Lecture 21: Differential Privacy

CS 138

Spring 2026

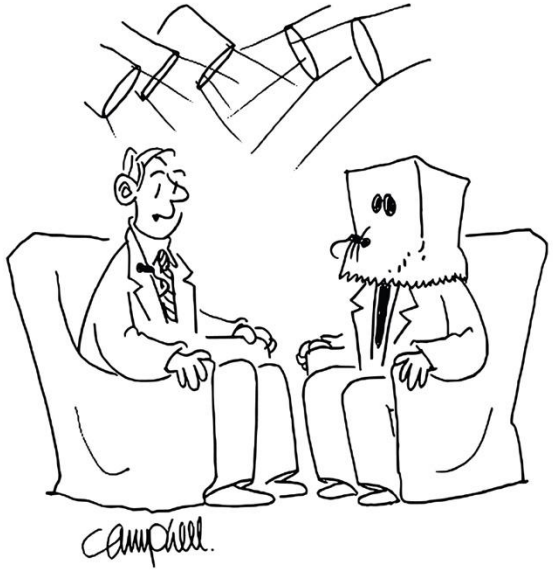

"Our guest has asked that we obscure his identity."

What is Privacy?

Privacy

Privacy concerns information about individuals (people, organizations, etc.)

- Often construed as legal right
- *Privacy* is not a synonym for confidentiality or for secrecy

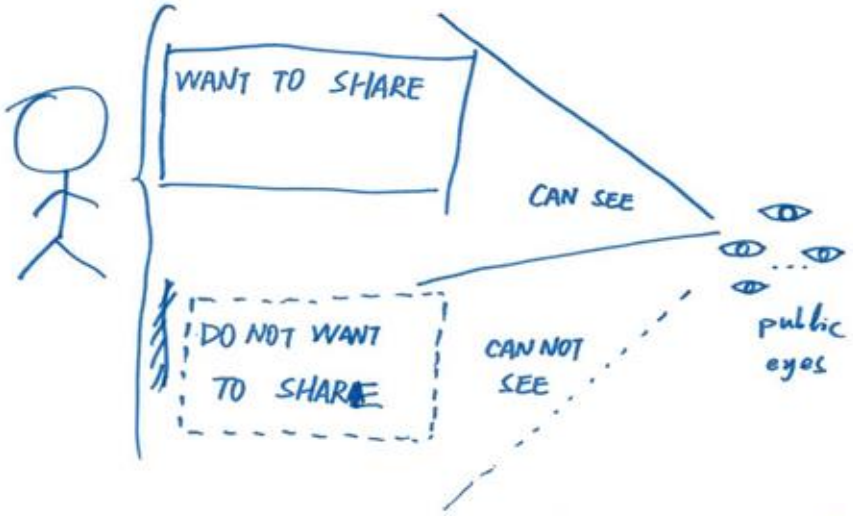

People have right to keep what they do not want to share invisible.

– AC, age 24

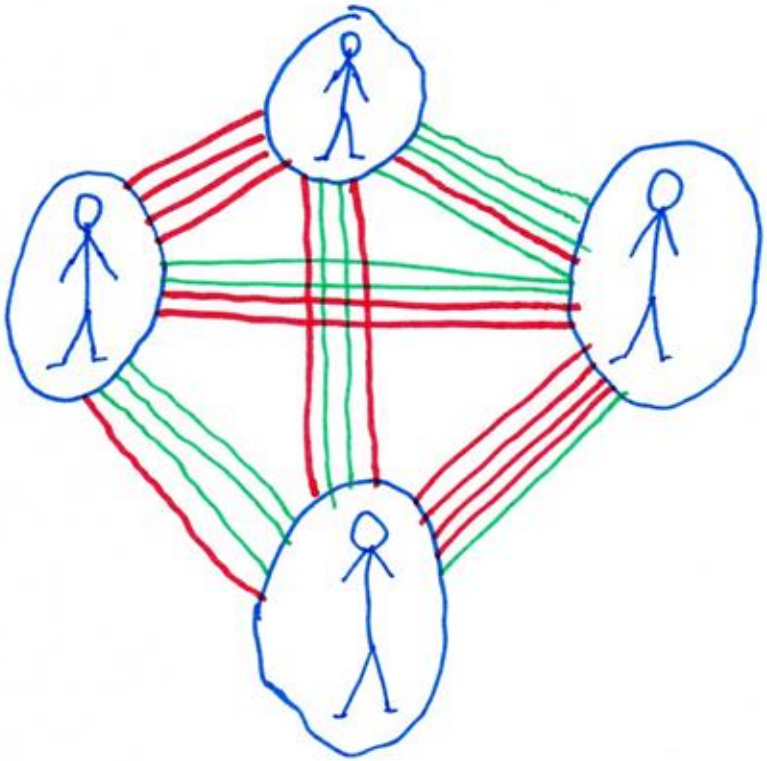

Privacy is a network: I share what I want with whom I want and trust and what matches with those in the network….

Green = share.
Red = don't.

Privacy and Freedom (1967)

Privacy is the claim of individuals, groups or institutions to determine for themselves when, how, and to what extent information about them is communicated to others.

– Alan Westin

Privacy Settings

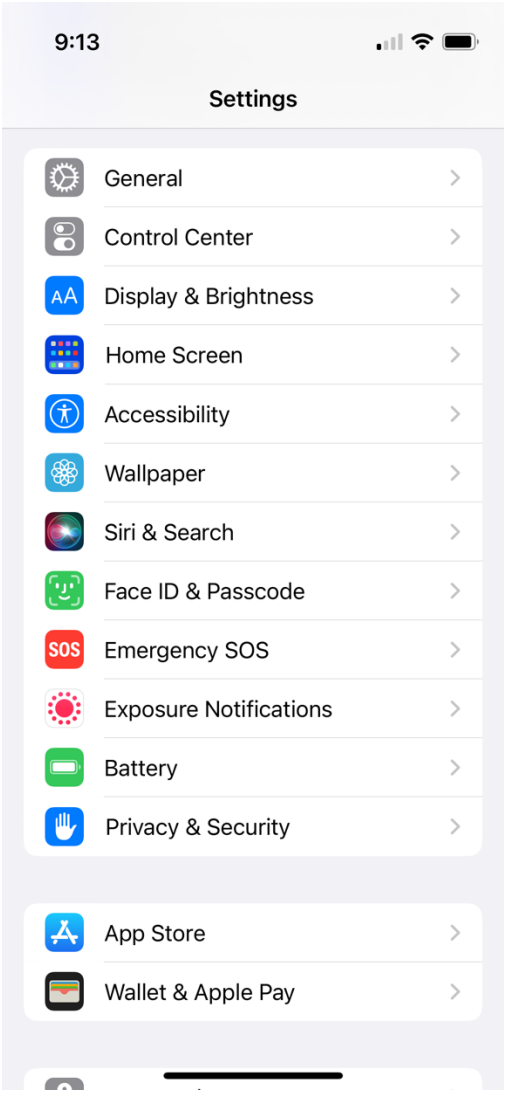

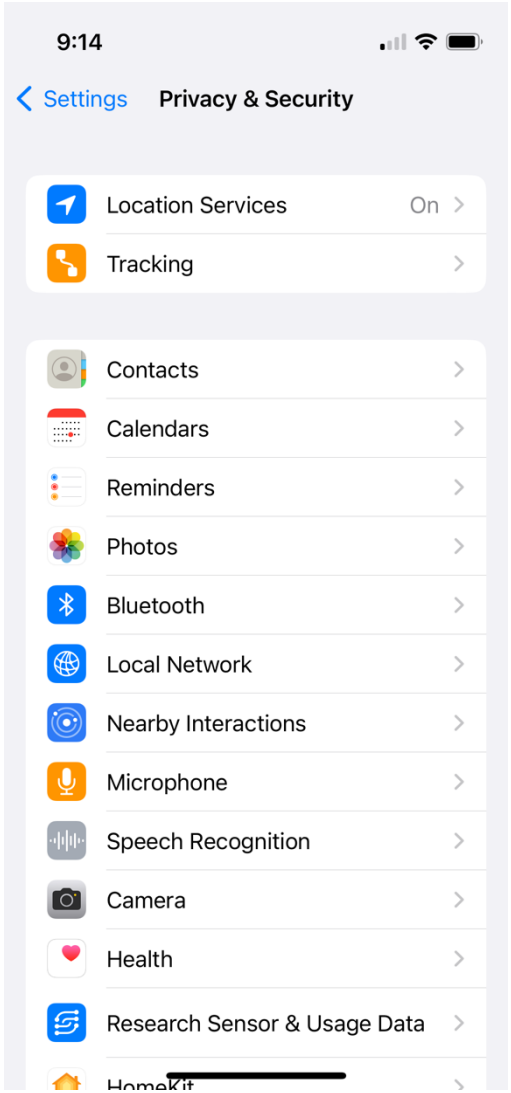

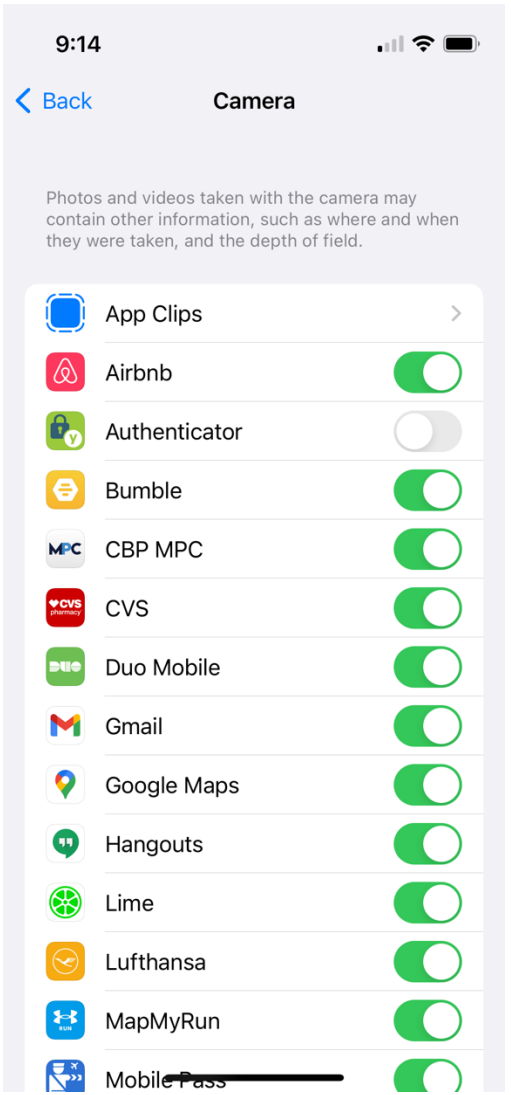

Opt-outs

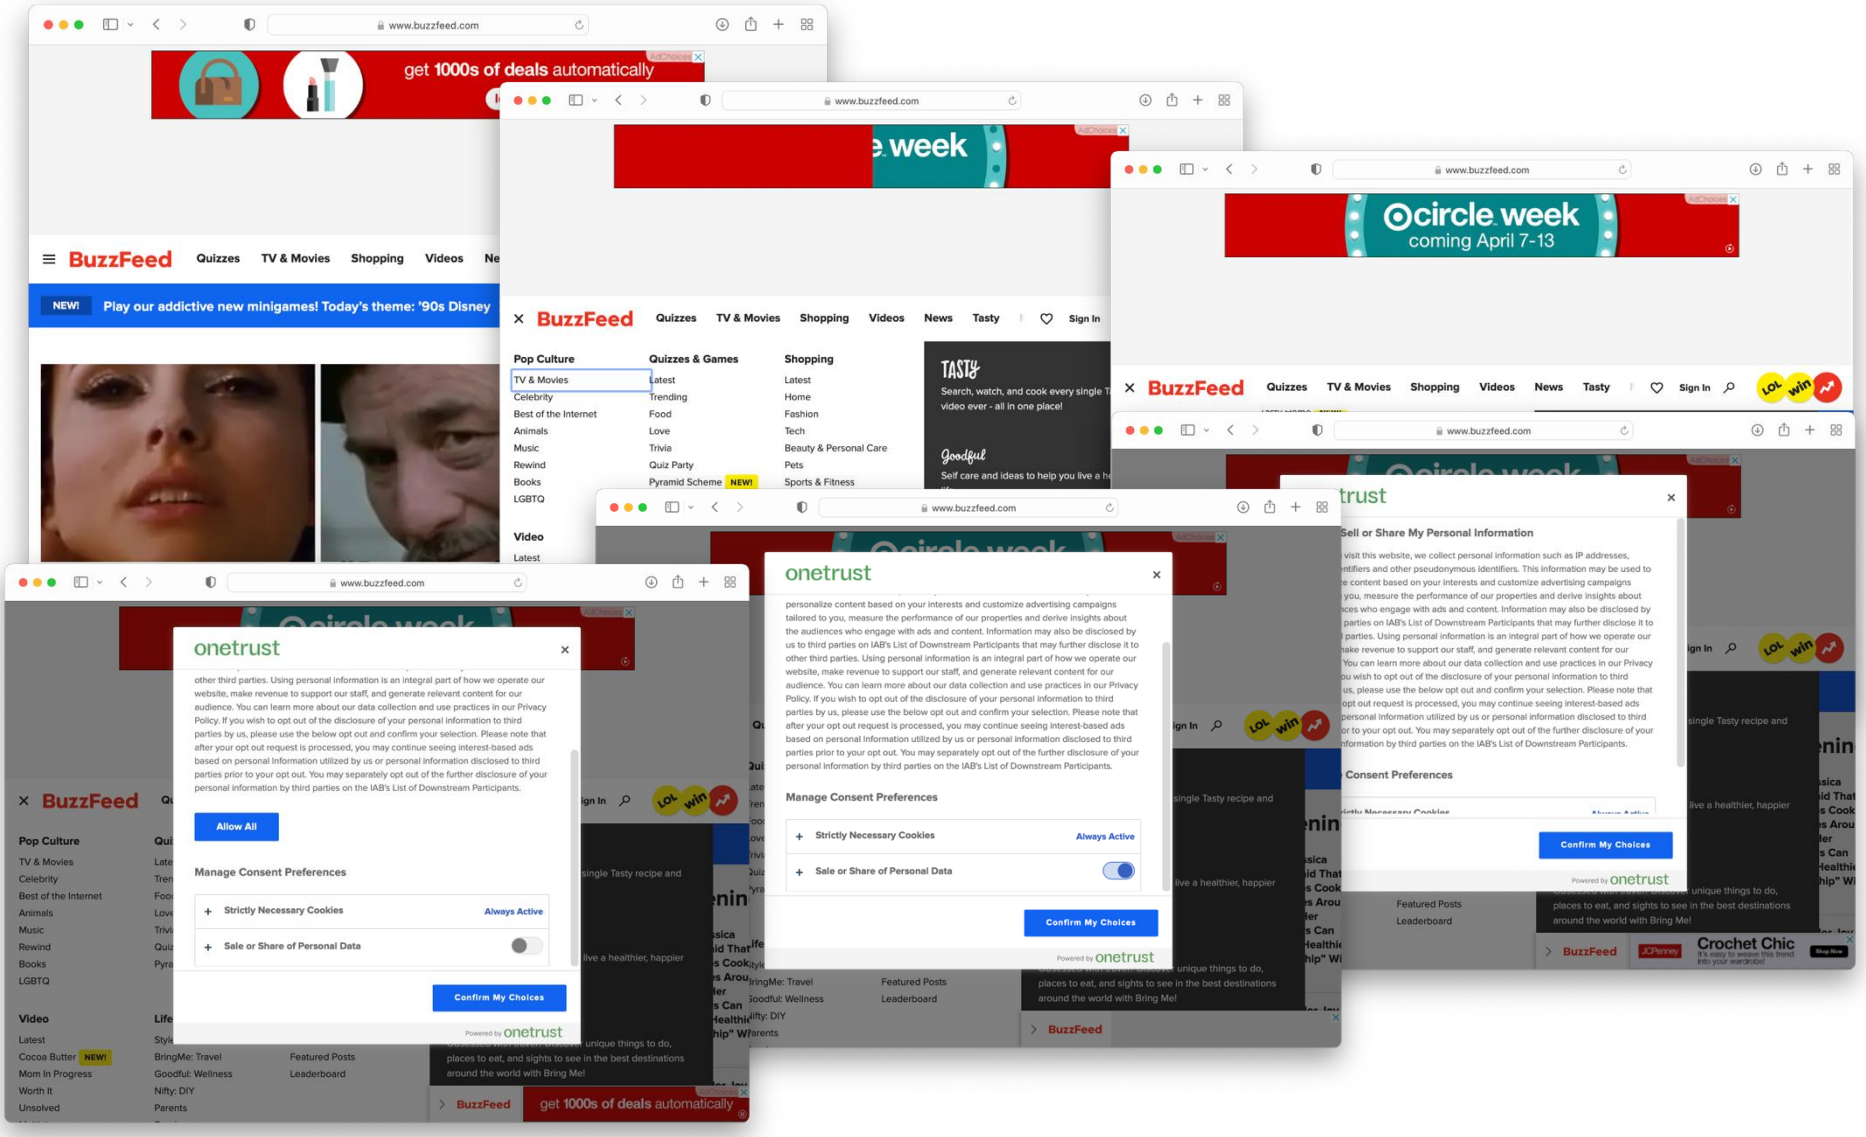
www.buzzfeed.com
get 1000s of deals automatically
BuzzFeed
Quizzes
TV & Movies
Shopping
Videos
NEW!
Play our addictive new minigames! Today's theme: '90s Disney
week
News
Tasty
Sign In
Pop Culture
TV & Movies
Celebrity
Best of the Internet
Animals
Music
Rewind
Books
LGBTQ
Quizzes & Games
Latest
Trending
Food
Love
Trivia
Quiz Party
Pyramid Scheme
NEW!
Shopping
Latest
Home
Fashion
Tech
Beauty & Personal Care
Pets
Sports & Fitness
TASTY
Search, watch, and cook every single
video ever - all in one place!
Goodful
Self care and ideas to help you live a
circle week
coming April 7-13
Video
Latest
onetrust
Allow All
Manage Consent Preferences
Strictly Necessary Cookies
Always Active
Sale or Share of Personal Data
Confirm My Choices
Powered by onetrust
Sell or Share My Personal Information
Consent Preferences
Life
BringMe: Travel
Goodful: Wellness
Nifty: DIY
Parents
Featured Posts
Leaderboard
Cocoa Butter
Mom In Progress
Worth It
Unsolved
places to eat, and sights to see in the best destinations
around the world with Bring Me!
Crochet Chic

Opt-outs

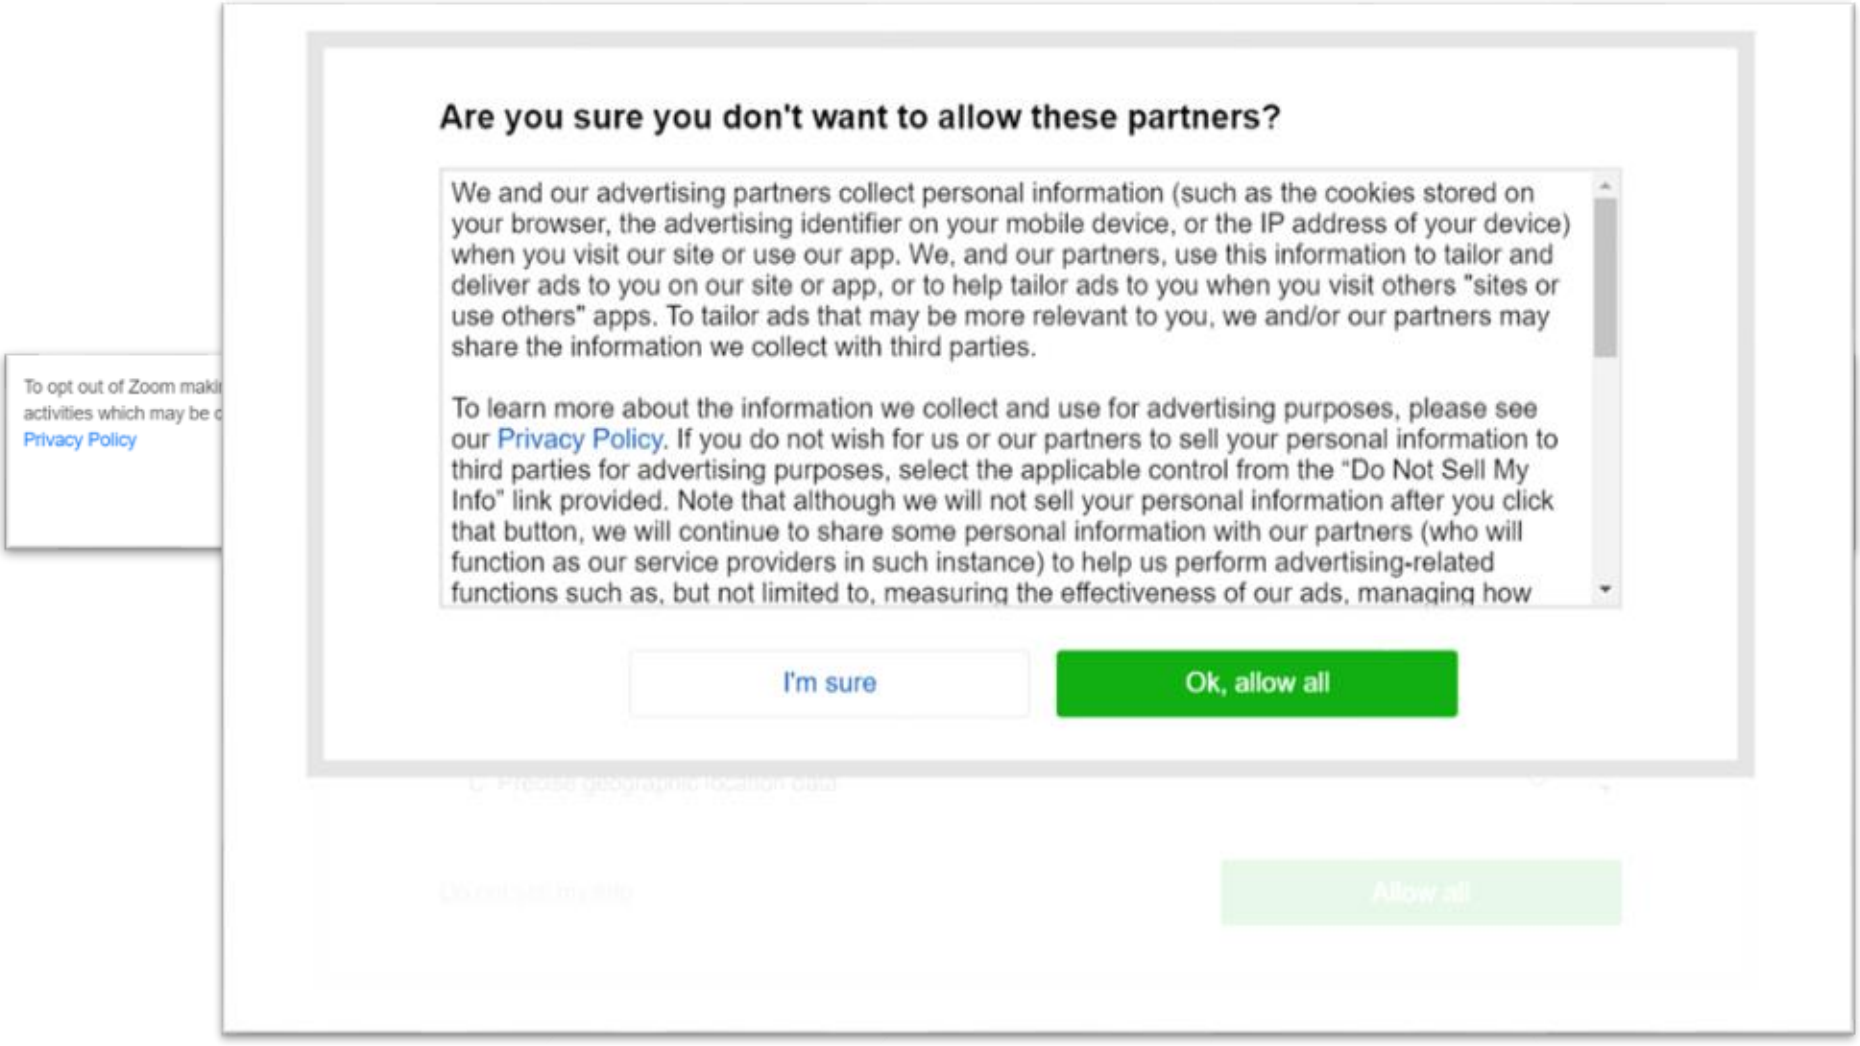
To opt out of Zoom maki
activities which may be c
Privacy Policy
Are you sure you don't want to allow these partners?
We and our advertising partners collect personal information (such as the cookies stored on your browser, the advertising identifier on your mobile device, or the IP address of your device) when you visit our site or use our app. We, and our partners, use this information to tailor and deliver ads to you on our site or app, or to help tailor ads to you when you visit others "sites or use others" apps. To tailor ads that may be more relevant to you, we and/or our partners may share the information we collect with third parties.
To learn more about the information we collect and use for advertising purposes, please see our Privacy Policy. If you do not wish for us or our partners to sell your personal information to third parties for advertising purposes, select the applicable control from the "Do Not Sell My Info" link provided. Note that although we will not sell your personal information after you click that button, we will continue to share some personal information with our partners (who will function as our service providers in such instance) to help us perform advertising-related functions such as, but not limited to, measuring the effectiveness of our ads, managing how
I'm sure
Ok, allow all

Your Choices Regarding Cookies on this Site

Ex Libris
a ProQuest Company

Please choose whether this site may use Functional and/or Advertising cookies, as described below:

REQUIRED COOKIES

These cookies are required to enable core site functionality.

FUNCTIONAL COOKIES

These cookies allow us to analyze site usage so we can measure and improve performance.

ADVERTISING COOKIES

These cookies are used by advertising companies to serve ads that are relevant to your interests.

Functionality Allowed

- Provide secure log-in
- Remember how far you are through an order
- Remember your log-in details
- Remember what is in your shopping cart
- Make sure the website looks consistent
- Allow you to share pages with social networks
- Allow you to post comments
- Serve ads relevant to your interests

This page transmits information using HTTPS protocol. Some vendors cannot support HTTPS opt-out requests. TrustArc will submit your preferences through HTTP in a pop-up window.

CANCEL | SUBMIT PREFERENCES | ADVANCED SETTINGS

Manipulation of privacy behavior

A few common examples of manipulative design in privacy-related interfaces

- Defaults are not privacy protective
- Buttons have confusing labels
- Framing - including wording
  that shames users to influence their decisions
  or makes them feel like they will be missing out
- Highlighting – visually emphasizing opt-in
- Cumbersome privacy choices - more
  difficult to choose privacy options

Your room is private.

– Alexia, age 11

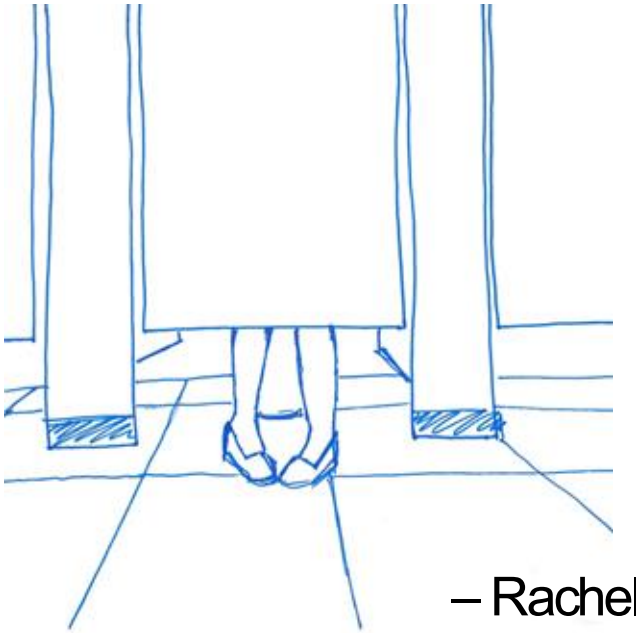

– Rachel, age 20

Contextual Integrity

- defines privacy relative to appropriate context
- considers information type, time, location, purpose, principals involved (subject, sender, receiver)
- dependent on social norms
- norms can change over time

General Guidelines

The FTC's Fair Information Practice Principals (FIPPs) are the most broadly recognized guidelines for handling private data in information systems

- Seek consent
- Minimize data use
- Limit storage
- Avoid linking

General Data Protection Regulation

Goal: Codify fundamental right to protection of personal data.

- Introduced individual rights
  1. The right to transparency
  2. The right to access
  3. The right to correct
  4. The right to delete
  5. The right to data portability
  6. The right to withdraw consent
  7. The right to object
- Additional obligations
  - Legal basis for processing
  - Purpose limitation
  - Data Minimization
  - Storage limitation
  - Security requirements
  - Privacy by design

- Adopted: April 14, 2016
- Effective: May 25, 2018

Anonymity

Deanonymization

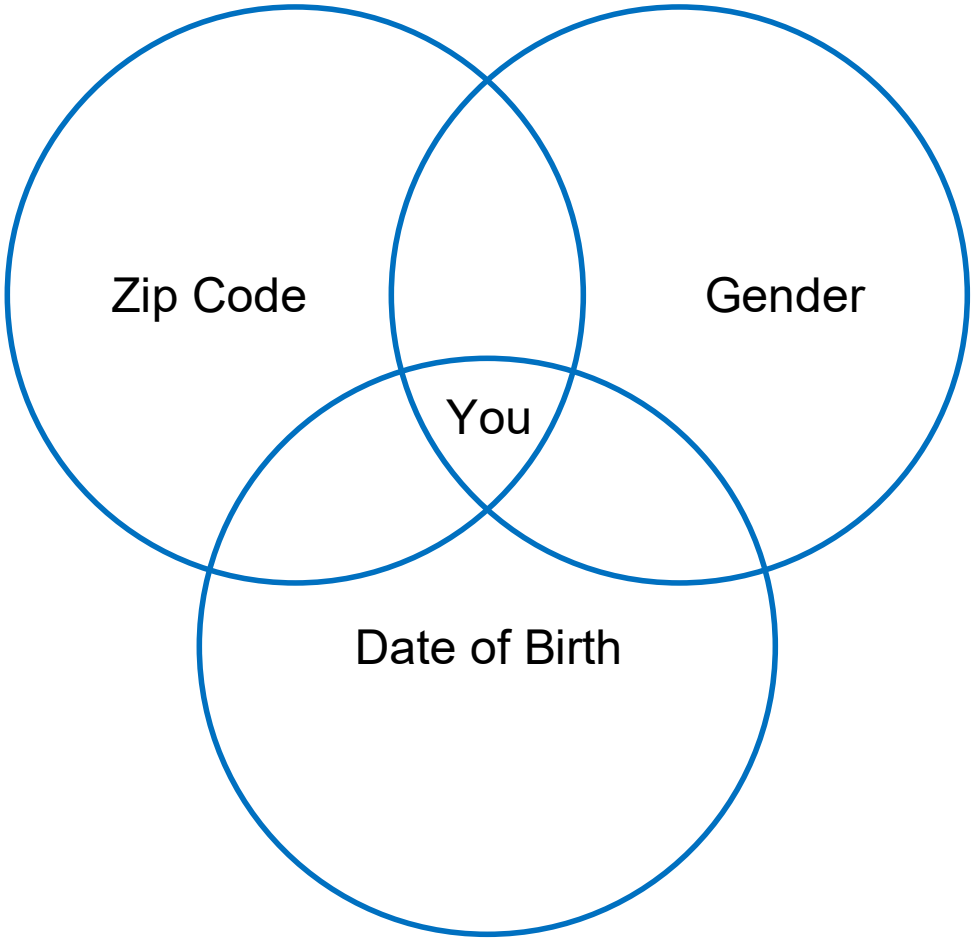

Deanonymization

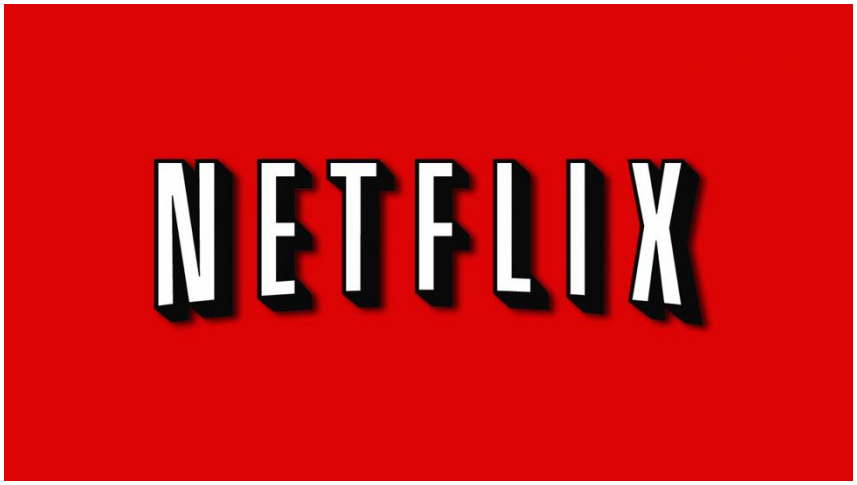

k-Anonymity

Name	Pronouns	Year	Grade
Alice	she/her	2026	95
Bob	he/him	2026	80
Charlie	they/them	2026	95
David	he/him	2026	60
Edward	he/him	2027	80
Flora	she/her	2027	99
Georgia	she/her	2027	60

- **Quasi-identifiers (QIs)** are sets of attributes that can be exploited for linking
- A database is **k-anonymous** if each QI maps to at least k different individuals
- Techniques: suppression and generalization

Exercise 1: k-anonymity

- Modify this dataset to make it 2-anonymous with respect to Zipcode/DOB/Sex

Zipcode	DOB	Gender	Marital Status	Health Issues
91711	9/27/00	female	divorced	hypertension
91711	9/30/00	female	divorced	obesity
91711	4/18/00	male	married	chest pain
91711	4/15/00	male	married	obesity
91767	3/13/99	male	married	hypertension
91767	3/18/99	male	married	shortness of breath
91767	9/13/00	female	married	shortness of breath
91767	9/07/00	female	married	obesity
91101	5/14/01	male	single	chest pain
91101	4/08/01	male	single	obesity
91101	9/15/01	female	married	shortness of breath

Database Privacy

Offline Privacy

Online Privacy

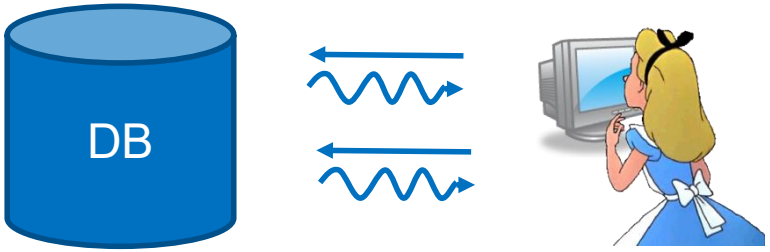

Defining Privacy: Try #1

- You don't know anything more after interacting with the database then you already knew

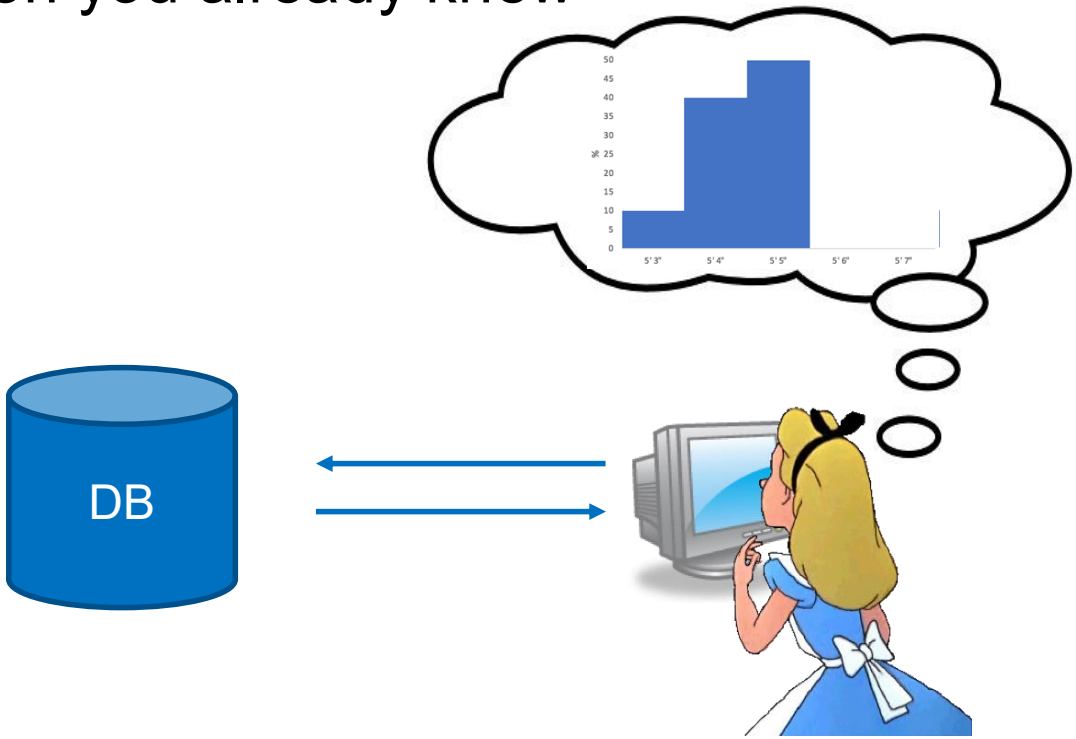

Differential Privacy

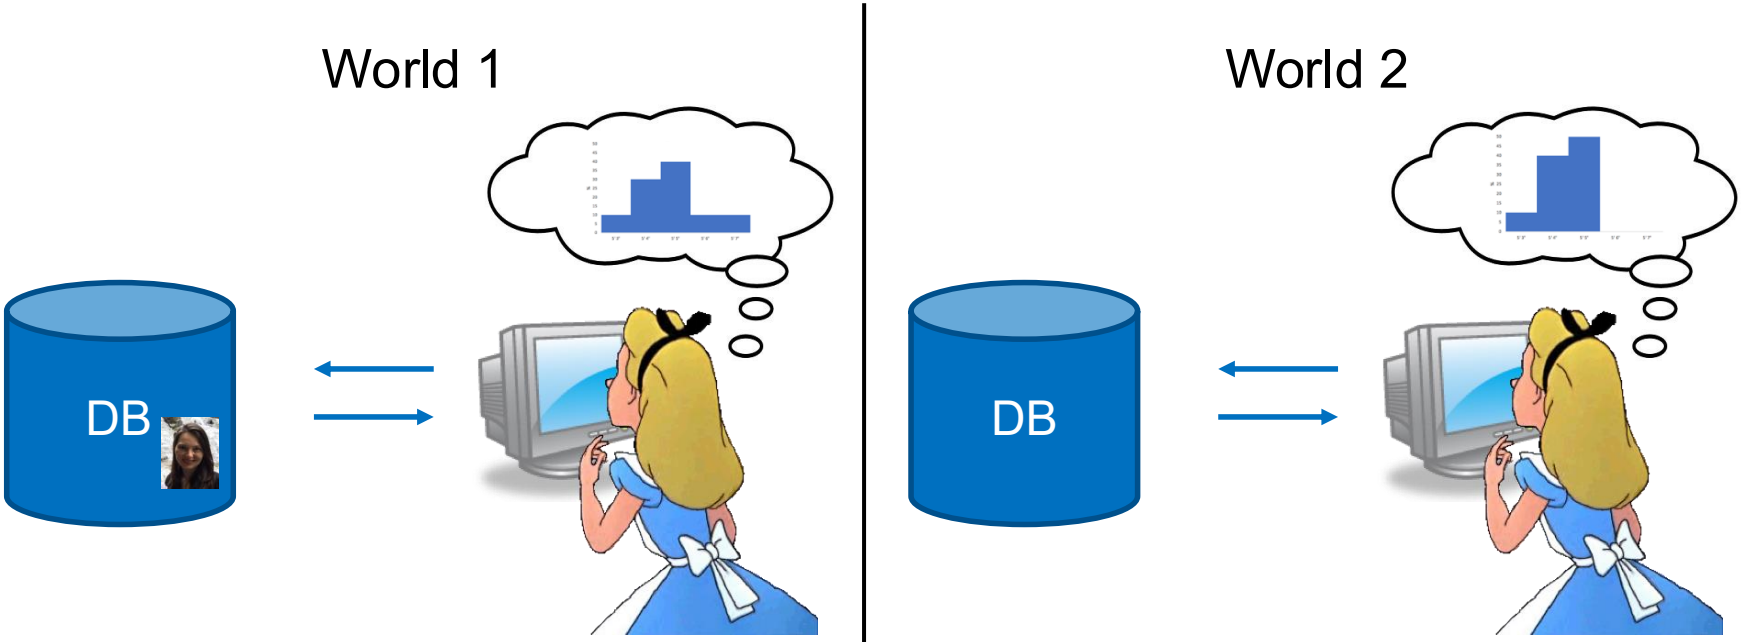

A query Q is ϵ-differentially private if $\forall D, r \in D$,
$\Pr[Q(D) = x] \leq e^{\epsilon} \cdot \Pr[Q(D - r) = x]$

Sensitivity

- The sensitivity Δ of a query Q is the maximum the answer to Q can possibly change between two databases that differ only by one person

- Q = number of people taller than 6 ft $\Delta = 1$
- Q = maximum height of a person $\Delta = 48$

Exercise 2: Sensitivity

- Assume you have a database containing the heights of 100 users specified in inches. You may assume that all heights are between 48 in and 96in.

- What is the sensitivity of the following queries?
  1. The number of people who are 5' 4"
  2. The median height in the dataset
  3. The mean height in the dataset

Laplacian Distribution

- Lap(b) is the probability distribution with the property that

$$\Pr[\, Lap(b) = x] = \frac{1}{2b} \cdot e^{-\frac{|x|}{b}}$$

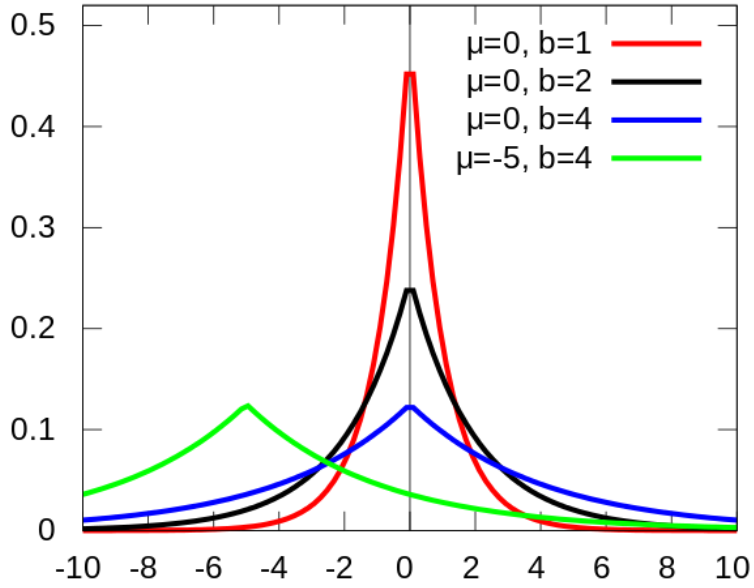

Laplacian Mechanism

- Given a query Q on a database D that has sensitivity Δ, respond with Q(D)+Y where Y is drawn from the distribution $Lap(\frac{\Delta}{\epsilon})$
- Theorem: this mechanism satisfies ϵ-differential privacy

$$\frac{\Pr[\,Q(D)\,+\,Y\,=\,x]}{\Pr[Q(D-r)\,+\,Y\,=\,x]} = \frac{\Pr[\,Y\,=\,x\,-\,Q(D)]}{\Pr[Y\,=\,x\,-\,Q(D-r)]} = \frac{\frac{1}{2(\Delta/\epsilon)}\cdot e^{-\frac{|x-Q(D)|}{\Delta/\epsilon}}}{\frac{1}{2(\Delta/\epsilon)}\cdot e^{-\frac{|x-Q(D-r)|}{\Delta/\epsilon}}} = \frac{e^{-\frac{|x-Q(D)|}{\Delta/\epsilon}}}{e^{-\frac{|x-Q(D-r)|}{\Delta/\epsilon}}}$$

$$= e^{\frac{|x-Q(D-r)|}{\Delta/\epsilon}-\frac{|x-Q(D)|}{\Delta/\epsilon}} = e^{\left(\frac{\epsilon}{\Delta}\right)\cdot(|x-Q(D-r)|-|x-Q(D)|)}$$

$$\leq e^{\left(\frac{\epsilon}{\Delta}\right)\cdot(|x-Q(D-r)-x+Q(D)|)} = e^{\left(\frac{\epsilon}{\Delta}\right)\cdot(|Q(D)-Q(D-r)|)}$$

$$\leq e^{\left(\frac{\epsilon}{\Delta}\right)\cdot\Delta} = \mathrm{e}^{\epsilon}$$

Randomized Response

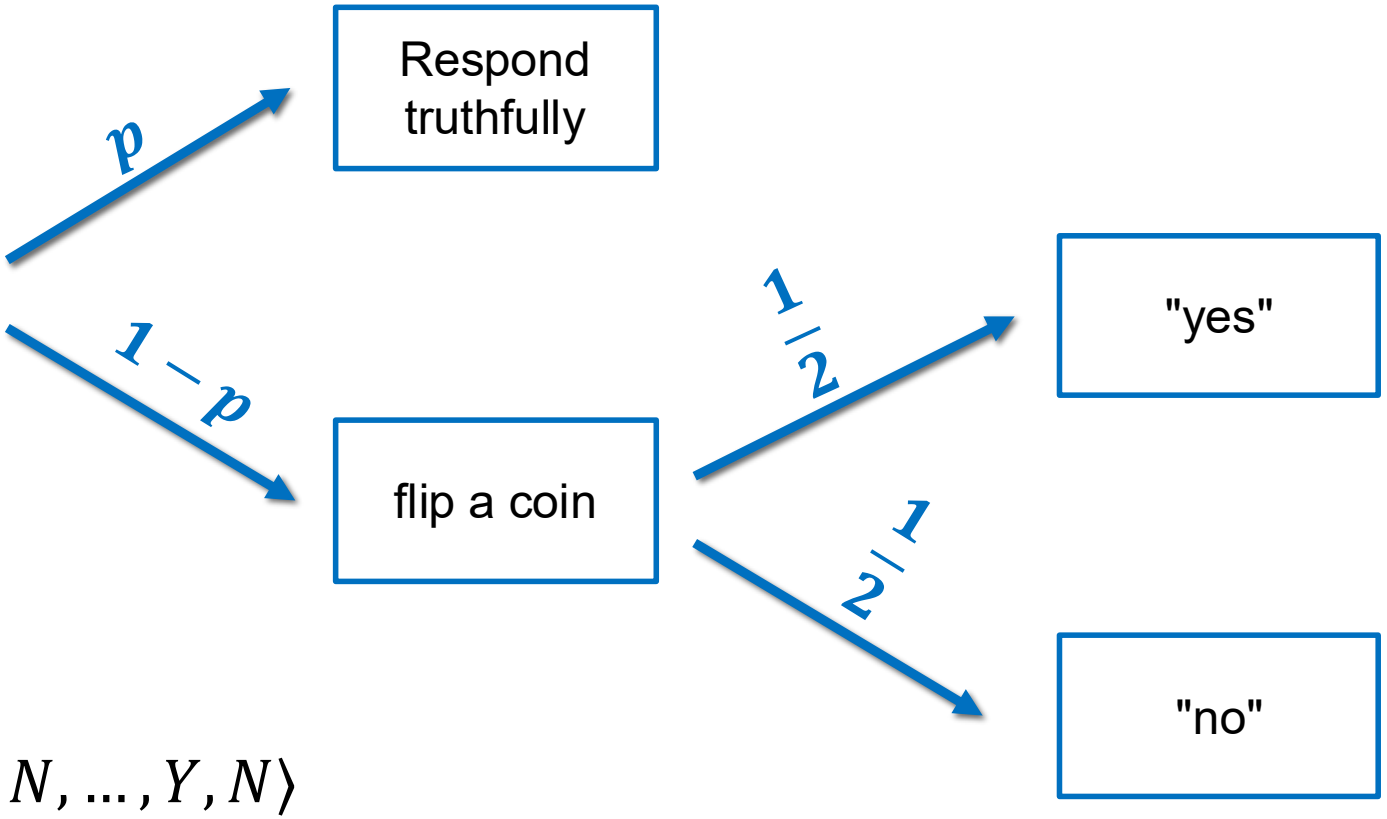

$\langle Y, N, \boldsymbol{Y}, Y, Y, N, \ldots, Y, N \rangle$

Theorem: this mechanism satisfies ϵ-differential privacy

Randomized Response

- Theorem: this mechanism satisfies ϵ-differential privacy

$$\frac{\Pr[\langle Y, N, \boldsymbol{Y}, Y, Y, N, \ldots, Y, N\rangle \mid f(\text{Bob}) = Y]}{\Pr[\langle Y, N, \boldsymbol{Y}, Y, Y, N, \ldots, Y, N\rangle \mid f(\text{Bob}) = N]}$$

$$= \frac{\Pr[Y \mid f(P_0)] \cdot \Pr[\text{N} \mid f(P_1)] \cdot \Pr[Y \mid f(Bob) = Y] \cdot \ldots \cdot \Pr[\text{N} \mid f(P_{n-1})]}{\Pr[Y \mid f(P_0)] \cdot \Pr[\text{N} \mid f(P_1)] \cdot \Pr[Y \mid f(Bob) = N] \cdot \ldots \cdot \Pr[\text{N} \mid f(P_{n-1})]}$$

$$= \frac{\Pr[Y \mid f(Bob) = Y]}{\Pr[Y \mid f(Bob) = N]}$$

$$= \frac{p \cdot 1 + (1-p) \cdot \frac{1}{2}}{p \cdot 0 + (1-p) \cdot \frac{1}{2}} = \frac{(1+p) \cdot \frac{1}{2}}{(1-p) \cdot \frac{1}{2}} = \frac{(1+p)}{(1-p)}$$

$$= \mathrm{e}^{\ln(\frac{1+p}{1-p})}$$

DP in action…

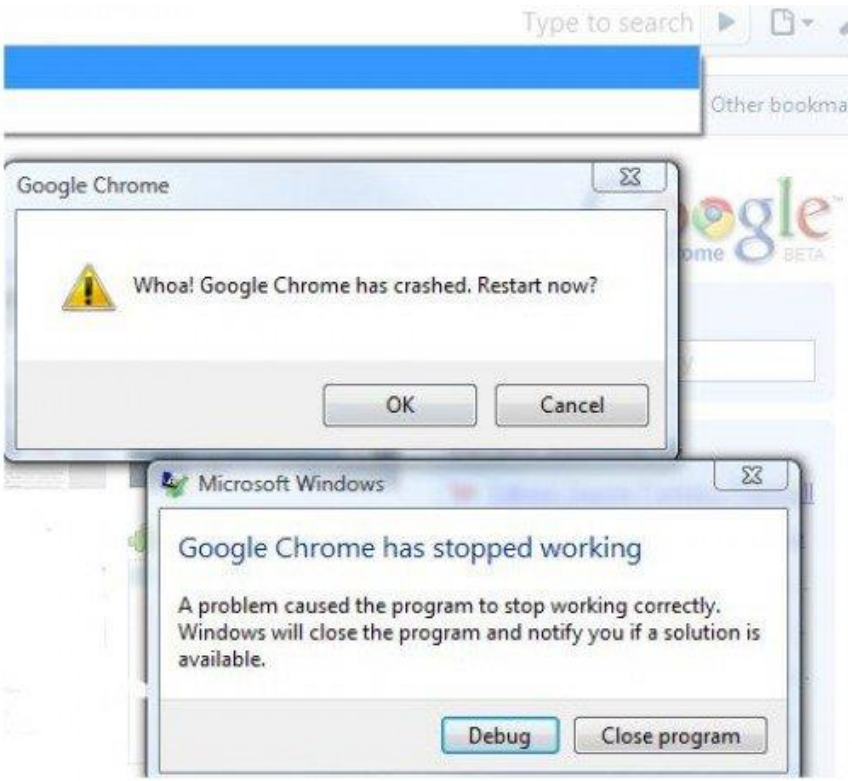

Linked in

United States™ Census Bureau

Differential Privacy

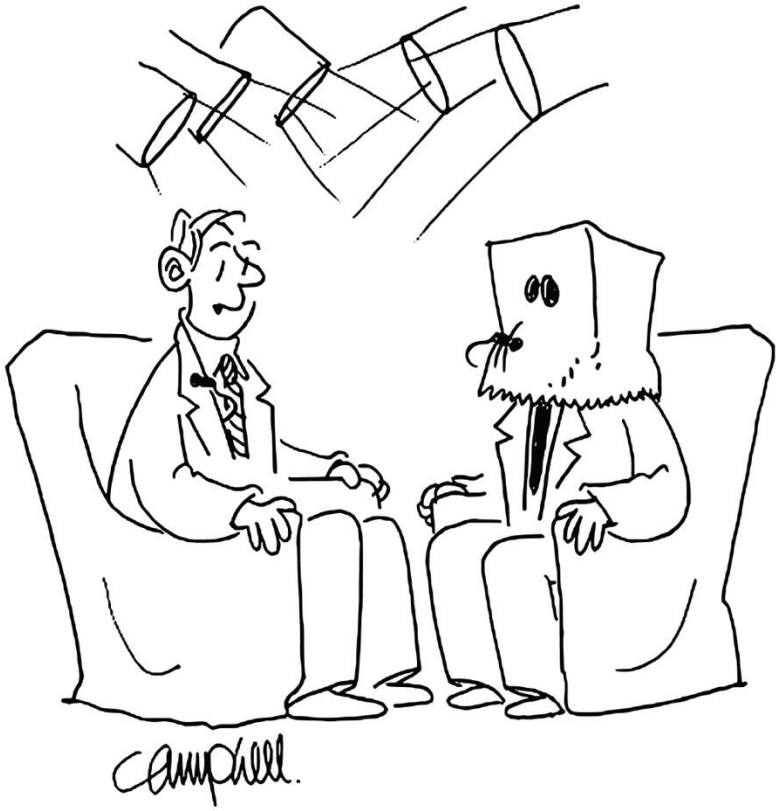

"Our guest has asked that we obscure his identity."