

# Leveraging Privacy in Data Analysis

**Ryan Rogers**


Applied Mathematics and Computational Sciences Department

Dissertation Defense

Advisors: Michael Kearns and Aaron Roth

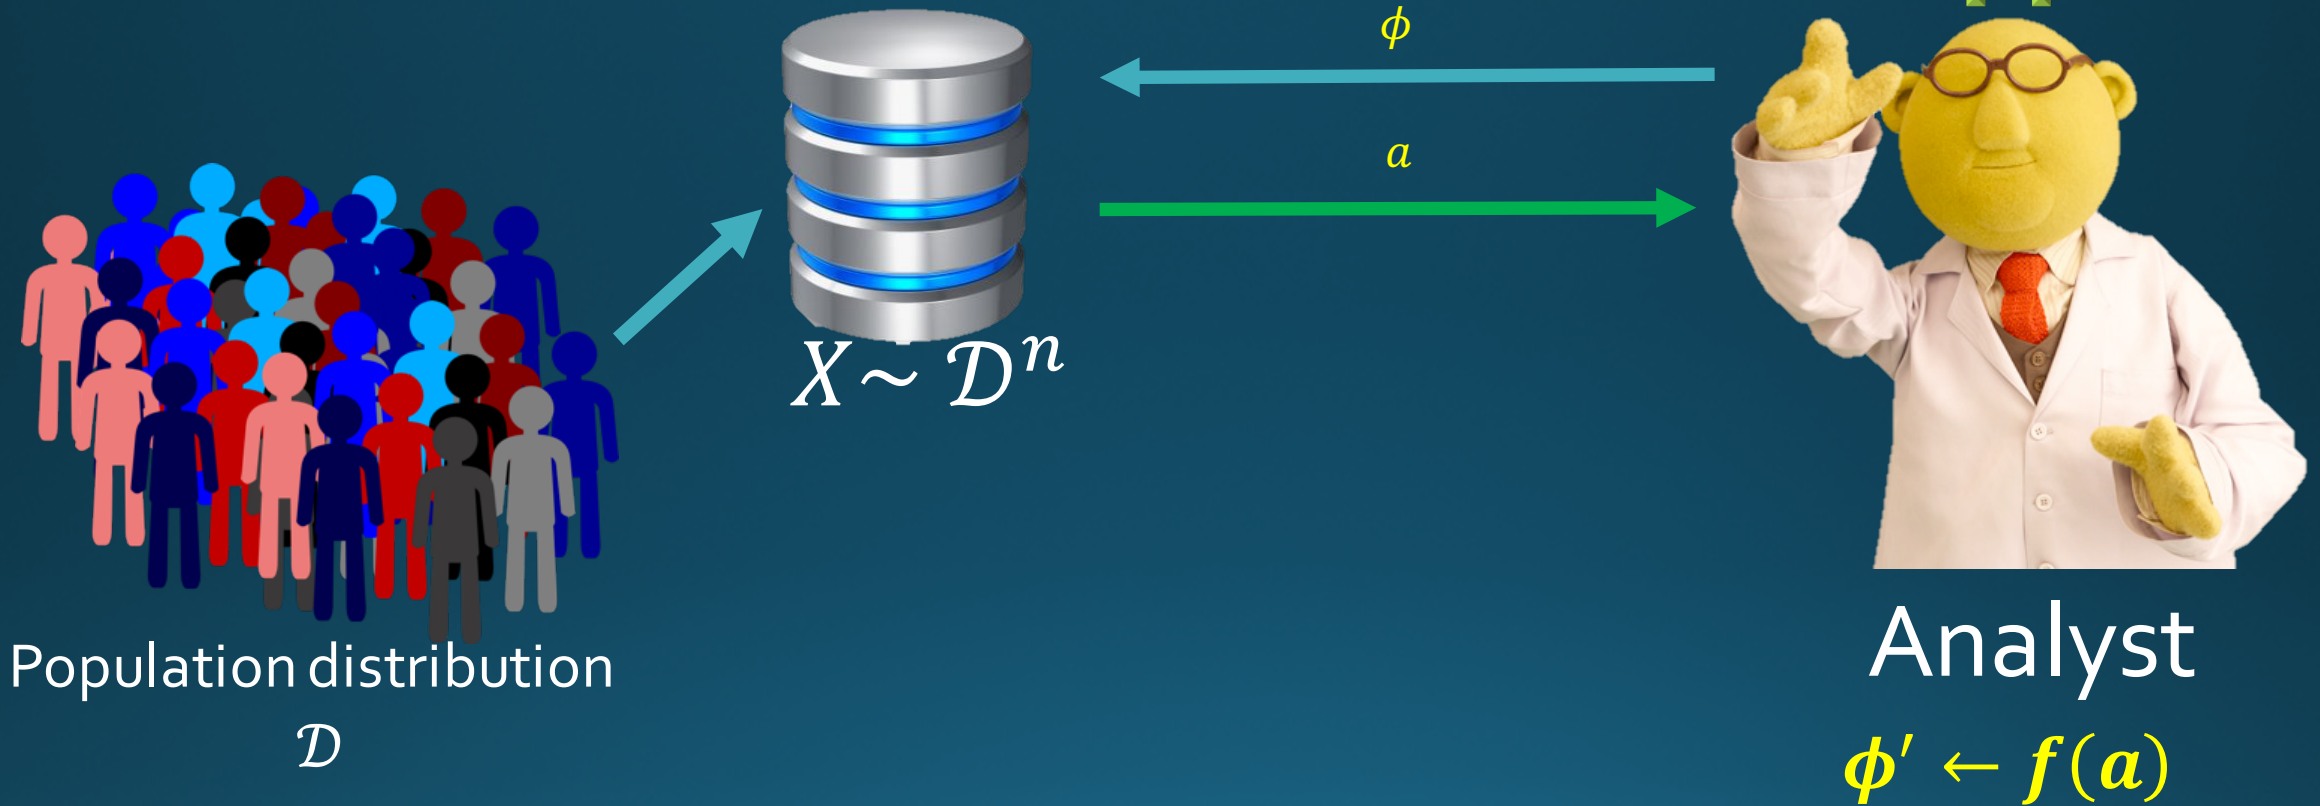


# Data Analysis






# Data Analysis

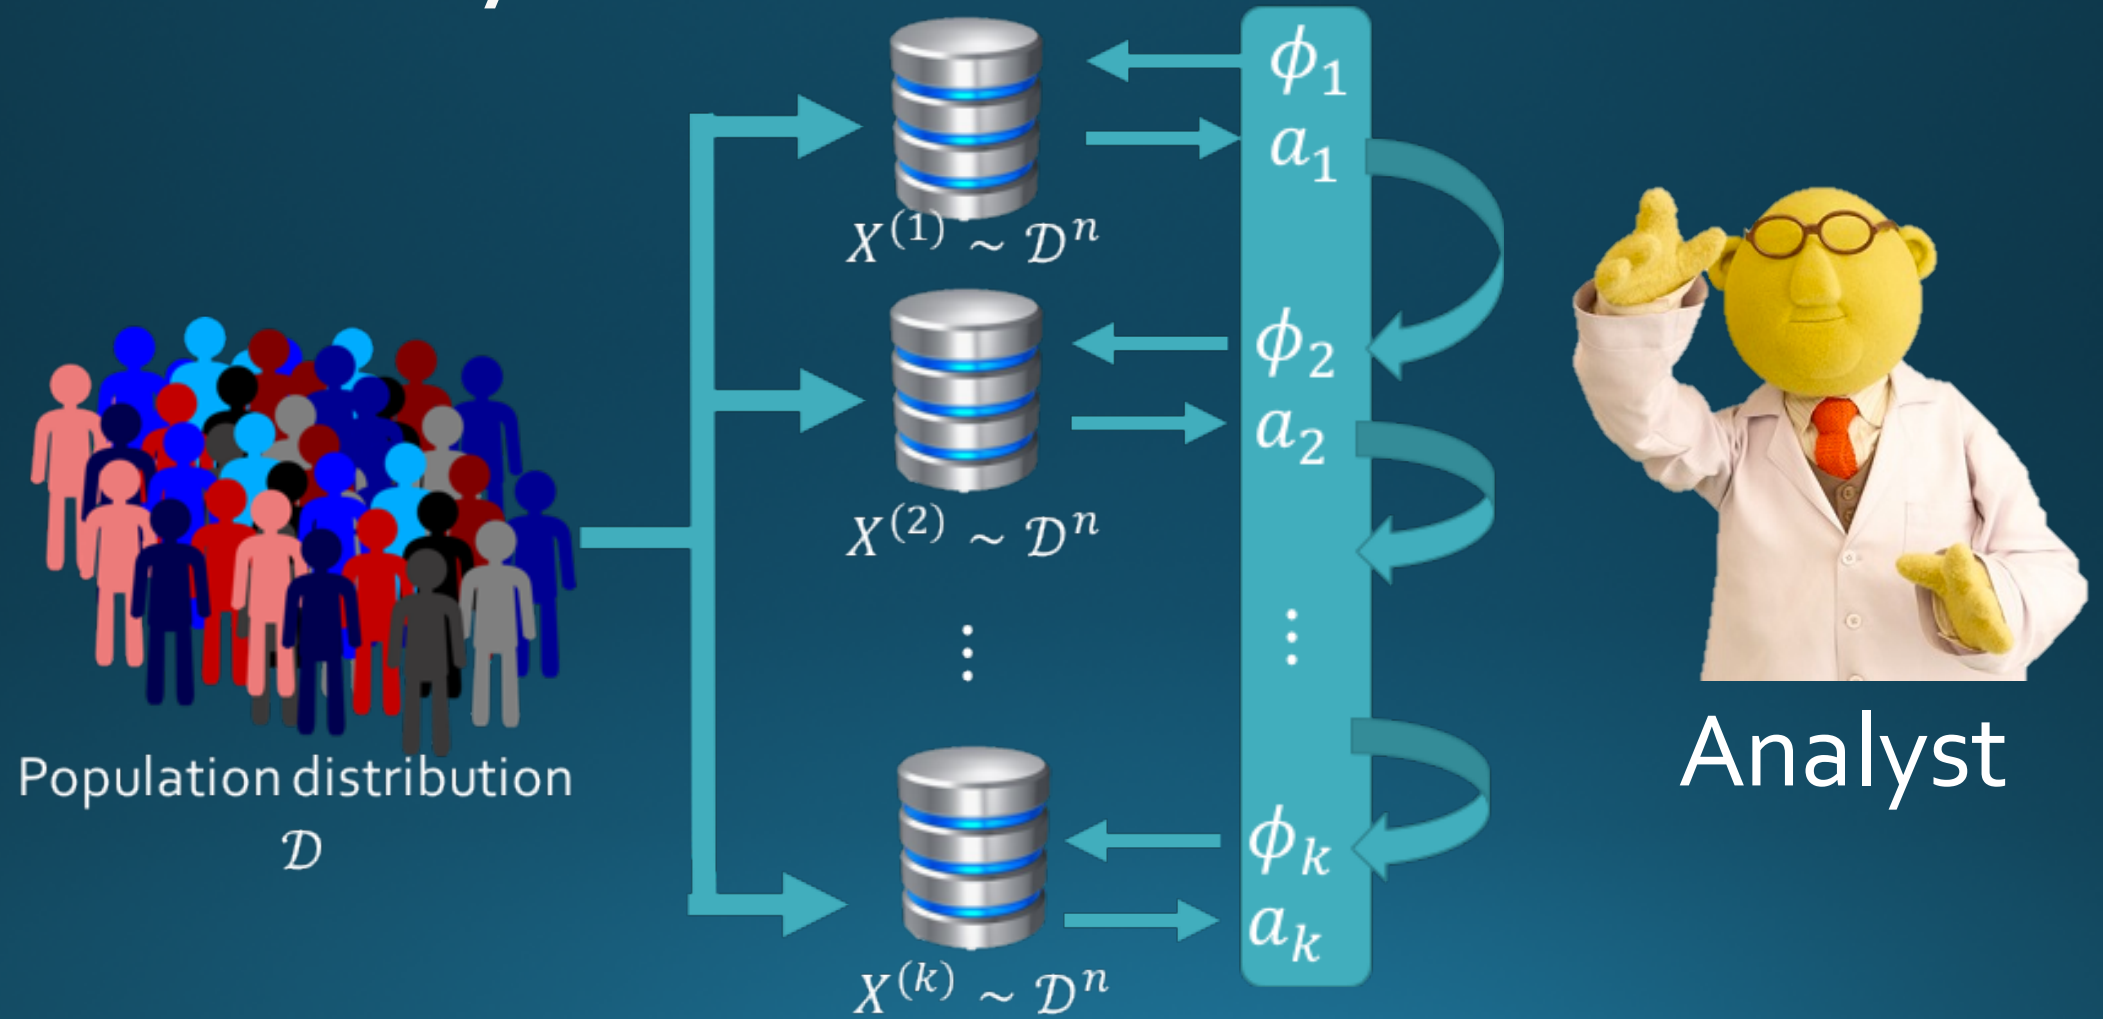


# Data Analysis - Ideal

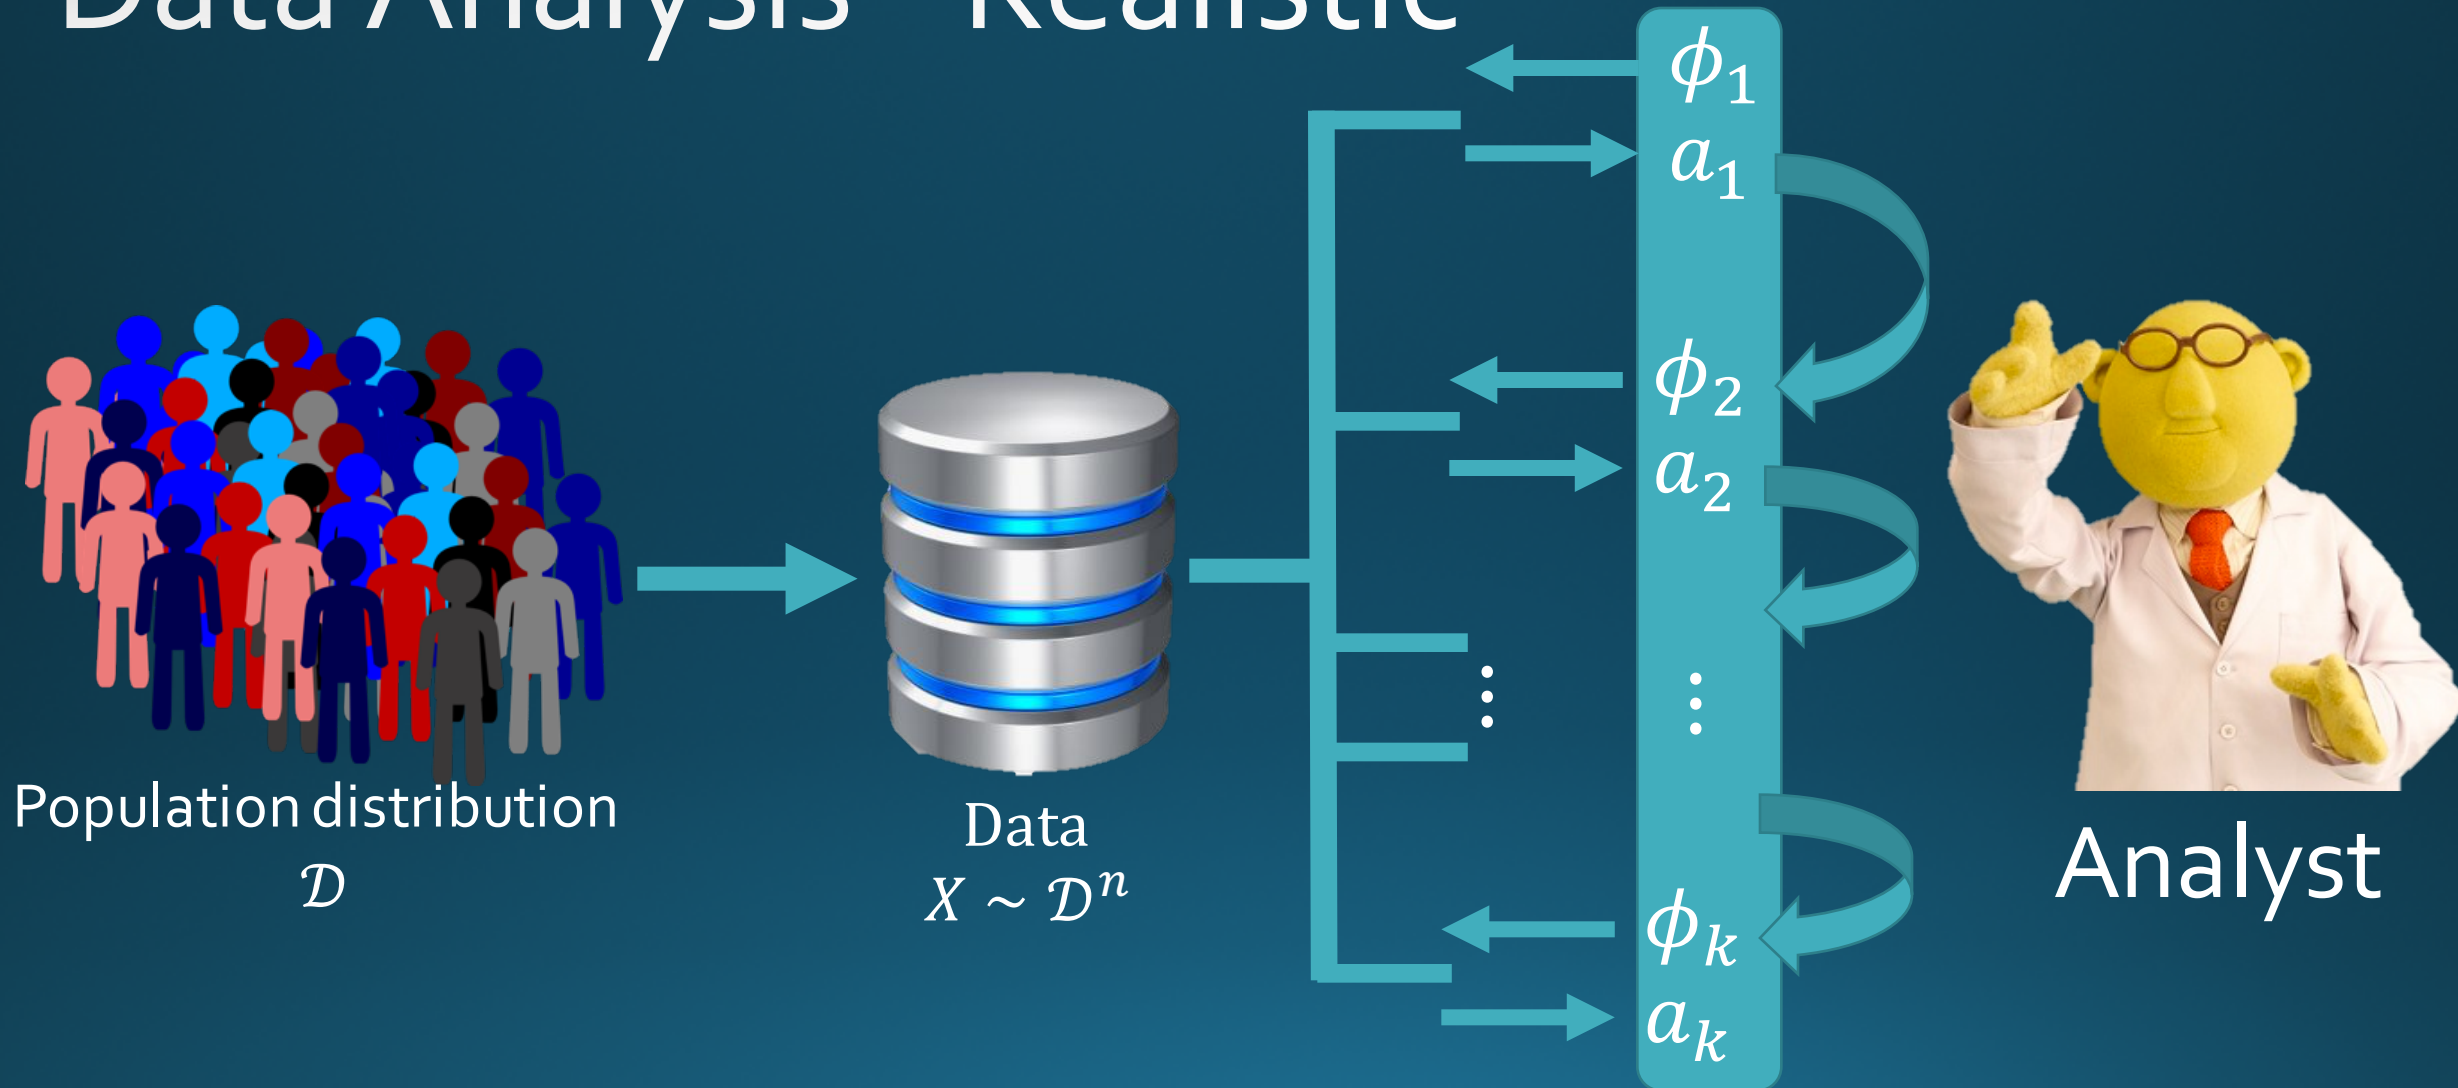


A lot of existing theory assumes tests are selected independently of the data.

# Data Analysis - Ideal



# Data Analysis - Realistic





# Data Analysis - Realistic

NOBA  
Browse Content / The Replication Crisis in Psychology

## The Replication Crisis in Psychology

By Edward Diener and Robert Biswas-Diener  
University of Illinois at Urbana-Champaign, Scientists like to think of science as self-correcting. To an alarming degree, it is

Unreliable research  
Trouble at the lab

World politics Business & finance Economics Science

## P-HACKING

OPEN ACCESS  
ESSAY

### Why Most Published Research Findings Are False

John P. A. Ioannidis  
Published: August 30, 2005 • DOI: 10.1371/journal.pmed.0020124

Article Authors Metrics

Abstract  
Modeling the Framework for False Positive Findings



## A Look into Macroeconomics: First Cheating Scandal

BY DANIEL WALTER  
DECEMBER 08, 2015  
DECEMBER 08, 2015  
3 COMMENTS

At the end of May, the Image (currently the word's top competition policies by the Chinese search giant) has been banned from submission review called it "Ma

A colorful cartoon illustration of a laboratory scene. Several scientists are depicted in various states of activity, some looking at equipment, others talking. The scene is filled with lab equipment like flasks, beakers, and a microscope. The artist's signature 'Jason Ford' is visible at the bottom right.

## The Statistical Crisis in Science

Data-dependent analysis—a “garden of forking paths”—explains why many statistically significant comparisons don’t hold up.

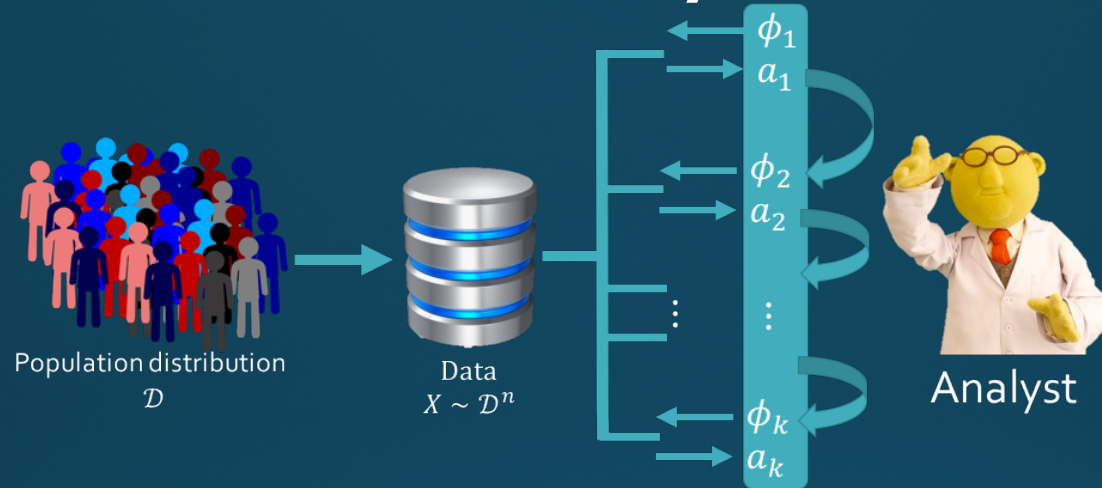
Andrew Gelman and Eric Loken

There is a growing realization that reported “statistically significant” claims in scientific publications are routinely mis-

a short mathematics test when it is expressed in two different contexts, involving either healthcare or the military. The question may be framed in terms of tested relationships; where the effect sizes are small, and analytical modes: when the

This multiple comparisons issue is well known in statistics and has been called “p-hacking” in an influential 2011 paper by the psychology researcher John Ioannidis. When effect sizes are small, the power of the test is low, and the number of tests conducted is large, the probability of finding a statistically significant result by chance alone is high.

# Adaptive Data Analysis




**Question:** How can we provide statistically valid answers to adaptively chosen analyses?

1. Traditional method – split the dataset into  $k$  chunks.
  - Requires  $k \ll n$ .
2. Limit the info learned about the dataset with each analysis [Dwork, Feldman, Hardt, Pitassi, Reingold, Roth'15].
  - Can handle  $k \gg n$ .



# Adaptive Data Analysis [DFHPRR'15]



# Example – Statistical Queries

- Let each analysis  $\phi_j: \mathcal{X} \rightarrow [0,1]$  for  $j = 1, \dots, k$ . We want to estimate  $\phi_j(\mathcal{D}) = \mathbb{E}_{Y \sim \mathcal{D}}[\phi_j(Y)]$


- Want to design algorithms  $\mathcal{M}_j$  such that  $\mathcal{M}_j(X) = a_j$  and

$$\max_{j \in \{1, \dots, k\}} \{ |a_j - \phi_j(\mathcal{D})| \} \leq \tau \quad \text{w.p.} \geq 1 - \beta$$

- [DFHPRR'15]:  $\mathcal{M}_j(X) = \frac{1}{n} \sum_{i=1}^n \phi_j(X_i) + N(0, \sigma^2)$

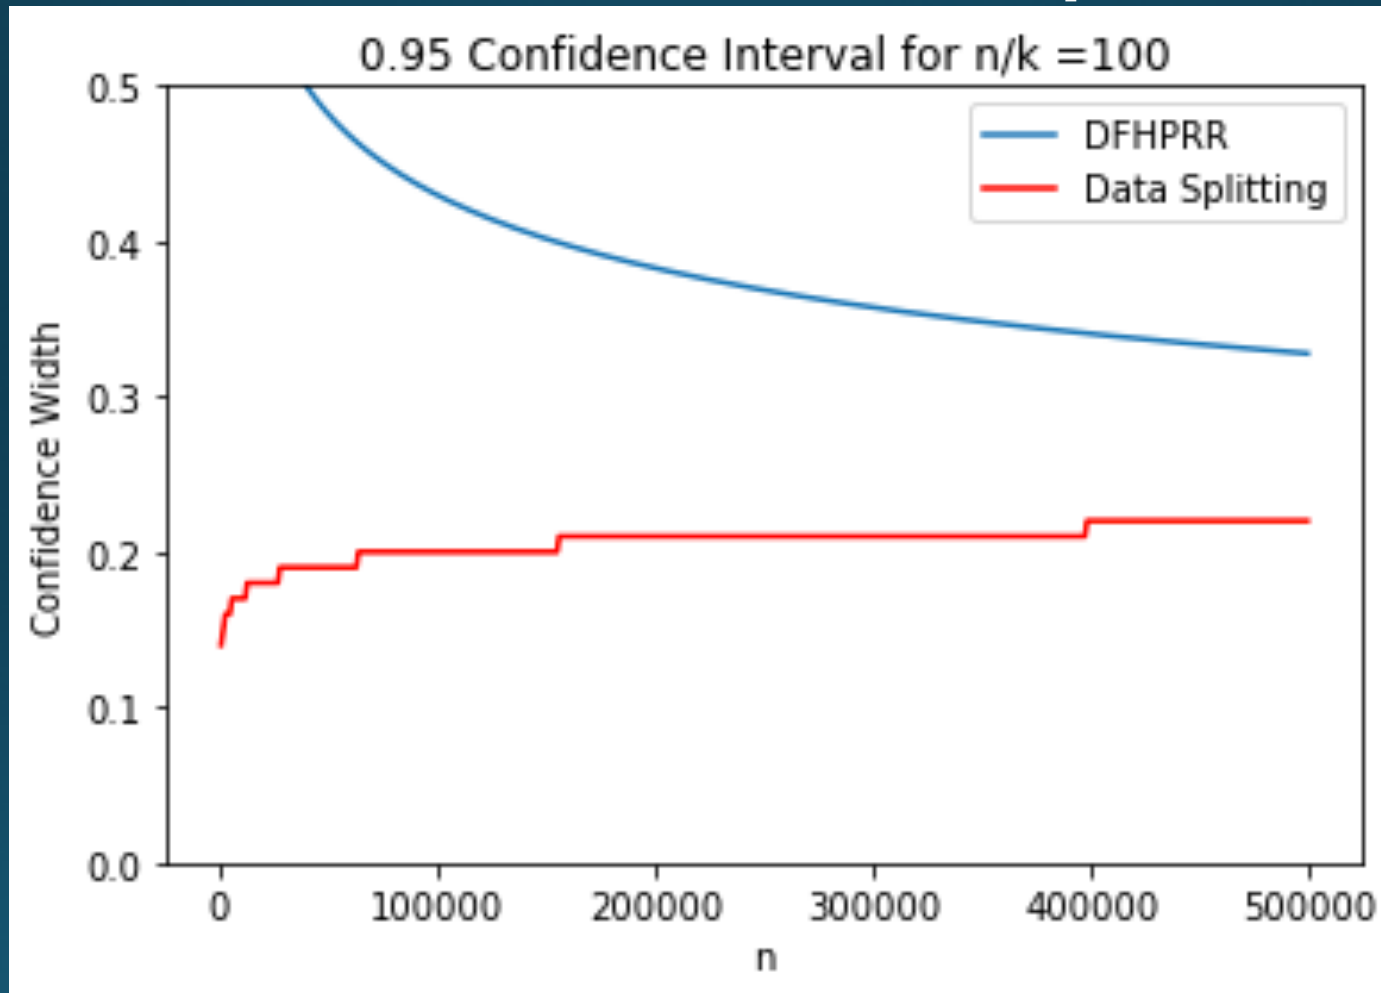
so for properly chosen  $\sigma$ , we can have  $k \sim n^2$

# Comparison with Data Splitting




For each of the  $k$  queries, bound accuracy for dataset of size =  $n/k$ , then union bound.


# Comparison with Data Splitting



# Comparison with Data Splitting



# Comparison with Data Splitting



Ongoing work  
with Roth,  
Smith, Thakkar



# Questions

- What types of algorithms can we use to ensure validity?
- Can we do more general types of analyses, beyond linear queries?
- What existing techniques can we use, and what needs to be modified for adaptivity and the algorithms we use?



# Related Work

- Lots of work in statistics community on selective inference  
[Freedman'83],[Leeb,Potscher'06],[Berk,Brown,Buja,Zhang,Zhao'13], ...
  - Specific to type of analyses performed
- [DFHPRR](STOC'15,NIPS'15,Science'15)
  - Initial connections between information, privacy and adaptive analysis
- Accuracy for specific queries
  - [DFHPRR] (STOC'15,Science'15)
  - [Bassily,Nissim,Smith,Steinke,Stemmer,Ullman'16]
  - [Cummings,Ligett,Nissim,Roth,Wu'16]
  - [Russo,Zou'16]
  - [Wang,Lei,Fienberg'16]
- Impossibility results
  - [Hardt,Ullman'14], [Steinke,Ullman'15]

# Outline

- Post Selection Hypothesis Testing [R,Roth,Smith,Thakkar FOCs'16]
  - Connection between Max-Info and Differential Privacy
- DP composition with adaptively selected parameters [R,Roth,Ullman,Vadhan NIPS'16]
  - Privacy Odometers and Filters
- Private Hypothesis Tests [Gaboardi,Lim,R,Vadhan ICML'16],[Kifer,R AISTATS'16]
  - Chi-Square Tests
- Directions for Future Work

# Max-Information [DFHPRR'15]

- Algorithm  $\mathcal{M}$  has small max-info  
 $\Rightarrow \mathcal{M}(X)$  and  $X$  are “close” to **independent**.
- The  **$\beta$ -approximate max-info** between  $\mathcal{M}(X)$  and  $X$

$$I_{\infty}^{\beta}(\mathcal{M}(X); X) = \log \left( \sup_{\mathcal{O}} \frac{\mathbb{P}[(\mathcal{M}(X), X) \in \mathcal{O}] - \beta}{\mathbb{P}[(\mathcal{M}(X'), X) \in \mathcal{O}]} \right)$$



Real World



Ideal World

# Max-Information of Algorithms [DFHPRR'15]

$$I_{\infty}^{\beta}(\mathcal{M}(X); X) = \log \left( \sup_o \frac{\mathbb{P}[(\mathcal{M}(X), X) \in o] - \beta}{\mathbb{P}[(\mathcal{M}(X'), X) \in o]} \right)$$

The  $\beta$ -approximate max-info of an algorithm  $\mathcal{M}$  on data sets of size  $n$  is:

$$I_{\infty, \Pi}^{\beta}(\mathcal{M}; n) = \sup_{\mathcal{D}: X \sim \mathcal{D}^n} \left\{ I_{\infty}^{\beta}(\mathcal{M}(X); X) \right\}$$

# Hypothesis Testing

- Hypothesis test  $t: \mathcal{X}^n \rightarrow \{Inconclusive, Reject H_0\}$  is defined by
  - null hypothesis  $H_0 \subseteq \Delta(\mathcal{X})$  and
  - statistic:

$$g: \mathcal{X}^n \rightarrow \mathbb{R}$$

	$H_0$ is True	$H_0$ is False
$t$ rejects	$\alpha = \text{False Discovery}$	Power
$t$ fails to reject	$1-\alpha = \text{Significance}$	Type II Error

- Want to bound  $\mathbb{P}[\text{False Discovery}] \leq \alpha$  and get good Power.




# $p$ -Values

- The  $p$ -value associated with a value  $y$  and a distribution  $\mathcal{D} \in H_0$  is given as

$$p(y) = \mathbb{P}_{X \sim \mathcal{D}^n} [g(X) > y]$$

- Denotes the prob of observing a value of the test statistic that is at least as extreme as  $y$ .



- Note that  $p(g(X)) \sim U[0,1]$  if  $X \sim \mathcal{D}^n$  where  $\mathcal{D} \in H_0$ .
- If we reject the model when  $p(g(X)) < \alpha$  then
$$\mathbb{P}[\text{False Discovery}] < \alpha.$$
- No longer true when  $t \leftarrow \mathcal{M}(X)$ !

# $p$ -Value Corrections

- Even when we use the data to determine a test, we still want to be able to control the  $\mathbb{P}[\text{False Discovery}]$ .
- A function  $\gamma: [0,1] \rightarrow [0,1]$  is a **valid  $p$ -value correction** function for a selection procedure  $\mathcal{M}: \mathcal{D}^n \rightarrow \mathcal{O}$  if for every  $\alpha$  the procedure:
  1. Select test  $t \leftarrow \mathcal{M}(X)$
  2. Reject  $H_0$  if  $p(g(X)) < \gamma(\alpha)$has probability at most  $\alpha$  of false discovery.

# Max-Info gives Valid $p$ -Value Corrections

- If we have selection procedure  $\mathcal{M}$  such that  $I_{\infty, \Pi}^{\beta}(\mathcal{M}, n) \leq m$  then a valid  $p$ -value correction function is

$$\gamma(\alpha) = \frac{\alpha - \beta}{2^m}$$

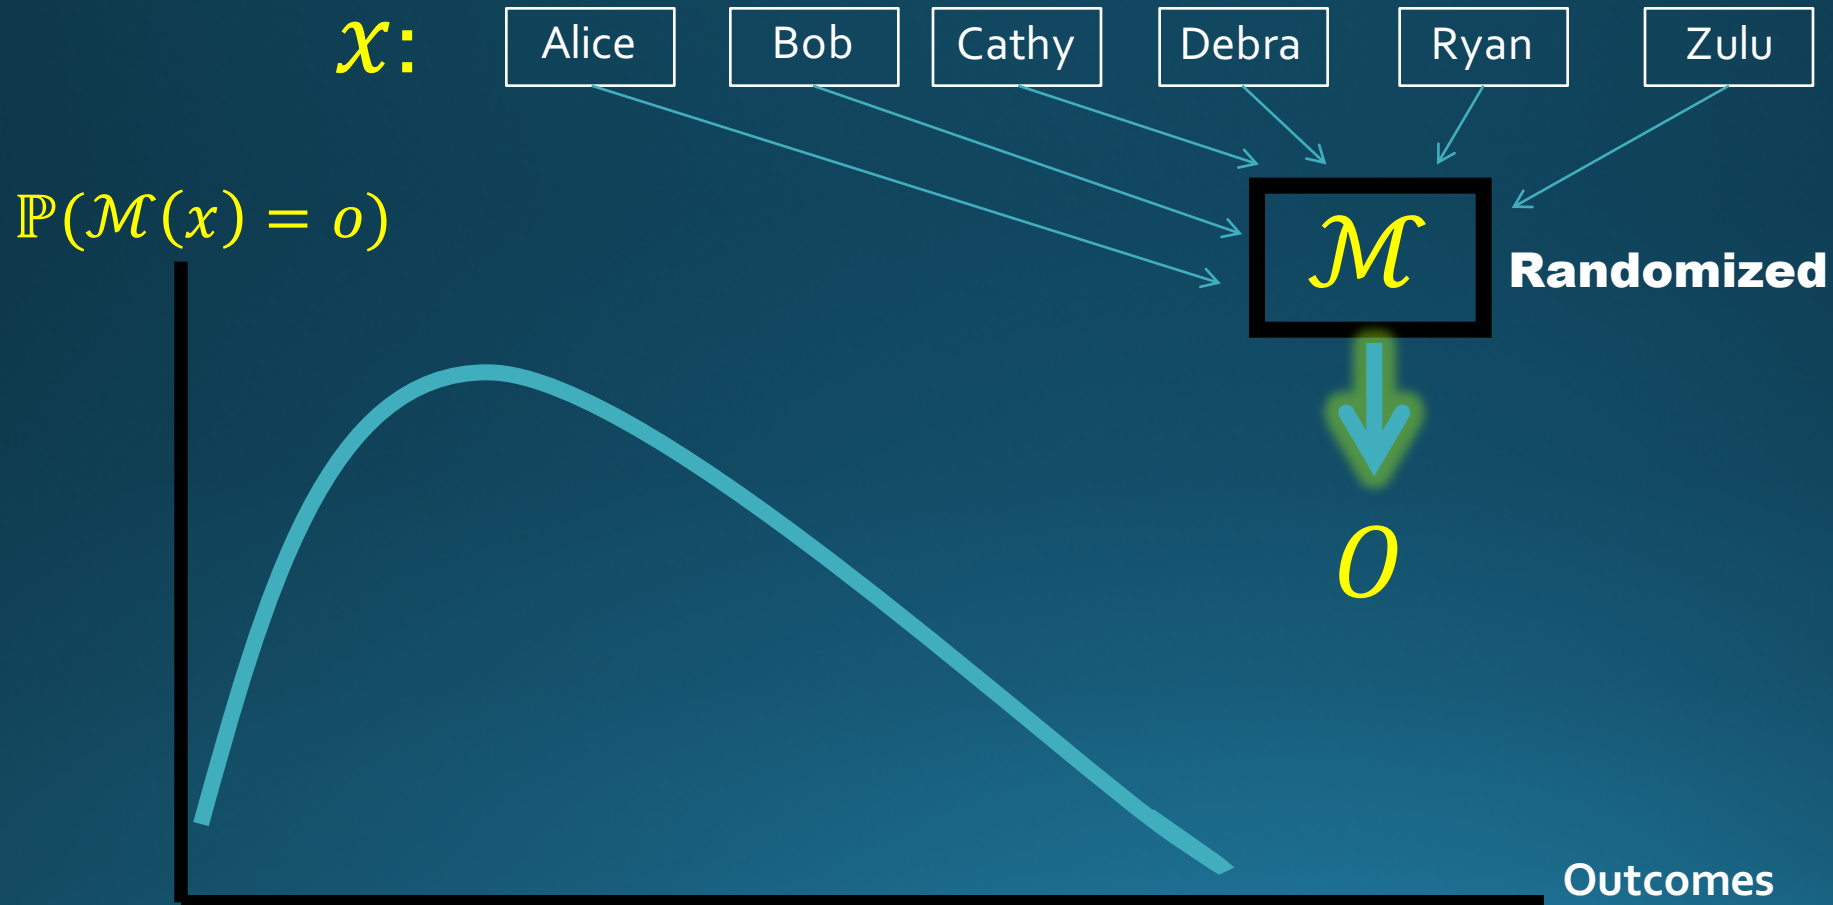
- Proof: Let  $S \subseteq \mathcal{X}^n \times \mathcal{O}$  be the event that  $\mathcal{M}$  selects a test statistic where the  $p$ -value is at most  $\gamma(\alpha)$ , but the null is true.

$$\begin{aligned} & \mathbb{P}[p(g(X)) \leq \gamma(\alpha) \cap t \leftarrow \mathcal{M}(X) | H_0] \\ &= \mathbb{P}[(X, \mathcal{M}(X)) \in S | H_0] \\ &\leq 2^m \underbrace{\mathbb{P}[(X, \mathcal{M}(X')) \in S | H_0]}_{\leq \gamma(\alpha)} + \beta \end{aligned}$$

# What procedures $\mathcal{M}$ have bounded max-info?

- [DFHPRR'15] Max-information bounds for:
  - (Pure) Differential Privacy – algorithmic stability condition.
  - Description Length –  $\log(\text{image size of } \mathcal{M} )$

# Differential Privacy [Dwork, McSherry, Nissim, Smith'o6]



# Differential Privacy [Dwork, McSherry, Nissim, Smith'o6]

$\mathcal{X}$ : Alice Bob Cathy Debra Ryan Zulu

$$\mathbb{P}(\mathcal{M}(x) = o)$$



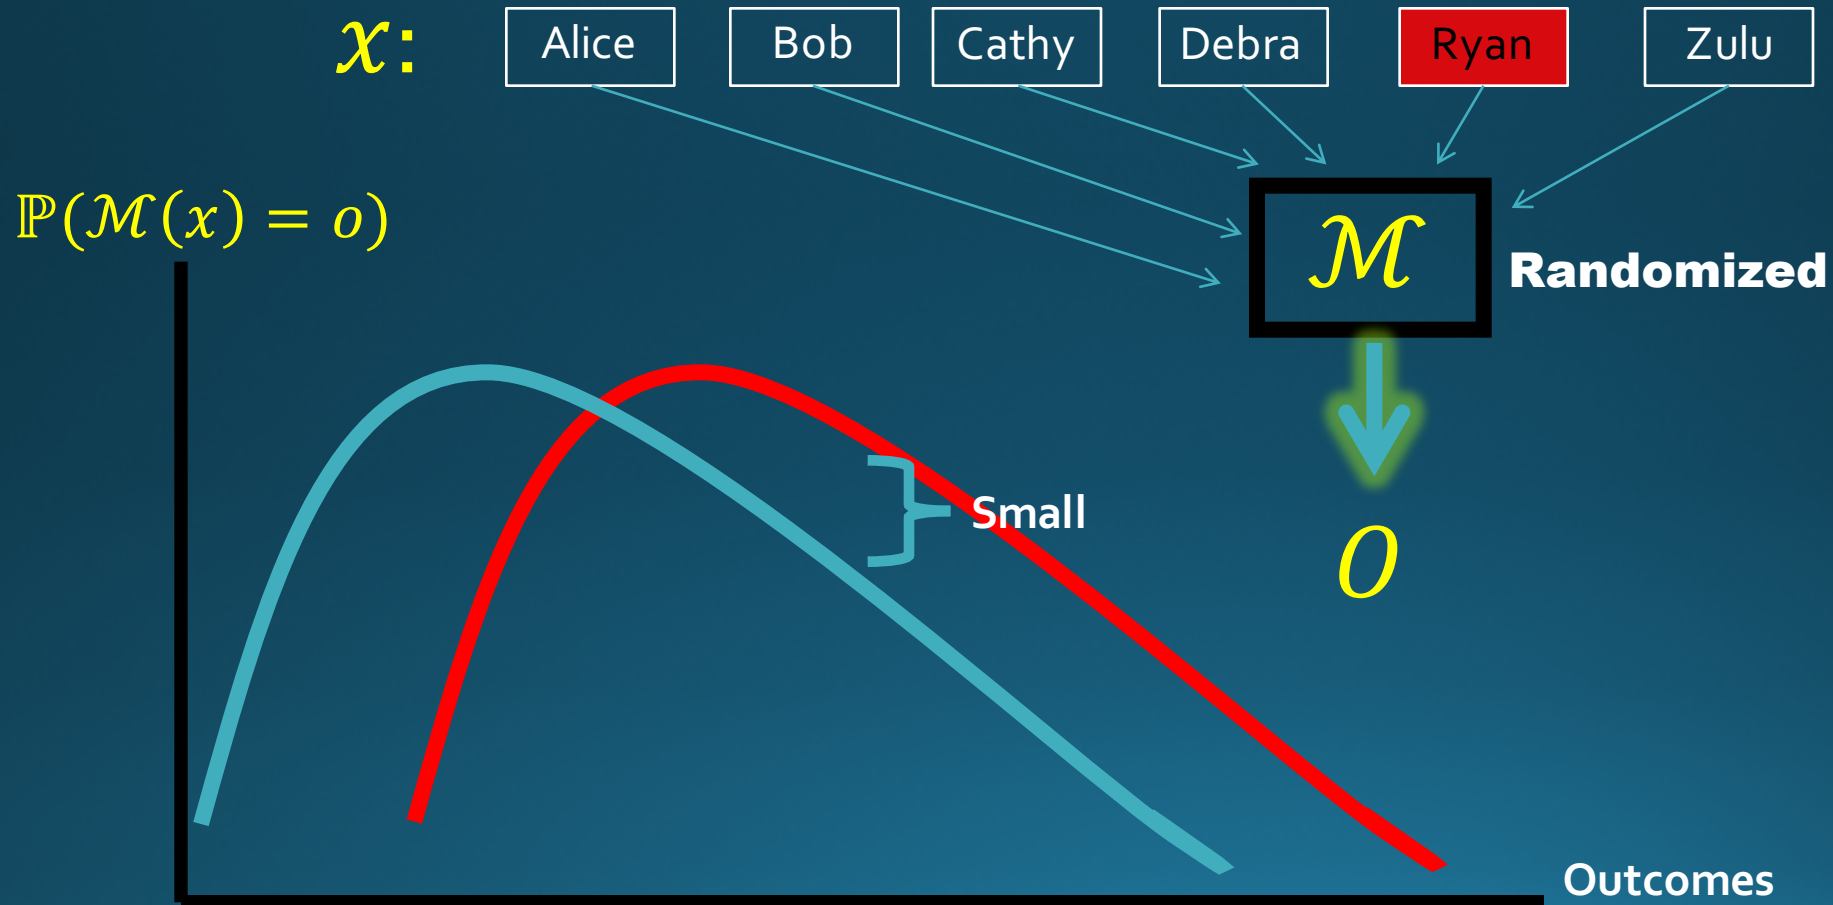
**Randomized**

$o$

Outcomes



# Differential Privacy [Dwork, McSherry, Nissim, Smith'o6]



# Differential Privacy [Dwork,McSherry,Nissim,Smith'o6]

- A randomized algorithm  $\mathcal{M}: \mathcal{D}^n \rightarrow \mathcal{O}$  is  $(\epsilon, \delta)$ -differentially private if for any neighboring data sets  $x, x' \in \mathcal{X}^n$  and any outcome  $S \subseteq \mathcal{O}$  we have

$$\mathbb{P}(\mathcal{M}(x) \in S) \leq e^\epsilon \mathbb{P}(\mathcal{M}(x') \in S) + \delta$$

If  $\delta = 0$  we say **pure** DP, and otherwise **approximate** DP.


# DP Composition

- If we run  $k$  many  $(\epsilon, 0)$ -DP algorithms  $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_k$  on the same data set, then:
  - [DMNS'06]: The composed algorithm  $\mathcal{M} = \mathcal{M}_k \circ \dots \circ \mathcal{M}_1$  is:  
 $(\epsilon k, 0)$ -DP.
  - [Dwork,Rothblum,Vadhan'10]: The composed algorithm  $\mathcal{M} = \mathcal{M}_k \circ \dots \circ \mathcal{M}_1$  is:  
 $(O(\epsilon\sqrt{k \log(1/\delta)}), \delta)$ -DP.

Quadratic  
Improvement  
with small  $\delta > 0$ !

# Post-selection Hypothesis Testing

- [RRST'16] Connection between Max-Information and (approx)-differential privacy



# Technical Contribution

- [DFHPRR'15]: If  $\mathcal{M}: \mathcal{X}^n \rightarrow \mathcal{O}$  is  $(\epsilon, 0)$ -DP, then for  $\beta > \frac{1}{\epsilon}$

$$I_{\infty, \Pi}^{\beta}(\mathcal{M}; n) \leq \tilde{O}(\epsilon^2 n)$$

Similar Max-Info bounds

- [RRST'16]: If  $\mathcal{M}: \mathcal{X}^n \rightarrow \mathcal{O}$  is  $(\epsilon, \delta)$ -DP, then

$$I_{\infty, \Pi}^{\beta}(\mathcal{M}; n) \leq \tilde{O}(\epsilon^2 n) \text{ where } \beta \approx n \sqrt{\frac{\delta}{\epsilon}}$$



# Consequences of Positive Result

Theorem: If  $\mathcal{M}: \mathcal{X}^n \rightarrow \mathcal{O}$  is  $(\epsilon, \delta)$ -DP, then

$$I_{\infty, \Pi}^{\beta}(\mathcal{M}; n) \leq \tilde{O}(\epsilon^2 n) \text{ where } \beta \approx n \sqrt{\frac{\delta}{\epsilon}}$$

- Recover (optimal) results of [BNSSSU'16] for low sensitive queries.
  - However, our bounds apply more generally (e.g. adaptive hypothesis tests).
- Composition of  $k$  adaptively selected  $(\epsilon, 0)$ -DP procedures:  $\mathcal{M}_1, \dots, \mathcal{M}_k$ 
  - [DFHPRR'15]:  $I_{\infty, \Pi}^{\beta}(\mathcal{M}_k \circ \dots \circ \mathcal{M}_1; n) \leq \tilde{O}(n\epsilon^2 k^2)$
  - [RRST'16]:  $I_{\infty, \Pi}^{\beta}(\mathcal{M}_k \circ \dots \circ \mathcal{M}_1; n) \leq \tilde{O}(n\epsilon^2 k)$

Via strong composition theorem from [DRV'10]


# Outline

- Post Selection Hypothesis Testing [R,Roth,Smith,Thakkar FOCs'16]
  - Connection between Max-Info and Differential Privacy
- DP composition with adaptively selected parameters [R,Roth,Ullman,Vadhan NIPS'16]
  - Privacy Odometers and Filters
- Private Hypothesis Tests [Gaboardi,Lim,R,Vadhan ICML'16],[Kifer,R AISTATS'17]
  - Chi-Square Tests
- Directions for Future Work

# DP Composition



Data:  $x$



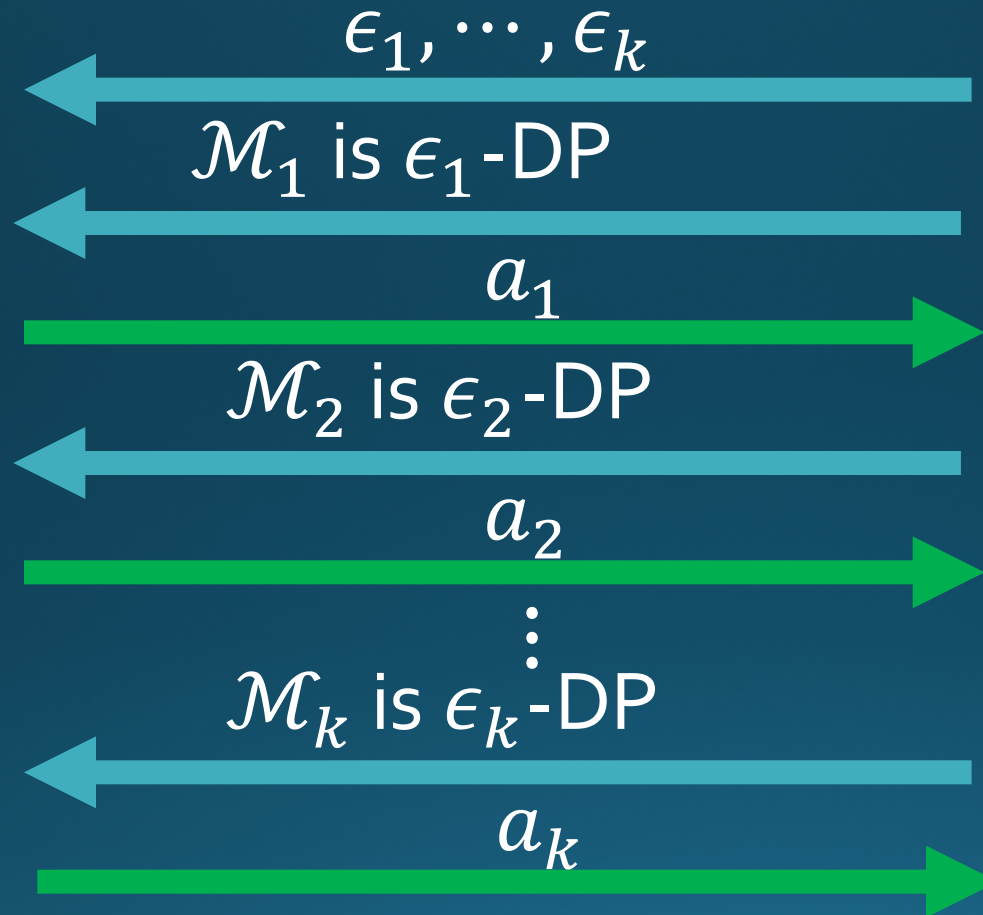
$$f(\epsilon_1, \dots, \epsilon_k) \leq \epsilon_g$$

$$\Rightarrow \mathcal{M}_k \circ \dots \circ \mathcal{M}_1 \text{ is } \epsilon_g\text{-DP}$$

# DP Composition



Data:  $x$



$$f(\epsilon_1, \dots, \epsilon_k) \leq \epsilon_g$$

$$\Rightarrow \mathcal{M}_k \circ \dots \circ \mathcal{M}_1 \text{ is } \epsilon_g\text{-DP}$$

# DP Composition - Adaptive Privacy Parameters

- **Our focus:** Allow the analyst to allocate his privacy budget adaptively – also adaptively select the number of analyses.
  - Natural to allow the analyst to select parameters AND analyses adaptively – different DP analyses have different utility vs. privacy tradeoffs.
- **Questions:**
  - Which composition theorems still apply when we can select the parameters adaptively?
  - How can we even define differential privacy in this adaptively parameter setting?



# Privacy Loss Random Variable

- Privacy loss for neighboring  $x, x'$  and for an algorithm  $\mathcal{M}: \mathcal{X}^n \rightarrow \mathcal{O}$ :

$$L(o) = \log \left( \frac{\mathbb{P}(\mathcal{M}(x)=o)}{\mathbb{P}(\mathcal{M}(x')=o)} \right) \text{ where } o \sim \mathcal{M}(x)$$

- Each round  $i = 1, 2, \dots, k$  the analyst selects  $\epsilon_i \geq 0$  and an  $\epsilon_i$ -DP algorithm  $\mathcal{M}_i$  based on previous outcomes in an adversarial way.

$$L(o_1, \dots, o_k) = \sum_{i=1}^k L_i(o_i) = \sum_{i=1}^k \log \left( \frac{\mathbb{P}(\mathcal{M}_i(x) = o_i \mid o_1, \dots, o_{i-1})}{\mathbb{P}(\mathcal{M}_i(x') = o_i \mid o_1, \dots, o_{i-1})} \right)$$

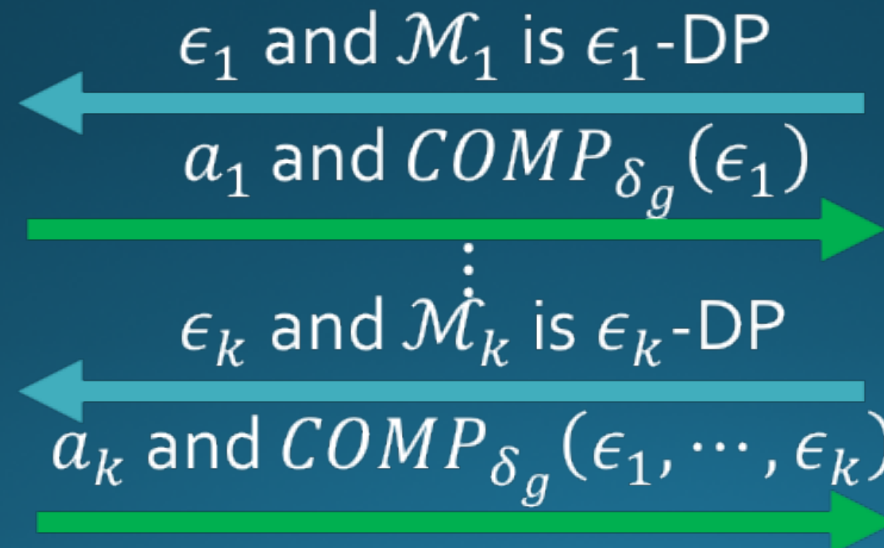
# Privacy Odometer [Roth, Ullman, Vadhan'16]

- Privacy odometer provides a running upper bound on privacy loss.
- A valid privacy odometer  $COMP_{\delta_g} : \mathbb{R}^* \rightarrow \mathbb{R}$  where an analyst selects  $\epsilon_1, \dots, \epsilon_k$  adaptively and w.p.  $\geq 1 - \delta_g$

$$|L(o_1, \dots, o_k)| \leq COMP_{\delta_g}(\epsilon_1, \dots, \epsilon_k)$$



Data:  $x$



Analyst

# Privacy Odometer Results [RRUV'16]

- Basic composition applies – for any  $\delta_g \geq 0$ , the following is a valid privacy odometer:

$$COMP_{\delta_g}(\epsilon_1, \dots, \epsilon_k) = \sum_{i=1}^k \epsilon_i$$

- For  $\sum_i^k \epsilon_i^2 \geq \frac{1}{n^2}$ , the following is a valid privacy odometer for  $\delta_g > 0$ :

$$COMP_{\delta_g}(\epsilon_1, \dots, \epsilon_k) = O\left(\sqrt{\sum_{i=1}^k \epsilon_i^2 \log\left(\frac{\log(n)}{\delta_g}\right)}\right)$$

- There is no valid privacy odometer with  $k > n$  and


$$COMP_{\delta_g}(\epsilon_1, \dots, \epsilon_k) = O\left(\sqrt{\sum_{i=1}^k \epsilon_i^2 \log\left(\frac{\log(n)}{\delta_g}\right)}\right)$$

# Proof Sketch

- Privacy Loss is a “biased” random walk with step size  $\pm\epsilon$ :

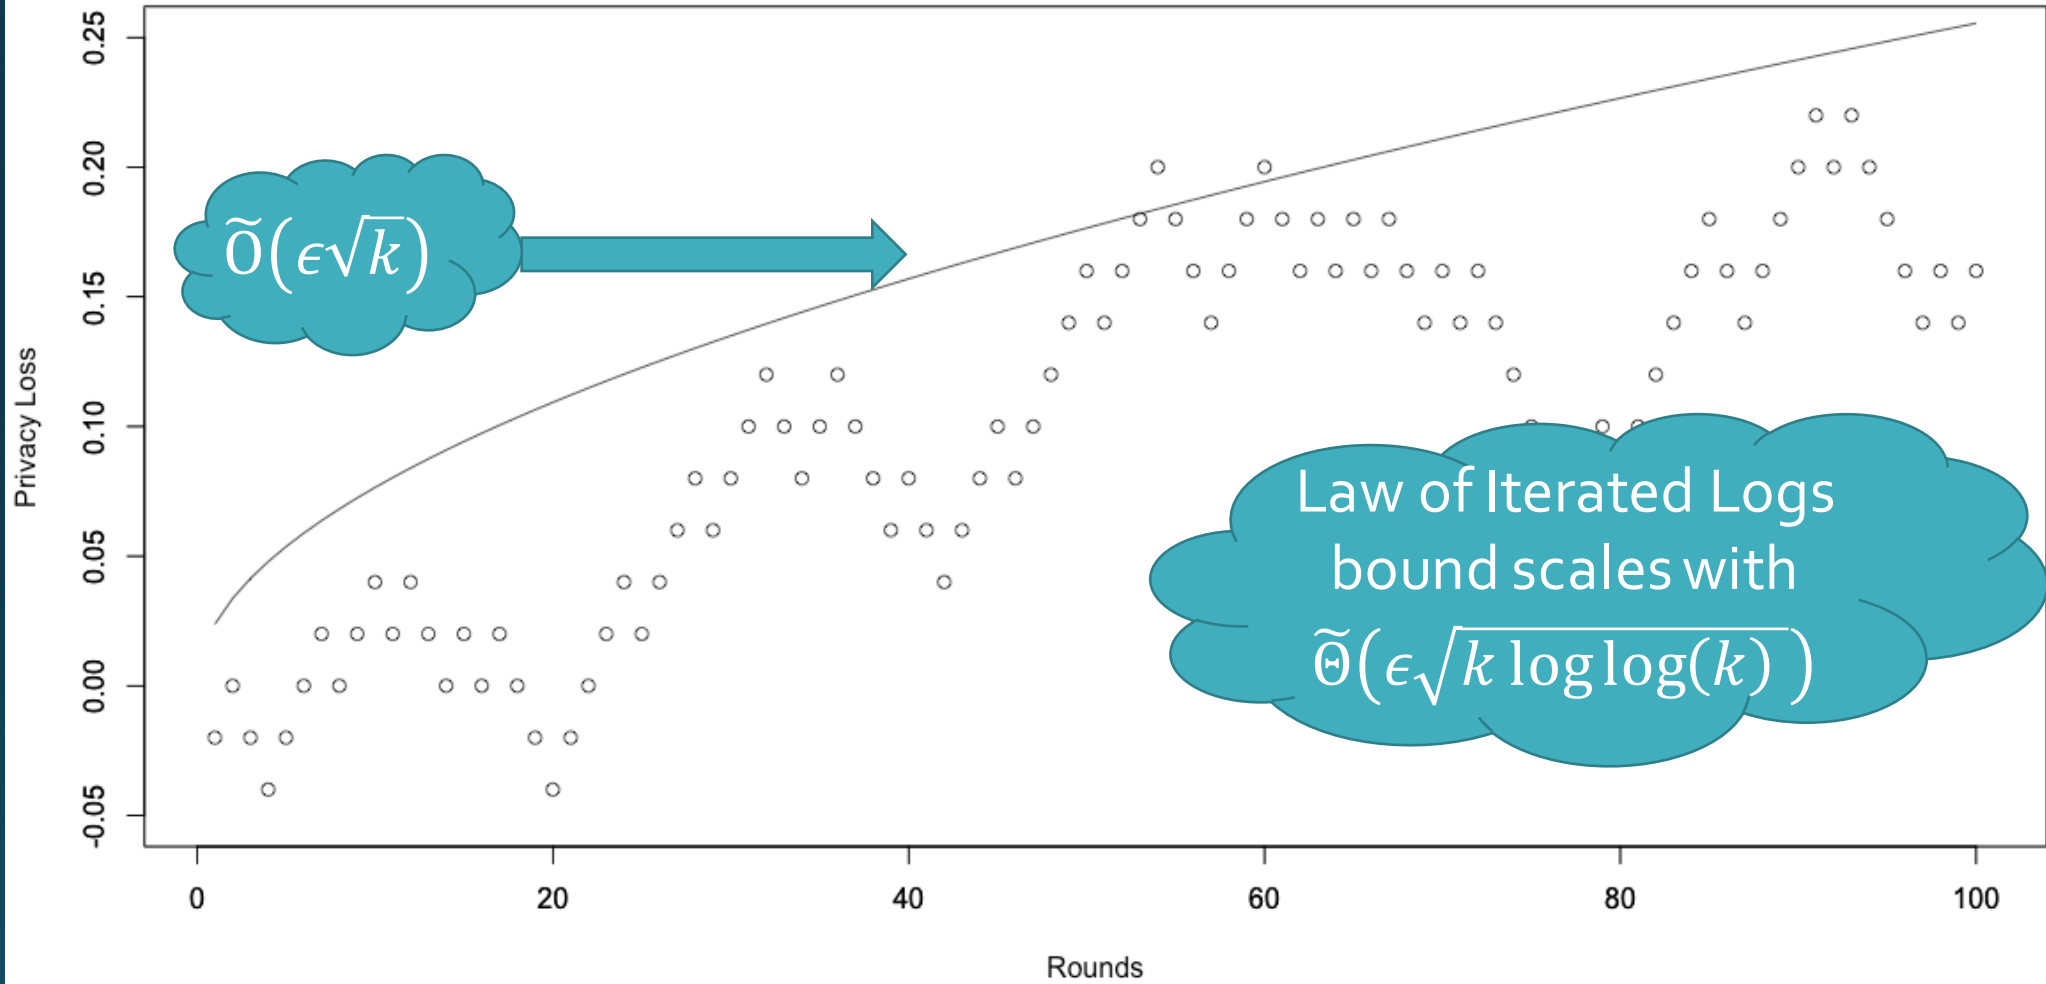
$$L(o_1, \dots, o_k) = \sum_{i=1}^k L_i(o_i) = \sum_{i=1}^k \log \left( \frac{\mathbb{P}(\mathcal{M}_i(x) = o_i \mid o_1, \dots, o_{i-1})}{\mathbb{P}(\mathcal{M}_i(x') = o_i \mid o_1, \dots, o_{i-1})} \right)$$

Bound on Privacy Loss





## Bound on Privacy Loss



# Outline

- Post Selection Hypothesis Testing [R,Roth,Smith,Thakkar FOCs'16]
  - Connection between Max-Info and Differential Privacy
- DP composition with adaptively selected parameters [R,Roth,Ullman,Vadhan NIPS'16]
  - Privacy Odometers and Filters
- Private Hypothesis Tests [Gaboardi,Lim,R,Vadhan ICML'16],[Kifer,R AISTATS'17]
  - Chi-Square Tests
- Directions for Future Work

# Differentially Private Hypothesis Tests

- DP analyses ensure statistical validity over adaptive sequences of analyses.
- Thus, we aim to develop analyses which are each DP and produce valid conclusions.
- **GOAL:** Valid hypothesis testing while preserving privacy.

# Classical Hypothesis Testing

- Want to design a test  $t: \mathcal{X}^n \rightarrow \{Inconclusive, Reject H_0\}$  s.t.:

	$H_0$ is True	$H_0$ is False
$t$ rejects	$\alpha =$ <b>False Discovery</b>	<b>Power</b>
$t$ fails to reject	$1-\alpha =$ <b>Significance</b>	<b>Type II Error</b>

- Want to ensure test has  $\mathbb{P}[\mathbf{False Discovery}] \leq \alpha$  and has good **Power**.

# Chi-Square Tests

- Categorical data  $X$ . Histogram:  $D = (D_1, \dots, D_d) \sim \text{Multinomial}(n, \vec{p})$ .
- 1. Goodness of Fit:**  $H_0: \vec{p} = \vec{p}^0$ 
  - Simple Test - data distribution completely determined
- 2. Independence Test:**  $H_0: Y^{(1)} \perp Y^{(2)}$ 
  - Composite Test – data distribution not completely determined
- Both classical tests use the Chi-Squared Statistic:

$$Q^2 = \sum \frac{(\text{Observed}_i - \text{Expected}_i)^2}{\text{Expected}_i}$$



# Private Goodness of Fit

- Add noise to each cell count to preserve differential privacy.
- Form the **private chi-squared statistic**:

$$Q_{DP}^2 = \sum_{i=1}^d \frac{(D_i + Z_i - np_i^0)^2}{np_i^0} \text{ where we use either:}$$

$$Z_i \sim N\left(\mathbf{0}, \mathbf{O}\left(\frac{\log\left(\frac{1}{\delta}\right)}{\epsilon^2}\right)\right) \text{ for } (\epsilon, \delta) - DP$$

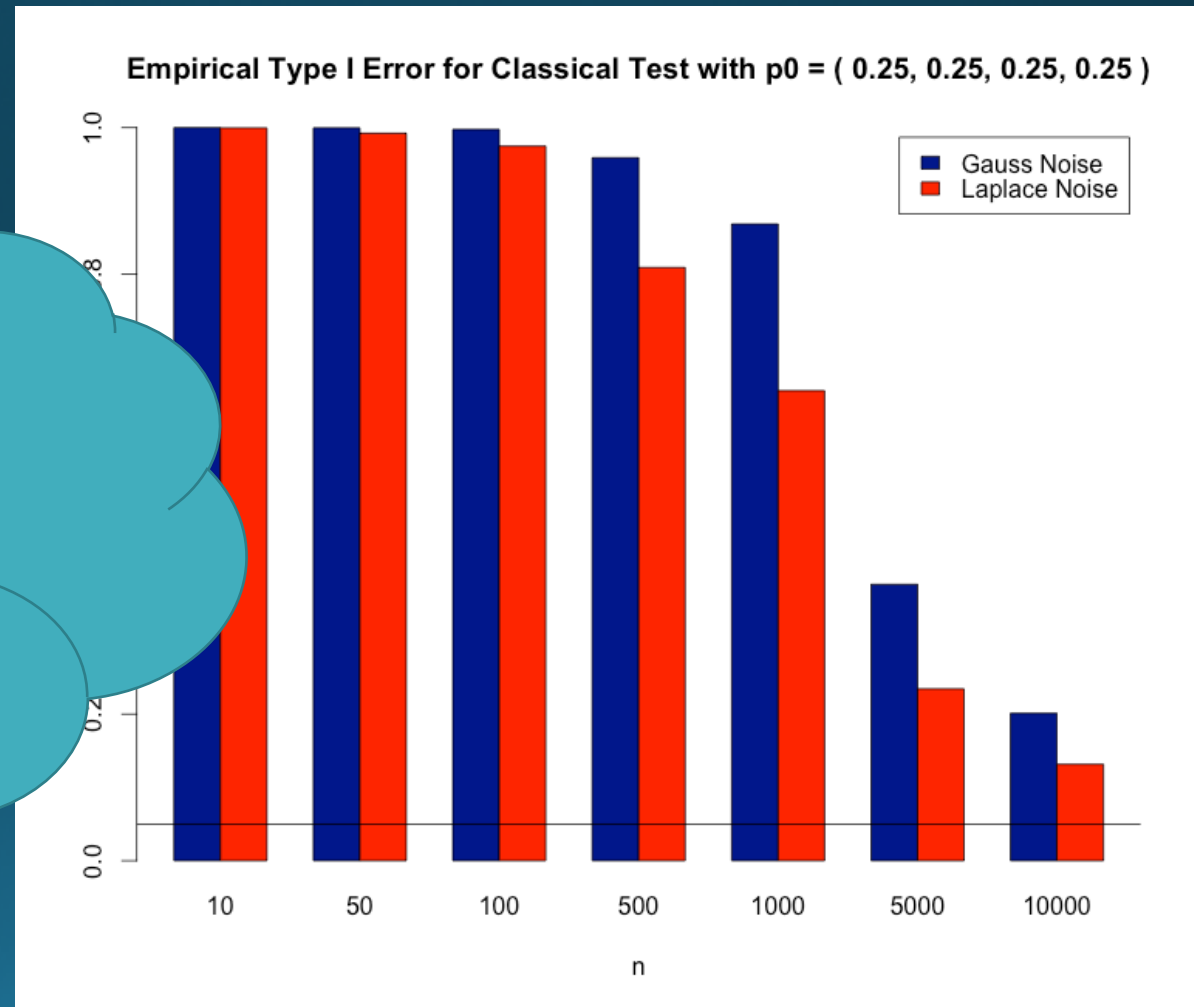
$$\text{or } Z_i \sim \text{Lap}\left(\mathbf{0}, \left(\frac{1}{\epsilon}\right)\right) \text{ for } (\epsilon, 0) - DP$$

# Naïve Approach – Use Classical Test

Noise is small as  $n \rightarrow \infty$

[Johnson and Shmatikov '13],  
[Vu and Slavkovic '09]

See Similar findings shown  
in [Fienberg, Rinaldo,  
Yang '10], [Karwa and  
Slavkovic '12, '16]

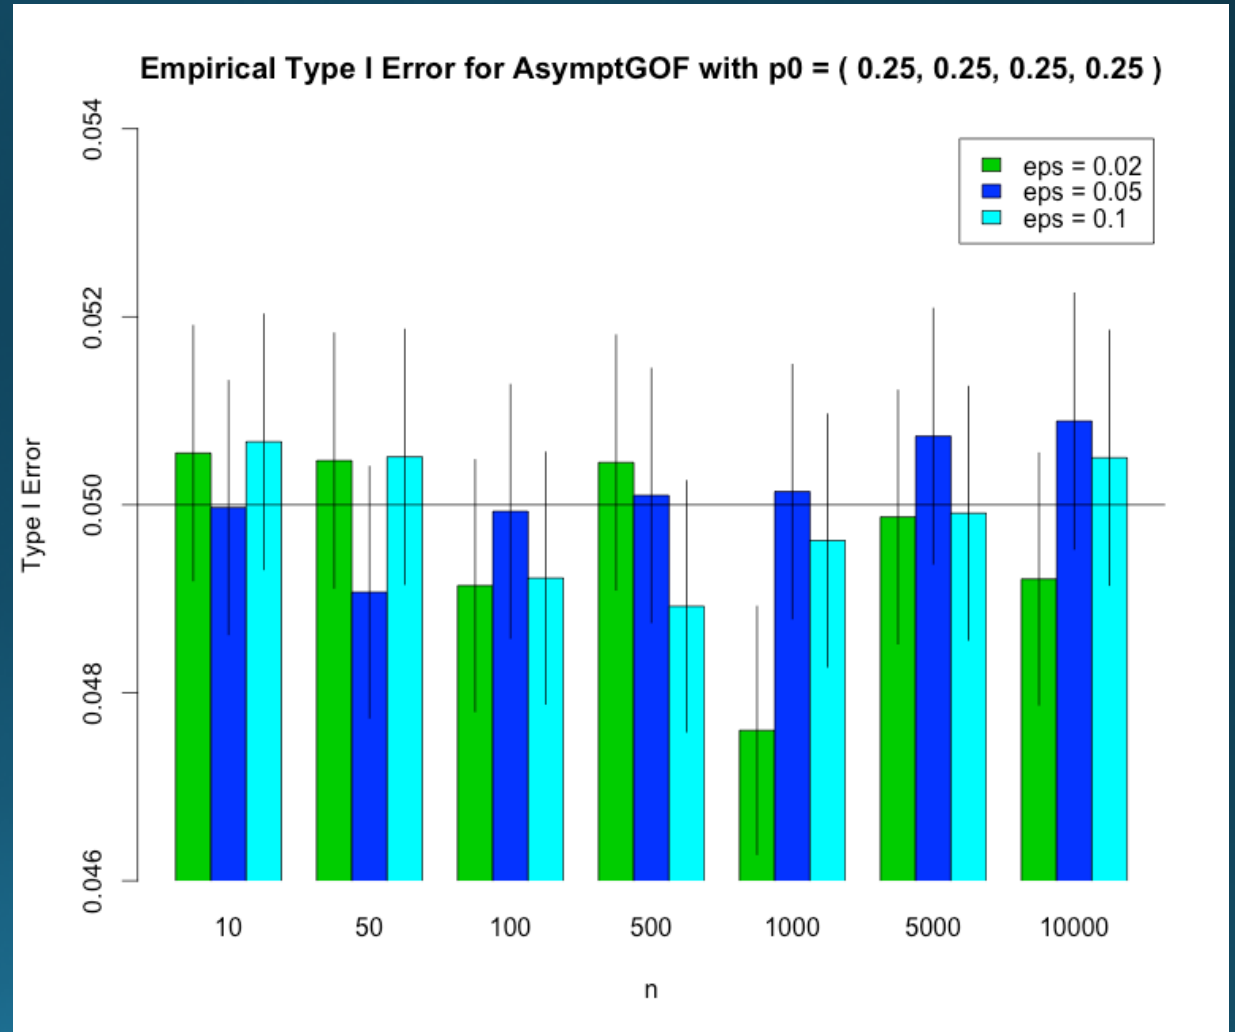


# AsymptGOF – Private GOF Test [GLRV'16]

Take the (Gaussian) noise into account

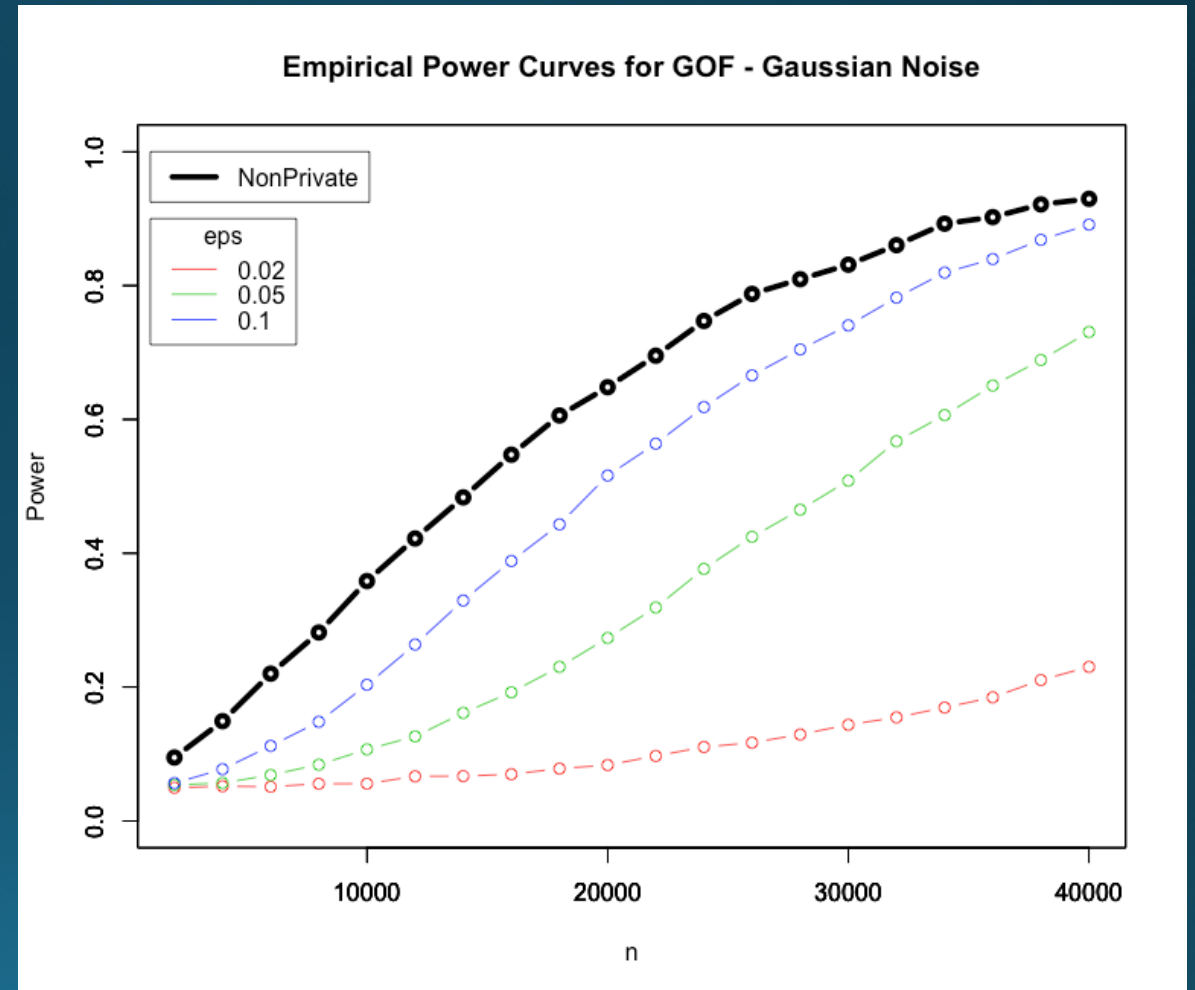
Set  $\delta = 10^{-6}$

Plotting the proportion of 100,000 trials that rejected  $H_0$ , despite it being true.



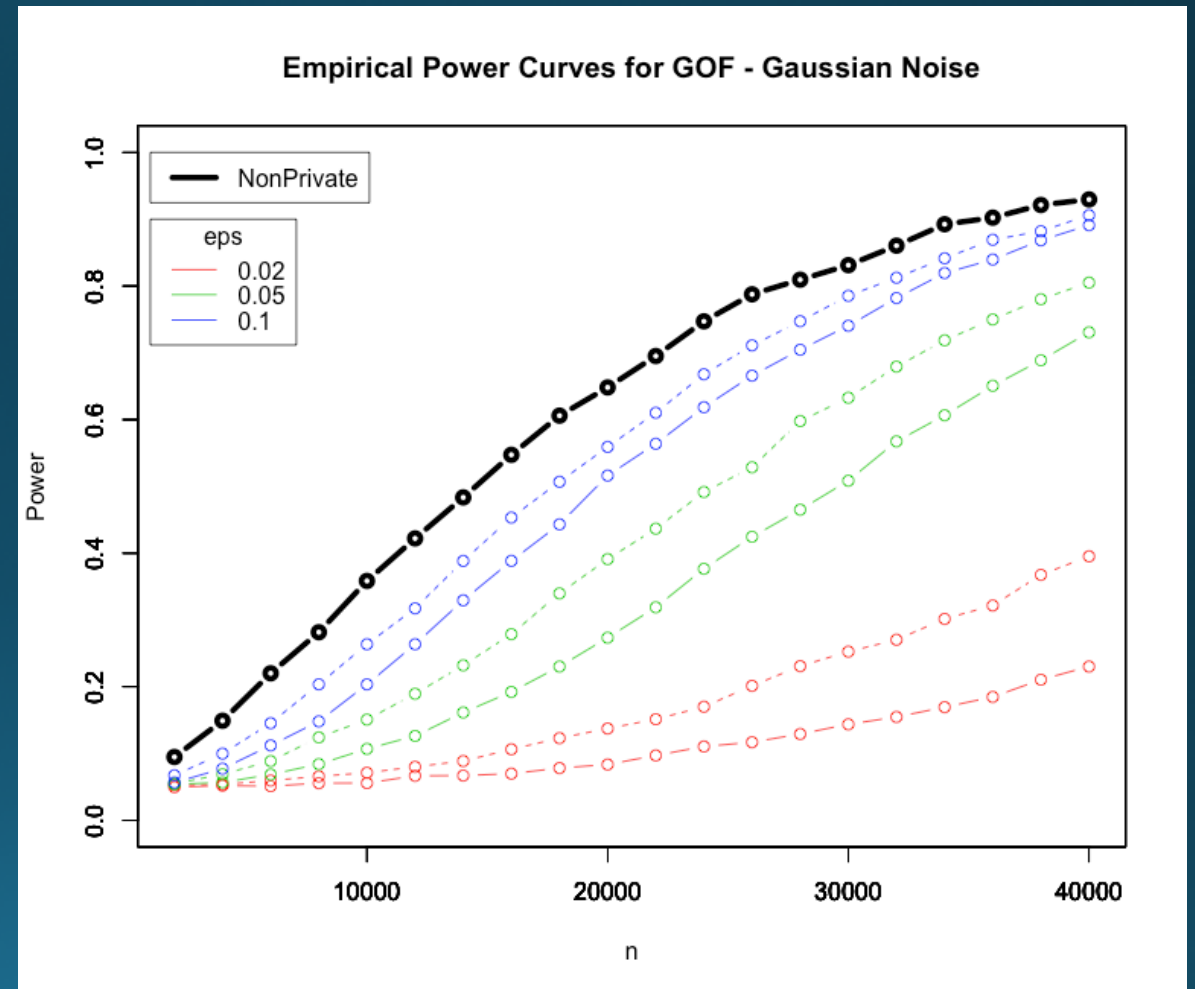
# AsymptGOF – Private GOF Test [GLRV'16]

- Set  $\delta = 10^{-6}$
- Test  $\vec{p}^0 = \left(\frac{1}{2}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}\right)$
- Generate data from  $\vec{p}^1 = \vec{p}^0 + 0.01 \left(1, -\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}\right)$
- Plot the proportion of 10,000 trials that correctly rejected null.



# AsymptGOF vs NewStatAsymptGOF [KR'17]

- Set  $\delta = 10^{-6}$
- Test  $\vec{p}^0 = \left(\frac{1}{2}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}\right)$
- Generate data from  $\vec{p}^1 = \vec{p}^0 + 0.01 \left(1, -\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}\right)$
- Plot the proportion of 10,000 trials that correctly rejected null.



# Outline

- Post Selection Hypothesis Testing [R,Roth,Smith,Thakkar-FOCS'16]
  - Connection between Max-Info and Differential Privacy
- DP composition with adaptively selected parameters [R,Roth,Ullman,Vadhan-NIPS'16]
  - Privacy Odometers and Filters
- Private Hypothesis Tests [Gaboardi,Lim,R,Vadhan-ICML'16],[Kifer,R-AISTATS'17]
  - Chi-Square Tests
- Directions for Future Work



# Directions for Future Work

- Adaptive Data Analysis:
  - We do not fully understand adaptive data analysis.
  - Is there a unifying measure in adaptive data analysis?
  - Is differential privacy the right approach?
- Develop private hypothesis tests that incorporate the noise:
  - Local model.
  - Other tests, e.g. ANOVA and Regression.

# Contributions

Thanks!

- **Adaptive Data Analysis:**

- [RRST- FOCS'16] – Information and privacy
- [RRUV- NIPS'16] – Adaptive parameter composition
- [GLRV- ICML'16], [KR- AISTATS'17] – Private hypothesis tests

- **Algorithmic Game Theory:**

- Incorporate privacy as a constraint:
  - [Kannan,Morgenstern,R,Roth- EC'15] - private allocations in kidney exchanges.
- Leverage stability of DP to solve new problems
  - [Kearns,Pai,R,Roth,Ullman'15a], [R,Roth,Ullman,Wu- EC'15] – Coordinate agents with incomplete information to desirable strategies
- [Hsu,Morgenstern,R,Roth,Vohra- STOC'16] - Prices as coordination devices
- [Jabbari,R,Roth,Wu- NIPS'16] - Revealed preferences
- [Dudik,Lahaie,R,Vaughan'17] – Prediction markets