Differential Privacy
the Mathematical Bulwark against Reidentification and Reconstruction

Cynthia Dwork
Harvard University
Radcliffe Institute for Advanced Study
Microsoft Research

This Talk in a Nutshell

Population as a Whole vs Needle in a Haystack

United States™ Census Bureau

BROWSE BY TOPIC EXPLORE DATA

We're the CFPB

The Consumer Financial Protection Bureau is a U.S. government agency that makes sure banks, lenders, and other financial companies treat you fairly.

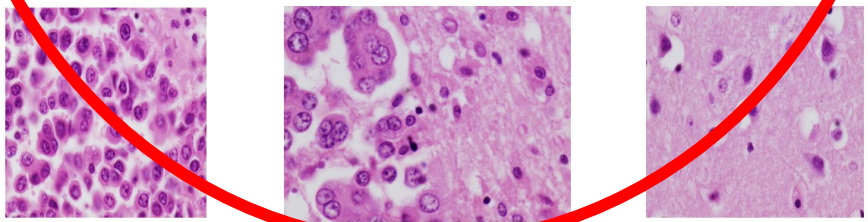

In November 2002, the New York Times reported tha
(DARPA) was developing a tracking system called "To
intended to detect terrorists through analyzing trove

Cell image credit: Andrew Dwork

Haystack image credit: Hackernoon

Statistics "Feel" Private

- A quantity computed from a sample, tells us about the population as a whole

On the right track but needs help.
Differential Privacy provides this help.

- Sense of privacy derived from this fact
  - "No one knows I am in the sample … I can claim I opted out"
  - "It's not about *me*"

"Statistical" Privacy for All Computations

Differential privacy preserves "I could have opted out" privacy for every computation, including total population counts

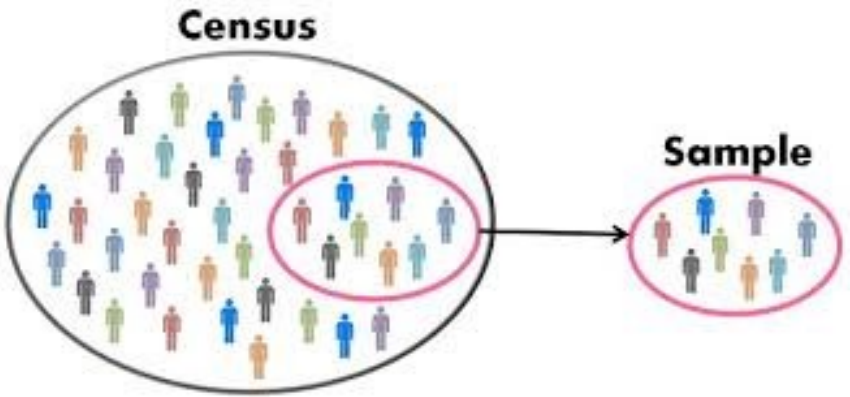

Image credit: Keydifferences.com

Abstracting The Problem

Privacy-Preserving Data Analysis

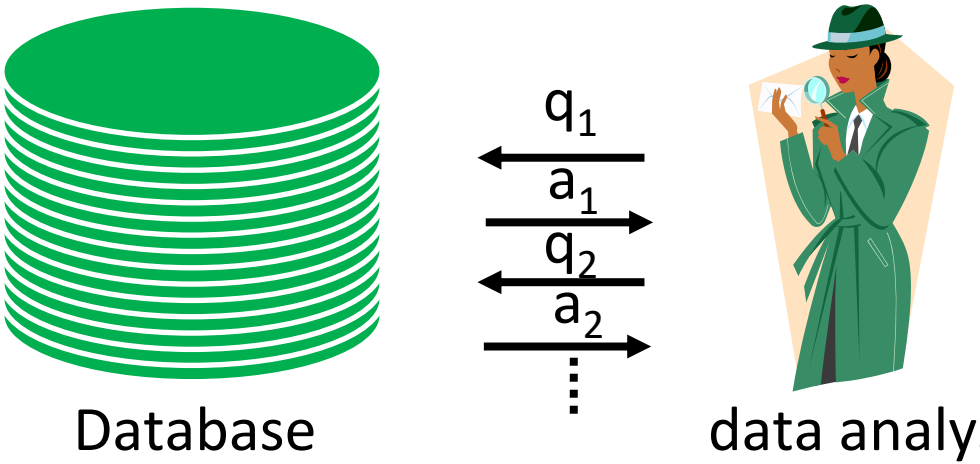

- 55 year old problem
- Driving scenario: analysis of US Census data

Image credit: The New Yorker Magazine

"Just" Statistics

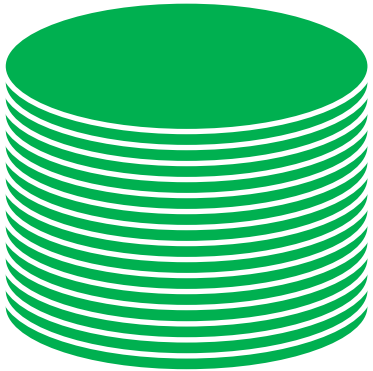

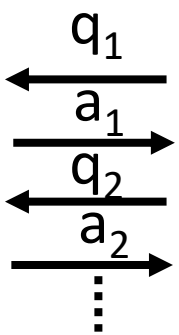

Database

data analyst

- How many members of the House floss regularly?
- How many members, other than the speaker, floss regularly?

Fundamental Law

- "Overly accurate" estimates of "too many" statistics is blatantly non-private
- Applies equally to non-interactive systems

Dinur and Nissim '03

The Definition of Differential Privacy
Motivation and Meaning

Privacy-Preserving Data Analysis?

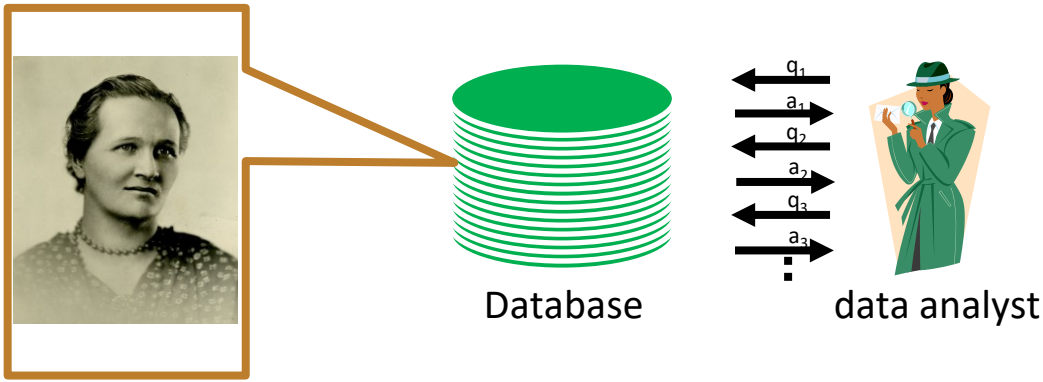

Privacy-Preserving Data Analysis?

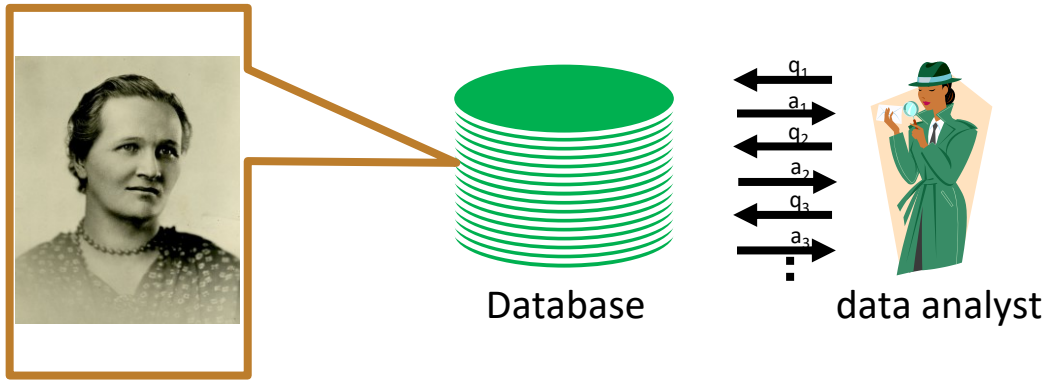

- “Can’t learn anything new about Payne”?
- Dalenius, 1977

Privacy-Preserving Data Analysis?

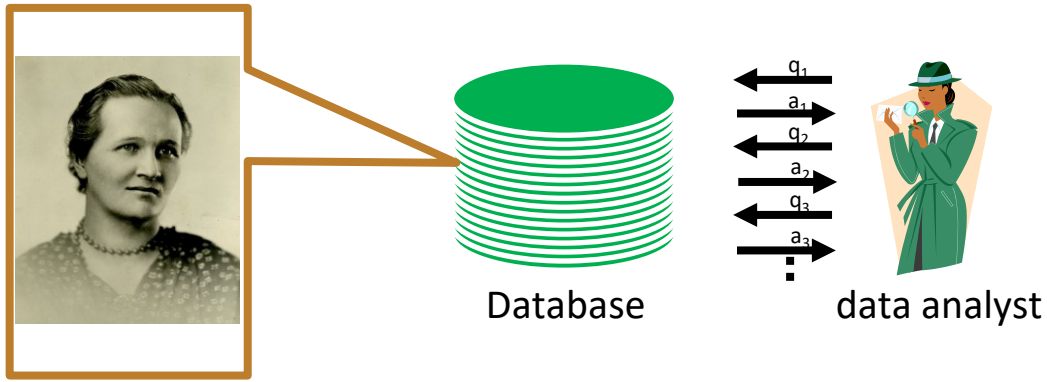

- "Can't learn anything new about Payne"?
- Then what is the point?

Privacy-Preserving Data Analysis?

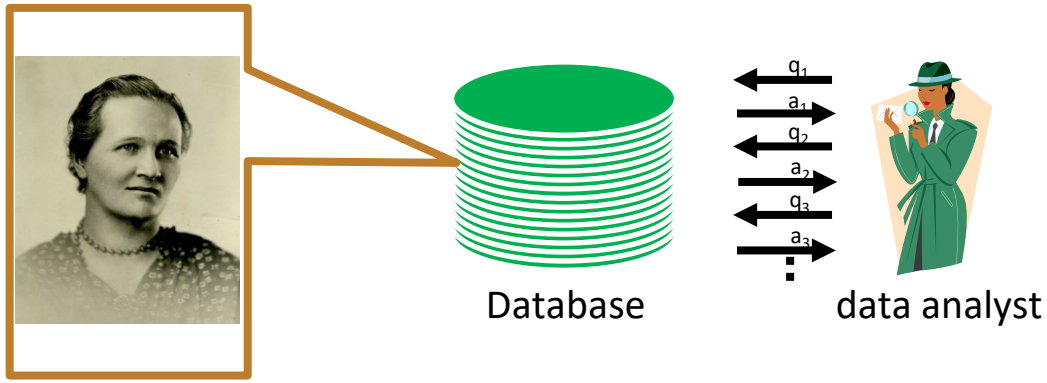

- "Can't learn anything new about Payne"?
- Then what is the point?

Privacy-Preserving Data Analysis?

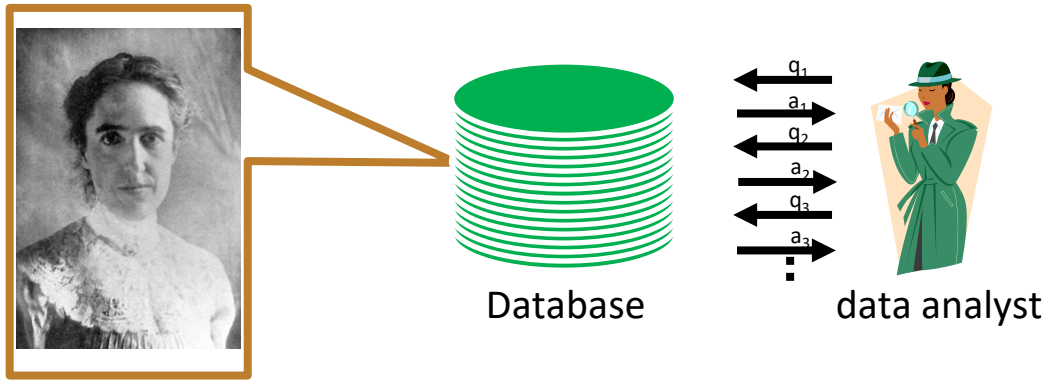

- Ideally: learn same things if Payne is replaced by another random member of the population

Differential Privacy

The outcome of any analysis is essentially equally likely, independent of whether any individual joins, or refrains from joining, the dataset.

Differential Privacy

The outcome of any analysis is essentially equally likely, independent of whether any individual joins, or refrains from joining, the dataset.

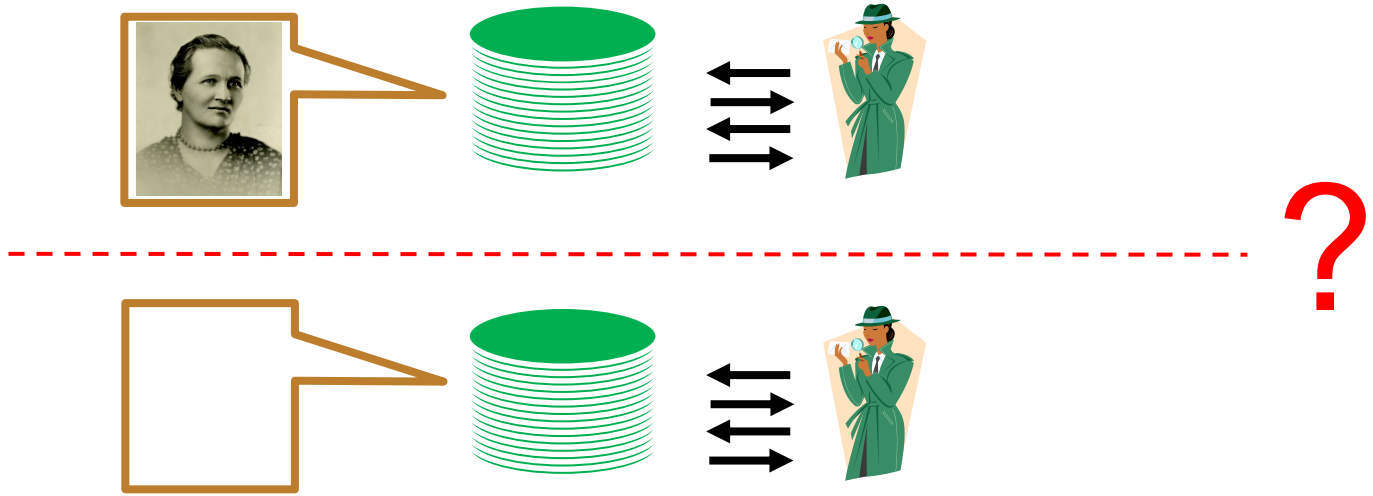

Differential Privacy

The outcome of any analysis is essentially equally likely, independent of whether any individual joins, or refrains from joining, the dataset

- An adv… …eir entirety can never disti…
- What is the source of uncertainty in “likely”?

Algorithm will flip coins

Privacy-Preserving Data Analysis?

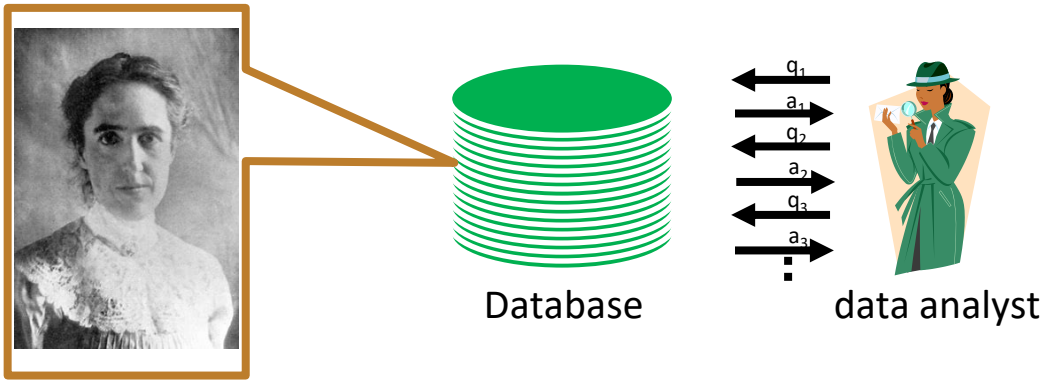

Stability preserves Payne's privacy AND prevents over-fitting

Privacy and Generalization are aligned!

Differential Privacy

M gives ϵ-differential privacy if for all pairs of adjacent data sets x, y, and all output events S

$$\Pr[\text{see } S \text{ on } M(x)] \leq e^{\epsilon} \Pr[\text{see } S \text{ on } M(y)]$$

Bound on "Privacy Loss"

Randomness introduced by M

Dwork, McSherry, Nissim, Smith 2006

Differential Privacy

M gives ϵ-differential privacy if for all pairs of adjacent data sets x, y, and all output events S

$$\Pr[\text{see } S \text{ on } M(x)] \leq e^{\epsilon} \Pr[\text{see } S \text{ on } M(y)]$$

Statement is about behavior of M.
Doesn't care who knows what. Now or in the future.

Dwork, McSherry, Nissim, Smith 2006

Differential Privacy

M gives ϵ-differential privacy if for all pairs of adjacent data sets x, y, and all output events S

$$\Pr[\text{see } S \text{ on } M(x)] \leq e^{\epsilon} \Pr[\text{see } S \text{ on } M(y)]$$

You can learn about Payne

You can only learn things you can learn without Payne

Dwork, McSherry, Nissim, Smith 2006

Differential Privacy

M gives ϵ-differential privacy if for all pairs of adjacent data sets x, y, and all output events S

$$\Pr[\text{see } S \text{ on } M(x)] \,/\, \Pr[\text{see } S \text{ on } M(y)] \leq e^{\epsilon}$$

"Bounded Ratio"

Dwork, McSherry, Nissim, Smith 2006

Differential Privacy Hides the Needle

Haystack vs Haystack sans needle

x y

Key Properties

- Future-Proof

  Resilient to present/future information from other sources

- Composes Gracefully and Automatically

  Understand cumulative privacy loss over multiple comps

  At worst, the losses add up.

- Differential privacy is programmable

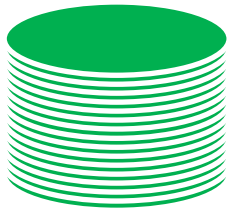

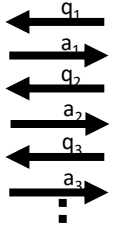

Database

data analyst

One Technique:
Laplace Noise Addition

The Laplace Mechanism

1 counting query: noise roughly 1

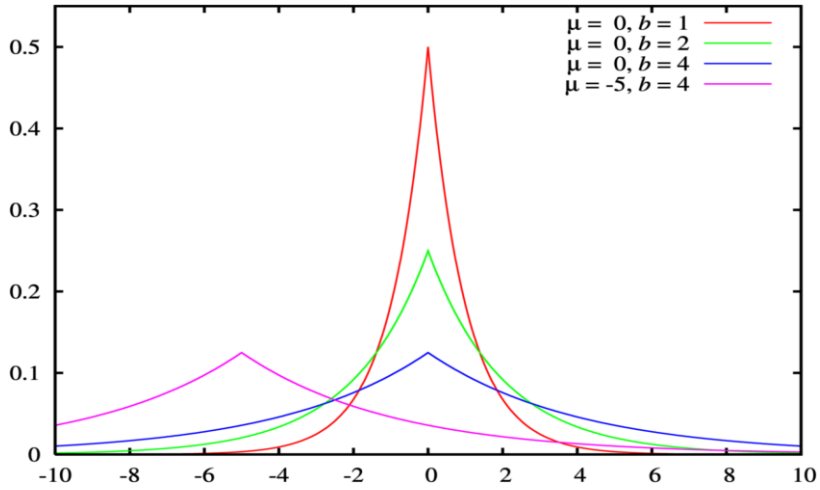

$$\Delta_1 = \max_{\text{adj } x,y} |f(x) - f(y)|$$

Theorem: Let $f: U^n \to R$.

$f(x) + \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$ yields ϵ-differential privacy.

The Laplace Mechanism: Proof

$$\frac{e^{|f(x)-t|/b}}{e^{|f(y)-t|/b}} \leq e^{\frac{|f(x)-f(y)|}{b}} \leq e^{\Delta_1/(\frac{\Delta_1}{\epsilon})}$$

$$\leq e^{\epsilon}$$

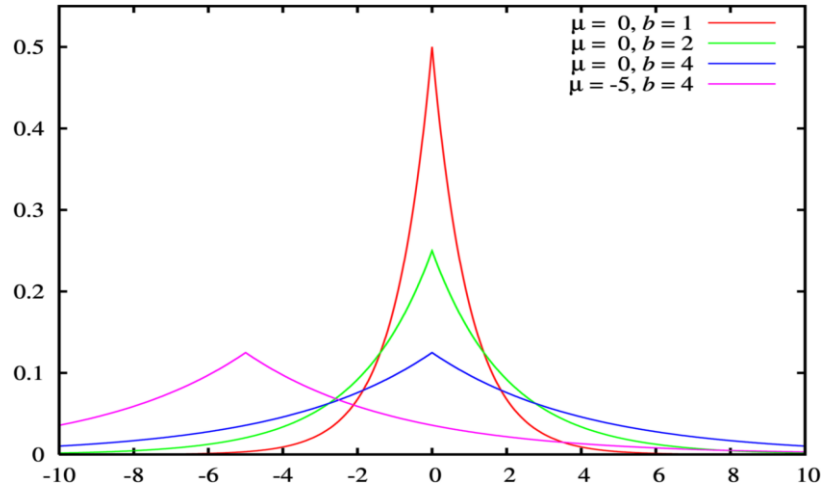

Theorem: Let $f: U^n \rightarrow R$.

$f(x) + \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$ yields ϵ-differential privacy.

The Laplace Mechanism

k counting queries: noise $\approx k$

$$\Delta_1 = \max_{\text{adj } x,y} ||f(x) - f(y)||_1$$

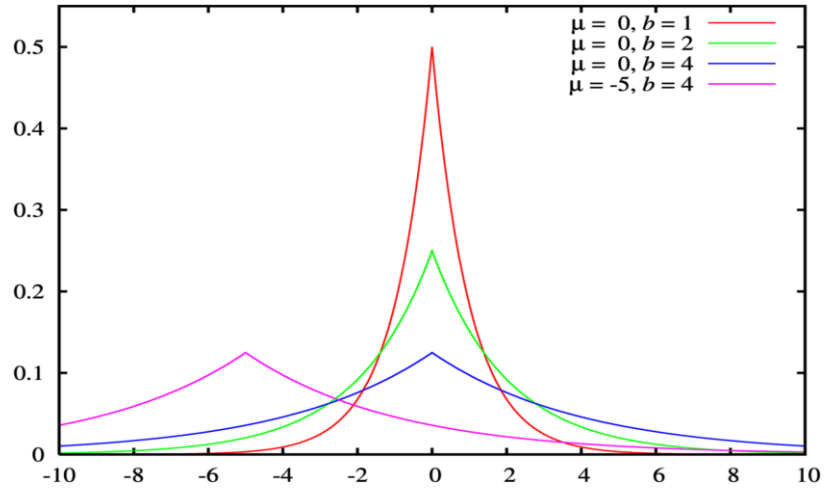

Theorem: Let $f: U^n \to R^k$.

$f(x) + \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)^k$ yields ϵ-differential privacy.

The Laplace Mechanism

k-cell histogram: noise roughly 1!

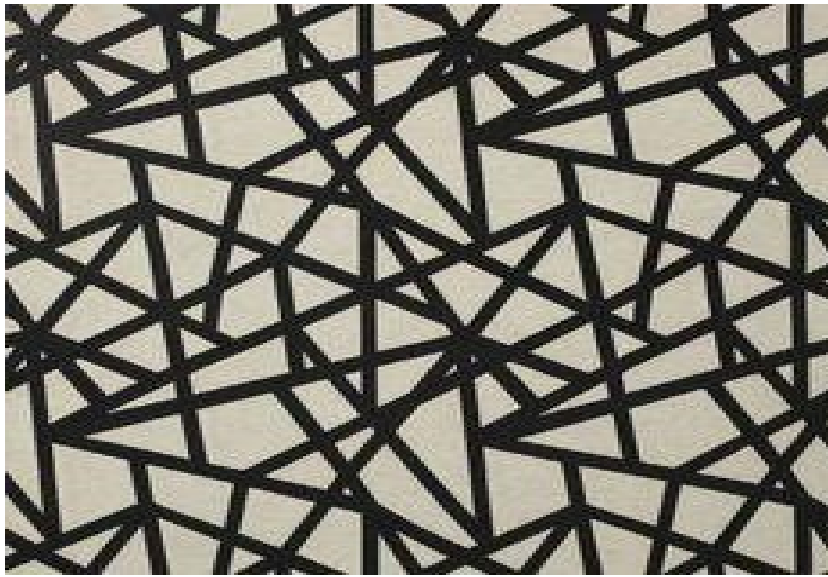

$$\Delta_1 = \max_{\text{adj } x,y} ||f(x) - f(y)||_1$$

Theorem: Let $f: U^n \to R^k$.

$f(x) + \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)^k$ yields ϵ-differential privacy.

Approximate Differential Privacy

M gives (ϵ, δ)-differential privacy if for all pairs of adjacent data sets x, y, and all output events S

$$\Pr[\text{see } S \text{ on } M(x)] \leq e^{\epsilon} \Pr[\text{see } S \text{ on } M(y)] + \delta$$

Dwork, Kenthapadi, McSherry, Mironov, and Naor 06

Why Relax?

- Lets us use Gaussian/Binomial noise
  - Advantage: noise depends on L2 sensitivity $\left|\left|f(x) - f(y)\right|\right|_2$; improve by factor of $\approx \sqrt{k}$
- Exploit that privacy loss is a random variable
  - Advanced composition: privacy loss $\leq \sqrt{k \log\left(\frac{1}{\delta}\right)}\ \epsilon$ (rather than $k\epsilon$) with probability $\geq 1 - \delta$
  - Concentrated DP (CDP), zCDP, Renyi DP: privacy loss rv is subgaussian
- General escape hatch
  - Use cryptography to simulate a trusted center
  - Make some nice connections to robust statistics (DP test for validity of assumptions)
  - Get rid of dependence on number of cells in L_∞ error for histograms

The Privacy Loss Random Variable

Fix adjacent x, y, draw $c \leftarrow M(x)$

$$\text{PrivacyLoss}(c) = \ln\left[\frac{\Pr[M(x) = c]}{\Pr[M(y) = c]}\right]$$

- ϵ-differential privacy ⇒ magnitude always bounded by ϵ
- Privacy loss will sometimes be negative!

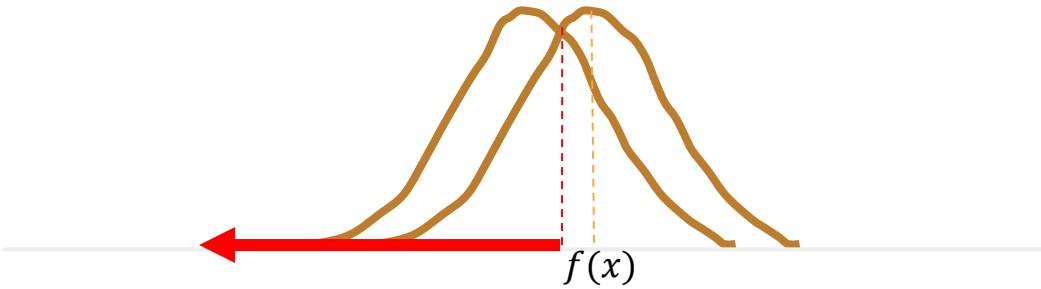

Gaussian Mechanism: $N(0, \sigma^2)$

- |Privacy Loss| $\leq \epsilon \Rightarrow |\ln \frac{e^{-z^2/2\sigma^2}}{e^{-(z+\Delta f)^2/2\sigma^2}}| \leq \epsilon$
- Suppose we take $\sigma = \Delta f / \epsilon$
  - $|z| < \sigma(1 - \frac{\epsilon}{2}) \Rightarrow$ |Privacy Loss| $\leq \epsilon$
  - To achieve (ϵ, δ)-DP, set $\sigma = \Delta f \sqrt{2 \ln 1/\delta}/\epsilon$
- What happens outside range $|z| < \sigma(1 - \frac{\epsilon}{2})$?
  - Death and destruction? No! Most likely, |Privacy Loss| is at most 2ϵ
- $\Pr[|\text{Privacy Loss}| \geq t\epsilon] \leq \Pr[|z| > \sigma(t - \frac{\epsilon}{2})] \approx \frac{1}{t} e^{-t^2/2}$

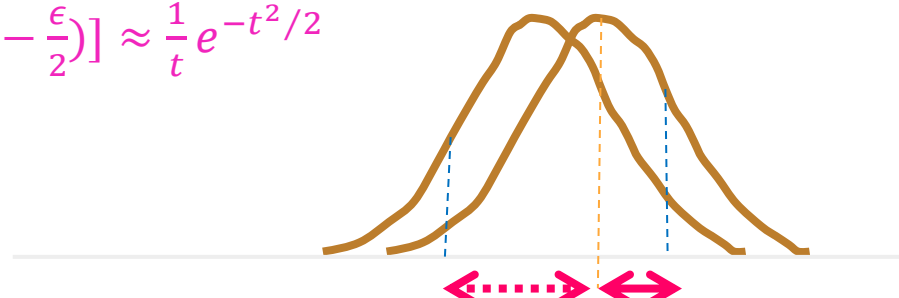

Composition

Adversary's Goal: Guess b

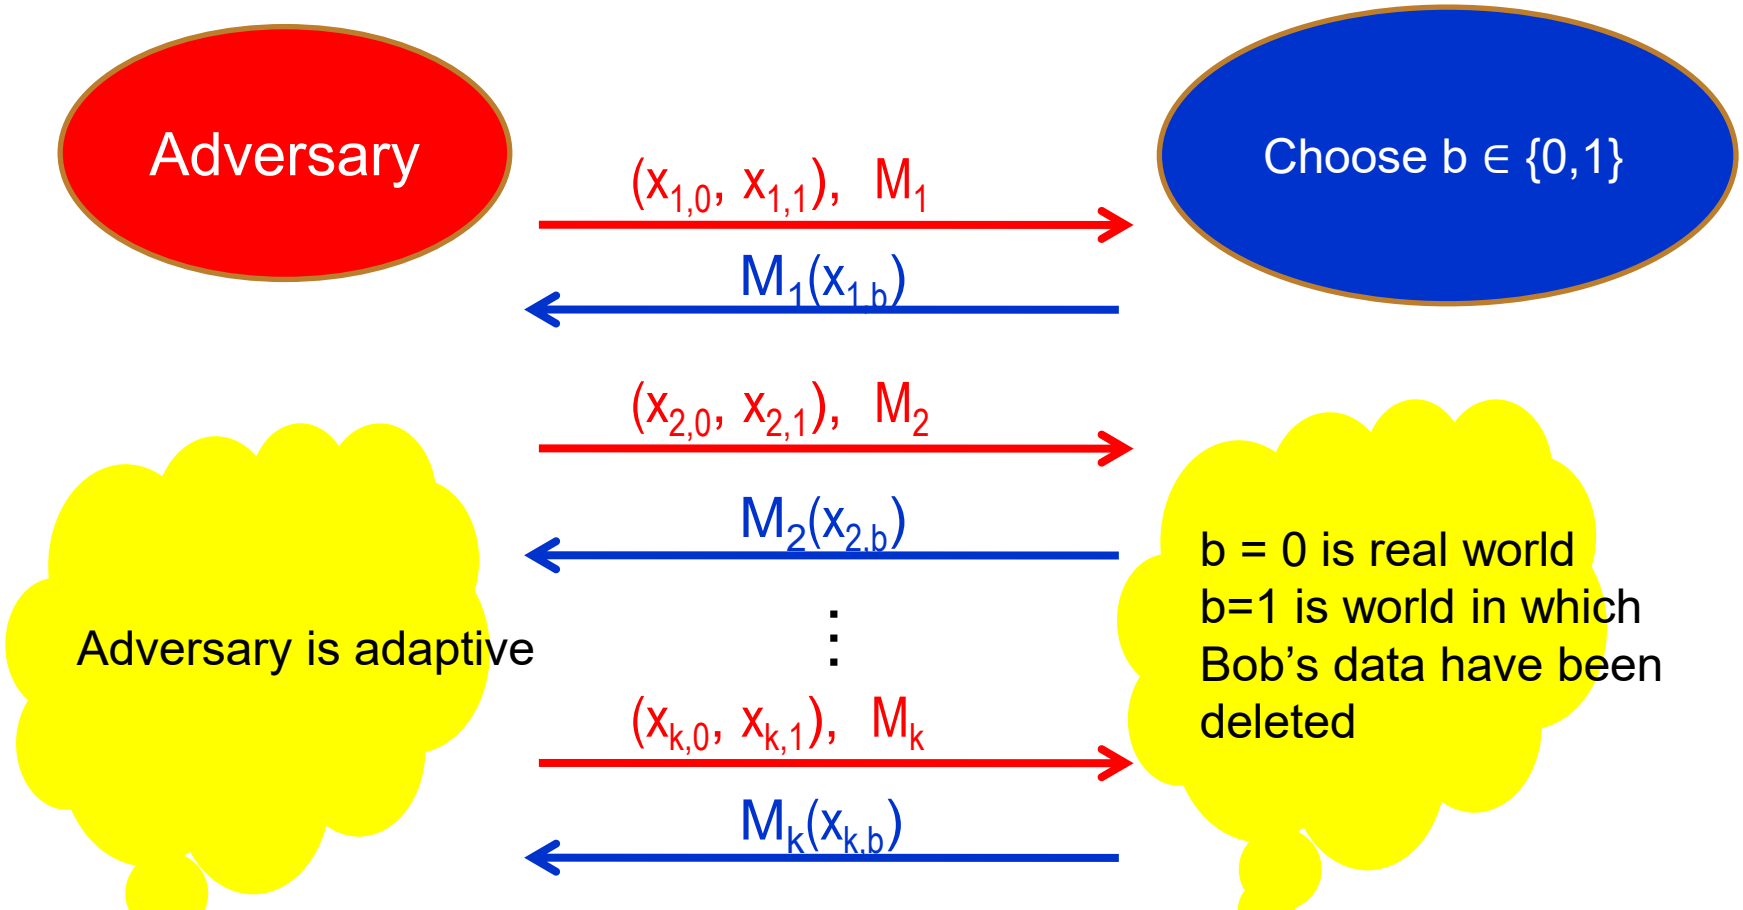

Advanced Composition Theorem

- Recall: privacy loss is sometimes negative ⇒ *there is cancellation!*
- For all ϵ, $\forall\delta$ simultaneously
  - $(k\epsilon^2 + \sqrt{k\log(1/\delta)}\,\epsilon, \delta)$-DP: loss is $O(\sqrt{k})\epsilon$ rather than $k\epsilon$
  - $(k\epsilon^2 + \sqrt{k\log(1/\delta)}\,\epsilon, k\delta' + \delta)$-DP
- Proof in 2 steps. For each ϵ-DP mechanism:
  - Expected loss is $\leq \epsilon^2$
  - Magnitude of loss is bounded by ϵ
    - Martingale, Azuma-Hoeffding
- "$\forall t > 0$: $\Pr[t$ standard deviations beyond expectation$] \leq e^{-t^2/2}$"

Dwork, Rothblum, Vadhan 2010

Advanced Composition Theorem

- Recall: privacy loss is sometimes negative ⇒ *there is cancellation!*
- For all $\epsilon, \forall\delta$ simultaneously
  - $(k\epsilon^2 + \sqrt{k\log(1/\delta)}\,\epsilon, \delta)$-DP: loss is $O(\sqrt{k})\epsilon$ rather than $k\epsilon$
  - $(k\epsilon^2 + \sqrt{k\log(1/\delta)}\,\epsilon, k\delta' + \delta)$-DP
- Proof in 2 steps. For each ϵ-DP mechanism:
  - Expected loss is $\leq \epsilon^2$
  - Magnitude of loss is bounded by ϵ
    - Martingale, Azuma-Hoeffding
- "$\forall t > 0$: $\Pr[t$ standard deviations beyond expectation$] \leq e^{-t^2/2}$"
- Starting point for Concentrated Differential Privacy

Dwork, Rothblum, Vadhan 2010

Concentrated Differential Privacy: Intuition

- To apply the composition theorem and have cumulative loss ϵ, force each "little" ϵ' to be $\leq \epsilon/\sqrt{k}$ with very high probability
- Concentration: under high degrees of composition, can be less conservative in each computation – no death and destruction if we "miss" some of the bounds a little bit, and we still have cancellation
- Bottom line: can get rid of the $\sqrt{\ln 1/\delta}$ factor in the standard deviation for the Gaussian mechanism

DworkRothblum'15, improved by BunSteinke'15

Rich Algorithmic Literature

- Counts, linear queries, histograms, contingency tables (marginals)
- Auctions for digital goods
- Location and spread (eg, median, high dimensional median, interquartile range)
- Dimension reduction (PCA, SVD), clustering
- Sparse regression/LASSO, logistic and linear regression
- Boosting, Multiplicative Weights, PAC learning
- Combinatorial optimization, mechanism design
- Privacy Under Continual Observation, Pan-Privacy, heavy hitters
- Finite sample confidence intervals
- Synthetic data generation
- Graph analyses
- Deep learning & gradient descent
- New models: distributed, federated, shuffle
- Hypothesis selection
- ...

Applications Beyond Privacy

- Econ/CS: First collusion-resilient solution for online auctions
- General solution for validity of exploratory data analysis
  - The “reusable holdout”
- Gentle measurements in quantum tomography

McSherry and Talwar
Dwork, Feldman, Hardt, Pitassi, Reingold, Roth
Aaronson and Rothblum

Differential Privacy In Practice

- Apple, Google, Microsoft, Uber, Facebook, LinkedIn
- OnTheMap
- Census-based research

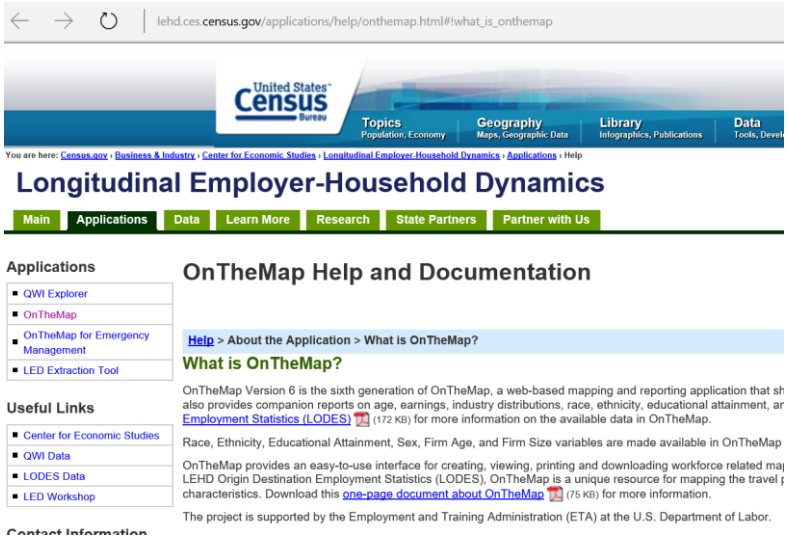

ChettyFriedmanHendranJonesPorter2018
MachanavajjhalaKiferAbowdGehrkeVilhuber2008

The 2020 Census

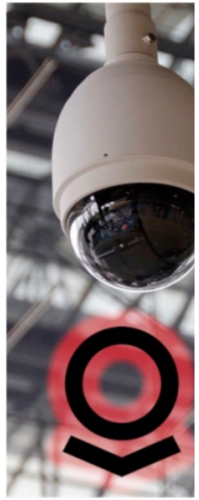

DP and the US 2020 Decennial Census

- The techniques used in 2010 do not suffice

"… technical advances revealed a new vulnerability, allowing people to reconstruct data from tables that were previously assumed to be privacy preserving…"

John Abowd, Chief Scientist and Associate Director of Research and Methodology, US Census Bureau

Staring Down the Database Reconstruction Theorem

John M. Abowd
Chief Scientist and Associate Director for Research and Methodology
U.S. Census Bureau
American Association for the Advancement of Science
Annual Meeting Saturday, February 16, 2019 3:30-5:00

United States Census Bureau
U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov

The views expressed in this talk are my own and not those of the U.S. Census Bureau.

What we did

- Database reconstruction for all 308,745,538 people in 2010 Census
- Link reconstructed records to commercial databases: acquire PII
- Successful linkage to commercial data: putative re-identification
- Compare putative re-identifications to confidential data
- Successful linkage to confidential data: confirmed re-identification
- Harm: attacker can learn self-response race and ethnicity

Slide: John Abowd, 2019


Image credit: Klauss Budloff

We fixed this for the 2020 Census by implementing differential privacy

Slide: John Abowd, 2019

Challenges

Privacy Budget and Allocation

- How (and who) to choose the “privacy budget”?
  - *Data consumers* do not advocate for privacy!
- How to allocate the budget?

Techniques Known and Not Yet known

- Data consumers – demographers, gerrymanderers & voting rights advocates, urban planners, social scientists – are not trained to interact with data in a differentially private way
- Basic tools of statistical inference in the presence of noisy statistics still need to be developed
- A first, and easiest step, in learning to deal with uncertainty

Thank you!

Harvard University, October 5, 2021