Data Mining 2017: Privacy and Differential Privacy
Slides stolen from Kobbi Nissim (with permission)

Also used: The Algorithmic Foundations of Differential Privacy by Dwork and Roth

Data Privacy – The Problem

- Given:
  - a dataset with sensitive information
    - Health records, census data, financial data, …
- How to:
  - Compute and release functions of the dataset,
  - Without compromising individual privacy.

Data Privacy – The Problem

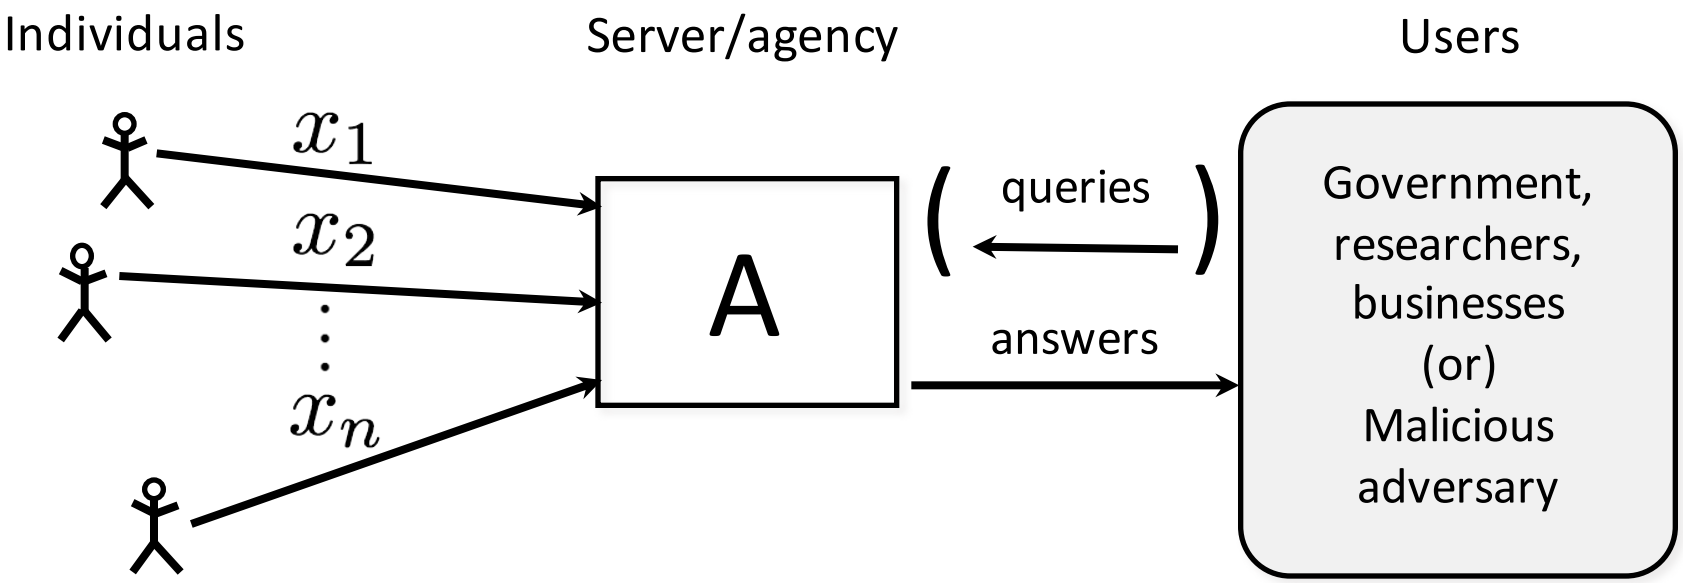

A Real Problem

Typical examples:

- Census
- Civic archives
- Medical records
- Search information
- Communication logs
- Social networks
- Genetic databases
- …

Benefits:

- New discoveries
- Improved medical care
- National security

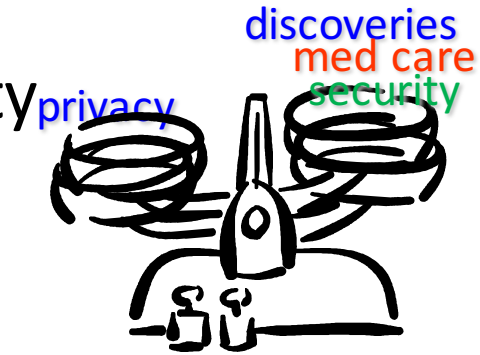

The Anonymization Dream

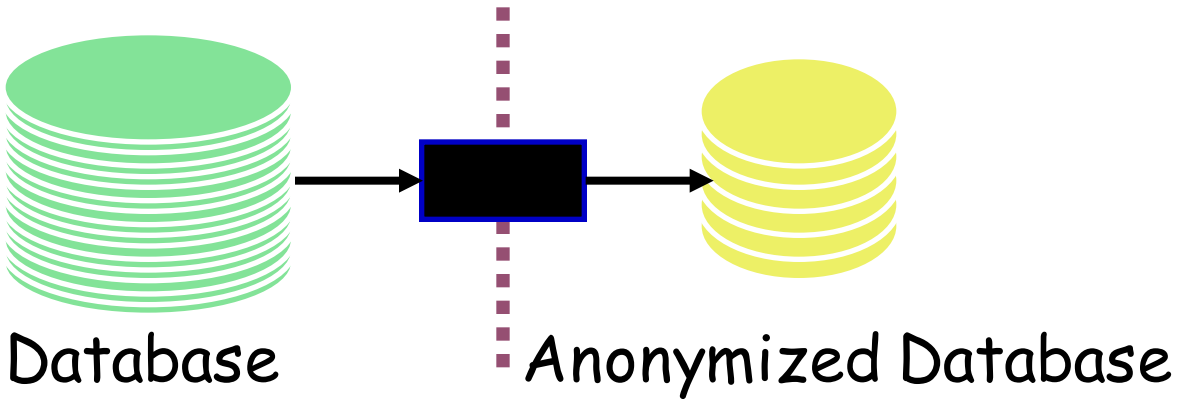

- Trusted curator:
  - Removes identifying information (name, address, ssn, …).
  - Replaces identities with random identifiers.
- Idea hard wired into practices, regulations, …, thought.
  - Many uses.
- Reality: series failures.
  - Pronounced both in academic and public literature.

Data Privacy

- Studied (at least) from the 60s
- Approaches: De-identification, redaction, auditing, noise addition, synthetic datasets …
  - Focus on *how* to provide privacy, not on *what* privacy protection is
  - May have been suitable for the pre-internet era
- Re-identification [Sweeney '00, …]
  - GIS data, health data, clinical trial data, DNA, health privacy data, text data, registry information, …
- Blatant non-privacy [Dinur, Nissim '03], …
- Auditors [Kenthapadi, Mishra, Nissim '05]
- AOL Debacle '06
- Genome-wide association studies (GWAS) [Homer et al. '08]
- Netflix award [Narayanan, Shmatikov '09]
  - Netflix canceled second contest
- Social networks [Backstrom, Dwork, Kleinberg '11]
- Genetic research studies [Gymrek, McGuire, Golan, Halperin, Erlich '13]
- Microtargeted advertising [Korolova 11]
- Recommendation Systems [Calandrino, Kilzer, Narayanan, Felten, Shmatikov 11]
- Israeli CBS [Mukatren, Nissim, Salman, Tromer '11]
- Attack on statistical aggregates [Homer et al. '08] [Dwork, Smith, Steinke, Vadhan '15]
- …

Slide idea stolen shamelessly from Or Sheffet

Linkage Attacks [Sweeney 2000]

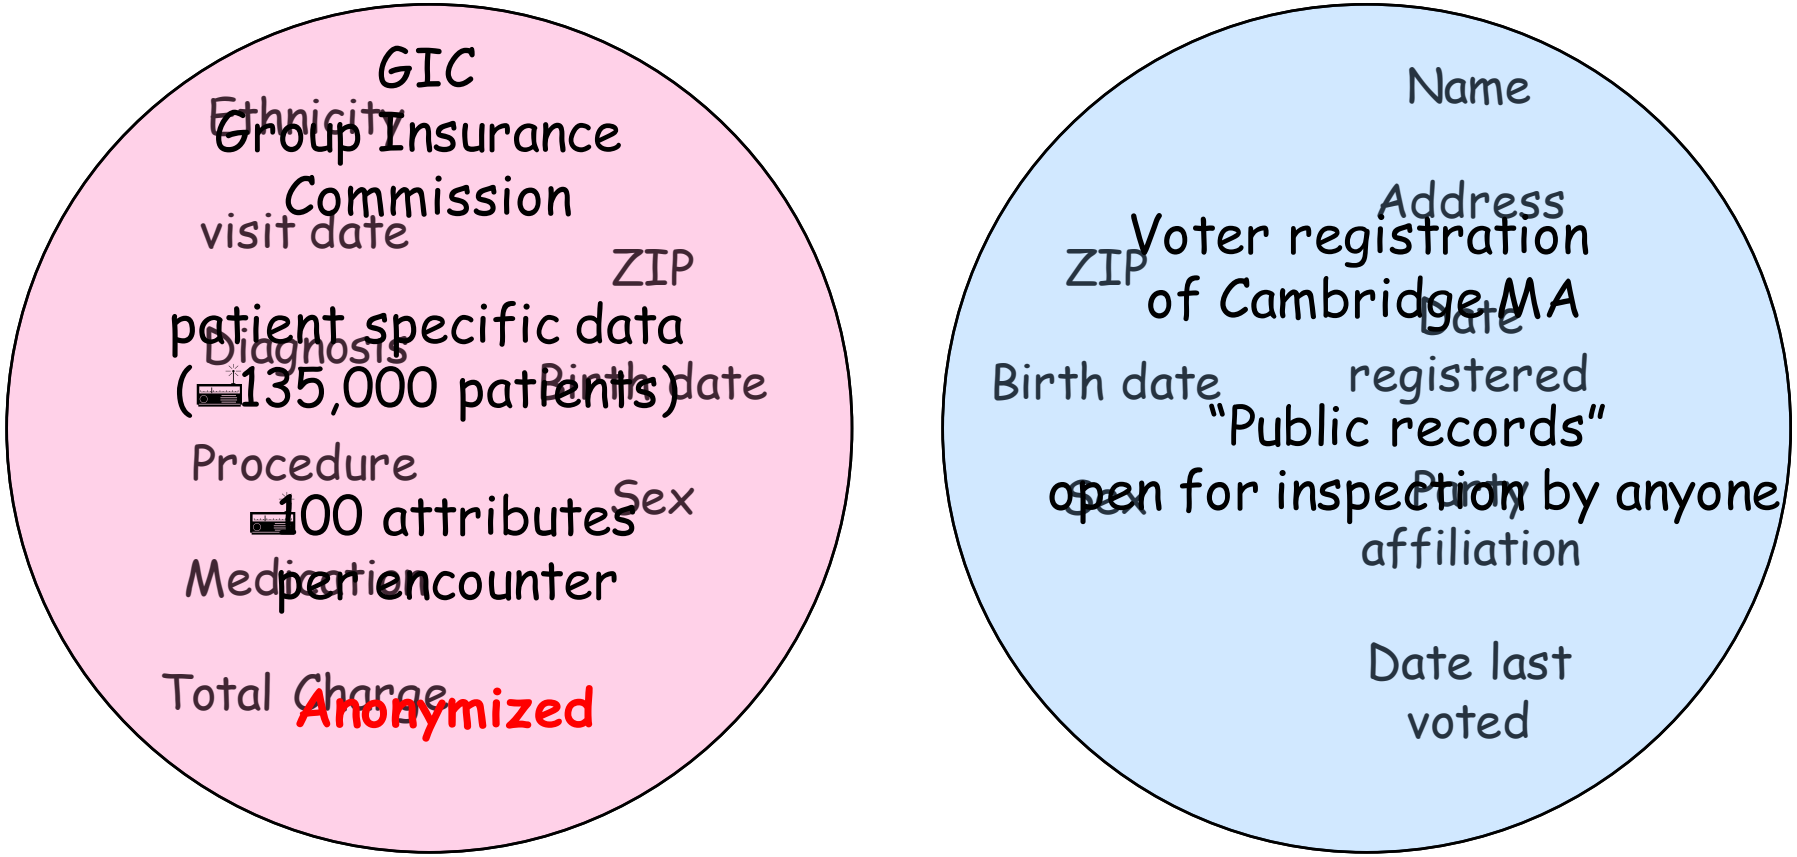

Anonymized GIC data

Voter registration

Linkage Attacks [Sweeney 2000]

- Quasi identifiers → re-identification
  - Not a coincidence:
    - dob+5zip → 69%
    - dob+9zip → 97%

- William Weld (governor of Massachusetts at the time)
  - According to the Cambridge Voter list:
    - Six people had his particular birth date
    - Of which three were men
    - He was the only one in his 5-digit ZIP code!

Search the Web - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Search Favorites

Address http://search.aol.com/aolcom/webhome Go Links

Google Search 3 blocked Check AutoLink AutoFill Options

AOL Main | Mail | AIM | Join AOL

AOL & AIM Sign In

make this my homepage

saved searches

AOL search

web pictures video audio news local shopping

privacy search

did you know:
you can search local businesses and information by clicking on 'local'

browse & discover: AOL Search Features CityGuide Government Guide MapQuest Movies
Pinpoint™ Shopping Mobile Search White Pages Yellow Pages AOL.com Pinpoint™ Travel Singingfish
Link to us Legal Web Directory Search Tools Advanced Search Help

AOL International AOL, Inc. Download AOL Accessibility Policy GetNetWise Terms of Use
Revised Member Privacy Policy Trademarks AOL Unsolicited Bulk E-mail Policy
Advertise With Us Feedback Help

Internet

start IPAM tut... Search t... cseg123 ... Microsof... untitled -... EN 17:07

AOL Data release (2006)

- AOL released search data
  - A sample of ~20M web queries collected from ~650k users over three months
- Goal: provide real query log data that is based on real users
  - *"It could be used for personalization, query reformulation or other types of search research"*
- The data set:

AnonID	Query	QueryTime	ItemRank	ClickURL

4417749	best dog for older owner	3/6/2006	11:48:24	1	http://www.canismajor.com
4417749	best dog for older owner	3/6/2006	11:48:24	5	http://dogs.about.com
4417749	landscapers in lilburn ga.	3/6/2006	18:37:26		
4417749	effects of nicotine	3/7/2006	19:17:19	6	http://www.nida.nih.gov
4417749	best retirement in the world	3/9/2006	21:47:26	4	http://www.escapeartist.com
4417749	best retirement place in usa	3/9/2006	21:49:37	10	http://www.clubmarena.com
4417749	best retirement place in usa	3/9/2006	21:49:37	9	http://www.committment.com
4417749	bi polar and heredity	3/13/2006	20:57:11		
4417749	adventure for the older american	3/17/2006	21:35:48		
4417749	nicotine effects on the body	3/26/2006	10:31:15	3	http://www.geocities.com
4417749	nicotine effects on the body	3/26/2006	10:31:15	2	http://health.howstuffworks.com
4417749	wrinkling of the skin	3/26/2006	10:38:23		
4417749	mini strokes	3/26/2006	14:56:56	1	http://www.ninds.nih.gov
4417749	panic disorders	3/26/2006	14:58:25		
4417749	jarrett t. arnold eugene oregon	3/23/2006	21:48:01	2	http://www2.eugeneweekly.com
4417749	jarrett t. arnold eugene oregon	3/23/2006	21:48:01	3	http://www2.eugeneweekly.com
4417749	plastic surgeons in gwinnett county	3/28/2006	15:04:23	1	http://www.wedalert.com
4417749	plastic surgeons in gwinnett county	3/28/2006	15:04:23	4	http://www.implantinfo.com
4417749	plastic surgeons in gwinnett county	3/28/2006	15:31:00		
441774	9 60 single men	3/29/2006	20:11:52	6	http://www.adultlovecompass.com
441774	9 60 single men	3/29/2006	20:14:14		
4417749	clothes for 60 plus age	4/19/2006	12:44:03		
4417749	clothes for age 60	4/19/2006	12:44:41	10	http://www.news.cornell.edu
4417749	clothes for age 60	4/19/2006	12:45:41		
4417749	lactose intolerant	4/21/2006	20:53:51	2	http://digestive.niddk.nih.gov
4417749	lactose intolerant	4/21/2006	20:53:51	10	http://www.netdoctor.co.uk
4417749	dog who urinate on everything	4/28/2006	13:24:07	6	http://www.dogdaysusa.com
4417749	fingers going numb	5/2/2006	17:35:47		

A Face Is Exposed for AOL Searcher No. 4417749

By MICHAEL BARBARO and TOM ZELLER Jr.
Published: August 9, 2006

The New York Times

Buried in a list of 20 million Web search queries collected by AOL and recently released on the Internet is user No. 4417749. The number was assigned by the company to protect the searcher's anonymity, but it was not much of a shield.

No. 4417749 conducted hundreds of searches over a three-month period on topics ranging from "numb fingers" to "60 single men" to "dog that urinates on

Name: Thelma Arnold
Age: 62
Widow
Residence: Lilburn, GA

Other Re-Identification Examples

[partial and unordered list]

- Netflix award [Narayanan, Shmatikov 08].
- Social networks [Backstrom, Dwork, Kleinberg 07, NS 09].
- Computer networks [Coull, Wright, Monrose, Collins, Reiter ’07, Ribeiro, Chen, Miklau, Townsley 08].
- Genetic data (GWAS) [Homer, Szelinger, Redman, Duggan, Tembe, Muehling, Pearson, Stephan, Nelson, Craig 08, ...].
- Microtargeted advertising [Korolova 11].
- Recommendation Systems [Calandrino, Kiltzer, Naryanan, Felten, Shmatikov 11].
- Israeli CBS [Mukatren, N, Salman, Tromer].
- …

k-Anonymity [SS98,S02]

- Prevent re-identification:
  - Make every individual's identity unidentifiable from other k-1 individuals

ZIP	Age	sex	Disease
23456	55	Female	Heart
12345	30	Male	Heart
12346	33	Male	Heart
13144	45	Female	Breast Cancer
13155	42	Male	Hepatitis
23456	42	Male	Viral

ZIP	Age	sex	Disease
23456	**	*	Heart
1234*	3*	Male	Heart
1234*	3*	Male	Heart
131**	4*	*	Breast Cancer
131**	4*	*	Hepatitis
23456	**	*	Viral

Both guys from zip 1234* that are in their thirties have heart problems

My (male) neighbor from zip 13155 has hepatitis!

Bugger! I Cannot tell which disease for the patients from zip 23456

Auditing

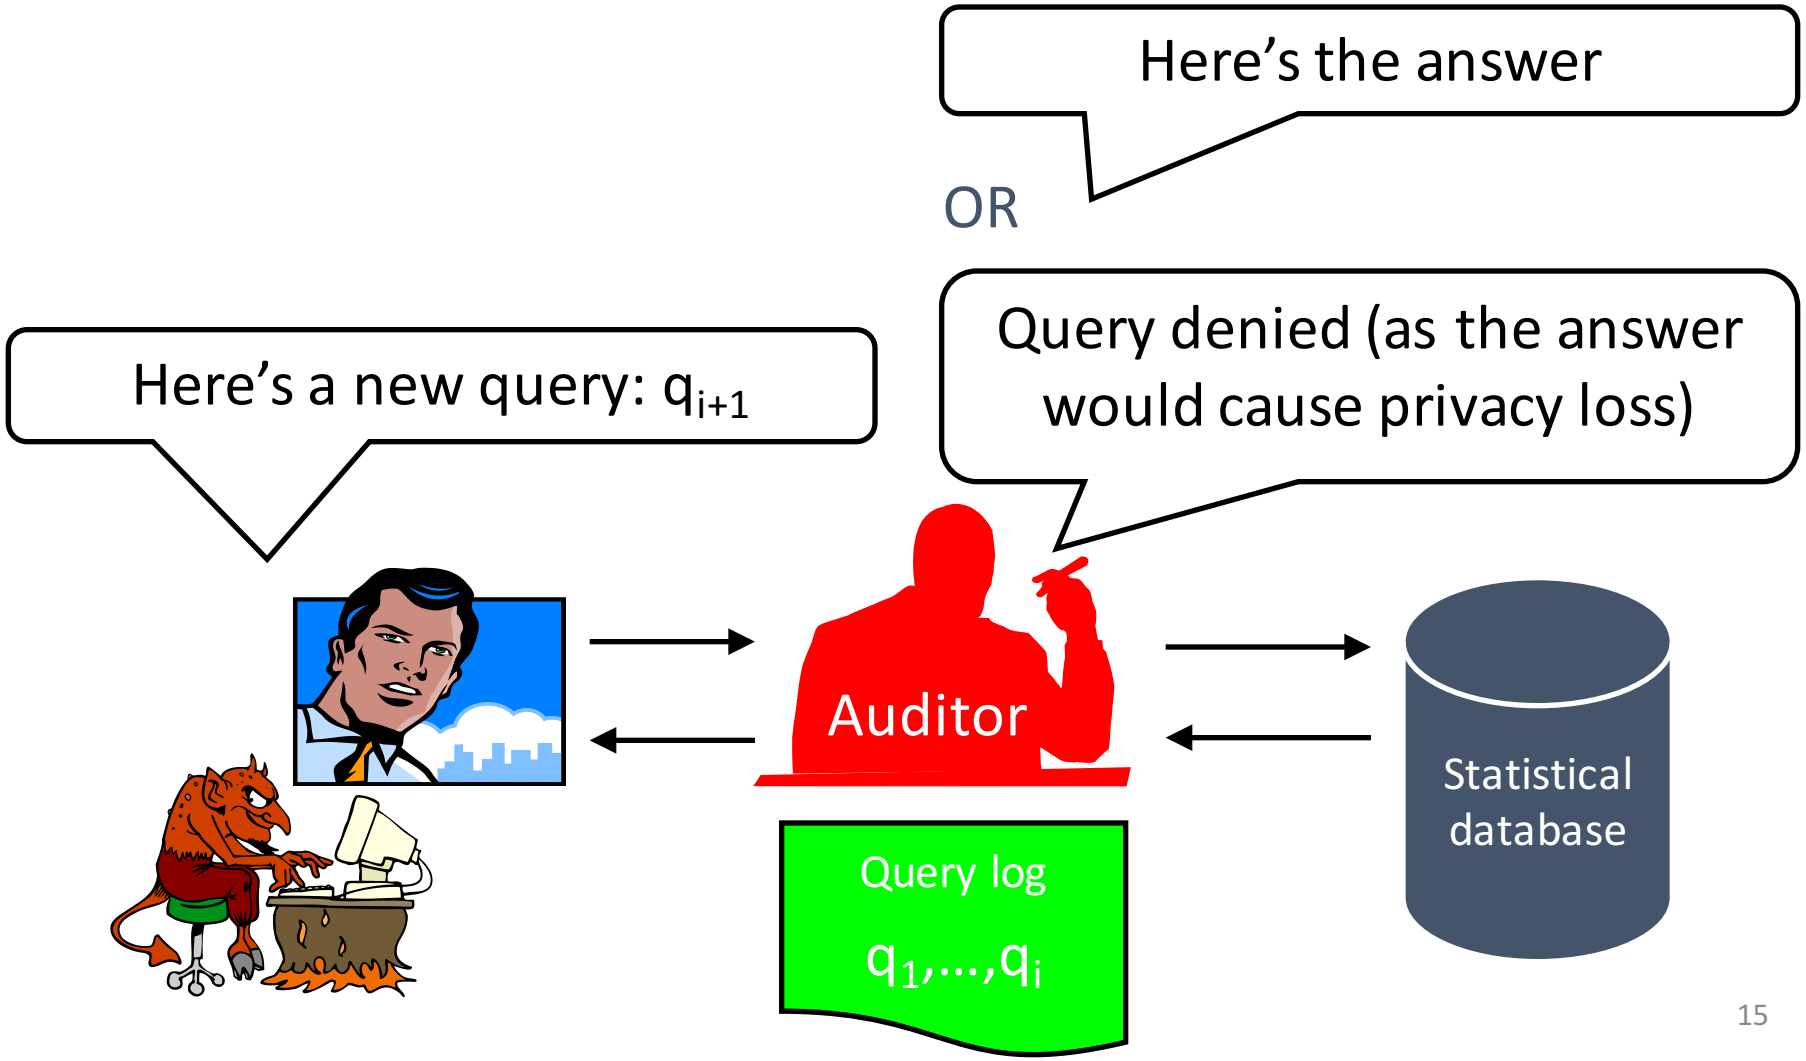
Here's the answer
OR
Query denied (as the answer would cause privacy loss)
Here's a new query: q_{i+1}
Auditor
Statistical database
Query log
$q_1,\ldots,q_i$

Example 1: Sum/Max auditing

d_i real, sum/max queries, privacy breached if some d_i learned

q1 = sum(d1,d2,d3)

sum(d1,d2,d3) = 15

q2 = max(d1,d2,d3)

Denied (the answer would cause privacy loss)

Oh well…

… After Two Minutes …

d_i real, sum/max queries, privacy breached if some d_i learned

q1 = sum(d1,d2,d3)

sum(d1,d2,d3) = 15

q2 = max(d1,d2,d3)

Denied (the answer would cause privacy loss)

Oh well…

There must here reason for the denial…

q2 is denied if d1=d2=d3

I win!

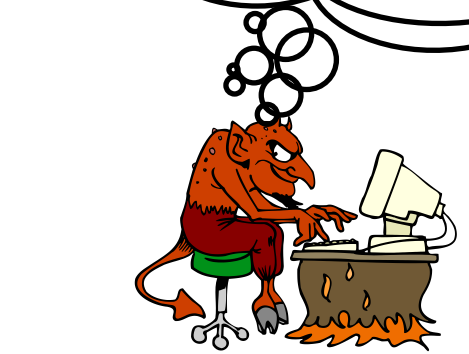

Example 2: Interval Based Auditing

d_i [0,100], sum queries, error =1

q1 = sum(d1,d2)

Sorry, denied

q2 = sum(d2,d3)

sum(d2,d3) = 50

Max Auditing

d1 d2 d3 d4 d5 d6 d7 d8 … d_{n-1} d_n

d_i real

q1 = max(d1,d2,d3,d4)

q2 = max(d1,d2,d3)

If denied: d4=M_{1234}

q2 = max(d1,d2)

If denied: d3=M_{123}

M_{1234}

M_{123} / denied

M_{12} / denied

Adversary's Success

q1 = max(d1,d2,d3,d4)

q2 = max(d1,d2,d3)

If denied: $d4=M_{1234}$

Denied with probability 1/4

q2 = max(d1,d2)

If denied: $d3=M_{123}$

Denied with probability 1/3

Success probability: $1/4 + (1- 1/4)\cdot 1/3 = 1/2$

Recover 1/8 of the database!

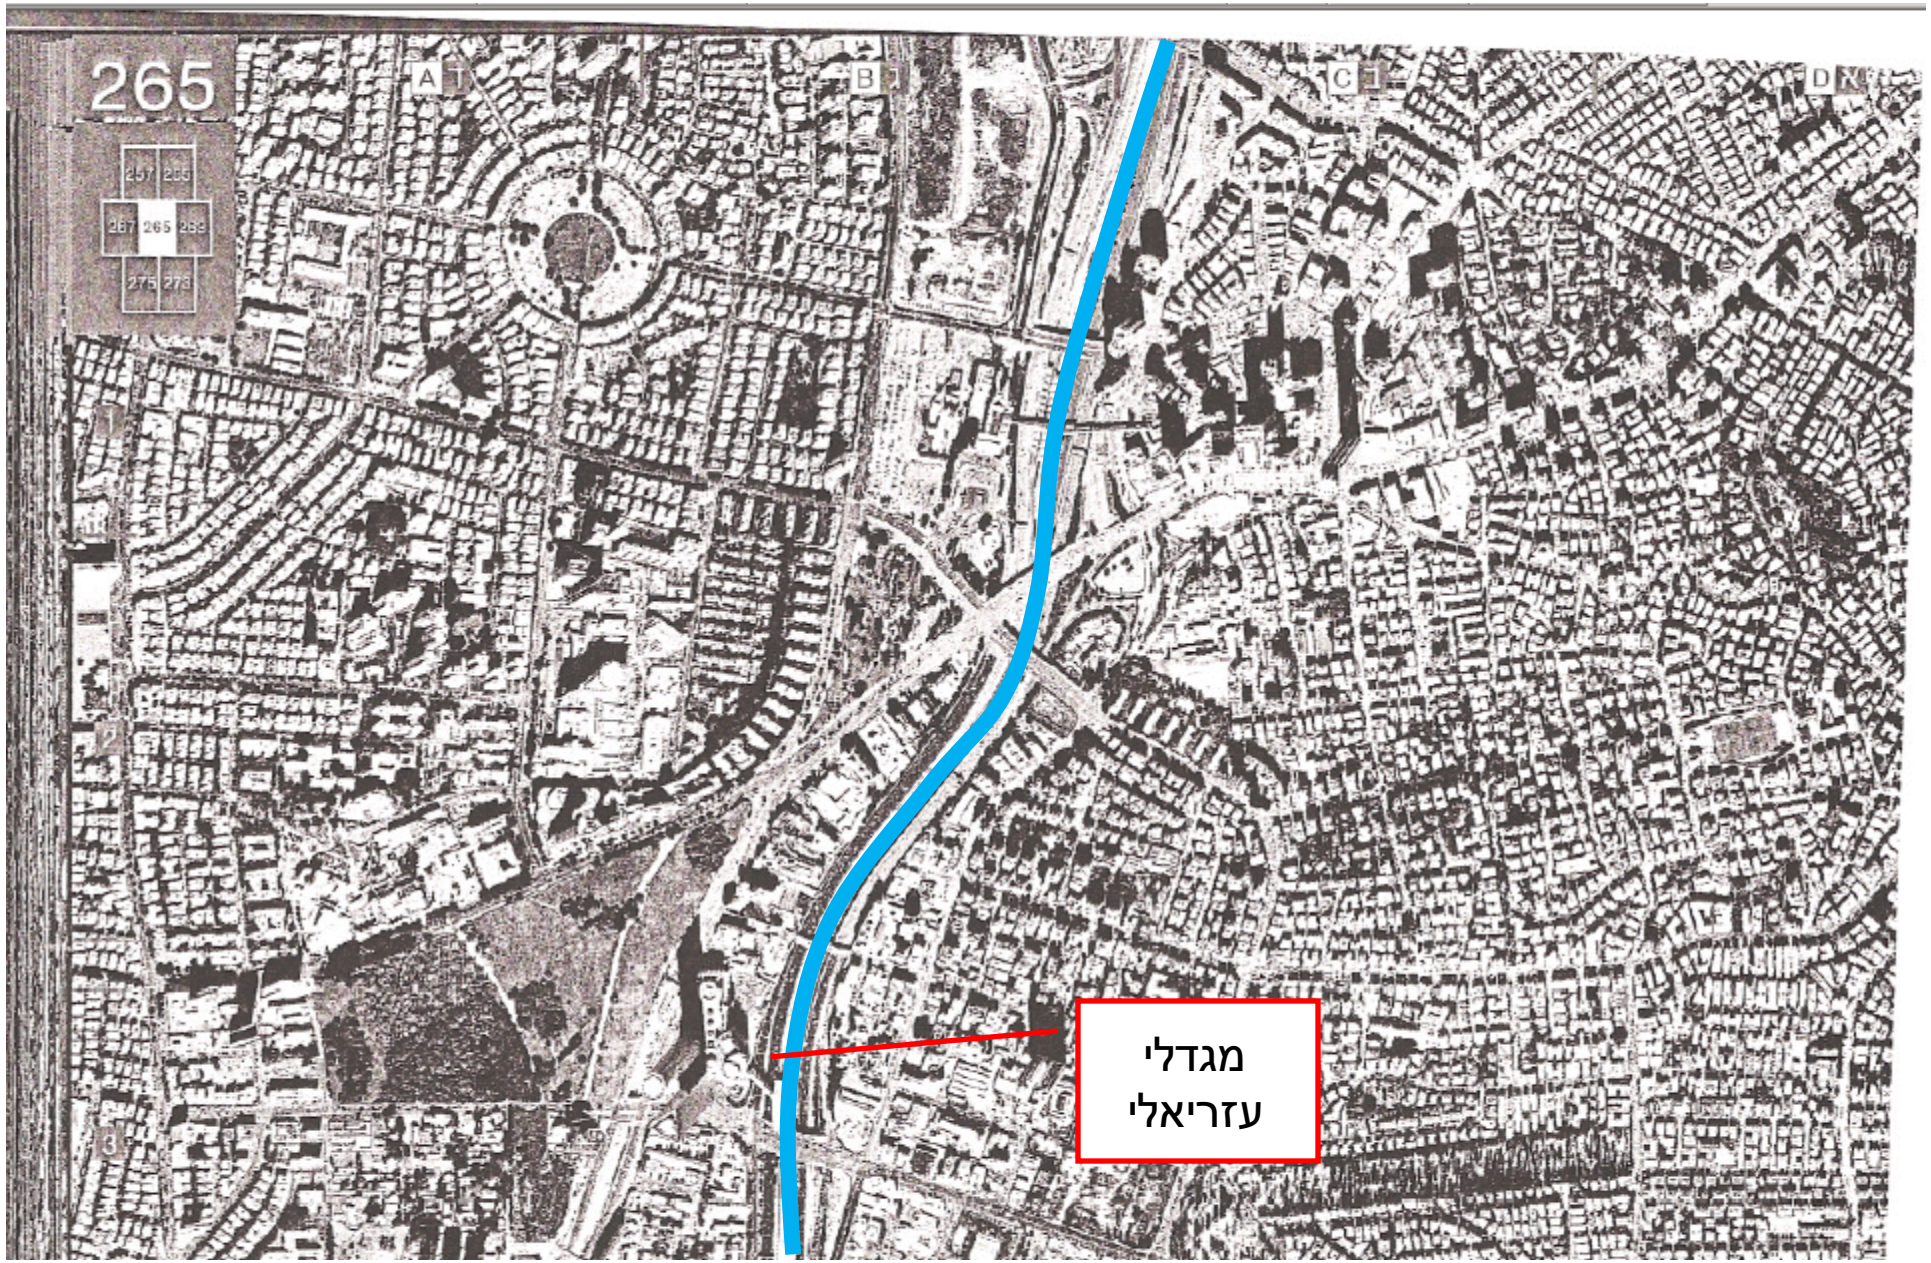

Thanks to Miri Regev

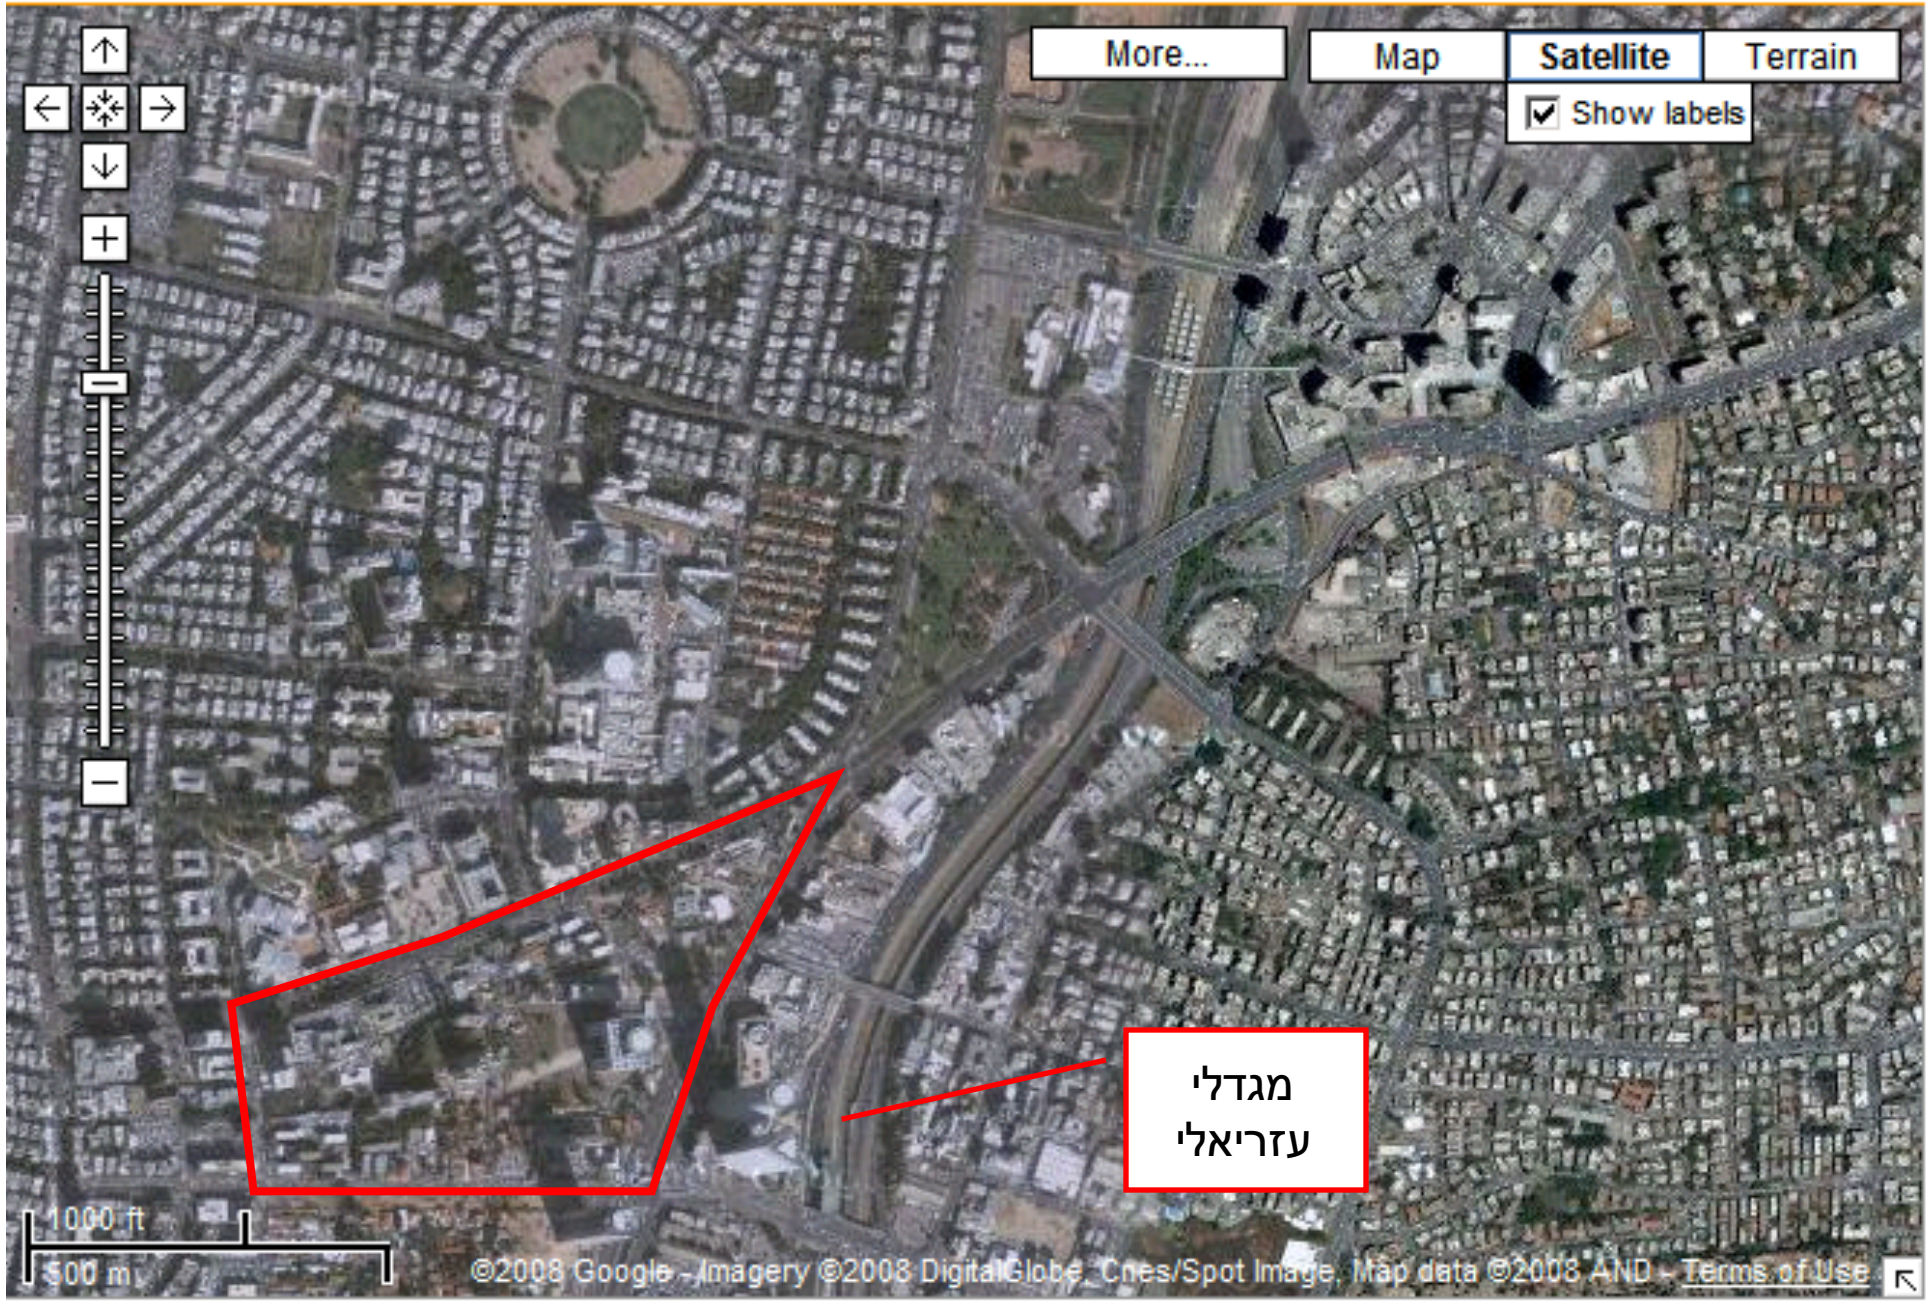

Thanks to Google

Randomization

The Scenario

- Users provide modified values of sensitive attributes
- Dataminer develops models about aggregated data

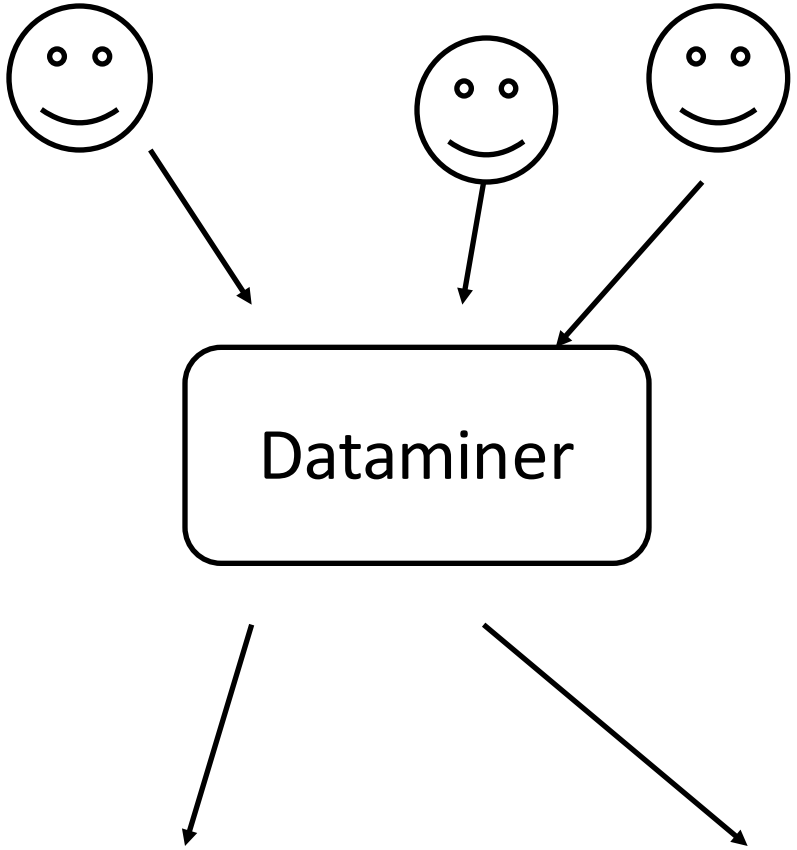

Can we develop accurate models without access to precise individual information?

Preserving Privacy

- Value distortion – return x_i+r_i
  - Uniform noise: $r_i \sim U(-a,a)$
  - Gaussian noise: $r_i \sim N(0,stdev)$
  - Perturbation of an entry is fixed
    - So that repeated queries do not reduce noise
- Privacy quantification: interval of confidence [AS 2000]
  - With c% confidence x_i is in the interval $[a_1, a_2]$. a_2-a_1 define the amount of privacy at c% confidence level
  - Examples:

	50%	95%	99.9%
Uniform	0.5 x 2a	0.95 x 2a	0.999 x 2a
Gaussian	1.34 stdev	3.92 stdev	6.8 stdev

  - Intuition: the larger the interval is, the better privacy is preserved.

Prior knowledge affects privacy

- Let X=age $r_i \sim U(-50,50)$
- [AS]: Privacy 100 at 100% confidence
- Seeing a measurement -10
  - Facts of life: Bob's age is between 0 and 40
- Assume you also know Bob has two children
  - Bob's age is between 15 and 40
- a-priori information may be used in attacking individual data

Privacy in Computation – The Problem

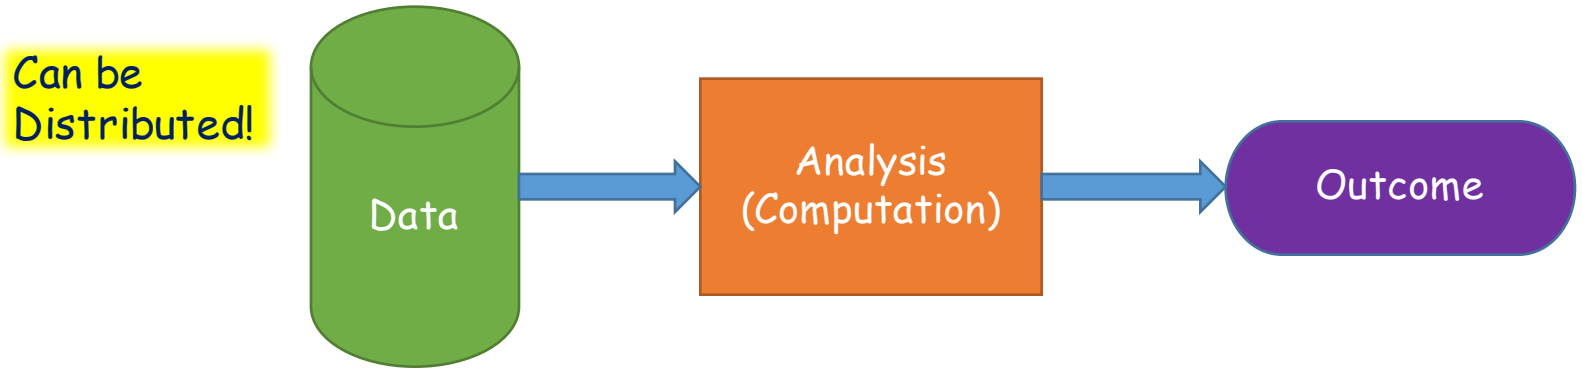

- Given a dataset with sensitive personal information

Health, social n/w, location, communication, …

- How to compute and release functions of the dataset

Academic research, informed policy, national security, …

- While protecting individual privacy

How many Trump fans in class?
Sensitive information …

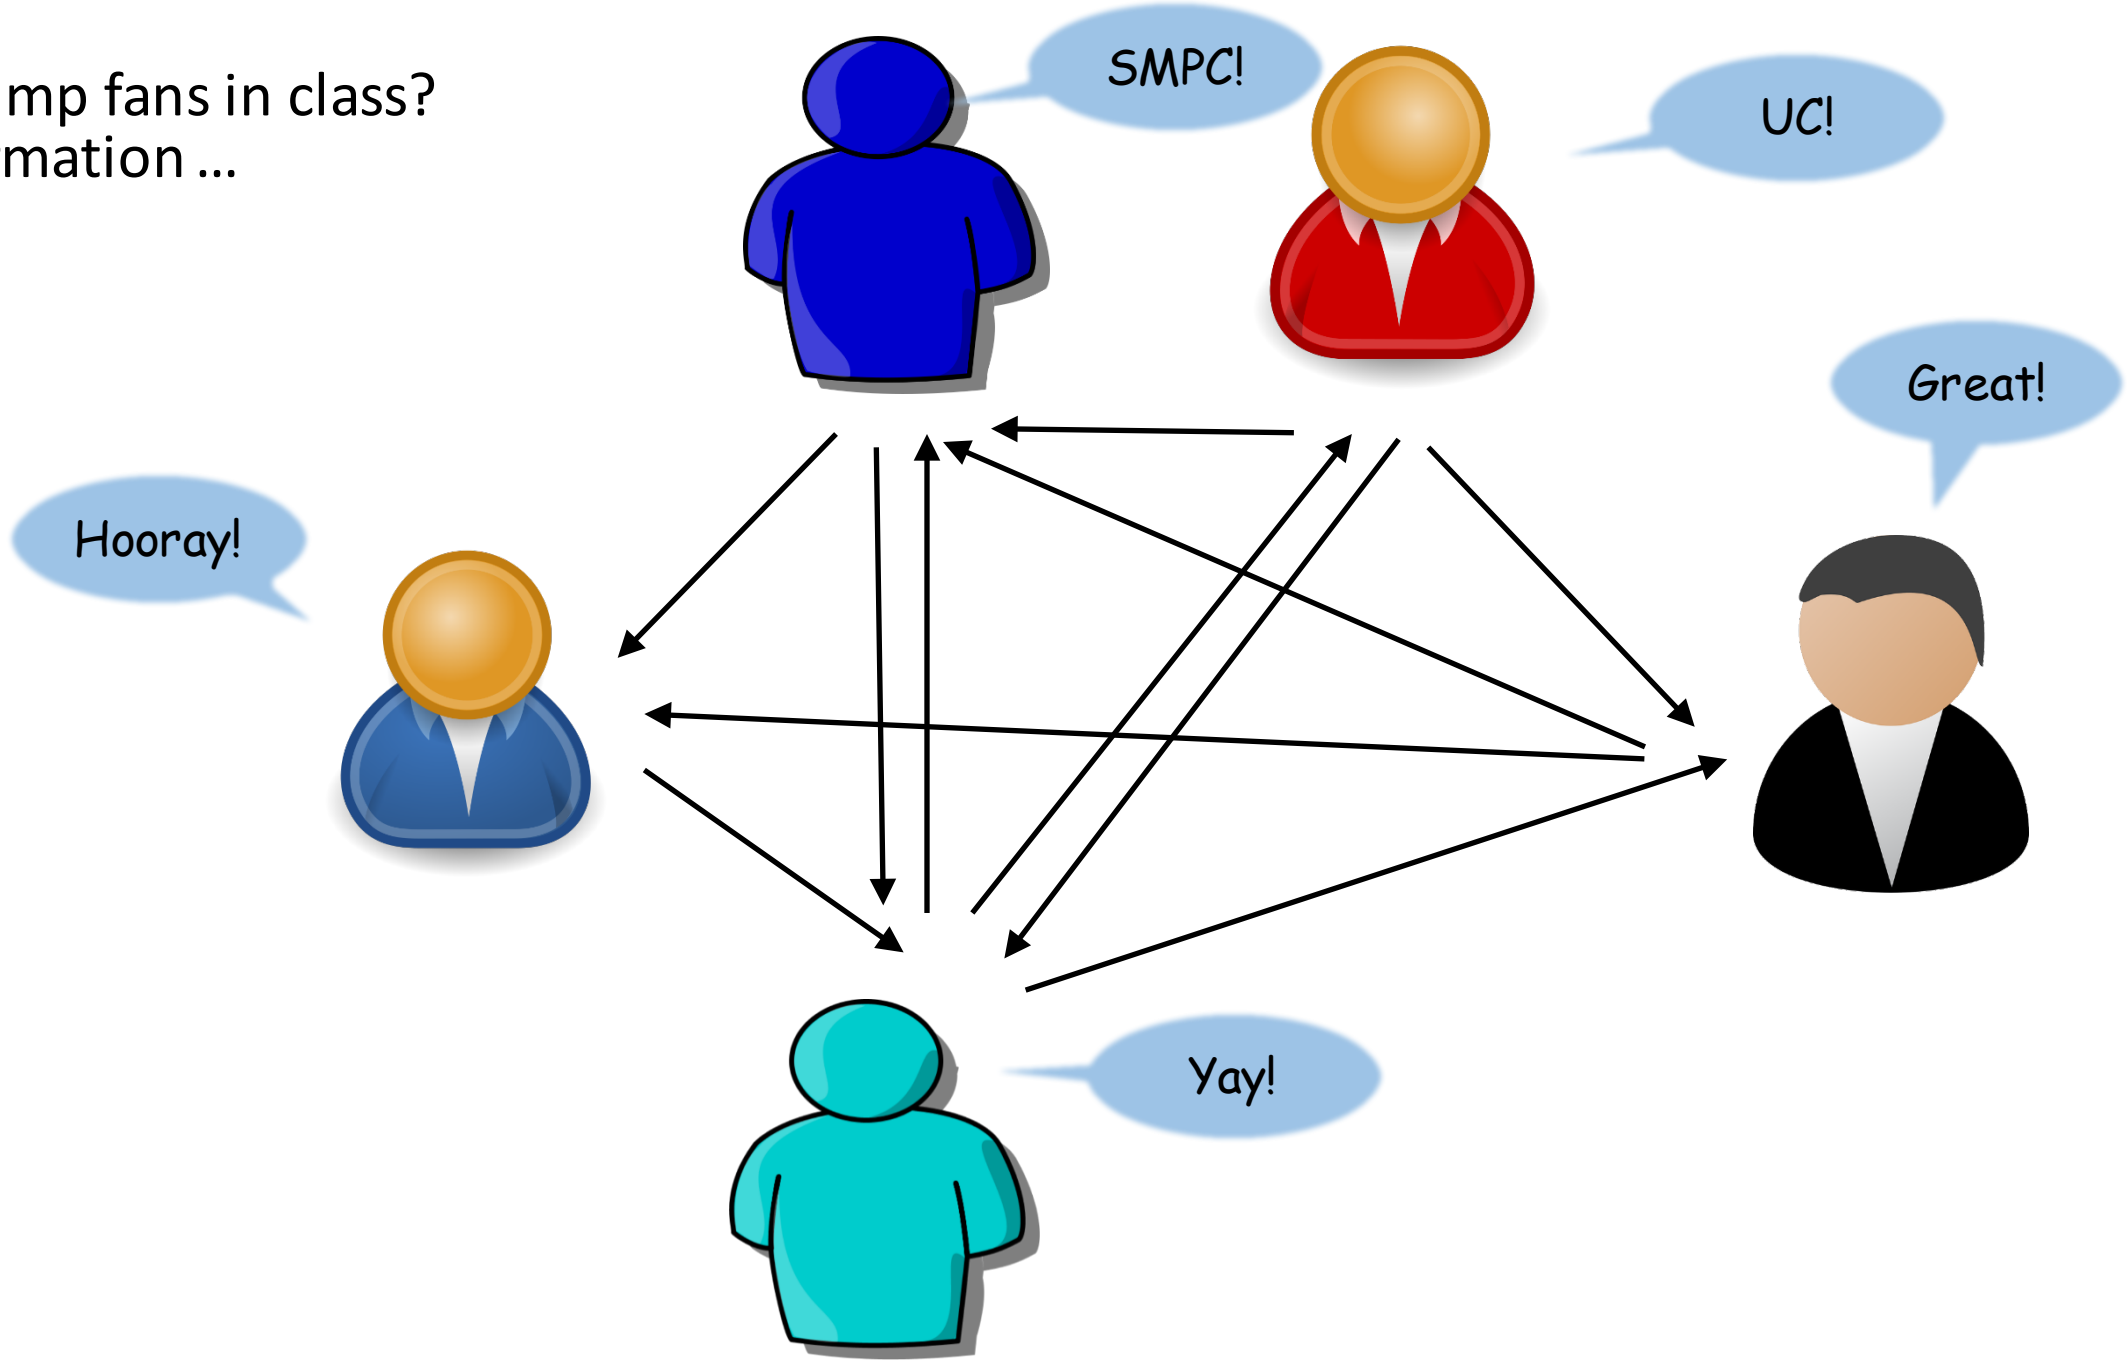

3 Trump fans in class

A few minutes later –
Kobbi join the class

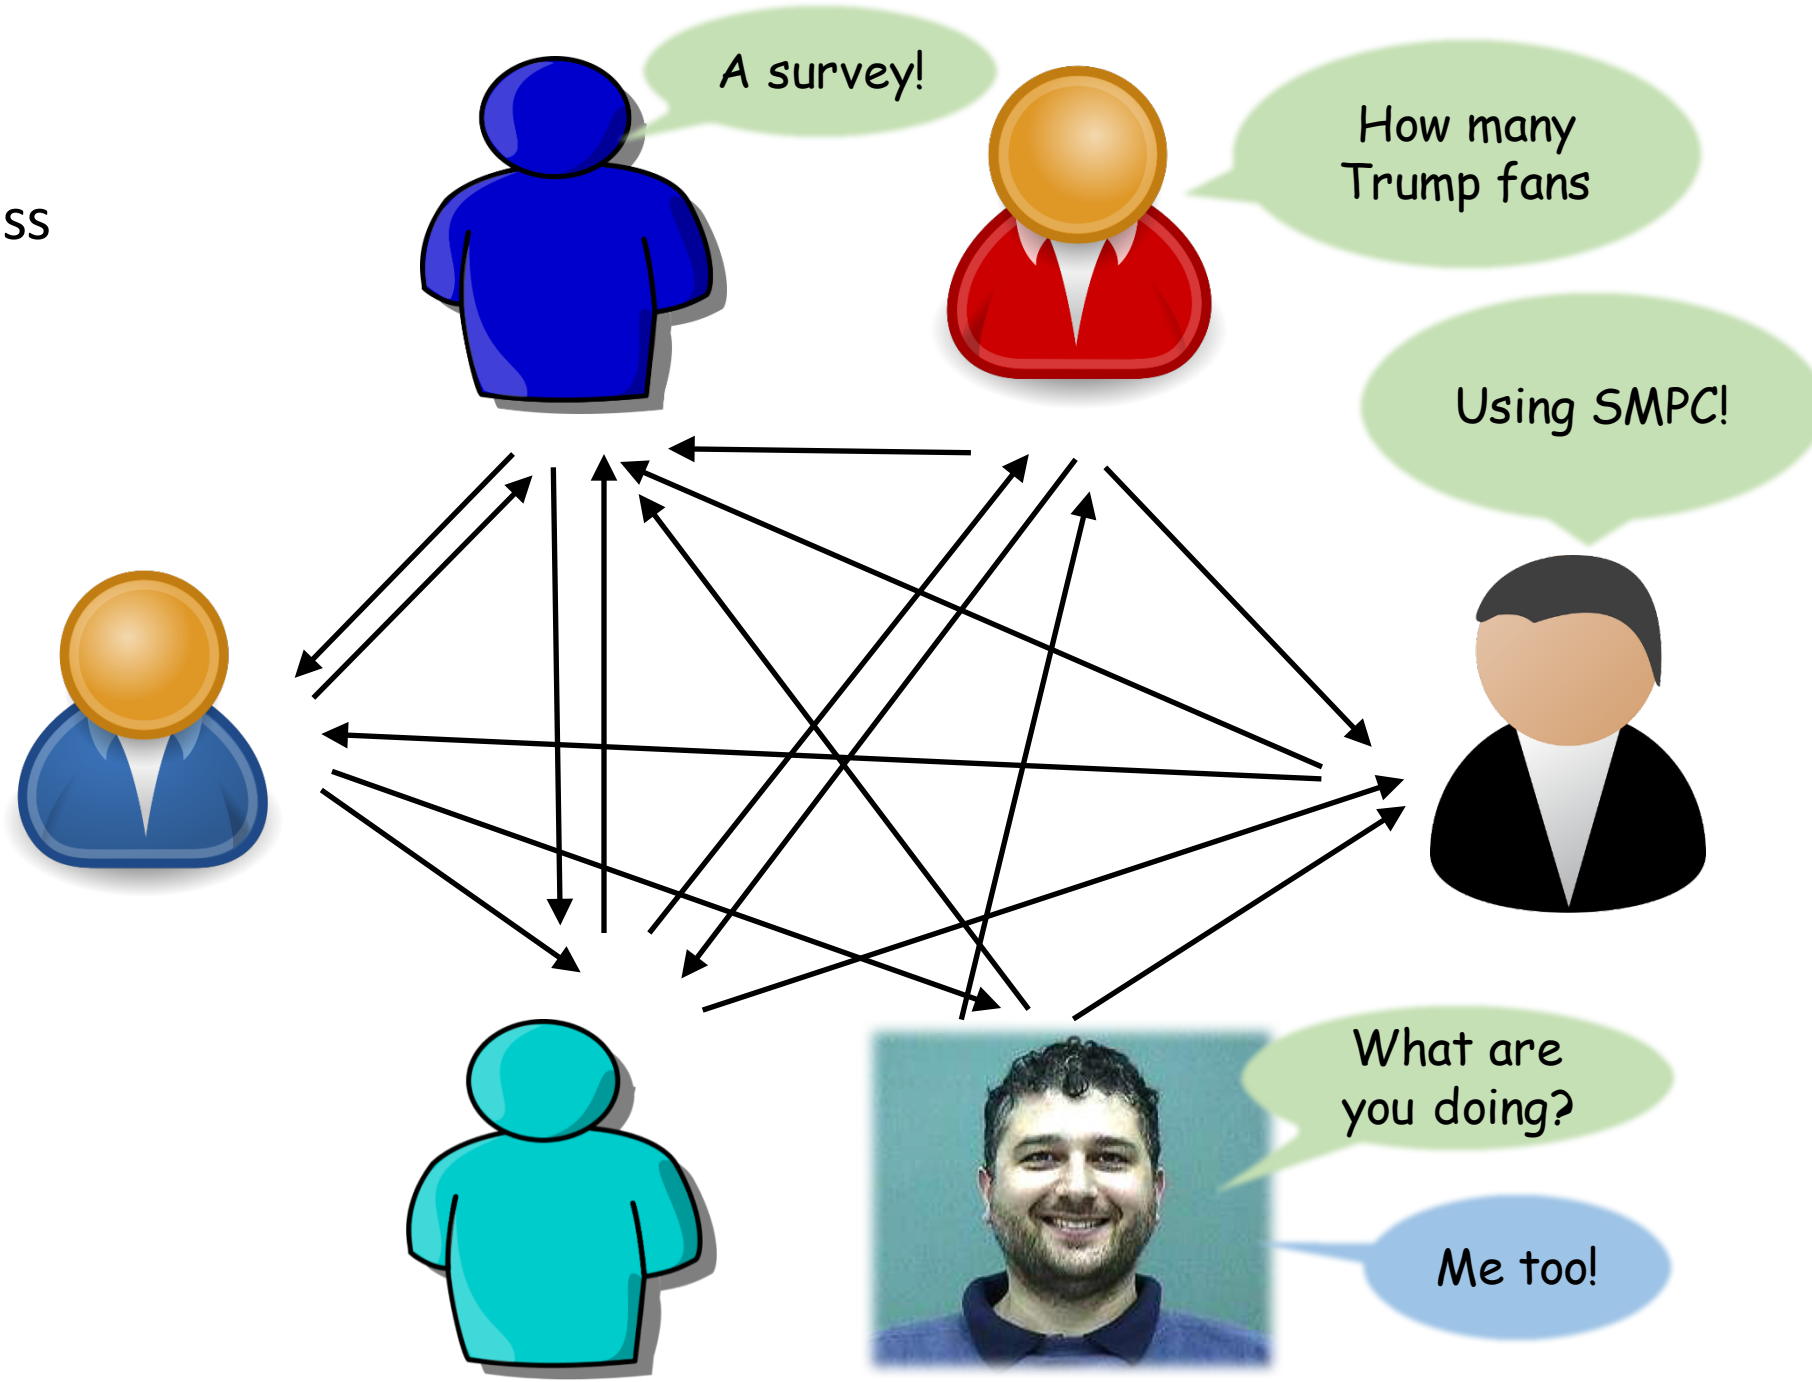

3 Trump fans

(after @Kobbi joins): 4 Trump fans

Composition

- Differencing attack:
  - How is my privacy affected when an attacker sees analysis before and after I join/leave?
- More generally: Composition
  - How is my privacy affected when an attacker combines results from two or more privacy preserving analyses?
- Fundamental law of information: the more information we extract from our data, the more is learned about individuals!
  - So, privacy *will* deteriorate as we use our data more and more
  - Best desiderata:
    - Deterioration is quantifiable and controllable
    - Deterioration not abrupt

My Privacy Desiderata

Real world:

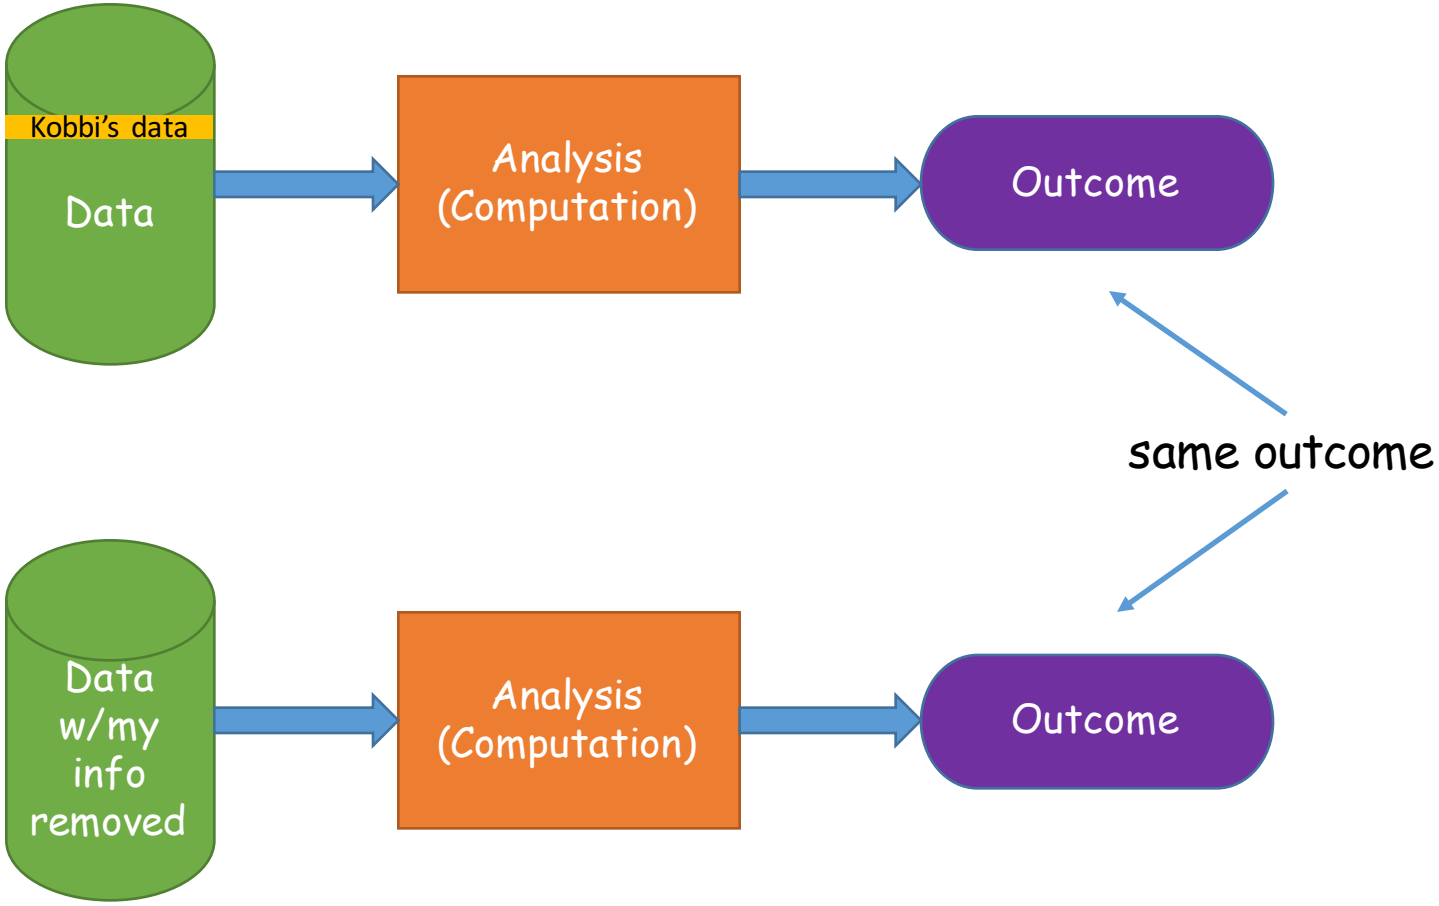

My ideal world:

Things to Note

- We only consider the *intended outcome* of analyses
  - Security flaws, hacking, implementation errors, human factors, …
    - Very important but very *different* questions
- My ideal world would hide whether I'm Trump fan
  - Resilient to differencing attacks
- Does not mean I'm fully protected
  - I'm only protected to the extent I'm protected in my ideal world
  - Some harm could happen to me even in my ideal world
    - Bob smokes in public
    - Study teaches that smoking causes cancer
    - Bob's health insurer raises his premium
    - Bob is harmed even if he does not participate in the study!

Our Privacy Desiderata

Real world:

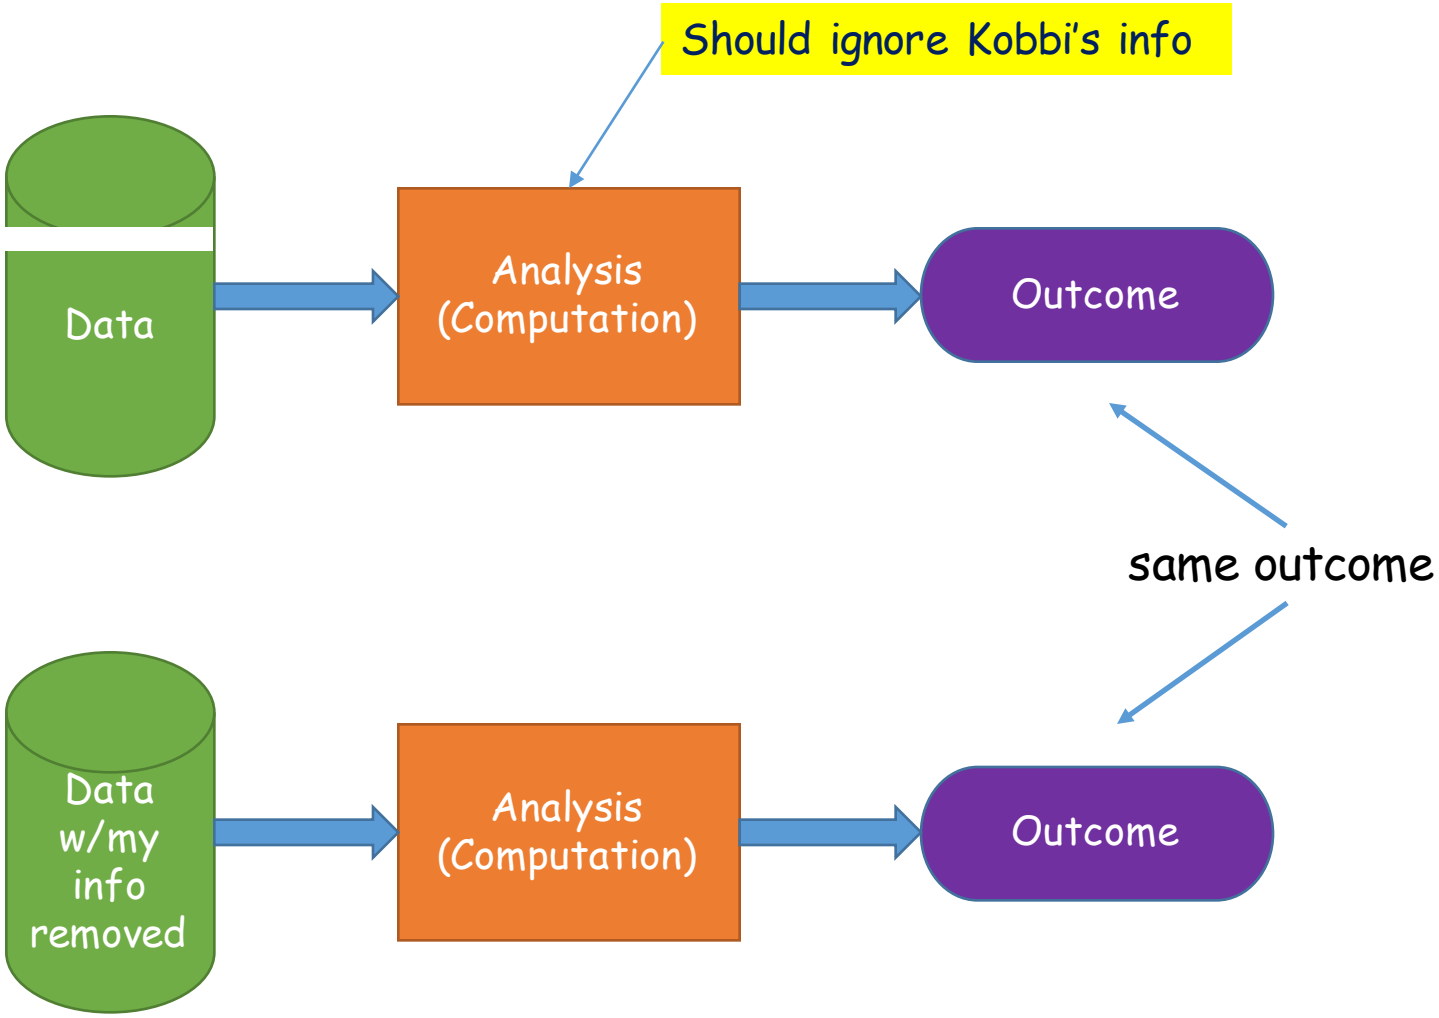

My ideal world:

Our Privacy Desiderata

Real world:

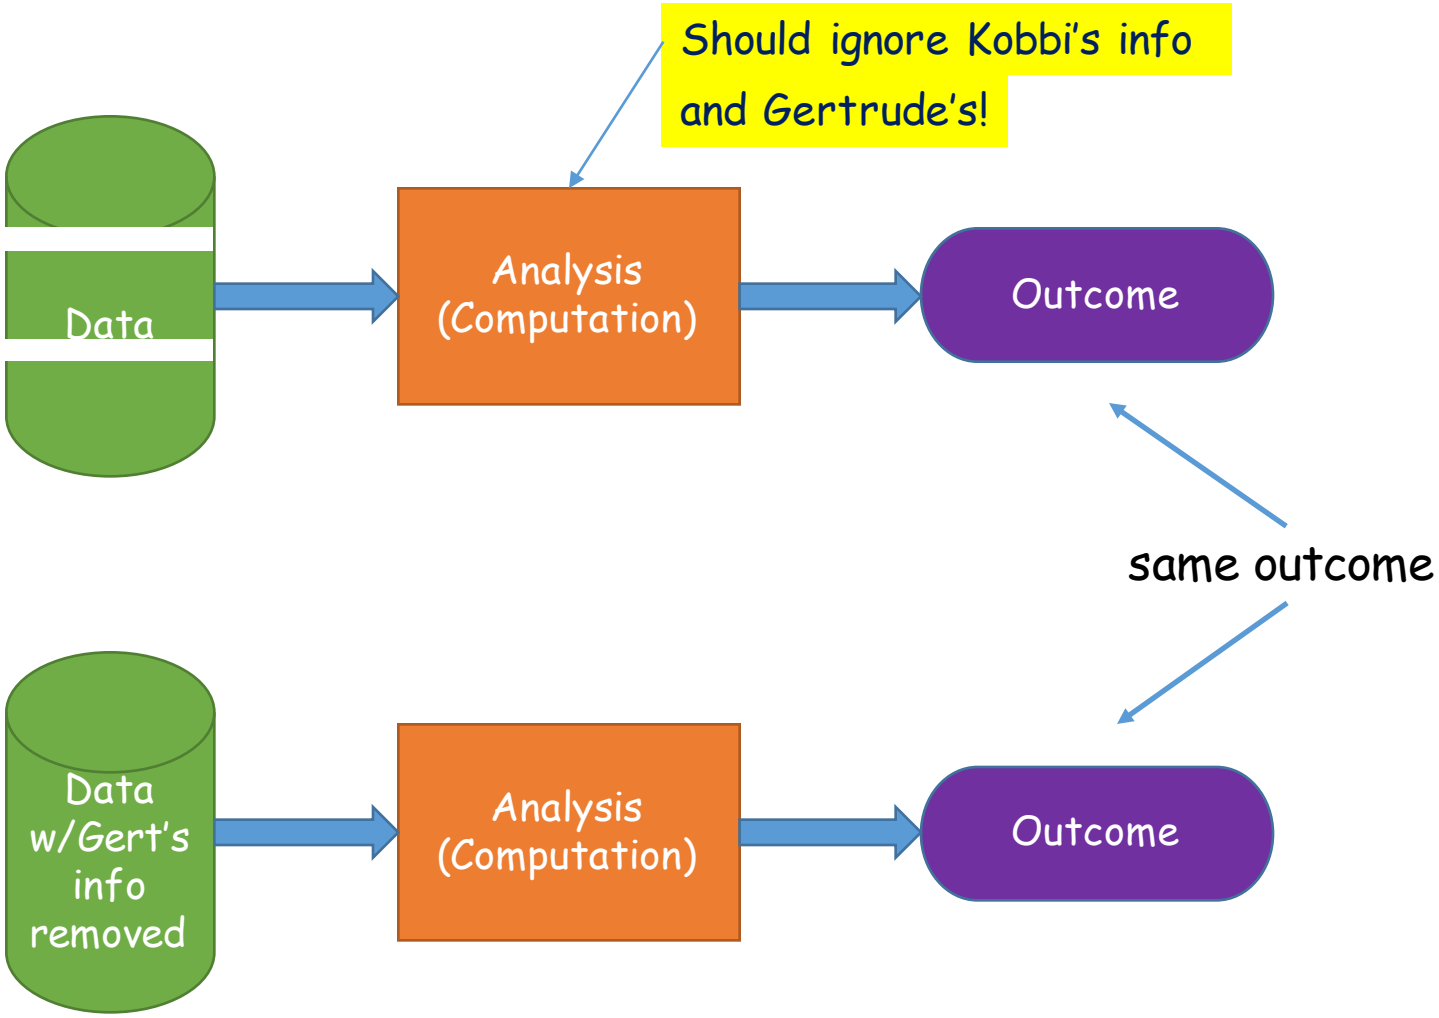

Gert's ideal world:

Our Privacy Desiderata

Real world:

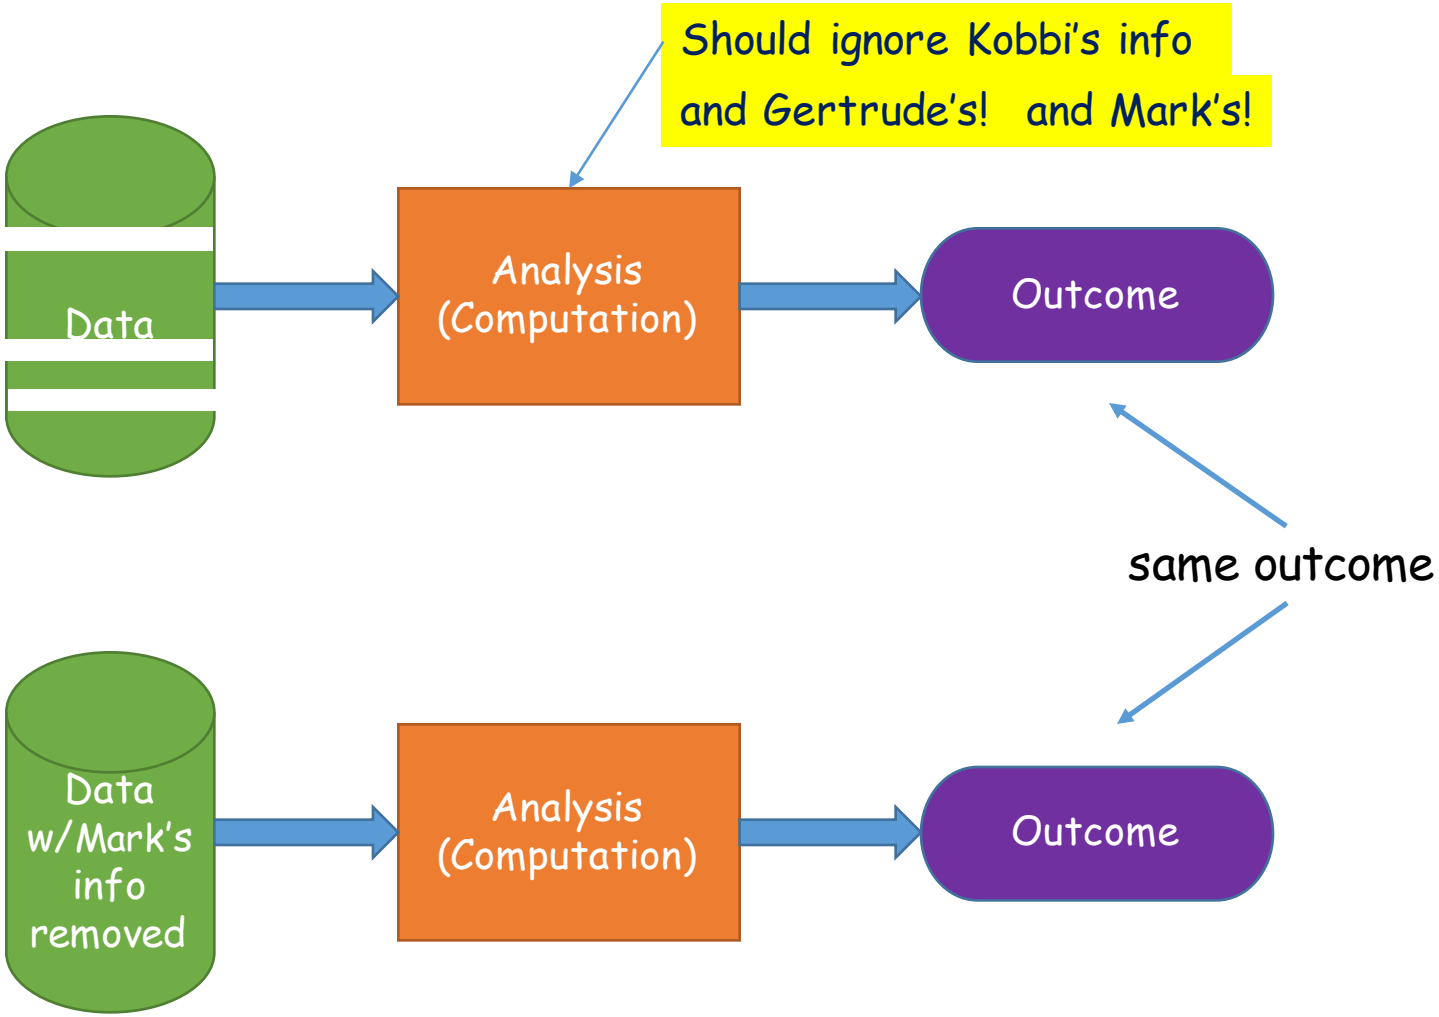

Mark's ideal world:

Our Privacy Desiderata

Real world:

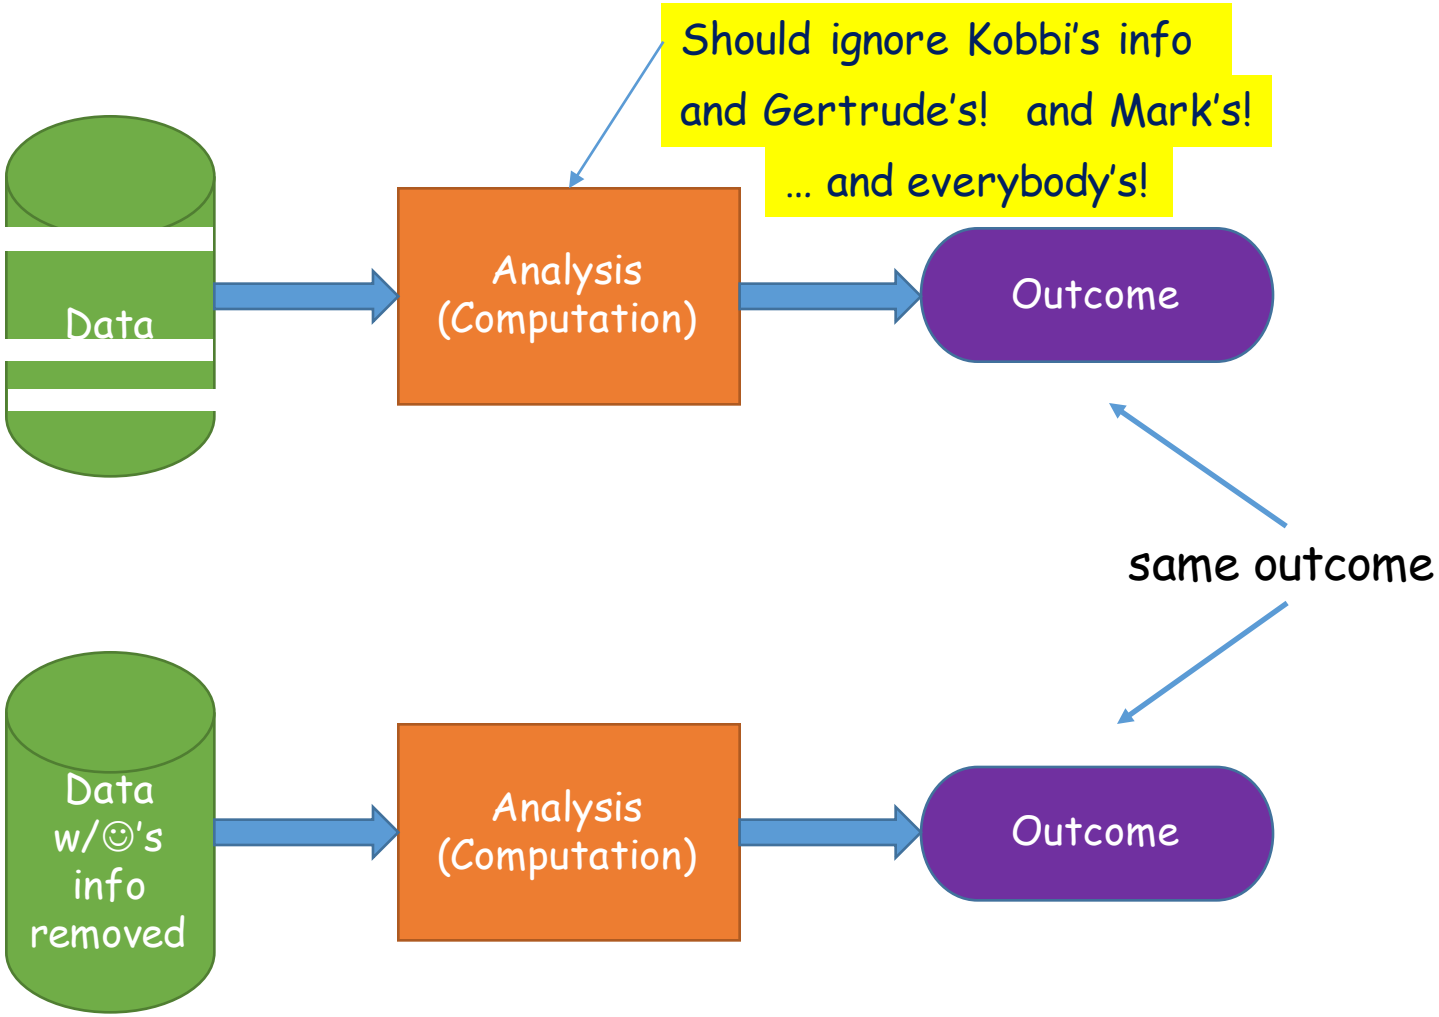

☺'s ideal world:

A Realistic Privacy Desiderata

Real world:

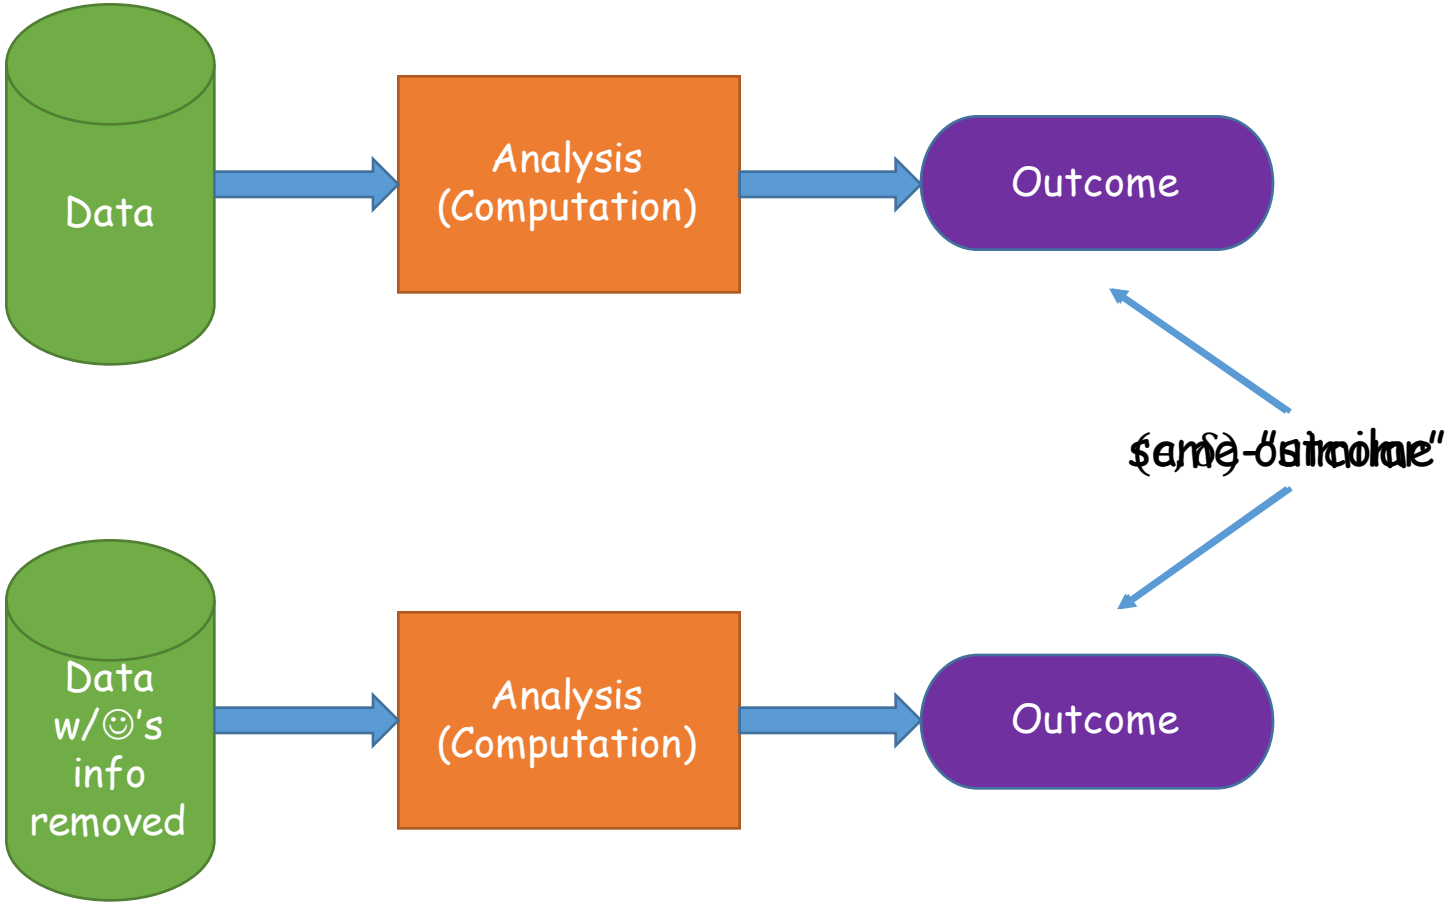

☺'s ideal world:

Differential Privacy [Dwork McSherry N Smith 06]

Real world:

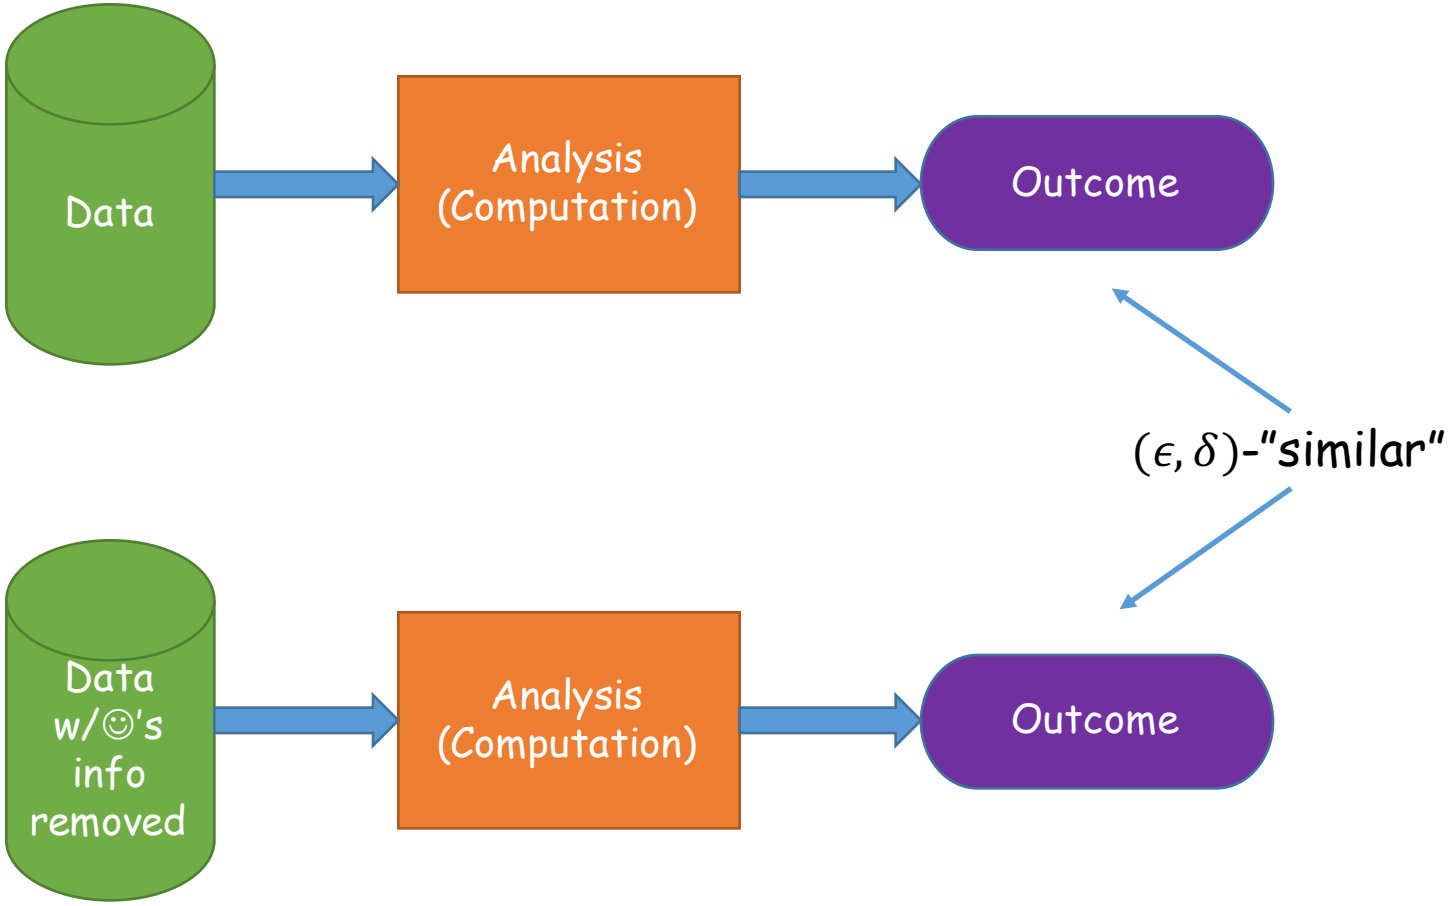

☺'s ideal world:

Differential Privacy [Dwork McSherry N Smith 06]

A (randomized) algorithm $M: X^n \to T$ satisfies (ϵ, δ)-differential privacy if for all datasets $x, x' \in X^n$ that differ on one entry,
For all subsets S of the outcome space T,

$$\Pr_{\mathrm{M}}[M(x) \in S] \leq e^{\epsilon} \Pr_{\mathrm{M}}[M(x') \in S] + \delta.$$

- The parameter ϵ measures 'leakage' or 'harm'
- For small ϵ: $e^{\epsilon} \approx 1 + \epsilon \approx 1$. Think $\epsilon \approx \frac{1}{100}$ or $\epsilon \approx \frac{1}{10}$ or $\epsilon \approx \frac{1}{\sqrt{n}}$ but not $\epsilon < \frac{1}{n}$
- Meaningful δ should be small ($\delta \ll 1/n$). Think $\delta \approx 2^{-60}$, or cryptographically small, or zero

Why Differential Privacy?

- DP: Strong, quantifiable, composable mathematical privacy guarantee
- A definition/standard/requirement, not a specific algorithm
  - Opens door to algorithmic development
- Provably resilient to known and unknown attack modes!
- Has a natural interpretation: I am protected (almost) to the extent I'm protected in my privacy-ideal scenario
- Theoretically, DP enables many computations with personal data while preserving personal privacy
  - Practicality in first stages of validation

Interpreting Differential Privacy

- ~~A naïve hope: Your beliefs about me are the same after you see the output as they were before.~~
- Suppose I smoke in public
  - A public health study could teach that I am at risk for cancer.
  - But it didn't matter whether or not my data was part of it.

Interpreting Differential Privacy

- Compare $x = (x_1, x_2, \dots, x_i, \dots, x_n)$
  to $x_{-i} = (x_1, x_2, \dots, \perp, \dots, x_n)$

- A is ε-differentially private if for all vectors x and for all i: $A(x) \approx_{\varepsilon} A(x_{-i})$.

- No matter what you know ahead of time, you learn (almost) the same things about me whether or not my data are used.

Interpreting Differential Privacy

- Compare $x = (x_1, x_2, \dots, x_i, \dots, x_n)$
  to $x_{-i} = (x_1, x_2, \dots, \perp, \dots, x_n)$

- A is ε-differentially private if for all vectors x and for all i: $A(x) \approx_{\varepsilon} A(x_{-i})$.

- Bayesian interpretation [DM'06, KS'08]:
  - A is ε-differentially private if for all distributions X, for all i:

$$X_i \mid A(X) = t \approx_{\varepsilon} X_i \mid A(X_{-i}) = t$$

Under any reasonable distance metric

Are you HIV positive?

- Toss Coin
  - If heads tell the truth
  - If tails, toss another coin:
    - If heads say yes
    - If tails say no

Randomized response [Warner 65]

- $x \in \{0,1\}$

$\alpha = ¼$

- $RR_\alpha(x) = f(x) = \begin{cases} x & w.p.\ \frac{1}{2} + \alpha \\ \neg x & w.p.\ \frac{1}{2} - \alpha \end{cases}$

w.p. ¾

w.p. ¼

- Claim: setting $\alpha = \frac{1}{2}\frac{e^\epsilon - 1}{e^\epsilon + 1}$, $RR_\alpha(x)$ is ϵ −differentially priv

small ϵ: $e^\epsilon \approx 1 + \epsilon$.
Get $\alpha \approx \frac{\epsilon}{4}$

- Proof:
  - Neighboring databases: $x = 0; x' = 1$
  - $\frac{\Pr[RR(0)=0]}{\Pr[RR(1)=0]} = \frac{\frac{1}{2}(1+\frac{e^\epsilon-1}{e^\epsilon+1})}{\frac{1}{2}(1-\frac{e^\epsilon-1}{e^\epsilon+1})} = e^\epsilon$

$\epsilon = \ln 3$

Can we make use of randomized response?

- $(x_1, x_2, \dots, x_n) \in \{0,1\}^n$; want to estimate $\sum x_i$
- $Y_i = RR_\alpha(x_i)$
- $E[Y_i] = x_i\left(\frac{1}{2} + \alpha\right) + (1 - x_i)\left(\frac{1}{2} - \alpha\right) = \frac{1}{2} + \alpha(2x_i - 1)$
- $x_i = \frac{E[Y_i] - \frac{1}{2} + \alpha}{2\alpha}$ suggesting estimate $\widehat{x_i} = \frac{y_i - \frac{1}{2} + \alpha}{2\alpha}$
- $E[\widehat{x_i}] = x_i$ by construction but $Var[\widehat{x_i}] = \frac{\frac{1}{4} - \alpha^2}{4\alpha^2} \approx \frac{1}{\epsilon^2}$ high!
- $E[\sum \widehat{x_i}] = \sum x_i$ and $Var[\sum \widehat{x_i}] = n\frac{\frac{1}{4} - \alpha^2}{4\alpha^2} \approx \frac{n}{\epsilon^2}$; stdev $\approx \frac{\sqrt{n}}{\epsilon}$
- Useful when $\frac{\sqrt{n}}{\epsilon} \ll n$ (i.e. $n \gg \frac{1}{\epsilon^2}$)

$\alpha =$ ¼

$$E[Y_i] = \frac{1}{4} + \frac{x_i}{2}$$

$$\widehat{x_i} = 2y_i - \frac{1}{2}$$

$$\text{Stdev} = \sqrt{\frac{3}{4}n}$$

Lots of noise?

Compare with sampling noise $\approx \sqrt{n}$

Basic properties of differential privacy: post processing

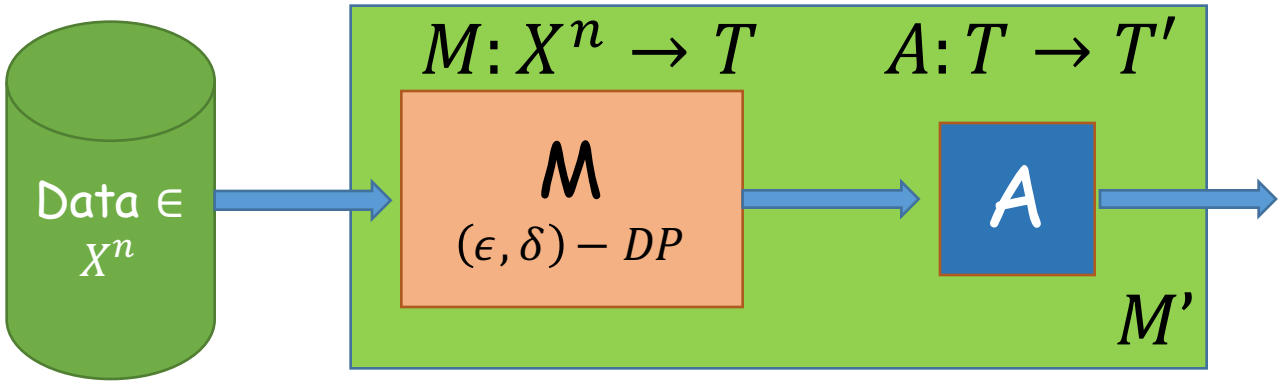

- Claim: M' is (ϵ, δ)-differentially private
- Proof:
  - Let x, x' be neighboring databases and S' a subset of T'
  - Let $S = \{z \in T: A(z) \in S'\}$ be the preimage of S' under A

$$\Pr[M'(x) \in S'] = \Pr[M(x) \in S]$$

$$\leq e^{\epsilon} \Pr[M(x') \in S] + \delta \;\; = e^{\epsilon} \Pr[M'(x') \in S'] + \delta$$

Basic properties of differential privacy: Basic composition [DMNS06, DKMMN06, DL09]

- Claim: M' is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$-differentially private
- Proof (for the case $\delta_1 = \delta_2 = 0$)
  - Let x, x' be neighboring databases and S a subset of $(T_1 \times T_2)$

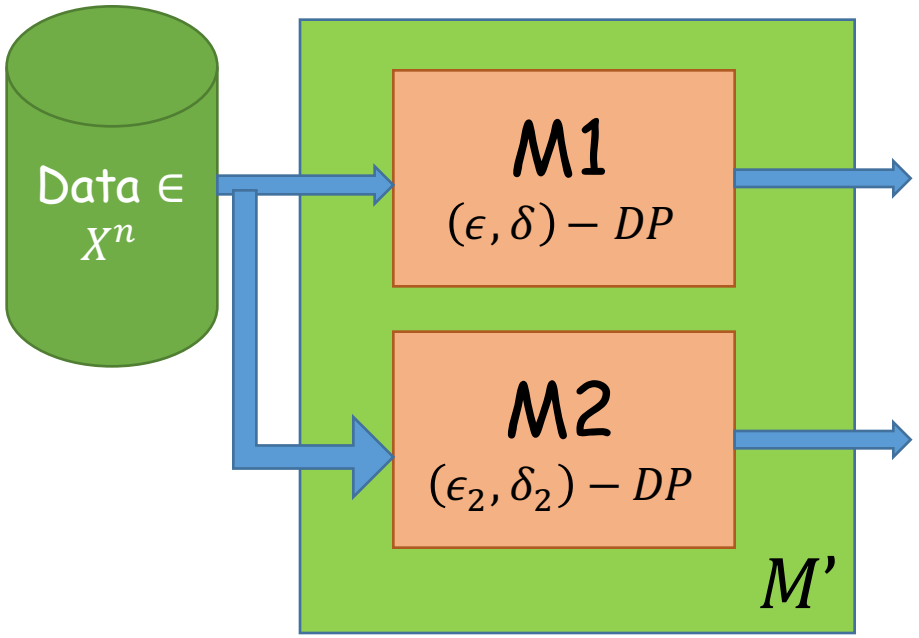

$M_i : X^n \to T_i$

$$\Pr[M'(x) \in S] = \sum_{(z_1,z_2)\in S} \Pr[M_1(x) = z_1 \wedge M_2(x) = z_2]$$

$$= \sum_{(z_1,z_2)\in S} \Pr[M_1(x) = z_1]\ \Pr[M_2(x) = z_2]$$

$$\leq \sum_{(z_1,z_2)\in S} e^{\epsilon_1}\Pr[M_1(x') = z_1]\ e^{\epsilon_2}\Pr[M_2(x') = z_2] = e^{\epsilon_1+\epsilon_2}\Pr[M'(x') \in S]$$

Basic properties of DP: Group privacy

- Let M be (ϵ, δ)-differentially private:
  - For all datasets $x, x' \in X^n$ that differ on one entry, for all subsets S of the outcome space T:

$$\Pr_{\mathrm{M}}[M(x) \in S] \leq e^{\epsilon} \Pr_{\mathrm{M}}[M(x') \in S] + \delta.$$

- Claim: for all databases $x, x' \in X^n$ that differ on t entries, for all subsets S of the outcome space T:

$$\Pr_{\mathrm{M}}[M(x) \in S] \leq e^{t\epsilon} \Pr_{\mathrm{M}}[M(x') \in S] + t\delta e^{t\epsilon}.$$

Neighboring Inputs

[What Should Be Protected?]

- Inputs are neighboring if they differ on the data of a single individual
  - Record privacy: Databases X, X' neighboring if differ on one record

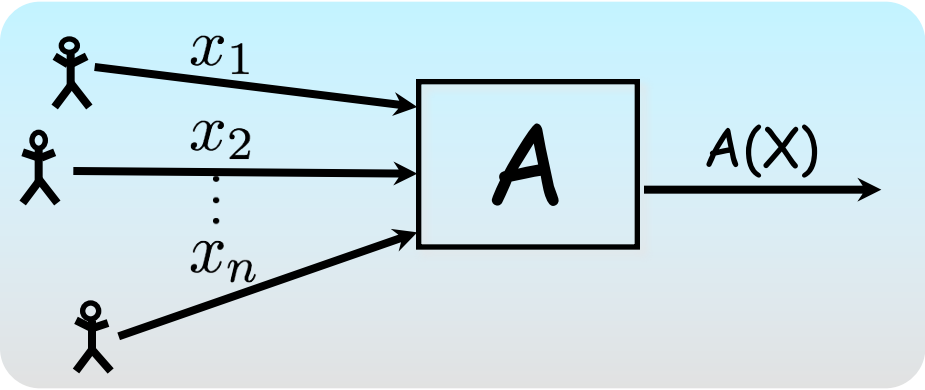

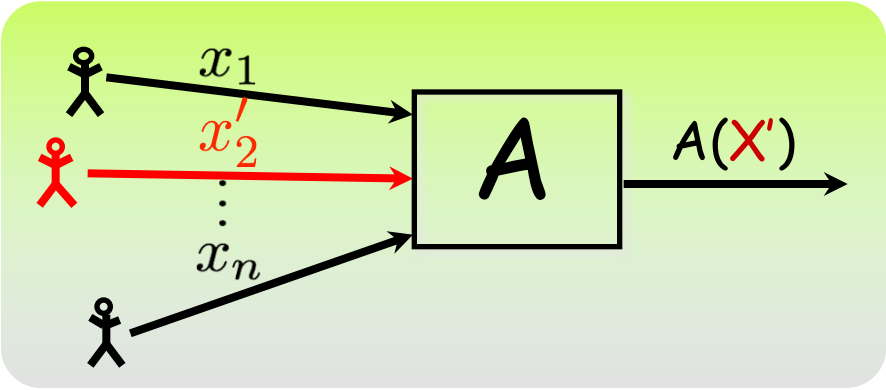

Neighboring Inputs

[What Should Be Protected?]

- Inputs are neighboring if they differ on the data of a single individual
  - Record privacy: Databases X, X' neighboring if differ on one record
  - Edge privacy: graphs G, G' neighboring if differ on one edge

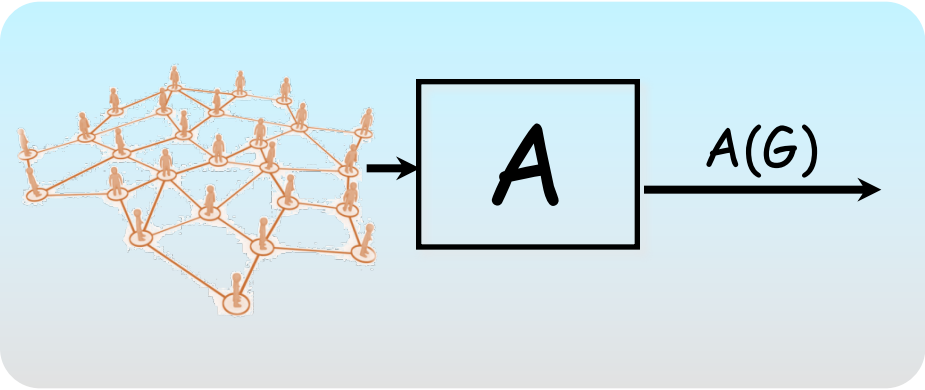

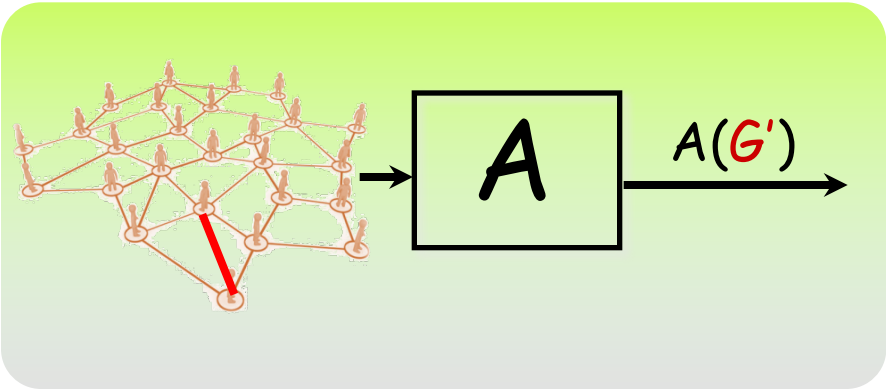

Image credit: www.perey.com

Neighboring Inputs

[What Should Be Protected?]

- Inputs are neighboring if they differ on the data of a single individual
  - Record privacy: Databases X, X' neighboring if differ on one record
  - Edge privacy: graphs G, G' neighboring if differ on one edge
  - Node privacy: graphs G, G' neighboring if differ on one node and its adjacent edges

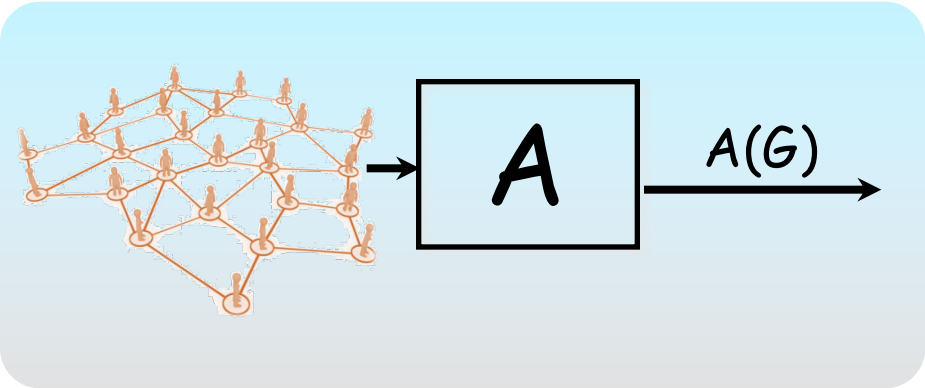

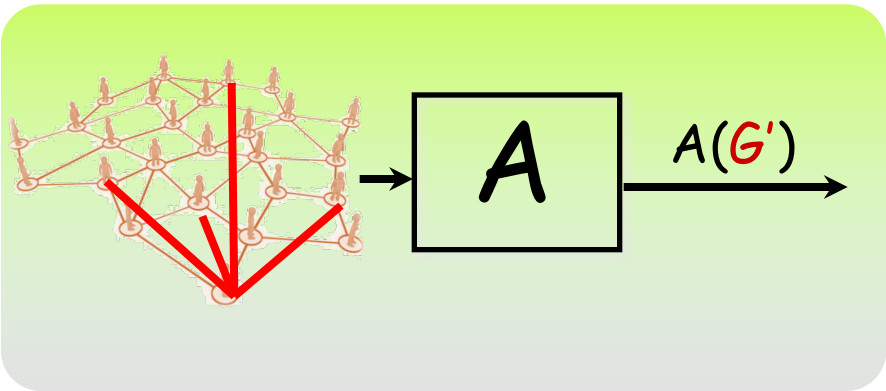

Image credit: www.perey.com

Laplacian: $Lap(b)$

- $Lap(b)$ is a distribution over the reals with density function
  - $h(x) = \frac{1}{2b}\exp(-\frac{|x|}{b})$
  - Expectation zero
  - Variance $2b^2$

$Lap(1/\epsilon)$:

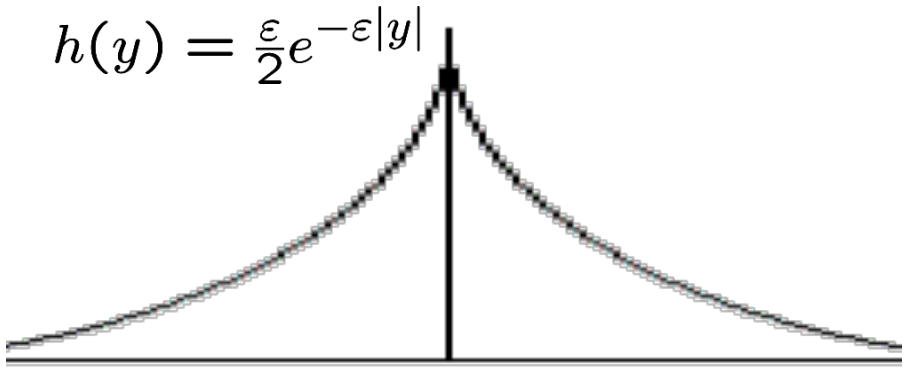

Laplacian: $Lap(b)$

- $Lap(b)$ is a distribution over the reals with density function
  - $h(x) = \frac{1}{2b}\exp(-\frac{|x|}{b})$
  - Expectation zero
  - Variance $2b^2$
  - if $Y \sim Lap(b)$ then $\Pr(|Y| > t \cdot b) = \exp(-t)$

Example: Counting Edges

[the basic technique]

- function $f(G) = \sum e_{ij}$ where $e_{ij} \in \{0,1\}$

Example: Counting Edges

[the basic technique]

- $f(G) = \sum e_{ij}$ where $e_{ij} \in \{0,1\}$
- Algorithm: On input G return $f(G) + Y$, where $Y \sim Lap(\frac{1}{\epsilon})$

  - $E[Y] = 0; \sigma[Y] = \sqrt{2}/\varepsilon$

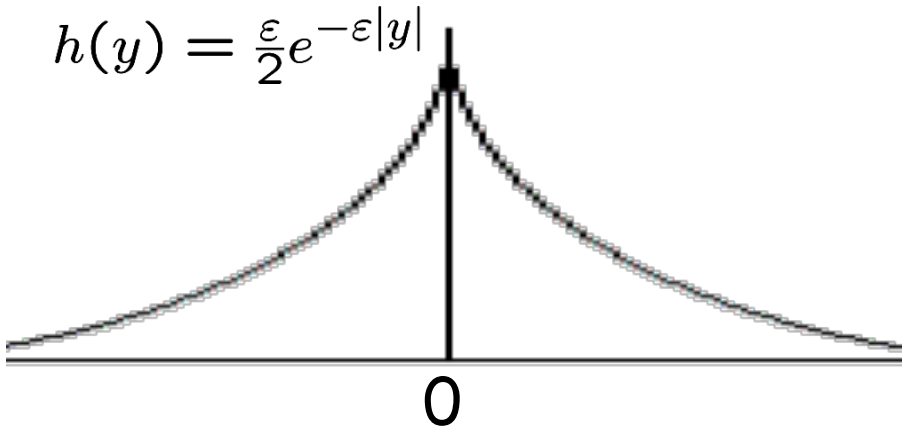

Example: Counting Edges

[the basic technique]

- $f(G) = \sum e_{ij}$ where $e_{ij} \in \{0,1\}$
- Algorithm: On input G return $f(G) + Y$, where $Y \sim Lap(\frac{1}{\epsilon})$

- Laplace Distribution:
  - $E[Y] = 0; \sigma[Y] = \sqrt{2}/\varepsilon$
  - Sliding property: $\frac{h(y)}{h(y+1)} \leq e^{\epsilon}$

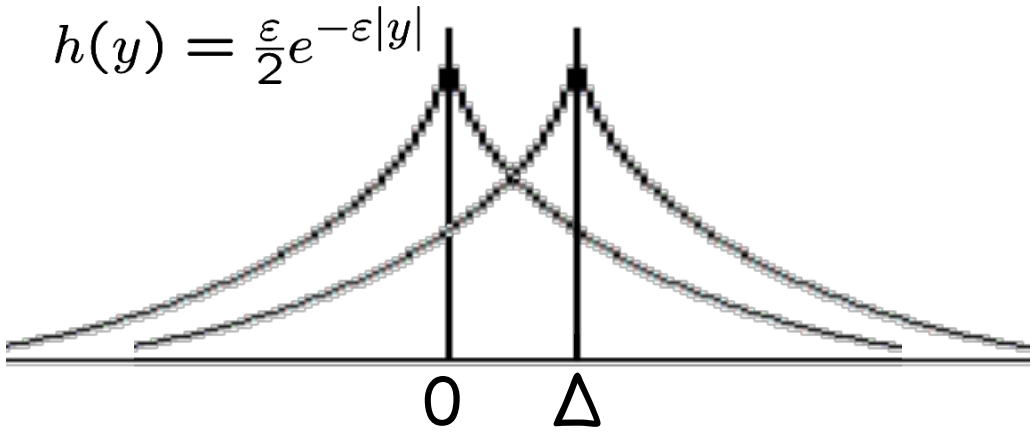

Example: Counting Edges

[the basic technique]

- $f(G) = \sum e_{ij}$ where $e_{ij} \in \{0,1\}$
- Algorithm: On input G return $f(G) + Y$, where $Y \sim Lap(\frac{1}{\epsilon})$

- Laplace Distribution:
  - $E[Y] = 0; \sigma[Y] = \sqrt{2}/\varepsilon$
  - Sliding property: $\frac{h(y)}{h(y+1)} \leq e^{\epsilon}$

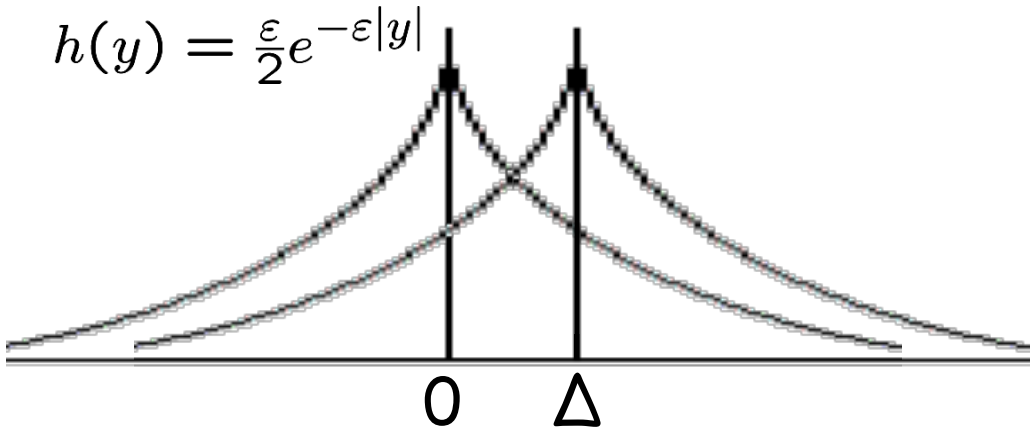

- For G, G' edge neighboring:

$$|f(G) - f(G')| = \left| \sum_{ij} e_{ij} - \sum_{ij} e'_{ij} \right| \leq 1$$

Neighboring: ℓ_1 distance at most one

differential privacy

$$\Pr[M(D_1) \in C] \leq e^{\varepsilon} \Pr[M(D_2) \in C]$$

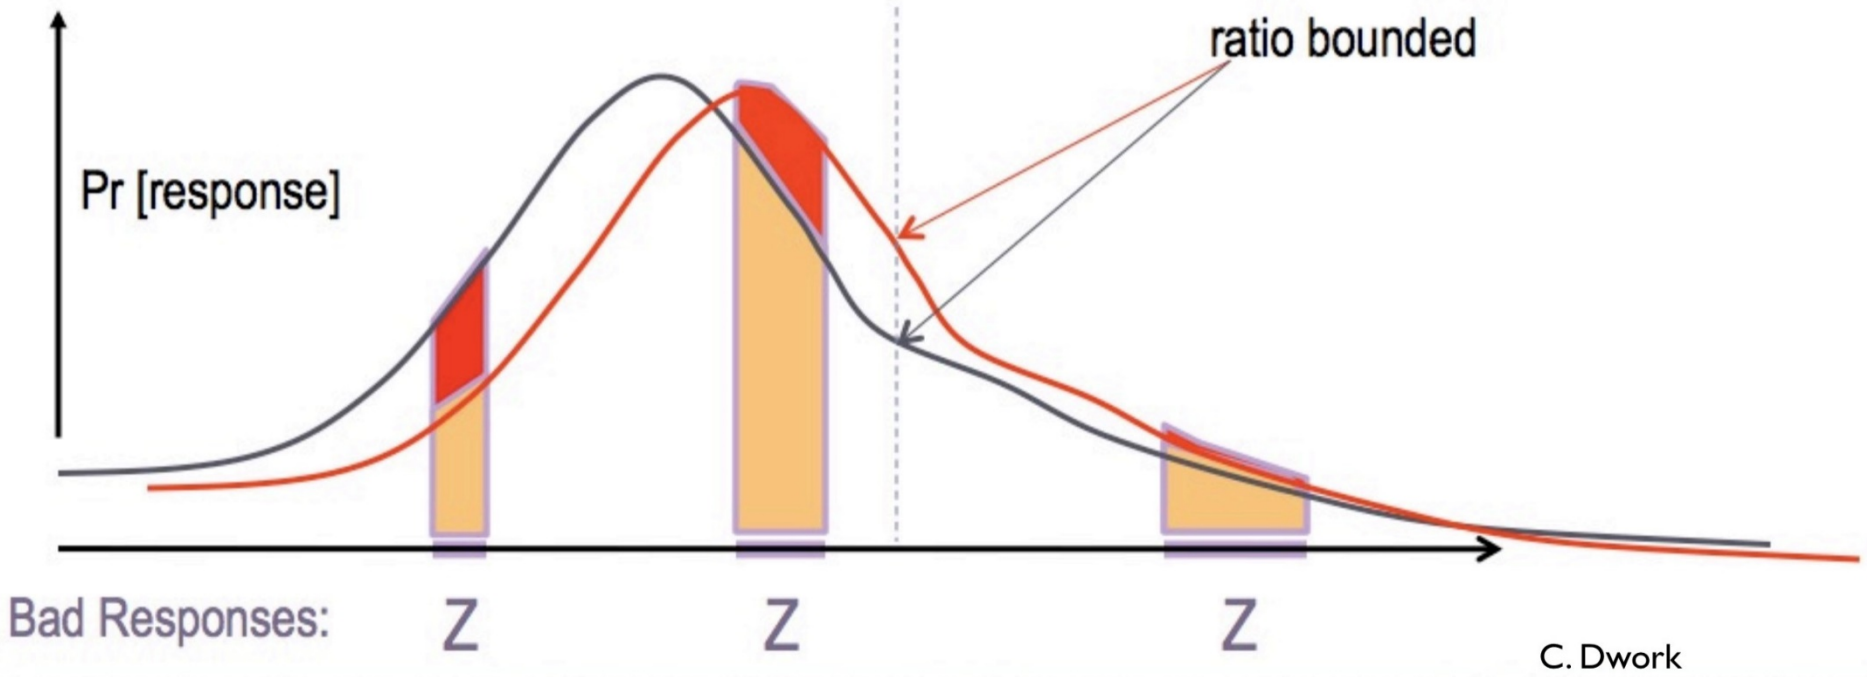

Framework of Global Sensitivity

$$GS_f = \max |f(G) - f(G')|_1 \text{ taken over neighboring } G, G'$$

$$A(G) = f(G) + Lap^d\left(\frac{GS_f}{\epsilon}\right)$$

- Many natural functions have low global sensitivity
  - e.g., histogram, mean,
  - Histogram: split input into bins, count how many in bin

The Laplace Mechanism [DMNS06]

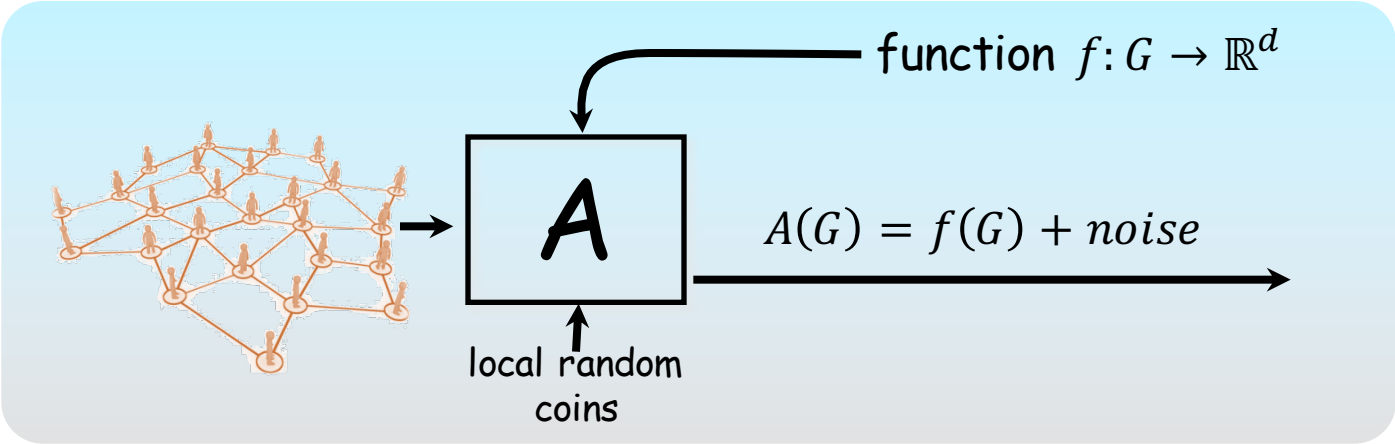

- $GS_f = \max |f(G) - f(G')|_1$ taken over neighboring G, G'
- Theorem [DMNS06]:
  - $A(G) = f(G) + Lap^d\left(\frac{GS_f}{\epsilon}\right)$ is 🎁 differentially private

The Laplace Mechanism: Proof

Density for Lap(b): $h(x) = \frac{1}{2b}\exp(-\frac{|x|}{b})$

- Neighboring inputs X and Y
- $f(X) \in R^d, f(Y) \in R^d$
- $f(X) + Lap^d(\frac{GS_f}{\epsilon})$ gives probability density function p_X
- $f(Y) + Lap^d(\frac{GS_f}{\epsilon})$ gives probability density function p_Y
- Choose any point $z \in R^d$
- $p_X(z)/p_Y(z) = \prod_{i=1}^{k}\left(\exp\left(-\frac{\epsilon|f(X)_i - z_i|}{GS(f)}\right) / \exp\left(-\frac{\epsilon|f(Y)_i - z_i|}{GS(f)}\right)\right) =$

$$\prod_{i=1}^{k}\left(\exp\left(\frac{\epsilon(|f(Y)_i - z_i| - |f(X)_i - z_i|)}{GS(f)}\right)\right)$$

The Laplace Mechanism: Proof

Density for Lap(b): $h(x) = \frac{1}{2b}\exp(-\frac{|x|}{b})$

- Choose any point $z \in R^d$
- $p_X(z)/p_Y(z) = \prod_{i=1}^{k}\left(\exp\left(-\frac{\epsilon|f(X)_i - z_i|}{GS(f)}\right) / \exp\left(-\frac{\epsilon|f(Y)_i - z_i|}{GS(f)}\right)\right) =$

$$\prod_{i=1}^{k}\left(\exp\left(\frac{\epsilon(|f(Y)_i - z_i| - |f(X)_i - z_i|)}{GS(f)}\right)\right) \leq$$

$$\prod_{i=1}^{k}\left(\exp\left(\frac{\epsilon(|f(X)_i - f(Y)_i|)}{GS(f)}\right)\right) =$$

$$\exp\left(\frac{\epsilon|f(X) - f(Y)|_1}{GS(f)}\right) \leq \exp(\epsilon)$$

Better Composition: Answering all threshold queries

- Data domain: $D = \{1, \dots, T\}$ (ordered domain with T elements)
- Database: $X \in D^n$
- Want (approx.) answers to all queries of the form: $q_t(X) = |\{i \ : \ 1 \leq x_i \leq t\}|$
- $GS(q_t)$ =1 (changing a data point in X can increase/decrease $q_t(X)$ by at most one)
- Idea: answer all T queries by adding noise $Lap(\frac{1}{\epsilon'})$ where $\epsilon' = \frac{\epsilon}{T}$
  - Using (simple) composition, this provides ϵ-differential privacy
  - Problem: noise magnitude linear in T; can we do better?

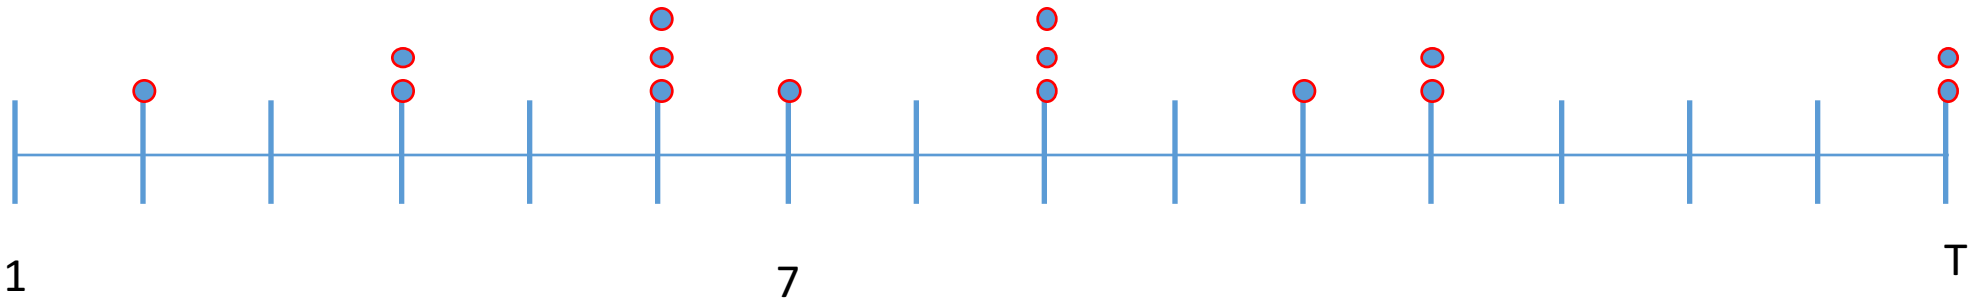

$n = 15$

$q_7(d) = \frac{7}{15}$

Answering all threshold queries

idea: compute $\log T$ histograms

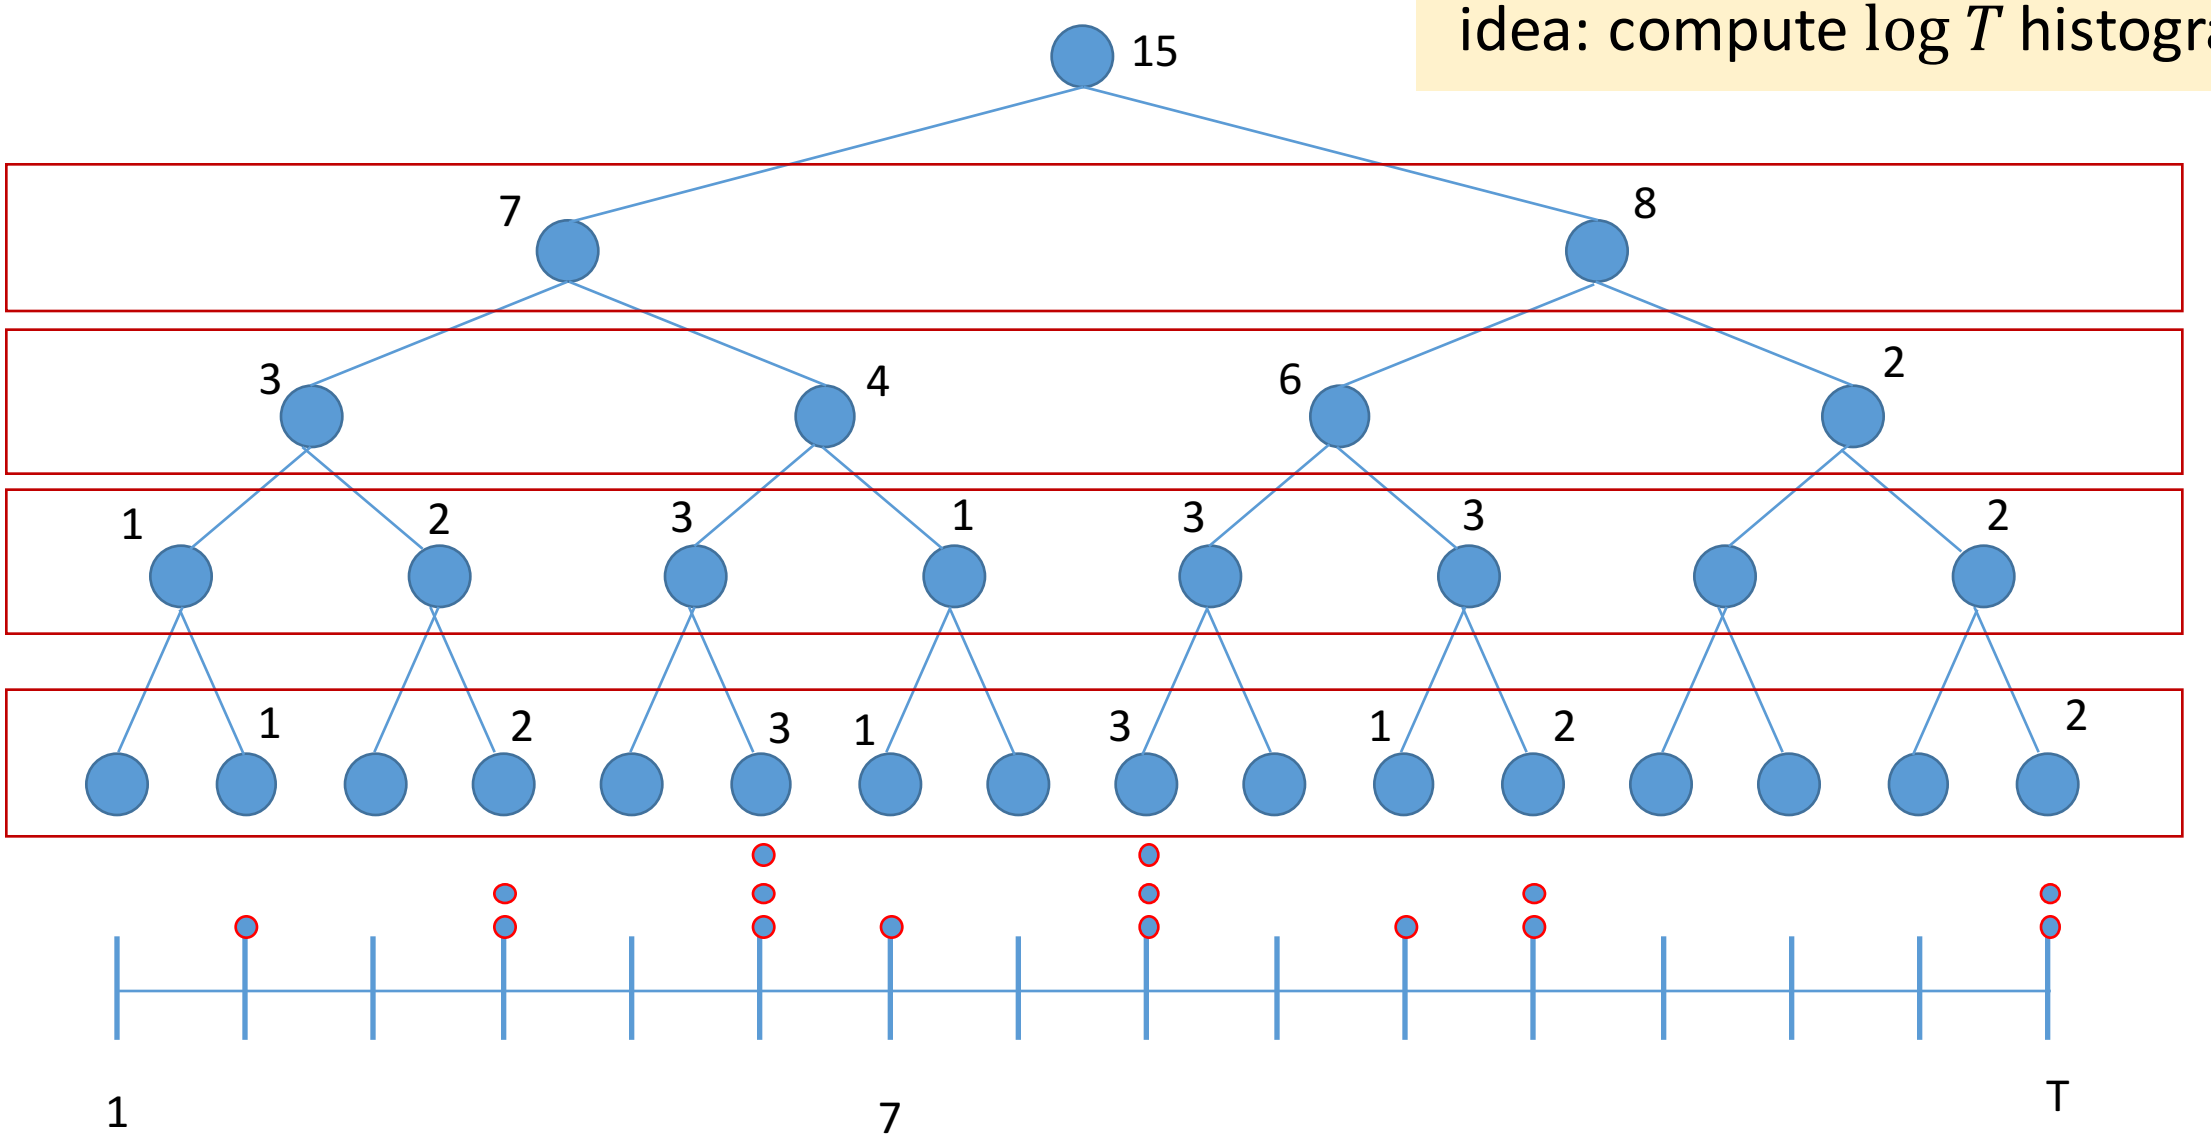

Answering all threshold queries

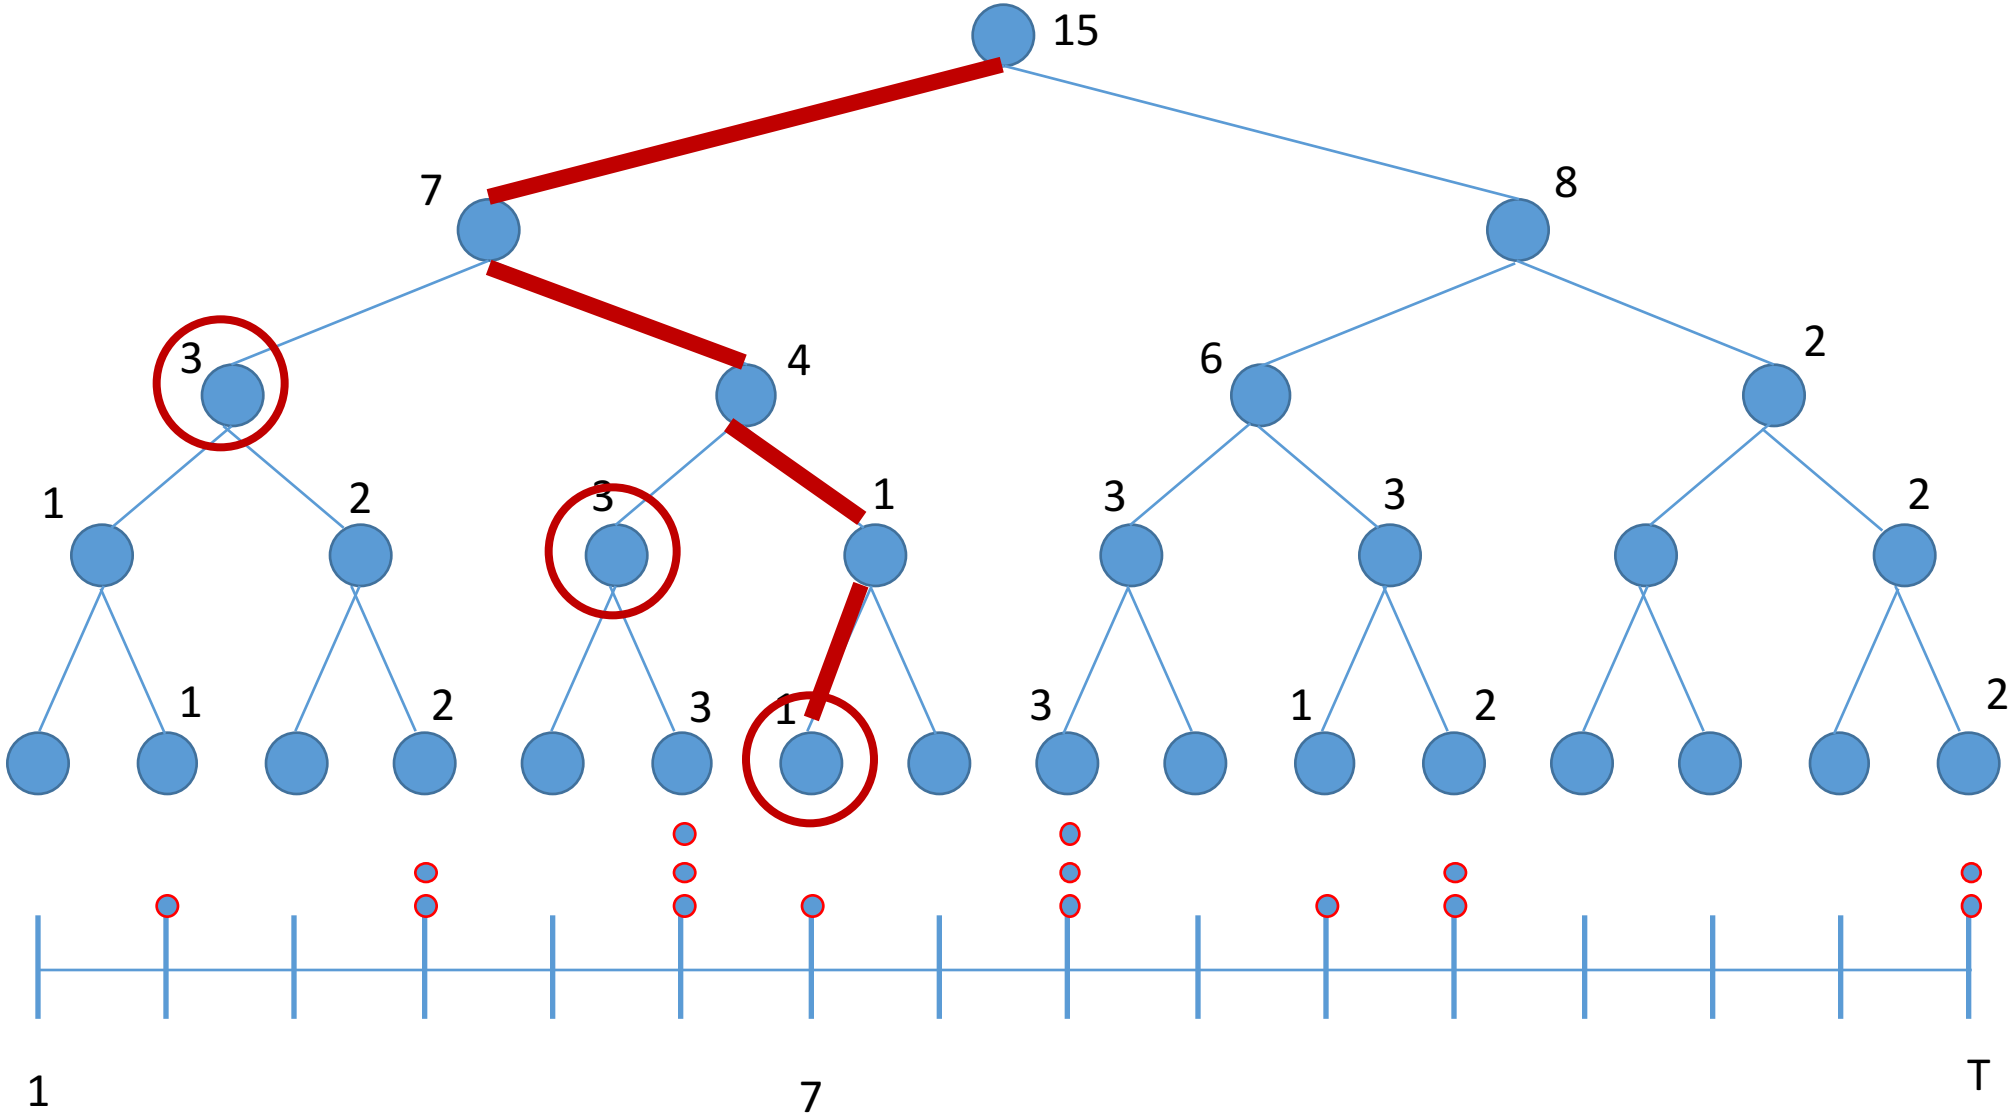

$n = 15$

$q_7(D) = \frac{7}{15}$

Answering all threshold queries

- What we get (using basic composition):
  - Computing $\log T$ histograms, each with $\epsilon' = \frac{\epsilon}{\log T}$
    - E.g., add noise $\mathrm{Lap}(2\epsilon/\log T)$ to each count
      - Noise variance $\sim \left(\frac{\log T}{\epsilon}\right)^2$
  - Each answer to threshold query is sum of (at most) $\log T$ noisy estimates
  - Overall noise variance $\sim \log T \left(\frac{\log T}{\epsilon}\right)^2$
    - Whp noise magnitude = $\frac{polylog(T)}{\epsilon}$

Edge vs. Node Privacy - Counting Edges

$$GS_f = \max |f(G) - f(G')|_1 \text{ taken over neighboring } G, G'$$

$$A(G) = f(G) + Lap^d\left(\frac{GS_f}{\epsilon}\right)$$

- Counting edges: $f(G) = \sum e_{ij}$ where $e_{ij} \in \{0,1\}$
- Edge privacy: $GS_f = 1$, noise $\sim \frac{1}{\epsilon}$
- Node privacy: $GS_f = n$, noise $\sim \frac{n}{\epsilon}$

Exponential Sampling [MT07]

- x_i= {books read by i this year}, Y = {book names}
- "Score" of $y \in Y$: $q(y, x) = \#\{i: y \in x_i\}$
- Goal: output book read by most

- Mechanism: given x, output book name y with probability prop to
  - $\exp(\epsilon \cdot q(y, x))$

- Claim: Mechanism is ε-differentially private

- Claim: If most popular website has score $T = \max_{y \in Y} q(y, x)$, then $E[q(y_0, x)] \geq T - O(\frac{\log|Y|}{\epsilon})$

Exponential Sampling

- Database X, function values in R, utility function from database and function value to the reals
- X=who reads what
- R=names of books
- u(X,r) – the closer to the real max the better
- $\Delta u = \max_{r \in R} \max_{X,Y:|X-Y|_1 \leq 1} |u(X,r) - u(Y,r)|$
- Output each possible $r \in R$ with probability proportional to
  - $\exp(\epsilon u(X,r)/\Delta u)$
- $\exp\left(\frac{\epsilon u(X,r)}{\Delta u}\right) / \exp\left(\frac{\epsilon u(Y,r)}{\Delta u}\right) = \exp\left(\frac{\epsilon(u(X,r)-u(Y,r))}{\Delta u}\right)$

Wrong!!

The correct proof

$$
\begin{aligned}
\frac{\Pr[\mathcal{M}_E(x,u,\mathcal{R})=r]}{\Pr[\mathcal{M}_E(y,u,\mathcal{R})=r]} &= \frac{\left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)}{\left(\frac{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}\right)} \\
&= \left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}\right)\cdot\left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right) \\
&= \exp\left(\frac{\varepsilon(u(x,r')-u(y,r'))}{2\Delta u}\right) \\
&\quad\cdot\left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right) \\
&\leq \exp\left(\frac{\varepsilon}{2}\right)\cdot\exp\left(\frac{\varepsilon}{2}\right)\cdot\left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right) \\
&= \exp(\varepsilon).
\end{aligned}
$$

Exponential mechanism gives good utility

- Prob(utility of exp mechanism $\leq Opt\ utility\ - \left(\frac{2\Delta u}{\epsilon}\right)(\ln(|R|) + t)$ $\leq \exp(-t)$

Next week: Uri Stemmer

- Will talk about practical local differential privacy
- Local? “users retain their data and only send the server randomizations that are safe for publication.
- The random response could be viewed as local differential privacy
- Talk will be in the context of heavy hitters : items that appear many times.