Differential Privacy

Some other basic mechanisms

Marco Gaboardi
Boston University

The opinions expressed in this course are mine and they do not not reflect those of the National Science Foundation or the US. Census Bureau.

(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 $S \subseteq R$:

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

Laplace Mechanism

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

Proof: Intuitively

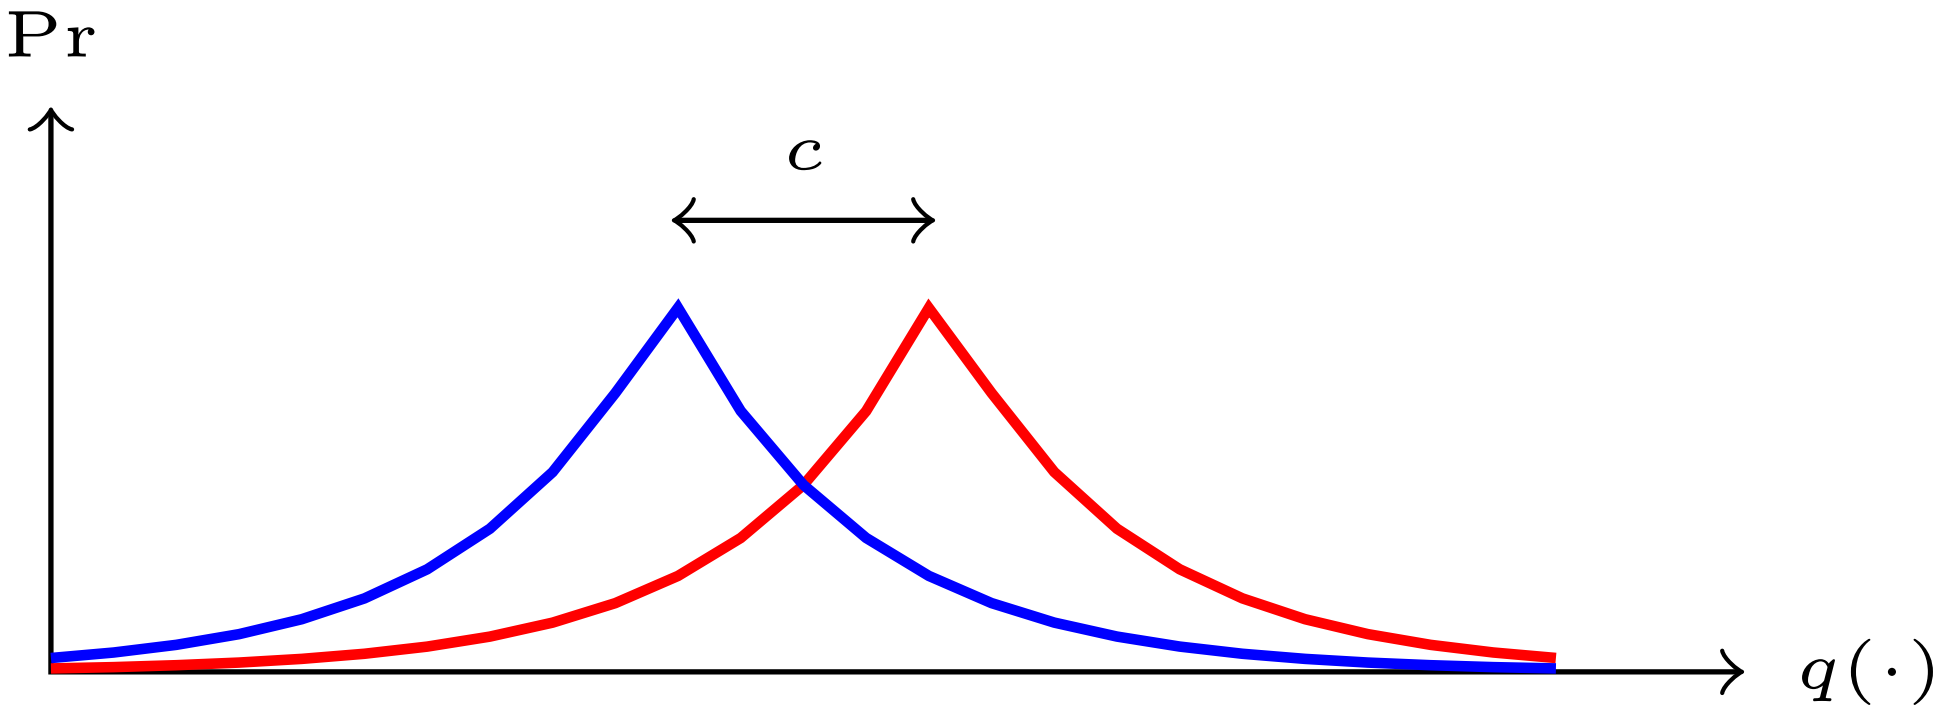

Another mechanism?

Not all the queries are numeric.

Auctions

¥1M

¥1M

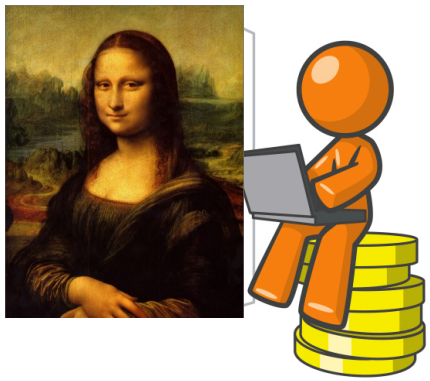

¥3M

¥1M

Auctions

¥1M

¥1M

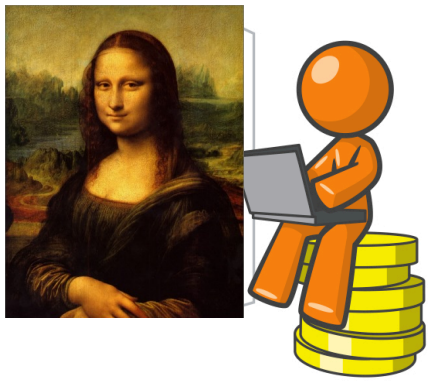

¥3M

¥1M

How shall we set the price under differential privacy?

Max response

(a)

(b)

(c)

(d)

Suppose that each one of us can vote for one star, and we want to say who is the star that receives most votes.

Max response

(a)

(b)

(c)

(d)

Suppose that each one of us can vote for one star, and we want to say who is the star that receives most votes.

The answer here is not a number, how can we release it under differential privacy?

Differentially private selection

- We want to select some element that maximize some value,
- The noise added for privacy should not destroy the utility,
- If we cannot return the maximal one, with high probability, we want to return one close to it.

Max response

(a)

(b)

(c)

(d)

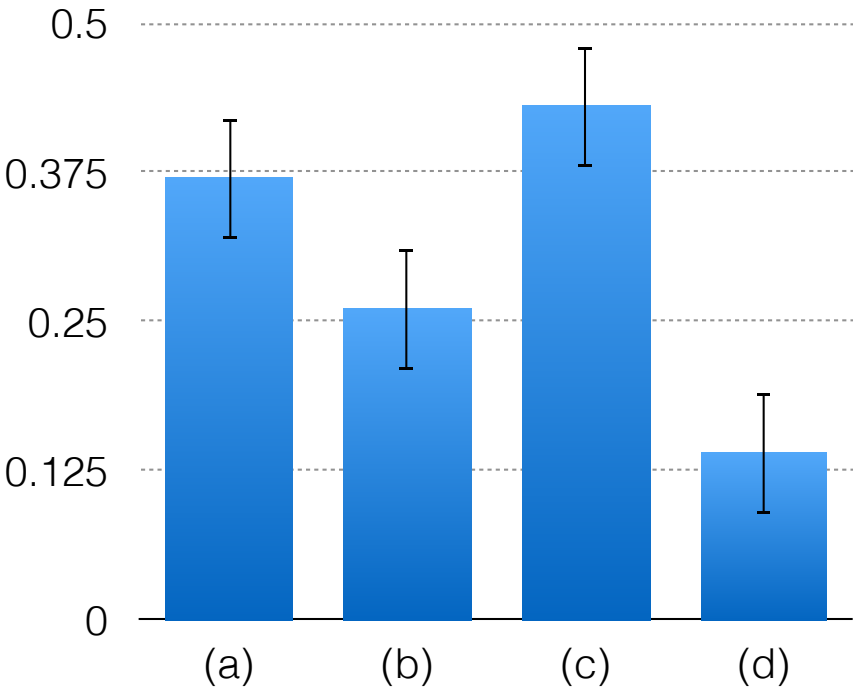

Intuition:
We could compute the histogram add Laplace noise to each score and then select the maximal noised score.

Report Noisy Max

Given a set of queries with sensitivity l, return the index of the noised query with the max value

$q_1(D)$+noise

$q_2(D)$+noise

$q_3(D)$+noise

.....

$q_k(D)$+noise

Report Noisy Max

Given a set of queries with sensitivity l, return the index of the noised query with the max value

$q_1(D)$+noise

$q_2(D)$+noise

$q_3(D)$+noise

.....

$q_k(D)$+noise

A naive analysis gives $k\varepsilon$-differentially private

Report Noisy Max - intuition

We can prove this algorithm
ε-differentially private

Report Noisy Max - intuition

We can prove this algorithm
ε-differentially private

Databases differing
in one individual

Report Noisy Max - intuition

$q_1(D)$+noise $q_1(D')$+noise

$q_2(D)$+noise $q_2(D')$+noise

$q_3(D)$+noise $q_3(D')$+noise

....

$q_k(D)$+noise $q_k(D')$+noise

Databases differing in one individual

We can prove this algorithm ε-differentially private

Report Noisy Max - intuition

l sensitive queries

$q_1(D)$+noise $q_1(D')$+noise

$q_2(D)$+noise $q_2(D')$+noise

$q_3(D)$+noise $q_3(D')$+noise

....

$q_k(D)$+noise $q_k(D')$+noise

Databases differing in one individual

We can prove this algorithm ε-differentially private

Report Noisy Max - intuition

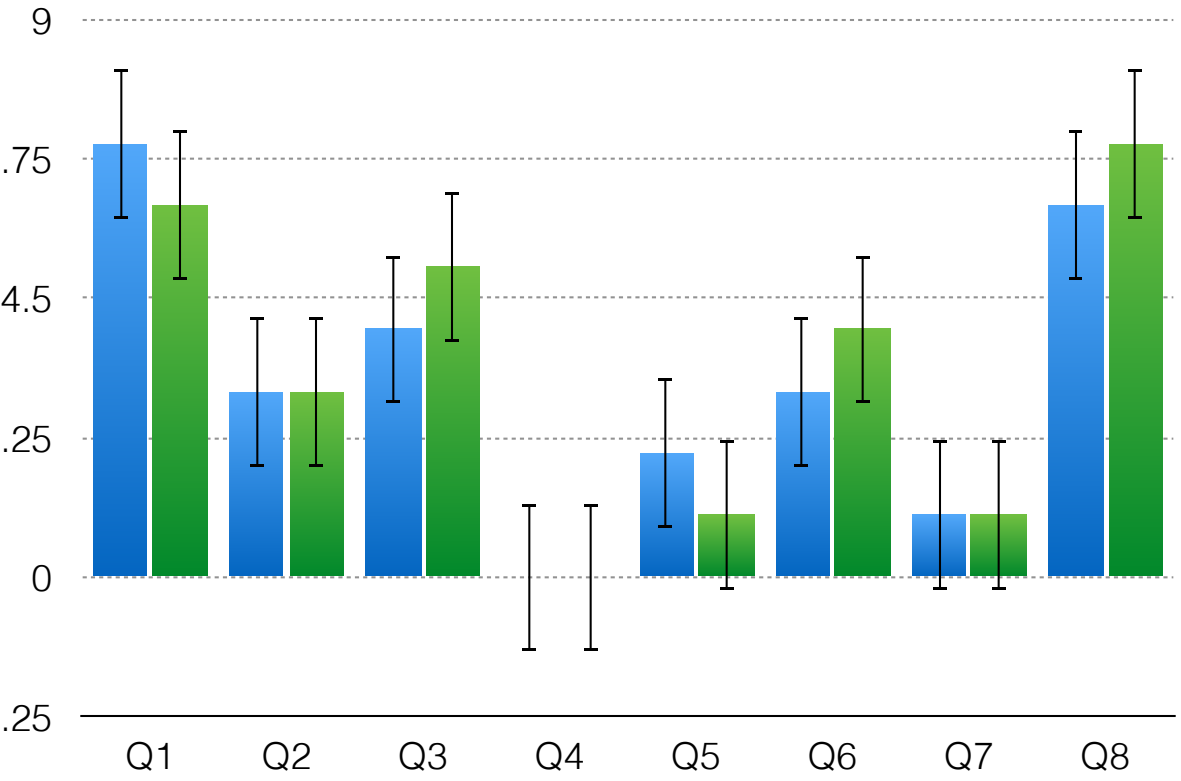

I sensitive queries

$q_1(D)$+noise $q_1(D')$+noise

$q_2(D)$+noise $q_2(D')$+noise

$q_3(D)$+noise $q_3(D')$+noise

....

$q_k(D)$+noise $q_k(D')$+noise

Databases differing in one individual

We can prove this algorithm ε-differentially private

Report Noisy Max - intuition

We need to coordinate noises

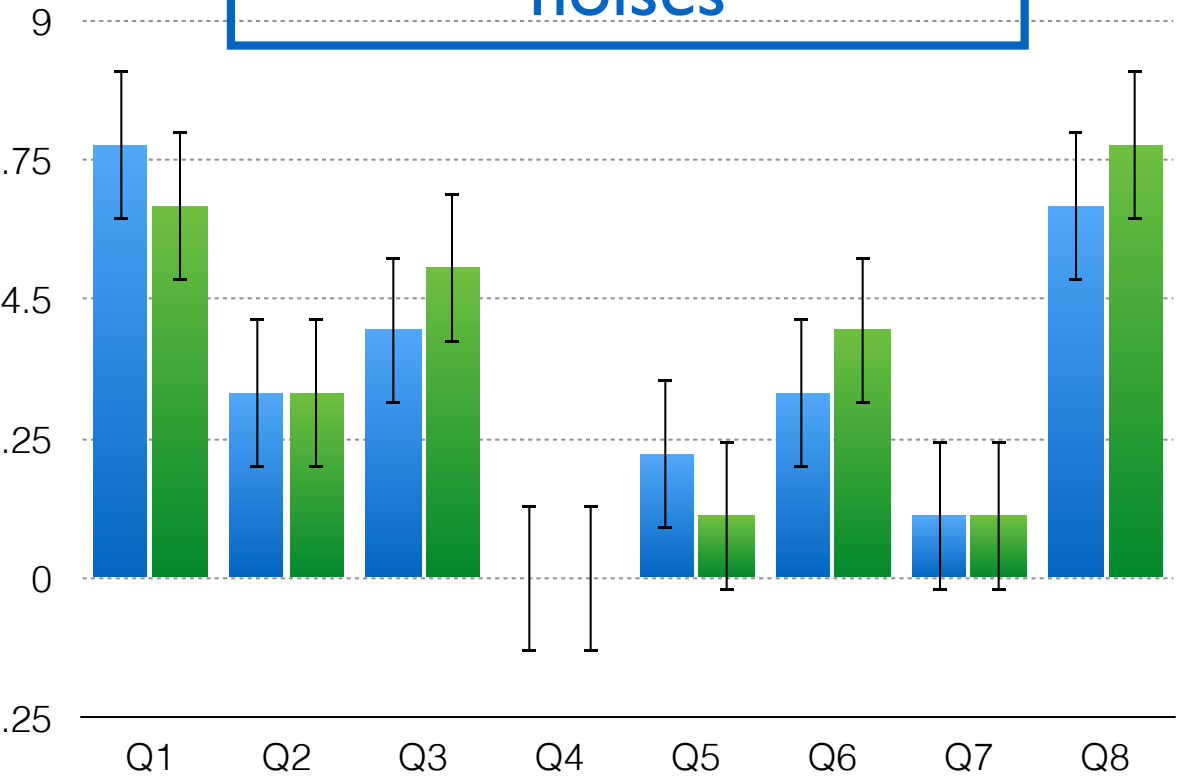

We can prove this algorithm ε-differentially private

I sensitive queries

$q_1(D)$+noise $q_1(D')$+noise

$q_2(D)$+noise $q_2(D')$+noise

$q_3(D)$+noise $q_3(D')$+noise

....

$q_k(D)$+noise $q_k(D')$+noise

Databases differing in one individual

Report Noisy Max

Algorithm 8 Pseudo-code for Report Noisy Max

1: **function** RNM(D, $q_1, \ldots, q_m$, ϵ)
2: **for** $i \leftarrow 1, \ldots, m$ **do**
3: $c_i \leftarrow \mathsf{LapMech}(D, q_i, \epsilon)$
4: **end for**
5: **return** $\mathsf{argmax}_i c_i$
6: **end function**

Report Noisy Max

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Without loss of generality, let us assume that $D = D' \cup \{x\}$.

Let us use: $c_k = q_k(D)$ and $c_k' = q_k(D')$

Report Noisy Max

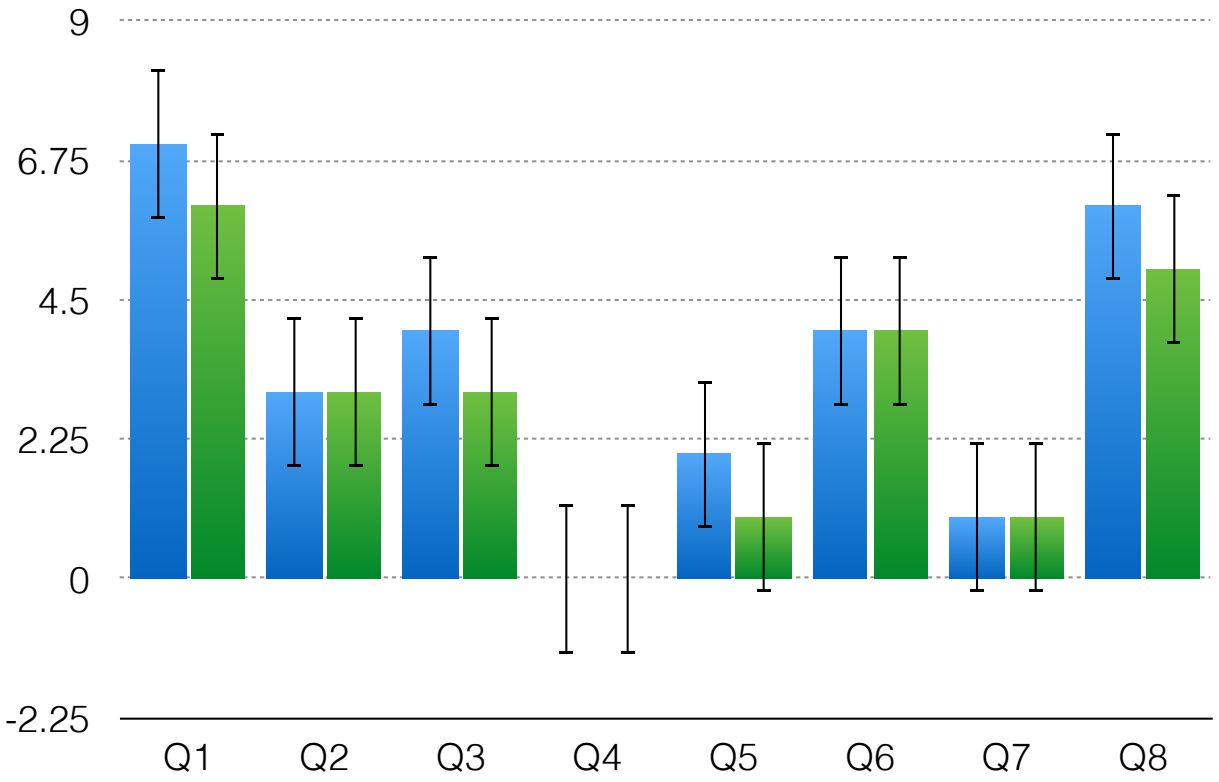

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Without loss of generality, let us assume that $D = D' \cup \{x\}$.

Let us use: $c_k = q_k(D)$ and $c_k' = q_k(D')$

Report Noisy Max

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k'

noise added at round i.

Report Noisy Max

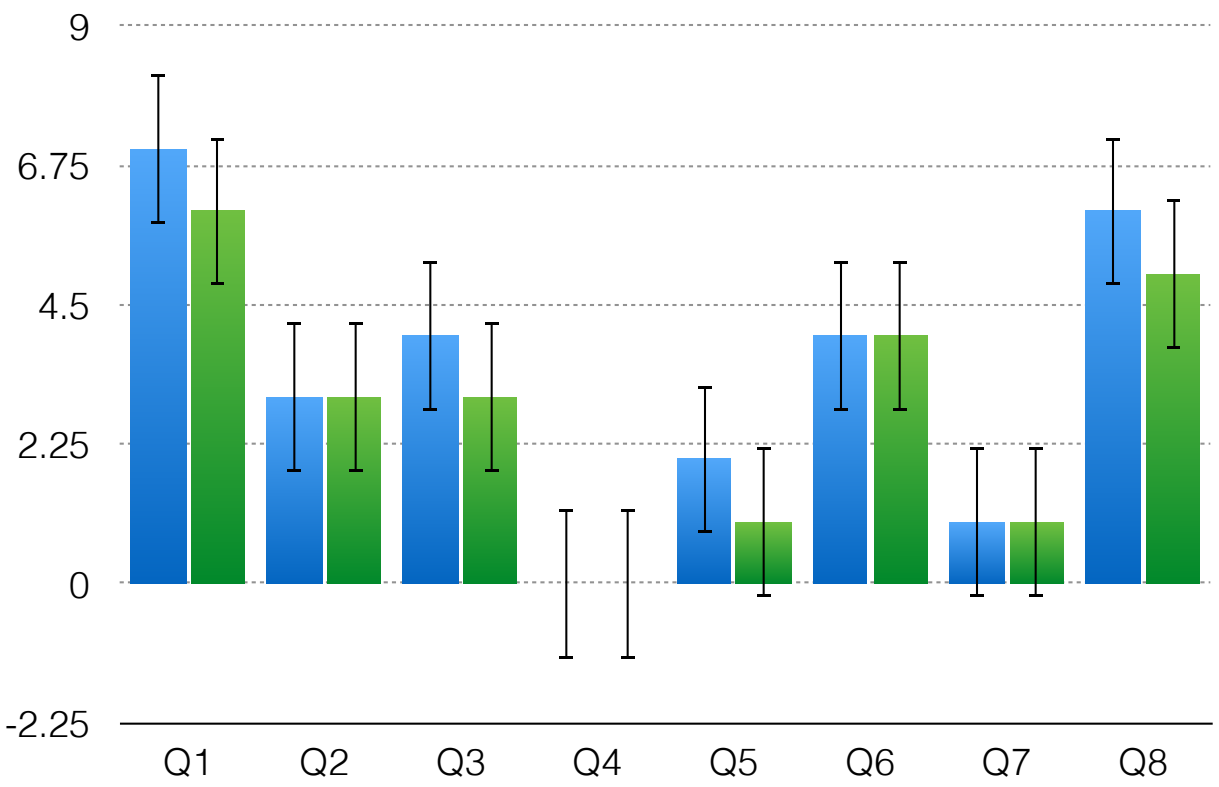

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k'

noise added at round i.

Report Noisy Max

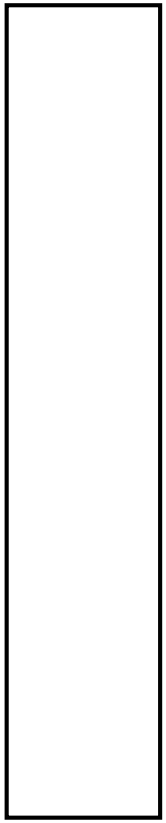

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k'

noise added at round i.

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

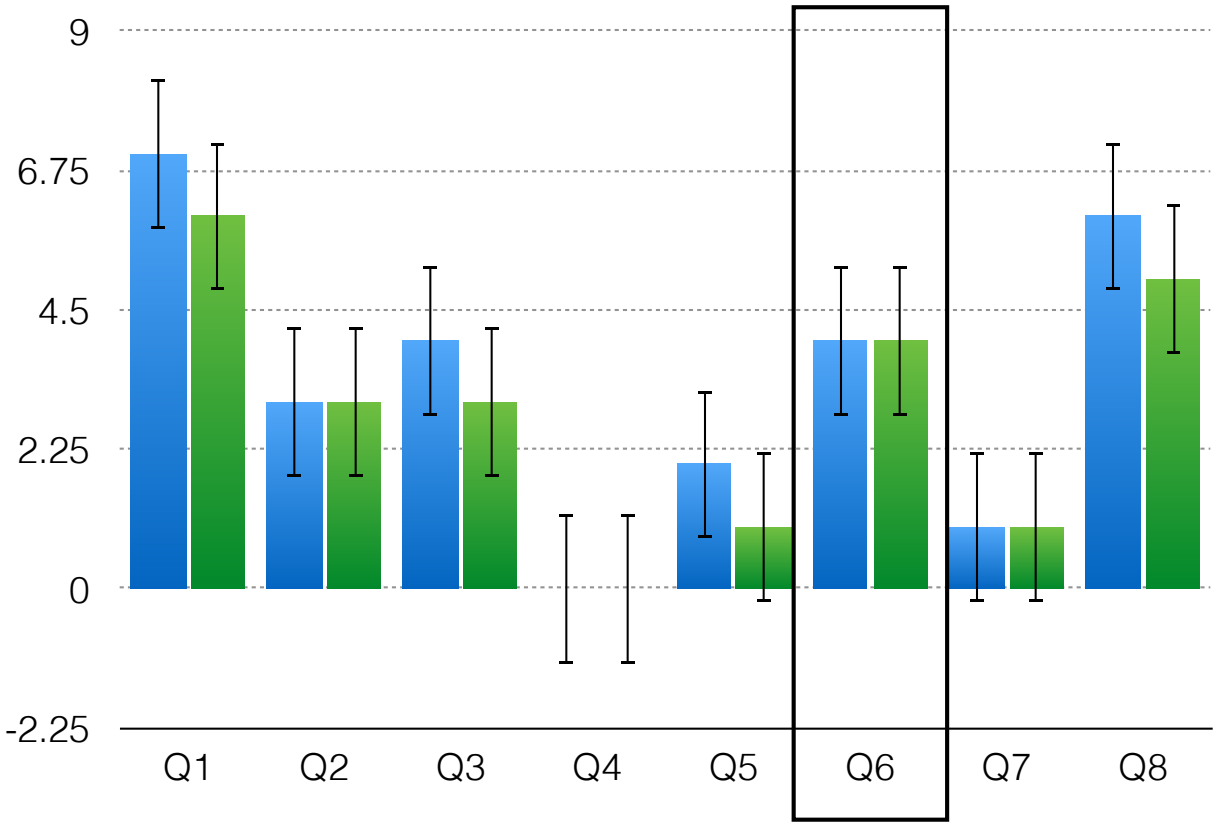

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

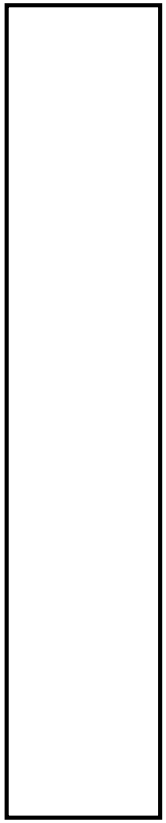

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

We want to show:

$$\Pr_{x \sim RNM(D)}[x = i \,|\, r_{-i}] \leq e^{\epsilon} \Pr_{x \sim RNM(D')}[x = i \,|\, r_{-i}]$$

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

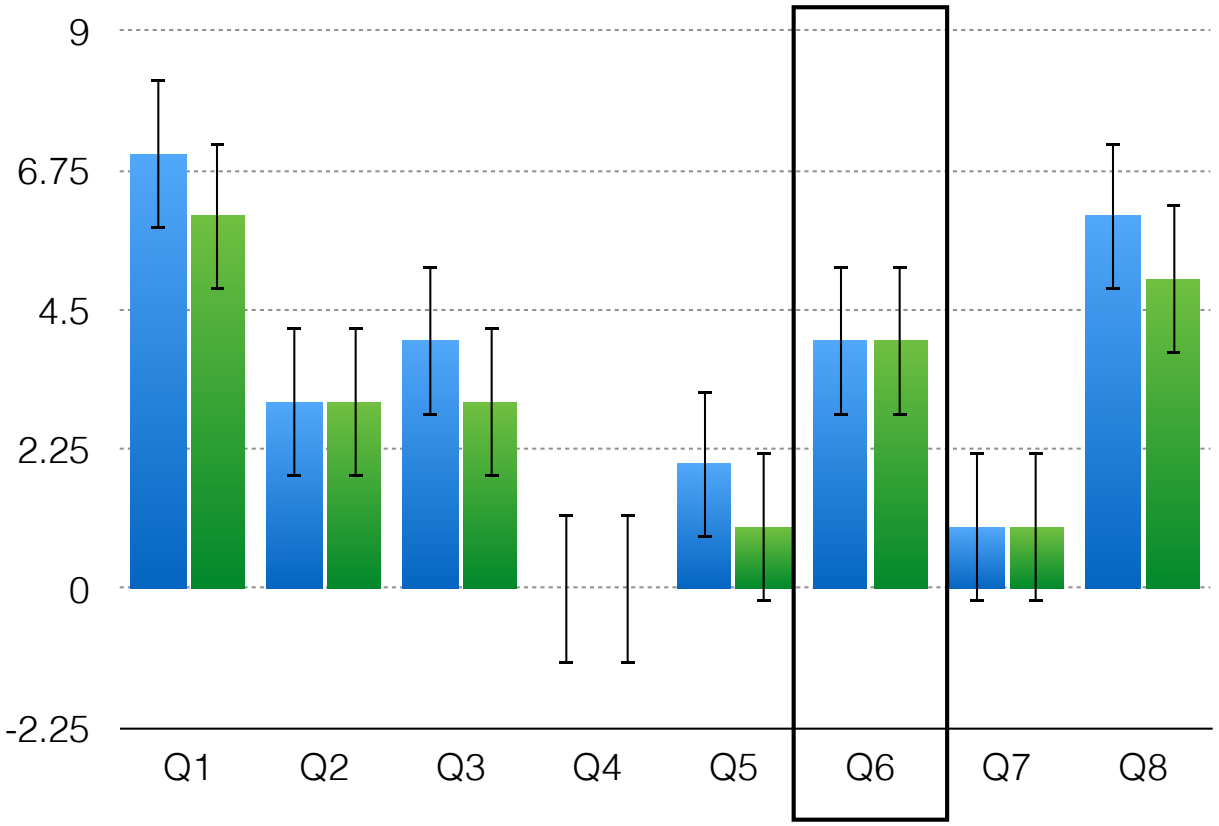

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

We want to show:

$$\Pr_{x \sim RNM(D)}[x = i \,|\, r_{-i}] \leq e^{\epsilon} \Pr_{x \sim RNM(D')}[x = i \,|\, r_{-i}]$$

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

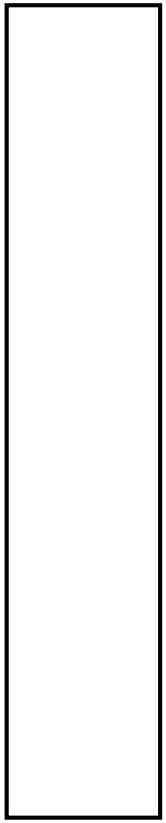

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

By fixing the noises r_j for all $j \neq i$ we can compute the following

$r^* = \min_r \; c_i + r \geq c_j + r_j$ for all j

Report Noisy Max

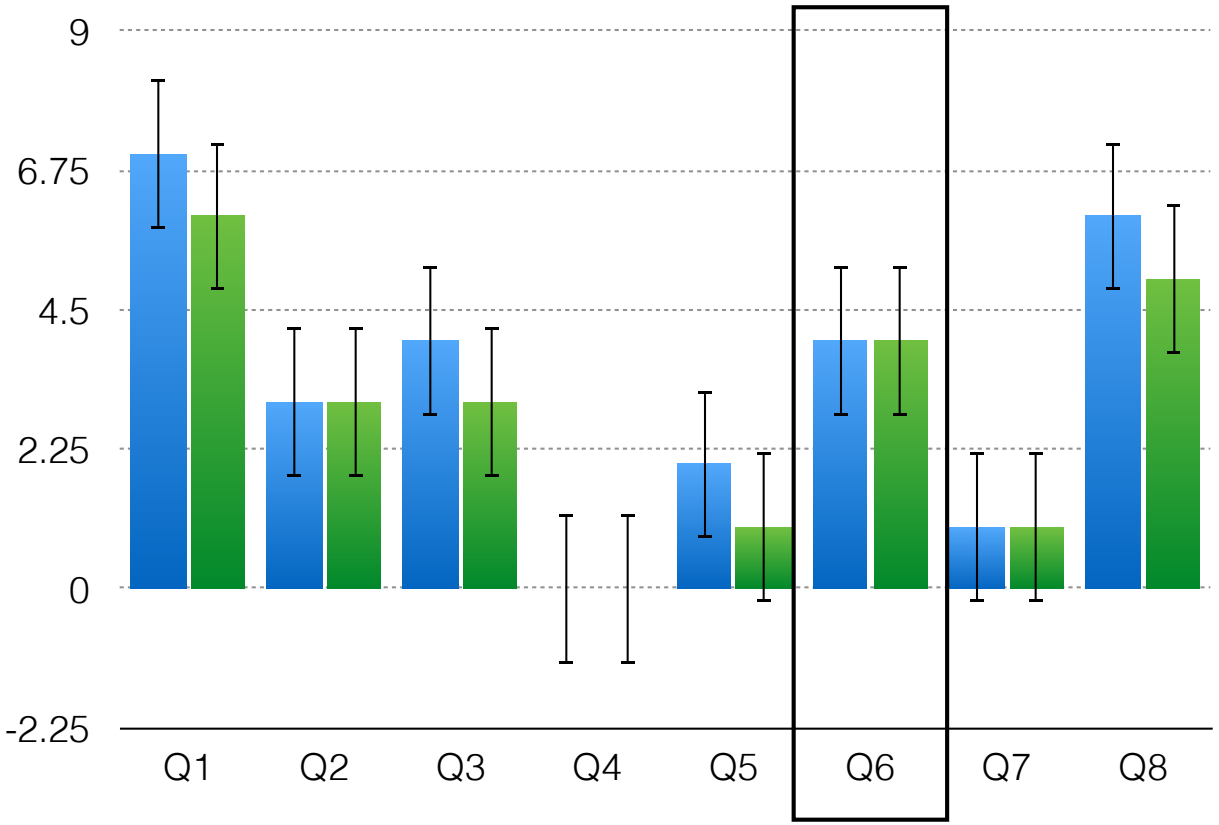

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

By fixing the noises r_j for all $j \neq i$ we can compute the following

$r^* = \min_r \; c_i + r \geq c_j + r_j$ for all j

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

$r^* = \min_r\ c_i + r \geq c_j + r_j$ for all j

Notice that

$$\Pr_{x \sim RNM(D)}[x = i \mid r_{-i}] = \Pr_{r \sim Lap}[r \geq r^*]$$

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

$r^* = \min_r\ c_i + r \geq c_j + r_j$ for all j

Notice that since

$$c_i + r^* \geq c_j + r_j$$

we also have for all j

$$c_i' + 1 + r^* \geq c_j' + r_j$$

and using this we have

$$\Pr_{x \sim RNM(D')}[x = i \mid r_{-i}] \geq \Pr_{r \sim Lap}[r \geq 1 + r^*]$$

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

$r^* = \min_r\ c_i + r \geq c_j + r_j$ for all j

Summarizing we have:

$$\Pr_{x \sim RNM(D)}[x = i \mid r_{-i}] = \Pr_{r \sim Lap}[r \geq r^*]$$

And

$$\Pr_{x \sim RNM(D')}[x = i \mid r_{-i}] \geq \Pr_{r \sim Lap}[r \geq 1 + r^*]$$

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Report Noisy Max

$r^* = \min_r\ c_i + r \geq c_j + r_j$ for all j

Summarizing we have:

$$\Pr_{x \sim RNM(D)}[x = i \mid r_{-i}] = \Pr_{r \sim Lap}[r \geq r^*]$$

And

$$\Pr_{x \sim RNM(D')}[x = i \mid r_{-i}] \geq \Pr_{r \sim Lap}[r \geq 1 + r^*]$$

How can we connect them?

Simplifying assumptions

$c_k \geq c_k'$

$c_k' + 1 \geq c_k$

Notation r_k, r_k' noise added at round i.

Let's focus on the iteration i and let's fix the noises r_j for all $j \neq i$

Laplace Distribution

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

b regulates the skewness of the curve,

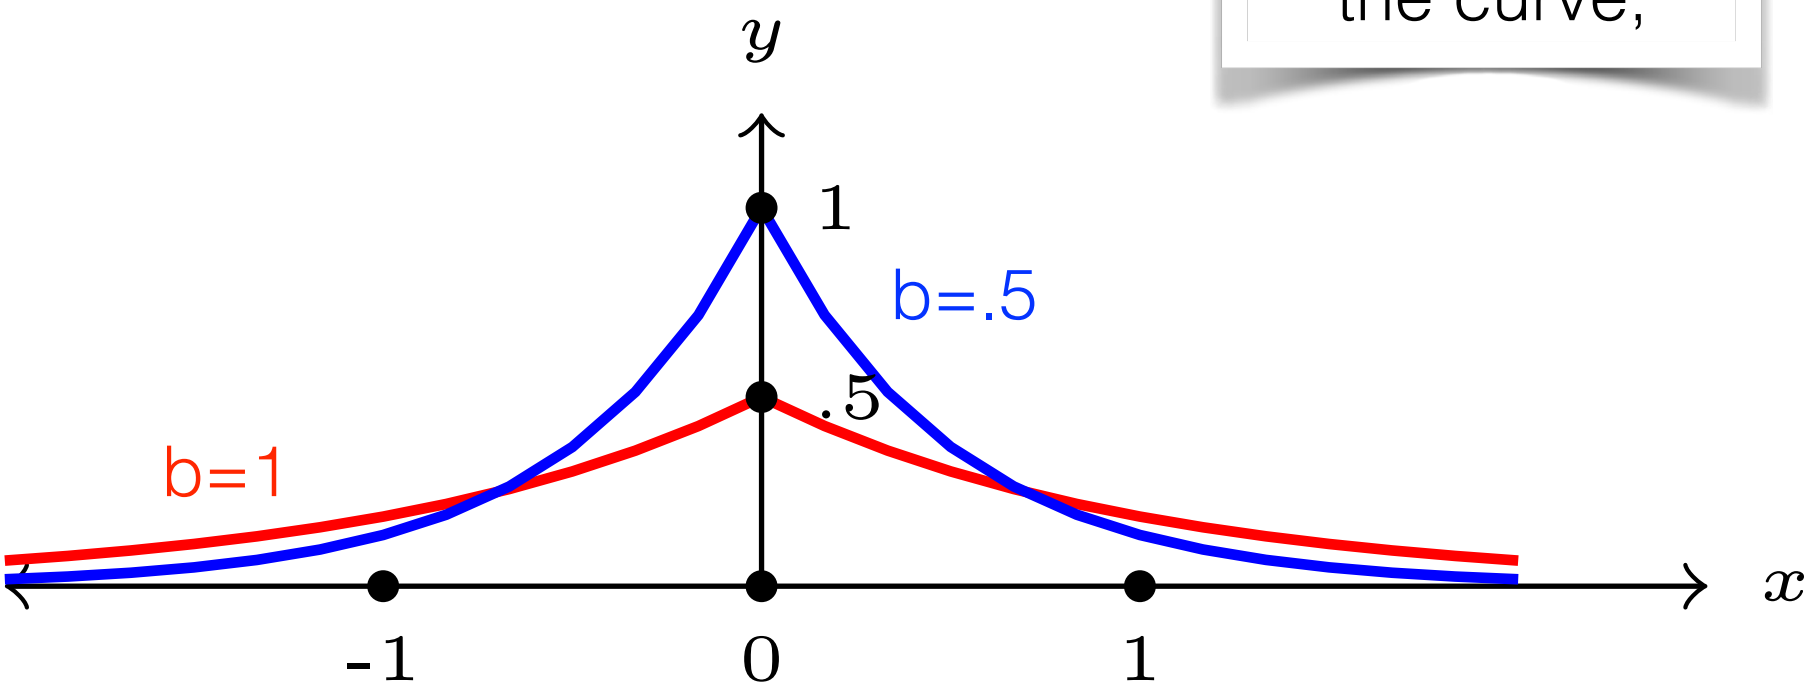

Sliding property of the Laplace Distribution

$$\Pr_{x \sim Lap(\frac{1}{\epsilon}, \mu)}[k \leq x] \leq e^{c\epsilon} \Pr_{x \sim Lap(\frac{1}{\epsilon}, \mu)}[k + c \leq x]$$

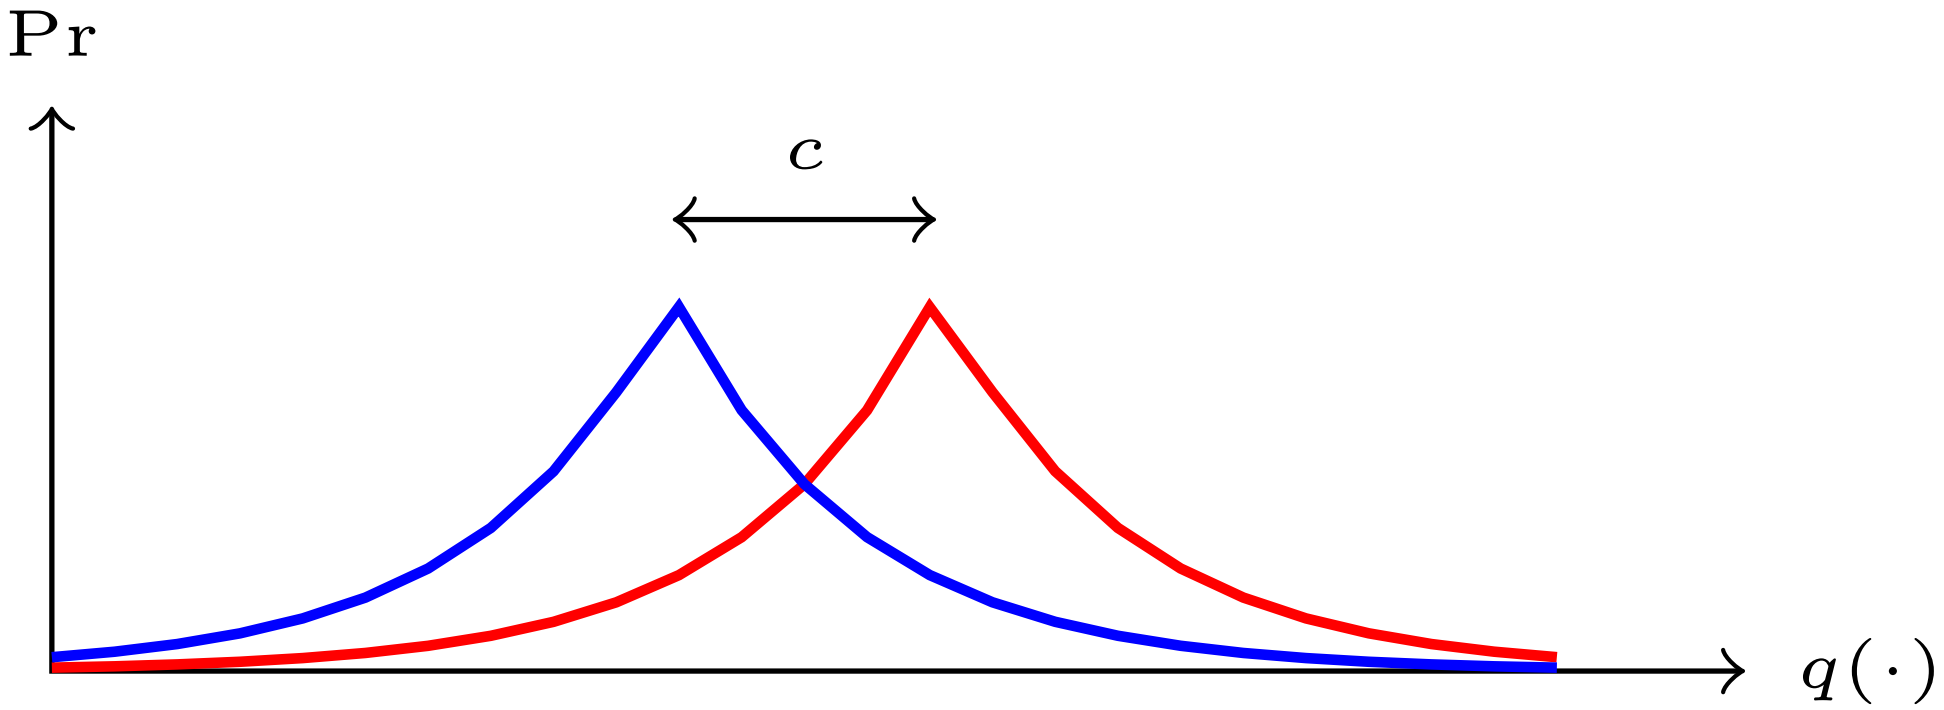

Report Noisy Max

Summarizing we have:

$$\Pr_{x \sim RNM(D)}[x = i \mid r_{\_i}]$$

$$= \Pr_{r \sim Lap}[r \geq r^*] \leq e^{\epsilon} \Pr_{r \sim Lap}[r \geq 1 + r^*]$$

$$\leq e^{\epsilon} \Pr_{x \sim RNM(D')}[x = i \mid r_{\_i}]$$

Report Noisy Max

In a similar way we can prove:

$$\Pr_{x \sim RNM(D')}[x = i \,|\, r\_{i}] \leq e^{\epsilon} \Pr_{x \sim RNM(D)}[x = i \,|\, r\_{i}]$$

Exponential Mechanism

The Exponential Mechanism generalize this approach.

Suppose that we have a scoring function u(D,o) that to each pair (database, potential output) assign a score (a negative real number).

We want to output approximately the element with the max score.

Exponential Mechanism

Exponential Mechanism:

$\mathcal{M}_E(x, u, \mathcal{R})$

return $r \in \mathcal{R}$ with prob. $\dfrac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\sum_{r' \in \mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}$

where

$$\Delta u = \max_{r \in \mathcal{R}} \max_{x \sim_1 y} \left| u(x, r) - u(y, r) \right|$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

The proof is very similar to the one for the Laplace Mechanism.

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

The proof is very similar to the one for the Laplace Mechanism.

$$\frac{\Pr[\mathcal{M}_E(x,u,\mathcal{R})=r]}{\Pr[\mathcal{M}_E(y,u,\mathcal{R})=r]} = \frac{\left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)}{\left(\frac{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}\right)}$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

The proof is very similar to the one for the Laplace Mechanism.

$$\frac{\Pr[\mathcal{M}_E(x,u,\mathcal{R})=r]}{\Pr[\mathcal{M}_E(y,u,\mathcal{R})=r]} = \frac{\left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)}{\left(\frac{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}\right)}$$

$$= \left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}\right) \cdot \left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

The proof is very similar to the one for the Laplace Mechanism.

$$\frac{\Pr[\mathcal{M}_E(x,u,\mathcal{R})=r]}{\Pr[\mathcal{M}_E(y,u,\mathcal{R})=r]} = \frac{\left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)}{\left(\frac{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}\right)}$$

$$= \left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}\right)\cdot\left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

$$= \exp\left(\frac{\varepsilon(u(x,r')-u(y,r'))}{2\Delta u}\right)\cdot\left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

Continuing

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

Continuing

$$= \exp\left(\frac{\varepsilon(u(x,r') - u(y,r'))}{2\Delta u}\right) \cdot \left(\frac{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

Continuing

$$= \exp\left(\frac{\varepsilon(u(x,r') - u(y,r'))}{2\Delta u}\right) \cdot \left(\frac{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

$$\leq \exp\left(\frac{\varepsilon}{2}\right) \cdot \exp\left(\frac{\varepsilon}{2}\right) \cdot \left(\frac{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

Continuing

$$= \exp\left(\frac{\varepsilon(u(x,r') - u(y,r'))}{2\Delta u}\right) \cdot \left(\frac{\sum_{r' \in \mathcal{R}} \exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r' \in \mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

$$\leq \exp\left(\frac{\varepsilon}{2}\right) \cdot \exp\left(\frac{\varepsilon}{2}\right) \cdot \left(\frac{\sum_{r' \in \mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}{\sum_{r' \in \mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

Here we change y with x by paying exp(ε/2).

Exponential Mechanism

Exponential Mechanism Accuracy theorem:

Let $\mathrm{OPT}_u(x) = \max_{r \in \mathcal{R}} u(x, r)$. Then

$$\Pr\left[OPT_u(x) - u(x, \mathcal{M}_E(x, u, \mathcal{R})) \geq \left(\frac{2\Delta u}{\epsilon}\right) \ln\left(\frac{|\mathcal{R}|}{\beta}\right)\right] \leq \beta$$

Exponential Mechanism

Exponential Mechanism Accuracy theorem:
Let $\mathrm{OPT}_u(x) = \max_{r \in \mathcal{R}} u(x, r)$. Then

$$\Pr\left[OPT_u(x) - u(x, \mathcal{M}_E(x, u, \mathcal{R})) \geq \left(\frac{2\Delta u}{\epsilon}\right) \ln\left(\frac{|\mathcal{R}|}{\beta}\right)\right] \leq \beta$$

It follows from this lemma

$$\Pr\left[u(\mathcal{M}_E(x, u, \mathcal{R})) \leq \mathrm{OPT}_u(x) - \frac{2\Delta u}{\varepsilon}\left(\ln\left(\frac{|\mathcal{R}|}{|\mathcal{R}_{\mathrm{OPT}}|}\right) + t\right)\right] \leq e^{-t}$$

Proof.

$$\Pr[u(\mathcal{M}_E(x, u, \mathcal{R})) \leq c] \leq \frac{|\mathcal{R}| \exp(\varepsilon c / 2\Delta u)}{|\mathcal{R}_{\mathrm{OPT}}| \exp(\varepsilon \mathrm{OPT}_u(x) / 2\Delta u)}$$

$$= \frac{|\mathcal{R}|}{|\mathcal{R}_{\mathrm{OPT}}|} \exp\left(\frac{\varepsilon(c - \mathrm{OPT}_u(x))}{2\Delta u}\right).$$

Exponential Mechanism

Exponential Mechanism Accuracy theorem:

Let $\mathrm{OPT}_u(x) = \max_{r \in \mathcal{R}} u(x, r)$. Then

$$\Pr\left[OPT_u(x) - u(x, \mathcal{M}_E(x, u, \mathcal{R})) \geq \left(\frac{2\Delta u}{\epsilon}\right) \ln\left(\frac{|\mathcal{R}|}{\beta}\right)\right] \leq \beta$$

Exponential Mechanism

Exponential Mechanism Accuracy theorem:
Let $\mathrm{OPT}_u(x) = \max_{r \in \mathcal{R}} u(x, r)$. Then

$$\Pr\left[OPT_u(x) - u(x, \mathcal{M}_E(x, u, \mathcal{R})) \geq \left(\frac{2\Delta u}{\epsilon}\right) \ln\left(\frac{|\mathcal{R}|}{\beta}\right)\right] \leq \beta$$

Let's compare it with the accuracy of the Laplace Mechanism.

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

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

Exponential Mech

Here we have a dependency on the size of the output space

Exponential Mechanism Accuracy the

Let $\mathrm{OPT}_u(x) = \max_{r \in \mathcal{R}} u(x, r)$. Then

$$\Pr\left[OPT_u(x) - u(x, \mathcal{M}_E(x, u, \mathcal{R})) \geq \left(\frac{2\Delta u}{\epsilon}\right) \ln\left(\frac{|\mathcal{R}|}{\beta}\right)\right] \leq \beta$$

Let's compare it with the accuracy of the Laplace Mechanism.

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

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

Exponential Mechanism

The Exponential Mechanism is a very general mechanism. It can actually be used as a kind of universal mechanism.

Unfortunately, when the output space is big it can be very costly to sample from it - the best option is to enumerate all the possibilities.

Moreover, when the output space is big also the accuracy get worse.