Differential Privacy
What is it and Where is it?

Cynthia Dwork
Harvard University
Radcliffe Institute for Advanced Study

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.

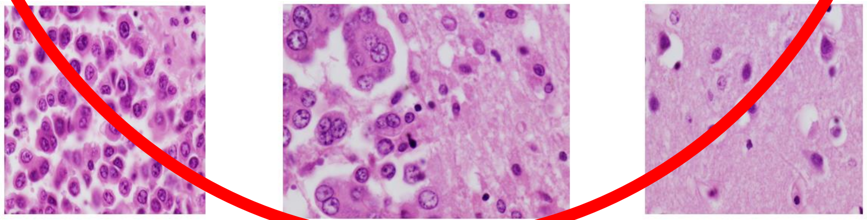

In November 2002, the New York Times reported tha
(DARPA) was developing a tracking system called "To
intended to detect terrorists through analyzing trove

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 <u>every</u> computation, including total population counts

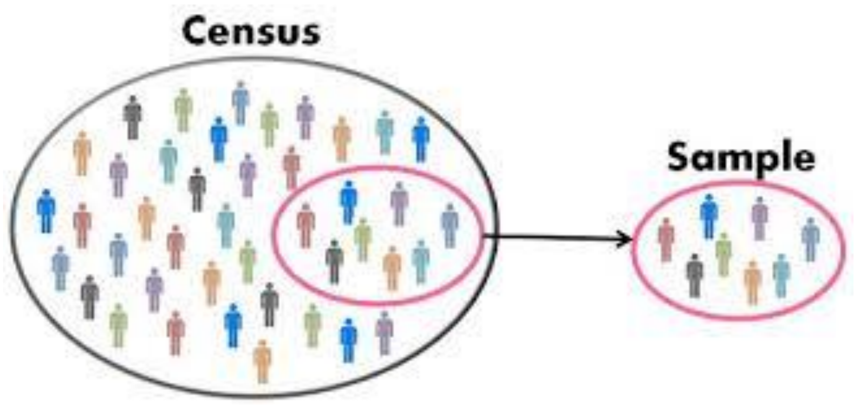

Image credit: Keydifferences.com

Abstracting The Problem

Privacy-Preserving Data Analysis

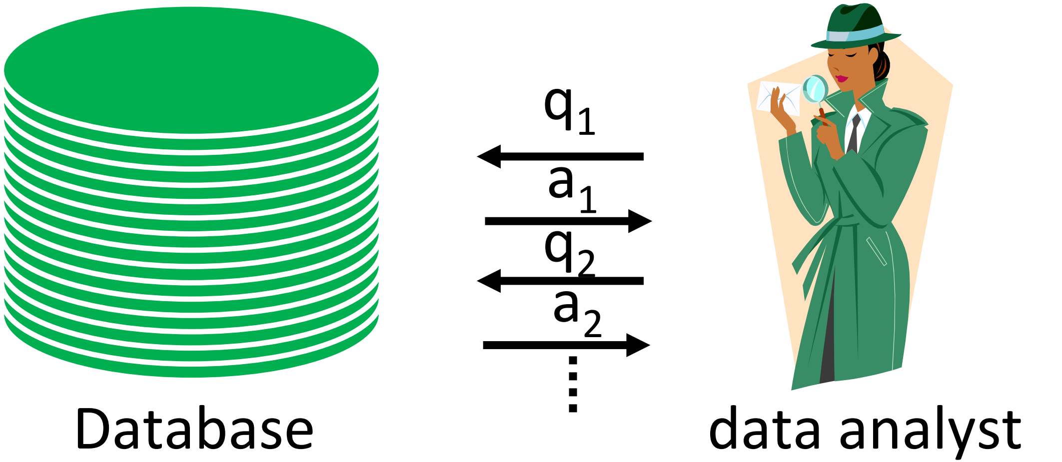

- Driving scenario: analysis of US Census data
- 55+ year old problem

Image credit: The New Yorker Magazine

"Just" Statistics

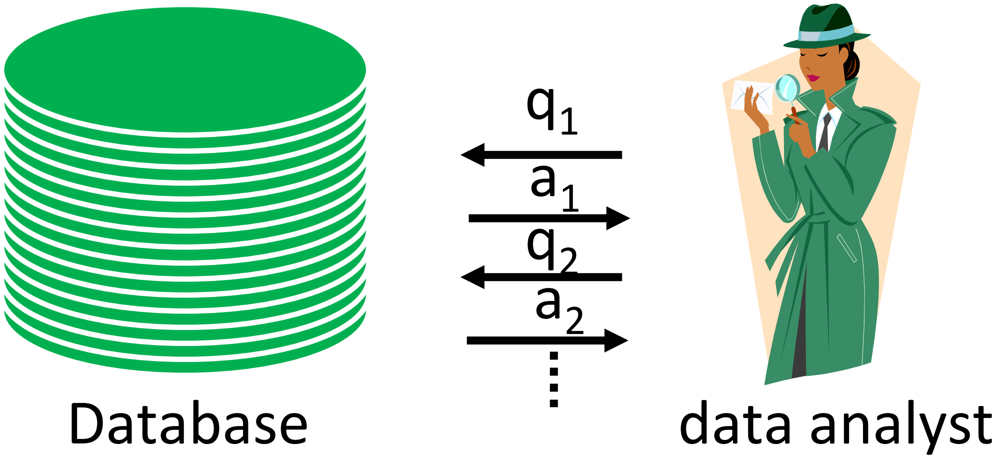

- How many living physics Nobel Laureates floss regularly?
- How many male living physics Nobel Laureates floss regularly?

Fundamental Law of Info Recovery

- “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?

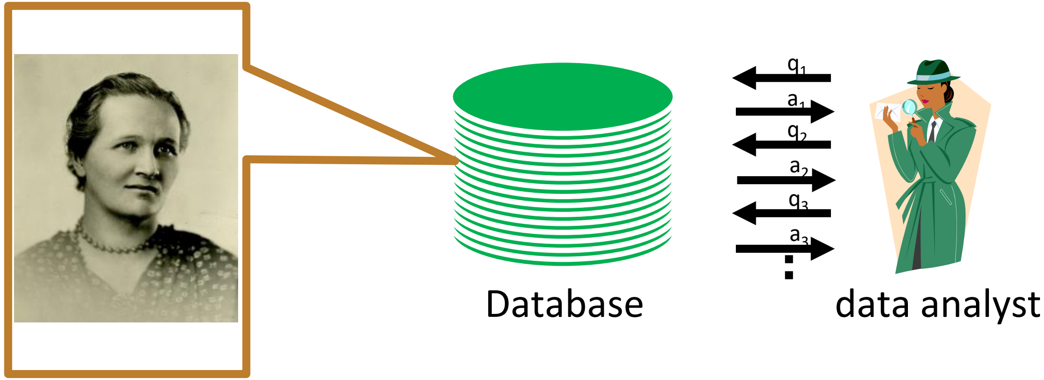

Privacy-Preserving Data Analysis?

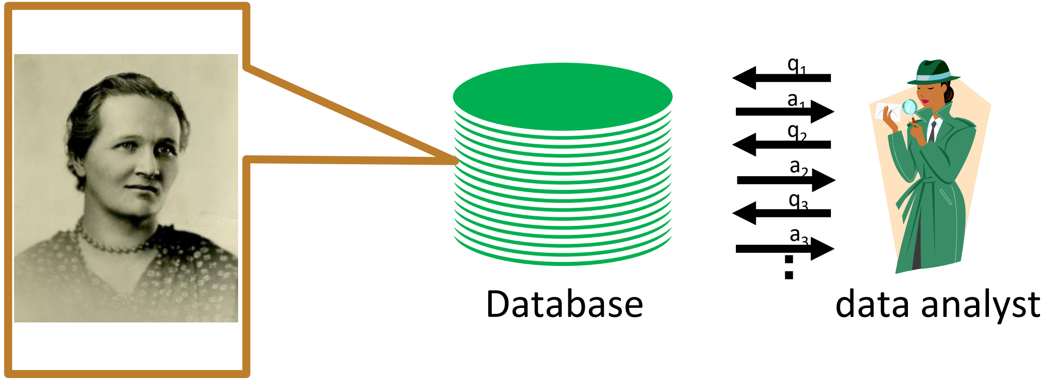

- “Can’t learn anything new about Payne”?
- Dalenius, 1977

Privacy-Preserving Data Analysis?

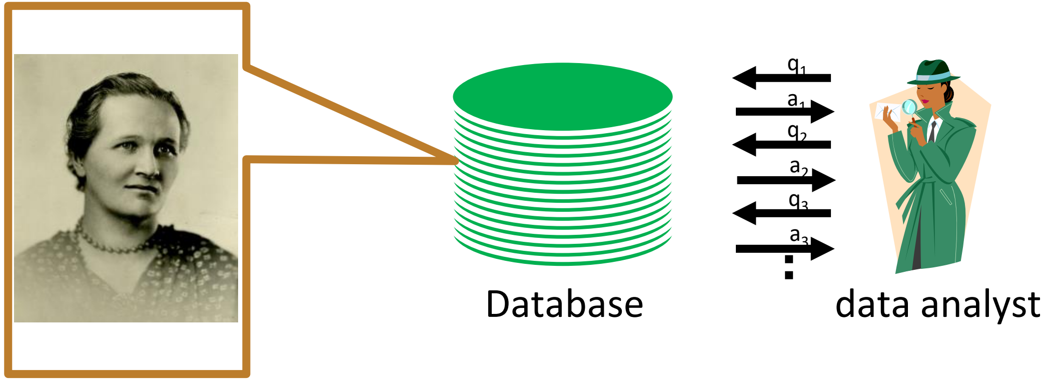

- "Can't learn anything new about Payne"?
- Then what is the point?

Privacy-Preserving Data Analysis?

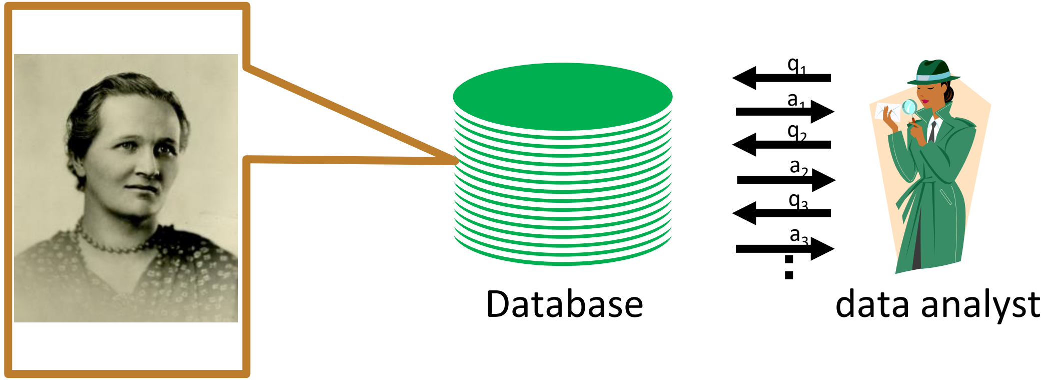

- “Can’t learn anything new about Payne”?
- Then what is the point?

Privacy-Preserving Data Analysis?

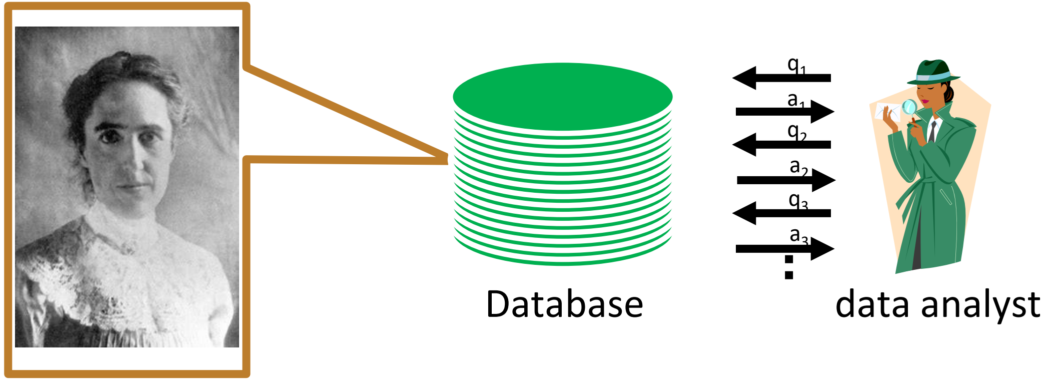

- Ideally: learn same things if Payne is replaced by another random member of the population

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.

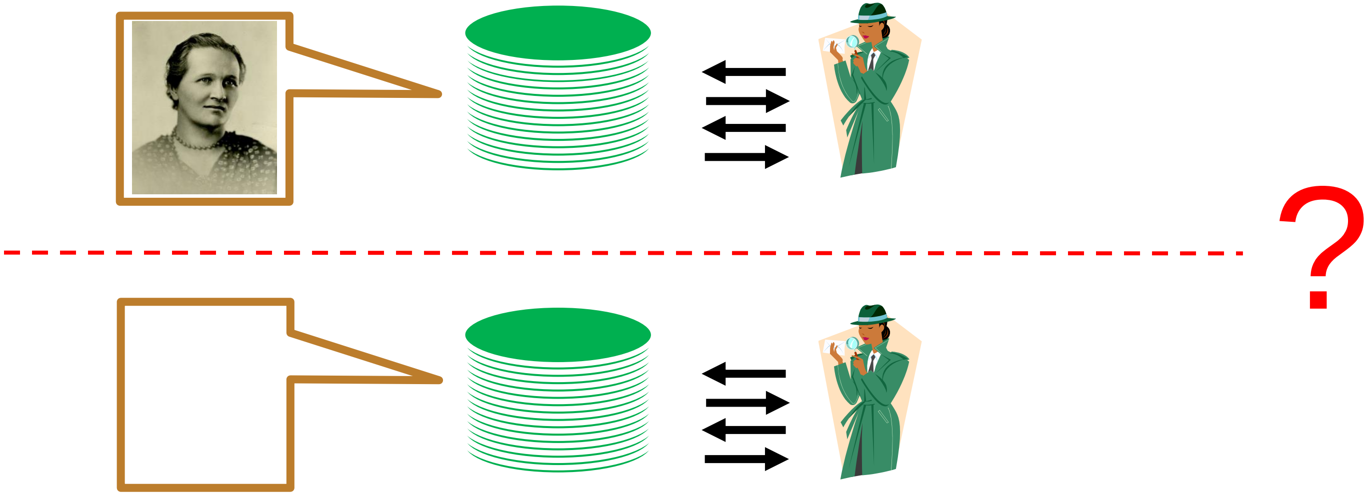

Differential Privacy

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

- An adversary knowing both datasets in their entirety can never distinguish
- What is the source of uncertainty in “likely”?

Differential Privacy

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

- An adv heir entirety can never disting

Algorithm will flip coins

- What is the source of uncertainty in “likely”?

Differential Privacy

- "Essentially equally likely"
  - Fair coin vs slightly biased coin, say, 500/1000 vs 501/1000
  - Given one of these coins, the "flipping algorithm" produces either Heads or Tails; the probability distribution on outcomes is nearly the same for the two coins
  - Given one (or even several) flip outcomes, can never determine which was true coin

Privacy-Preserving Data Analysis?

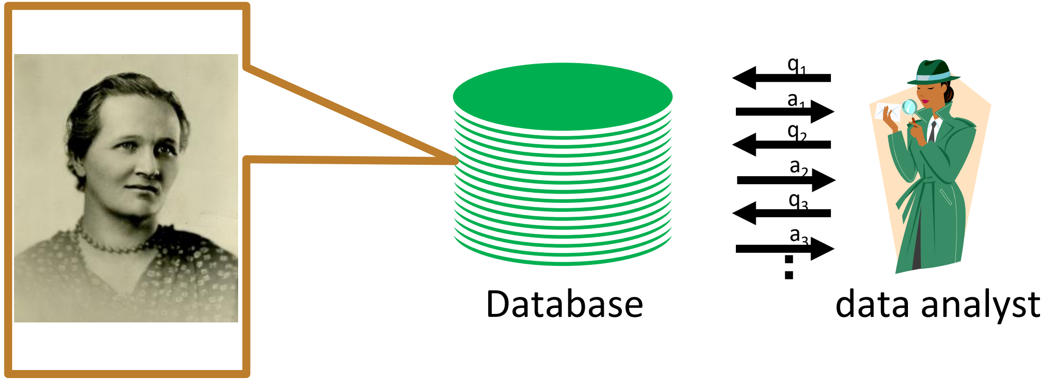

Stability preserves Payne's privacy AND prevents over-fitting

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

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

"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

$$\Pr[\text{see } S \text{ on } M(x)] \leq (1+\epsilon) \Pr[\text{see } S \text{ on } M(y)]$$

$e^{\epsilon} \approx 1 + \epsilon$ when ϵ is small

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 <u>can</u> learn about Payne

You can only learn things you can learn without Payne

Dwork, McSherry, Nissim, Smith 2006

Teachings vs Participation

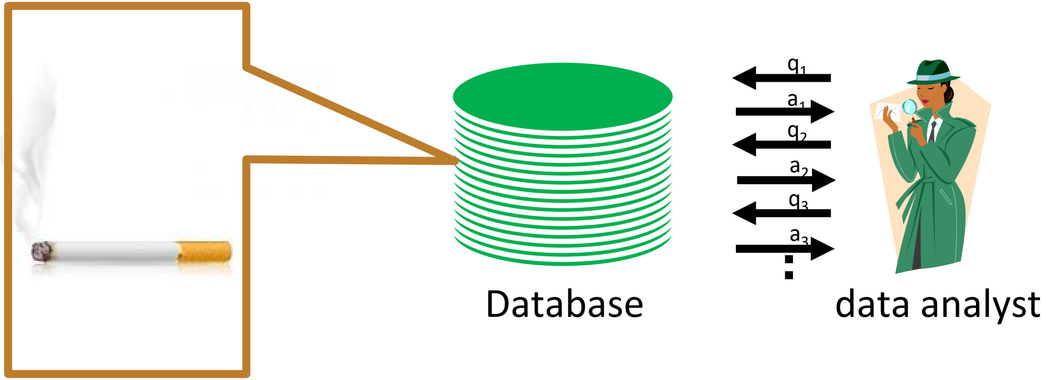

SURGEON GENERAL'S WARNING: Smoking Causes Lung Cancer, Heart Disease, Emphysema, and May Complicate Pregnancy.

Differential Privacy Hides the Needle

$\underbrace{\text{Haystack}}_{x}$ vs $\underbrace{\text{Haystack sans needle}}_{y}$

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 <u>programmable</u>

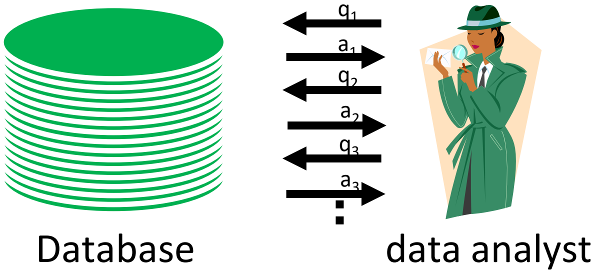

One Technique:
Laplace Noise Addition

What is the Population of Washington, DC?

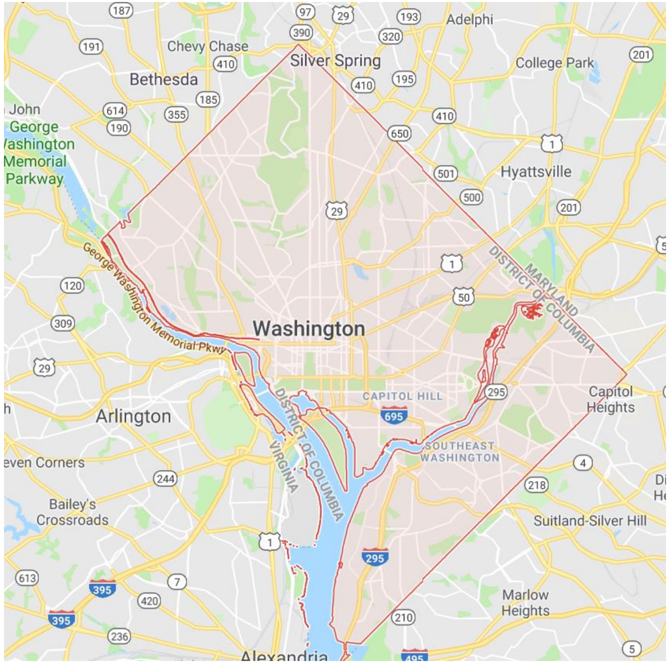

Population estimates, July 1, 2017, (V2017)	**693,972**
PEOPLE	
Population	
Population estimates, July 1, 2017, (V2017)	**693,972**
Population estimates base, April 1, 2010, (V2017)	601,766
Population, percent change - April 1, 2010 (estimates base) to July 1, 2017, (V2017)	15.3%
Population, Census, April 1, 2010	601,723
Age and Sex	
Persons under 5 years, percent	6.5%
Persons under 18 years, percent	17.9%
Persons 65 years and over, percent	12.1%
Female persons, percent	52.6%
Race and Hispanic Origin	
White alone, percent (a)	45.1%
Black or African American alone, percent (a)	47.1%
American Indian and Alaska Native alone, percent (a)	0.6%
Asian alone, percent (a)	4.3%
Native Hawaiian and Other Pacific Islander alone, percent (a)	0.1%
Two or More Races, percent	2.7%
Hispanic or Latino, percent (b)	11.0%
White alone, not Hispanic or Latino, percent	36.8%
Population Characteristics	
Veterans, 2012-2016	27,754
Foreign born persons, percent, 2012-2016	14.0%

https://www.census.gov/quickfacts/fact/table/dc/PST045217

True Count 601,723

True Count 601,723

- One person opting out moves the count to 601,722

True Count 601,723

- One person opting out moves the count to 601,722
- Add random noise to obscure the difference
    - 601,722 vs 601,723

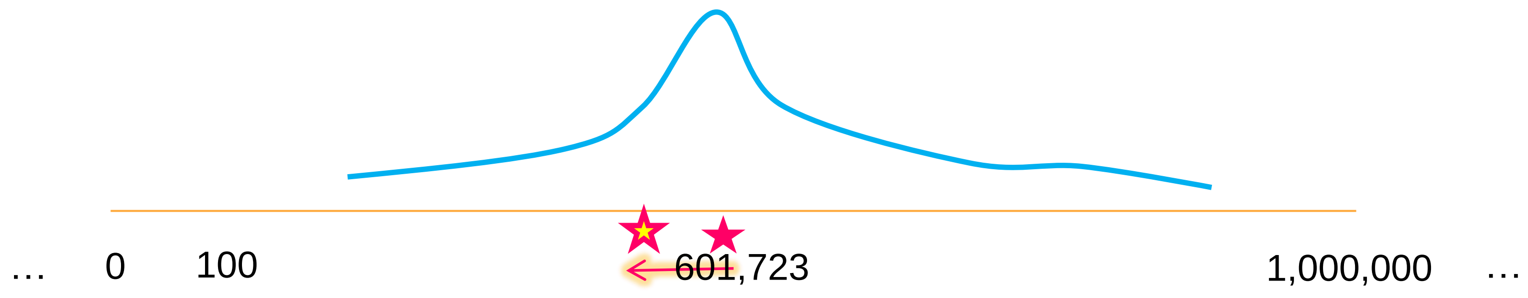

True Count 601,723

- One person opting out moves the count to 601,722
- Add random noise to obscure the difference
  - 601,722 vs 601,723
- How "fat" should these curves be?
  - How much can one person affect the count? $\Delta = 1$
  - How tightly do we want to restrict the privacy loss? $\epsilon = 0.01?$ $\epsilon = 0.1?$
  - Answer: Δ/ϵ

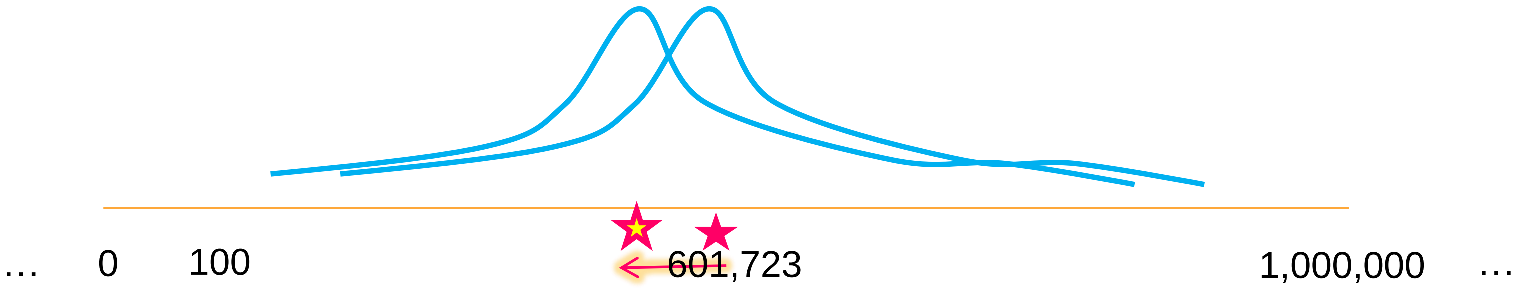

Laplace Noise with Parameter Δ/ϵ

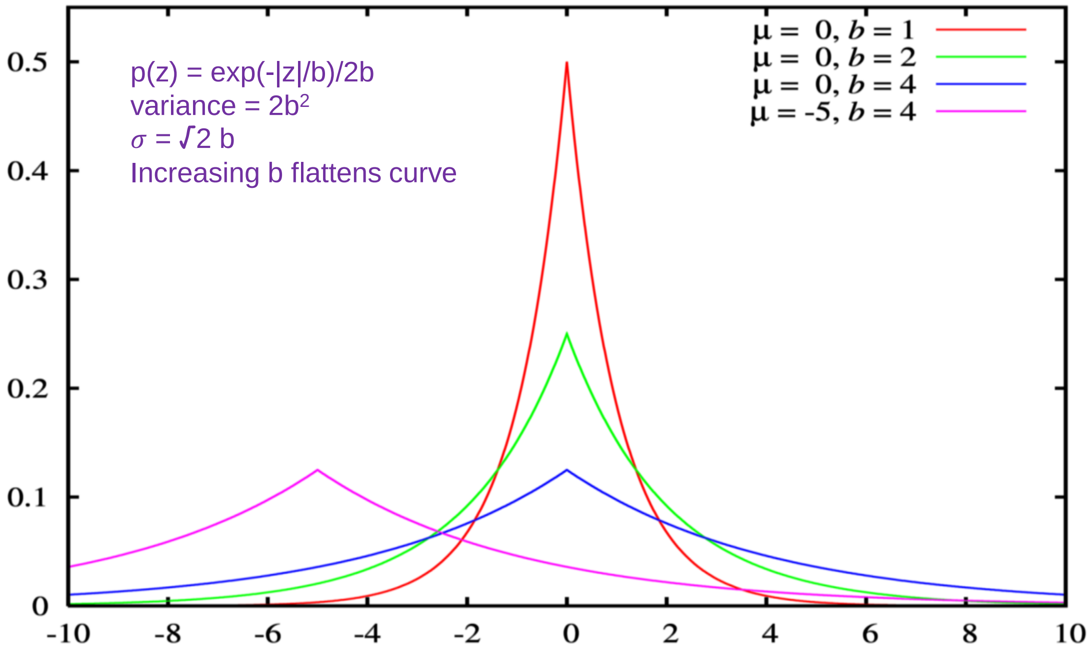

The Local Model
Privacy "rolled in" before collection

Did You Floss Last Night?

- Flip a fair coin.
  - Heads: Flip again and respond "Yes" if heads, "No" if otherwise
  - Tails: Answer honestly

- Privacy Analysis:
  - Pr [say "Y" given that truth = Y] / Pr [say "Y" given that truth = N] = 3
    - If truth is Y, will say "Y" if first coin is tails (probability ½) or first coin is heads and second coin is heads (probability ½ × ½ = ¼), total probability ¾
    - If truth is N, will say "Y" only if first and second coins are heads, probability ¼
  - Pr [say "N" given that truth = N] / Pr [say "N" given that truth = Y] = 3

- Reverse engineer the noise to learn approximate flossing fraction:
  - Key observation: [#True "Y" among the ≈ n/2 answering honestly] ≈ #"Y" – n/4

Warner 1965

Differential Privacy Deployed

Centralized Model

- Google, Facebook: Covid-19 mobility data, e.g., response to work-from-home, stay-at-home
- Microsoft: Windows reporting and ML; Office predictive text language models, Office Workplace analytics
- LinkedIn: publisher tools
- FaceBook: database of popular URLs
- Uber: determining average trip distance
- Research: Opportunity Atlas

Local Model

- Google: RAPPOR / Fuschia.cobalt hundreds of data analytics metrics
- Microsoft: Windows telemetry for user experiences; Machine learning of predictive models
- Apple: new words, emojis, deeplinks, lookup hints inside notes; Health type usage; Safari Autoplay Intent Detection; energy-draining and crashing domains

Industry Platforms:

- Private TensorFlow (Google)
- Microsoft + Harvard open source platform
- Private SQL (Google)
- Flex (Uber)
- DiffPrivLib (IBM)

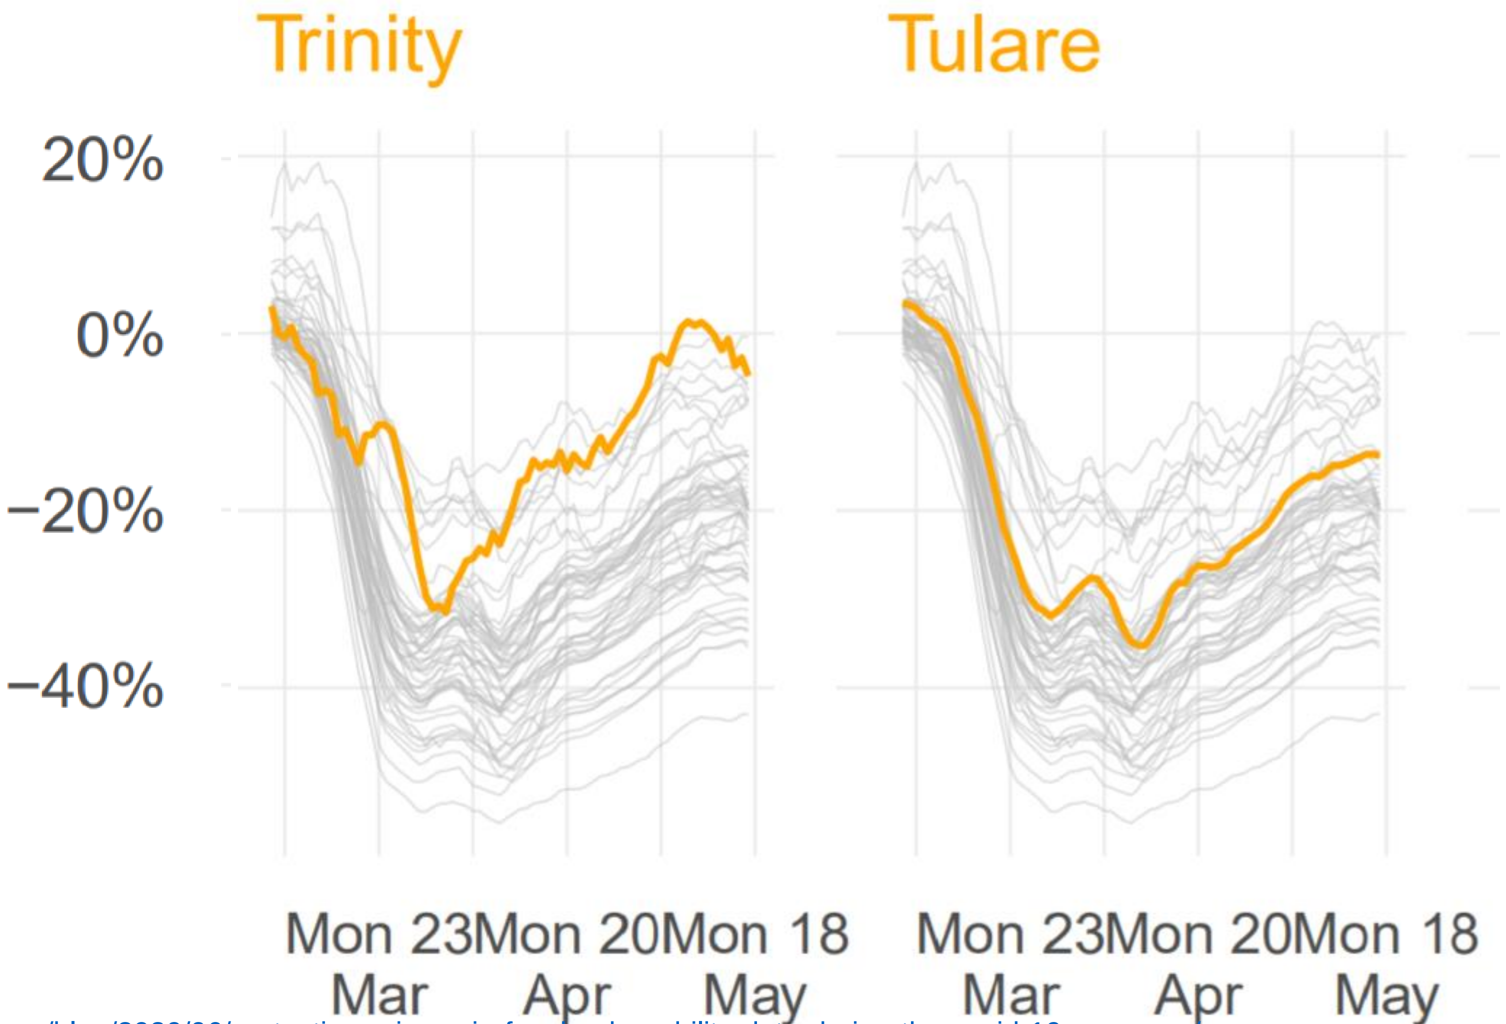

https://research.fb.com/blog/2020/06/protecting-privacy-in-facebook-mobility-data-during-the-covid-19-response/

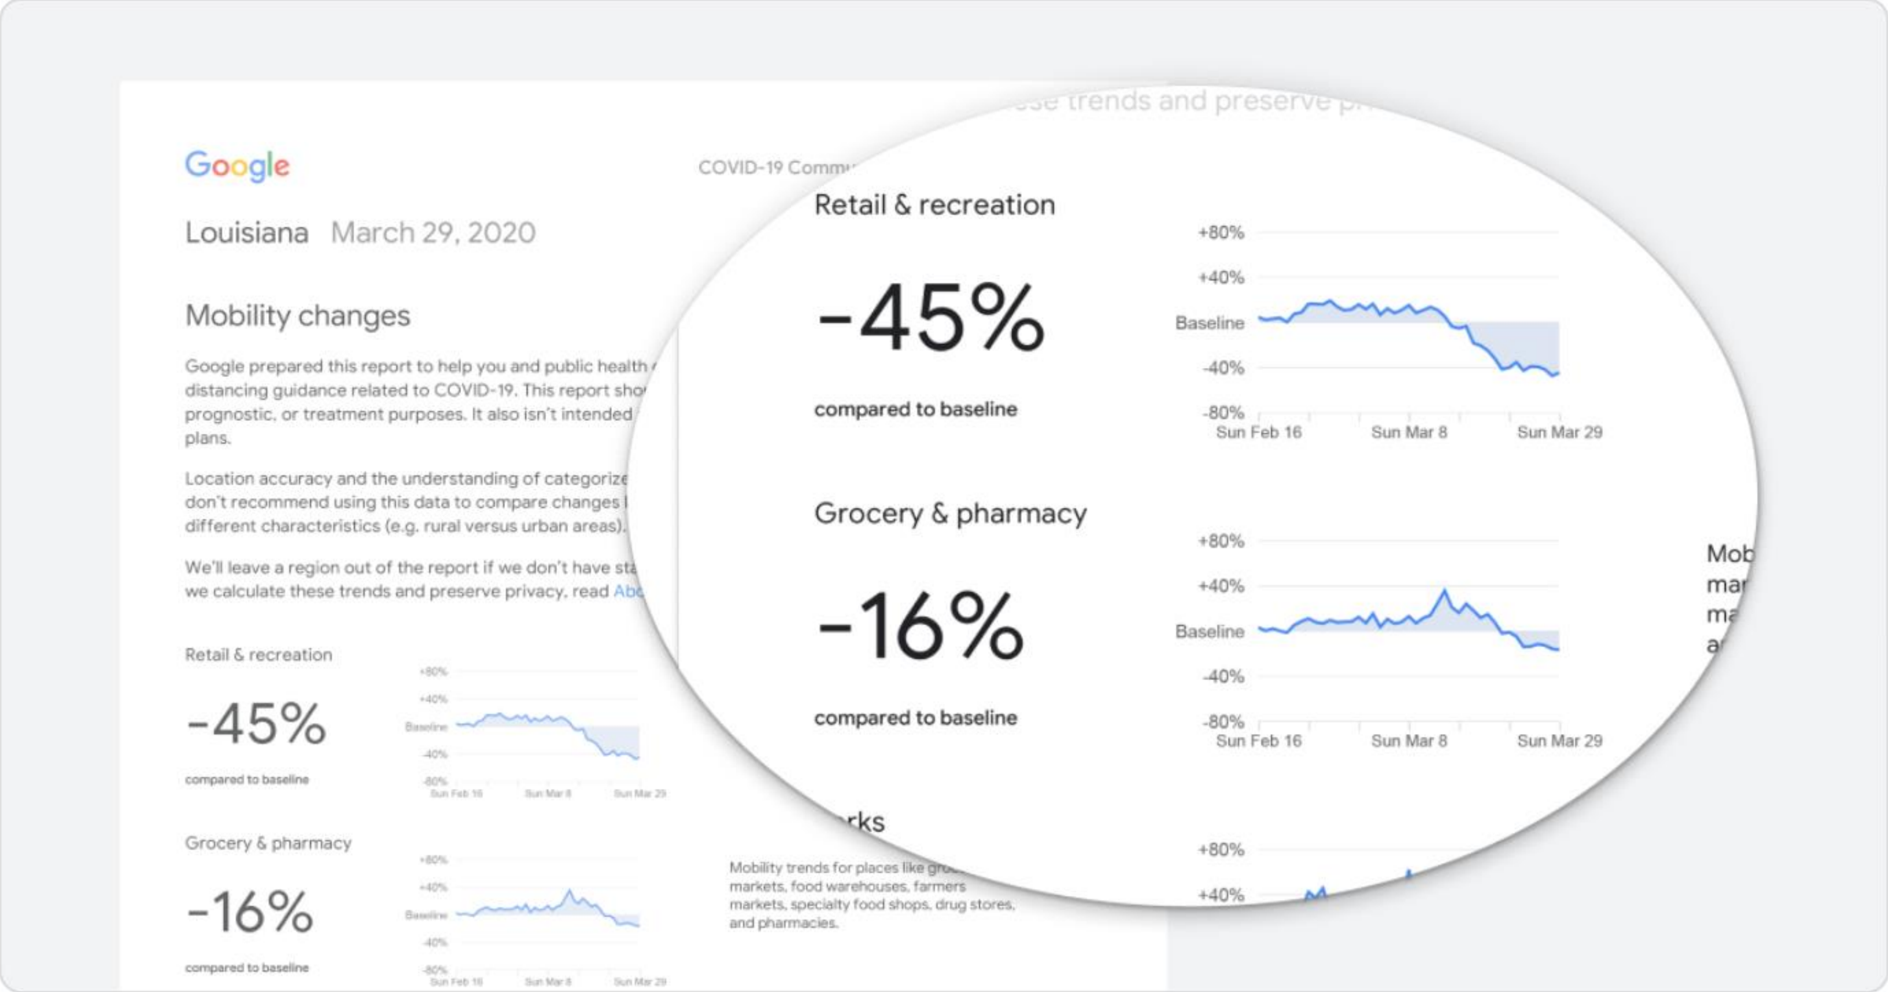

https://www.google.com/covid19/mobility/

The Census

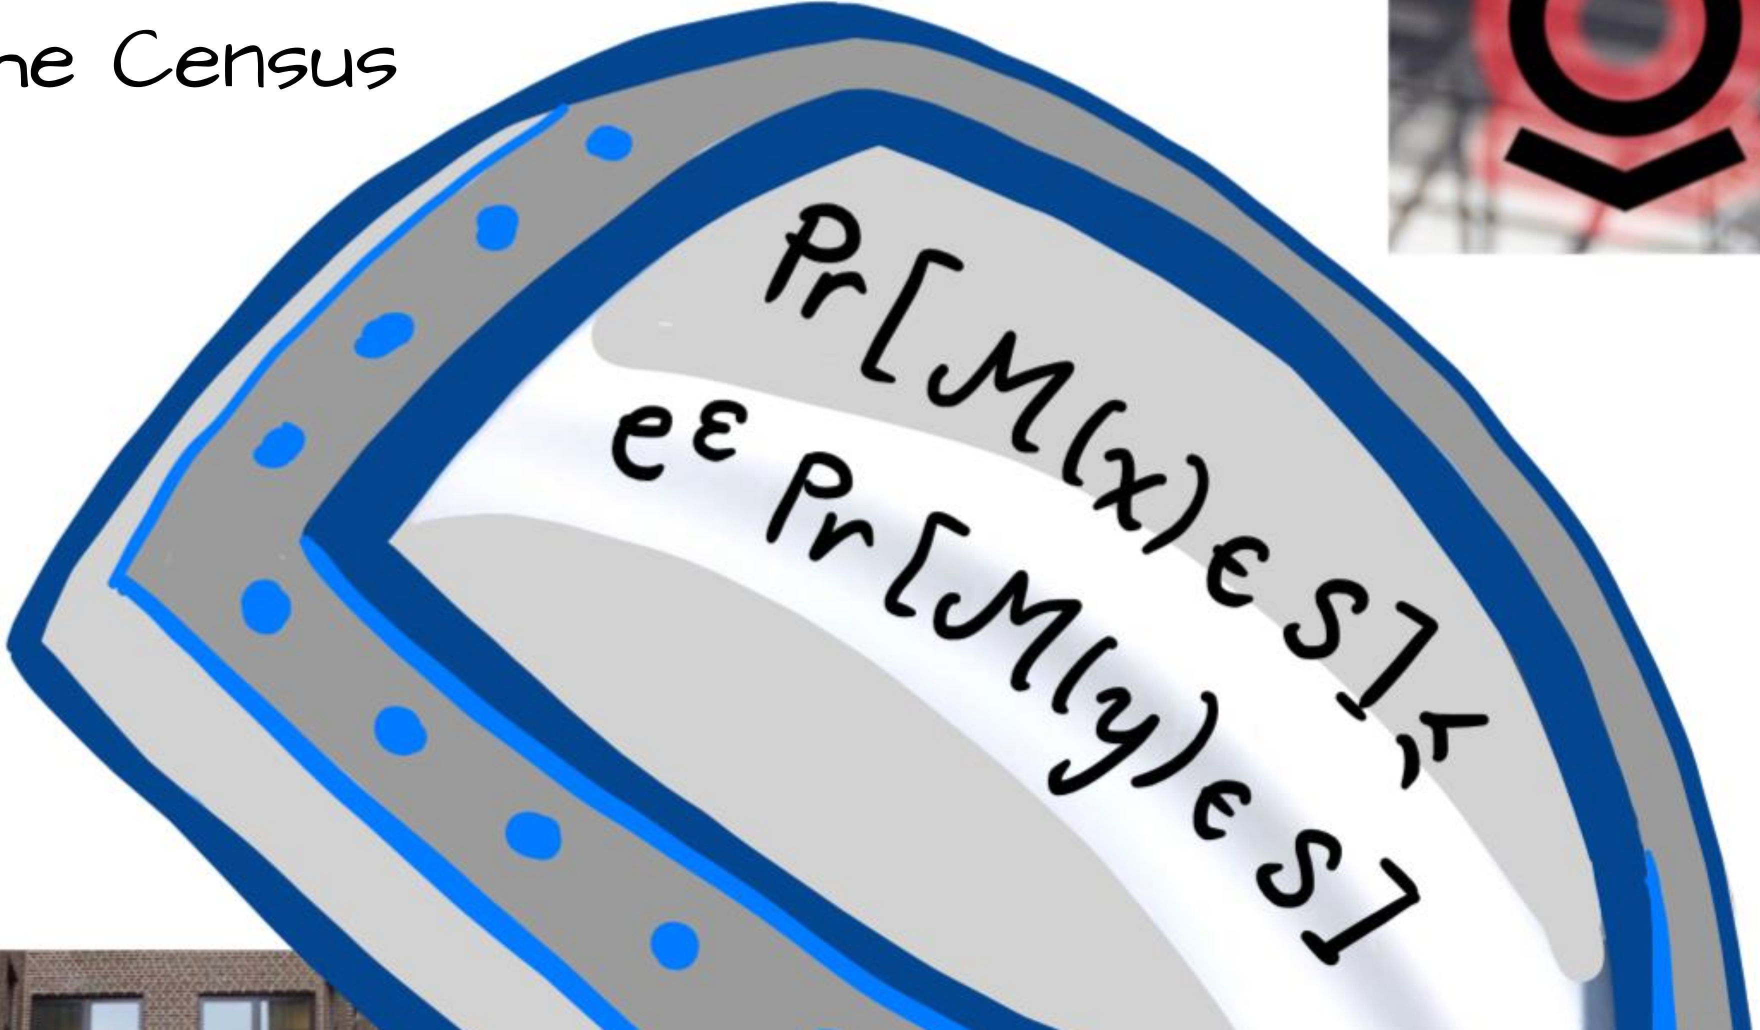

The US 2010 Decennial Census

- 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

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.

Slide: John Abowd, 2019

What they did

- Database reconstruction for all 308,745,538 people in 2010 Census
- Link reconstructed records to 2010 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
    - 38% of putative (52 million; 17% of population)
- Harm: attacker can learn self-response race and ethnicity

Trust

Every Census Bureau employee takes a lifetime oath to protect your personal identification. Disclosing ANY information that could identify you or your family means 5 years in prison, or $250,000 in fines, or both.

https://www.census.gov/programs-surveys/acs/about/acs-and-census.html

We fixed this for the 2020 Census by implementing differential privacy

Slide: John Abowd, 2019

Jubilation?

Not

Challenges

- Researchers – historians, sociologists, demographers, economists – are not trained to interact with data in a differentially private way
  - Post-processing to create synthetic data, with non-negative, integer, marginals, introduces statistical bias
  - Analysts must be trained to adjust for noise
    With DP, <u>can</u> adjust: the generation mechanism is known!
  - The Fundamental Law persists (it is a law)

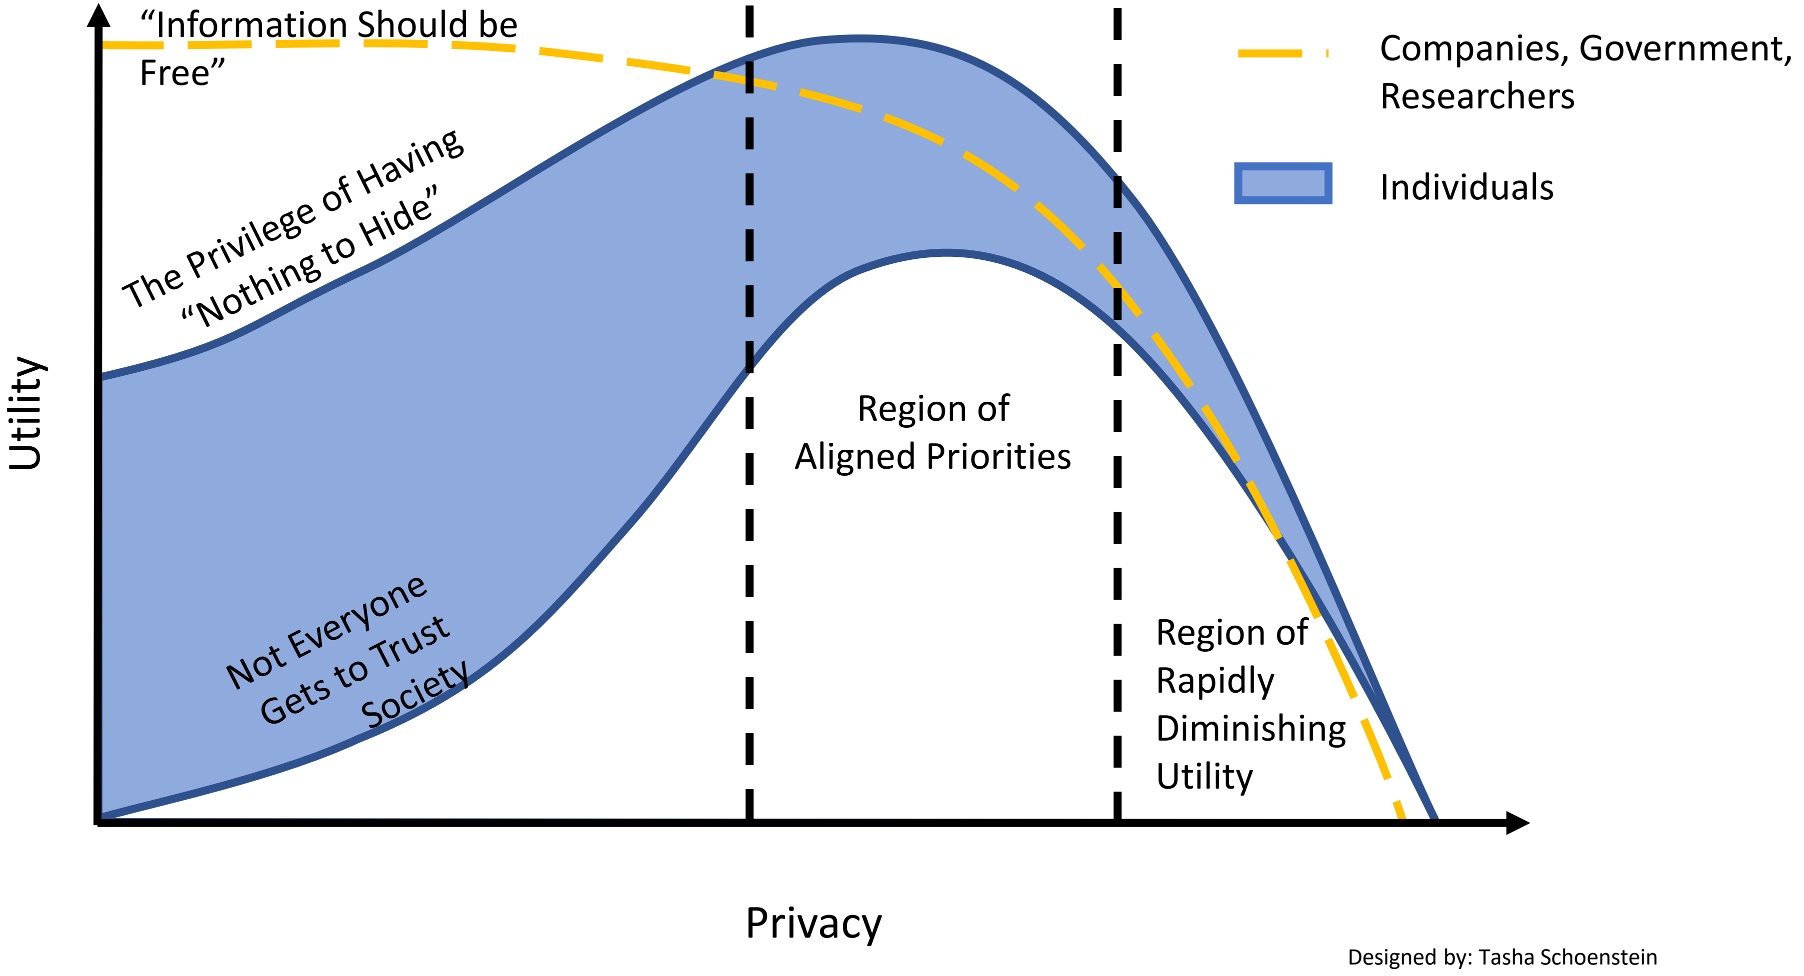

“Information Should be Free”
The Privilege of Having “Nothing to Hide”
Not Everyone Gets to Trust Society
Region of Aligned Priorities
Region of Rapidly Diminishing Utility
Companies, Government, Researchers
Individuals
Utility
Privacy
Designed by: Tasha Schoenstein

Allocation of the Privacy Resource

- Privacy Budget: Cap on cumulative privacy loss
  - How (and who) to choose the privacy budget?
  - How to prioritize the spending on different queries?

Reprise: This Talk in a Nutshell

- The intuition behind privacy of a sample is on the right track but needs help
- Differential privacy provides that help
- Differential privacy turns every calculation into a statistic with “opt-out” semantics.

- DP hides the needle, reveals the population as a whole

Thank You!

Seoul National University and Cyberspace
September 21-22, 2020