From last week,...

Name	Weight
Alice	40
Bob	60
Charles	80
Doe	60

private information

⟹ Average Weight = 60

public information

- If Alice, Bob, and Doe reveal their information, Charles' weight will be also revealed.

Name	Weight
Alice	40
Bob	60
Charles	40
Doe	60

Name	Weight
Alice	40
Bob	60
Charles	80
Doe	60

cannot distinguish

Average weight + noise = 60 publicity

Definition. Two tables are neighbors if they are different just by one record.

Goal: Publicity cannot distinguish any neighboring tables.

Idea: Add noise to the published information.

Laplacian Distribution

Lap(b) has distribution $\underbrace{p(x; b)}_{\text{prob. that we have } x \text{ from } Lap(b)} = \frac{1}{2b} \cdot \exp\left(\frac{-|x|}{b}\right)$

Example Lap 1

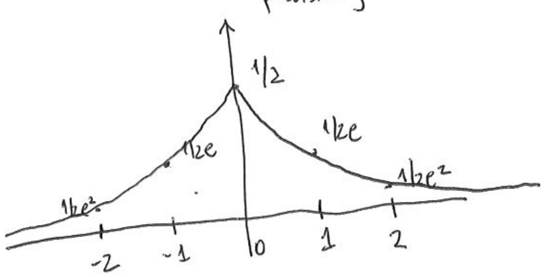

Expected Value = 0

Variance = $2b^2$

larger b = wider probability distribution.

Average Weight $= 60$ $\longrightarrow$ Average Weight $= 60 +$ some noise drawn from Laplacian Distribution

---

$f(T)$:= statistical conclusion from T $\Rightarrow$ (in example) $f(T)$:= average weight obtained from T.

$$GS(f) := \max_{T, T' : \text{neighboring tables}} |f(T) - f(T')|$$

maximum differene in statistical conclusion obtained from two neighboring table.

Published Information $= f(T) + Lap\left[\frac{GS(f)}{\varepsilon}\right]$ ε is a parameter between 0 and 1.

Larger $GS(f)$ $\rightarrow$ fatter probability distribution

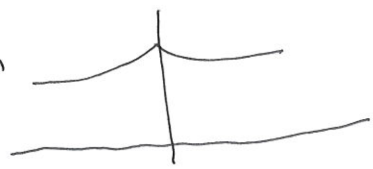

larger noise

Larger ε $\rightarrow$ thinner probability distribution

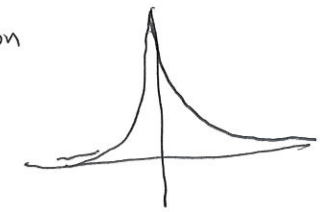

smaller noise

Example Let assume that weights are always between 30 and ~~120~~ 150

When we have 4 persons

$$f(T) = \frac{w_1 + w_2 + w_3 + w_4}{4}$$ $\rightarrow$ change from 30 to 150

- sum changed by 120
- average changed by 30

Assume that $\varepsilon = 0.1$

Noise drawn from $Lap\left(\frac{30}{0.1}\right) = Lap(300)$

large noise

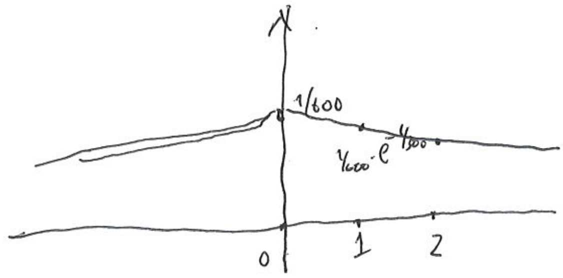

$$\frac{1}{600} e^{-1/300} \approx \frac{1}{600}\left(1 - \frac{1}{300}\right)$$

When we have 1,000 persons $f(T) = \frac{w_1 + \dots + w_{1000}}{1000}$

- Sum changed by 120
- average changed by 0.12

Assume that $\varepsilon = 0.1$

Noise drawn from $Lap\left(\frac{0.12}{0.1}\right) = Lap(1.2)$

Small noise

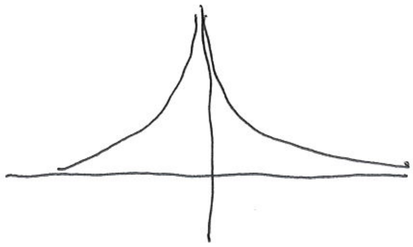

* Noise is usually smaller when we have more people in the table

---

Chebyshev's Inequality "The probability that a random sampling is far from the expected value by more than $k \cdot [S.D]$ is no larger

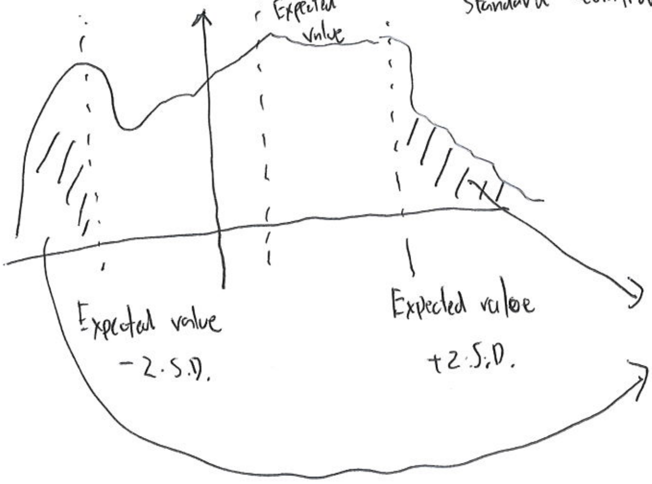

standard deviation

$$\Pr\left[|X - \mu| \geq k \cdot S.D\right] \leq 1/k^2$$

Probability $\leq 1/4$.

Lap (1.2) has variance $= 2 \cdot 1.2^2$

standard deviation $= \sqrt{2} \cdot 1.2 \approx 1.7$

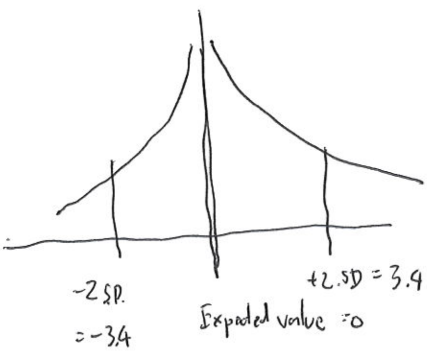

Prob. that we add noise more than 3.4 is no larger than 3.4

Prob. that we publish the value between 56.6 and 63.4 is no less than 3/4.

<u>Open question</u>. Discuss why Laplacian Distribution will not work when

$$f(T) = \text{minimum weight} \quad \text{or} \quad f(T) = \text{maximum weight}$$

<u>Theorem</u>. Let T and T' be two neighboring tables. $Out(T)$ and $Out(T')$ are outputs obtained from Laplacian mechanism. For all y,

$$1 - \varepsilon \leq \frac{\Pr[Out(T) = y]}{\Pr[Out(T') = y]} \leq 1 + \varepsilon$$

[Note: when we see y as a published information it is hard to guess if it comes from T or T']

<u>Proof</u> The noise added ~~to~~ when $Out(T) = y$ is $\cancel{\{y - f(T)\}}$. $y - f(T)$.

We have that noise with probability

$$p(x; b) = \frac{1}{2b} \cdot \exp\left(-\frac{x}{b}\right) \qquad x \to y - f(T) \qquad \text{when } b = \frac{GS(f)}{\varepsilon}$$

$$\Pr[Out(T) = y] = \frac{\varepsilon}{2GS(f)} \cdot \exp\left(-\varepsilon \frac{|y - f(T)|}{GS(f)}\right)$$

$$\Pr[Out(T') = y] = \frac{\varepsilon}{2GS(f)} \exp\left(-\varepsilon \frac{|y - f(T')|}{GS(f)}\right)$$

$$\frac{\Pr[Out(T) = y]}{\Pr[Out(T') = y]} = \frac{\frac{\varepsilon}{2\cancel{GS(f)}} \exp\left(-\varepsilon \frac{|y - f(T)|}{GS(f)}\right)}{\frac{\varepsilon}{2\cancel{GS(f)}} \exp\left(-\varepsilon \frac{|y - f(T')|}{GS(f)}\right)} \qquad |p| - |q| \leq |p - q|$$

$$= \exp\left[-\frac{\varepsilon}{GS(f)} |y - f(T)| + \frac{\varepsilon}{GS(f)} |y - f(T')|\right]$$

$$= \exp\left[\frac{\varepsilon}{GS(f)} \left[|y - f(T')| - |y - f(T)|\right]\right]$$

$$\leq \exp\left[\frac{\varepsilon}{GS(f)} |y - f(T') - (y - f(T))|\right]$$

$$= \exp\left[\frac{\varepsilon}{GS(f)} |f(T) - f(T')|\right]$$

$$\leq \exp\left[\frac{\varepsilon}{\cancel{GS(f)}} \cdot \cancel{GS(f)}\right]$$

$$= e^{\varepsilon} \approx 1 + \varepsilon$$

$\cancel{GS(f)}$

$$|f(T) - f(T')| \leq \max_{T, T' : \text{neighboring tables}} |f(T) - f(T')| = GS(f)$$

By Similar proof, we have

$$\frac{\Pr[\,Out\,(T') = y\,]}{\Pr[\,Out\,(T) = y\,]} \leq e^{\varepsilon}$$

$$\frac{\Pr[\,Out\,(T) = y\,]}{\Pr[\,Out\,(T') = y\,]} \geq e^{-\varepsilon} \simeq 1-\varepsilon$$

Definition A scheme is ε-differential private if, for all possible value y and neighboring tables T, T',

$$e^{-\varepsilon} \leq \frac{\Pr[\,Out\,(T) = y\,]}{\Pr[\,Out\,(T') = y\,]} \leq e^{\varepsilon}$$

Collorary Laplacian mechanism is ε-differential private.

Larger ε → Smaller noise / less privacy.

trade-off