PRIVATE DENSEST SUBMATRIX SELECTION: PRIVACY IS NEARLY FREE

Anonymous authors
Paper under double-blind review

ABSTRACT

We study the densest $k \times k$ submatrix problem for an $n \times n$ Gaussian random matrix under differential privacy, where a record is either one entry or one row. Without privacy, the optimum is $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$, and the best known polynomial-time algorithm attains a $\frac{2\sqrt{2}}{3}$ fraction of it. This constant is optimal among online algorithms and conjectured optimal among all polynomial-time ones. Our main result is an online polynomial-time ε-DP mechanism that attains the same $\frac{2\sqrt{2}}{3}$ fraction once $k \gg \log n \cdot \log(k \log n)/\varepsilon^2$. We also show that no ε-DP mechanism, efficient or not, achieves a per-entry average above a constant depending on ε alone. In fact the recoverable fraction of the optimum at fixed ε is $\Theta(\min\{1, \varepsilon\sqrt{k/\log n}\})$, so the threshold of our main result is within a $\log \log$ factor of optimal. Under row-level privacy, no mechanism recovers a constant fraction of the optimum unless $\varepsilon = \Omega(\log n)$, regardless of k.

1 INTRODUCTION

Finding a dense block in a large matrix is a basic primitive in data analysis. Given $G \in \mathbb{R}^{n \times n}$ and $k \leq n$, the densest $k \times k$ submatrix problem asks for row and column sets $I, J \subset [n]$ of size k maximizing the average $\mathrm{Ave}(G_{I,J}) = \frac{1}{k^2}\sum_{i \in I, j \in J} G_{ij}$. It is the real-valued version of the maximum-edge biclique and densest-k-subgraph problems, and it is the primitive underlying biclustering (Madeira & Oliveira, 2004; Shabalin et al., 2009), in which one looks for a block of genes and experimental conditions with jointly elevated expression. The same primitive appears in co-clustering for collaborative filtering (Dhillon, 2001; George & Merugu, 2005) and in dense-block detection for fraud in bipartite activity data (Hooi et al., 2016). Such data are typically individual-level, and the block that an algorithm returns reveals information about the entries that made it dense. In this work we ask what can be achieved for this problem under differential privacy, in the random-matrix model in which its algorithmic theory has been developed, and how the answer depends on the privacy budget ε, the block size k, and on whether a record is a single entry or a whole row (entry-level versus row-level privacy).

The random-matrix model is that of Bhamidi et al. (2017), Gamarnik & Li (2018) and Bhamidi et al. (2026), in which G has i.i.d. standard Gaussian entries. Without privacy, the problem is well understood. The optimum concentrates at $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$. The best known polynomial-time algorithm, an incremental greedy procedure, achieves a constant fraction $\frac{2\sqrt{2}}{3} \approx 0.943$ of it. Bhamidi et al. (2026) recently proved that this constant is optimal among all *online* algorithms, and conjectured it to be optimal over all polynomial-time algorithms for growing k. We show that this picture survives entry-level privacy.

Our primary contribution is an online polynomial-time ε-DP mechanism that attains $(\frac{2\sqrt{2}}{3} - o(1))\eta_{\mathrm{OPT}}$ with high probability once k exceeds $\log n \cdot \log(k \log n)/\varepsilon^2$. It runs a batched version of the incremental greedy procedure, with each selection made by an exponential mechanism rather than an arg-max. Selecting many rows at once costs no more privacy than selecting one, because changing a single entry changes the block sum of any candidate set by a bounded amount regardless of its size. Choosing the batch sizes carefully removes the asymptotic effect of privacy on the utility. We also show that the exponential mechanism over all $\binom{n}{k}^2$ blocks is ε-DP and attains

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.

Table 1: The densest $k \times k$ submatrix of a Gaussian matrix without and with entry-level ε-DP, in units of $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$.

	non-private	**ε-DP (this work)**
statistical optimum	1 (Bhamidi et al., 2017)	1 once $k \gg (\log n)^2/\varepsilon^2$, not efficient (Thm. F.1)
best achievable at $k = t \log n$	1	$\Theta(\min\{1, \varepsilon\sqrt{t}\})$ (Cor. 4.3)
best achievable at $k = o(\log n)$	1	$o(1)$ (Cor. 4.3)
best known polynomial time	$\frac{2\sqrt{2}}{3}$ (Gamarnik & Li, 2018)	$\frac{2\sqrt{2}}{3}$ once $k \gg \log n \cdot \log(k \log n)/\varepsilon^2$ (Thm. 3.2)
optimal among online algorithms	$\frac{2\sqrt{2}}{3}$ (Bhamidi et al., 2026)	$\frac{2\sqrt{2}}{3}$, by an online mechanism (Thm. 3.2, Prop. E.2)

$(1 - o(1))\eta_{\mathrm{OPT}}$ once $k \gg (\log n)^2/\varepsilon^2$ (Theorem F.1). For sufficiently large k, this reproduces the statistical-to-computational gap of the non-private problem.

Our second contribution says what happens when k is too small for the first. We prove that no ε-DP mechanism, efficient or not, can achieve a per-entry average larger than a constant depending on ε alone. We give a simple mechanism which looks only at the signs of the entries that gets within a constant factor of that bound. Together, these establish the asymptotic recoverable fraction of the optimum at fixed ε is $\Theta(\min\{1, \varepsilon\sqrt{k/\log n}\})$. Nothing is recoverable when $k = o(\log n)$, while a constant fraction is once $k \asymp \log n/\varepsilon^2$. The threshold $\log n \cdot \log(k \log n)/\varepsilon^2$ of our main result is therefore within a $\log \log$ factor of the best possible.

Our third contribution is the row-level picture, where a private record is a whole row rather than one entry. We show that no ε-row-DP mechanism recovers a constant fraction of the optimum unless $\varepsilon = \Omega(\log n)$, whatever k is, so a larger block no longer helps. A one-shot mechanism that fixes the columns and selects rows by their clipped row sums recovers a constant fraction of η_{OPT} once $\varepsilon = \omega(\log n)$, so $\log n$ is the correct threshold. The two regimes diverge because under entry-level privacy a large block averages many more records, while under row-level privacy each row is one record and its contribution to the block is as large as the signal itself.

While our privacy guarantees hold for all real inputs, the utility guarantees in this work are for the Gaussian ensemble and are asymptotic in n. Our core focus is a theoretical study of phase transitions and computational gaps, rather than the design of an immediately deployable method. Appendix J extends this work to general unit-variance sub-Gaussian entries. Crucially, the mechanism transfers and still attains $(\frac{2\sqrt{2}}{3} - o(1)) \cdot 2\sqrt{\log n/k}$ with the same threshold for k.

1.1 Further related work

The exponential mechanism of McSherry & Talwar (2007) is our basic tool, and its loss of $O(\Delta \log N/\varepsilon)$ against the best of N candidates is the source of every penalty term here. Private top-k selection (Bhaskar et al., 2010; Durfee & Rogers, 2019; Qiao et al., 2021; McKenna & Sheldon, 2020; Gillenwater et al., 2022) and its lower bounds (Bafna & Ullman, 2017; Steinke & Ullman, 2017) concern the setting where a single record can shift every candidate's score, which is what the entry-level model avoids. Private densest subgraph (Nguyen & Vullikanti, 2021; Farhadi et al., 2022; Dhulipala et al., 2022; Dinitz et al., 2025; Khayatian et al., 2025; Zhou, 2025) is edge-private and worst-case, with additive error guarantees that never see the extreme-value scale $\sqrt{\log n/k}$. Private recovery in planted models (Mohamed et al., 2022; Chen et al., 2023) targets a planted signal, whereas here the objective is an extremal statistic of pure noise. Our mechanism is a private greedy algorithm in the sense of Gupta et al. (2010), and it uses the same size-independent sensitivity of subset scores as private submodular maximization (Mitrovic et al., 2017), but the batching argument of Section 3.4 is new. Appendix A has a fuller discussion.

1.2 Outline

Section 2 sets up the problem, the privacy model, and the one lemma on sensitivity that the rest of the paper depends on. Section 3 presents the batched mechanism, states the main theorem, and explains why batching and constant clipping work. Section 4 gives the information-theoretic upper bound, the sign mechanism that matches it, and the resulting phase diagram. Section 5 treats row-level privacy, Section 6 reports experiments, and the appendices contain the full proofs, the extended related work, and the extension to sub-Gaussian entries.

2 PROBLEM FORMULATION

2.1 THE DENSEST SUBMATRIX PROBLEM

We write $[n] = \{1, \dots, n\}$, $\mathrm{Sum}(G_{I,J}) = \sum_{i \in I, j \in J} G_{ij}$ and $\mathrm{Ave} = \mathrm{Sum}/(|I||J|)$. All logarithms are natural, and $\phi, \Phi, \bar{\Phi}$ are the standard normal density, distribution function and upper tail. Throughout, G has i.i.d. $N(0,1)$ entries, $k = k(n)$ satisfies $k = \exp(o(\log n))$, and $n \geq 4k$. All limits are as $n \to \infty$, and "whp" means with probability $1 - o(1)$.

For this ensemble, Bhamidi et al. (2017) showed that the optimum concentrates at $\eta_{\mathrm{OPT}} := 2\sqrt{\log n/k}$. The upper bound is a union bound over the $\binom{n}{k}^2$ candidate blocks (Lemma C.1) and holds throughout our regime. The lower bound is their second-moment analysis, which we use only for the benchmark of Appendix F. Gamarnik & Li (2018) began the algorithmic study. They showed that the Large Average Submatrix (LAS) search of Shabalin et al. (2009), which alternately re-optimizes columns against the current rows and rows against the current columns, converges for constant k to a block of average $\sqrt{2\log n/k} = \eta_{\mathrm{OPT}}/\sqrt{2}$, and that an *incremental greedy procedure* (IGP) does better. The IGP starts from a single row and alternately adds the best column against the current rows and the best row against the current columns, each drawn from a fresh group of n/k indices, until both sets have size k. It achieves

$$\eta_{\mathrm{ALG}} := \tfrac{4}{3}\sqrt{2\log n/k} = \tfrac{2\sqrt{2}}{3}\,\eta_{\mathrm{OPT}} \approx 0.943\,\eta_{\mathrm{OPT}}.$$

The factor $\frac{4}{3}$ comes from scoring each new column against only the r rows present when it is added, since summing $\sqrt{r}$ over both sides gives $2\int_0^1 \sqrt{x}\,dx = \frac{4}{3}$. Bhamidi et al. (2026) proved via a branching overlap-gap property (Gamarnik, 2021; Gamarnik et al., 2020) that η_{ALG} is optimal among *online* algorithms, those whose output is committed in k stages with the s-th stage seeing only the principal $\lfloor sn/k \rfloor \times \lfloor sn/k \rfloor$ block of G. They conjecture that no polynomial-time algorithm does better for growing k.

2.2 DIFFERENTIAL PRIVACY AND THE SELECTION MECHANISM

Definition 2.1 (Entry-level differential privacy)**.** Two matrices $G, G' \in \mathbb{R}^{n \times n}$ are *neighbors*, written $G \sim G'$, if they agree in every entry but one, which may take any two real values. A randomized mechanism $\mathcal{A}$ is *ε-differentially private* (ε-DP) if

$$\Pr[\mathcal{A}(G) \in S] \;\leq\; e^{\varepsilon}\,\Pr[\mathcal{A}(G') \in S]$$

for all neighbors $G \sim G'$ and all measurable sets S of outputs.

This is entry-level privacy, the analogue of edge privacy for graphs. It is the right notion when each entry is a separate record. It does not by itself protect a subject who contributes an entire row, for which one needs row-level privacy, treated in Section 5.

Gaussian entries are unbounded, so every mechanism in this work first clips. For $\theta > 0$ let $\widetilde{G}_{ij} = \max(-\theta, \min(\theta, G_{ij}))$ and $\Delta := 2\theta$. Two functions of θ appear in the bounds, the clipped variance $\sigma_\theta^2 := \mathbb{E}[\widetilde{G}_{11}^2] \in (0,1]$ and the mean loss from clipping below, $\psi(\theta) := \mathbb{E}[(-\theta - G_{11})_+] \leq e^{-\theta^2/2}$. Both tend to their limits 1 and 0 as $\theta \to \infty$ (Lemma C.3).

Every method in this work builds on a selection mechanism. A set of rows (or columns) I is already in hand, a *pool* $P \subseteq [n]$ of candidate columns (or rows) is available, and the step chooses a subset $S \subseteq P$ of prescribed size κ, scoring each candidate subset by the clipped block sum $\mathrm{Sum}(\widetilde{G}_{I,S})$. The choice is made by the exponential mechanism of McSherry & Talwar (2007) over the κ-subsets of the pool.

Definition 2.2 (Subset selection mechanism)**.** For an index set I, a pool P, an integer $\kappa \leq |P|$ and a budget $\varepsilon' > 0$, the mechanism $\mathcal{M}_{\mathrm{sel}}(I, P, \kappa, \varepsilon')$ outputs $S \subseteq P$ with $|S| = \kappa$, drawn from $\Pr[S] \propto \exp\left(\frac{\varepsilon'}{2\Delta}\,\mathrm{Sum}(\widetilde{G}_{I,S})\right)$, with I indexing rows and S columns or the reverse.

The score is additive over the members of S, so the weight $\exp\left(\frac{\varepsilon'}{2\Delta}\,\mathrm{Sum}(\widetilde{G}_{I,S})\right)$ of S is a product $\prod_{j \in S} w_j$ of per-item weights w_j. A distribution of this form is a conditional Poisson design (Hájek,

1964; Chen et al., 1994). It can be sampled exactly by a dynamic program in $O(|P|\kappa)$ time, or by including each item independently with a probability increasing in w_j and keeping the first draw of size κ, which takes $O(\sqrt{\kappa})$ draws in expectation and $O(|P|)$ space (Proposition C.15). Both samplers are exact, so the privacy guarantee matches the exponential mechanism itself.

Lemma 2.3 (Sensitivity, Lemmas C.4–C.7)**.** *If $G \sim G'$ differ at (i_0, j_0) then $|\operatorname{Sum}(\widetilde{G}_{I,J}) - \operatorname{Sum}(\widetilde{G}'_{I,J})| \le \Delta \cdot \mathbf{1}[i_0 \in I,\, j_0 \in J]$ for all I, J. Consequently $\mathcal{M}_{\mathrm{sel}}(I, P, \kappa, \varepsilon')$ is ε'-DP, and it selects a set whose score is within $\frac{2\Delta}{\varepsilon'}(\log\binom{|P|}{\kappa} + t)$ of the best with probability $1 - e^{-t}$.*

Composing these mechanisms with ε_τ privacy budgets results in a $(\sum_\tau \varepsilon_\tau)$-DP mechanism.

3 The Batched Incremental Greedy Mechanism

We present our main mechanism and its guarantee. We begin with the simplest private mechanism for the problem, a single call to $\mathcal{M}_{\mathrm{sel}}$ that selects one side of the block, which already exhibits the composition-free selection that Lemma 2.3 allows. We then batch the incremental greedy procedure, state the main theorem, and explain the ideas behind it. Full proofs are in Appendices D and E.

3.1 Warm-up: the one-shot mechanism

Fix $I_0 = [k]$. The one-shot mechanism $\mathcal{M}_{\mathrm{OS}}$ returns (I_0, J^*) with $J^* = \mathcal{M}_{\mathrm{sel}}(I_0, [n] \setminus I_0, k, \varepsilon)$. By Lemma 2.3 it is ε-DP with no composition at all.

Theorem 3.1 (One-shot mechanism, proved in Appendix D)**.** *Let $\theta = \theta_n := \max\{1, \sqrt{2\log(k \log n)}\}$. Suppose $\varepsilon\sqrt{k}/(\theta_n\sqrt{\log n}) \to \infty$ and $k \gg \theta_n^2 \log n$, the latter automatic when $\varepsilon = O(1)$. Then $\mathcal{M}_{\mathrm{OS}}$ is ε-DP, runs in $O(nk)$ time, and satisfies $\operatorname{Ave}(G_{I_0,J^*}) \ge (1/\sqrt{2} - o(1))\eta_{\mathrm{OPT}}$ whp and in expectation. The finite-n rate behind it is Proposition D.1.*

Whatever the fixed rows are, the column scores are i.i.d. sums of k clipped Gaussians. Thus the non-private benchmark, the best k of them, is a one-sided extreme-value quantity of order $k \cdot \sigma_\theta\sqrt{k}\sqrt{2\log n}$, which after division by k^2 is $\sigma_\theta\eta_{\mathrm{OPT}}/\sqrt{2}$. By Lemma 2.3 the selection loses at most $\frac{2\Delta}{\varepsilon}\log\binom{n}{k}$, which is $O(\theta \log n/(\varepsilon k))$ per entry. Section 3.5 shows that clipping changes neither of these quantities by more than a $1 - o(1)$ factor. The constant $1/\sqrt{2}$ is the LAS fixed-point value of Gamarnik & Li (2018), and $\mathcal{M}_{\mathrm{OS}}$ falls short of η_{ALG} because it adapts on one side only.

3.2 The mechanism

Our main mechanism, $\mathcal{M}_{\mathrm{BIG}}$, runs the IGP in B batches rather than one row or column at a time, selecting each batch by one call to $\mathcal{M}_{\mathrm{sel}}$. Fix a batch count B with $1 \le B \le \sqrt{k/3}$ and set $\kappa := \lfloor k/(B+1) \rfloor$. Batch sizes are $\kappa_b := \kappa$ for $b < B$ and $\kappa_B := k - B\kappa \in [\kappa, 2\kappa]$, so that $\kappa + \sum_b \kappa_b = k$. We define prefix points $p_0 := \kappa$ and $p_b := \lfloor (b\kappa + 1)n/k \rfloor$. These serve as endpoints for our index groups $Q_b := \{p_{b-1} + 1, \ldots, p_b\}$, of size between 2κ and $2\kappa n/k$. We allow a per-step privacy budget of $\varepsilon_0 := \varepsilon/(2B)$.

Algorithm 1 $\mathcal{M}_{\mathrm{BIG}}$: batched incremental greedy mechanism

Require: $G \in \mathbb{R}^{n\times n}$; block size k; batches B; budget ε; clipping level θ

$\widetilde{G} \leftarrow$ entrywise clipping of G to $[-\theta, \theta]$; $\quad \Delta \leftarrow 2\theta$; $\quad \varepsilon_0 \leftarrow \varepsilon/(2B)$

$I_0 \leftarrow [\kappa]$, $J_0 \leftarrow [\kappa]$ ▷ data-independent initial slot

for $b = 1, \ldots, B$ **do**

  $S_b \leftarrow \mathcal{M}_{\mathrm{sel}}(I_{b-1}, Q_b, \kappa_b, \varepsilon_0)$ ▷ select columns of Q_b against the current rows

  $J_b \leftarrow J_{b-1} \cup S_b$

  $R_b \leftarrow \mathcal{M}_{\mathrm{sel}}(J_b, Q_b, \kappa_b, \varepsilon_0)$ ▷ select rows of Q_b against the current columns

  $I_b \leftarrow I_{b-1} \cup R_b$

end for

return (I_B, J_B)

Table 2: Parameters of $\mathcal{M}_{\mathrm{BIG}}$ and their values in Theorem 3.2.

	meaning	**value for the guarantee**
n, k	matrix dimension, block size	$k = \exp(o(\log n))$; for fixed ε, $k \gg \log n \cdot \log(k \log n)$
ε	privacy budget	any $\varepsilon > 0$ with $\varepsilon\sqrt{k} \gg \theta\sqrt{\log n}$
θ	clipping level	$\theta = \max\{1, \sqrt{2\log(k \log n)}\}$
B	number of batches	$B \to \infty$ with $B \ll \varepsilon\sqrt{k}/(\theta\sqrt{\log n})$ and $B \le \sqrt{k/3}$
$\kappa, \varepsilon_0, \Delta$	batch size, step budget, sensitivity	$\lfloor k/(B+1) \rfloor$, $\varepsilon/(2B)$, 2θ (derived)

Each step costs $O(|Q_b|k)$ to compute the per-item scores plus the sampling cost of Section 2, so the whole algorithm runs in $O(nk)$ time. The data-reveal schedule is chosen to align with the online class of Bhamidi et al. (2026), as each batch b reads only indices in Q_b. This costs one data-independent initial slot, an $O(1/B)$ loss in performance. Table 2 collects the parameters and the values under which our guarantees hold.

3.3 Main result

We are now ready to present our main result.

Theorem 3.2 (Batched mechanism)**.** *Let $\theta = \theta_n := \max\{1, \sqrt{2\log(k \log n)}\}$ and suppose $\varepsilon\sqrt{k}/(\theta_n\sqrt{\log n}) \to \infty$. For fixed ε this holds whenever $k \gg \log n \cdot \log(k \log n)$, for example for any $k \ge (\log n)^{1+\gamma}$ with $\gamma > 0$. Choose $B \to \infty$ with $B \ll \varepsilon\sqrt{k}/(\theta_n\sqrt{\log n})$, $B \le \sqrt{k/3}$, and $\kappa = \lfloor k/(B+1) \rfloor \gg \theta_n^2 \log n$ (the last is automatic when $\varepsilon = O(1)$). Then $\mathcal{M}_{\mathrm{BIG}}$ is ε-DP for every real input and for every fixed $\delta > 0$*

$$\Pr\Big[\,\mathrm{Ave}(G_{I_B,J_B}) \;\ge\; \Big(\tfrac{2\sqrt{2}}{3} - \delta\Big)\eta_{\mathrm{OPT}}\Big] \;\longrightarrow\; 1,$$

with the same bound holding in expectation.

For these scalings, privacy is asymptotically free relative to the best known polynomial-time algorithm, and, under the conjecture of Bhamidi et al. (2026), relative to the best possible one. We also give a finite version of the result, which accounts for all of the lower-order terms.

Proposition 3.3 (Rate for Theorem 3.2, proved in Appendix E)**.** *Let $\zeta \in (0, \frac{1}{2}]$, $1 \le B \le \sqrt{k/3}$, and $\theta \ge 1$ with $\theta\sqrt{\log n} \le \frac{\zeta}{60}\sigma_\theta\sqrt{\kappa}$. Then for all large n, with probability at least $1 - 1/n$,*

$$\mathrm{Ave}(G_{I_B,J_B}) \;\ge\; (1-\zeta)\,\sigma_\theta\Big(1-\frac{2}{B}\Big)\frac{2\sqrt{2}}{3}\,\eta_{\mathrm{OPT}} - \frac{32\,\theta B\big(k\log(2en/k) + B\log n\big)}{\varepsilon k^2} - e^{-\theta^2/2} - \frac{6\log n}{Bk},$$

and the same bound holds in expectation up to an additive $o(\eta_{\mathrm{OPT}})$.

The first term is the value the batched greedy procedure captures, the second is the privacy penalty, which is of order B times that of a single exponential-mechanism draw, and the last two are the costs of clipping and of the initial slot. The ratio of the penalty to η_{OPT} is $O(\theta B\sqrt{\log n}/(\varepsilon\sqrt{k}))$, which the choice of B in Theorem 3.2 sends to zero. With $\theta = \theta_n$ the clipping term is $1/(k \log n)$, and the last term is $O(\log n/(Bk))$. The choices of B and θ are optimal only asymptotically, and at finite sizes a smaller B is better, since the early batches score candidates on few rows and are the first to be overwhelmed by the privacy noise (Section 6).

3.4 Why batching

Consider first privatizing the IGP directly, replacing each of its $2k-1$ arg-max steps by a call to $\mathcal{M}_{\mathrm{sel}}$ with $\kappa = 1$ and budget $\varepsilon/(2k-1)$. Each step loses about $\frac{2\Delta(2k-1)}{\varepsilon}\log n$ on its score, a total of order $\theta \log n/\varepsilon$ on the average, which is negligible relative to η_{OPT} only when $\varepsilon \gg \theta\sqrt{k \log n}$. Advanced composition (Dwork et al., 2010b) removes the $\sqrt{k}$ but leaves a factor $\sqrt{\log n \cdot \log(1/\delta)}$, so fixed ε remains out of reach.

Now let the IGP add rows and columns in B batches of $\kappa \approx k/(B+1)$, and privatize each batch by a single call to $\mathcal{M}_{\mathrm{sel}}$ on its pool of size m, with budget $\varepsilon/(2B)$. By Lemma 2.3 the score of a

κ-subset has sensitivity Δ regardless of κ, so the loss of the call is

$$\frac{2\Delta}{\varepsilon/(2B)} \log \binom{m}{\kappa} \;\approx\; 4\theta \cdot \frac{2B}{\varepsilon} \cdot \frac{k}{B} \log n \;=\; \Theta\Big(\frac{\theta k \log n}{\varepsilon}\Big),$$

so the factor k/B from the candidate count cancels the factor $2B$ from the budget split. Every batch costs what a single selection of k items costs, so the total over $2B$ steps is of order B times the one-shot penalty, while the value captured is a Riemann sum for the IGP's integral, $\frac{4}{3}k^{3/2}\sqrt{2\log n}(1 - O(1/B))$. Letting $B \to \infty$ slowly recovers the IGP constant at a privacy cost that is $o(1)$ relative to the signal.

3.5 What clipping costs

Clipping at a constant θ changes two things: the mechanism optimizes the clipped block sum rather than the true one, and the largest of many clipped column sums need not behave like the largest of Gaussian ones. Two lemmas handle these.

Lemma 3.4 (Clipping loss under monotone selection, Lemma C.12)**.** *Let I be a fixed row set, P a pool of columns, and $G_{I\times P}$ i.i.d. $N(0,1)$. Let a selection rule choose $S \subseteq P$ so that $\Pr[j \in S \mid G]$ is nondecreasing in the entries of column j. Then $\mathbb{E}\big[\sum_{j\in S}\sum_{i\in I}(-\theta - G_{ij})_+\big] \le \psi(\theta)\,|I|\,\mathbb{E}|S|$. $\mathcal{M}_{\mathrm{sel}}$ and the non-private top-κ rule both satisfy the hypothesis.*

The first lemma says that clipping can only hurt on entries clipped from below, since $G \ge \widetilde{G} - (-\theta - G)_+$ entrywise, and that a selection preferring large clipped sums does not prefer blocks with many such entries. The proof is Harris's inequality (Harris, 1960), since the selection probability of a column is nondecreasing in its entries and the below-clipping loss is nonincreasing. The loss per entry is thus at most $\psi(\theta) \le e^{-\theta^2/2}$, exponentially small in θ^2, while the sensitivity saved is linear in θ.

Lemma 3.5 (Top κ of m clipped sums, Lemma C.9)**.** *Let $S_1, \dots, S_m$ be i.i.d. sums of r clipped Gaussians, and let $\zeta \in (0, \frac{1}{2}]$, $\kappa = \exp(o(\log m))$, and $\theta\sqrt{\log(m/\kappa)} \le \frac{\zeta}{60}\sigma_\theta\sqrt{r}$. Then whp the sum of the κ largest S_j is at least $(1 - \frac{\zeta}{2})\kappa\sigma_\theta\sqrt{r}\sqrt{2\log(m/\kappa)}$.*

The second lemma says that the largest of m sums of r bounded terms are as large as if the terms were Gaussian with the same variance, provided $r \gg \theta^2 \log(m/\kappa)$. This is the moderate-deviation regime, in which the tail of a bounded sum has the Gaussian exponent to first order. The condition $r \gg \theta^2 \log n$ is not an extra requirement for fixed ε. For the one-shot mechanism the privacy penalty is negligible if and only if $k \gg \theta^2 \log n/\varepsilon^2$, which is the same condition. Appendix B discusses this further and explains why the same device fails under row-level privacy. Each step reads only the entries in its own group Q_b, which no earlier step has seen, so conditionally on the past its scores are sums of fresh clipped Gaussians and both lemmas apply (Lemma E.3). Every entry of the final block is read by one and only one step (Lemma E.4), so the block sum is the sum of the per-step scores and the argument of Section 3.4 goes through to give Proposition 3.3.

4 DP Bounds and the Small-Block Regime

Theorem 3.2 says nothing below its threshold $k \asymp \log n \cdot \log(k \log n)/\varepsilon^2$. This section determines, up to constants, what is achievable there. The answer has two halves, an information-theoretic bound that applies to every ε-DP mechanism, and a mechanism that attains it up to a constant. Together they show that the recoverable fraction of the optimum at fixed ε is $\Theta(\min\{1, \varepsilon\sqrt{k/\log n}\})$, so that nothing is recoverable for $k = o(\log n)$ and a constant fraction is recoverable once $k \asymp \log n/\varepsilon^2$.

4.1 The upper bound

We first establish an upper bound, a constant depending on ε, on the performance of any DP submatrix recovery algorithm. Let

$$C^*_\varepsilon := \sup_{a\in(0,1)} \frac{(e^\varepsilon - 1)\,\phi(\Phi^{-1}(1-a))}{1 + a(e^\varepsilon - 1)},$$

so that $C^*_\varepsilon = \frac{\varepsilon}{\sqrt{2\pi}}(1 + O(\varepsilon)) \approx 0.4\varepsilon$ as $\varepsilon \to 0$ and $C^*_\varepsilon = \Theta(\sqrt{\varepsilon})$ as $\varepsilon \to \infty$ (Lemma I.5).

Theorem 4.1 (Upper bound for every ε-DP mechanism, proved in Appendix G)**.** *For every ε-DP mechanism $\mathcal{A}$ outputting (I, J) with $|I| = |J| = k$, with no restriction on computation, $\mathbb{E}[\mathrm{Ave}(G_{\mathcal{A}(G)})] \le C^*_\varepsilon$.*

Notice that C^*_ε does not depend on n or k, only ε, while $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$ grows as k shrinks. The bound is therefore below η_{OPT} once $k \lesssim \log n/\varepsilon^2$, and in particular every ε-DP mechanism recovers a vanishing fraction of the optimum when $k \ll \log n/\varepsilon^2$.

The bound is proved one entry at a time. Fix an entry (i, j) and let $p(x)$ be the probability that it ends up in the output block when $G_{ij} = x$, averaged over the other entries. Changing this one entry produces a neighboring matrix, so privacy forces p to vary by at most a factor e^ε over all x (Dwork et al., 2015a; Bassily et al., 2016). In other words, how much the mechanism wants an entry can depend on the entry's value only a little. The entry's contribution to the expected block sum is $\mathbb{E}[Xp(X)]$ with $X \sim N(0, 1)$, and the best such p can do is to take a high value above some threshold and a low one below it, which gives exactly C^*_ε times the mean of p (Lemma I.4). Summing over the n^2 entries, of which k^2 are selected, gives $\mathbb{E}[\mathrm{Sum}] \le C^*_\varepsilon k^2$. The threshold shape is the same one that appears in local privacy (Kairouz et al., 2016; Duchi et al., 2018), where each record is likewise handled on its own.

4.2 The sign mechanism

The matching lower bound comes from the simplest mechanism in the paper. Let $Y_{ij} := \mathbf{1}[G_{ij} > 0]$ be the sign matrix and fix $I_0 = [k]$. The *sign mechanism* $\mathcal{M}_{\mathrm{sign}}$ is $\mathcal{M}_{\mathrm{sel}}(I_0, [n] \setminus I_0, k, \varepsilon)$ with the clipped block sum replaced by the sign count $\mathrm{Sum}(Y_{I_0,S})$, so that $\Delta = 1$ and $\Pr[J^* = S] \propto \exp(\frac{\varepsilon}{2}\mathrm{Sum}(Y_{I_0,S}))$. Changing one entry moves one sign count by at most one, so it is ε-DP by Lemma 2.3, and it runs in $O(nk)$ time. It is the $\theta \to 0$ limit of the one-shot mechanism.

Theorem 4.2 (Sign mechanism, proved in Appendix D)**.** *Let $\varepsilon' := \min\{\varepsilon, \sqrt{\log n/k}\}$. Then $\mathcal{M}_{\mathrm{sign}}$ run at budget ε' is ε-DP and*

$$\mathbb{E}[\mathrm{Ave}(G_{I_0,J^*})] \;\ge\; (1 - o(1))\sqrt{2/\pi}\,\tanh(\varepsilon'/4).$$

In particular, for $\varepsilon \le 1$, $\mathbb{E}[\mathrm{Ave}(G_{I_0,J^*})]/\eta_{\mathrm{OPT}} \ge (0.09 - o(1))\min\{\varepsilon\sqrt{k/\log n}, 1\}$. More generally the first bound holds at any budget $\varepsilon' \le \varepsilon$ with $\varepsilon'^2 k \le (4 - \delta)\log n$, and a finite-$n$ form is Lemma D.2.

Below the threshold $\varepsilon^2 k \asymp \log n$ the mechanism spends its whole budget and the guarantee is about 0.2ε, independent of k and $\log n$, against the upper bound $C^*_\varepsilon \approx 0.4\varepsilon$. Above it the mechanism deliberately spends only $\sqrt{\log n/k}$, since spending more buys nothing in order. The reason is that while a column's inclusion probability is small, flipping one of its signs multiplies it by about $e^{\varepsilon/2}$, so each entry of a selected column is reweighted by $e^{\pm\varepsilon/4}$ according to its sign, and a standard Gaussian reweighted this way has mean $\sqrt{2/\pi}\tanh(\varepsilon/4)$. Above the threshold the selected columns have inclusion probabilities near one, a flipped sign no longer moves them, and the value is capped at the non-private sign-selection value $\sqrt{2/\pi}\,\eta_{\mathrm{OPT}}/\sqrt{2}$, which decreases in k (Appendix H).

4.3 The phase diagram

Combining the two halves, and using Theorem 3.2 above its threshold, gives the following.

Corollary 4.3 (Phase diagram, proved in Appendix G)**.** *For fixed $\varepsilon \le 1$, the supremum V^* over ε-DP mechanisms of the expected value satisfies*

$$0.09\,\min\left\{\varepsilon\sqrt{k/\log n},\,1\right\} - o(1) \;\le\; V^*/\eta_{\mathrm{OPT}} \;\le\; \min\left\{0.35\,\varepsilon\sqrt{k/\log n},\,1 + o(1)\right\},$$

and $V^*/\eta_{\mathrm{OPT}} \ge \frac{2\sqrt{2}}{3} - o(1)$ once $\varepsilon^2 k \gg \log n \cdot \log(k \log n)$.

Thus $V^*/\eta_{\mathrm{OPT}} \asymp \min\{1, \varepsilon\sqrt{k/\log n}\}$, and both halves of the lower bound are attained by polynomial-time mechanisms, $\mathcal{M}_{\mathrm{sign}}$ and $\mathcal{M}_{\mathrm{BIG}}$. Asymptotically, nothing is recoverable for $k = o(\log n)$, a constant fraction $\Theta(\min\{1, \varepsilon\sqrt{t}\})$ is recoverable at $k = t\log n$, and the full online-optimal constant is recoverable once k exceeds $\log n$ by a $\log(k \log n)/\varepsilon^2$ factor.

In the language of statistical-to-computational gaps, privacy shifts the statistical threshold to $k \asymp \log n/\varepsilon^2$ while leaving the constant-factor computational gap of Gamarnik & Li (2018) untouched. Privacy does not close the gap from above either: we establish that the exponential mechanism over all submatrices attains η_{OPT} once $k \gg (\log n)^2/\varepsilon^2$ (Theorem F.1). So for sufficiently large k the statistical optimum is still reachable privately, just not efficiently, and the gap is the same as without privacy.

5 Row-Level Privacy

Now let neighbors differ in an entire row, so that a record is a subject contributing one row of G. This is the analogue of node privacy in the graph literature (Kasiviswanathan et al., 2013; Blocki et al., 2013). We use a one-shot style mechanism.

Unlike before, we select rows against fixed columns, since each row is one record. Fix the columns $J_0 := [k]$ and let $t_i := \sum_{j \in J_0} G_{ij}$ be the row sums. The *reverse one-shot* mechanism $\mathcal{M}_{\mathrm{RO}}$ is $\mathcal{M}_{\mathrm{sel}}(J_0, [n], k, \varepsilon)$ with the clipped block sum replaced by the row sums t_i clipped at $\tau := c\sqrt{k \log n}$ for a constant $c > \sqrt{2}$, so that whp no row sum is clipped. Changing one record moves one clipped score by at most 2τ, so $\Delta = 2\tau$ whatever k is, and $\mathcal{M}_{\mathrm{RO}}$ is ε-row-DP for every real input and runs in $O(nk)$ time. The analysis of Section 3.1 then gives the following.

Theorem 5.1 (Reverse one-shot mechanism, proved in Appendix I)**.** *If $\varepsilon = \omega(\log n)$, then $\mathcal{M}_{\mathrm{RO}}$ recovers $(1/\sqrt{2} - o(1))\eta_{\mathrm{OPT}}$ whp.*

This is a much stricter privacy condition, so it is unsurprising that the same style of mechanism needs a much larger ε for the same guarantee. One might hope that a more clever mechanism could improve the dependence on ε, but this turns out not to be the case.

Theorem 5.2 (Row-level upper bound, proved in Appendix I)**.** *Every ε-row-DP mechanism has $\mathbb{E}[\mathrm{Ave}] \le C^*_\varepsilon/\sqrt{k}$. Hence for every k, no mechanism recovers a constant fraction of η_{OPT} unless $\varepsilon = \Omega(\log n)$.*

The proof is the argument of Section 4 run on rows instead of entries. Privacy now constrains how the inclusion of row i depends on its own data, and what row i contributes to the block sum is $\langle X_i, \mathbf{1}_J \rangle \sim N(0, k)$, a Gaussian of scale $\sqrt{k}$ in place of 1, which is where the $1/\sqrt{k}$ comes from.

6 Experiments

We run four mechanisms from the paper along the two axes of the phase diagram: the sign mechanism $\mathcal{M}_{\mathrm{sign}}$, the one-shot $\mathcal{M}_{\mathrm{OS}}$ with clipping level $\theta = 2$, and the batched mechanism $\mathcal{M}_{\mathrm{BIG}}$ with $B = 2$ and with $B = 8$, both at $\theta = 2$. The left panel of Figure 1 fixes $n = 10^5$ and $k = 1000$ and sweeps the budget from $\varepsilon = 1/4$ to 64. The right panel fixes $\varepsilon = 2.5$ and sweeps the block size from $k = 100$ to 12800, which needs $n = 10^6$ to keep the block small against the matrix, since the theory has $k = \exp(o(\log n))$. The two sizes are comparable because every value is divided by that of the non-private incremental greedy procedure at the same n, k and θ, and this ratio depends on n only weakly, through $\log n$. Each mechanism is also run without privacy, with the top-κ arg-max in place of $\mathcal{M}_{\mathrm{sel}}$ on the same clipped scoring, which gives the ceiling it would reach with unlimited budget at that size. Every point is a mean over three independent matrices. Since each mechanism reads each entry at most once, on an index set that does not depend on the data (Lemma E.3), entries are drawn at the moment they are first read rather than the full matrix in advance, and the outcome has the same distribution either way. Further runs, including a sweep in k at $\varepsilon = 1$, the sign mechanism across k, the baselines and a check on the scaling of the penalty, are in Appendix H.

Fixed size, varying budget. The gap between each curve and the IGP splits into two parts. The gap up to the mechanism's own dashed line is the cost of privacy. It closes as ε grows for every mechanism, but at different rates: the sign mechanism and the one-shot spend their budget on a single selection and are near their ceilings at moderate ε, while $\mathcal{M}_{\mathrm{BIG}}$ splits its budget over $2B$ selections and needs a budget about $2B$ times larger for the same relative penalty, which is the price B pays in Proposition 3.3. The gap from the dashed line up to the IGP is structural and does not

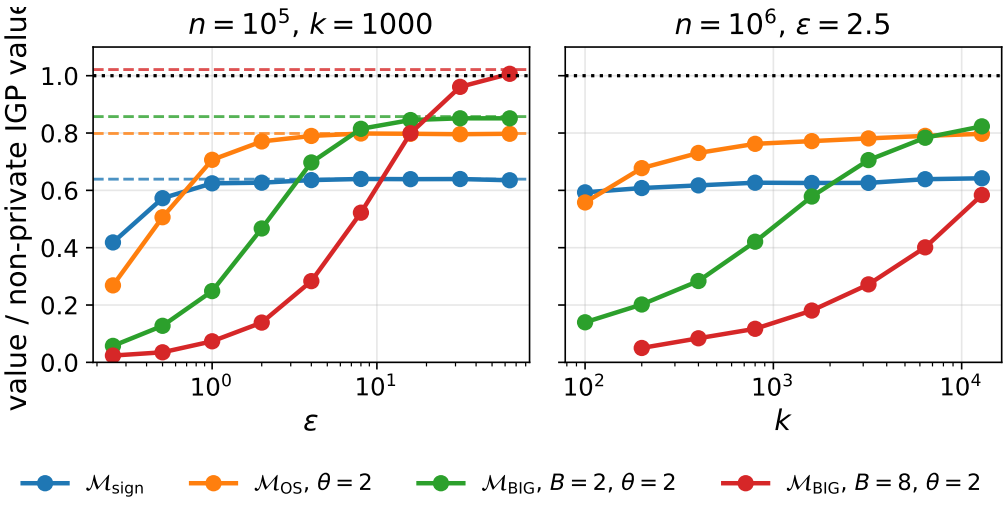


Figure 1: Value of four mechanisms relative to the non-private incremental greedy procedure at the same n, k and clipping level. Left: against the privacy budget at a fixed size. Right: against the block size at a fixed budget. Dashed lines on the left are the same mechanisms without privacy, the ceilings each would reach with unlimited budget. Dotted line is the IGP itself.

depend on ε. For the sign mechanism and the one-shot it is the price of adapting on one side only, the $1/\sqrt{2}$ of Theorem 3.1 against the $\frac{2\sqrt{2}}{3}$ of the IGP, plus the loss from ignoring magnitudes in the sign case. For $\mathcal{M}_{\text{BIG}}$ it is the data-independent initial slot, which shrinks with B as the $(1 - \frac{2}{B})$ factor of Proposition 3.3 says, and at $B = 8$ it has closed, since a batch that takes the best κ of one pool slightly beats κ separate best-of-one picks at finite n. The two costs pull in opposite directions in B, so the curves cross. At small budgets the sign mechanism is best, as Section 4 predicts for the small-block regime, and the one-shot takes over at intermediate budgets. The batched mechanisms win once ε is large enough for their penalty to vanish, first $B = 2$ and then $B = 8$, which ends at the IGP.

Fixed budget, varying size. The right panel tests the main theorem in its own direction. The privacy penalty of a batch does not grow with k while the signal does. So at a fixed budget every private curve should rise toward its non-private ceiling as k grows, and a mechanism with more selection steps should get there at a larger k. This is what the panel shows. The sign mechanism is at its ceiling throughout. The one-shot reaches its own, higher, ceiling within the first few points, since at this budget there is enough signal for magnitudes to matter. The batched mechanisms are still climbing at the end of the sweep, $B = 2$ ahead of $B = 8$, because the same relative penalty needs k larger by a factor of B^2. Over the sweep $B = 2$ overtakes first the sign mechanism and then the one-shot. $B = 8$ only catches up with the sign mechanism, and ends far below its ceiling at the IGP. By the B^2 scaling, reaching that ceiling would take k about sixteen times larger than for $B = 2$, beyond the sizes we can run. The crossovers of the left panel therefore also happen in k, at sizes that scale as $1/\varepsilon^2$. Both experiments are far from the limit $B \to \infty$, $n \to \infty$ of the theorem. We read them as a check on the relations the analysis predicts rather than on any constant.

7 Conclusion

For the densest $k \times k$ submatrix of a Gaussian matrix, entry-level privacy at a fixed budget costs nothing relative to the best known polynomial-time algorithm once $k \gg \log n \cdot \log(k \log n)/\varepsilon^2$, by a mechanism that is itself online and hence optimal in that class, and the recoverable fraction at fixed ε is $\Theta(\min\{1, \varepsilon\sqrt{k/\log n}\})$. The enabling ideas, batching and constant clipping, do not help under row-level privacy. The constants in the small-block regime and the picture under approximate differential privacy remain open.

AI USE STATEMENT

Generative AI tools were used extensively in this work, in an interactive setting in which the authors set the research questions, chose the mechanisms and the model, decided what to prove and what to run, and directed every revision. Within that setting the tools were used for the following tasks with required disclosure: assisting with the derivation and drafting of proofs, writing and editing the manuscript text, writing the experiment, figure and verification code, and conducting literature search. They were also used to audit the manuscript for internal consistency and to suggest experiments, from which the authors selected. We have not used generative AI tools for generating experimental data or figures from anything other than the provided code, and the remaining tasks with required disclosure are not applicable. We have reviewed all AI-assisted work. Every proof was checked line by line by the authors, the moment-generating-function and tilted-moment inequalities in the moderate-deviation lemma, the two identities used in Appendix I, and the decomposition of Lemma E.4 were additionally verified numerically, the code was tested against brute-force enumeration for the samplers and against a materialized matrix for the lazy simulation, and every citation was verified against the cited source or its venue listing. We take responsibility for the final content of this work, including text, claims and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This is a theoretical study on synthetic Gaussian random matrices. No human-subject data, personal data, or real datasets were used. Its subject is differential privacy, and we have been careful to state the privacy model precisely. The main results protect individual entries, which does not protect a subject who contributes a whole row, and we analyze row-level privacy separately and state plainly that the guarantees are asymptotic and do not constitute a deployable private biclustering method. We see no potential for harmful application beyond the general dual-use character of privacy analysis, and no conflicts of interest.

REPRODUCIBILITY STATEMENT

All theoretical claims are proved in full in Appendices C–J, with assumptions stated in Section 2. The experiments use no external data. Instances are i.i.d. Gaussian matrices generated from fixed seeds, the exact simulation scheme is justified in Section 6 and Appendix H, and all numerical values in the paper are read from the JSON result files produced by the supplied code (`dp_submatrix.py`, with a README giving the commands and approximate running times). The supplementary material contains the code, the result files, and the figure-generation scripts.

REFERENCES

K. Amin, A. Kulesza, A. Muñoz Medina, and S. Vassilvitskii. Bounding user contributions: a bias-variance trade-off in differential privacy. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, PMLR 97, 2019.

M. Bafna and J. Ullman. The price of selection in differential privacy. In *Proceedings of the 30th Conference on Learning Theory (COLT)*, PMLR 65, 2017.

B. Balle and Y.-X. Wang. Improving the Gaussian mechanism for differential privacy: analytical calibration and optimal denoising. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, PMLR 80, pages 394–403, 2018.

R. Bassily, K. Nissim, A. Smith, T. Steinke, U. Stemmer, and J. Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1046–1059, 2016.

A. Beimel, K. Nissim, and U. Stemmer. Private learning and sanitization: pure vs. approximate differential privacy. In *Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM)*, LNCS 8096, pages 363–378, 2013.

S. Bhamidi, P. S. Dey, and A. B. Nobel. Energy landscape for large average submatrix detection problems in Gaussian random matrices. *Probability Theory and Related Fields*, 168:919–983, 2017.

S. Bhamidi, D. Gamarnik, and S. Gong. Finding a dense submatrix of a random matrix. Sharp bounds for online algorithms. *Electronic Communications in Probability*, 31, arXiv:2507.19259, 2026.

R. Bhaskar, S. Laxman, A. Smith, and A. Thakurta. Discovering frequent patterns in sensitive data. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 503–512, 2010.

J. Blocki, A. Blum, A. Datta, and O. Sheffet. The Johnson–Lindenstrauss transform itself preserves differential privacy. In *53rd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 410–419, 2012.

J. Blocki, A. Blum, A. Datta, and O. Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 87–96, 2013.

C. Borgs, J. Chayes, and A. Smith. Private graphon estimation for sparse graphs. In *Advances in Neural Information Processing Systems 28 (NeurIPS)*, 2015.

C. Borgs, J. Chayes, A. Smith, and I. Zadik. Revealing network structure, confidentially: improved rates for node-private graphon estimation. In *59th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 533–543, 2018.

M. Bun, J. Ullman, and S. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1–10, 2014.

H. Chen, V. Cohen-Addad, T. d'Orsi, A. Epasto, J. Imola, D. Steurer, and S. Tiegel. Private estimation algorithms for stochastic block models and mixture models. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, 2023.

H. Chen, J. Ding, T. d'Orsi, Y. Hua, C.-H. Liu, and D. Steurer. Private graphon estimation via sum-of-squares. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1112–1123, 2024a.

H. Chen, J. Ding, Y. Hua, and D. Steurer. Private edge density estimation for random graphs: optimal, efficient and robust. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, 2024b.

S. Chen and S. Zhou. Recursive mechanism: towards node differential privacy and unrestricted joins. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pages 653–664, 2013.

X.-H. Chen, A. P. Dempster, and J. S. Liu. Weighted finite population sampling to maximize entropy. *Biometrika*, 81(3):457–469, 1994.

R. Cummings, V. Feldman, A. McMillan, and K. Talwar. Mean estimation with user-level privacy under data heterogeneity. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022.

W.-Y. Day, N. Li, and M. Lyu. Publishing graph degree distribution with node differential privacy. In *Proceedings of the 2016 ACM SIGMOD International Conference on Management of Data*, pages 123–138, 2016.

Inderjit S. Dhillon. Co-clustering documents and words using bipartite spectral graph partitioning. In *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 269–274, 2001.

L. Dhulipala, Q. C. Liu, S. Raskhodnikova, J. Shi, J. Shun, and S. Yu. Differential privacy from locally adjustable graph algorithms: k-core decomposition, low out-degree ordering, and densest subgraphs. In *63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 754–765, 2022.

L. Dhulipala, G. Z. Li, and Q. C. Liu. Near-optimal differentially private k-core decomposition. arXiv:2312.07706, 2023.

M. Dinitz, S. Kale, S. Lattanzi, and S. Vassilvitskii. Almost tight bounds for differentially private densest subgraph. In *Proceedings of the 2025 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2908–2950, 2025.

J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521):182–201, 2018.

D. Durfee and R. Rogers. Practical differentially private top-k selection with pay-what-you-get composition. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019.

C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference (TCC)*, LNCS 3876, pages 265–284, 2006.

C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 715–724, 2010a.

C. Dwork, G. N. Rothblum, and S. Vadhan. Boosting and differential privacy. In *51st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 51–60, 2010b.

C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth. Preserving statistical validity in adaptive data analysis. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pages 117–126, 2015a.

C. Dwork, A. Smith, T. Steinke, J. Ullman, and S. Vadhan. Robust traceability from trace amounts. In *56th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 650–669, 2015b.

Marek Eliáš, Michael Kapralov, Janardhan Kulkarni, and Yin Tat Lee. Differentially private release of synthetic graphs. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 560–578, 2020.

A. Farhadi, M. Hajiaghayi, and E. Shi. Differentially private densest subgraph. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR 151:11581–11597, 2022.

D. Gamarnik. The overlap gap property: a topological barrier to optimizing over random structures. *Proceedings of the National Academy of Sciences*, 118(41):e2108492118, 2021.

D. Gamarnik and Q. Li. Finding a large submatrix of a Gaussian random matrix. *The Annals of Statistics*, 46(6A):2511–2561, 2018.

D. Gamarnik, A. Jagannath, and A. S. Wein. Low-degree hardness of random optimization problems. In *61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 131–140, 2020.

Thomas George and Srujana Merugu. A scalable collaborative filtering framework based on co-clustering. In *Proceedings of the 5th IEEE International Conference on Data Mining (ICDM)*, pp. 625–628, 2005.

K. Georgiev and S. B. Hopkins. Privacy induces robustness: information-computation gaps and sparse mean estimation. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022.

B. Ghazi, R. Kumar, and P. Manurangsi. User-level differentially private learning via correlated sampling. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, 2021.

J. Gillenwater, M. Joseph, A. Muñoz Medina, and M. R. Diaz. A joint exponential mechanism for differentially private top-k. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, PMLR 162, 2022.

A. Gupta, K. Ligett, F. McSherry, A. Roth, and K. Talwar. Differentially private combinatorial optimization. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1106–1125, 2010.

A. Gupta, A. Roth, and J. Ullman. Iterative constructions and private data release. In *Theory of Cryptography Conference (TCC)*, LNCS 7194, pages 339–356, 2012.

J. Hájek. Asymptotic theory of rejective sampling with varying probabilities from a finite population. *The Annals of Mathematical Statistics*, 35(4):1491–1523, 1964.

M. Hardt and K. Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 705–714, 2010.

T. E. Harris. A lower bound for the critical probability in a certain percolation process. *Mathematical Proceedings of the Cambridge Philosophical Society*, 56(1):13–20, 1960.

M. Hay, C. Li, G. Miklau, and D. Jensen. Accurate estimation of the degree distribution of private networks. In *Ninth IEEE International Conference on Data Mining (ICDM)*, pages 169–178, 2009.

Bryan Hooi, Hyun Ah Song, Alex Beutel, Neil Shah, Kijung Shin, and Christos Faloutsos. FRAUDAR: Bounding graph fraud in the face of camouflage. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 895–904, 2016.

P. Kairouz, S. Oh, and P. Viswanath. Extremal mechanisms for local differential privacy. *Journal of Machine Learning Research*, 17(17):1–51, 2016.

G. Kamath, A. Mouzakis, and V. Singhal. New lower bounds for private estimation and a generalized fingerprinting lemma. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022.

S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. Smith. Analyzing graphs with node differential privacy. In *Theory of Cryptography Conference (TCC)*, LNCS 7785, pages 457–476, 2013.

A. Khayatian, A. Vullikanti, and A. Konar. Differentially private Densest-k-Subgraph. arXiv:2505.03858, 2025.

D. Levy, Z. Sun, K. Amin, S. Kale, A. Kulesza, M. Mohri, and A. T. Suresh. Learning with user-level privacy. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, 2021.

Y. Liu, A. T. Suresh, F. Yu, S. Kumar, and M. Riley. Learning discrete distributions: user vs. item-level privacy. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020.

S. C. Madeira and A. L. Oliveira. Biclustering algorithms for biological data analysis: a survey. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 1(1):24–45, 2004.

R. McKenna and D. Sheldon. Permute-and-flip: a new mechanism for differentially private selection. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020.

F. McSherry and I. Mironov. Differentially private recommender systems: building privacy into the Netflix Prize contenders. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 627–636, 2009.

F. McSherry and K. Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 94–103, 2007.

M. Mitrovic, M. Bun, A. Krause, and A. Karbasi. Differentially private submodular maximization: data summarization in disguise. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, PMLR 70, pages 2478–2487, 2017.

M. S. Mohamed, D. Nguyen, and A. Vullikanti. Differentially private community detection for stochastic block models. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, PMLR 162, 2022.

S. Narayanan. Better and simpler lower bounds for differentially private statistical estimation. arXiv:2310.06289, 2023.

D. Nguyen and A. Vullikanti. Differentially private densest subgraph detection. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, PMLR 139:8140–8151, 2021.

K. Nissim, S. Raskhodnikova, and A. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 75–84, 2007.

N. Peters, E. Tsfadia, and J. Ullman. Smooth lower bounds for differentially private algorithms via padding-and-permuting fingerprinting codes. In *Proceedings of the 37th Conference on Learning Theory (COLT)*, PMLR 247, 2024.

G. Qiao, W. J. Su, and L. Zhang. Oneshot differentially private top-k selection. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, PMLR 139:8672–8681, 2021.

S. Raskhodnikova and A. Smith. Lipschitz extensions for node-private graph statistics and the generalized exponential mechanism. In *57th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 495–504, 2016.

A. Sealfon and J. Ullman. Efficiently estimating Erdős–Rényi graphs with node differential privacy. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 3770–3780, 2019.

A. A. Shabalin, V. J. Weigman, C. M. Perou, and A. B. Nobel. Finding large average submatrices in high dimensional data. *The Annals of Applied Statistics*, 3(3):985–1012, 2009.

T. Steinke and J. Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2):3–22, 2016.

T. Steinke and J. Ullman. Tight lower bounds for differentially private selection. In *58th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 552–563, 2017.

F. Zhou. Continual release of densest subgraphs: privacy amplification and sublinear space via subsampling. arXiv:2510.11640, 2025.

A Extended Related Work

This appendix expands the related work of the introduction, and it connects the row-level results of Section 5 and Appendix I to the literature on private graph analysis and user-level privacy. We have tried to give enough context that it can be read on its own.

Differentially private selection The exponential mechanism of McSherry & Talwar (2007) is the basic tool throughout this work. Its utility guarantee, a loss of order $\Delta \log N/\varepsilon$ against the best of N candidates, is where every penalty term in the paper comes from. The closest classical problem is private top-k selection, where one wants the k items with the largest counts out of d. Bhaskar et al. (2010) introduced a peeling mechanism that runs the exponential mechanism k times, Durfee & Rogers (2019) showed that peeling can be done in one shot with Gumbel noise, Qiao et al. (2021) studied one-shot Laplace noise, McKenna & Sheldon (2020) proposed permute-and-flip, and Gillenwater et al. (2022) sample directly from an exponential mechanism over sequences of length k. The lower bounds of Bafna & Ullman (2017) and Steinke & Ullman (2017) show that a loss of order $\log d/\varepsilon$ is unavoidable for private selection, and our penalties have this form too.

What differs in our setting is the sensitivity. In top-k selection a single record can move every count, so the score of a k-subset has sensitivity growing with k, and one pays for k separate selections one way or another. Under entry-level privacy a record is one matrix entry, which appears in a single term of the block sum, so the score of a subset has the same sensitivity however large the

subset is. This is why the one-shot mechanism needs no composition at all, and it is what makes the batching of Section 3 work. Sampling the resulting distribution over subsets is the classical problem of conditional Poisson design, which goes back to Hájek (1964) and Chen et al. (1994), and Proposition C.15 gives a self-contained sampler for completeness. Composition itself is standard (Dwork et al., 2006; Dwork & Roth, 2014; Dwork et al., 2010b). Advanced composition would not have been enough for a step-by-step privatization of the greedy procedure, since it still leaves a $\sqrt{\log n \cdot \log(1/\delta)}$ factor in the threshold, whereas batching removes the dependence on k entirely and stays within pure differential privacy.

Private greedy algorithms $\mathcal{M}_{\mathrm{BIG}}$ is a private greedy algorithm in the style of Gupta et al. (2010), who replace each choice of a greedy procedure by an exponential mechanism. It shares with the private submodular maximization of Mitrovic et al. (2017) the feature that a subset's score has sensitivity independent of the subset's size under element-level privacy. In both of these works the privacy budget is still spent one greedy step at a time. The observation that a whole batch of greedy choices can be made at the price of one, so that the composition cost scales with the number of batches rather than the number of steps, does not appear there, and we believe it is new.

Where the upper bound comes from The inequality $\sup_x p(x) \le e^{\varepsilon} \inf_x p(x)$ behind both of our upper bounds is a very simple use of the privacy definition. It is the one-record stability fact of Dwork et al. (2015a) and Bassily et al. (2016), where it shows that a private mechanism's expected interaction with any one record can change by at most a factor e^{ε}. What we add is the rearrangement argument that makes the constant sharp for a Gaussian record. The maximizer in that argument is a two-level function, which is the shape of the staircase mechanisms that Kairouz et al. (2016) show to be optimal for local differential privacy. This is no coincidence. In the small-block regime the exponential mechanism is essentially reweighting each entry on its own, so entry-level privacy behaves like local privacy there, and the one-record correlation bounds of Duchi et al. (2018) are the natural point of comparison. We suspect that the sign mechanism is in fact optimal among one-shot mechanisms in this regime, and that the remaining gap in Corollary 4.3 lies in the upper bound of Theorem 4.1 rather than in the mechanism, but we have not pursued this.

Our upper bounds are bounds on the correlation between the output and a single record of the actual random instance. In this sense they are closer to the tracing and fingerprinting arguments behind most approximate-DP lower bounds (Bun et al., 2014; Dwork et al., 2015b; Steinke & Ullman, 2016; Kamath et al., 2022) than to packing arguments (Hardt & Talwar, 2010), though we use the pure-DP constraint in the most elementary way possible, which is what makes the constant explicit. The bounds do not see n or the number of candidates at all, so under entry-level privacy they say something only for k up to about $\log n/\varepsilon^2$, where the sign mechanism matches them. Under row-level privacy this blindness costs nothing, since the threshold there does not depend on k. Extending the row-level bound to (ε, δ)-privacy would presumably need fingerprinting-type machinery (Narayanan, 2023; Peters et al., 2024) in place of the pointwise inequality, since an additive δ per output cannot simply be summed over all $\binom{n}{k}^2$ outputs.

Hardness for private mechanisms The constant $\frac{2\sqrt{2}}{3}$ is believed to be optimal for polynomial-time algorithms, on the basis of the overlap-gap property (Gamarnik, 2021) and low-degree hardness (Gamarnik et al., 2020), and Bhamidi et al. (2026) prove it for online algorithms. One might hope that privacy allows a hardness result of its own, for instance through the privacy-implies-robustness reduction of Georgiev & Hopkins (2022). It does not seem to. Our failure probabilities are polynomially rather than exponentially small, so the reduction gives only a small robustness radius, and in any case there is no planted signal for an adversary to attack. Stability-based arguments do not help either. The full-space exponential mechanism of Appendix F is ε-DP for every ε and still attains the optimum, so differential privacy is too weak a notion of stability to trigger overlap-gap obstructions. The best private polynomial-time constant therefore cannot be determined without resolving the non-private conjecture.

Private graph algorithms and planted models Edge-level privacy for graph statistics goes back to Nissim et al. (2007), and node-level privacy to Kasiviswanathan et al. (2013) and Blocki et al. (2013). The problem closest to ours is differentially private densest subgraph. Nguyen & Vullikanti (2021) gave the first edge-private algorithms, Farhadi et al. (2022) gave upper and lower bounds with

additive error, Dhulipala et al. (2022) worked in the local edge-privacy model, Dinitz et al. (2025) obtained almost tight purely additive bounds, Dhulipala et al. (2023) treated k-core decomposition, Zhou (2025) continual release, and Khayatian et al. (2025) the size-constrained densest-k-subgraph problem through the leading eigenvector of the adjacency matrix. All of these are worst-case results on $0/1$ inputs with additive error. None engages with the extreme-value scale $\sqrt{\log n/k}$ that governs the random-matrix problem, where an additive error of Δ/ε per entry is already as large as the signal. We are not aware of node-private results for densest subgraph. Private recovery in statistical graph models has been studied for the stochastic block model by Mohamed et al. (2022) and Chen et al. (2023). Those works combine a random model with edge-level privacy as we do, but their goal is to recover a planted structure from noise, while here there is no planted signal and the objective is an extremal statistic of pure noise. We are not aware of prior work on differential privacy for extremal submatrix or biclustering problems in random matrices.

Row-level privacy The rest of this appendix concerns the row-level model of Section 5 and Appendix I, in which two matrices are neighbors when they differ in a whole row, so that a record is a subject who contributes a row. Two things are shown there. First, every ε-row-DP mechanism has expected value at most $C^*_\varepsilon/\sqrt{k}$, with the same per-record constant C^*_ε as in the entry-level bound (Theorem 5.2). Since η_{OPT} also scales like $1/\sqrt{k}$, a constant fraction of the optimum is out of reach unless ε is at least of order $\log n$, whatever k is. Second, this threshold is the right one up to constants (Theorem 5.1): the reverse one-shot mechanism, which fixes the columns, clips each row's sum over those columns at the scale $\sqrt{k \log n}$, and picks the rows with one exponential mechanism, recovers a $1/\sqrt{2}$ fraction of the optimum once $\varepsilon \gg \log n$. The contrast with entry-level privacy, where the threshold is $k \asymp \log n/\varepsilon^2$ and so large k helps, is the point of what follows.

Edge privacy versus node privacy View G as a weighted complete bipartite graph between rows and columns. Entry-level privacy is then edge privacy, the notion under which private graph analysis began (Nissim et al., 2007; Hay et al., 2009) and under which private cut and synthetic graph release have been developed (Gupta et al., 2012; Blocki et al., 2012; Eliáš et al., 2020), and row-level privacy is node privacy for the nodes on the row side (Kasiviswanathan et al., 2013; Blocki et al., 2013; Chen & Zhou, 2013). The graph literature has known from the start that node privacy is much more expensive than edge privacy, because a single node touches many edges and even a subgraph count then has sensitivity of the order of the maximum degree. Appendix I quantifies this for one problem in one random model. The two upper bounds make the source of the gap explicit. Both bound the correlation between the output and a single record, and the per-record bound is the same in the two models. What differs is how many records make up the output and how much each carries. Under edge privacy there are k^2 entries of unit scale, giving a bound of $k^2 C^*_\varepsilon$ on the expected block sum, while under node privacy there are k rows of scale $\sqrt{k}$, giving $k^{3/2} C^*_\varepsilon$. The factor of $\sqrt{k}$ between the two is what lets a large block size rescue a fixed budget in one model and not the other.

Degree bounds and score clipping The usual way to obtain node-private algorithms is to force the sensitivity down. One projects the graph onto a graph of bounded degree D, or extends the statistic from bounded-degree graphs to all graphs by a Lipschitz extension (Kasiviswanathan et al., 2013; Blocki et al., 2013; Day et al., 2016; Raskhodnikova & Smith, 2016). The sensitivity then scales with D, and the nodes of degree larger than D are lost. Clipping the row sums in the reverse one-shot mechanism is the same idea in our setting, and in the user-level privacy literature the same operation is called contribution bounding (Amin et al., 2019). What the row-level upper bound adds is that here the projection threshold is forced on us. The rows that carry the objective are the extreme ones, with row sums of order $\sqrt{k \log n}$, so the clipping level must be at that scale for the mechanism to see them, and the upper bound says that no cleverer mechanism avoids paying for it. The $\sqrt{\log n}$ gap between this scale and the typical row sum $\sqrt{k}$ is what becomes the threshold $\varepsilon \asymp \log n$. This is a sharper form of the familiar fact that node-private algorithms pay for the degree bound and lose the high-degree nodes. In our problem the high-degree nodes are the whole point.

Node privacy in random graphs, and user-level privacy Node-private estimation in random graph models is probably the closest existing setting. Borgs et al. (2015) and Borgs et al. (2018) gave node-private graphon estimators and optimal rates for $G(n, p)$ and block models, Sealfon & Ullman (2019) gave an efficient estimator of the edge density, and Chen et al. (2024a;b) obtained efficient and robust estimators through sum-of-squares. In all of these the target is a global parameter, so the

influence of one node is diluted by averaging over all n nodes, and node privacy costs only a lower-order term. The densest-submatrix objective is the opposite kind of statistic. It is decided by the k most extreme rows, each of which is one record, and no averaging is available. The same distinction shows up wherever a matrix has people on one side and items on the other. It is the distinction between rating-level and user-level privacy in private recommender systems (McSherry & Mironov, 2009), between event-level and user-level privacy under continual observation (Dwork et al., 2010a), and it is the subject of a recent line of work on user-level learning (Liu et al., 2020; Levy et al., 2021; Ghazi et al., 2021; Cummings et al., 2022), where user-level guarantees are consistently found to be much more expensive unless each user's contribution is bounded. In our terms rating-level is entry-level and user-level is row-level, and the row-level threshold is a case where the extra cost is not a constant or logarithmic factor but a shift of the threshold by a factor of roughly $\sqrt{k \log n}$. The reverse one-shot mechanism, which chooses the items first and then privately chooses the users, is the natural way to use the fact that the items are not records. It has no analogue under edge privacy, where the two sides are symmetric.

Why only the rows are chosen adaptively Under row-level privacy the column steps of $\mathcal{M}_{\mathrm{BIG}}$ become the classical private top-k selection problem, since a single row now moves every column score, just as a single user moves every count in Bhaskar et al. (2010); Durfee & Rogers (2019); Qiao et al. (2021); Gillenwater et al. (2022). The lower bounds of Bafna & Ullman (2017) and Steinke & Ullman (2017) for that problem, which descend from the packing arguments of Hardt & Talwar (2010) and the separations of Beimel et al. (2013), can be read as saying that adaptive column selection cannot be made cheap under node privacy. This is why Appendix I selects only the rows adaptively and fixes the columns. Whether adapting on both sides can beat the constant $1/\sqrt{2}$ under pure row-level privacy is open. A direct privatization of the batched mechanism would need a budget growing with k for its column steps, whose sensitivity is now $2\theta\kappa$, and we do not know whether this is necessary.

Node-private densest subgraph Finally, the row-level results suggest what to expect for node-private densest subgraph in a random model, a problem we have not seen studied: a threshold of order $\log n$ in the budget that does not depend on k, matched by a mechanism that fixes one side of the bipartition and privately selects the other, with an upper bound obtained by bounding the correlation between the output and a single node. Whether this is what happens in $G(n, p)$ or in the planted models of Mohamed et al. (2022) and Chen et al. (2023), where a node's degree rather than its row sum is the relevant statistic, is a natural question that our argument does not address.

B Why Constant Clipping Suffices

This appendix expands on Section 3.5. It explains in words, without new lemmas, why the constant clipping level behind Theorems 3.1 and 3.2 works, why it is perhaps surprising that it does, and why the same trick is not available under the row-level privacy of Section 5. All of our mechanisms clip the entries of G at some level θ and then run an exponential mechanism whose sensitivity is 2θ. The obvious choice is θ a little above $\sqrt{2 \log n}$. Then no entry is clipped whp, the mechanism can be analyzed as if it ran on G itself, and the utility argument is the non-private one plus a penalty term. We instead take θ constant, or at most of order $\sqrt{\log(k \log n)}$. The notation is that of Sections 2 to 3, and the lemmas we point to are for the reader who wants the proofs.

The mismatch The exponential mechanism pays roughly $\Delta \log N/\varepsilon$ for a choice among N candidates, where Δ is the most a single record can change any candidate's score. The score of a column is a sum of k entries, and a typical entry is of constant size, so the typical influence of one record on a score is a constant. But privacy is a worst-case guarantee. The mechanism has to be charged for the largest entry it could ever see, which after clipping is θ. Clipping at $\sqrt{2 \log n}$ sets θ to the largest entry that will in fact be seen among the n^2 entries of a Gaussian matrix, so the mechanism is charged a factor $\sqrt{\log n}$ more than the typical influence, on each of the roughly $k \log n$ units of selection it performs. Corollary 4.3 shows what this would forfeit. It would move the threshold for full recovery from about $k \asymp \log n$ up to $k \asymp (\log n)^2$, and in between, where the sign mechanism already recovers a constant fraction of the optimum, the exponential mechanism at that clipping level would have no guarantee at all.

What changes with a constant θ, and why it does not matter With θ constant a constant fraction of the entries are clipped, and two things change. The mechanism now optimizes the wrong objective, the clipped block sum instead of the true one. And the extreme-value analysis, which asks how large the largest of n column sums is, was carried out for Gaussian sums, while clipped sums are not Gaussian.

The first change is harmless, because clipping only ever helps on the entries that matter. An entry clipped from above loses value in the score but not in the truth, since the true block sum is larger than the clipped one on such entries. Entries clipped from below are the only ones where the truth is worse than the score, and there the loss is $(-\theta - G_{ij})_+$, whose mean $\psi(\theta)$ is at most $e^{-\theta^2/2}$. The worry is that the mechanism, in maximizing clipped sums, might somehow be drawn toward blocks with many of these entries. It is not, and the reason is monotonicity. The probability that a column is selected is a nondecreasing function of its entries, the loss from below-clipping is a nonincreasing function of the same entries, and for independent entries two such functions are negatively correlated. This is Harris's inequality, Lemma C.11. So in expectation the selected block has no more below-clipped entries than a random block would, and the loss is at most $\psi(\theta)$ per entry (Lemma C.12). This is exponentially small in θ^2, while the sensitivity we save is linear in θ. With θ of order $\sqrt{2\log(k\log n)}$ the loss per entry is $1/(k\log n)$, which is negligible against $\sqrt{\log n/k}$, and the sensitivity is of order $\sqrt{\log(k\log n)}$, which is $\sqrt{\log\log n}$ for polylogarithmic k, rather than $\sqrt{\log n}$.

The second change is the substantive one, and it is where the structure of the problem enters. The objective is determined by the largest column sums, which sit about $\sqrt{2\log n}$ standard deviations above the mean of a sum of k terms. A sum of k bounded terms of size θ can reach that many standard deviations only if $\sqrt{2\log n}$ is small compared to $\sqrt{k}/\theta$, that is, only if k is large compared to $\theta^2\log n$. This is the moderate-deviation regime, in which sums of bounded variables still have Gaussian tails to first order in the exponent. Lemma C.8 is the version we use. Below this regime the tails are governed by the bound θ rather than by the variance, and the extreme columns are simply not reachable by clipped sums. So the analysis needs $k \gg \theta^2\log n$. The point is that privacy needs this condition anyway, since the privacy penalty is negligible if and only if $k \gg \theta^2\log n/\varepsilon^2$. The two conditions coincide because they are the same requirement. A budget of ε can only distinguish scores that differ by about θ/ε per unit of $\log N$, and the extreme columns are distinguishable from typical ones by about $\sqrt{2k\log n}$, so both conditions say that θ has to be small compared to the standard deviation of a column sum divided by $\sqrt{\log n}$. What we found surprising is that nothing else is needed. One might have thought that clipping a constant fraction of the entries would distort which columns are extreme, but at the scale $\sqrt{2k\log n}$ a column's rank is decided by the bulk of its entries and not by a few large ones, and the bulk is untouched.

The sign mechanism as a limiting case Push θ to zero and rescale. Each entry then contributes only its sign, and we obtain the sign mechanism of Section 4.2. Below the condensation threshold $\varepsilon^2 k \lesssim 4\log n$ its value is $\sqrt{2/\pi}\tanh(\varepsilon/4)$, with no slack in ε and independent of k, because in that regime the exponential mechanism reweights each entry by $e^{\pm\varepsilon/4}$ according to its sign, and a standard Gaussian reweighted this way has precisely that mean. The sign mechanism cannot reach the extreme-value regime. Its score has variance k, but a sum of signs only reaches $\sqrt{2\log n}$ standard deviations when $k \gg \log n$, and once the mechanism concentrates its value saturates at $\sqrt{2/\pi}$ times the non-private one-shot value, which is the correlation between an entry and its sign. So the sign mechanism owns the small-block regime and clipping at a slowly growing θ owns the extreme regime. The two meet at $k \asymp \log n/\varepsilon^2$, where the values $\sqrt{2/\pi}\tanh(\varepsilon/4)$ and $\sqrt{2/\pi}\sqrt{2\log n/k}$ are of the same order. Figure 4a shows this crossover.

Why the trick fails under row-level privacy In the row model a record is an entire row, and the natural mechanism scores rows by their sums over k fixed columns. A row score is now a single record's worth of data rather than a sum over many records. If we clip the row score at $\theta\sqrt{k}$ for a constant θ, the sensitivity becomes $2\theta\sqrt{k}$, but the extreme rows have scores of order $\sqrt{2k\log n}$, which is much larger, and they are all clipped to the same value. The mechanism can then no longer tell which of the many rows above the threshold are the extreme ones. In the language of the previous paragraphs, there is no sum whose moderate deviations can carry the extreme value, and the condition $k \gg \theta^2\log n$ has nothing to act on. Constant clipping in the row model only

gives the small-block regime, a value of order $\varepsilon/\sqrt{k}$, the row analogue of the sign mechanism's ε. Reaching the extreme regime requires clipping at the extreme scale $\sqrt{k \log n}$, a sensitivity that is $\sqrt{\log n}$ above the typical row score, and this is why the row-level threshold is $\varepsilon \asymp \log n$ and does not depend on k. Theorem 5.2 shows that this is not an artifact of the mechanism. With k records each carrying a signal of scale $\sqrt{k}$, the per-record correlation bound caps the value at $C^*_\varepsilon/\sqrt{k}$ however one clips. This is the whole difference between the two models. Under entry-level privacy the records are individually small and there are many of them, so their sum can be extreme while each one is bounded. Under row-level privacy the record is itself the extreme object.

What this means in practice The penalty is linear in θ, so relative to clipping at $c\sqrt{\log n}$ the saving is a factor of $c\sqrt{\log n}/\theta$, which is about 3.4 at $n = 10^5$ and about 4 at $n = 10^7$ for $c = 2$ and $\theta = 2$. This is a factor by which the effective privacy budget is multiplied, or equivalently a factor of its square by which the block size needed for a given recovery is divided. There is no cost in privacy, since the guarantee is pure ε-DP either way, no cost in running time, and by Lemma C.12 a loss in utility of at most $\psi(\theta)$ per entry, which is already below 0.01 at $\theta = 2$. The experiments of Section 6 run at $\theta = 2$.

C Preliminary Results

This appendix proves the results stated in Section 2, namely the upper bound on the optimum, the properties of σ_θ and ψ, the privacy facts behind Lemma 2.3, and the samplers of Proposition C.15. It also proves the two clipping lemmas of Section 3.5, Lemmas 3.4 and 3.5, and collects a few technical lemmas that the later appendices use. Two conventions are in force throughout. A statement of the form "for every fixed $\zeta > 0$, whp [bound depending on ζ]" means that for each fixed ζ the bound holds with probability tending to one as $n \to \infty$. When we write $o(1)$ in place of ζ, we mean that ζ may be taken arbitrarily small at the cost of taking n larger. For $\theta > 0$ we write $\mathrm{clip}_\theta(x) := \max(-\theta, \min(\theta, x))$, and $Y_\theta := \mathrm{clip}_\theta(G_{11})$ is a clipped standard Gaussian. It is symmetric, satisfies $|Y_\theta| \le \theta$, and has variance σ^2_θ.

C.1 The optimum

The upper bound on the non-private optimum quoted in Section 2 is a union bound.

Lemma C.1 (Optimum, upper bound)**.** *For every fixed $\zeta > 0$, whp* $\max_{|I|=|J|=k} \mathrm{Ave}(G_{I,J}) \le (1+\zeta)\eta_{\mathrm{OPT}}$.

Proof. Fix a pair (I, J) with $|I| = |J| = k$. The sum $\mathrm{Sum}(G_{I,J})$ is a sum of k^2 independent standard Gaussians, so $\mathrm{Sum}(G_{I,J}) \sim N(0, k^2)$, and therefore

$$\Pr\big[\,\mathrm{Ave}(G_{I,J}) > (1+\zeta)\eta_{\mathrm{OPT}}\big] = \Pr\big[N(0,1) > 2(1+\zeta)\sqrt{k\log n}\big] \le \exp\big(-2(1+\zeta)^2 k \log n\big).$$

There are $\binom{n}{k}^2 \le (en/k)^{2k}$ pairs (I, J). Since $k = \exp(o(\log n))$ we have $\log(en/k) = (1 + o(1))\log n$, so a union bound gives a failure probability of at most

$$\exp\big(2k\log(en/k) - 2(1+\zeta)^2 k\log n\big) = \exp\big(-2k\log n\,((1+\zeta)^2 - 1 - o(1))\big),$$

which tends to zero. □

The matching lower bound is due to Bhamidi et al. (2017). It was stated there for fixed k, and Bhamidi et al. (2026) use it in the form below, for $k = \exp(o(\log n))$. We need it only for the benchmark of Appendix F.

Fact C.2 (Optimum, lower bound (Bhamidi et al., 2017; 2026))**.** For every fixed $\zeta > 0$, whp $\max_{|I|=|J|=k} \mathrm{Ave}(G_{I,J}) \ge (1-\zeta)\eta_{\mathrm{OPT}}$.

C.2 The clipping functions

The functions σ_θ and ψ were introduced in Section 2. Their elementary properties are the following.

Lemma C.3 (Properties of σ_θ and ψ)**.** *For every $\theta > 0$,*

$$\psi(\theta) = \phi(\theta) - \theta\bar{\Phi}(\theta) \le \frac{\phi(\theta)}{1+\theta^2} \le e^{-\theta^2/2}, \qquad 1 - \sigma_\theta^2 \le 2\phi(\theta)\Big(\theta + \frac{1}{\theta}\Big),$$

and $\sigma_\theta^2 \ge \sigma_1^2 > 1/2$ whenever $\theta \ge 1$.

Proof. Directly from the definition,

$$\psi(\theta) = \mathbb{E}\big[(-\theta - G_{11})_+\big] = \int_\theta^\infty (x-\theta)\phi(x)\,dx = \phi(\theta) - \theta\bar{\Phi}(\theta),$$

by the symmetry of the Gaussian and the identity $\int_\theta^\infty x\phi(x)dx = \phi(\theta)$. The Mills-ratio bound $\bar{\Phi}(\theta) \ge \frac{\theta}{1+\theta^2}\phi(\theta)$ then gives $\psi(\theta) \le \phi(\theta)\big(1 - \frac{\theta^2}{1+\theta^2}\big) = \frac{\phi(\theta)}{1+\theta^2}$, and $\phi(\theta) \le e^{-\theta^2/2}$ gives the last inequality. For the variance, symmetry gives

$$1 - \sigma_\theta^2 = \mathbb{E}\big[G_{11}^2 - \mathrm{clip}_\theta(G_{11})^2\big] = 2\,\mathbb{E}\big[(G_{11}^2 - \theta^2)\mathbf{1}_{G_{11}>\theta}\big] = 2\big(\theta\phi(\theta) + (1-\theta^2)\bar{\Phi}(\theta)\big),$$

using $\int_\theta^\infty x^2\phi(x)dx = \theta\phi(\theta) + \bar{\Phi}(\theta)$. For $\theta < 1$, applying $\bar{\Phi}(\theta) \le \phi(\theta)/\theta$ to the second term yields $1 - \sigma_\theta^2 \le 2\phi(\theta)(\theta + 1/\theta)$. For $\theta \ge 1$ the second term is nonpositive and the bound holds a fortiori. Finally, σ_θ^2 is nondecreasing in θ, and a direct computation gives $\sigma_1^2 = 2\Phi(1) - 1 - 2\phi(1) + 2\bar{\Phi}(1) \approx 0.516$. □

C.3 Privacy tools

The next four lemmas are the standard differential privacy facts our mechanisms rely on, and together they make up Lemma 2.3 of the main text. Throughout, $\Delta = 2\theta$.

Lemma C.4 (Sensitivity)**.** *If $G \sim G'$ are neighbors differing at the entry (i_0, j_0), then for all $I, J \subseteq [n]$,*

$$\big|\,\mathrm{Sum}(\widetilde{G}_{I,J}) - \mathrm{Sum}(\widetilde{G}'_{I,J})\big| \le \Delta \cdot \mathbf{1}[i_0 \in I,\ j_0 \in J] \le \Delta.$$

Proof. The clipped matrices $\widetilde{G}$ and $\widetilde{G}'$ differ only in the entry (i_0, j_0), where both take values in $[-\theta, \theta]$. So the two block sums agree unless $(i_0, j_0) \in I \times J$, in which case they differ by at most 2θ. □

Lemma C.5 (Exponential mechanism over subsets)**.** *Fix a row set I, a pool $P \subseteq [n]$, an integer $\kappa \le |P|$, and $\varepsilon' > 0$. The mechanism which outputs $S \subseteq P$ with $|S| = \kappa$ and $\Pr[S] \propto \exp\big(\frac{\varepsilon'}{2\Delta}\,\mathrm{Sum}(\widetilde{G}_{I,S})\big)$ is ε'-DP. This is $\mathcal{M}_{\mathrm{sel}}(I, P, \kappa, \varepsilon')$ of Definition 2.2. The same holds with the roles of rows and columns exchanged.*

Proof. Let $u(S) = \mathrm{Sum}(\widetilde{G}_{I,S})$ and $u'(S) = \mathrm{Sum}(\widetilde{G}'_{I,S})$ for neighboring $G \sim G'$, and let $Z = \sum_{|S|=\kappa} e^{\varepsilon' u(S)/2\Delta}$ and Z' be the corresponding normalizing constants. By Lemma C.4, $|u(S) - u'(S)| \le \Delta$ for every S, so each weight satisfies $e^{\varepsilon'(u(S)-u'(S))/2\Delta} \in [e^{-\varepsilon'/2}, e^{\varepsilon'/2}]$, and hence Z'/Z lies in the same interval. Therefore

$$\frac{\Pr[S]}{\Pr'[S]} = \frac{e^{\varepsilon' u(S)/2\Delta}}{e^{\varepsilon' u'(S)/2\Delta}} \cdot \frac{Z'}{Z} \in \big[e^{-\varepsilon'}, e^{\varepsilon'}\big],$$

which is the definition of ε'-DP. □

Lemma C.6 (Utility of the exponential mechanism (McSherry & Talwar, 2007))**.** *Let Y be sampled from $\Pr[Y = y] \propto e^{\beta u(y)}$ over N candidates, and let $u^* = \max_y u(y)$. Then for every $t > 0$,*

$$\Pr\left[u(Y) < u^* - \frac{\log N + t}{\beta}\right] \le e^{-t}.$$

This is the standard utility bound for the exponential mechanism (McSherry & Talwar, 2007), in the form of Dwork & Roth (2014, Theorem 3.11), and we omit the proof.

Lemma C.7 (Adaptive composition (Dwork et al., 2006; Dwork & Roth, 2014)). *Suppose mechanisms $\mathcal{M}_1, \dots, \mathcal{M}_T$ are run in sequence, where $\mathcal{M}_\tau$ may depend on the outputs $y_{<\tau}$ of the earlier mechanisms and is ε_τ-DP for every fixed history $y_{<\tau}$. Then the sequence $(\mathcal{M}_1, \dots, \mathcal{M}_T)$ is $(\sum_\tau \varepsilon_\tau)$-DP.*

This is basic composition (Dwork et al., 2006), in its adaptive form (Dwork & Roth, 2014, Section 3.5), and we omit the proof.

C.4 Moderate deviations for clipped sums

The extreme-value analysis of the main text needs a lower bound on the upper tail of a sum of independent clipped Gaussians, at the scale of $\sqrt{2\log n}$ standard deviations. The lemma below is a non-asymptotic form of Cramér's moderate-deviation theorem with explicit constants. We prove it directly so that its dependence on the clipping level θ is transparent, and Appendix J reuses it for sub-Gaussian entries. In the statement, $1/b$ is the number of standard deviations at which the threshold x lies, and a_0 is the tilt $x/(r\sigma^2)$ measured against the scale $1/\theta$ set by the bound.

Lemma C.8 (Moderate-deviation lower bound). *Let $Y_1, \dots, Y_r$ be i.i.d. with $\mathbb{E}Y = 0$, $|Y| \le \theta$ and $\operatorname{Var} Y = \sigma^2 > 0$, and let $S = \sum_{i \le r} Y_i$. Fix $x > 0$ and define*

$$b := \frac{\sigma\sqrt{r}}{x}, \qquad a_0 := \frac{\theta x}{r\sigma^2}.$$

If $b \le \frac{1}{10}$ and $a_0 \le \frac{1}{24}$, then

$$\Pr[S \ge x] \;\ge\; \frac{3}{4}\exp\Big(-\frac{x^2}{2r\sigma^2}\big(1 + 37b + 16a_0\big)\Big).$$

Proof. The plan is this. We define a reweighted law Q_λ under which $\{S \ge x\}$ is a typical event rather than a tail event, lower bound its probability there by Chebyshev's inequality, and convert back to the original law by an explicit change of measure. The price of the conversion is controlled by the cumulant generating function of Y, which we bound first.

Two elementary inequalities. For all real z,

$$1 + z \le e^z, \tag{1}$$

and, since $\exp$ is convex and so lies below its chord from $(0,1)$ to $(1,e)$ on $[0,1]$, with $e - 1 \le 2$,

$$e^u \le 1 + (e-1)u \le 1 + 2u \qquad (0 \le u \le 1). \tag{2}$$

Step 1: cumulant bounds. Fix $\lambda \ge 0$ with $\lambda\theta \le 1$ and let $M(\lambda) = \mathbb{E}e^{\lambda Y}$. We claim that

$$\frac{\lambda^2\sigma^2}{2}(1 - \lambda\theta) \;\le\; \log M(\lambda) \;\le\; \frac{\lambda^2\sigma^2}{2}(1 + 2\lambda\theta). \tag{3}$$

For the upper bound, apply $e^w \le 1 + w + \frac{w^2}{2}e^{|w|}$ with $w = \lambda Y$. Since $\mathbb{E}Y = 0$ the linear term vanishes, and since $|Y| \le \theta$ we may bound $e^{|w|} \le e^{\lambda\theta}$, so that

$$M(\lambda) \le 1 + \frac{\lambda^2}{2}\mathbb{E}\big[Y^2 e^{\lambda|Y|}\big] \le 1 + \frac{\lambda^2\sigma^2}{2}e^{\lambda\theta}.$$

Now apply (1) with $z = \frac{\lambda^2\sigma^2}{2}e^{\lambda\theta}$, and then (2) with $u = \lambda\theta \in [0,1]$ inside the exponent, which is legitimate since the prefactor $\frac{\lambda^2\sigma^2}{2}$ is nonnegative:

$$M(\lambda) \le \exp\Big(\frac{\lambda^2\sigma^2}{2}e^{\lambda\theta}\Big) \le \exp\Big(\frac{\lambda^2\sigma^2}{2}(1 + 2\lambda\theta)\Big).$$

For the lower bound, define

$$u_0 := \frac{\lambda^2\sigma^2}{2}\Big(1 - \frac{\lambda\theta}{3}\Big). \tag{4}$$

Since $\lambda\theta \le 1$ the bracket lies in $[\frac{2}{3}, 1]$, and since $\sigma \le \theta$ we have $\lambda^2\sigma^2 \le \lambda\theta$, so

$$0 \le u_0 \le \frac{\lambda^2\sigma^2}{2} \le \frac{\lambda\theta}{2} \le \frac{1}{2}. \tag{5}$$

Applying $e^w \geq 1 + w + \frac{w^2}{2} - \frac{|w|^3}{6}$ with $w = \lambda Y$, together with $\mathbb{E}Y = 0$ and $\mathbb{E}|Y|^3 \leq \theta\sigma^2$, gives $M(\lambda) \geq 1 + \frac{\lambda^2\sigma^2}{2} - \frac{\lambda^3\theta\sigma^2}{6} = 1 + u_0$. Both sides are at least 1, so applying the increasing function log gives $\log M(\lambda) \geq \log(1 + u_0)$. Combining this with $\log(1 + t) \geq t - t^2/2$ for $t \geq 0$ (the difference vanishes at $t = 0$ and has derivative $t^2/(1 + t) \geq 0$), and with $u_0^2/2 \leq \frac{\lambda^2\sigma^2}{2} \cdot \frac{\lambda\theta}{2}$ from (5), we get

$$\log M(\lambda) \geq u_0 - \frac{u_0^2}{2} \geq \frac{\lambda^2\sigma^2}{2}\Big(1 - \frac{\lambda\theta}{3} - \frac{\lambda\theta}{2}\Big) \geq \frac{\lambda^2\sigma^2}{2}(1 - \lambda\theta),$$

which is (3).

Step 2: the reweighted law and its moments. For $\lambda \geq 0$ define a probability law Q_λ on $\mathbb{R}$ by

$$Q_\lambda(A) := \frac{\mathbb{E}\big[e^{\lambda Y}\mathbf{1}_{Y\in A}\big]}{M(\lambda)} \qquad \text{for measurable } A \subseteq \mathbb{R}, \tag{6}$$

that is, the law of Y reweighted by $e^{\lambda Y}$ and renormalized. $M(\lambda)$ is finite since $|Y| \leq \theta$, and Q_0 is the law of Y. Write $m(\lambda) := \mathbb{E}_{Q_\lambda}[Y] = \mathbb{E}[Ye^{\lambda Y}]/M(\lambda)$ and $v(\lambda) := \mathbb{E}_{Q_\lambda}[Y^2] = \mathbb{E}[Y^2e^{\lambda Y}]/M(\lambda)$. We use two facts about Q_λ. First, if $Y_1, \ldots, Y_r$ are independent with law Q_λ then, directly from (6) applied to each coordinate, for every measurable $A \subseteq \mathbb{R}^r$,

$$\Pr\big[(Y_1, \ldots, Y_r) \in A\big] = M(\lambda)^r\, \mathbb{E}_{Q_\lambda^{\otimes r}}\big[e^{-\lambda S}\mathbf{1}_A\big]. \tag{7}$$

Second, m is continuous and nondecreasing in λ with $m(0) = 0$, so λ can be chosen to place the $Q_\lambda^{\otimes r}$-mean of S at any prescribed value in the range of rm. This is the purpose of Q_λ: under a suitable Q_{λ^*} the event $\{S \geq x\}$ is central rather than extreme, and a crude lower bound on its probability suffices.

Next we bound m and v for $\lambda\theta \leq 1$. From $|e^w - 1 - w| \leq 1.4w^2$ on $[-1, 1]$ (the maximum of $(e^w - 1 - w)/w^2$ on that interval is attained at $w = 1$ and equals $e - 2 < 1.4$), together with $\mathbb{E}|Y|^3 \leq \theta\sigma^2$,

$$\mathbb{E}[Ye^{\lambda Y}] = \lambda\sigma^2 + \mathbb{E}\big[Y(e^{\lambda Y} - 1 - \lambda Y)\big] \in \big[\lambda\sigma^2(1 - 1.4\lambda\theta),\ \lambda\sigma^2(1 + 1.4\lambda\theta)\big].$$

Dividing by $M(\lambda) \in [1, e^{3\lambda^2\sigma^2/2}]$ (the lower bound by Jensen's inequality, since $\mathbb{E}Y = 0$, and the upper bound from (3) and $\lambda\theta \leq 1$), and using $e^{-3\lambda^2\sigma^2/2} \geq 1 - 1.5\lambda^2\sigma^2 \geq 1 - 1.5\lambda\theta$, we get

$$\lambda\sigma^2(1 - 3\lambda\theta) \leq m(\lambda) \leq \lambda\sigma^2(1 + 1.4\lambda\theta), \qquad v(\lambda) \leq \sigma^2 e^{\lambda\theta} \leq 3\sigma^2. \tag{8}$$

Step 3: choice of λ. Set $x' := x + 2\sqrt{3r}\,\sigma = x(1 + 2\sqrt{3}b)$. Since $x' \leq 1.4x$ and $a_0 \leq \frac{1}{24}$, we have $\theta x'/(r\sigma^2) \leq \frac{1}{12}$, and then by (8),

$$m\Big(\frac{2x'}{r\sigma^2}\Big) \geq \frac{2x'}{r}\Big(1 - \frac{6\theta x'}{r\sigma^2}\Big) \geq \frac{x'}{r}.$$

By continuity and monotonicity of m there is therefore $\lambda^* \leq 2x'/(r\sigma^2)$ with $rm(\lambda^*) = x'$, and $\lambda^*\theta \leq 2\theta x'/(r\sigma^2) \leq \frac{1}{6} \leq 1$, so all bounds above apply at λ^*. Write $a := \theta x'/(r\sigma^2) \leq 2a_0$. Then (8) pins λ^* down: from $x'/r \leq \lambda^*\sigma^2(1 + 2.8a)$ we get $\lambda^* \geq \frac{x'}{r\sigma^2}(1 - 2.8a)$, and from $x'/r \geq \lambda^*\sigma^2(1 - 6a)$ together with $1/(1 - 6a) \leq 1 + 12a$ for $a \leq \frac{1}{12}$ we get $\lambda^* \leq \frac{x'}{r\sigma^2}(1 + 12a)$. Thus

$$\lambda^* = \frac{x'}{r\sigma^2}(1 + \eta) \qquad \text{for some } |\eta| \leq 12a.$$

Step 4: change of measure. Under $Q_{\lambda^*}^{\otimes r}$ the sum S has mean $rm(\lambda^*) = x'$ and variance at most $rv(\lambda^*) \leq 3r\sigma^2$, so by Chebyshev's inequality

$$Q_{\lambda^*}^{\otimes r}\big[\,|S - x'| \leq 2\sqrt{3r}\,\sigma\,\big] \geq \frac{3}{4}.$$

On that event, $S \geq x' - 2\sqrt{3r}\sigma = x$ and $e^{-\lambda^* S} \geq e^{-\lambda^*(x' + 2\sqrt{3r}\sigma)}$. Applying (7) with $A = \{S \geq x\}$ and restricting the expectation to that event, we get

$$\Pr[S \geq x] = M(\lambda^*)^r\, \mathbb{E}_{Q_{\lambda^*}^{\otimes r}}\big[e^{-\lambda^* S}\mathbf{1}_{S\geq x}\big] \geq \frac{3}{4}\exp\Big(r\log M(\lambda^*) - \lambda^* x' - 2\lambda^*\sqrt{3r}\,\sigma\Big).$$

Step 5: evaluation. By (3) and $\lambda^*\theta \le 2a$, and writing $\lambda^* = \frac{x'}{r\sigma^2}(1+\eta)$,

$$r\log M(\lambda^*) - \lambda^* x' \ge \frac{x'^2}{2r\sigma^2}\Big[(1+\eta)^2(1-2a) - 2(1+\eta)\Big] \ge -\frac{x'^2}{2r\sigma^2}(1+8a),$$

using $(1+\eta)^2 - 2(1+\eta) \ge -1$ and $(1+\eta)^2 \le 4$. Also $2\lambda^*\sqrt{3r}\sigma \le 2\sqrt{3}(1+12a)\frac{x'}{\sigma\sqrt{r}} \le \frac{8x'}{\sigma\sqrt{r}}$. Substituting $x' = x(1+2\sqrt{3}b)$, $\frac{x^2}{2r\sigma^2} = \frac{1}{2b^2}$ and $\frac{x'}{\sigma\sqrt{r}} = \frac{1+2\sqrt{3}b}{b}$, the exponent is at least

$$-\frac{1}{2b^2}\Big[(1+2\sqrt{3}b)^2(1+8a) + 16b(1+2\sqrt{3}b)\Big].$$

For $b \le \frac{1}{10}$ and $a \le \frac{1}{12}$ we have $(1+2\sqrt{3}b)^2 \le 1+9b$, $72ab \le 6b$ and $16b(1+2\sqrt{3}b) \le 22b$, so the bracket is at most $1+37b+8a \le 1+37b+16a_0$. This completes the proof. □

The extreme-value lemma of the main text, Lemma 3.5, follows. Its content is that, once the number of summands is large compared to $\theta^2\log(m/\kappa)$, the largest of m clipped sums behave as if the summands were Gaussian with the clipped variance.

Lemma C.9 (Top κ of m clipped sums)**.** *Let $S_1,\dots,S_m$ be i.i.d., each the sum of r i.i.d. copies of Y_θ, and let $S_{(1)} \ge \cdots \ge S_{(m)}$ denote their order statistics. Let $\zeta \in (0,\frac{1}{2}]$, let $1 \le \kappa = \exp(o(\log m))$, write $L := \log(m/\kappa)$, and suppose that $\theta\sqrt{L} \le \frac{\zeta}{60}\sigma_\theta\sqrt{r}$. Then, for all sufficiently large m,*

$$\Pr\Big[\sum_{l\le\kappa} S_{(l)} < \Big(1-\frac{\zeta}{2}\Big)\kappa\,\sigma_\theta\sqrt{r}\sqrt{2L}\Big] \le \exp\big(-(m/\kappa)^{\zeta/4}\big).$$

Proof. Let $x := (1-\frac{\zeta}{2})\sigma_\theta\sqrt{r}\sqrt{2L}$ and let K be the number of indices j with $S_j \ge x$. If $K \ge \kappa$ then the κ largest sums are each at least x and the claim holds, so it suffices to bound $\Pr[K<\kappa]$. Now $K \sim \mathrm{Bin}(m,p)$ with $p = \Pr[S_1 \ge x]$, and we lower bound p using Lemma C.8. In the notation of that lemma, $b = 1/((1-\frac{\zeta}{2})\sqrt{2L})$, which is at most $\frac{1}{10}$ for large m since $L\to\infty$ under $\kappa = \exp(o(\log m))$, and $a_0 = (1-\frac{\zeta}{2})\theta\sqrt{2L}/(\sigma_\theta\sqrt{r}) \le \sqrt{2}\zeta/60 \le \frac{\zeta}{40}$ by hypothesis, so that $16a_0 \le \frac{2\zeta}{5}$. Hence, once m is large enough that $37b \le \frac{\zeta}{20}$,

$$p \ge \frac{3}{4}\exp\Big(-\Big(1-\frac{\zeta}{2}\Big)^2 L\Big(1+\frac{\zeta}{20}+\frac{2\zeta}{5}\Big)\Big) = \frac{3}{4}\Big(\frac{\kappa}{m}\Big)^{(1-\zeta/2)^2(1+9\zeta/20)} \ge \frac{3}{4}\Big(\frac{\kappa}{m}\Big)^{1-\zeta/2},$$

where the last step uses $(1-\frac{\zeta}{2})^2(1+\frac{9\zeta}{20}) = 1-\frac{11\zeta}{20}-\frac{\zeta^2}{5}+\frac{9\zeta^3}{80} \le 1-\frac{\zeta}{2}$ for $\zeta \le \frac{1}{2}$. Thus $\mu := mp \ge \frac{3}{4}\kappa(m/\kappa)^{\zeta/2} \ge 2\kappa$ for large m, and $\{K<\kappa\} \subseteq \{K \le \mu/2\}$. The multiplicative Chernoff bound gives $\Pr[K \le \mu/2] \le e^{-\mu/8} \le \exp(-(m/\kappa)^{\zeta/4})$ for large m. □

The slack between $1-\frac{\zeta}{2}$ in the conclusion and $1-\zeta$ in the statements that use the lemma absorbs the factor $1-o(1)$ from $L = (1-o(1))\log n$ and the rounding of κ. The same statement holds for Gaussian summands, where no clipping level enters. This version handles the row sums in Appendix I.

Lemma C.10 (Top κ of m Gaussians)**.** *Let $Z_1,\dots,Z_m$ be i.i.d. $N(0,1)$ with order statistics $Z_{(1)} \ge \cdots \ge Z_{(m)}$, let $\zeta \in (0,\frac{1}{2}]$, $1 \le \kappa = \exp(o(\log m))$ and $L := \log(m/\kappa)$. Then for all sufficiently large m, $\Pr[\sum_{l\le\kappa} Z_{(l)} < (1-\zeta)\kappa\sqrt{2L}] \le \exp(-(m/\kappa)^{\zeta/2})$.*

Proof. Let $x := (1-\zeta)\sqrt{2L}$, which is at least 1 for large m, and $p := \bar\Phi(x)$. The Mills-ratio bound $\bar\Phi(x) \ge \frac{x}{1+x^2}\phi(x) \ge \frac{\phi(x)}{2x}$ for $x \ge 1$ gives

$$p \ge \frac{e^{-(1-\zeta)^2 L}}{2x\sqrt{2\pi}} = \Big(\frac{\kappa}{m}\Big)^{(1-\zeta)^2}\frac{1}{2\sqrt{2\pi}(1-\zeta)\sqrt{2L}} \ge \frac{3}{4}\Big(\frac{\kappa}{m}\Big)^{1-\zeta},$$

where the last step uses $(1-\zeta)^2 \le 1-\zeta-\frac{\zeta}{2}$ for $\zeta \le \frac{1}{2}$, so that $(\kappa/m)^{(1-\zeta)^2} \ge (\kappa/m)^{1-\zeta}(m/\kappa)^{\zeta/2}$, and $(m/\kappa)^{\zeta/2}$ exceeds $\frac{3}{2}\sqrt{2\pi}\sqrt{2L}$ for large m. The rest of the proof of Lemma C.9 applies with ζ in place of $\zeta/2$: $\mu = mp \ge \frac{3}{4}\kappa(m/\kappa)^{\zeta} \ge 2\kappa$ and $\Pr[K \le \mu/2] \le e^{-\mu/8} \le \exp(-(m/\kappa)^{\zeta/2})$ for large m. □

C.5 Clipping loss under monotone selection

Clipping an entry from below reduces the true value of any block containing it. The lemmas of this subsection show that a selection rule which prefers large clipped sums does not, in expectation, select blocks with more such entries than a random block does. This is Lemma 3.4 of the main text, and Appendices D and E use it to pass from clipped scores to true block sums. The first ingredient is the classical association inequality of Harris (1960).

Lemma C.11 (Harris's inequality)**.** *Let $X = (X_1, \dots, X_d)$ have independent real coordinates, and let $f, g : \mathbb{R}^d \to \mathbb{R}$ be measurable with f bounded and g integrable, f nondecreasing and g nonincreasing in each coordinate. Then $\mathbb{E}[f(X)g(X)] \le \mathbb{E}[f(X)]\,\mathbb{E}[g(X)]$.*

Proof. We induct on d. For $d = 1$, let X' be an independent copy of X. For every pair of values, $(f(X) - f(X'))(g(X) - g(X')) \le 0$, since f and g are ordered oppositely. Taking expectations and expanding,

$$0 \ge \mathbb{E}\big[(f(X) - f(X'))(g(X) - g(X'))\big] = 2\mathbb{E}[f(X)g(X)] - 2\mathbb{E}[f(X)]\mathbb{E}[g(X)],$$

which is the claim. For $d > 1$, condition on $X_2, \dots, X_d$. By the case $d = 1$ applied to the coordinate X_1,

$$\mathbb{E}[fg \mid X_{2:d}] \le \mathbb{E}[f \mid X_{2:d}]\,\mathbb{E}[g \mid X_{2:d}] =: F(X_{2:d})\,H(X_{2:d}).$$

Now F is nondecreasing and H nonincreasing in each of the remaining coordinates, so the inductive hypothesis gives $\mathbb{E}[FH] \le \mathbb{E}[F]\mathbb{E}[H] = \mathbb{E}[f]\mathbb{E}[g]$, and combining the two displays completes the induction. □

Lemma C.12 (Clipping loss)**.** *Let I be a fixed row set and P a pool of columns, with the entries $G_{I\times P}$ i.i.d. $N(0,1)$. Let a selection rule, possibly randomized with randomness independent of G, choose $S \subseteq P$, and suppose that for each $j \in P$ the inclusion probability $\pi_j(G) := \Pr[j \in S \mid G]$ is nondecreasing in each entry of the column $G_{I,j}$. Let $N_j := \sum_{i\in I}(-\theta - G_{ij})_+$ be the clipping loss of column j. Then*

$$\mathbb{E}\Big[\sum_{j\in S} N_j\Big] \le \psi(\theta)\,|I|\,\mathbb{E}|S|.$$

The exponential mechanism over κ-subsets of P, $\Pr[S] \propto \exp(\beta\,\mathrm{Sum}(\widetilde{G}_{I,S}))$ for any $\beta > 0$, and the non-private rule which selects the κ columns of largest $\mathrm{Sum}(\widetilde{G}_{I,\{j\}})$, both satisfy the hypothesis.

Proof. We have $\mathbb{E}[\sum_{j\in S} N_j] = \sum_{j\in P}\mathbb{E}[\pi_j(G)N_j]$. Fix j and condition on the columns other than j, which we denote G_{-j}. The loss N_j is a function of the column $G_{I,j}$ alone and is nonincreasing in each of its entries, while $\pi_j(\cdot, G_{-j})$ is nondecreasing in each of them by hypothesis. Since the entries of $G_{I,j}$ are independent, Lemma C.11 gives

$$\mathbb{E}\big[\pi_j(G)N_j \,\big|\, G_{-j}\big] \le \mathbb{E}\big[\pi_j(G) \,\big|\, G_{-j}\big]\,\mathbb{E}[N_j].$$

Taking expectations over G_{-j} and summing over j gives $\sum_j \mathbb{E}[\pi_j]\,\mathbb{E}[N_j] = \mathbb{E}|S| \cdot |I|\,\psi(\theta)$, since $\mathbb{E}[N_j] = |I|\psi(\theta)$ and $\sum_j \mathbb{E}[\pi_j] = \mathbb{E}|S|$.

It remains to verify the hypothesis for the two rules. For the exponential mechanism, Lemma C.14(a) below expresses the inclusion probability as $\pi_j = w_j c_j/(1 + w_j c_j)$, where $w_j = \exp(\beta\sum_{i\in I}\widetilde{G}_{ij})$ and c_j depends only on the other columns. This is increasing in w_j, and w_j is nondecreasing in each G_{ij} because clip_θ is nondecreasing. For the top-κ rule, raising an entry of column j can only increase its clipped sum, and hence can only move j into the selected set, whatever rule independent of G breaks ties. □

Lemma C.12 bounds an expectation. Its high-probability companion below is what the whp statements of Appendices D and E use. The same monotonicity argument, applied to the exponential of the loss, shows that the loss of the selected set is dominated in the moment-generating-function sense by the loss of a fixed set of the same size.

Lemma C.13 (Clipping loss, high probability)**.** *In the setting of Lemma C.12, suppose that the rule always selects κ columns and that for every κ-set S the probability $\pi_S(G) := \Pr[S \text{ is selected} \mid G]$ is nondecreasing in each entry of the columns in S (both rules of Lemma C.12 satisfy this). Then for every $0 \le \lambda \le \theta/2$,*

$$\mathbb{E}\Big[\exp\Big(\lambda \sum_{j\in S} N_j\Big)\Big] \;\le\; \exp\Big(|I|\kappa\Big(\lambda\psi(\theta) + \frac{2\phi(\theta)}{\theta^3}\lambda^2\Big)\Big).$$

Let $N = N^{(1)} + \cdots + N^{(T)}$, where conditionally on the history before step s the loss $N^{(s)}$ is of this form with a fixed set I_s, a selected size κ_s and i.i.d. $N(0,1)$ entries, and suppose $\sum_s |I_s|\kappa_s \le k^2$. Then for $\theta \ge 1$ and every $t > 0$,

$$\Pr\Big[N \ge \Big(\psi(\theta) + \frac{\phi(\theta)}{\theta^2}\Big)k^2 + \frac{2t}{\theta}\Big] \le e^{-t}, \qquad \textit{and} \qquad \psi(\theta) + \frac{\phi(\theta)}{\theta^2} \le e^{-\theta^2/2}.$$

Proof. Let $X := (-\theta - Z)_+$ with $Z \sim N(0,1)$. Then $\mathbb{E}X = \psi(\theta)$, and for $p \ge 2$,

$$\mathbb{E}X^p = \int_\theta^\infty (z-\theta)^p \phi(z)\, dz = \phi(\theta)\int_0^\infty x^p e^{-\theta x - x^2/2}\, dx \le \phi(\theta)\int_0^\infty x^p e^{-\theta x}\, dx = \frac{p!\,\phi(\theta)}{\theta^{p+1}},$$

so for $0 \le \lambda \le \theta/2$,

$$\mathbb{E}e^{\lambda X} \le 1 + \lambda\psi(\theta) + \frac{\phi(\theta)}{\theta}\sum_{p\ge 2}\Big(\frac{\lambda}{\theta}\Big)^p \le 1 + \lambda\psi(\theta) + \frac{2\phi(\theta)}{\theta^3}\lambda^2 \le \exp\Big(\lambda\psi(\theta) + \frac{2\phi(\theta)}{\theta^3}\lambda^2\Big).$$

For the first claim, write $N_S := \sum_{j\in S} N_j$, so that $\mathbb{E}[e^{\lambda N_S}] = \sum_{|S|=\kappa} \mathbb{E}[\pi_S(G) e^{\lambda N_S}]$. Fix S and condition on the columns outside S. As a function of the $|I|\kappa$ entries of the columns in S, which are independent, $e^{\lambda N_S}$ is nonincreasing in each entry, while π_S is nondecreasing in each. For the exponential mechanism this is because $\pi_S = w_S / \sum_{S'} w_{S'} = \big(1 + \sum_{S'\neq S} w_{S'}/w_S\big)^{-1}$ with $w_S = \prod_{j\in S} w_j$, and each ratio $w_{S'}/w_S = \prod_{j\in S'\setminus S} w_j \big/ \prod_{j\in S\setminus S'} w_j$ is nonincreasing in w_j for every $j \in S$, while w_j is nondecreasing in the entries of column j. For the top-κ rule, raising an entry of a selected column keeps it selected. Lemma C.11 gives $\mathbb{E}[\pi_S e^{\lambda N_S} \mid G_{-S}] \le \mathbb{E}[\pi_S \mid G_{-S}]\,(\mathbb{E}e^{\lambda X})^{|I|\kappa}$, and summing over S after taking expectations gives the claim, since $\sum_S \pi_S = 1$.

For the second claim, the first claim applied conditionally at each step and the tower property give $\mathbb{E}e^{\lambda N} \le \exp(k^2(\lambda\psi(\theta) + 2\phi(\theta)\lambda^2/\theta^3))$ for $\lambda \le \theta/2$. Take $\lambda = \theta/2$ in the Chernoff bound $\Pr[N \ge u] \le e^{-\lambda u}\mathbb{E}e^{\lambda N}$ with $u = (\psi(\theta) + \phi(\theta)/\theta^2)k^2 + 2t/\theta$. The exponent is $-\frac{\theta}{2}\psi(\theta)k^2 - \frac{\phi(\theta)}{2\theta}k^2 - t + \frac{\theta}{2}\psi(\theta)k^2 + \frac{\phi(\theta)}{2\theta}k^2 = -t$. Finally, by Lemma C.3, $\psi(\theta) \le \phi(\theta)/(1+\theta^2)$, so $\psi(\theta) + \phi(\theta)/\theta^2 \le \phi(\theta)(\frac{1}{1+\theta^2} + \frac{1}{\theta^2}) \le \frac{3}{2}\phi(\theta) \le e^{-\theta^2/2}$ for $\theta \ge 1$, as $\frac{3}{2}\phi(0) < 1$. □

C.6 Sampling

The exponential mechanism over κ-subsets with additive scores has product form, $\Pr[S] \propto \prod_{j\in S} w_j$ with $w_j = \exp(\beta\,\mathrm{Sum}(\widetilde{G}_{I,\{j\}}))$, and is known as conditional Poisson design (Hájek, 1964; Chen et al., 1994). This subsection proves Proposition C.15 of the main text, which gives two exact samplers, and first records a formula for the inclusion probabilities that Lemmas C.12 and C.13 and Appendix D use. For weights $w \in (0,\infty)^N$ we write $e_r(w) := \sum_{|T|=r}\prod_{j\in T} w_j$ for the r-th elementary symmetric polynomial, with $e_0 := 1$.

Lemma C.14 (Inclusion probabilities)**.** *Let R be sampled from $\Pr[R] \propto \prod_{j\in R} w_j$ over κ-subsets of $[N]$.*

(a) For every j and $\kappa < N$ (for $\kappa = N$ trivially $\Pr[j \in R] = 1$), with w_{-j} the weights with w_j removed and $c_j := e_{\kappa-1}(w_{-j})/e_\kappa(w_{-j})$,

$$\Pr[j \in R] = \frac{w_j c_j}{1 + w_j c_j}.$$

(b) If $w_l \ge 1$ for every l, then $\kappa\, e_\kappa(w) \ge (N - \kappa + 1)\, e_{\kappa-1}(w)$, i.e. $e_{\kappa-1}(w)/e_\kappa(w) \le \kappa/(N-\kappa+1)$.

Proof. (a) Splitting the sum $e_\kappa(w)$ according to whether $j \in T$ gives $e_\kappa(w) = e_\kappa(w_{-j}) + w_j e_{\kappa-1}(w_{-j})$, while the total weight of subsets containing j is $w_j e_{\kappa-1}(w_{-j})$. Hence $\Pr[j \in R] = w_j e_{\kappa-1}(w_{-j})/e_\kappa(w)$, and dividing numerator and denominator by $e_\kappa(w_{-j})$ gives the claim.

(b) Counting each κ-subset S once for each of its elements,

$$\kappa\, e_\kappa(w) = \sum_{|S|=\kappa} \sum_{j\in S} \prod_{l\in S} w_l = \sum_{|T|=\kappa-1} \prod_{l\in T} w_l \sum_{j\notin T} w_j \geq \sum_{|T|=\kappa-1} \prod_{l\in T} w_l \cdot (N-\kappa+1),$$

since each of the $N-\kappa+1$ terms w_j with $j \notin T$ is at least 1. □

Proposition C.15 (Exact sampling)**.** *Given $w_1, \dots, w_m > 0$ and $\kappa \leq m$, a set $T \subseteq [m]$ with $|T| = \kappa$ and $\Pr[T = T_0] \propto \prod_{j\in T_0} w_j$ can be sampled from this law without approximation*

- *(a) in $O(m\kappa)$ time and space, by a dynamic program, and*
- *(b) in $O(m)$ space, by rejective sampling, in expected time $O(m \cdot \mathbb{E}[\text{number of trials}])$, where the number of trials is geometric with success probability $\Pr[\mathrm{PB} = \kappa]$ for a Poisson-binomial variable PB with mean κ and variance at most κ, which is at least of order $\kappa^{-1/2}$ (Hájek, 1964).*

Proof. (a) For $0 \leq \ell \leq m$ and $0 \leq r \leq \kappa$ let $Z_\ell(r)$ denote the r-th elementary symmetric polynomial of the suffix $w_{\ell+1}, \dots, w_m$, with $Z_\ell(0) = 1$ and $Z_m(r) = 0$ for $r \geq 1$. Splitting on whether $\ell+1$ belongs to the subset gives the recursion $Z_\ell(r) = Z_{\ell+1}(r) + w_{\ell+1} Z_{\ell+1}(r-1)$, so the whole table can be filled in $O(m\kappa)$ time, in log-space to avoid overflow. We then scan $j = 1, \dots, m$, maintaining the number r of slots remaining (initially κ), and include j with probability $w_j Z_j(r-1)/Z_{j-1}(r)$, decrementing r if so. Correctness is an induction on j. Conditional on the decisions for $1, \dots, j-1$ having left r slots, the target law of $T \cap \{j, \dots, m\}$ is proportional to $\prod_{i\in T_0} w_i$ over r-subsets T_0 of $\{j, \dots, m\}$. Under this law $\Pr[j \in T] = w_j Z_j(r-1)/Z_{j-1}(r)$ by the same splitting as in Lemma C.14(a), and conditioning on the decision at j leaves the target law on $\{j+1, \dots, m\}$ with r or $r-1$ slots.

(b) For any $t > 0$, include each j independently with probability $p_j := w_j t/(1 + w_j t)$, and let S be the resulting set. Then

$$\Pr[S = T_0] = \prod_{j\in T_0} \frac{w_j t}{1 + w_j t} \prod_{j\notin T_0} \frac{1}{1 + w_j t} = \frac{t^{|T_0|} \prod_{j\in T_0} w_j}{\prod_{j\leq m}(1 + w_j t)},$$

so conditionally on $|S| = \kappa$ the law of S is proportional to $\prod_{j\in T_0} w_j$, for every value of t. Choosing t by bisection so that $\sum_j p_j = \kappa$, and drawing S repeatedly until $|S| = \kappa$, therefore produces an exact sample. The number of trials is geometric with the stated success probability, since $|S|$ is Poisson-binomial with mean κ and variance $\sum_j p_j(1-p_j) \leq \kappa$. □

All experiments in Section 6 use the rejective sampler.

D Proofs for the One-Shot Mechanisms

This appendix proves the two results of the main text that rest on a single call to $\mathcal{M}_{\mathrm{sel}}$: Proposition D.1, the finite-n rate behind Theorem 3.1 of Section 3.1, and Theorem 4.2 of Section 4.2. Throughout, $I_0 = [k]$ and $P = [n] \setminus I_0$.

D.1 The one-shot mechanism

Proposition D.1 (Rate for Theorem 3.1)**.** *Let $\zeta \in (0, \frac{1}{2}]$ and $\theta \geq 1$ with $\theta\sqrt{\log n} \leq \frac{\zeta}{60}\sigma_\theta\sqrt{k}$. Then for all large n, with probability at least $1 - 2/n$,*

$$\mathrm{Ave}(G_{I_0,J^*}) \geq (1-\zeta)\sigma_\theta \frac{\eta_{\mathrm{OPT}}}{\sqrt{2}} - \frac{4\theta(k\log(en/k) + \log n)}{\varepsilon k^2} - e^{-\theta^2/2} - \frac{4\log n}{k^2}.$$

Proof. Write $\tilde{s}_j = \sum_{i \in I_0} \widetilde{G}_{ij}$ for $j \in P$. The rows of I_0 are fixed in advance, so the scores $\tilde{s}_j$ are i.i.d. sums of k copies of Y_θ. We need three bounds, one for the non-private benchmark, one for the exponential mechanism, and one for what clipping costs. We combine them, pass to expectations, and then specialize to θ_n to obtain Theorem 3.1.

The benchmark. We apply Lemma C.9 with $(m, r, \kappa) = (n-k, k, k)$. Here $L = \log((n-k)/k) = (1 - o(1)) \log n$, so the hypothesis $\theta\sqrt{L} \le \frac{\zeta}{60}\sigma_\theta\sqrt{k}$ of the lemma follows from that of the proposition, and $k = \exp(o(\log m))$. So with probability at least $1 - \exp(-((n-k)/k)^{\zeta/4})$,

$$u^* := \max_{|S|=k} \sum_{j \in S} \tilde{s}_j \ge \Big(1 - \frac{\zeta}{2}\Big)(1 - o(1))\, \sigma_\theta k^{3/2}\sqrt{2\log n} \ge (1-\zeta)\sigma_\theta k^{3/2}\sqrt{2\log n}$$

for large n, where the factor $1 - o(1)$ is $\sqrt{L/\log n}$.

The mechanism. The one-shot mechanism is the exponential mechanism of Lemma C.6 with $\binom{n-k}{k} \le (en/k)^k$ candidates and $\beta = \varepsilon/(4\theta)$. With $t = \log n$ the lemma gives, with probability at least $1 - 1/n$,

$$\sum_{j \in J^*} \tilde{s}_j \ge u^* - \frac{4\theta}{\varepsilon}\big(k \log(en/k) + \log n\big).$$

The clipping loss. Every entry satisfies $G_{ij} \ge \widetilde{G}_{ij} - (-\theta - G_{ij})_+$, so the true block sum is at least the clipped one minus the loss from clipping below,

$$\mathrm{Sum}(G_{I_0,J^*}) \ge \sum_{j \in J^*} \tilde{s}_j - N, \qquad N := \sum_{j \in J^*} \sum_{i \in I_0} (-\theta - G_{ij})_+.$$

Lemma C.13 with $I = I_0$, $\kappa = k$ and $t = 2\log n$ gives $N \le e^{-\theta^2/2}k^2 + 4\log n/\theta \le e^{-\theta^2/2}k^2 + 4\log n$ with probability at least $1 - n^{-2}$.

Combining. Since $k = \exp(o(\log n))$, the benchmark fails with probability at most n^{-2} for large n, so the three bounds hold together with probability at least $1 - 1/n - 2n^{-2} \ge 1 - 2/n$. On that event we divide by k^2, and $k^{-1/2}\sqrt{2\log n} = \eta_{\mathrm{OPT}}/\sqrt{2}$ turns the benchmark into the first term of the proposition.

The expectation. Let F be the event that the benchmark bound or the mechanism bound fails, so that $\Pr[F] \le 1/n + n^{-2}$. Each clipped score is at least $-\theta k$, so $\mathrm{Sum}(G_{I_0,J^*}) \ge \sum_{j \in J^*} \tilde{s}_j - N \ge -\theta k^2 - N$ always, and on F^c the bound of the proposition holds with the clipping terms $e^{-\theta^2/2} + 4\log n/k^2$ replaced by N/k^2. Now take expectations. The term $-\theta$ contributes at most $\theta \Pr[F] = o(\eta_{\mathrm{OPT}})$, since $\theta \le \sqrt{k} = n^{o(1)}$ by hypothesis. The clipping loss contributes $\mathbb{E}N/k^2 \le \psi(\theta) \le e^{-\theta^2/2}$ by Lemma C.12. The remaining terms lose at most $\sigma_\theta \eta_{\mathrm{OPT}} \Pr[F] = o(\eta_{\mathrm{OPT}})$. So the bound of the proposition holds in expectation up to an additive $o(\eta_{\mathrm{OPT}})$.

The choice θ_n. With $\theta_n = \max\{1, \sqrt{2\log(k \log n)}\}$ we have $e^{-\theta_n^2/2} \le 1/(k \log n) = o(\sqrt{\log n/k})$ and $\log n/k^2 = o(\sqrt{\log n/k})$, and $\sigma_{\theta_n} \to 1$ since $\theta_n \to \infty$. The penalty term is $\Theta(\theta_n \log n/(\varepsilon k))$, and its ratio to η_{OPT} is $\Theta(\theta_n\sqrt{\log n}/(\varepsilon\sqrt{k}))$. This tends to zero if and only if $\varepsilon\sqrt{k}/(\theta_n\sqrt{\log n}) \to \infty$. The hypothesis $\theta_n\sqrt{\log n} \le \frac{\zeta}{60}\sigma_{\theta_n}\sqrt{k}$ follows from this condition when $\varepsilon = O(1)$, and otherwise it is the assumption $k \gg \theta_n^2 \log n$ of Theorem 3.1. For fixed ε the condition reads $k/(\log n \cdot \log(k \log n)) \to \infty$. Privacy is Lemma C.5, and the $O(nk)$ runtime is the cost of the $n - k$ scores plus the sampler of Proposition C.15. This proves Theorem 3.1. □

D.2 The sign mechanism

The sign mechanism depends on G only through the indicators $y_{ij} = \mathbf{1}[G_{ij} > 0]$, and its analysis uses the Gaussian law through two numbers only, the probability of a positive entry and the mean of a positive part. So we prove the finite-n bound for an arbitrary mean-zero law, in a form that Appendix J reuses, and specialize to the Gaussian afterwards. The argument computes in closed form what a single sign does to a column's inclusion probability. The one point that needs care is that this effect is a factor close to $e^{\varepsilon/2}$ only while the inclusion probability is small. Once the

mechanism has concentrated on a few columns, their inclusion probabilities are near one and a single sign no longer moves them. The hypothesis $\varepsilon^2 k \le (4-\delta)\log n$ is what keeps every inclusion probability small.

Lemma D.2 (Sign mechanism, general entries)**.** *Let the entries of G be i.i.d. copies of a random variable X with $\mathbb{E}X = 0$ and $\mathbb{E}|X| < \infty$, let $p := \Pr[X > 0]$ and $m^+ := \mathbb{E}X_+$, and let J^* be the output of the sign mechanism with fixed rows $I_0 = [k]$, that is, $\Pr[J^* = S] \propto \exp(\frac{\varepsilon}{2}\sum_{i\in I_0, j\in S} y_{ij})$ over k-subsets S of P. Suppose $n \ge 10k$, and define*

$$q := \frac{4k}{n}\exp\Big(\frac{\varepsilon}{2}\big(1+\sqrt{k}+\sqrt{k\log(n/k)}\big)\Big), \qquad r := \frac{k}{n} + \frac{n}{k}\,e^{-(n-k-1)/24}.$$

Then

$$\mathbb{E}\big[\,\mathrm{Ave}(G_{I_0,J^*})\big] \;\ge\; m^+\big(e^{\varepsilon/2}-1\big)(1-q)_+\Big(\frac{1}{1+p\,(e^{\varepsilon/2}-1)} - r\Big).$$

Proof. Write $P = [n]\setminus I_0$, $y = (y_{ij})_{i\in I_0, j\in P}$ for the sign indicators, which are i.i.d. Bernoulli(p), $u_j := \sum_{i\in I_0} y_{ij}$ for the number of positive entries in column j, and $w_j := e^{(\varepsilon/2)u_j}$. The mechanism depends on G only through y, so the inclusion probability $\pi_j(y) := \Pr[j \in J^* \mid G]$ is a function of y, and by Lemma C.14(a),

$$\pi_j(y) = \frac{w_j c_j}{1 + w_j c_j}, \qquad c_j = c_j(y_{-j}) := \frac{e_{k-1}(w_{-j})}{e_k(w_{-j})}, \tag{9}$$

where c_j depends only on the columns other than j.

Effect of one sign. Fix a pair (i,j) and write $\pi_j(1)$ and $\pi_j(0)$ for the values of π_j when y_{ij} is set to 1 or 0 with all other signs held fixed. Both are functions of $y_{-(ij)}$, the signs other than y_{ij}. Flipping y_{ij} from 0 to 1 multiplies w_j by $e^{\varepsilon/2}$ and leaves c_j alone, so with $z := w_j c_j$ evaluated at $y_{ij} = 0$,

$$\pi_j(0) = \frac{z}{1+z}, \qquad \pi_j(1) = \frac{e^{\varepsilon/2}z}{1+e^{\varepsilon/2}z},$$
$$\pi_j(1) - \pi_j(0) = \frac{(e^{\varepsilon/2}-1)\,z}{(1+z)(1+e^{\varepsilon/2}z)} = (e^{\varepsilon/2}-1)\,\pi_j(0)\big(1-\pi_j(1)\big).$$

We use this in two ways. First, $\pi_j(1) - \pi_j(0) \ge 0$, and $\pi_j(1) \le e^{\varepsilon/2}\pi_j(0)$ since $1 - \pi_j(1) \le 1$. Second, the difference is a factor $(e^{\varepsilon/2}-1)(1-\pi_j(1))$ times $\pi_j(0)$, which is close to $(e^{\varepsilon/2}-1)\pi_j(0)$ only when $\pi_j(1)$ is small. This is the point made before the lemma.

Expected value of one entry. The entries are independent, so G_{ij} is independent of $y_{-(ij)}$, and $\mathbb{E}[G_{ij}\mathbf{1}_{y_{ij}=1}] = \mathbb{E}[X\mathbf{1}_{X>0}] = m^+$ while $\mathbb{E}[G_{ij}\mathbf{1}_{y_{ij}=0}] = \mathbb{E}[X\mathbf{1}_{X\le 0}] = -m^+$, the second because $\mathbb{E}X = 0$. Conditioning on $y_{-(ij)}$ therefore gives

$$\mathbb{E}\big[G_{ij}\pi_j(y)\big] = m^+\,\mathbb{E}\big[\pi_j(1)\big] - m^+\,\mathbb{E}\big[\pi_j(0)\big] = m^+(e^{\varepsilon/2}-1)\,\mathbb{E}\big[\pi_j(0)\big(1-\pi_j(1)\big)\big], \tag{10}$$

where the expectations on the right are over $y_{-(ij)}$.

A lower bound on $\mathbb{E}[\pi_j(0)]$. Since $\Pr[y_{ij} = 1] = p$, we have $\Pr[j \in J^*] = p\,\mathbb{E}[\pi_j(1)] + (1-p)\,\mathbb{E}[\pi_j(0)] \le \big(1 + p(e^{\varepsilon/2}-1)\big)\mathbb{E}[\pi_j(0)]$ by the bound $\pi_j(1) \le e^{\varepsilon/2}\pi_j(0)$. On the other hand $\Pr[j \in J^*] = k/(n-k)$, because the columns are exchangeable and $\sum_{j\in P}\Pr[j \in J^*] = k$. Hence

$$\mathbb{E}[\pi_j(0)] \ge \frac{k}{(n-k)\big(1+p(e^{\varepsilon/2}-1)\big)}. \tag{11}$$

Keeping $\pi_j(1)$ small. It remains to bound $\pi_j(1)$ on an event of high probability. By the identity in the proof of Lemma C.14(b), $k\,e_k(w_{-j}) = \sum_{|T|=k-1}\prod_{l\in T} w_l \sum_{l\notin T} w_l \;\ge\; e_{k-1}(w_{-j})\cdot\Sigma$, where Σ is the sum of the $n-2k$ smallest of the $n-k-1$ weights w_{-j}. So $c_j \le k/\Sigma$. Let $t_0 := kp - \sqrt{k}$ and let M be the number of columns $l \in P\setminus\{j\}$ with $u_l \ge t_0$. Each such column has weight at least $e^{(\varepsilon/2)t_0}$, and at most $k-1$ of them are excluded from the $n-2k$ smallest weights, so $\Sigma \ge (M-k+1)e^{(\varepsilon/2)t_0}$. Now $u_l \sim \mathrm{Bin}(k,p)$, so Chebyshev's inequality gives $\Pr[u_l < t_0] \le kp(1-p)/k \le \frac{1}{4}$, and M is binomial with $N := n-k-1$ trials and success

probability at least $\frac{3}{4}$. The multiplicative Chernoff bound gives $\Pr[M \le N/2] \le e^{-N/24}$. On the event $\{M > N/2\}$ we have $M - k + 1 \ge (n-k-1)/2 - k + 1 \ge n/4$, using $n \ge 10k$, and hence

$$c_j \le \frac{4k}{n}\, e^{-(\varepsilon/2)t_0}.$$

Next, with $y_{ij} = 1$ the count u_j equals $1 + u'_j$ with $u'_j \sim \mathrm{Bin}(k-1, p)$, and Hoeffding's inequality gives $\Pr[u'_j \ge (k-1)p + t] \le e^{-2t^2/(k-1)} \le e^{-2t^2/k}$ for $k \ge 2$ and $t > 0$ (for $k = 1$ the event is empty). Take $t := \sqrt{k \log(n/k)}$, so that this probability is at most $(k/n)^2$, and let Ω be the intersection of $\{M > N/2\}$ and $\{u'_j < (k-1)p + t\}$. Then Ω is an event of $y_{-(ij)}$ with $\Pr[\Omega^c] \le e^{-(n-k-1)/24} + (k/n)^2$. On Ω, $u_j \le 1 + kp + t$, and therefore

$$w_j c_j\big|_{y_{ij}=1} \le e^{(\varepsilon/2)(1+kp+t)} \cdot \frac{4k}{n}\, e^{-(\varepsilon/2)(kp-\sqrt{k})} = \frac{4k}{n}\, e^{(\varepsilon/2)(1+\sqrt{k}+t)} = q,$$

so that $1 - \pi_j(1) = 1/(1 + w_j c_j) \ge 1/(1+q) \ge 1 - q$ on Ω.

Combining. The integrand in (10) is nonnegative, so restricting to Ω and using $\pi_j(0) \le 1$,

$$\mathbb{E}\big[\pi_j(0)\big(1 - \pi_j(1)\big)\big] \ge (1-q)_+\, \mathbb{E}\big[\pi_j(0)\mathbf{1}_\Omega\big] \ge (1-q)_+\big(\mathbb{E}[\pi_j(0)] - \Pr[\Omega^c]\big).$$

Summing (10) over the $k(n-k)$ pairs (i,j) and using (11),

$$\mathbb{E}\big[\,\mathrm{Sum}(G_{I_0,J^*})\big] \ge m^+(e^{\varepsilon/2}-1)(1-q)_+\Big(\frac{k^2}{1 + p(e^{\varepsilon/2}-1)} - k(n-k)\Pr[\Omega^c]\Big),$$

and $(n-k)\Pr[\Omega^c]/k \le (n/k)\big((k/n)^2 + e^{-(n-k-1)/24}\big) = r$. Dividing by k^2 gives the claim. □

Proof of Theorem 4.2. For the Gaussian, $p = \frac{1}{2}$ and $m^+ = 1/\sqrt{2\pi}$, so that $m^+(e^{\varepsilon/2}-1)/(1 + \frac{1}{2}(e^{\varepsilon/2}-1)) = \sqrt{2/\pi}\,(e^{\varepsilon/2}-1)/(e^{\varepsilon/2}+1) = \sqrt{2/\pi}\tanh(\varepsilon/4)$, and Lemma D.2 reads

$$\mathbb{E}\big[\,\mathrm{Ave}(G_{I_0,J^*})\big] \ge \sqrt{\frac{2}{\pi}}\,(1-q)_+\Big(\tanh\frac{\varepsilon}{4} - \frac{e^{\varepsilon/2}-1}{2}\,r\Big) = \sqrt{\frac{2}{\pi}}\tanh\frac{\varepsilon}{4}\,(1-q)_+\Big(1 - \frac{e^{\varepsilon/2}+1}{2}\,r\Big),$$

using $(e^{\varepsilon/2}-1)/\tanh(\varepsilon/4) = e^{\varepsilon/2}+1$. The condition $n \ge 10k$ holds for large n since $k = \exp(o(\log n))$. What is left is to show that $q \to 0$ and $e^{\varepsilon/2}r \to 0$ under $\varepsilon^2 k \le (4-\delta)\log n$, uniformly in ε. For q: since $1 \le \sqrt{k}$ and $\log(n/k) \le \log n$, the exponent satisfies

$$\begin{aligned}\frac{\varepsilon}{2}\big(1 + \sqrt{k} + \sqrt{k\log(n/k)}\big) &\le \varepsilon\sqrt{k} + \frac{\varepsilon\sqrt{k}}{2}\sqrt{\log n} \\ &\le \sqrt{(4-\delta)\log n} + \frac{\sqrt{4-\delta}}{2}\log n \le 2\sqrt{\log n} + \Big(1 - \frac{\delta}{8}\Big)\log n,\end{aligned}$$

where the last step uses $\sqrt{1-\delta/4} \le 1 - \delta/8$. Hence $q \le 4k\, n^{-\delta/8} e^{2\sqrt{\log n}} \to 0$. For r: the hypothesis gives $\varepsilon \le 2\sqrt{\log n/k} \le 2\sqrt{\log n}$, so $e^{\varepsilon/2} \le e^{\sqrt{\log n}} = n^{o(1)}$, while $r \le 2k/n$ for large n. So $e^{\varepsilon/2}r \to 0$. Both rates depend only on n, k and δ, which gives the bound at any budget $\varepsilon' \le \varepsilon$ with $\varepsilon'^2 k \le (4-\delta)\log n$, and running at $\varepsilon' \le \varepsilon$ is ε-DP by Lemma C.5. The budget $\varepsilon' = \min\{\varepsilon, \sqrt{\log n/k}\}$ satisfies $\varepsilon'^2 k \le \log n$, so the first bound of the theorem holds with $\delta = 3$. For the second, let $\varepsilon \le 1$, so that $\varepsilon' \le 1$ and $\tanh(\varepsilon'/4) \ge 0.97\,\varepsilon'/4$, whence $\mathbb{E}[\mathrm{Ave}] \ge (1-o(1))\,0.19\,\varepsilon'$. Dividing by $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$ gives $(1-o(1))\,0.095\,\varepsilon'\sqrt{k/\log n} = (1-o(1))\,0.095\min\{\varepsilon\sqrt{k/\log n}, 1\}$. □

The constant 4 in the hypothesis comes from a crude bound on c_j, which uses only the number of columns with an above-average count and not the full weight distribution, and we have not tried to optimize it. In the experiments of Section 6 the value falls below $\sqrt{2/\pi}\tanh(\varepsilon/4)$ from $\varepsilon^2 k \approx 5\log n$ on. Beyond the concentration point the bound of Lemma D.2 is vacuous, and rightly so. Once the selected columns have inclusion probabilities near one, (10) shows that a single sign contributes almost nothing, and the value is that of the best k columns by sign count, $\sqrt{2/\pi}$ times the non-private one-shot value, which decreases in k.

E PROOFS FOR THE BATCHED MECHANISM

This appendix gives the full analysis of $\mathcal{M}_{\mathrm{BIG}}$ from Section 3. We prove Proposition E.2, that the mechanism is online in the sense of Bhamidi et al. (2026), private, and runs in $O(nk)$ time, then Proposition 3.3, the finite-n rate, and from it Theorem 3.2. Throughout, B, κ, κ_b, p_b, Q_b and ε_0 are as in Section 3.2. A few arithmetic facts about the schedule come up repeatedly, so we record them here. Since $B \le \sqrt{k/3}$ we have $\kappa \ge B+1$ for $B \ge 2$, which gives $\kappa \le k - Bk/(B+1) \le \kappa_B = k - B\kappa < k/(B+1) + B < \kappa + B + 1 \le 2\kappa$ (for $B = 1$, $\kappa_1 = \lceil k/2 \rceil \le 2\lfloor k/2 \rfloor$ directly). Since $B\kappa + 1 \le k$ we have $p_B \le n$. Finally, under $n \ge 4k$ every group satisfies $|Q_b| \ge \kappa(n/k - 1) - 1 \ge \kappa n/(2k) \ge 2\kappa$ and $|Q_b| \le (\kappa+1)n/k \le 2\kappa n/k$. For the online model we also need to fix positions. Within a batch, the elements of R_b and of S_b are placed at the positions $b\kappa+1, \dots, b\kappa+\kappa_b$ in increasing order of index, though any rule that depends only on the set would do. The initial slot occupies positions $1, \dots, \kappa$ with $i_s = j_s = s$.

E.1 ONLINE MEMBERSHIP

Section 2 described the model of Bhamidi et al. (2026) in one sentence, so we restate it here for matrices. An algorithm in that model returns a $k \times k$ submatrix of G together with an ordering of its rows and columns: there are row indices $i_1, \dots, i_k$ and column indices $j_1, \dots, j_k$, and the output is the matrix $(G_{i_a, j_{a'}})_{a,a' \le k}$. We call a the *position* of row i_a. For such an output M we write $M_{\le s}$ for its principal $s \times s$ block, and we let $\mathcal{P}_s := \{G_{ij} : i, j \le \lfloor sn/k \rfloor\}$ denote the principal data prefix.

Definition E.1 (Online algorithms, (Bhamidi et al., 2026, Definition 2.2))**.** A deterministic algorithm $\mathcal{A}$ mapping G to a positioned $k \times k$ submatrix of G is *online* if (a) $\mathcal{A}(G)_{\le 1}$ is constructed based on $\mathcal{P}_1$, and (b) for each $2 \le s \le k$, given the blocks $\mathcal{A}(G)_{\le r} \setminus \mathcal{A}(G)_{\le r-1}$ for $r \le s-1$, the block $\mathcal{A}(G)_{\le s} \setminus \mathcal{A}(G)_{\le s-1}$ is determined by $\mathcal{P}_s$ and those earlier blocks. We denote the class of such algorithms by OLA. A randomized algorithm is online if $\mathcal{A}(\cdot; \omega) \in \mathrm{OLA}$ for every fixed value ω of its internal randomness, which is drawn independently of G.

Proposition E.2 (Online membership, privacy and runtime)**.** *$\mathcal{M}_{\mathrm{BIG}}$ is a randomized online algorithm in the sense of Definition E.1. It is ε-DP for every real input, and it runs in $O(nk)$ time, in $O(n)$ working space with rejective sampling.*

We prove the three claims in turn, starting with online membership.

Proof of online membership. Let $\mathcal{A}$ be the algorithm which runs $\mathcal{M}_{\mathrm{BIG}}$ on G and returns the positioned submatrix $(G_{i_a, j_{a'}})_{a,a' \le k}$, with the positions fixed at the start of this appendix. Clipping does not affect the argument, since $\widetilde{G}_{ij}$ is a function of G_{ij} alone. Fix the internal randomness ω. Two facts about the prefix points do the work. First, $p_b = \lfloor (b\kappa+1)n/k \rfloor \le \lfloor sn/k \rfloor$ whenever $s \ge b\kappa + 1$. Second, $\kappa \le p_1 \le p_2 \le \cdots \le p_B \le n$. We check clauses (a) and (b) of Definition E.1 position by position.

Position 1. The row and column at position 1 both have index 1, chosen without looking at G, so $\mathcal{A}(G)_{\le 1} = (G_{11})$, and $1 \le \lfloor n/k \rfloor$ since $n \ge k$. This is clause (a).

Positions $2 \le s \le \kappa$. These carry the constant indices $i_s = j_s = s$. The new block at position s consists of the entries $G_{s, j_{a'}}$ and $G_{i_a, s}$ with $a, a' \le s$, and all of their indices are at most $s \le \lfloor sn/k \rfloor$, so the block is a function of $\mathcal{P}_s$.

Positions in batch b. Let $b\kappa < s \le b\kappa + \kappa_b$. The indices at these positions are the elements of R_b and S_b. For fixed ω these are functions of the clipped scores $\tilde{s}_j = \sum_{i \in I_{b-1}} \widetilde{G}_{ij}$, $j \in Q_b$, and $\tilde{t}_i = \sum_{j \in J_b} \widetilde{G}_{ij}$, $i \in Q_b$, hence of G restricted to $I_{b-1} \times Q_b$ and $Q_b \times J_b$, together with the sets I_{b-1} and J_{b-1} determined by the earlier positions. All of I_{b-1}, J_b and Q_b lie in $[p_b]$, because $I_{b-1} \subseteq [\kappa] \cup Q_1 \cup \cdots \cup Q_{b-1}$, $J_b \subseteq [\kappa] \cup Q_1 \cup \cdots \cup Q_b$, and $Q_a \subseteq [p_a] \subseteq [p_b]$ for every $a \le b$. Since $s \ge b\kappa + 1$ gives $p_b \le \lfloor sn/k \rfloor$, the indices i_s, j_s are determined by $\mathcal{P}_s$ and the earlier positions. The entries of the new block are $G_{i_s, j_{a'}}$ and G_{i_a, j_s} with $a, a' \le s$. Each of these indices is either in $[\kappa]$ or in some $Q_{b'}$ with $b' \le b$, so all lie in $[p_b] \subseteq [\lfloor sn/k \rfloor]$, and the block lies in $\mathcal{P}_s$.

So for every s the block at position s is determined by $\mathcal{P}_s$ and the indices at earlier positions. Clause (b) is stated in terms of the earlier *blocks*, that is, the earlier output values, rather than the earlier

indices, but the argument gives more. By induction on s, every index at a position $\leq s$, and hence every entry of $\mathcal{A}(G)_{\leq s}$, is a function of $\mathcal{P}_s$ and ω alone, since the indices at positions $< s$ are functions of $\mathcal{P}_{s-1} \subseteq \mathcal{P}_s$. So the block at position s is determined by $\mathcal{P}_s$ already, which is stronger than clause (b), and $\mathcal{A}(\cdot;\omega) \in \mathrm{OLA}$. □

E.2 Privacy and runtime

Proof of the privacy and runtime claims of Proposition E.2. Conditional on the outputs of all earlier steps, the current index set (I_{b-1} at a column step, J_b at a row step) is fixed, and the step is the mechanism of Lemma C.5 itself, with budget ε_0. So each step is ε_0-DP for every fixed history, and Lemma C.7 over the $2B$ steps gives $2B\varepsilon_0 = \varepsilon$.

For the runtime, a step computes $|Q_b|$ scores, each a sum over at most k entries, at cost $O(|Q_b| \cdot k) = O(\kappa n)$. Over $2B$ steps this is $O(2B\kappa n) = O(nk)$. The sample at each step is the distribution of Proposition C.15 on a pool of size $O(\kappa n/k)$ with batch size at most 2κ, and costs $O(\kappa^2 n/k) = O(nk/B^2)$ time and space with sampler (a), or $O(\kappa n/k)$ space with sampler (b). Either way the score computation dominates. Since each entry index lies in at most one of the sets E_τ of Lemma E.3 below, the entries may be consumed as they are read, which gives $O(n)$ working space with sampler (b). □

E.3 Freshness and decomposition

The two lemmas of this subsection are the ones Section 3.5 appeals to when it says that the one-shot analysis applies step by step despite the mechanism's adaptivity. We number the steps $\tau = 1, \ldots, 2B$ in execution order, so that column step b is $\tau = 2b - 1$ and row step b is $\tau = 2b$, and we define the index sets

$$E_{2b-1} := \big([\kappa] \cup Q_1 \cup \cdots \cup Q_{b-1}\big) \times Q_b, \qquad E_{2b} := Q_b \times \big([\kappa] \cup Q_1 \cup \cdots \cup Q_b\big), \qquad E_0 := [\kappa] \times [\kappa].$$

We let ω denote the internal randomness and $\mathcal{F}_\tau := \sigma\big(\omega,\, G|_{E_1 \cup \cdots \cup E_\tau}\big)$.

Lemma E.3 (Freshness). *The sets $E_0, E_1, \ldots, E_{2B}$ are pairwise disjoint. Step τ reads its input only on E_τ. Its output, and hence the sets I_b, J_b produced by the steps completed by time τ, are $\mathcal{F}_\tau$-measurable, and $G|_{E_{\tau+1}}$ is independent of $\mathcal{F}_\tau$. Consequently, conditionally on $\mathcal{F}_{2b-2}$, the set I_{b-1} is fixed and the entries $G_{I_{b-1} \times Q_b}$ are i.i.d. $N(0,1)$, so the scores $\tilde{s}_j$, $j \in Q_b$, of column step b are i.i.d. sums of $|I_{b-1}|$ copies of Y_θ. Conditionally on $\mathcal{F}_{2b-1}$ the same holds for $G_{Q_b \times J_b}$ and the scores $\tilde{t}_i$ with $|J_b|$ copies, and $\mathrm{Sum}(G_{E_0}) \sim N(0, \kappa^2)$ independently of $\mathcal{F}_{2B}$.*

Proof. First, what each step reads. Column step b reads $I_{b-1} \times Q_b$, and row step b reads $Q_b \times J_b$. Since $I_{b-1}, J_b \subseteq [\kappa] \cup Q_1 \cup \cdots \cup Q_b$ and $I_{b-1} \cap Q_b = \emptyset$, these are contained in E_{2b-1} and E_{2b} respectively.

Next, disjointness, by cases. Two column steps $E_{2b-1}, E_{2b'-1}$ with $b \neq b'$ have the disjoint column sets $Q_b, Q_{b'}$. Two row steps $E_{2b}, E_{2b'}$ have the disjoint row sets $Q_b, Q_{b'}$. For a column step E_{2b-1} against a row step $E_{2b'}$: if $b' \geq b$ the row sets $[\kappa] \cup Q_{<b}$ and $Q_{b'}$ are disjoint, and if $b' < b$ the column sets Q_b and $[\kappa] \cup Q_{\leq b'}$ are disjoint. Finally $E_0 = [\kappa]^2$ has both its row set and its column set disjoint from every Q_b.

Measurability follows by induction on τ: the output of step τ is a function of ω, of $G|_{E_\tau}$, and of the outputs of earlier steps. Independence holds because the entries of G are independent, the index sets are disjoint, and ω is independent of G. For the conditional laws, given $\mathcal{F}_{2b-2}$ the set I_{b-1} is fixed, and $\tilde{s}_j = \sum_{i \in I_{b-1}} \widetilde{G}_{ij}$ for $j \in Q_b$ are sums over disjoint sets of entries of $G|_{E_{2b-1}}$, which are i.i.d. $N(0,1)$ and independent of $\mathcal{F}_{2b-2}$. The row case and the corner block are identical. □

Lemma E.4 (Decomposition). *For any matrix $H \in \mathbb{R}^{n \times n}$,*

$$\mathrm{Sum}(H_{I_B, J_B}) = \mathrm{Sum}(H_{E_0}) + \sum_{b=1}^{B} \sum_{j \in S_b} \sum_{i \in I_{b-1}} H_{ij} + \sum_{b=1}^{B} \sum_{i \in R_b} \sum_{j \in J_b} H_{ij}.$$

In particular, taking $H = \widetilde{G}$, the clipped block sum is the corner sum plus the sum of the selected scores, each score taken at the step in which its index was selected.

Proof. Let the entry (i, j) have its row i in stage a and its column j in stage a', where stage 0 is $[\kappa]$ and stage $b \geq 1$ is R_b for rows and S_b for columns. The entry is counted in the column-step term for $b = a'$ if and only if $i \in I_{a'-1}$, that is, if and only if $a \leq a' - 1$. It is counted in the row-step term for $b = a$ if and only if $j \in J_a$, that is, if and only if $a' \leq a$, and it is counted in the corner if and only if $a = a' = 0$. For $(a, a') \neq (0, 0)$ one and only one of $a < a'$ and $a' \leq a$ holds, and for $(a, a') = (0, 0)$ the condition $a' \leq a$ holds but there is no row step $b = 0$, so the entry is counted only in the corner. So every entry of $I_B \times J_B$ is counted once and only once. The identity is combinatorial and holds for any matrix. □

E.4 Proof of Proposition 3.3

Proof. The plan is the one of Section 3.4: bound the benchmark, the mechanism's loss and the clipping loss at each step, add up over steps through Lemma E.4, and then simplify. Fix a step, and let r denote the number of entries summed in each of its scores ($r = |I_{b-1}| \geq b\kappa$ at column step b, $r = |J_b| \geq (b+1)\kappa$ at row step b, the latter because J_b already contains S_b), let $\kappa' \in [\kappa, 2\kappa]$ be its batch size, let $m \in [2\kappa, 2\kappa n/k]$ be its pool size, and let u^* be the maximum of $\sum_{j \in S} \tilde{s}_j$ over κ'-subsets of the pool.

The benchmark. By Lemma E.3, conditionally on the history the scores of the step are i.i.d. sums of r copies of Y_θ, so Lemma C.9 applies with (m, r, κ'). Its hypotheses hold: $\kappa' \leq 2\kappa \leq 2k = \exp(o(\log m))$ since $m \geq \kappa n/(2k) \geq n^{1-o(1)}$, and $L = \log(m/\kappa')$ satisfies $\log(n/(4k)) \leq L \leq \log n$, so $L = (1 - o(1)) \log n$, while $\theta\sqrt{L} \leq \frac{\zeta}{60}\sigma_\theta\sqrt{r}$ follows from the assumption $\theta\sqrt{\log n} \leq \frac{\zeta}{60}\sigma_\theta\sqrt{\kappa}$ together with $r \geq \kappa$. Hence, with probability at least $1 - \exp(-(n/4k)^{\zeta/4})$,

$$u^* \geq \Big(1 - \frac{\zeta}{2}\Big)(1 - o(1))\,\sigma_\theta\sqrt{r}\,\kappa'\sqrt{2\log n},$$

the factor $1 - o(1)$ being $\sqrt{L/\log n}$. We carry this factor along and dispose of it at the end.

The mechanism. The step is the exponential mechanism of Lemma C.6 with $\binom{m}{\kappa'} \leq (em/\kappa')^{\kappa'} \leq (2en/k)^{2\kappa}$ candidates, since $m/\kappa' \leq 2n/k$ and $\kappa' \leq 2\kappa$, and with $\beta = \varepsilon_0/(4\theta) = \varepsilon/(8B\theta)$. With $t = 2\log n$, the selected set has score at least

$$u^* - \frac{8\theta B}{\varepsilon}\big(2\kappa\log(2en/k) + 2\log n\big)$$

with probability at least $1 - n^{-2}$.

The clipping loss. Let N denote the total clipping loss $\sum(-\theta - G_{ij})_+$ over the entries of $I_B \times J_B$ outside E_0. Lemma E.4 applied to the nonnegative matrix $H_{ij} = (-\theta - G_{ij})_+$ splits N into per-step contributions, and at each step Lemma C.12 applies conditionally on the history, with the fixed index set I_{b-1} or J_b and the fresh entries of Lemma E.3. Summing,

$$\mathbb{E}[N] \leq \psi(\theta)\sum_{b=1}^{B}\big(|I_{b-1}| + |J_b|\big)\kappa_b = \psi(\theta)\big(k^2 - \kappa^2\big) \leq \psi(\theta)k^2,$$

which is the expectation bound. For the high-probability bound, the per-step losses are in the setting of Lemma C.13, conditionally on the history, with $\sum_b(|I_{b-1}| + |J_b|)\kappa_b \leq k^2$, so with $t = 2\log n$ that lemma gives $N \leq e^{-\theta^2/2}k^2 + 4\log n/\theta \leq e^{-\theta^2/2}k^2 + 4\log n$ with probability at least $1 - n^{-2}$.

The corner. Since $\mathrm{Sum}(G_{E_0}) \sim N(0, \kappa^2)$, we have $|\,\mathrm{Sum}(G_{E_0})| \leq 2\kappa\sqrt{\log n}$ with probability at least $1 - 2n^{-2}$.

Combining. Since $k = \exp(o(\log n))$, each benchmark event fails with probability at most n^{-2} for large n, so a union bound over the $4B + 2$ events (two per step, the clipping loss and the corner) shows that all of the above hold with probability at least $1 - (4B+3)n^{-2} \geq 1 - 1/n$ for large n (the corner contributing $2n^{-2}$), as $B \leq \sqrt{k/3} = \exp(o(\log n))$. By Lemma E.4 applied to G, together

with $G_{ij} \geq \widetilde{G}_{ij} - (-\theta - G_{ij})_+$, we then have

$$\mathrm{Sum}(G_{I_B,J_B}) \geq \Big(1 - \frac{\zeta}{2}\Big)(1 - o(1))\,\sigma_\theta \kappa^{3/2}\sqrt{2\log n}\Big(\sum_{b=1}^{B}\sqrt{b} + \sum_{b=2}^{B+1}\sqrt{b}\Big) - \frac{32\theta B^2}{\varepsilon}\big(\kappa\log(2en/k) + \log n\big) - e^{-\theta^2/2}k^2 - 4\log n - 2\kappa\sqrt{\log n},$$

where in the first term we have used $\sqrt{r} \geq \sqrt{b\kappa}$ at column step b, $\sqrt{r} \geq \sqrt{(b+1)\kappa}$ at row step b, and $\kappa' \geq \kappa$. The first term is a Riemann sum for the IGP's integral, and we simplify it next. Since $\sqrt{b} \geq \int_{b-1}^{b}\sqrt{x}\,dx$,

$$\sum_{b=1}^{B}\sqrt{b} \geq \frac{2}{3}B^{3/2} \qquad \text{and} \qquad \sum_{b=2}^{B+1}\sqrt{b} \geq \frac{2}{3}\big((B+1)^{3/2} - 1\big) \geq \frac{2}{3}\big(B^{3/2} - 1\big),$$

so the bracket is at least $\frac{4}{3}B^{3/2}(1 - \frac{1}{2B^{3/2}})$. Also $\kappa = \lfloor k/(B+1)\rfloor \geq \frac{k}{B+1}(1 - \frac{B+1}{k})$ with $(B+1)/k \leq (\sqrt{k/3}+1)/k \to 0$, so that

$$\kappa^{3/2} \geq (1 - o(1))\Big(\frac{k}{B+1}\Big)^{3/2} = (1 - o(1))\,\frac{k^{3/2}}{B^{3/2}}\Big(\frac{B}{B+1}\Big)^{3/2} \geq (1 - o(1))\,\frac{k^{3/2}}{B^{3/2}}\Big(1 - \frac{3}{2B}\Big),$$

and $(1-\frac{3}{2B})(1-\frac{1}{2B^{3/2}}) \geq 1 - \frac{2}{B}$ for every $B \geq 1$. So the first term is at least $(1-\frac{\zeta}{2})(1-o(1))\sigma_\theta(1-\frac{2}{B})\frac{4}{3}k^{3/2}\sqrt{2\log n}$, which is at least $(1-\zeta)\sigma_\theta(1-\frac{2}{B})\frac{4}{3}k^{3/2}\sqrt{2\log n}$ for large n when $B \geq 2$ (for $B = 1$ the first term is nonnegative while $1 - \frac{2}{B} < 0$, so the bound of the proposition holds trivially). For the penalty, $B\kappa \leq k$ gives $\frac{32\theta B^2}{\varepsilon}(\kappa\log(2en/k) + \log n) \leq \frac{32\theta B}{\varepsilon}(k\log(2en/k) + B\log n)$, and for the last two terms $4\log n + 2\kappa\sqrt{\log n} \leq 4k\log n/B + 2k\log n/B = 6k\log n/B$, using $B \leq k$, $\kappa \leq k/B$ and $\sqrt{\log n} \leq \log n$. Dividing by k^2 and using $\frac{4}{3}k^{-1/2}\sqrt{2\log n} = \frac{2\sqrt{2}}{3}\eta_{\mathrm{OPT}}$, we obtain the bound of the proposition.

For the expectation, we argue as in the proof of Proposition D.1, with one addition for the corner. Let F be the event that the benchmark, mechanism or corner bound fails, so that $\Pr[F] \leq n^{-1}$. By Lemma E.4, $\mathrm{Sum}(G_{I_B,J_B}) \geq \mathrm{Sum}(G_{E_0}) - \theta(k^2 - \kappa^2) - N$ always, since each selected clipped score is at least $-\theta$ times its number of summands. The clipping term contributes $\mathbb{E}N/k^2 \leq \psi(\theta)$ in all cases. The term $-\theta$ contributes at most $\theta\Pr[F] \leq \theta/n = o(\eta_{\mathrm{OPT}})$ on F, as $\theta \leq \sqrt{\kappa}$ by hypothesis. The corner contributes $\mathbb{E}[|\mathrm{Sum}(G_{E_0})|\mathbf{1}_F]/k^2 \leq \kappa\sqrt{\Pr[F]}/k^2 = o(\eta_{\mathrm{OPT}})$ on F by the Cauchy–Schwarz inequality, since $\mathrm{Sum}(G_{E_0}) \sim N(0,\kappa^2)$. On F^c the bound of the proposition holds with the clipping term $e^{-\theta^2/2}$ replaced by N/k^2, and $\mathbb{E}N/k^2 \leq \psi(\theta) \leq e^{-\theta^2/2}$. So the expectation bound holds as stated, up to $o(\eta_{\mathrm{OPT}})$. □

E.5 Proof of Theorem 3.2

Proof of Theorem 3.2. Privacy is Proposition E.2, which also gives online membership and the $O(nk)$ runtime. For the utility, we take $\theta = \theta_n = \max\{1, \sqrt{2\log(k\log n)}\}$ in Proposition 3.3 and check its terms one at a time. We have $e^{-\theta_n^2/2} \leq 1/(k\log n) = o(\eta_{\mathrm{OPT}})$, and $\sigma_{\theta_n} \to 1$ by Lemma C.3, so the first term is $(1 - \zeta - o(1))(1 - \frac{2}{B})\frac{2\sqrt{2}}{3}\eta_{\mathrm{OPT}}$. The last term is $O(\log n/k) = o(\eta_{\mathrm{OPT}})$. The penalty term is $\Theta(\theta_n B\log n/(\varepsilon k))$, whose ratio to $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$ is $\Theta(\theta_n B\sqrt{\log n}/(\varepsilon\sqrt{k}))$, and this tends to zero under the stated choice $B \ll \varepsilon\sqrt{k}/(\theta_n\sqrt{\log n})$. Since also $B \to \infty$, the factor $1 - \frac{2}{B}$ tends to one. What is left is the hypothesis $\theta_n\sqrt{\log n} \leq \frac{\zeta}{60}\sigma_{\theta_n}\sqrt{\kappa}$ of the proposition. It holds once $\kappa \gg \theta_n^2\log n$, which the theorem assumes. When $\varepsilon = O(1)$ this assumption is automatic, because the choice of B gives $\kappa \approx k/B \gg \theta_n\sqrt{k\log n}/\varepsilon$, and the right side is at least $\theta_n^2\log n$ as soon as $\varepsilon \leq \sqrt{k}/(\theta_n\sqrt{\log n})$, which holds for bounded ε under the condition $\varepsilon\sqrt{k}/(\theta_n\sqrt{\log n}) \to \infty$. Since ζ is arbitrary, the whp bound follows, and the expectation bound follows in the same way from the expectation form of the proposition.

For fixed ε, the condition $\varepsilon\sqrt{k}/(\theta_n\sqrt{\log n}) \to \infty$ reads $k/(\log n \cdot \log(k\log n)) \to \infty$. If $k \geq (\log n)^{1+\delta}$ then $\log n \leq k^{1/(1+\delta)}$ and $\log(k\log n) \leq 2\log k$, so $k/(\log n \cdot \log(k\log n)) \geq k^{\delta/(1+\delta)}/(2\log k) \to \infty$. □

F THE FULL-SPACE EXPONENTIAL MECHANISM

This appendix proves Theorem F.1, the statistical benchmark of Table 1 and of the discussion after Corollary 4.3. The exponential mechanism over all candidate blocks attains $(1-o(1))\eta_{\mathrm{OPT}}$, but we do not know how to sample it in polynomial time. Let $\mathcal{M}_{\mathrm{EM}}$ output (I^*, J^*) with

$$\Pr\big[(I^*, J^*) = (I, J)\big] \propto \exp\Big(\frac{\varepsilon}{4\theta}\,\mathrm{Sum}(\widetilde{G}_{I,J})\Big)$$

over all pairs with $|I| = |J| = k$.

Theorem F.1 (Private recovery of the optimum)**.** *$\mathcal{M}_{\mathrm{EM}}$ is ε-DP for every real input. Take $\theta = c\sqrt{\log n}$ with a constant $c > 2$. If $\varepsilon = \omega(\log n/\sqrt{k})$, which for fixed ε reads $k \gg (\log n)^2/\varepsilon^2$, then for every fixed $\zeta > 0$,*

$$\Pr\big[\,\mathrm{Ave}(G_{I^*,J^*}) \geq (1-\zeta)\eta_{\mathrm{OPT}}\big] \longrightarrow 1.$$

Proof. Privacy is Lemma C.5 with the pairs (I, J) as the candidates. Each entry appears in at most one term of $\mathrm{Sum}(\widetilde{G}_{I,J})$, so the score has sensitivity 2θ, and the exponent has the form $\frac{\varepsilon}{2\Delta}$ with $\Delta = 2\theta$.

For utility, a union bound over the n^2 entries gives $\Pr[\widetilde{G} \neq G] \leq 2n^{2-c^2/2} \to 0$. So whp the mechanism runs on G itself, and the non-private optimum of Fact C.2 is available as the benchmark. We apply Lemma C.6 with $N = \binom{n}{k}^2 \leq (en/k)^{2k}$ candidates, $\beta = \varepsilon/(4\theta)$ and $t = \log n$, which gives, whp,

$$\mathrm{Ave}(G_{I^*,J^*}) \geq (1-\zeta/2)\eta_{\mathrm{OPT}} - \frac{4\theta\big(2k\log(en/k) + \log n\big)}{\varepsilon k^2}.$$

The penalty is $O((\log n)^{3/2}/(\varepsilon k))$, which is $o(\eta_{\mathrm{OPT}})$ precisely when $\varepsilon = \omega(\log n/\sqrt{k})$. Under that condition it is at most $\frac{\zeta}{2}\eta_{\mathrm{OPT}}$ for large n, and the claim follows. □

A constant clipping level would cut the penalty by the factor $c\sqrt{\log n}/\theta$ and bring the threshold down to that of Theorem 3.2. The clipping loss is not the obstacle. The argument of Lemma C.12, applied entry by entry, controls it for this mechanism too, since the probability that an entry (i, j) is selected is nondecreasing in G_{ij}. What we lack is a lower bound on the *clipped* optimum, the analogue of Fact C.2 for bounded entries, and we do not pursue it.

What separates $\mathcal{M}_{\mathrm{EM}}$ from $\mathcal{M}_{\mathrm{BIG}}$ is computation, not privacy. Sampling $\mathcal{M}_{\mathrm{EM}}$ is a $\binom{n}{k}^2$-way selection, and we know of no polynomial-time procedure for it.

G THE EXPECTATION BOUND: ONE-SIDED VERSION AND SHARPENING

This appendix supports Section 4. We prove the expectation bound of Theorem 4.1 in a one-sided form, with the constant $C_\varepsilon := e^\varepsilon \phi(\Phi^{-1}(1-e^{-\varepsilon}))$, which is simpler to state and is enough for every qualitative conclusion of the paper. We then explain how the two-sided constant C^*_ε of the main text sharpens it, and we prove Corollary 4.3. Lemma I.5 compares the two constants.

The one-sided bound rests on the following lemma, which is used in Step 2 of the proof of Theorem G.2 below.

Lemma G.1 (Bathtub)**.** *Let $X \sim N(0,1)$ and let $q : \mathbb{R} \to [0, U]$ be measurable with $\mathbb{E}[q(X)] = \bar{q} \in (0, U]$. Then*

$$\mathbb{E}[Xq(X)] \leq U\,\phi\big(\Phi^{-1}(1-\bar{q}/U)\big).$$

Proof. If $\bar{q} = U$ then $q = U$ almost surely and $\mathbb{E}[Xq(X)] = 0 = U\phi(-\infty)$, so assume $\bar{q} < U$. Let $\tau := \Phi^{-1}(1-\bar{q}/U)$, so that $U\Pr[X > \tau] = \bar{q}$, and let $q^* := U\mathbf{1}_{X>\tau}$. Then $\mathbb{E}[q^*(X)] = \bar{q} = \mathbb{E}[q(X)]$. For every x we have $(x-\tau)(q^*(x) - q(x)) \geq 0$, since $x > \tau$ gives $q^*(x) = U \geq q(x)$ and $x \leq \tau$ gives $q^*(x) = 0 \leq q(x)$. Taking expectations, $\mathbb{E}[X(q^* - q)] \geq \tau\mathbb{E}[q^* - q] = 0$, so

$$\mathbb{E}[Xq(X)] \leq \mathbb{E}[Xq^*(X)] = U\int_\tau^\infty x\phi(x)\,dx = U\phi(\tau),$$

by the identity $\int_\tau^\infty x\phi(x)dx = \phi(\tau)$. □

Theorem G.2 (One-sided expectation bound)**.** *For every ε-DP mechanism $\mathcal{A}$ outputting (I, J) with $|I| = |J| = k$, $\mathbb{E}[\mathrm{Ave}(G_{\mathcal{A}(G)})] \le C_\varepsilon$.*

Proof. Fix an entry (i, j), write $X = G_{ij}$, and define

$$p(x) := \Pr\big[(i,j) \in I \times J \,\big|\, X = x\big], \qquad \bar{p} := \mathbb{E}[p(X)] = \Pr\big[(i,j) \in I \times J\big],$$

where the conditional probability averages over the other entries of G and the internal randomness of $\mathcal{A}$.

Step 1: privacy as a constraint on one entry. Take any $x, x' \in \mathbb{R}$ and any fixed value g_- of the other entries. The matrices (x, g_-) and (x', g_-) are neighbors, so ε-DP gives $\Pr[(i,j) \text{ selected} \mid X = x, g_-] \le e^\varepsilon \Pr[(i,j) \text{ selected} \mid X = x', g_-]$. The law of g_- does not depend on x or x', because the entries are independent, so averaging over g_- gives $p(x) \le e^\varepsilon p(x')$ for all x, x', and in particular $p(x) \le e^\varepsilon \bar{p}$ for all x.

Step 2: the correlation. By the tower property, $\mathbb{E}[G_{ij}\mathbf{1}[(i,j) \in I \times J]] = \mathbb{E}[Xp(X)]$. If $\bar{p} = 0$ this is 0. Otherwise Lemma G.1 with $U = e^\varepsilon \bar{p}$ gives

$$\mathbb{E}[Xp(X)] \le e^\varepsilon \bar{p}\, \phi\big(\Phi^{-1}(1 - e^{-\varepsilon})\big) = C_\varepsilon\, \bar{p}.$$

Step 3: summation. We sum over the n^2 entries and use $\sum_{i,j} \bar{p}_{ij} = \mathbb{E}|I \times J| = k^2$, which gives

$$\mathbb{E}\big[\,\mathrm{Sum}(G_{\mathcal{A}(G)})\big] = \sum_{i,j} \mathbb{E}\big[G_{ij}\mathbf{1}[(i,j) \in I \times J]\big] \le C_\varepsilon \sum_{i,j} \bar{p}_{ij} = C_\varepsilon k^2.$$

Dividing by k^2 gives the claim. □

Step 1 in fact yields the two-sided constraint $\sup_x p \le e^\varepsilon \inf_x p$, and the proof above uses only its upper half. Replacing Lemma G.1 by Lemma I.4, which uses both halves, gives $\mathbb{E}[\mathrm{Ave}(G_{\mathcal{A}(G)})] \le C^*_\varepsilon$ with no other change, and this is the form stated as Theorem 4.1. The sharpening matters most for small ε, where $C^*_\varepsilon = \Theta(\varepsilon)$ while $C_\varepsilon = \Theta(\varepsilon\sqrt{\log(1/\varepsilon)})$, so that $C^*_\varepsilon / C_\varepsilon \to 0$. Both constants are tabulated after Lemma I.5. With $\tau_\varepsilon := \Phi^{-1}(1 - e^{-\varepsilon})$, the Mills ratio gives $\tau_\varepsilon \sim \sqrt{2\varepsilon}$ and $C_\varepsilon \sim \sqrt{2\varepsilon}$ as $\varepsilon \to \infty$, and $C_\varepsilon \to 0$ as $\varepsilon \to 0$, since a mechanism with no budget cannot correlate its output with the data at all. The bound is correspondingly blunt. C_ε does not depend on n or k, so relative to η_{OPT} it is informative only for $k = O(\log n/\varepsilon^2)$. Corollary 4.3 shows that in that range it is nonetheless tight up to constants.

The cap $\mathbb{E}[\mathrm{Ave}] \le (1 + o(1))\eta_{\mathrm{OPT}}$ in Corollary 4.3 follows from Lemma C.1 together with the trivial bound $\mathrm{Ave} \le \max_{i,j} |G_{ij}|$ on its failure event F. By the Cauchy–Schwarz inequality the contribution of that event is $\sqrt{\Pr[F]} \cdot O(\sqrt{\log n}) = n^{-\Omega(1)} = o(\eta_{\mathrm{OPT}})$.

Proof of Corollary 4.3. Fix $\varepsilon \le 1$ and write $y := \varepsilon\sqrt{k/\log n}$. The upper bound is Theorem 4.1 together with Lemma I.5(b) and the cap just established. By convexity of e^ε, $C^*_\varepsilon \le (e^\varepsilon - 1)/\sqrt{2\pi} \le (e-1)\varepsilon/\sqrt{2\pi} \le 0.69\,\varepsilon$ for $\varepsilon \le 1$, and dividing by $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$ gives $0.35\,y$. The lower bound is the "in particular" clause of Theorem 4.2, since $\mathcal{M}_{\mathrm{sign}}$ at budget $\min\{\varepsilon, \sqrt{\log n/k}\}$ is one ε-DP mechanism. The last statement is the expectation form of Theorem 3.2, whose hypothesis $\varepsilon\sqrt{k}/(\theta_n\sqrt{\log n}) \to \infty$ reads $\varepsilon^2 k \gg \log n \cdot \log(k \log n)$ for fixed ε. □

H Experimental Details

This appendix supports Section 6. It gives the details behind the two sweeps there and reports four further experiments that did not fit in the main text: the sign mechanism in the small-block regime, the phase diagram in k at several budgets, a comparison with three baselines from the private selection literature, and the convergence of the non-private procedures in n. For each we say what was run, what it shows, and why it is there. All code, result files and figure scripts are in the supplementary material, and the README gives the commands and running times.

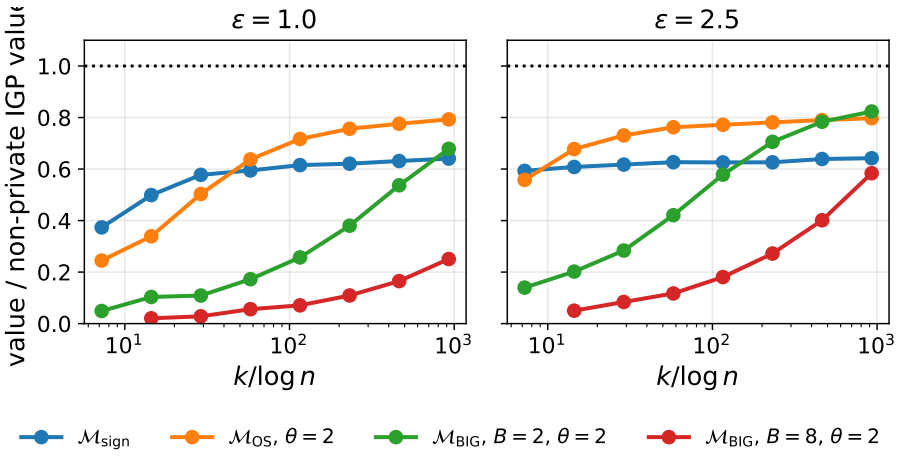


Figure 2: Experiment N4 at both budgets, $n = 10^6$: value relative to the non-private IGP at the same k, against $k/\log n$. The right panel is the right panel of Figure 1.

How the simulation works. No matrix is ever materialized. Every mechanism we run reads each entry of G at most once, on an index set that is disjoint from those of all earlier steps and does not depend on the history. For $\mathcal{M}_{\mathrm{BIG}}$ this is Lemma E.3, and the others are either one-step or use the group partition of Gamarnik & Li (2018). The joint law of scores, outputs and final average is therefore the same whether G is generated in advance or its entries are drawn at the moment they are first read. We do the latter, so n is limited by time rather than memory. Each block is drawn once, the clipped or sign column sums are computed for the scores and the true column sums for the value from the same draw, and the block is discarded. The final average is assembled from the per-step true sums of the selected rows and columns (Lemma E.4) plus the never-examined initial corner, drawn at the end. We checked this decomposition against a materialized 6000×6000 matrix for $B \in \{1, 2, 3, 4\}$, private and non-private, and checked both subset samplers against exhaustive enumeration on a five-element instance (20,000 draws, maximum deviation 0.004, within Monte Carlo error). All experiments use the rejective sampler, whose number of trials never exceeded 300 at selection sizes up to 3200 and 500 at 12800. Throughput is about 10^8 entries per second on one core.

Non-private twins and the benchmark. Every private configuration is paired with the same mechanism run without privacy, with the top-κ arg-max in place of $\mathcal{M}_{\mathrm{sel}}$ on the same clipped scoring, on independent matrices. This is the ceiling the mechanism reaches with unlimited budget at that size, and dividing by it isolates the cost of privacy from everything else. The benchmark of Section 6 is the incremental greedy procedure itself, run non-privately on $\theta = 2$ scores so that it is on the same footing as the mechanisms. Its value is the non-private twin of the per-row DP-IGP baseline below. All values are normalized by $\sqrt{2 \log n / k}$ unless marked absolute, so that 1 is the one-shot's limit and $\frac{4}{3}$ the greedy's.

The sweeps of Section 6 (experiments N3 and N4). N3 is at $n = 10^5$ ($\log n = 11.5$), $k = 1000$, $\varepsilon \in \{0.25, 0.5, 1, 2, 4, 8, 16, 32, 64\}$, with $\mathcal{M}_{\mathrm{OS}}$ at $\theta \in \{\mathrm{sign}, 1, 2\}$, $\mathcal{M}_{\mathrm{BIG}}$ with $B \in \{2, 4, 8\}$ at $\theta = 2$, the three baselines, and the IGP. N4 is at $n = 10^6$ ($\log n = 13.8$), $\varepsilon \in \{1, 2.5\}$, $k \in \{100, 200, \ldots, 12800\}$, with $\mathcal{M}_{\mathrm{OS}}$ at $\theta \in \{\mathrm{sign}, 2\}$, $\mathcal{M}_{\mathrm{BIG}}$ with $B \in \{2, 8\}$ at $\theta = 2$ where $B \le \sqrt{k/3}$, and the IGP. Both use three independent runs per configuration. Standard errors of the normalized values are at most 0.008 in N3 and 0.016 in N4, small against every gap the figures are read for. Section 6 shows N3 and the $\varepsilon = 2.5$ half of N4. Figure 2 shows both halves of N4 side by side, and the comparison is the point. At $\varepsilon = 1$ every batched curve sits where the $\varepsilon = 2.5$ curve sits at a k smaller by about $(2.5)^2$, so the crossovers that the main text shows within the sweep at $\varepsilon = 2.5$ move past its end at $\varepsilon = 1$. That is the $1/\varepsilon^2$ in the threshold of Theorem 3.2 seen directly.

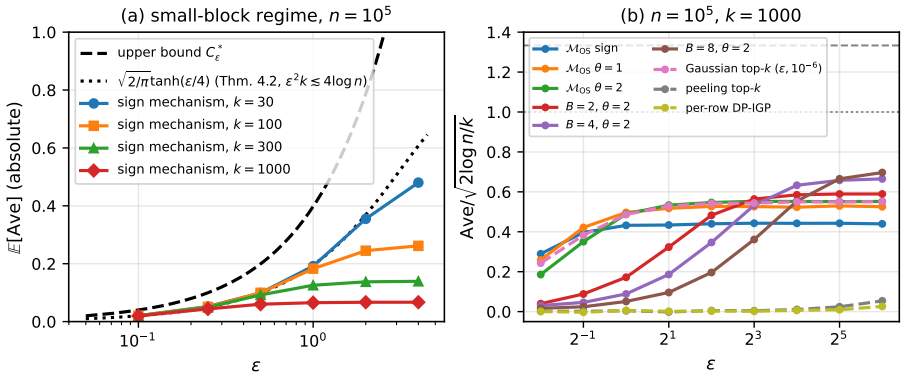


Figure 3: (a) Experiment N1: absolute expected value of the sign mechanism at $n = 10^5$ against ε, with the curve $\sqrt{2/\pi}\tanh(\varepsilon/4)$ of Theorem 4.2, which is a lower bound for $\varepsilon^2 k \lesssim 4\log n$, and the bound C^*_ε of Theorem 4.1. (b) Experiment N3 with baselines: normalized values at $n = 10^5$, $k = 1000$ for the one-shot mechanism at $\theta \in \{\mathrm{sign}, 1, 2\}$, the batched mechanism at $\theta = 2$, and the three baselines (dashed).

The sign mechanism in the small-block regime (experiment N1, Figure 3a). This is the one experiment on the constants of Section 4. We run the sign mechanism at $n = 10^5$, $k \in \{30, 100, 300, 1000\}$, $\varepsilon \in \{0.1, 0.25, 0.5, 1, 2, 4\}$, and plot its absolute expected value against ε together with the guarantee $\sqrt{2/\pi}\tanh(\varepsilon/4)$ of Theorem 4.2 and the bound C^*_ε of Theorem 4.1. The per-run standard deviation of a block average is about $1/k$ here, so small k needs many runs: 200 at $k \in \{30, 100\}$ and 20 at $k \in \{300, 1000\}$, giving standard errors at most 0.004. Two things are being tested. First, below the concentration point the value should not depend on k or n at all, only on ε, and should sit on the tanh curve, which is what the theorem's reweighting picture says. Below their concentration points the four curves coincide and track the tanh curve to within noise. Second, once $\varepsilon^2 k$ passes about $5\log n$ the mechanism concentrates on the columns with the most positive signs and its value stops at the non-private sign-selection value, which falls like $1/\sqrt{k}$, so the larger k peel off the common curve earlier and lower. Against C^*_ε the mechanism sits at a factor of about two below the bound at small budgets, growing slowly with ε, which is the $2 + O(\varepsilon)$ slack the main text leaves open.

The phase diagram in k (experiment N2, Figure 4a). Corollary 4.3 says the recoverable fraction at fixed ε behaves like $\min\{1, \varepsilon\sqrt{k/\log n}\}$ up to constants. To see the shape we fix $\varepsilon \in \{0.5, 1, 2\}$ at $n = 10^5$ and sweep $k \in \{25, 50, \ldots, 3200\}$, running the sign mechanism, $\mathcal{M}_{\mathrm{OS}}$ at $\theta = 2$, and $\mathcal{M}_{\mathrm{BIG}}$ with $B = 4$ where allowed, three runs each. Standard errors grow at small k, to at most 0.04 at $k = 25$. The sign mechanism's normalized value rises like $\varepsilon\sqrt{k/\log n}$, peaks near $k \asymp \log n/\varepsilon^2$, and then declines along its ceiling, and the peak moves left as ε grows, by a factor between two and four per doubling on this grid, against the four the corollary predicts. The $\theta = 2$ one-shot rises with the same slope, later, and then follows the non-private one-shot ceiling, which sits above the sign mechanism's, which is the crossover of Appendix B between signs and magnitudes. This is the phase diagram at one n, with ε as a parameter. The $\mathcal{M}_{\mathrm{BIG}}$ runs at these budgets are still far from their ceilings at $k = 3200$, which is why N4 goes to $n = 10^6$ and larger k.

Baselines (Figure 3b, Table 3). The natural alternatives to $\mathcal{M}_{\mathrm{sel}}$ are the private top-k selection methods, which pick one item at a time and pay for each. We run three at $n = 10^5$, $k = 1000$, all at $\theta = 2$. Peeling top-k (Bhaskar et al., 2010; Durfee & Rogers, 2019) selects k columns against the fixed rows by k sequential single-item exponential mechanisms with budget ε/k each. Per-row DP-IGP privatizes the incremental greedy procedure step by step, $2k - 1$ single-item exponential mechanisms with budget $\varepsilon/(2k - 1)$ each, which is the direct privatization that Section 3.4 argues against. Gaussian top-k adds Gaussian noise to every column score and takes the top k of the noisy scores, with the noise calibrated by the analytic Gaussian mechanism of Balle & Wang (2018) to be $(\varepsilon, 10^{-6})$-DP. This is a weaker privacy notion than ours, and it is the strongest alternative we

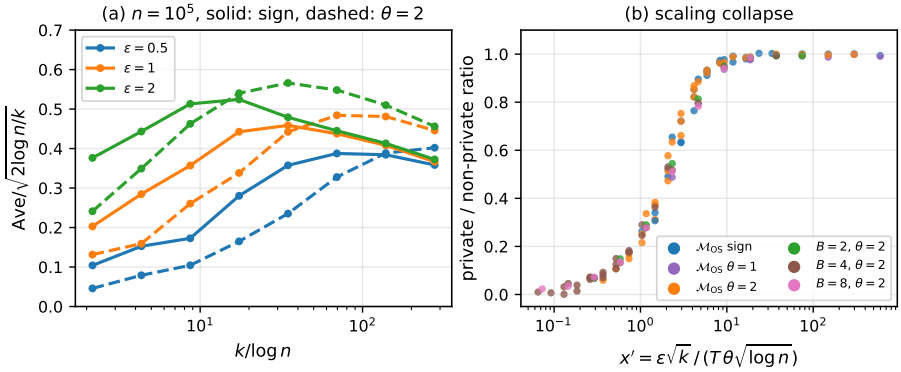


Figure 4: (a) Experiment N2: normalized value against $k/\log n$ at fixed ε, $n = 10^5$, for the sign mechanism and the one-shot with $\theta = 2$. (b) Scaling collapse: private-to-non-private ratio of all 123 private configurations of N2 and N3 against $x' = \varepsilon\sqrt{k}/(T\theta\sqrt{\log n})$, with T the number of private steps.

know of. The table records what happens. The two composition-based baselines stay within a few hundredths of zero at every budget, since their $T = \Theta(k)$ steps put them at $x' \leq 0.3$ in the scaling variable below, while their non-private twins are as good as anything in the paper. That is the cost of composition at this size, and it is what batching removes. Gaussian top-k, despite its δ, is indistinguishable from the pure-DP one-shot at the same clipping level and shares its ceiling, so the approximate-DP relaxation buys nothing here.

Table 3: Experiment N3 in numbers: normalized values at $n = 10^5$, $k = 1000$ (means over three runs, s.e. ≤ 0.01). "∞" is the non-private twin of the same mechanism, and T is the number of private steps.

mechanism	ε=0.25	ε=1	ε=4	ε=16	ε=64	∞
sign (T=1)	0.29	0.43	0.44	0.44	0.44	0.44
θ=1 (T=1)	0.26	0.50	0.53	0.52	0.53	0.53
θ=2 (T=1)	0.19	0.49	0.55	0.55	0.55	0.55
B=2, θ=2 (T=4)	0.04	0.17	0.48	0.59	0.59	0.59
B=4, θ=2 (T=8)	0.03	0.09	0.35	0.63	0.67	0.67
B=8, θ=2 (T=16)	0.02	0.05	0.20	0.55	0.70	0.71
Gaussian top-k, (ε, 10^{-6}), θ=2	0.24	0.49	0.54	0.55	0.55	0.55
peeling, θ=2 (T=k)	0.00	0.00	0.01	0.01	0.05	0.55
per-row DP-IGP, θ=2 (T=2k−1)	0.00	0.01	0.00	0.01	0.03	0.69

The scaling collapse (Figure 4b). Propositions D.1 and 3.3 say the relative privacy penalty is of order $T\theta\sqrt{\log n}/(\varepsilon\sqrt{k})$, with T the number of private steps, 1 for the one-shot mechanisms and $2B$ for $\mathcal{M}_{\mathrm{BIG}}$. If that is right, the private-to-non-private ratio of any configuration should depend on $(n, k, \varepsilon, \theta, B)$ only through $x' = \varepsilon\sqrt{k}/(T\theta\sqrt{\log n})$. We test this by plotting the ratio of all 123 private configurations of N2 and N3 against x', taking $\theta = 1$ for the sign mechanism (in the equivalent ± 1 formulation its sensitivity is 2 and its scores have variance k, as for clipping at $\theta = 1$). They fall on one curve. Binned by x', the mean ratio and within-bin standard deviation are 0.02 and 0.02 on $[0.05, 0.2)$, 0.07 and 0.01 on $[0.2, 0.5)$, 0.15 and 0.03 on $[0.5, 1)$, 0.30 and 0.05 on $[1, 2)$, 0.59 and 0.09 on $[2, 4)$, 0.86 and 0.06 on $[4, 8)$, 0.96 and 0.02 on $[8, 16)$, and 0.99 and 0.01 beyond. The variable is the theorems' own, so this is a consistency check rather than independent evidence. It is a useful one all the same: it says the analysis has the right combination of parameters, and it gives a rule of thumb. Recovering 90% of a mechanism's non-private value needs $x' \gtrsim 5$, so for $\mathcal{M}_{\mathrm{OS}}$ with $\theta = 2$ at $n = 10^5$ this is $k \gtrsim 1200/\varepsilon^2$, and for $\mathcal{M}_{\mathrm{BIG}}$ it is $(2B)^2$ times that.

Non-private convergence in n (Figure 5). The last experiment has no privacy in it. We run the one-shot, the batched greedy with $B \in \{2, 4, 8, 16, 32\}$ (the larger B exceed the $\sqrt{k/3}$ of Theorem 3.2, but the schedule is still well defined), and the IGP with arg-max on unclipped scores

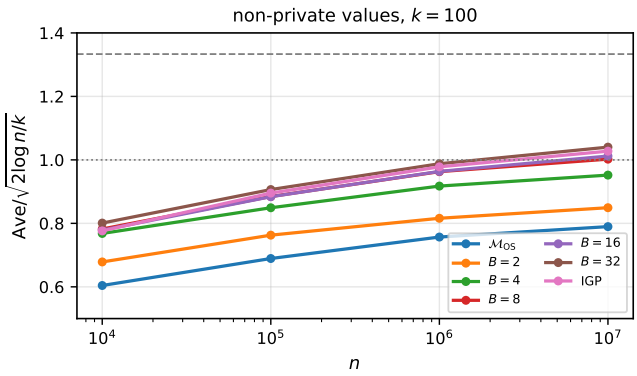


Figure 5: Non-private values (arg-max selection on the unclipped scores) as n grows from 10^4 to 10^7 at $k = 100$, for the one-shot and batched procedures and the IGP. The dotted and dashed lines are the limits 1 and $\frac{4}{3}$. Two runs per point.

at $k = 100$ and n from 10^4 to 10^7, two runs per point, and plot the normalized value against n. The limits are 1 and $\frac{4}{3}$, and nothing gets near them: at $n = 10^7$ the one-shot is at about 0.8 and the IGP at about 1.03. The approach is logarithmic in n, since each step takes the best of about n/k Gaussian sums and the expected maximum of m Gaussians trails $\sqrt{2 \log m}$ by a correction of order $\log \log m / \sqrt{\log m}$. This is why the main text measures ratios to the non-private procedures at the same n rather than distances to the limiting constants, which no simulation reaches. The ratios converge much faster because the slow factor cancels: $\mathcal{M}_{\mathrm{BIG}}(B{=}32)/\mathcal{M}_{\mathrm{OS}}$ is already 1.32 at $n = 10^7$ against the limit $\frac{4}{3}$. The spread across B at fixed n is the $(1 - \frac{2}{B})$ initial-slot loss of Proposition 3.3, and $B = 32$ sitting just above the IGP is the finite-n advantage of taking the best κ from one pool over κ separate best-of-one picks, noted in Section 6.

I Row-Level Privacy: Matching Thresholds

Here we prove Theorems 5.1 and 5.2 of Section 5, where a single record is an entire row of G. The upper bound is $\mathbb{E}[\mathrm{Ave}(G_{\mathcal{A}(G)})] \le C^*_\varepsilon/\sqrt{k}$ for every ε-row-DP mechanism $\mathcal{A}$, with C^*_ε depending on ε alone (Theorem I.6), so no mechanism recovers a constant fraction of η_{OPT} unless $\varepsilon = \Omega(\log n)$, whatever k is (Corollary I.7). The matching mechanism recovers $(1/\sqrt{2} - o(1))\eta_{\mathrm{OPT}}$ whenever $\varepsilon = \omega(\log n)$ (Theorem I.10), so the threshold $\varepsilon \asymp \log n$ is sharp up to constants. We use Lemmas C.5, C.6 and C.10 and Proposition C.15 of Appendix C, and Lemma G.1 of Appendix G. The constant C^*_ε comes from a two-sided version of the bathtub lemma, which uses both inequalities of the privacy definition, and the same sharpening gives the constant of Theorem 4.1 in the main text. Appendix A relates these results to the private graph, user-level privacy and lower-bound literatures.

I.1 Model

Definition I.1 (Row neighbors and ε-row-DP)**.** Matrices $G, G' \in \mathbb{R}^{n \times n}$ are *row neighbors*, written $G \sim_{\mathrm{r}} G'$, if they agree in every row but one. The differing row may take any two values in $\mathbb{R}^n$. A randomized mechanism $\mathcal{A}$ is *ε-row-DP* if $\Pr[\mathcal{A}(G) \in S] \le e^\varepsilon \Pr[\mathcal{A}(G') \in S]$ for all $G \sim_{\mathrm{r}} G'$ and all measurable S.

A single-entry change is a row change, so every ε-row-DP mechanism is ε-DP in the sense of Section 2. The converse fails, and the mechanisms of the main text are not ε-row-DP. Throughout this section G has i.i.d. $N(0,1)$ entries, $X_i \in \mathbb{R}^n$ is the i-th row, mechanisms output index pairs (I, J) with $|I| = |J| = k$ (a finite output space), $k = \exp(o(\log n))$, and the internal randomness of a mechanism is independent of G.

The privacy lemma below is the row-level analogue of Lemma C.5, and it is what makes $\mathcal{M}_{\mathrm{RO}}$ private. The sensitivity of a subset score is again that of a single summand, since a row change alters a single summand.

Lemma I.2 (Exponential mechanism over subsets, row-level)**.** *Let $f_1, \dots, f_n : \mathbb{R}^n \to \mathbb{R}$ be functions with $\sup_{x,x'} |f_i(x) - f_i(x')| \le \Delta'$ for every i, let $P \subseteq [n]$, $\kappa \le |P|$ and $\varepsilon' > 0$. The mechanism which outputs $R \subseteq P$ with $|R| = \kappa$ and $\Pr[R] \propto \exp\left(\frac{\varepsilon'}{2\Delta'} \sum_{i \in R} f_i(X_i)\right)$ is ε'-row-DP.*

Proof. If $G \sim_{\mathrm{r}} G'$ differ in row i_0, the score $u(R) = \sum_{i \in R} f_i(X_i)$ changes by $|f_{i_0}(X_{i_0}) - f_{i_0}(X'_{i_0})| \cdot \mathbf{1}[i_0 \in R] \le \Delta'$ for every R. The rest is the computation in the proof of Lemma C.5, with Δ replaced by Δ'. □

I.2 The two-sided bathtub lemma and the constant C^*_ε

The bathtub lemma of Appendix G uses only the upper half of the privacy constraint on a single record. Using both halves gives a smaller constant, and that constant is the C^*_ε of Section 4.

Definition I.3. For $\varepsilon > 0$ let

$$C^*_\varepsilon := \sup_{a \in (0,1)} \frac{(e^\varepsilon - 1)\,\phi\big(\Phi^{-1}(1-a)\big)}{1 + a(e^\varepsilon - 1)}.$$

Lemma I.4 (Two-sided bathtub)**.** *Let $Z \sim N(0,1)$ and let $q : \mathbb{R} \to [0, \infty)$ be measurable with $\mathbb{E}[q(Z)] = \mu > 0$ and $\sup q \le e^\varepsilon \inf q$. Then $\mathbb{E}[Zq(Z)] \le C^*_\varepsilon\, \mu$.*

Proof. Set $L := \inf q$. Since $\mu > 0$ and $\sup q \le e^\varepsilon L$, we have $L > 0$, so $L \le q \le e^\varepsilon L$ and $L \le \mu \le e^\varepsilon L$. Pick $a \in [0,1]$ with $L(1 + a(e^\varepsilon - 1)) = \mu$, which exists because the left side runs over $[L, e^\varepsilon L]$ as a runs over $[0,1]$. If $a \in \{0,1\}$ then $\mu \in \{L, e^\varepsilon L\}$, which forces q to be constant almost surely, so $\mathbb{E}[Zq(Z)] = 0$ and there is nothing to prove. So assume $a \in (0,1)$, let $s := \Phi^{-1}(1-a) \in \mathbb{R}$, and let q^* be the two-level function equal to $e^\varepsilon L$ on $\{Z > s\}$ and to L on $\{Z \le s\}$, so that $\mathbb{E}[q^*(Z)] = e^\varepsilon La + L(1-a) = \mu$. For every z we have $(z-s)(q^*(z) - q(z)) \ge 0$: if $z > s$ then $q^*(z) = e^\varepsilon L \ge q(z)$, and if $z \le s$ then $q^*(z) = L \le q(z)$. Taking expectations gives $\mathbb{E}[Z(q^*(Z) - q(Z))] \ge s\,\mathbb{E}[q^*(Z) - q(Z)] = 0$. Since $\mathbb{E}[Z\mathbf{1}_{Z>s}] = \phi(s)$ and $\mathbb{E}[Z\mathbf{1}_{Z \le s}] = -\phi(s)$,

$$\mathbb{E}[Zq(Z)] \le \mathbb{E}[Zq^*(Z)] = (e^\varepsilon L - L)\phi(s) = \mu \cdot \frac{(e^\varepsilon - 1)\phi(\Phi^{-1}(1-a))}{1 + a(e^\varepsilon - 1)} \le C^*_\varepsilon\, \mu.$$

□

Lemma I.5 (Properties of C^*_ε)**.** *(a) C^*_ε is nondecreasing in ε.*

(b) $\dfrac{2\tanh(\varepsilon/2)}{\sqrt{2\pi}} \le C^*_\varepsilon \le \dfrac{e^\varepsilon - 1}{\sqrt{2\pi}}$. *In particular $C^*_\varepsilon = \dfrac{\varepsilon}{\sqrt{2\pi}}(1 + O(\varepsilon))$ as $\varepsilon \to 0$.*

(c) $C^*_\varepsilon \le C_\varepsilon = e^\varepsilon \phi(\Phi^{-1}(1 - e^{-\varepsilon}))$, the constant of Theorem G.2.

(d) For $\varepsilon \ge 2$, $C^*_\varepsilon \le \sqrt{2\varepsilon} + 1$.

(e) $C^*_\varepsilon \ge \frac{1 - e^{-\varepsilon}}{2} C_\varepsilon$. Together with (c) and $C_\varepsilon = (1 + o(1))\sqrt{2\varepsilon}$ as $\varepsilon \to \infty$, this gives $C^*_\varepsilon = \Theta(\sqrt{\varepsilon})$ for large ε.

Proof. (a) For each fixed a the ratio $(e^\varepsilon - 1)/(1 + a(e^\varepsilon - 1))$ is increasing in ε, hence so is the supremum.

(b) The upper bound uses $\phi \le \phi(0) = 1/\sqrt{2\pi}$ and a denominator of at least 1. For the lower bound take $a = \frac{1}{2}$, so that $\Phi^{-1}(1-a) = 0$ and the ratio is $\frac{e^\varepsilon - 1}{1 + (e^\varepsilon - 1)/2}\phi(0) = \frac{2(e^\varepsilon - 1)}{(e^\varepsilon + 1)\sqrt{2\pi}}$. Both bounds equal $\frac{\varepsilon}{\sqrt{2\pi}}(1 + O(\varepsilon))$ as $\varepsilon \to 0$.

(c) The constraint $\sup q \le e^\varepsilon \inf q$ together with $\mathbb{E}q = \mu$ implies $q \le e^\varepsilon \mu$, which is the hypothesis of Lemma G.1 with $U = e^\varepsilon \mu$. Hence $\mathbb{E}[Zq] \le U\phi(\Phi^{-1}(1 - \mu/U)) = C_\varepsilon \mu$ for every admissible q. Each term of the supremum defining C^*_ε is $\mathbb{E}[Zq^*]/\mathbb{E}[q^*]$ for the admissible two-level function $q^* = e^\varepsilon \mathbf{1}_{Z>s} + \mathbf{1}_{Z \le s}$ with $s = \Phi^{-1}(1-a)$, as computed in the proof of Lemma I.4. Hence $C^*_\varepsilon \le C_\varepsilon$.

(d) By (c) it suffices to bound C_ε. Let $\tau_\varepsilon := \Phi^{-1}(1 - e^{-\varepsilon})$, so that $1 - \Phi(\tau_\varepsilon) = e^{-\varepsilon}$. For $\varepsilon \geq 2$ we have $\tau_\varepsilon \geq \Phi^{-1}(1 - e^{-2}) > 1$. The tail bound $1 - \Phi(\tau) \leq e^{-\tau^2/2}$ for $\tau \geq 0$ gives $\tau_\varepsilon \leq \sqrt{2\varepsilon}$, and the Mills-ratio bound $1 - \Phi(\tau) \geq \frac{\tau}{1+\tau^2}\phi(\tau)$ gives $\phi(\tau_\varepsilon) \leq e^{-\varepsilon}(1 + \tau_\varepsilon^2)/\tau_\varepsilon$. Hence $C_\varepsilon = e^\varepsilon \phi(\tau_\varepsilon) \leq \tau_\varepsilon + 1/\tau_\varepsilon \leq \sqrt{2\varepsilon} + 1$.

(e) Take $a = e^{-\varepsilon}$ in Definition I.3. Then $\Phi^{-1}(1 - a) = \tau_\varepsilon$, the numerator is $(e^\varepsilon - 1)\phi(\tau_\varepsilon) = (1 - e^{-\varepsilon})C_\varepsilon$, and the denominator is $1 + (e^\varepsilon - 1)e^{-\varepsilon} \leq 2$. For the asymptotics of C_ε, the bound $1 - \Phi(\tau) \leq \phi(\tau)/\tau$ gives $\phi(\tau_\varepsilon) \geq \tau_\varepsilon e^{-\varepsilon}$, so $C_\varepsilon \geq \tau_\varepsilon$, and the Mills-ratio bound gives $\tau_\varepsilon^2 \geq 2\varepsilon - \log(2\pi) - 2\log\frac{1+\tau_\varepsilon^2}{\tau_\varepsilon} = 2\varepsilon - O(\log\varepsilon)$. □

The supremum in Definition I.3 is a one-dimensional maximization and is easy to evaluate numerically: $C^*_{0.1} = 0.040$, $C^*_{0.25} = 0.100$, $C^*_1 = 0.398$, $C^*_2 = 0.787$, $C^*_4 = 1.516$, $C^*_8 = 2.743$ and $C^*_{16} = 4.561$, against $C_{0.1} = 0.187$, $C_{0.25} = 0.381$, $C_1 = 1.024$, $C_2 = 1.607$, $C_4 = 2.453$, $C_8 = 3.658$ and $C_{16} = 5.358$. The maximizing a decreases from about $\frac{1}{2}$ at small ε to 0 as $\varepsilon \to \infty$, where $C^*_\varepsilon = \Theta(\sqrt{\varepsilon})$ by (e). The numerical ratio $C^*_\varepsilon/\sqrt{2\varepsilon}$ is 0.69 at $\varepsilon = 8$ and 0.81 at $\varepsilon = 16$, and we have not determined whether it tends to 1.

I.3 The upper bound

The bound is the argument of Theorem G.2 run on rows instead of entries. Privacy now constrains how the inclusion of row i depends on its own data, and what row i contributes to the block sum is a Gaussian of scale $\sqrt{k}$ rather than 1.

Theorem I.6 (Expectation bound for every ε-row-DP mechanism, Theorem 5.2)**.** *Let $\mathcal{A}$ be any ε-row-DP mechanism outputting (I, J) with $|I| = |J| = k$. Then*

$$\mathbb{E}\big[\operatorname{Ave}(G_{\mathcal{A}(G)})\big] \;\leq\; \frac{C^*_\varepsilon}{\sqrt{k}}.$$

Proof. Since $\operatorname{Sum}(G_{I,J}) = \sum_{i\in I}\langle X_i, \mathbf{1}_J\rangle$, we have

$$\mathbb{E}\big[\operatorname{Sum}(G_{\mathcal{A}(G)})\big] = \sum_{i=1}^{n} \mathbb{E}\big[\mathbf{1}[i \in I]\,\langle X_i, \mathbf{1}_J\rangle\big], \tag{12}$$

where $(I, J) = \mathcal{A}(G)$. Each expectation is finite, since $|\langle X_i, \mathbf{1}_J\rangle| \leq \|X_i\|_1$ is integrable. Fix i. We bound the i-th term in three steps.

Step 1: privacy as a constraint on one record. For a fixed output (I, J) define

$$q_{IJ}(x) := \Pr\big[\mathcal{A}(x, G_{-i}) = (I, J)\big] \qquad (x \in \mathbb{R}^n),$$

the probability over the internal randomness and over the other rows G_{-i}, and let $\mu_{IJ} := \mathbb{E}[q_{IJ}(X_i)] = \Pr[\mathcal{A}(G) = (I, J)]$. For any $x, x' \in \mathbb{R}^n$ and any fixed value g_{-i} of the other rows, $(x, g_{-i}) \sim_{\mathrm{r}} (x', g_{-i})$, so ε-row-DP gives $\Pr_{\mathcal{A}}[\mathcal{A}(x, g_{-i}) = (I, J)] \leq e^\varepsilon \Pr_{\mathcal{A}}[\mathcal{A}(x', g_{-i}) = (I, J)]$. The rows are independent, so the law of G_{-i} does not depend on x or x', and averaging over $g_{-i} \sim G_{-i}$ gives

$$q_{IJ}(x) \leq e^\varepsilon q_{IJ}(x') \quad \text{for all } x, x' \in \mathbb{R}^n, \qquad \text{i.e.} \qquad \sup q_{IJ} \leq e^\varepsilon \inf q_{IJ}. \tag{13}$$

In particular, either $q_{IJ} \equiv 0$, so that $\mu_{IJ} = 0$, or $q_{IJ} > 0$ everywhere.

Step 2: reduction to a scalar. Fix (I, J) with $i \in I$ and $\mu_{IJ} > 0$. Since $|J| = k$ and the entries of X_i are i.i.d. $N(0, 1)$, the scalar $Z := \langle X_i, \mathbf{1}_J\rangle/\sqrt{k}$ is standard normal. Let $\tilde{q}(z) := \mathbb{E}[q_{IJ}(X_i) \mid Z = z]$. Then $\tilde{q}$ takes values in $[\inf q_{IJ}, \sup q_{IJ}]$ almost surely, so $\sup\tilde{q} \leq e^\varepsilon \inf\tilde{q}$ after modification on a null set, $\mathbb{E}[\tilde{q}(Z)] = \mu_{IJ}$, and by the tower property $\mathbb{E}[Z q_{IJ}(X_i)] = \mathbb{E}[Z\tilde{q}(Z)]$. Lemma I.4 therefore gives

$$\mathbb{E}\big[\langle X_i, \mathbf{1}_J\rangle\, q_{IJ}(X_i)\big] = \sqrt{k}\,\mathbb{E}[Z\tilde{q}(Z)] \leq \sqrt{k}\, C^*_\varepsilon\, \mu_{IJ}. \tag{14}$$

When $\mu_{IJ} = 0$ the same inequality holds trivially.

Step 3: summation. Conditionally on $X_i = x$, the event $\{\mathcal{A}(G) = (I, J)\}$ has probability $q_{IJ}(x)$. By the tower property and Fubini's theorem, justified by $\sum_{(I,J)} \mathbb{E}[|\langle X_i, \mathbf{1}_J\rangle| q_{IJ}(X_i)] \le \mathbb{E}\|X_i\|_1 < \infty$,

$$\mathbb{E}\big[\mathbf{1}[i \in I]\langle X_i, \mathbf{1}_J\rangle\big] = \sum_{(I,J):\, i\in I} \mathbb{E}\big[\langle X_i, \mathbf{1}_J\rangle q_{IJ}(X_i)\big] \le \sqrt{k}\, C_\varepsilon^* \sum_{(I,J):\, i\in I} \mu_{IJ} = \sqrt{k}\, C_\varepsilon^* \Pr[i \in I].$$

Substituting into (12) and using $\sum_i \Pr[i \in I] = \mathbb{E}|I| = k$ gives $\mathbb{E}[\mathrm{Sum}(G_{\mathcal{A}(G)})] \le k^{3/2} C_\varepsilon^*$. Dividing by k^2 completes the proof. □

Corollary I.7 (Vanishing fraction unless $\varepsilon = \Omega(\log n)$, for every k)**.** *For every ε-row-DP mechanism $\mathcal{A}$ and every k,*

$$\frac{\mathbb{E}[\mathrm{Ave}(G_{\mathcal{A}(G)})]}{\eta_{\mathrm{OPT}}} \;\le\; \frac{C_\varepsilon^*}{2\sqrt{\log n}} \;\le\; \frac{\sqrt{2\varepsilon}+1}{2\sqrt{\log n}} \quad (\varepsilon \ge 2).$$

Hence if $\varepsilon = \varepsilon(n) = o(\log n)$ then $\mathbb{E}[\mathrm{Ave}(G_{\mathcal{A}(G)})]/\eta_{\mathrm{OPT}} \to 0$, uniformly over all ε-row-DP mechanisms and all sequences $k = k(n)$, and if $\varepsilon \le (2-\delta)\log n$ for a fixed $\delta \in (0,2)$ then, for large n, no ε-row-DP mechanism has $\mathbb{E}[\mathrm{Ave}] \ge (1 - \delta/4)\eta_{\mathrm{OPT}}$.

Proof. The first inequality is Theorem I.6 divided by $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$, and the second is Lemma I.5(d). For the two remaining claims, when $\varepsilon < 2$ we first replace ε by 2 using Lemma I.5(a). If $\varepsilon = o(\log n)$ the right side tends to 0. If $2 \le \varepsilon \le (2-\delta)\log n$ then $(\sqrt{2\varepsilon}+1)/(2\sqrt{\log n}) \le \sqrt{1-\delta/2} + o(1) < 1 - \delta/4$ for large n, since $(1-\delta/4)^2 = 1 - \delta/2 + \delta^2/16 > 1 - \delta/2$. □

The contrast with Corollary 4.3 is the point. Under entry-level privacy the bound C_ε^* becomes vacuous relative to $\eta_{\mathrm{OPT}} = 2\sqrt{\log n/k}$ as soon as $k \ge 4\log n/(C_\varepsilon^*)^2$, which for fixed ε is $k = \Omega(\log n)$, and Sections 3.1 and 3 show that large k does rescue a fixed ε. Under row-level privacy the extra factor $1/\sqrt{k}$ cancels the $1/\sqrt{k}$ in η_{OPT}, and no k helps.

I.4 The reverse one-shot mechanism

This is the mechanism of Theorem 5.1. It fixes the columns and selects the rows, so that every score is a function of one record. A row can be arbitrarily large, so the scores must be clipped. The clipping level is the high-probability scale of a row sum over k columns, $\sqrt{k\log n}$, and its ratio to the typical scale $\sqrt{k}$ of a row sum is the factor $\sqrt{\log n}$ that sets the threshold.

Fix $J_0 := [k]$ and a constant $c > \sqrt{2}$, and set $\tau := c\sqrt{k\log n}$. For $i \in [n]$ let $t_i := \mathrm{Sum}(G_{\{i\},J_0}) = \sum_{j\le k} G_{ij}$ and $\tilde{t}_i := \max(-\tau, \min(\tau, t_i))$.

Definition I.8 (Reverse one-shot mechanism $\mathcal{M}_{\mathrm{RO}}$)**.** Output $R \subseteq [n]$ with $|R| = k$ and $\Pr[R] \propto \exp\big(\frac{\varepsilon}{4\tau}\sum_{i\in R}\tilde{t}_i\big)$, and return (R, J_0).

Proposition I.9 (Privacy and runtime)**.** *$\mathcal{M}_{\mathrm{RO}}$ is ε-row-DP for all real inputs and runs in $O(nk)$ time.*

Proof. Apply Lemma I.2 with $f_i(x) = \max(-\tau, \min(\tau, \sum_{j\le k} x_j))$, which has range $[-\tau, \tau]$ and hence $\Delta' = 2\tau$, with $P = [n]$, $\kappa = k$ and $\varepsilon' = \varepsilon$. The mechanism's parameter $\frac{\varepsilon}{4\tau}$ equals $\frac{\varepsilon'}{2\Delta'}$. Computing all t_i costs $O(nk)$, and the sample is the distribution of Proposition C.15 with $m = n$, so either sampler runs in $O(nk)$ time. □

Theorem I.10 (Utility of $\mathcal{M}_{\mathrm{RO}}$, Theorem 5.1)**.** *For every fixed $\zeta \in (0, \frac{1}{2}]$, whp,*

$$\mathrm{Ave}(G_{R,J_0}) \;\ge\; \Big(\frac{1}{\sqrt{2}} - \zeta\Big)\eta_{\mathrm{OPT}} \;-\; \frac{4c\sqrt{k\log n}\,\big(k\log(en/k) + \log n\big)}{\varepsilon\, k^2}.$$

The penalty is $\Theta\big((\log n)^{3/2}/(\varepsilon\sqrt{k})\big)$, and its ratio to $\eta_{\mathrm{OPT}}/\sqrt{2} = \sqrt{2\log n/k}$ is $\Theta(\log n/\varepsilon)$, independent of k. Consequently, if $\varepsilon = \omega(\log n)$ then for every fixed $\delta > 0$, $\Pr[\mathrm{Ave}(G_{R,J_0}) \ge (1/\sqrt{2} - \delta)\eta_{\mathrm{OPT}}] \to 1$ and $\mathbb{E}[\mathrm{Ave}(G_{R,J_0})] \ge (1/\sqrt{2} - o(1))\eta_{\mathrm{OPT}}$, and if $\varepsilon = A\log n$ with a constant $A > 2\sqrt{2}c$, then whp $\mathrm{Ave}(G_{R,J_0}) \ge \frac{1}{\sqrt{2}}\big(1 - \frac{2\sqrt{2}c}{A} - \zeta\big)\eta_{\mathrm{OPT}}$.

Proof. The proof follows that of Proposition D.1, with the clipping handled by showing that whp it never occurs.

No clipping whp. Each $t_i \sim N(0,k)$, so $\Pr[|t_i| > \tau] \le 2e^{-\tau^2/(2k)} = 2n^{-c^2/2}$, and a union bound over the n rows gives $\Pr[\tilde{t} \neq t] \le 2n^{1-c^2/2} \to 0$ since $c > \sqrt{2}$. We work on the event $\tilde{t} = t$.

The benchmark. We have $\max_{|R|=k} \sum_{i\in R} t_i = \sqrt{k} \sum_{l\le k} Z_{(l)}$, where $Z_{(l)}$ are the order statistics of the i.i.d. standard normal variables $Z_i = t_i/\sqrt{k}$. Lemma C.10 with $(m,\kappa) = (n,k)$ and $\zeta/2$ in place of ζ applies since $k = \exp(o(\log n))$, and gives $\sum_{l\le k} Z_{(l)} \ge (1-\zeta/2)k\sqrt{2\log(n/k)}$ with probability $1 - \exp(-(n/k)^{\zeta/4})$. Since $\log(n/k) = (1-o(1))\log n$, for large n we get whp $\max_R \sum_{i\in R} t_i \ge (1-\zeta)k^{3/2}\sqrt{2\log n}$.

The mechanism. Lemma C.6 with $N = \binom{n}{k} \le (en/k)^k$, $\beta = \varepsilon/(4\tau)$ and $t = \log n$ gives, with probability at least $1 - 1/n$,

$$\sum_{i\in R} t_i \;\ge\; \max_{R'} \sum_{i\in R'} t_i - \frac{4\tau}{\varepsilon}\big(k\log(en/k) + \log n\big).$$

Combining. Dividing by k^2 and using $k^{-1/2}\sqrt{2\log n} = \eta_{\mathrm{OPT}}/\sqrt{2}$ and $\tau = c\sqrt{k\log n}$ gives the displayed bound. For the ratio, $4c\sqrt{k\log n}\cdot k\log(en/k)(1+o(1))/(\varepsilon k^2)$ divided by $\sqrt{2\log n/k}$ equals $\frac{2\sqrt{2}c\log n}{\varepsilon}(1+o(1))$, since $\log(en/k) = (1+o(1))\log n$. The $\log n$ term in the penalty is smaller by a factor of k. The whp statements for $\varepsilon = \omega(\log n)$ and $\varepsilon = A\log n$ follow.

The expectation. Let F be the failure event of the three whp bounds, so that $\Pr[F] \le 2n^{1-c^2/2} + \exp(-(n/k)^{\zeta/4}) + 1/n = n^{-\Omega(1)}$ as $c > \sqrt{2}$. Since $\mathrm{Ave}(G_{R,J_0}) = \sum_{i\in R} t_i/k^2 \ge -\max_i |t_i|/k$, the Cauchy–Schwarz inequality gives

$$\mathbb{E}\big[\,\mathrm{Ave}(G_{R,J_0})\mathbf{1}_F\big] \ge -\frac{\sqrt{\mathbb{E}[\max_i t_i^2]\Pr[F]}}{k} \ge -\frac{\sqrt{k(4\log n + 2\log 2)\Pr[F]}}{k},$$

where we have used $\mathbb{E}[\max_i Z_i^2] \le 4\log\big(\sum_i \mathbb{E}e^{Z_i^2/4}\big) = 4\log(n\sqrt{2})$, by Jensen's inequality applied to $\exp(\cdot/4)$ together with $\mathbb{E}e^{Z^2/4} = \sqrt{2}$. This is $O(\sqrt{\log n/k}\cdot n^{-\Omega(1)}) = o(\sqrt{\log n/k})$, while on F^c the displayed lower bound holds, so $\mathbb{E}[\mathrm{Ave}] \ge (1-o(1))\cdot(\text{bound}) - o(\eta_{\mathrm{OPT}})$, which under $\varepsilon = \omega(\log n)$ is $(1/\sqrt{2} - o(1))\eta_{\mathrm{OPT}}$. □

Corollary I.7 and Theorem I.10 together pin down, up to constants, the privacy budget at which a constant fraction of η_{OPT} becomes achievable under row-level privacy: it is $\varepsilon \asymp \log n$, with no dependence on k. The $1/\sqrt{k}$ of Theorem I.6 is matched too, since $\mathcal{M}_{\mathrm{RO}}$ attains $\Theta(\sqrt{\log n/k})$ once $\varepsilon \ge A\log n$ with a constant $A > 2\sqrt{2}c$. What is not pinned down is the constant. Corollary I.7 is informative for $\varepsilon < 2\log n$, and Theorem I.10 applies for $\varepsilon > 2\sqrt{2}c\log n > 4\log n$.

Theorem I.6 is for pure differential privacy. Under (ε,δ)-row-DP, (13) becomes $q_{IJ}(x) \le e^{\varepsilon} q_{IJ}(x') + \delta$, and the additive δ terms are summed over $\binom{n}{k}^2$ outputs in Step 3, which yields nothing. A planted-instance argument suggests the same threshold in the worst case. Plant k rows with row sum τ and set the rest to 0. For a fixed configuration of the other rows, changing row i from planted to null is a row change, so per-row (ε,δ)-privacy bounds the probability of selecting a planted row by e^{ε} times that of selecting a null row plus δ, and averaging over a uniformly random planted set gives $\mathbb{E}|R\cap P| \lesssim e^{\varepsilon}k^2/n + k\delta$, i.e. a vanishing fraction of the planted rows recovered when $\varepsilon \le (1-\delta')\log(n/k)$. We have not carried out the bookkeeping of the conditioning on the other rows and state this only as a heuristic. Whether the ensemble bound extends to approximate DP, presumably via fingerprinting-type arguments (Bun et al., 2014; Narayanan, 2023; Peters et al., 2024) rather than the single-record correlation bound used here, we leave open. The theorem also assumes $|I| = |J| = k$. The bound $\mathbb{E}[\mathrm{Sum}(G_{\mathcal{A}(G)})] \le k^{3/2}C^*_{\varepsilon}$ extends verbatim to outputs with $|I|, |J| \le k$, and the theorem extends to any row distribution with independent rows, with $C^*_{\varepsilon}\sqrt{k}$ replaced by the corresponding two-level rearrangement functional of the row-sum distribution.

J INDEPENDENT SUB-GAUSSIAN ENTRIES

Section 2 takes the entries of G to be i.i.d. $N(0,1)$. Here we ask, as promised in the introduction, which results survive when they are i.i.d. copies of a general mean-zero, unit-variance, sub-Gaussian random variable. The mechanisms are unchanged. The achievability results, Theorems 3.1 and 3.2, transfer with the same constant and the same threshold. The upper bound of Theorem 4.1 and the sign mechanism of Theorem 4.2 transfer with law-dependent constants, so the shape of the phase diagram of Corollary 4.3 is universal for $k \gg \log n$. What does not transfer is the identification of the optimum, the optimality of $\frac{2\sqrt{2}}{3}$ among online algorithms, and the small-k impossibility. Nothing in this section is used elsewhere in the paper.

J.1 SETTING AND NOTATION

Let X be a real random variable with $\mathbb{E}X = 0$, $\mathbb{E}X^2 = 1$, and

$$\mathbb{E}e^{\lambda X} \le e^{\nu^2\lambda^2/2} \qquad \text{for all } \lambda \in \mathbb{R}, \tag{15}$$

for some $\nu \ge 1$. Every such X has $\nu \ge 1$, since expanding both sides of (15) to second order at $\lambda = 0$ gives $\frac{1}{2} \le \frac{\nu^2}{2}$. Throughout this section $G \in \mathbb{R}^{n\times n}$ has i.i.d. entries distributed as X, and $k = \exp(o(\log n))$, $n \ge 4k$ as in the rest of the paper. The Gaussian benchmark

$$\eta_n := 2\sqrt{\log n/k}$$

is now just a number, and we do not identify it with the optimum. For $\theta > 0$ we define

$$\psi^-(\theta) := \mathbb{E}(-\theta - X)_+, \quad \psi^+(\theta) := \mathbb{E}(X - \theta)_+, \quad \mu_\theta := \mathbb{E}\,\mathrm{clip}_\theta(X), \quad \sigma_\theta^2 := \mathrm{Var}\left(\mathrm{clip}_\theta(X)\right),$$

and $p := \Pr[X > 0]$, $m^+ := \mathbb{E}X_+$. Since $\mathrm{clip}_\theta(X) = X - (X-\theta)_+ + (-\theta - X)_+$ and $\mathbb{E}X = 0$, we have $\mu_\theta = \psi^-(\theta) - \psi^+(\theta)$. In particular $\mu_\theta \ne 0$ in general, because clipping an asymmetric law shifts its mean. Also $\mathbb{E}[X\mathbf{1}_{X>0}] = m^+$ and $\mathbb{E}[X\mathbf{1}_{X\le 0}] = -m^+$, again by $\mathbb{E}X = 0$.

Much of the paper uses nothing about the Gaussian law. The privacy lemmas (Lemmas C.4–C.7) hold for every real input. The freshness and decomposition lemmas (Lemmas E.3 and E.4) and the proof of online membership in Appendix E are combinatorial statements about index sets, and freshness uses only the independence of the entries. The moderate-deviation lemma (Lemma C.8) is stated for arbitrary i.i.d. mean-zero bounded summands, and Harris's inequality (Lemma C.11) needs only independent coordinates. What does use the Gaussian law is the identification of the optimum with η_{OPT} (Lemma C.1 and Fact C.2), the constant C_ε^* of the upper bound, and the constant $\sqrt{2/\pi}\tanh(\varepsilon/4)$ of the sign mechanism, together with the Gaussian tail estimates behind Lemmas C.9 and C.13 and the corner term, which the proof of Theorem J.4 replaces. We treat each of these in turn.

J.2 TAIL FACTS

The two lemmas here replace Lemma C.3 and the Gaussian tail of a block sum.

Lemma J.1 (Tails and clipping functions)**.** *For every $t > 0$, $\Pr[X > t] \le e^{-t^2/2\nu^2}$ and $\Pr[X < -t] \le e^{-t^2/2\nu^2}$. Consequently, for every $\theta > 0$,*

$$\psi^\pm(\theta) \le \frac{\nu^2}{\theta}e^{-\theta^2/2\nu^2}, \qquad |\mu_\theta| \le \frac{2\nu^2}{\theta}e^{-\theta^2/2\nu^2}, \tag{16}$$

$$1 - \sigma_\theta^2 \le 2(\theta^2 + 2\nu^2)e^{-\theta^2/2\nu^2} + \frac{4\nu^4}{\theta^2}e^{-\theta^2/\nu^2}, \qquad \sigma_\theta^2 \le 1. \tag{17}$$

The right-hand sides of (16) *and* (17) *are decreasing in θ and tend to 0, so there is $\theta_0 = \theta_0(\nu) \ge 2\nu$ such that $\sigma_\theta^2 \ge \frac{1}{2}$ and $|\mu_\theta| \le \theta/2$ for all $\theta \ge \theta_0$.*

Proof. The tail bound is Chernoff's inequality with $\lambda = t/\nu^2$: $\Pr[X > t] \le e^{-\lambda t}\mathbb{E}e^{\lambda X} \le e^{-t^2/\nu^2 + t^2/2\nu^2} = e^{-t^2/2\nu^2}$. Since $-X$ also satisfies (15), the same holds for the lower tail.

For ψ^+ we integrate the tail. Since $(X-\theta)_+ = \int_\theta^\infty \mathbf{1}_{X>t}\,dt$, Fubini gives

$$\psi^+(\theta) = \int_\theta^\infty \Pr[X>t]\,dt \le \int_\theta^\infty e^{-\frac{t^2}{2\nu^2}}\,dt \le \int_\theta^\infty \frac{t}{\theta}\,e^{-\frac{t^2}{2\nu^2}}\,dt = \frac{\nu^2}{\theta}e^{-\frac{\theta^2}{2\nu^2}},$$

using $t/\theta \ge 1$ on the range of integration. The bound for ψ^- is the same with $-X$ in place of X, and $|\mu_\theta| = |\psi^- - \psi^+| \le \psi^- + \psi^+$.

For (17), write $Y := \mathrm{clip}_\theta(X)$. Since $|Y| \le |X|$ pointwise, $\mathbb{E}Y^2 \le \mathbb{E}X^2 = 1$, and so $\sigma_\theta^2 = \mathbb{E}Y^2 - \mu_\theta^2 \le 1$. For the lower bound, $Y = X$ on $\{|X| \le \theta\}$ and $Y^2 \ge 0$ elsewhere, so $\mathbb{E}Y^2 \ge \mathbb{E}[X^2\mathbf{1}_{|X|\le\theta}] = 1 - \mathbb{E}[X^2\mathbf{1}_{|X|>\theta}]$. By the layer-cake formula,

$$\begin{aligned}\mathbb{E}[X^2\mathbf{1}_{X>\theta}] &= \int_0^\infty \Pr\left[X^2\mathbf{1}_{X>\theta} > s\right] ds = \theta^2\Pr[X>\theta] + \int_{\theta^2}^\infty \Pr[X>\sqrt{s}]\,ds \\ &= \theta^2\Pr[X>\theta] + \int_\theta^\infty 2t\Pr[X>t]\,dt,\end{aligned}$$

and the tail bound gives $\mathbb{E}[X^2\mathbf{1}_{X>\theta}] \le \theta^2 e^{-\theta^2/2\nu^2} + 2\nu^2 e^{-\theta^2/2\nu^2}$, using $\int_\theta^\infty 2te^{-t^2/2\nu^2}dt = 2\nu^2 e^{-\theta^2/2\nu^2}$. The same bound holds for $\mathbb{E}[X^2\mathbf{1}_{X<-\theta}]$, and therefore

$$1 - \sigma_\theta^2 = 1 - \mathbb{E}Y^2 + \mu_\theta^2 \le 2(\theta^2 + 2\nu^2)e^{-\theta^2/2\nu^2} + \mu_\theta^2,$$

and $\mu_\theta^2 \le (4\nu^4/\theta^2)e^{-\theta^2/\nu^2}$ by (16). Finally, $\frac{d}{d\theta}\left[(\theta^2+2\nu^2)e^{-\theta^2/2\nu^2}\right] = -\frac{\theta^3}{\nu^2}e^{-\theta^2/2\nu^2} < 0$, and $\theta^{-1}e^{-\theta^2/2\nu^2}$ and $\theta^{-2}e^{-\theta^2/\nu^2}$ are decreasing, so both right-hand sides are decreasing in θ. They tend to 0, and $|\mu_\theta| \le 2\nu^2/\theta \le \theta/2$ for $\theta \ge 2\nu$, which gives θ_0. □

Lemma J.2 (Cumulant expansion)**.** *There is a constant C_ν depending only on ν such that for all $|\lambda| \le 1/\nu$,*

$$\log \mathbb{E}e^{\lambda X} \le \frac{\lambda^2}{2} + C_\nu|\lambda|^3.$$

Proof. By Taylor's theorem with Lagrange remainder, $e^u = 1 + u + \frac{u^2}{2} + \frac{u^3}{6}e^\xi$ for some ξ between 0 and u, so $e^u \le 1 + u + \frac{u^2}{2} + \frac{|u|^3}{6}e^{|u|}$. Applying this with $u = \lambda X$ and using $\mathbb{E}X = 0$ and $\mathbb{E}X^2 = 1$,

$$\mathbb{E}e^{\lambda X} \le 1 + \frac{\lambda^2}{2} + \frac{|\lambda|^3}{6}\,\mathbb{E}\left[|X|^3 e^{|\lambda X|}\right] \le 1 + \frac{\lambda^2}{2} + \frac{|\lambda|^3}{6}\left(\mathbb{E}X^6\right)^{1/2}\left(\mathbb{E}e^{2|\lambda||X|}\right)^{1/2},$$

by the Cauchy–Schwarz inequality. For the sixth moment, the layer-cake formula and Lemma J.1 give

$$\mathbb{E}X^6 = \int_0^\infty 6t^5\Pr[|X|>t]\,dt \le 12\int_0^\infty t^5 e^{-t^2/2\nu^2}dt = 96\nu^6,$$

the last integral being $8\nu^6$ after the substitution $s = t^2/2\nu^2$. For the exponential moment, $e^{2|\lambda||X|} \le e^{2\lambda X} + e^{-2\lambda X}$, so (15) gives $\mathbb{E}e^{2|\lambda||X|} \le 2e^{2\lambda^2\nu^2} \le 2e^2$ for $|\lambda| \le 1/\nu$. Thus $\mathbb{E}e^{\lambda X} \le 1 + \frac{\lambda^2}{2} + C_\nu|\lambda|^3$ with $C_\nu := \frac{1}{6}\sqrt{192e^2}\,\nu^3$, and $\log(1+z) \le z$ completes the proof. □

J.3 An upper bound on the optimum

The union bound of Lemma C.1 used the Gaussian tail of a block sum. For a general law, the tail of a sum of k^2 entries at the relevant scale is Gaussian to first order in the exponent only when $k \gg \log n$. Below that scale the answer depends on the law: for X a Rademacher variable and $k = 1$, the optimum is 1, not $2\sqrt{\log n}$.

Theorem J.3 (Upper bound on the optimum)**.** *Suppose $k/\log n \to \infty$. Then for every fixed $\zeta > 0$, whp $\max_{|I|=|J|=k} \mathrm{Ave}(G_{I,J}) \le (1+\zeta)\eta_n$.*

Proof. Fix (I, J) with $|I| = |J| = k$ and let $S := \mathrm{Sum}(G_{I,J})$, a sum of $r := k^2$ i.i.d. copies of X. Let $x := (1+\zeta)k^2\eta_n = 2(1+\zeta)k^{3/2}\sqrt{\log n}$ and $\lambda := x/r = 2(1+\zeta)\sqrt{\log n/k}$. Since $k/\log n \to \infty$ we have $\lambda \to 0$, so $\lambda \le 1/\nu$ for all large n. By Chernoff's inequality and Lemma J.2,

$$\begin{aligned}\Pr[S \ge x] &\le \exp\big(-\lambda x + r\log \mathbb{E}e^{\lambda X}\big) \le \exp\Big(-\lambda x + \frac{r\lambda^2}{2} + C_\nu r\lambda^3\Big)\\ &= \exp\Big(-\frac{x^2}{2r}\big(1 - 2C_\nu\lambda\big)\Big),\end{aligned}$$

using $\lambda x = x^2/r$, $r\lambda^2 = x^2/r$ and $r\lambda^3 = x^3/r^2 = \lambda\cdot x^2/r$. Now $\frac{x^2}{2r} = 2(1+\zeta)^2 k\log n$ and $2C_\nu\lambda \to 0$. There are $\binom{n}{k}^2 \le (en/k)^{2k} = \exp(2k\log n(1+o(1)))$ pairs (I,J), so a union bound gives a failure probability of at most

$$\exp\Big(-2k\log n\big[(1+\zeta)^2(1-2C_\nu\lambda) - 1 - o(1)\big]\Big),$$

which tends to zero. □

J.4 Achievability transfers with the same constant

This is the sub-Gaussian form of Proposition 3.3 and Theorem 3.2.

Theorem J.4 ($\mathcal{M}_{\mathrm{BIG}}$ with sub-Gaussian entries)**.** *Let $\zeta \in (0, \frac{1}{2}]$, $1 \le B \le \sqrt{k/3}$, $\kappa = \lfloor k/(B+1)\rfloor$, and $\theta \ge \theta_0(\nu)$ with $2\theta\sqrt{\log n} \le \frac{\zeta}{60}\sigma_\theta\sqrt{\kappa}$. Then $\mathcal{M}_{\mathrm{BIG}}$, with the schedule and budget split of Section 3.2, is ε-DP for every real input, and whp*

$$\begin{aligned}\mathrm{Ave}(G_{I_B,J_B}) \;\ge\; &(1-\zeta)\sigma_\theta\Big(1-\frac{2}{B}\Big)\frac{2\sqrt{2}}{3}\,\eta_n - \frac{32\theta B\big(k\log(2en/k) + B\log n\big)}{\varepsilon k^2}\\ &- \psi^-(\theta)\log n - |\mu_\theta| - \frac{2\nu\sqrt{\log n}}{Bk}.\end{aligned}$$

Proof. Privacy is Lemmas C.5 and C.7, which hold for every real input. For utility we follow the proof of Proposition 3.3 in Appendix E and say what changes.

Centering. Write $\widetilde{G}_{ij} = \mathrm{clip}_\theta(G_{ij}) = \mu_\theta + \widetilde{Z}_{ij}$, where the $\widetilde{Z}_{ij}$ are i.i.d. with mean zero and variance $\sigma_\theta^2 \ge \frac{1}{2}$, and $|\widetilde{Z}_{ij}| \le \theta + |\mu_\theta| \le 2\theta$ since $\theta \ge \theta_0$. At a column step with r rows present, the score of column j is $\tilde{s}_j = r\mu_\theta + \sum_i \widetilde{Z}_{ij}$, and the score of a κ'-subset S is $r\kappa'\mu_\theta + \sum_{j\in S}\sum_i \widetilde{Z}_{ij}$. The exponential-mechanism weights $\exp(\beta\sum_{j\in S}\tilde{s}_j)$ are therefore all multiplied by the same constant $e^{\beta r\kappa'\mu_\theta}$, so the law of the selected set is unchanged if the centered scores $\sum_i \widetilde{Z}_{ij}$ are used instead. We analyze the mechanism with centered scores. Row steps are the same.

The benchmark. By Lemma E.3, which uses only the independence of the entries, conditionally on the history the centered scores of a step are i.i.d. sums of r copies of $\widetilde{Z}$. Lemma C.9 was stated for clipped Gaussians, but its proof uses only Lemma C.8, which applies to any i.i.d. mean-zero summands bounded by 2θ with variance $\sigma_\theta^2 > 0$. The hypothesis of Lemma C.9 becomes $2\theta\sqrt{L} \le \frac{\zeta}{60}\sigma_\theta\sqrt{r}$, which follows from the theorem's hypothesis since $L \le \log n$ and $r \ge \kappa$. Hence, with the same probability as before, the best centered κ'-subset score is at least $(1-\frac{\zeta}{2})(1-o(1))\sigma_\theta\sqrt{r}\,\kappa'\sqrt{2\log n}$, and the factor $(1-\frac{\zeta}{2})(1-o(1))$ is disposed of as in Appendix E.

The mechanism. Lemma C.6 is distribution-free, so the selected centered score is within $\frac{8\theta B}{\varepsilon}(2\kappa\log(2en/k) + 2\log n)$ of the best with probability $1 - n^{-2}$, as before.

The clipping loss and the mean shift. By Lemma E.4 applied to G, together with $G_{ij} \ge \widetilde{G}_{ij} - (-\theta - G_{ij})_+$ on every entry charged at some step,

$$\mathrm{Sum}(G_{I_B,J_B}) \ge \mathrm{Sum}(G_{E_0}) + \sum_{\text{steps}}\big(\text{selected uncentered score}\big) - N^-,$$

where $N^- := \sum(-\theta - G_{ij})_+$ over $(i,j) \in I_B \times J_B \setminus E_0$. The selected uncentered score at a step is the selected centered score plus $r\kappa'\mu_\theta$, and by Lemma E.4 applied to the all-ones matrix the

shifts total $\mu_\theta(k^2-\kappa^2) \ge -|\mu_\theta|k^2$. For N^-, Lemma C.12 applies step by step. Its hypothesis, that a column's inclusion probability is nondecreasing in the column's entries, holds because clip_θ is nondecreasing, and its proof is Harris's inequality for the independent entries of the column. Thus $\mathbb{E}N^- \le \psi^-(\theta)(k^2-\kappa^2)$, with $\mathbb{E}(-\theta-X)_+ = \psi^-(\theta)$ in place of the Gaussian $\psi(\theta)$, and Markov's inequality gives $N^- \le \psi^-(\theta)k^2\log n$ with probability $1-1/\log n$.

The corner. $\mathrm{Sum}(G_{E_0})$ is a sum of κ^2 i.i.d. copies of X, so by (15) it satisfies $\mathbb{E}e^{\lambda\,\mathrm{Sum}(G_{E_0})} \le e^{\nu^2\kappa^2\lambda^2/2}$, and the Chernoff argument of Lemma J.1 gives $|\,\mathrm{Sum}(G_{E_0})| \le 2\nu\kappa\sqrt{\log n}$ with probability $1-2n^{-2}$.

Combining. The union bound and the arithmetic of Appendix E (the Riemann sums, the bound $\kappa^{3/2} \ge (1-o(1))k^{3/2}B^{-3/2}(1-\frac{3}{2B})$, and $B\kappa \le k$) are unchanged, and dividing by k^2 gives the displayed bound, with the terms $\psi^-(\theta)\log n$, $|\mu_\theta|$ and $2\nu\sqrt{\log n}/(Bk)$ in place of their Gaussian counterparts. □

Corollary J.5 (Recovery of the fraction $\frac{2\sqrt{2}}{3}$)**.** *Let $\theta_n := \max\{\theta_0(\nu),\, \nu\sqrt{2\log(k\log n)}\}$, and suppose $\varepsilon\sqrt{k}/(\theta_n\sqrt{\log n}) \to \infty$, $B \to \infty$ with $B \ll \varepsilon\sqrt{k}/(\theta_n\sqrt{\log n})$, $B \le \sqrt{k/3}$, and $\kappa \gg \theta_n^2\log n$. Then for every fixed $\delta > 0$,*

$$\Pr\Big[\,\mathrm{Ave}(G_{I_B,J_B}) \ge \Big(\tfrac{2\sqrt{2}}{3}-\delta\Big)\eta_n\Big] \to 1,$$
$$\Pr\Big[\,\mathrm{Ave}(G_{I_B,J_B}) \ge \Big(\tfrac{2\sqrt{2}}{3}-\delta\Big)\max_{|I|=|J|=k}\mathrm{Ave}(G_{I,J})\Big] \to 1.$$

For fixed ε and ν, the hypotheses hold whenever $k \gg \log n\cdot\log(k\log n)$.

Proof. With $\theta = \theta_n$ we have $e^{-\theta_n^2/2\nu^2} \le 1/(k\log n)$, so by Lemma J.1, $\psi^-(\theta_n)\log n \le \nu^2/(\theta_n k) = o(\eta_n)$, $|\mu_{\theta_n}| \le 2\nu^2/(\theta_n k\log n) = o(\eta_n)$, and $1-\sigma^2_{\theta_n} \le 2(\theta_n^2+2\nu^2)/(k\log n) + 4\nu^4/(\theta_n^2(k\log n)^2) \to 0$ because $\theta_n^2 = O(\log(k\log n)) \ll k\log n$. The penalty term of Theorem J.4 has ratio $\Theta(\theta_n B\sqrt{\log n}/(\varepsilon\sqrt{k}))$ to η_n, which tends to zero by the choice of B. The corner term is $O(\nu\sqrt{\log n}/k) = o(\eta_n)$, and $1-\frac{2}{B} \to 1$. The hypothesis $2\theta_n\sqrt{\log n} \le \frac{\zeta}{60}\sigma_{\theta_n}\sqrt{\kappa}$ holds for large n because $\kappa \gg \theta_n^2\log n$ and $\sigma^2_{\theta_n} \ge \frac{1}{2}$. Since ζ is arbitrary, the first statement follows. For the second, the hypotheses imply $k \gg \log n$, so Theorem J.3 gives $\max \mathrm{Ave} \le (1+\delta')\eta_n$ whp for every fixed $\delta' > 0$. On the intersection of the two events the ratio is at least $(\frac{2\sqrt{2}}{3}-\delta)/(1+\delta')$, and since δ, δ' are arbitrary the claim follows. The fixed-ε statement is as in Theorem 3.2, with θ_n larger by the factor ν. □

The same argument with κ replaced by k gives the one-shot mechanism $\mathrm{Ave}(G_{I_0,J^*}) \ge (\frac{1}{\sqrt{2}} - o(1))\eta_n$ under the same threshold. One thing differs from the Gaussian case. The guarantee is relative to an upper bound on the optimum. We have shown that $\mathrm{Ave}\,/\mathrm{OPT} \ge \frac{2\sqrt{2}}{3} - o(1)$, and that OPT lies in $[(\frac{2\sqrt{2}}{3}-o(1))\eta_n, (1+o(1))\eta_n]$ whp, but we have not identified OPT. The second-moment argument of Bhamidi et al. (2017) is Gaussian, and for a general law the constant of the optimum may lie strictly inside that interval. Likewise, the optimality of $\frac{2\sqrt{2}}{3}$ among online algorithms is a Gaussian statement of Bhamidi et al. (2026), and we do not know the online-optimal constant for a general law.

J.5 The upper bound and the sign mechanism

These are the sub-Gaussian forms of Theorem 4.1 and Theorem 4.2. The law enters only through $p = \Pr[X > 0]$ and $m^+ = \mathbb{E}X_+$.

Theorem J.6 (Upper bound for every ε-DP mechanism)**.** *For every ε-DP mechanism $\mathcal{A}$ outputting (I, J) with $|I| = |J| = k$, and every k,*

$$\mathbb{E}\big[\,\mathrm{Ave}(G_{\mathcal{A}(G)})\big] \le (e^\varepsilon - 1)\,m^+.$$

Proof. As in the proof of Theorem G.2, fix an entry (i,j), let $p(x) := \Pr[(i,j) \in I\times J \mid G_{ij} = x]$ and $\bar p := \mathbb{E}p(X)$. Step 1 of that proof uses only the independence of the entries and gives $\sup p \le e^\varepsilon \inf p$. If $\bar p = 0$ the entry contributes nothing to $\mathbb{E}[\mathrm{Sum}]$, so assume $\bar p > 0$ and let $L := \inf p > 0$.

Choose $a \in [0,1]$ with $L(1 + a(e^\varepsilon - 1)) = \bar{p}$, which is possible since $L \le \bar{p} \le e^\varepsilon L$. If $a \in \{0,1\}$ then $\bar{p} \in \{\inf p, \sup p\}$, so p is constant almost surely and $\mathbb{E}[Xp(X)] = \bar{p}\,\mathbb{E}X = 0$, which is at most the bound. Otherwise let $\tau := \inf\{t : \Pr[X > t] \le a\}$, so that $\Pr[X > \tau] \le a \le \Pr[X \ge \tau]$, and let U be the event $\{X > \tau\}$ together with a portion of the atom $\{X = \tau\}$, selected by an auxiliary uniform random variable independent of everything else, so that $\Pr[U] = a$. Define the randomized function $q^* := e^\varepsilon L$ on U and $q^* := L$ off U, so that $\mathbb{E}q^* = L(1-a) + e^\varepsilon La = \bar{p} = \mathbb{E}p(X)$.

We claim $(X - \tau)(q^* - p(X)) \ge 0$ pointwise. Where $X > \tau$ we have $q^* = e^\varepsilon L \ge \sup p \ge p(X)$. Where $X < \tau$ we have $q^* = L = \inf p \le p(X)$, and where $X = \tau$ the first factor vanishes. Taking expectations, $\mathbb{E}[X(q^* - p(X))] \ge \tau\,\mathbb{E}[q^* - p(X)] = 0$. Therefore, using $\mathbb{E}X = 0$,

$$\mathbb{E}[Xp(X)] \le \mathbb{E}[Xq^*] = L\,\mathbb{E}X + (e^\varepsilon - 1)L\,\mathbb{E}[X\mathbf{1}_U] = (e^\varepsilon - 1)L\,\mathbb{E}[X\mathbf{1}_U].$$

It remains to see that $0 \le \mathbb{E}[X\mathbf{1}_U] \le m^+$, after which $L \le \bar{p}$ gives the bound $(e^\varepsilon - 1)\bar{p}\,m^+$. If $\tau > 0$ then $U \subseteq \{X \ge \tau\} \subseteq \{X > 0\}$, so $0 \le X\mathbf{1}_U \le X\mathbf{1}_{X>0}$ pointwise. If $\tau \le 0$ then $U \supseteq \{X > \tau\} \supseteq \{X > 0\}$, so $X\mathbf{1}_U = X\mathbf{1}_{X>0} + X\mathbf{1}_{U\cap\{X\le 0\}} \le X\mathbf{1}_{X>0}$, and $\mathbb{E}[X\mathbf{1}_U] \ge \mathbb{E}[X] = 0$ since U contains every point where $X > 0$. In both cases $0 \le \mathbb{E}[X\mathbf{1}_U] \le \mathbb{E}[X\mathbf{1}_{X>0}] = m^+$.

Summing over the n^2 entries, with $\sum_{i,j} \bar{p}_{ij} = \mathbb{E}|I \times J| = k^2$, gives $\mathbb{E}[\mathrm{Sum}(G_{\mathcal{A}(G)})] \le (e^\varepsilon - 1)m^+k^2$, and dividing by k^2 completes the proof. □

For the Gaussian, $m^+ = 1/\sqrt{2\pi}$ and the bound is the upper bound on C^*_ε in Lemma I.5(b). The constant C^*_ε of the main text is sharper because it exploits the shape of the Gaussian density, which a general law does not permit.

Theorem J.7 (Sign mechanism with sub-Gaussian entries)**.** *Let J^* be the output of the sign mechanism of Section 4.2, run on $Y_{ij} = \mathbf{1}[G_{ij} > 0]$ with parameter $\varepsilon/2$. If $\varepsilon^2 k \le (4-\delta)\log n$ for a fixed $\delta > 0$, then*

$$\mathbb{E}\big[\,\mathrm{Ave}(G_{I_0,J^*})\big] \;\ge\; (1 - o(1))\,m^+ \cdot \frac{e^{\varepsilon/2} - 1}{1 + p\,(e^{\varepsilon/2} - 1)}.$$

Proof. Lemma D.2 was proved for an arbitrary law with $\mathbb{E}X = 0$ and $\mathbb{E}|X| < \infty$, so sub-Gaussianity is not needed here. It gives, for $n \ge 10k$,

$$\mathbb{E}\big[\,\mathrm{Ave}(G_{I_0,J^*})\big] \ge m^+ \frac{e^{\varepsilon/2} - 1}{1 + p(e^{\varepsilon/2} - 1)}\,(1 - q)_+\Big(1 - \big(1 + p(e^{\varepsilon/2} - 1)\big)r\Big),$$

with q and r as in the lemma. The proof of Theorem 4.2 shows that $q \to 0$ and $e^{\varepsilon/2}r \to 0$ under $\varepsilon^2 k \le (4-\delta)\log n$, using only $k = \exp(o(\log n))$. Since $1 + p(e^{\varepsilon/2} - 1) \le e^{\varepsilon/2}$, the last factor is $1 - o(1)$ as well. □

For the Gaussian, $p = \frac{1}{2}$ and $m^+ = 1/\sqrt{2\pi}$, and the constant is $\sqrt{2/\pi}\tanh(\varepsilon/4)$, recovering Theorem 4.2.

Corollary J.8 (The small-block regime, for every law)**.** *Fix $\varepsilon > 0$ and $\delta > 0$, and let $k \le (4 - \delta)\log n/\varepsilon^2$. Then*

$$(1 - o(1))\,m^+ \frac{e^{\varepsilon/2} - 1}{1 + p(e^{\varepsilon/2} - 1)} \;\le\; \sup_{\mathcal{A}\ \varepsilon\text{-DP}} \mathbb{E}\big[\,\mathrm{Ave}(G_{\mathcal{A}(G)})\big] \;\le\; (e^\varepsilon - 1)\,m^+,$$

and the ratio of the two constants is $(e^{\varepsilon/2} + 1)(1 + p(e^{\varepsilon/2} - 1)) = 2 + O(\varepsilon)$ as $\varepsilon \to 0$, for every law.

Proof. The two inequalities are Theorems J.7 and J.6. The ratio is $\frac{(e^\varepsilon - 1)(1 + p(e^{\varepsilon/2} - 1))}{e^{\varepsilon/2} - 1} = (e^{\varepsilon/2} + 1)(1 + p(e^{\varepsilon/2} - 1))$, which tends to 2 as $\varepsilon \to 0$ for every $p \in [0,1]$. □

J.6 What the phase diagram becomes

Combining Corollaries J.5 and J.8 with Theorem J.3 gives the analogue of Corollary 4.3. For fixed $\varepsilon \le 1$ and fixed ν, and $k \gg \log n$, writing $V^* := \sup_{\mathcal{A}\ \varepsilon\text{-DP}} \mathbb{E}[\mathrm{Ave}(G_{\mathcal{A}(G)})]$,

$$0.15\, m^+ \min\Big\{\varepsilon\sqrt{\frac{k}{\log n}},\, 1\Big\} - o(1) \;\le\; \frac{V^*}{\eta_n} \;\le\; \min\Big\{0.86\, m^+\varepsilon\sqrt{\frac{k}{\log n}},\; 1 + o(1)\Big\},$$

$$\frac{V^*}{\eta_n} \;\ge\; \tfrac{2\sqrt{2}}{3} - o(1) \quad \text{once } k \gg \frac{\log n \cdot \log(k \log n)}{\varepsilon^2}.$$

The left inequality is Theorem J.7 with $e^{\varepsilon/2} - 1 \ge \varepsilon/2$ and $1 + p(e^{\varepsilon/2} - 1) \le e^{1/2}$ for $\varepsilon \le 1$, applied at budget ε when $\varepsilon^2 k \le \log n$ and, as in the proof of Corollary 4.3 in Appendix G, at the smaller budget $\varepsilon' = \sqrt{\log n / k} \le \varepsilon$ otherwise, using that the supremum is nondecreasing in the budget. The first upper bound is Theorem J.6 with $e^{\varepsilon} - 1 \le 1.72\varepsilon$ on $(0, 1]$. The cap $1 + o(1)$ is Theorem J.3 together with the trivial bound $\mathrm{Ave} \le \max_{i,j} |G_{ij}|$ on its failure event, whose contribution is $O(\nu\sqrt{\log n}) \cdot n^{-\Omega(1)} = o(\eta_n)$, and the right statement is Corollary J.5. For $k \gg \log n \cdot \log(k \log n)$, where Corollary J.5 shows that whp $\mathrm{OPT} \in [(\frac{2\sqrt{2}}{3} - o(1))\eta_n, (1 + o(1))\eta_n]$, the same holds with η_n replaced by OPT up to the factor $\frac{2\sqrt{2}}{3}$ in the constants. In the band $\log n \ll k \lesssim \log n \cdot \log(k \log n)$ we know only $\mathrm{OPT} \le (1 + o(1))\eta_n$, and the statements are relative to η_n. So the shape $\Theta(\min\{1, \varepsilon\sqrt{k/\log n}\})$ and the transition at $k \asymp \log n/\varepsilon^2$ are universal for $k \gg \log n$, with the law entering only through $m^+ = \mathbb{E}X_+$ and ν.

Below $k \asymp \log n$ the picture depends on the law, and the analogue for general laws of the Gaussian statement in Corollary 4.3, that no mechanism recovers a constant fraction of the optimum for $k = o(\log n)$, fails. The absolute bound of Theorem J.6 still holds, but the optimum need not grow. For Rademacher entries and $k = 1$ the optimum is 1, and the sign mechanism recovers the fraction $\tanh(\varepsilon/4) \approx \varepsilon/4$ of it at fixed ε by Theorem J.7, since there $p = \frac{1}{2}$ and $m^+ = \frac{1}{2}$. What drives the Gaussian impossibility is that $\eta_{\mathrm{OPT}} \to \infty$ for small k, a property of unbounded tails. For bounded or light-tailed laws the optimum at $k = o(\log n)$ is governed by large deviations at the law's own scale, and one would need the law's rate function to say more. We have not examined the row-level results of Appendix I here. A row sum $\langle X_i, \mathbf{1}_J \rangle$ is a sum of k entries, so the same tools apply, but we have not carried the analysis out.