We say that algo A is ε-DP if for any pair of neighboring datasets D, D' and for any set of possible outputs S we have:

$$\frac{1}{e^{\varepsilon}} \Pr[A(D') \in S] \leq \Pr[A(D) \in S] \leq e^{\varepsilon} \Pr[A(D') \in S]$$

- D includes your data, D' does not
- S is set of outputs you are "worried" about
- $\Pr[\,]$ over only # of A
- For small x, $e^x \approx (1+x)$

Interpretations

- Suppose Eve sees the output of A and has to guess whether input was D or D'! Eve may have all kinds of other data, computations, etc. Let S = set of outputs for which Eve says "D". Then

$$\Pr[\text{Eve says "}D\text{"} \mid D] \approx_\varepsilon \Pr[\text{Eve says "}D\text{"} \mid D'].$$

- If $\Pr[A(D') \in S]$ is small, e.g. 1/million, then OK if ε is large, e.g. $\varepsilon = 10$ gives $e^{10}/\text{million} = 0.02$. But if $\Pr[A(D') \in S] = 1/4$, need $\varepsilon \approx 0$, since even $e^1/4 \not\approx 3/4$.

Randomized Response Revisited

roll "dice"

prob. p → answer truthfully

prob. $1-p$ → flip coin

H → "yes"

T → "no"

Utility Analysis:

$$\Pr[\text{"yes"} \mid \text{yes}] = p + (1-p)/2$$

$$\Pr[\text{"yes"} \mid \text{no}] = (1-p)/2$$

∴ If t = true frac of yes then

$$\text{exp. frac of "yes"} = t[p + (1-p)/2] + (1-t)(1-p)/2$$

$$\text{estimate} = tp + (1-p)/2$$

& solve for t

Claim:

$$|t - \hat{t}| \lesssim \frac{1}{\sqrt{pn}}$$ where

n = population size

Privacy Analysis:

At local or individual level.

$$\frac{Pr[\text{"yes"}|\text{yes}]}{Pr[\text{"yes"}|\text{no}]} = \frac{p+(1-p)/2}{(1-p)/2}$$

$$= \frac{p}{(1-p)/2} + 1$$

$$2\frac{p}{1-p} + 1$$

$$Set = e^{\varepsilon}, \quad \varepsilon = \ln\left(1 + 2\frac{p}{1-p}\right)$$

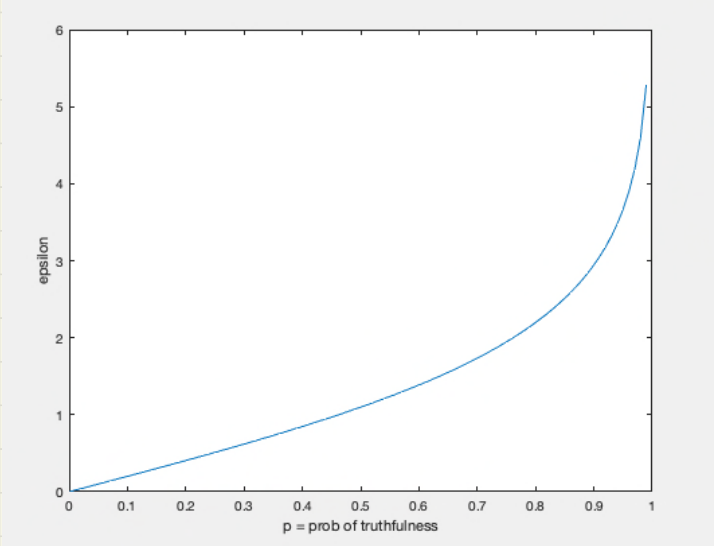

Privacy vs. Accuracy

For $n = 10, 100, 1000, 10K$

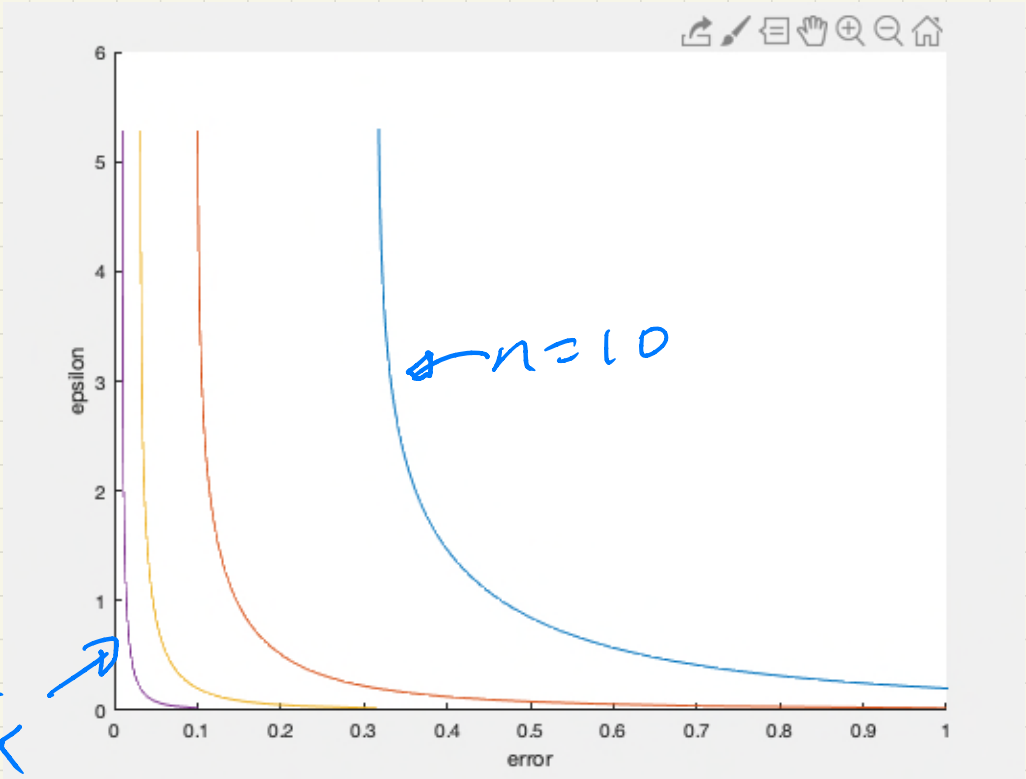

Theoretical Pareto curve of error vs. privacy!

A More General Tool: The Laplace Mechanism

- Consider DBs of the form $\bar{x} \in [0,1]^n$, so each person i is a number x_i
- Want to compute some function $f(\bar{x}) = f(x_1, x_2, \ldots x_n) \in \mathbb{R}$

  e.g. average, median, stdev, max, min, etc.
- Neighboring: $\bar{x}$ & $\bar{x}'$ differ in only <u>one</u> entry x_i

Define sensitivity of f as:

$$\Delta f = \max_{\text{neighboring } \bar{x}, \bar{x}'} \{ |f(\bar{x}) - f(\bar{x}')| \}$$

Examples

- $f = \text{avg}$: $\Delta f = 1/n$ (remember $x_i \in [0,1]$)
- $f = \text{stdev}$: $\sqrt{\sum_i (x_i - \mu)^2 / n} \rightarrow \Delta f \approx 1/\sqrt{n}$
- $f = \text{median}$:
- $f = \max$ or $\min$:

Intuition: Δf captures extent of "aggregation" of f

smaller $\Delta f \leftrightarrow$ more aggregation $\leftrightarrow$ easier to provide DP

The Laplace Distribution

- randomly draw a value v with probability

$$\propto \frac{1}{2b} e^{-|v|/b}$$

($\propto$ proportional to)

- mean = max = 0
- larger $|v|$ less likely
- b is a parameter
- $b \to 0$: rapid fall-off in $|v|$
- $b \to \infty$: slower fall-off

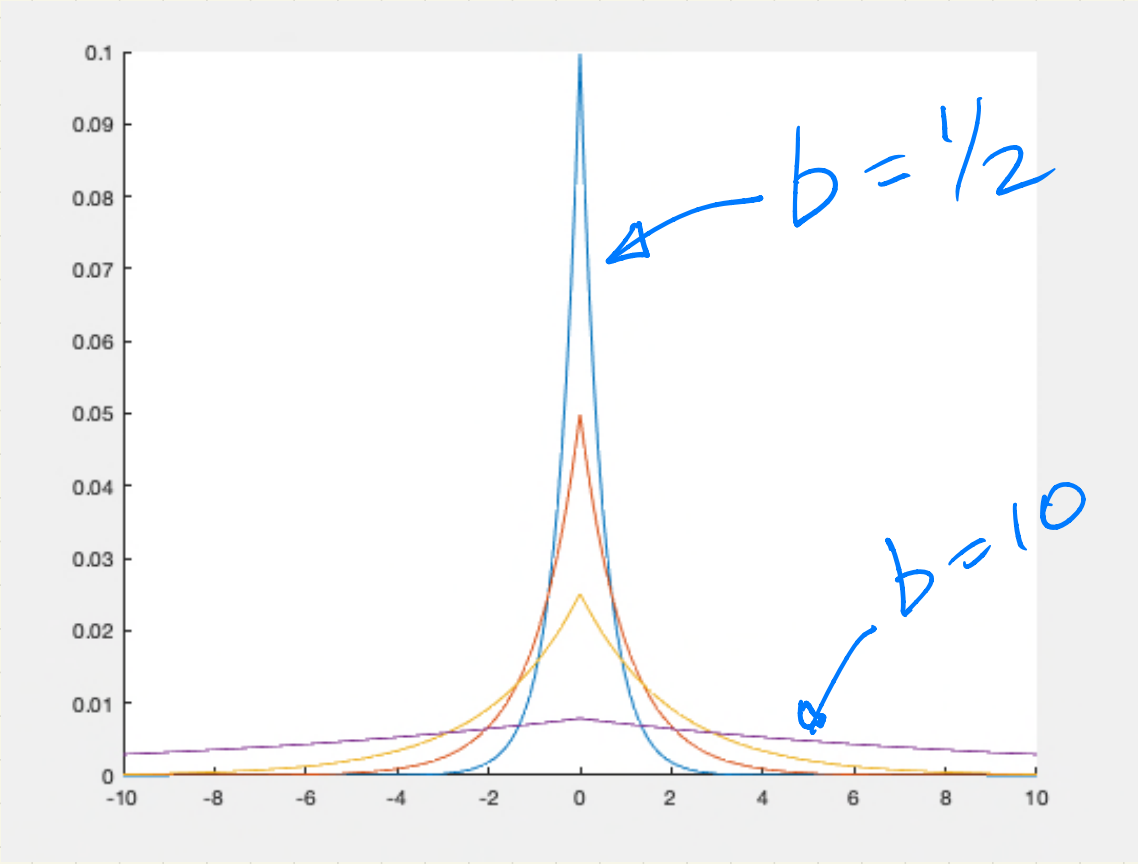
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
-10
-8
-6
-4
-2
0
2
4
6
8
10
b = 1/2
b = 10

Laplace Mechanism:

- compute $f(\bar{x})$ exactly
- output: $f(x) + v$

  where v is chosen randomly from Laplace distribution with $b = \Delta f / \varepsilon$
- smaller Δf/larger ε: less noise
- larger Δf/smaller ε: more noise

Privacy: Laplace Mech. satisfies ε-DP. (proof)

Utility:

$|\text{output} - f(\vec{x})| \approx$ stdev of $\nu \approx b = \frac{\Delta f}{\varepsilon}$

So if e.g. $\Delta f = 1/n$, we can set $\varepsilon = 1/\sqrt{n}$, then $b = \frac{1/n}{1/\sqrt{n}} = \frac{1}{\sqrt{n}}$.

As $n \to \infty$,

$\varepsilon \to 0$ perfect privacy

and $b \to 0$ perfect accuracy