Private optimization without constraint violations

Andrés Muñoz Medina, Umar Syed, Sergei Vassilvitskii, **Ellen Vitercik**



Google Research and UC Berkeley

AISTATS'21

Private linearly-constrained optimization

Goal: Privately find $\boldsymbol{x} \in \mathbb{R}^n$ maximizing $g(\boldsymbol{x})$ such that $A\boldsymbol{x} \leq \boldsymbol{b}(D)$

Lipschitz

Private database $D \subseteq \mathcal{X}$

Solution can't violate any constraint

Crucial in many applications, such as resource allocation

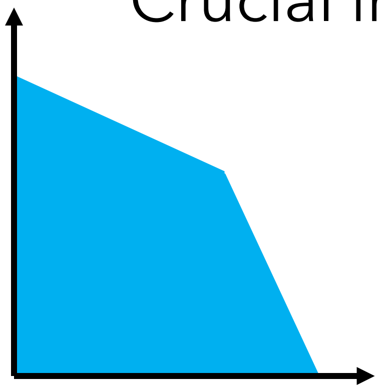

Example: linear programming

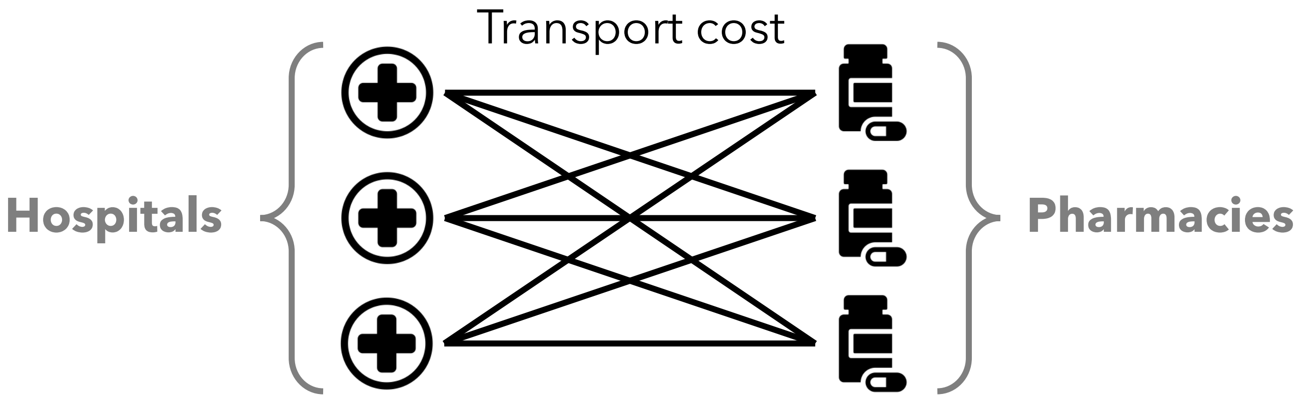

Goal: Decide which pharmacies should supply which hospitals

RHS of constraints is private:

Indicates number of patients with disease at each hospital

If constraints violated, hospital can't treat all patients

Our contributions

1. Differentially-private algorithm

2. **Main result:** Nearly-matching lower bound on loss
   Matching up to log factors

Most related prior research

Differentially private linear programming

Hsu et al., ICALP'14; Cummings et al., WINE'15

Primary distinctions:

- Specific to linear programming
- Allow constraints to be violated by bounded amount
- Constraints $(A, \boldsymbol{b})$ and objective function can be private

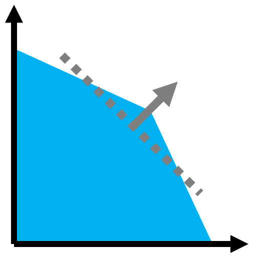

Outline

1. Introduction
2. **Background: Differential privacy**
3. Algorithm
4. Lower bound
5. Experiments
6. Conclusion

Differential privacy

$\boldsymbol{x}(D) \in \mathbb{R}^n$: algorithm's output given database D

Algorithm is **differentially private** if:

$\boldsymbol{x}(D)$ reveals (almost) nothing more about a record in D

than it would have if the record wasn't in D

Differential privacy

$\boldsymbol{x}(D) \in \mathbb{R}^n$: algorithm's output given database D

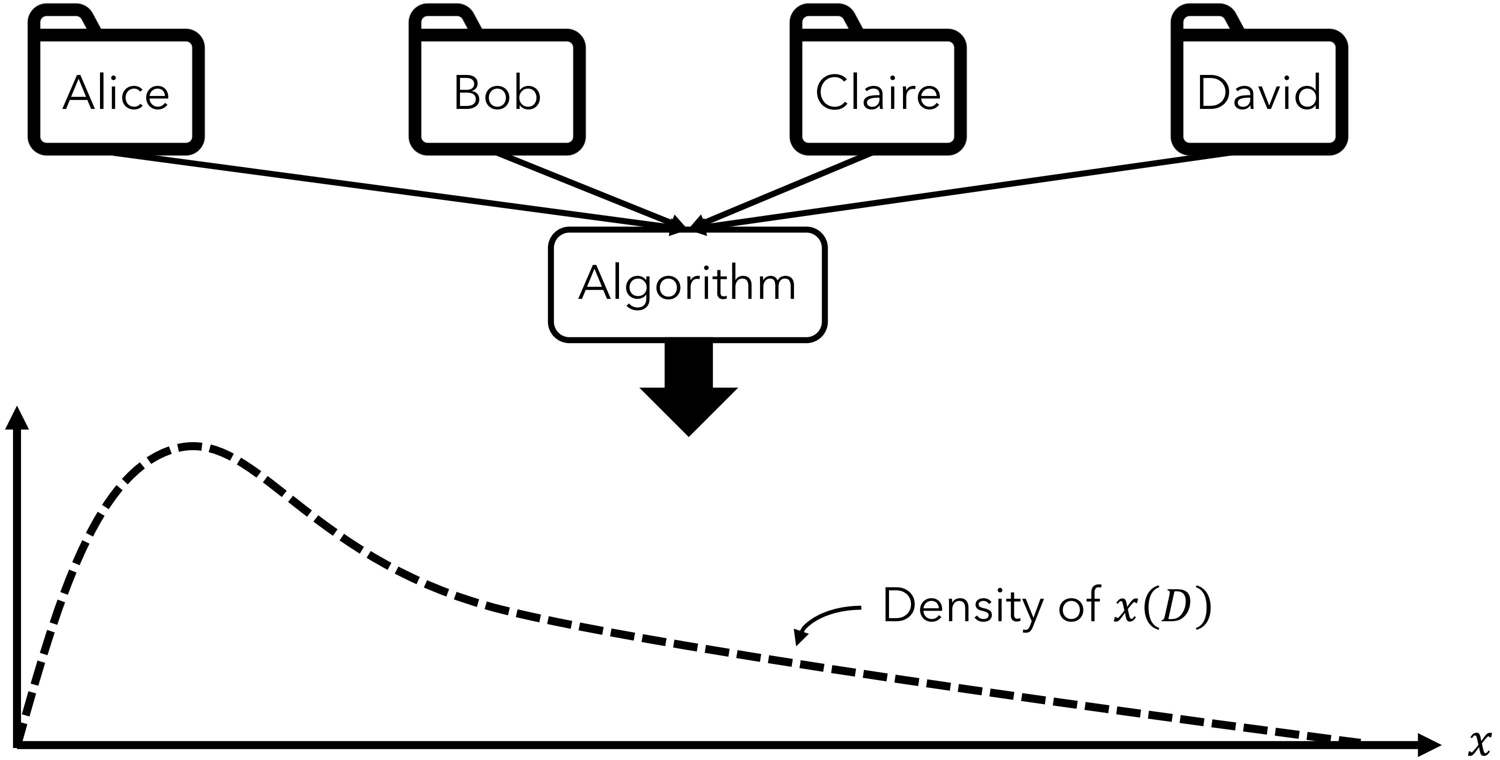

Differential privacy

$\boldsymbol{x}(D) \in \mathbb{R}^n$: algorithm's output given database D

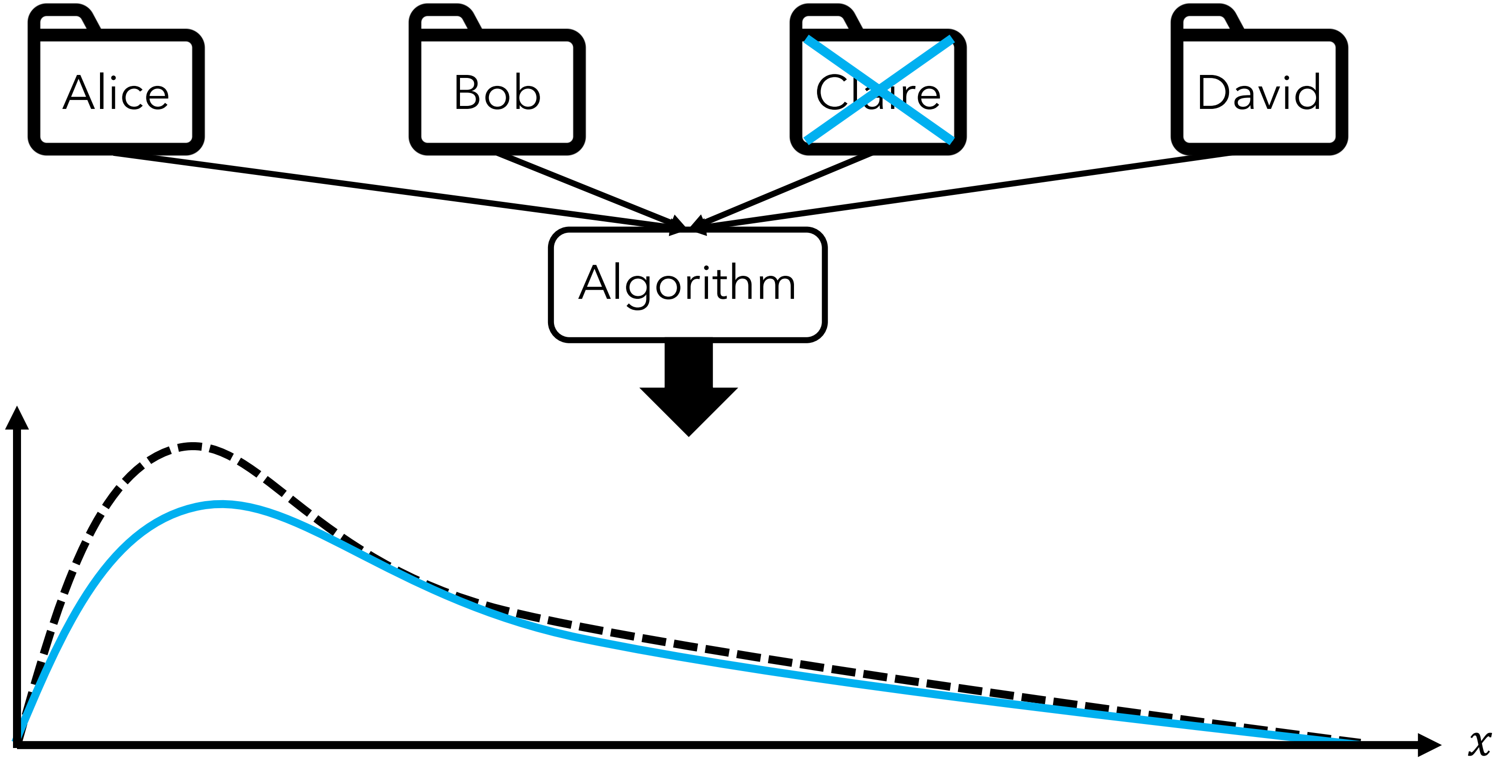

Differential privacy

Two databases D, D' are *neighboring* if differ on ≤ 1 element

Denoted $D \sim D'$

Algorithm is (ε, δ)-differentially private if:

For any $D \sim D'$ and $V \subseteq \mathbb{R}^n$, $\mathbb{P}[\boldsymbol{x}(D) \in V] \leq e^{\varepsilon}\mathbb{P}[\boldsymbol{x}(D') \in V] + \delta$

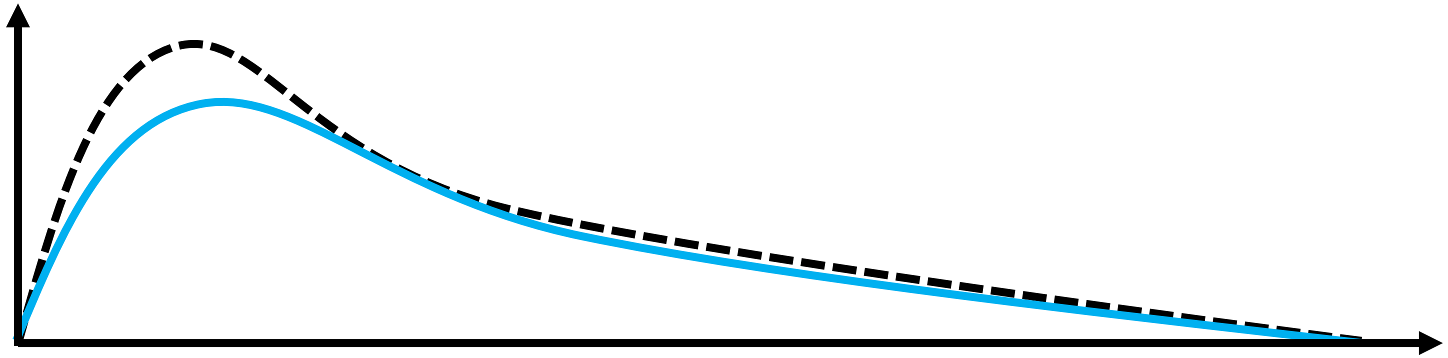

Outline

1. Introduction
2. Background: Differential privacy
3. **Algorithm**
4. Lower bound
5. Experiments
6. Conclusion

Feasibility assumption

If feasible region changes too much between databases:
    Private optimization w/o constraint violations is **impossible**

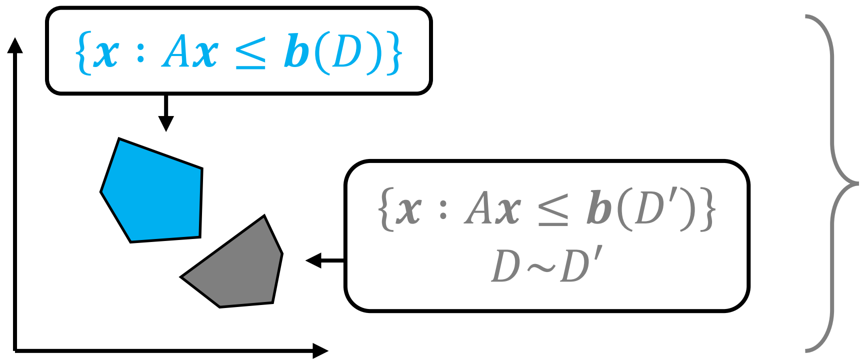

There's no (ϵ, δ)-DP algorithm with $\delta < 1$

Assumption: $\cap_{D \subseteq \mathcal{X}} \{\boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}(D)\} \neq \emptyset$
    E.g., it includes the origin

In particular, $\cap_{D \subseteq \mathcal{X}} \{\boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}(D)\} = \{\boldsymbol{x} : A\boldsymbol{x} \leq (b_1^*, \ldots, b_m^*)\}$

Algorithm

1. Map constraint vector $\boldsymbol{b}(D) \mapsto \overline{\boldsymbol{b}}(D)$ such that $\overline{\boldsymbol{b}}(D) \leq \boldsymbol{b}(D)$ using the Truncated Laplace Mechanism

Algorithm

1. Map constraint vector $\boldsymbol{b}(D) \mapsto \overline{\boldsymbol{b}}(D)$ such that $\overline{\boldsymbol{b}}(D) \leq \boldsymbol{b}(D)$:
   - Sensitivity: $\Delta = \max_{D \sim D'} \|\boldsymbol{b}(D) - \boldsymbol{b}(D')\|_1$
   - $s = \frac{\Delta}{\epsilon} \ln\left(\frac{m(e^{\epsilon}-1)}{\delta} + 1\right)$
   - η_i = Truncated Laplace noise with scale $\frac{\Delta}{\epsilon}$ and support $[-s, s]$
   - $\overline{\boldsymbol{b}}(D)_i = \max\{\boldsymbol{b}(D)_i - s + \eta_i, b_i^*\}$

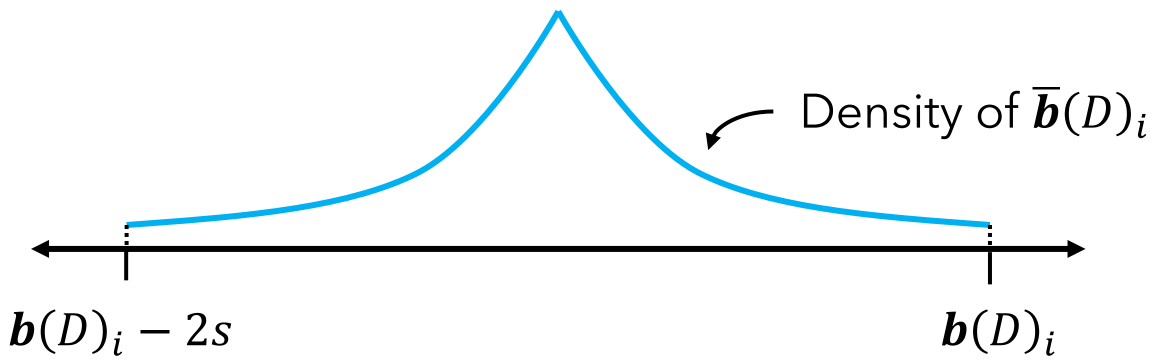

Algorithm

1. Map constraint vector $\boldsymbol{b}(D) \mapsto \overline{\boldsymbol{b}}(D)$ such that $\overline{\boldsymbol{b}}(D) \leq \boldsymbol{b}(D)$ using the Truncated Laplace Mechanism
2. Return $\boldsymbol{x} \in \mathbb{R}^n$ maximizing $g(\boldsymbol{x})$ such that $A\boldsymbol{x} \leq \overline{\boldsymbol{b}}(D)$

Important properties:

1. Satisfies **constraints** with probability 1
   $A\boldsymbol{x} \leq \overline{\boldsymbol{b}}(D) \leq \boldsymbol{b}(D)$

2. Satisfies **(ε, δ)-DP**
   Truncated Laplace is private

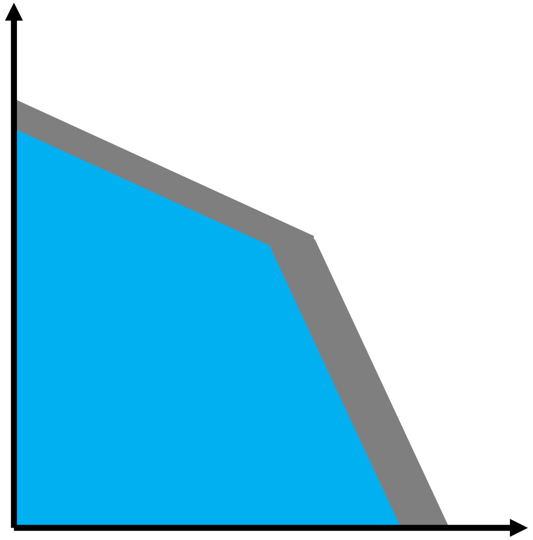

Outline

1. Introduction
2. Background: Differential privacy
3. Algorithm
   - **Quality guarantee**
4. Lower bound
5. Experiments
6. Conclusion

Linear system condition number

$$\alpha_{p,q}(A) = \sup_{\boldsymbol{u} \geq 0} \left\{ \|\boldsymbol{u}\|_{p^*} : \begin{array}{c} \|A^T \boldsymbol{u}\|_{q^*} = 1 \\ \text{\& the rows of } A \text{ corresponding to nonzero} \\ \text{components of } \boldsymbol{u} \text{ are linearly independent} \end{array} \right\}$$

E.g., when $p = q = 2$ and A is nonsingular, $\alpha_{p,q}(A) = \sigma_{\min}^{-1}(A)$

Theorem [Li, '93]**:**

- Let $S = \{\boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}\}$ and $S' = \{\boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}'\}$
- For all $\boldsymbol{x} \in S$, $\inf_{\boldsymbol{x}' \in S'} \|\boldsymbol{x} - \boldsymbol{x}'\|_q \leq \alpha_{p,q}(A) \|\boldsymbol{b} - \boldsymbol{b}'\|_p$

Quality guarantee

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$g(\boldsymbol{x}^*)$

Optimal solution

Quality guarantee

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$$g(\boldsymbol{x}^*) - g(\boldsymbol{x}(D))$$

Algorithm's output

Quality guarantee

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$$g(\boldsymbol{x}^*) - g(\boldsymbol{x}(D)) \leq \Delta$$

Constraints' sensitivity

Quality guarantee

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$$g(\boldsymbol{x}^*) - g(\boldsymbol{x}(D)) \leq \Delta \cdot L \cdot \inf_{p \geq 1}\left\{\alpha_{p,q}(A) \sqrt[p]{m}\right\}$$

Number of constraints

Quality guarantee

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$$g(\boldsymbol{x}^*) - g(\boldsymbol{x}(D)) \leq \Delta \cdot L \cdot \inf_{p \geq 1}\left\{\alpha_{p,q}(A)\sqrt[p]{m}\right\} \cdot \frac{2}{\varepsilon} \cdot \ln\left(\frac{m(e^{\varepsilon} - 1)}{\delta} + 1\right)$$

(ε, δ)-differential privacy

Nearly-matching lower bound

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$$g(\boldsymbol{x}^*) - g(\boldsymbol{x}(D)) \le \Delta \cdot L \cdot \inf_{p \ge 1} \left\{ \alpha_{p,q}(A) \sqrt[p]{m} \right\} \cdot \frac{2}{\varepsilon} \cdot \ln\left(\frac{m(e^\varepsilon - 1)}{\delta} + 1 \right)$$

Lower bnd (informal)**:** Exist problems s.t. for any (ε, δ)-DP alg,

$$g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \ge \Delta \cdot L \cdot \inf_{p \ge 1} \left\{ \alpha_{p,1}(A) \sqrt[p]{m} \right\} \cdot \frac{1}{4\varepsilon} \cdot \ln\left(\frac{e^\varepsilon - 1}{2\delta} + 1 \right)$$

Nearly-matching lower bound

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$$g(\boldsymbol{x}^*) - g(\boldsymbol{x}(D)) \leq \Delta \cdot L \cdot \inf_{p \geq 1}\left\{\alpha_{p,q}(A)\sqrt[p]{m}\right\} \cdot \frac{2}{\varepsilon} \cdot \ln\left(\frac{m(e^{\varepsilon} - 1)}{\delta} + 1\right)$$

Lower bnd (informal)**:** Exist problems s.t. for any (ε, δ)-DP alg,

$$g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \Delta \cdot L \cdot \inf_{p \geq 1}\left\{\alpha_{p,1}(A)\sqrt[p]{m}\right\} \cdot \frac{1}{4\varepsilon} \cdot \ln\left(\frac{e^{\varepsilon} - 1}{2\delta} + 1\right)$$

Takeaway: Matching up to $\boldsymbol{O(\ln m)}$

Quality upper bound

Upper bound: Suppose g is L-Lipschitz under $\|\cdot\|_q$. Then

$$g(\boldsymbol{x}^*) - g(\boldsymbol{x}(D)) \leq \Delta \cdot L \cdot \inf_{p \geq 1}\left\{\alpha_{p,q}(A)\sqrt[p]{m}\right\} \cdot \underbrace{\frac{2}{\varepsilon} \cdot \ln\left(\frac{m(e^{\varepsilon} - 1)}{\delta} + 1\right)}_{2s}$$

Proof:

- $\boldsymbol{b}$: Arbitrary vector in support of $\bar{\boldsymbol{b}}(D)$ and $S = \{\boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}\}$
- From Li ['93]: $\inf_{\boldsymbol{x} \in S} \|\boldsymbol{x}^* - \boldsymbol{x}\|_q \leq \alpha_{p,q}(A)\|\boldsymbol{b}(D) - \boldsymbol{b}\|_p$
- $\|\boldsymbol{b}(D) - \boldsymbol{b}\|_p \leq 2s\sqrt[p]{m}$

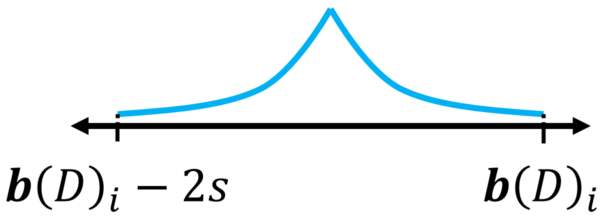

Outline

1. Introduction
2. Background: Differential privacy
3. Algorithm
4. **Lower bound**
5. Experiments
6. Conclusion

Nearly-matching lower bound

Theorem (more details):

- A: arbitrary diagonal matrix
- $g(x) = \langle 1, x \rangle$
- For any $\Delta > 0$, exists mapping from databases D to $\boldsymbol{b}(D)$ s.t.:
  1. Sensitivity of $\boldsymbol{b}(D)$ is Δ
  2. For any $\epsilon > 0, \delta \in (0, ½]$ and any (ε, δ)-DP algorithm,

$$g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A) \sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$$

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

- For all $i \in \mathbb{Z}$, let D_i be a database with $D_i \sim D_{i+1}$ & $b(D_i) = \Delta i$

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

- For all $i \in \mathbb{Z}$, let D_i be a database with $D_i \sim D_{i+1}$ & $b(D_i) = \Delta i$
- For any $V \subseteq \mathbb{R}$, $\mathbb{P}[x(D_i) \in V] \leq e^{\varepsilon}\mathbb{P}[x(D_{i-1}) \in V] + \delta$

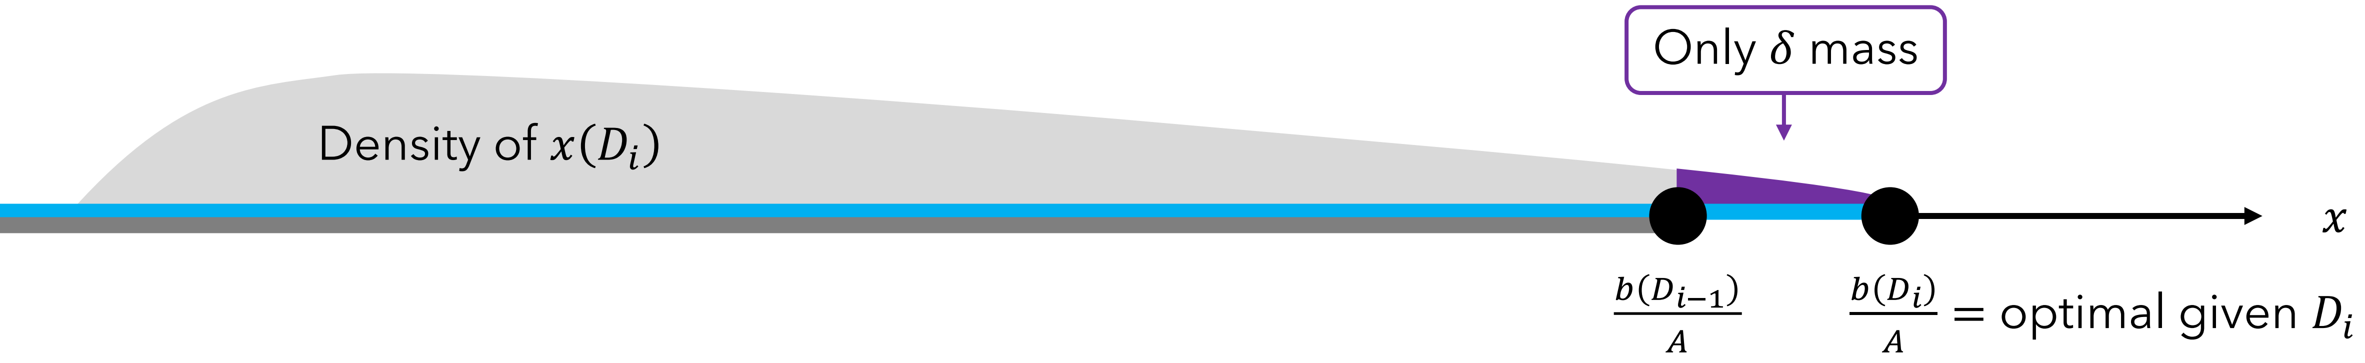

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

- For all $i \in \mathbb{Z}$, let D_i be a database with $D_i \sim D_{i+1}$ & $b(D_i) = \Delta i$
- For any $V \subseteq \mathbb{R}$, $\mathbb{P}[x(D_i) \in V] \leq e^{\varepsilon}\mathbb{P}[x(D_{i-1}) \in V] + \delta$

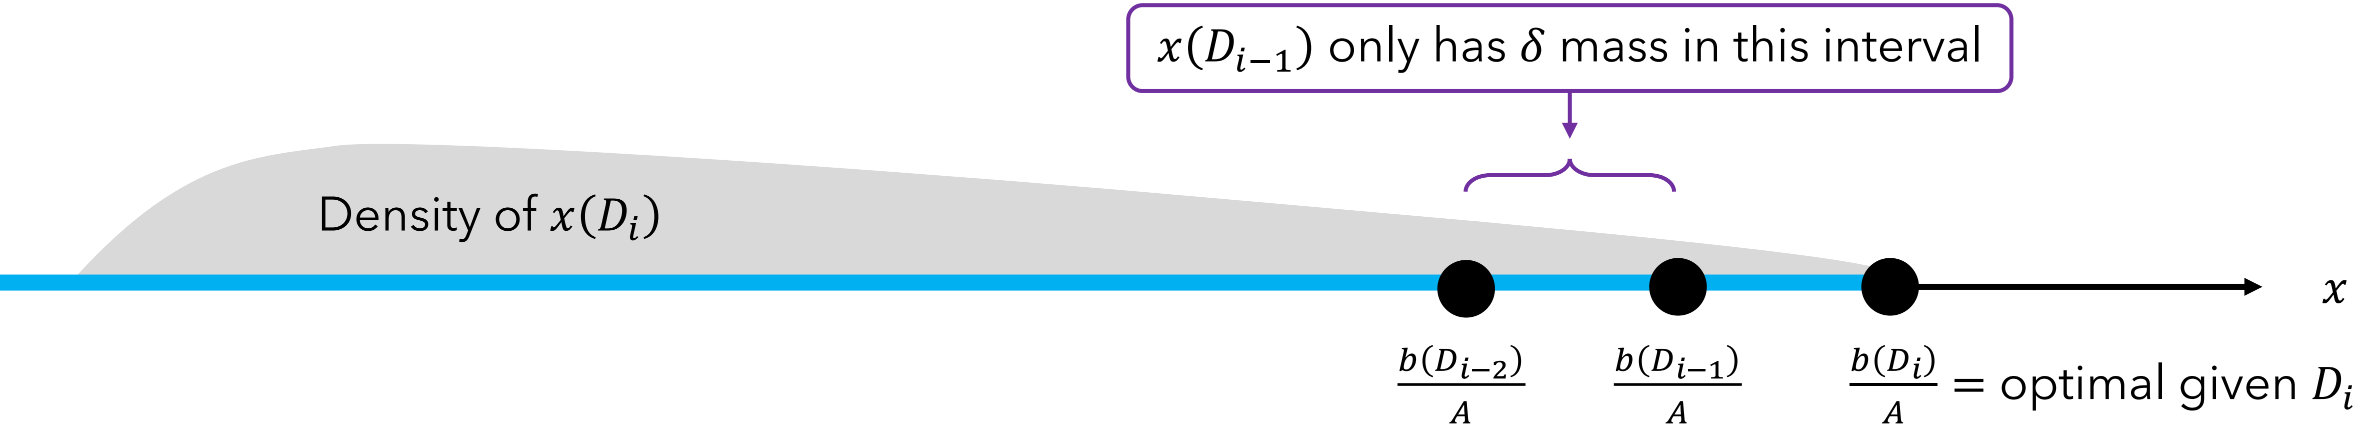

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

- For all $i \in \mathbb{Z}$, let D_i be a database with $D_i \sim D_{i+1}$ & $b(D_i) = \Delta i$
- For any $V \subseteq \mathbb{R}$, $\mathbb{P}[x(D_i) \in V] \leq e^{\varepsilon}\mathbb{P}[x(D_{i-1}) \in V] + \delta$

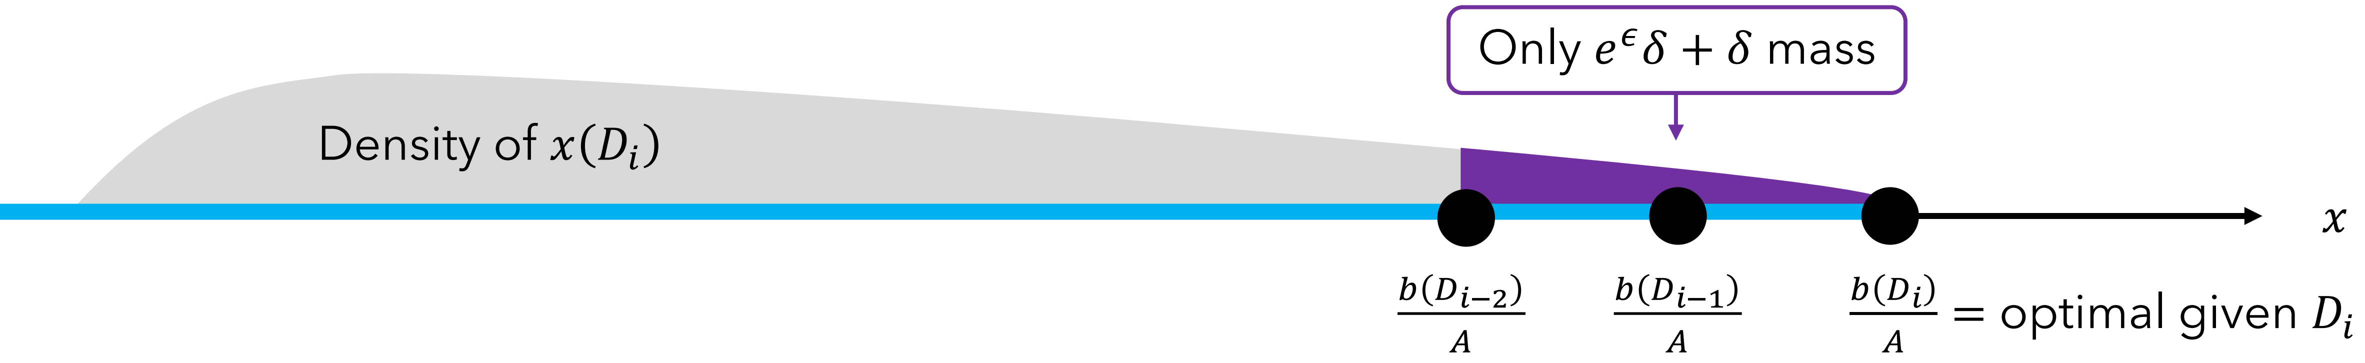

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

- For all $i \in \mathbb{Z}$, let D_i be a database with $D_i \sim D_{i+1}$ & $b(D_i) = \Delta i$
- For any $V \subseteq \mathbb{R}$, $\mathbb{P}[x(D_i) \in V] \leq e^{\varepsilon}\mathbb{P}[x(D_{i-1}) \in V] + \delta$

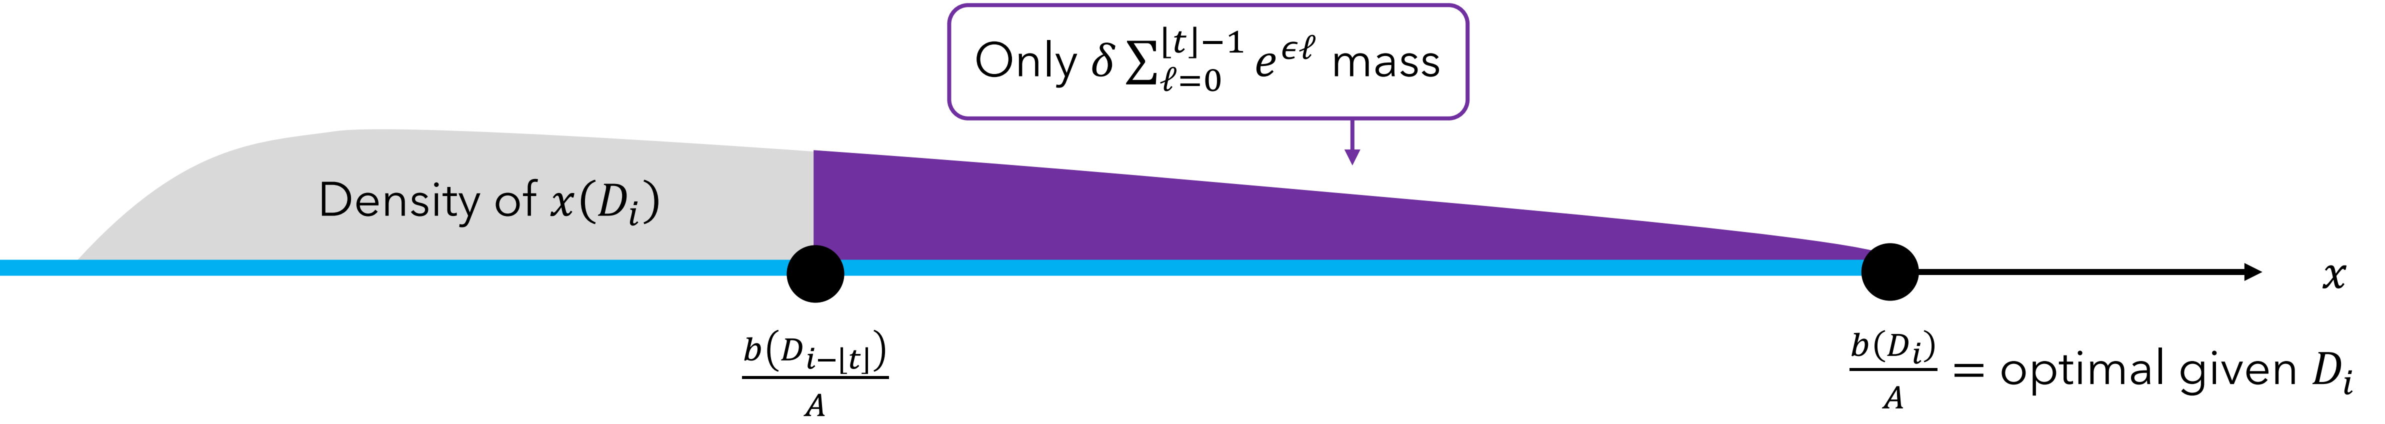

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

- For all $i \in \mathbb{Z}$, let D_i be a database with $D_i \sim D_{i+1}$ & $b(D_i) = \Delta i$
- For any $V \subseteq \mathbb{R}$, $\mathbb{P}[x(D_i) \in V] \leq e^{\varepsilon}\mathbb{P}[x(D_{i-1}) \in V] + \delta$

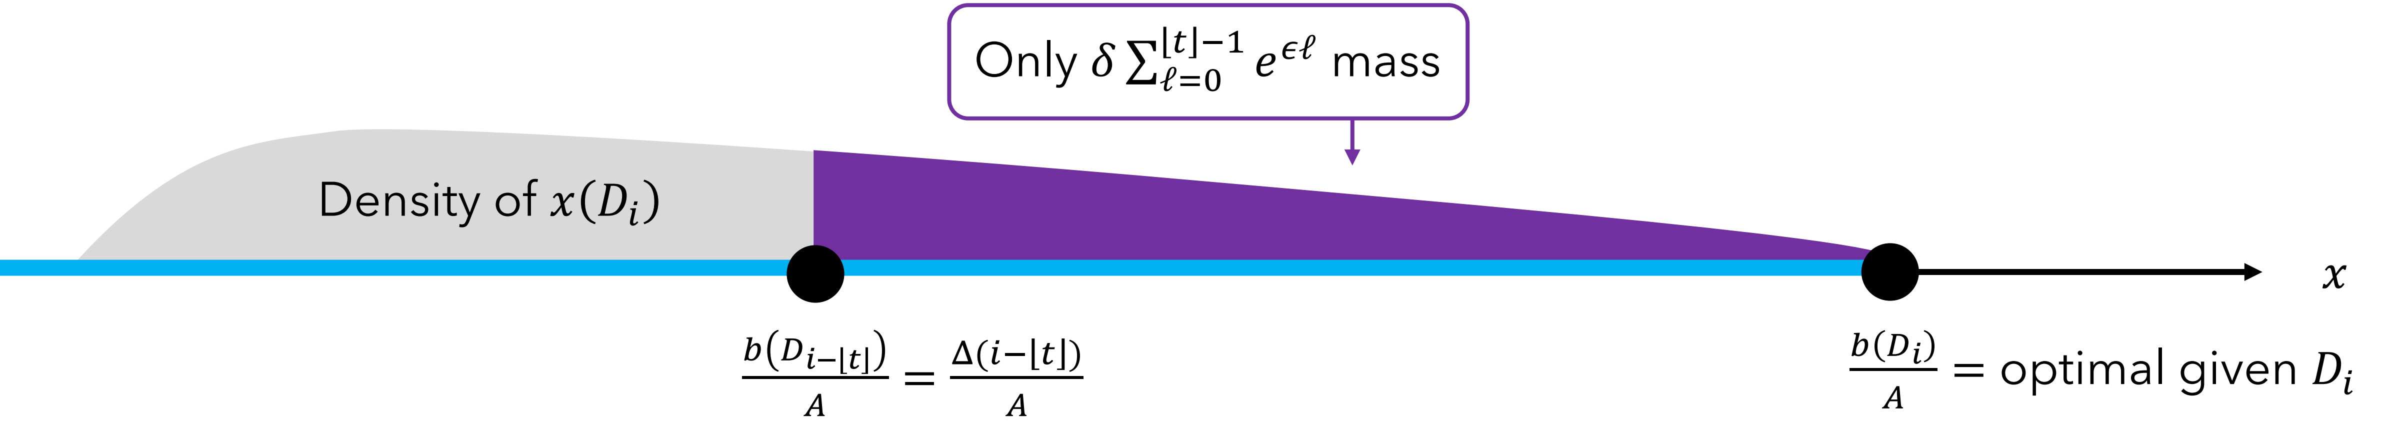

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

- For all $i \in \mathbb{Z}$, let D_i be a database with $D_i \sim D_{i+1}$ & $b(D_i) = \Delta i$
- For any $V \subseteq \mathbb{R}$, $\mathbb{P}[x(D_i) \in V] \leq e^{\varepsilon}\mathbb{P}[x(D_{i-1}) \in V] + \delta$

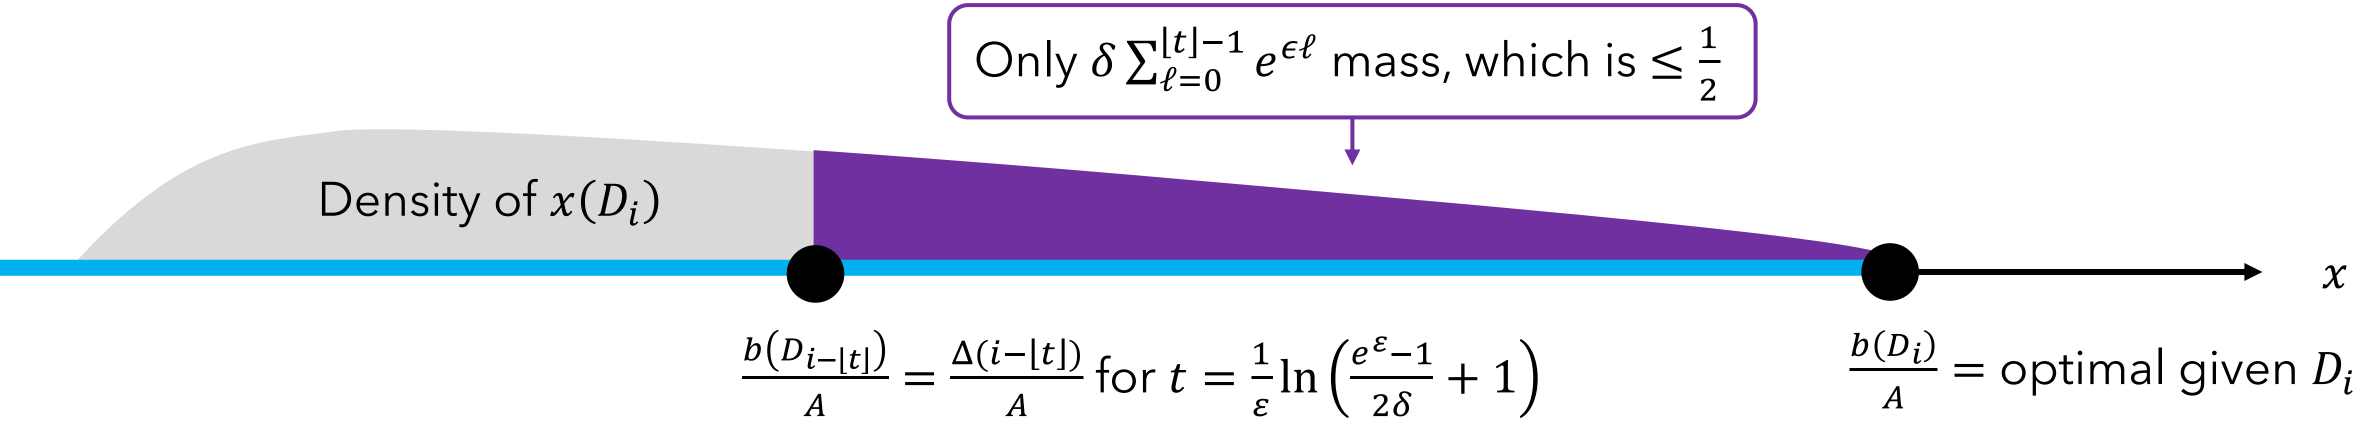

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p\geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon}\ln\left(\frac{e^{\varepsilon}-1}{2\delta}+1\right)$

Proof sketch for 1D special case (max $g(x) = x$ s.t. $Ax \leq b(D)$):

Law of total exp.: $\mathbb{E}[g(x(D_i))] \leq \frac{\Delta i}{A} - \frac{\Delta\lfloor t\rfloor}{A} \cdot \mathbb{P}[x(D_i) \leq \lfloor t\rfloor]$

$$\leq g(x^*) - \frac{\Delta t}{4A}$$

$$= g(x^*) - \frac{\Delta t}{4} \cdot \alpha_{q,1}(A) \qquad (\forall q)$$

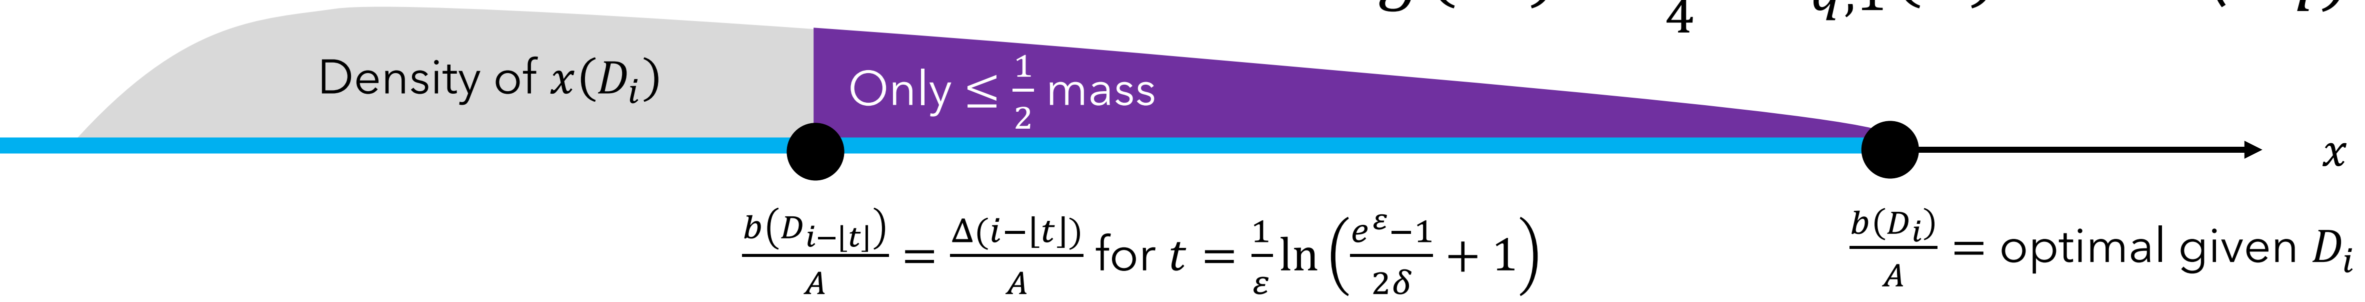

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch: Diagonal matrix A with entries $a_1, \ldots, a_m > 0$

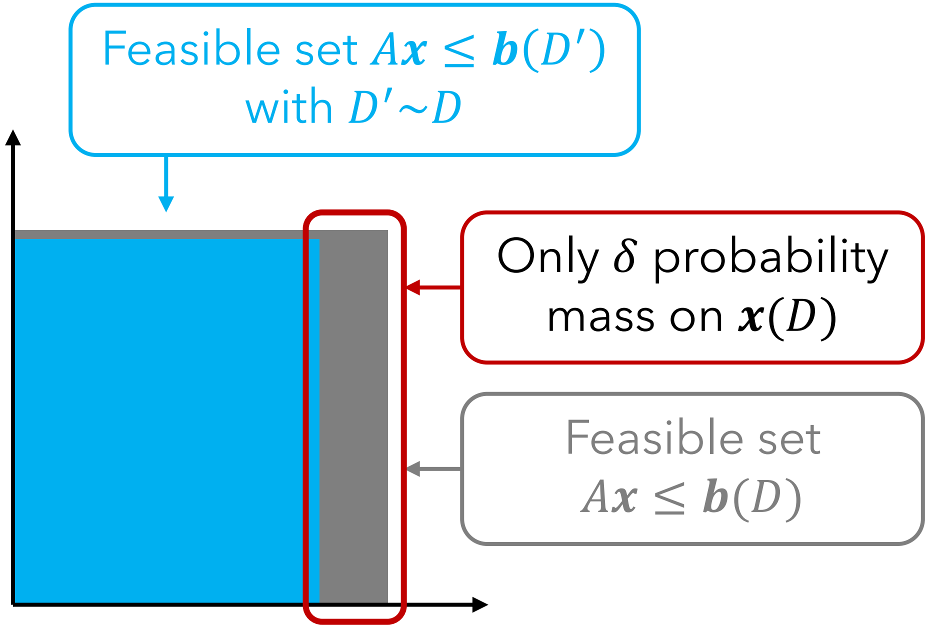

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch: Diagonal matrix A with entries $a_1, \ldots, a_m > 0$

- $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \left(\sum \frac{1}{a_i}\right) \cdot \frac{\Delta}{4\varepsilon} \cdot \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

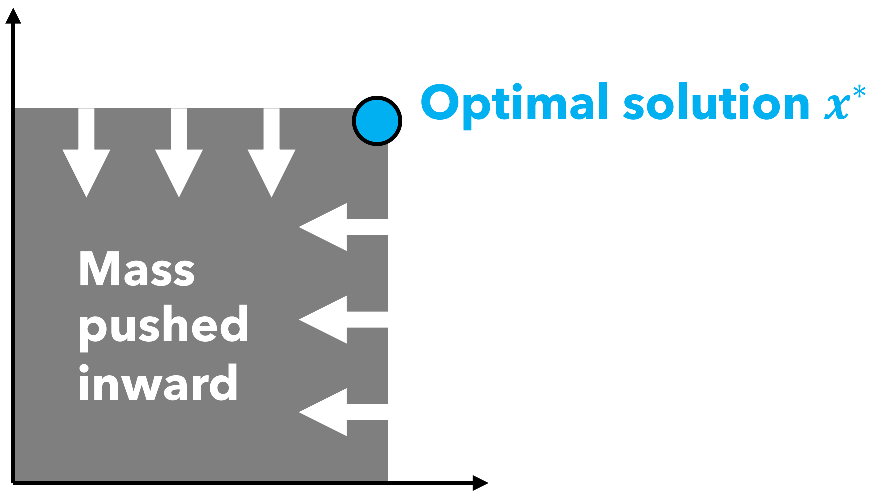

Lower bound: Proof sketch

Theorem: $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \inf_{p \geq 1}\{\alpha_{p,1}(A)\sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Proof sketch: Diagonal matrix A with entries $a_1, \ldots, a_m > 0$

- $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \left(\sum \frac{1}{a_i}\right) \cdot \frac{\Delta}{4\varepsilon} \cdot \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$
- $\alpha_{\infty,1}(A) = \sup_{\boldsymbol{u} \geq 0}\{\|\boldsymbol{u}\|_1 : \|A^T \boldsymbol{u}\|_\infty = 1\} = \sum \frac{1}{a_i}$
- $g(\boldsymbol{x}^*) - \mathbb{E}[g(\boldsymbol{x}(D))] \geq \alpha_{\infty,1}(A) \cdot \sqrt[\infty]{m} \cdot \frac{\Delta}{4\varepsilon} \cdot \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

  $\geq \inf_{p \geq 1}\{\alpha_{p,1}(A) \cdot \sqrt[p]{m}\} \cdot \frac{\Delta}{4\varepsilon} \cdot \ln\left(\frac{e^{\varepsilon}-1}{2\delta} + 1\right)$

Outline

1. Introduction
2. Background: Differential privacy
3. Algorithm
4. Lower bound
5. **Experiments**
6. Conclusion

Experiments with Dow Jones data

Individuals pool money to invest

- Amount private except to investment manager

Goal: Minimize variance subject to minimum expected return

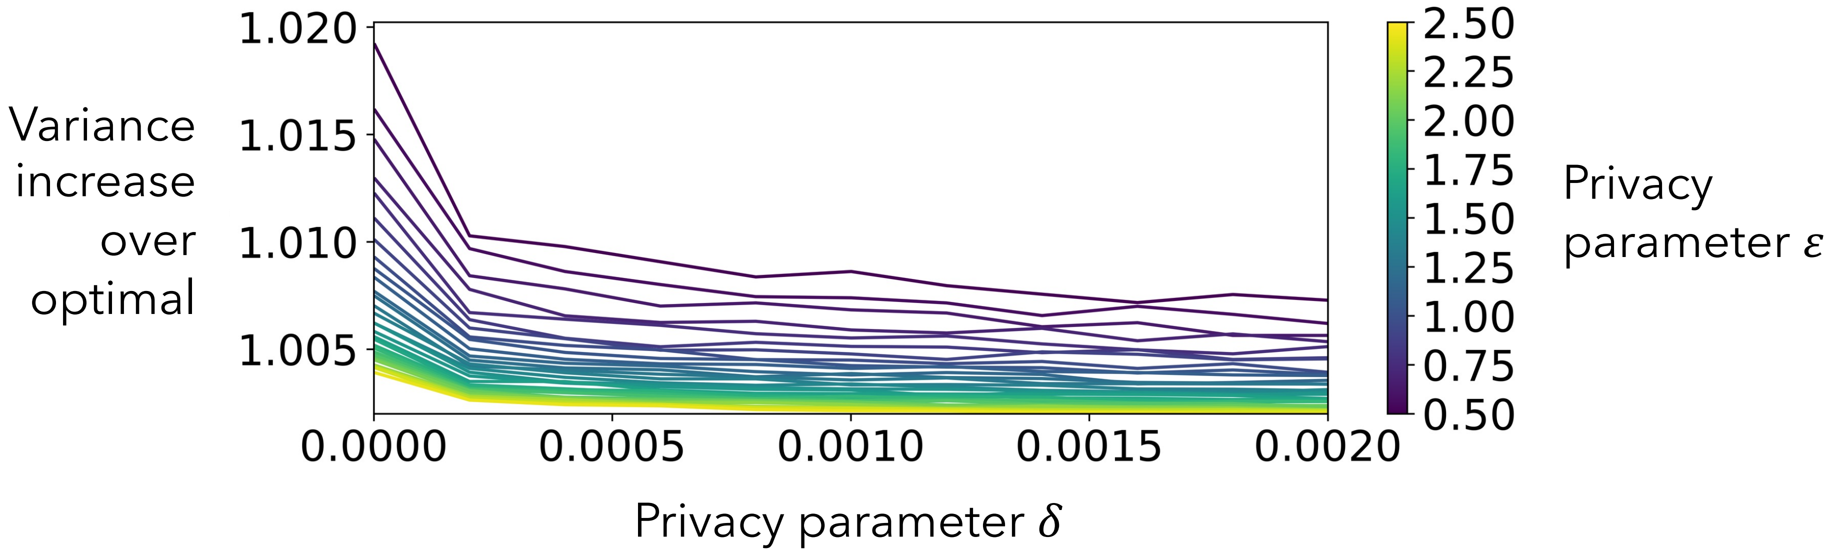

Outline

1. Introduction
2. Background: Differential privacy
3. Algorithm
4. Lower bound
5. Experiments
6. **Conclusion**

Conclusions

Algorithm for linearly-constrained optimization

Solution never violates the constraints

Algorithm's loss is optimal up to log factors

Future research: What if matrix A is private?