NODE-PRIVATE LEARNING FROM MULTI-ATTRIBUTION DATA VIA STABLE CONTRIBUTION BOUNDING

Anonymous authors
Paper under double-blind review

ABSTRACT

We study node-level differential privacy for multi-attribution learning, where each example involves several users. A standard pipeline caps each user's contributions and then runs an example-level private learner on the selected examples. Removing one user can change which other examples are selected, and this recourse, not the cap, determines the privacy of the pipeline. For deterministic structure-only selectors, bounded recourse is necessary and sufficient to transfer example-level privacy to node-level privacy, and exact contribution bounding can have recourse linear in the dataset size even at degree two. We give a stable selector, based on quadratic regularization and feasible rounding, that retains a constant fraction of the feasible optimum with a privacy factor growing at most as the square root of the dataset size, and structural classes on which the factor is independent of the size. Population-risk guarantees and experiments on public attribution structures show when the reduced privacy noise outweighs the loss of training examples.

1 INTRODUCTION

Modern machine learning systems increasingly train on data where individual examples contain information about multiple users. In email networks, each thread involves senders and recipients; in collaborative authorship graphs, each paper has multiple authors; in social media, each post tags multiple participants. This *multi-attribution* structure creates a fundamental challenge for differential privacy (Dwork et al., 2006; Dwork & Roth, 2014) as protecting a single user requires controlling the influence of *all* examples they appear in, not just examples they directly contributed.

Recent work by Ganesh et al. (2025) formalized this challenge through *fixed-graph differential privacy*. In their framework, the data consists of a set of users V and a set of training examples E, where each example may involve multiple users. This multi-attribution structure forms a hypergraph where users are vertices and examples are hyperedges. The key modeling assumption is that this attribution structure, i.e., which users appear in which examples, is public knowledge, but the example contents (the actual data values) are private.

We study node-level privacy, which also protects whether a user participates. A central challenge arises when one user is added to or removed from the network. Not only does the set of users change, but all examples involving that user must also be added or removed to maintain privacy. This affects up to Δ examples per user, where Δ denotes the maximum number of examples any user appears in (the maximum degree in the hypergraph).

Ganesh et al. (2025)'s key algorithmic contribution is *contribution bounding*, i.e., selecting a subset of examples such that each user appears in at most k examples, then applying standard differentially private learning algorithms to this subset. Under node adjacency, limiting a user's selected contributions does not control how many other examples the selector replaces when that user is removed. We ask when this two-phase approach gives node privacy and how to design contribution bounding to retain useful data while controlling these changes. We call the number of selected examples that change the selector's *recourse*.

Our contributions. We characterize the role of selection stability in node-private learning and develop a stable contribution-bounding method. The selector uses only the attribution structure; the learner receives the selected contents, and only the learner's output is released.

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.

1. Instability of exact contribution bounding. Removing one user can change an optimal selection throughout the dataset, even when every user appears in only two examples. We prove this for the fractional optimum and for every exact integral selector, including randomized selectors. The corresponding privacy cost can grow linearly with dataset size (Section 3).

2. Stable selection with retention guarantees. Quadratic regularization and feasible randomized rounding give a selector that retains a constant fraction of the feasible optimum in expectation for bounded attribution size and constant regularization. At a fixed contribution cap, its privacy factor grows at most as the square root of dataset size, up to logarithmic terms. We further obtain factors independent of dataset size on two structural classes with arbitrarily large user degrees (Section 4).

3. A characterization of the privacy factor. For deterministic selectors, the worst-case recourse is exactly the factor needed to transfer pure example-level privacy to pure node-level privacy uniformly over learners and contents. For randomized selectors, we prove sufficiency from high-probability recourse and a converse under an output-separation condition. We also give a learning task on which every learner meeting a concrete accuracy requirement incurs privacy loss proportional to the selection changes (Section 5).

4. Learning guarantees and experiments. We give population-risk bounds that account for both the retained sample size and the privacy factor. Paired experiments on synthetic and public attribution structures, including datasets with millions of examples, examine when reduced privacy noise offsets the loss of training data (Sections 5.4 and 6).

2 PRELIMINARIES

We first describe the data and what it means to protect a user, then formulate the contribution-bounding problem. The distinction between changing a user's data and removing that user is important for the two-phase pipeline.

Notation. We use $[n]$ to denote the set $\{1, 2, \ldots, n\}$. For a vector $v \in \mathbb{R}^d$, we write $\|v\|_2$ for its Euclidean norm. We also use the ℓ_1 norm to measure the total change in fractional weights.

A hypergraph $H = (V, E)$ has vertices V (representing users) and hyperedges E (representing examples), where each $e \in E$ is a subset of V. We treat H as an indexed multihypergraph: each example has an identity, and different examples may involve the same users. Let $n = |E|$, $E_v = \{e \in E : v \in e\}$, $\deg_H(v) = |E_v|$, and $\Delta(H) = \max_v \deg_H(v)$. A public bound r limits how many users are associated with one example, so $|e| \le r$.

An *attributed dataset* is a tuple $\mathcal{D} = (V, E, \{x_e\}_{e \in E})$ where (V, E) is a hypergraph and $x_e \in \mathcal{X}$ is the feature vector for example e. The hypergraph structure (V, E) encodes which users appear in which examples. More generally, x_e may be any example content supplied to the learner, such as a feature-label pair. Example identities, attributions, and contents are preserved on shared examples when comparing node neighbors.

We define two distinct adjacency relations on attributed datasets, for different threat models.

Definition 1 (Fixed-hypergraph adjacency)**.** *Datasets $\mathcal{D} = (V, E, \{x_e\})$ and $\mathcal{D}' = (V, E, \{x'_e\})$ are fixed-hypergraph adjacent if there is a user $v^* \in V$ such that $x_e = x'_e$ for every $e \notin E_{v^*}$. The structure is identical on the two inputs; contents may change on examples involving v^*.*

Definition 2 (Node adjacency)**.** *Datasets $\mathcal{D} = (V, E, \{x_e\})$ and $\mathcal{D}' = (V', E', \{x'_e\})$ are* node-adjacent *if they differ by adding or removing one user and all their incident examples. Surviving examples retain their identities, attributions, and contents.*

We distinguish three privacy guarantees by the change they protect against. Example-level privacy protects the addition or removal of one example. Fixed-graph privacy protects changes to a user's example contents while holding the attribution structure fixed. Node privacy protects the addition or removal of the user together with all incident examples. We use add/remove adjacency for the downstream learner, so replacing one example counts as two changes.

Definition 3 (Example-level differential privacy)**.** *Let $\mathcal{X}^*$ denote the finite multisets of example contents, including the empty multiset. A randomized mechanism $\mathcal{A} : \mathcal{X}^* \to \mathcal{Y}$ is (ε, δ)-example-DP if, for every pair Z, Z' differing by one insertion or deletion and every measurable $F \subseteq \mathcal{Y}$,*

$$\Pr[\mathcal{A}(Z) \in F] \le e^{\varepsilon} \Pr[\mathcal{A}(Z') \in F] + \delta. \tag{1}$$

For two output distributions, the notation $\mathcal{A}(Z) \approx_{\varepsilon,\delta} \mathcal{A}(Z')$ means that (1) holds in both directions for every measurable event. Pure differential privacy means $\delta = 0$.

Definition 4 (Fixed-graph differential privacy (Ganesh et al., 2025))**.** *A randomized mechanism $\mathcal{B}$ on a fixed domain of attributed datasets is (ε, δ)-fixed-graph DP if $\mathcal{B}(\mathcal{D}) \approx_{\varepsilon,\delta} \mathcal{B}(\mathcal{D}')$ for every fixed-hypergraph adjacent pair in that domain.*

Definition 5 (Node differential privacy)**.** *A randomized mechanism $\mathcal{B}$ on a fixed domain of attributed datasets is (ε, δ)-node-DP if $\mathcal{B}(\mathcal{D}) \approx_{\varepsilon,\delta} \mathcal{B}(\mathcal{D}')$ for every node-adjacent pair in that domain.*

Contribution bounding. Given a hypergraph (V, E) and a budget $k \in \mathbb{N}$, the *contribution bounding problem* is to select a sub-hypergraph (V, S) with $S \subseteq E$ that maximizes $|S|$ subject to the *contribution constraint* that $|\{e \in S : v \in e\}| \leq k$, for all users $v \in V$. We use an integer cap $k \geq 1$. With an indicator y_e for whether example e is selected, the problem is

$$\max_{y \in \{0,1\}^E} \sum_{e \in E} y_e \quad \text{subject to} \quad \sum_{e \ni v} y_e \leq k \quad (v \in V). \tag{2}$$

We write $\mathrm{OPT}_{\mathrm{IP}}$ for this optimum and $\mathrm{OPT}_{\mathrm{LP}}$ for the optimum after relaxing $y_e \in [0, 1]$. The indicator $y_e = 1$ means that example e is retained; a fractional value in the relaxation assigns it a weight between zero and one. We want to retain many examples, but also to keep the selection from changing too much when one user is removed.

The two-phase pipeline. Ganesh et al. (2025) employ a two-phase private learning approach, wherein a contribution-bounded dataset $(V, S, \{x_e\}_{e \in S})$ is first obtained, followed by a standard DP-SGD or DP-MF algorithm. Group privacy guarantees can then be applied to translate to fixed-graph DP (Ganesh et al., 2025). In our pipeline, a possibly randomized selector $\mathcal{M}$ uses the structure H to choose a feasible set $S(H)$ without examining the contents. The learner $\mathcal{A}$ then receives $\mathsf{Out}(\mathcal{D}) = \{x_e : e \in S(H)\}$ and uses randomness independent of the selector. Only the learner's output is released; the selected set and attribution structure remain internal. We write $\mathcal{A} \circ \mathcal{M}$ for this composition. We require one example-DP learner defined on all finite multisets, including the empty multiset, because the selected size depends on the private input.

For fixed-hypergraph neighbors, we can use the same selection on both sides. At most k selected contents then change, corresponding to at most $2k$ insertions and deletions under our learner convention. Under node adjacency, removing a user can also change the selection of surviving examples. We call $|S(H) \triangle S(H')|$ the selection *recourse*. With shared contents on surviving examples, $d_{\mathrm{Ham}}(\mathsf{Out}(\mathcal{D}), \mathsf{Out}(\mathcal{D}')) \leq |S(H) \triangle S(H')|$, where d_{Ham} is the minimum number of multiset insertions and deletions. This is the quantity we control to obtain node privacy.

Public parameters and the input domain. The privacy guarantee compares neighboring inputs under the same mechanism. We therefore fix the parameters used to choose the learner's noise independently of the private input. When using a public size bound N, degree bound $\bar{\Delta}$, or structural bound $\bar{\rho}$, we require that bound to hold throughout the input class. The realized values n and $\Delta(H)$ may still differ between neighbors.

Formally, let $\mathcal{H}$ be a fixed class of attribution structures, and let $\mathbb{D}_{\mathcal{H}}$ contain every assignment of contents from $\mathcal{X}$ to a structure in that class. Privacy is required for every neighboring pair in $\mathbb{D}_{\mathcal{H}}$. We write $\mathbb{D}$ for the full space of attributed datasets.

Direct learning and utility. With a public degree bound $\bar{\Delta}$, we can instead apply the learner to all examples. Node neighbors then differ by at most $\bar{\Delta}$ examples, and group privacy gives parameters $\varepsilon = \bar{\Delta}\varepsilon_0$ and $\delta = \delta_0 \sum_{i=0}^{\bar{\Delta}-1} e^{i\varepsilon_0}$. We use this direct approach as the raw-data baseline. Contribution bounding can reduce the required privacy noise, while also reducing the training set. For learning guarantees, we use a loss $\ell(\theta; x) \in [0, 1]$, parameters $\theta \in \Theta \subseteq \mathbb{R}^d$, and population risk $L(\theta) = \mathbb{E}_{x \sim P} \ell(\theta; x)$ with minimizer θ^*. Assumptions on how contents are sampled are stated in the utility results; the privacy guarantees apply to every content assignment in the input domain.

Related work. Contribution bounding for multi-attribution learning was studied by Ganesh et al. (2025), who introduced fixed-graph privacy and analyzed greedy selection, and by Cohen-Addad et al. (2025), who developed a distributed selection algorithm. In the fixed-graph model, a user's contents can change while the attribution structure remains fixed. We protect the addition or removal of the user and all associated examples, which can also change the selection of other users' examples.

The N2E framework of Hu et al. (2025) is closely related in its use of stable graph preprocessing to transfer an edge-level privacy guarantee to node privacy. Our analysis concerns contribution bounding on hypergraphs, the amount of data retained, and the privacy factor required for composition with an arbitrary example-DP learner. Algorithmically, the use of regularization and rounding connects to stable graph optimization under weight perturbations (Kumabe & Yoshida, 2023; Liu et al., 2024); here the perturbation is a user deletion. These connections and broader survey is in Appendix A.

3 Instability of Unregularized Contribution Bounding

The predominant approach to multi-attribution privacy is a two-phase pipeline, wherein, we first apply contribution bounding to limit how many examples each user appears in, and then run standard differentially private learning on the selected subset (Ganesh et al., 2025). Under fixed-graph adjacency, the attribution structure and hence the selected identities can stay the same when a user's contents change. Under node adjacency, removing that user also changes the structure used to make the selection. Other examples can then be selected in place of those previously retained.

Why stability is crucial for node-DP. Before presenting our technical results, we build intuition for why stability matters. Under node-level differential privacy, removing one user should cause only small changes to the algorithm's output distribution. If many selected examples change, the learner receives substantially different training data. An example-private learner controls the effect of one example change; group privacy accounts for all the changes made by selection. A bound of t changes therefore gives a factor t in the pure-privacy parameter. We will show that this is the exact factor needed for guarantees uniform over learners when selection is deterministic, and give corresponding lower bounds for exact randomized selectors (Section 5).

Theorem 1 (Instability under node adjacency, uniform cap)**.** *Let $H = C_{2m+1}$ be the odd cycle on $2m+1$ vertices with uniform cap $k = 1$, so that $\Delta = r = 2$ and $|V| = |E| = 2m+1$. The LP relaxation of contribution bounding has a unique optimal solution on H, namely $y_e = 1/2$ on every edge, and for every vertex v a unique optimal solution on $H \setminus \{v\} = P_{2m}$, namely the indicator of the path's unique perfect matching. The zero-extended distance between them is*

$$\left\| y^*(H) - \widetilde{y}^*(H \setminus \{v\}) \right\|_1 \;=\; m + \tfrac{1}{2} \;=\; |V|/2.$$

Corollary 2 (Every optimal selector is unstable)**.** *On $H = C_{2m+1}$ with cap $k = 1$, every rule that returns a maximum matching, deterministic or randomized, has a vertex deletion for which the expected selection recourse is at least m.*

The fractional optimum is unique before and after deletion, so its instability does not arise from choosing between tied optima. For integral selection, the conclusion applies to every rule that returns a maximum matching, including randomized rules. Even a user with only two examples can therefore cause a global change in the data passed to the learner.

Appendix B.1 gives the proofs and additional constructions. These include linear recourse for a fixed-order greedy selector and instability caused by tie-breaking on an even cycle. This motivates our next question: can we design contribution bounding algorithms that are stable under node adjacency? We next show how regularization and feasible rounding provide such a method.

4 Regularization Restores Stability

We now design a selector whose output changes less when a user is removed. We first assign fractional weights to the examples, using regularization to control how those weights respond to deletion. We then turn the weights into a feasible subset while controlling the changes introduced by rounding. This separates the stability of the optimization problem from the stability of selecting training examples.

Regularized formulation. Consider the regularized optimization problem:

$$\max_{y \in [0,1]^E} \left(\sum_{e \in E} y_e - \frac{\lambda}{2} \|y\|_2^2 \right) \quad \text{subject to} \quad \sum_{e \ni v} y_e \le k \quad \text{for every } v \in V. \tag{3}$$

Here $\lambda > 0$ controls the strength of the penalty, and $y^\lambda(H)$ denotes the unique optimizer. The objective is λ-strongly concave. We use this curvature to turn a bound on the objective loss from deleting a user into a bound on the movement of the optimizer.

Sample size vs. regularization tradeoff. Without binding constraints, the first-order optimality condition gives $y_e \approx 1/\lambda$ (capped at 1 by the box constraint). Thus, with strong regularization $\lambda \gg 1$, each example receives weight $\approx 1/\lambda$, and the total selected mass is $\sum_e y_e \approx n/\lambda$ rather than $\approx n$. Our guarantees allow a constant choice of λ. In particular, for $0 < \lambda \leq 1$, the fractional mass is at least half of the unregularized LP optimum. The retention analysis below also accounts for the examples lost in rounding.

4.1 Stability of the fractional solution under node deletion

We now analyze how the optimal solution changes when a node is removed. The key challenge is that node-adjacent hypergraphs $H = (V, E)$ and $H' = (V', E')$ have different edge sets, so the optimization variables live in different spaces ($\mathbb{R}^E$ vs $\mathbb{R}^{E'}$). We handle this by extending both solutions to a common domain. Specifically, we assign weight zero to the deleted examples when comparing the two optimizers.

Theorem 3 (Fractional stability under node deletion)**.** *Let $H = (V, E)$ be a hypergraph with per-vertex budget $k \geq 1$, and let $H' = (V', E')$ be obtained by deleting a vertex $v^* \in V$ and all incident hyperedges, with $n = |E|$. For the optimizers of* (3) *with the same $\lambda > 0$, let $\widetilde{y}^\lambda(H') \in \mathbb{R}^E$ denote the zero-extension of $y^\lambda(H')$. Then*

$$\left\|y^\lambda(H) - \widetilde{y}^\lambda(H')\right\|_2^2 \leq \frac{2k}{\lambda}, \qquad \left\|y^\lambda(H) - \widetilde{y}^\lambda(H')\right\|_1 \leq \sqrt{\frac{2nk}{\lambda}}.$$

The bound controls all weight changes, including both the removed examples and the surviving ones. Its dependence on the deleted user's contributions is through the cap k, even if that user appears in many examples. To see why, zero-extend the optimizer on the smaller dataset. This gives a feasible point on the original dataset, and the objective gap is at most k. Strong concavity then bounds the distance between the two optimizers. See Appendix B.2 for the full proof and extension to user-specific budgets.

Approximate optimization. A feasible solver with a uniform public objective-gap tolerance has a corresponding stability guarantee. In particular, a gap at most $\lambda/(8N)$ on every admissible input adds at most one to the fractional recourse bound. Remark 6 gives the general correction and its effect on retention.

4.2 Feasible rounding with bounded recourse

The previous section bounds the change in the *fractional* optimizer y^λ under node deletion. To use this in the node-DP pipeline, we must convert fractional weights into an *integer* subset $S \subseteq E$ while *preserving feasibility* (each user appears in at most k retained examples). Sampling examples independently from these weights can exceed a user's cap. We therefore keep only the first k sampled examples at each user according to shared random priorities. An example is retained only if it is kept at every user to whom it is attributed.

Algorithm 1 Priority-truncated sampling

Require: Hypergraph $H = (V, E)$, integer cap $k \geq 1$, weights $w \in [0, 1]^E$

1: Draw mutually independent $U_e, P_e \sim \mathrm{Unif}[0, 1]$ for every $e \in E$
2: Set $R \leftarrow \{e \in E : U_e < w_e\}$
3: For each user v, form $R(v) = \{e \in R : v \in e\}$ and let $K(v)$ contain the first $\min\{k, |R(v)|\}$ examples in increasing order of P_e, with ties broken by a fixed order of example identities
4: Return $S \leftarrow \{e \in R : e \in K(v) \text{ for every } v \in e\}$

We couple the two runs so that the only source of randomness is a shared pair (U_e, P_e) per edge. Under this coupling, the tentative sampled sets differ *only* on edges e for which the shared uniform U_e falls between the two inclusion thresholds w_e and $\widetilde{w}'_e$. The number of such edges is a sum of independent Bernoulli variables with mean exactly $\mu = \|w - \widetilde{w}'\|_1$, so it concentrates around μ with a Bernstein bound.

Inserting or deleting one tentative example can change the final selection only through that example and at most one displaced example at each of its users. The amplification is at most $1 + r$, indepen-

dently of the cap and user degrees (Lemma 13). Each user's priority list is computed once from the tentative set, so a rejection at one user does not trigger another round of selection at other users.

Theorem 4 (Stability of feasible rounding)**.** *Let H' be obtained from $H = (V, E)$ by one user deletion. Choose weights $w \in [0,1]^E$ and $w' \in [0,1]^{E'}$, and let $\widetilde{w}'$ be the zero-extension of w'. Couple the two runs of Algorithm 1 using the same sampling uniforms and priorities on shared examples. Both outputs are feasible. With $\mu = \|w - \widetilde{w}'\|_1$, for every $\eta \in (0,1)$, with probability at least $1 - \eta$,*

$$|S(H) \triangle S(H')| \le (1+r)\left(\mu + \sqrt{2\mu \log(1/\eta)} + \frac{2}{3}\log(1/\eta)\right). \tag{4}$$

For use below, let $\tau(a, \eta)$ be the ceiling of the right-hand side of (4) with $\mu = a$. Any uniform bound a on the sampling disagreement mean gives recourse at most $\tau(a, \eta)$ with probability at least $1 - \eta$.

The concentration and truncation arguments are proved in Appendix B.3. For retention, we use half of the regularized weights in the sampling step. This leaves capacity at each user while retaining a constant fraction of the fractional mass when the cap is at least logarithmic in the attribution size.

Half-weight rounding retains at least a $1/(4r)$ fraction of the fractional mass in expectation for every cap, and at least a quarter when $k \ge 6\ln(2r)$ (Lemma 15). For $0 < \lambda \le 1$, regularization retains at least half of the unregularized LP optimum. Thus for $0 < \lambda \le 1$ and $k \ge 6\ln(2r)$, the expected selection size is at least $\mathrm{OPT}_{\mathrm{LP}}/8$. Under the coupling in Theorem 4, half-weight sampling also halves the disagreement mean: for the regularized optimizers, $\mu \le \frac{1}{2}\sqrt{2nk/\lambda}$. We convert this stability bound into node privacy using the composition theorem in Section 5.

The retention proof is in Appendix B.4. The ℓ_2 regularizer in (3) yields a unique optimizer and enables a clean stability analysis, but it can shrink the total fractional mass $\sum_e y_e$ when many contribution constraints are slack. This shrinkage is not a privacy requirement in itself. It is an artifact of pulling all coordinates toward zero, which reduces the amount of data passed to the downstream learner. Alternative regularizers can preserve more mass, but their response to node deletion must be analyzed separately. We give these analyses in Appendix B.6 and discuss the choice of λ in Remark 7.

4.3 Stability under structural assumptions

The general bound allows changes to propagate through the dataset. On some attribution structures, however, the effect of a deletion remains controlled even when the dataset and user degrees are large. To understand when this happens, consider an example shared by the deleted user and a surviving user. Removing that example frees some of the surviving user's contribution capacity. The optimizer can use this capacity to increase other weights, which may in turn affect further users.

We analyze this effect on the surviving structure by varying the users' capacities. The total released fractional mass is at most $(r-1)k$: the deleted user has at most k units of mass, and each removed example involves at most $r-1$ other users. The remaining question is how much the optimizer moves per unit of released capacity. We measure this through a *response norm* and use a uniform bound on it to obtain stability independent of dataset size.

For a fixed structure $G = (V, E)$ and caps $c \in \mathbb{R}^V_{\ge 0}$, let $y_G(c)$ denote the unique optimizer of (3) with constraints $\sum_{e \ni u} y_e \le c_u$; we keep $\lambda > 0$ fixed. For a deletion of v^*, let b_u be the fractional mass of removed examples also attributed to surviving user u. We have $0 \le b \le k\mathbf{1}$, $\|b\|_1 \le (r-1)k$, and the original optimizer restricted to surviving examples equals $y_{H'}(k\mathbf{1} - b)$ (Lemma 16).

Thus, on the surviving structure, deleting v^* increases the available caps from $k\mathbf{1} - b$ to $k\mathbf{1}$. To describe the response to this change, write $y = y_G(c)$ and let $U = \{e : 0 < y_e < 1\}$ and $T = \{u : \sum_{e \ni u} y_e = c_u\}$. Here U contains the fractional examples and T contains all saturated users. Let $B_{U,T} \in \{0,1\}^{U \times T}$ be their incidence matrix. We define

$$\rho(G, c) := \left\|(B_{U,T}^\top)^+\right\|_{1\to 1} = \max_{u \in T} \left\|(B_{U,T}^\top)^+ e_u\right\|_1, \tag{5}$$

where $^+$ denotes the Moore–Penrose pseudoinverse and e_u is the u^{th} unit vector. We set $\rho = 0$ if U or T is empty. Fractional examples are the coordinates whose weights can move within the interval $(0,1)$, and the saturated users are those whose capacity constraints are tight. Their incidence matrix describes which moving weights share a tight constraint. The pseudoinverse gives the linear response of these weights to a change in the tight capacities. Its $1 \to 1$ norm measures the largest total weight change

per unit capacity change while these sets remain fixed. For example, if each fractional example contains at most one saturated user, the nonzero columns of $B_{U,T}$ have disjoint supports and $\rho \le 1$.

Theorem 5 (Stability along a cap path)**.** *Let $0 \le b \le k\mathbf{1}$ and $y(t) = y_G(k\mathbf{1} - tb)$ for $t \in [0,1]$. The solution path is continuous and piecewise affine, and*

$$\|y(1) - y(0)\|_1 \le \|b\|_1 \sup_{t\in[0,1]} \rho(G, k\mathbf{1} - tb).$$

The proof identifies the linear response on each piece of the path and integrates it across changes in the saturated constraints (Appendix B.5). Together with Lemma 16, this bounds the movement among surviving examples by $(r-1)k$ times the largest response norm along the path. Adding the mass of the deleted examples gives the following calibration for the rounded selector.

Corollary 6 (Structural calibration)**.** *Fix $k \ge 1$, $\lambda > 0$, and an input class $\mathcal{H}$ whose examples have size at most r. Suppose a public constant $\bar\rho \ge 1$ bounds ρ along the cap path of Lemma 16 for every node-adjacent pair in $\mathcal{H}$. Then every such pair satisfies $\|y^\lambda(H) - \widetilde{y}^\lambda(H')\|_1 \le \mu_\rho := k + (r-1)k\bar\rho$. For a target budget $\varepsilon > 0$, $\delta \in (0,1)$, set $\eta = \delta/2$ and $t_\rho = \tau(\mu_\rho/2, \eta)$. Run Algorithm 1 with weights $y^\lambda/2$, followed by an $(\varepsilon_0, \delta_0)$-example-DP learner, where $\varepsilon_0 = \varepsilon/t_\rho$ and $\delta_0 \le \delta e^{-\varepsilon}/(2t_\rho)$. The pipeline is (ε, δ)-node-DP on $\mathcal{H}$, with $t_\rho = O(r^2 k\bar\rho + r\log(1/\eta))$.*

The class and the constants in this statement are fixed across neighboring inputs; the guarantee covers every adjacent pair in that class. For bounded $\bar\rho$, the calibration depends on r and k but has no dependence on N or maximum degree. The next result gives two classes with explicit uniform bounds. Call a user *heavy* if its degree exceeds k.

Proposition 7 (Classes with bounded response norm)**.** *For every $\lambda > 0$, the following bounds hold along every cap path associated with a node deletion from the stated class.*

- *(i) If each example contains at most one heavy user, then $\rho \le 1$, and the total fractional movement is at most rk.*
- *(ii) If the structure is a disjoint union of simple cliques, then $\rho \le 3/2$, and the total fractional movement is at most $5k/2$.*

Both classes allow arbitrarily large degrees. In the first, fractional examples do not couple saturated users. In the second, users are densely connected, but the incidence matrices still have bounded response norm. These give privacy guarantees with a group size independent of the number of users or examples. Appendix B.5 contains the proofs, exact clique and odd-cycle calculations, and a discussion of numerical evaluation.

5 From Stability to Node-Level Privacy

Under node adjacency the contents of shared examples are identical, so the learner sees insertions and deletions only from the symmetric difference of the selected sets, and group privacy accounts for them through the learner's example-level guarantee.

5.1 General composition theorem

We state the composition result in general form, applicable to any preprocessing algorithm with bounded instability. For randomized selection we compare outputs through a coupling, a joint distribution with the correct marginal on each input; sharing the uniforms and priorities in the rounding supplies one.

Theorem 8 (High-probability stability implies node-DP)**.** *Let $\mathcal{M}$ be a preprocessing algorithm that outputs a multiset of examples $\mathsf{Out}(\mathcal{D})$ to be fed to a downstream learner $\mathcal{A}$. Suppose that for every node-adjacent pair in the input domain there is a coupling (Z, Z') of $\mathsf{Out}(\mathcal{D})$ and $\mathsf{Out}(\mathcal{D}')$ such that $\Pr[d_{\mathrm{Ham}}(Z, Z') \le t] \ge 1 - \eta$, with the integer $t \ge 0$ and $\eta \in [0,1)$ not depending on the pair. If $\mathcal{A}$ is one $(\varepsilon_0, \delta_0)$-example-DP mechanism on all finite multisets, run with independent randomness, then $\mathcal{A} \circ \mathcal{M}$ is (ε, δ)-node-DP with*

$$\varepsilon = t\,\varepsilon_0, \qquad \delta = \eta \;+\; \delta_0 \textstyle\sum_{i=0}^{t-1} e^{i\varepsilon_0} \;\le\; \eta \;+\; t\,e^{(t-1)\varepsilon_0}\delta_0.$$

On coupled pairs at distance at most t we apply group privacy; the remaining pairs have total probability at most η, which is why η enters δ additively (proof in Appendix C.1).

5.2 A CHARACTERIZATION FOR DETERMINISTIC SELECTORS

For deterministic selection, no factor smaller than the worst-case recourse works for every example-private learner and every assignment of contents. We state the equivalence for pure privacy.

Theorem 9 (Characterization for deterministic selection)**.** *Let $\mathcal{M}$ be a deterministic structure-only selector on $\mathcal{H}$, and suppose $\mathcal{X}$ contains two distinguishable values. Fix $\varepsilon_0 > 0$ and an integer $t \geq 0$. The following are equivalent: (i) every node-adjacent pair $H, H' \in \mathcal{H}$ satisfies $|S(H) \triangle S(H')| \leq t$; (ii) for every $(\varepsilon_0, 0)$-example-DP learner $\mathcal{A}$, the composition $\mathcal{A} \circ \mathcal{M}$ is $(t\varepsilon_0, 0)$-node-DP on all attributed datasets in $\mathbb{D}_{\mathcal{H}}$.*

Sufficiency follows from group privacy. For necessity, we assign two content values so that the change in a bounded sum equals the selection recourse, and Laplace noise on that sum gives an example-private learner that witnesses the factor (proof in Appendix C.3).

Randomized selectors. Appendix C.3 gives the converse for randomized selectors. When one input has a fixed selected set and its neighbor moves at least T selected examples with probability at least η, node privacy requires $\varepsilon \geq T\varepsilon_0 + \log(\eta - 2\delta)$ for $\delta < \eta/2$ (Lemma 19); on the odd cycle every exact selector therefore pays $\varepsilon \geq m\varepsilon_0 - \log 4$ at $\delta \leq 1/8$ (Theorem 20), and every learner that matches a benchmark accuracy on an explicit classification task pays a comparable cost (Theorem 21).

5.3 INSTANTIATION FOR REGULARIZED CONTRIBUTION BOUNDING

We now apply Theorem 8 to the regularized selector, with the recourse bound of Theorem 4 and the public size bound.

Corollary 10 (Node-DP via regularized contribution bounding)**.** *Fix public bounds N and r, a cap $1 \leq k \leq N$, and $\lambda > 0$. Consider the input class with at most N examples, each of size at most r. Solve* (3)*, apply Algorithm 1 with weights $y^{\lambda}/2$, and run a single $(\varepsilon_0, \delta_0)$-example-DP learner on the selected contents. For $\eta \in (0, 1)$, set $\mu_N = k + \sqrt{2Nk/\lambda}$ and $t_N = \tau(\mu_N/2, \eta)$. The resulting pipeline is node-private with $\varepsilon = t_N\varepsilon_0$ and $\delta = \eta + \delta_0 \sum_{i=0}^{t_N - 1} e^{i\varepsilon_0}$. At constant $\lambda > 0$, $t_N = O(r(\sqrt{Nk} + \log(1/\eta))) = \tilde{O}(r\sqrt{Nk})$.*

For a target budget $\varepsilon > 0$, $\delta \in (0, 1)$, we set $\eta = \delta/2$, $\varepsilon_0 = \varepsilon/t_N$, and $\delta_0 \leq \delta e^{-\varepsilon}/(2t_N)$; these parameters are fixed across neighboring inputs. Under the structural assumption of Corollary 6 we use t_ρ instead (proof for t_N in Appendix C.2).

5.4 LEARNING UTILITY

Contribution bounding keeps fewer training examples, but a stable selection lets the learner use less privacy noise. This tradeoff can be expressed through the retained size $M = |S|$ and the recourse factor t. Suppose the contents are independent draws from a distribution P, independently of the structure; since selection uses only the structure, the retained contents are still independent draws from P. For a learner with expected excess population risk $\tilde{O}(m^{-1/2} + \sqrt{d}/(\varepsilon_0 m))$ on m examples, calibration at $\varepsilon_0 = \varepsilon/t$ gives, conditional on $M \geq 1$,

$$\mathbb{E}\Big[L_P(\widehat{\theta}) - L_P^{\star} \mid M\Big] \leq \tilde{O}\Big(\frac{1}{\sqrt{M}} + \frac{\sqrt{d}\,t}{\varepsilon M}\Big),$$

where L_P is population risk and $L_P^{\star}$ its optimum. Learner must satisfy the variable-size example-DP assumption (Sec. 2), with δ_0 calibrated (Corollary 10); private stochastic convex optimization attains the stated rate under convexity, Lipschitz, and bounded-domain assumptions (Bassily et al., 2019).

The first term is the statistical cost of using the selected examples; the second is the privacy cost. Direct learning on all n examples has privacy term proportional to $\bar{\Delta}/n$, while the pipeline's term is proportional to t/M; if a constant fraction of the data is retained, it suffices for t to be smaller than the public degree bound $\bar{\Delta}$. The full guarantee, retention conditions, the comparison with direct learning, variable sample sizes, and distribution shift are in Appendix D.

6 EXPERIMENTS

We study how selection stability and retention affect learning on real attribution structures. The main comparison uses DBLP coauthorship and the Ask Ubuntu, Math Stack Exchange, and Stack Overflow

Table 1: Learning on real attribution structures with synthetic contents, at $\varepsilon = 4$ and $\delta = 10^{-5}$. Means over ten trials. The pipeline uses either conservative (t_N) or structural (t_ρ) calibration; direct learning uses all examples. Accuracy gains are paired differences in percentage points, with one standard error. Structural calibration uses the empirical response-bound scenarios stated below.

Logistic-regression accuracy (%; higher is better)				
Structure	Conservative	Structural	Direct	Gain over direct
DBLP	72.30	90.32	90.36	-0.04 ± 0.07
Ask Ubuntu	58.02	87.59	85.23	2.36 ± 0.86
Math Stack Exchange	58.57	88.75	81.68	7.06 ± 1.02
Stack Overflow	72.86	90.35	90.10	0.25 ± 0.11
Mean-estimation RMSE (lower is better)				
Structure	Conservative	Structural	Direct	Direct / structural
DBLP	0.1912	0.00360	0.00116	$0.32\times$
Ask Ubuntu	0.7770	0.04563	0.04222	$0.93\times$
Math Stack Exchange	0.6908	0.02326	0.05829	$2.51\times$
Stack Overflow	0.1413	0.00146	0.01036	$7.08\times$

thread datasets of Benson et al. (2018). We keep examples involving at most three users, giving 2.76 million papers, 187 thousand Ask Ubuntu threads, 646 thousand Math Stack Exchange threads, and 10.03 million Stack Overflow threads. Features and labels are generated synthetically on these fixed structures, so the comparisons isolate the effect of attribution on learning.

Setup. We compare the same regularized selection under conservative calibration t_N and structural calibration t_ρ, together with direct learning on all examples using the degree-based factor $\bar{\Delta}$. We use $\lambda = 1$, half-weight priority rounding, cap $k = 3$ for DBLP and $k = 10$ for the thread datasets, and target budget $(\varepsilon, \delta) = (4, 10^{-5})$. Within each of ten trials, the pipeline arms share the selected examples, contents, test data, and base noise draws; only the calibration changes. The structural arm uses $\bar{\rho} = 4.5$ for DBLP and $\bar{\rho} = 3$ for the thread datasets. These are empirical calibration scenarios: the corresponding node-DP guarantee requires a uniform response bound over the input class, which pointwise measurements do not establish. Appendix E.1 gives the mechanisms, public-bound conventions, and numerical details.

Learning on real attribution structures. Table 1 shows that structural calibration substantially improves the utility of the selected data. Relative to the conservative pipeline, it increases logistic-regression accuracy by 17.5 to 30.2 percentage points across the four structures, with no change in retention. Comparison with direct learning shows where this improvement also compensates for the examples removed by contribution bounding. On Math Stack Exchange, accuracy increases from 81.68% to 88.75%, a paired gain of 7.06 ± 1.02 percentage points, and mean-estimation RMSE is reduced by a factor of 2.51. On Ask Ubuntu, accuracy gain is 2.36 ± 0.86 points. On Stack Overflow, structural pipeline reaches 90.35% accuracy, close to the nonprivate reference of 90.36%, while reducing mean-estimation RMSE by a factor of 7.08 relative to direct learning. DBLP has a smaller maximum user degree relative to its size. Direct learning and structural calibration have similar accuracy, and direct learning gives smaller mean-estimation error. These comparisons reflect noise reduction from stable selection and the amount of data retained, the two parts of the utility tradeoff.

Stability and scaling. The appendix also compares selection rules under user deletion on email and coauthorship structures, tests the instability constructions, and studies synthetic structures with up to three million users and twelve million examples. These experiments connect the learning results to the movement and retention of the selector. Figures, complete comparisons, and response-norm diagnostics are collected in Appendix E.

7 Conclusion

Selection stability under node adjacency transfers example-level privacy to node-level privacy for any downstream learner, and regularization with feasible rounding gives a selector with explicit stability and retention guarantees. Bounding the capacity released by a deletion yields structural classes on which the privacy factor is independent of dataset size and degree; the learning bounds and experiments show when the reduced noise outweighs the retained data lost.

AI USE STATEMENT

In this work, we used generative AI tools to help design and provide feedback on the research methodology and experiments, clean and reformat data, and interpret results. We have not used generative AI tools to develop the theoretical framework, formulate mathematical claims, provide critical ingredients for proving them, assist in the writing of proofs, implement methods, generate synthetic data sets or to propose or refine hypotheses, and assisting with translation and supporting qualitative and thematic data analysis are not applicable to this work. Additionally, we used generative AI tools to summarize and analyze existing literature, format references, suggest a structure for the paper, and propose a title and keywords. We have reviewed all AI-assisted work. The authors checked the literature summaries and bibliographic entries against the original papers. AI-written scripts converted the saved trial records of the learning experiments into the tables and figures, and the authors checked the converted values against the saved results. AI tools also drafted interpretations of the experimental results. The authors compared each interpretation with the saved results and removed or qualified any claim the data did not support, such as when a sampled-column ridge diagnostic had been described as the response norm itself, or when the pipeline was said to match non-private accuracy at a size where it was still below it. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work is theoretical, and its experiments use public attribution structures with synthetic contents, so no private data is processed. Each privacy guarantee is stated with the public bounds it assumes.

REPRODUCIBILITY STATEMENT

Complete proofs appear in the appendix with the assumptions of each result stated. The supplementary material contains the code, the saved results, and a notebook that regenerates the tables and figures.

REFERENCES

Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016. URL https://arxiv.org/abs/1607.00133.

Sushant Agarwal, Gautam Kamath, Mahbod Majid, Argyris Mouzakis, Rose Silver, and Jonathan Ullman. Private mean estimation with person-level differential privacy. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025. URL https://arxiv.org/abs/2405.20405.

Kareem Amin, Alex Kulesza, Andres Munoz, and Sergei Vassilvitskii. Bounding user contributions: A bias-variance trade-off in differential privacy. In *International Conference on Machine Learning*, pp. 263–271. PMLR, 2019. URL https://proceedings.mlr.press/v97/amin19a.html.

Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pp. 394–403. PMLR, 2018. URL https://proceedings.mlr.press/v80/balle18a.html.

Nikhil Bansal, Nitish Korula, Viswanath Nagarajan, and Aravind Srinivasan. Solving packing integer programs via randomized rounding with alterations. *Theory of Computing*, 8(1):533–565, 2012. URL https://theoryofcomputing.org/articles/v008a024/.

Raef Bassily and Ziteng Sun. User-level private stochastic convex optimization with optimal rates. In *International Conference on Machine Learning*, pp. 1838–1851. PMLR, 2023. URL https://proceedings.mlr.press/v202/bassily23a.html.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Differentially private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pp. 464–473. IEEE, 2014. URL https://arxiv.org/abs/1405.7085.

Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 1046–1059, 2016. URL https://arxiv.org/abs/1511.02513.

Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Guha Thakurta. Private stochastic convex optimization with optimal rates. *Advances in Neural Information Processing Systems*, 32, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/3bd8fdb090f1f5eb66a00c84dbc5ad51-Abstract.html.

Alberto Bemporad, Manfred Morari, Vivek Dua, and Efstratios N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, 2002. doi: 10.1016/S0005-1098(01)00174-1.

Austin R. Benson, Rediet Abebe, Michael T. Schaub, Ali Jadbabaie, and Jon Kleinberg. Simplicial closure and higher-order link prediction. *Proceedings of the National Academy of Sciences*, 115 (48):E11221–E11230, 2018. URL https://arxiv.org/abs/1802.06916.

Piotr Berman. A $d/2$ approximation for maximum weight independent set in d-claw free graphs. In *Scandinavian Workshop on Algorithm Theory (SWAT)*, volume 1851 of *Lecture Notes in Computer Science*, pp. 214–219, 2000. URL https://link.springer.com/chapter/10.1007/3-540-44985-X_19.

Piotr Berman and Martin Fürer. Approximating maximum independent set in bounded degree graphs. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 365–371, 1994. URL https://pure.psu.edu/en/publications/approximating-maximum-independent-set-in-bounded-degree-graphs/.

Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pp. 87–96, 2013. URL https://arxiv.org/abs/1208.4586.

Christian Borgs, Jennifer Chayes, and Adam Smith. Private graphon estimation for sparse graphs. In *Advances in Neural Information Processing Systems*, volume 28, 2015. URL https://papers.nips.cc/paper/2015/hash/7250eb93b3c18cc9daa29cf58af7a004-Abstract.html.

Christian Borgs, Jennifer Chayes, Adam Smith, and Ilias Zadik. Revealing network structure, confidentially: Improved rates for node-private graphon estimation. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 533–543, 2018. URL https://arxiv.org/abs/1810.02183.

Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002. URL https://jmlr.org/papers/v2/bousquet02a.html.

Zachary Charles, Arun Ganesh, Ryan McKenna, H Brendan McMahan, Nicole Mitchell, Krishna Pillutla, and Keith Rush. Learning with user-level differential privacy under fixed compute budgets. In *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 2025. URL https://arxiv.org/abs/2407.07737. Preprint title: Fine-tuning large language models with user-level differential privacy, arXiv:2407.07737; accepted-papers listing at https://satml.org/2025/accepted-papers/.

Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011. URL https://jmlr.org/papers/v12/chaudhuri11a.html.

Shixi Chen and Shuigeng Zhou. Recursive mechanism: Towards node differential privacy and unrestricted joins. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pp. 653–664, 2013. URL https://arxiv.org/abs/1304.4795.

Lynn Chua, Badih Ghazi, Yangsibo Huang, Pritish Kamath, Ravi Kumar, Daogao Liu, Pasin Manurangsi, Amer Sinha, and Chiyuan Zhang. Mind the privacy unit! User-level differential privacy for language model fine-tuning. In *First Conference on Language Modeling (COLM)*, 2024. URL https://arxiv.org/abs/2406.14322.

Vincent Cohen-Addad, Alessandro Epasto, Jason Lee, and Morteza Zadimoghaddam. Scalable contribution bounding to achieve privacy. *arXiv preprint arXiv:2507.23432*, 2025. URL https://arxiv.org/abs/2507.23432.

Rachel Cummings, Vitaly Feldman, Audra McMillan, and Kunal Talwar. Mean estimation with user-level privacy under data heterogeneity. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://papers.nips.cc/paper_files/paper/2022/hash/bbba680acd2826f23928c6675a19f0e7-Abstract-Conference.html.

Marek Cygan. Improved approximation for 3-dimensional matching via bounded pathwidth local search. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 509–518. IEEE, 2013. URL https://arxiv.org/abs/1304.1424.

Ameya Daigavane, Gagan Madan, Aditya Sinha, Abhradeep Guha Thakurta, Gaurav Aggarwal, and Prateek Jain. Node-level differentially private graph neural networks. In *ICLR 2022 Workshop on Privacy, Accountability, Interpretability, Robustness, Reasoning on Structured Data (PAIR^{2}Struct)*, 2022. URL https://arxiv.org/abs/2111.15521.

Michael Dinitz, George Z. Li, Quanquan C. Liu, and Felix Zhou. Differentially private matchings. *arXiv preprint arXiv:2501.00926*, 2025. URL https://arxiv.org/abs/2501.00926.

Andy Dong and Arun Ganesh. Privacy amplification for BandMF via b-min-sep subsampling. *arXiv preprint arXiv:2602.09338*, 2026. URL https://arxiv.org/abs/2602.09338.

Wei Dong and Ke Yi. Residual sensitivity for differentially private multi-way joins. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD)*, 2021. URL https://www.cse.ust.hk/~yike/ResidualSensitivity-full.pdf.

Cynthia Dwork and Aaron Roth. *The algorithmic foundations of differential privacy*. Now Publishers, Inc., 2014. URL https://www.nowpublishers.com/article/Details/TCS-042.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284. Springer, 2006. URL https://link.springer.com/chapter/10.1007/11681878_14.

Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Generalization in adaptive data analysis and holdout reuse. In *Advances in Neural Information Processing Systems*, volume 28, 2015. URL https://proceedings.neurips.cc/paper/2015/hash/bad5f33780c42f2588878a9d07405083-Abstract.html.

Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17: 449–467, 1965. URL https://www.cambridge.org/core/journals/canadian-journal-of-mathematics/article/paths-trees-and-flowers/08B492B72322C4130AE800C0610E0E21.

Mourad El Ouali, Antje Fretwurst, and Anand Srivastav. Inapproximability of b-matching in k-uniform hypergraphs. In *International Workshop on Algorithms and Computation (WALCOM)*, volume 6552 of *Lecture Notes in Computer Science*, pp. 57–69. Springer, 2011. URL https://link.springer.com/chapter/10.1007/978-3-642-19094-0_8.

Alessandro Epasto, Mohammad Mahdian, Jieming Mao, Vahab Mirrokni, and Lijie Ren. Smoothly bounding user contributions in differential privacy. In *Advances in Neural Information Processing Systems*, volume 33, pp. 13999–14010, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/a0dc078ca0d99b5ebb465a9f1cad54ba-Abstract.html.

Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy amplification by iteration. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 521–532. IEEE, 2018. URL https://arxiv.org/abs/1808.06651.

Arun Ganesh, Ryan McKenna, H. Brendan McMahan, Adam Smith, and Fan Wu. It's my data too: Private ML for datasets with multi-user training examples. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, volume 267 of *Proceedings of Machine Learning Research*, pp. 18189–18205, 2025. URL https://proceedings.mlr.press/v267/ganesh25a.html.

Badih Ghazi, Ravi Kumar, and Pasin Manurangsi. User-level differentially private learning via correlated sampling. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/a89cf525e1d9f04d16ce31165e139a4b-Abstract.html.

Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Raghu Meka, and Chiyuan Zhang. User-level differential privacy with few examples per user. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/3d57795f0e263aa69577f1bbceade46b-Abstract-Conference.html.

Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. Differentially private combinatorial optimization. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1106–1125, 2010. URL https://arxiv.org/abs/0903.4510.

Johan Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, 182(1):105–142, 1999. URL https://link.springer.com/article/10.1007/BF02392825.

Elad Hazan, Shmuel Safra, and Oded Schwartz. On the complexity of approximating k-set packing. *Computational Complexity*, 15(1):20–39, 2006. URL https://link.springer.com/article/10.1007/s00037-006-0205-6.

Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, 2014a. URL https://arxiv.org/abs/1311.2828.

Justin Hsu, Aaron Roth, Tim Roughgarden, and Jonathan Ullman. Privately solving linear programs. In *Automata, Languages, and Programming (ICALP)*, volume 8572 of *Lecture Notes in Computer Science*, pp. 612–624. Springer, 2014b. URL https://link.springer.com/chapter/10.1007/978-3-662-43948-7_51.

Yihua Hu, Hao Ding, and Wei Dong. N2e: A general framework to reduce node-differential privacy to edge-differential privacy for graph analytics. *Proceedings of the ACM on Management of Data*, 2025. URL https://arxiv.org/abs/2511.20125.

Palak Jain, Adam Smith, and Connor Wagaman. Time-aware projections: Truly node-private graph statistics under continual observation. In *2024 IEEE Symposium on Security and Privacy (SP)*, 2024. URL https://arxiv.org/abs/2403.04630.

Iden Kalemaj, Sofya Raskhodnikova, Adam Smith, and Charalampos E Tsourakakis. Node-differentially-private estimation of the number of connected components. In *Proceedings of the 42nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, pp. 183–194, 2023. URL https://arxiv.org/abs/2304.05890.

Shiva Prasad Kasiviswanathan, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Analyzing graphs with node differential privacy. In *Theory of Cryptography Conference*, pp. 457–476. Springer, 2013. URL https://link.springer.com/chapter/10.1007/978-3-642-36594-2_26.

Daniel Kifer and Ashwin Machanavajjhala. No free lunch in data privacy. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pp. 193–204, 2011. URL https://pure.psu.edu/en/publications/no-free-lunch-in-data-privacy.

Daniel Kifer and Ashwin Machanavajjhala. Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems*, 39(1), 2014. URL https://users.cs.duke.edu/~ashwin/pubs/pufferfish_TODS.pdf.

Soh Kumabe and Yuichi Yoshida. Lipschitz continuous algorithms for graph problems. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, 2023. URL https://arxiv.org/abs/2211.04674.

Daniel Levy, Ziteng Sun, Kareem Amin, Satyen Kale, Alex Kulesza, Mehryar Mohri, and Ananda Theertha Suresh. Learning with user-level privacy. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/67e235e7f2fa8800d8375409b566e6b6-Abstract.html.

Yang Li, Michael Purcell, Thierry Rakotoarivelo, David Smith, Thilina Ranbaduge, and Kee Siong Ng. Private graph data release: A survey. *ACM Computing Surveys*, 2023. doi: 10.1145/3569085. URL https://arxiv.org/abs/2107.04245.

Changchang Liu, Supriyo Chakraborty, and Prateek Mittal. Dependence makes you vulnerable: Differential privacy under dependent tuples. In *Network and Distributed System Security Symposium (NDSS)*, 2016. URL https://www.princeton.edu/~pmittal/publications/ddp-ndss16.pdf.

Quanquan C. Liu, Grigoris Velegkas, Yuichi Yoshida, and Felix Zhou. Pointwise lipschitz continuous graph algorithms. *arXiv preprint arXiv:2405.08938*, 2024. URL https://arxiv.org/abs/2405.08938.

Suqing Liu, Xuan Bi, and Tianxi Li. Grand: Graph release with assured node differential privacy. *arXiv preprint arXiv:2507.00402*, 2025. URL https://arxiv.org/abs/2507.00402.

Yuhan Liu, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Michael Riley. Learning discrete distributions: User vs item-level privacy. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/f06edc8ab534b2c7ecbd4c2051d9cb1e-Abstract.html.

Yuhan Liu, Ananda Theertha Suresh, Wennan Zhu, Peter Kairouz, and Marco Gruteser. Algorithms for bounding contribution for histogram estimation under user-level privacy. In *International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*. PMLR, 2023. URL https://proceedings.mlr.press/v202/liu23ae.html.

Meike Neuwohner. Passing the limits of pure local search for weighted k-set packing. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1090–1137, 2023. URL https://arxiv.org/abs/2202.01248.

Sofya Raskhodnikova and Adam Smith. Lipschitz extensions for node-private graph statistics and the generalized exponential mechanism. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 495–504. IEEE, 2016. URL https://cs-people.bu.edu/sofya/pubs/Lipschitz-deg-ext-FOCS16.pdf.

Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003. URL https://homepages.cwi.nl/~lex/co/.

Shuang Song, Yizhen Wang, and Kamalika Chaudhuri. Pufferfish privacy mechanisms for correlated data. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 2017. URL https://arxiv.org/abs/1603.03977.

Aravind Srinivasan. Improved approximation guarantees for packing and covering integer programs. *SIAM Journal on Computing*, 29(2):648–670, 1999. URL https://www.cs.umd.edu/~srin/papers.html.

Maxim Sviridenko and Justin Ward. Large neighborhood local search for the maximum set packing problem. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 792–803, 2013. URL https://link.springer.com/chapter/10.1007/978-3-642-39206-1_67.

Theophile Thiery and Justin Ward. An improved approximation for maximum weighted k-set packing. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1138–1162. SIAM, 2023. URL https://arxiv.org/abs/2301.07537.

Jonathan Ullman and Adam Sealfon. Efficiently estimating Erdős–Rényi graphs with node differential privacy. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/hash/955cb567b6e38f4c6b3f28cc857fc38c-Abstract.html.

Yu-Xiang Wang, Jing Lei, and Stephen E. Fienberg. Learning with differential privacy: Stability, learnability and the sufficiency and necessity of ERM principle. *Journal of Machine Learning Research*, 17(183):1–40, 2016. URL https://jmlr.org/papers/v17/15-313.html.

GUIDE TO THE APPENDIX

The appendix brings together the proofs, extensions, and experimental details. The analysis follows the progression from contribution bounding to privacy and learning, with each proof grouped with the extensions of the result it supports.

- **Appendix A: Related work.** Multi-attribution privacy, node-private graph analysis, stable optimization, and approximation background.
- **Appendix B: Contribution bounding and stability.** Instability constructions; regularization and approximate optimization; feasible rounding and retention; structural response bounds; alternative regularizers.
- **Appendix C: Privacy transfer and lower bounds.** Composition, calibration, the deterministic characterization, randomized lower bounds, and privacy extensions.
- **Appendix D: Learning utility and extensions.** Population-risk guarantees, comparison with direct learning, variable-size learners, and distribution shift.
- **Appendix E: Experiments.** Reproducible methods, complete learning comparisons, scaling studies, selector comparisons, and response diagnostics.

A RELATED WORK

We build on prior work in differential privacy, algorithmic stability, and approximation algorithms.

Contribution bounding and multi-attribution privacy. In user-level private learning, contribution bounds control how many records from one user can affect training. Prior work studies the statistical tradeoffs of capping contributions, smooth contribution limits, and the choice of cap (Amin et al., 2019; Epasto et al., 2020; Liu et al., 2023). For examples associated with several users, Ganesh et al. (2025) introduced the fixed-graph privacy model and studied greedy contribution bounding, including its utility in logistic regression and transformer training. Cohen-Addad et al. (2025) developed a distributed algorithm in which users propose examples subject to their remaining capacities. Dong & Ganesh (2026) extended subsampling and privacy amplification for banded matrix-factorization mechanisms to multi-attribution learning under fixed-graph adjacency. We study how contribution bounding changes when a user and all incident examples are added or removed, and how those selection changes affect a downstream example-DP learner.

Node-private graph analysis. Node privacy has been studied through restricted sensitivity, graph projections, Lipschitz extensions, and mechanisms for specific graph statistics (Kasiviswanathan et al., 2013; Blocki et al., 2013; Chen & Zhou, 2013; Raskhodnikova & Smith, 2016; Borgs et al., 2015; 2018; Ullman & Sealfon, 2019; Kalemaj et al., 2023; Jain et al., 2024). Related applications include graph release and graph neural networks (Liu et al., 2025; Daigavane et al., 2022); Li et al. (2023) survey this literature. Particularly close to our pipeline is the N2E framework of Hu et al. (2025), which combines clipping that controls edge distance with private degree estimation and then applies an edge-DP mechanism. We likewise use control of changes after preprocessing to transfer a record-level privacy guarantee. Our analysis concerns feasible contribution bounding on hypergraphs, its retention objective, and an exact privacy characterization for deterministic structure-only selectors. The structural response bounds further give calibrations independent of maximum degree on specified input classes.

Stable combinatorial optimization. Kumabe & Yoshida (2023) study Lipschitz continuity of graph algorithms under weight perturbations and use an entropy-regularized LP for pointwise-stable bipartite matching. Liu et al. (2024) develop an analysis of regularized LP relaxations, including minimum cuts and b-matching. This line of work provides a close algorithmic connection between regularization, rounding, and stability. Here a node deletion removes variables and changes the contribution constraints on surviving examples. Our fractional proof compares optimizers through a feasible zero extension and the lost objective contribution. The structural analysis expresses deletion as a change in capacities on the surviving structure, identifying the incidence-based response norm

used for node-privacy calibration. The approximation background for hypergraph b-matching is reviewed in Appendix A.

Private optimization and downstream learning. Private matching and allocation algorithms aim to provide privacy for the released combinatorial solution (Gupta et al., 2010; Hsu et al., 2014a;b; Dinitz et al., 2025). In our pipeline, the feasible selection is internal and only the downstream output is released. Our guarantees apply to a learner through its example-DP property, so they can be combined with private empirical risk minimization and stochastic optimization (Chaudhuri et al., 2011; Bassily et al., 2014; 2019) or other learners with the required adjacency guarantee, such as suitably calibrated DP-SGD (Abadi et al., 2016). User-level learning in the single-attribution setting has been studied for mean estimation, learning, and model training (Liu et al., 2020; Levy et al., 2021; Ghazi et al., 2021; Cummings et al., 2022; Bassily & Sun, 2023; Charles et al., 2025; Ghazi et al., 2023; Chua et al., 2024; Agarwal et al., 2025). We focus on the additional selection changes caused by shared attribution.

Other connections to privacy and stability. Stability also appears in generalization, adaptive data analysis, and privacy amplification by iteration (Bousquet & Elisseeff, 2002; Dwork et al., 2015; Bassily et al., 2016; Feldman et al., 2018; Wang et al., 2016). The stability used here measures changes to the selected dataset before learning. Privacy frameworks for shared or correlated records address related questions about which changes to protect (Kifer & Machanavajjhala, 2011; 2014; Liu et al., 2016; Song et al., 2017; Dong & Yi, 2021). We fix node addition and deletion as the adjacency relation and analyze the resulting two-phase mechanism.

Algorithmic stability and privacy. In learning theory, stability commonly measures how predictions or losses change after a training example is modified (Bousquet & Elisseeff, 2002). Differential privacy also yields stability guarantees useful for generalization under adaptive data analysis (Dwork et al., 2015; Bassily et al., 2016). In the other direction, sensitivity bounds support private releases through calibrated noise (Dwork et al., 2006; Chaudhuri et al., 2011), while contractive iterations can amplify privacy (Feldman et al., 2018). Our stability quantity is the distance between selected datasets under a user addition or deletion. We use this distance to transfer privacy through an arbitrary example-DP learner, and characterize the exact transfer factor for deterministic selectors.

Approximation algorithms for contribution bounding. The contribution bounding optimization problem is the maximum unweighted b-matching problem in hypergraphs, a prototypical column-sparse packing integer program. For uniform budget $k = 1$ it reduces to r-Set Packing, which is NP-hard to approximate within $\Omega(r/\log r)$ (Hazan et al., 2006) and admits a $(r+1+\varepsilon)/3$-approximation via local search (Cygan, 2013) (see also Sviridenko & Ward (2013), and Berman (2000); Neuwohner (2023); Thiery & Ward (2023) for the weighted case); randomized rounding with alterations achieves an $(er + o(r))$-approximation for r-column-sparse packing (Bansal et al., 2012), with guarantees improving as capacities grow (Srinivasan, 1999); on the hardness side, El Ouali et al. (2011) state a lower bound of $\Omega(r/(k \log r))$ for r-uniform hypergraphs in the range $2 \le k \le r/\log r$; we note that at $k = 2$ this exceeds the polynomial-time $O(\sqrt{r})$ ratio that follows from Bansal et al. (2012) for capacities at least two, so the two statements cannot both hold for large r unless $\mathrm{P} = \mathrm{NP}$, and we do not rely on the hardness claim for $k \ge 2$, while for graphs ($r = 2$) the problem is exactly solvable for every k via simple b-matching (Schrijver, 2003). Related is maximum independent set, hard within $\Omega(n^{1-\epsilon})$ in general (Håstad, 1999) and $O(\Delta)$-approximable in bounded-degree graphs (Berman & Fürer, 1994). Our instability result (Theorem 1) shows that even *optimal* solutions to these problems are unstable under node adjacency, establishing that approximation quality alone does not control the privacy cost; our regularized selector provides explicit stability and retention guarantees.

Remark 1 (Approximation hardness)**.** *The approximation complexity depends on the realized maximum hyperedge size $r_{\mathrm{real}} = \max_{e \in E} |e| \le r$, not the maximum degree $\Delta = \max_{v \in V} \deg(v)$. For uniform budget $k = 1$, the problem is equivalent to r-Set Packing, which is NP-hard to approximate within $\Omega(r/\log r)$ (Hazan et al., 2006). When $r = 2$ (graphs), the problem becomes polynomial-time solvable via maximum matching (Edmonds, 1965). The maximum degree Δ affects privacy and stability rather than approximation hardness.*

B CONTRIBUTION BOUNDING AND STABILITY

We first prove the instability results, then analyze regularization, rounding, and retention. The final two subsections develop the structural response bounds and compare alternative regularizers.

B.1 PROOFS FOR INSTABILITY

B.1.1 PROOF OF THEOREM 1 AND COROLLARY 2

Proposition 11 (The greedy baseline is $\Omega(|V|)$-unstable on the odd cycle)**.** *Let $H = C_{2m+1}$ with vertices $0, 1, \dots, 2m$, hyperedges $e_i = \{i-1, i\}$ for $i \in [2m+1]$ (indices modulo $2m+1$), uniform cap $k = 1$, and let the greedy rule of Ganesh et al. (2025) process the hyperedges in the fixed order $e_1 \prec e_2 \prec \cdots \prec e_{2m+1}$, accepting each hyperedge that keeps every vertex within its cap. Then greedy returns $R = \{e_1, e_3, \dots, e_{2m-1}\}$ on H and $A = \{e_2, e_4, \dots, e_{2m}\}$ on the instance H' obtained by deleting vertex 0, so $|R \triangle A| = 2m = |V| - 1$ deterministically, and, for every $\varepsilon_0 > 0$, there are contents and an $(\varepsilon_0, 0)$-example-DP mechanism for which every (ε, δ)-node-DP guarantee for the composition requires $\varepsilon \geq 2m\varepsilon_0 + \log(1 - 2\delta)$ when $\delta < 1/2$.*

Proof. On H, e_1 is accepted; e_2 shares vertex 1 with e_1 and is rejected; e_3 is accepted; and so on, so the odd-indexed hyperedges $e_1, e_3, \dots, e_{2m-1}$ are accepted while each even-indexed hyperedge e_{2i} is rejected because vertex $2i-1$ is already saturated by e_{2i-1}. Finally $e_{2m+1} = \{2m, 0\}$ is rejected because vertex 0 is saturated by e_1. Deleting vertex 0 removes e_1 and e_{2m+1} and leaves the path $e_2, \dots, e_{2m}$; in the same order, e_2 is now accepted, e_3 is rejected at vertex 2, e_4 is accepted, and so on, giving the even-indexed hyperedges. The two selections are disjoint and each has m elements. The deleted instance's greedy output is deterministic, so the separation premise of Lemma 19 holds with $T = 2m$ and $\eta = 1$. □

Proof of Theorem 1. Label the vertices $0, 1, \dots, 2m$ and the edges $e_i = \{i-1, i\}$ for $i \in [2m+1]$, indices modulo $2m+1$. With uniform cap $k = 1$ the LP is $\max \sum_e y_e$ subject to $y_{e_i} + y_{e_{i+1}} \leq 1$ for every vertex (each vertex is incident to exactly two consecutive edges) and $y \in [0,1]^E$.

Value and uniqueness on H. Summing the $2m+1$ vertex constraints counts each edge exactly twice, so $2\sum_e y_e \leq 2m+1$ and the LP value is at most $m + \frac{1}{2}$; the all-halves vector is feasible and attains it. Conversely, any optimal y attains equality in the summed bound, hence in every vertex constraint: $y_{e_i} + y_{e_{i+1}} = 1$ for all i. Thus $y_{e_{i+1}} = 1 - y_{e_i}$, and iterating this relation once around the cycle returns to y_{e_1} after an odd number of alternations, giving $y_{e_1} = 1 - y_{e_1}$. Hence $y_{e_1} = \frac{1}{2}$, and by propagation every coordinate equals $\frac{1}{2}$: the optimum is unique.

Value and uniqueness on $H \setminus \{v\}$. By symmetry take $v = 0$; the deletion removes edges e_1 and e_{2m+1} and leaves the path on vertices $1, \dots, 2m$ with edges $e_2, \dots, e_{2m}$. The path is bipartite, so the degree-constrained LP is integral and its optimal face is the convex hull of maximum matchings. A path on $2m$ vertices has exactly one perfect matching, $A = \{e_2, e_4, \dots, e_{2m}\}$, of size m: the leaf vertex 1 must be matched by e_2, which forces vertex 3 to be matched by e_4, and so on. Hence the LP value is m and the optimal face is the single point $\mathbf{1}_A$.

Distance. The zero-extension of $\mathbf{1}_A$ equals 0 on the two deleted edges and on the $m-1$ unmatched path edges, and 1 on the m matched edges, so

$$\left\| y^*(H) - \widetilde{y}^*(H \setminus \{v\}) \right\|_1 = 2 \cdot \tfrac{1}{2} + m \cdot \tfrac{1}{2} + (m-1) \cdot \tfrac{1}{2} = m + \tfrac{1}{2} = \frac{|V|}{2}. \qquad \square$$

Proof of Corollary 2. Delete $u = 0$ or $u' = 1$. As above, $H \setminus \{0\}$ is the path on $1, \dots, 2m$ with unique perfect matching $A = \{e_2, e_4, \dots, e_{2m}\}$, and $H \setminus \{1\}$ is the path on $2, \dots, 2m, 0$ with unique perfect matching $B = \{e_3, e_5, \dots, e_{2m+1}\}$; the two index sets are disjoint. Every maximum matching R of C_{2m+1} has $|R| = m$, and

$$|R \triangle A| = |R| + |A| - 2|R \cap A| = 2m - 2|R \cap A|,$$

likewise for B. Adding the two identities and using $|R \cap A| + |R \cap B| \leq |R| = m$ (disjointness of A and B) gives $|R \triangle A| + |R \triangle B| \geq 2m$. For a randomized selection rule the identity holds pointwise, so the two expected distances also sum to at least $2m$, and one of the two deletions moves

the selection by at least m in expectation. The bound is tight: any maximum matching $R \subseteq A \cup B$ attains $2m$ exactly. □

B.1.2 THE HUB INSTANCE WITH HETEROGENEOUS BUDGETS (REMARK 2)

Remark 2 (A hub instance with heterogeneous budgets)**.** *Instability scaling with the maximum degree can also be exhibited by a degree-Δ hub when per-user budgets are allowed to differ, a natural extension of formulation* (2)*. Take $V = \{u\} \cup \{l_1, \dots, l_\Delta\}$ with hyperedges $e_i = \{u, l_i\}$ and singletons $f_i = \{l_i\}$, and budgets $k_u = \Delta$, $k_{l_i} = 1$. The optimal face on H is $\{y_{e_i} + y_{f_i} = 1 \; \forall i\}$; under sequential lexicographic optimization with $e_1 \prec \cdots \prec e_\Delta \prec f_1 \prec \cdots \prec f_\Delta$ the selected solution takes all e_i, while on $H \setminus \{u\}$ the optimum is unique and takes all f_i, giving $\|y^* - \widetilde{y}'^*\|_1 = 2\Delta$ (Appendix B.1).*

We provide a construction showing $\|y^*(H) - y^*(H')\|_1 = 2\Delta$.

Construction. Let $V = \{u\} \cup L$ where $L = \{l_1, \dots, l_\Delta\}$ and u is a hub user. Create hyperedges:

- **Type A (with hub):** For each $i \in [\Delta]$, create hyperedge $e_i = \{u, l_i\}$ of size 2.
- **Type B (without hub):** For each $i \in [\Delta]$, create hyperedge $f_i = \{l_i\}$ of size 1.

Set contribution budgets:

- Hub u: budget $k_u = \Delta$
- All vertices in L: budget $k_{l_i} = 1$ for each $i \in [\Delta]$

The maximum degree is Δ (achieved by u). In total there are $|E| = 2\Delta$ hyperedges and $|V| = \Delta + 1$ vertices.

Lexicographic tie-breaking rule. When multiple optimal solutions exist, we select the solution obtained by *lexicographic optimization* under the ordering $e_1 \prec e_2 \prec \cdots \prec e_\Delta \prec f_1 \prec f_2 \prec \cdots \prec f_\Delta$: over the set of optimal solutions, first maximize y_{e_1}, then, subject to that, maximize y_{e_2}, and so on through all coordinates. Since each step maximizes a linear function over a nonempty compact convex set, the rule is well defined, and after all coordinates are processed the selected solution is unique.

Analysis on $H = (V, E)$. The linear program is:

$$\max \sum_{e \in E} y_e \quad \text{s.t.} \quad y_{e_1} + \cdots + y_{e_\Delta} \le \Delta \quad \text{(constraint for } u\text{)}$$
$$y_{e_i} + y_{f_i} \le 1 \quad \forall i \in [\Delta] \quad \text{(constraint for each } l_i\text{)}$$
$$y_e \in [0, 1] \quad \forall e \in E$$

Claim: The unique optimal solution under lexicographic tie-breaking is $y^*_{e_i} = 1$ for all $i \in [\Delta]$ and $y^*_{f_i} = 0$ for all $i \in [\Delta]$.

Proof of claim: The maximum objective value is Δ, which can be achieved by either:

- **Solution A:** $y_{e_i} = 1, y_{f_i} = 0$ for all i (select all Type A edges)
- **Solution B:** $y_{e_i} = 0, y_{f_i} = 1$ for all i (select all Type B edges)

Both are feasible and achieve objective Δ, as does every point of the optimal face $\{y : y_{e_i} + y_{f_i} = 1 \; \forall i\}$. Lexicographic optimization resolves this degeneracy: maximizing y_{e_1} over the optimal face gives $y_{e_1} = 1$, which forces $y_{f_1} = 0$; proceeding through $e_2, \dots, e_\Delta$ in order forces $y_{e_i} = 1$ and $y_{f_i} = 0$ for every i. Therefore the selected solution is exactly Solution A: $y^*_{e_i} = 1, y^*_{f_i} = 0$.

Analysis on $H' = H \setminus \{u\}$ (removing user u). After removing u, all Type A hyperedges $\{e_1, \dots, e_\Delta\}$ are also removed (they contain u). The remaining graph has vertices $V' = L$ and edges $E' = \{f_1, \dots, f_\Delta\}$.

The linear program on H' is:

$$\max \sum_{i=1}^{\Delta} y_{f_i} \quad \text{s.t.} \quad y_{f_i} \leq 1, \quad y_{f_i} \in [0,1] \quad \forall i.$$

The unique optimal solution is $y'^*_{f_i} = 1$ for all $i \in [\Delta]$ with objective Δ.

Computing ℓ_1 distance. To compare y^* and y'^*, we extend y'^* to the full domain E by setting $y'^*_{e_i} = 0$ for all i (the removed hyperedges). Then:

$$\begin{aligned}\|y^* - y'^*\|_1 &= \sum_{i=1}^{\Delta} |y^*_{e_i} - y'^*_{e_i}| + \sum_{i=1}^{\Delta} |y^*_{f_i} - y'^*_{f_i}| \\ &= \sum_{i=1}^{\Delta} |1-0| + \sum_{i=1}^{\Delta} |0-1| \\ &= \Delta + \Delta = 2\Delta.\end{aligned}$$

This completes the proof. □

B.1.3 Constant-degree graphs with linear instability

The hub construction of Remark 2 exhibits instability scaling with the maximum degree under heterogeneous budgets. We now present the even-cycle construction referenced in Remark 4: constant-degree graphs with uniform caps exhibit instability linear in the number of vertices under deterministic tie-breaking. Together with the odd-cycle results of Appendix B.1.1, this shows the problem is endemic to the unregularized formulation.

Theorem 12 (Linear instability at constant degree)**.** *There exist node-adjacent graphs $H = (V, E)$ and $H' = (V', E')$ with maximum degree $\Delta = 2$ and maximum edge size $r = 2$ such that under lexicographic tie-breaking, the selected solutions satisfy*

$$\|y^*(H) - \widetilde{y}^*(H')\|_1 = \Theta(|V|),$$

where $\widetilde{y}^*(H')$ is the zero-extension of $y^*(H')$ to $\mathbb{R}^E$.

Proof. Let H be the even cycle on $2m$ vertices $v_1, \ldots, v_{2m}$ with edges $e_i = (v_i, v_{i+1})$ for $i = 1, \ldots, 2m-1$ and $e_{2m} = (v_{2m}, v_1)$. Order the edges lexicographically: $e_1 \prec e_2 \prec \cdots \prec e_{2m}$. Set the budget $k = 1$ for all vertices (maximum matching).

The cycle has exactly two perfect matchings:

$$\begin{aligned}M_{\text{odd}} &= \{e_1, e_3, e_5, \ldots, e_{2m-1}\}, \\ M_{\text{even}} &= \{e_2, e_4, e_6, \ldots, e_{2m}\}.\end{aligned}$$

Both achieve the maximum objective value of m (matching all vertices), so both are optimal solutions to the LP relaxation. Under lexicographic tie-breaking, the solution containing the smallest edge e_1 is selected, giving $y^*(H) = M_{\text{odd}}$.

Now construct $H' = H \setminus \{v_1\}$ by removing vertex v_1 and its incident edges e_1 and e_{2m}. The remaining graph is a path on vertices $v_2, \ldots, v_{2m}$ with edges $e_2, \ldots, e_{2m-1}$. The maximum matching on this path has size $m - 1$.

The lexicographically first (inclusion-preferring) maximum matching of the path is

$$M' = \{e_2, e_4, e_6, \ldots, e_{2m-2}\}.$$

This is the unique optimal solution under lexicographic tie-breaking on H'.

To compute the symmetric difference, we extend $y^*(H')$ to the full domain E by setting $y^*_{e_1}(H') = 0$ and $y^*_{e_{2m}}(H') = 0$ (the removed edges). The matchings M_{odd} and M' are nearly disjoint:

- M_{odd} contains $\{e_1, e_3, e_5, \ldots, e_{2m-1}\}$ (all odd-indexed edges).
- M' contains $\{e_2, e_4, e_6, \ldots, e_{2m-2}\}$ (all even-indexed edges except e_{2m}).

These sets share no common edges. Therefore:

$$\begin{aligned}\|y^*(H) - \widetilde{y}^*(H')\|_1 &= |M_{\text{odd}} \setminus M'| + |M' \setminus M_{\text{odd}}| \\ &= |M_{\text{odd}}| + |M'| \\ &= m + (m-1) \\ &= 2m - 1 \\ &= \Theta(|V|).\end{aligned}$$

Despite the maximum degree being $\Delta = 2$, the instability is linear in the network size $|V| = 2m$. □

Remark 3. *This construction demonstrates that degeneracy in the unregularized optimization landscape creates instability even for sparse graphs. The issue is not merely that high-degree vertices create many dependencies, but that tie-breaking among multiple optima can cause global solution changes under local perturbations. This motivates our regularization approach, which enforces uniqueness and eliminates the need for tie-breaking entirely.*

Remark 4 (Tie-breaking alone already causes global flips)**.** *A complementary mechanism operates when the optimal face is degenerate. On an even cycle on $2m$ vertices with cap $k = 1$, the LP has two disjoint optimal perfect matchings, and the inclusion-preferring lexicographic rule flips from one to the other when a single vertex is deleted, changing all selected edges (Appendix B.1.3, Theorem 12). Corollary 2 strengthens this from one rule to every rule; we retain the even-cycle instance because tie-breaking-driven flips cleanly illustrate the degenerate-face mechanism; the necessity construction of Section 5 uses the odd cycle.*

B.2 Fractional stability and approximate optimization

Proof of Theorem 3. Write $f_H(y) = \sum_{e \in E} y_e - \frac{\lambda}{2}\|y\|_2^2$ for the objective of (3) on H, and likewise $f_{H'}$ on H'. Write $y = y^\lambda(H)$ and $y' = y^\lambda(H')$. Let E_* be the examples removed with the deleted user, let $a = y|_{E'}$, and let $z \in \mathbb{R}^E$ be the zero-extension of y'. The restriction a is feasible on H'. The extension z is feasible on H: it has zero load at the deleted user and the same loads as y' at every other user. Consequently,

$$f_H(z) = f_{H'}(y') \geq f_{H'}(a), \qquad f_H(y) \geq f_H(z).$$

The mass on removed examples is at most k, so

$$0 \leq f_H(y) - f_H(z) \leq f_H(y) - f_{H'}(a) = \sum_{e \in E_*} \left(y_e - \frac{\lambda}{2} y_e^2 \right) \leq \sum_{e \in E_*} y_e \leq k. \tag{6}$$

For any feasible x on H, strong concavity and first-order optimality at y imply

$$f_H(x) \leq f_H(y) + \langle \nabla f_H(y), x - y \rangle - \frac{\lambda}{2}\|x - y\|_2^2 \leq f_H(y) - \frac{\lambda}{2}\|x - y\|_2^2.$$

Apply this inequality with $x = z$ and use (6) to obtain

$$\|y - z\|_2^2 \leq \frac{2k}{\lambda}.$$

Finally, Cauchy–Schwarz gives $\|y - z\|_1 \leq \sqrt{n}\,\|y - z\|_2$, proving the stated bound over all coordinates, including the removed examples. □

Remark 5 (Heterogeneous budgets)**.** *Suppose each user v has an integer cap $k_v \geq 1$, with the caps of surviving users unchanged after deletion. The restriction and zero-extension remain feasible. Replacing k by the deleted user's cap in* (6) *gives*

$$\left\|y^\lambda(H) - \widetilde{y}^\lambda(H')\right\|_1 \leq \sqrt{\frac{2nk_{v^*}}{\lambda}} \leq \sqrt{\frac{2nk_{\max}}{\lambda}}.$$

This includes the hub instance of Remark 2. With the first k_v tentative examples kept at user v, Lemma 13 and Theorem 4 hold with the same constants. The unconditional retention bound in Lemma 15 also holds, and its stronger bound applies when $\min_v k_v \geq 6\ln(2r)$.

Remark 6 (Inexact solutions of the regularized program)**.** *If $\mu_\star$ bounds the movement of the exact optimizers and the solver's objective gap is at most ϵ_{opt} on every input, the corrected bound is*

$$\widehat{\mu} = \mu_\star + 2\sqrt{\frac{2N\epsilon_{\mathrm{opt}}}{\lambda}}. \tag{7}$$

Fix a public tolerance $\epsilon_{\mathrm{opt}} \geq 0$. Consider a solver that returns a feasible point $\hat{y}(H)$ satisfying

$$f_H(y^\lambda(H)) - f_H(\hat{y}(H)) \leq \epsilon_{\mathrm{opt}}$$

on every admissible input. Quadratic growth gives

$$\|\hat{y}(H) - y^\lambda(H)\|_2 \leq \sqrt{\frac{2\epsilon_{\mathrm{opt}}}{\lambda}}.$$

For neighboring inputs with at most N examples, zero-extend both vectors from the smaller input to the larger example set. The triangle inequality then gives

$$\|\hat{y}(H) - \widetilde{\hat{y}}(H')\|_1 \leq \|y^\lambda(H) - \widetilde{y}^\lambda(H')\|_1 + 2\sqrt{\frac{2N\epsilon_{\mathrm{opt}}}{\lambda}}.$$

This proves (7) *for any uniform exact-solution bound $\mu_\star$, including the size-based and structural bounds. Use the corrected bound before halving the weights in the rounding and privacy formulas. The correction is at most one when $\epsilon_{\mathrm{opt}} \leq \lambda/(8N)$.*

A feasible primal point $\hat{y}$ and a dual-feasible point α give the certificate

$$f_H(y^\lambda(H)) - f_H(\hat{y}) \leq g_H(\alpha) - f_H(\hat{y}),$$

where g_H is the dual objective in Appendix E.1. Thus a solver that returns a feasible point with certified gap at most ϵ_{opt} on every input meets the hypothesis. The certificate is evaluated at the returned feasible point, including any feasibility correction. The same public tolerance is used on every input.

The objective comparison in the proof of Lemma 15 also gives

$$\sum_e \hat{y}_e \geq f_H(\hat{y}) \geq \phi(\lambda)\,\mathrm{OPT}_{\mathrm{LP}} - \epsilon_{\mathrm{opt}}.$$

Since the rounding part of that lemma holds for every feasible fractional vector, half-weight rounding of $\hat{y}$ satisfies

$$\mathbb{E}|S| \geq \frac{[\phi(\lambda)\,\mathrm{OPT}_{\mathrm{LP}} - \epsilon_{\mathrm{opt}}]_+}{4r},$$

with denominator 4 when $k \geq 6\ln(2r)$, where $[a]_+ = \max\{a, 0\}$. Approximate optimization therefore has explicit effects on both stability and retention.

B.3 Proofs for priority-truncated sampling

Lemma 13 (Recourse of priority truncation)**.** *Fix the priorities, their tie-breaking order, and the cap. Let $F(R)$ be the output of the truncation step of Algorithm 1 on a tentative set R. If each example is attributed to at most r users, then for any $R, R' \subseteq E$,*

$$|F(R) \triangle F(R')| \leq (1+r)|R \triangle R'|.$$

Proof of Lemma 13. Fix the priorities and their tie-breaking order. First suppose $R' = R \cup \{e_0\}$ with $e_0 \notin R$. Only the lists at users in e_0 can change. At each such user, the list either stays the same, gains e_0, or gains e_0 and loses one example that previously occupied its last slot. Hence the final output can change only on e_0 and on at most one displaced example at each user in e_0. All other examples have the same membership in every relevant priority list. It follows that

$$|F(R) \triangle F(R')| \leq 1 + |e_0| \leq 1 + r.$$

The same bound holds for a deletion by exchanging R and R'. For arbitrary tentative sets, remove the examples in $R \setminus R'$ and insert those in $R' \setminus R$, one at a time. The truncation is defined on every intermediate set, so the triangle inequality gives

$$|F(R) \triangle F(R')| \leq (1+r)|R \triangle R'|.$$

□

Lemma 14 (Tentative disagreements under shared uniforms)**.** *Let $w, w' \in [0,1]^E$, and use independent $U_e \sim \mathrm{Unif}[0,1]$ to form $R = \{e : U_e < w_e\}$ and $R' = \{e : U_e < w'_e\}$. Then $D = |R \triangle R'|$ is a sum of independent Bernoulli variables with mean $\mu = \|w - w'\|_1$. For every $\eta \in (0,1)$,*

$$\Pr\left[D > \mu + \sqrt{2\mu \log(1/\eta)} + \frac{2}{3}\log(1/\eta)\right] \le \eta. \tag{8}$$

Proof. The membership indicators of example e disagree precisely when its uniform falls between w_e and w'_e, an event of probability $|w_e - w'_e|$. These events are independent across examples. Therefore $\mathbb{E}D = \mu$ and $\mathrm{Var}(D) \le \mu$. For $\mu > 0$, Bernstein's inequality gives

$$\Pr[D - \mu \ge t] \le \exp\left(-\frac{t^2}{2(\mu + t/3)}\right), \qquad t > 0.$$

Set $t = \sqrt{2\mu L} + 2L/3$ with $L = \log(1/\eta)$ to obtain (8). If $\mu = 0$, the two tentative sets agree almost surely and the claim follows directly. □

Proof of Theorem 4. The output is feasible because every selected example at user v belongs to $K(v)$, which has size at most k. This holds for any weights in $[0,1]^E$.

Use the coupling in the theorem and extend the weights on H' by zeros. Its tentative set R' contains no example incident to the deleted user. Applying the truncation on H to R' therefore gives exactly the same output as applying it on H': the deleted user's list is empty, and the lists of all surviving users are unchanged. Thus both outputs can be written using the same map F on subsets of E.

Lemma 14 bounds $D = |R \triangle R'|$ with mean $\mu = \|w - \widetilde{w}'\|_1$. For every realization of the shared priorities, Lemma 13 gives

$$|S(H) \triangle S(H')| \le (1 + r)D.$$

Combining this deterministic inequality with (8) proves (4). □

B.4 Fractional mass and rounded retention

Lemma 15 (Fractional mass and rounded retention)**.** *Let $\mathrm{OPT}_{\mathrm{LP}}$ be the optimum of the unregularized contribution-bounding relaxation. For every $\lambda > 0$,*

$$\sum_{e \in E} y_e^\lambda \ge \phi(\lambda)\,\mathrm{OPT}_{\mathrm{LP}}, \qquad \phi(\lambda) = \begin{cases} 1 - \lambda/2, & 0 < \lambda \le 1, \\ 1/(2\lambda), & \lambda > 1. \end{cases}$$

More generally, for any feasible fractional vector y, Algorithm 1 with weights $w = y/2$ satisfies

$$\mathbb{E}|S| \ge \frac{1}{4r}\sum_e y_e \quad \textit{for every } k \ge 1, \qquad \mathbb{E}|S| \ge \frac{1}{4}\sum_e y_e \quad \textit{if } k \ge 6\ln(2r).$$

Proof of Lemma 15. **Fractional mass.** Let $y^\star$ be an optimizer of the unregularized LP and set $a = \min\{1, 1/\lambda\}$. The vector $ay^\star$ is feasible. Since $\|y^\star\|_2^2 \le \sum_e y_e^\star = \mathrm{OPT}_{\mathrm{LP}}$, optimality of y^λ gives

$$f_H(y^\lambda) \ge f_H(ay^\star) \ge \left(a - \frac{\lambda a^2}{2}\right)\mathrm{OPT}_{\mathrm{LP}} = \phi(\lambda)\,\mathrm{OPT}_{\mathrm{LP}}.$$

The claimed mass bound follows from $\sum_e y_e^\lambda \ge f_H(y^\lambda)$.

Retention when $k \ge 6\ln(2r)$. Let y be any feasible fractional vector and sample with weights $w = y/2$. Fix an example e with $y_e > 0$, and condition on its inclusion in the tentative set R. At each user $v \in e$, the number of other tentative examples is

$$X_v = \sum_{h \ni v,\, h \ne e} \mathbf{1}\{U_h < y_h/2\}, \qquad \mathbb{E}X_v \le k/2.$$

The conditioning does not change these independent sampling indicators. If $X_v < k$, example e is in the first k positions at v regardless of its priority. The Chernoff bound gives

$$\Pr[X_v \ge k] \le (e/4)^{k/2} = e^{-ck}, \qquad c = \frac{2\ln 2 - 1}{2} > \frac{1}{6}.$$

Taking a union bound over at most r users,

$$\Pr[e \in S \mid e \in R] \geq 1 - re^{-ck} \geq \frac{1}{2} \quad \text{when } k \geq 6\ln(2r).$$

Since $\Pr[e \in R] = y_e/2$, we obtain $\Pr[e \in S] \geq y_e/4$. Summing over examples proves the stronger retention bound.

Retention for every $k \geq 1$. Condition on $e \in R$ and on $P_e = p$. At user $v \in e$, let $N_v(p)$ count the other sampled examples with priority less than p. Independence of sampling and priorities gives

$$\mathbb{E}N_v(p) = \frac{p}{2} \sum_{h \ni v,\, h \neq e} y_h \leq \frac{pk}{2}.$$

Markov's inequality yields $\Pr[N_v(p) \geq k] \leq p/2$. A union bound shows that e survives at every one of its users with probability at least $1 - rp/2$, which is at least $1/2$ for $0 \leq p \leq 1/r$. Integrating over this interval gives

$$\Pr[e \in S] \geq \frac{y_e}{2} \int_0^{1/r} \frac{1}{2}\, dp = \frac{y_e}{4r}.$$

Examples with zero weight contribute zero to either bound. Summing proves the unconditional guarantee. □

Remark 7 (Choice of regularization strength)**.** *The fractional stability bound decreases as $\lambda^{-1/2}$. For $0 < \lambda \leq 1$, the retained fractional mass is at least half of $\mathrm{OPT}_{\mathrm{LP}}$, and hence at least one eighth of $\mathrm{OPT}_{\mathrm{LP}}$ remains in expectation after rounding when $k \geq 6\ln(2r)$. For $\lambda > 1$, the mass guarantee decreases as $1/\lambda$. These bounds support a constant choice of λ when balancing stability and retention. The empirical calibration ratio $t_N/\mathbb{E}|S|$ is evaluated across values of λ in Appendix E.*

B.5 The response norm: proofs and examples

Throughout this subsection, $\lambda > 0$ is fixed. For a structure $G = (V, E)$ and caps $c \in \mathbb{R}^V_{\geq 0}$, write

$$P_G(c): \quad \max_{y \in [0,1]^E} \left\{ f_G(y) := \sum_{e \in E} y_e - \frac{\lambda}{2}\|y\|_2^2 : \sum_{e \ni u} y_e \leq c_u \text{ for every } u \in V \right\}.$$

The feasible set is nonempty and compact, and the objective is strictly concave, so its optimizer $y_G(c)$ is unique. The sets U and T in (5) always denote the fractional examples and all saturated users, respectively. Saturated users with no fractional incident example give zero columns of $B_{U,T}$ and zero columns of $(B_{U,T}^\top)^+$.

B.5.1 Cap perturbations and the solution path

Lemma 16 (Node deletion as a cap perturbation)**.** *Let $H' = H \setminus \{v^*\}$, with surviving examples E', and define*

$$b_u := \sum_{\substack{e \in E_{v^*} \\ u \in e}} y_e^\lambda(H), \qquad u \neq v^*.$$

Then $0 \leq b \leq k\mathbf{1}$,

$$y^\lambda(H)|_{E'} = y_{H'}(k\mathbf{1} - b), \qquad \|b\|_1 = \sum_{e \in E_{v^*}} (|e| - 1) y_e^\lambda(H) \leq (r-1)k.$$

Proof of Lemma 16. Fix the weights on the deleted examples E_{v^*} at their values in $y^\lambda(H)$. The constraint at v^* then depends only on these fixed weights. At each remaining user u, they consume b_u units of capacity, leaving cap $k - b_u$ for the surviving examples. In particular, $0 \leq b_u \leq k$. The objective is separable across examples, so optimizing over the remaining coordinates is exactly $P_{H'}(k\mathbf{1} - b)$. If the restriction of $y^\lambda(H)$ were not optimal for this program, replacing that restriction by a better feasible solution would improve the objective on H. Uniqueness therefore gives

$$y^\lambda(H)|_{E'} = y_{H'}(k\mathbf{1} - b).$$

Summing the released load counts each deleted example once for each of its users other than v^*:

$$\|b\|_1 = \sum_{u \neq v^*} \sum_{\substack{e \in E_{v^*} \\ u \in e}} y_e^\lambda(H) = \sum_{e \in E_{v^*}} (|e| - 1) y_e^\lambda(H) \leq (r-1) \sum_{e \in E_{v^*}} y_e^\lambda(H) \leq (r-1)k.$$

A node addition is the same pair read in the opposite direction, with b defined from the optimizer on the larger structure. □

Proof of Theorem 5. Continuity and affine pieces. The piecewise affine form is a property of strictly convex parametric quadratic programs; see, for example, Bemporad et al. (2002). We give the argument here to include dependent constraints and zero caps.

Minimize $-f_G$ and include the box constraints among the linear inequalities. The Karush–Kuhn–Tucker conditions are necessary and sufficient, since the feasible set is polyhedral and the objective is convex. Choose a subset of inequalities to be binding and set the multipliers of the remaining inequalities to zero. For this choice, primal feasibility, nonnegative multipliers, stationarity, and the chosen equalities form a closed polyhedron in (t, y, α), where $c(t) = k\mathbf{1} - tb$ and α here includes all inequality multipliers. Its projection onto (t, y) is a closed polyhedron. Every point in this projection satisfies the optimality conditions, so uniqueness makes it the graph of a single-valued function. Convexity of the graph implies that this function is affine on its parameter interval.

There are finitely many choices of binding inequalities, and their projected graphs cover the solution path. They agree wherever their parameter intervals overlap. These intervals are closed, so the finitely many affine functions join continuously, including at $t = 0$ and $t = 1$. Subdivide further at any point where a coordinate reaches 0 or 1, or a cap slack reaches zero. On every resulting open interval the sets U and T are constant and $y(t)$ is affine. This argument does not require unique multipliers or linearly independent saturated constraints.

Response on an open interval. Fix such an interval and set $B = B_{U,T}$. Let $\alpha(t) \geq 0$ now denote optimal multipliers of the cap constraints. Complementary slackness gives $\alpha_u(t) = 0$ for every $u \notin T$. On fractional coordinates, the box multipliers vanish, and stationarity reads

$$\lambda y_U(t) = \mathbf{1} - B\alpha_T(t).$$

Taking differences between two points of the interval shows that $\dot{y}_U$ lies in $\mathrm{range}(B)$, even if the multipliers themselves are not uniquely determined. Every cap in T is an equality throughout the interval, and the coordinates outside U are constant at 0 or 1. Differentiating these equalities gives

$$B^\top \dot{y}_U = -b_T.$$

There is exactly one vector in $\mathrm{range}(B)$ with this image under $B^\top$: if two such vectors existed, their difference would lie in both $\mathrm{range}(B)$ and $\ker(B^\top)$, which are orthogonal. The Moore–Penrose pseudoinverse therefore gives

$$\dot{y}_U = -(B^\top)^+ b_T, \qquad \dot{y}_e = 0 \quad (e \notin U). \tag{9}$$

In particular, b_T belongs to $\mathrm{range}(B^\top)$ on this interval. Although λ does not appear explicitly in (9), it determines the optimizer and hence the sets U and T.

Integration along the path. On each open interval,

$$\|\dot{y}(t)\|_1 = \left\|(B^\top)^+ b_T\right\|_1 \leq \rho(G, c(t))\, \|b_T\|_1 \leq \rho(G, c(t))\, \|b\|_1.$$

The finitely many breakpoints do not affect the integral. Continuity and piecewise affinity yield

$$\begin{aligned} \|y(1) - y(0)\|_1 &\leq \int_0^1 \|\dot{y}(t)\|_1 \, dt \\ &\leq \int_0^1 \rho(G, c(t))\, \|b_{T(t)}\|_1 \, dt \\ &\leq \|b\|_1 \sup_{t \in [0,1]} \rho(G, c(t)). \end{aligned}$$

The definition of ρ at a breakpoint uses the saturated and fractional sets at that point; no derivative is asserted there. □

B.5.2 From the response bound to node privacy

Proof of Corollary 6. Consider a deletion pair $H, H' \in \mathcal{H}$. Lemma 16 and Theorem 5, applied on $G = H'$, give

$$\|y^\lambda(H)|_{E'} - y^\lambda(H')\|_1 \le \|b\|_1 \bar{\rho} \le (r-1)k\bar{\rho}.$$

The mass on the deleted examples is at most k. After extending the smaller solution by zero on these examples, the full distance is therefore at most $\mu_\rho = k + (r-1)k\bar{\rho}$. The same bound covers additions by symmetry.

For half-weight rounding, the number of tentative disagreements under shared uniforms is a sum of independent indicators with mean at most $\mu_\rho/2$. Bernstein's inequality and Lemma 13 give

$$\Pr\left[|S(H) \triangle S(H')| > t_\rho\right] \le \eta.$$

Apply Theorem 8 to the fixed downstream learner. The resulting parameters satisfy

$$\varepsilon_{\text{node}} = t_\rho \varepsilon_0 = \varepsilon, \qquad \delta_{\text{node}} \le \eta + t_\rho e^{(t_\rho - 1)\varepsilon_0}\delta_0 \le \frac{\delta}{2} + \frac{\delta}{2} = \delta.$$

Finally, $\mu_\rho \le rk\bar{\rho}$ because $\bar{\rho} \ge 1$. Using $\sqrt{ab} \le (a+b)/2$ in the expression for t_ρ gives $t_\rho = O(r^2 k\bar{\rho} + r\log(1/\eta))$. □

The input class and public calibration. The input class $\mathcal{H}$ and the constants $r, k, \lambda, \bar{\rho}$ are fixed before the private input is observed. The assumption concerns every cap path induced by an adjacent pair in this class, including the active sets between its endpoints. For a class closed under node deletion, every deletion of an admissible input is covered; additions are covered whenever the enlarged input also belongs to the class. Both classes in Proposition 7 are closed under node deletion. A measured norm on one input or a finite set of sampled paths can serve as a diagnostic, while a uniform class bound supplies the assumption used in the privacy theorem.

Combining bounds and using approximate optimizers. If a public size bound N is also available, both fractional stability bounds apply to the same selector. One may replace μ_ρ by $\min\{\mu_\rho, \mu_N\}$ in the rounding formula, where μ_N is defined in Corollary 10. For feasible approximate optimizers with a common certified objective-gap tolerance ϵ_{opt} on every input, Remark 6 gives the valid replacement

$$\widehat{\mu} = \min\{\mu_\rho, \mu_N\} + 2\sqrt{2N\epsilon_{\text{opt}}/\lambda}.$$

In particular, the public choice $\epsilon_{\text{opt}} \le \lambda/(8N)$ adds at most one to the fractional bound. This transfers the structural calibration to an implementation that satisfies the stated tolerance.

B.5.3 Response norms of example classes

Proof of Proposition 7. (i) At most one heavy user per example. Consider a deletion pair with larger structure H, surviving structure $G = H \setminus \{v^*\}$, and caps $c(t) = k\mathbf{1} - tb$. Let $d^*_u = |\{e \in E_{v^*} : u \in e\}|$ count the deleted examples containing u. Since every weight is at most one, $b_u \le d^*_u$. If u is not heavy in H, then

$$\sum_{\substack{e \in E(G) \\ u \in e}} y_e(t) \le \deg_G(u) = \deg_H(u) - d^*_u \le k - d^*_u \le k - b_u \le c_u(t).$$

The first inequality is strict whenever u belongs to a fractional example. Thus a saturated user with a nonzero column in $B_{U,T}$ must be heavy in H. By assumption, each row of $B = B_{U,T}$ contains at most one nonzero entry. Its nonzero columns therefore have disjoint supports, so $D = B^\top B$ is diagonal and

$$(B^\top)^+ = B(B^\top B)^+ = BD^+.$$

For a column with $d^U_u > 0$ entries, this operator puts weight $1/d^U_u$ on each of those entries, giving column ℓ_1 norm one. Zero columns contribute zero. Hence $\rho \le 1$ throughout the path, and the full fractional movement is at most $k + (r-1)k = rk$.

(ii) Disjoint simple cliques. By symmetry and uniqueness, the optimizer has the same weight on every edge within each clique. Deleting a vertex from a clique K_m releases the same load at each of

the remaining $m-1$ vertices. Consequently, along the cap path every component is a clique with a uniform cap and a common edge weight. A component contributes a nonzero block to $B_{U,T}$ only when its edges are fractional and its vertices are saturated. In that case all its vertices belong to T, and the block is the full edge–vertex incidence matrix B_p of a clique K_p.

For $p = 2$, $B_2 = (1\ \ 1)$ and $(B_2^\top)^+ = (\frac{1}{2}\ \ \frac{1}{2})$, whose column norms are $1/2$. For $p \geq 3$,

$$B_p^\top B_p = (p-2)I + J, \qquad (B_p^\top B_p)^{-1} = \frac{1}{p-2}\left(I - \frac{J}{2(p-1)}\right),$$

where J is the all-ones matrix. The column indexed by u of $(B_p^\top)^+ = B_p(B_p^\top B_p)^{-1}$ has value

$$\begin{cases} 1/(p-1), & \text{on each edge containing } u, \\ -1/((p-1)(p-2)), & \text{on each other edge.} \end{cases}$$

Its ℓ_1 norm is

$$(p-1)\frac{1}{p-1} + \frac{(p-1)(p-2)}{2}\frac{1}{(p-1)(p-2)} = \frac{3}{2}.$$

The pseudoinverse respects the decomposition into components, so the norm of the full operator is the largest block norm. Therefore $\rho \leq 3/2$ on every cap path. Since $r = 2$, the full fractional movement is at most $k + (3/2)k = 5k/2$. □

Exact movement for a clique. Suppose $\lambda \leq 1$ and a user is deleted from K_m with $m \geq k+2$. The edge weight before deletion is $k/(m-1)$ and the released load at each surviving user is $b_u = k/(m-1)$. On the surviving clique K_{m-1}, the cap path has edge weight

$$y_e(t) = \frac{k - tk/(m-1)}{m-2}, \qquad t \in [0,1].$$

For $m \geq \max\{4, k+2\}$ and $t > 0$, this component has at least three vertices, every edge is fractional, and its response norm is exactly $3/2$. In the boundary case $m = 3$, $k = 1$, the surviving component is a single edge and its response norm is $1/2$ for $t > 0$. In all cases $m \geq k+2$, the surviving movement is

$$\binom{m-1}{2}\left(\frac{k}{m-2} - \frac{k}{m-1}\right) = \frac{k}{2}.$$

This also illustrates the distinction between the operator norm, which covers all cap directions, and the movement along the particular direction produced by a node deletion.

Odd cycles. On C_n with n odd, $k = 1$, and $\lambda \in (0,2)$, the point $y_e = 1/2$ on every edge satisfies the optimality conditions with cap multipliers $\alpha_u = (2-\lambda)/4$. Thus all edges are fractional, all users are saturated, and B is the full edge–vertex incidence matrix of the cycle. To solve $B^\top z = e_u$, start at u and propagate the equations around the cycle. At each vertex other than u, the two incident coordinates sum to zero, so their signs alternate. Since the cycle is odd, the two coordinates at u have the same sign and each equals $1/2$. Hence every coordinate of the unique solution has absolute value $1/2$, and

$$\rho(C_n, \mathbf{1}) = \left\|(B^\top)^{-1}\right\|_{1\to 1} = \frac{n}{2}.$$

This is a response to an infinitesimal cap change at the stated solution. The movement over a full deletion path also depends on how long each saturated pattern persists, as expressed by the integral in the proof of Theorem 5.

B.6 Alternative regularizers and the sample size tradeoff

The ℓ_2 penalty in (3) is analytically convenient because it makes the objective strongly concave and yields a unique optimizer with quantitative stability guarantees under node deletion. The downside is that this penalty can shrink the total fractional mass $\sum_e y_e$ when many constraints are slack. In the extreme case where the load constraints are inactive, the objective decouples across coordinates and the maximizer satisfies $y_e = \min\{1, 1/\lambda\}$, so the total mass scales like n/λ. This motivates alternative regularizers that retain the stability benefits of strict concavity while better preserving mass on unconstrained coordinates.

Throughout, we write $A \in \{0,1\}^{|V|\times|E|}$ for the vertex–edge incidence matrix of the hypergraph, where $A_{v,e} = \mathbf{1}\{v \in e\}$. Then the contribution constraints can be written compactly as $Ay \leq k\mathbf{1}$.

A simple fix that preserves mass. A minimal modification is to regularize around the all-ones vector rather than around the origin,

$$\max_{y\in[0,1]^E}\left\{\sum_{e\in E} y_e - \frac{\lambda}{2}\|y-\mathbf{1}\|_2^2\right\} \quad \text{s.t.} \quad Ay \le k\mathbf{1}.$$

This objective is still λ-strongly concave, so the optimizer is unique and the value-gap argument of Theorem 3 applies; the resulting bound, however, is strictly weaker. The zero-extension places each deleted coordinate at distance one from the reference point, so the value gap acquires an additive $\frac{\lambda}{2}|E_{v^*}| \le \frac{\lambda\Delta}{2}$ term and the stability bound becomes

$$\left\|y^\lambda(H) - \widetilde{y}^\lambda(H')\right\|_1 \;\le\; \sqrt{n\big(2k/\lambda + \Delta\big)} \;\le\; \sqrt{2nk/\lambda} + \sqrt{n\Delta},$$

whose second term does not decay as λ grows. The obstruction is real and λ-independent: on the triangle C_3 with uniform cap $k = 1$, the centered optimizer is $y_e = 1/2$ on all three edges for every $\lambda > 0$ (the optimizer is unique by strict concavity, hence inherits the triangle's edge-transitive symmetry, and along the symmetric direction the objective increases up to the cap), while deleting any vertex leaves a single edge whose centered optimizer is 1; the zero-extended ℓ_1 change is therefore exactly $\frac{1}{2}+\frac{1}{2}+\frac{1}{2} = \frac{3}{2}$ at every λ, whereas the origin-centered penalty on the same pair gives $2/\lambda \to 0$ for $\lambda \ge 2$. The mechanism is simple: node deletion zero-fills coordinates, a small perturbation from the origin's viewpoint and a maximal one from the all-ones viewpoint. Unlike (3), unconstrained coordinates are pushed toward $y_e \approx 1$ rather than $y_e \approx 1/\lambda$, so when the hypergraph admits a large feasible subset the centered penalty preserves $\sum_e y_e = \Theta(n)$; but at the constant-λ operating point of our analysis the origin-centered penalty already retains $\Omega(\mathrm{OPT}_{\mathrm{LP}})$ mass (Lemma 15), so the centered variant may improve retention constants while lacking a worst-case stability guarantee that improves with λ.

Entropy regularization. Another natural choice is an entropy-type barrier that keeps coordinates in the interior and introduces strict concavity:

$$\max_{y\in(0,1)^E}\left\{\sum_e y_e - \frac{1}{\tau}\sum_e \big(y_e\log y_e + (1-y_e)\log(1-y_e)\big)\right\} \quad \text{s.t.} \quad Ay \le k\mathbf{1}.$$

The entropy term contributes curvature $-(1/\tau)\cdot(y_e(1-y_e))^{-1}$ on coordinate e, so the objective is strictly (and in fact strongly) concave on $(0,1)^E$. This again guarantees a unique optimizer and rules out tie-breaking-driven instability. Unlike ℓ_2 regularization, entropy regularization does not directly shrink all coordinates toward 0, but it does introduce its own mass tradeoff. For a completely unconstrained coordinate, the stationary point satisfies

$$1 - \frac{1}{\tau}\log\left(\frac{y_e}{1-y_e}\right) = 0, \qquad \text{so} \qquad y_e = \frac{1}{1+e^{-\tau}}.$$

Thus, larger τ pushes y_e closer to 1, while smaller τ stabilizes more aggressively but pulls y_e toward $1/2$. In other words, entropy regularization can preserve substantial mass without forcing $y_e \approx 1/\lambda$, but it still exhibits a stability–retention tradeoff through the parameter τ.

From the stability perspective the ℓ_2 template does transfer, with one important difference: under node deletion the value-gap comparison acquires the entropy of the deleted coordinates, which is bounded by $|E_{v^*}|\ln 2/\tau$ rather than by the budget k, and this term does not disappear as the regularization strengthens.

Proposition 17 (Node-adjacency stability of entropy regularization)**.** *Fix $\tau > 0$. Let $y^\tau(H)$ and $y^\tau(H')$ be the optimizers of the entropy-regularized program on H and on $H' = H \setminus \{v^*\}$, and let $\widetilde{y}^\tau(H')$ be the zero-extension of the latter. Then*

$$\left\|y^\tau(H) - \widetilde{y}^\tau(H')\right\|_1 \;\le\; \sqrt{\frac{n}{2}\big(\tau k + |E_{v^*}|\ln 2\big)} \;\le\; \sqrt{\frac{n}{2}\big(\tau k + \Delta\ln 2\big)}.$$

Proof. Write $F_H(y) = \sum_e y_e + \tau^{-1}\sum_e h(y_e)$ with $h(y) = -y\ln y - (1-y)\ln(1-y)$, so that $h(0) = 0$, $0 \le h \le \ln 2$, and $h'' = -1/(y(1-y)) \le -4$; hence F_H is $(4/\tau)$-strongly concave. Let $y = y^\tau(H)$, $y' = y^\tau(H')$, and $\widehat{y}'$ the zero-extension of y'. The point $\widehat{y}'$ is feasible for H: the

load at v^* is zero, and at every other vertex it equals the load of y' on H', which is at most k. Since $h(0) = 0$, $F_H(\widehat{y}') = F_{H'}(y')$. The restriction of y to E' is feasible for H', so optimality of y' gives

$$F_{H'}(y') \;\geq\; F_{H'}(y|_{E'}) \;=\; F_H(y) - \sum_{e \in E_{v^*}} \big(y_e + \tau^{-1} h(y_e)\big) \;\geq\; F_H(y) - k - |E_{v^*}| \ln 2/\tau,$$

using $\sum_{e \ni v^*} y_e \leq k$. Hence $F_H(y) - F_H(\widehat{y}') \leq k + |E_{v^*}| \ln 2/\tau$. The maximizer y lies in $(0,1)^E$ because h' diverges at the endpoints, so the first-order optimality condition $\langle \nabla F_H(y), z - y \rangle \leq 0$ holds for every feasible z, and strong concavity gives $F_H(y) - F_H(\widehat{y}') \geq (2/\tau)\|y - \widehat{y}'\|_2^2$. Therefore $\|y - \widehat{y}'\|_2^2 \leq \tau k/2 + |E_{v^*}| \ln 2/2$, and Cauchy–Schwarz over the n coordinates gives the claim. □

The bound does not vanish as $\tau \to 0$: strong entropy regularization drives every unconstrained coordinate toward $1/2$, so the deleted user's hyperedges retain mass whose entropy enters the value gap. On the triangle C_3 with $k = 1$ the optimizer is $y_e = 1/2$ on all three edges for every τ (by symmetry and strict concavity, as for the centered quadratic), while the single surviving edge after a deletion moves only to $1/(1 + e^{-\tau})$; the total zero-extended change $1 + (1/(1 + e^{-\tau}) - 1/2)$ is therefore dominated by the unavoidable deleted mass, and the movement on surviving edges is $O(\tau)$. Whether the $\sqrt{n\Delta}$ term of the proposition is attained by the movement on surviving edges on some family is open; what is settled is that entropy regularization admits a quantitative node-adjacency stability theorem, with the same composition consequences as the ℓ_2 penalty, but no bound that improves without limit as the regularization strengthens.

Penalizing only load violations. A different direction is to avoid shrinking unconstrained coordinates by penalizing only budget violations,

$$\max_{y \in [0,1]^E} \left\{ \sum_e y_e - \frac{\lambda}{2} \sum_{v \in V} \left(\max \left\{ 0, \sum_{e \ni v} y_e - k \right\} \right)^2 \right\}.$$

When the solution has slack with respect to the cap constraints, the penalty is inactive and the optimizer stays close to the unregularized objective, which helps preserve mass. When loads approach the cap, the penalty activates smoothly and discourages solutions that rely on nearly-saturated vertices, which reduces sensitivity to perturbations.

Unlike ℓ_2 regularization, however, this objective is not strictly concave, and the degeneracy is not hypothetical: on even cycles the optimal face is a segment and every optimizer violates the cap.

Proposition 18 (Violation penalties have nonunique, cap-violating optimizers)**.** *On the even cycle C_{2m} with hyperedges $e_i = \{i-1, i\}$ (indices modulo $2m$), uniform cap $k = 1$, and any $\lambda > 1/2$, the violation-penalty objective attains its maximum value $m\big(1 + \frac{1}{4\lambda}\big)$ exactly on the one-parameter family*

$$y_{e_i} = a \text{ for odd } i, \qquad y_{e_i} = 1 + \tfrac{1}{2\lambda} - a \text{ for even } i, \qquad a \in \big[\tfrac{1}{2\lambda},\, 1\big].$$

Hence the optimizer is not unique, and every optimizer places load $1 + \frac{1}{2\lambda} > k$ on every vertex.

Proof. For $y \in [0,1]^E$ let $s = \sum_e y_e$ and $\ell_v = y_{e_v} + y_{e_{v+1}}$ be the load at vertex v, so $\sum_v \ell_v = 2s$. By convexity of $t \mapsto (t-1)_+^2$ and Jensen's inequality, $\sum_v (\ell_v - 1)_+^2 \geq 2m\,(s/m - 1)_+^2$, so the objective is at most $g(u) := m\big[u - \lambda(u-1)_+^2\big]$ with $u = s/m \in [0,2]$. The function g is maximized at $u^* = 1 + 1/(2\lambda)$, which lies in $(1,2]$ exactly when $\lambda \geq 1/2$, with $g(u^*) = m(1 + 1/(4\lambda))$. Equality in Jensen's inequality requires every load to equal u^*, and on an even cycle the vectors with all loads equal to u^* are exactly the alternating vectors $y_{e_{i+1}} = u^* - y_{e_i}$, parametrized by $a = y_{e_1}$ with $0 \leq a \leq 1$ and $0 \leq u^* - a \leq 1$, that is, $a \in [1/(2\lambda), 1]$. Conversely every such vector attains $g(u^*)$, since its penalty is $\frac{\lambda}{2} \cdot 2m\,(u^* - 1)^2 = m/(4\lambda)$. For $\lambda > 1/2$ the interval is nondegenerate. □

Consequently the violation penalty reproduces the degenerate optimal face that drives Theorem 12, and in addition never returns a cap-feasible point on this family. A strictly concave tie-breaker and a feasibility-restoring truncation are both required, at which point the analysis reduces to the ℓ_2 case.

Dual regularization. Finally, one can regularize in the dual. Let $\alpha \geq 0$ be Lagrange multipliers for the constraints $Ay \leq k\mathbf{1}$. Adding a proximal penalty $\frac{\gamma}{2}\|\alpha\|_2^2$ to the dual objective makes the dual strongly convex and stabilizes how the shadow prices α change when a vertex and its incident edges are added or removed. Strong duality alone, however, does not transfer stability of α to stability of the primal optimizer: the primal objective is linear, its optimal face need not be a singleton, and recovering a stable primal point requires additional strict concavity in the primal or a separate argument. In fact dual regularization alone cannot help on the odd cycle: there the primal optimal face is a singleton on both sides (Theorem 1), so any procedure that returns a primal optimum of the unregularized program, however its dual is regularized, moves by $|V|/2$. Dual regularization can therefore only be useful in combination with a primal perturbation, which returns the analysis to the primal variants above.

Summary. The ℓ_2 penalty in (3) and, by Proposition 17, entropy regularization are the two variants with quantitative node-adjacency stability theorems; only the ℓ_2 bound improves without limit as the regularization strengthens, the entropy bound retaining a $\sqrt{n\Delta}$ term whose tightness is open. The centered quadratic preserves mass on unconstrained coordinates but provably loses the improving-with-λ stability, as the triangle obstruction shows. The violation penalty has a nonunique, cap-violating optimal face on even cycles (Proposition 18), and dual regularization alone cannot repair the odd cycle. Whenever a variant supplies a high-probability recourse bound, our stability-to-node-DP composition theorem applies unchanged.

C Privacy transfer and lower bounds

The composition and calibration proofs are followed by the converse results. We then record the dataset formulation of stability and the extensions to multiple deletions and fixed-graph privacy.

C.1 Composition from selection stability

Proof. Fix a node-adjacent pair $\mathcal{D}, \mathcal{D}'$ in the input domain and a measurable event F in the learner's output space. Let γ be a coupling of the selected multisets (Z, Z') with

$$\Pr_{\gamma}[d_{\mathrm{Ham}}(Z, Z') > t] \leq \eta.$$

Write $G = \{d_{\mathrm{Ham}}(Z, Z') \leq t\}$, and set

$$a_F(z) = \Pr[\mathcal{A}(z) \in F], \qquad \delta_t = \delta_0 \sum_{i=0}^{t-1} e^{i\varepsilon_0}.$$

The probability in a_F is over the learner's independent randomness. For a fixed pair $(z, z') \in G$, a sequence of at most t example insertions and deletions joins z to z'. Repeatedly applying example-level DP along this sequence gives

$$a_F(z) \leq e^{t\varepsilon_0} a_F(z') + \delta_t.$$

For $t = 0$, the multisets are identical and this inequality holds with $\delta_t = 0$.

The marginals of γ are the actual selected-data distributions. Therefore

$$\begin{aligned}
\Pr[(\mathcal{A} \circ \mathcal{M})(\mathcal{D}) \in F] &= \mathbb{E}_{\gamma}[a_F(Z)] \\
&\leq \mathbb{E}_{\gamma}[\mathbf{1}_G a_F(Z)] + \Pr_{\gamma}(G^c) \\
&\leq e^{t\varepsilon_0} \mathbb{E}_{\gamma}[\mathbf{1}_G a_F(Z')] + \delta_t \Pr_{\gamma}(G) + \eta \\
&\leq e^{t\varepsilon_0} \Pr[(\mathcal{A} \circ \mathcal{M})(\mathcal{D}') \in F] + \delta_t + \eta.
\end{aligned}$$

Interchanging the two inputs gives the reverse inequality. The coupling may be chosen separately for each adjacent pair; it is used only in the analysis. Finally, $\delta_t \leq te^{(t-1)\varepsilon_0}\delta_0$, which gives the stated bound. □

A bound on expected recourse can also give a tail bound, for example through Markov's inequality, but direct high-probability control permits tighter calibration.

C.2 Calibration of the regularized pipeline

Proof. Fix a node-adjacent pair in the stated class and orient it so that H' is obtained by deleting a vertex from H. Let $n = |E(H)| \le N$. By Theorem 3, the exact regularized optimizers satisfy

$$\|y^\lambda(H) - \widetilde{y}^\lambda(H')\|_1 \le \sqrt{2nk/\lambda} \le \mu_N,$$

where the smaller vector is extended by zero on deleted examples.

Couple the rounding runs by using the same uniforms and priorities for shared examples. With half weights, the tentative disagreement count D is a sum of independent indicators with mean

$$\nu = \tfrac{1}{2}\|y^\lambda(H) - \widetilde{y}^\lambda(H')\|_1 \le \mu_N/2.$$

The Bernstein bound of Lemma 14 gives, with probability at least $1 - \eta$,

$$D \le \frac{\mu_N}{2} + \sqrt{\mu_N \log(1/\eta)} + \frac{2}{3}\log(1/\eta).$$

Here we use that the displayed tail bound is increasing in the upper bound on ν. Lemma 13 then gives

$$|S(H) \triangle S(H')| \le (1 + r)D \le t_N.$$

The selected contents of shared examples agree, so their multiset distance is at most this symmetric difference. Theorem 8 yields

$$\varepsilon = t_N \varepsilon_0, \qquad \delta = \eta + \delta_0 \sum_{i=0}^{t_N - 1} e^{i\varepsilon_0}.$$

Both $k + \sqrt{2nk/\lambda}$ and $\sqrt{2nk/\lambda}$ are increasing in n, so the argument also applies with the sharper choice $\mu_N = \sqrt{2Nk/\lambda}$. Since both inputs have at most N examples, the same public upper bound μ_N and hence the same t_N apply to every adjacent pair in the class. At constant $\lambda > 0$ and $1 \le k \le N$, we have $\mu_N = O(\sqrt{Nk})$. Applying $\sqrt{ab} \le (a + b)/2$ to the square-root term gives

$$t_N = O\big(r(\sqrt{Nk} + \log(1/\eta))\big).$$

For the target-budget choice in the main text,

$$\eta + t_N e^{(t_N - 1)\varepsilon_0}\delta_0 \le \delta/2 + t_N e^{\varepsilon}\frac{\delta e^{-\varepsilon}}{2t_N} = \delta.$$

□

C.3 Privacy lower bounds and deterministic characterization

Overview. For randomized selection, we prove a converse when one input has a fixed selected set and its neighbor produces a set at distance at least T with probability at least η. For suitable contents and an example-private learner, node privacy then requires $\varepsilon \ge T\varepsilon_0 + \log(\eta - 2\delta)$ when $\delta < \eta/2$ (Lemma 19). The unique matching left by a vertex deletion on the odd cycle provides this fixed output for every exact selector.

On the odd cycle C_{2m+1} with cap $k = 1$ and $m \ge 2$, for every exact selector and every $\varepsilon_0 > 0$, there are contents and a pure ε_0-example-private learner such that node privacy with $\delta \le 1/8$ requires $\varepsilon \ge m\varepsilon_0 - \log 4$ (Theorem 20). Thus a unit contribution cap and degree two can coexist with a privacy factor linear in dataset size.

The lower bound extends to a concrete accuracy requirement. For a lexicographic exact selector, we construct a classification task on which every learner matching an example-private benchmark's accuracy must incur a node-privacy cost proportional to the number of selection changes (Theorem 21). The task and proof follow.

Lemma 19 (Privacy loss from separated selections)**.** *Let $\mathcal{M}$ be a possibly randomized structure-only selector. Suppose a node-adjacent pair H, H' satisfies $\mathcal{M}(H') = S'$ almost surely and*

$$\Pr[|\mathcal{M}(H) \triangle S'| \ge T] \ge \eta$$

for an integer $T \geq 1$ and $\eta \in (0, 1]$. Assume that $\mathcal{X}$ contains two distinguishable values. For every $\varepsilon_0 > 0$, there are contents agreeing on all shared examples and an $(\varepsilon_0, 0)$-example-DP mechanism $\mathcal{A}$ for which any (ε, δ)-node-DP guarantee for the composition must satisfy

$$e^{\varepsilon - T\varepsilon_0} + 2\delta \geq \eta.$$

In particular, for $\delta < \eta/2$,

$$\varepsilon \geq T\varepsilon_0 + \log(\eta - 2\delta).$$

Proof of Lemma 19. Fix a bounded statistic $\phi : \mathcal{X} \to [-1, 1]$ with $\phi(x_+) = 1$ and $\phi(x_-) = -1$, and assign contents

$$x_e = \begin{cases} x_-, & e \in S', \\ x_+, & e \notin S', \end{cases}$$

to every hyperedge of either instance. Since $\mathcal{M}$ selects as a function of the structure alone, this assignment does not affect its selected sets, and shared hyperedges receive the same content in both datasets, as node adjacency requires.

For a selected multiset Z define $q(Z) = \sum_{z \in Z} \phi(z)$. On the planted contents, writing S for the selected set, $q(S) = |S \setminus S'| - |S \cap S'|$ and $q(S') = -|S'|$, so

$$q(S) - q(S') \;=\; |S \setminus S'| - |S \cap S'| + |S'| \;=\; |S \setminus S'| + |S' \setminus S| \;=\; |S \triangle S'|.$$

Inserting or deleting one example changes q by at most $\sup |\phi| = 1$ on every dataset, so the Laplace mechanism $\mathcal{A}(Z) = q(Z) + \mathrm{Lap}(1/\varepsilon_0)$ is $(\varepsilon_0, 0)$-example-DP under add/remove adjacency, with no restriction to the planted instances.

Consider the event $\mathcal{E} = [q(S') + T, \infty)$. On H' the composed output is $q(S') + \mathrm{Lap}(1/\varepsilon_0)$ almost surely, so $\Pr[\mathcal{A} \circ \mathcal{M}(H') \in \mathcal{E}] = \frac{1}{2} e^{-T\varepsilon_0}$. On H, conditionally on the event $|\mathcal{M}(H) \triangle S'| \geq T$, which has probability at least η, the center of the Laplace distribution is $q(\mathcal{M}(H)) \geq q(S') + T$, so the output lands in $\mathcal{E}$ with conditional probability at least $\frac{1}{2}$; hence $\Pr[\mathcal{A} \circ \mathcal{M}(H) \in \mathcal{E}] \geq \eta/2$. If $\mathcal{A} \circ \mathcal{M}$ is (ε, δ)-node-DP, then

$$\frac{\eta}{2} \;\leq\; e^{\varepsilon} \cdot \frac{1}{2} e^{-T\varepsilon_0} + \delta,$$

which rearranges to $e^{\varepsilon - T\varepsilon_0} + 2\delta \geq \eta$, and for $\delta < \eta/2$ to $\varepsilon \geq T\varepsilon_0 + \log(\eta - 2\delta)$. For deterministic $\mathcal{M}$, take $\eta = 1$ and $T = |S(H) \triangle S(H')|$. □

The equivalence identifies the smallest privacy factor that works for every example-private learner and every content assignment. The argument also gives the approximate-privacy lower bound $\varepsilon \geq T\varepsilon_0 + \log(1 - 2\delta)$ for a neighboring pair with recourse T and $\delta < 1/2$.

Proof of Theorem 9. If every selected-set distance is at most t, the selected-content distance is also at most t. Apply Theorem 8 with $\eta = 0$ and $\delta_0 = 0$ to obtain $(t\varepsilon_0, 0)$-node-DP for every $(\varepsilon_0, 0)$-example-DP learner.

Conversely, suppose some adjacent pair has selected-set distance $T > t$. Determinism gives the premise of Lemma 19 with $\eta = 1$. Because the domain contains all content assignments on its admissible structures, the lemma's planted contents are admissible. For the constructed $(\varepsilon_0, 0)$-example-DP mechanism, a pure node-DP guarantee must satisfy $\varepsilon \geq T\varepsilon_0$. Since $\varepsilon_0 > 0$, this excludes $(t\varepsilon_0, 0)$-node-DP. The argument also covers $t = 0$. □

Theorem 20 (Privacy loss for exact selection on the odd cycle)**.** *Let $H = C_{2m+1}$, with cap $k = 1$, $m \geq 2$, and cyclic edge order $e_1, \ldots, e_{2m+1}$. Fix $\varepsilon_0 > 0$ and a content domain with two distinguishable values. For each selector below, there are a vertex deletion, contents, and an $(\varepsilon_0, 0)$-example-DP mechanism such that every (ε, δ)-node-DP guarantee for their composition satisfies the stated bound.*

(i) For the lexicographically first maximum-matching rule, with inclusion of earlier edges preferred, delete the vertex shared by e_1 and e_{2m+1}. The recourse is $2m$, and

$$\varepsilon \geq 2m\varepsilon_0 + \log(1 - 2\delta), \qquad \delta < 1/2.$$

(ii) For every rule that returns a maximum matching, deterministic or randomized, one of two adjacent vertex deletions gives

$$\varepsilon \geq m\varepsilon_0 + \log(1/2 - 2\delta), \qquad \delta < 1/4.$$

Proof of Theorem 20. Label the vertices $0, \ldots, 2m$ and edges $e_i = \{i-1, i\}$ cyclically with the order $e_1 \prec \cdots \prec e_{2m+1}$.

(i) The lexicographically first maximum-matching rule includes e_1, then every subsequent edge whose inclusion extends to a maximum matching, yielding $R = \{e_1, e_3, \ldots, e_{2m-1}\}$ on H. Deleting vertex 0, shared by e_1 and e_{2m+1}, leaves the path on $1, \ldots, 2m$, whose unique perfect matching is $A = \{e_2, e_4, \ldots, e_{2m}\}$ (Theorem 1); uniqueness makes the rule's output on the deleted instance deterministic regardless of its tie-breaking details. The index sets of R and A are disjoint, so $|R \triangle A| = 2m$, and Lemma 19 with $\eta = 1$, $T = 2m$ gives $\varepsilon \geq 2m\varepsilon_0 + \log(1 - 2\delta)$ for $\delta < 1/2$.

(ii) Let A and B be the unique perfect matchings of the paths obtained by deleting vertex 0 or vertex 1; as computed in the proof of Corollary 2, $A = \{e_2, e_4, \ldots, e_{2m}\}$ and $B = \{e_3, e_5, \ldots, e_{2m+1}\}$ are disjoint, and every maximum matching R of H has $|R| = m$ and $|R \triangle A| + |R \triangle B| = 4m - 2(|R \cap A| + |R \cap B|) \geq 2m$. This inequality holds pointwise over the randomness of any rule $\mathcal{M}$ required to output a maximum matching, so the events $\{|R \triangle A| \geq m\}$ and $\{|R \triangle B| \geq m\}$ cover the probability space and

$$\Pr\left[|R \triangle A| \geq m\right] + \Pr\left[|R \triangle B| \geq m\right] \;\geq\; 1.$$

For at least one of the two deletions, then, $\Pr[|\mathcal{M}(H) \triangle S'| \geq m] \geq \frac{1}{2}$, while $\mathcal{M}$ on that deleted instance outputs its unique maximum matching almost surely. Lemma 19 with $T = m$ and $\eta = \frac{1}{2}$ gives $\varepsilon \geq m\varepsilon_0 + \log(\frac{1}{2} - 2\delta)$ for $\delta < 1/4$. In particular $\delta \leq 1/8$ gives $\varepsilon \geq m\varepsilon_0 - \log 4$, which is $\Omega(m\varepsilon_0)$ whenever $m\varepsilon_0 \geq 2 \log 4$. □

Theorem 21 (Privacy lower bound under an accuracy requirement)**.** *Let $H = C_{2m+1}$ with cap $k = 1$, and let $\mathcal{M}$ be the lexicographic selector of Theorem 20. Deleting the vertex shared by e_1 and e_{2m+1} changes its selection from a matching R of size m to a disjoint matching A of size m.*

Fix $K \geq 2$ labels in addition to a null label 0. For $j \in [K]$, define $\mathcal{D}_j$ by assigning label j to the examples in R and label 0 to all other examples. Let $\mathcal{F} = \{h_0, \ldots, h_K\}$ be the constant classifiers, and let P_j put unit mass on label j. Under zero-one loss, $L_j(h_a) = \mathbf{1}[a \neq j]$. Identify the learner's output h_a with its label a.

Suppose $\mathcal{A}$ is a function of the selected multiset alone and satisfies the accuracy requirement

$$\mathbb{E}[L_j(\mathcal{A}(j^m))] \leq \tfrac{1}{4}, \quad \textit{equivalently} \quad \Pr[\mathcal{A}(j^m) = j] \geq \tfrac{3}{4}, \qquad j \in [K].$$

If $\mathcal{A} \circ \mathcal{M}$ is (ε, δ)-node-DP and $\delta < 3/4$, then

$$\varepsilon \geq \log K + \log(3/4 - \delta).$$

For $\varepsilon_0 > 0$ and $K = \lfloor e^{m\varepsilon_0}/3 \rfloor \geq 2$, the counting-score exponential mechanism is $(\varepsilon_0, 0)$-example-DP and meets this accuracy requirement. Consequently, if $m\varepsilon_0 \geq \log 6$ and $\delta \leq 1/4$, every learner meeting the same requirement must satisfy

$$\varepsilon \geq m\varepsilon_0 - \log 12 = \frac{T\varepsilon_0}{2} - \log 12, \qquad T = 2m.$$

Proof of Theorem 21. Fix $j \in [K]$ and let $\mathcal{D}_j$ and $\mathcal{D}'_j$ be the node-adjacent pair in which the vertex shared by e_1 and e_{2m+1} is deleted. The surviving hyperedges keep their contents, so the pair is node-adjacent in the sense of Definition 2. On the intact instance $\mathcal{M}$ selects R, whose contents are j^m; on the deleted instance $\mathcal{M}$ selects the unique perfect matching A of the remaining path, whose contents are 0^m under every $\mathcal{D}_j$, because A is disjoint from R and all examples outside R carry label 0. Since the downstream learner sees only the selected multiset, the composed output distribution on the deleted side is one distribution Q, identical for all j.

Let $E_j = \{\mathcal{A} = j\}$. By hypothesis $\Pr[\mathcal{A} \circ \mathcal{M}(\mathcal{D}_j) \in E_j] \geq \frac{3}{4}$, so node-DP across the adjacent pair gives

$$\tfrac{3}{4} \;\leq\; e^{\varepsilon}\, Q(E_j) + \delta \qquad \text{for every } j \in [K].$$

The events $E_1, \dots, E_K$ are disjoint, so summing,

$$1 \;\geq\; \sum_{j=1}^{K} Q(E_j) \;\geq\; K e^{-\varepsilon}\Big(\tfrac{3}{4} - \delta\Big),$$

which rearranges to $\varepsilon \geq \log K + \log(\frac{3}{4} - \delta)$.

For the benchmark, consider the exponential mechanism over the label set $\{0, 1, \dots, K\}$ with counting scores $c_a(Z) = |\{z \in Z : z = a\}|$:

$$\Pr[\mathcal{A}_{\exp}(Z) = a] \;=\; \frac{\exp(\varepsilon_0\, c_a(Z))}{\sum_b \exp(\varepsilon_0\, c_b(Z))}.$$

Adding one example labeled a^* raises c_{a^*} by one and leaves every other count unchanged. Writing $W = \sum_b \exp(\varepsilon_0 c_b)$ and W' for the normalizer after the addition, $W' = W + (e^{\varepsilon_0} - 1)\exp(\varepsilon_0 c_{a^*})$, so $W \leq W' \leq e^{\varepsilon_0} W$. For any outcome a the probability ratio is $\exp(\varepsilon_0(c'_a - c_a)) \cdot W/W'$, whose first factor lies in $\{1, e^{\varepsilon_0}\}$ and whose second lies in $[e^{-\varepsilon_0}, 1]$; hence the ratio lies in $[e^{-\varepsilon_0}, e^{\varepsilon_0}]$, and removal is symmetric. The mechanism is therefore $(\varepsilon_0, 0)$-example-DP under add/remove adjacency, with no factor two, because the counting scores are monotone. On input j^m the counts are $c_j = m$ and $c_a = 0$ otherwise, so

$$\Pr[\mathcal{A}_{\exp}(j^m) = j] \;=\; \frac{e^{m\varepsilon_0}}{e^{m\varepsilon_0} + K},$$

which is at least $3/4$ when $K \leq e^{m\varepsilon_0}/3$. For $m\varepsilon_0 \geq \log 6$ the floor loses at most a factor of two, $K = \lfloor e^{m\varepsilon_0}/3 \rfloor \geq e^{m\varepsilon_0}/6$, so any learner meeting this benchmark accuracy satisfies, for $\delta \leq 1/4$,

$$\varepsilon \;\geq\; \log K + \log \tfrac{1}{2} \;\geq\; m\varepsilon_0 - \log 6 - \log 2 \;=\; m\varepsilon_0 - \log 12. \qquad \square$$

C.4 End-to-end node-DP from dataset stability

We use the following notion of stability for the *actual dataset* fed to the learner.

Definition 6 (High-probability dataset stability)**.** *A preprocessing map $\mathcal{M}$ is (t, η)-stable under node adjacency on its input domain if, for every adjacent pair $\mathcal{D}, \mathcal{D}'$, there is a coupling (Z, Z') of its output multisets such that*

$$\Pr[d_{\mathrm{Ham}}(Z, Z') \leq t] \geq 1 - \eta.$$

The parameters t and η are common to the domain, and d_{Ham} counts example insertions and deletions. For the regularized selector, we use the coupling obtained from shared sampling uniforms and priorities.

Theorem 22 (Dataset stability implies node-DP)**.** *This is Theorem 8 restated for Definition 6: if $\mathcal{M}$ is (t, η)-stable under node adjacency and $\mathcal{A}$ is $(\varepsilon_0, \delta_0)$-example-DP under add/remove adjacency, then $\mathcal{A} \circ \mathcal{M}$ is $(t\varepsilon_0,\ \eta + \delta_0 \sum_{i=0}^{t-1} e^{i\varepsilon_0})$-node-DP. The proof is that of Theorem 8 in Appendix C.1.*

Corollary 23 (Parameter setting for a target node-level budget)**.** *Assume $t \geq 1$. Fix a target node-level budget (ε, δ) and choose any $\eta \in (0, \delta)$. If $\mathcal{M}$ is (t, η)-stable and $\mathcal{A}$ is $(\varepsilon_0, \delta_0)$ example-DP with*

$$\varepsilon_0 = \varepsilon/t, \qquad \delta_0 \leq \frac{\delta - \eta}{t\, e^{(t-1)\varepsilon/t}} \;\leq\; \frac{\delta - \eta}{\sum_{i=0}^{t-1} e^{i\varepsilon/t}},$$

then $\mathcal{A} \circ \mathcal{M}$ is (ε, δ) node-DP (the first, conservative choice of δ_0 suffices, since every summand of the geometric sum is at most $e^{(t-1)\varepsilon/t}$).

Proof. Plug $\varepsilon_0 = \varepsilon/t$ into Theorem 22 and choose δ_0 so that the resulting δ is at most the target δ. $\square$

C.5 Multiple deletions and fixed-graph privacy

Remark 8 (Multiple deletions)**.** *For a change of s nodes, chaining the coupling over the s adjacent steps gives at most $s\,t_N$ record changes on the good event and failure probability at most $s\eta$, so a single release is $(\varepsilon_s, \delta_s)$-node-DP at group distance s with $\varepsilon_s = s\,t_N\,\varepsilon_0$ and $\delta_s \leq s\eta + \delta_0 \sum_{i<st_N} e^{i\varepsilon_0}$, under the single public calibration t_N valid across the whole chain; for this, N must*

bound the number of examples of every dataset in the chain, that is, of every dataset within node distance s of the input (for instance $N \geq n + s\bar{\Delta}$ when at most s users with at most $\bar{\Delta}$ examples each may be added), since t_N bounds the recourse of a step only when both of its instances have at most N examples. Continual settings with adaptive releases additionally require privacy composition across releases; quantifying cumulative retention over long deletion sequences is an open direction.

Remark 9 (Comparison with fixed-graph DP)**.** *Fixed-graph DP (Ganesh et al., 2025) conditions on the attribution structure: neighboring datasets share the hypergraph and only contents change, so contribution bounding yields group size k under replacement adjacency, or $2k$ under the add/remove adjacency used by our learner assumption, since one replacement is one removal plus one addition. Node adjacency protects participation itself, and the group size becomes the recourse bound $t_N = \tilde{O}(r\sqrt{Nk})$ at constant λ; node changes are pure removals or additions, so no doubling arises on that side. The group-size multiplier of node privacy over fixed-graph privacy for this preprocessor is therefore $t_N/(2k) = \tilde{O}(r\sqrt{N/k})$, and allowing the raw-data fallback of Remark 10, at most $\min\{t_N, \bar{\Delta}\}/(2k)$ in the dense regime. The corresponding ε-group-size terms are $2k\varepsilon_0$ and $t_N\varepsilon_0$, whose ratio is $t_N/(2k)$; the node-DP guarantee additionally incurs the coupling failure probability η in δ. This compares ε-group sizes, not end-to-end risk, since the pipelines may retain different numbers of examples. One caveat: if the hypergraph is separately published, node DP of the learning pipeline cannot by itself conceal participation, so the comparison is meaningful when the structure is not released.*

D Learning utility and extensions

This appendix gives the full utility guarantee for structure-independent contents, its comparison with direct learning, and extensions to variable sample sizes and distribution shift.

D.1 Utility of the selected data

Contribution bounding changes both the number of training examples and the noise needed for privacy. Keeping fewer examples increases statistical error, but a more stable selection lets the learner use less privacy noise. To compare these effects, we state the learning guarantee in terms of the retained sample size and the recourse factor used to calibrate the learner.

For this statistical analysis, we assume that contents are independent draws from a common distribution P, independently of the attribution structure. Because selection uses only the structure, the retained contents remain independent draws from P conditional on their identities. This assumption concerns learning utility; the privacy guarantee holds for every content assignment in the input domain.

Let $L_P(\theta) = \mathbb{E}_{x\sim P}\ell(\theta; x)$ and $L_P^\star = \inf_{\theta\in\Theta} L_P(\theta)$. Assume $\ell \in [0, 1]$. Write $B(m; \varepsilon_0, \delta_0)$ for an upper bound on the learner's expected excess population risk on m independent examples. The learner is one $(\varepsilon_0, \delta_0)$-example-DP mechanism on all finite datasets, including the empty dataset, as specified in Section 2.

Theorem 24 (Utility of the two-phase pipeline)**.** *Fix public parameters and a valid recourse calibration $t \geq 1$ with failure probability $\eta = \delta/2$. Use half-weight priority rounding of y^λ and a learner with the guarantee B, calibrated with $\varepsilon_0 = \varepsilon/t$ and $\delta_0 \leq \delta e^{-\varepsilon}/(2t)$. The pipeline is (ε, δ)-node-DP on its input class. Under the content assumption above, for every realized size $M = |S| \geq 1$,*

$$\mathbb{E}\Big[L_P(\widehat{\theta}) - L_P^\star \mid M\Big] \leq B(M; \varepsilon/t, \delta_0).$$

At $M = 0$, excess risk is at most one. In particular, for a learner satisfying $B(m; \varepsilon_0, \delta_0) = \tilde{O}(1/\sqrt{m} + \sqrt{d}/(\varepsilon_0 m))$, the conditional excess risk is

$$\tilde{O}\left(\frac{1}{\sqrt{M}} + \frac{\sqrt{d}\,t}{\varepsilon M}\right).$$

If also $0 < \lambda \leq 1$, $k \geq 6\ln(2r)$, and $\mathrm{OPT}_{\mathrm{LP}} \geq cn$ for a constant $c > 0$, then

$$\mathbb{E}[L_P(\widehat{\theta})] - L_P^\star \leq \tilde{O}\left(\frac{1}{\sqrt{n}} + \frac{\sqrt{d}\,t}{\varepsilon n}\right) + \exp\left(-\Omega\left(\frac{n}{(1+r)^2}\right)\right). \tag{10}$$

The choice $t = t_N$ is valid on the size-bounded class of Corollary 10; $t = t_\rho$ is valid under the structural assumption of Corollary 6.

The two terms have a direct interpretation. The term $1/\sqrt{M}$ is the statistical error from learning on the retained examples. The term $\sqrt{d}\,t/(\varepsilon M)$ accounts for privacy noise: it increases with the recourse factor and decreases with the number of retained examples. When the feasible optimum is proportional to the full dataset size, the retention guarantee lets us express both terms in terms of n.

The displayed learner rate is the private stochastic convex optimization rate for convex Lipschitz losses over a bounded domain (Bassily et al., 2019); here it is an assumption on the chosen learner, including its handling of variable sample sizes. At constant λ, the size-based calibration gives $t_N = \tilde{O}(r\sqrt{Nk})$. When $N = \Theta(n)$ and the retention condition in (10) holds, the privacy term is therefore $\tilde{O}(\sqrt{d}\,r\sqrt{k}/(\varepsilon\sqrt{n}))$. With a uniform structural bound, the same analysis gives the privacy term $\tilde{O}(\sqrt{d}\,t_\rho/(\varepsilon n))$.

Proof of Theorem 24. Fix an attribution structure H with n examples. The selector uses only H and randomness independent of the contents. Conditional on any selected identity set S, its M retained contents therefore have law P^M. The learner uses independent randomness, so its assumed expected-risk bound gives

$$\mathbb{E}[L_P(\widehat{\theta}) - L_P^\star \mid S] \le B(M; \varepsilon_0, \delta_0), \qquad M \ge 1.$$

Averaging over sets of the same size proves the conditional claim. At size zero the bounded loss gives excess risk at most one. Separately, Theorem 8 and the public calibration yield

$$\eta + \delta_0 \sum_{i=0}^{t-1} e^{i\varepsilon_0} \le \frac{\delta}{2} + te^{\varepsilon}\frac{\delta e^{-\varepsilon}}{2t} = \delta.$$

Thus privacy holds on the full input class, without a distributional assumption on contents and without conditioning on the selected size.

To obtain an unconditional utility bound, write $\bar{m} = \mathbb{E}[M \mid H]$. The selection size is a function of the independent pairs $(U_e, P_e)_{e\in E}$. Changing one pair changes the selected set by at most $1 + r$ examples. If tentative membership changes, this follows from Lemma 13. If the example stays tentative and only its priority changes, each of its at most r users can exchange that example with at most one other example in its priority list. No other list changes. Hence bounded differences gives

$$\Pr[M < \bar{m} - a \mid H] \le \exp\left(-\frac{2a^2}{n(1+r)^2}\right), \qquad a > 0. \tag{11}$$

Under the theorem's retention assumptions, Lemma 15 gives $\bar{m} \ge cn/8$. It follows that

$$\Pr[M < cn/16 \mid H] \le \exp\left(-\frac{c^2 n}{128(1+r)^2}\right).$$

On the complementary event, substitute $M \ge cn/16$ into the assumed learner rate. On the failure event, use the upper bound one on excess risk. This proves (10), with the trivial bound one covering sample sizes for which $cn/16 < 1$. □

D.2 Retention and comparison with direct learning

Comparison with learning on all examples. At a public degree bound $\bar{\Delta} \ge 1$, direct application of an example-DP learner to all n examples gives $\tilde{O}(1/\sqrt{n} + \sqrt{d}\,\bar{\Delta}/(\varepsilon n))$ under the same learner-rate assumption. If the selection retains a constant fraction of the data with high probability, contribution bounding reduces the privacy term when t is smaller than $\bar{\Delta}$. More generally, if $M \ge m_*$ with high probability, the relevant privacy terms are proportional to t/m_* and $\bar{\Delta}/n$, while the statistical terms scale as $m_*^{-1/2}$ and $n^{-1/2}$. Thus the comparison depends on retention as well as stability. The condition $\mathrm{OPT}_{\mathrm{LP}} \ge cn$ requires at least cn/k users, because $\mathrm{OPT}_{\mathrm{LP}} \le k|V|$.

Remark 10 (Raw-data baseline and public algorithm choice)**.** *Suppose the input class has maximum user degree at most a public $\bar{\Delta} \geq 1$. A node deletion removes at most $\bar{\Delta}$ examples, so group privacy gives a direct-learning baseline with $\varepsilon_0 = \varepsilon/\bar{\Delta}$ and $\delta_0 \leq \delta e^{-\varepsilon}/\bar{\Delta}$. Under the learner-rate assumption of Theorem 24, its risk is*

$$\tilde{O}\left(\frac{1}{\sqrt{n}} + \frac{\sqrt{d}\,\bar{\Delta}}{\varepsilon n}\right).$$

If every structure in the class satisfies $\mathrm{OPT}_{\mathrm{LP}} \geq cn$, with $0 < \lambda \leq 1$ and $k \geq 6\ln(2r)$, both procedures have sample sizes of order n for the utility analysis. Choosing the procedure with the smaller public factor, t or $\bar{\Delta}$, gives

$$\tilde{O}\left(\frac{1}{\sqrt{n}} + \frac{\sqrt{d}\,\min\{t, \bar{\Delta}\}}{\varepsilon n}\right) + \exp\left(-\Omega\left(\frac{n}{(1+r)^2}\right)\right).$$

For more general input classes, compare risk bounds using public lower bounds on the raw and retained sample sizes, together with a uniform retention-failure bound. This fixes the choice across neighboring inputs. Instance-specific retention measurements can also be reported as empirical comparisons; using them to select the released procedure requires a separate privacy analysis.

General retention levels. For any $\lambda > 0$ and integer $k \geq 1$, define

$$q = \begin{cases} 1, & k \geq 6\ln(2r), \\ r, & k < 6\ln(2r), \end{cases} \qquad a_H = \frac{\phi(\lambda)\,\mathrm{OPT}_{\mathrm{LP}}(H)}{4q}.$$

Then $\bar{m} \geq a_H$. If $a_H \geq 2$ and B is nonincreasing in its sample-size argument, (11) gives the explicit bound

$$\mathbb{E}[L_P(\widehat{\theta})] - L_P^\star \leq B(\lfloor a_H/2 \rfloor; \varepsilon/t, \delta_0) + \exp\left(-\frac{a_H^2}{2n(1+r)^2}\right).$$

For smaller a_H, the exact conditional guarantee remains available by averaging $B(M; \varepsilon/t, \delta_0)$ over $M \geq 1$ and adding $\Pr[M = 0]$. Here a_H is used to analyze utility; it is not used to choose private parameters. For a deterministic approximate solver satisfying Remark 6, the same argument applies with

$$a_H = \frac{[\phi(\lambda)\,\mathrm{OPT}_{\mathrm{LP}}(H) - \epsilon_{\mathrm{opt}}]_+}{4q}$$

and the recourse calibration corrected for optimization error.

D.3 Learners on variable-size datasets

Remark 11 (Private count and public normalization)**.** *A learner may adapt to sample size while remaining one mechanism on all finite inputs. One construction uses a private count. For a fixed $\beta \in (0,1)$, form*

$$\widetilde{M} = |Z| + \mathrm{Lap}(1/(\beta\varepsilon_0)), \qquad m_0 = \max\{1, \mathrm{round}(\widetilde{M})\}.$$

For each fixed value of m_0, let $\mathcal{A}_{m_0}$ be an $((1-\beta)\varepsilon_0, \delta_0)$-example-DP mechanism on all finite datasets. Running $\mathcal{A}_{m_0}(Z)$ after the private count is then $(\varepsilon_0, \delta_0)$-example-DP by adaptive composition. For any $\zeta \in (0,1)$,

$$\Pr\left[|\widetilde{M} - |Z|| > \frac{\log(1/\zeta)}{\beta\varepsilon_0}\right] \leq \zeta.$$

This controls the count error in a utility analysis of the chosen family $\mathcal{A}_{m_0}$.

For example, suppose $\ell_0(\theta; x)$ is convex and L-Lipschitz, and define, for a fixed $m_0 > 0$,

$$\theta_{m_0}(Z) = \arg\min_{\theta \in \Theta}\left\{\frac{1}{m_0}\sum_{x \in Z}\ell_0(\theta; x) + \frac{\kappa}{2}\|\theta\|_2^2\right\},$$

on a closed convex domain where the minimum exists. Under example add/remove adjacency, its Euclidean sensitivity is at most $L/(\kappa m_0)$, independently of $|Z|$. To see this, add one example x to Z,

let the two minimizers be u, v, and apply strong convexity at both minima. Adding the two inequalities gives

$$\kappa\|u - v\|_2^2 \le \frac{\ell_0(u; x) - \ell_0(v; x)}{m_0} \le \frac{L}{m_0}\|u - v\|_2.$$

Gaussian output perturbation calibrated to this sensitivity, followed by projection onto Θ if needed, provides a private family indexed by m_0. One may use the private count above or fix m_0 publicly. This gives a privacy construction; its population-risk bound includes the normalization, regularization, and perturbation errors of that particular learner. In Theorem 24, these are included in the assumed bound B.

Population-risk assumptions. The transfer theorem accepts any established population-risk bound for the chosen learner. The convex rate displayed in Section 5.4 is achieved by private stochastic convex optimization algorithms (Bassily et al., 2019). For output perturbation and an L-Lipschitz population loss, the general estimate is

$$\mathbb{E}[L_P(\Pi_\Theta(\theta_Z + \xi)) - L_P(\theta_Z)] \le L\,\mathbb{E}\|\xi\|_2, \qquad \theta_Z \in \Theta.$$

A bound quadratic in the noise scale requires additional control on the objective or a different algorithmic analysis. The risk guarantee used in the two-phase pipeline is the one proved for the actual learner.

D.4 Distribution shift after selection

When contents depend on the attribution structure, selection can change their distribution. We first state a comparison that does not require independence, then give the learning guarantee under a conditional sampling assumption.

Lemma 25 (Selection-bias decomposition)**.** *Let P and Q be distributions on examples, and let $\ell \in [0, 1]$. For every θ,*

$$|L_P(\theta) - L_Q(\theta)| \le \mathrm{TV}(P, Q),$$

and consequently

$$L_P(\theta) - \inf_\vartheta L_P(\vartheta) \le L_Q(\theta) - \inf_\vartheta L_Q(\vartheta) + 2\,\mathrm{TV}(P, Q).$$

Both inequalities hold pointwise when Q and θ are random, and hence also after taking expectations.

Proof. The first inequality is the variational bound for total variation applied to the bounded function $x \mapsto \ell(\theta; x)$. It holds uniformly over θ, so the optimal risks also differ by at most $\mathrm{TV}(P, Q)$. Adding and subtracting the risk and optimal risk under Q proves the second inequality. □

Theorem 26 (Utility under a conditional selected-data distribution)**.** *Let W record the attribution structure and selected identities, with M the selected size. Suppose that, conditional on W and $M \ge 1$, the selected contents are independent draws from a distribution Q_W. Let one learner satisfy an expected excess-risk bound $B(m; \varepsilon_0, \delta_0)$ uniformly over these distributions and sample sizes, using randomness independent of selection. Then, for a target distribution P,*

$$\mathbb{E}\Big[L_P(\widehat{\theta}) - L_P^\star \mid W\Big] \le B(M; \varepsilon_0, \delta_0) + 2\,\mathrm{TV}(P, Q_W), \qquad M \ge 1.$$

The unconditional bound follows by averaging this expression on $M \ge 1$ and adding $\Pr[M = 0]$. With the example-DP and calibration assumptions of Theorem 24, the composed learner is also node-DP.

Proof. Conditional on W, apply the learner's guarantee to the product distribution Q_W^M, then apply Lemma 25. The loss bound handles $M = 0$. Privacy follows separately from the recourse guarantee and holds for every content assignment. □

Under the structure-independent content assumption in the main text, $Q_W = P$ and the distribution-shift term is zero. For selection among structural strata with different content distributions, a common selected-example marginal alone does not establish the conditional independence needed for the learner's i.i.d. guarantee; dependence or nonidentical distributions require their own learning bound.

E EXPERIMENTS: METHODS AND COMPLETE RESULTS

This appendix begins with the experimental protocol, then presents the complete learning comparisons, synthetic scaling and instability experiments, deletion experiments on smaller real structures, and numerical response diagnostics.

E.1 EXPERIMENTAL SETUP AND IMPLEMENTATION

We first give the datasets, learning mechanisms, and numerical procedures. The following subsections collect the learning results, scaling experiments, deletion comparisons, and response diagnostics.

E.1.1 LARGE-STRUCTURE LEARNING COMPARISONS

Large structures and paired learning trials. The synthetic structures are configuration-model hypergraphs. We draw degrees from a discrete power law with exponent 2.3, truncate them at 2 percent of the requested number of examples, and rescale to the required number of incidences. We shuffle stubs and group them into examples with sizes uniform in $\{2, \dots, r\}$, remove duplicate users within an example, and discard examples with fewer than two users. The requested number of examples is four times the number of users. The scaling series and the fixed-size synthetic comparisons use different structure seeds and are reported separately.

For each large structure, we solve the plain regularized program at $\lambda = 1$ and apply half-weight priority rounding. We use five trials for each synthetic structure and ten for each real structure. In the runs with saved paired trials, the pipeline arms share contents, test data, rounding uniforms, priorities, and base noise draws, with scales determined by their respective calibrations. The raw and nonprivate references use the same trial's full data and test examples. We compute the standard error of a paired difference from its trial-wise values using sample variance with denominator one less than the trial count. The paired comparisons use the saved trial records. At one million and three million users, the scaling files contain aggregate means but no trial records; those points are reported as means without paired standard errors. These repetitions measure variation over contents and algorithmic randomness conditional on a fixed structure.

Calibration conventions. We use the conservative choice $\mu_N = k + \sqrt{2Nk/\lambda}$, which remains valid under Corollary 10, and the structural choice $\mu_\rho = k + (r-1)k\bar{\rho}$. Both are substituted into the rounding formula with $\eta = \delta/2$. In these experiments, the realized size, largest attribution size, and maximum degree specify the illustrated public-bound scenario: $N = n$, $r = \max_e |e|$, and $\bar{\Delta} = \max_v \deg(v)$. A privacy application must fix these bounds over its admissible input class. The structural choices are $\bar{\rho} = 3$ for $r = 2$ and for the thread datasets, and $\bar{\rho} = 4.5$ for $r = 3$ and DBLP. Pointwise response measurements and sampled columns are numerical diagnostics, not uniform certificates for these choices.

Mean estimation and logistic regression. For mean estimation, the pipeline releases a jointly noisy sum and count at example-level ℓ_2 sensitivity $\sqrt{2}$, divides by the noisy count floored at one, and measures Euclidean RMSE relative to the population mean. The reported large-structure estimates are not projected onto the unit ball. The raw baseline adds Gaussian noise to the full sum at sensitivity $\bar{\Delta}$ and divides by N. For logistic regression, the feature dimension is 10, the regularization strength is $\kappa = 1/2$, and the test set has $20{,}000$ fresh examples. The pipeline uses one tenth of ε_0 for a Laplace count release, normalizes the objective by that count rounded and floored at one, and uses the remaining budget for Gaussian output perturbation. The ERM is computed with L-BFGS-B, gradient tolerance 10^{-8}, and at most 300 iterations. The task distribution is the one described with Table 8. Gaussian scales are checked using the analytic Gaussian privacy condition. The nonprivate reference is computed on every trial and then averaged.

Numerical solutions. The large-structure solver uses accelerated projected gradient on the dual and checks its stopping conditions every 100 iterations. We restore feasibility by dividing each example weight by the largest load-to-cap ratio among its users. Examples incident to a zero-cap user are set to zero before this correction. The primal-dual gap is evaluated on the corrected point. Across the intact and busiest-user-deleted instances used here, the saved gaps are numerical diagnostics; the $500{,}000$-user synthetic instance with $r = 3$ has an intact gap of $5.38 \cdot 10^{-7}$. On the two $100{,}000$-user

$r = 3$ comparisons, the intact-instance gaps are below $2.2 \cdot 10^{-11}$, while the saved deleted-instance gaps are $1.24 \cdot 10^{-6}$ and $4.73 \cdot 10^{-7}$. The one-million- and three-million-user scaling files do not include primal-dual gaps. The reported calibration factors use the exact-optimizer formulas. For a finite solver to inherit the theoretical privacy guarantee, a uniform public objective-error bound and the correction in Remark 6 must also be included. Observing a small gap on the tested inputs establishes numerical accuracy on those inputs.

E.1.2 Small-structure and construction experiments

Synthetic experiments. Figures 3, 6, 4, and 5 solve the regularized program with a warm-started SLSQP solver (convergence tolerance 10^{-12}; maximum cap-constraint violation across all reported solves $6.7 \cdot 10^{-13}$), with trial counts 120 (Figure 6), 200 per point (Figure 4), 250 per point and 50 for the concentrated instance (Figure 5). The concentrated instance has $n = 100{,}000$ examples of size two over 150,001 users, namely 50,000 examples $\{v_{\mathrm{hub}}, u_i\}$ joining the hub to distinct users u_i and 50,000 vertex-disjoint pairs of further users. With $k = 1$ and $\lambda = 1$ the regularized optimizer puts weight 1 on every disjoint pair and $1/50{,}000$ on every hub example, so the half-weight rounding retains $\mathrm{Binomial}(50{,}000, 1/2)$ disjoint pairs and at most one hub example; the mean-estimation run of Figure 5 samples this selection directly with $m_0 = 25{,}000$, and the logistic-regression run solves the program and rounds explicitly ($\mathbb{E}|S| = 24{,}998$).

Real structures. The structures of Table 5 and Figure 7 are loaded from the timestamped simplex files of Benson et al. (2018), namely email-Enron in full (10,885 emails, duplicate vertices within a hyperedge merged), email-Eu restricted to emails with at most three distinct addresses (217,012 of 235,263), and a uniform random sample without replacement (seed 2052) of 40,000 of the 146,176 DBLP papers dated 2017. The exact LP is solved with HiGHS; greedy processes hyperedges in file order unless stated otherwise. Public bounds are set to the realized $N = n$, $\bar{\Delta}$, and r (except $r = 3$ on email-Eu, which is enforced by the truncation).

Solver for the regularized program on the real structures. With $A \in \{0,1\}^{V \times E}$ the incidence matrix, the program (3) has the Lagrangian dual

$$\min_{\alpha \in \mathbb{R}^V_{\geq 0}} \; g(\alpha), \qquad g(\alpha) = k \sum_{v \in V} \alpha_v + \sum_{e \in E} \max_{y_e \in [0,1]} \Big\{ \big(1 - s_e(\alpha)\big)\, y_e - \frac{\lambda}{2} y_e^2 \Big\}, \qquad s_e(\alpha) = \sum_{v \in e} \alpha_v,$$

whose inner maximum is attained at $y_e(\alpha) = \min\{1, \max\{0, (1 - s_e(\alpha))/\lambda\}\}$, so that $\nabla g(\alpha)_v = k - \sum_{e \ni v} y_e(\alpha)$ is the cap slack at v and g is convex and differentiable. Strong duality holds for this concave quadratic program with linear constraints, and the primal optimizer is $y^\lambda = y(\alpha^*)$. We minimize g by projected gradient descent onto $\alpha \geq 0$ with Nesterov momentum, restarted whenever the gradient at the extrapolated point and the last step form an acute angle, and with the diagonal step $\eta_v = \lambda/(r \deg v)$, which is valid because the Hessian of g is $\lambda^{-1} A \,\mathrm{diag}\big(\mathbf{1}[0 < (1 - s_e(\alpha))/\lambda < 1]\big) A^\top$, entrywise nonnegative with row sums at most $\lambda^{-1} \sum_{e \ni v} |e| \leq r \deg(v)/\lambda$, so that $\mathrm{diag}(r \deg v/\lambda)$ majorizes it. Every 50 iterations the cap violation $\max_v \big(\sum_{e \ni v} y_e(\alpha) - k \big)$ is evaluated and the iteration stops once it is below $10^{-7} k$ and the iterate has moved by less than $10^{-9}(1 + \|\alpha\|_\infty)$ (at most 20,000 iterations; a few hundred to two thousand are used). The returned point is $y(\alpha)$ scaled, on each edge, by the largest violation ratio $\max_{v \in e} \max\{1, \sum_{e' \ni v} y_{e'}(\alpha)/k\}$ among its vertices, which restores exact feasibility (maximum pre-scaling violation $9.6 \cdot 10^{-7}$). The duality gap $g(\alpha) - f_H(y)$ at termination certifies the suboptimality of the returned feasible point; over the intact and deleted instances of the three structures at $\lambda \in \{1/4, 1, 4\}$ the gaps lie between $6.5 \cdot 10^{-9}$ and $6.2 \cdot 10^{-6}$, which Remark 6 converts into an increase of μ of at most 2.2.

Rounding, recourse, and the end-to-end arms. Rounded recourse and retention use 30 shared-seed trials (10 for email-Eu). The end-to-end arms use 30 trials (10 for email-Eu). The pipeline releases the clipped sum and the count of the selected examples jointly with Gaussian noise at add/remove ℓ_2-sensitivity $\sqrt{2}$ and budget $(\varepsilon_0, \delta_0)$, $\varepsilon_0 = \varepsilon/t_N$, $\delta_0 = \frac{\delta}{2t_N} e^{-\varepsilon}$, and outputs the clipped ratio; raw group privacy releases the noisy sum over all examples at node-level ℓ_2-sensitivity $\bar{\Delta}$ and budget (ε, δ), normalized by N; the fixed-graph reference releases the selected sum and count as the pipeline does but at budget $(\varepsilon/(2k),\ \frac{\delta}{4k} e^{-\varepsilon})$. All Gaussian arms use the classical calibration $\sigma = \Delta_2 \sqrt{2 \ln(1.25/\delta)}/\varepsilon$, whose theorem covers $\varepsilon < 1$; for the arms with larger budgets (raw group

privacy at $\varepsilon = 4$, the direct mechanism of Figure 4 at $\varepsilon = 2$, and the fixed-graph reference of the supervised task at $\varepsilon = 1.8$) we verified against the analytic Gaussian mechanism condition of Balle & Wang (2018) that this σ satisfies (ε, δ)-DP (it exceeds the analytic minimum by 12, 21, and 17 percent).

Other selectors (Table 7). Random-order greedy uses 10 uniformly random permutations (5 on email-Eu) shared by the intact and deleted instances; size-sorted greedy orders examples by increasing number of users with ties in file order; the distributed rule draws one shared uniform ordering per trial (10 trials, 5 on email-Eu) and runs rounds until no example is added (5 to 10 rounds); decentralized truncation and the regularized rounding use 30 shared priority draws (10 on email-Eu). The odd-cycle check uses 2,000 random orders per length.

Supervised task (Table 8). The ERM solutions are computed by L-BFGS to gradient tolerance 10^{-9}. The pipeline's learner releases $\tilde{M} = M + \mathrm{Lap}(10/\varepsilon_0)$, sets $m_0 = \max\{1, \mathrm{round}(\tilde{M})\}$, minimizes $\frac{1}{m_0}\sum_{e \in S} \ell(\theta; x_e, y_e) + \frac{\kappa}{2}\|\theta\|_2^2$ with the logistic loss ℓ and $\kappa = 1/2$, and adds Gaussian noise at sensitivity $1/(m_0\kappa)$ calibrated to $(0.9\varepsilon_0, \delta_0)$; raw group privacy minimizes the objective normalized by N over all examples and adds noise at sensitivity $\Delta/(N\kappa)$ calibrated to (ε, δ); the fixed-graph reference runs the pipeline's learner at $(\varepsilon/(2k),\ \frac{\delta}{4k}e^{-\varepsilon})$. The DBLP sample and the concentrated instance are those above; the selection is the regularized rounding at $\lambda = 1$ with half weights.

E.1.3 Grid and response measurements

Grid movement and response measurements. For Figure 1(a), we solve the intact and central-user-deleted $L \times L$ grids with $k = 1$ and measure fractional movement on surviving examples. For the response diagnostics, we also evaluate the cap path of Lemma 16 at $t \in \{0, 1/4, 1/2, 3/4, 1\}$. At a fixed point, the canonical calculation uses all numerically saturated constraints and the rank-aware pseudoinverse described in Appendix E.5. The separate sparse calculation uses positive dual multipliers and a diagonal ridge; its maximum is taken over all active columns when there are at most 20,000, and otherwise over a random sample supplemented by high-degree columns. A sampled maximum is a lower estimate of the maximum for that ridge calculation, and five cap-path samples do not bound the continuous path.

On full DBLP, the sparse pointwise diagnostic is 3.467 over 3,492 evaluated columns out of 170,733 active users. On Math Stack Exchange it is 2.224 over all 6,895 active columns. On a one-million-thread Stack Overflow sample it is 2.135 over all 17,861 active columns. These quantities describe the evaluated instances and the ridge procedure; they are distinct from a uniform bound on the canonical response norm.

E.2 Learning comparisons on large attribution structures

Tables 2 and 3 give the full mean-estimation and logistic-regression comparisons behind the main-paper results, together with the fixed-size synthetic comparisons. The four real-structure rows include every dataset used in Table 1; the synthetic rows use separate structure seeds from the scaling series.

Table 2: Learning with synthetic contents on fixed attribution structures. Means over five trials on each synthetic structure and ten on each real structure. The pipeline arms share the selected data; t_ρ uses the empirical structural bound specified in Section 6. The nonprivate reference is averaged over the same trials.

	Mean-estimation RMSE			Logistic accuracy			
Structure	t_N	t_ρ	Raw	t_N	t_ρ	Raw	Nonpriv.
Synthetic 100k, $r = 2$	0.4578	0.0140	0.1559	0.6450	0.8995	0.7286	0.9035
Synthetic 100k, $r = 3$	1.1279	0.0509	0.1927	0.4755	0.8798	0.6487	0.9048
Synthetic 500k, $r = 2$	0.2631	0.0037	0.1500	0.6813	0.9035	0.7222	0.9022
Synthetic 500k, $r = 3$	0.3391	0.0078	0.1842	0.6216	0.9057	0.7273	0.9061
DBLP	0.1912	0.0036	0.0012	0.7230	0.9032	0.9036	0.9035
Ask Ubuntu	0.7770	0.0456	0.0422	0.5802	0.8759	0.8523	0.9055
Math SX	0.6908	0.0233	0.0583	0.5857	0.8875	0.8168	0.9031
Stack Overflow	0.1413	0.0015	0.0104	0.7286	0.9035	0.9010	0.9036

Table 3: Paired logistic-accuracy differences, in percentage points, with one standard error. Each difference is computed within trial before averaging.

Structure	Structural minus conservative	Structural minus raw
Synthetic 100k, $r = 2$	25.45 ± 1.98	17.093 ± 4.134
Synthetic 100k, $r = 3$	40.44 ± 5.72	23.117 ± 5.159
Synthetic 500k, $r = 2$	22.22 ± 3.24	18.126 ± 1.964
Synthetic 500k, $r = 3$	28.41 ± 5.16	17.843 ± 2.021
DBLP	18.02 ± 1.74	-0.037 ± 0.068
Ask Ubuntu	29.57 ± 3.67	2.363 ± 0.857
Math SX	30.18 ± 2.90	7.064 ± 1.016
Stack Overflow	17.49 ± 1.62	0.249 ± 0.106

E.3 Synthetic scaling and instability experiments

We generate heavy-tailed attribution structures with $r \in \{2, 3\}$, cap $k = 5$, and approximately four examples per user. Table 4 covers 10^4 to $3 \cdot 10^6$ users. The largest structures contain approximately twelve million examples. Structural calibration improves mean accuracy over the conservative calibration throughout this series. Relative to direct learning, it has lower mean accuracy at $10{,}000$ users with $r = 3$ and higher mean accuracy at every larger tested size. The scaling series uses different structure seeds from the fixed-size comparisons in Table 2.

For fixed r, k, and $\bar{\rho}$, t_ρ is constant, so $t_\rho/\widehat{m}$ decreases as retention grows. Figure 1(b) shows this behavior through the largest saved runs. The paired accuracy comparisons in Figure 2 use only runs with saved trial records; Table 4 also includes the larger aggregate means.

Grid and pendant-cycle diagnostics. The grid experiments in Figure 1(a) examine how the fractional movement varies with the size of the structure. On the tested $L \times L$ grids with $k = 1$ and $L \in \{11, 21, 41, 81, 161\}$, the measured surviving movements at $\lambda = 1$ are 9.0, 17.6, 27.9, 46.0, and 77.8. At $\lambda = 1/4$, the measured movement equals $L - 2$ at each tested size. Since the grid has $n = 2L(L - 1)$ edges, these measurements give empirical evidence of movement on the $\sqrt{n}$ scale.

For comparison, attaching one pendant user to an odd cycle, then deleting that user at $k = 1$ and $\lambda = 1$, gives a measured surviving movement of $5/6$ for cycle lengths 21, 101, and 1001. The measured pendant-edge weight is $1/3$, while the response norm at the unperturbed cycle is $n/2$. The supremum bound is therefore at least $n/6$ on these instances. The change of saturated sets along the path explains why a large response at one endpoint can coexist with a much smaller total movement.

E.3.1 Calibration, retention, and the supervised task

Calibration scaling. Figure 4 (Appendix E) illustrates the predicted t_N/m_0 scaling on clustered networks ($k \in \{2, 4, 8\}$, parameters in the caption). The fitted slopes of the pipeline error against n are -0.64, -0.62, and -0.61 for $k = 2, 4, 8$ against the asymptotic $-1/2$, and doubling the cap

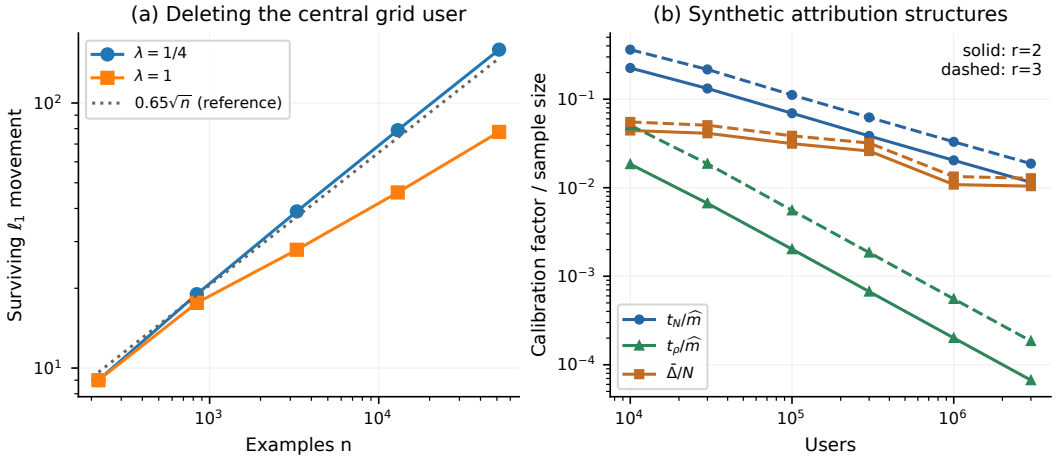


Figure 1: Structure and calibration. (a) Surviving fractional movement after deleting the central user of an $L \times L$ grid with $k = 1$. (b) Calibration factor divided by sample size on synthetic structures through three million users; $\widehat{m}$ is mean retained size. Structural factors use the empirical response-bound scenarios.

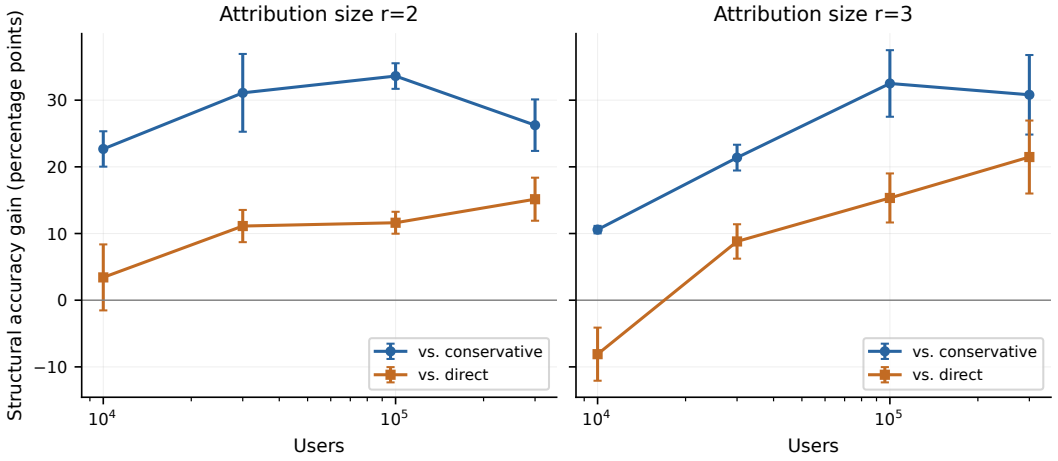


Figure 2: Paired logistic-accuracy gains on synthetic structures, separated by attribution size. Each point is a mean difference with one standard error across five trials on a fixed structure. Positive values favor structural calibration.

lowers the error by 0.66 to 0.67 against the prediction $2^{-1/2}$. The realized-size ratios $\mathbb{E}M/m_0$ lie in $[0.87, 1.00]$, so the public normalizer m_0 carries a bias of at most 0.04. The mechanism meets the single-mechanism requirement of Theorem 24 but is not claimed to attain its assumed rate, and at these sizes the noise exceeds $\|\theta^*\|_2$; on a concentrated network with $n = 10^5$ examples, half through one hub ($\bar{\Delta} = 50{,}000$, $r = 2$, $k = 1$), the pipeline attains RMSE 0.196 ± 0.006 against 1.95 ± 0.06 for raw group privacy and 0.3 for the constant predictor (Appendix E). The direct mechanism decays as $1/n$ and dominates the pipeline by orders of magnitude, because this family admits a task-specific sensitivity bound, whereas the generic pipeline is calibrated using t_N.

A supervised task. On private logistic regression with output perturbation (Table 8, Appendix E), where the pipeline's learner releases its sample size with a tenth of the example-level budget and normalizes by the release, the picture repeats. On the concentrated instance the pipeline reaches test accuracy 0.68 against 0.51 for raw group privacy (0.91 non-private and with the structure public), and on DBLP raw group privacy matches the non-private 0.90 while the pipeline sits at chance.

Table 4: Complete synthetic scaling series. Accuracy means are shown for every saved size; the last column gives a paired structural-minus-direct difference, in percentage points with one standard error, when trial records are available. The dash denotes an unavailable paired standard error, not a missing mean. Structural calibration uses $\bar{\rho} = 3$ for $r = 2$ and 4.5 for $r = 3$.

Users	r	Conservative	Structural	Direct	Nonprivate	Paired gain
10,000	2	0.5248	0.7516	0.7174	0.9072	3.42 ± 4.95
30,000	2	0.5645	0.8754	0.7642	0.8974	11.12 ± 2.42
100,000	2	0.5633	0.8994	0.7833	0.9018	11.62 ± 1.64
300,000	2	0.6385	0.9010	0.7495	0.9031	15.15 ± 3.23
1,000,000	2	0.7501	0.9047	0.8582	0.9073	—
3,000,000	2	0.8233	0.9065	0.8697	0.9084	—
10,000	3	0.5363	0.6422	0.7231	0.9013	-8.09 ± 3.98
30,000	3	0.5699	0.7838	0.6956	0.9059	8.81 ± 2.58
99,999	3	0.5595	0.8845	0.7312	0.9027	15.33 ± 3.68
300,000	3	0.5853	0.8934	0.6787	0.8971	21.47 ± 5.47
999,998	3	0.7185	0.9051	0.8490	0.9059	—
2,999,991	3	0.7260	0.9012	0.8432	0.8999	—

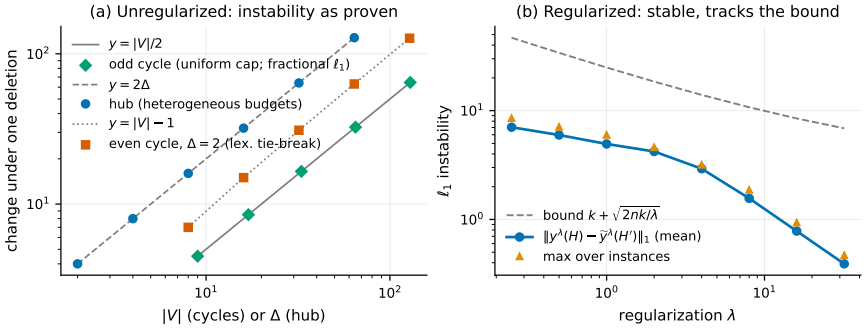


Figure 3: Instability and its repair. **(a)** The three lower-bound constructions realized exactly: the odd cycle (fractional distance $|V|/2$, unique optima on both sides), the hub with heterogeneous budgets (2Δ), and the even cycle under lexicographic selection ($|V| - 1$). **(b)** Fractional instability of the regularized program on a random hypergraph (25 vertices, 80 hyperedges of sizes two and three, $k = 3$) under deletion of the maximum-degree or a random vertex, for $\lambda \in \{0.25, \ldots, 32\}$; the plotted bound is $k + \sqrt{2nk/\lambda}$, and the largest ratio of measured movement to this bound is 0.25.

The trade-off across λ, decomposed. Figure 5 reports the per-λ study on a sparse random instance. Sampling halves the fractional mass by design and the cap-enforcing truncation removes only about three percent more ($\mathbb{E}|R| = 22.6$ versus $\mathbb{E}|S| = 21.9$ at $\lambda = 1$). The unavoidable part of the deletion change sits at its maximum, the cap $k = 4$, until the per-edge weights fall below the cap constraint at $\lambda = 8$, while the instability part decays from 5.8 to 0 as λ grows; retention falls as well, so the calibration ratio $t_N/\mathbb{E}|S|$ is U-shaped with its minimum at $\lambda \in [1, 2]$ (8.1 and 8.1 against 10.7 and 18.3 at the extremes). An ablation at $\lambda = 1$ compares the rounding with weighted ERM on the fractional weights (an empirical arm with no formal guarantee) and with uncapped sampling, which violates the contribution cap in 46 percent of draws and is reported only as an infeasible reference; all three arms release the clipped sum under the same sensitivity-one mechanism and common public normalizer, and at this size ($n = 100$) the injected noise dominates, so their errors agree to within 7 percent and the arms are separated by feasibility rather than accuracy; agreement here cannot certify that truncation is intrinsically mild. Exact uniform-cap verifications accompany the sweep: the odd-cycle fractional distance equals $|V|/2$ exactly at $m \in \{3, 6, 10\}$ with both optima certified unique by coordinate-wise minimization and maximization over the optimal face, and the triangle's centered-regularizer change equals $3/2$ at every λ tested, as Appendix B.6 predicts (the origin-centered change there is $2/\lambda$). Finally, on a concentrated network with $n = 10^5$ examples, half through a single hub user ($\bar{\Delta} = 50{,}000$, $r = 2$, $k = 1$; the instance is specified in Appendix E.1), the pipeline achieves RMSE 0.196 ± 0.006 against the raw group-privacy baseline's 1.95 ± 0.06 at the same node budget (means $\pm$ standard error over 50 trials, on the biased-Rademacher family with $\|\theta^*\|_2 = 0.3$), a $10\times$ improvement whose error also falls below the constant-zero predictor's 0.3; the

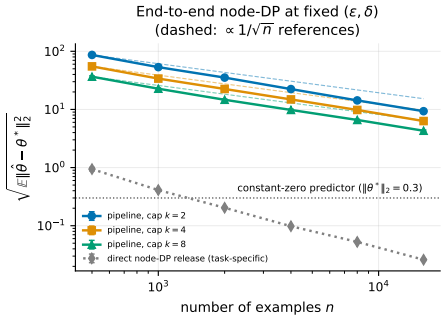


Figure 4: End-to-end node-DP mean estimation on clustered networks ($n_i = 50$ users per cluster, $r = 2$, $d = 10$, $\varepsilon = 4$, $\delta = 10^{-5}$, $k \in \{2, 4, 8\}$, 200 trials per point, two-standard-error bars). The pipeline uses $\lambda = 1$, half-weight priority rounding, and one mechanism for all sizes: Gaussian noise on the clipped example sum at add/remove sensitivity one with $\varepsilon_0 = \varepsilon/t_N$, $\eta = \delta/2$, $\delta_0 = \frac{\delta}{2t_N} e^{-\varepsilon}$, $N = n$, normalized by the public proxy $m_0 = Nk/(2n_i)$. Contents are a trial-randomized biased-Rademacher family with $\|x\|_2 = 1$ and $\|\theta^*\|_2 = 0.3$ (the constant-zero predictor's error, dotted line); the gray curve is the task-specific direct mechanism at node-level sensitivity.

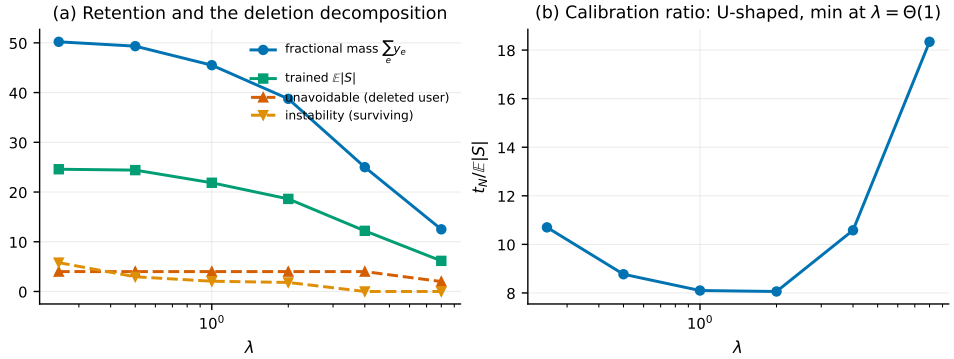


Figure 5: The stability–utility trade-off across λ on a sparse random hypergraph with $|V| = 30$, $n = 100$, hyperedge sizes two and three, cap $k = 4$, deleting the maximum-degree user ($\Delta = 16$; 250 trials per point). **(a)** Retention at each stage (fractional mass, trained $\mathbb{E}|S|$) and the decomposition of the ℓ_1 change under node deletion into the unavoidable part (mass on the deleted user's edges, pinned at the cap while that user's constraint is tight) and the instability part (movement on surviving edges, which regularization drives to zero). **(b)** The calibration ratio $t_N/\mathbb{E}|S|$, the utility axis of the trade-off once the learner is calibrated, U-shaped with minimum at $\lambda = \Theta(1)$. Monte Carlo standard errors are below marker size.

realized-size ratio is $\mathbb{E}M/m_0 = 1.00$ for both arms and the non-private sample-mean RMSE is 0.006. On the sparse instance the baseline wins, as the algorithm-choice criterion of Remark 10 predicts. On the clustered family of Figure 4 the realized-size ratios $\mathbb{E}M/m_0$ are $[0.87, 0.92]$, $[0.93, 0.98]$, and $[0.99, 1.00]$ for $k = 2, 4, 8$.

Retention and the choice of λ. Figure 6 traces both sides of the tradeoff that Section 5 resolves. Retention is flat up to $\lambda = 1$ and then decays as n/λ, and the rounded retention satisfies the Lemma 15 floor at every value of λ, with roughly a factor two to spare, consistent with the lower bound. The calibration ratio t_N/m is U-shaped with its minimum exactly at $\lambda = 1$ (values 6.0, 4.7, 3.7, 6.3 at $\lambda = \frac{1}{4}, \frac{1}{2}, 1, 2$), consistent with the stability and retention tradeoff in Remark 7, below $\lambda = 1$ stability degrades at full retention, above it retention shrinks faster than stability improves, and the observed minimum over the tested grid occurs at $\lambda = 1$. The real structures of Table 5 lie outside the condition $k \geq 6\ln(2r)$ of Lemma 15 ($k = 10$ against $6\ln 74 \approx 26$ on Enron, $k = 3$ against $6\ln 6 \approx 11$ on email-Eu and $6\ln 182 \approx 31$ on DBLP), so only the unconditional floor $\frac{1}{4r}\sum_e y_e^\lambda$ applies to them, yet their rounded retention at $\lambda = 1$ is 0.47 to 0.50 of the fractional mass (349 of 710, 880 of 1,871, and 17,496 of 35,146), so the observed retention on these structures is substantially higher than the unconditional guarantee; truncation removes little beyond the half-weight sampling.

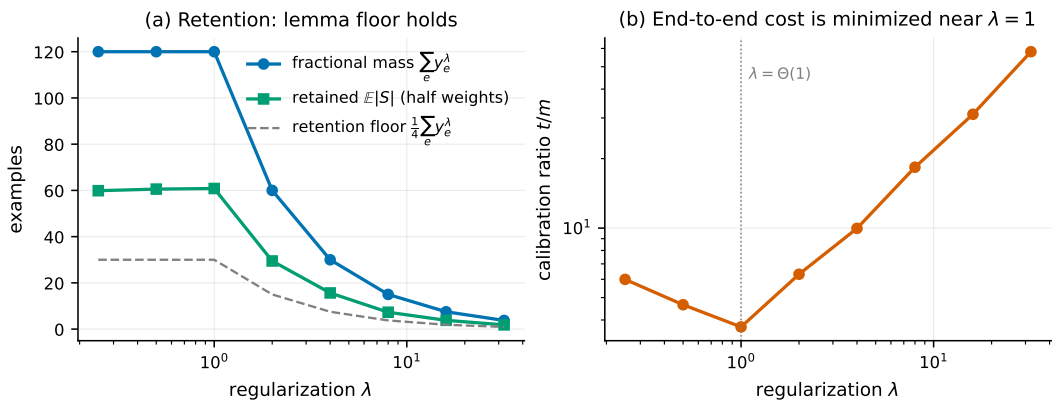


Figure 6: The stability–retention frontier on a random hypergraph with 50 vertices, 120 hyperedges of sizes two and three, and cap $k = 11 \geq 6\ln(2r)$, the regime of Lemma 15. **(a)** Fractional mass $\sum_e y_e^\lambda$, rounded retention $\mathbb{E}|S|$ at half weights (120 trials per point), and the guaranteed floor $\frac{1}{4}\sum_e y_e^\lambda$; Monte Carlo standard errors are below marker size. **(b)** The calibration ratio t_N/m that governs the privacy term of the end-to-end risk, with t_N the public factor of Corollary 10 at $\eta = 10^{-3}$ and $m = \mathbb{E}|S|$.

E.4 Deletion experiments on smaller real structures

Figure 7 plots the movement and the noise-to-data ratios across λ.

We repeat the measurements on three public structures with synthetic contents, so that no private text is processed, namely the email-Enron and email-Eu hypergraphs (an email is a hyperedge over its sender and recipients) and a fixed-seed sample of 40,000 papers from the DBLP co-authorship hypergraph for 2017 (a paper is a hyperedge over its authors), all from Benson et al. (2018); for email-Eu we keep emails with at most three addresses, a structure-only step with zero recourse that fixes $r = 3$. Table 5 lists the parameters; N, $\bar{\Delta}$, and r are set to the realized values and treated as public. We delete the maximum-degree user and measure the change of the exact LP optimum (HiGHS, deterministic tie-breaking), of greedy in file order, and of the regularized optimizer and its rounding, with contents from the family of Figure 4 at $\varepsilon = 4$, $\delta = 10^{-5}$, $\eta = \delta/2$; the pipeline has no design constant to normalize by, so it releases a noisy count with the noisy sum (one mechanism at add/remove sensitivity $\sqrt{2}$), whereas raw group privacy divides by the public N.

Three findings emerge. First, instability is not confined to the worst-case constructions. The exact LP moves 30 units of mass on surviving emails of email-Eu for a deleted user whose own emails carried 3, and 6 on DBLP for a user carrying 3; greedy in file order cascades on both email structures (17 and 6 surviving examples beyond the user's own), and the order matters, since random-order greedy and the distributed rule of Cohen-Addad et al. (2025) change 22 to 23, 9, and 1 to 2 surviving examples on Enron, email-Eu, and DBLP, decentralized truncation 19, 6, and 6, and the size-sorted pass of Ganesh et al. (2025) none on these deletions (Table 7); the regularized optimizer moves at most 0.3 units on email-Eu and DBLP at $\lambda \leq 1$ and its rounding changes at most 1.7 surviving examples on all three (1.3 to 6.7 with the user's own, Table 5), while on Enron the LP is stable under this deletion and the regularized solution redistributes a few units. Second, $t_N/\mathbb{E}|S|$ lies between 1.75 and 35 at $\lambda = 1$ while $\bar{\Delta}/N$ lies between 0.0008 and 0.12, so with the worst-case calibration raw group privacy has the smaller calibration ratio on each of these three structures, and both end-to-end arms select it (RMSE 0.005 on DBLP and 0.12 on email-Eu; on Enron, with 143 users, none of the estimators we run beats the constant predictor); the pipeline's RMSE in Table 6 is dominated by its noisy count, whose noise exceeds $\mathbb{E}|S|$ on all three structures. Third, the fixed-graph reference quantifies the price of protecting the structure, since with the structure public the selection is worth using on email-Eu (0.055 against 0.12), and node-level protection of the same selection multiplies the group size by $t_N/(2k) = 466$. Under the structural calibration the factor on email-Eu drops from 2,794 to $t_\rho = 179$ at $\bar{\rho} = 5$, but retention is at most $k|V|$ with only 976 users, and raw group privacy still wins there; on Enron, with the size bound $r = 37$, $(r-1)k\bar{\rho}$ exceeds μ_N, so the structural calibration would need a

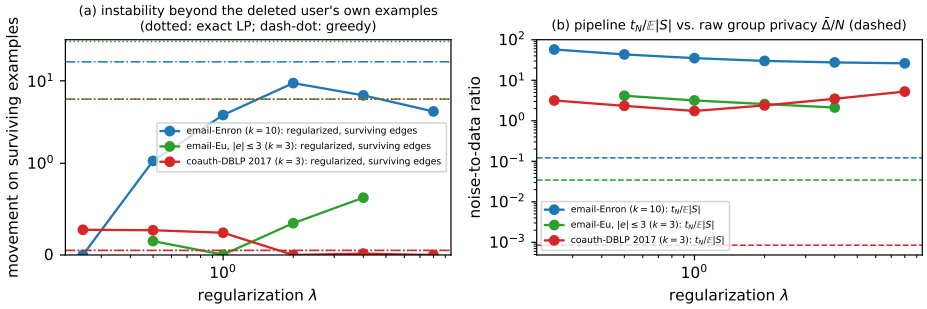


Figure 7: Real structures. **(a)** Movement of the regularized selection on surviving examples across λ under deletion of the maximum-degree user, against the exact LP (dotted) and greedy (dash-dotted). **(b)** The pipeline's noise-to-data ratio $t_N/\mathbb{E}|S|$ against raw group privacy's $\bar{\Delta}/N$ (dashed); the pipeline wins only where the solid curve lies below the dashed line. Standard errors of the recourse are below marker size.

Table 5: Sizes and deletion recourse for the smaller real structures. The deleted user has maximum degree $\bar{\Delta}$. LP, greedy, and regularized movement counts only surviving examples; rounded recourse includes the deleted user's examples. The regularized and rounded results use $\lambda = 1$.

Structure	Examples	Users	r	$\bar{\Delta}$	k
email-Enron	10,885	143	37	1,327	10
email-Eu, $\lvert e\rvert \le 3$	217,012	976	3	7,453	3
DBLP 2017 sample	40,000	100,902	91	34	3

Structure	LP	Greedy	Regularized	Rounded	$\mathbb{E}\lvert S\rvert$
email-Enron	0	17	3.9	6.7	349
email-Eu, $\lvert e\rvert \le 3$	30	6	0.0	2.0	880
DBLP 2017 sample	6	0	0.2	1.3	17,496

smaller public size bound. These examples illustrate the cost of low retention and large attribution sets. The large-structure comparisons in Section 6 examine settings with many users and concentrated contributions.

Other selectors and processing orders. Table 7 repeats the deletion of Table 5 for greedy in three orders, for the distributed rule of Cohen-Addad et al. (2025), in which every unsaturated user proposes its remaining budget of top eligible examples under a shared random ordering and an example is kept when all its users propose it in the same round, and for decentralized truncation, in which every user keeps its k incident examples of smallest shared random priority and an example is kept if all its users keep it, which is the first round of the distributed rule and Algorithm 1 with all-ones weights. Randomizing the order does not remove the cascade on the real structures, where random-order greedy and the distributed rule change 22 to 23 surviving examples on Enron and 9 on email-Eu, more than file order, and 1 to 2 on DBLP, where file order changes none; the size-sorted single pass of Ganesh et al. (2025) changes none on these three deletions, which shows that the realized recourse of greedy is order- and instance-specific, not that it is small. What no order provides is a public bound on the recourse. Decentralized truncation has recourse at most $(1 + r)\deg(v^*)$ by Lemma 13 applied with $R = E$, so under node privacy its calibration needs a degree bound and its group size $(1 + r)\bar{\Delta}$ is never below the raw baseline's $\bar{\Delta}$ while it retains fewer examples (it is a fixed-graph tool, with group size k there); for greedy in a fixed order no degree-based bound can hold at all (Proposition 11 has $\Delta = 2$ and recourse $|V| - 1$), and for random-order greedy and the distributed rule we know no nontrivial public bound. The regularized selection is calibrated by $t_N = \tilde{O}(r\sqrt{Nk})$ with no degree bound and changes at most 1.7 surviving examples here; its half weights halve retention by design (Lemma 15). On the odd cycle itself the cascade of Proposition 11 is order-specific, since under a uniformly random shared order greedy on C_{2m+1} with $k = 1$ changes 0.55 to 0.59 surviving edges in expectation for $m \in \{10, 50, 500\}$ (2,000 orders each) while returning $0.87m$ to $0.91m$ edges, so

Table 6: Calibration and mean-estimation error for the structures in Table 5, at $\varepsilon = 4$. RMSE uses 30 trials (10 for email-Eu); standard errors are 12–32% of the pipeline values and at most 9% of the other values. The pipeline jointly releases a noisy selected sum and count; direct learning uses the public size N. The fixed-graph reference uses the same selection with group size $2k$. The constant-zero predictor has RMSE 0.3.

Structure	t_N	$t_N/\mathbb{E}\lvert S\rvert$	$\bar{\Delta}/N$
email-Enron	12,263	35.2	0.122
email-Eu, $\lvert e\rvert \le 3$	2,794	3.18	0.034
DBLP 2017 sample	30,558	1.75	0.0008

Structure	Pipeline RMSE	Direct RMSE	Fixed-graph RMSE
email-Enron	58,233	0.467	0.395
email-Eu, $\lvert e\rvert \le 3$	7,275	0.123	0.055
DBLP 2017 sample	145,916	0.005	0.007

Table 7: Other selectors on the real structures under the deletion of Table 5, reporting the surviving examples whose selection changes (mean $\pm$ standard error over the shared orderings or priority draws of Appendix E.1) and the retention $|S|$. Greedy in file order, in uniformly random order, and in the increasing-size order of Ganesh et al. (2025); the distributed rule of Cohen-Addad et al. (2025) with its default shared random ordering, run to convergence; decentralized truncation (Algorithm 1 with all-ones weights); the regularized selection at $\lambda = 1$ with half weights (different random draws than Table 5).

	email-Enron ($k = 10$)		email-Eu, $\lvert e\rvert \le 3$ ($k = 3$)		DBLP 2017 sample ($k = 3$)	
selector	change	$\lvert S\rvert$	change	$\lvert S\rvert$	change	$\lvert S\rvert$
greedy, file order	17	597	6	1,379	0	34,909
greedy, random order	22.4 ± 2.3	570	8.8 ± 0.4	1,335	1.7 ± 0.4	34,916
greedy, size-sorted order	0	689	0	1,882	0	35,187
distributed rule	22.8 ± 1.5	568	9.2 ± 1.4	1,338	0.9 ± 0.2	34,922
decentralized truncation	19.1 ± 0.9	296	6.2 ± 0.8	637	6.2 ± 0.4	34,284
regularized, $\lambda = 1$	1.7 ± 0.2	351	0.0 ± 0.0	894	0.1 ± 0.1	17,490

random-order greedy is stable there but suboptimal, and Corollary 2 does not apply to it. All edges of the cycle have the same size, so the size-sorted order of Ganesh et al. (2025) with index tie-breaking is the order of Proposition 11.

A supervised task. Table 8 reports ℓ_2-regularized logistic regression ($d = 10$, regularization $\kappa = 1/2$) on the DBLP sample and on the concentrated instance, with features $x_e = s \circ b_e/\sqrt{d}$ for a trial-level sign vector s and independent Rademacher b_e, so that $\|x_e\|_2 = 1$ and the logistic loss is 1-Lipschitz, labels $y_e = \operatorname{sign}(\langle \theta_{\mathrm{true}}, x_e \rangle + \xi_e)$ with $\|\theta_{\mathrm{true}}\|_2 = 3$ and $\xi_e \sim \mathcal{N}(0, 0.09)$, and test accuracy on 20,000 fresh examples over 50 trials. All private arms are output perturbation (Chaudhuri et al., 2011) of the ERM solution normalized by a size that is public conditionally on the arm's own releases. The pipeline's learner is construction (i) of Remark 11. It releases the selected count with Laplace noise at budget $\varepsilon_0/10$, normalizes the objective by $m_0 = \max\{1, \operatorname{round}(\tilde{M})\}$, and adds Gaussian noise at add/remove sensitivity $1/(m_0\kappa)$ and budget $(0.9\varepsilon_0, \delta_0)$, so it is $(\varepsilon_0, \delta_0)$-example-DP on every input and the pipeline is (ε, δ)-node-DP at the public t_N by Corollary 10. Raw group privacy normalizes by the public N and adds Gaussian noise at node-level sensitivity $\bar{\Delta}/(N\kappa)$ and budget (ε, δ); the fixed-graph reference runs the pipeline's learner at budget $(\varepsilon/(2k),\ \frac{\delta}{4k}e^{-\varepsilon})$. The picture matches mean estimation. On the concentrated instance ($t_N = 919$, $\mathbb{E}|S| = 24{,}998$) the pipeline reaches accuracy 0.68 where raw group privacy, whose noise scale $\bar{\Delta}/(N\kappa) = 1$ swamps the signal, stays at 0.51; with the structure public the same selection gives 0.91, the non-private accuracy. On DBLP ($t_N = 30{,}558$) raw group privacy is noiseless at this scale and matches the non-private 0.90, while the pipeline's $\varepsilon_0 = 1.3 \cdot 10^{-4}$ leaves it at chance.

Table 8: Private ℓ_2-regularized logistic regression by output perturbation at $\varepsilon = 4$, $\delta = 10^{-5}$, with test accuracy on $20{,}000$ fresh examples (mean $\pm$ standard error over 50 trials; chance is 0.5). The pipeline's learner releases its sample size privately (Remark 11); the fixed-graph reference runs the same learner at group size $2k$ with the structure public.

structure	n	$\bar{\Delta}$	k	t_N	non-private	pipeline	raw group privacy	fixed-graph
concentrated instance	100,000	50,000	1	919	0.906 ± 0.002	0.677 ± 0.010	0.510 ± 0.015	0.906 ± 0.002
DBLP 2017 sample	40,000	34	3	30,558	0.901 ± 0.001	0.515 ± 0.014	0.901 ± 0.001	0.901 ± 0.001

E.5 Response-norm computation and diagnostics

Evaluating the defined norm. For specified fractional and saturated sets, a reduced singular value decomposition $B = Q\Sigma R^\top$, retaining its nonzero singular values, gives

$$(B^\top)^+ = Q\Sigma^{-1}R^\top, \qquad \rho = \max_{u \in T} \|Q\Sigma^{-1}R^\top e_u\|_1.$$

Zero columns can be removed before the computation. Singular blocks with positive rank contribute through their nonzero singular values. Numerical evaluation should report the tolerances used to identify the fractional and saturated sets and to determine the rank. On large matrices, the same columns can be obtained by rank-aware least-squares methods. The largest norm over a sampled set of columns is a lower estimate of the maximum over all columns.

An isotropic ridge satisfies $B(B^\top B + \epsilon I)^{-1} \to (B^\top)^+$ as $\epsilon \downarrow 0$. A finite ridge and numerical solution error still need to be accounted for if the resulting value is to be used as an upper bound. A diagonal ridge $B(B^\top B + \epsilon D)^{-1}$ generally has a different limit on coordinate right-hand sides when $B^\top B$ is singular. Agreement on compatible perturbations in (9) is insufficient to identify the operator norm, which is computed from all coordinate vectors.

Numerical diagnostics. For the numerical diagnostics we use the positive-multiplier set T_+ and its incidence matrix B_+, with

$$Q_\epsilon = B_+\big(B_+^\top B_+ + \epsilon\,\mathrm{diag}(d^U)\big)^{-1}, \qquad \epsilon = 10^{-8},$$

where d_u^U is the fractional degree. We evaluate all columns when $|T_+| \le 20{,}000$ and otherwise sample $3{,}000$ random columns together with the 500 columns of largest fractional degree. Write q_ϵ for the largest evaluated column norm, maximized over the five sampled points of a cap path. Table 9 records these diagnostics and their products with the released load. The quantity q_ϵ depends on this matrix choice, regularization, and sampling scheme; ρ in (5) is defined using all saturated users and the Moore–Penrose pseudoinverse.

Pointwise diagnostics on the large structures. Table 10 distinguishes the sparse ridge calculation from the canonical pseudoinverse calculation. The latter uses all saturated constraints and all columns, evaluated component by component when their size permits. Both describe the optimizer at the measured point; neither establishes a uniform bound over an input class or all points of a deletion path.

Table 9: Numerical diagnostics at $\lambda = 1$ for the three busiest deletions in each structure, with one deletion per row. Here q_ϵ is the maximum evaluated column norm of the diagonal-ridge operator over the sampled cap-path points. The product column is $\|b\|_1 q_\epsilon$, and the last column is the measured ℓ_1 movement on surviving examples. These are instance-level diagnostics; the uniform class bounds are stated in Proposition 7.

structure	k	$\lvert T_+ \rvert$	degree	$\|b\|_1$	q_ϵ	product	movement
email-Enron	10	109	1,327	1.87	3.15	5.9	3.85
			1,282	9.69	3.09	30.0	15.2
			1,097	9.73	3.05	29.7	19.7
email-Eu, $\lvert e \rvert \leq 3$	3	759	7,453	1.46	4.71	6.9	0.003
			5,475	0.50	4.71	2.4	0.000
			5,061	2.29	4.71	10.8	7.43
DBLP 2017 sample	3	2,277	34	9.25	3.53	32.7	0.24
			31	9.44	3.53	33.3	0.91
			28	9.51	3.53	33.6	0.23

Table 10: Pointwise response diagnostics at $\lambda = 1$. Sparse values use the positive-multiplier ridge calculation; the column count is evaluated / active. A dash in the canonical column means the computation exceeds the component-size limit of the dense routine.

Structure	Examples	Sparse	Columns	Canonical
DBLP prefix	100,000	2.705	5,891 / 5,891	2.705
DBLP prefix	300,000	2.951	3,447 / 26,146	—
DBLP full	2,761,543	3.467	3,492 / 170,733	—
Ask Ubuntu	186,881	1.924	2,031 / 2,031	1.924
Math Stack Exchange prefix	100,000	2.031	1,596 / 1,596	2.031
Math Stack Exchange full	646,131	2.224	6,895 / 6,895	—
Stack Overflow prefix	1,000,000	2.135	17,861 / 17,861	—