Differential Privacy: An Introduction

Salil Vadhan
Center for Research on Computation & Society
School of Engineering & Applied Sciences
Harvard University

Privacy Tools for Sharing Research Data
Summer 2015 Orientation

Data Privacy: The Problem

Given a dataset with sensitive information, such as:

- Census data
- Health records
- Social network activity
- Telecommunications data

- Academic research
- Informing policy
- Identifying subjects for drug trial
- Searching for terrorists
- Market analysis
- …

How can we:

- enable “desirable uses” of the data
- while protecting the “privacy” of the data subjects?

????

Approach 1: Encrypt the Data

Name	Sex	Blood	…	HIV?
Chen	F	B	…	Y
Jones	M	A	…	N
Smith	M	O	…	N
Ross	M	O	…	Y
Lu	F	A	…	N
Shah	M	B	…	Y

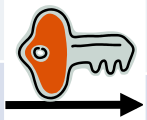

Name	Sex	Blood	…	HIV?
100101	001001	110101	…	110111
101010	111010	111111	…	001001
001010	100100	011001	…	110101
001110	010010	110101	…	100001
110101	000000	111001	…	010010
111110	110010	000101	…	110101

Problems?

Approach 2: Anonymize the Data

Name	Sex	Blood	…	HIV?
Chen	F	B	…	Y
Jones	M	A	…	N
Smith	M	O	…	N
Ross	M	O	…	Y
Lu	F	A	…	N
Shah	M	B	…	Y

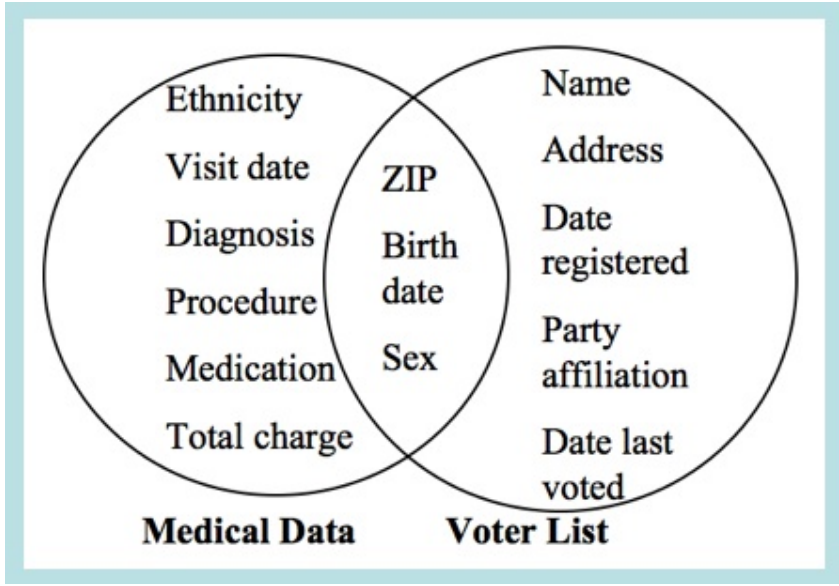

[Sweeney `97]

"re-identification" often easy

Problems?

Approach 3: Mediate Access

Name	Sex	Blood	…	HIV?
Chen	F	B	…	Y
Jones	M	A	…	N
Smith	M	O	…	N
Ross	M	O	…	Y
Lu	F	A	…	N
Shah	M	B	…	Y

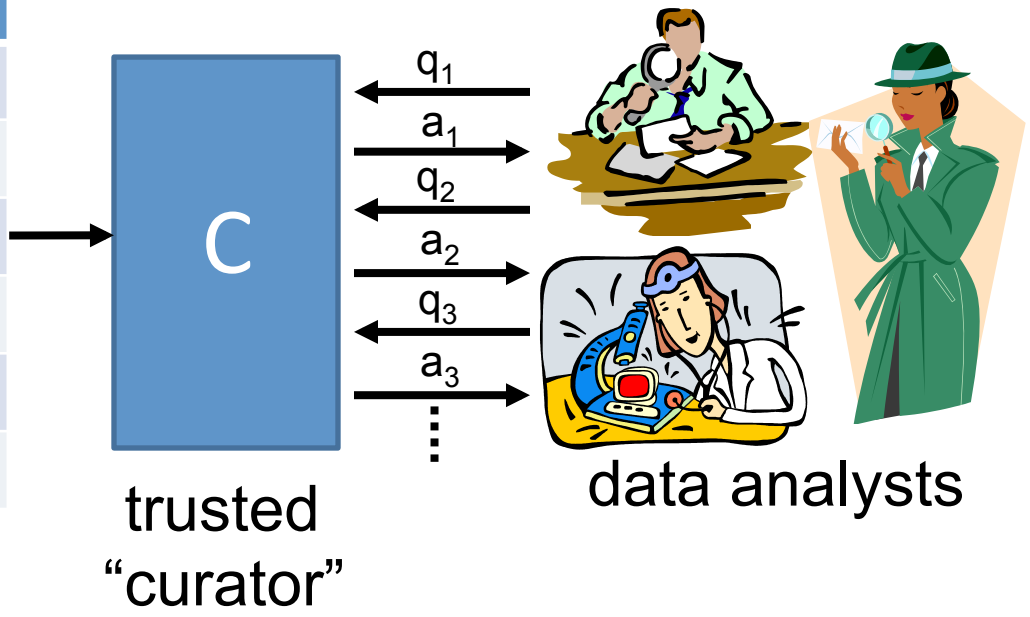

Problems?

Privacy Models from CS

Model	Utility	Privacy	Who Holds Data?
Differential Privacy	statistical analysis of dataset	individual-specific info	trusted curator
Secure Function Evaluation	any query desired	everything other than result of query	original users (or semi-trusted delegates)
Fully Homomorphic (or Functional) Encryption	any query desired	everything (except possibly result of query)	untrusted server

For other two topics, see Shafi Goldwasser's talk
at White House-MIT Big Data Privacy Workshop 3/3/14

Differential privacy

[Dinur-Nissim ’03+Dwork, Dwork-Nissim ’04, Blum-Dwork-McSherry-Nissim ’05, Dwork-McSherry-Nissim-Smith ’06]

Sex	Blood	…	HIV?
F	B	…	Y
M	A	…	N
M	O	…	N
M	O	…	Y
F	A	…	N
M	B	…	Y

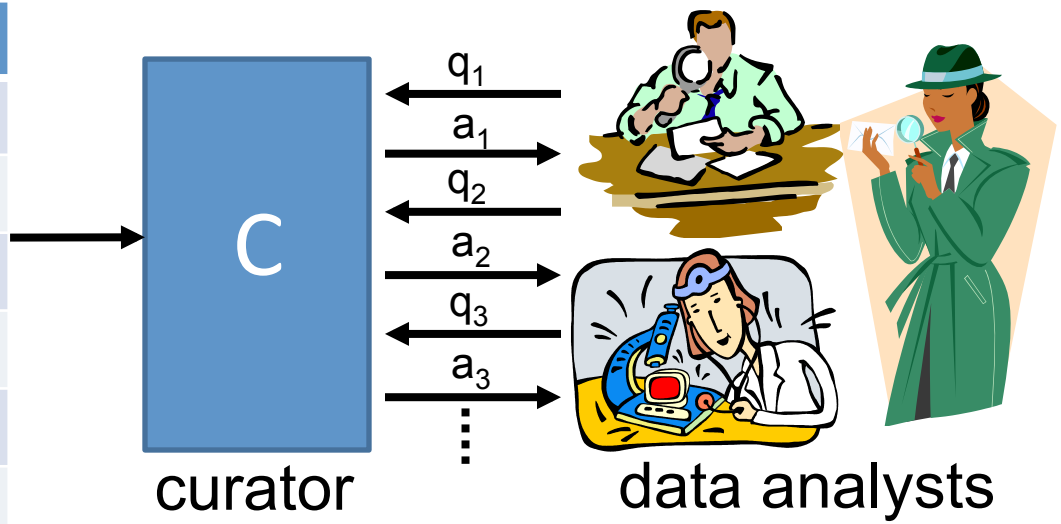

Requirement: effect of each individual should be “hidden”

Differential privacy

[Dinur-Nissim ’03+Dwork, Dwork-Nissim ’04, Blum-Dwork-McSherry-Nissim ’05, Dwork-McSherry-Nissim-Smith ’06]

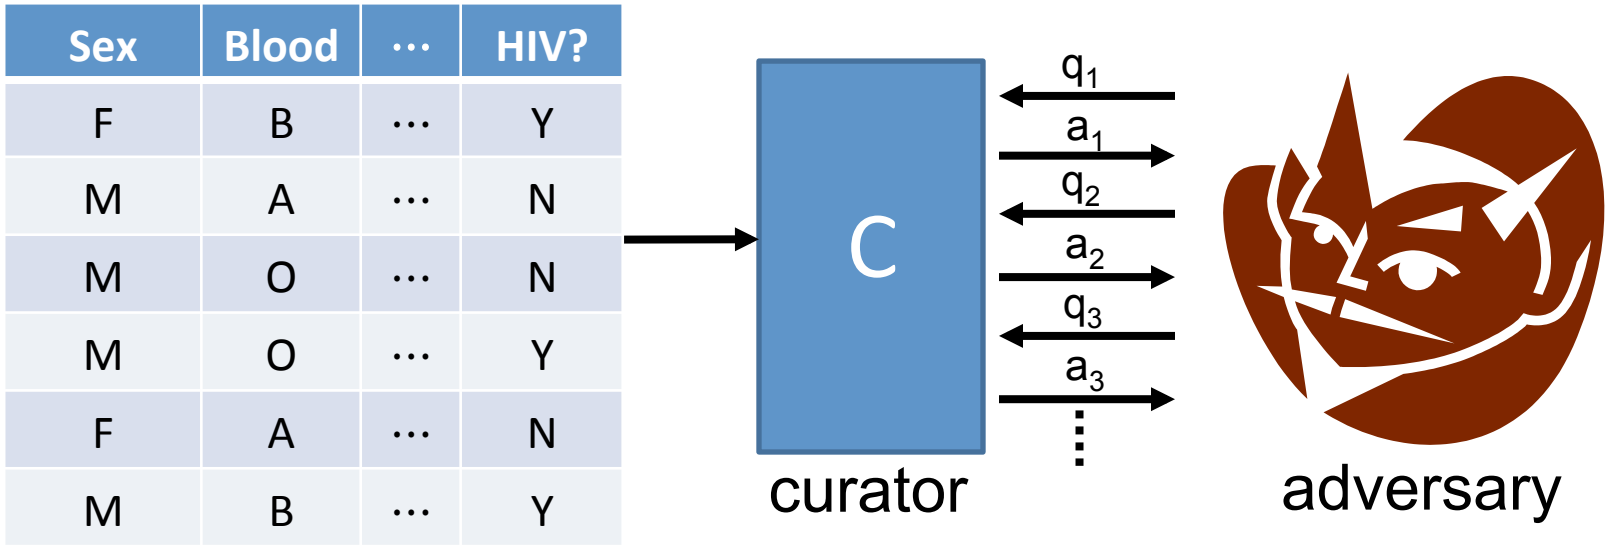

Differential privacy

[Dinur-Nissim ’03+Dwork, Dwork-Nissim ’04, Blum-Dwork-McSherry-Nissim ’05, Dwork-McSherry-Nissim-Smith ’06]

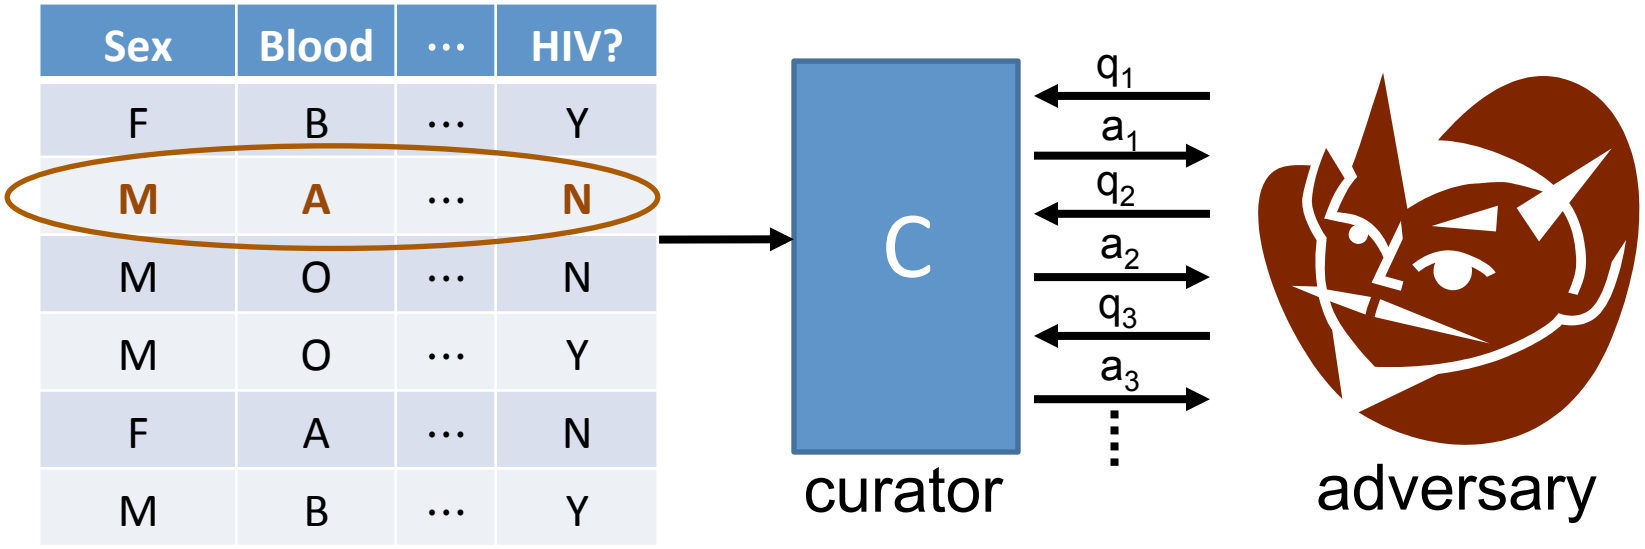

Requirement: an adversary shouldn’t be able to tell if any one person’s data were changed arbitrarily

Differential privacy

[Dinur-Nissim ’03+Dwork, Dwork-Nissim ’04, Blum-Dwork-McSherry-Nissim ’05, Dwork-McSherry-Nissim-Smith ’06]

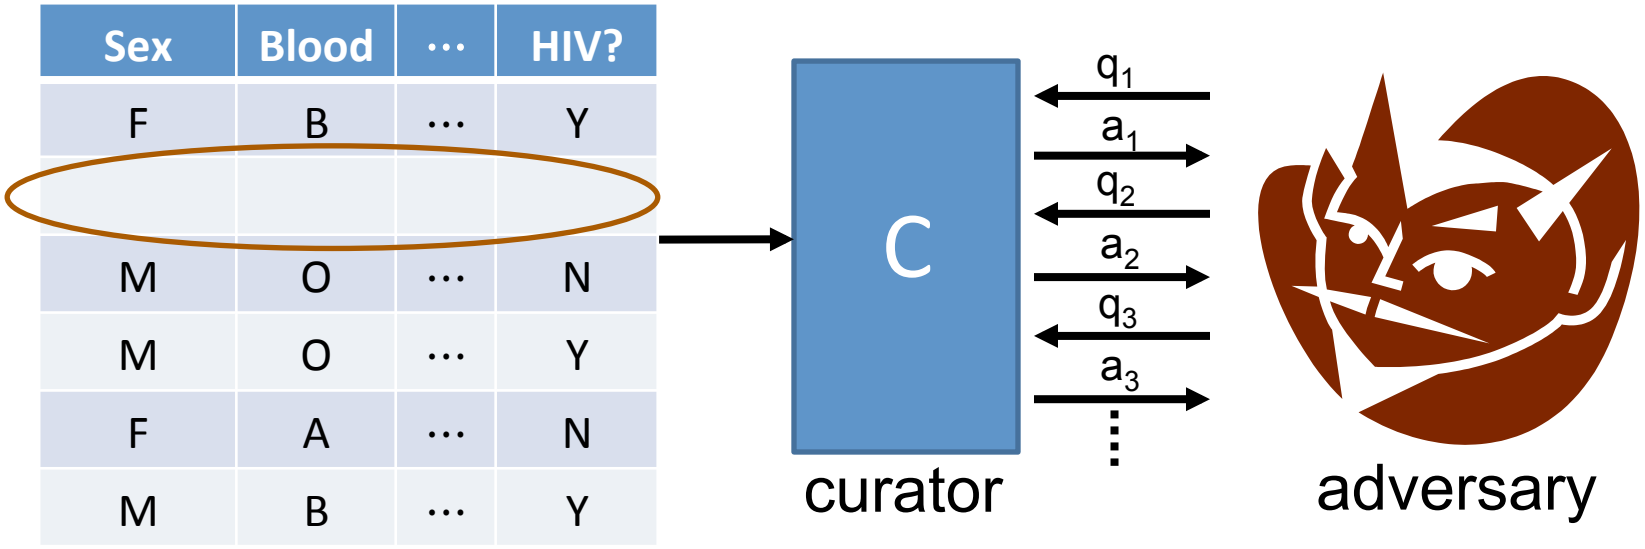

Requirement: an adversary shouldn’t be able to tell if any one person’s data were changed arbitrarily

Differential privacy

[Dinur-Nissim ’03+Dwork, Dwork-Nissim ’04, Blum-Dwork-McSherry-Nissim ’05, Dwork-McSherry-Nissim-Smith ’06]

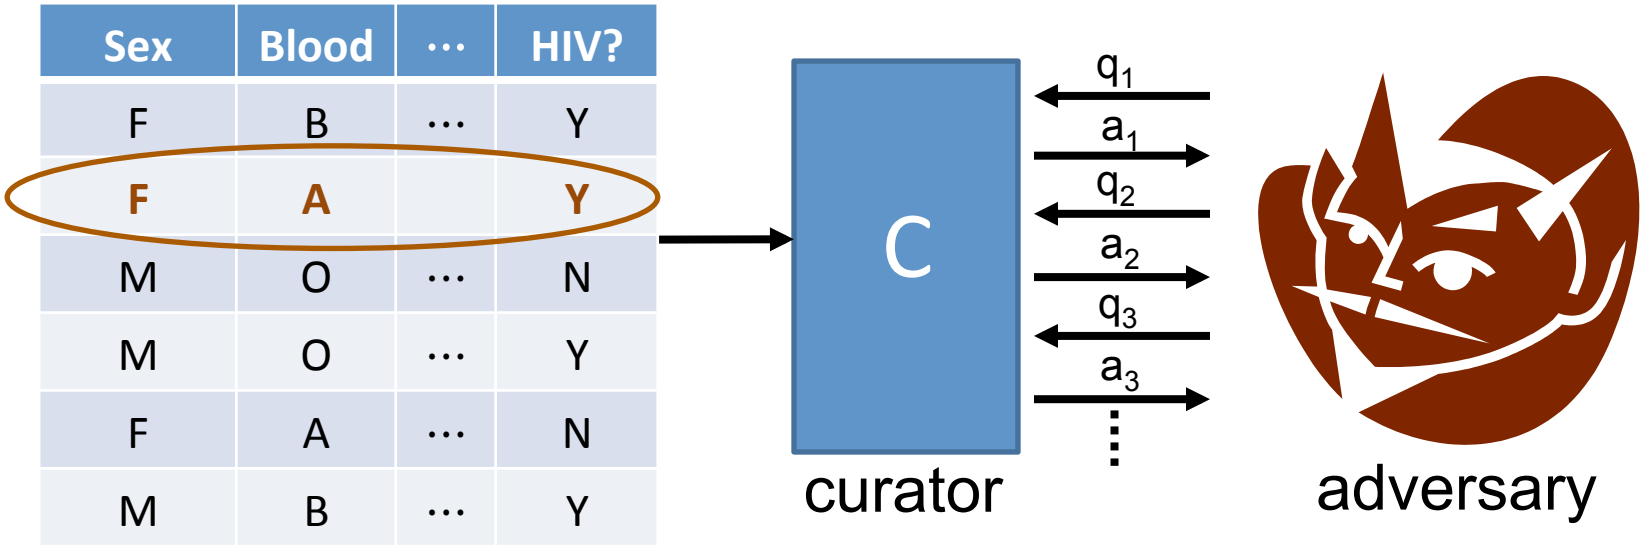

Requirement: an adversary shouldn’t be able to tell if any one person’s data were changed arbitrarily

Simple approach: random noise

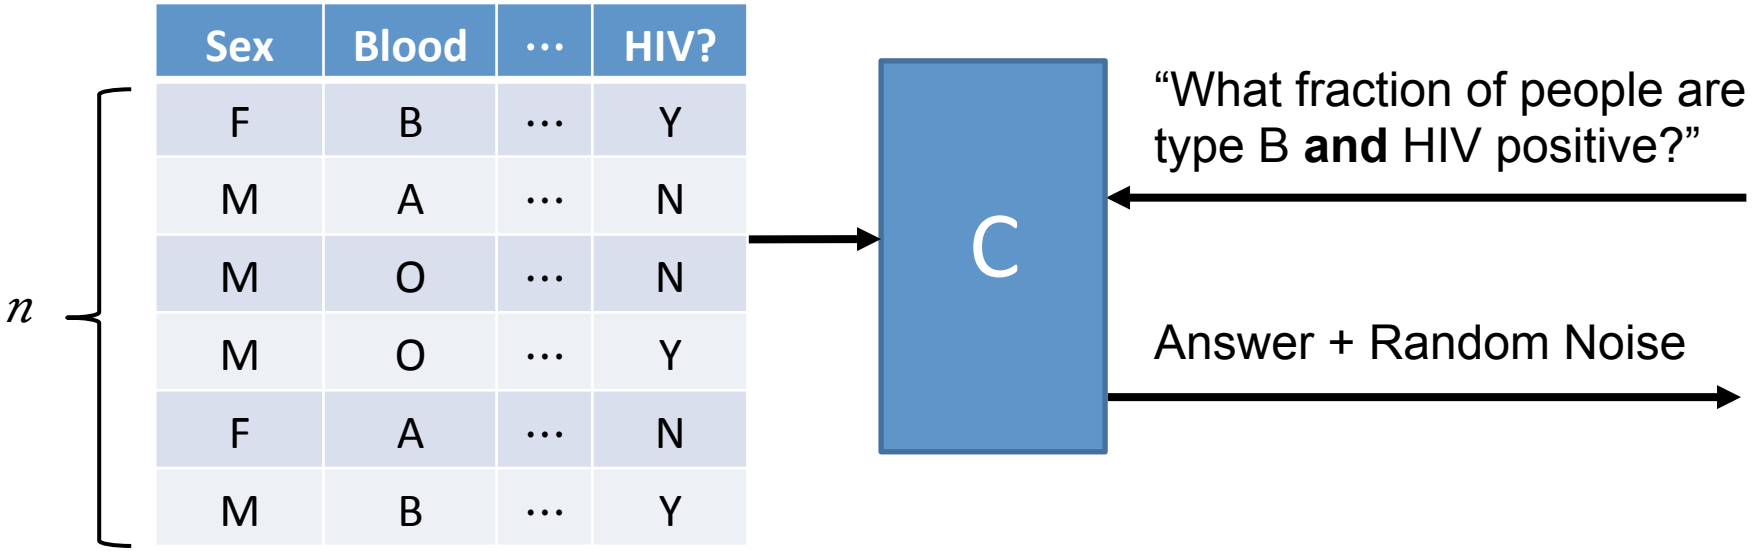

- Very little noise needed to hide each person as $n \to \infty$.
- Limited to answering $\approx n^2$ queries [Dwork-Naor-Vadhan '12]

Differential privacy

[Dinur-Nissim '03+Dwork, Dwork-Nissim '04, Blum-Dwork-McSherry-Nissim '05, Dwork-McSherry-Nissim-Smith '06]

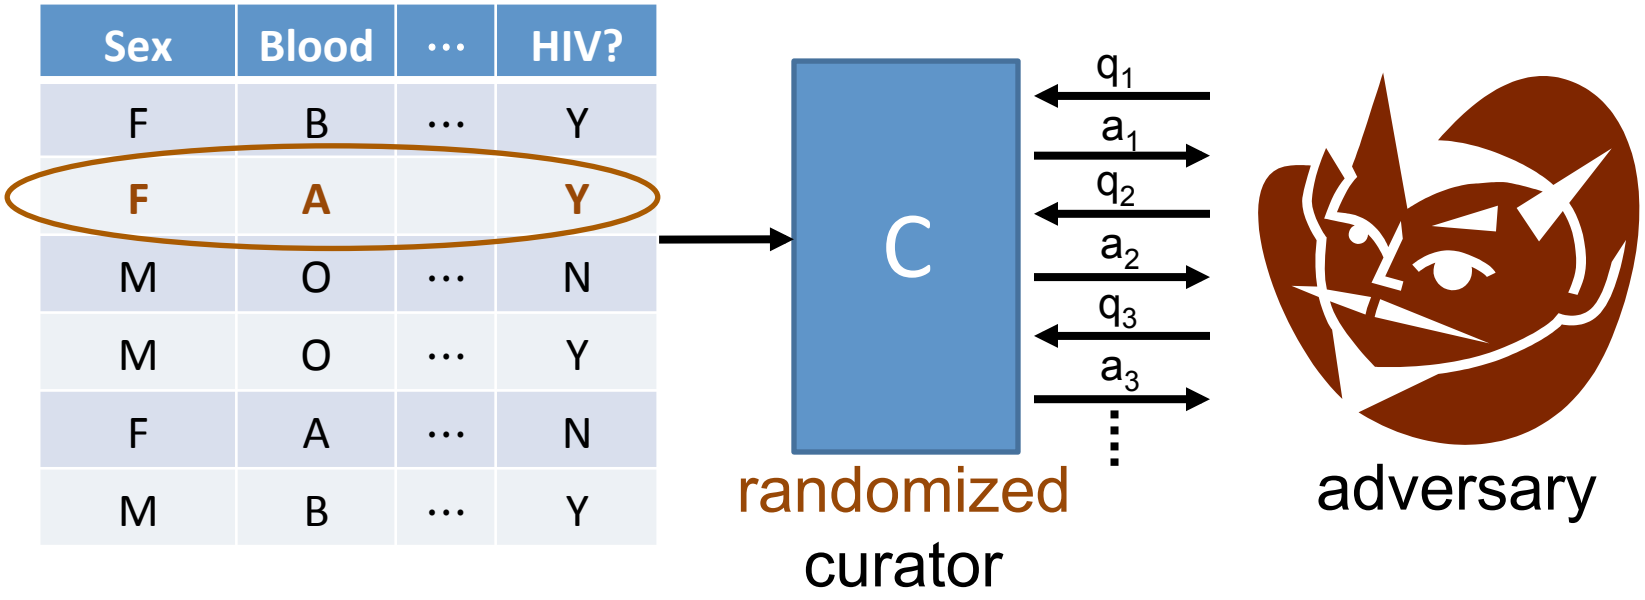

Requirement: for all D, D' differing on one row, and all $q_1,\ldots,q_t$

Distribution of $C(D,q_1,\ldots,q_t) \approx_{\varepsilon}$ Distribution of $C(D',q_1,\ldots,q_t)$

Some Differentially Private Algorithms

- histograms [DMNS06]
- contingency tables [BCDKMT07, GHRU11, TUV12, DNT14],
- machine learning [BDMN05,KLNRS08],
- regression & statistical estimation [CMS11,S11,KST11,ST12,JT13]
- clustering [BDMN05,NRS07]
- social network analysis [HLMJ09,GRU11,KRSY11,KNRS13,BBDS13]
- approximation algorithms [GLMRT10]
- singular value decomposition [HR12, HR13, KT13, DTTZ14]
- streaming algorithms [DNRY10,DNPR10,MMNW11]
- mechanism design [MT07,NST10,X11,NOS12,CCKMV12,HK12,KPRU12]
- …

See Simons Institute Workshop on Big Data & Differential Privacy 12/13

Differential Privacy: Interpretations

Distribution of $C(D,q_1,\ldots,q_t) \approx_{\downarrow\varepsilon}$ Distribution of $C(D',q_1,\ldots,q_t)$

- Whatever an adversary learns about me, it could have learned from everyone else's data.
- Mechanism cannot leak "individual-specific" information.
- Above interpretations hold regardless of adversary's auxiliary information.
- Composes gracefully (k repetitions **)** kε differentially private)

But

- No protection for information that is not localized to a few rows.
- No guarantee that subjects won't be "harmed" by results of analysis.

Simple approach: random noise

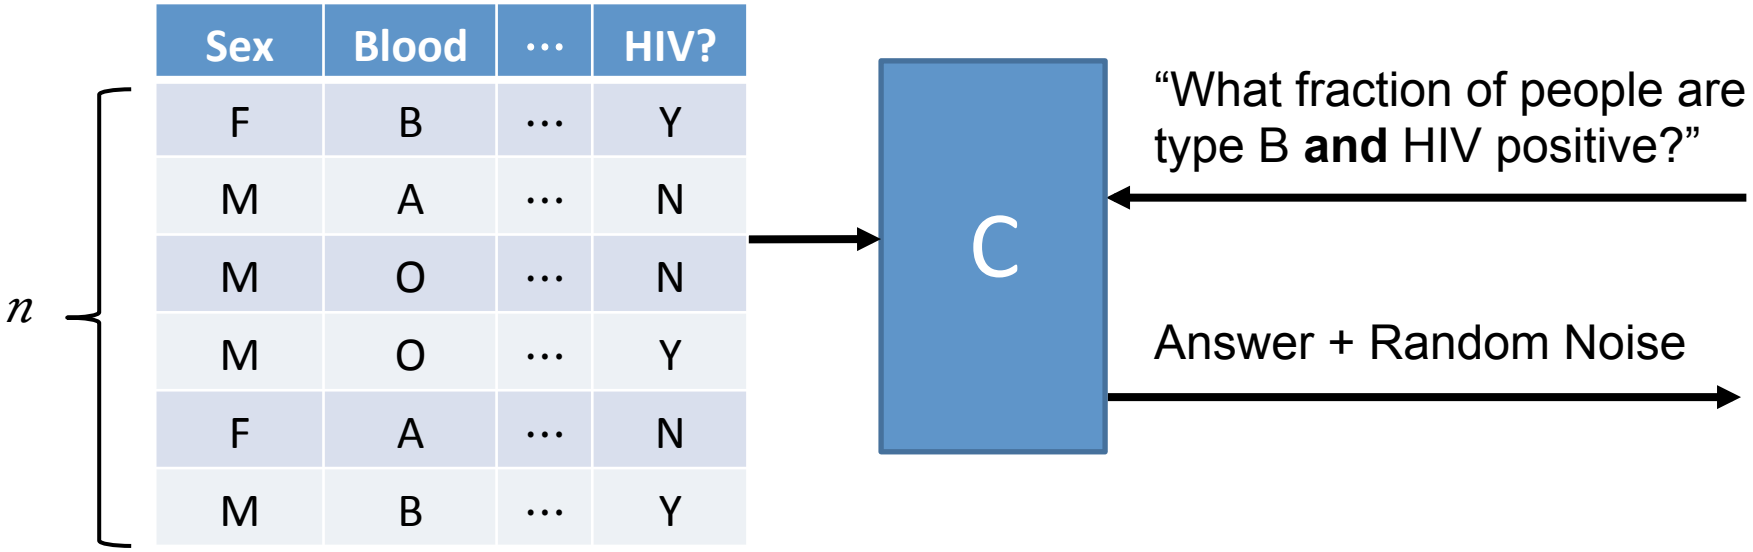

- Very little noise needed to hide each person as $n \to \infty$.
- Limited to answering $\approx n^2$ queries [Dwork-Naor-Vadhan '12]

Amazing possibility: synthetic data

[Blum-Ligett-Roth '08]

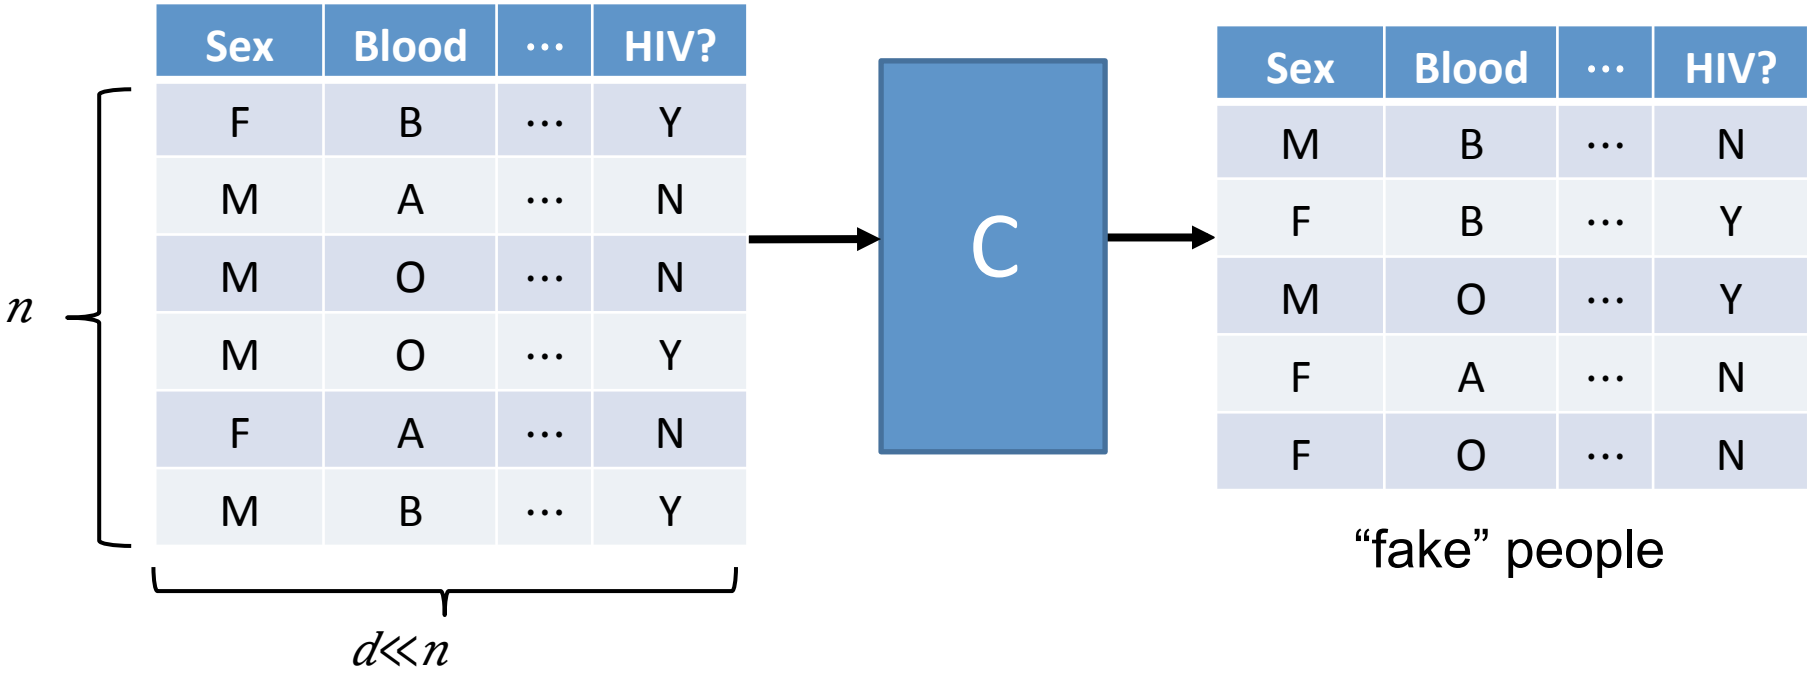

Utility: preserves fraction of people with *every* set of attributes!

Challenge: make this computationally feasible for high-dimensional datasets

Amazing Possibility II: Statistical Inference & Machine Learning

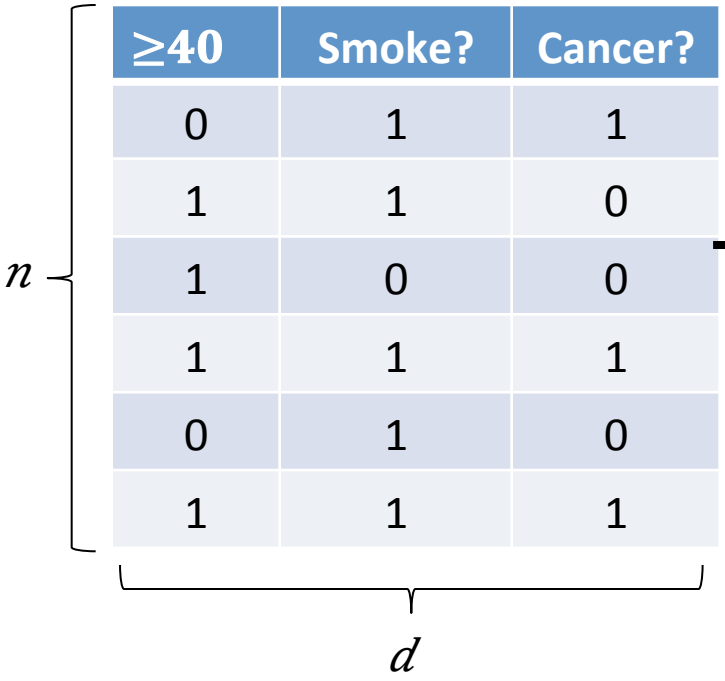

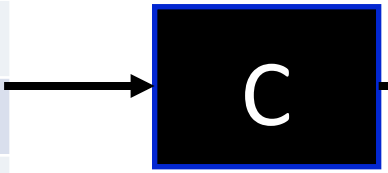

Hypothesis h or model θ about world, e.g. rule for predicting disease from attributes

Theorem [KLNRS08,S11]: Differential privacy for vast array of machine learning and statistical estimation problems with little loss in convergence rate as $n \to \infty$.

- Optimizations & practical implementations for logistic regression, ERM, LASSO, SVMs in [RBHT09,CMS11,ST13,JT14].

Challenges for DP in Practice

- Accuracy for "small data" (moderate values of n)
- Modelling & managing privacy loss over time
  - Especially over many different analysts & datasets
- Analysts used to working with raw data
  - One approach: "Tiered access"
  - DP for wide access, raw data only by approval with strict terms of use (cf. Census PUMS vs. RDCs)
- Cases where privacy concerns are not "local" (e.g. privacy for large groups) or utility is not "global" (e.g. targeting)
- Matching guarantees with privacy law & regulation
- …

Some Efforts to Bring DP to Practice

- CMU-Cornell-PennState “Integrating Statistical and Computational Approaches to Privacy”
  - See http://onthemap.ces.census.gov/
- UCSD “Integrating Data for Analysis, Anonymization, and Sharing” (iDash)
- UT Austin “Airavat: Security & Privacy for MapReduce”
- UPenn “Putting Differential Privacy to Work”
- Stanford-Berkeley-Microsoft “Towards Practicing Privacy”
- Duke-NISSS “Triangle Census Research Network”
- Harvard “Privacy Tools for Sharing Research Data”
- …

Goals for our DP Tools

- General-purpose: applicable to most datasets uploaded to Dataverse.
- Automated: no differential privacy expert optimizing algorithms for a particular dataset or application
- Tiered access: DP interface for wide access to rough statistical information, helping users decide whether to apply for access to raw data (cf. Census PUMS vs RDCs)

Demo of our DP Tools

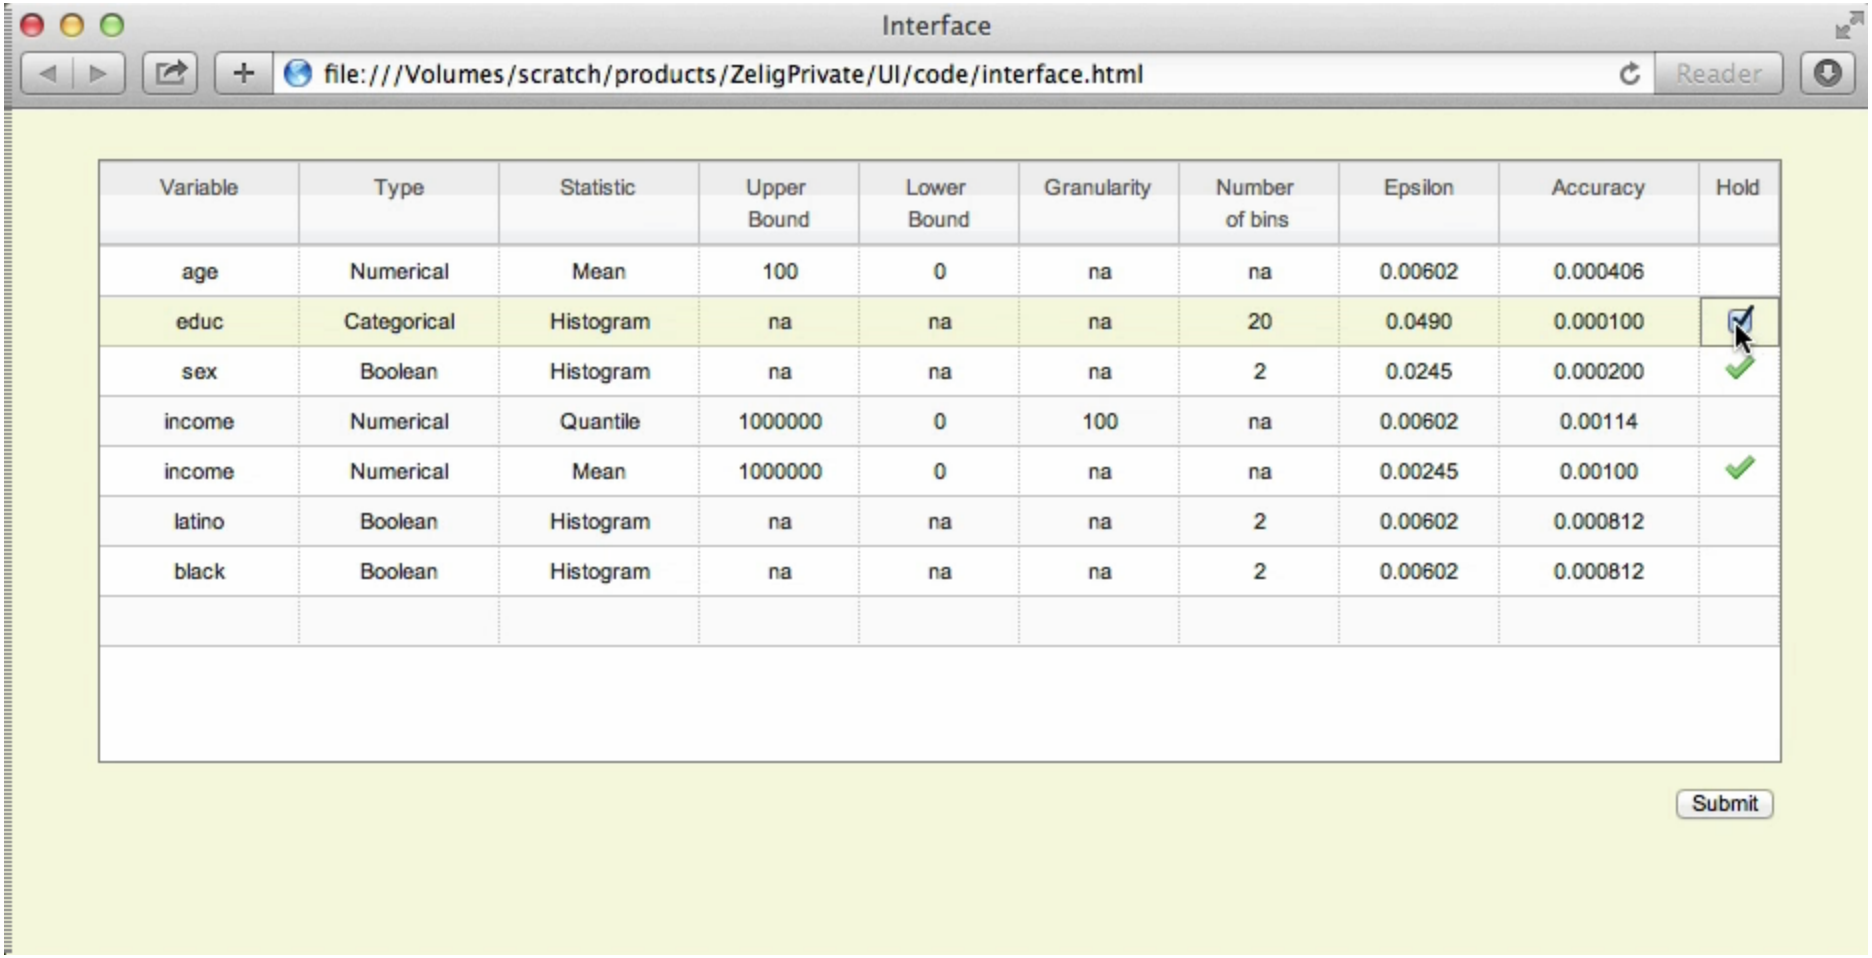

Variable	Type	Statistic	Upper Bound	Lower Bound	Granularity	Number of bins	Epsilon	Accuracy	Hold
age	Numerical	Mean	100	0	na	na	0.00602	0.000406	
educ	Categorical	Histogram	na	na	na	20	0.0490	0.000100	☑
sex	Boolean	Histogram	na	na	na	2	0.0245	0.000200	✓
income	Numerical	Quantile	1000000	0	100	na	0.00602	0.00114	
income	Numerical	Mean	1000000	0	na	na	0.00245	0.00100	✓
latino	Boolean	Histogram	na	na	na	2	0.00602	0.000812	
black	Boolean	Histogram	na	na	na	2	0.00602	0.000812	

Submit

Where we're headed

Current Tool

- Means, quantiles, CDFs, histograms
- Computed at time of dataset deposit
- Depositor decides how to allocate "privacy budget"

This Summer

- Interactive queries by data analyst
- Visualization of error introduced by DP in TwoRavens
- Utility testing for social science
- Least-squares regression
- Contingency tables
- Integration with Dataverse
- Attacks on aggregate data

Differential Privacy: Summary

Differential Privacy offers

- Strong, scalable privacy guarantees
- Compatibility with many types of “big data” analyses
- Amazing possibilities for what can be achieved in principle

There are some challenges, but reasons for optimism

- Intensive research effort from many communities
- Some successful uses in practice already
- Differential privacy easier as $n \to \infty$