


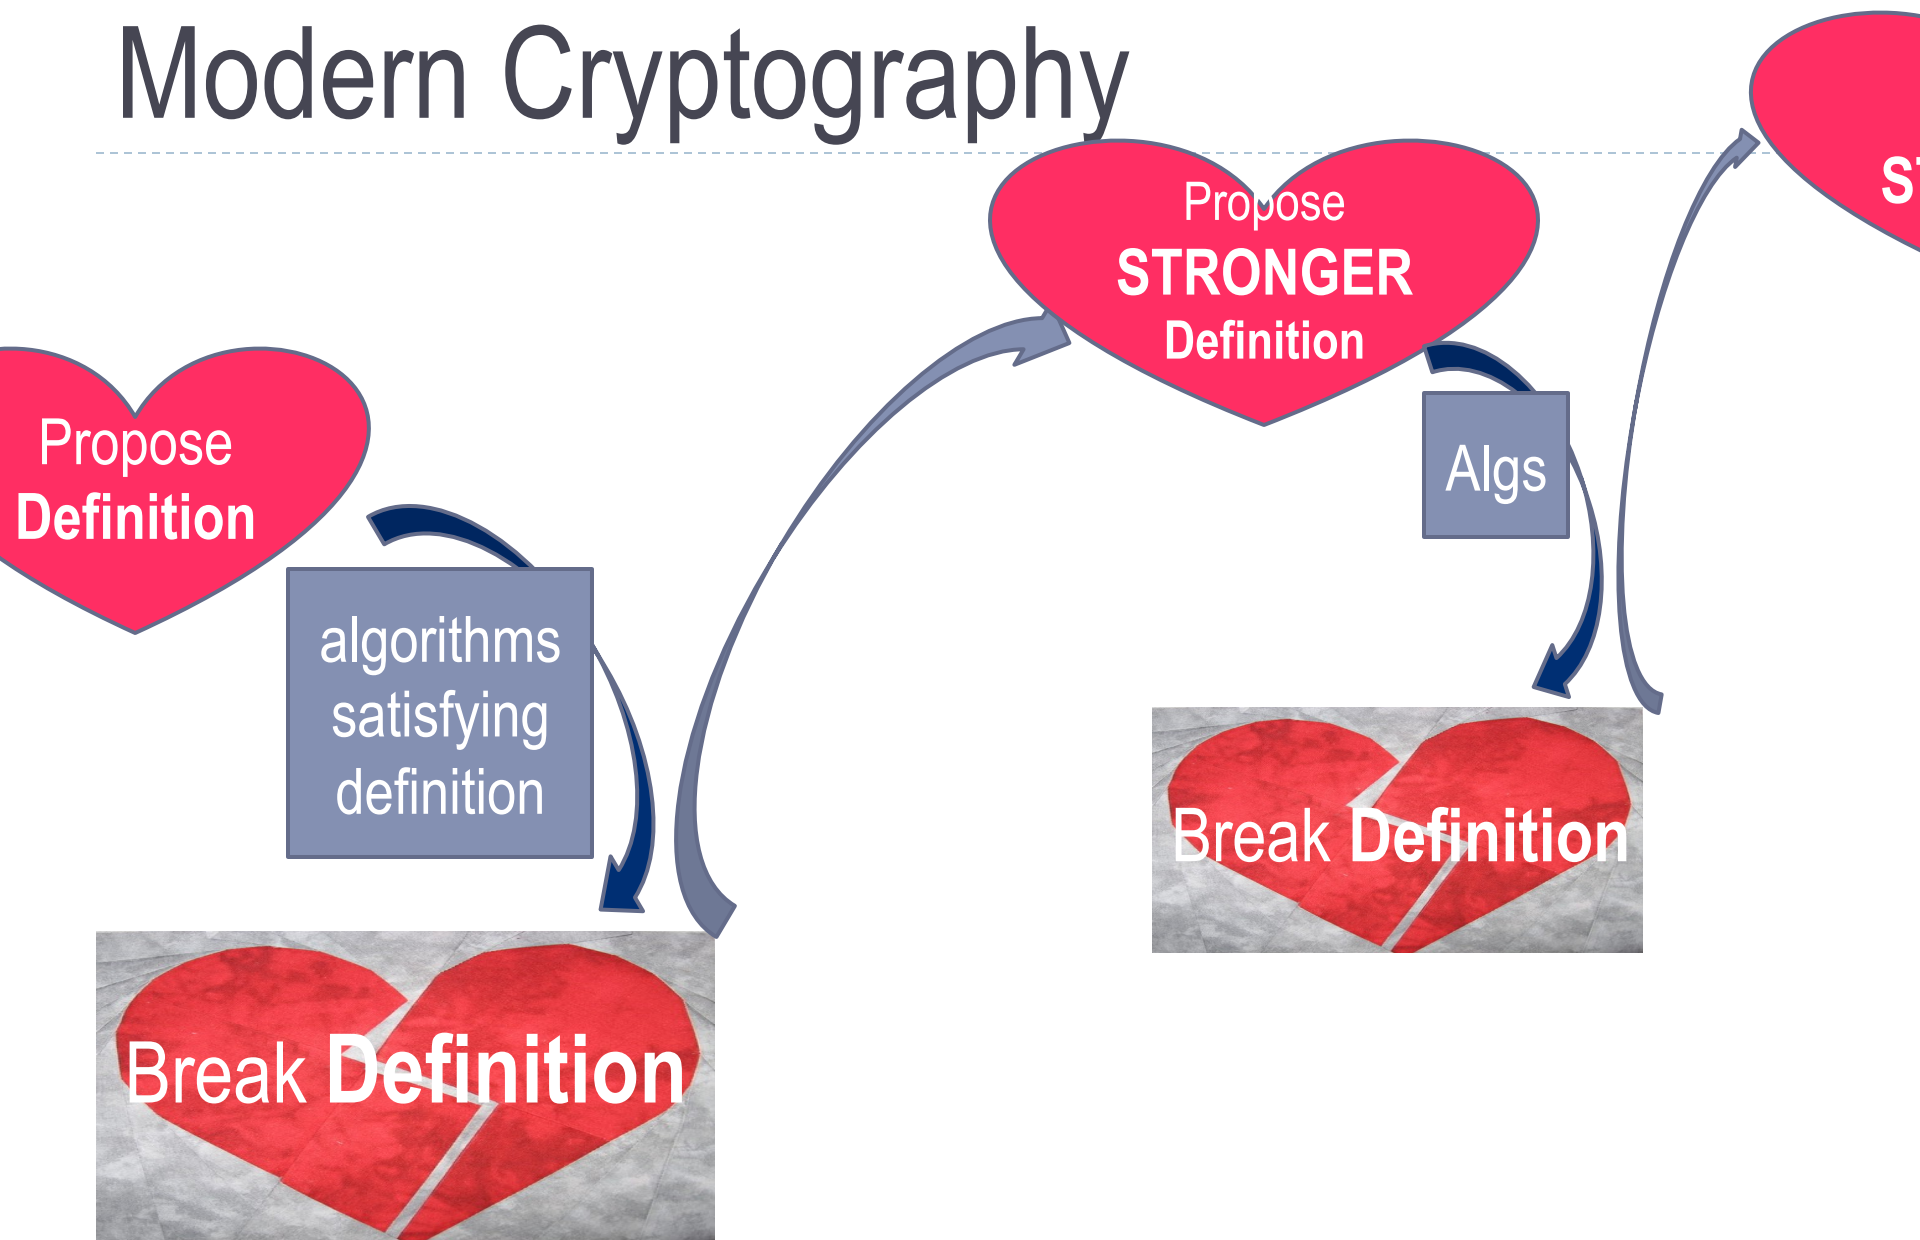
The State of the Art

Cynthia Dwork, Microsoft Research

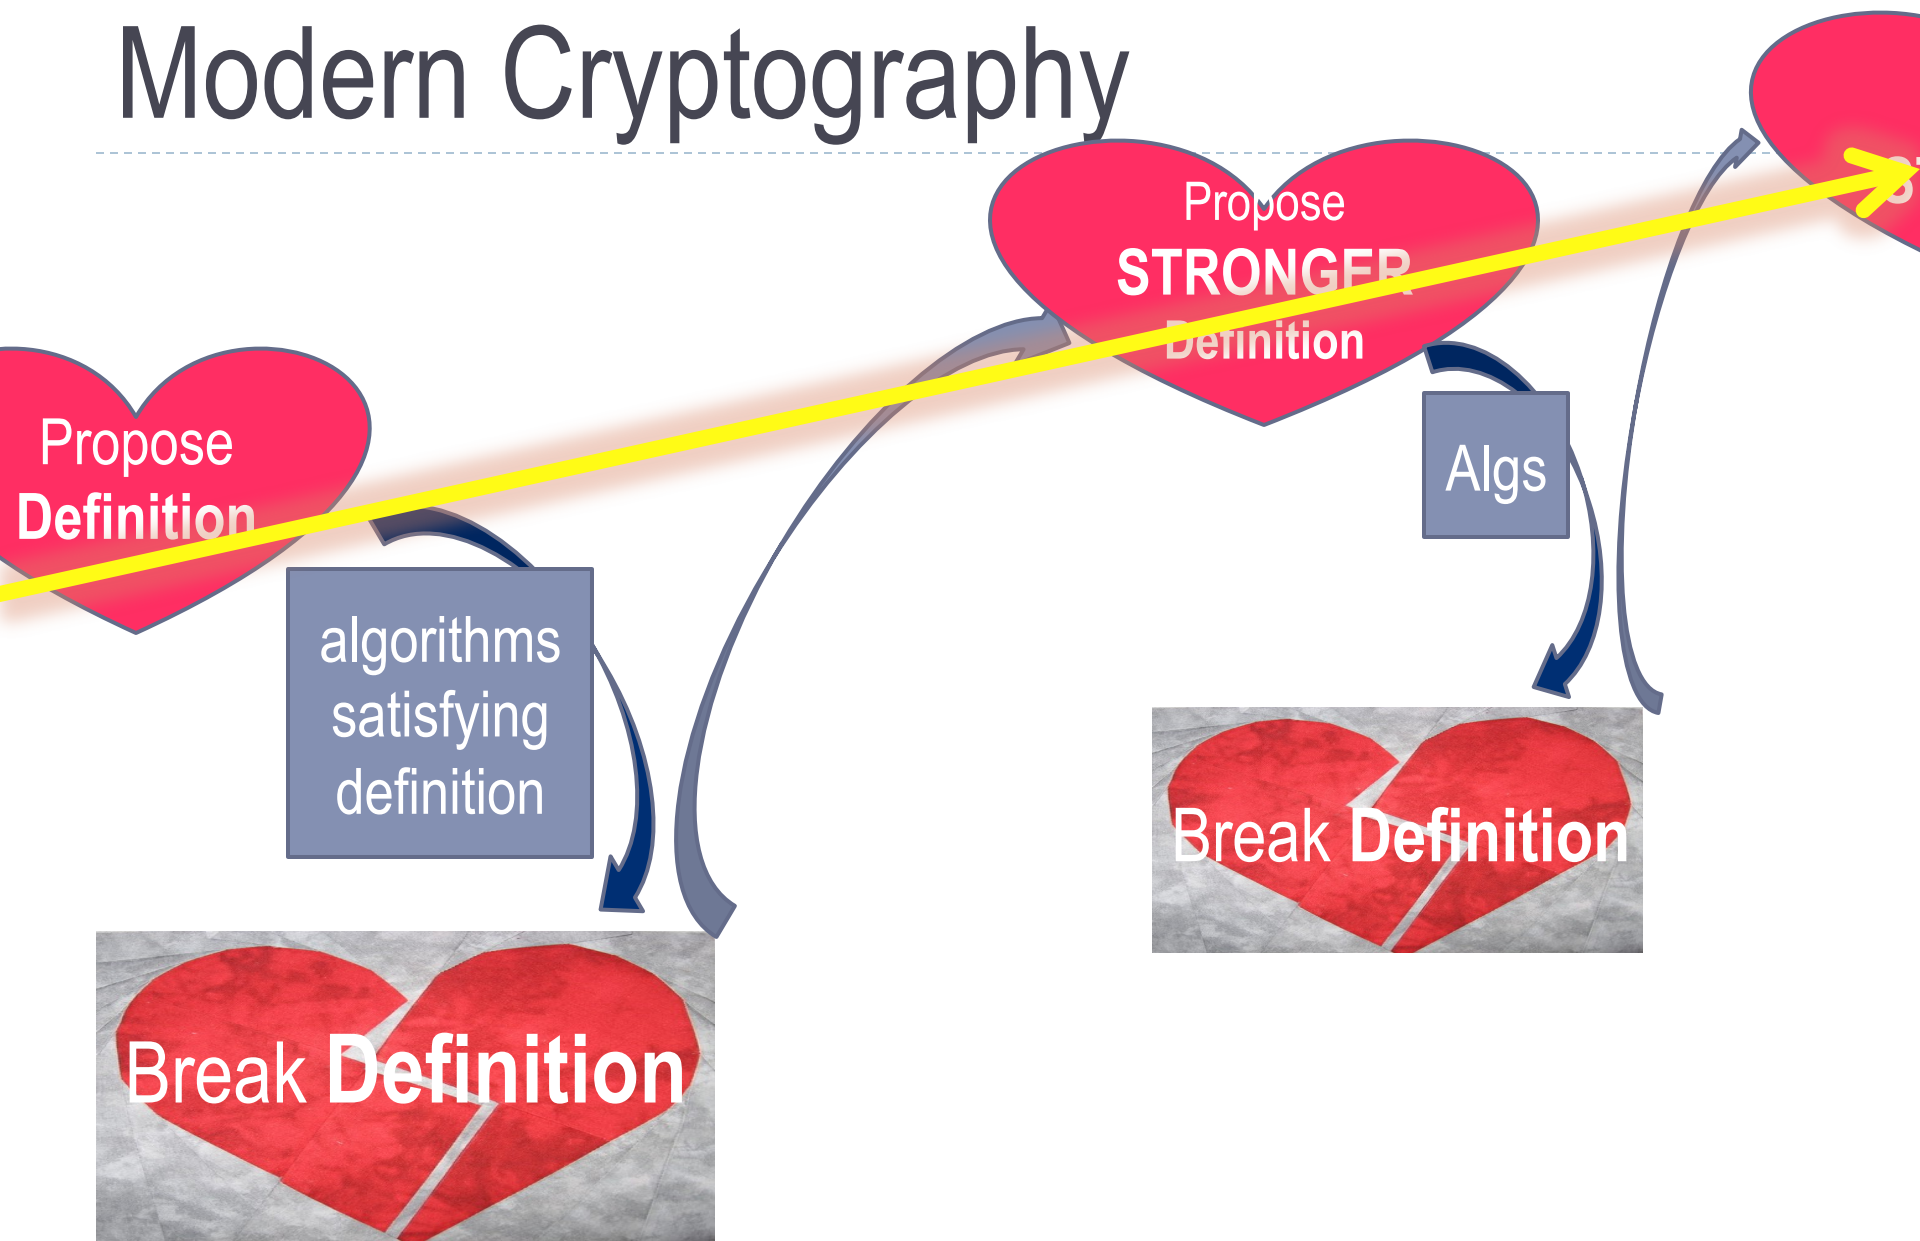
Pre-Modern Cryptography



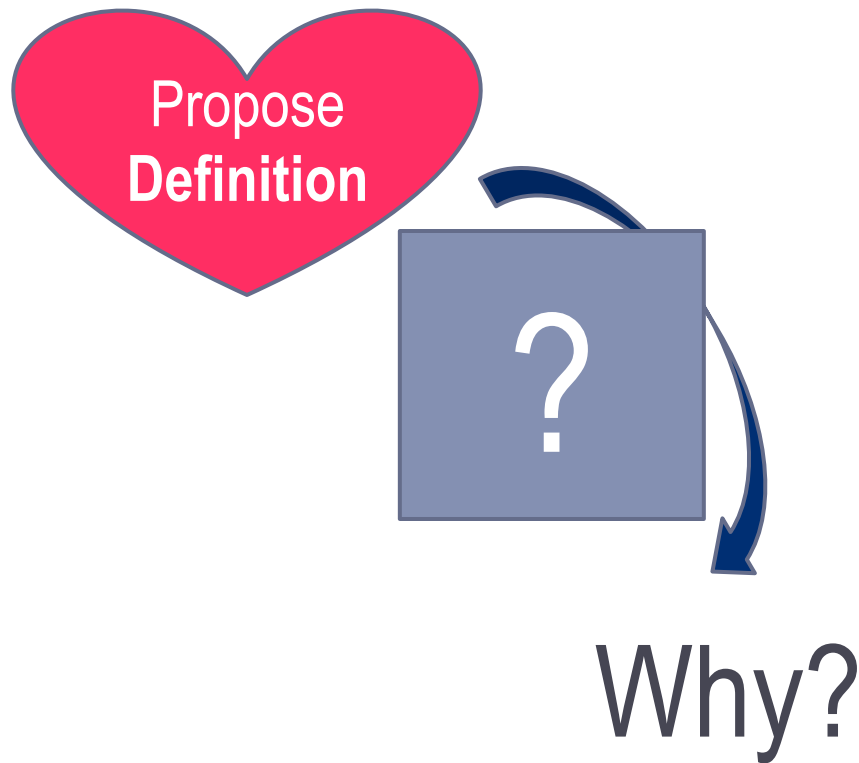
Modern Cryptography



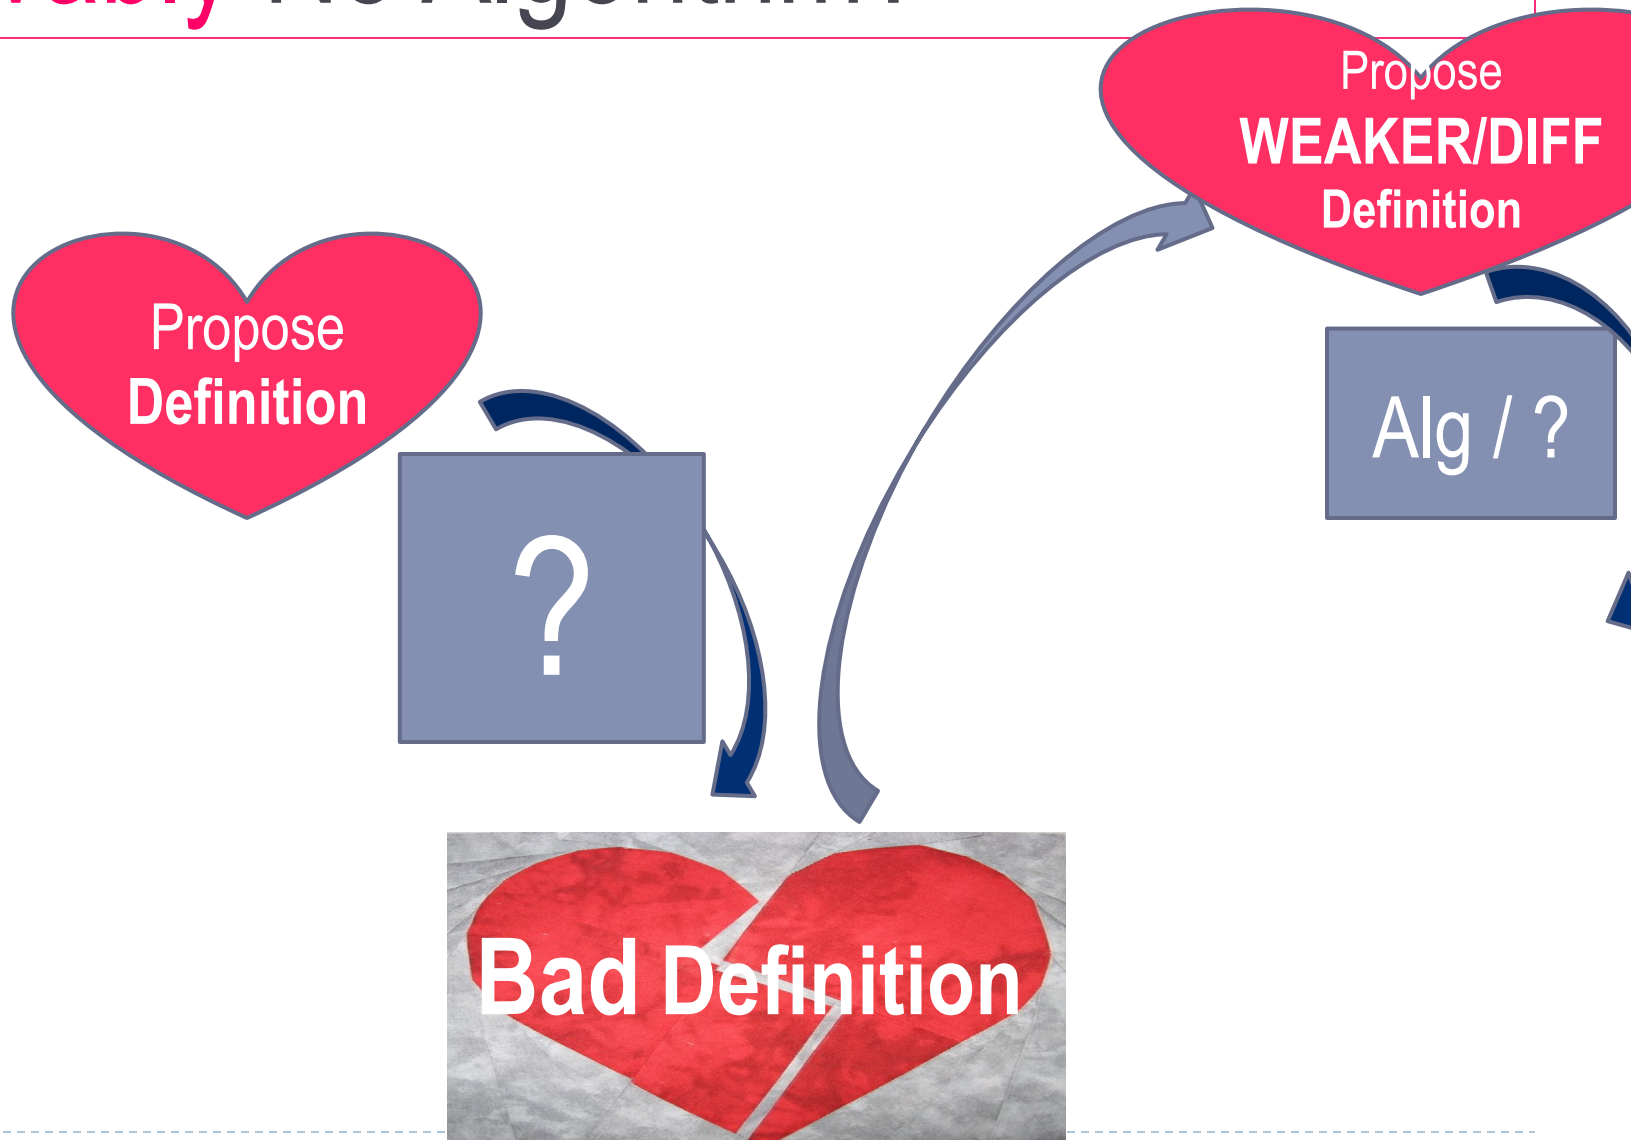
Modern Cryptography




No Algorithm?



Provably No Algorithm?




The Privacy Dream



- ▶ Census, financial, medical data; OTC drug purchases; social networks; MOOCs data; call and text records; energy consumption; loan, advertising, and applicant data; ad clicks product correlations, query logs,...




Fundamental Law of Info Recovery



“Overly accurate” estimates of “too many” statistics is blatantly non-private

Anonymization (aka De-Identification)

- ▶ Remove “personally identifying information”




Name sex DOB zip symptoms previous admissions medications family history

...



Anonymization (aka De-Identification)

- ▶ Remove “personally identifying information”




Name sex DOB zip symptoms previous admissions medications family history

...



Anonymization (aka De-Identification)

- ▶ Remove “personally identifying information”




Name sex DOB zip symptoms previous admissions medications family history

...



Anonymization (aka De-Identification)

- ▶ Remove “personally identifying information”



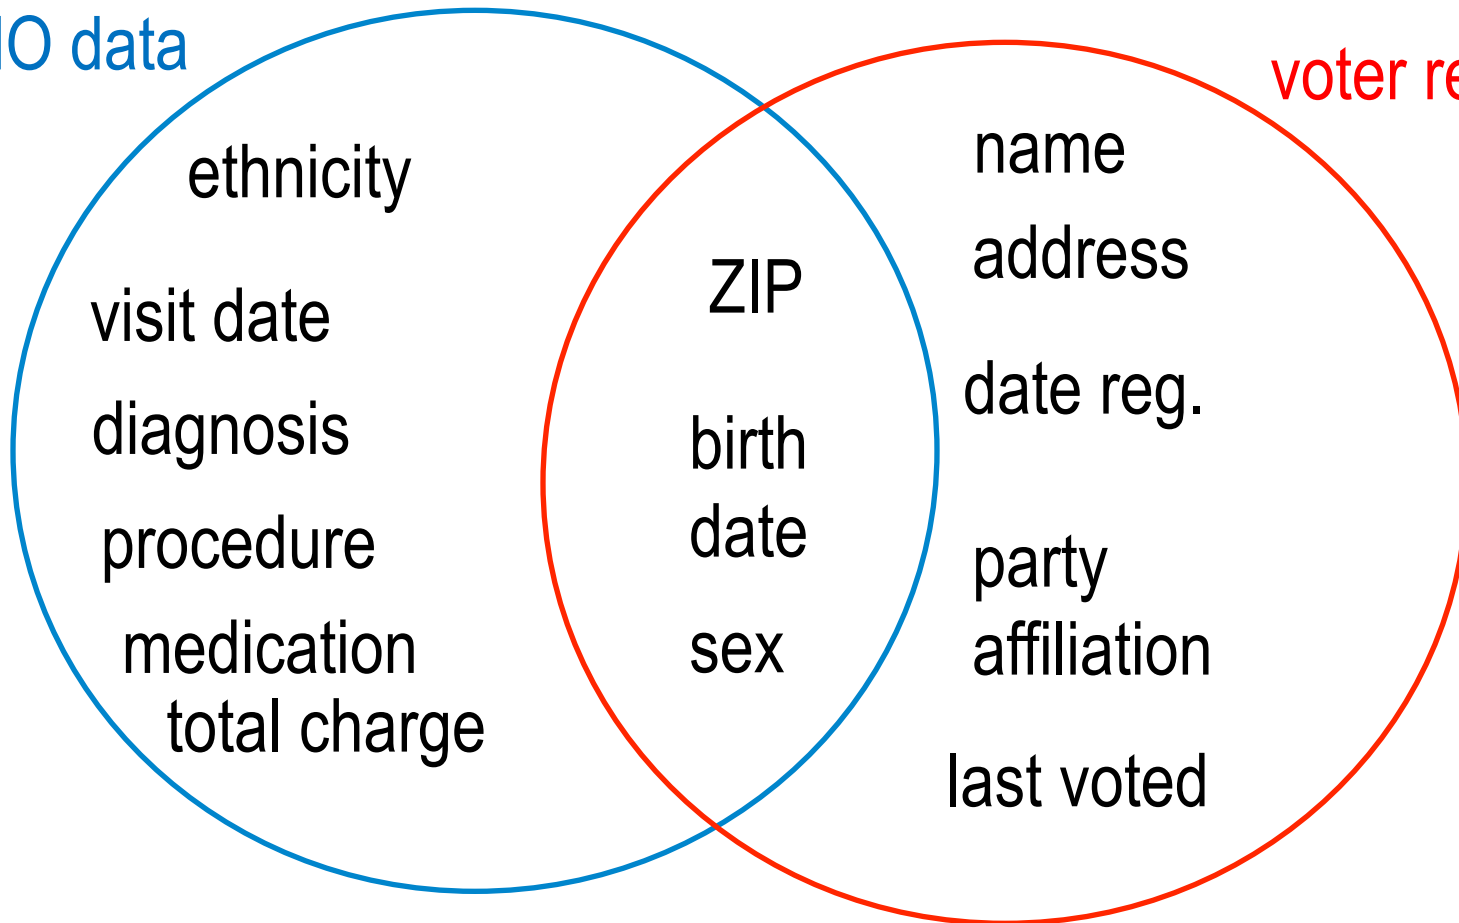
Name **sex DOB zip** symptoms previous admissions medications family history

...



William Weld's Medical Records

HMO data

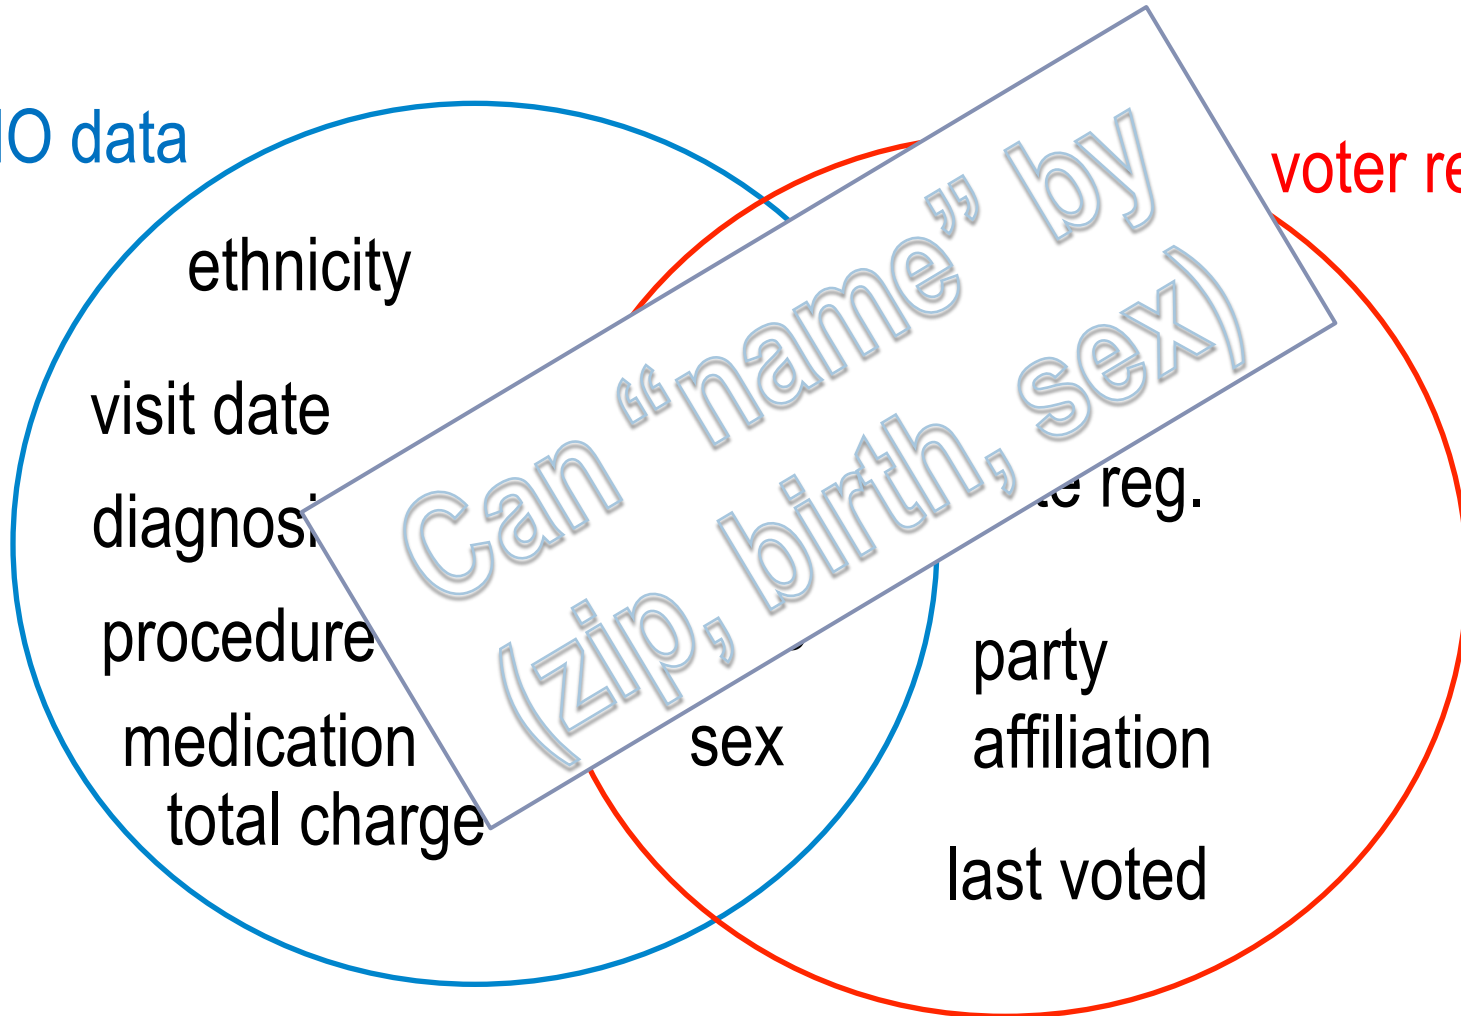


voter registration
data

William Weld's Medical Records


HMO data

voter registration
data



Anonymization (aka De-Identification)

- ▶ Remove “personally identifying information”




Name sex DOB **zip** symptoms **previous admissions** medications family history

...



Anonymization (aka De-Identification)

- ▶ Remove “personally identifying information”



Name sex DOB zip symptoms **previous admissions** medications **family history**

...



NETFLIX

Netflix Prize

Home Rules Leaderboard Register Update Submit Download

NETFLIX

Browse Recommendations Friends Queue Buy DVDs
Home Genres New Releases Previews Netflix Top 100 Crit

Movies For You

Randy, the following movies were chosen based on your interest in:
[Bowling for Columbine](#)
[Carnivale: Season 1](#)
[Fahrenheit 9/11](#)

The Big One

★★★★☆

...er subversive
...ly from
...n /
...ichael

Carnivale: Season 2

Disc Series

★★★★☆

Daniel Kraus
...rivingly cre
...series contr
...document t

Carnivale: Season 1

Disc Series

★★★★☆

Daniel Kraus
...rivingly cre
...series contr
...document t

Recover

★★★★☆

...entures of a motley cre
...nies who've made the O
...stbowl their ... [Read Mo](#)

All Discs Guaranteed

You really liked it...

Now own it for just \$5.99

Shop as low

titles

Original art

Welcome!

The Netflix Prize seeks to substantially improve the accuracy of predictions about how much someone is going to love a movie based on their movie preferences. Improve it enough and you win one (or more) Prizes. Winning the Netflix Prize improves our ability to connect people to the movies they love.

Read the [Rules](#) to see what is required to win the Prizes. If you are interested in joining the quest, you should [register a team](#).

You should also read the [frequently-asked questions](#) about the Prize. And check out how various teams are doing on the [Leaderboard](#).

Good luck and thanks for helping!

NETFLIX

Netflix Prize

Home Rules Leaderboard Register

Can "name" by 3
(title, approx date)
pairs

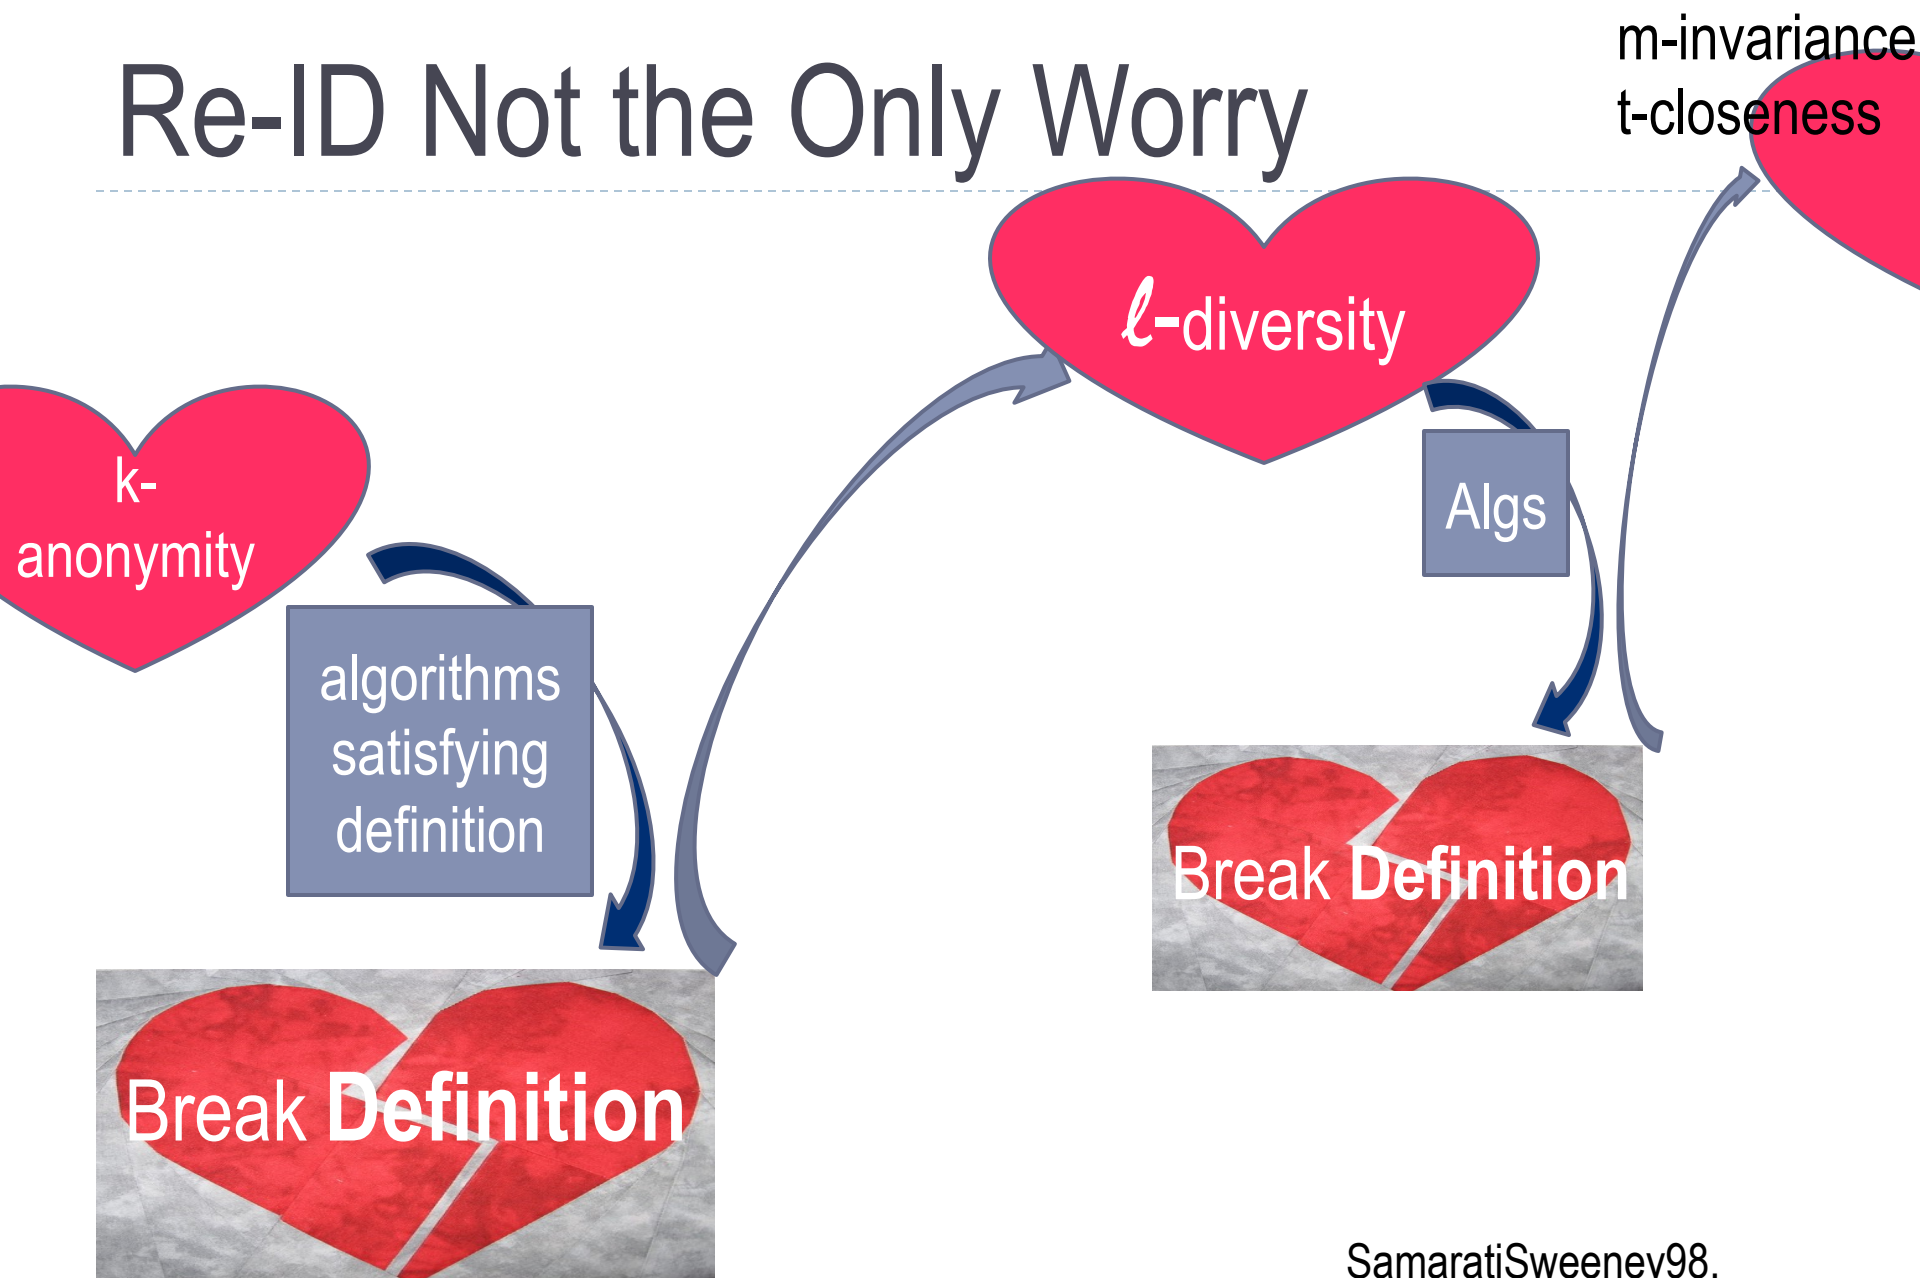
The Netflix Prize seeks to substantially improve the accuracy of predictions about how much someone is going to love a movie based on their movie preferences. Improve it enough and you win one (or more) Prizes. Winning the Netflix Prize improves our ability to connect people to the movies they love.

Read the [Rules](#) to see what is required to win the Prizes. If you are interested in joining the quest, you should [register a team](#).

You should also read the [frequently-asked questions](#) about the Prize. And check out how various teams are doing on the [Leaderboard](#).

Good luck and thanks for helping!





Re-ID Not the Only Worry



SamaratiSweeney98,

▶ MakanvajhalaGehrkeKiferVenkitasubramaniam06,XiaoTao07,LiLiVenkatasubmraminan07

Culprit: Diverse Background Info



Billing for Targeted Advertisements



+




+




Product Recommendations

- ▶ X's preferences influence Y's experience
 - ▶ Combining evolving *similar items lists* with a little knowledge (from your blog) of what you bought, an adversary can infer purchases you did not choose to publicize




People who bought this also bought...

GWAS Statistics




SNP: Single Nucleotide (A,C,G,T) polymorphism





SNP statistics of Case Group

“Can” Test Case Group Membership
using target’s DNA and HapMap

Culprit: Diverse Background Info



Voter Weld: M, DOB, zip




Billing

Blog


People who bought this...



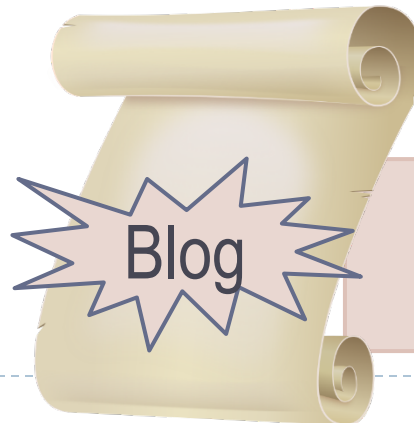


Culprit: Diverse Background Info



Voter Weld: M, DOB, zip



HapMap



People who bought this...

▶ Billing

How Should We Approach Privacy?

- ▶ “Computer science got us into this mess, can computer science get us out of it?” (Sweeney, 2012)



How Should We Approach Privacy?

- ▶ “Computer science got us into this mess, can computer science get us out of it?” (Sweeney, 2012)
- ▶ Complexity of this type requires a mathematically rigorous theory of privacy and its loss.



How Should We Approach Privacy?

- ▶ “Computer science got us into this mess, can computer science get us out of it?” (Sweeney, 2012)
- ▶ Complexity of this type requires a mathematically rigorous theory of privacy and its loss.
 - ▶ We cannot discuss tradeoffs between privacy and statistical utility without a measure that captures cumulative harm over multiple uses.
 - ▶ Other fields -- economics, ethics, policy -- cannot be brought to bear without a “currency,” or measure of privacy, with which to work.



Useful Databases that Teach

- ▶ Database teaches that smoking causes cancer.
 - ▶ Smoker S's insurance premiums rise.
 - ▶ **Premiums rise even if S not in database!**
- ▶ Learning that smoking causes cancer is the whole point.
 - ▶ Smoker S enrolls in a smoking cessation program.
- ▶ **Differential privacy: limit harms to the teachings, not participation**

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



Useful Databases that Teach

- ▶ Database teaches that smoking causes cancer.
 - ▶ Smoker S's insurance premiums rise.
 - ▶ **Premiums rise even if S not in database!**
- ▶ Learning that smoking causes cancer is the whole point.
 - ▶ Smoker S enrolls in a smoking cessation program.
- ▶ **Differential privacy: limit harms to the teachings, not participation**

The likelihood of any possible harm to ME is essentially independent of whether I join, or refrain from joining, the database.



Useful Databases that Teach

- ▶ Database teaches that smoking causes cancer.
 - ▶ Smoker S's insurance premiums rise.
 - ▶ **Premiums rise even if S not in database!**
- ▶ Learning that smoking causes cancer is the whole point.
 - ▶ Smoker S enrolls in a smoking cessation program.
- ▶ **Differential privacy: limit harms to the teachings, not participation**

High premiums, busted, purchases revealed to co-worker...
Essentially equally likely when I'm in as when I'm out



Differential Privacy [D., McSherry, Nissim, Smith '06]

M gives ϵ -differential privacy if for all pairs of data sets D, D' differing in one element, and all subsets \mathcal{C} of possible outputs

$$\Pr[M(D) \in \mathcal{C}] \leq e^{\epsilon} \Pr[M(D') \in \mathcal{C}]$$

Randomness introduced by M



Differential Privacy [D., McSherry, Nissim, Smith '06]

M gives ϵ -differential privacy if for all pairs of data sets D, D' differing in one element, and all subsets \mathcal{C} of possible outputs

$$\Pr[M(D) \in \mathcal{C}] \leq (1 + \epsilon) \Pr[M(D') \in \mathcal{C}]$$



Randomness introduced by M

If a bad event is very unlikely when I'm not in dataset (D')
then it is still very unlikely when I am (D)



Differential Privacy [D., McSherry, Nissim, Smith '06]

\mathcal{C} gives ϵ -differential privacy if for all pairs of data sets D, D' differing in one element, and all subsets \mathcal{S} of possible outputs

$$\Pr[\mathcal{C}(D) \in \mathcal{S}] \leq (1 + \epsilon) \Pr[\mathcal{C}(D') \in \mathcal{S}]$$



Randomness introduced by M

If a bad event is very unlikely when I'm not in dataset (D')
then it is still very unlikely when I am (D)

Impossible to know the actual probabilities of bad events.
Can still control change in risk due to joining the database.



Differential Privacy [D., McSherry, Nissim, Smith '06]

\mathcal{C} gives ϵ -differential privacy if for all pairs of data sets D, D' differing in one element, and all subsets \mathcal{S} of possible outputs

$$\Pr[\mathcal{C}(D) \in \mathcal{S}] \leq (1 + \epsilon) \Pr[\mathcal{C}(D') \in \mathcal{S}]$$

“Privacy Loss”

If a bad event is very unlikely when I'm not in dataset (D')
then it is still very unlikely when I am (D)

Impossible to know the actual probabilities of bad events.
Can still control change in risk due to joining the database.



Differential Privacy [D., McSherry, Nissim, Smith '06]

C gives ϵ -differential privacy if for all pairs of data sets D, D' differing in one element, and all subsets S of possible outputs

$$\Pr[C(D) \in S] \leq (1 + \epsilon) \Pr[C(D') \in S]$$

“Privacy Loss”

If a bad event is very unlikely when I'm not in dataset (D')
then it is still very unlikely when I am (D)




Differential Privacy

- ▶ **Nuanced measure of privacy loss**
 - ▶ Captures cumulative harm over multiple uses, multiple databases
- ▶ **Adversary's background knowledge is irrelevant**
 - ▶ Immune to re-identification attacks, etc.
- ▶ **“Programmable”**
 - ▶ Construct complicated private analyses from simple private building blocks

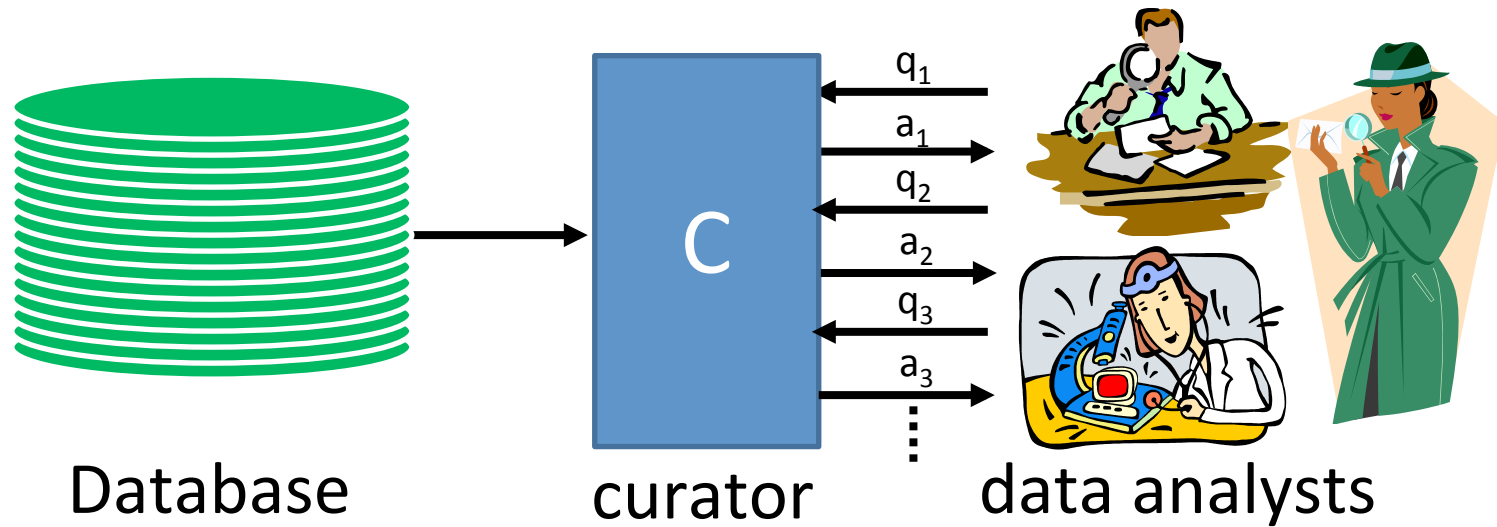


Recall: Fundamental Law



“Overly accurate” estimates of “too many” statistics is blatantly non-private


Answer Only Questions Asked




Intuition

- ▶ Want to compute $f(D)$
- ▶ Adding \bullet pulls $f(D^\dagger)$
 - ▶ Add random noise to obscure difference

$f(D)$ vs $f(D^\dagger)$




Intuition

- ▶ Want to compute $f(D)$
- ▶ Adding  pulls $f(D^\dagger)$
 - ▶ Add random noise to obscure difference

$$|f(D) - f(D^\dagger)| / \epsilon$$

*Algorithms, geometry,
learning theory,
complexity theory,
cryptography, statistics,
machine learning,
programming languages,
verification, databases,
economics, ...*



Not a Panacea

- ▶ Fundamental Law of Information Recovery still holds



Challenge: The Meaning of Loss

- ▶ Sometimes the theory gives exactly the right answer
 - ▶ Large loss in differential privacy translates to “obvious” real life privacy breach, under circumstances known to be plausible
- ▶ Other times?
 - ▶ Do all large losses translate to such realizable privacy breaches, or is the theory too pessimistic?




Policy Recommendation

- ▶ **Publish all Epsilons!**
 - ▶ Penalize when $\epsilon = \infty$

Combines motivation for data breach notification statutes and environmental laws requiring disclosures of toxic releases with an incentive to start using (minimal) differential privacy

“Just a Few”?




Randomly choose a few rows;
Publish in entirety.

OK?



Fundamental Law

- ▶ There is a (LARGE) set of statistics, S
 - ▶ An analyst having even a **remotely** accurate estimate of EVERY statistic in S can completely violate any reasonable notion of privacy



Must be very inaccurate on some statistic in S

“Overly accurate” estimates of “too many” statistics is blatantly non-private

- ▶ There is a very simple way of designing small sets of statistics, T
 - ▶ An analyst having estimates of 80% of the statistics in T that beat the sampling error can completely violate any reasonable notion of privacy