □

Lemma 12 (Analyzer-computable debiasing)**.** *All constants in the deflation pipeline are public. If $\tau_q^2 = \mathrm{Var}(E_{qq})$, then*

$$t_q := \mathbb{E}[(E^2)_{qq}] = (n-1)\sigma^2 + \tau_q^2$$

is used in the noise-trace subtraction. The off-diagonal scale σ, the diagonal variances τ_q^2, the Perron second-order subtraction, and the randomized-response variance are determined by the published mechanism parameters. Heterogeneous diagonal contributions are carried explicitly through σ_d rather than identified with degree noise.

Lemma 13 (High-probability to RMSE). *Let $\widehat{N}'$ satisfy $|\widehat{N}' - N_5| \le R$ with probability $1 - n^{-13}$ under $\mu_\pm$, and let $\widehat{N} := \min\{\max\{\widehat{N}', 0\}, n^5\}$. Then $\mathbb{E}_{\mu_\pm}[(\widehat{N} - N_5)^2]^{1/2} \le R + n^{-1.5}$.*

Proof. $N_5 \in [0, n^5]$ always, so $|\widehat{N} - N_5| \le n^5$ on the bad event and $\le |\widehat{N}' - N_5| \le R$ otherwise; hence $\mathbb{E}[(\widehat{N} - N_5)^2] \le R^2 + n^{10} \cdot n^{-13}$. □

A.7 Full proof of Theorem 7(ii)

Standing assumptions. (N*) Write $E = H + Z$, where H is symmetric and hollow and $Z = \mathrm{diag}(\zeta_1, \dots, \zeta_n)$. The entries $\{H_{qr}\}_{q<r}$ are mutually independent, centered, independent of the graph, have public variance σ^2, and satisfy $\Pr(|H_{qr}| \ge t\sigma) \le 2e^{-t/K_0}$. The variables ζ_q are mutually independent, centered, independent of H and of the graph, with public variances $\tau_q^2 \le C_0\sigma_d^2$ and tails $\Pr(|\zeta_q| \ge t\sigma_d) \le 2e^{-t/K_0}$. The diagonal scale satisfies $\sigma_d \le C_d\sigma$. Finally E is independent of the randomized-response noise M. The Gaussian oracle has $\sigma_d = \sigma$; the Fang–Yi instantiation has $\sigma = \widetilde{O}(\sqrt{n}/\varepsilon)$ and $\sigma_d = \Theta(1/\varepsilon)$. (F) $A \sim \mu_\pm = \mathrm{Bern}(\frac{1}{2} \pm \gamma^*)^{\otimes\binom{n}{2}}$, $\gamma^* = \Theta(1/(\varepsilon n)) = o(1)$; $W = A^2$; $\lambda_1 \ge \cdots \ge \lambda_n$ are the eigenvalues of A, $\{v_i\}$ is a common orthonormal eigenbasis of A and W, and μ_i are the eigenvalues of W in decreasing order. (R) $\sigma \le cn/\varepsilon$ and $\varepsilon \in [n^{-1/4}, \varepsilon_0]$, so $\sigma \le cn^{5/4}$.

Lemma 14 (Family spectral profile). *There are constants such that with probability $\ge 1 - n^{-14}$ under $\mu_\pm$ the event $\mathcal{B}$ holds, on which: (a) all degrees $d_i \in [0.49n, 0.51n]$, and every two distinct vertices have at least $n/5$ common neighbours; (b) A is connected and non-bipartite, so λ_1 is simple, $v_1 > 0$ entrywise, and $Av_1 = +\sqrt{\mu_1}v_1$ with $\mu_1 = \lambda_1^2$ the (simple) top eigenvalue of W; (c) $\lambda_1 = \Theta(n)$ and $\max_{i\ge2}|\lambda_i| \le C_A\sqrt{n}$, hence $\mu_1 = \Theta(n^2)$, $\mu_2 \le C_A^2 n$, and the spectral gap $g := \mu_1 - \mu_2 \ge n^2/8$; (d) $\|v_1\|_\infty \le C/\sqrt{n}$; (e) $c'n^3 \le \|A^3\|_F \le Cn^3$, $\|PA^3P\|_F \le Cn^2$, $\|PAP\|_{\mathrm{op}} \le C_A\sqrt{n}$, and $\|\delta\|_F \le Cn^{2.5}$ for $\delta := W^2 - \mu_1^2\Pi_1$; (f) $\sum_i (d_i - \frac{n}{2})^2 \le C(n^2 \log n + n/\varepsilon^2) \le Cn^2 \log n \cdot (1 + n^{-1/2})$; (g) $\|PWP\|_F \le Cn^{1.5}$ and $\mathrm{tr}\, W = 2|E(G)| \le n^2$.*

Proof. (a): Chernoff for $d_i \sim \mathrm{Bin}(n-1, \frac{1}{2} \pm \gamma^*)$ plus a union bound, failure $ne^{-\Omega(n)}$; likewise, for each fixed pair $u \ne v$ the $n-2$ joint-adjacency indicators are independent under the product prior, so $|N(u) \cap N(v)| \sim \mathrm{Bin}\big(n-2, (\frac{1}{2} \pm \gamma^*)^2\big)$ and Chernoff with a union bound over the $\binom{n}{2}$ pairs gives $\min_{u\ne v}|N(u) \cap N(v)| \ge n/5$, failure $n^2e^{-\Omega(n)}$. (The min-degree event alone would not suffice here: $K_{n/2,n/2}$ has min degree $n/2$ and is bipartite.) (b): by (a), any two vertices are at distance ≤ 2, so A is connected; any present edge uv (one exists: degrees are positive by (a)) together with a common neighbour w of u and v forms a triangle, so A is non-bipartite; Perron–Frobenius gives simplicity and positivity for A, and $W = A^2$ has top eigenvalue λ_1^2 with the same eigenvector, simple on the event of (c) since $\max_{i\ge2}\lambda_i^2 < \lambda_1^2$. (c): write $A = \bar{A} + (A - \bar{A})$, $\bar{A} := \mathbb{E}A = (\frac{1}{2} \pm \gamma^*)(J - I)$, a rank-one matrix plus $O(1)$-norm shift; $\|A - \bar{A}\|_{\mathrm{op}} \le C\sqrt{n}$ with failure $e^{-\Omega(n)}$ by the bounded-entry matrix concentration used in Lemma 11 (Füredi–Komlós; entries are centered, independent, bounded by 1). Weyl then pins $\lambda_1 = (\frac{1}{2} \pm \gamma^*)n + O(\sqrt{n}) = \Theta(n)$ and $|\lambda_i| \le C_A\sqrt{n}$ for $i \ge 2$; thus $\mu_1 = \Theta(n^2)$, $\mu_2 = \max(\lambda_2^2, \lambda_n^2) \le C_A^2 n$, and $g \ge n^2/8$ for n large. (d): Davis–Kahan against $u := \mathbf{1}/\sqrt{n}$, the top eigenvector of $\bar{A} + \frac{1}{2}I$: $\|v_1 - u\| \le 2\|A - \bar{A}\|_{\mathrm{op}}/\,\mathrm{gap}(\bar{A}) = O(n^{-1/2})$, so $\|v_1\|_\infty \le \|u\|_\infty + \|v_1 - u\| = O(n^{-1/2})$. (e): $\|A^3\|_F \le \|A\|_{\mathrm{op}}^2\|A\|_F = O(n^2 \cdot n)$ and $\|A^3\|_F \ge \lambda_1^3 = \Omega(n^3)$; $\|PA^3P\|_F^2 = \sum_{i\ge2}\lambda_i^6 \le n \cdot C^6n^3 = O(n^4)$; $\|PAP\|_{\mathrm{op}} = \max_{i\ge2}|\lambda_i|$; $\|\delta\|_F^2 = \sum_{i\ge2}\mu_i^4 \le n \cdot C^8n^4 = O(n^5)$. (f): $d_i - \mathbb{E}d_i$ is a sum of $n-1$ independent $\pm$ bits: Hoeffding gives $(d_i - \mathbb{E}d_i)^2 \le Cn\log n$ for all i w.p. $1 - n^{-15}$; and $|\mathbb{E}d_i - \frac{n}{2}| = O(1 + n\gamma^*) = O(1/\varepsilon)$, so $\sum_i (d_i - \frac{n}{2})^2 \le 2Cn^2 \log n + 2n \cdot O(1/\varepsilon^2)$, and $n/\varepsilon^2 \le n^{3/2}$ for $\varepsilon \ge n^{-1/4}$. (g): $\|PWP\|_F^2 = \sum_{i\ge2}\mu_i^2 \le n \cdot C^4n^2$. □

Lemma 15 (Truncation). *Let C_T be a constant larger than $30K_0 \max\{1, C_d\}$ and set*

$$\mathcal{T} := \left\{ \max_{q<r} |H_{qr}| \le C_T \sigma \log n, \quad \max_q |\zeta_q| \le C_T \sigma \log n \right\}.$$

Then $\Pr(\mathcal{T}^c) \le 2n^{-28}$. Truncating and recentering the active noise entries changes every mean used below by $O(n^{-10})$.

Proof. Use the two tail bounds and a union bound over $O(n^2)$ entries. The tail mean of each entry is $O(\sigma n^{-30})$ because $\sigma_d \le C_d\sigma$. The coefficient multiplying one noise entry in any displayed mean below is at most n^{12}, so the total truncation shift is $O(n^2 n^{12} \sigma n^{-30}) = O(n^{-10})$ under (R). □

Lemma 16 (Conditional linear forms). *Let $\mathcal{F}$ be independent of the active wedge-noise block and let C be a symmetric, $\mathcal{F}$-measurable matrix. Conditionally on $\mathcal{F}$,*

$$\mathrm{Var}(\langle E, C\rangle \mid \mathcal{F}) \le C\big(\sigma^2 \|C\|_F^2 + \sigma_d^2 \|\operatorname{diag} C\|_2^2\big) \le C'\sigma^2 \|C\|_F^2,$$

and, on $\mathcal{T}$, with conditional probability at least $1 - n^{-15}$,

$$|\langle E, C\rangle| \le C'\sigma \log^2 n \, \|C\|_F.$$

If a is a coefficient vector independent of Z, then similarly

$$\left| \sum_i a_i \zeta_i \right| \le C\sigma_d \log^2 n \, \|a\|_2.$$

Proof. Condition on $\mathcal{F}$. Each display is a sum of independent centered sub-exponential variables. The variance bound is direct, and Bernstein's inequality gives the high-probability bounds. □

Lemma 17 (Image of a fixed vector). *For a fixed unit vector v measurable with respect to the graph,*

$$\mathbb{E}\|Ev\|^2 = (n-1)\sigma^2 + \sum_q \tau_q^2 v_q^2 = O(n\sigma^2),$$

and on $\mathcal{T}$, with probability at least $1 - n^{-14}$, $\|Ev\| \le C_2\sigma\sqrt{n}\log^2 n$.

Proof. The qth coordinate has variance $\sum_{r\neq q} \sigma^2 v_r^2 + \tau_q^2 v_q^2$. Sum over q and use $\sigma_d \le C_d\sigma$. The coordinate bounds follow from Lemma 16 and a union bound. □

Lemma 18 (Operator norm; heterogeneous interface). *With probability at least $1 - n^{-14}$,*

$$\|E\|_{\mathrm{op}} \le C_3(\sigma\sqrt{n} + \sigma_d)\log n = \widetilde{O}(\sigma\sqrt{n}).$$

Consequently, on $\mathcal{B}$, $\|E\|_{\mathrm{op}} \le g/8$ throughout (R), for c sufficiently small and n sufficiently large.

Proof. Apply the truncated Füredi–Komlós trace-moment bound to the hollow matrix H and the entrywise tail bound to $\|Z\|_{\mathrm{op}} = \max_i |\zeta_i|$. This is the only spectral-gap argument; no Frobenius shortcut is used. □

Lemma 19 (Conditional quadratic forms). *Let $\mathcal{F}$ be independent of E and let B_1, B_2 be symmetric, $\mathcal{F}$-measurable matrices. On $\mathcal{T}$, conditionally on $\mathcal{F}$, with probability at least $1 - n^{-15}$,*

$$|\operatorname{tr}(EB_1EB_2) - \mathbb{E}[\operatorname{tr}(EB_1EB_2) \mid \mathcal{F}]| \le C_4\sigma^2 \log^4 n \, \|B_1\|_F \|B_2\|_F.$$

Moreover

$$\mathbb{E}[\operatorname{tr}(EB_1EB_2) \mid \mathcal{F}] = \sigma^2\big(\operatorname{tr} B_1 \operatorname{tr} B_2 + \operatorname{tr}(B_1B_2)\big) + \theta,$$

where

$$\theta = \sum_q (\tau_q^2 - 2\sigma^2)(B_1)_{qq}(B_2)_{qq}, \qquad |\theta| \le C\sigma^2 \|B_1\|_F \|B_2\|_F.$$

Proof. After truncation, the independent upper-triangular entries form a sub-Gaussian vector and the trace is a quadratic form whose coefficient Frobenius norm is $O(\|B_1\|_F\|B_2\|_F)$. Hanson–Wright gives the absolute-value deviation. The mean follows from the covariance identity for a symmetric matrix with off-diagonal variance σ^2 and diagonal variance τ_q^2. □

Definition 5 (Blockwise good event)**.** *Let $\mathcal{G}$ be the intersection of $\mathcal{B}$, $\mathcal{T}$, the operator-norm event, and the $O(1)$ conditional concentration events invoked below. Whenever a coefficient contains one noise block, we condition on that block and apply Lemma 16 or 19 to an independent active block; readouts against M are conditioned on (H, Z). The tower property and a finite union bound give* $\Pr(\mathcal{G}) \geq 1 - n^{-13}$.

Inputs. The analyzer holds the wedge oracle $\widehat{W}_\sigma = W + E = W + H + Z$ under (N*), with $\hat{d}_i = (\widehat{W}_\sigma)_{ii} = d_i + \zeta_i$. It also holds the debiased blind randomized-response release $\widehat{A} = A + M$: M is symmetric and hollow, its upper-triangular entries are independent and centered with $|M_{jk}| \leq C_M/\varepsilon$ and variance $v_A = \Theta(1/\varepsilon^2)$, and M is independent of (H, Z).

Spectral-deflation estimator.

(S1) Top eigenpair. Compute the largest eigenvalue $\hat{\mu}_1$ of $\widehat{W}_\sigma$ and a unit eigenvector $\hat{v}_1$ (sign fixed by $\mathbf{1}^\top \hat{v}_1 > 0$); set $\widehat{\Pi}_1 := \hat{v}_1 \hat{v}_1^\top$, $\widehat{P} := I - \widehat{\Pi}_1$.

(S2) Second-order Perron mean subtraction (eigenvalue). $\hat{\mu}_1^c := \hat{\mu}_1 - (n-1)\sigma^2/\hat{\mu}_1$.

(S3) Noise-trace subtraction. $\widehat{W^2} := \widehat{W}_\sigma^2 - \mathrm{diag}(t_1, \ldots, t_n)$ with the public constants $t_q := (n-1)\sigma^2 + \tau_q^2 = \mathbb{E}[(E^2)_{qq}]$.

(S4) Deflated coefficient. $\hat{\delta} := \widehat{W^2} - (\hat{\mu}_1^c)^2 \widehat{\Pi}_1$.

(S5) Closed-walk estimate with projector-tilt mean subtraction.

$$\widehat{F} := (\hat{\mu}_1^c)^{5/2} - (n-1)\sigma^2(\hat{\mu}_1^c)^{1/2} + \langle \hat{\delta}, \widehat{A} \rangle,$$

the middle term being the analyzer-computable mean of the eigenvector-tilt bias $\mu_1^{5/2}\beta^2$ (Lemma 22); the readout $\langle \cdot, \widehat{A} \rangle$ touches off-diagonal entries only, matching $A_{jj} = 0$.

(S6) Degeneracy corrections (Perron-block readout of $D_5 = \langle 5W + 5B, A \rangle$, $B_{ij} = \frac{1}{2}(d_i + d_j - 4)W_{ij}$; Lemma 6). With $\widehat{D}' := \mathrm{diag}(\hat{d}_i - 2)$ and $\widehat{B} := \frac{1}{2}(\widehat{D}'\widehat{W}_\sigma + \widehat{W}_\sigma \widehat{D}')$, set

$$\widehat{D}_5 := 5\Big[\underbrace{(\hat{\mu}_1^c)^{3/2}}_{\approx \sqrt{\mu_1}\, v_1^\top W v_1} + \langle \widehat{P}\widehat{W}_\sigma \widehat{P}, \widehat{A} \rangle \Big] + 5\Big[\underbrace{(\hat{\mu}_1^c)^{3/2} \sum_i (\hat{d}_i - 2)\hat{v}_{1,i}^2}_{\approx \sqrt{\mu_1}\, v_1^\top B v_1} + \langle \widehat{P}\widehat{B}\widehat{P}, \widehat{A} \rangle \Big].$$

Here we used $v_1^\top W v_1 = \mu_1$ and $v_1^\top B v_1 = \mu_1 v_1^\top D' v_1$ (see the exact split before Lemma 27); the split $\langle X, A \rangle = \sqrt{\mu_1}\, v_1^\top X v_1 + \langle PXP, A \rangle$ is exact for any symmetric X because $A = \sqrt{\mu_1}\Pi_1 + PAP$ on $\mathcal{B}$ and Π_1 commutes with A.

(S7) Assembly and clipping (Lemma 13). $\widehat{N}_5' := \frac{1}{10}(\widehat{F} - \widehat{D}_5)$, $\widehat{N}_5 := \min\{\max\{\widehat{N}_5', 0\}, n^5\}$.

Remarks on well-definedness and publicness. Every constant used—σ, σ_d, τ_q^2, t_q, $(n-1)\sigma^2$, and v_A—is a public function of the mechanism parameters; nothing uses the unknown μ_1, v_1, or A. On $\mathcal{G}$, $\|E\|_{\mathrm{op}} \leq g/8$ (Lemma 18) so $\hat{\mu}_1$ is simple by Weyl ($\hat{\mu}_1 - \hat{\mu}_2 \geq g - 2\|E\|_{\mathrm{op}} \geq 3g/4 > 0$) and $\hat{v}_1$ is well-defined up to the fixed sign; also $\hat{\mu}_1 \geq \mu_1 - g/8 > 0$, so (S2), (S5), (S6) are well-defined ($\hat{\mu}_1^c \geq \hat{\mu}_1 - o(\hat{\mu}_1) > 0$ since $(n-1)\sigma^2/\hat{\mu}_1^2 = O(n\sigma^2/n^4) = o(1)$). Off $\mathcal{G}$ the estimate is arbitrary; clipping (S7) caps its MSE contribution (Lemma 13). Throughout Lemmas 20–28 we work conditionally on a graph A in $\mathcal{B}$; all noise probabilities are over (H, Z, M), which are jointly independent of A.

Lemma 20 (Deterministic eigenpair expansion)**.** *Condition on $A \in \mathcal{B}$ and suppose $\|E\|_{\mathrm{op}} \leq g/8$. Write $S := \sum_{i \geq 2} \frac{v_i v_i^\top}{\mu_1 - \mu_i}$, $e_{1j} := v_j^\top E v_1$. Then $\hat{\mu}_1$ is simple, and the top unit eigenvector can be written $\hat{v}_1 = (v_1 + w)/\sqrt{1 + \|w\|^2}$ with $w \perp v_1$ and: (i) $\|w\| \leq 2\|Ev_1\|/g$; (ii) $\|w - SEv_1\| \leq 3\|E\|_{\mathrm{op}}\|Ev_1\|/g^2$; (iii) $\hat{\mu}_1 = \mu_1 + e_{11} + v_1^\top ESEv_1 + R_3$ with $|R_3| \leq 3\|E\|_{\mathrm{op}}\|Ev_1\|^2/g^2$.*

Proof. Weyl gives $|\hat{\mu}_1 - \mu_1| \le \|E\|_{\mathrm{op}} \le g/8$ and $\hat{\mu}_2 \le \mu_2 + g/8$, so $\hat{\mu}_1 - \hat{\mu}_2 \ge 3g/4 > 0$: simple. Let $\hat{v}$ be a top unit eigenvector, $P := I - v_1 v_1^\top$. Projecting $(W+E)\hat{v} = \hat{\mu}_1\hat{v}$ onto $v_1^\perp$ (which W preserves): $(\hat{\mu}_1 - W)|_{v_1^\perp}(P\hat{v}) = PE\hat{v}$. For $i \ge 2$, $\hat{\mu}_1 - \mu_i \ge \mu_1 - g/8 - \mu_2 = 7g/8$, so $R_{\hat{\mu}} := \big((\hat{\mu}_1 - W)|_{v_1^\perp}\big)^{-1} P$ exists with $\|R_{\hat{\mu}}\|_{\mathrm{op}} \le 8/(7g)$, and $\|P\hat{v}\| \le \frac{8}{7g}\|E\|_{\mathrm{op}} \le \frac{1}{7} < 1$; hence $\alpha := v_1^\top \hat{v} \neq 0$, and after a sign flip $\alpha > 0$; set $w := P\hat{v}/\alpha$, so $\hat{v} = (v_1 + w)/\sqrt{1+\|w\|^2}$. The eigen-equation for $v_1 + w$ projected onto $v_1^\perp$ reads $w = R_{\hat{\mu}} PE(v_1 + w)$, whence $\|w\| \le \frac{8}{7g}(\|Ev_1\| + \|E\|_{\mathrm{op}}\|w\|)$ and, using $\|E\|_{\mathrm{op}} \le g/8$, $\|w\| \le \frac{4}{3g}\|Ev_1\| \le 2\|Ev_1\|/g$: (i). Iterating once, $w - SEv_1 = (R_{\hat{\mu}} - S)PEv_1 + R_{\hat{\mu}} PEw$; on $v_1^\perp$ the operator $R_{\hat{\mu}} - S$ has eigenvalues $\frac{\mu_1 - \hat{\mu}_1}{(\hat{\mu}_1 - \mu_i)(\mu_1 - \mu_i)}$, of magnitude $\le \frac{8\|E\|_{\mathrm{op}}}{7g^2}$, so

$$\begin{aligned}\|w - SEv_1\| &\le \frac{8}{7g^2}\|E\|_{\mathrm{op}}\|Ev_1\| + \frac{8}{7g}\|E\|_{\mathrm{op}} \cdot \frac{4}{3g}\|Ev_1\| \\ &\le 3\|E\|_{\mathrm{op}}\|Ev_1\|/g^2 :\end{aligned}$$

(ii). Projecting the eigen-equation onto v_1 gives exactly $\hat{\mu}_1 = \mu_1 + e_{11} + v_1^\top Ew$ (as $v_1^\top W w = \mu_1 v_1^\top w = 0$); then $v_1^\top Ew = v_1^\top ESEv_1 + v_1^\top E(w - SEv_1)$ and (ii) give (iii). On $\mathcal{B}$ the estimator's sign convention $\mathbf{1}^\top \hat{v}_1 > 0$ selects this $\hat{v}$: $\mathbf{1}^\top v_1 = \Theta(\sqrt{n})$ by Lemma 14(d) with positivity, while $|\mathbf{1}^\top w| \le \sqrt{n}\|w\| = o(\sqrt{n})$. □

Lemma 21 (Eigenvalue recovery). *On $\mathcal{G}$: $\Delta := \hat{\mu}_1^c - \mu_1$ satisfies $\Delta = e_{11} + Z_\Delta$ with*

$$\begin{aligned}|Z_\Delta| &\le C\log^4 n\big(\sigma^2/n^{1.5} + \sigma^3/n^{2.5}\big) =: \rho_2, \\ |e_{11}| &\le C_1\sigma\log^2 n, \qquad |\Delta| \le C\sigma\log^4 n.\end{aligned}$$

Proof. By Lemma 20(iii), $\hat{\mu}_1 = \mu_1 + e_{11} + Q_S + R_3$, $Q_S := v_1^\top ESEv_1$. Mean of Q_S: Lemma 19 with $B_1 = v_1 v_1^\top$, $B_2 = S$ gives $\mathbb{E}Q_S = \sigma^2[\mathrm{tr}\, S + v_1^\top S v_1] + \theta = \sigma^2 \mathrm{tr}\, S + \theta$ (as $Sv_1 = 0$), $|\theta| \le 3C_0\sigma^2 \max_q |S_{qq}| \le 4C_0\sigma^2/g = O(\sigma^2/n^2)$. Moreover $\mathrm{tr}\, S = \frac{n-1}{\mu_1} + r$ with $0 \le r = \sum_{i\ge 2} \frac{\mu_i}{\mu_1(\mu_1 - \mu_i)} \le \frac{\mathrm{tr}\, W}{\mu_1 g} \le O(n^{-2})$ (all $\mu_i \ge 0$ since $W = A^2 \succeq 0$; $\mathrm{tr}\, W \le n^2$). Fluctuation: Lemma 19 with $\|B_1\|_F = 1$, $\|S\|_F \le \sqrt{n} \cdot \frac{8}{7g} = O(\sqrt{n}/n^2)$ gives $|Q_S - \mathbb{E}Q_S| \le C_4\sigma^2\log^4 n \cdot O(n^{-1.5})$. Remainder: $|R_3| \le 3\|E\|_{\mathrm{op}}\|Ev_1\|^2/g^2 \le C\log^5 n\, \sigma^3 n^{1.5}/n^4 = \widetilde{O}(\sigma^3/n^{2.5})$ by Lemmas 17, 18. Subtraction: $\hat{\mu}_1^c = \hat{\mu}_1 - (n-1)\sigma^2/\hat{\mu}_1$ and

$$\begin{aligned}\left|\frac{(n-1)\sigma^2}{\hat{\mu}_1} - \frac{(n-1)\sigma^2}{\mu_1}\right| &\le \frac{n\sigma^2|\hat{\mu}_1 - \mu_1|}{\mu_1(\mu_1 - g/8)} \\ &\le C\log^2 n \cdot n\sigma^3\sqrt{n}/n^4 = \widetilde{O}(\sigma^3/n^{2.5}).\end{aligned}$$

Collecting: $Z_\Delta = \sigma^2 r + \theta + (Q_S - \mathbb{E}Q_S) + R_3 +$ (subtraction error), of the stated size; note $\sigma^2/n^2 \le \sigma^2/n^{1.5}$. Finally $\rho_2 \le \sigma\log^4 n\,(\sigma/n^{1.5} + \sigma^2/n^{2.5}) \le C\sigma\log^4 n$ by (R), giving $|\Delta| \le |e_{11}| + \rho_2 \le C\sigma\log^4 n$; the leading fluctuation is e_{11} (which is $\Theta(\sigma)$ in second moment, with no log), while $\rho_2 = \widetilde{O}(\sigma \cdot n^{-\Omega(1)})$ is a genuinely lower-order correction (smaller than e_{11} by a power of n, not merely a log). □

Lemma 22 (Eigenvector tilt). *On $\mathcal{G}$, write $\beta^2 := 1 - (v_1^\top \hat{v}_1)^2 = \|w\|^2/(1+\|w\|^2)$. Then*

$$\beta^2 = \frac{(n-1)\sigma^2}{\mu_1^2} + Z_\beta, \qquad |Z_\beta| \le C\log^5 n\left(\frac{\sigma^2}{n^{3.5}} + \frac{\sigma^3}{n^{4.5}} + \frac{\sigma^4}{n^6}\right),$$

$$\textit{and } \beta^2 \le C\log^4 n\,\frac{\sigma^2 n}{g^2} = \widetilde{O}(\sigma^2/n^3).$$

Proof. $Q_2 := \|SEv_1\|^2 = v_1^\top ES^2 Ev_1$. Lemma 19 ($B_1 = v_1 v_1^\top$, $B_2 = S^2$, $\|S^2\|_F \le \sqrt{n}(8/7g)^2 = O(\sqrt{n}/n^4)$): $\mathbb{E}Q_2 = \sigma^2\,\mathrm{tr}(S^2) + \sigma^2\|Sv_1\|^2 + \theta_2 = \sigma^2\,\mathrm{tr}(S^2) + \theta_2$, $|\theta_2| \le 3C_0\sigma^2\|S^2\|_{\mathrm{op}} = O(\sigma^2/n^4)$, and $|Q_2 - \mathbb{E}Q_2| \le C_4\sigma^2\log^4 n \cdot O(n^{-3.5})$. Also $\mathrm{tr}(S^2) = \frac{n-1}{\mu_1^2} + r_2$,

$$0 \le r_2 = \sum_{i\ge 2}\left(\frac{1}{(\mu_1 - \mu_i)^2} - \frac{1}{\mu_1^2}\right) \le \frac{2\,\mathrm{tr}\, W}{\mu_1(7g/8)^2} \cdot C = O(n^{-4}).$$

Next, $\|w\|^2 - Q_2 = 2\langle SEv_1, w - SEv_1\rangle + \|w - SEv_1\|^2$, bounded via Lemma 20(i)–(ii) and Lemmas 17, 18 by

$$C(\|Ev_1\|/g)(\|E\|\|Ev_1\|/g^2) + C(\|E\|\|Ev_1\|/g^2)^2$$
$$= \widetilde{O}(\sigma^3 n/n^{5.5}) + \widetilde{O}(\sigma^4 n^2/n^8) = \widetilde{O}(\sigma^3/n^{4.5} + \sigma^4/n^6).$$

Finally $|\beta^2 - \|w\|^2| \le \|w\|^4 \le (2\|Ev_1\|/g)^4 = \widetilde{O}(\sigma^4 n^2/n^8)$. Sum the pieces. The crude bound is Lemma 20(i) squared. □

Lemma 23 (Perron aggregate)**.** *Define* $\mathcal{P} := (\hat{\mu}_1^c)^{5/2} - (n-1)\sigma^2(\hat{\mu}_1^c)^{1/2} - (\hat{\mu}_1^c)^2 \hat{v}_1^\top A \hat{v}_1$ *(the totality of Perron-carrying terms in* $\widehat{F} - F$*). On* $\mathcal{G}$*,*

$$\left|\mathcal{P} - \frac{1}{2}\mu_1^{3/2} e_{11}\right| \le C \log^5 n \big(n^{1.5}\sigma^2 + n^{0.5}\sigma^3\big) \le C' c \log^5 n \cdot n^3 \sigma,$$

so $\mu_1^{5/2}$ *is recovered to* $\widetilde{O}(n^3\sigma)$*:* $|\mathcal{P}| \le C \log^5 n \, n^3 \sigma$.

Proof. Decompose $\hat{v}_1^\top A \hat{v}_1 = (1-\beta^2)\sqrt{\mu_1} + \beta^2 \xi$ with $\xi := u^\top A u$, $u := w/\|w\|$, $|\xi| \le \|PAP\|_{\mathrm{op}} \le C_A\sqrt{n}$ (Lemma 14(e); $u \perp v_1$). Then

$$\mathcal{P} = (\hat{\mu}_1^c)^2\big[((\hat{\mu}_1^c)^{1/2} - \sqrt{\mu_1}) + \beta^2\sqrt{\mu_1} - \beta^2\xi\big] - (n-1)\sigma^2(\hat{\mu}_1^c)^{1/2}.$$

On $\mathcal{G}$, $\hat{\mu}_1^c \in [\mu_1/2, 2\mu_1]$ and $(\hat{\mu}_1^c)^{1/2} - \sqrt{\mu_1} = \frac{\Delta}{2\sqrt{\mu_1}} + R_{1/2}$, $|R_{1/2}| \le C\Delta^2/\mu_1^{3/2}$. Writing $(\hat{\mu}_1^c)^2 = \mu_1^2 + E_2$, $|E_2| \le 3\mu_1|\Delta|$:

(1) $(\hat{\mu}_1^c)^2 \frac{\Delta}{2\sqrt{\mu_1}} = \frac{1}{2}\mu_1^{3/2}\Delta + O(\sqrt{\mu_1}\Delta^2) = \frac{1}{2}\mu_1^{3/2} e_{11} + \frac{1}{2}\mu_1^{3/2} Z_\Delta + \widetilde{O}(n\sigma^2)$, and $\mu_1^{3/2}|Z_\Delta| \le C\log^4 n(n^{1.5}\sigma^2 + n^{0.5}\sigma^3)$ by Lemma 21;

(2) $(\hat{\mu}_1^c)^2 R_{1/2} = \widetilde{O}(\sqrt{\mu_1}\,\sigma^2) = \widetilde{O}(n\sigma^2)$;

(3)
$$(\hat{\mu}_1^c)^2\beta^2\sqrt{\mu_1} - (n-1)\sigma^2(\hat{\mu}_1^c)^{1/2}$$
$$= \mu_1^{5/2}\beta^2 - (n-1)\sigma^2\sqrt{\mu_1} + E_2\beta^2\sqrt{\mu_1} - (n-1)\sigma^2[(\hat{\mu}_1^c)^{1/2} - \sqrt{\mu_1}];$$

here $\mu_1^{5/2}\beta^2 - (n-1)\sigma^2\sqrt{\mu_1} = \mu_1^{5/2} Z_\beta + (n-1)\sigma^2\big(\mu_1^{5/2}/\mu_1^2 - \sqrt{\mu_1}\big) = \mu_1^{5/2} Z_\beta$ exactly, with $\mu_1^{5/2}|Z_\beta| \le C\log^5 n(n^{1.5}\sigma^2 + n^{0.5}\sigma^3 + \sigma^4/n)$ (Lemma 22), while $|E_2\beta^2\sqrt{\mu_1}| = \widetilde{O}(\sigma^3)$ and $(n-1)\sigma^2|(\hat{\mu}_1^c)^{1/2} - \sqrt{\mu_1}| = \widetilde{O}(n\sigma^2 \cdot \sigma/n) = \widetilde{O}(\sigma^3)$;

(4) $|(\hat{\mu}_1^c)^2\beta^2\xi| \le 2\mu_1^2 \cdot \widetilde{O}(\sigma^2/n^3) \cdot C_A\sqrt{n} = \widetilde{O}(n^{1.5}\sigma^2)$. Collecting, $|\mathcal{P} - \frac{1}{2}\mu_1^{3/2} e_{11}| \le C\log^5 n(n^{1.5}\sigma^2 + n^{0.5}\sigma^3 + n\sigma^2 + \sigma^3 + \sigma^4/n)$, and the first two terms dominate. Domination by $n^3\sigma$: $n^{1.5}\sigma^2 \le n^3\sigma$ iff $\sigma \le n^{1.5}$ (true: $\sigma \le cn^{5/4}$); $n^{0.5}\sigma^3 \le c^2 n^3\sigma$ since $\sigma^2 \le c^2 n^{2.5}$. With $|e_{11}| \le C_1\sigma\log^2 n$ and $\mu_1^{3/2} = O(n^3)$ the final display follows. □

Augmented good event. Let $\mathcal{G}_M$ be the intersection of the following, each holding with probability $\ge 1 - n^{-13}$ over the RR noise M (independent of A, H, Z; entries $|M_{jk}| \le C_M/\varepsilon$, $\mathrm{Var}(M_{jk}) = v_A = \Theta(1/\varepsilon^2)$, mean zero, independent for $j < k$, zero diagonal): (M1) $\|M\|_F \le C_5 n/\varepsilon$; (M2) $\|M\|_{\mathrm{op}} \le C_5\sqrt{n}\log n/\varepsilon$ (Lemma 11 applied to M: bounded ⇒ sub-exponential); (M3)

$$\mathrm{tr}(W^2M^2) \le 2(n-1)v_A\,\mathrm{tr}(W^2) + Cn^{4.5}\log^4 n/\varepsilon^2 \le C_5 n^5/\varepsilon^2$$

(the deviation $\widetilde{O}(n^{4.5}/\varepsilon^2)$ from Lemma 19 with $B_1 = W^2$, $B_2 = I$ is o of the mean $\Theta(n^5/\varepsilon^2)$), hence $\|WM\|_F = \|MW\|_F \le C_5 n^{2.5}/\varepsilon$ (Hanson–Wright for the quadratic form $M \mapsto \mathrm{tr}(W^2M^2)$ in the bounded independent entries of M; mean $(n-1)v_A\,\mathrm{tr}(W^2)$ since $\mathbb{E}[M^2] = (n-1)v_A I$ and $\mathrm{tr}(W^2) = \sum_i \mu_i^2 = \Theta(n^4)$). Set $\mathcal{G}^+ := \mathcal{G} \cap \mathcal{G}_M$; a union bound gives $\Pr[\mathcal{G}^+] \ge 1 - n^{-12}$. Linear and quadratic forms in M obey Lemmas 16, 19 verbatim with σ replaced by $\sqrt{v_A} = \Theta(1/\varepsilon)$ (bounded ⇒ the same sub-exponential input). Throughout, "$\widetilde{O}$" absorbs $\mathrm{polylog}(n)$; the estimand's leading term is polylog-free (Remark 9).

Error decomposition. On $\mathcal{B}$, $F := \operatorname{tr}(A^5) = \mu_1^{5/2} + \langle \delta, A\rangle$ with $\delta = W^2 - \mu_1^2\Pi_1$ (Theorem 6). Expand $\widehat{W^2} = \widehat{W}_\sigma^2 - \operatorname{diag}(t_q) = W^2 + (WE+EW) + E_c^2$, where $E_c^2 := E^2 - \operatorname{diag}(t_q)$ is mean zero (the choice $t_q = \mathbb{E}[(E^2)_{qq}]$ kills the diagonal mean; off-diagonal, $\mathbb{E}[(E^2)_{jk}] = \sum_l \mathbb{E}[E_{jl}E_{lk}] = 0$ for $j \neq k$). Substituting into (S5) and using $\langle W^2, A\rangle = \langle \delta, A\rangle + \mu_1^2 v_1^\top A v_1 = \langle \delta, A\rangle + \mu_1^{5/2}$ ($v_1^\top A v_1 = \sqrt{\mu_1}$ on $\mathcal{B}$), regrouping yields the exact identity

$$\widehat{F} - F = \mathcal{P} + \mathcal{M} + \mathcal{L} + \mathcal{Q},$$
$$\mathcal{M} = \langle W^2 - (\hat{\mu}_1^c)^2\hat{\Pi}_1, M\rangle, \quad \mathcal{L} = \langle WE + EW, \widehat{A}\rangle, \quad \mathcal{Q} = \langle E_c^2, \widehat{A}\rangle,$$

and $\mathcal{P}$ the Perron aggregate of Lemma 23. The Perron re-subtraction $-(\hat{\mu}_1^c)^2\hat{\Pi}_1$, read against $\widehat{A} = A + M$, contributes $-(\hat{\mu}_1^c)^2\hat{v}_1^\top A\hat{v}_1$ to $\mathcal{P}$ and $-(\hat{\mu}_1^c)^2\hat{v}_1^\top M\hat{v}_1$ to $\mathcal{M}$; it does not cancel the first-order $WE + EW$ terms, which are bounded directly and are already within budget (Lemma 25). No cancellation is needed, because $\langle WE + EW, A\rangle = \widetilde{O}(n^3\sigma)$ is itself of the leading order $n^3\sigma$.

Lemma 24 (Seam / RR read-out $\mathcal{M}$)**.** *On $\mathcal{G}^+$, $|\mathcal{M}| \le \widetilde{O}(n^{2.5}/\varepsilon + n^{2.5}\sigma/\varepsilon) \le \widetilde{O}(n^{2.5}/\varepsilon + n^3\sigma)$.*

Proof. $\Gamma := W^2 - (\hat{\mu}_1^c)^2\hat{\Pi}_1$ is a function of (A, E), independent of M, so by Lemma 16 for M, $|\mathcal{M}| = |\langle \Gamma, M\rangle| \le C_1\varepsilon^{-1}\log^2 n\,\|\Gamma\|_F$. Write $\Gamma = \delta - [(\hat{\mu}_1^c)^2\hat{\Pi}_1 - \mu_1^2\Pi_1]$; then

$$\begin{aligned}\|(\hat{\mu}_1^c)^2\hat{\Pi}_1 - \mu_1^2\Pi_1\|_F &\le |(\hat{\mu}_1^c)^2 - \mu_1^2|\,\|\hat{\Pi}_1\|_F + \mu_1^2\|\hat{\Pi}_1 - \Pi_1\|_F \\ &\le 3\mu_1|\Delta| + \mu_1^2\sqrt{2}\,\beta = \widetilde{O}(n^2\sigma) + \widetilde{O}(n^{2.5}\sigma)\end{aligned}$$

by Lemmas 21 ($|\Delta| \le C\sigma\log^4 n$) and 22 ($\beta \le \widetilde{O}(\sigma/n^{1.5})$), with $\|\hat{\Pi}_1\|_F = 1$, $\|\hat{\Pi}_1 - \Pi_1\|_F = \sqrt{2}\,\beta$. With $\|\delta\|_F \le Cn^{2.5}$ (Lemma 14(e)), $\|\Gamma\|_F \le Cn^{2.5}(1 + \widetilde{O}(\sigma))$; multiply by $C_1\varepsilon^{-1}\log^2 n$ and use $n^{2.5}\sigma/\varepsilon \le n^3\sigma$ (as $\varepsilon \ge n^{-1/4}$ gives $1/(n^{1/2}\varepsilon) \le n^{-1/4}$). □

Lemma 25 (First-order wedge $\mathcal{L}$)**.** *On $\mathcal{G}^+$, $\mathcal{L} = 2\langle E, A^3\rangle + \langle E, WM + MW\rangle$ with $|2\langle E, A^3\rangle| \le \widetilde{O}(\sigma n^3)$ and $|\langle E, WM + MW\rangle| \le \widetilde{O}(\sigma n^{2.5}/\varepsilon)$; hence $|\mathcal{L}| \le \widetilde{O}(n^3\sigma)$.*

Proof. $\langle WE + EW, A\rangle = 2\operatorname{tr}(EWA) = 2\langle E, WA\rangle = 2\langle E, A^3\rangle$ ($WA = AW = A^3$ commute on $\mathcal{B}$). Lemma 16 with $C = A^3$ (fixed given A, independent of E), $\|A^3\|_F \le Cn^3$: $|2\langle E, A^3\rangle| \le \widetilde{O}(\sigma n^3)$. Next $\langle WE + EW, M\rangle = \langle E, WM + MW\rangle$ (cyclic; symmetry). Conditioning on M, $WM + MW$ is independent of E; Lemma 16 gives $|\langle E, WM + MW\rangle| \le C_1\sigma\log^2 n\,\|WM + MW\|_F \le C_1\sigma\log^2 n\cdot 2\|WM\|_F \le \widetilde{O}(\sigma n^{2.5}/\varepsilon)$ by (M3). Finally $\sigma n^{2.5}/\varepsilon \le n^3\sigma$ ($\varepsilon \ge n^{-1/4}$), so the first term dominates. The crude bound $\|WM + MW\|_F \le 2\|W\|_{\mathrm{op}}\|M\|_F = \Theta(n^3/\varepsilon)$ would lose a factor $\sqrt{n}$ and exit the budget; (M3), which uses that the random M is not aligned with the Perron direction of W, is what keeps $\mathcal{L}$ in budget. □

Lemma 26 (Second-order wedge $\mathcal{Q}$)**.** *On $\mathcal{G}^+$, $\mathcal{Q} = \operatorname{tr}(E^2A) + \operatorname{tr}(E^2M)$ with $|\operatorname{tr}(E^2A)| \le \widetilde{O}(\sigma^2 n^{1.5})$ and $|\operatorname{tr}(E^2M)| \le \widetilde{O}(\sigma^2 n^{1.5}/\varepsilon)$; hence $|\mathcal{Q}| \le \widetilde{O}(\sigma^2 n^{1.5}/\varepsilon) \le \widetilde{O}(n^3\sigma)$.*

Proof. $\mathcal{Q} = \langle E_c^2, A\rangle + \langle E_c^2, M\rangle = \operatorname{tr}(E^2A) + \operatorname{tr}(E^2M)$ (both A, M have zero diagonal, so $\operatorname{diag}(t_q)$ drops). By Lemma 19 with $B_1 = I$, $B_2 = A$: $\mathbb{E}\operatorname{tr}(E^2A) = \sigma^2[\operatorname{tr}A + \operatorname{tr}A] + \theta = \theta$ ($\operatorname{tr}A = 0$), $|\theta| \le 3C_0\sigma^2\sum_q |A_{qq}| = 0$, and the fluctuation is $\le C_4\sigma^2\log^4 n\,\|I\|_F\|A\|_F = \widetilde{O}(\sigma^2 n^{1.5})$. Condition on M (on $\mathcal{G}_M$): Lemma 19 with $B_1 = I$, $B_2 = M$ gives conditional mean $\sigma^2[\operatorname{tr}M + \operatorname{tr}M] + \theta_M = 0$ and fluctuation $\le C_4\sigma^2\log^4 n\,\|I\|_F\|M\|_F \le \widetilde{O}(\sigma^2 n^{1.5}/\varepsilon)$ by (M1). Domination: $\sigma^2 n^{1.5}/\varepsilon \le n^3\sigma \iff \sigma \le n^{1.5}\varepsilon$, and $\sigma \le cn/\varepsilon \le n^{1.5}\varepsilon \iff \varepsilon^2 \ge c/n^{1/2}$, true for $\varepsilon \ge n^{-1/4}$. □

Exact split and deflation gains. On $\mathcal{B}$, Π_1 commutes with A, so for any symmetric X,

$$\langle X, A\rangle = \sqrt{\mu_1}\,v_1^\top X v_1 + \langle PXP, A\rangle. \tag{5}$$

Applying (5) to $X = W$ (with $v_1^\top W v_1 = \mu_1$) and to $X = B$ (with $v_1^\top B v_1 = \mu_1\sum_i (d_i - 2)v_{1i}^2$, since $Wv_1 = \mu_1 v_1$) gives the true correction

$$D_5 = 5[\mu_1^{3/2} + \langle PWP, A\rangle] + 5\big[\mu_1^{3/2}\textstyle\sum_i (d_i - 2)v_{1i}^2 + \langle PBP, A\rangle\big],$$

matching (S6) term by term. The deflation gains, from Lemma 14 and degree-centering, are

$$\|PWP\|_F \le Cn^{1.5}, \qquad \|PBP\|_F \le \widetilde{O}(n^{2.5}),$$

the latter because $B = (\frac{n}{2} - 2)W + \frac{1}{2}(\Delta_D W + W\Delta_D)$ with $\Delta_D = \mathrm{diag}(d_i - \frac{n}{2})$, so $\|PBP\|_F \le (\frac{n}{2})\|PWP\|_F + \|\Delta_D W\|_F$, and

$$\|\Delta_D W\|_F^2 = \sum_i \left(d_i - \frac{n}{2}\right)^2 \|W_{i\cdot}\|_2^2 \le n^3 \sum_i \left(d_i - \frac{n}{2}\right)^2 = \widetilde{O}(n^5)$$

($W_{ij} \le n$; Lemma 14(f)). Without the centering, $\|PBP\|_F$ would be $\Theta(n^3)$ and the RR readout below would exit the budget by $\sqrt{n}$; the deflation of the Perron component ($\|PWP\|_F = O(n^{1.5}) \ll \|W\|_F = \Theta(n^2)$) is what makes both corrections affordable.

Lemma 27 (Projector swap)**.** *Let Y be symmetric and measurable with respect to (A, H, Z), hence independent of M. On $\mathcal{G}^+$,*

$$\|\hat{P}Y\hat{P} - PYP\|_F \le 2\beta\|Y\|_F, \qquad \beta := \|\hat{P} - P\|_{\mathrm{op}} = \widetilde{O}(\sigma/n^{1.5}).$$

Consequently, for every symmetric R fixed given (A, H, Z),

$$|\langle \hat{P}Y\hat{P}, R\rangle - \langle PYP, R\rangle| \le 2\beta\|Y\|_F\|R\|_F,$$

and, with conditional probability at least $1 - n^{-15}$ over M,

$$|\langle \hat{P}Y\hat{P}, M\rangle - \langle PYP, M\rangle| \le \widetilde{O}(\beta\|Y\|_F/\varepsilon).$$

Proof. Use $\hat{P}Y\hat{P} - PYP = (\hat{P} - P)Y\hat{P} + PY(\hat{P} - P)$, then apply Cauchy–Schwarz for R and the conditional linear-form bound for M. □

Lemma 28 (Corrections; unified and complete)**.** *On $\mathcal{G}^+$,*

$$|\widehat{D}_5 - D_5| \le \widetilde{O}\left(n^3\sigma + n^{2.5}\sigma_d + \frac{n^{2.5}}{\varepsilon}\right).$$

Proof. Write $\widehat{D}_5 - D_5 = 5(a+b) + 5(c+d)$ following the four terms of (S6) minus the exact split (5).

(a) W-Perron. $|a| = |(\hat{\mu}_1^c)^{3/2} - \mu_1^{3/2}| = \widetilde{O}(n\sigma)$ by Lemma 21.

(b) W-bulk. For $S_Y(R) := \langle \hat{P}Y\hat{P} - PYP, R\rangle$, the exact expansion is

$$\begin{aligned} b = {} & \langle PWP, M\rangle + \langle PEP, A\rangle + \langle PEP, M\rangle \\ & + S_W(A) + S_W(M) + S_E(A) + S_E(M). \end{aligned}$$

The three direct terms are respectively $\widetilde{O}(n^{1.5}/\varepsilon)$, $\widetilde{O}(\sigma n)$, and $\widetilde{O}(\sigma n/\varepsilon)$. Lemma 27, with $\|W\|_F = O(n^2)$, $\|E\|_F = \widetilde{O}(\sigma n)$, and $\|A\|_F = O(n)$, gives

$$\begin{aligned} |S_W(A)| &= \widetilde{O}(\sigma n^{1.5}), \quad |S_W(M)| = \widetilde{O}(\sigma n^{0.5}/\varepsilon), \\ |S_E(A)| &= \widetilde{O}(\sigma^2 n^{0.5}), \quad |S_E(M)| = \widetilde{O}(\sigma^2/(n^{0.5}\varepsilon)). \end{aligned}$$

(c) B-Perron. Put $q_i = d_i - 2$, $r = \sum_i q_i v_{1i}^2$, and $\alpha = (\hat{\mu}_1^c)^{3/2}$. Exactly,

$$\begin{aligned} c = {} & (\alpha - \mu_1^{3/2})r + \alpha \sum_i \zeta_i v_{1i}^2 + \alpha \sum_i \zeta_i(\hat{v}_{1i}^2 - v_{1i}^2) \\ & + \alpha \sum_i q_i(\hat{v}_{1i}^2 - v_{1i}^2). \end{aligned}$$

The first term is $\widetilde{O}(n^2\sigma)$. Since $\sum_i v_{1i}^4 = O(1/n)$, the second is $\widetilde{O}(n^{2.5}\sigma_d)$. Using $\|Z\|_F = \widetilde{O}(\sqrt{n}\sigma_d)$ and $\|\hat{v}_1^{\circ 2} - v_1^{\circ 2}\|_2 = \widetilde{O}(\sigma/n^{1.75})$, the third is $\widetilde{O}(n^{1.75}\sigma\sigma_d)$. In the fourth term, the constant part of q_i cancels because both eigenvectors are unit; Lemma 14(f) then gives $\widetilde{O}(n^{2.25}\sigma)$.

(d) B-bulk, including all projector swaps. Let $D' = \mathrm{diag}(d_i - 2)$ and define $C := \widehat{B} - B$. Since $E = H + Z$ and Z is diagonal,

$$C = C_1 + C_2 + C_3 + C_4 + C_5,$$
$$C_1 = \tfrac{1}{2}(D'H + HD'), \quad C_2 = D'Z, \quad C_3 = \tfrac{1}{2}(ZW + WZ),$$
$$C_4 = \tfrac{1}{2}(ZH + HZ), \qquad C_5 = Z^2.$$

There is no residual term. The exact bulk expansion is

$$d = \langle PBP, M\rangle + \langle PCP, A\rangle + \langle PCP, M\rangle + S_B(A) + S_B(M) + S_C(A) + S_C(M).$$

The final two terms are the projector swaps involving $\widehat{B} - B$.

On the blockwise good event,

$$\|C_1\|_F = \widetilde{O}(\sigma n^2), \quad \|C_2\|_F = \widetilde{O}(\sigma_d n^{1.5}), \quad \|C_3\|_F = \widetilde{O}(\sigma_d n^2),$$
$$\|C_4\|_F = \widetilde{O}(\sigma\sigma_d n), \qquad \|C_5\|_F = \widetilde{O}(\sigma_d^2\sqrt{n}).$$

Blockwise conditioning with Lemmas 16–19 gives

$$|\langle PC_1P, A\rangle| = \widetilde{O}(\sigma n^2),$$
$$|\langle PC_2P, A\rangle| + |\langle PC_3P, A\rangle| = \widetilde{O}(\sigma_d n^2),$$
$$|\langle PC_4P, A\rangle| = \widetilde{O}(\sigma\sigma_d n), \qquad |\langle PC_5P, A\rangle| = \widetilde{O}(\sigma_d^2 n).$$

Conditionally on (H, Z),

$$|\langle PCP, M\rangle| \le \frac{\widetilde{O}(1)}{\varepsilon}\big(\sigma n^2 + \sigma_d n^2 + \sigma\sigma_d n + \sigma_d^2\sqrt{n}\big).$$

The B swaps satisfy

$$|S_B(A)| = \widetilde{O}(\sigma n^{2.5}), \qquad |S_B(M)| = \widetilde{O}(\sigma n^{1.5}/\varepsilon),$$

using $\|B\|_F = O(n^3)$. The swaps involving $C = \widehat{B} - B$ are explicitly

$$|S_C(A)| \le 2\beta\|C\|_F\|A\|_F, \qquad |S_C(M)| \le \widetilde{O}(\beta\|C\|_F/\varepsilon),$$

where $\beta = \widetilde{O}(\sigma/n^{1.5})$. Substitution yields, against A,

$$\widetilde{O}(\sigma^2 n^{1.5} + \sigma\sigma_d n^{1.5} + \sigma^2\sigma_d n^{0.5} + \sigma\sigma_d^2),$$

and the same expression divided by $n\varepsilon$ against M. All displayed terms are within the target budget under (R) and $\sigma_d \le C_d\sigma$. Finally, $|\langle PBP, M\rangle| = \widetilde{O}(n^{2.5}/\varepsilon)$ is the leading correction readout. Summing (a)–(d) proves the lemma. □

Error budget. Collecting Lemmas 23, 24, 25, 26, and 28, every term is bounded by

$$R := C\log^5 n\left(n^3\sigma + n^{2.5}\sigma_d + \frac{n^{2.5}}{\varepsilon}\right)$$

on $\mathcal{G}^+$. In particular, the leading channels are $\mu_1^{3/2}e_{11} = O_{\mathbb{P}}(n^3\sigma)$, the diagonal B-Perron term $O_{\mathbb{P}}(n^{2.5}\sigma_d)$, and the seam and B-bulk randomized-response readouts $O_{\mathbb{P}}(n^{2.5}/\varepsilon)$. The second-order terms $n^{1.5}\sigma^2$, $n^{0.5}\sigma^3$, and $\sigma^2 n^{1.5}/\varepsilon$, as well as all explicit C_i and projector-swap terms in Lemma 28, are absorbed using $\varepsilon \ge n^{-1/4}$, $\sigma \le cn/\varepsilon$, and $\sigma_d \le C_d\sigma$. Thus, conditionally on every $A \in \mathcal{B}$,

$$\Pr\big[|\widehat{F} - F| + |\widehat{D}_5 - D_5| \le R\big] \ge 1 - n^{-12}.$$

Remark 9 (Where the polylog sits)**.** The leading off-diagonal Perron fluctuation has standard deviation $\Theta(n^3\sigma)$ without a logarithmic loss. The diagonal Perron readout has standard deviation $O(n^{2.5}\sigma_d)$, and the two principal randomized-response readouts have standard deviation $\Theta(n^{2.5}/\varepsilon)$. Polylogarithms enter only through the simultaneous tail envelopes for strictly lower-order terms and the operator-norm event.

Proof of Theorem 7(ii). Fix $A \in \mathcal{B}$ and take probability over the independent blocks (H, Z, M). On $\mathcal{G}^+$, the estimator is well-defined and

$$|\widehat{N}_5' - N_5| = \frac{1}{10}|(\widehat{F} - F) - (\widehat{D}_5 - D_5)| \le \frac{1}{5}R.$$

Clipping to $[0, n^5]$ and Lemma 13 (applied with its n^{-13} replaced by the n^{-12} of $\mathcal{G}^+$; the residual $n^{10} \cdot n^{-12} = n^{-2}$ is absorbed under the root) convert the conditional high-probability bound to

$$\mathbb{E}_{H,Z,M}\left[(\widehat{N}_5 - N_5)^2\right]^{1/2} \le \widetilde{O}\left(n^3\sigma + n^{2.5}\sigma_d + \frac{n^{2.5}}{\varepsilon}\right).$$

The complement of $\mathcal{B}$ has superpolynomially small probability, so taking the maximum over μ_+ and μ_- gives the same Bayes-risk bound. For the Gaussian oracle, $\sigma_d = \sigma$ and the middle term is absorbed by $n^3\sigma$. For the Fang–Yi release, $\sigma = \widetilde{O}(\sqrt{n}/\varepsilon)$ and $\sigma_d = \Theta(1/\varepsilon)$, yielding $\widetilde{O}(n^{3.5}/\varepsilon)$. □

A.8 Proof of Corollary 1

Proof sketch of Corollary 1. **Privacy and communication.** The three components carry their own guarantees—the off-diagonal wedge release $(\varepsilon/3, \delta)$-edge-DP (Lemma 10 run at the $\varepsilon/3$ share); the degree release, each report through the geometric mechanism (Fang & Yi, 2025, Lem. 3.1) at budget $\varepsilon/6$, so that the two reports one edge changes compose to $\varepsilon/3$-edge-DP; and randomized response at $\varepsilon/3$—each valid on all graphs at $d = n - 1$; basic composition gives (ε, δ)-edge-DP. Communication is dominated by the wedge release's $\widetilde{O}(n^2)$ bits per user at $q = 1$, $d = n-1$ (Fang & Yi, 2025, Thm. 4.3). **Utility.** Apply Theorem 7(ii) with $\sigma = \sigma_{\mathrm{FY}} = \widetilde{O}(\sqrt{n}/\varepsilon)$ off the diagonal and $\sigma_d = \Theta(1/\varepsilon)$ on the degree diagonal. Lemmas 11 and 12 verify the heterogeneous sub-exponential interface. Substitution into $\widetilde{O}(n^3\sigma + n^{2.5}\sigma_d + n^{2.5}/\varepsilon)$ gives $\widetilde{O}(n^{3.5}/\varepsilon)$ for $\varepsilon \ge n^{-1/4}$. Graph-independence and mutual independence of the two noise blocks are exactly what the perturbation proof uses. The lower bound is Lemma 5 exactly as instantiated in the proof of Theorem 5: the bound there is proved for the prior risk on $\mu_\pm$ against arbitrary estimators over per-edge product transcripts, which is the quantity the corollary compares. Note the wedge release is not a per-edge product channel—one edge influences the distribution of $\Theta(n)$ released pairs—so the corollary refutes no theorem; it locates the protocol on the far side of the model boundary that Theorem 5 draws. □

A.9 Additional discussion

The boundary, in context. The parity split organizes a boundary already visible in known rates: shuffling beats the best known local 4-cycle upper bound ($\widetilde{O}(n^{2.75})$ vs. $O(n^3)$) but, as Section 4 shows, so far only matches local for 5-cycles ($\Theta(n^4/\varepsilon)$ vs. local $O(n^4)$). Both certified witnesses live outside $\mathcal{C}$ (Remark 1): what is proven is the algebraic parity split and the shuffle-model rates, not a proved utility dichotomy wholly inside $\mathcal{C}$—there the even side is conditional, the odd side Conjecture 2. Corollary 1 buys the separation back with $\Theta(n)$ messages per edge on a hard family. Our lower bound also closes the clique-versus-cycle gap of Suppakitpaisarn et al. (2025): Theorem 5 raises the cycle bound to $\Omega(n^{k-1}/\varepsilon)$ in the non-interactive per-edge local model—cycles exactly as hard as cliques there—leaving the residual $\sqrt{n}$ as a question about row-level randomizers and interaction, the latter open even for triangles (Eden et al., 2023).

The triangle as Rosetta stone. The apparent $\sqrt{n}$ cycle gap is a conflation of models; the triangle tells the whole story. Non-interactively, dense triangle counting is $\Theta(n^2/\varepsilon)$ (Imola et al., 2022; Eden et al., 2023), Theorem 5 at $k = 3$ supplying the per-edge ε-dependence; with interaction the known floor drops to $\Omega(n^{3/2}/\varepsilon)$, matched by a two-round protocol only on C_4-sparse instances, and whether interaction beats n^2 on dense graphs is open (Imola et al., 2021; Eden et al., 2023)—the same interactivity question that survives, for per-edge channels by Theorem 5, at every k. Betzer et al. (2024) publish per-vertex walk counts and Katz centrality over several local rounds, but walk counts are neither closed-walk traces nor simple-cycle counts, so they leave this question open.

A.10 Experimental details (§6)

All shuffle mechanisms are simulated through their asymptotic noise laws with normalized constants—a noise-model simulation validating the estimator algebra and rates, not an end-to-

end message-level implementation—and all polylog calibration constants are set to 1, so the plots validate scaling exponents, not privacy-calibrated practical performance. Graphs: $G(n, \frac{1}{2})$ and the hard family $\mu_\pm = \mathrm{Bern}(\frac{1}{2} \pm \gamma^*)^{\otimes\binom{n}{2}}$, $\gamma^* = 1/(\varepsilon n)$, alternating signs across repetitions; $n \in \{200, 400, 800, 1600\}$, 16 repetitions per configuration, $\varepsilon = 1$ for the n-sweeps and $\varepsilon \in \{0.5, 1, 2, 4\}$ at $n = 800$. Ground truth N_5 by the exact identity

$$\mathrm{tr}(A^5) = 10N_5 + 5\,\mathrm{tr}(A^3) + 5\sum_i (d_i - 2)(A^3)_{ii},$$

validated against brute-force enumeration at $n = 8$. Our protocol: wedge releases with i.i.d. symmetric discrete-Laplace noise at scale $3\sqrt{n}/\varepsilon$ (the certified pair-indexed instantiation at the $\varepsilon/3$ share; polylog constants and the sub-split constants of Lemma 10(e) set to 1), diagonals from two independent degree releases at scale $3/\varepsilon$ (one per wedge factor), blind RR at $\varepsilon/3$, degeneracy corrections by post-processing $\widehat{A}$. The Corollary 1 protocol is simulated through its asymptotic noise law (same scales) with the exact (S1)–(S7) estimator of the Theorem 7(ii) proof—top eigenpair, second-order Perron mean subtraction with the $(n-1)\sigma^2$ constant, noise-trace debiasing of $\widehat{W}^2$ with the heterogeneous diagonal t_q, deflated readout, projector-tilt mean subtraction, exact projected contractions $\langle \hat{P}\widehat{W}\hat{P}, \widehat{A}\rangle$ and $\langle \hat{P}\widehat{B}\hat{P}, \widehat{A}\rangle$ for the corrections, clipping to $[0, n^5]$—not a simplified surrogate. The separation slope difference carries a paired bootstrap 95% CI (resampling repetitions, shared graph instances across mechanisms), stored with the results. Fitted slopes carry standard errors of ± 0.03–0.25, computed as ordinary least-squares residual standard errors of the four-point log–log fits (log RMSE vs. log n, or vs. log ε for the ε-sweep); the same seeded graph instances are shared across mechanisms at each n (so between-mechanism slope comparisons are paired; for the deflation-vs-RR gap the paired bootstrap CI below is the relevant uncertainty), and RMSE at each point aggregates 16 independent graph-and-noise repetitions (a fresh graph per repetition, shared across mechanisms within it); slopes and standard errors are computed and stored by the released script. In particular the separation slopes are 3.35 ± 0.08 (deflation) vs. 3.98 ± 0.12 (per-edge RR), slope difference 0.63 (bootstrap mean 0.64) with paired-bootstrap 95% CI [0.27, 1.01]. All plug-in estimators (the D_5 correction and the local-RR baseline) use the unbiased multilinear extensions of Lemma 6—the raw binary matrix-power identities are biased on noisy debiased matrices, since monomials reusing an edge pick up $\mathbb{E}[\widehat{A}_e^2] \neq A_e$; the released script verifies multilinearity (vanishing second differences) and, for the local-RR plug-in, Monte-Carlo unbiasedness at startup. Local RR: full-budget randomized response with the debiased plug-in. Central: $N_5 + \mathrm{Lap}(\Delta/\varepsilon)$, $\Delta = (n-2)(n-3)(n-4)$—the number of simple 5-cycles through a fixed edge $\{u, v\}$ of K_n, one per ordered internal path u–x–y–z–v (at $n = 5$: 6 per edge, consistent with K_5's 12 cycles over 10 edges). Figure 1 shows the four displays of §6; its pivot-baseline line is the theoretical $n^{4.5}$ rate with the constant matched to our RMSE at $n = 200$, so only its slope is meaningful. Relative RMSEs quoted in §6 divide the stored RMSE by the mean count $\mathbb{E}N_5 = 12\binom{n}{5}2^{-5}$ at $p = \frac{1}{2}$ (on $\mu_\pm$ the mean differs by a factor $(1 \pm 2\gamma^*)^5$, under 1% at $n = 1600$).

Remark 10. The failure of the free-exact-W strengthening is not in the local-channel step but in the packing: $\mathrm{tr}(W) = 2\sum_e A_e$ leaks the density with mean gap $\Theta(n^2\gamma)$ over pooled sd $\Theta(n)$, equal to $n\gamma^* = \Theta(1/\varepsilon) \gg 1$ at the threshold, so the free scalar $\mathrm{tr}(W)$ separates μ_+ from μ_- a factor $1/\varepsilon$ before the private transcript can. Equivalently, for symmetric A_0 and $\delta W = A_0\Delta + \Delta A_0$, $\delta\,\mathrm{tr}(A^5) = \frac{5}{2}\langle A_0^3, \delta W\rangle$: the entire first-order signal is a noiseless linear read-out of the free δW. We therefore state Theorem 5 in the model where W is recovered from the same ε-local transcript, which is the class the minimax bound should target.

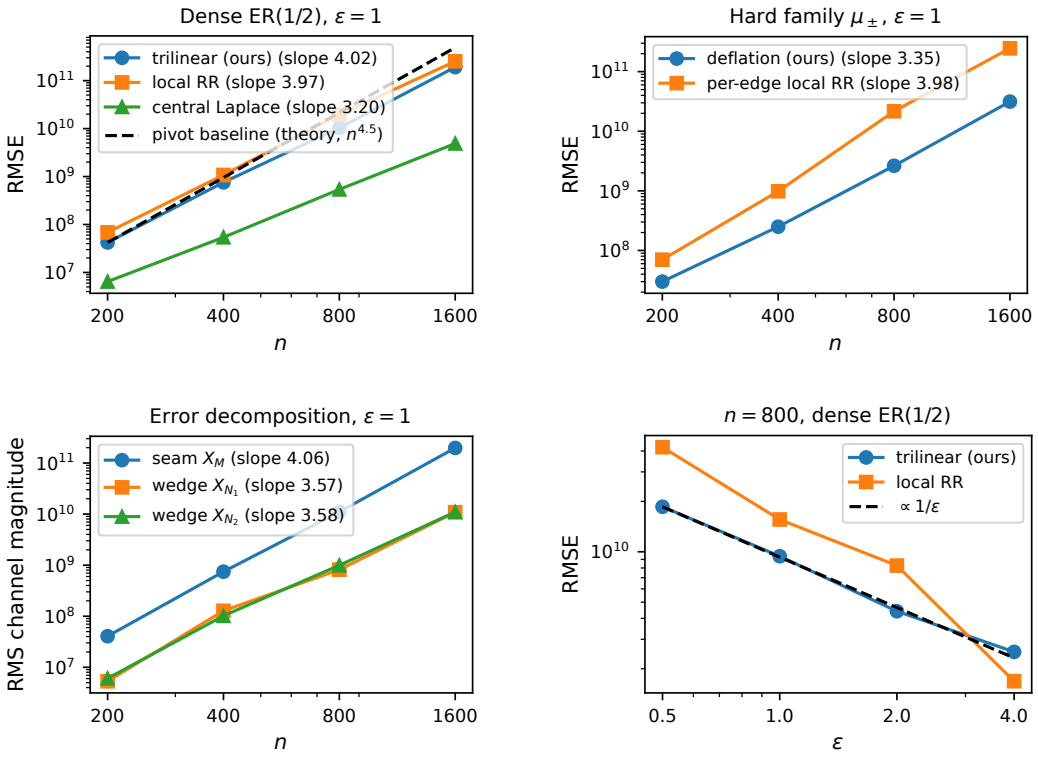

Figure 1: Top left: RMSE scaling on $G(n, \frac{1}{2})$, $\varepsilon = 1$—ours ties local RR at slope ≈ 4, central at ≈ 3, pivot baseline ($n^{4.5}$, theoretical) above. Top right: the Corollary 1 separation on $\mu_{\pm}$—deflation at slope 3.35 vs. per-edge RR at 3.98. Bottom left: error decomposition of Theorem 3—the seam channel (slope 4.06) dominates the wedge channels (≈ 3.6) by a growing factor. Bottom right: $1/\varepsilon$ dependence at $n = 800$.

A.11 Proof sketches for Theorems 1–4

Below are the full sketches of Theorems 1 and 3 (the main text gives shorter versions), an aside on the privacy argument of Theorem 2, and the sketch of Theorem 4, which the main text omits.

Theorem 1.

Proof sketch. Since $A^{2m} = (A^2)^m = W^m$, the count carries no bare factor of A. By (1), $\mathrm{tr}(W^m) = \sum_{k_1,\dots,k_m} \prod_t \langle a_{k_t}, a_{k_{t+1}} \rangle$ (indices mod m) is a U-statistic over m-tuples of user-indexed, bounded-influence contributions. Privacy of the post-processed polynomial $P(\widehat{W}) = \mathrm{tr}(\widehat{W}^m)$ is immediate; the utility content of the theorem is the $m = 2$ instance, which is the proven negative-binomial rate of Fang & Yi (2025) at $\widetilde{O}(n^{2.75})$—a pair-indexed release that, like ours, exits $\mathcal{C}$ (Remark 1); in-$\mathcal{C}$ amplifiability (Definition 4) is conditional on the within-row ledger of Remark 3. For general m, the linear noise tier costs $\widetilde{O}\big(v_0^{1/2}\,\mathrm{tr}(W^{2m-1})^{1/2}\big) = \widetilde{O}(n^{2m-1.5}/\varepsilon)$ by the per-user variance computation of Theorem 4—again $\sqrt{n}$ below the local frontier—but we have not carried out the higher tiers, so we scope the utility claim to $m = 2$ and record the general-m rate as a flagged extension, not a claim. Identity alone is not utility. □

Theorem 2 (why no amplification shortcut applies). No amplification shortcut certifies the wedge releases: an edge touches only two rows, but a row outer product is not low-sensitivity—one edge flips $O(n)$ of its entries; the $O(1)$-influence framing of §3 describes the conditional in-$\mathcal{C}$ variant only (Remark 1).

Theorem 3.

Proof sketch. By Lemma 1 the error is $\widehat{T} - \mathrm{tr}(A^5) = X_{N_1} + X_{N_2} + X_M +$ (higher order), the leading part being linear in each noise; distinct noise monomials are L^2-orthogonal, so the seam variance below lower-bounds the total unconditionally, and the nonlinear tiers total $\widetilde{O}(n^7/\varepsilon^2)$ for $\varepsilon \geq n^{-1/2}$ (Lemma 7). The adjacency (seam) channel is

$$X_M = \sum_{i,j,k} W_{ij} M_{jk} W_{ki} = \sum_{j,k} M_{jk} (A^4)_{kj},$$

using $\sum_i (A^2)_{ij}(A^2)_{ki} = (W^2)_{kj} = (A^4)_{kj}$. The entries $\{M_{jk}\}_{j<k}$ are independent, mean-zero, with variance $v_A = \Theta(1/\varepsilon^2)$ for blind randomized response, and M is mirrored, so

$$\mathrm{Var}[X_M] = 4 v_A \sum_{j<k} (A^4)_{jk}^2 = \Theta\big(v_A \|A^4\|_F^2\big).$$

In the dense regime $\|A^4\|_F = \Theta(n^4)$—the average degree $\Theta(n)$ gives $\lambda_1 = \Omega(n)$, so $\|A^4\|_F \geq \lambda_1^4 = \Omega(n^4)$, while $\|A^4\|_F \leq \|A\|_2^3 \|A\|_F = O(n^4)$ (this spectral route, unlike an entrywise one, survives dense bipartite graphs, where half the entries of A^4 vanish)—whence $\mathrm{Var}[X_M] = \Theta(n^8/\varepsilon^2)$, giving an RMSE contribution $\Theta(n^4/\varepsilon)$. The wedge channels $X_{N_i} = \mathrm{tr}(N_i A W) = \mathrm{tr}(N_i A^3)$ deserve care. Under the certified pair-indexed instantiation the off-diagonal entries of N_i are i.i.d. with per-entry variance $\widetilde{O}(n/\varepsilon^2)$ (Lemma 10), so $\mathrm{Var}[X_{N_i}] = \widetilde{O}(n/\varepsilon^2) \cdot \|A^3\|_F^2 = \widetilde{O}(n^7/\varepsilon^2)$, a factor $\widetilde{O}(n)$ below the seam (the degree-noise diagonal adds $O(n^5/\varepsilon^2)$, lower order still; the conditional per-user variant lands at the same bound by a different route, Remark 7). Either way the wedge channels are lower order; the bare edge dominates this estimator's error. □

The rate is therefore pinned by the bare-edge channel $\widehat{A}$—precisely the seam of §3 (Conjecture 1): the two wedge factors are low-noise (pair-indexed or amplified) but wash out, while the single seam edge conjecturally cannot be amplified in the one-round sum-based class $\mathcal{C}$ and sets $\Theta(n^4/\varepsilon)$.

Theorem 4.

Proof sketch. Independence of the m wedge releases and the blind edge release gives unbiasedness, and basic composition gives the fixed ledger. The bare-edge term $X_M = \langle M, A^{k-1} \rangle$ has variance $\Theta_k(n^{2k-2}/\varepsilon^2)$; the wedge and higher-order tiers fit the matching upper budget for $\varepsilon \geq n^{-1/2}$

($\varepsilon \geq \sqrt{\log n/n}$ if $k = 3$; Lemma 7). Lemma 6 gives $\|\widehat{D}_k - D_k\|_2 = O_k(n^{k-1}/\varepsilon)$, so Minkowski's inequality yields the simple-count upper bound. When the correction is little-o of the trace error, the reverse triangle inequality gives the stated Θ_k rate. Appendix A.2 supplies the full tier computation. □