CSE660 Differential Privacy

September 18, 2017

Marco Gaboardi
Room: 338-B
gaboardi@buffalo.edu
http://www.buffalo.edu/~gaboardi

(ε,δ)-Differential Privacy

Definition

Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon) \Pr[Q(b_2) \in S] + \delta$$

Randomized Response

Algorithm 1 Pseudo-code for Randomized Response

function RandomizedResponse(D, q, ϵ)

  for $k \leftarrow 1$ to $|D|$ **do**

$$S_i \leftarrow \begin{cases} q(d_i) & \text{with probability } \frac{e^\epsilon}{1+e^\epsilon} \\ \neg q(d_i) & \text{with probability } \frac{1}{1+e^\epsilon} \end{cases}$$

  end for

  return $\frac{(\textbf{sum } S)}{|D|}$

end function

Example

Let's consider a medical dataset containing informations on whether each patient has a disease.

We can have $q(d_i)=1$ if patient i has the disease and $q(d_i)=0$ otherwise.

We can use randomized response to estimate the proportion of patient that have the disease.

The noise that each individual adds protect his/her value.

Randomized Response

Privacy Theorem:
Randomized response is ε-differentially private.

Accuracy Theorem:

$$\Pr_{r \leftarrow RR(D,q,\epsilon)}\left[\left|\frac{1+e^{\epsilon}}{e^{\epsilon}-1}\left(r-\frac{1}{1+e^{\epsilon}}\right)-q(D)\right| \geq \frac{1+e^{\epsilon}}{(e^{\epsilon}-1)}\sqrt{\frac{\log(2/\beta)}{2n}}\right] \leq \beta$$

Noise on Input vs Noise on Output

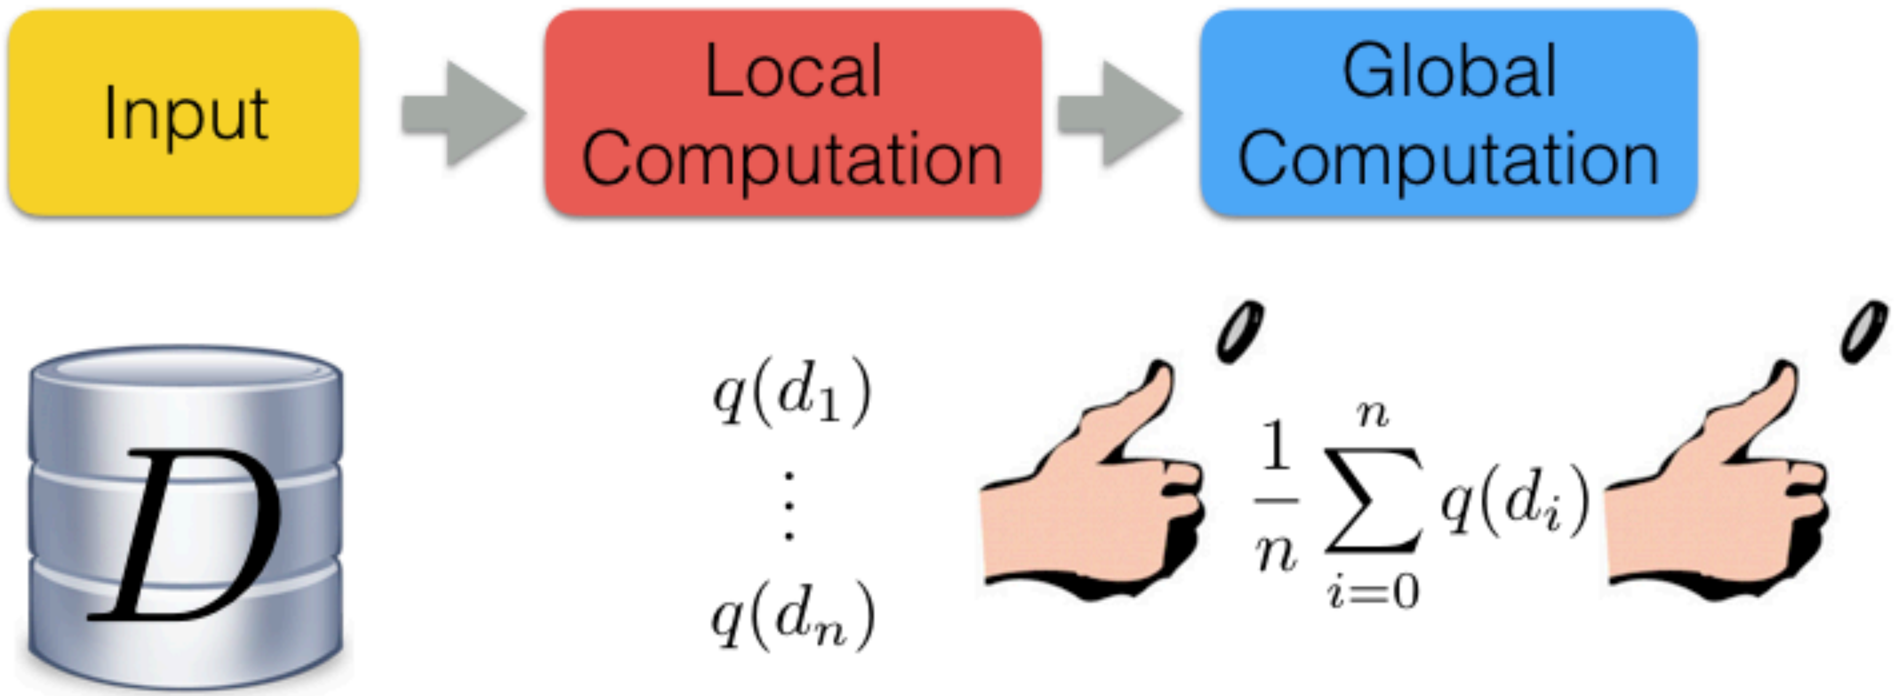

Noise on the output

Question: What is a good way to add noise to the output of a statistical query?

Intuitive answer: it depends on ε or the accuracy we want to achieve, and on the scale that a change of an individual can have on the output.

Global Sensitivity

Definition 1.8 (Global sensitivity). The *global sensitivity* of a function $q : \mathcal{X}^n \to \mathbb{R}$ is:

$$\Delta q = \max \left\{ |q(D) - q(D')| \;\middle|\; D \sim_1 D' \in \mathcal{X}^n \right\}$$

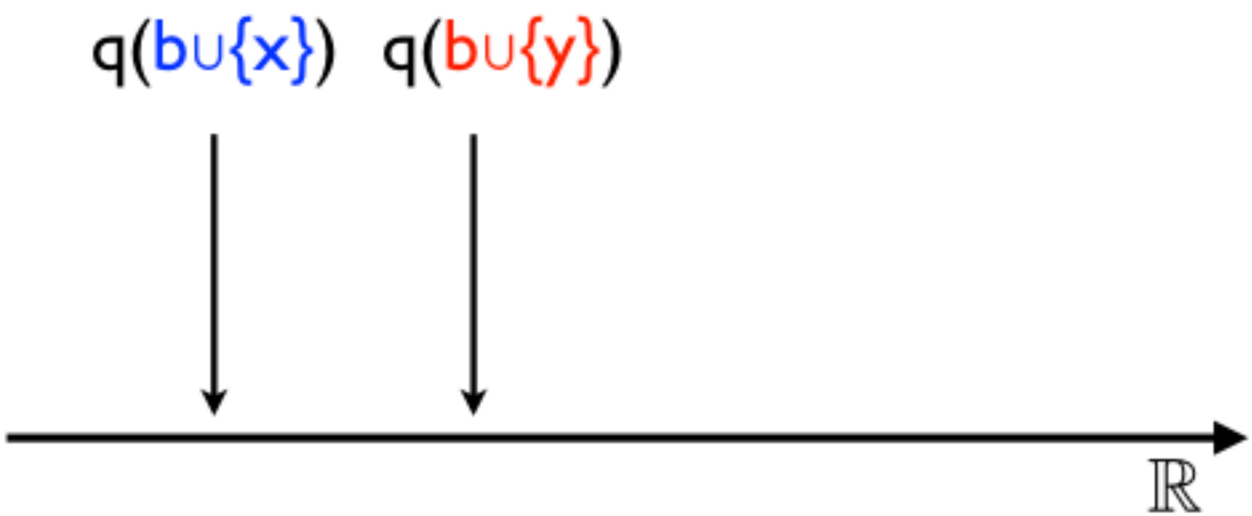

Laplace Distribution

$$\text{Lap}(b, \mu)(X) = \frac{1}{2b} \exp\left(-\frac{|\mu - X|}{b}\right)$$

b regulates the skewness of the curve,

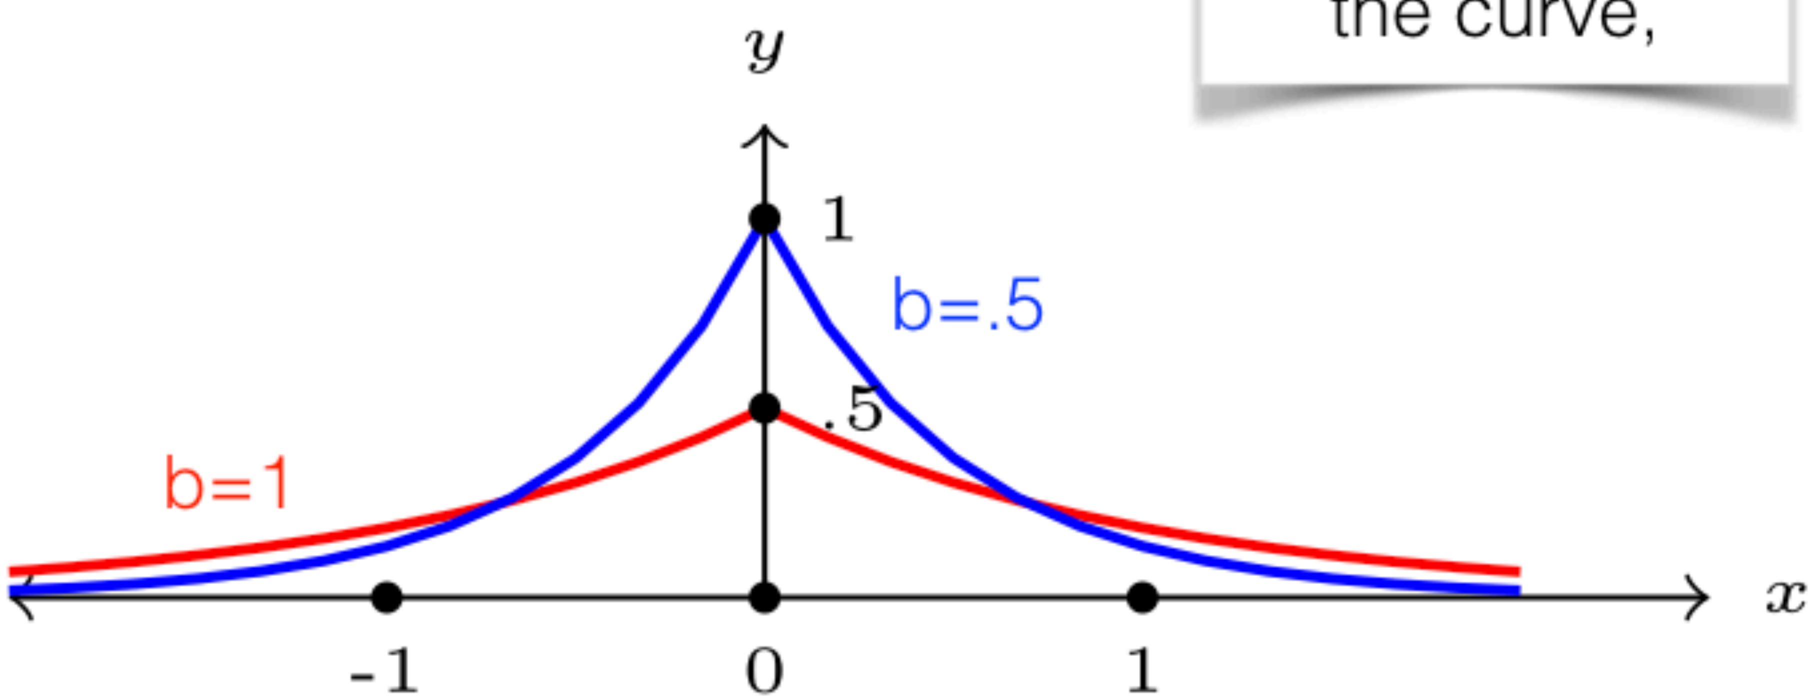

Laplace Mechanism

Algorithm 2 Pseudo-code for the Laplace Mechanism

```
function LapMech(D, q, ε)
    Y ←$ Lap(Δq/ε)(0)
    return q(D) + Y
end function
```

Laplace Mechanism

Theorem (Privacy of the Laplace Mechanism)
The Laplace mechanism is ε-differentially private.

Proof: Intuitively

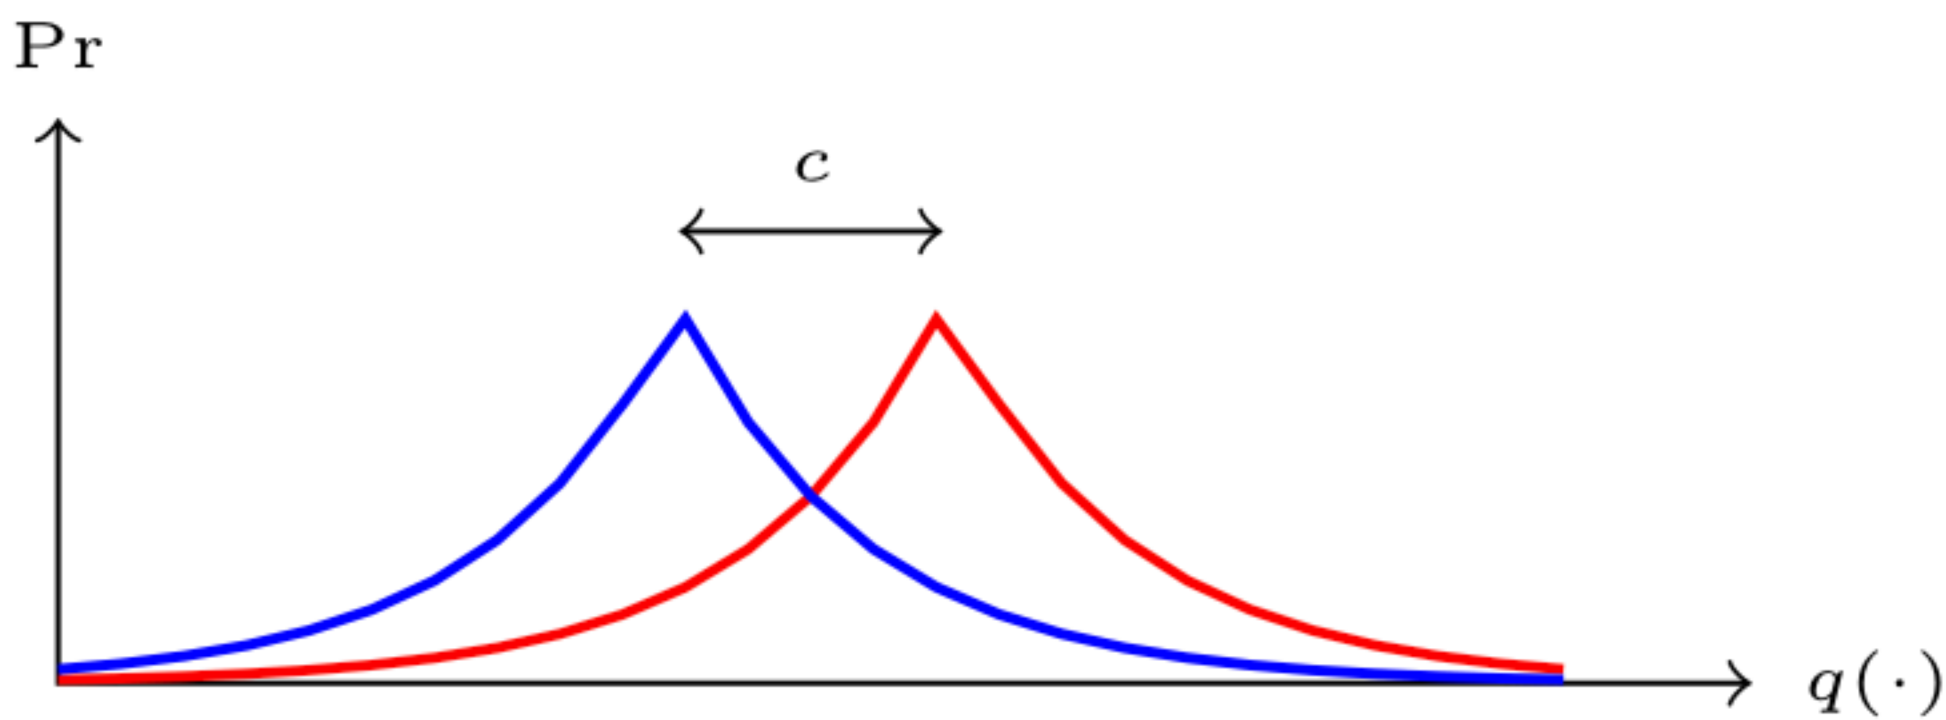

Laplace Mechanism

Theorem (Privacy of the Laplace Mechanism)
The Laplace mechanism is ε-differentially private.

Proof:

Consider $D \sim_1 D' \in \mathcal{X}^n$, $q : \mathcal{X}^n \to \mathbb{R}$, and let p and p' denote the probability density function of $\mathsf{LapMech}(D, q, \epsilon)$ and $\mathsf{LapMech}(D', q, \epsilon)$ We compare them at an arbitrary point $z \in \mathbb{R}$.

$$\frac{p(z)}{p'(z)} = \frac{\exp\left(-\frac{\epsilon|q(D)-z|}{\Delta q}\right)}{\exp\left(-\frac{\epsilon|q(D')-z|}{\Delta q}\right)}$$

Laplace Mechanism

Theorem (Privacy of the Laplace Mechanism)
The Laplace mechanism is ε-differentially private.

Continued proof:

$$\frac{p(z)}{p'(z)} = \frac{\exp\left(-\frac{\epsilon|q(D)-z|}{\Delta q}\right)}{\exp\left(-\frac{\epsilon|q(D')-z|}{\Delta q}\right)}$$

$$= \exp\left(\frac{\epsilon(|q(D')-z|-|q(D)-z|)}{\Delta q}\right)$$

$$\leq \exp\left(\frac{\epsilon(|q(D')-q(D)|)}{\Delta q}\right)$$

$$\leq \exp(\epsilon)$$

Similarly, we can prove that $\exp(-\epsilon) \leq \frac{p(z)}{p'(z)}$

Example Revisited

Let's consider again a medical dataset containing informations on whether each patient has a disease.
We can have $q(d_i)=1$ if patient i has the disease and $q(d_i)=0$ otherwise.

We first compute the proportion of patients that have the disease. Notice that this has sensitivity 1/n.

Then we can add Laplace noise proportional to $1/\varepsilon n$.

The noise protects each individual value.

Laplace Mechanism

Question: How accurate is the answer that we get from the Laplace Mechanism?

Laplace Mechanism

Accuracy Theorem: let $r = \text{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right) \ln\left(\frac{1}{\beta}\right)\right] = \beta$$

This represents the variable measuring the difference between the noised answer and the non-noised one.

This is our alpha. Notice that we express it in terms of beta.

Laplace Mechanism

Accuracy Theorem: let $r = \mathsf{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right)\ln\left(\frac{1}{\beta}\right)\right] = \beta$$

Proof: By definition of the Laplace mechanism we have:

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right)\ln\left(\frac{1}{\beta}\right)\right] = \Pr\left[|Y| \geq \left(\frac{\Delta q}{\epsilon}\right)\ln\left(\frac{1}{\beta}\right)\right]$$

where Y is drawn from Lap(Δq/ε)(0)

Tail bound for the Laplace Distribution

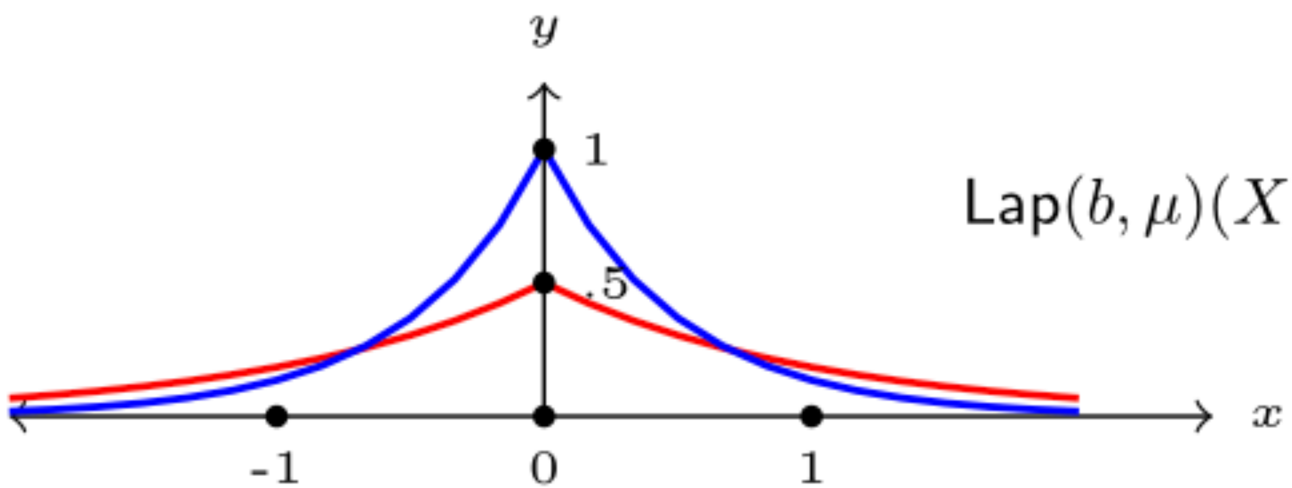

$$\text{Lap}(b,\mu)(X) = \frac{1}{2b}\exp\left(-\frac{|\mu - X|}{b}\right)$$

$$\Pr\Big[|X| \geq b\,t\Big] = \exp(-t)$$

Laplace Mechanism

Accuracy Theorem: let $r = \mathsf{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right)\ln\left(\frac{1}{\beta}\right)\right] = \beta$$

Continued proof: applying this bound we get:

$$\Pr\left[|Y| \geq \left(\frac{\Delta q}{\epsilon}\right)\ln\left(\frac{1}{\beta}\right)\right] = \exp\left(-\ln\left(\frac{1}{\beta}\right)\right) = \beta$$

Laplace Mechanism

Accuracy Theorem: let $r = \mathsf{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right) \ln\left(\frac{1}{\beta}\right)\right] = \beta$$

Intuitive reading: with high probability we have:

$$\Big|q(D) - r\Big| \leq O\Big(\frac{1}{n}\Big)$$

Example revisited

Accuracy Theorem: let $r = \mathsf{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right)\ln\left(\frac{1}{\beta}\right)\right] = \beta$$

Let's consider again the example of the medical dataset containing information on whether each patient has a disease or not ($q(d_i)=1$ or $q(d_i)=0$).

We can use the Laplace Mechanism to estimate the proportion of patients that have the disease.
We need to fix the parameters.

Example revisited

Accuracy Theorem: let $r = \mathsf{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right) \ln\left(\frac{1}{\beta}\right)\right] = \beta$$

Let's fix the following values for the parameters:

n=1,000,000
ε=1
β=0.05
$\ln(\frac{1}{\beta}) = 2.99$

Question: What is the sensitivity?

$$\Delta q = 10^{-6}$$

Example revisited

Accuracy Theorem: let $r = \mathsf{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right) \ln\left(\frac{1}{\beta}\right)\right] = \beta$$

With this set of parameters we have with 95% confidence

$$r - 0.0000299 \leq q(D) \leq r + 0.0000299.$$

Randomized Response vs Laplace

Accuracy for Randomize response: with high probability we have:

$$\Big|r - q(D)\Big| \leq O\Big(\frac{1}{\sqrt{n}}\Big)$$

Accuracy for Laplace: with high probability we have:

$$\Big|q(D) - r\Big| \leq O\Big(\frac{1}{n}\Big)$$

Some important properties

- Resilience to post-processing
- Group privacy
- Composition

Resilience to Post-processing

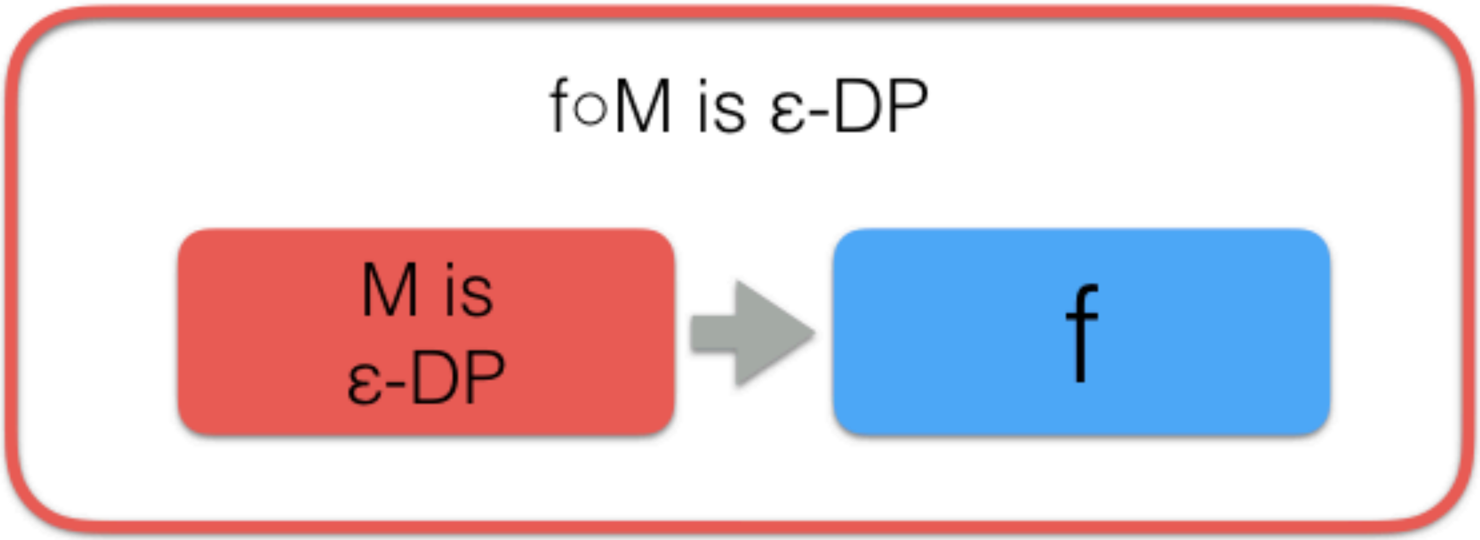

Resilience to Post-processing

Proposition 1.1 (Post-processing)**.** Let $\mathcal{M} : \mathcal{X}^n \to R$ be a randomized algorithm that is ϵ-differentially private. Let $f : R \to R'$ be an arbitrary deterministic mapping. Then $f \circ \mathcal{M} : \mathcal{X}^n \to R'$ is also ϵ-differentially private.

Proof. Fix any pair of neighboring databases $D \sim_1 D'$, and fix any event $S \subseteq R'$. Let $T = \{r \in R : f(r) \in S\}$. We have

$$
\begin{aligned}
\Pr[f(\mathcal{M}(D)) \in S] &= \Pr[\mathcal{M}(D) \in T] \\
&\leq \exp(\epsilon) Pr[\mathcal{M}(D') \in T] \\
&= \exp(\epsilon) Pr[f(\mathcal{M}(D')) \in S]
\end{aligned}
$$

Resilience to Post-processing

Question: Why is resilience to post-processing important?

Answer: Because it is what allows us to publicly release the result of a differentially private analysis!

Group Privacy

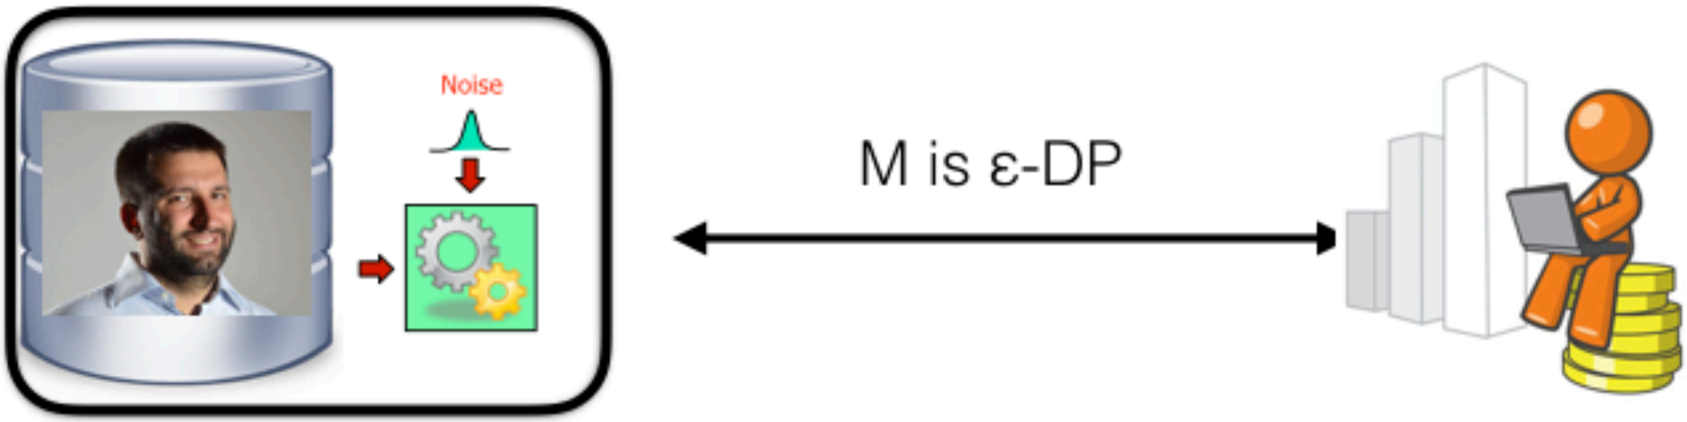

$$\Pr[\mathcal{M}(D) = r] \leq e^{\epsilon} \Pr[\mathcal{M}(D') = r]$$

Group Privacy

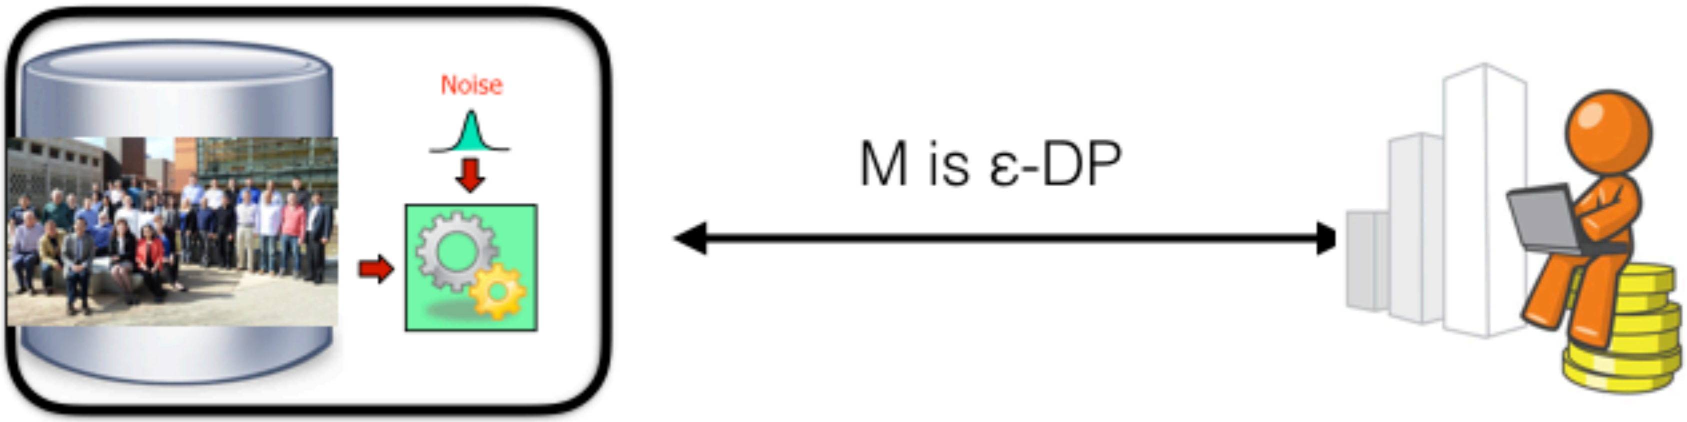

$$\Pr[\mathcal{M}(D) \in S] \leq \exp(k\epsilon) \Pr[\mathcal{M}(D') \in S]$$

Group Privacy

Proposition 1.2 (Group Privacy). Let $\mathcal{M} : \mathcal{X}^n \to R$ be a randomized algorithm that is ϵ-differentially private. Then, $\mathcal{M}$ is $k\epsilon$-differentially private for groups of size k. That is, for datasets $D, D' \in \mathcal{X}^n$ such that $D \Delta D' \leq k$ and for all $S \subseteq R$ we have

$$\Pr[\mathcal{M}(D) \in S] \leq \exp(k\epsilon) \Pr[\mathcal{M}(D') \in S]$$

Proof. Fix any pair of databases D, D' with $D \Delta D' \leq k$. Then, we have databases $D_0, D_1, \ldots, D_k$ such that $D_0 = D$, $D_k = D'$ and $D_i \Delta D_{i+1} \leq 1$. Fix also any event $S \subseteq R'$. Then, we have have

$$\begin{aligned}
\Pr[\mathcal{M}(D) \in S] &= \Pr[\mathcal{M}(D_0) \in S] \\
&\leq \exp(\epsilon) \Pr[\mathcal{M}(D_1) \in S] \\
&\leq \exp(\epsilon)(\exp(\epsilon) \Pr[\mathcal{M}(D_2) \in S]) = \exp(2\epsilon) \Pr[\mathcal{M}(D_2) \in S] \\
&\leq \cdots \\
&\leq \exp(k\epsilon) \Pr[\mathcal{M}(D_k) \in S] = \exp(k\epsilon) \Pr[\mathcal{M}(D') \in S]
\end{aligned}$$

Group Privacy

Question: Why is group privacy important?

Answer: Because it allows to reason about privacy at different level of granularities!