

Responsible Data Science

Anonymity and privacy

April 2 & 9, 2024

Prof. Julia Stoyanovich

Center for Data Science &
Computer Science and Engineering
New York University

Reading for weeks 10 & 11

Robust De-anonymization of Large Sparse Datasets

Arvind Narayanan and Vitaly Shmatikov

The University of Texas at Austin

Abstract

We present a new class of statistical de-anonymization attacks against high-dimensional micro-data, such as individual preferences, recommendations, transaction records and so on. Our techniques are robust to perturbation in the data and tolerate some mistakes in the adversary's background knowledge.

We apply our de-anonymization methodology to the Netflix Prize dataset, which contains anonymous movie ratings of 300,000 subscribers of Netflix, the world's largest online movie rental service. We demonstrate that an adversary who knows only a little bit about an individual subscriber can easily identify that subscriber's record in the dataset. Using the Internet Movie Database as the source of background knowledge, we successfully identified the Netflix records of known users, uncovering their apparent political preferences and other potentially sensitive information.

1 Introduction

Datasets containing micro-data, that is, information about specific individuals, are increasingly becoming public in response to "open government" laws and to support data mining research. Some datasets include legally protected information such as health histories; others contain individual preferences and transactions, which many people may view as private or sensitive.

Privacy risks of publishing micro-data are well-known. Even if identifiers such as names and Social Security numbers have been removed, the adversary can use background knowledge and cross-correlation with other databases to re-identify individual data records. Famous attacks include de-anonymization of a Massachusetts hospital discharge database by joining it with a public votes database [25] and privacy breaches caused by (ostensibly anonymized) AOL search data [16].

Micro-data are characterized by high dimensionality

and sparsity. Each record contains many attributes (i.e., columns in a database schema), which can be viewed as dimensions. Sparsity means that for the average record, there are no "similar" records in the multi-dimensional space defined by the attributes. This sparsity is empirically well-established [7, 4, 19] and related to the "fat tail" phenomenon: individual transaction and preference records tend to include statistically rare attributes.

Our contributions. Our first contribution is a formal model for privacy breaches in anonymized micro-data (section 2). We present two definitions, one based on the probability of successful de-anonymization, the other on the amount of information recovered about the target. Unlike previous work [25], we do not assume a priori that the adversary's knowledge is limited to a fixed set of "quasi-identifiers" attributes. Our model thus encompasses a much broader class of de-anonymization attacks than simple cross-database correlation.

Our second contribution is a very general class of de-anonymization algorithms, demonstrating the fundamental limits of privacy in public micro-data (section 4). Under very mild assumptions about the distribution from which the records are drawn, the adversary with a small amount of background knowledge about an individual can use it to identify, with high probability, this individual's record in the anonymized dataset and to learn all anonymously released information about him or her, including sensitive attributes. For sparser datasets, such as most real-world datasets of individual transactions, preferences, and recommendations, very little background knowledge is needed (as few as 5-10 attributes in our case study). Our de-anonymization algorithm is robust to the imprecision of the adversary's background knowledge and to perturbation that may have been applied to the data prior to release. It works even if only a subset of the original dataset has been published.

Our third contribution is a practical analysis of the Netflix Prize dataset, containing anonymized movie ratings of 300,000 Netflix subscribers (section 5). Netflix—the world's largest online DVD rental

COMMUNICATIONS OF THE ACM

What does it mean to preserve privacy?

BY CYNTHIA DWORK

A Firm Foundation for Private Data Analysis

IN THE INFORMATION realm, loss of privacy is usually associated with failure to control access to information, to control the flow of information, or to control the purposes for which information is employed. Differential privacy arose in a context in which ensuring privacy is a challenge even if all these control problems are solved: privacy-preserving statistical analysis of data.

The problem of *statistical disclosure control*—revealing accurate statistics about a set of respondents while preserving the privacy of individuals—has a venerable history, with an extensive literature spanning statistics, theoretical computer science, security, databases, and cryptography (see, for example, the excellent survey of Adam and Wortmann,¹ the discussion of related work in Blum et al.,² and the *Journal of Official Statistics* dedicated to confidentiality and disclosure control).

This long history is a testament to the importance of the problem. Statistical databases can be of enormous social value: they are used for apportioning resources, evaluating medical therapies, understanding the spread of disease, improving economic utility, and informing us about ourselves as a species.

The data may be obtained in diverse ways. Some data, such as census, tax, and other sorts of official data, is compelled; other data is collected opportunistically, for example, from traffic on the Internet, transactions on Amazon, and search engine query logs; other data is provided altruistically, by respondents who hope that sharing their information will help others to avoid a specific misfortune, or more generally, to increase the public good. Altruistic data donors are typically promised their individual data will be kept confidential—in short, they are promised "privacy." Similarly, medical data and legally compelled data, such as census data and tax return data, have legal privacy

Key insights

- In analyzing private data, only by focusing on rigorous privacy guarantees can we convert the cycle of "propose-break-propose again" into a path of progress.
- A natural approach to defining privacy is to require that accessing the database teaches the analyst nothing about any individual, but this is problematic: the whole point of a statistical database is to teach general truths, for example, that smoking causes cancer. Learning this fact teaches the data analyst something about the likelihood with which certain individuals, not necessarily in the database, will develop cancer. We therefore need a definition that separates the utility of the database (learning that smoking causes cancer) from the increased risk of harm due to joining the database. This is the intuition behind differential privacy.
- This can be achieved, often with low distortion. The key idea is to randomize responses so as to effectively hide the presence or absence of the data of any individual, over the course of the lifetime of the database.

Reading for weeks 10 & 11

data security

1 of 26

TEXT ONLY

Understanding **DATABASE RECONSTRUCTION ATTACKS** on Public Data

**THESE ATTACKS
ON STATISTICAL
DATABASES ARE
NO LONGER A
THEORETICAL
DANGER.**

SIMSON GARFINKEL,
JOHN M. BROWD, AND
CHRISTIAN MARTINDALE
U.S. CENSUS BUREAU

In 2020 the U.S. Census Bureau will conduct the Constitutionally mandated decennial Census of Population and Housing. Because a census involves collecting large amounts of private data under the promise of confidentiality, traditionally statistics are published only at high levels of aggregation. Published statistical tables are vulnerable to DRAs [database reconstruction attacks], in which the underlying microdata is recovered merely by finding a set of microdata that is consistent with the published statistical tabulations. A DRA can be performed by using the tables to create a set of mathematical constraints and then solving the resulting set of simultaneous equations. This article shows how such an attack can be addressed by adding noise to the published tabulations, so that the reconstruction no longer results in the original data. This has implications for the 2020 Census.

The goal of the census is to count every person once,

csisqjwv | september-october 2016 | 1



Can a set of equations keep U.S. census data private?

By Jeffrey Morris | Jan 4, 2019, 2:42 PM

The U.S. Census Bureau is mulling moves among social scientists with what it calls a "big change" in how it plans to safeguard the confidentiality of data it releases from the decennial census.

The agency announced in December 2018 that it will apply a mathematical concept called differential privacy to its release of 2020 census data after conducting experiments that suggest current approaches don't assure confidentiality. But critics of the new policy believe the Census Bureau is moving too quickly to fix a system that isn't broken. They also fear the changes will degrade the quality of the information used by the worlds of researchers, business, and government agencies.

The move has implications that extend far beyond the research community. Proponents of differential privacy say a fence, ongoing legal battle over plans to add a citizenship question to the 2020 census has only underscored the need to assure people that the government will protect their privacy.

A noisy conflict

The Census Bureau's job is to collect, analyze, and disseminate useful information about the U.S. population. And there's a lot of it: The agency generated some 7.8 billion statistics about the 388 million people counted in the 2010 census, for example.

At the same time, the bureau is prohibited by law from releasing any information "for which the data furnished by any particular establishment or individual can be identified."

Once upon a time, meeting that requirement meant simply removing the names and addresses of respondents. Over the past several decades, however, census officials have developed a bag of statistical tricks aimed at providing additional protection without undermining the quality of the data.

Reading for weeks 7 & 8

DataSynthesizer: Privacy-Preserving Synthetic Datasets

Haryue Ping
Drexel University, USA
hp354@drexel.edu

Julia Stoyanovich*
Drexel University, USA
stoyanovich@drexel.edu

Bill Howe†
University of Washington, USA
billhowe@cs.washington.edu

ABSTRACT

To facilitate collaboration over sensitive data, we present DataSynthesizer, a tool that takes a sensitive dataset as input and generates a structurally and statistically similar synthetic dataset with strong privacy guarantees. The data owners need not release their data, while potential collaborators can begin developing models and methods with some confidence that their results will work similarly on the real dataset. The distinguishing feature of DataSynthesizer is its usability — the data owner does not have to specify any parameters to start generating and sharing data safely and effectively.

DataSynthesizer consists of three high-level modules — DataDescriber, DataGenerator, and ModelInspector. The first, DataDescriber, investigates the data types, correlations and distributions of the attributes in the private dataset, and produces a data summary, adding noise to the distributions to preserve privacy. DataGenerator samples from the summary computed by DataDescriber and outputs synthetic data. ModelInspector shows an intuitive description of the data summary that was computed by DataDescriber, allowing the data owner to evaluate the accuracy of the summarization process and adjust any parameters, if desired.

We describe DataSynthesizer and illustrate its use in an urban science context, where sharing sensitive, legally encumbered data between agencies and with outside collaborators is regarded as the primary obstacle to data-driven governance.

The code implementing all parts of this work is publicly available at <https://github.com/DataResponsibly/DataSynthesizer>.

CCS CONCEPTS

• Security and privacy → Data anonymization and sanitization; Privacy protections; Usability in security and privacy;

KEYWORDS

Data Sharing, Synthetic Data, Differential Privacy

*This work was supported in part by NSF Grants No. 180432 and 183856 and NSF Grant No. 201091.

†This work was supported by the University of Washington Information School Microsoft, the Gordon and Betty Moore Foundation (Award #2015-10126) and the Alfred P. Sloan Foundation (Award #2018) through the Data Science Environment program.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted, without fee, provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyright for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
989662/35, 07 pages, 0. 096

© 2022 ACM. 978-1-5002-5282-6/22/07...\$15.00
DOI: <https://doi.org/10.1145/3525354.3525377>

ACM Reference format

Haryue Ping, Julia Stoyanovich, and Bill Howe. 2022. DataSynthesizer: Privacy-Preserving Synthetic Datasets. In *Proceedings of ACM CCS ’22*, Chicago, IL, USA, June 27–30, 2022, 7 pages.
DOI: <https://doi.org/10.1145/3525354.3525377>

1 INTRODUCTION

Collaborative projects in the social and health sciences increasingly require sharing sensitive, privacy-encumbered data. Social scientists, government agencies, health workers, and non-profits are eager to collaborate with data scientists, but formal data sharing agreements are too slow and expensive to create in ad hoc situations — our colleagues report that 18 months is a typical timeframe to establish such agreements! As a result, many promising collaborations can fail before they even begin. Data scientists require access to the data before they can understand the problem or even determine whether they can help. But data owners cannot share data without significant legal protections in place. Beyond legal concerns, there is a general reluctance to share sensitive data with non-experts before they have “proven themselves,” since they do not understand the context in which the data was collected and may be distracted by spurious results.

To bootstrap these collaborations without incurring the cost of formal data sharing agreements, we aim to generate datasets that are structurally and statistically similar to the real data but that are 1) obviously synthetic to put the data owners at ease, and 2) offer strong privacy guarantees to prevent adversaries from extracting any sensitive information. These two requirements are not redundant: strong privacy guarantees are not always sufficient to convince data owners to release data, and even seemingly random datasets may not prevent subtle privacy attacks. With this approach, data scientists can begin to develop models and methods with synthetic data, but maintain some degree of confidence that their work will remain relevant when applied to the real data once proper data sharing agreements are in place.

We propose a tool named DataSynthesizer to address this problem. Assume that the private dataset contains one table with n attributes and k tuples, and that the values in each attribute are homogeneous, that is, they are all of the same data type. We are interested in producing a synthetic dataset such that summary statistics of all numerical, categorical, string, and datetime attributes are similar to the private dataset. What statistics we preserve depends on the data type, as we will discuss in Section 3.

DataSynthesizer infers the domains of each attribute and derives a description of the distribution of attribute values in the private dataset. This information is saved in a dataset description file, to which we refer as data summary. Then DataSynthesizer is able to generate synthetic datasets of arbitrary size by sampling from the probabilistic model in the dataset description file.

Epistemic Parity: Reproducibility as an Evaluation Metric for Differential Privacy

Lucas Fieserblatt*
New York University
New York, NY, USA
luc.fieserblatt@nyu.edu

Bernadese Hertrian
University of Washington
Seattle, WA, USA
bernadese@uw.edu

Anastasia Holovenko
Ukrainian Catholic University
Lviv, Ukraine
anastasia.holovenko@ucu.edu.ua

Wonkwon Lee
New York University
New York, NY, USA
w2733@nyu.edu

Joshua Loftus
London School of Economics
London, UK
J.L.Loftus@lse.ac.uk

Elizabeth McKinnie
University of Colorado
Boulder, CO, USA
elizabeth.mckinnie@colorado.edu

Taras Rumozhak
Ukrainian Catholic University
Lviv, Ukraine
taras.rumozhak@ucu.edu.ua

Andrii Stachnik
Ukrainian Catholic University
Lviv, Ukraine
andrii.stachnik@ucu.edu.ua

Bill Howe
University of Washington
Seattle, WA, USA
billhowe@uw.edu

Julia Stoyanovich†
New York University
New York, NY, USA
stoyanovich@nyu.edu

ABSTRACT

Differential privacy (DP) data synthesizers are increasingly proposed to afford public release of sensitive information, offering theoretical guarantees for privacy (and, in some cases, utility), but limited empirical evidence of utility in practical settings. Utility is typically measured as the error on representative proxy tasks, such as descriptive statistics, multivariate correlations, the accuracy of trained classifiers, or performance over a query workload. The ability for these results to generalize to practitioners’ experience has been questioned in a number of settings, including the U.S. Census. In this paper, we propose an evaluation methodology for synthetic data that avoids assumptions about the representativeness of proxy tasks, instead measuring the likelihood that published conclusions would change had the authors used synthetic data, a condition we call epistemic parity. Our methodology consists of reproducing empirical conclusions of peer-reviewed papers on real, publicly available data, then re-running these experiments a second time on DP synthetic data and comparing the results.

We instantiate our methodology over a benchmark of recent peer-reviewed papers in the social sciences. We express the authors’ claims computationally to automate the experiment, generate DP synthetic datasets using multiple state-of-the-art mechanisms, then estimate the likelihood that these conclusions will hold. We find that, for measurable

*Fieserblatt is the first author, Howe and Stoyanovich are the senior authors.

Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. This is a minor revision of the paper entitled “Epistemic Parity: Reproducibility as an Evaluation Metric for Differential Privacy,” published in *VLDB*, Vol. 16, No. 11, ISSN 2307-4977. DOI: <https://doi.org/10.14778/3611479.3611117>

privacy regimes, DP synthesizers can achieve high epistemic parity for several papers in our benchmark. However, some papers, and particularly some specific findings, are difficult to reproduce for any of the synthesizers. Given these results, we recommend a new class of mechanisms that offer stronger utility guarantees (as measured by epistemic parity) and more nuanced privacy protection using application-specific risks and threat models.

1. INTRODUCTION

Differential privacy (DP) has been studied intensely for over a decade, and has recently enjoyed uptake in both the private and public sectors. In situations where the downstream analysis is known, one can design specialized mechanisms with high utility [25, 28]. But an active research area is to design general DP data synthesizers (henceforth, synthesizers) that model the entire data distribution, inject noise, then sample the noisy model to generate synthetic datasets intended to be broadly usable in a variety of anticipated applications. Evidence to support claims of general utility is typically presented as results on proxy tasks over common public datasets (e.g., the ubiquitous Adult dataset [13]). Proxy tasks may include descriptive statistics, queries involving one or two variables [25, 24, 28, 5], classification accuracy [13, 26, 29], and information theoretic measures [24]. Although these proxy tasks are generally representative of real tasks, the implicit claim of generalization to practice is rarely explored.

Limited empirical evidence on relevant tasks undermines trust in the practical use of DP. The US Census Bureau adopted DP for disclosure avoidance in the 2020 census, interpreting federal law (the Census Act, 13 U.S.C. § 214, and the Confidential Information Protection and Statistical Policy Act of 2002) as a mandate to use advanced methods

motivation

Truth or dare?

Did you go out drinking over the weekend?

let's call this property **P** (Truth=Yes) and estimate **p**, the fraction of the class for whom **P** holds

1. flip a coin **C1**

1. if **C1** is tails, then **respond truthfully**

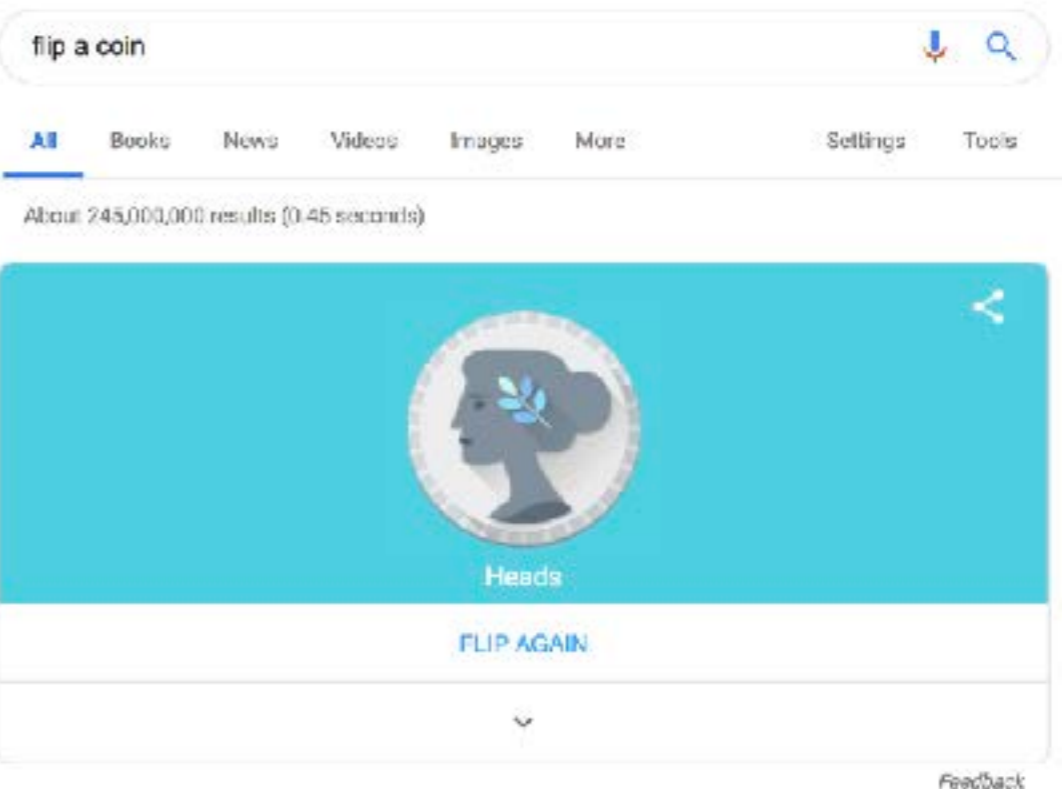
2. if **C1** is heads, then flip another coin **C2**

1. if **C2** is heads then **Yes**

2. else **C2** is tails then respond **No**

the expected fraction of **Yes** answers is:

$$A = \frac{3}{4}p + \frac{1}{4}(1-p) = \frac{1}{4} + \frac{p}{2}$$



thus, we estimate **p** as:

$$\tilde{p} = 2A - \frac{1}{2}$$

Randomized response

Did you go out drinking over the weekend?

let's call this property **P** (Truth=Yes) and estimate **p**, the fraction of the class for whom **P** holds

1. flip a coin **C1**

1. if **C1** is tails, then **respond truthfully**

2. if **C1** is heads, then flip another coin **C2**

1. if **C2** is heads then **Yes**

2. else **C2** is tails then respond **No**

randomization - adding noise - is what gives plausible deniability a process privacy method

the expected number of **Yes** answers is:

$$A = \frac{3}{4}p + \frac{1}{4}(1-p) = \frac{1}{4} + \frac{p}{2}$$

privacy comes from plausible deniability

Privacy: two sides of the coin

protecting an individual

plausible deniability



learning about the population

noisy estimates

do we really
need
randomization?

Some other options

- Data release approaches that fail to protect privacy (these are prominent classes of methods, there are others):
 - **sampling** (“just a few”) - release a small subset of the database
 - **aggregation** (e.g., **k-anonymity** - each record in the release is indistinguishable from at least $k-1$ other records)
 - **de-identification** - mask or drop personal identifiers
 - **query auditing** - stop answering queries when they become unsafe

Sampling (“just a few”)

- Suppose that we take a random small sample D' of D and release it without any modification
- If D' is much smaller than D , then every respondent is unlikely to appear in D'
- This technique provides protection for “the typical” (or for “most”) members of the dataset
- But it may be argued that **atypical** individuals are the ones needing stronger protection!
- In any case, this method is problematic because a respondent who does appear has **no plausible deniability!**
- Suppose next that appearing in the sample D' has terrible consequences. Then, every time subsampling occurs - some individual suffers horribly!

Aggregation without randomization

- Alice and Bob are professors at State University.
- In March, Alice publishes an article: “.... the current freshman class at State U is **3,005** students, **202** of whom are from families earning over \$1M per year.”
- In April, Bob publishes an article: “... **201** families in State U’s freshman class of **3,004** have household incomes exceeding \$1M per year.”
- Neither statement discloses the income of the family of any one student. But, taken together, they state that **John, a student who dropped out at the end of March**, comes from a family that earns \$1M. Anyone who has this **auxiliary information** — that John dropped out at the end of March — will be able to learn about the income of John’s family.

this is known as a problem of **composition**, and can be seen as a kind of a **differencing attack**

A basic differencing attack

- **X**: count the number of HIV-positive people in **D**
- **Y**: count the number of HIV-positive people in **D** not named *Freddie*;
- **X - Y** tells you whether *Freddie* is HIV-positive

what if $X - Y > 1$, do we still have a problem?

Reconstruction: death by a 1000 cuts

- Another serious issue for aggregation without randomization, or with an insufficient amount of randomization: **reconstruction attacks**
- **The Fundamental Law of Information Recovery** (starting with the seminal results by Irit Dinur & Kobbi Nissim, PODS 2003): overly accurate estimates of too many statistics can completely destroy privacy
- Under what conditions can an adversary reconstruct a candidate database \mathbf{D}' that agrees with the real database \mathbf{D} in **99%** of the entries?
- Suppose that \mathbf{D} has n tuples, and that noise is bounded by some quantity \mathbf{E} . Then there exists an adversary that can reconstruct \mathbf{D} to within $4\mathbf{E}$ positions, issuing all possible 2^n queries
- Put another way: if the magnitude of the noise is less than $n/401$, then 99% of \mathbf{D} can be reconstructed by the adversary. Really, any number higher than 401 will work
- **There are also reconstruction results under a limited number of queries**

$$4E = \frac{4n}{401} < \frac{n}{100}$$

Reconstruction: death by a 1000 cuts

Privacy-Preserving Data Analysis for the Federal Statistical Agencies

January 2017



John Abowd, Lorenzo Alvisi, Cynthia Dwork, Sampath Kannan, Ashwin Machanavajjhala, and Jerome Reiter

we'll discuss the use of differential privacy by the 2020 US Census later today

The Fundamental Law of Information Recovery has troubling implications for the publication of large numbers of statistics by a statistical agency: it says that the confidential data may be vulnerable to database reconstruction attacks based entirely on the data published by the agency itself. **Left unattended, such risks threaten to undermine, or even eliminate, the societal benefits inherent in the rich data collected by the nation's statistical agencies.** The most pressing immediate problem for any statistical agency is how to modernize its disclosure limitation methods in light of the Fundamental Law.

De-identification

- Also known as **anonymization**
- Mask or drop identifying attribute or attributes, such as social security number (SSN), name, mailing address
- Turns out that this also doesn't work because **auxiliary information** is available
- Fundamentally, this is due to **the curse of dimensionality**: high-dimensional data is sparse, the more you know about individuals, the less likely it is that two individuals will look alike

de-identified data can be re-identified with a linkage attack

A linkage attack: Governor Weld

In 1997, Massachusetts Group Insurance Commission released "anonymized" data on state employees that showed every single hospital visit!

Latanya Sweeney, a grad student, sought to show the ineffectiveness of this "anonymization."

She knew that Governor Weld resided in Cambridge, Massachusetts, a city of 54,000 residents and seven ZIP codes.

For twenty dollars, she purchased the complete voter rolls from the city of Cambridge, a database containing, among other things, the name, address, ZIP code, birth date, and sex of every voter.

Only six people in Cambridge shared his birth date, only three of them men, and of them, only he lived in his ZIP code.

Follow up: ZIP code, birthdate, and sex sufficient to identify 87% of Americans!

<https://arstechnica.com/tech-policy/2009/09/your-secrets-live-online-in-databases-of-ruin/>

The Netflix prize linkage attack

- In 2006, Netflix released a dataset containing ~100M **movie ratings** by ~500K users (about 1/8 of the Netflix user base at the time)
- **FAQ:** “Is there any customer information in the dataset that should be kept private?”

*“No, all customer identifying information has been removed; all that remains are ratings and dates. This follows our privacy policy, which you can review here. Even if, for example, you knew all your own ratings and their dates you probably couldn’t identify them reliably in the data because only **a small sample** was included (less than one-tenth of our complete dataset) and that **data was subject to perturbation**. Of course, since you know all your own ratings that really isn’t a privacy problem is it?”*

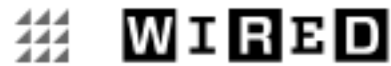
The real question: How much does the adversary need to know about a Netflix subscriber to identify her record in the dataset, and thus learn her complete movie viewing history?

The Netflix prize linkage attack

- Very little auxiliary information is needed to de-anonymize an average subscriber record from the Netflix Prize dataset
- **Perturbation, you say?** With 8 movie ratings (of which 2 may be completely wrong) and dates that may have a 14-day error, 99% of records be uniquely identified in the dataset
- For 68%, two ratings and dates (with a 3-day error) are sufficient
- **Even without any dates, a substantial privacy breach occurs, especially when the auxiliary information consists of movies that are not blockbusters:** Two movies are no longer sufficient, but 84% of subscribers can be uniquely identified if the adversary knows 6 out of 8 moves outside the top 500

We cannot assume a priori that any data is harmless!

The Netflix prize linkage attack



An in-the-closet lesbian mother is suing Netflix for privacy invasion, alleging the movie rental company made it possible for her to be outed when it disclosed insufficiently anonymous information about nearly half-a-million customers as part of its \$1 million contest to improve its recommendation system.

The suit known as [Doe v. Netflix \(.pdf\)](#) was filed in federal court in California on Thursday, alleging that Netflix violated fair-trade laws and a federal privacy law protecting video rental records, when it launched its popular contest in September 2006.

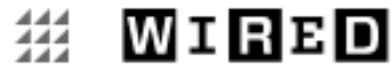
The suit seeks more than \$2,500 in damages for each of more than 2 million Netflix customers.

RYAN SINGEL SECURITY 12.17.09 04:29 PM

NETFLIX SPILLED YOUR BROKEBACK MOUNTAIN SECRET, LAWSUIT CLAIMS



The Netflix prize linkage attack



RYAN BINDEL SECURITY 03.12.10 02:48 PM

NETFLIX CANCELS RECOMMENDATION CONTEST AFTER PRIVACY LAWSUIT



Netflix is canceling its second \$1 million Netflix Prize to settle a legal challenge that it breached customer privacy as part of the first contest's race for a better movie-recommendation engine.

Query auditing

- Monitor queries: each query is granted or denied depending on what other queries were answered in the past
- If this method were to work, it could be used to detect that a differencing attack is about to take place
- Unfortunately, it doesn't work:
 - **Query auditing is computationally infeasible**
 - Refusal to respond to a query may itself be disclosive
 - We refuse to execute a query, then what? No information access at all?

Query auditing

- We have a set of (secret) Boolean variables \mathbf{X} and the result of some *statistical queries* over this set
- A *statistical query* \mathbf{Q} specifies a subset \mathbf{S} of the variables in \mathbf{X} , and returns the sum of the values of all variables in \mathbf{S}

Example:

Relation Employees (name, age, salary)

Query `select sum(salary) from Employees where age > 35`

Suppose that Employees (name, age) is public, but salary is confidential

Query auditing

- We have a set of (secret) Boolean variables \mathbf{X} and the result of some *statistical queries* over this set
- A *statistical query* \mathbf{Q} specifies a subset \mathbf{S} of the variables in \mathbf{X} , and returns the sum of the values of all variables in \mathbf{S}
- **The auditing problem:** Decide whether the value of any Boolean variable is determined by the results of the queries
- **Main result:** The Boolean auditing problem is **coNP-complete**
 - coNP-complete is the hardest class of problems in coNP: all coNP problems can be formulated as a special case of any coNP-complete problem
 - if P does not equal NP , then there does not exist a polynomial time algorithm that solves this problem

**privacy-
preserving data
analysis**

Privacy: two sides of the coin

protecting an individual

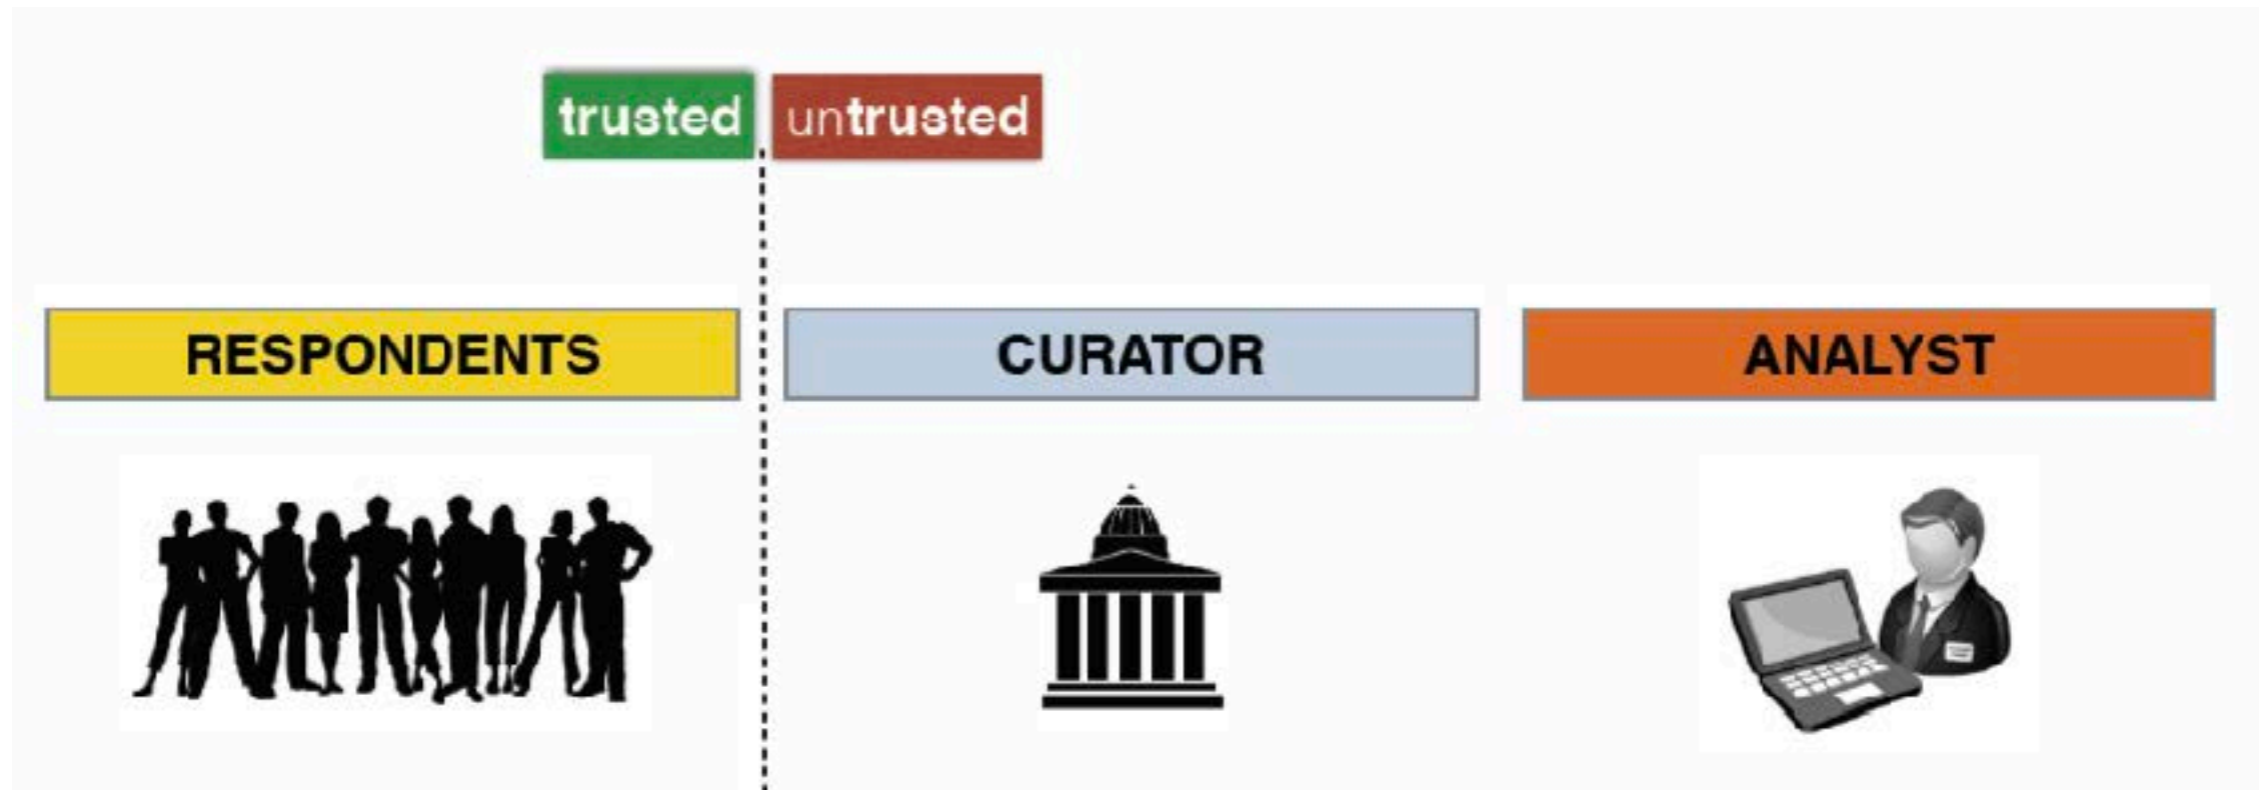
plausible deniability



learning about the population

noisy estimates

Privacy-preserving data analysis

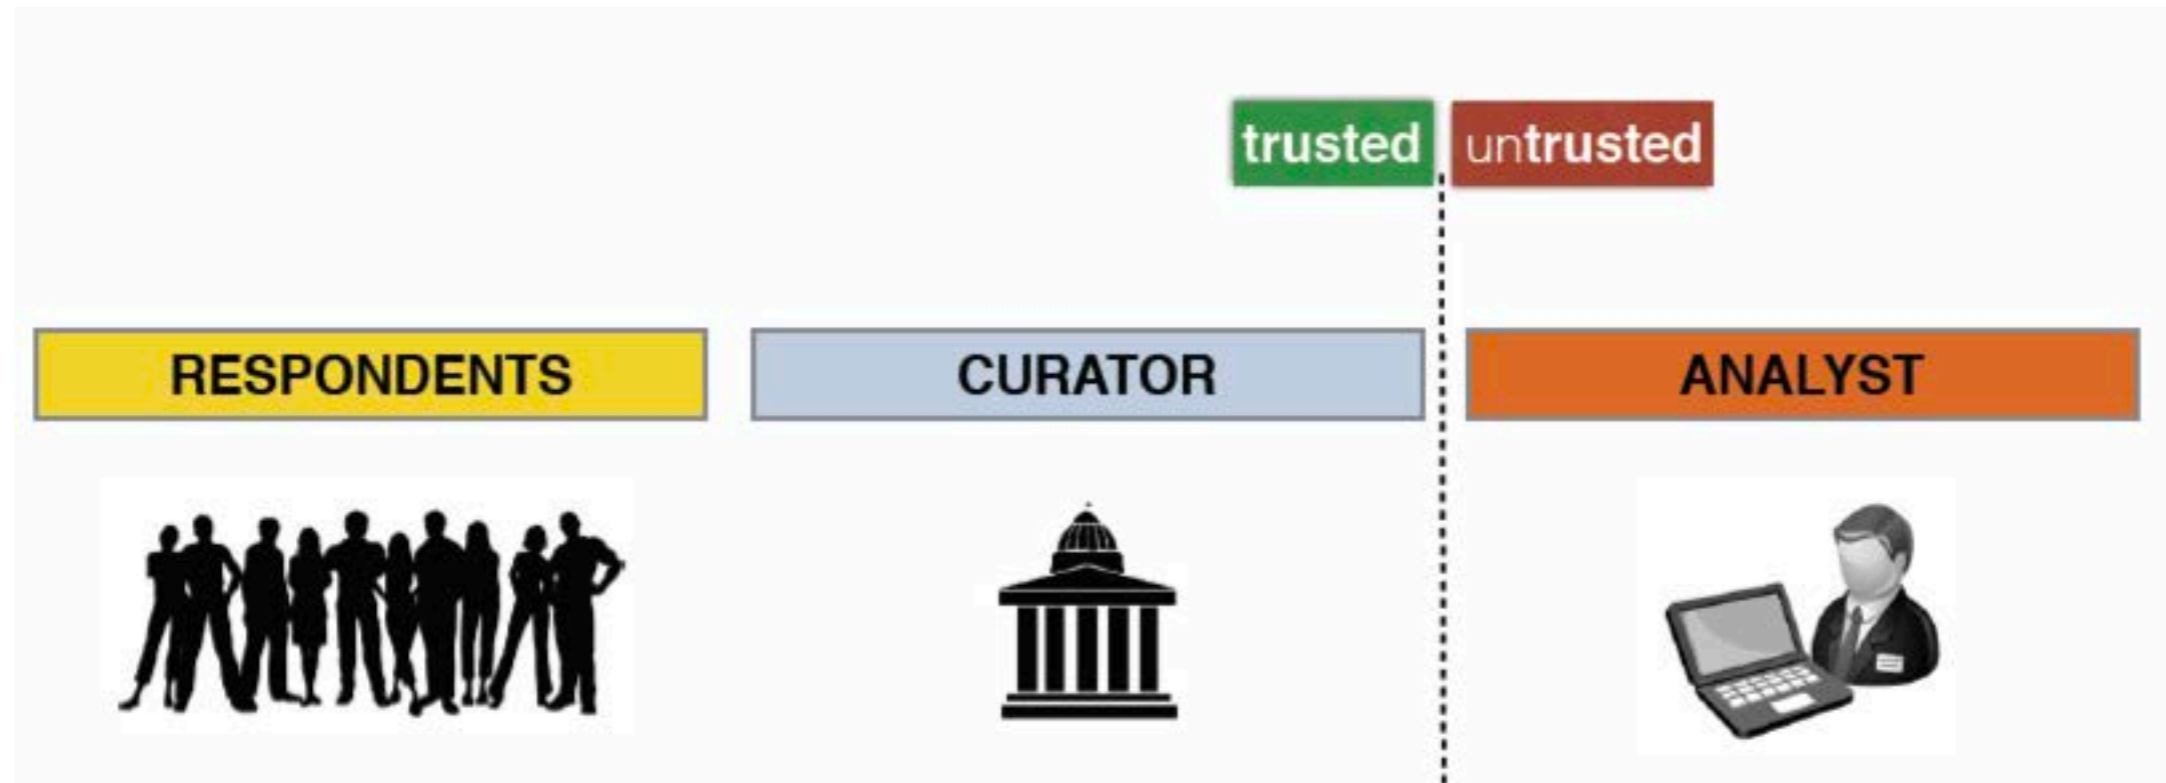


respondents contribute their personal data

the **curator** is **untrusted**, collects data, releases it to analysts

the **analyst** is **untrusted**, extracts value from data

Privacy-preserving data analysis

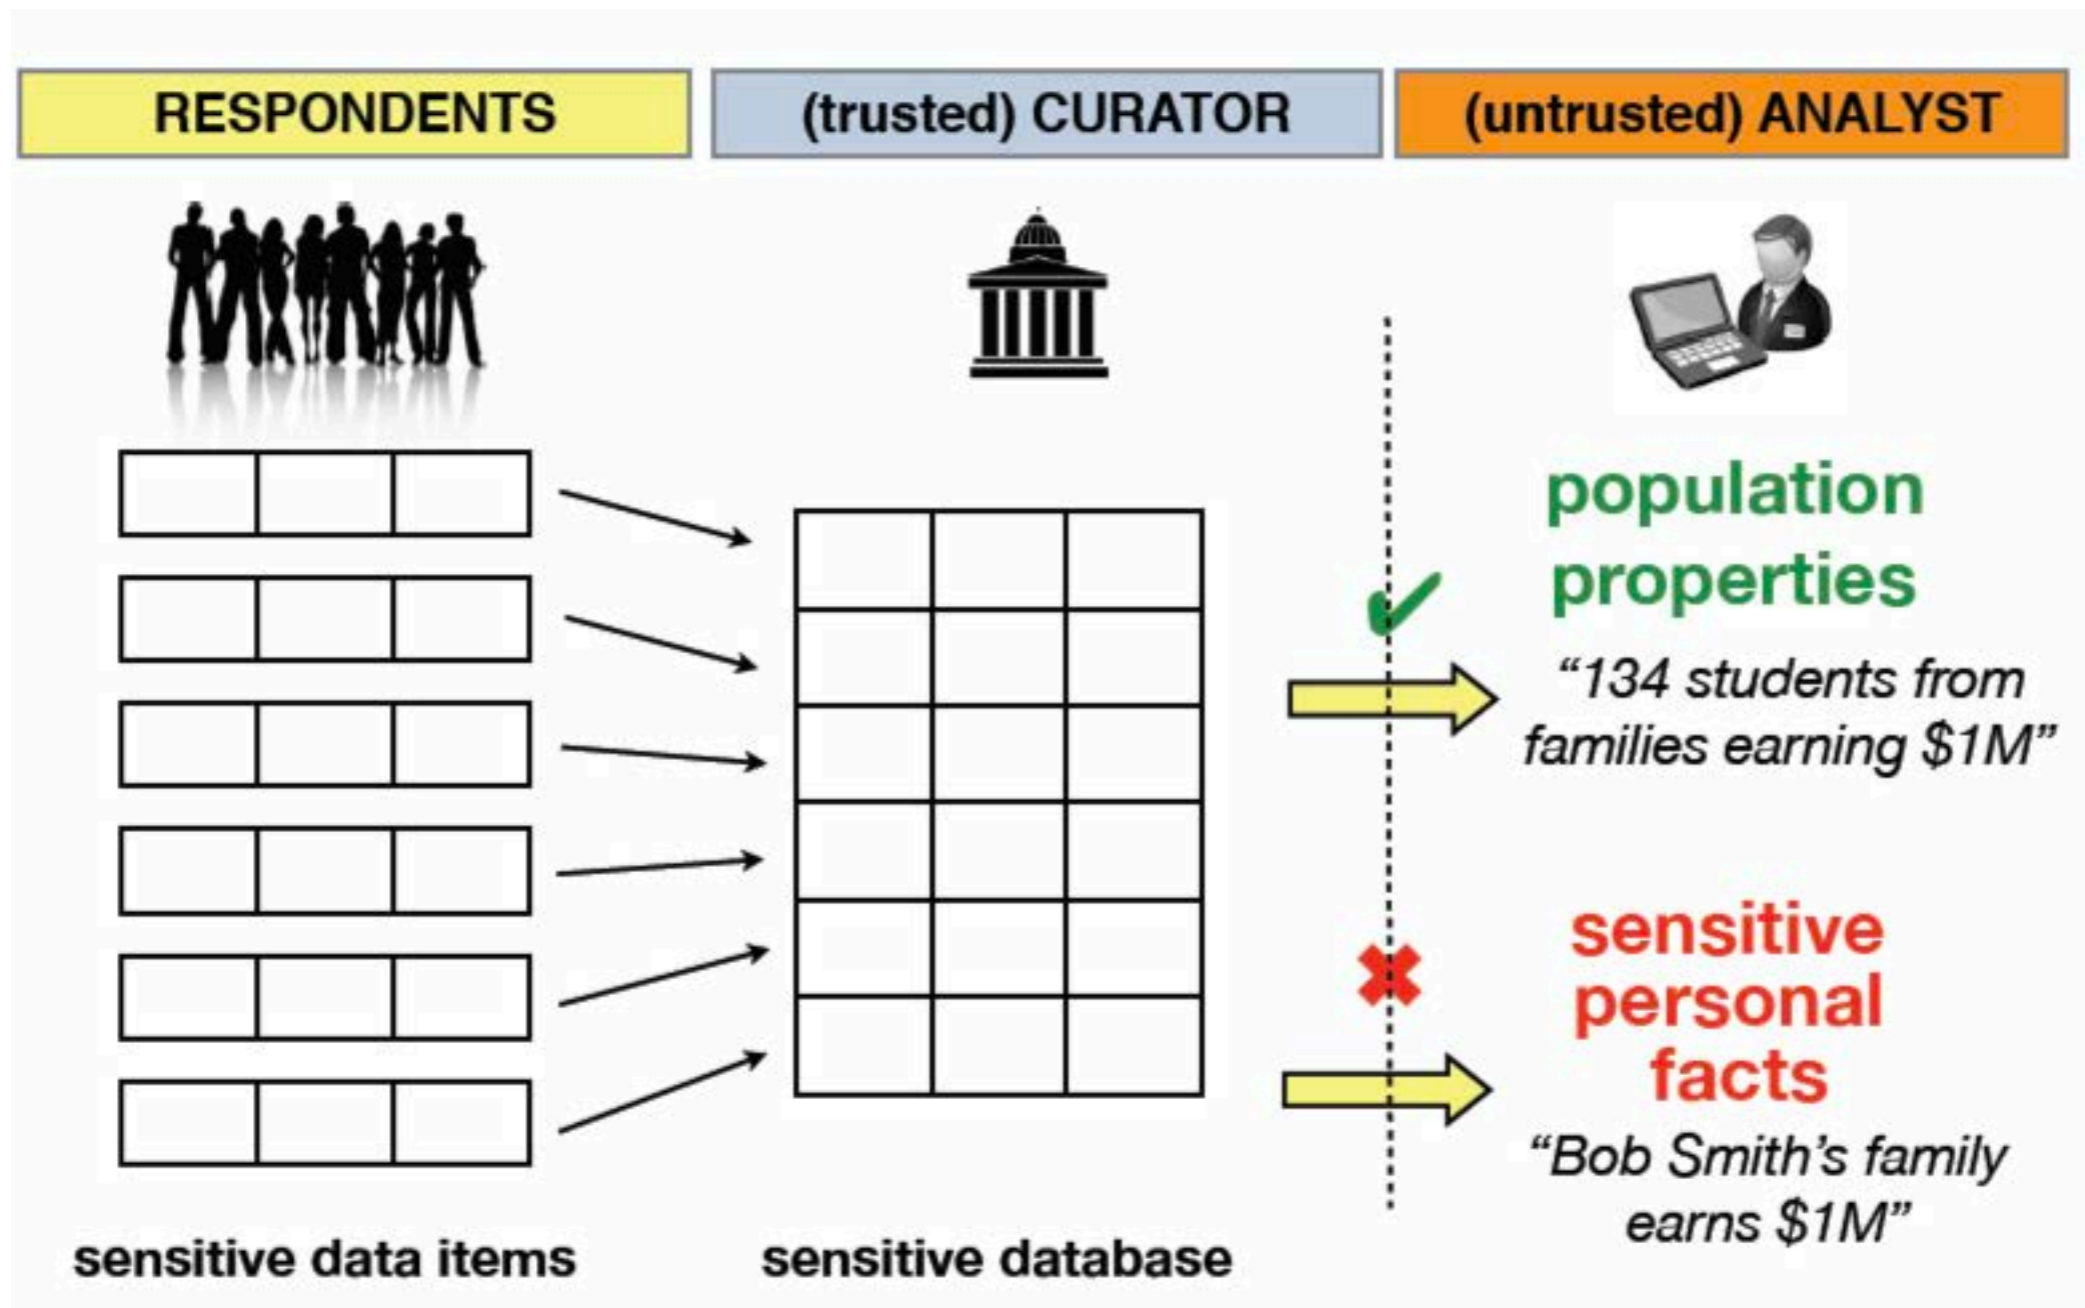


respondents in the population seek protection of their personal data

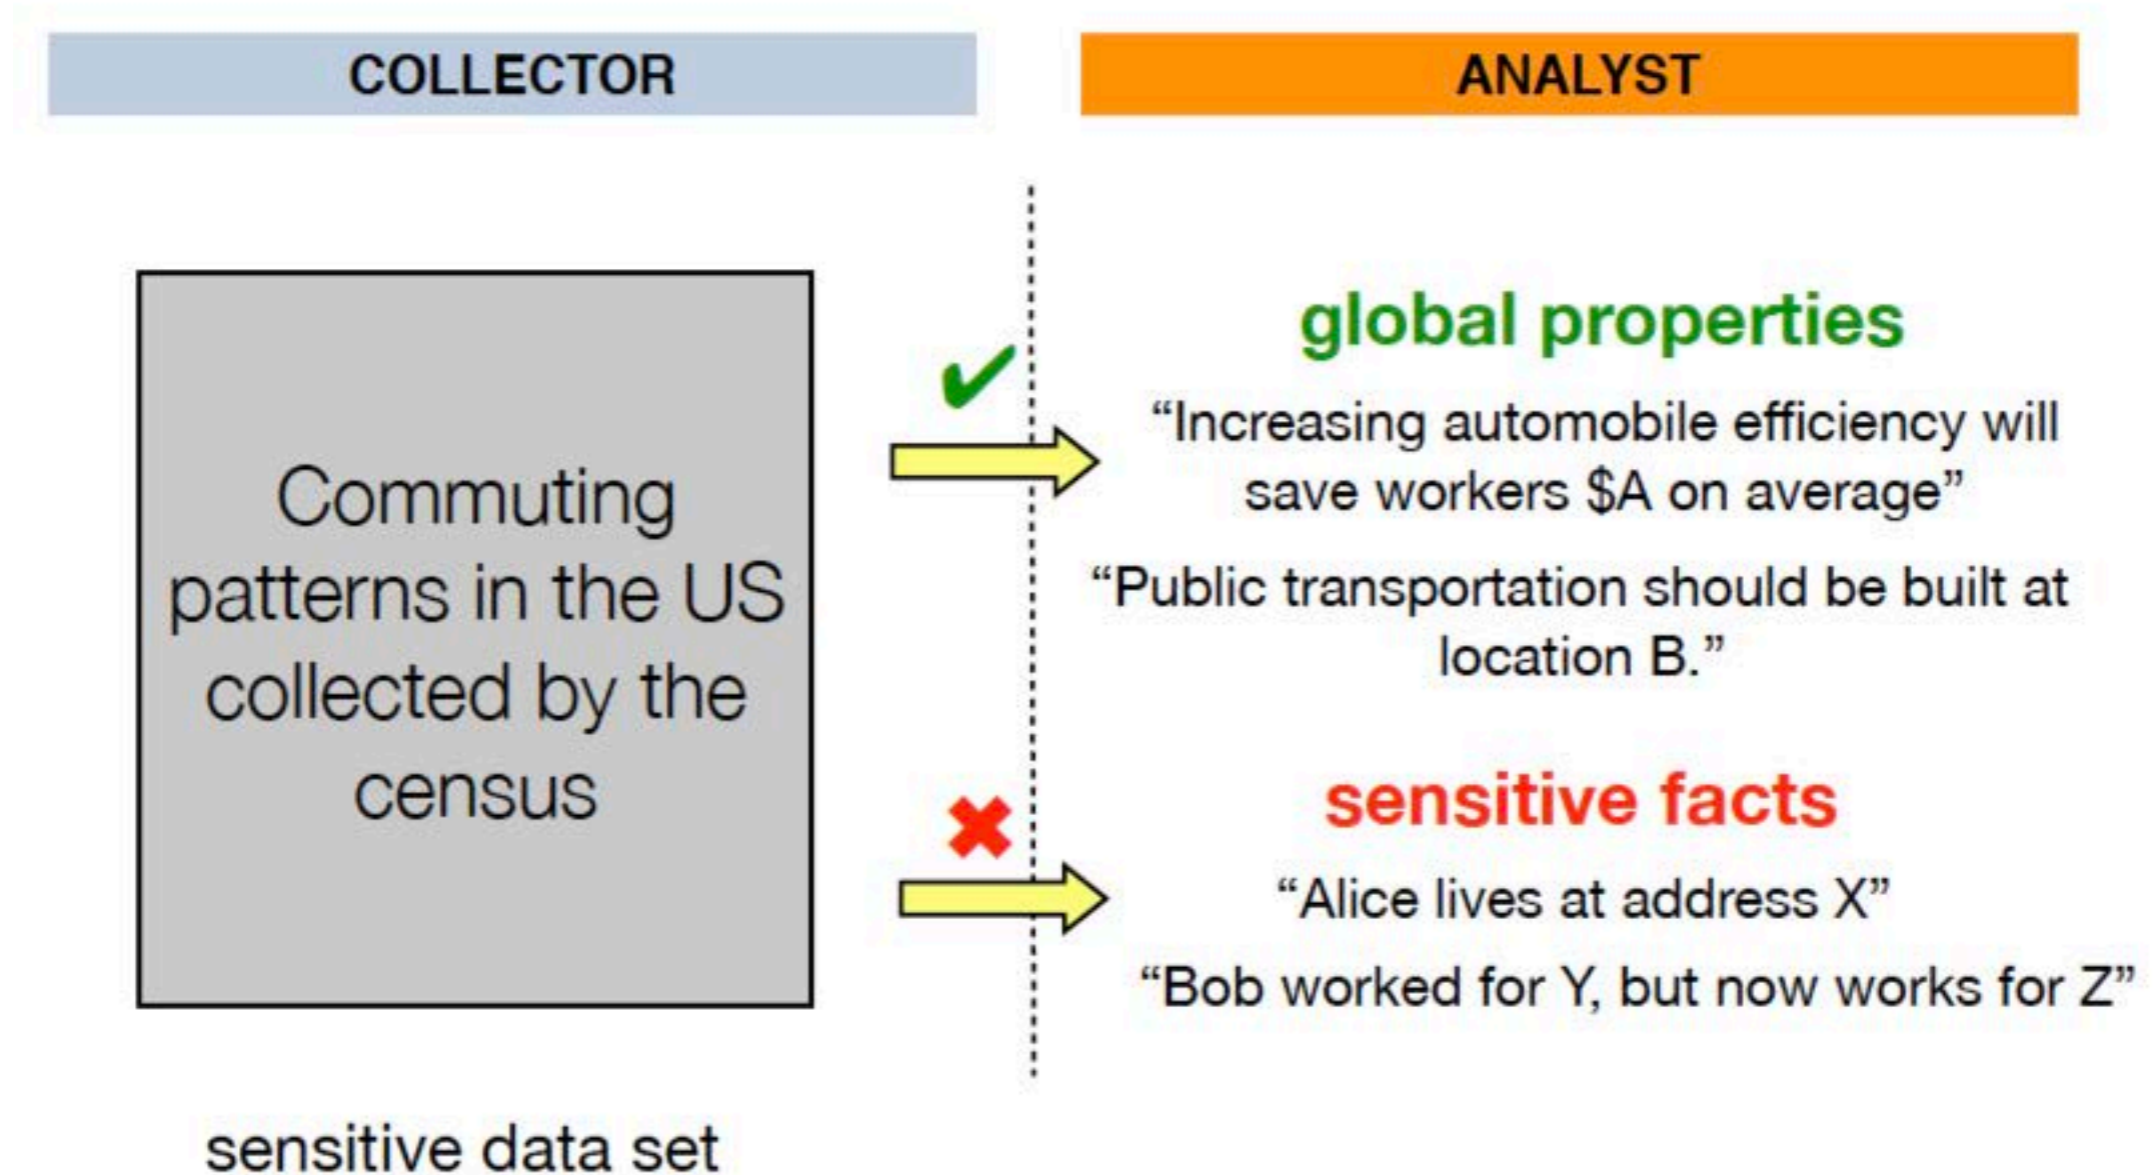
the **curator** is **trusted** to collect data and is responsible for safely releasing it

the **analyst** is **untrusted** and wants to gain the most accurate insights into the population

Privacy-preserving data analysis



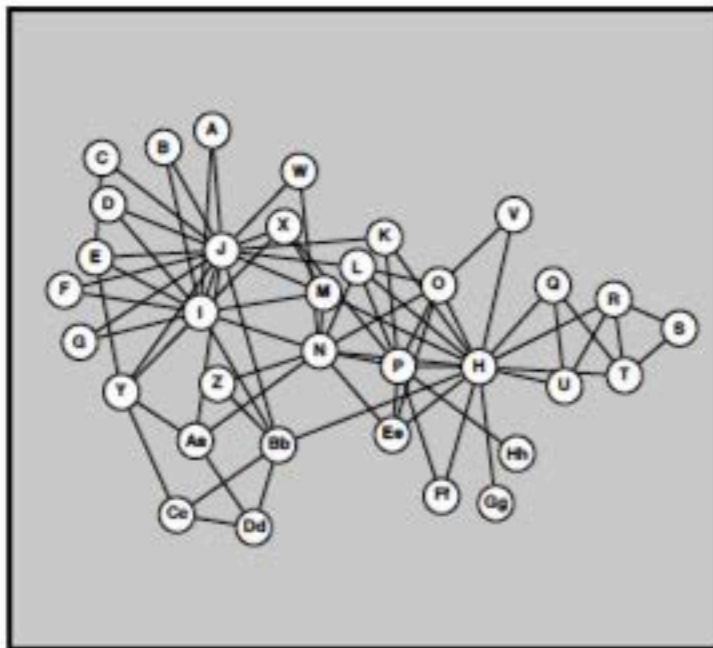
Example: US Census



Example: Social networks

COLLECTOR

ANALYST



sensitive data set



global properties

“How rapidly do rumors spread in this network?”

“Are people most likely to form friendships with those who share their attributes?”

sensitive facts

“Alice is present in this network”

“Alice and Bob are connected”

Defining private data analysis

- Take 1: If **nothing is learned** about any individual in the dataset, then no individual can be harmed by analysis.
- **Dalenius' Desideratum**: an *ad omnia* (Latin: "for all") privacy goal for statistical databases, as opposed to *ad hoc* (Latin: "for this"). Anything that can be learned about a respondent from the statistical database should be learnable without access to the database.
- Put another way, the adversary's prior and posterior views about an individual should not be different.
- This objective is **unachievable** because of auxiliary information.
- **Example**: Alice knows that John smokes. She read a medical research study that found a causal relationship between smoking and lung cancer. Alice concludes, based on study results and her prior knowledge about John, that he has a heightened risk of developing lung cancer.
- Further, the risk is to everyone in a particular group (smokers, in this example), **irrespective of whether they participated in the study**.

Defining private data analysis

- Take 1: If **nothing is learned** about any individual in the dataset, then no individual can be harmed by analysis.
- **Dalenius' Desideratum**: an "*ad omnia*" (opposed to *ad hoc*) privacy goal for statistical databases: Anything that can be learned about a respondent from the statistical database should be learnable without access to the database.
- Put another way, the adversary's prior and posterior views about an individual should not be different.
- Take 2: The information released about the sensitive dataset is virtually indistinguishable **whether or not a respondent's data is in the dataset**. This is an informal statement of **differential privacy**: that no information **specific to an individual** is revealed.

Defining private data analysis

DOI:10.1145/1886739.1886759

What does it mean to preserve privacy?

BY CYNTHIA DWORK

A Firm Foundation for Private Data Analysis

IN THE INFORMATION realm, loss of privacy is usually associated with failure to control access to information, to control the flow of information, or to control the purposes for which information is employed. Differential privacy arose in a context in which ensuring privacy is a challenge even if all these control problems are solved: privacy-preserving statistical analysis of data.

The problem of *statistical disclosure control*—revealing accurate statistics about a set of respondents while preserving the privacy of individuals—has a venerable history, with an extensive literature spanning statistics, theoretical computer science, security, databases, and cryptography (see, for example, the excellent survey of Adam and Wortmann,¹ the discussion of related work in Blum et al.,² and the *Journal of Official Statistics* dedicated to confidentiality and disclosure control).

This long history is a testament to the importance of the problem. Statistical databases can be of enormous social value; they are used for apportioning resources, evaluating medical therapies, understanding the spread of disease, improving economic utility, and informing us about ourselves as a species.

The data may be obtained in diverse ways. Some data, such as census, tax, and other sorts of official data, is compelled; other data is collected opportunistically, for example, from traffic on the Internet, transactions on Amazon, and search engine query logs; other data is provided altruistically, by respondents who hope that sharing their information will help others to avoid a specific misfortune, or more generally, to increase the public good. Altruistic data donors are typically promised their individual data will be kept confidential—in short, they are promised “privacy.” Similarly, medical data and legally compelled data, such as census data and tax return data, have legal privacy

» Key Insights

- In analyzing private data, only by focusing on rigorous privacy guarantees can we convert the cycle of “propose-break-propose again” into a path of progress.
- A natural approach to defining privacy is to require that accessing the database teaches the analyst nothing about any individual. But this is problematic: the whole point of a statistical database is to teach general truths, for example, that smoking causes cancer. Learning this fact teaches the data analyst something about the likelihood with which certain individuals, not necessarily in the database, will develop cancer. We therefore need a definition that separates the utility of the database (learning that smoking causes cancer) from the increased risk of harm due to joining the database. This is the intuition behind differential privacy.
- This can be achieved, often with low distortion. The key idea is to randomize responses so as to effectively hide the presence or absence of the data of any individual over the course of the lifetime of the database.

“A natural approach to defining privacy is to require that accessing the database teaches the analyst nothing about any individual. But this is problematic: **the whole point of a statistical database is to teach general truths**, for example, that smoking causes cancer. Learning this fact teaches the data analyst something about the likelihood with which certain individuals, not necessarily in the database, will develop cancer. We therefore **need a definition that separates the utility of the database** (learning that smoking causes cancer) **from the increased risk of harm due to joining the database. This is the intuition behind differential privacy.**”

***differential
privacy (DP)***

Differential privacy: the formalism

We will define privacy with respect to a database \mathbf{D} that is made up of rows (equivalently, tuples) representing individuals. Tuples come from some universe of datatypes (the set of all possible tuples).

The ℓ_1 norm of a database \mathbf{D} , denoted $\|\mathbf{D}\|_1$ is the number of tuples in \mathbf{D} .

The ℓ_1 distance between databases \mathbf{D}_1 and \mathbf{D}_2 represents the number of tuples on which they differ. $\|\mathbf{D}_1 - \mathbf{D}_2\|_1$

We refer to a pair of databases that differ in at most 1 tuple as **neighboring databases** $\|\mathbf{D}_1 - \mathbf{D}_2\|_1 \leq 1$

Of these \mathbf{D}_1 and \mathbf{D}_2 , one, say \mathbf{D}_2 , is a subset of the other, and, when a proper subset, the larger database \mathbf{D}_2 contains 1 extra tuple.

Differential privacy: the formalism

The information released about the sensitive dataset is virtually indistinguishable **whether or not a respondent's data is in the dataset**. This is an informal statement of **differential privacy**. That is, no information **specific to an individual** is revealed.

A randomized algorithm M provides **ϵ -differential privacy** if, for all neighboring databases D_1 and D_2 , and for any set of outputs S :

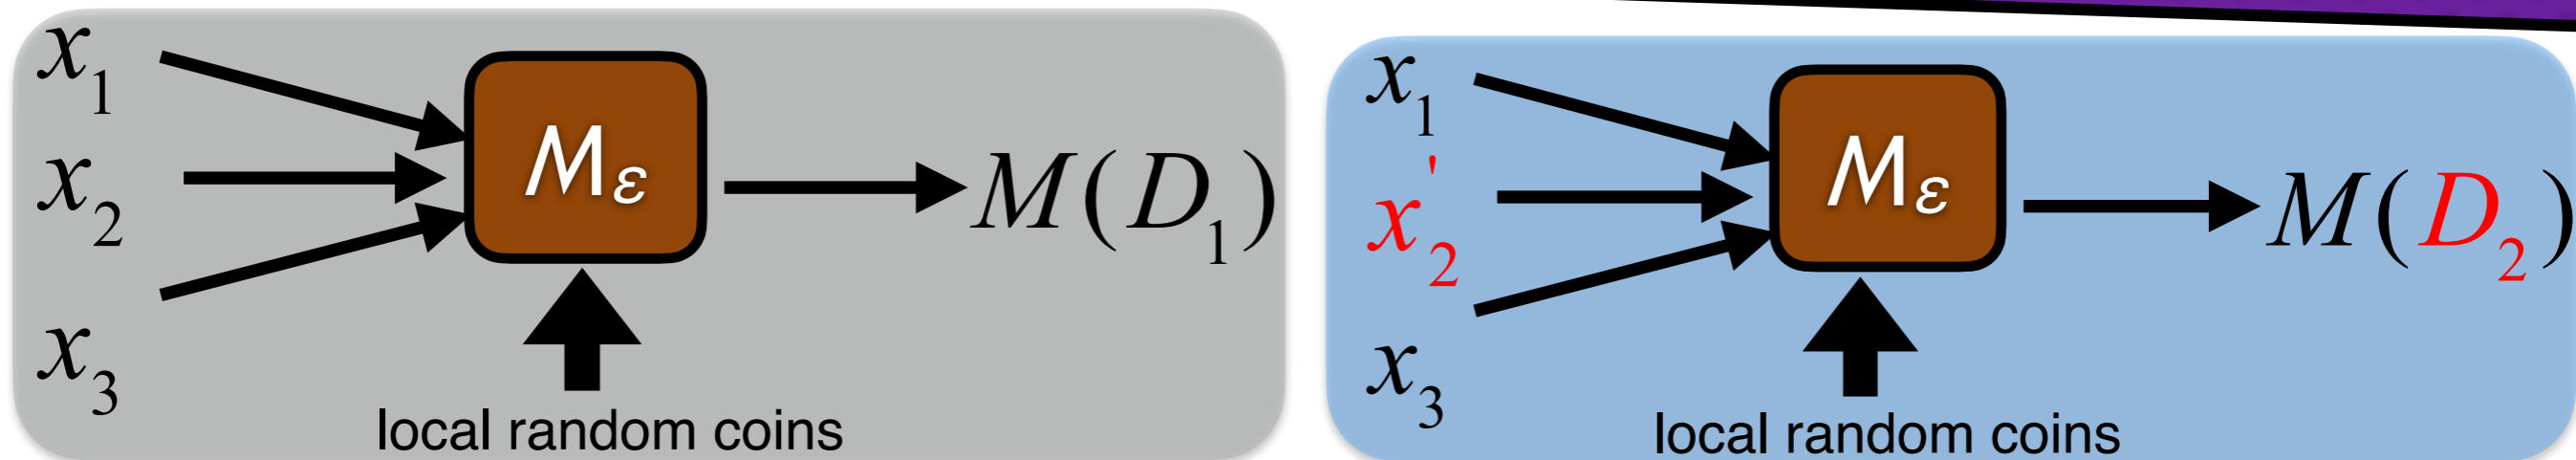
$$\Pr[M(D_1) \in S] \leq e^\epsilon \Pr[M(D_2) \in S]$$

ϵ (epsilon) is a privacy parameter

↓ **lower ϵ = stronger privacy** ↑

The notion of **neighboring databases** is integral to plausible deniability: D_1 can represent a database with a particular respondent's data, D_2 can represent a neighboring database but without that respondent's data

Differential privacy: the formalism



A randomized algorithm M provides **ϵ -differential privacy** if, for all neighboring databases D_1 and D_2 , and for any set of outputs S :

$$\Pr[M(D_1) \in S] \leq e^\epsilon \Pr[M(D_2) \in S]$$

Think of database of respondents $D=(x_1, \dots, x_n)$ as **fixed** (not random),
 $M(D)$ is a random variable distributed over possible outputs

Neighboring databases induce **close distributions** on outputs

Back to randomized response

Did you go out drinking over the weekend?

1. flip a coin **C1**

1. if **C1** is tails, then **respond truthfully**

2. if **C1** is heads, then flip another coin **C2**

1. if **C2** is heads then **Yes**

2. else **C2** is tails then respond **No**

Denote:

• Truth=Yes by **P**

• Response=Yes by **A**

• **C1**=tails by **T**

• **C1**=heads and **C2**=tails by **HT**

• **C1**=heads and **C2**=heads by **HH**

A randomized algorithm **M** provides **ϵ -differential privacy** if, for all neighboring databases **D_1** and **D_2** , and for any set of outputs **S**:

$$\Pr[M(D_1) \in S] \leq e^\epsilon \Pr[M(D_2) \in S]$$

$$\Pr[A | P] = \Pr[T] + \Pr[HH] = \frac{3}{4}$$

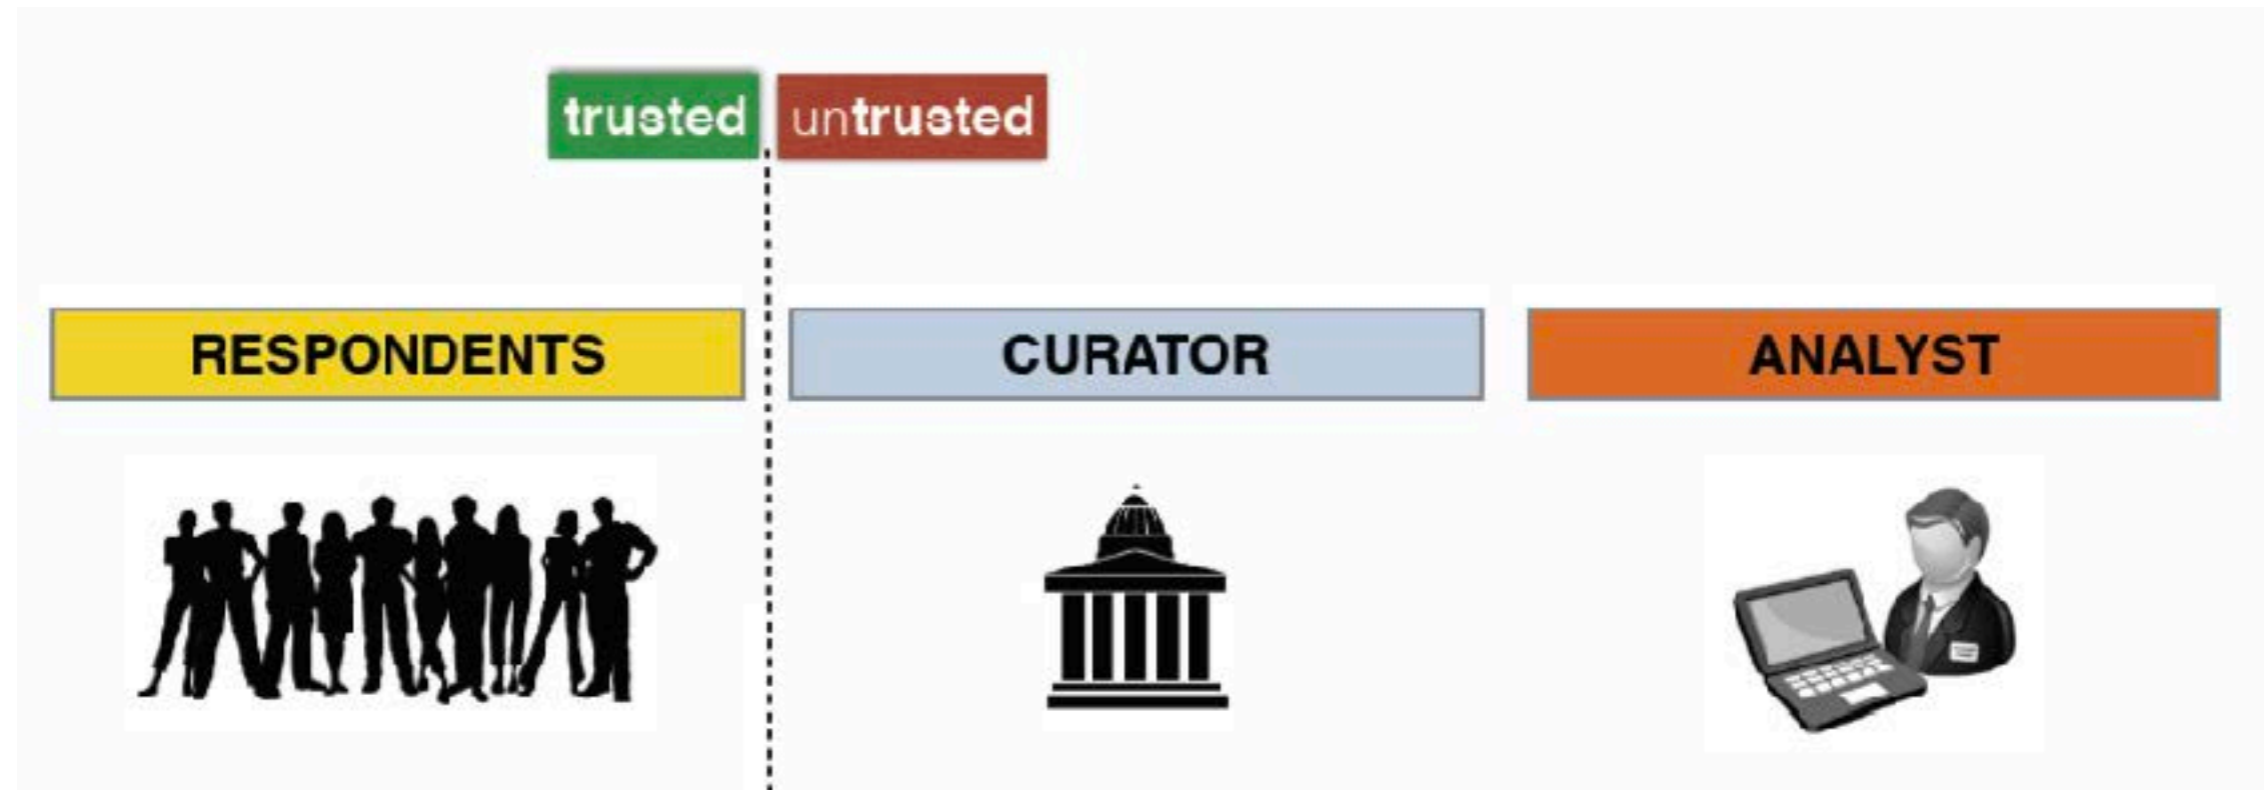
$$\Pr[A | \neg P] = \Pr[HH] = \frac{1}{4}$$

$$\Pr[A | P] = 3 \Pr[A | \neg P]$$

$$\Rightarrow \epsilon = \ln 3$$

our version of randomized response is
 $(\ln 3)$ -differentially private

Local differential privacy

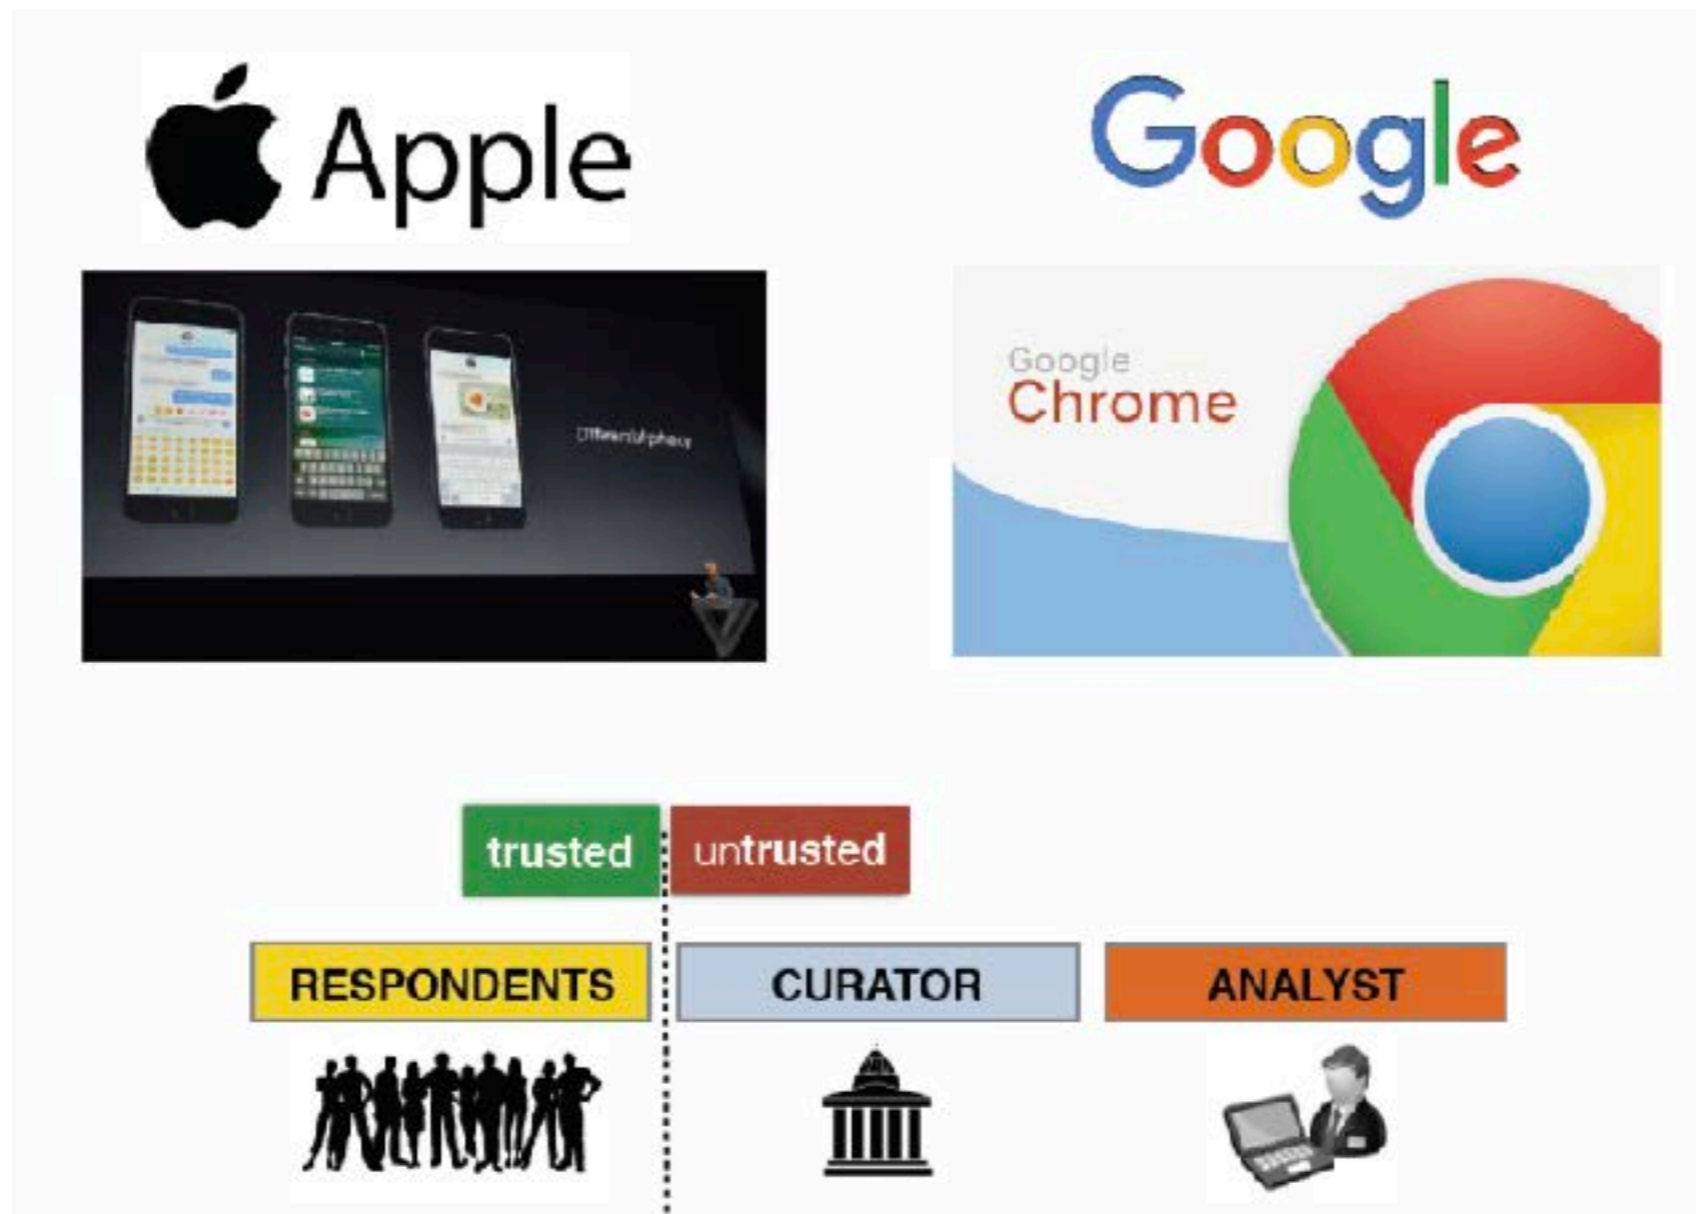


respondents contribute their personal data

the **curator** is **untrusted**, collects data, releases it to analysts

the **analyst** is **untrusted**, extracts value from data

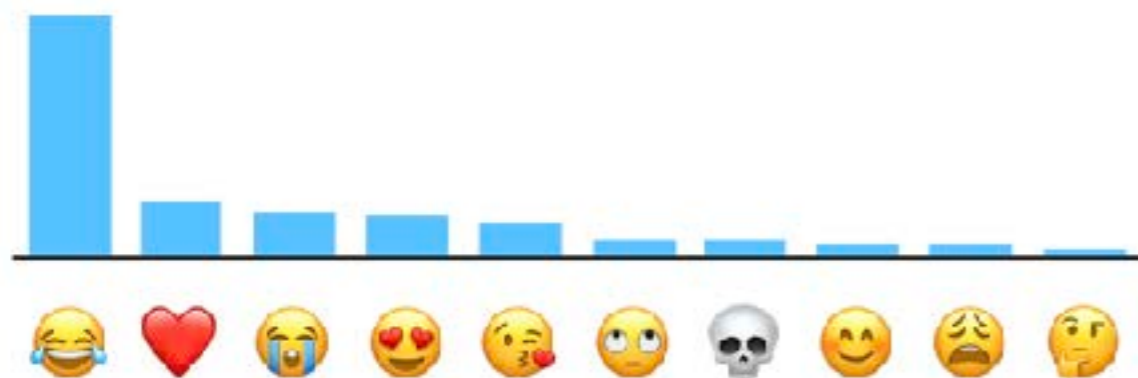
Differential privacy in the field



Example: What's your favorite emoji?

A privacy-preserving system

Apple has adopted and further developed a technique known in the academic world as *local differential privacy* to do something really exciting: gain insight into what many Apple users are doing, while helping to preserve the privacy of individual users. It is a technique that enables Apple to learn about the user community without learning about individuals in the community. Differential privacy transforms the information shared with Apple before it ever leaves the user's device such that Apple can never reproduce the true data.



Example: What's your favorite emoji?

Apple uses local differential privacy to help protect the privacy of user activity in a given time period, while still gaining insight that improves the intelligence and usability of such features as:

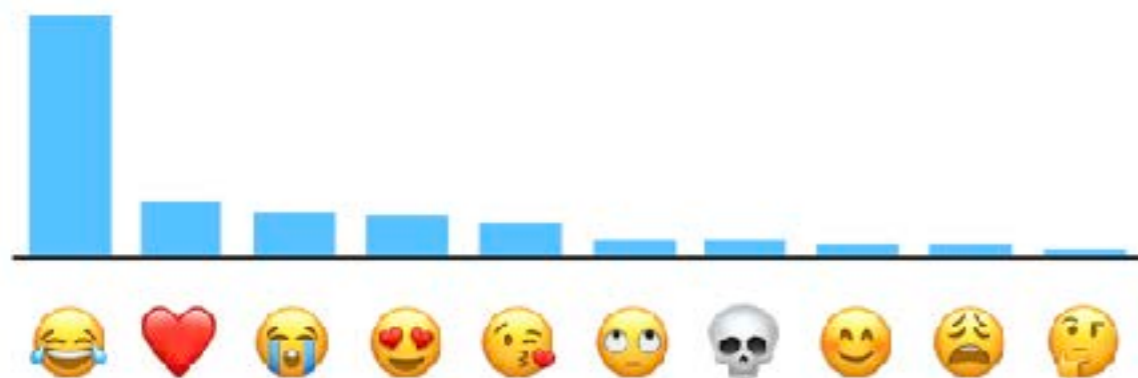
- QuickType suggestions
- Emoji suggestions
- Lookup Hints
- Safari Energy Draining Domains
- Safari Autoplay Intent Detection (macOS High Sierra)
- Safari Crashing Domains (iOS 11)
- Health Type Usage (iOS 10.2)



Example: What's your favorite emoji?

Privacy budget

The Apple differential privacy implementation incorporates the concept of a **per-donation *privacy budget*** (quantified by the parameter epsilon), and sets a strict limit on the number of contributions from a user in order to preserve their privacy. The reason is that the slightly-biased noise used in differential privacy tends to average out over a large numbers of contributions, making it theoretically possible to determine information about a user's activity over a large number of observations from a single user (though it's important to note that Apple doesn't associate any identifiers with information collected using differential privacy).



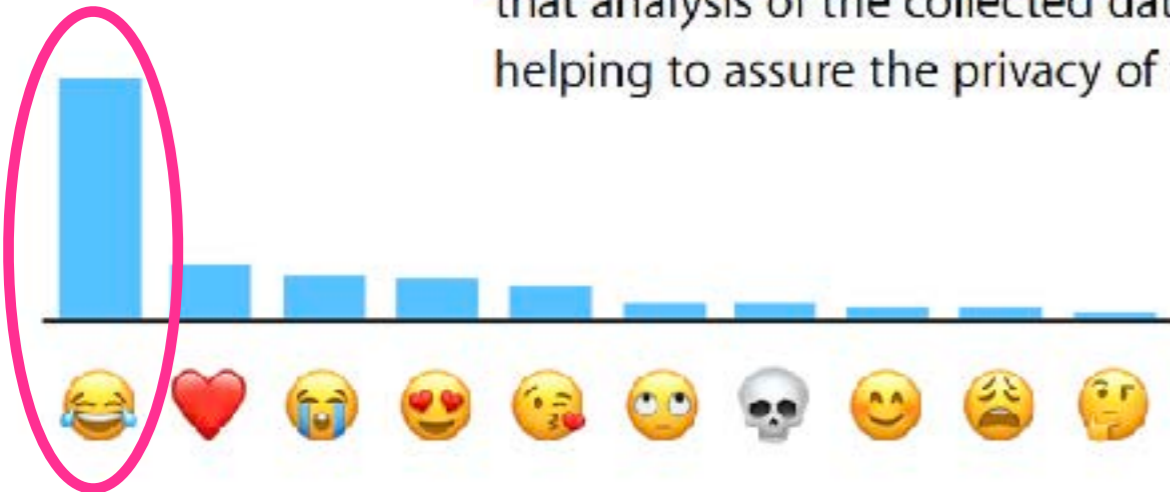
Example: What's your favorite emoji?

Count Mean Sketch

In our use of the Count Mean Sketch technique for differential privacy, the original information being processed for sharing with Apple is encoded using a series of mathematical functions known as *hash functions*, making it easy to represent data of varying sizes in a matrix of fixed size.

The data is encoded using variations of a SHA-256 hash followed by a privatization step and then written into the sketch matrix with its values initialized to zero.

The noise injection step works as follows: After encoding the input as a vector using a hash function, each coordinate of the vector is then flipped (written as an incorrect value) with a probability of $1/(1 + e^{\epsilon/2})$, where ϵ is the privacy parameter. This assures that analysis of the collected data cannot distinguish actual values from flipped values, helping to assure the privacy of the shared information.



Transparency is important!

ANDY GREENBERG

SECURITY 09.15.2017 09:28 AM

≡ WIRED

How One of Apple's Key Privacy Safeguards Falls Short

Apple has boasted of its use of a cutting-edge data science known as "differential privacy." Researchers say they're doing it wrong.

Epsilon, Epsilon

"...[Researchers] examined how Apple's software injects random noise into personal information—ranging from emoji usage to your browsing history to HealthKit data to search queries—before your iPhone or MacBook upload that data to Apple's servers.

Ideally, that obfuscation helps protect your private data from any hacker or government agency that accesses Apple's databases, advertisers Apple might someday sell it to, or even Apple's own staff. But **differential privacy's effectiveness depends on a variable known as the "privacy loss parameter," or "epsilon,"** which determines just how much specificity a data collector is willing to sacrifice for the sake of protecting its users' secrets. By taking apart Apple's software to determine the epsilon the company chose, the researchers found that **MacOS uploads significantly more specific data than the typical differential privacy researcher might consider private.** iOS 10 uploads even more. And perhaps most troubling, according to the study's authors, is that **Apple keeps both its code and epsilon values secret,** allowing the company to potentially change those critical variables and erode their privacy protections with little oversight...."

A closer look at differential privacy

A randomized algorithm M provides **ϵ -differential privacy** if, for all neighboring databases D_1 and D_2 , and for any set of outputs S :

$$\Pr[M(D_1) \in S] \leq e^\epsilon \Pr[M(D_2) \in S]$$

↓ **lower ϵ = stronger privacy** ↑

- The state-of-the-art in privacy technology, first proposed in 2006
- Has precise mathematical properties, captures cumulative privacy loss over multiple uses with the concept of a **privacy budget**
- Privacy guarantee encourages participation by respondents
- Robust against strong adversaries, with auxiliary information, including also **future auxiliary information!**
- Precise error bounds that can be made public

A closer look at differential privacy

A randomized algorithm ***M*** provides **ϵ -differential privacy** if, for all neighboring databases ***D*₁** and ***D*₂**, and for any set of outputs ***S***:

$$\Pr[M(D_1) \in S] \leq e^\epsilon \Pr[M(D_2) \in S]$$

 **lower ϵ = stronger privacy** 

ϵ (epsilon) cannot be too small: think 1/10, not 1/2⁵⁰

Differential privacy is a condition on the **algorithm *M*** (process privacy). Saying simply that “the output is safe” does not take into account how it was computed, and is insufficient.

**query sensitivity
& composition**

Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$



lower ϵ = stronger privacy



- Example 1: counting queries
 - “How many elements in D satisfy property P ?” **What’s Δq ?**
 - “What fraction of the elements in D satisfy property P ?”
- Example 2: max / min
 - “What is the maximum employee salary in D ?” **What’s Δq ?**

Intuition: for a given ϵ , the higher the sensitivity, the more noise we need to add to meet the privacy guarantee

Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$

query q	query sensitivity Δq
select count(*) from D	1
select count(*) from D where sex = Male and age > 30	?

Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$

query q	query sensitivity Δq
select count(*) from D	1
select count(*) from D where sex = Male and age > 30	1
select MAX(salary) from D	?

Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$

query q	query sensitivity Δq
select count(*) from D	1
select count(*) from D where sex = Male and age > 30	1
select MAX(salary) from D	$MAX(salary) - MIN(salary)$
select gender, count(*) from D group by gender	?

Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$

query q	query sensitivity Δq
select count(*) from D	1
select count(*) from D where sex = Male and age > 30	1
select MAX(salary) from D	$MAX(salary) - MIN(salary)$
select gender, count(*) from D group by gender	1 (disjoint groups, presence or absence of one tuple impacts only one of the counts)

Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$

query q

query sensitivity Δq

select gender, count(*)
from D group by gender

1 (**disjoint groups**, presence or absence of one tuple impacts only one of the counts)

an arbitrary list of m counting queries

?

Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$

query q

query sensitivity Δq

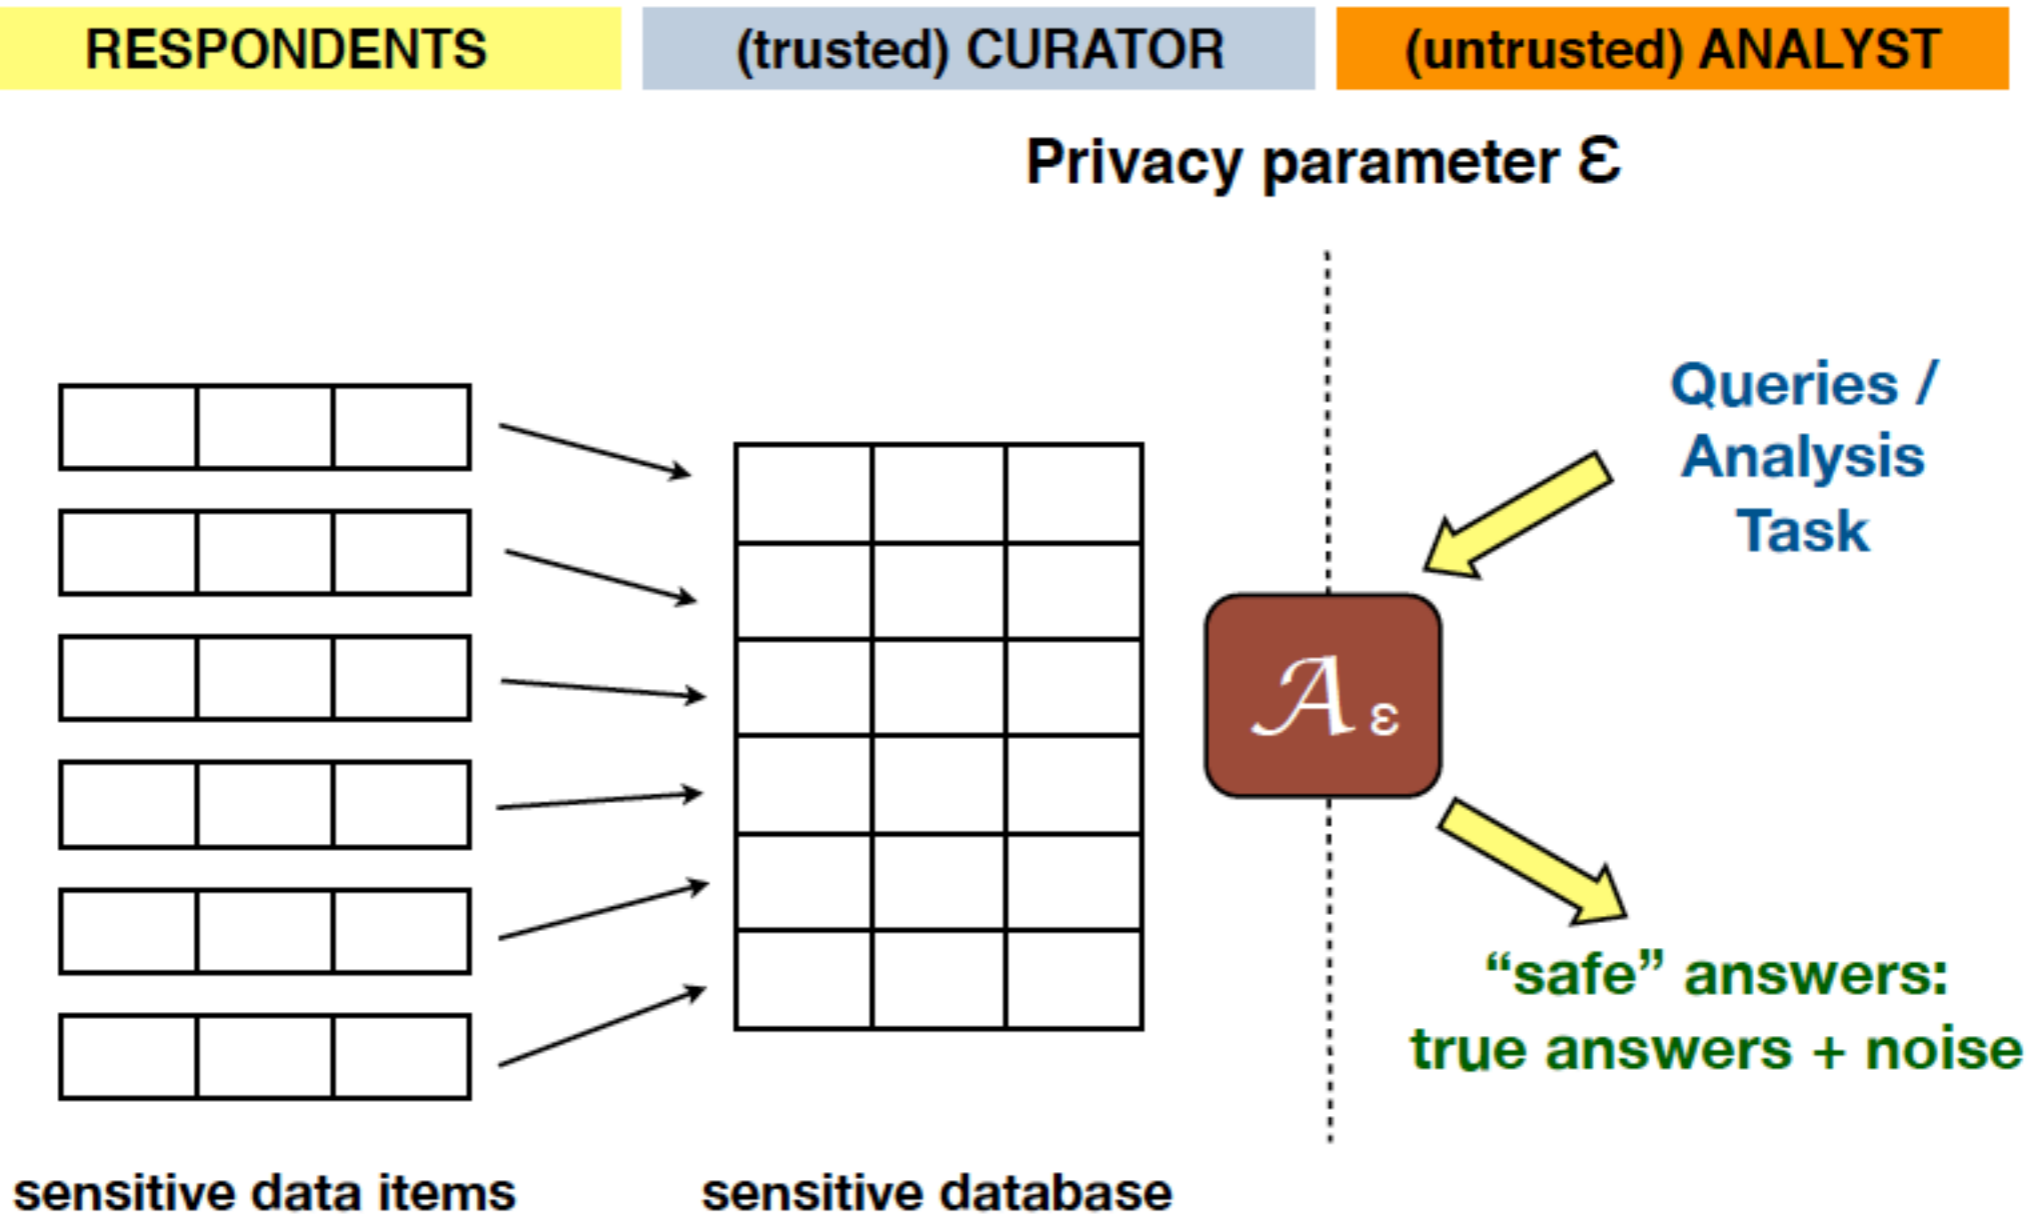
select gender, count(*)
from D group by gender

1 (**disjoint groups**, presence or absence of one tuple impacts only one of the counts)

an arbitrary list of m counting queries

m (no assumptions about the queries, and so a single individual may change the answer of **every query** by 1)

Adding noise

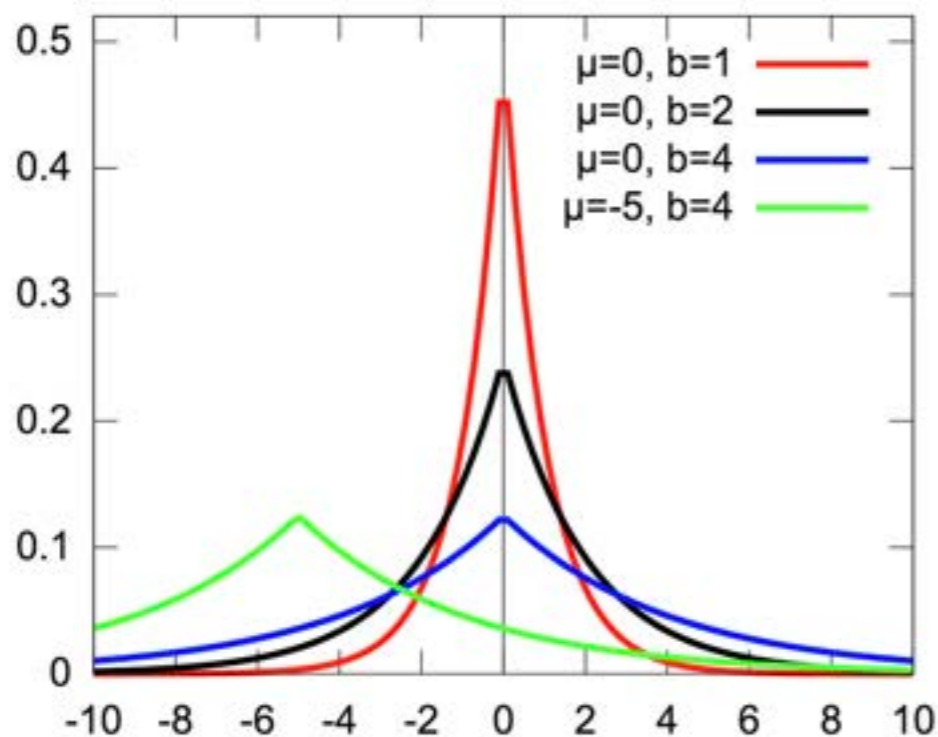


Adding noise

Use the **Laplace mechanism** to answer q in a way that's ϵ -differentially private

$$M(\epsilon) : q(D) + \text{Lap}\left(\frac{\Delta q}{\epsilon}\right)$$

The Laplace distribution, centered at 0 with scale b , denoted **Lap(b)**, is the distribution with probability density function:

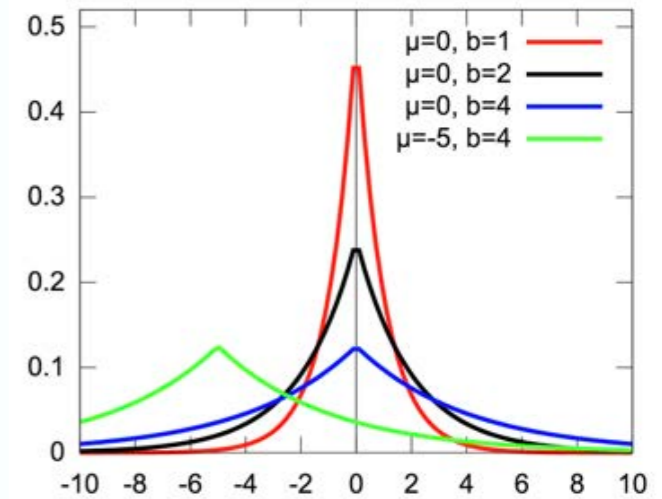
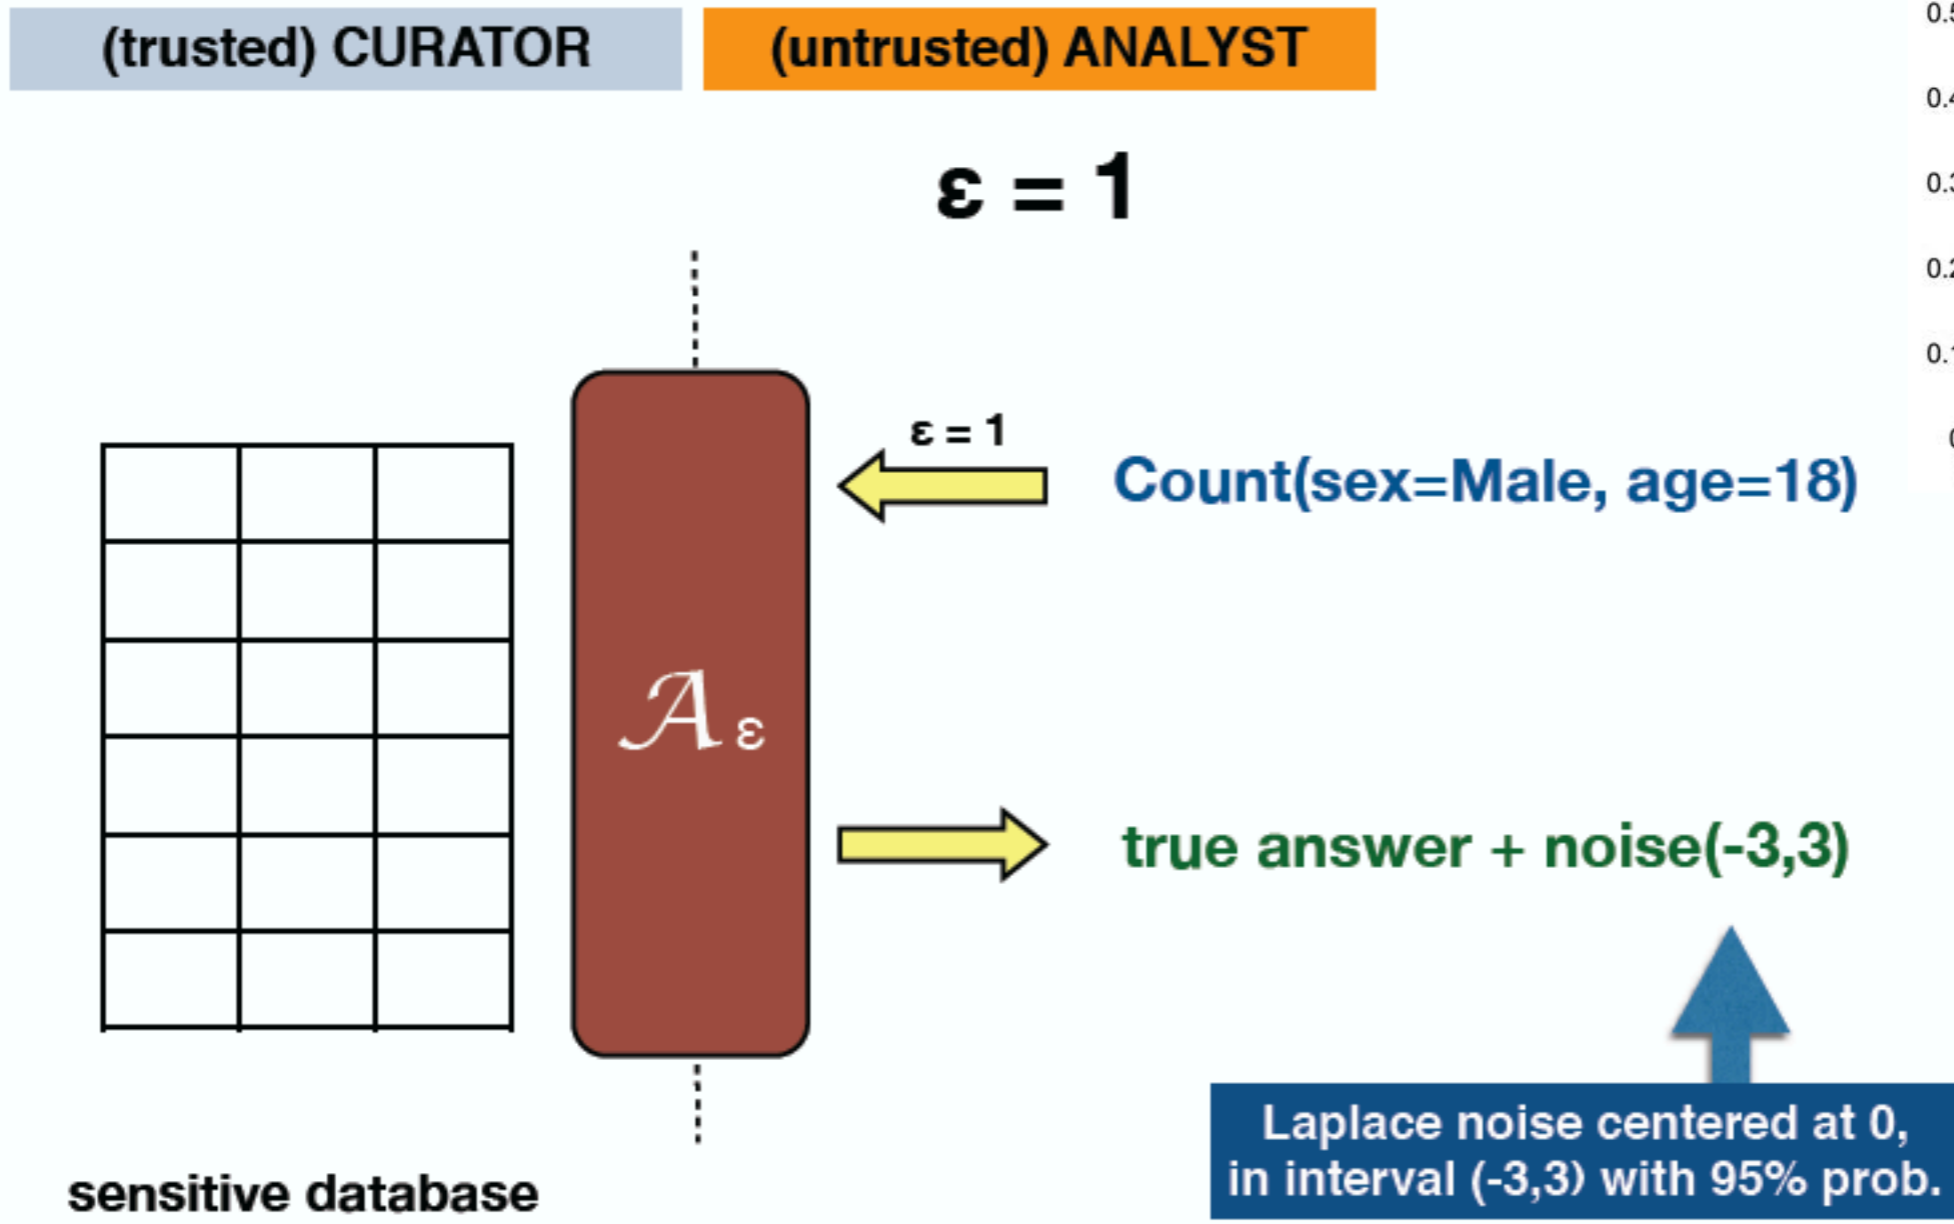


https://en.wikipedia.org/wiki/Laplace_distribution

fix sensitivity Δq , verify that more noise is added for lower ϵ

↓ lower ϵ = stronger privacy ↑

Adding noise



Query sensitivity

The ℓ_1 sensitivity of a query q , denoted Δq , is the maximum difference in the result of that query on a pair of neighboring databases

$$\Delta q = \max_{D, D'} |q(D) - q(D')|$$

query q

query sensitivity Δq

parallel composition

select gender, count(*)
from D group by gender

1 (**disjoint groups**, presence or absence of one tuple impacts only one of the counts)

sequential composition

an arbitrary list of m counting queries

m (no assumptions about the queries, and so a single individual may change the answer of **every query** by 1)

Sequential composition

- Consider 4 queries executed in sequence
 - Q1: select count(*) from D under $\epsilon_1 = 0.5$
 - Q2: select count(*) from D where sex = Male under $\epsilon_2 = 0.2$
 - Q3: select count(*) from D where sex = Female under $\epsilon_3 = 0.25$
 - Q4: select count(*) from D where age > 20 under $\epsilon_4 = 0.25$
- $\epsilon = \epsilon_1 + \epsilon_2 + \epsilon_3 + \epsilon_4 = 1.2$ That is: all queries together are ϵ -differentially private for $\epsilon = 1.2$. **Can we make a stronger guarantee?**
- This works because **Laplace noise is additive**

More generally: set a **cumulative privacy budget**, and split it between all queries, pre-processing, other data manipulation steps of the pipeline

Parallel composition

- If the inputs are disjoint, then the result is ϵ -differentially private for $\epsilon = \max(\epsilon_1, \dots, \epsilon_k)$
 - Q1: select count(*) from D under $\epsilon_1 = 0.5$
 - Q2: select count(*) from D where sex = Male under $\epsilon_2 = 0.2$
 - Q3: select count(*) from D where sex = Female under $\epsilon_3 = 0.25$
 - Q4: select count(*) from D where age > 20 under $\epsilon_4 = 0.25$
- $\epsilon = \epsilon_1 + \max(\epsilon_2, \epsilon_3) + \epsilon_4 = 1$ That is: all queries together are ϵ -differentially private for $\epsilon = 1$.

Composition and consistency

- Consider again 4 queries executed in sequence
 - Q1: select count(*) from D under $\varepsilon_1 = 0.5$ returns **2005**
 - Q2: select count(*) from D where sex = Male under $\varepsilon_2 = 0.2$ returns **1001**
 - Q3: select count(*) from D where sex = Female under $\varepsilon_3 = 0.25$ returns **995**
 - Q4: select count(*) from D where age > 20 under $\varepsilon_4 = 0.25$ returns **1789**

Assuming that there are 2 genders in D, Male and Female, there is **no database consistent with these statistics!**

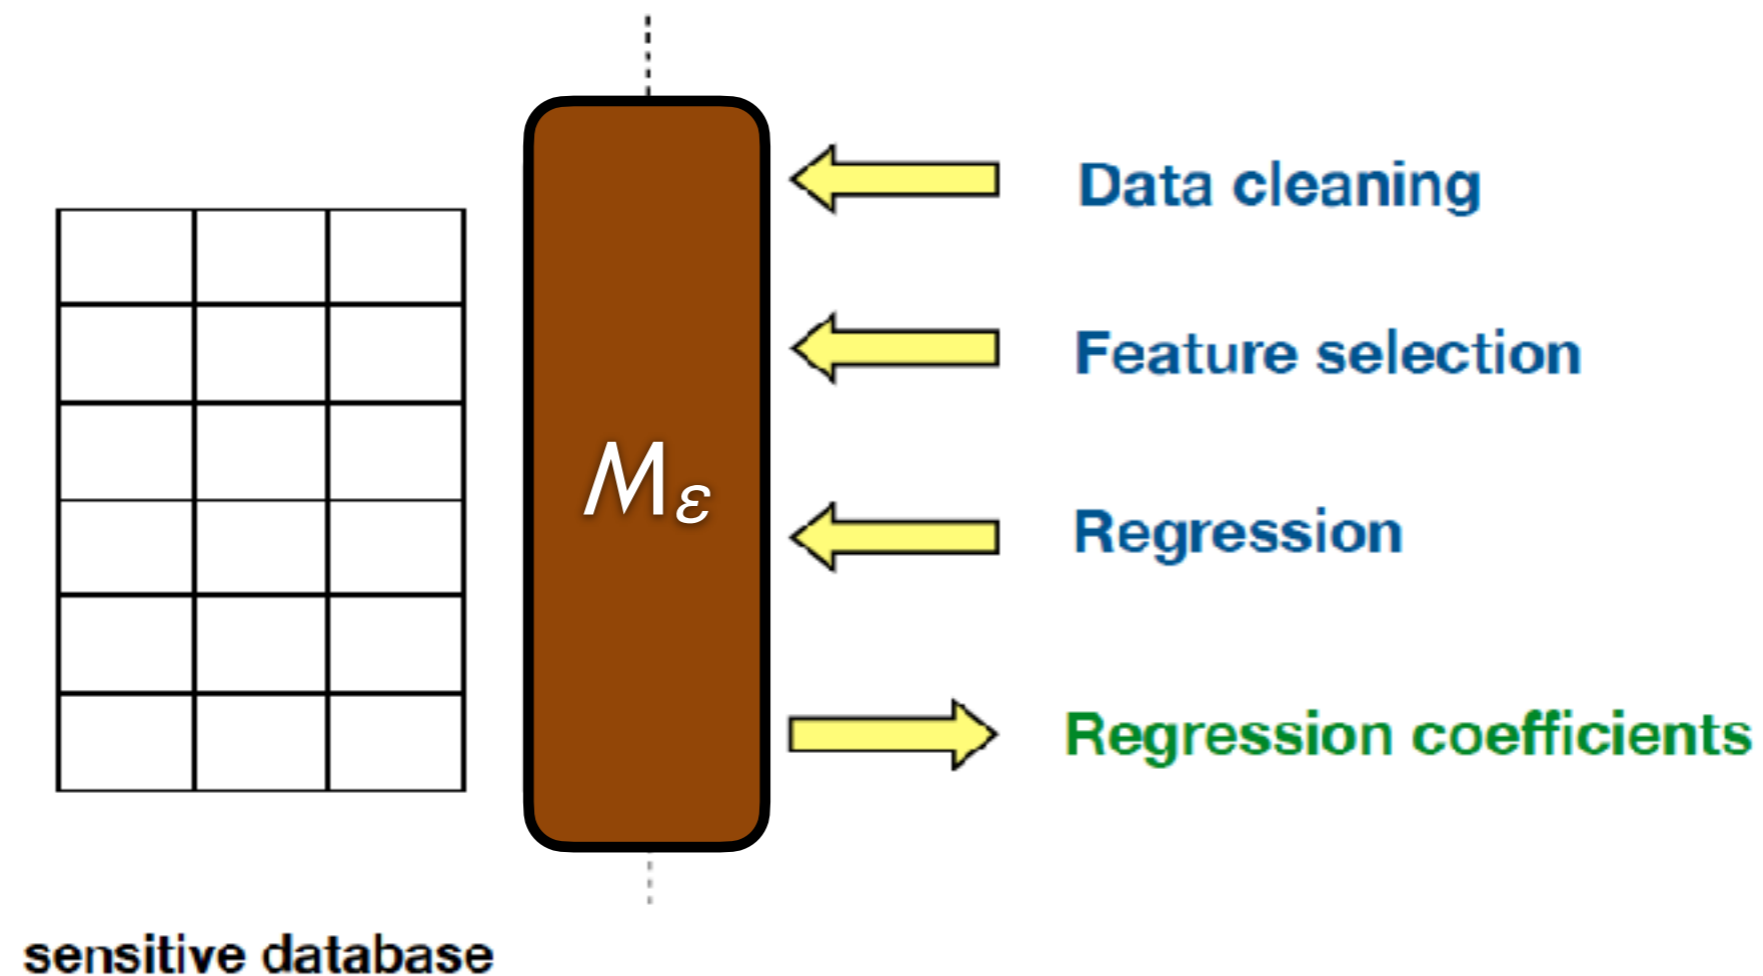
Also don't want any negative counts + may want to impose datatype checks, e.g., no working adults with age = 5 etc.

Entire workflow must be DP

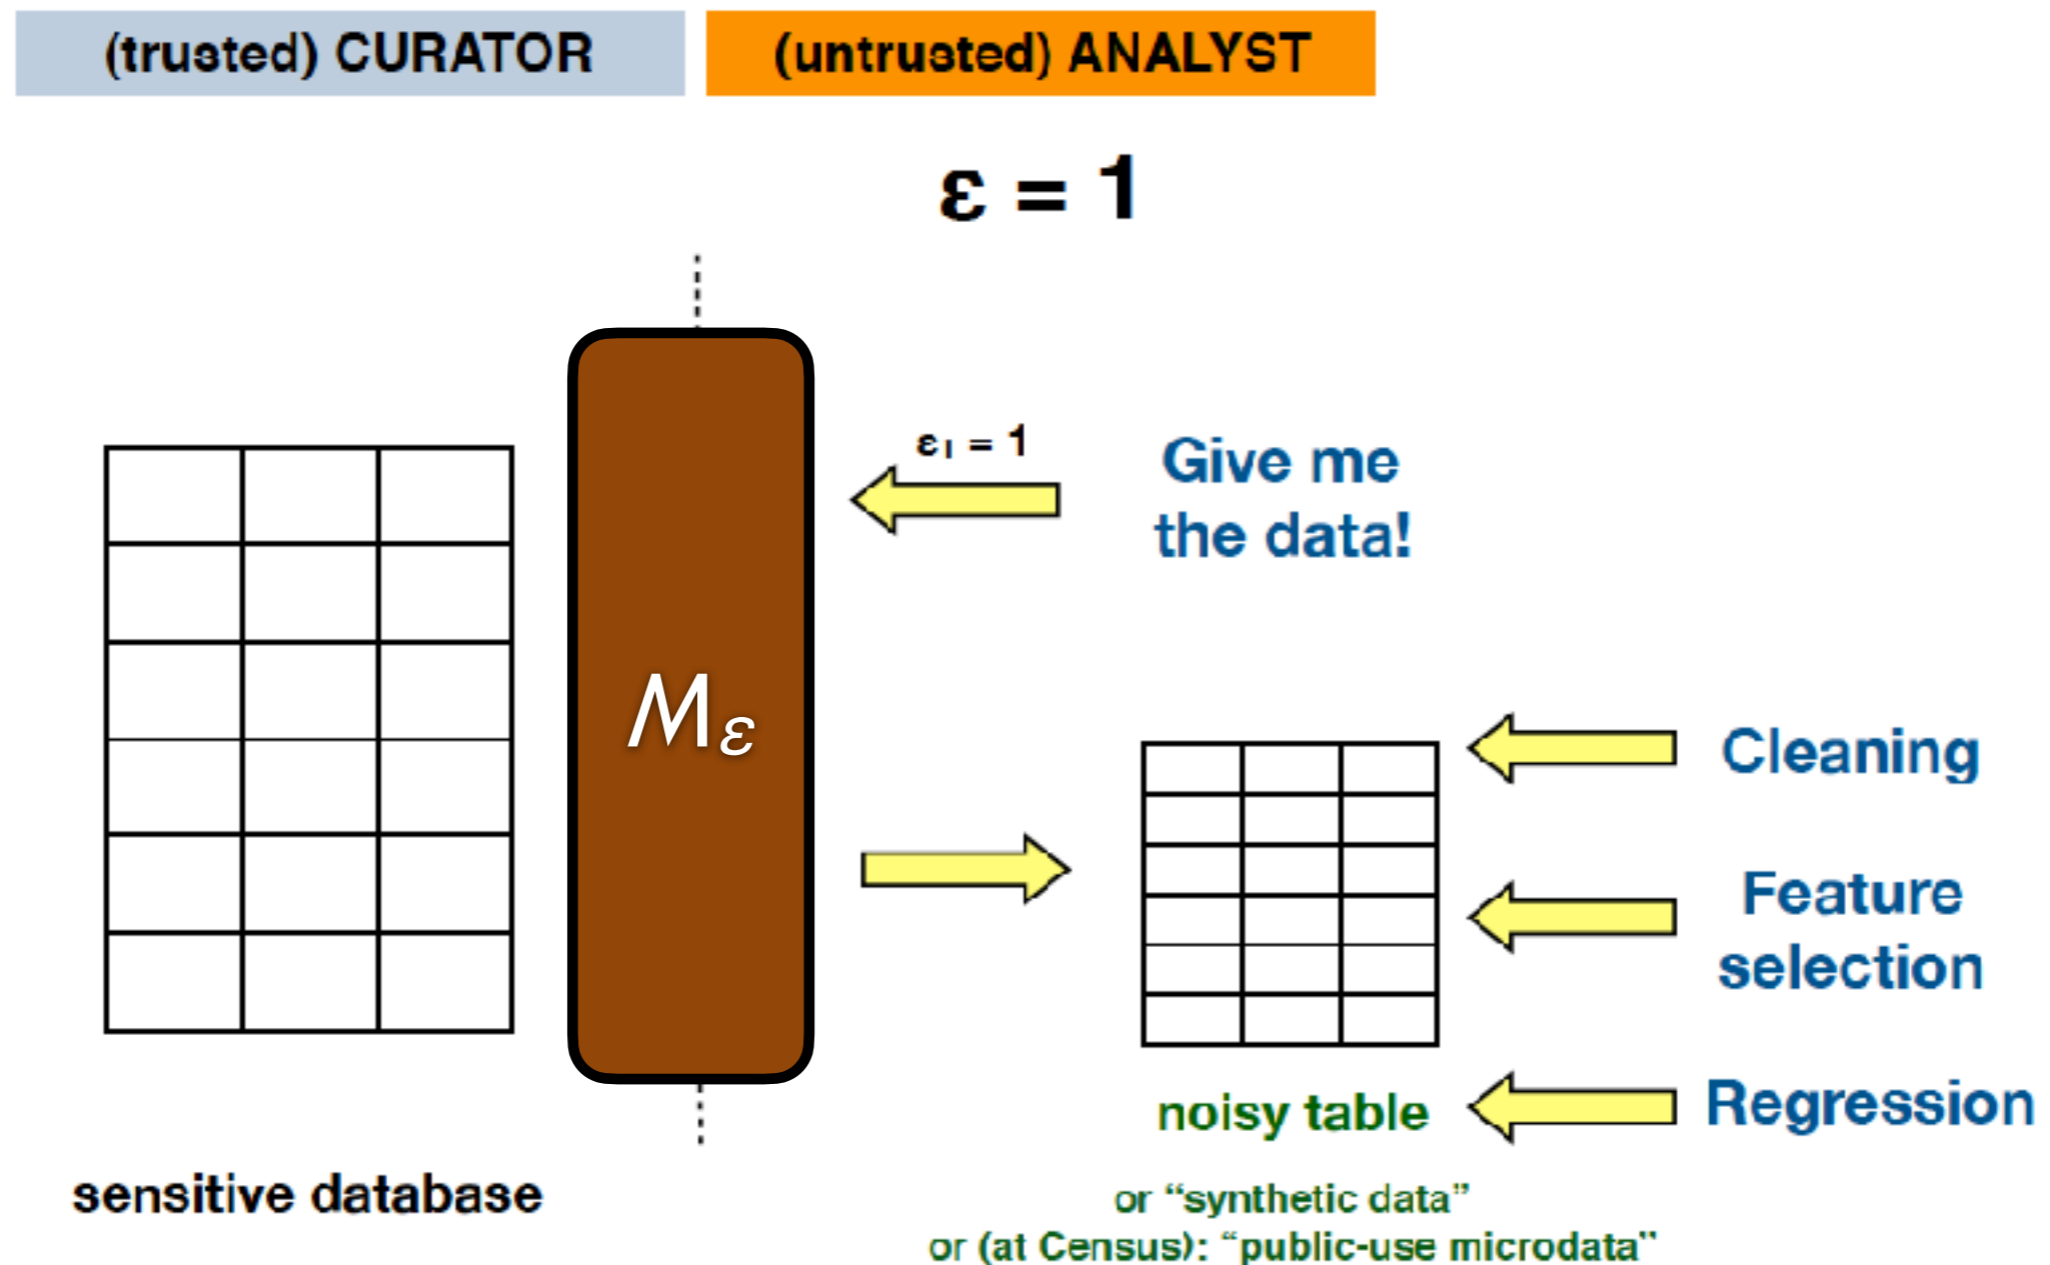
(trusted) CURATOR

(untrusted) ANALYST

$\epsilon = 1$



Privacy-preserving synthetic data



**DP *synthetic*
data *generation***

DP synthetic data

Lots of advantages

- Consistency is not an issue
- Analysts can treat synthetic data as a regular dataset, run existing tools
- No need to worry about the privacy budget
- Can answer as many queries as they want, and any kind of a query they want, including record-level queries

What's the catch?

Recall the Fundamental Law of Information Recovery. It tells us that we cannot answer all these queries accurately and still preserve privacy!

Therefore, when releasing synthetic data, we need to document it with which queries it supports well

Data Synthesizer



input

uid	sex	race	MarriageSta	DateOfBirth	age	lev	fel	conv	decile	score
1	1	0	1	4/18/47	69	0	0	0	1	
2	2	0	2	1/22/82	34	0	0	0	8	
3	3	0	2	5/14/91	24	0	0	0	4	
4	3	0	2	1/21/93	23	0	0	0	8	
5	4	0	1	2/22/73	43	0	0	0	1	
6	5	0	1	3/22/71	44	0	0	0	1	
7	6	0	1	7/23/74	41	0	0	0	4	
8	7	0	3	2/25/73	43	0	0	0	4	
9	8	0	1	8/10/94	21	0	0	0	8	
10	9	0	3	6/1/88	27	0	0	0	4	
11	10	0	3	8/22/78	37	0	0	0	1	
12	11	1	3	12/27/74	41	0	0	0	4	
13	12	0	2	6/14/68	47	0	0	0	1	
14	13	1	3	3/25/85	31	0	0	0	8	
15	14	0	2	1/25/79	37	0	0	0	1	
16	15	0	4	4/22/90	25	0	0	0	10	
17	16	0	2	12/24/84	31	0	0	0	5	
18	17	0	3	1/8/85	31	0	0	0	3	
19	18	0	2	6/28/51	64	0	0	0	6	
20	19	0	2	11/29/94	21	0	0	0	9	
21	20	0	2	8/6/88	27	0	0	0	2	
22	21	0	3	3/22/95	21	0	0	0	4	
23	22	1	3	1/23/92	24	0	0	0	4	
24	23	0	4	1/10/75	43	0	0	0	1	
25	24	0	1	8/24/83	32	0	0	0	3	
26	25	0	1	2/8/89	27	0	0	0	3	
27	26	0	2	5/3/79	36	0	0	0	3	
28	27	1	3	2/28/82	36	0	0	0	3	

Data
Describer



summary

age	int	min=23 max=60	32% mis	
name	str	length 10 to 98	no mis	
sex	str	cat	10% mis	

Data
Generator



output

uid	sex	race	MarriageSta	DateOfBirth	age	lev	fel	conv	decile	score
1	1	0	1	4/18/47	69	0	0	0	1	
2	2	0	2	1/22/82	34	0	0	0	8	
3	3	0	2	5/14/91	24	0	0	0	4	
4	3	0	2	1/21/93	23	0	0	0	8	
5	4	0	1	2/22/73	43	0	0	0	1	
6	5	0	1	3/22/71	44	0	0	0	1	
7	6	0	1	7/23/74	41	0	0	0	4	
8	7	0	3	2/25/73	43	0	0	0	4	
9	8	0	1	8/10/94	21	0	0	0	8	
10	9	0	3	6/1/88	27	0	0	0	4	
11	10	0	3	8/22/78	37	0	0	0	1	
12	11	1	3	12/27/74	41	0	0	0	4	
13	12	0	2	6/14/68	47	0	0	0	1	
14	13	1	3	3/25/85	31	0	0	0	8	
15	14	0	2	1/25/79	37	0	0	0	1	
16	15	0	4	4/22/90	25	0	0	0	10	
17	16	0	2	12/24/84	31	0	0	0	5	
18	17	0	3	1/8/85	31	0	0	0	3	
19	18	0	2	6/28/51	64	0	0	0	6	
20	19	0	2	11/29/94	21	0	0	0	9	
21	20	0	2	8/6/88	27	0	0	0	2	
22	21	0	3	3/22/95	21	0	0	0	4	
23	22	1	3	1/23/92	24	0	0	0	4	
24	23	0	4	1/10/75	43	0	0	0	1	
25	24	0	1	8/24/83	32	0	0	0	3	
26	25	0	1	2/8/89	27	0	0	0	3	
27	26	0	2	5/3/79	36	0	0	0	3	
28	27	1	3	2/28/82	36	0	0	0	3	

Model
Inspector



comparison

	before	after
age	int min=23 max=60 32% mis	int min=23 max=60 32% mis
name	str length 10 to 98 no mis	str length 10 to 98 no mis
sex	str cat 10% mis	str cat 10% mis

Data Synthesizer



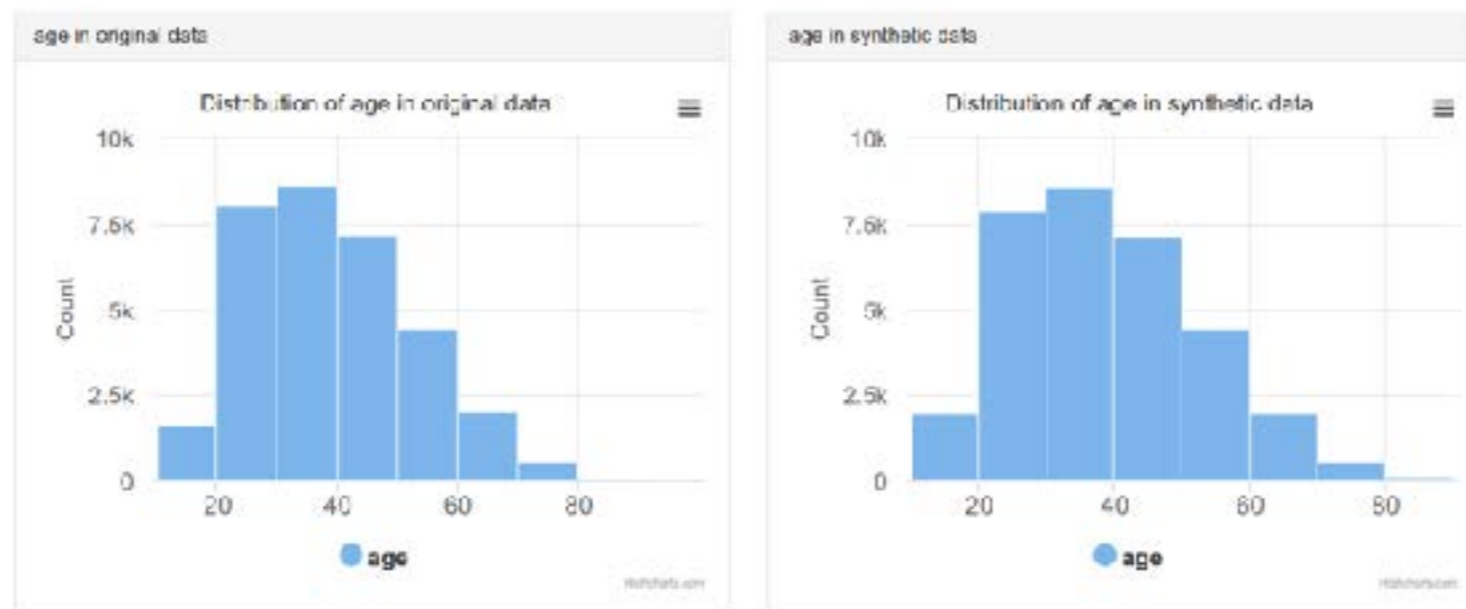
- Main goal: **usability first**
 - user is the data owner
 - the tool picks up data types from the input file: categorical / string / numerical (integer, float) / date-time
 - the tool computes the frequency of missing values per attribute
 - user can then inspect the result, over-ride what was learned about an attribute, e.g., whether it's categorical, or what its datatype is
- The tool generates an output dataset of a specified size, in one of three modes
 - **random** - type-consistent random output
 - **independent attribute** - learn a noisy histogram for each attribute
 - **correlated attribute** - learn a noisy Bayesian network (BN)

Data Synthesizer: Independent attributes



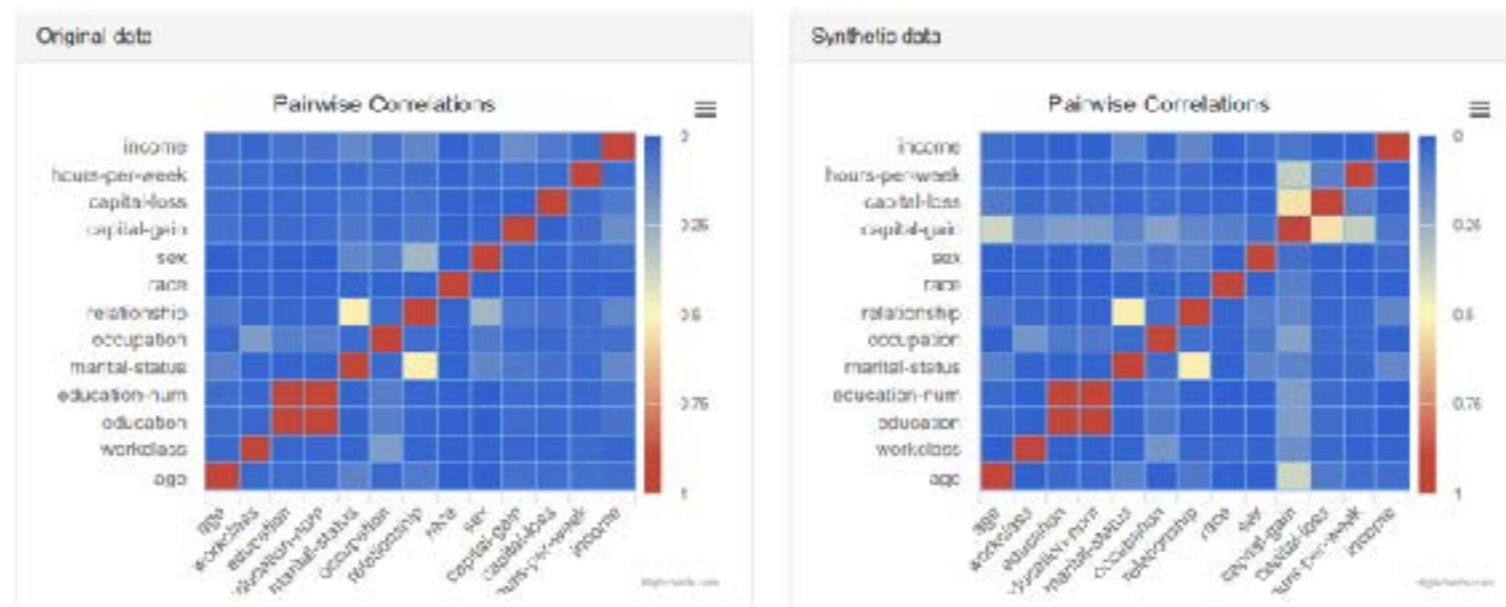
Given the over-all privacy budget ϵ , and an input dataset of size n . Allocate ϵ/d of the budget to each attribute \mathbf{A}_i in $\{\mathbf{A}_1, \dots, \mathbf{A}_d\}$. Then for each attribute:

- Compute the i th histogram with t bins ($t=20$ by default), with query q_i
- The sensitivity Δq_i of this (or any other) histogram query is $2/n$ **Why?**
- So, each bin's noisy probability is computed by adding $Lap\left(\frac{2d}{\epsilon n}\right)$



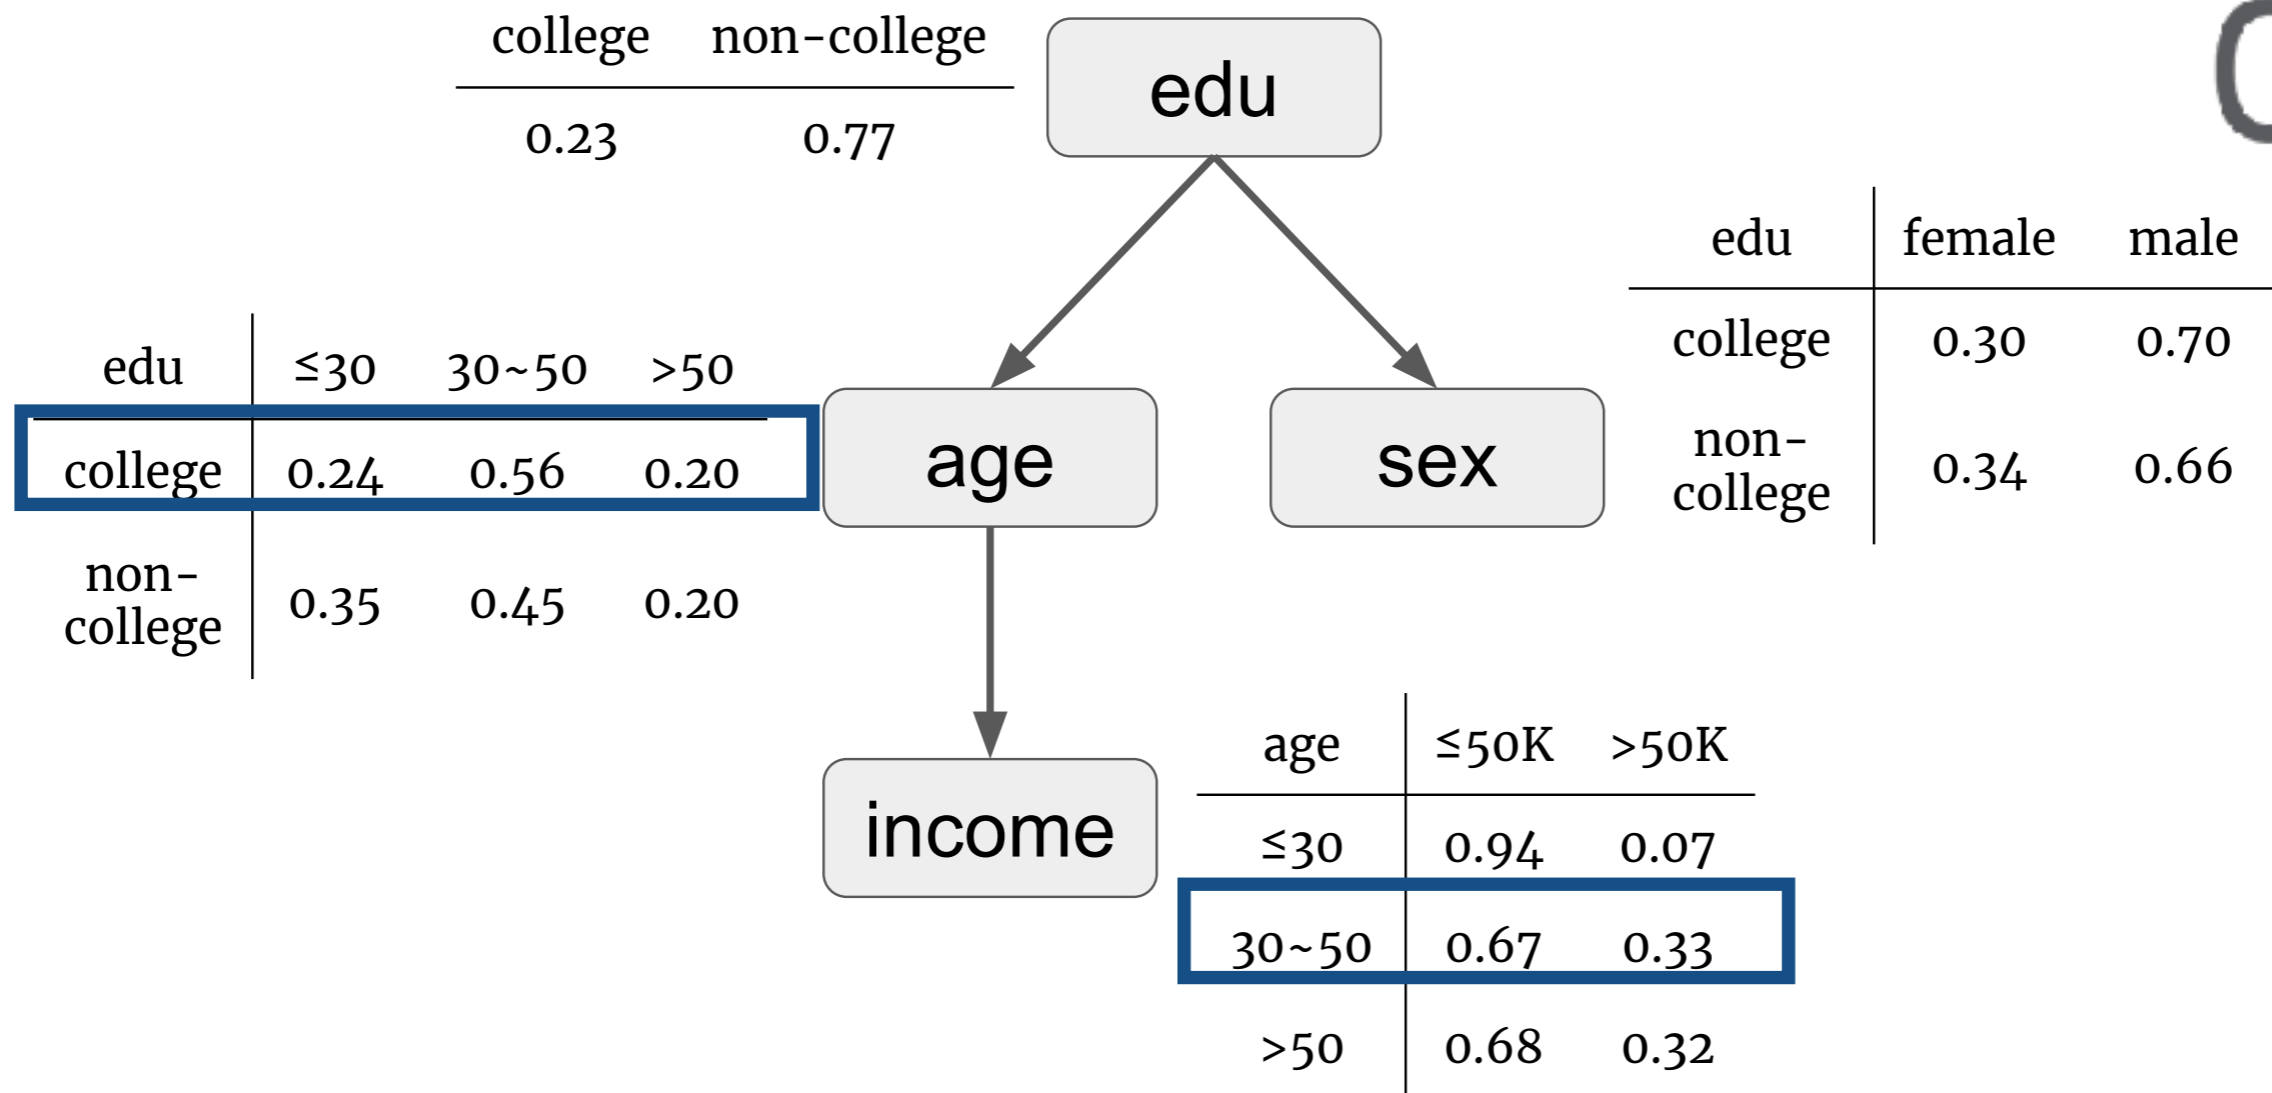
Data Synthesizer: Correlated attributes

- Learn a differentially private Bayesian network (BN)
- Use the method called **PrivBayes** [Zhang, Cormode, Procopiuc, Srivastava, Xiao, 2016]
- Privacy budget is split equally between (a) network structure computation and (b) populating the conditional probability tables of each BN node
- User inputs privacy budget ϵ and the maximum number of parents for a BN node k - you'll play with these settings as part of HW2
- The tool treats a missing attribute value as one of the values in the attribute's domain (not shown in the examples in the next two slides)



Data Synthesizer: Correlated attributes

$K=1$ not a causal DAG, a regular Bayesian network!



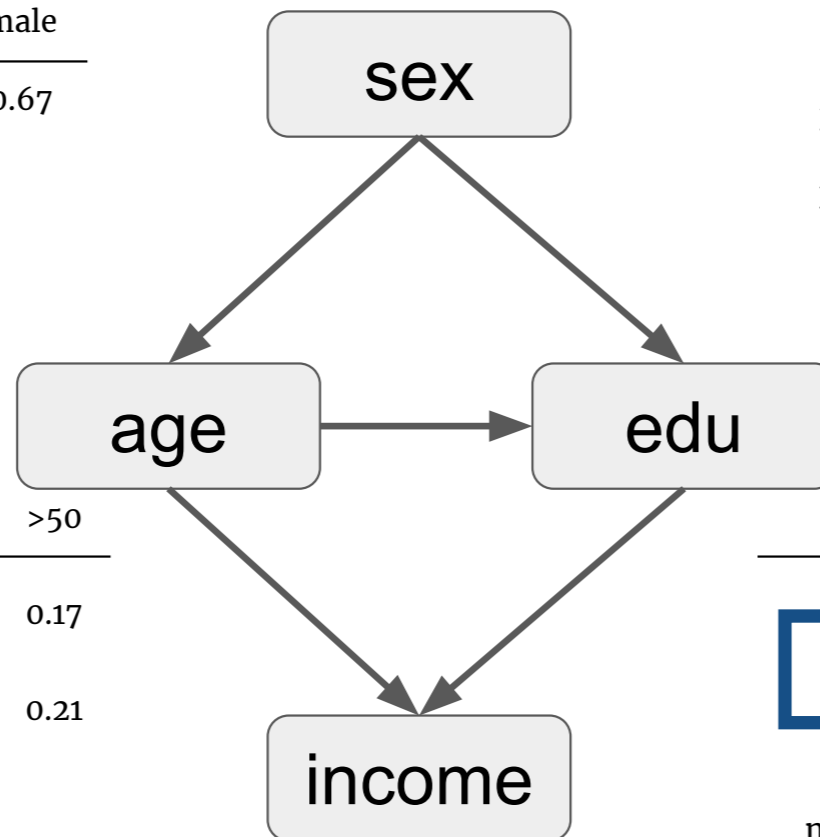
Data Synthesizer: Correlated attributes

not a causal DAG, a regular Bayesian network!

K=2



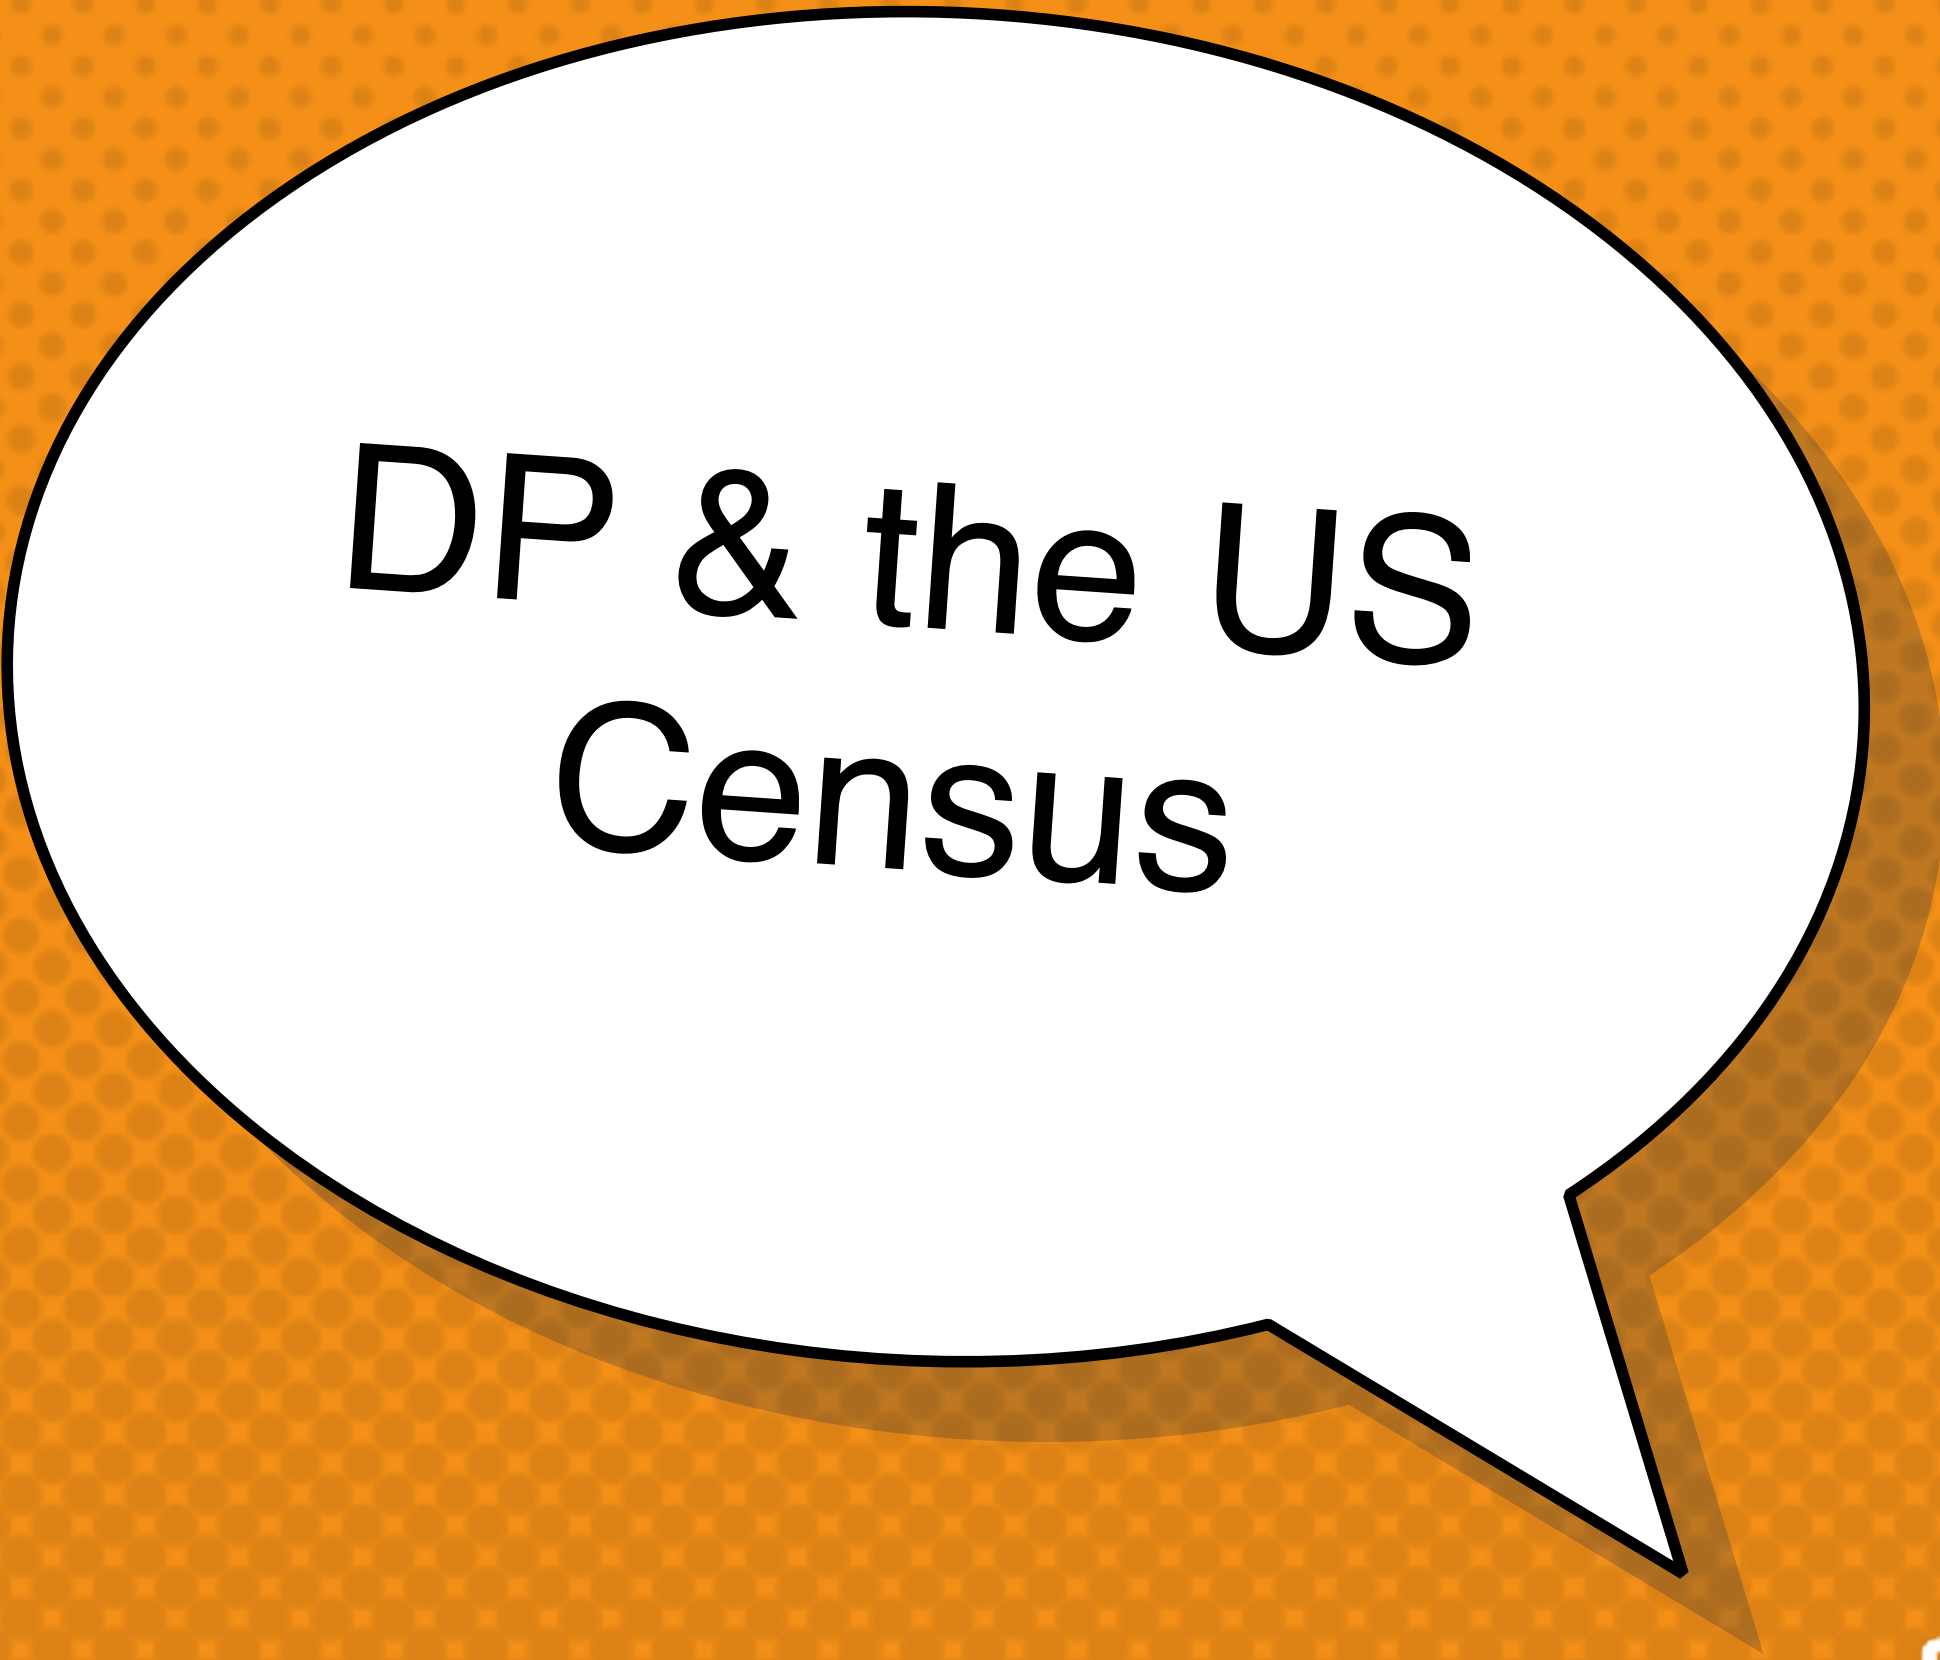
female	male
0.33	0.67



sex	≤30	30~50	>50
female	0.40	0.43	0.17
male	0.29	0.59	0.21

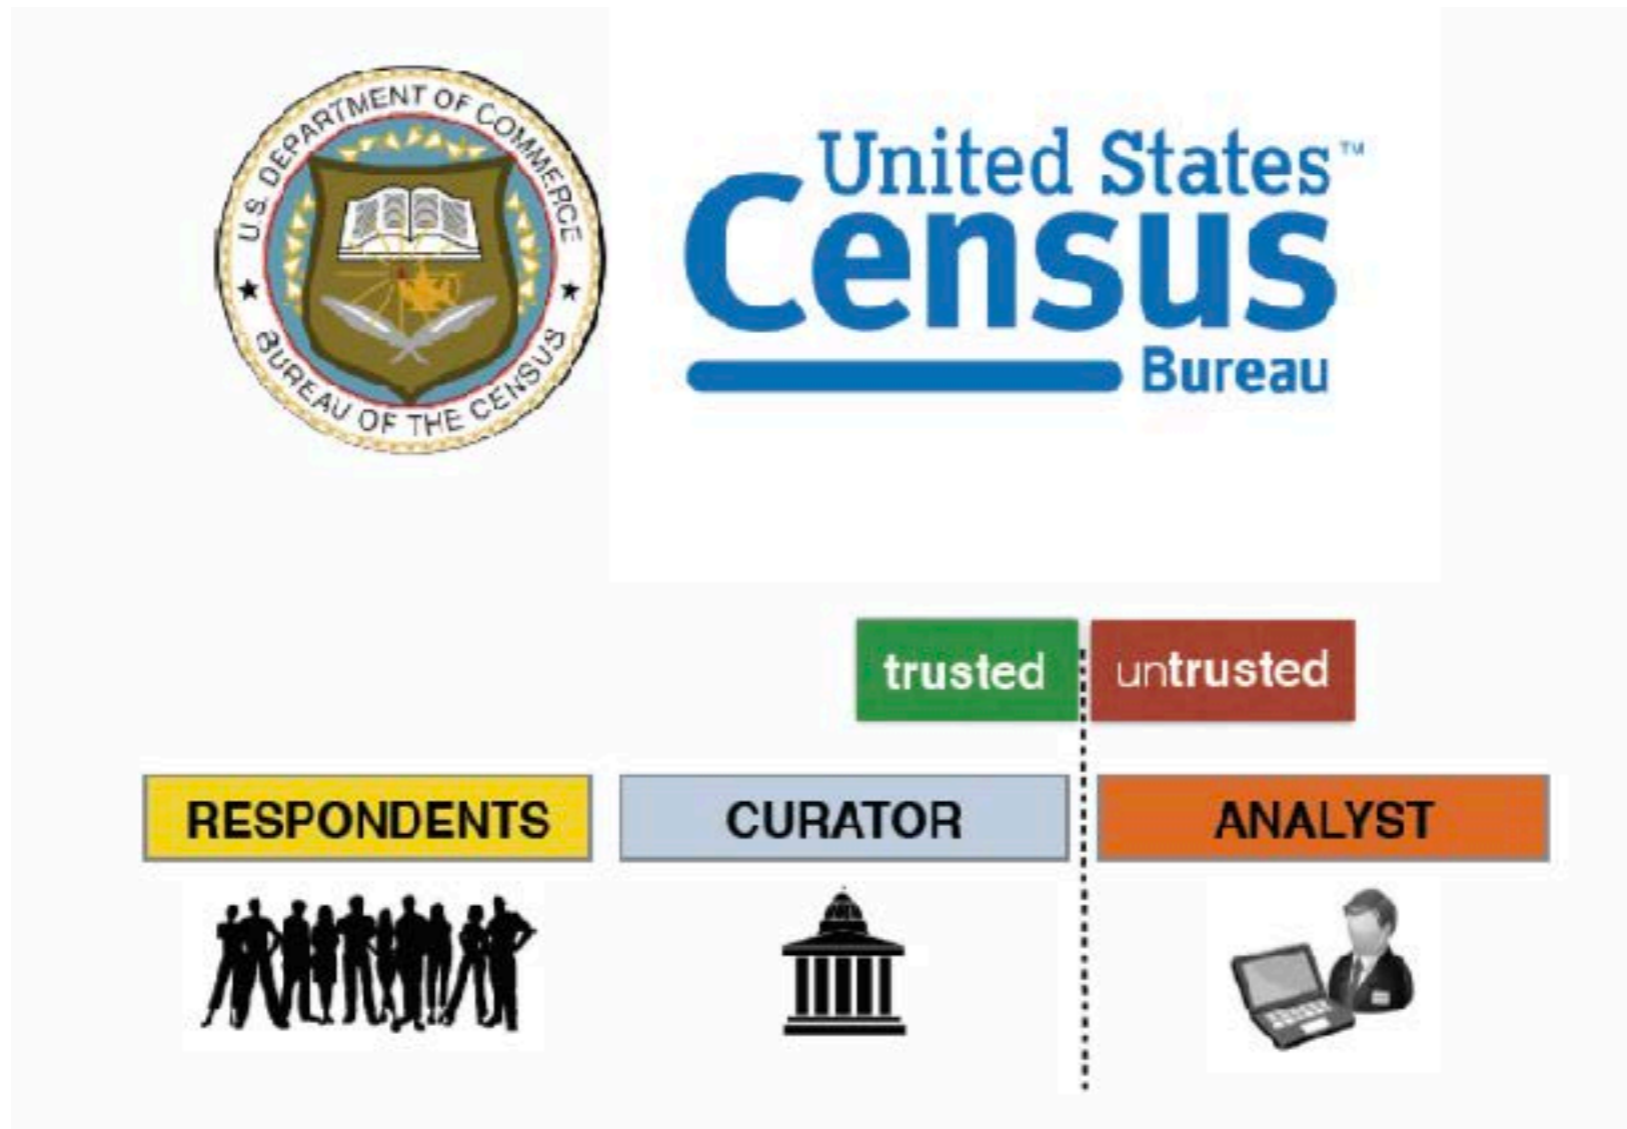
age	sex	college	non-college
≤30	female	0.18	0.82
≤30	male	0.16	0.84
30~50	female	0.25	0.75
30~50	male	0.28	0.72
>50	female	0.17	0.83
>50	male	0.25	0.75

edu	age	≤50K	>50K
college	≤30	0.83	0.17
college	30~50	0.45	0.55
college	>50	0.41	0.59
non-college	≤30	0.96	0.04
non-college	30~50	0.76	0.24
non-college	>50	0.75	0.25



**DP & the US
Census**

Differential privacy in the field

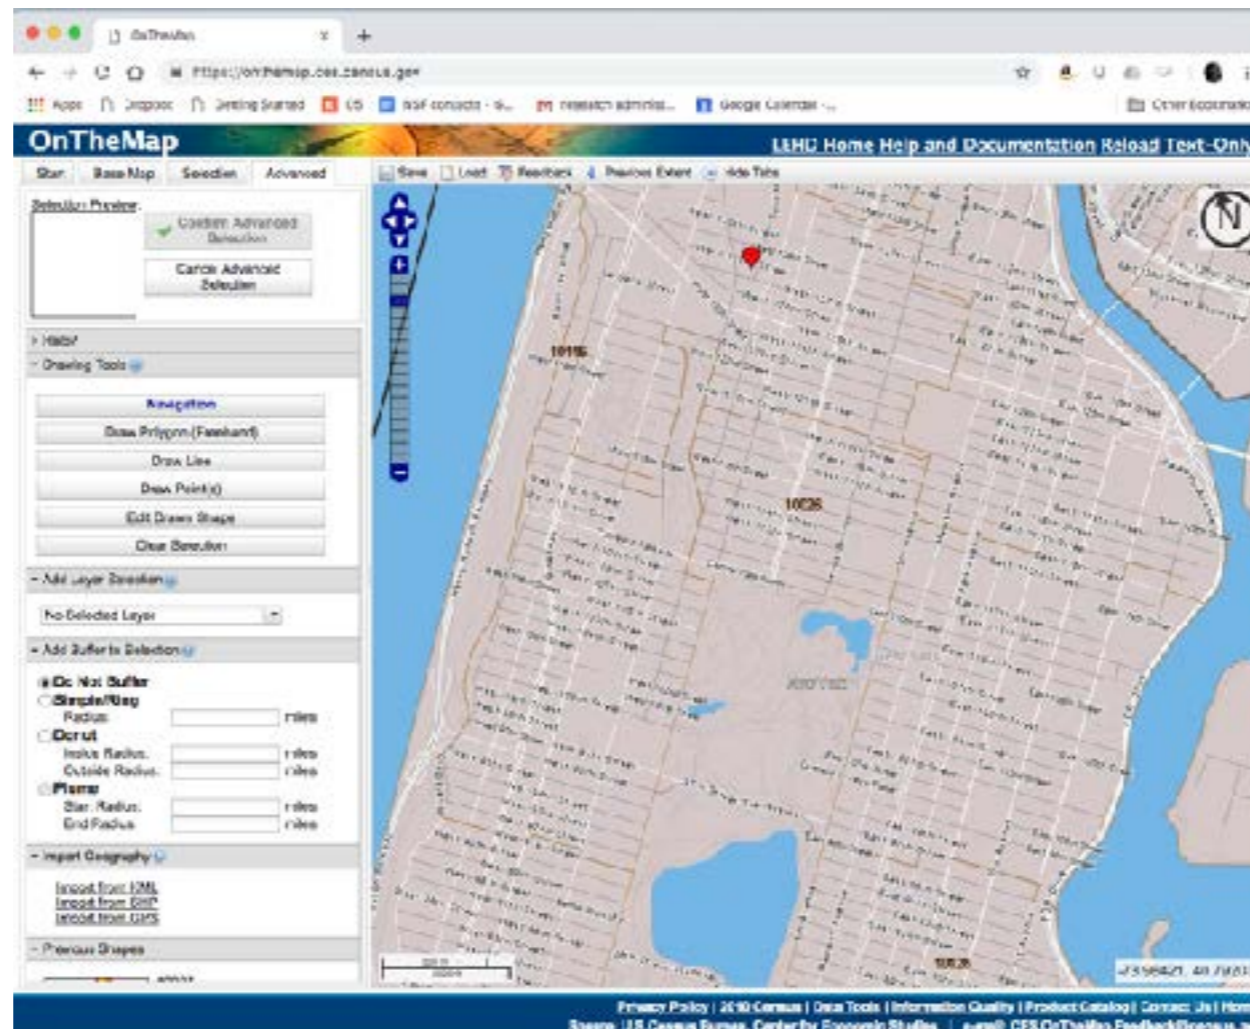


Decennial Census 2020

Differential privacy in the field

First adoption by the US Census Bureau:

OnTheMap (2008), synthetic data about where people in the US live and work



Differential privacy in the field

TheUpshot

To Reduce Privacy Risks, the Census Plans to Report Less Accurate Data

Guaranteeing people's confidentiality has become more of a challenge, but some scholars worry that the new system will ~~impede research.~~

The New York Times

By Mark Hansen

Dec. 5, 2018



A 2018 census test letter mailed to a resident in Providence, R.I. The nation's test run of the 2020 Census is in Rhode Island. Michelle R. Smith/Associated Press

At the root of the problem are the tables of aggregate statistics that the bureau publishes. There are hundreds of tables — sex by age, say, or ethnicity by race — summarizing the population at several levels of geography, from areas the size of a city block all the way up to the level of a state or the nation. In 2010, the bureau released tables with nearly eight billion numbers in all. That was about 25 numbers for each person living in the United States, even though Americans were asked only 10 questions about themselves. In other words, the tables were generated in so many ways that the Census Bureau ended up releasing more data in aggregate than it had collected in the first place.

Reconstruction attack: an example

TABLE 1: FICTIONAL STATISTICAL DATA FOR A FICTIONAL BLOCK

STATISTIC	GROUP	AGE		
		COUNT	MEDIAN	MEAN
1A	total population	7	30	38
2A	female	4	30	33.5
2B	male	3	30	44
2C	black or African American	4	51	48.5
2D	white	3	24	24
3A	single adults	(D)	(D)	(D)
3B	married adults	4	51	54
4A	black or African American female	3	36	36.7
4B	black or African American male	(D)	(D)	(D)
4C	white male	(D)	(D)	(D)
4D	white female	(D)	(D)	(D)
5A	persons under 5 years	(D)	(D)	(D)
5B	persons under 18 years	(D)	(D)	(D)
5C	persons 64 years or over	(D)	(D)	(D)

Note: Married persons must be 15 or over

Reconstruction attack: an example

Let's assume that the oldest person is 125 years old, and that everyone's age is different. How many possible age combinations are there?

$$\binom{125}{3} = 317,750$$

But only 40 combinations have median = 30 and mean = 44

Idea: extract all such constraints, represent them as a mathematical model, have an automated solver find a solution.

TABLE 2: POSSIBLE AGES FOR A MEDIAN OF 30 AND MEAN OF 44

A	B	C	A	B	C	A	B	C
1	30	101	11	30	91	21	30	81
2	30	100	12	30	90	22	30	80
3	30	99	13	30	89	23	30	79
4	30	98	14	30	88	24	30	78
5	30	97	15	30	87	25	30	77
6	30	96	16	30	86	26	30	76
7	30	95	17	30	85	27	30	75
8	30	94	18	30	84	28	30	74
9	30	93	19	30	83	29	30	73
10	30	92	20	30	82	30	30	72

What does the law say?

Title 13 of U.S. Code authorizes data collection and publication of statistics by the Census Bureau.

Section 9 of Title 13 requires privacy protections: “Neither the Secretary, nor any other officer or employee of the Department of Commerce or bureau or agency thereof, ... may ... make **any publication whereby the data furnished by any particular establishment or individual under this title can be identified**” (Title 13 U.S.C. § 9(a)(2), Public Law 87-813).

In 2002, Congress further clarified the concept of identifiable data: it is prohibited to publish “**any representation of information that permits the identity of the respondent to whom the information applies to be reasonably inferred by either direct or indirect means**” (Pub. L. 107–347, Title V, §502 (4), Dec. 17, 2002, 116 Stat. 2969).

Section 214 of Title 13 outlines penalties: fines up to \$5,000 or imprisonment up to 5 years or both per incident (data item), up to \$250,000 in total.

DP in the 2020 Census: pushback



Implications of Differential Privacy for Census Bureau Data and Research

Task Force on Differential Privacy for Census Data †
Institute for Social Research and Data Innovation (ISRDI)
University of Minnesota

November 2018
Version 2
Working Paper No. 2018-6

- noisy data - **impact on critical decisions**
- difficult to explain differential privacy / privacy budget to the public - **how do we set epsilon?**
- disagreement about whether using differential privacy is legally required
- messaging is difficult to get right “**the result doesn’t change whether or not you participate**” - might discourage participation!

Revealing **characteristics** of individuals vs. their **identity**, is there a distinction?

But the Census collects “generic” **harmless data**, is this really a big deal?

What sorts of trade-offs should we be aware of? Who should decide?

**beyond
differential privacy**

The Strava Heat Map

WIRED

JEREMY HSB SECURITY JAN 29, 2018 7:14 PM

The Strava Heat Map and the End of Secrets

The US military is reexamining security policies after fitness tracker data shared on social media revealed bases and patrol routes

”Over the weekend, researchers and journalists raised the alarm about how **anyone can identify secretive military bases and patrol routes** based on public data shared by a “social network for athletes” called Strava.

This past November, the San Francisco-based Strava announced a huge update to its global heat map of user activity that displays 1 billion activities—including running and cycling routes—undertaken by exercise enthusiasts wearing Fitbits or other wearable fitness trackers. [...]

But the biggest danger may come from potential adversaries figuring out “patterns of life,” by tracking and even identifying military or intelligence agency personnel as they go about their duties or head home after deployment. These digital footprints that echo the real-life steps of individuals underscore a greater challenge to governments and ordinary citizens alike: each person’s connection to online services and personal devices makes it increasingly difficult to keep secrets.”



Is genetic data your own?

NEWS | SCIENCE AND POLICY

We will find you: DNA search used to nab Golden State Killer can home in on about 60% of white Americans

Researchers call for limiting how ancestry databases can be used to protect privacy

11 OCT 2018 | BY JOCKLYN HALLER

Science

If you're white, live in the United States, and a distant relative has uploaded their DNA to a public ancestry database, there's a good chance **an internet sleuth can identify you from a DNA sample you left somewhere**. That's the conclusion of a new study, which finds that by combining an anonymous DNA sample with some basic information such as someone's rough age, researchers **could narrow that person's identity to fewer than 20 people by starting with a DNA database of 1.3 million individuals**. [...]

The study was sparked by **the April arrest of the alleged "Golden State Killer," a California man accused of a series of decades-old rapes and murders**. To find him—and more than a dozen other criminal suspects since then—law enforcement agencies first test a crime scene DNA sample, which could be old blood, hair, or semen, for hundreds of thousands of DNA markers—signposts along the genome that vary among people, but whose identity in many cases are shared with blood relatives. They then upload the DNA data to **GEDmatch, a free online database where anyone can share their data from consumer DNA testing companies such as 23andMe and Ancestry.com to search for relatives who have submitted their DNA**. Searching GEDMatch's nearly 1 million profiles revealed several relatives who were the equivalent to third cousins to the crime scene DNA linked to the Golden State Killer. Other information such as genealogical records, approximate age, and crime locations then allowed the sleuths to home in on a single person.

<https://www.science.org/content/article/we-will-find-you-dna-search-used-nab-golden-state-killer-can-home-about-60-white>

r/ai

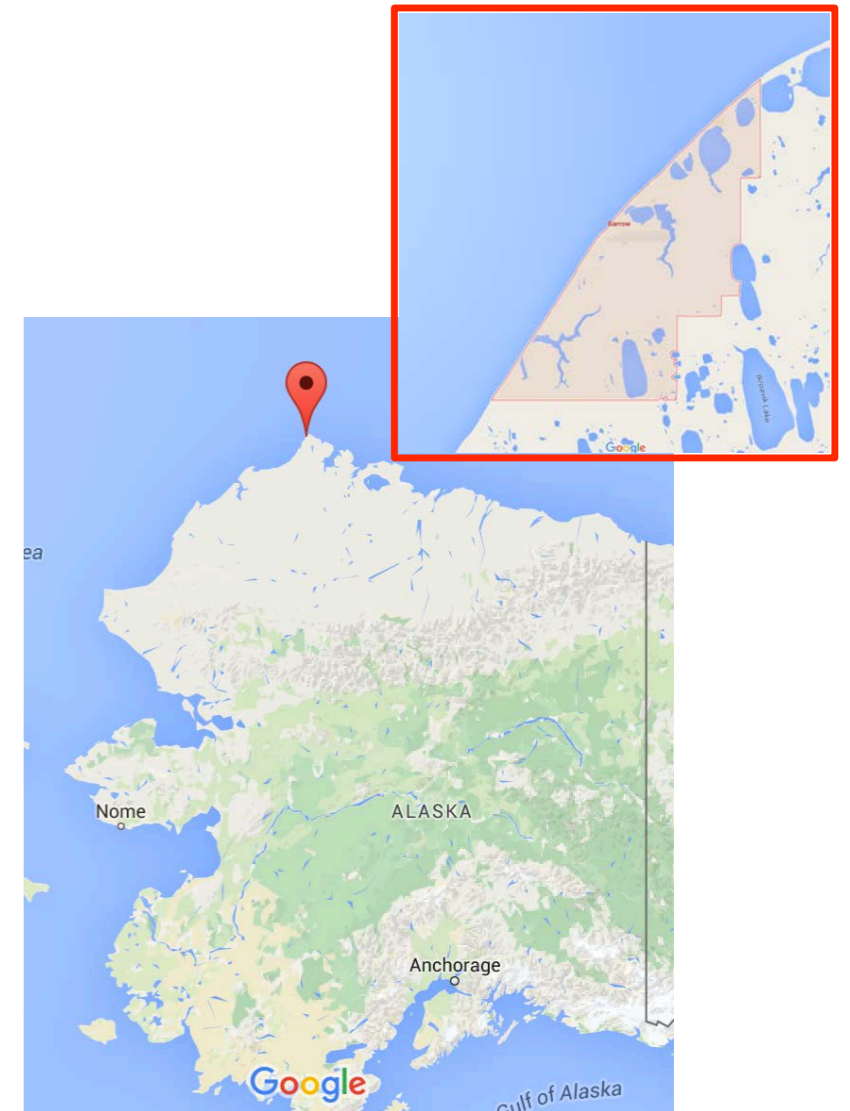
*the origins of
data protection*

Barrow, Alaska, 1979

Native leaders and city officials, worried about drinking and associated violence in their community, **invited a group of sociology researchers** to assess the problem and work with them to devise solutions.

Methodology:

- 10% representative sample (N=88) of everyone over the age of 15 using a 1972 demographic survey
- Interviewed on attitudes and values about use of alcohol
- Obtained psychological histories & drinking behavior
- Given the Michigan Alcoholism Screening Test
- Asked to draw a picture of a person (to determine cultural identity)



Study “results”

Alcohol Plagues Eskimos; Alcoholism Plagues Eskimo Village

DAVA SOBEL ();
January 22, 1980,
, Section Science Times, Page C1, Column , words

 PERMISSIONS

[DISPLAYING ABSTRACT]

THE Inupiat Eskimos of Alaska's North Slope, whose culture has been overwhelmed by energy development activities, are "practically committing suicide" by mass alcoholism, University of Pennsylvania researchers said here yesterday. The alcoholism rate is 72 percent among the 2,000 Eskimo men and women in the village of Barrow, where violence is becoming the ...

At the conclusion of the study researchers formulated a report entitled “**The Inupiat, Economics and Alcohol on the Alaskan North Slope**”, released **simultaneously** at a press release and to the Barrow community.

The press release was picked up by the New York Times, who ran a front page story entitled “**Alcohol Plagues Eskimos**”

Harms and backlash

Article Preview

Eskimos Irate Over Alcoholism Study

[DISPLAYING ABSTRACT]

BARROW, ALASKA HOT tempers and tension arising from a scientific report that found a high rate of alcoholism in this predominantly Eskimo community have abated somewhat after two days of meetings here at the northernmost point of Alaska.

 PERMISSIONS

Study **results were revealed** in the context of a press conference that was held far from the Native village, and **without the presence, much less the knowledge or consent**, of any community member who might have been able to present any context concerning the socioeconomic conditions of the village.

Study results suggested that **nearly all adults in the community were alcoholics**. In addition to the shame felt by community members, the town's Standard and Poor bond rating suffered as a result, which in turn decreased the tribe's ability to secure funding for much needed projects.

Problems

Edward F. Foulks, M.D., “Misalliances In The Barrow Alcohol Study”

Methodological

- “The authors once again met with the Barrow Technical Advisory Group, who stated their concern that only Natives were studied, and that outsiders in town had not been included.”
- “The **estimates of the frequency of intoxication based on association with the probability of being detained** were termed “ludicrous, both logically and statistically.”

Ethical

- Participants not in control of how their data is used
- Significant harm: social (stigmatization) and financial (bond rating)

can differential privacy help with this?

**need an ethical
framework!**

Responsible Data Science

Anonymity and privacy

Thank you!



NYU

TANDON SCHOOL
OF ENGINEERING



NYU

Center for
Data Science

r/ai center
for
responsible
ai