Global differential privacy

Material taken from here, here

Final project

- Sign up for a check-in meeting with me this week
- Final project deliverables
  - Presentation (50%)
    - Talk about problem setup/motivation, technique, evaluation
    - Peer grading
  - 6-page writeup (50%) due December 10

Last class

- How to publish data about someone *while protecting their privacy?*
- DP promises to protect an individual from any additional harm they might face due to their data being included in a database
- Mechanism: randomized response satisfies ϵ-differential privacy

Global differential privacy

- Single trusted party who collects data from users
- Trusted party will generate a noisy answer to a query to protect user privacy
- Examples:
  - U.S. Census
  - A large tech company releases a model computed on user data
- Less noise is needed compared to local differential privacy

Global DP mechanism: Laplace distribution

- Laplace distribution is commonly used in global DP
    - Probability density function: Lap$(x \mid b) = \frac{1}{2b} e^{-|x|/b}$, where b is the scale

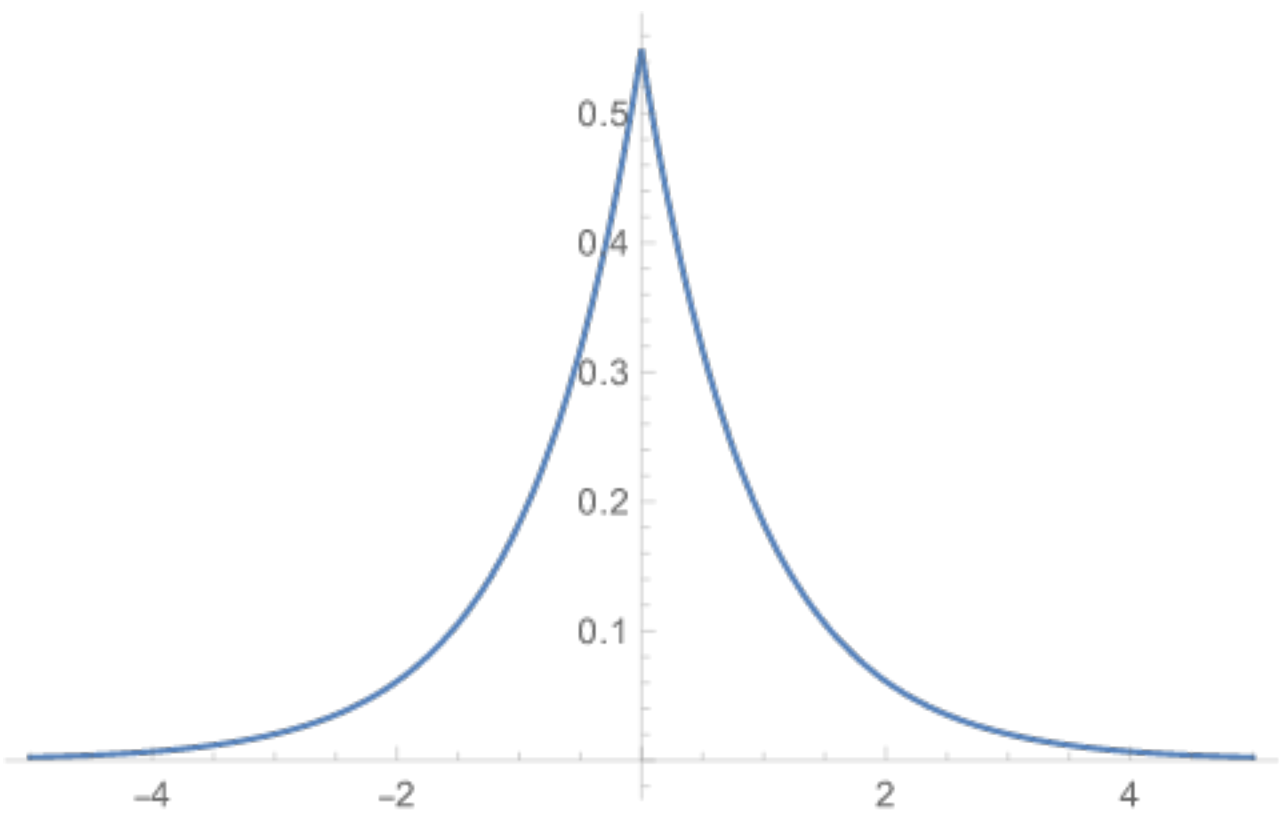

- **Definition:** Given any function $f : \mathbb{N}^{|X|} \rightarrow \mathbb{R}^k$, *Laplace mechanism* is defined as $M_L(x, f(\cdot), \epsilon) = f(x) + (Y_1, \cdots, Y_k)$ where Y_i are i.i.d. random variables drawn from Lap$(\Delta f / \epsilon)$
    - $\Delta f = ||x - y||_1$ is the *sensitivity*, or the magnitude by which a single individual's data can change the function f in the worst case

Example: counting queries

- Want to publish how many people in a database satisfies a given condition, e.g., how many wear glasses?
  - Do a normal count
  - Sensitivity of a counting query is 1, so add noise drawn from Lap$(1/\epsilon)$ to the result
- What if you want to publish the number of complaints received on a given day?
  - Someone very unhappy could send in five complaints
  - Sensitivity is higher, so need to add more noise: Lap$(5/\epsilon)$

Privacy of Laplace mechanism

- **Theorem:** The Laplace mechanism preserves ϵ-differential privacy.
- **Proof:**
  - Let $x \in \mathbb{N}^{|X|}$ and $y \in \mathbb{N}^{|X|}$ be such that $\|x - y\|_1 \leq 1$, and let $f(\cdot)$ be some function. Let p_x denote the probability density function of $M_L(x, f, \epsilon)$, and let p_y denote the same for y.
  - We compare the two at some arbitrary point $z \in \mathbb{R}^k$:

$$\frac{p_x(z)}{p_y(z)} = \prod_{i=1}^{k} \left(\frac{\exp(-\epsilon\,|f(x)_i - z_i|/\Delta f)}{\exp(-\epsilon\,|f(y)_i - z_i|/\Delta f}\right) = \prod_{i=1}^{k} (\exp(\frac{\epsilon(|f(y)_i - z_i| - |f(x)_i - z_i|)}{\Delta f})$$

$$\leq \prod_{i=1}^{k} \exp(\frac{\epsilon\,|f(x)_i - f(y)_i|}{\Delta f})$$ (triangle inequality)

$$= \exp(\frac{\epsilon \cdot \|f(x) - f(y)\|_1}{\Delta f}) \leq \exp(\epsilon)$$ (sensitivity definition)

What about other distributions?

- Can one use Gaussian distribution?
  - Yes, but only under a more general definition of DP
- A mechanism A is (ϵ, δ)-differentially private if for all neighboring databases D_1, D_2, and all sets S of outputs
$P[A(D_1) \in S] \leq e^{\epsilon} \cdot P[A(D_2) \in S] + \delta$
- Approximate DP instead of pure DP
  - Laplace is a $(\epsilon,0)$-DP mechanism

Approximate DP

- What does δ mean?
  - Intuitively, can think of it as the “probability that something goes wrong”
- Privacy loss: $\mathscr{L}_{D_1,D_2}(O) = \ln\left(\dfrac{P[A(D_1) = O)]}{P[A(D_1) = O)]}\right)$
- x-axis: all events according to their probabilities, y-axis: $\exp(\mathscr{L})$
  - Laplace: privacy loss always within ϵ
  - Gaussian: some chance for “bad events”, where the privacy loss to be greater than ϵ!

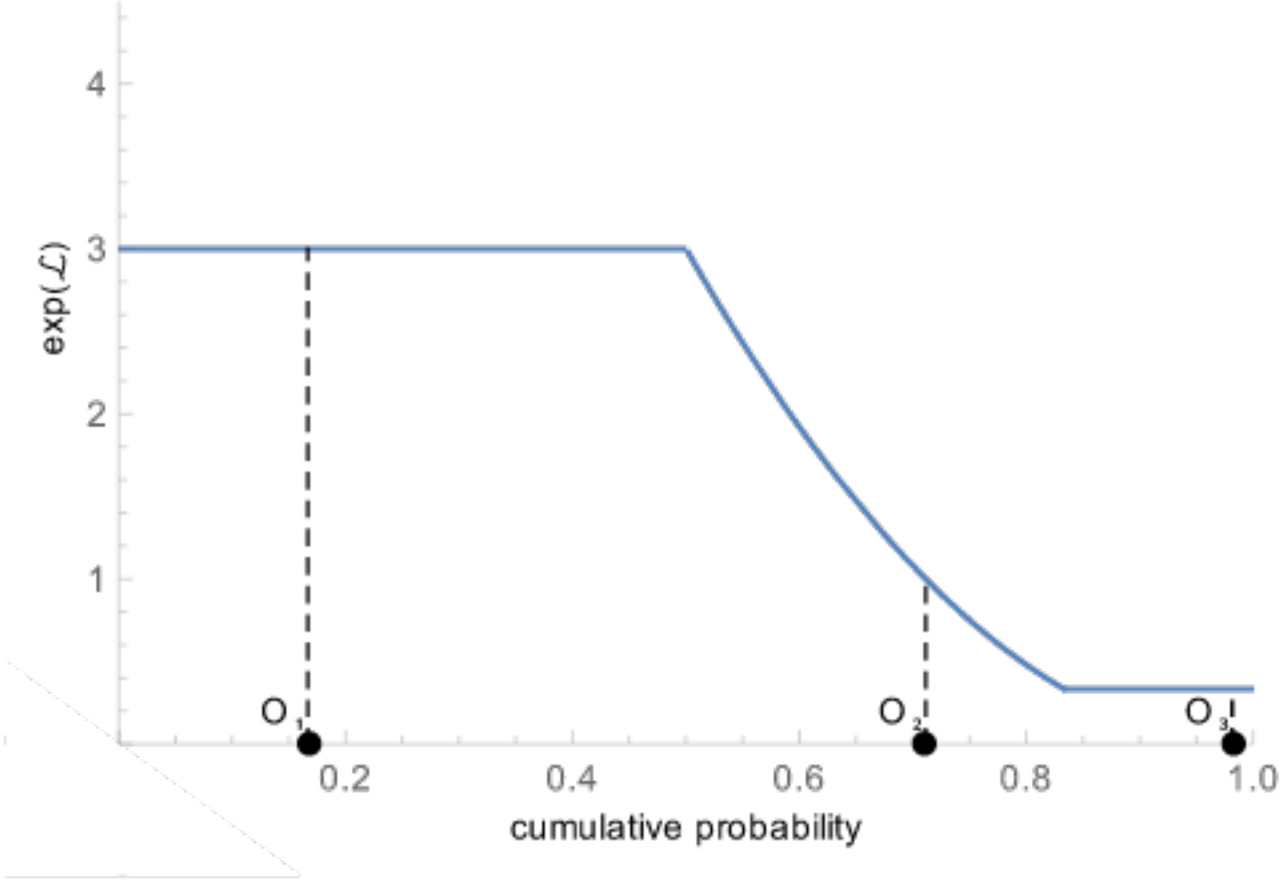

Laplace

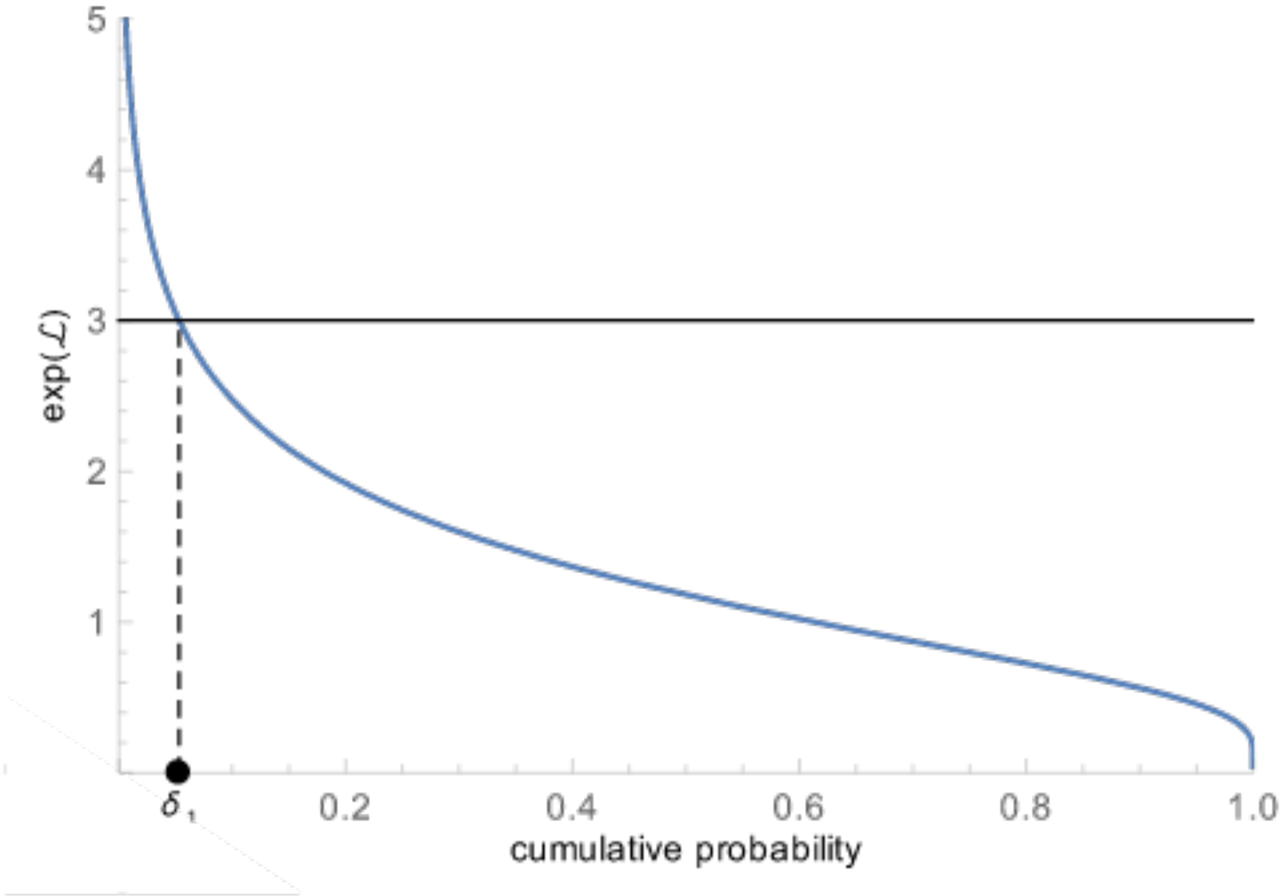

Gaussian

Approximate DP

- How to quantify the privacy loss when there are bad events?
- The blue area is the actual value of δ, which is the mass of all possible bad events
- If a mechanism causes a distinguishing event where there is no privacy, then the ratio is 0
- How to set δ?
  - Usually set $\delta < 1/n$
- Why use Gaussian distribution at all?
  - Noise scales well with sensitivity (square root instead of linear)

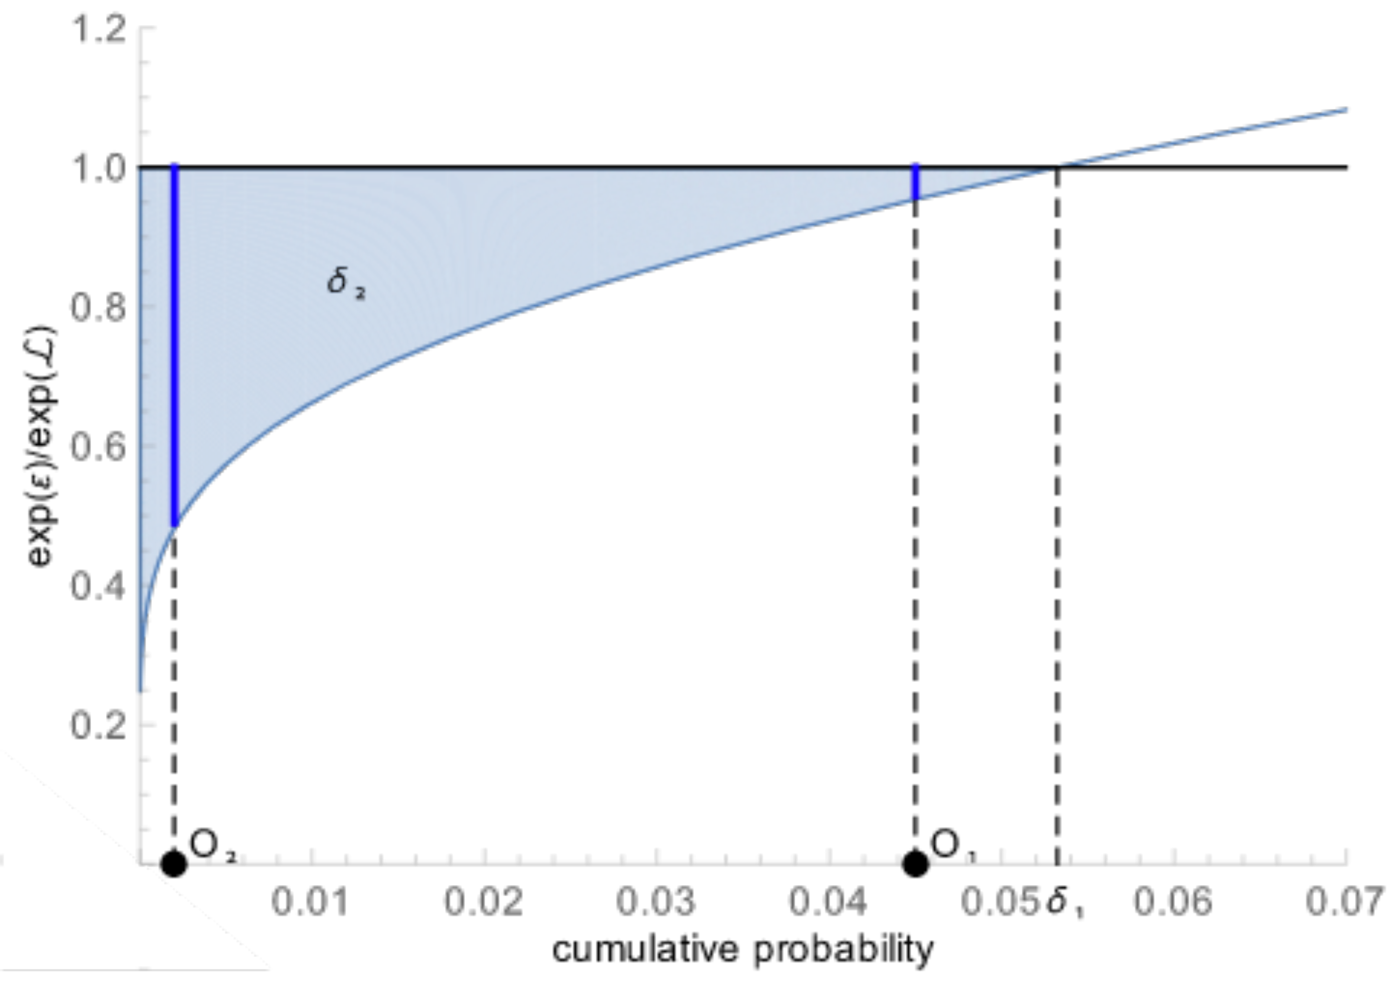

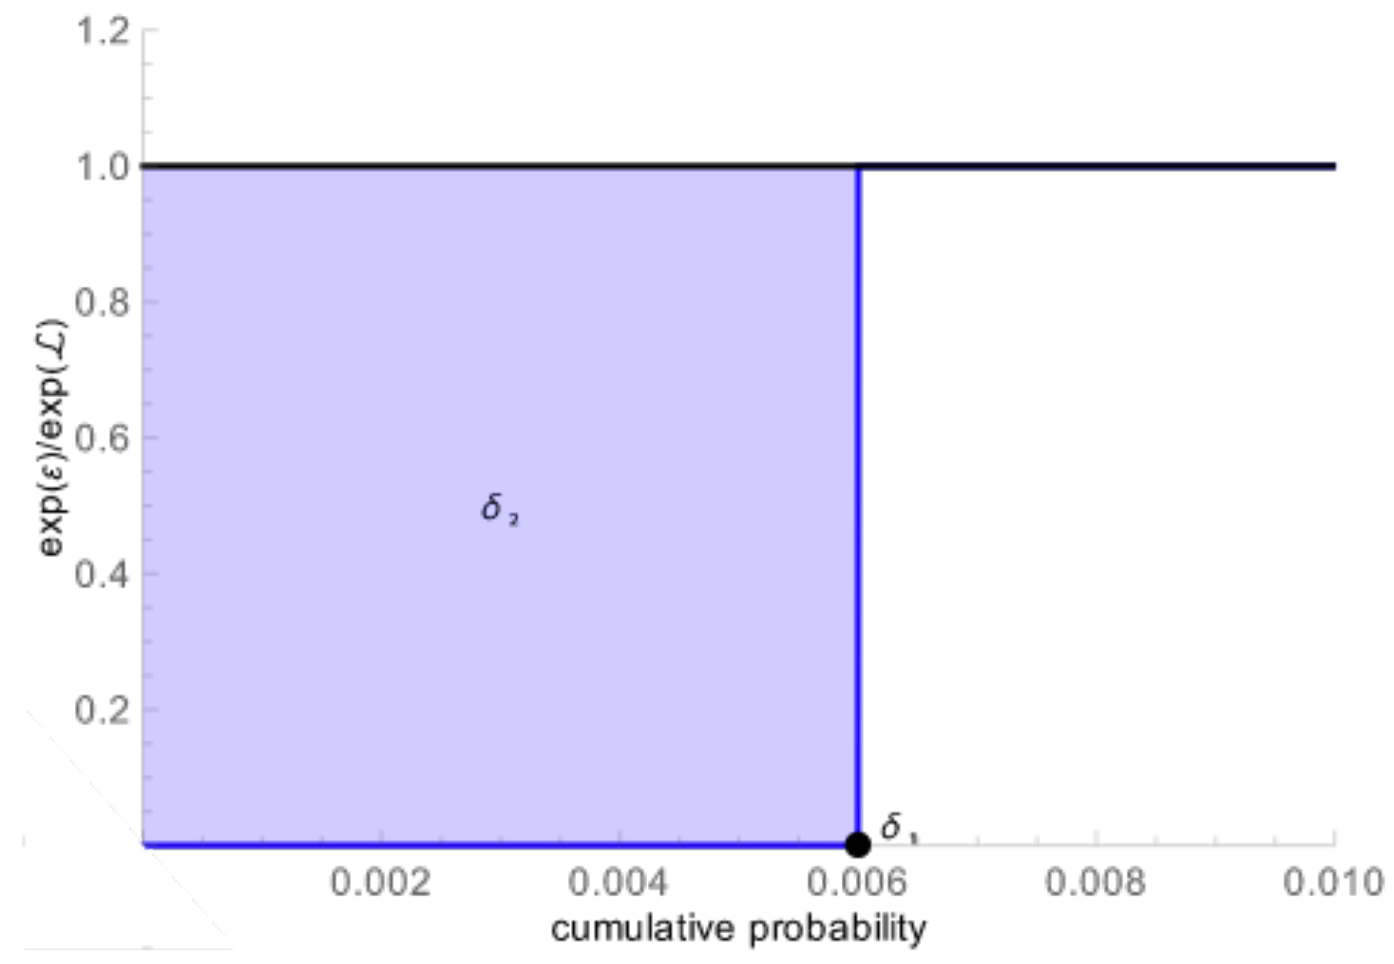

Today: Sage

Next class: Bitcoin & Ethereum

- Bitcoin was created by Satoshi Nakamoto in 2009
- Cryptographic currency to remove trust from institutions
- Two core components
  - Immutable & public ledger
  - Cryptographic transactions
- We will see some basic & hardcore crypto used!