*

Everything You Always Wanted to Know About Differential Privacy*

Kobbi Nissim
Department of Computer Science
Georgetown University

Alexandra Wood
Berkman Klein Center for Internet & Society
Harvard University

Challenges and New Approaches for Protecting Privacy in Federal Statistical Programs
National Academies of Sciences, June 6, 2019

(*But Were Afraid to Ask)

Data Privacy: The Problem

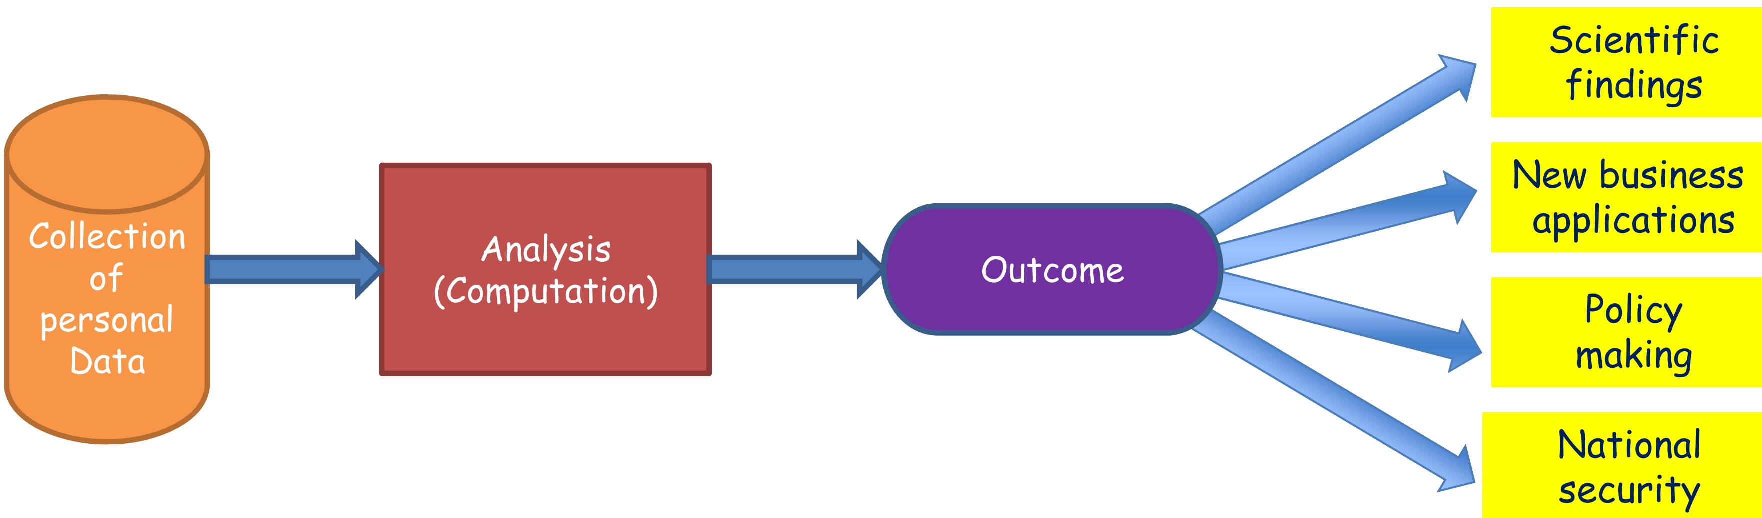

Given a dataset with sensitive personal information, how can one compute and release functions of the dataset while protecting individual privacy?

*

Attacks on SDL Techniques

- Re-identification [Sweeney '00, …]
  - GIC data, health data, clinical trial data, DNA, pharmacy data, text data, registry information, …
- Blatant non-privacy [Dinur, Nissim '03], …
- Auditors [Kenthapadi, Mishra, Nissim '05]
- AOL Debacle '06
- Genome-Wide association studies (GWAS) [Homer et al. '08]
- Netflix award [Narayanan, Shmatikov '09]
- Social networks [Backstrom, Dwork, Kleinberg '11]
- Genetic research studies [Gymrek, McGuire, Golan, Halperin, Erlich '11]
- Microtargeted advertising [Korolova 11]
- Recommendation Systems [Calandrino, Kiltzer, Naryanan, Felten, Shmatikov 11]
- Israeli CBS [Mukatren, Nissim, Salman, Tromer '14]
- Attack on statistical aggregates [Homer et al. '08] [Dwork, Smith, Steinke, Vadhan '15]
- Reconstruction attack on 2010 Census data

Slide idea stolen shamelessly from Or Sheffet

*

Takeaways from Privacy Failures

In setting clear meaningful privacy goals

In understanding how normative and technical conceptions of privacy interact

Lack of rigor leads to unanticipated privacy failures.

In analyzing resilience to future attacks

In taking auxiliary knowledge into account

In accounting for privacy loss across multiple releases

In scrutiny of privacy technology

*

Takeaways from Privacy Failures

- Specific findings:
  - Redaction of identifiers is insufficient for protecting privacy.
  - Similarly: aggregation, noise addition*, …
  - Auxiliary information needs to be taken into account.
  - Regulation and technology only considered a limited scope of privacy failures.
    - New failure modes: whether an individual participated in study, inferences
  - Any useful analysis of personal data must leak some information about individuals.
  - Leakages accumulate with multiple analyses/releases.

Mathematical facts, not matters of policy

A New Line of Work

Emerging from theoretical computer science (~2003).

Yields new concept: Differential privacy (2006):

- Rich theory and new privacy concepts.
- Mathematically provable privacy guarantees.
- In first stages of implementation and real-world use
  - US Census, Google, Apple, Uber, …

Yeah, Yeah ...

What is Differential Privacy?

Differential Privacy is …

… not a specific technique or algorithm!

Differential Privacy is …

… a definition (i.e., a standard) of privacy*

It expresses a specific desiderata of an analysis:

> Any information-related risk to a person should not change significantly as a result of that person's information being included, or not, in the analysis.

*More precisely, a family of related mathematical definitions: pure DP, approximate DP, concentrated DP, …

A Privacy Desiderata

Real world:

My ideal world:

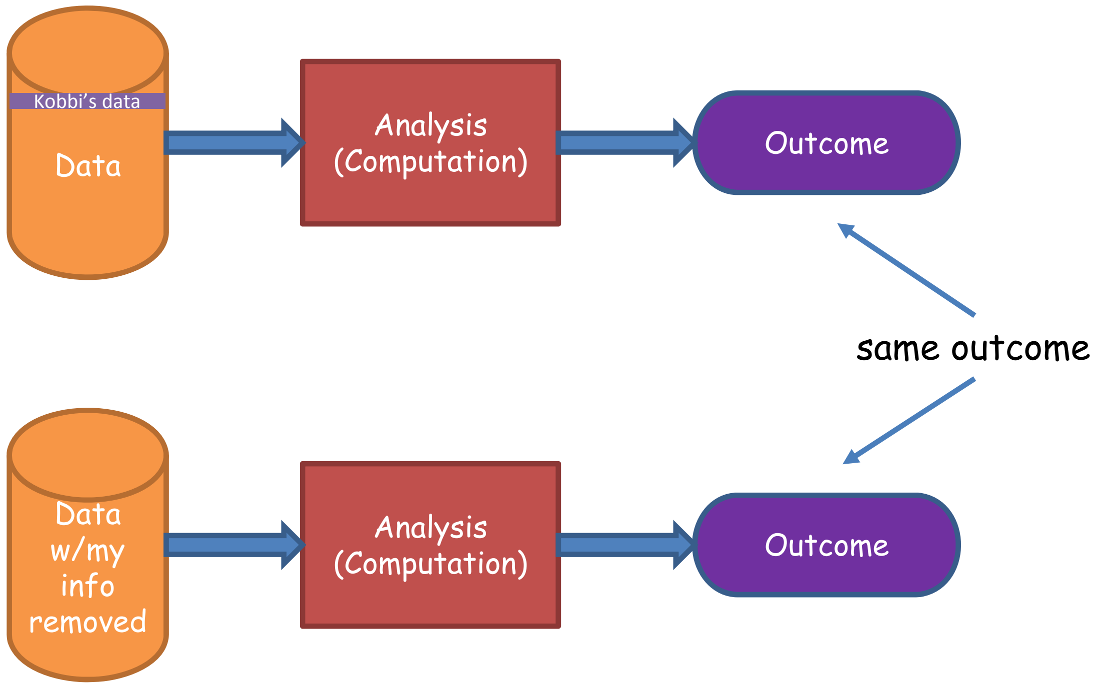

A Privacy Desiderata

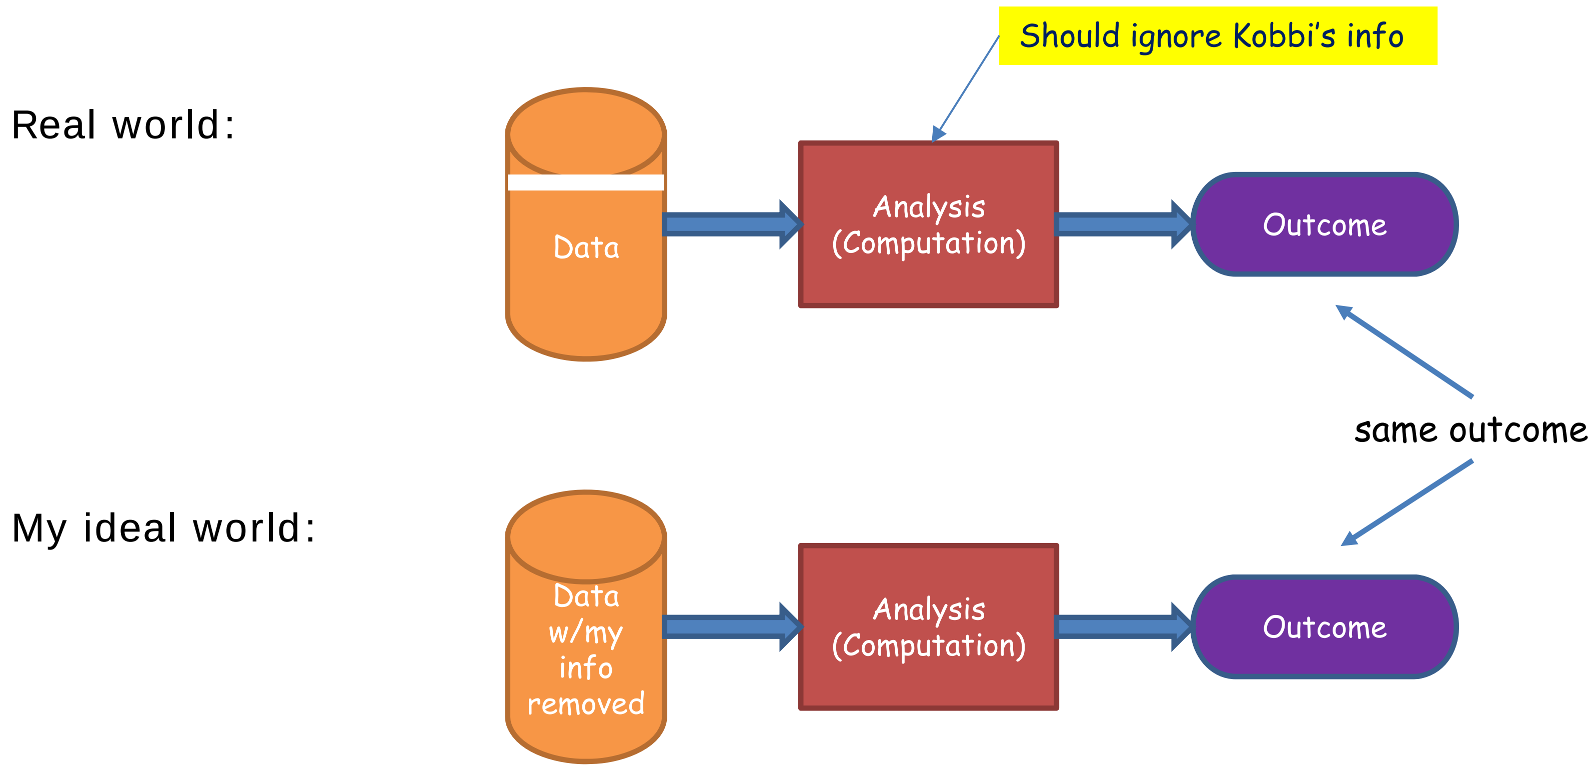

A Privacy Desiderata

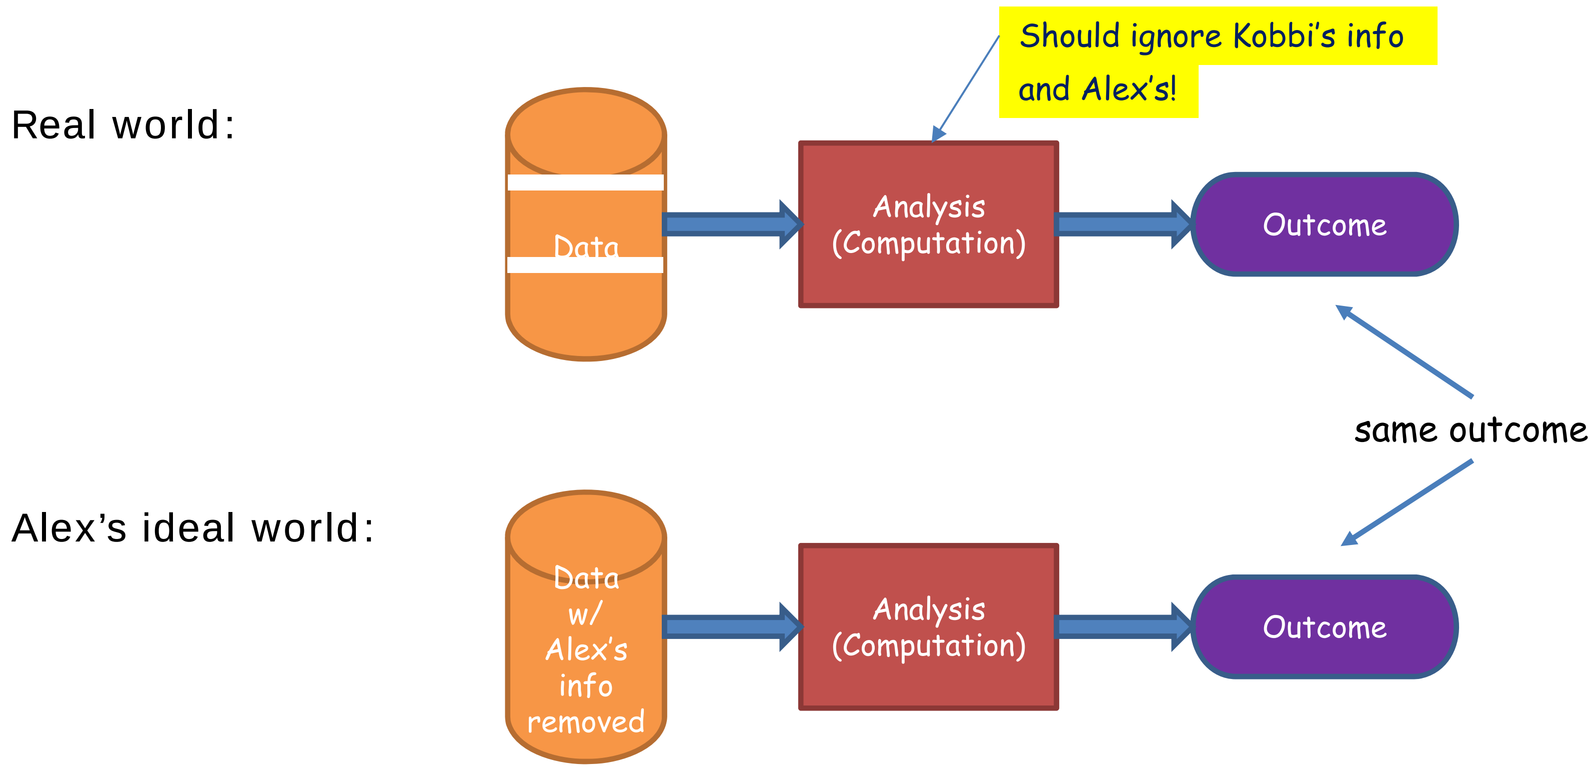

AW2

A Privacy Desiderata

Real world:

Sasha's ideal world:

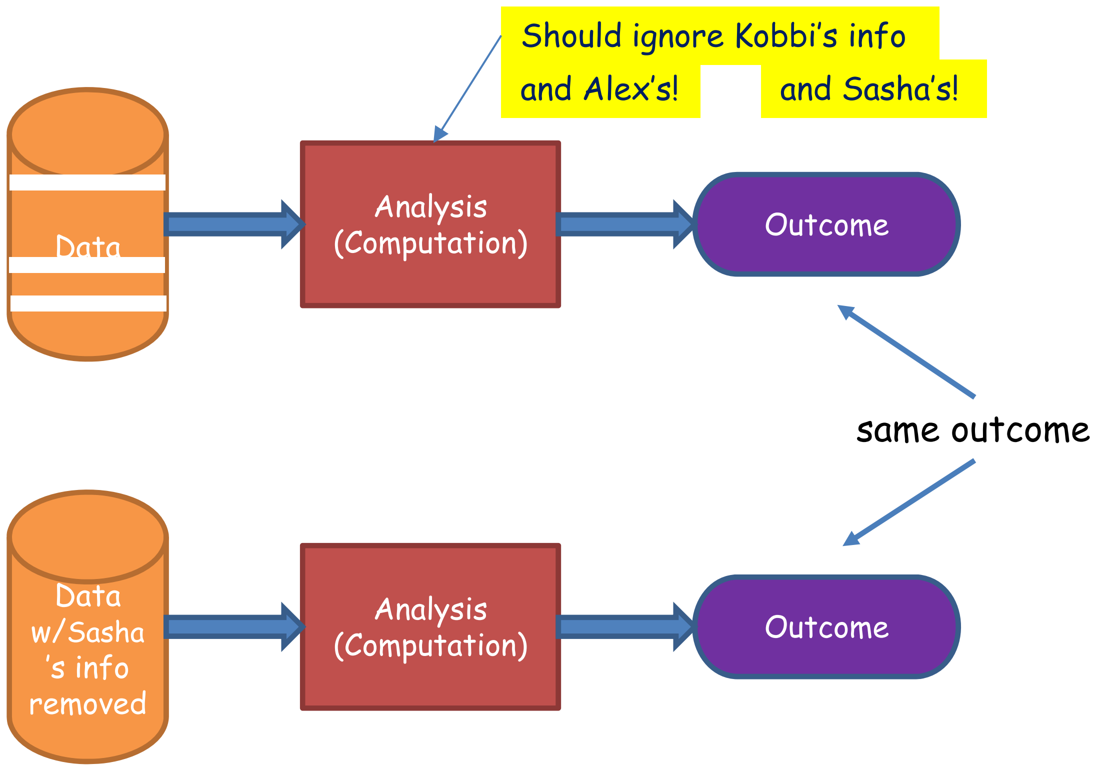

Slide 14

AW2 JULES

Alexandra Wood, 5/15/2018

A Privacy Desiderata

Real world:

☺'s ideal world:

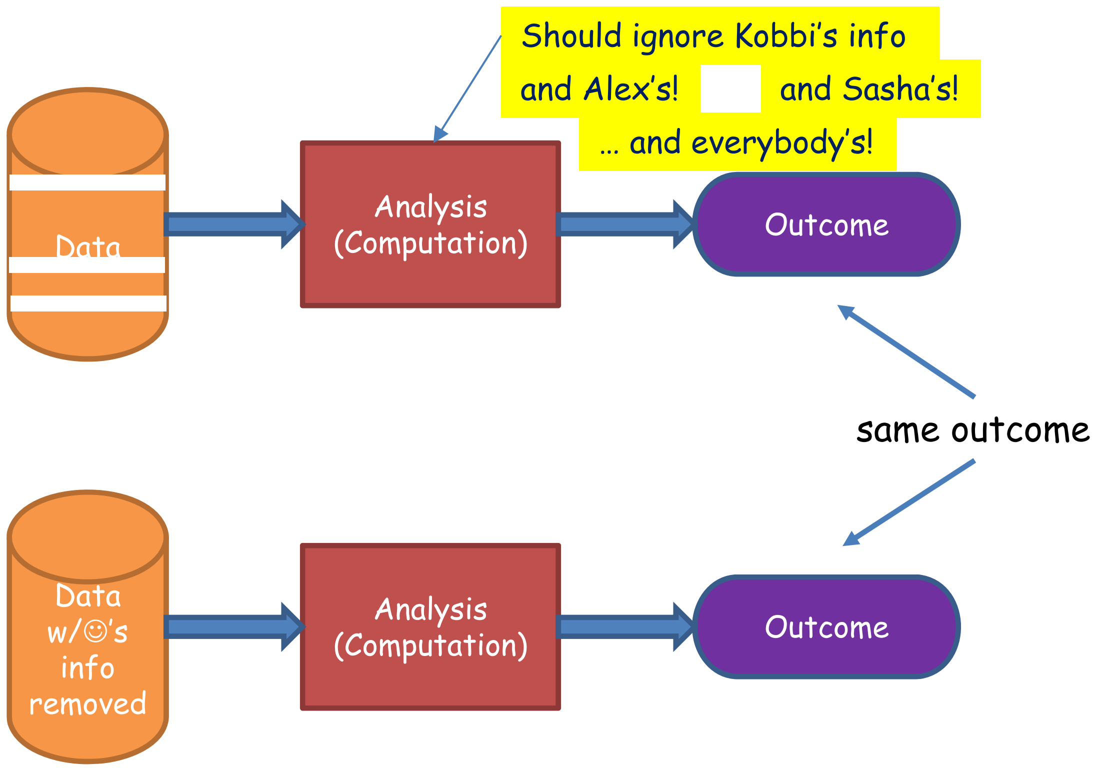

Oops!

That did not go so well …

A More Realistic Privacy Desiderata

Real world:

☺'s ideal world:

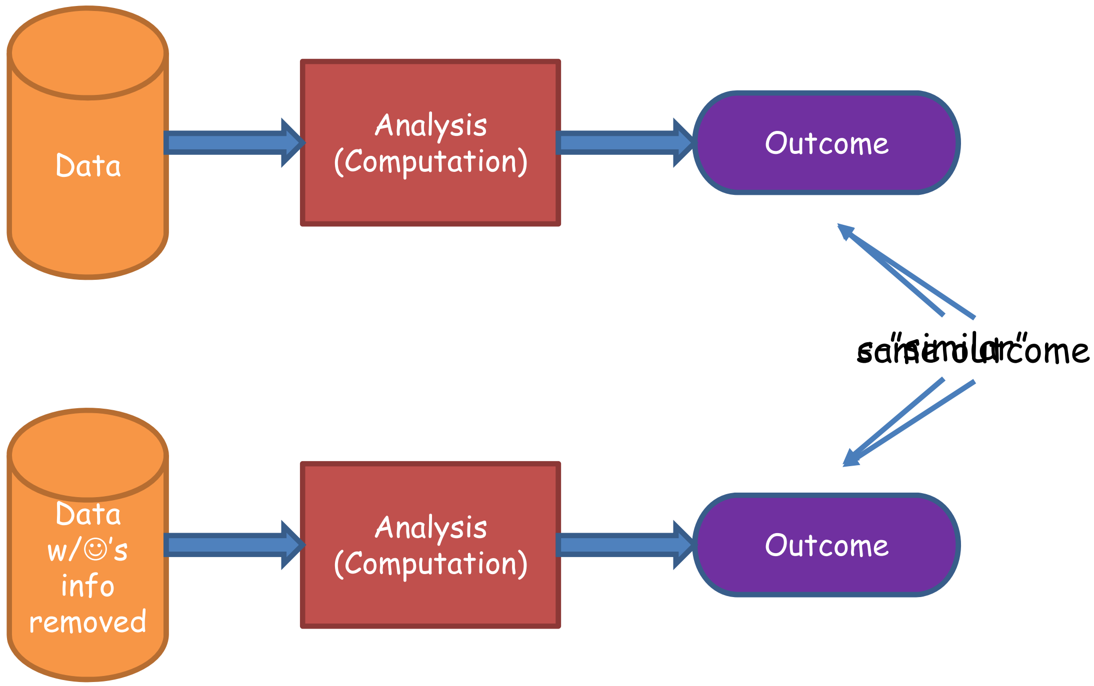

Differential Privacy [Dwork McSherry Nissim Smith '06]

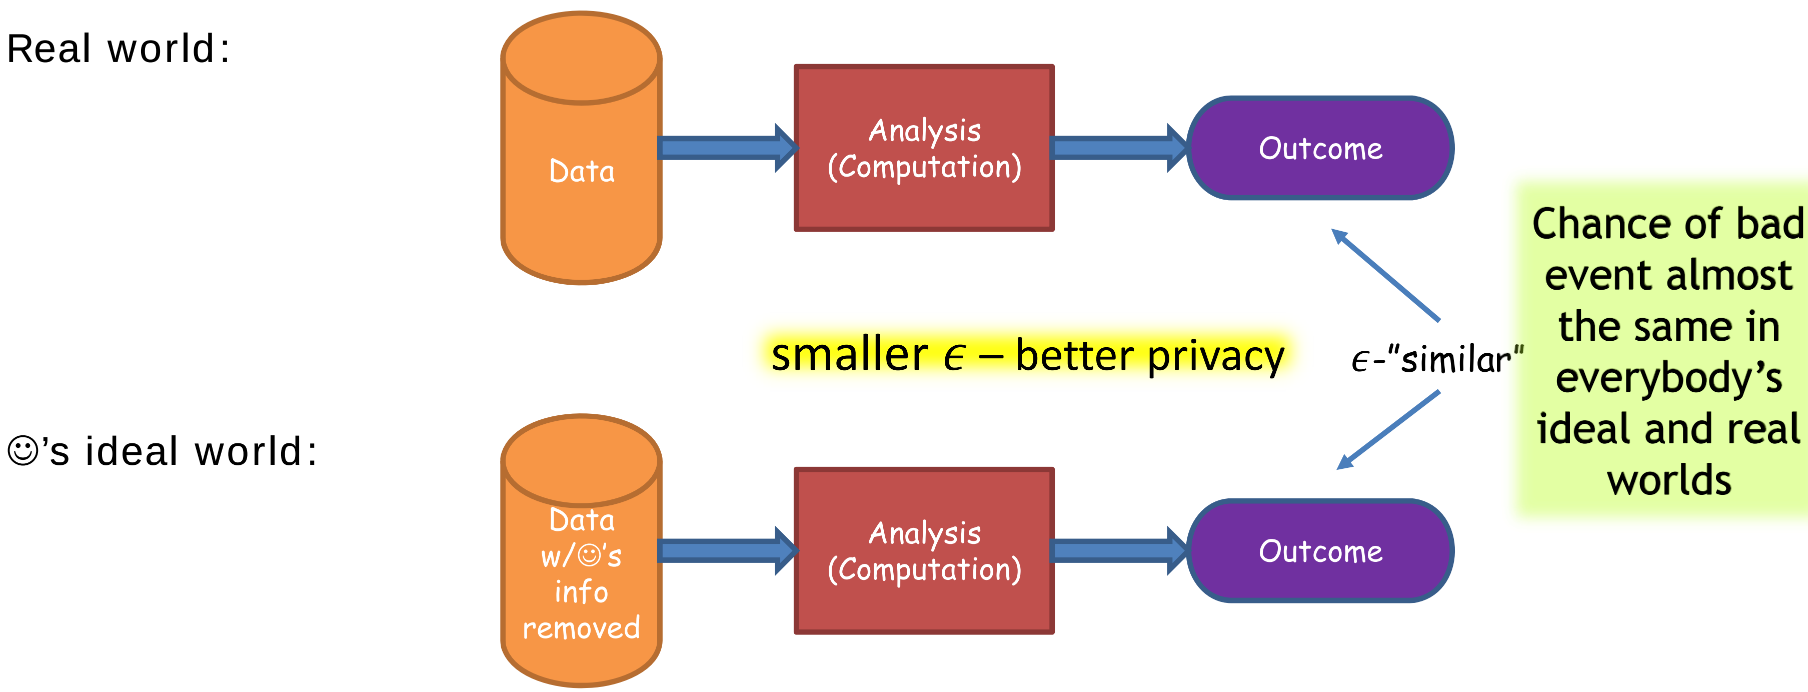

Understanding Differential Privacy

- “Automatic” opt-out: I am protected (almost) as if my info was not used at all.
- I incur limited risk: Contributing my real info can increase the probability I will be denied insurance by at most 1%.
  - When compared with not participating, or contributing fake info.
- These privacy guarantees are provided independent of the methods used by a potential attacker and in presence of arbitrary auxiliary information.
- Future proof: Avoids the “penetrate and patch” cycle.

A Privacy "Budget"

DP provides provable privacy guarantees with respect to the cumulative risk from successive data releases.

- Combination of ϵ-differentially private computations results in differential privacy (with larger ϵ).
- Can manage accumulated privacy loss.
  - Whereas other known definitions of privacy do not measure the cumulative risk from multiple analyses/releases.

- This is an important feature, not a bug!
  - Consider how ignoring the fuel gauge would not make your car run indefinitely without refueling.

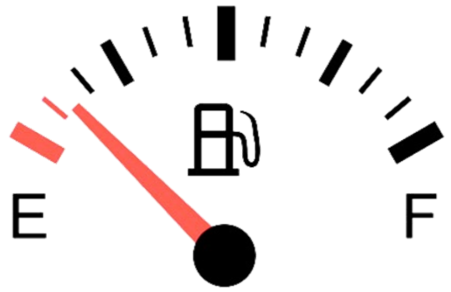

*

Transparency

DP has the benefit of transparency.

- It is not necessary to maintain secrecy around a differentially private computation or its parameters.
  - Whereas some traditional techniques relied on secret algorithms or parameters.
- Benefits of transparency include:
  - Knowledge accumulation.
  - Scrutiny by the scientific community.
  - Possibility of accounting for DP in statistical inference.

*

Application for Public Access to Data

DP can be used to provide broad, public access to data or data summaries in a privacy-preserving way.

- Can consider data publications that were otherwise impossible.
  - Whereas traditional techniques would require (more often) to apply controls in addition to de-identification.

Differential Privacy and Concepts from Privacy Law and Policy*

- PII: Differential privacy can be interpreted as ensuring that using an individual's data will (essentially) not reveal any **personally identifiable information** that is **specific** to her.
  - Here, **specific** refers to information that cannot be inferred unless the individual's information is used in the analysis.

AW3

Differential Privacy and Concepts from Privacy Law and Policy *

✓ PII

- Singling out:
  - This can be formalized mathematically.
    - DP protects against a specific notion of singling out ("predicate singling out").
    - Note: rigorous argument also wrt FERPA's concept of de-identification.

Slide 24

AW3 Singling out text
Alexandra Wood, 5/15/2018

Differential Privacy and Concepts from Privacy Law and Policy *

✓ PII

✓ Singling out

- Linkage: Microdata or contingency tables that allow the identification of population uniques **cannot be created** using statistics produced by a differentially private tool.
  - This can be formalized and proved mathematically.

Differential Privacy and Concepts from Privacy Law and Policy *

✓ PII

✓ Singling out

✓ Linkage

- Inference: Differential privacy masks the contribution of any single individual, (essentially) making it **impossible to infer any information specific to an individual**, including whether an individual's information was used at all.
  - But DP does not protect against all inferences.

Differential Privacy and Concepts from Privacy Law and Policy *

- ✓ PII
- ✓ Singling out
- ✓ Linkage
- ✓ Inference

Differential privacy provides protection (far) beyond "identifiability."

Example: Reasoning About Risk
Gertrude's Life Insurance

- Gertrude:
  - Age: 65
  - She has a $100,000 life insurance policy.
  - She is considering participating in a medical study but is concerned it may affect her insurance premium.

Example: Reasoning About Risk Gertrude's Life Insurance

- Based on her age and sex, she has a 1% chance of dying next year. Her life insurance premium is set at 0.01 x $100,000 = $1,000.
- Gertrude is a coffee drinker. If the medical study finds that 65-year-old female coffee drinkers have a 2% chance of dying next year, her premium would be set at $2,000.
  - This would be her baseline risk: Her premium would be set at $2,000 even if she were not to participate in the study.
- Can Gertrude's premium increase beyond her baseline risk?
  - She is worried that the study may reveal more about her, such as that she *specifically* has a 50% chance of dying next year. This can increase her premium from $2,000 to $50,000!

Example: Reasoning About Risk
Gertrude's Life Insurance

- Reasoning about Gertrude's risk
  - Imagine instead the study is performed using differential privacy with $\varepsilon = 0.01$.
  - The insurance company's estimate of Gertrude's risk of dying in the next year can increase to at most

$$(1+\varepsilon)\cdot 2\% = 2.02\%.$$

  - Her premium would increase to at most $2,020. Therefore, Gertrude's risk would be ≤ $2020 - $2000 = $20.

Example: Reasoning About Risk
Gertrude's Life Insurance

- Generally, calculating one's baseline is very complex (if possible at all).
  - In particular, in our example the 2% baseline depends on the potential outcome of the study.
  - The baseline may also depend on many other factors Gertrude does not know.
- However, differential privacy provides simultaneous guarantees for every possible baseline value.
  - The guarantee covers not only changes in Gertrude's life insurance premiums, but also her health insurance and more.

How is differential privacy achieved?

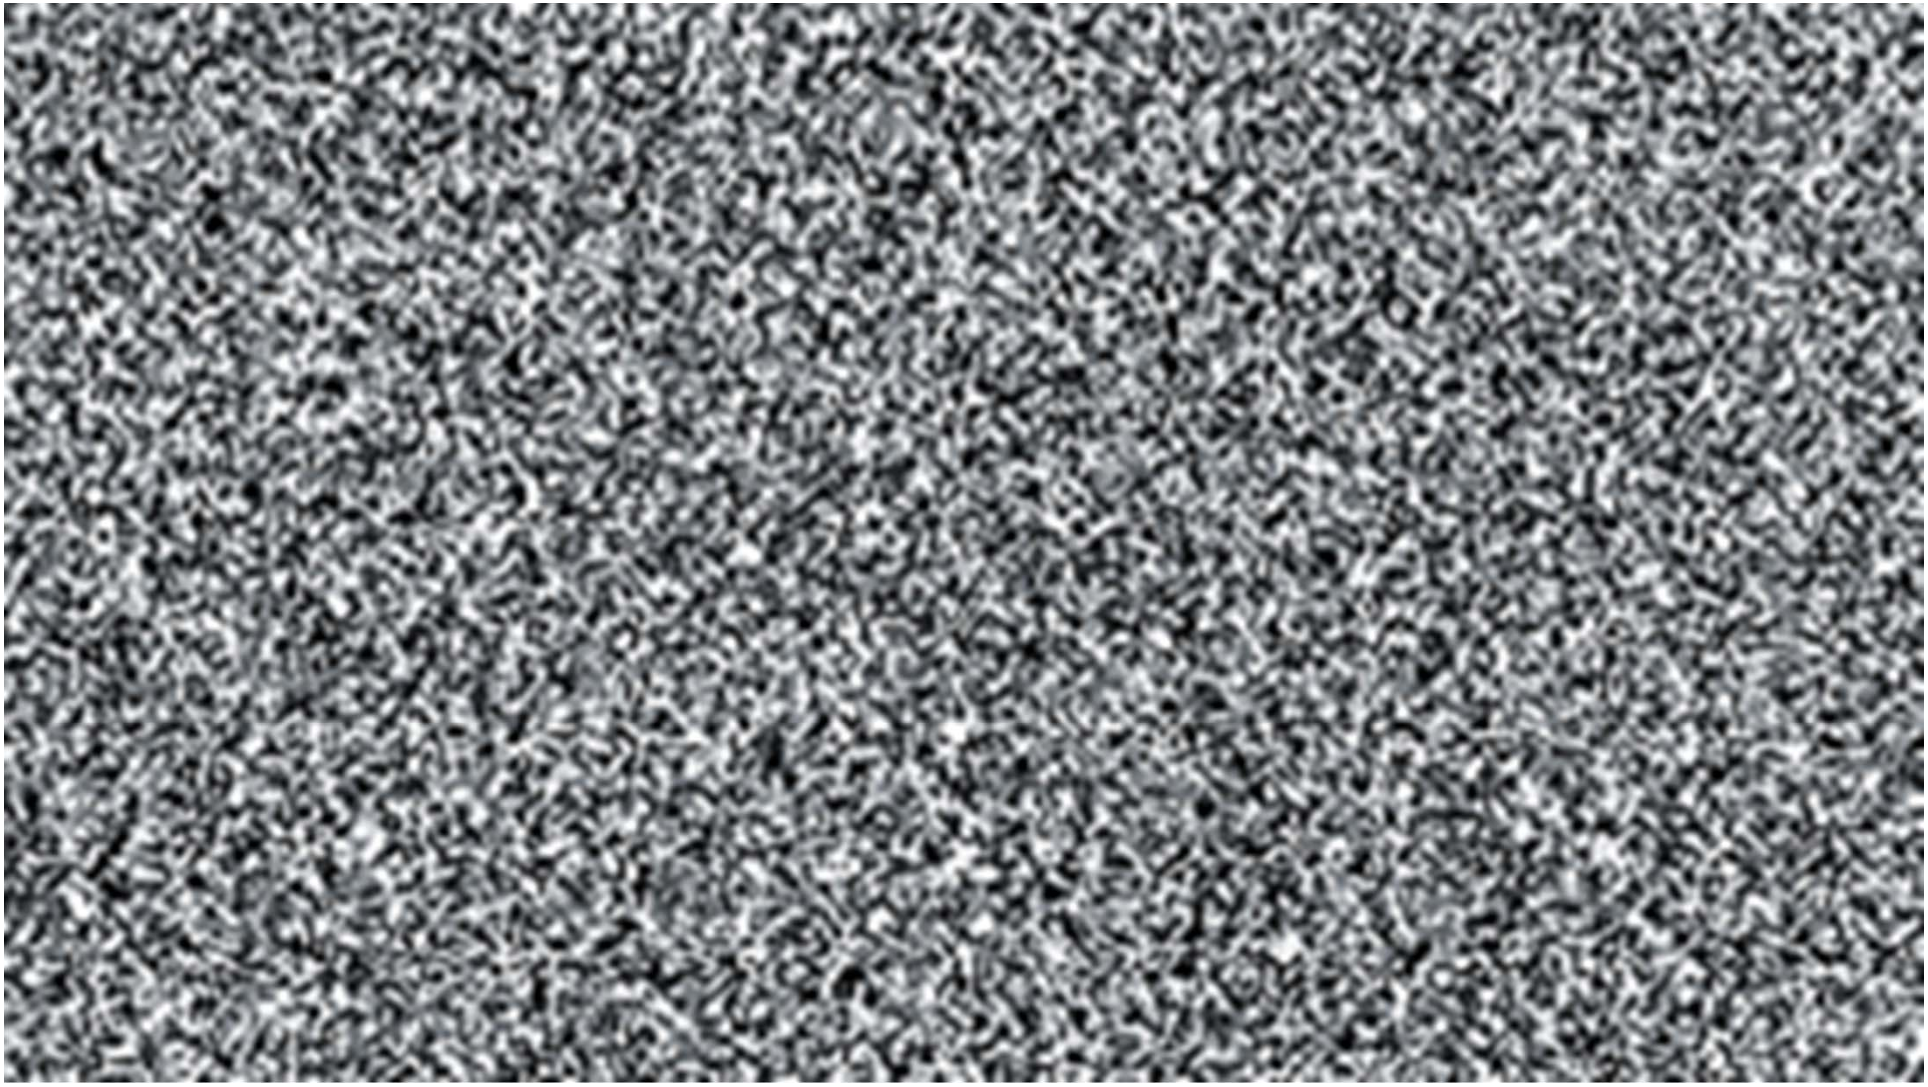

Differentially Private Computations

Algorithms maintain differential privacy via the introduction of *carefully crafted* random noise into the computation.

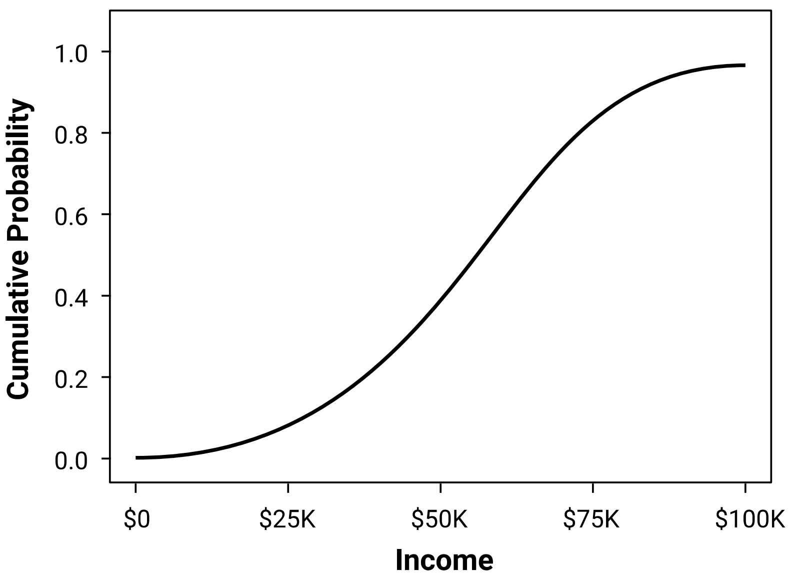

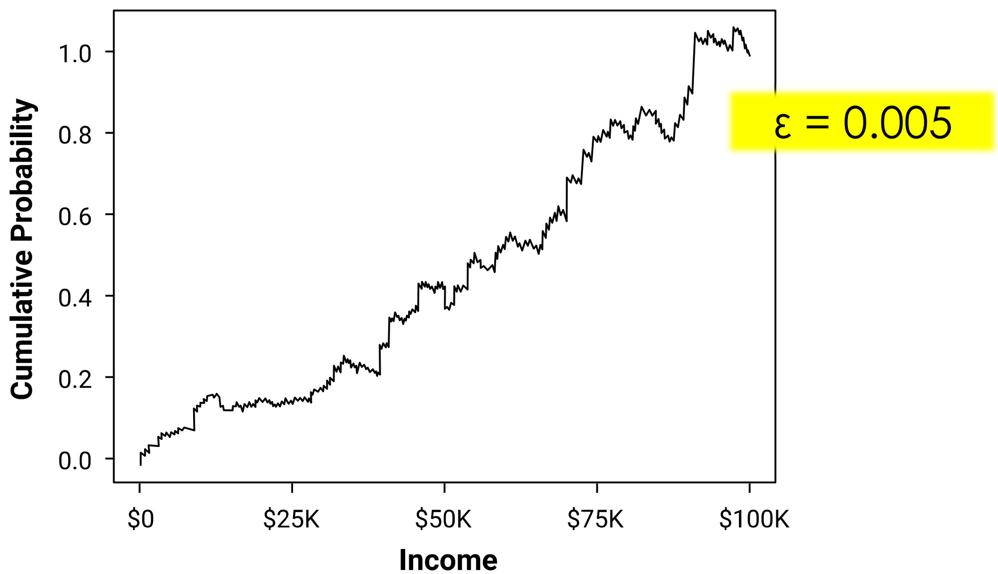

(These CDFs are stylized examples.)

Differentially Private Computations

Algorithms maintain differential privacy via the introduction of *carefully crafted* random noise into the computation.

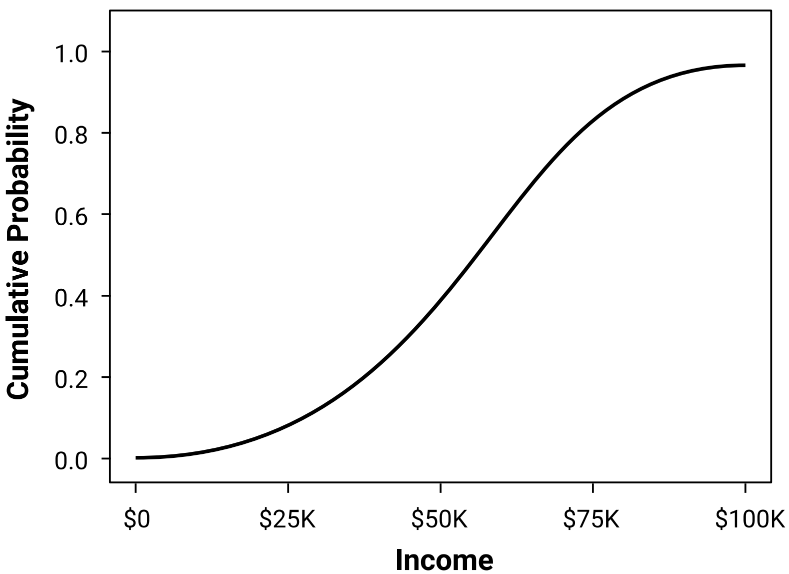

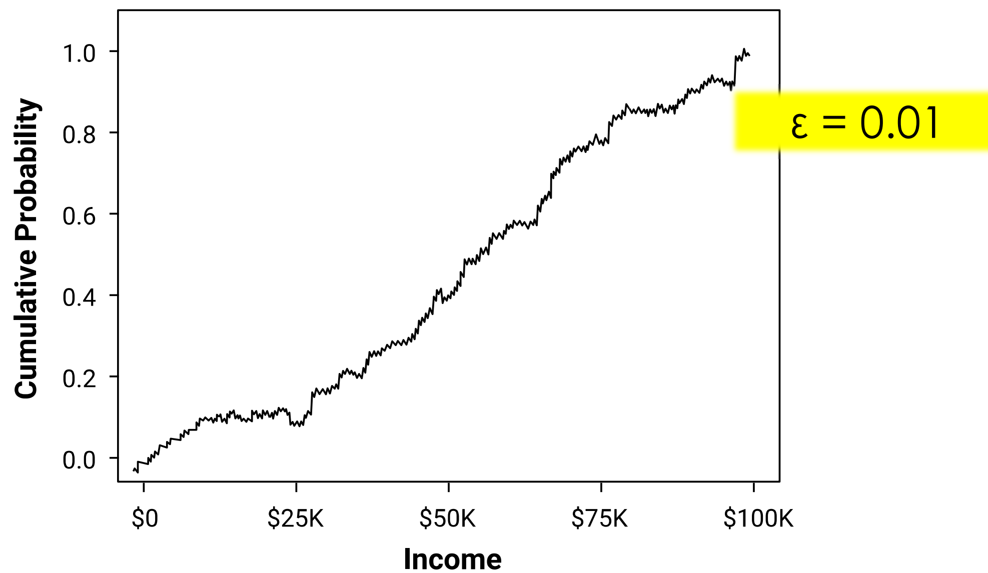

(These CDFs are stylized examples.)

Differentially Private Computations

Algorithms maintain differential privacy via the introduction of *carefully crafted* random noise into the computation.

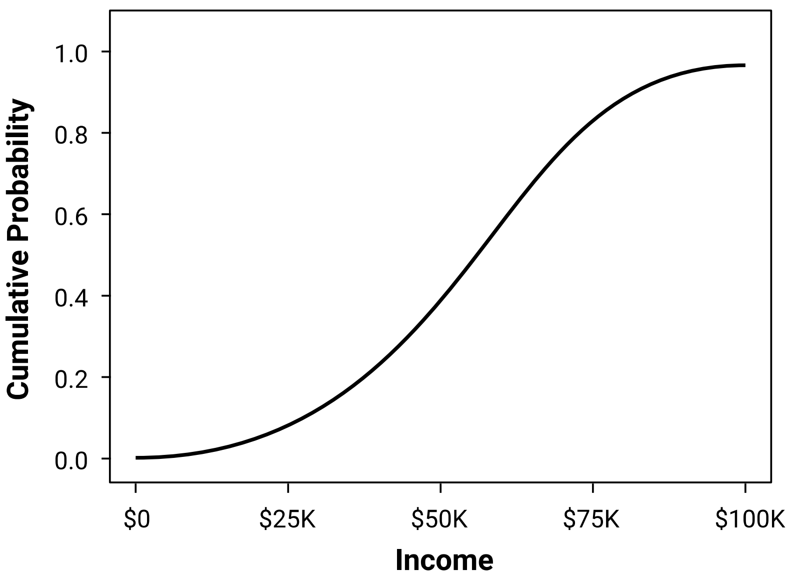

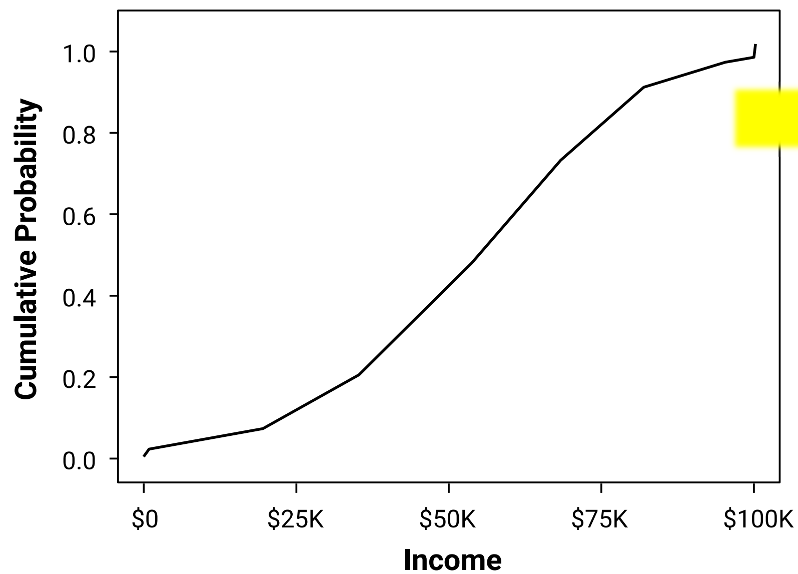

$\varepsilon = 0.1$

(These CDFs are stylized examples.)

What can be Computed with Differential Privacy?

- Descriptive statistics: counts, mean, median, histograms, boxplots, etc.
- Supervised and unsupervised ML tasks: classification, regression, clustering, distribution learning, etc.
- Generation of synthetic data

Because of noise addition, differentially private algorithms work best when the number of data records is large.

*

Applications

U.S. Census Bureau

http://onthemap.ces.census.gov

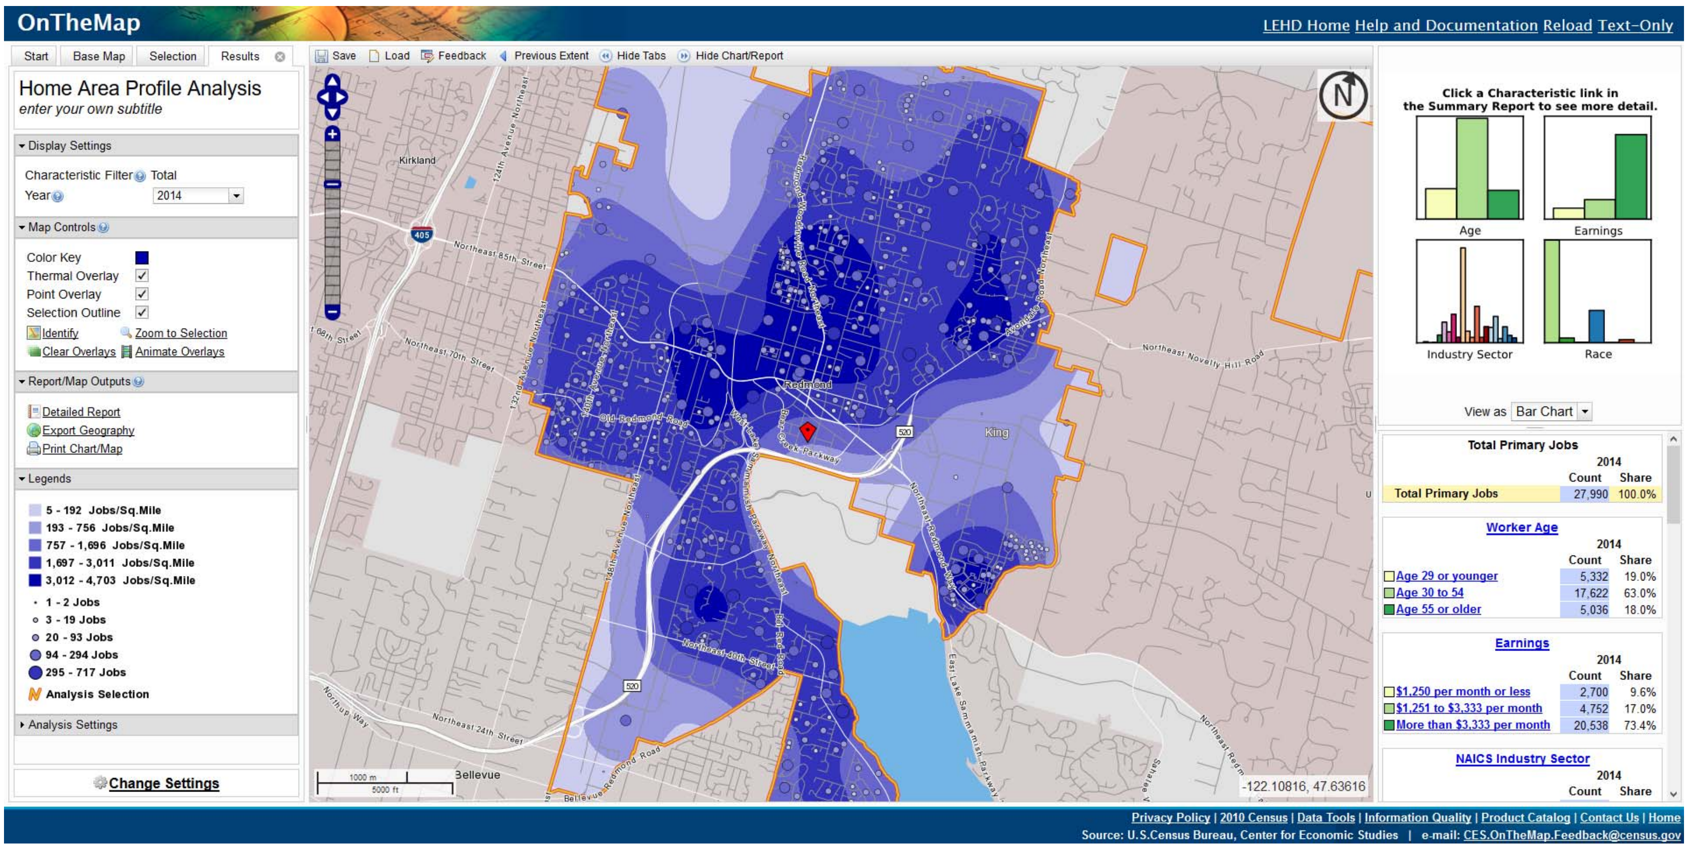

2008 AD

GOOGLE

RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response

Úlfar Erlingsson
Google, Inc.
ulfar@google.com

Vasyl Pihur
Google, Inc.
vpihur@google.com

Aleksandra Korolova
University of Southern California
korolova@usc.edu

ABSTRACT

Randomized Aggregatable Privacy-Preserving Ordinal Response, or RAPPOR, is a technology for crowdsourcing statistics from end-user client software, anonymously, with strong privacy guarantees. In short, RAPPORs allow the forest of client data to be studied, without permitting the possibility of looking at individual trees. By applying randomized response in a novel manner, RAPPOR provides the mechanisms for such collection as well as for efficient, high-utility analysis of the collected data. In particular, RAPPOR permits statistics to be collected on the population of client-side strings with strong privacy guarantees for each client, and without linkability of their reports.

This paper describes and motivates RAPPOR, details its differential-privacy and utility guarantees, discusses its practical deployment and properties in the face of different attack models, and, finally, gives results of its application to both synthetic and real-world data.

1 Introduction

Crowdsourcing data to make better, more informed decisions is becoming increasingly commonplace. For any such crowdsourcing, privacy-preservation mechanisms should be

asked to flip a fair coin, in secret, and answer "Yes" if it comes up heads, but tell the truth otherwise (if the coin comes up tails). Using this procedure, each respondent retains very strong deniability for any "Yes" answers, since such answers are most likely attributable to the coin coming up heads; as a refinement, respondents can also choose the untruthful answer by flipping another coin in secret, and get strong deniability for both "Yes" and "No" answers.

Surveys relying on randomized response enable easy computations of accurate population statistics while preserving the privacy of the individuals. Assuming absolute compliance with the randomization protocol (an assumption that may not hold for human subjects, and can even be nontrivial for algorithmic implementations [23]), it is easy to see that in a case where both "Yes" and "No" answers can be denied (flipping two fair coins), the true number of "Yes" answers can be accurately estimated by $2(Y - 0.25)$, where Y is the proportion of "Yes" responses. In expectation, respondents will provide the true answer 75% of the time, as is easy to see by a case analysis of the two fair coin flips.

Importantly, for one-time collection, the above randomized survey mechanism will protect the privacy of any specific respondent, irrespective of any attacker's prior knowl-

2014 AD

Apple

ANDY GREENBERG SECURITY 06.13.16 07:02 PM

APPLE'S 'DIFFERENTIAL PRIVACY' IS ABOUT COLLECTING YOUR DATA—BUT NOT *YOUR* DATA

2016 AD

HARVARD UNIVERSITY HARVARD.EDU

CONTACT OPEN POSITIONS

Harvard University Privacy Tools Project

Home Research ▾ News People ▾ Publications Software ▾ Outreach ▾

The Privacy Tools Project is a broad effort to advance a multidisciplinary understanding of data privacy issues and build computational, statistical, legal, and policy tools to help address these issues in a variety of contexts. It is a collaborative effort between Harvard's Center for Research on Computation and

LATEST NEWS & BLOG POSTS

Graduate Student Michael Bar-Sinai Presented at the 8th Annual ESPAnet Israel 2017

PI Salil Vadhan, PI Kobbi Nissim, and Senior Researcher Marco Gaboardi Presented at the Third Biennial Secure and Trustworthy CyberSpace Principal Investigators' Meeting (SaTC PI Meeting '17)

Berkman Klein Center Seeks Applications for 2017 Summer Internship Program

Harvard Magazine Highlights Privacy Tools Project in Article on Privacy and Security

George Kellaris Featured on CRCS Blog

Privacy Tools Project Featured in Harvard Law Review

Berkman Klein Center Seeks Fellow for Privacy

2018 AD

DP in Practice: Challenges

Transitioning to Practice

- A relatively new concept:
  - How to communicate its strengths and limitations?
  - What are the “right” use cases for implementation at this stage?
- Access to data:
  - Via a mechanism; Noise added
  - Limited by the ”privacy budget”
    - Setting the budget is a policy question
- Matching guarantees with privacy law & regulation

Conclusion

*

*

Main Takeaways

- Accumulating failures: anonymization & traditional SDL techniques
- Differential privacy:
  - A standard providing a rigorous framework for developing privacy technologies with provable quantifiable guarantees
  - Rich theoretical work, now transitioning to practice
    - First real-world applications and use
  - Not a panacea; to be combined (wisely!) with other technical and policy tools

Resources

*

*

Learning More About Differential Privacy

- [Page et al, 2018] Differential Privacy: An Introduction For Statistical Agencies, UK ONS.
- [Wood et al, 2019] Differential Privacy: A Primer for a Non-technical Audience, Vanderbilt JETLaw.
- [Nissim et al, 2018] Bridging the gap between computer science and legal approaches to privacy, Harvard JOLT.
- [Dwork 2011] A Firm Foundation for Private Data Analysis, CACM January 2011.
- [Heffetz & Ligett, 2014] Privacy and Data-Based Research, Journal of Economic Perspectives.
- [Dwork & Roth, 2014] The Algorithmic Foundations of Differential Privacy, Now publishers.
- \+ [Vadhan, 2017] The Complexity of Differential Privacy

less technical

technical

*

Projects, Software Tools [Partial List]

[Microsoft Research] PINQ

[UT Austin] Airavat: Security & Privacy for MapReduce

[UC Berkeley] GUPT

[CMU-Cornell-PennState] Integrating Statistical and Computational Approaches to Privacy

[US Census] OnTheMap

[Google] Rappor, TensorFlow Privacy

[UCSD] Integrating Data for Analysis, Anonymization, and Sharing (iDash)

[UPenn] Putting Differential Privacy to Work

[Stanford-Berkeley-Microsoft] Towards Practicing Privacy

[Duke-NISS] Triangle Census Research Network

[Harvard] Privacy Tools

[Georgetown-Harvard-BU] Formal Privacy Models and Title 13

[Harvard-Georgetown-Buffalo] Computing over Distributed Sensitive Data

Backup Slides

The Privacy Tools Project

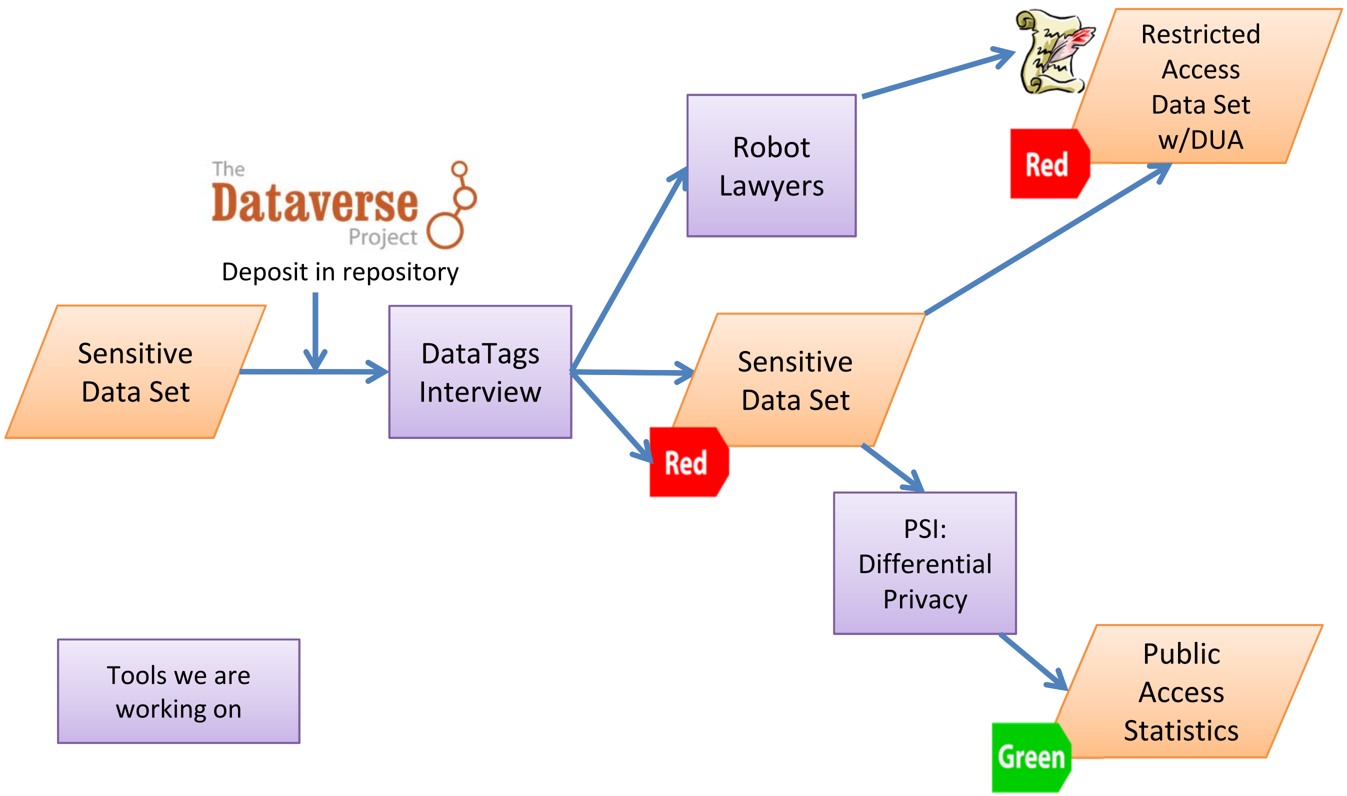

*

The Privacy Tools Project: Robot Lawyers

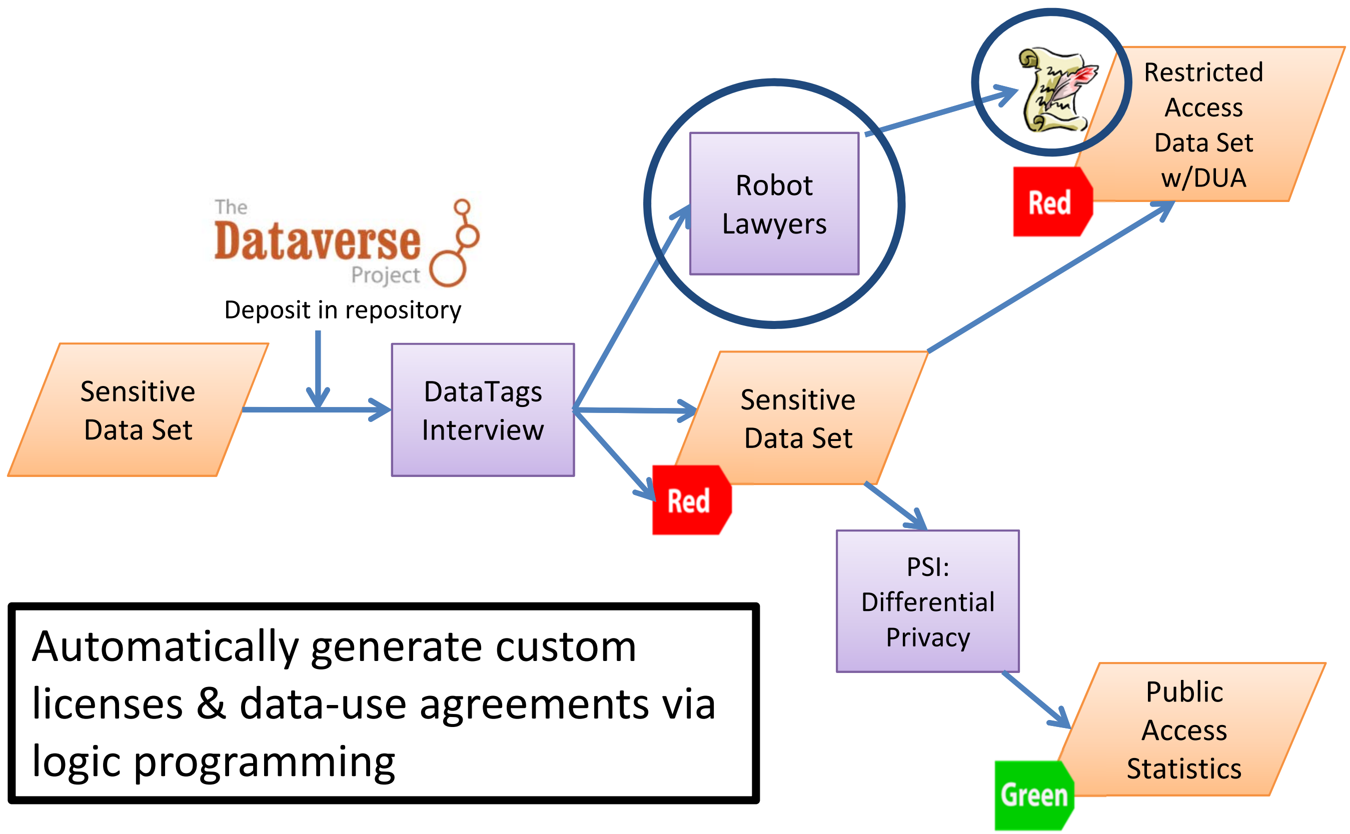

Automatically generate custom licenses & data-use agreements via logic programming

*

The Privacy Tools Project: PSI

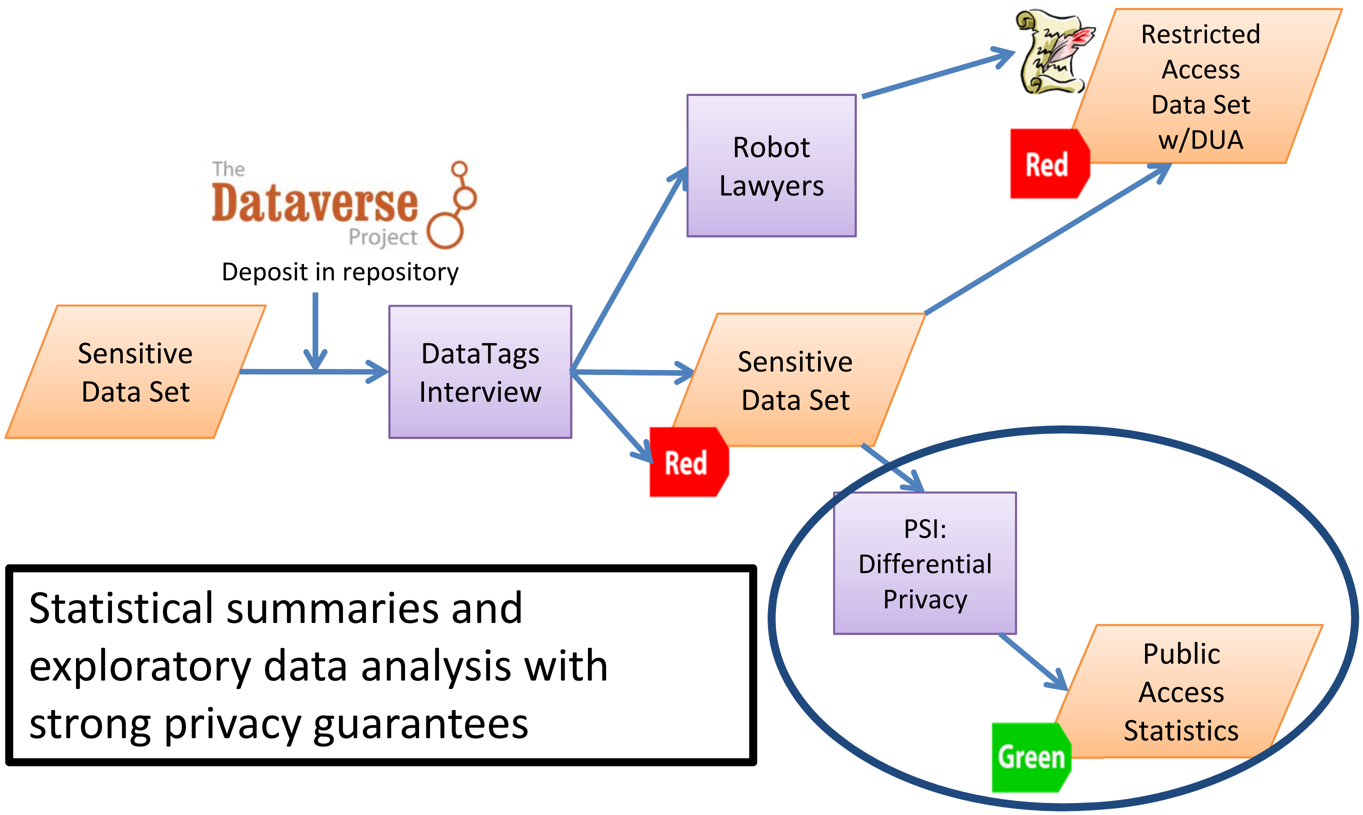

Statistical summaries and exploratory data analysis with strong privacy guarantees

*

The Privacy Tools Project: Bridging Defs

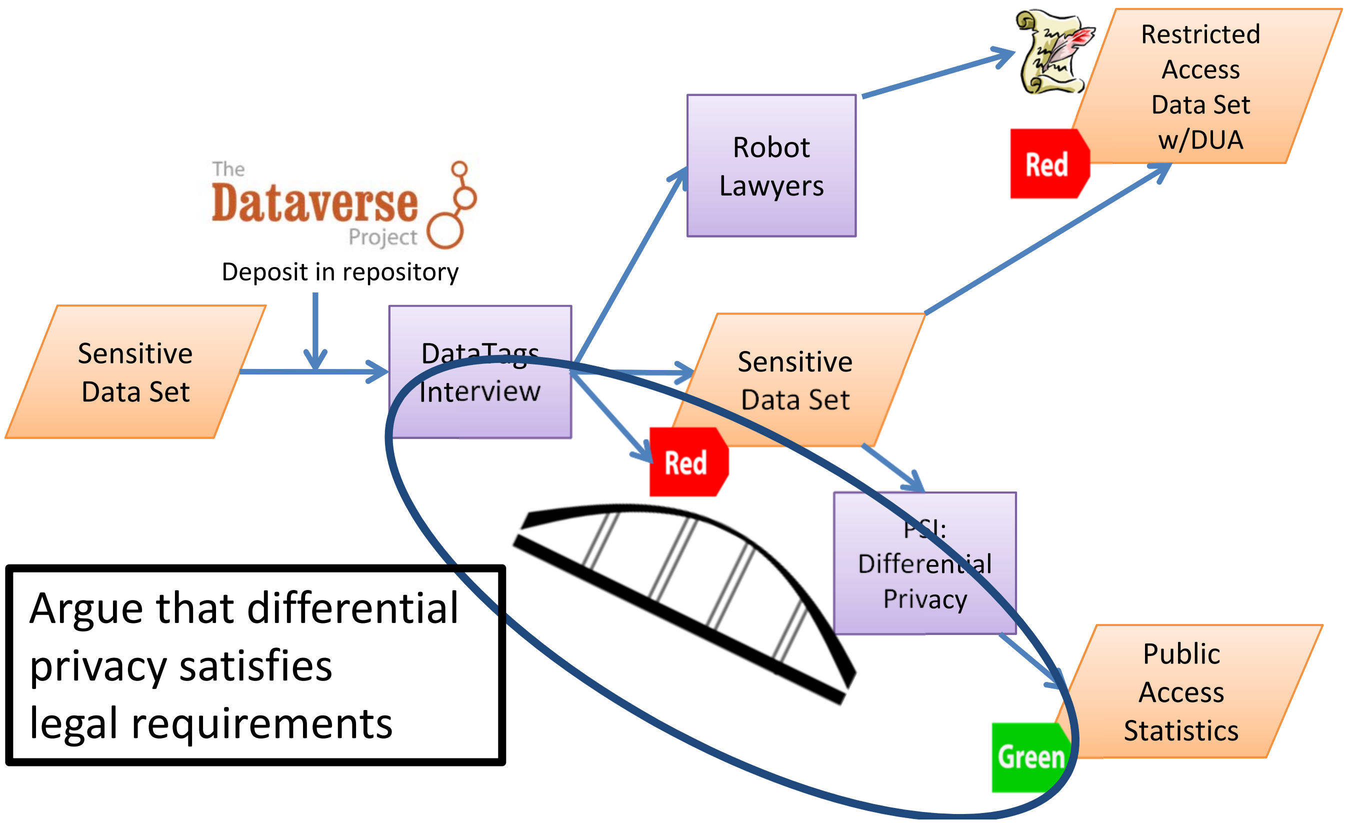

Argue that differential privacy satisfies legal requirements