Statistically Valid Inferences from Privacy Protected Data[1]

Gary King[2]

Institute for Quantitative Social Science
Harvard University

University of Chicago, 11/8/2019

[1] Joint work with Georgina Evans, Margaret Schwenzfeier, Abhradeep Thakurta.
[2] GaryKing.org/dp

Solving Political Problems Technologically

Differential Privacy & Inferential Validity

A General Purpose, Statistically Valid DP Algorithm

The Algorithm in Practice

Convincing Facebook to Make Data Available

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via “constitutional design”)

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?"

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica.

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)
- We iterate, and I propose a 2-part solution

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)
- We iterate, and I propose a 2-part solution
  - Outside academics: send proposals, no company veto

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)
- We iterate, and I propose a 2-part solution
  - Outside academics: send proposals, no company veto
  - Trusted 3rd party: Commission at Social Science One signs NDAs, agree not to publish from the data, chooses datasets, makes final decisions; can report publicly if Facebook reneges

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)
- We iterate, and I propose a 2-part solution
  - Outside academics: send proposals, no company veto
  - Trusted 3rd party: Commission at Social Science One signs NDAs, agree not to publish from the data, chooses datasets, makes final decisions; can report publicly if Facebook reneges
- Problem solved, without balancing $\rightsquigarrow$ agreements, announcements, funding, 30+ people assigned at Facebook

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)
- We iterate, and I propose a 2-part solution
  - Outside academics: send proposals, no company veto
  - Trusted 3rd party: Commission at Social Science One signs NDAs, agree not to publish from the data, chooses datasets, makes final decisions; can report publicly if Facebook reneges
- Problem solved, without balancing $\rightsquigarrow$ agreements, announcements, funding, 30+ people assigned at Facebook
- Just one issue:

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)
- We iterate, and I propose a 2-part solution
  - Outside academics: send proposals, no company veto
  - Trusted 3rd party: Commission at Social Science One signs NDAs, agree not to publish from the data, chooses datasets, makes final decisions; can report publicly if Facebook reneges
- Problem solved, without balancing $\rightsquigarrow$ agreements, announcements, funding, 30+ people assigned at Facebook
- Just one issue: Facebook's implementation plan was illegal!

Convincing Facebook to Make Data Available

Solving a Political Problem Technologically (via "constitutional design")

- Gary visits Facebook to persuade them to make data available
- In my hotel room packing, email arrives: "Hey what do we do about this?" This was Cambridge Analytica. (The worst timed lobby effort in history! Time to go home.)
- 3 days later: "Could you do a study of the 2016 election?"
- I'd love to, but I need 2 things & you'll only give me 1:
  - Complete access to data, people, etc. (like employees)
  - No pre-publication approval (like NO employees ever)
- We iterate, and I propose a 2-part solution
  - Outside academics: send proposals, no company veto
  - Trusted 3rd party: Commission at Social Science One signs NDAs, agree not to publish from the data, chooses datasets, makes final decisions; can report publicly if Facebook reneges
- Problem solved, without balancing $\rightsquigarrow$ agreements, announcements, funding, 30+ people assigned at Facebook
- Just one issue: Facebook's implementation plan was illegal!
- New Problem: Sharing data without it leaving Facebook

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation,

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing,

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms,

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements,

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing,

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models,

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)

- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data;

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer,

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal:

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy; & possible to discover population level patterns

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy; & possible to discover population level patterns
  - $\approx$ differential privacy

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy; & possible to discover population level patterns
  - $\approx$ differential privacy (seems to satisfy regulators et al.)

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy; & possible to discover population level patterns
  - $\approx$ differential privacy (seems to satisfy regulators et al.)
  - New Problem:

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy; & possible to discover population level patterns
  - $\approx$ differential privacy (seems to satisfy regulators et al.)
  - New Problem: Most DP algorithms are statistically invalid!

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy; & possible to discover population level patterns
  - $\approx$ differential privacy (seems to satisfy regulators et al.)
  - New Problem: Most DP algorithms are statistically invalid!
    - *unknown* statistical properties (usually *biased*)

Data Sharing Regime $\rightsquigarrow$ Data Access Regime

Solving Another Political Problem Technologically (via CS & Statistics)

- Data Sharing Regime: I give you data (maybe you sign DUA)
  - Venerable, but failing
  - Increasing public concern with privacy
  - Scholars discovered: de-identification doesn't work!
  - Nor does aggregation, query auditing, data clean rooms, legal agreements, restricted viewing, paired programmer models, etc.
  - Trusting researchers fails spectacularly at times (C.A.!)
  - Even trusting a researcher known to be trustworthy can fail
- Data Access Regime
  - Trusted server holds data; researchers as adversaries, can run any method $\rightsquigarrow$ noisy answer, a limited number of times
  - Goal: impossible to violate individual privacy; & possible to discover population level patterns
  - $\approx$ differential privacy (seems to satisfy regulators et al.)
  - New Problem: Most DP algorithms are statistically invalid!
    - *unknown* statistical properties (usually *biased*)
    - *no* uncertainty estimates

Solving Political Problems Technologically

Differential Privacy & Inferential Validity

A General Purpose, Statistically Valid DP Algorithm

The Algorithm in Practice

Theories of Inference: Statistics vs. CS

Theories of Inference: Statistics vs. CS

Population
⋮
Mary
Andrey
Georgie
Gary
Meg
Abhradeep
Ella
Anya
Greg
Max
Mean income: $48

Quantity of Interest

Theories of Inference: Statistics vs. CS

Population	Sample
⋮	✗
Mary	✓
Andrey	✓
Georgie	✓
Gary	✓
Meg	✓
Abhradeep	✓
Ella	✓
Anya	✓
Greg	✓
Max	✓
Mean income:	$48

Quantity of Interest

Theories of Inference: Statistics vs. CS

Population	Sample	$
⋮	✗	
Mary	✓	76
Andrey	✓	96
Georgie	✓	145
Gary	✓	122
Meg	✓	86
Abhradeep	✓	127
Ella	✓	72
Anya	✓	132
Greg	✓	95
Max	✓	134
Mean income:	$48	$108

Quantity of Interest

Usually no direct relevance

Theories of Inference: Statistics vs. CS

Population	Sample	$
⋮	✗	
Mary	✓	76
Andrey	✓	96
Georgie	✓	145
Gary	✓	122
Meg	✓	86
Abhradeep	✓	127
Ella	✓	72
Anya	✓	132
Greg	✓	95
Max	✓	134
Mean income:	$48	$108

Quantity of Interest

Usually no direct relevance

Theories of Inference: Statistics vs. CS

Population	Sample	$	+Privacy
⋮	✗		
Mary	✓	76	
Andrey	✓	96	
Georgie	✓	145	Noise & Censoring
Gary	✓	122	
Meg	✓	86	
Abhradeep	✓	127	
Ella	✓	72	
Anya	✓	132	
Greg	✓	95	
Max	✓	134	
Mean income:	$48	$108	

Quantity of Interest

Usually no direct relevance

Theories of Inference: Statistics vs. CS

Population	Sample	$	+Privacy	=dp$
⋮	✗			
Mary	✓	76		85
Andrey	✓	96		103
Georgie	✓	145		75
Gary	✓	122	Noise & Censoring	113
Meg	✓	86		125
Abhradeep	✓	127		97
Ella	✓	72		101
Anya	✓	132		128
Greg	✓	95		83
Max	✓	134		201
Mean income:	$48	$108		$111

Quantity of Interest

Usually no direct relevance

No direct relevance

Theories of Inference: Statistics vs. CS

Population	Sample	$	+Privacy	=dp$
⋮	✗			
Mary	✓	76		85
Andrey	✓	96		103
Georgie	✓	145		75
Gary	✓	122	Noise & Censoring	113
Meg	✓	86		125
Abhradeep	✓	127		97
Ella	✓	72		101
Anya	✓	132		128
Greg	✓	95		83
Max	✓	134		201
Mean income:	$48	$108		$111

Differential Privacy and its Inferential Challenges

Differential Privacy and its Inferential Challenges

- Estimators

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy
    - Fundamental problems for statistical inference

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy
    - Fundamental problems for statistical inference
- The DP Standard

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy
    - Fundamental problems for statistical inference
- The DP Standard
  - Including (D) or excluding (D') you doesn't change conclusions

$$\frac{\Pr[M(s, D) = m]}{\Pr[M(s, D') = m]} \in (1 \pm \epsilon)$$

for all D, D', m

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy
    - Fundamental problems for statistical inference
- The DP Standard
  - Including (D) or excluding (D') you doesn't change conclusions

$$\frac{\Pr[M(s, D) = m]}{\Pr[M(s, D') = m]} \in (1 \pm \epsilon)$$

  for all D, D', m
- Examples all proven to protect the biggest possible outlier

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy
    - Fundamental problems for statistical inference
- The DP Standard
  - Including (D) or excluding (D') you doesn't change conclusions

$$\frac{\Pr[M(s, D) = m]}{\Pr[M(s, D') = m]} \in (1 \pm \epsilon)$$

  for all D, D', m
- Examples all proven to protect the biggest possible outlier
  - $M(\text{mean}, D) = \frac{1}{n} \sum_{i=1}^{n} c(y_i, \Lambda) + N\left(0, \frac{4\Lambda}{n\epsilon}\right)$

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy
    - Fundamental problems for statistical inference
- The DP Standard
  - Including (D) or excluding (D') you doesn't change conclusions

$$\frac{\Pr[M(s, D) = m]}{\Pr[M(s, D') = m]} \in (1 \pm \epsilon)$$

  for all D, D', m
- Examples all proven to protect the biggest possible outlier
  - $M(\text{mean}, D) = \frac{1}{n} \sum_{i=1}^{n} c(y_i, \Lambda) + N\left(0, \frac{4\Lambda}{n\epsilon}\right)$
  - Or: mess with gradients, $X_i' X_i$, data, QOIs, etc.

Differential Privacy and its Inferential Challenges

- Estimators
  - Classical Statistics: Apply statistic s to dataset D, $s(D)$
  - DP Mechanism: $M(s, D)$, with noise & censoring
    - Essential components of ensuring privacy
    - Fundamental problems for statistical inference
- The DP Standard
  - Including (D) or excluding (D') you doesn't change conclusions

$$\frac{\Pr[M(s, D) = m]}{\Pr[M(s, D') = m]} \in (1 \pm \epsilon)$$

    for all D, D', m
- Examples all proven to protect the biggest possible outlier
  - $M(\text{mean}, D) = \frac{1}{n} \sum_{i=1}^{n} c(y_i, \Lambda) + N\left(0, \frac{4\Lambda}{n\epsilon}\right)$
  - Or: mess with gradients, $X_i' X_i$, data, QOIs, etc.
- Statistical properties: usually biased, no uncertainty estimates

Properties of Differential Privacy

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Average privacy loss $\ll$ maximum privacy loss

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Average privacy loss $\ll$ maximum privacy loss
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Average privacy loss $\ll$ maximum privacy loss
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
- Risk for small groups (k) drops linearly, $k\epsilon$

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Average privacy loss $\ll$ maximum privacy loss
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
- Risk for small groups (k) drops linearly, $k\epsilon$
- Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Average privacy loss $\ll$ maximum privacy loss
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
- Risk for small groups (k) drops linearly, $k\epsilon$
- Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
- Privacy Budget

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Average privacy loss $\ll$ maximum privacy loss
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
- Risk for small groups (k) drops linearly, $k\epsilon$
- Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
- Privacy Budget
  - Can sum and limit risks across analyses & researchers

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Average privacy loss $\ll$ maximum privacy loss
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
- Risk for small groups (k) drops linearly, $k\epsilon$
- Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
- Privacy Budget
  - Can sum and limit risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run

DP: Completely Changes Statistical Best Practices

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems
  - Researcher biases

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.
  - Lower probability of serendipitous discovery

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.
  - Lower probability of serendipitous discovery
  - Higher probability of being fooled by data

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.
  - Lower probability of serendipitous discovery
  - Higher probability of being fooled by data
  - Must plan data analyses carefully!

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.
  - Lower probability of serendipitous discovery
  - Higher probability of being fooled by data
  - Must plan data analyses carefully!
- Risks

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.
  - Lower probability of serendipitous discovery
  - Higher probability of being fooled by data
  - Must plan data analyses carefully!
- Risks
  - No differential privacy: no data access or privacy at risk

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.
  - Lower probability of serendipitous discovery
  - Higher probability of being fooled by data
  - Must plan data analyses carefully!
- Risks
  - No differential privacy: no data access or privacy at risk
  - No inferential validity: incorrect scientific conclusions, medical & policy advice; society and individuals at risk

DP: Completely Changes Statistical Best Practices

- Normally we try to avoid being fooled by:
  - Data problems — by running every possible diagnostic, data exploration and visualization, and conducting numerous statistical checks
  - Researcher biases — avoiding p-hacking via preregistration or "multiple comparison" corrections
- With DP: tips the scales
  - p-hacking avoided almost automatically
  - Little opportunity to explore data, run diagnostics, etc.
  - Lower probability of serendipitous discovery
  - Higher probability of being fooled by data
  - Must plan data analyses carefully!
- Risks
  - No differential privacy: no data access or privacy at risk
  - No inferential validity: incorrect scientific conclusions, medical & policy advice; society and individuals at risk
  - $\rightsquigarrow$ We need both DP and inferential validity

Solving Political Problems Technologically

Differential Privacy & Inferential Validity

A General Purpose, Statistically Valid DP Algorithm

The Algorithm in Practice

A Differentially Private Estimator

A Differentially Private Estimator

Private data

A Differentially Private Estimator

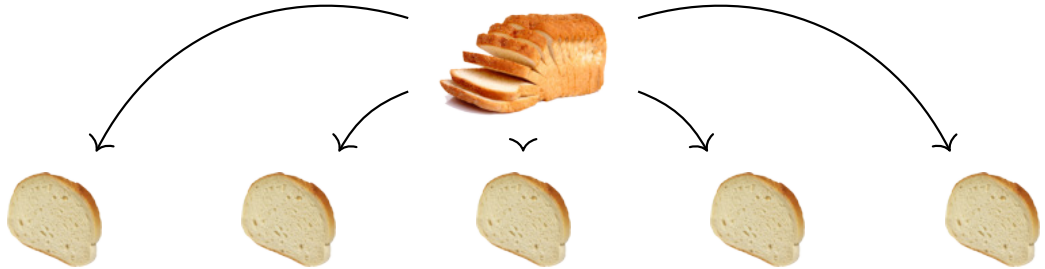

Private data

Partition

A Differentially Private Estimator

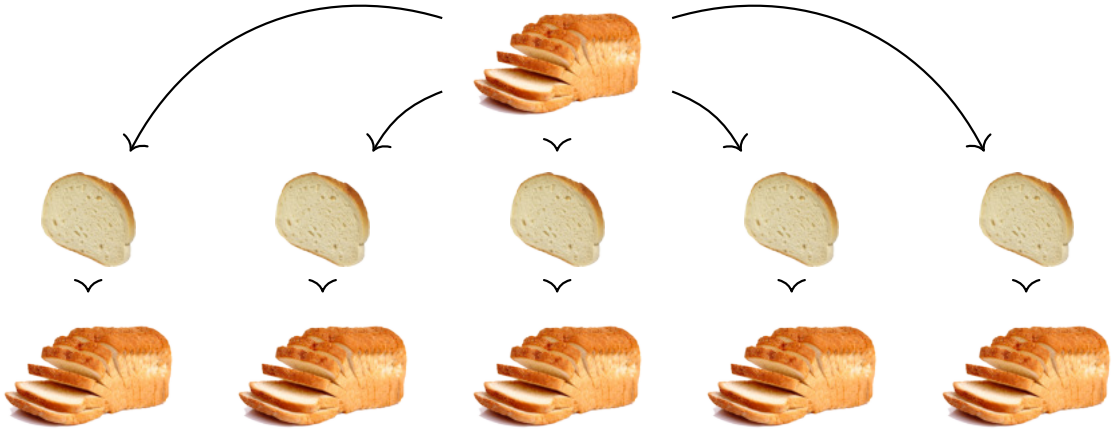

Private data

Partition

Bag of little bootstraps

A Differentially Private Estimator

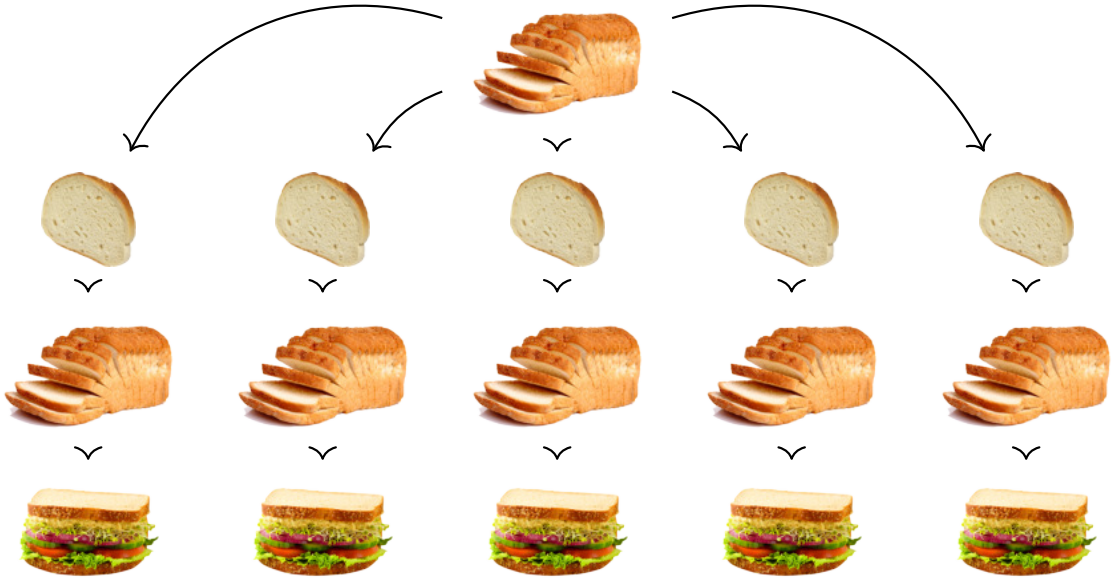

Private data

Partition

Bag of little bootstraps

Estimator

A Differentially Private Estimator

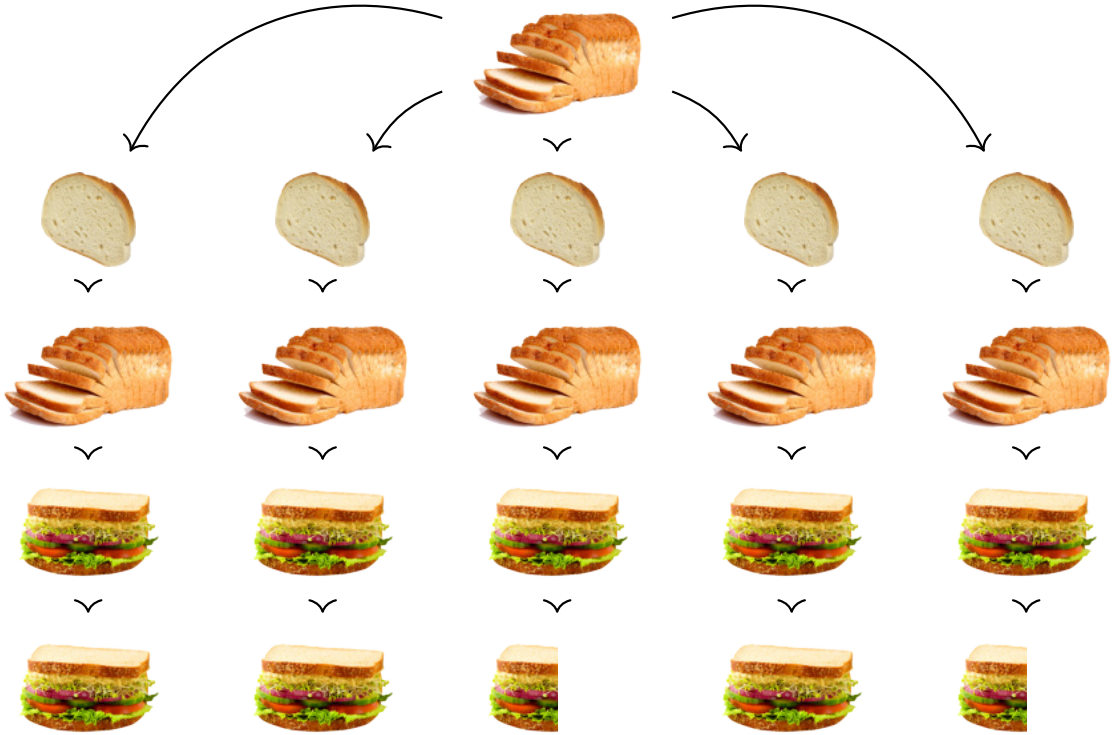

Private data

Partition

Bag of little bootstraps

Estimator

Censor

A Differentially Private Estimator

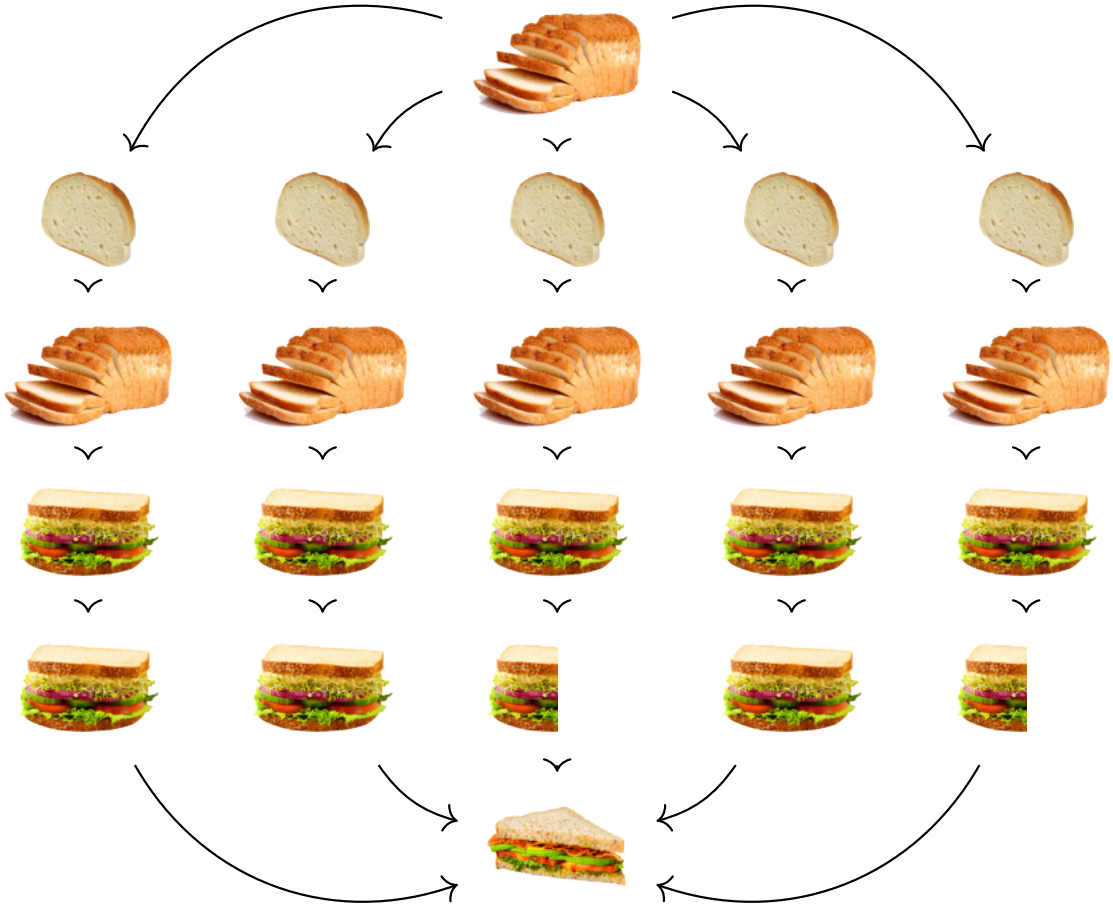

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

A Differentially Private Estimator

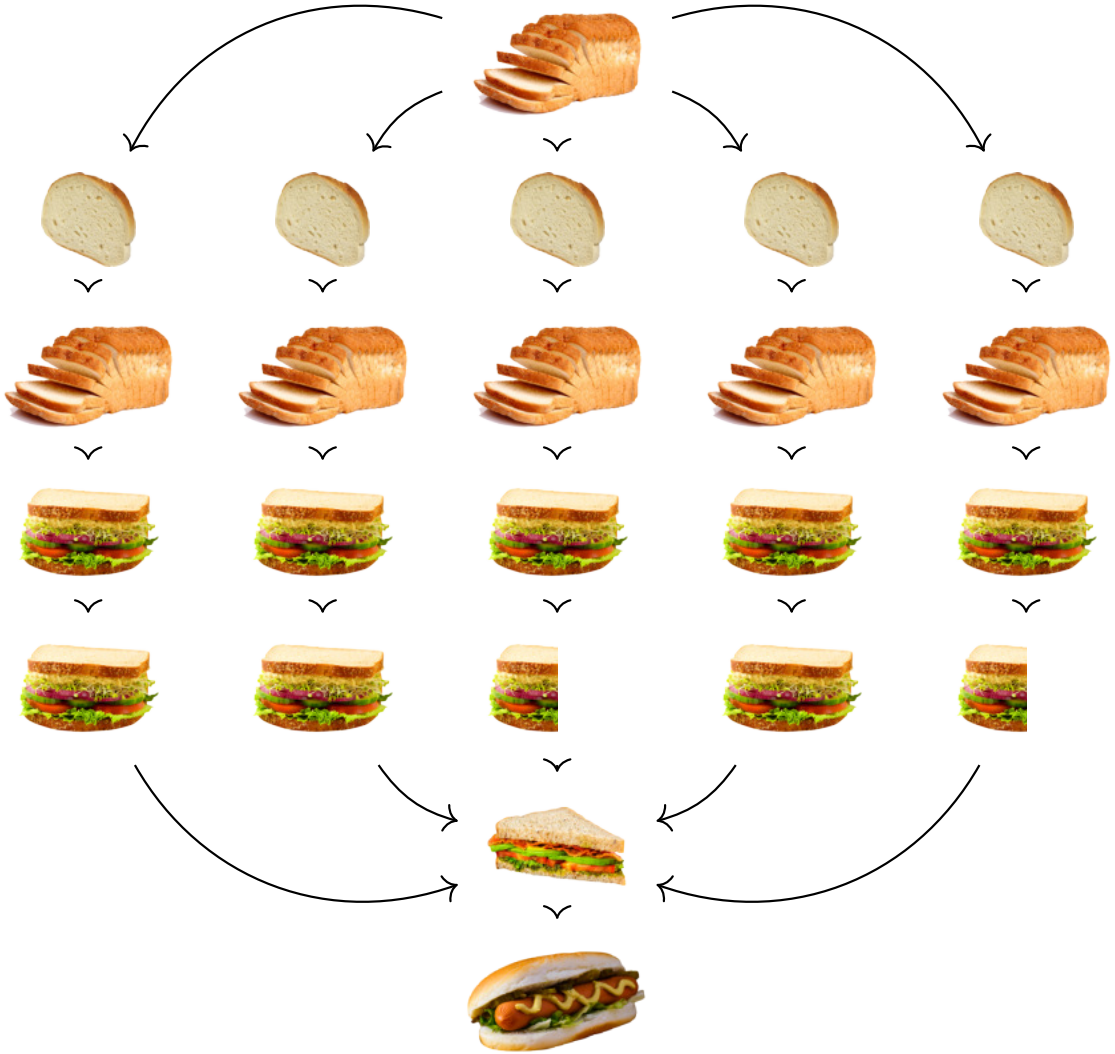

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

A Differentially Private Estimator

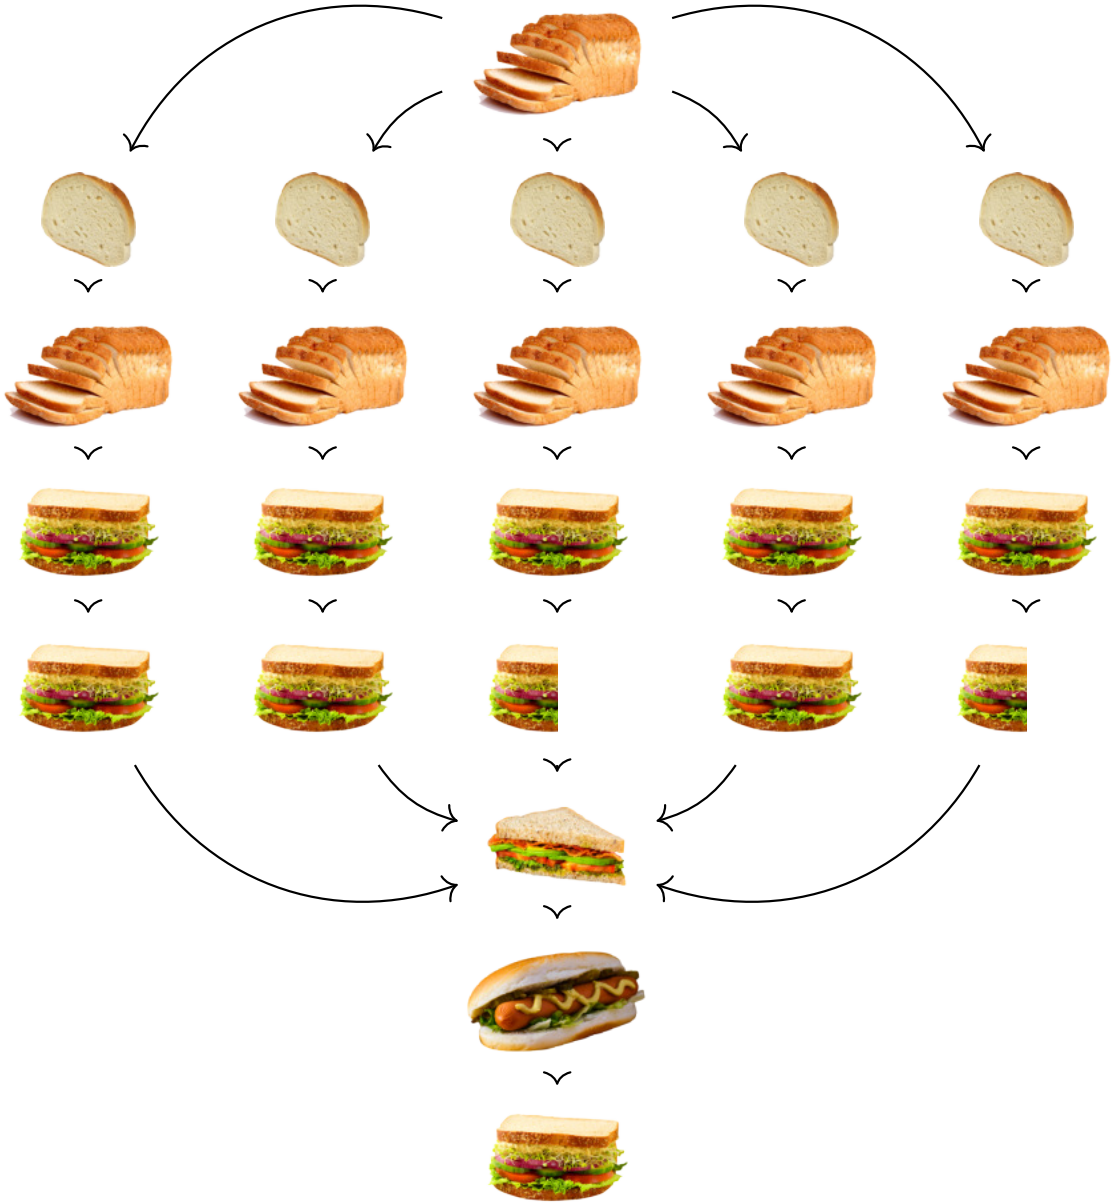

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction

A Differentially Private Estimator

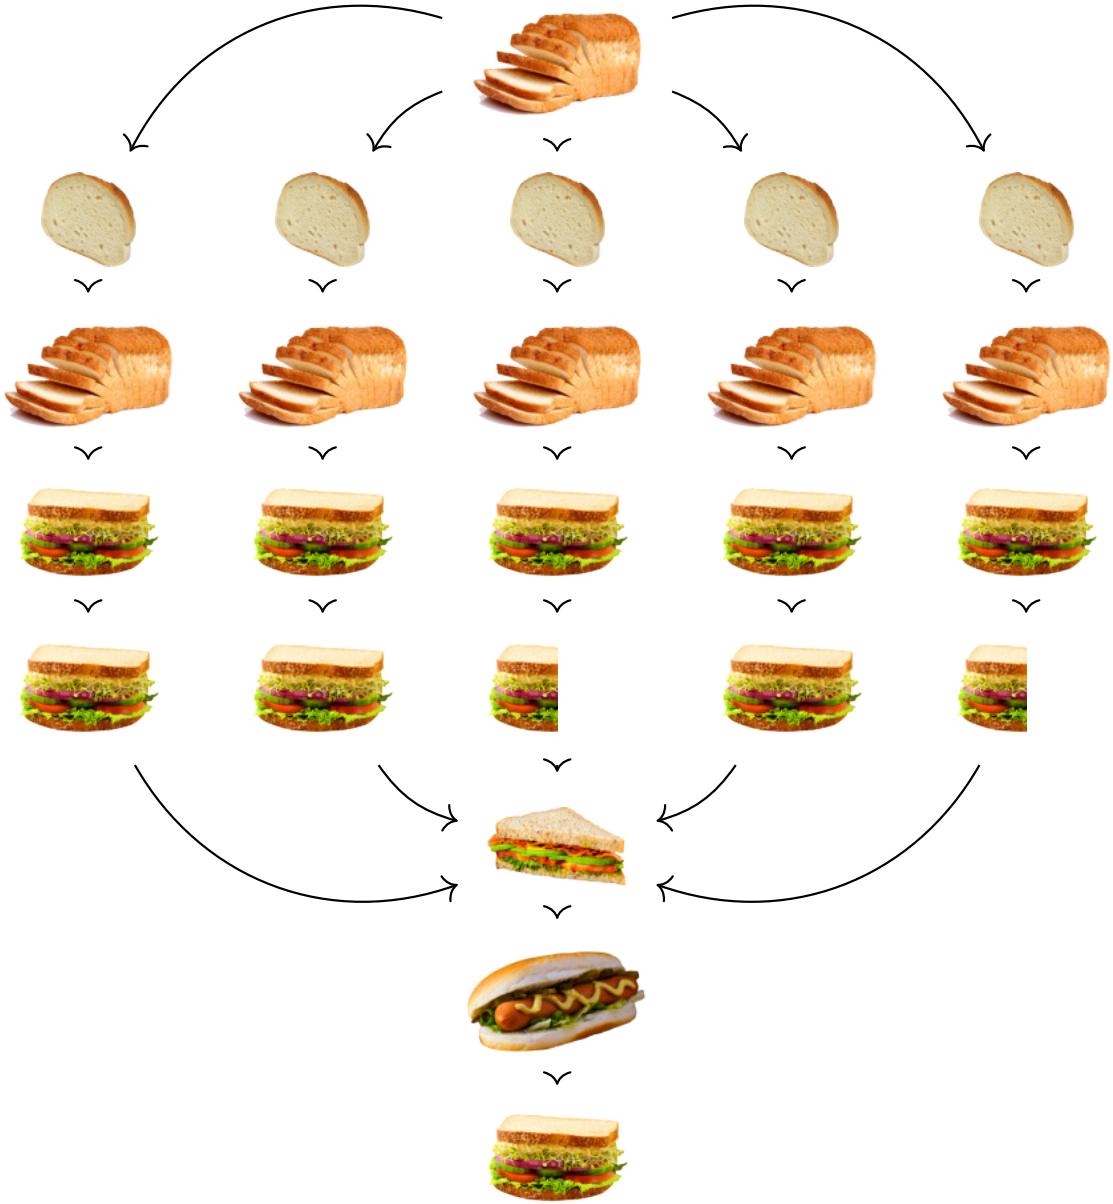

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

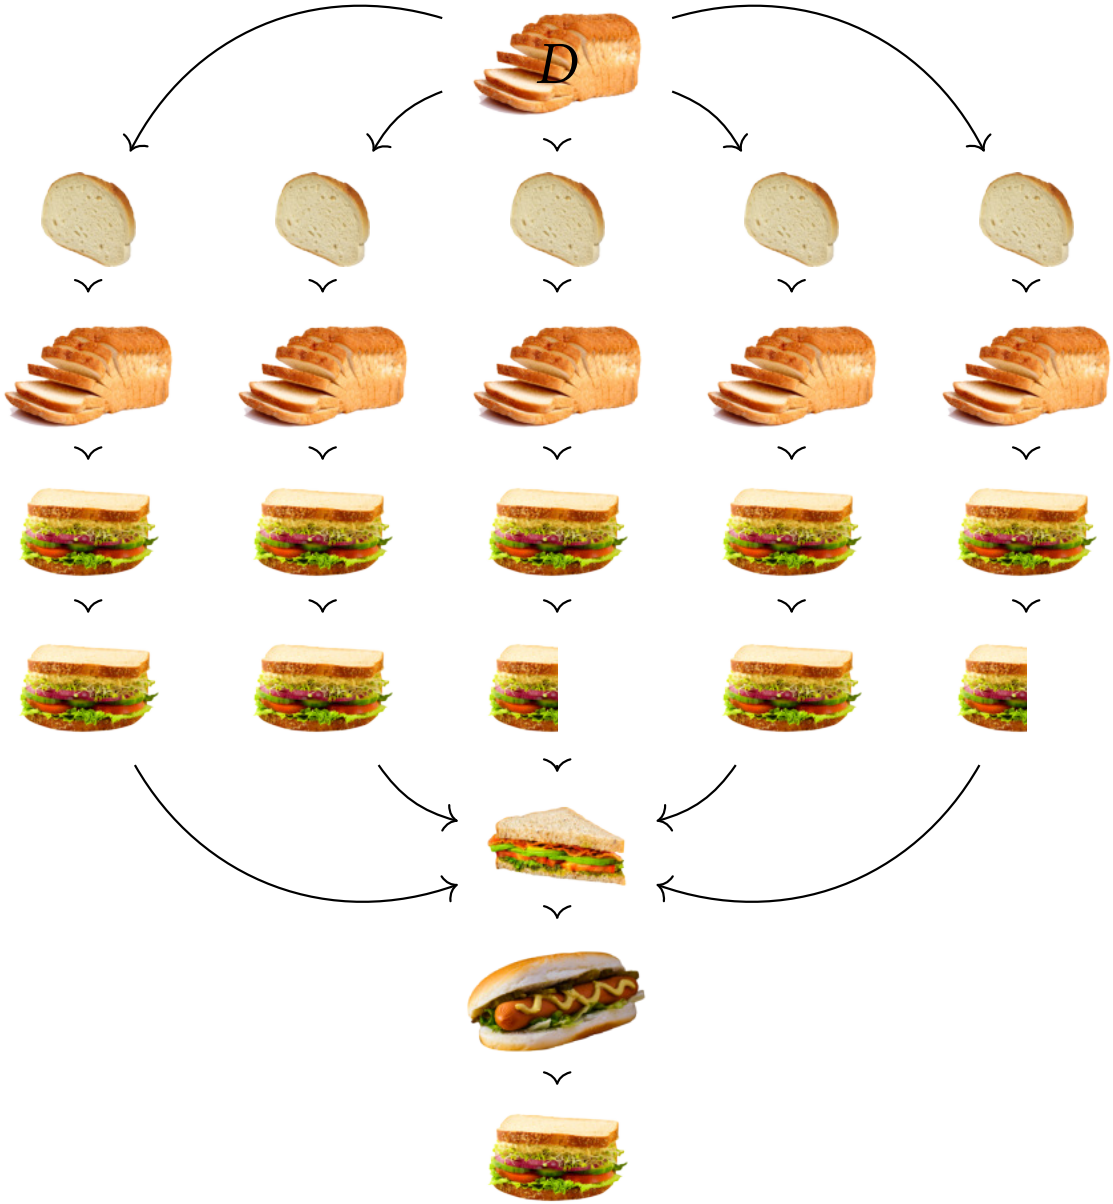

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

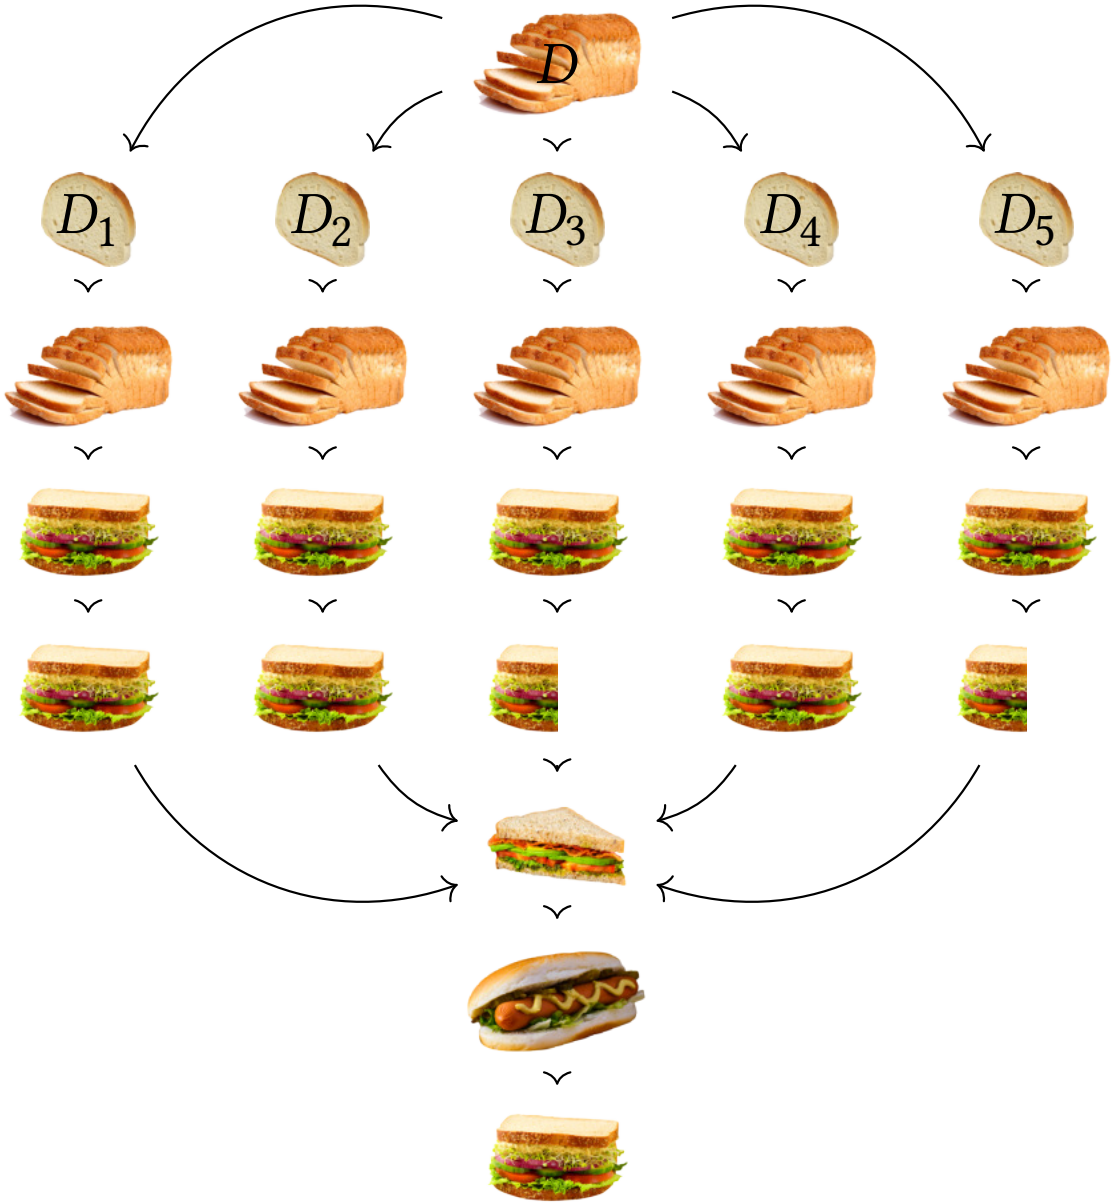

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

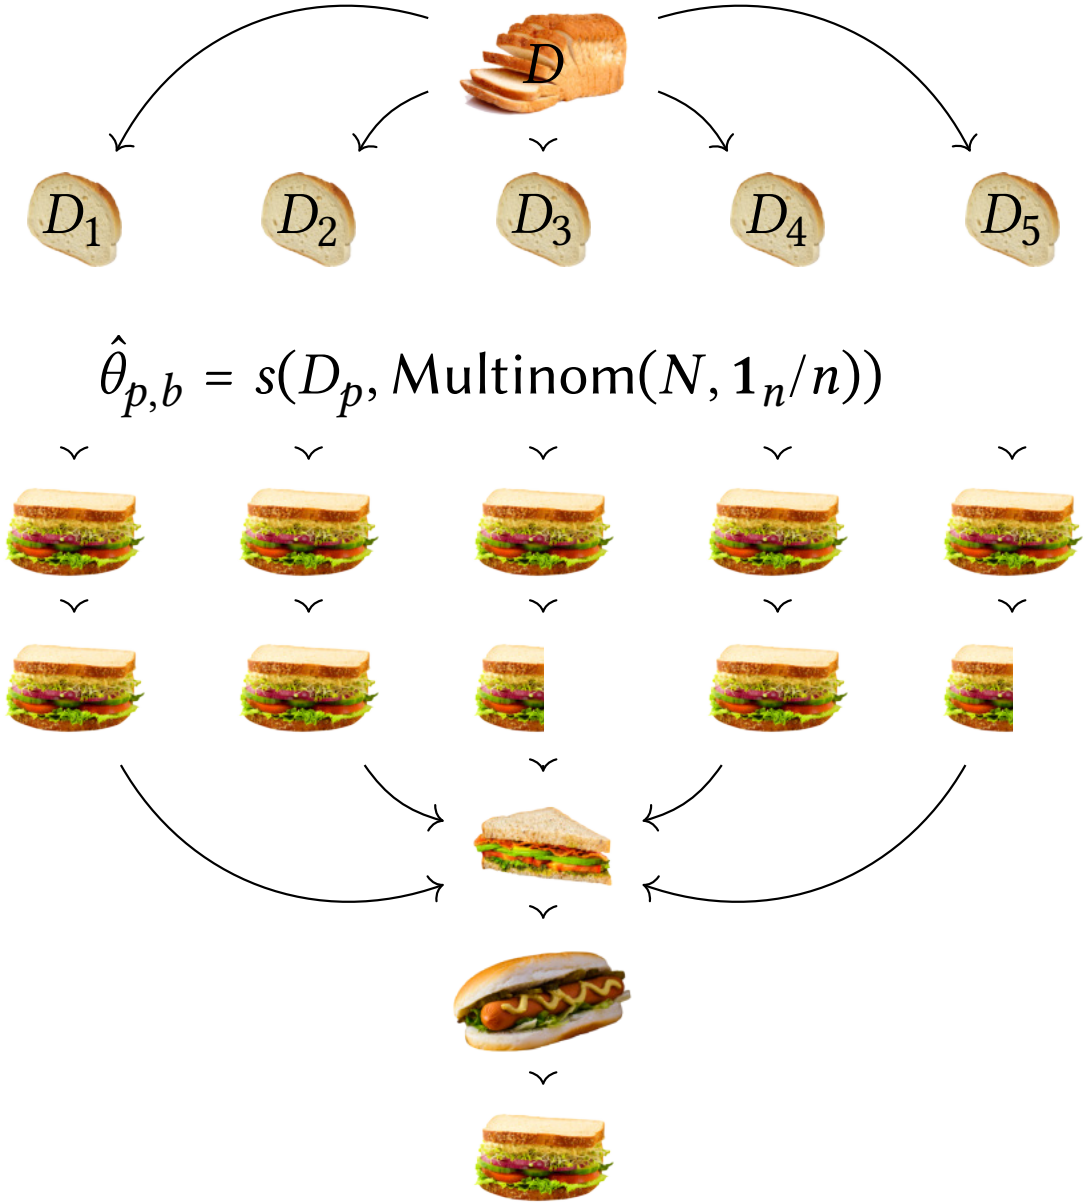

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

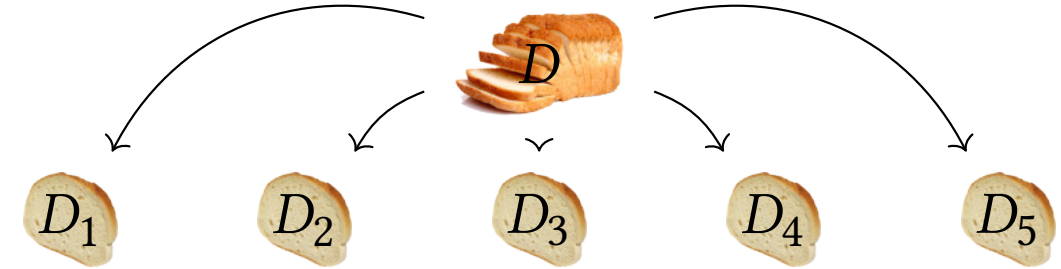

Private data

Partition

$$\hat{\theta}_{p,b} = s(D_p, \text{Multinom}(N, \mathbf{1}_n/n))$$

Bag of little bootstraps

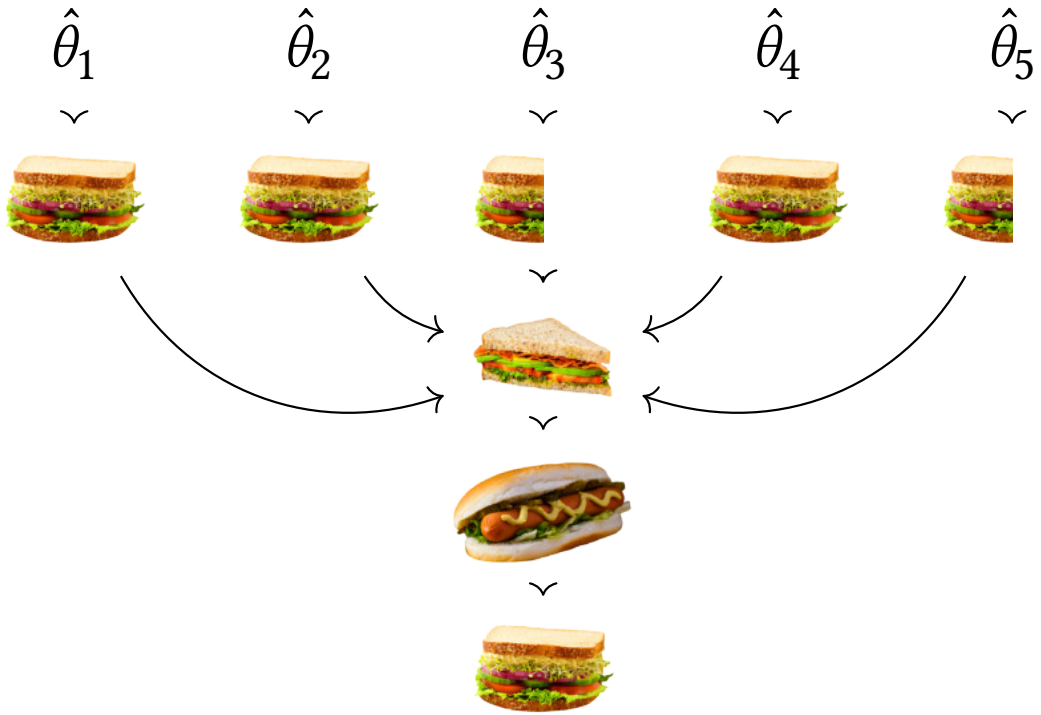

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

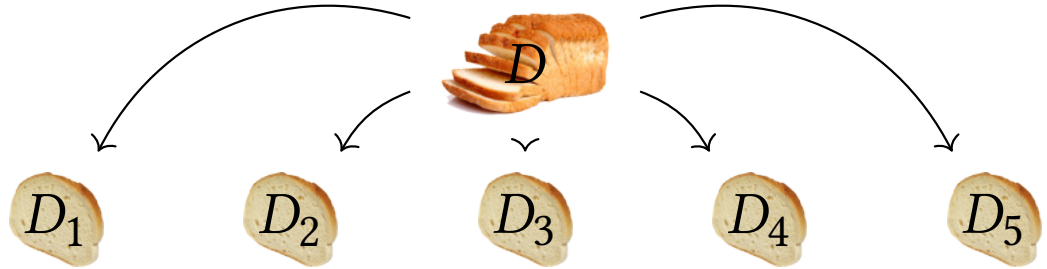

Private data

Partition

$$\hat{\theta}_{p,b} = s(D_p, \text{Multinom}(N, \mathbf{1}_n/n))$$

Bag of little bootstraps

$\hat{\theta}_1$ $\hat{\theta}_2$ $\hat{\theta}_3$ $\hat{\theta}_4$ $\hat{\theta}_5$

Estimator

$$\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$$

Censor

Average

Noise

Bias Correction
(& variance estimation)

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

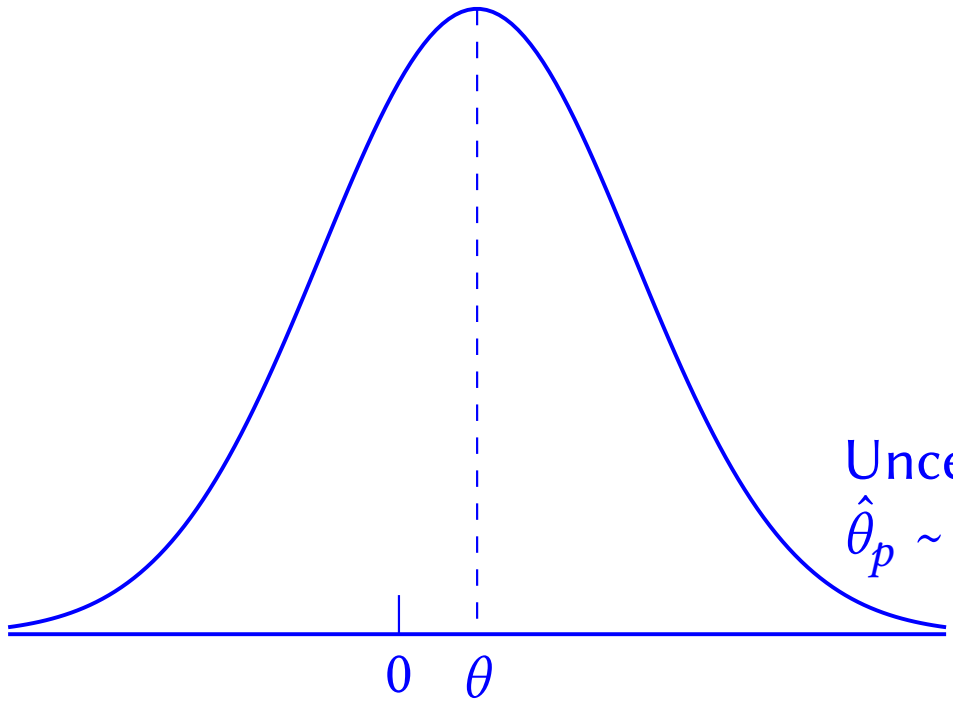

Uncensored distribution

$\hat{\theta}_p \sim N(\theta, \sigma^2)$

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

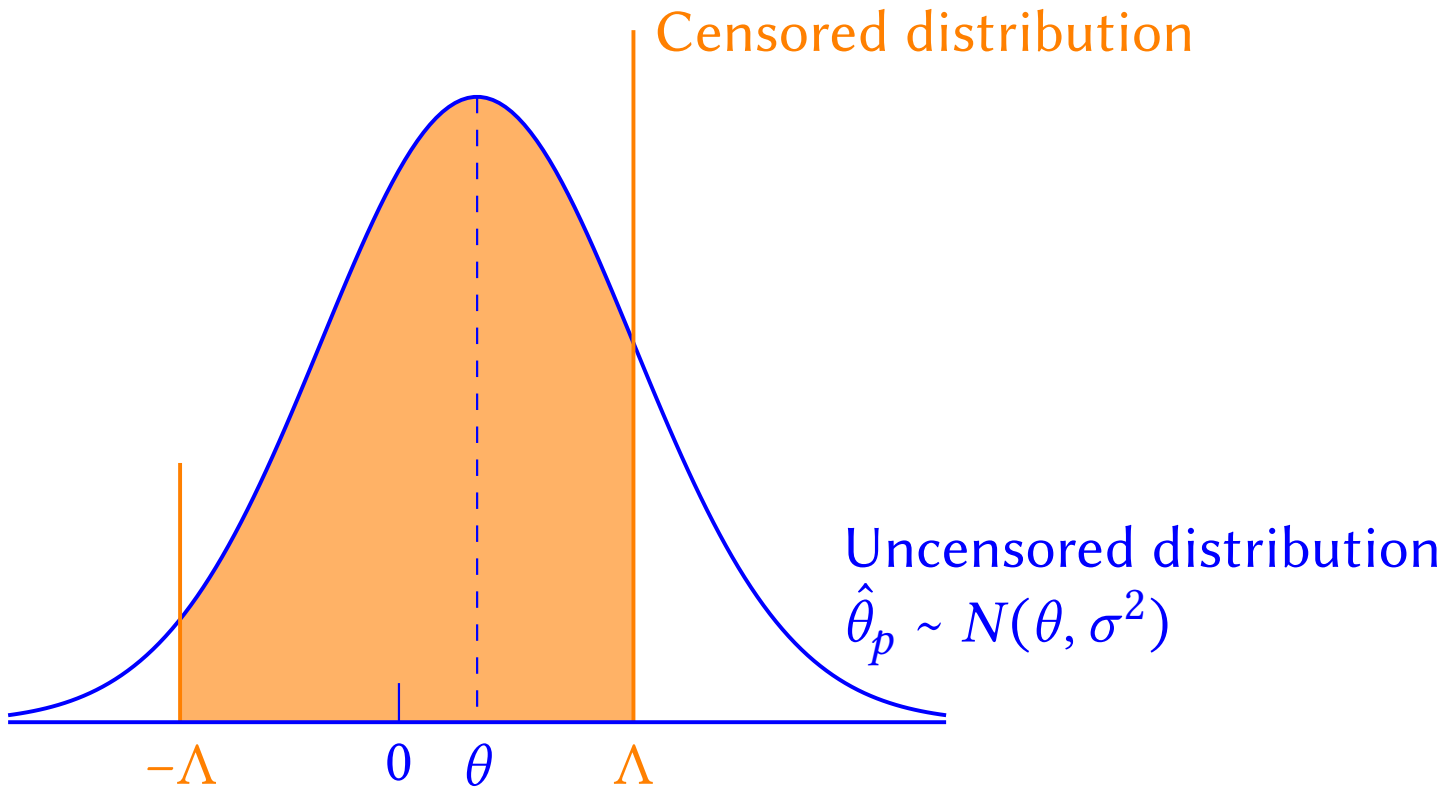

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

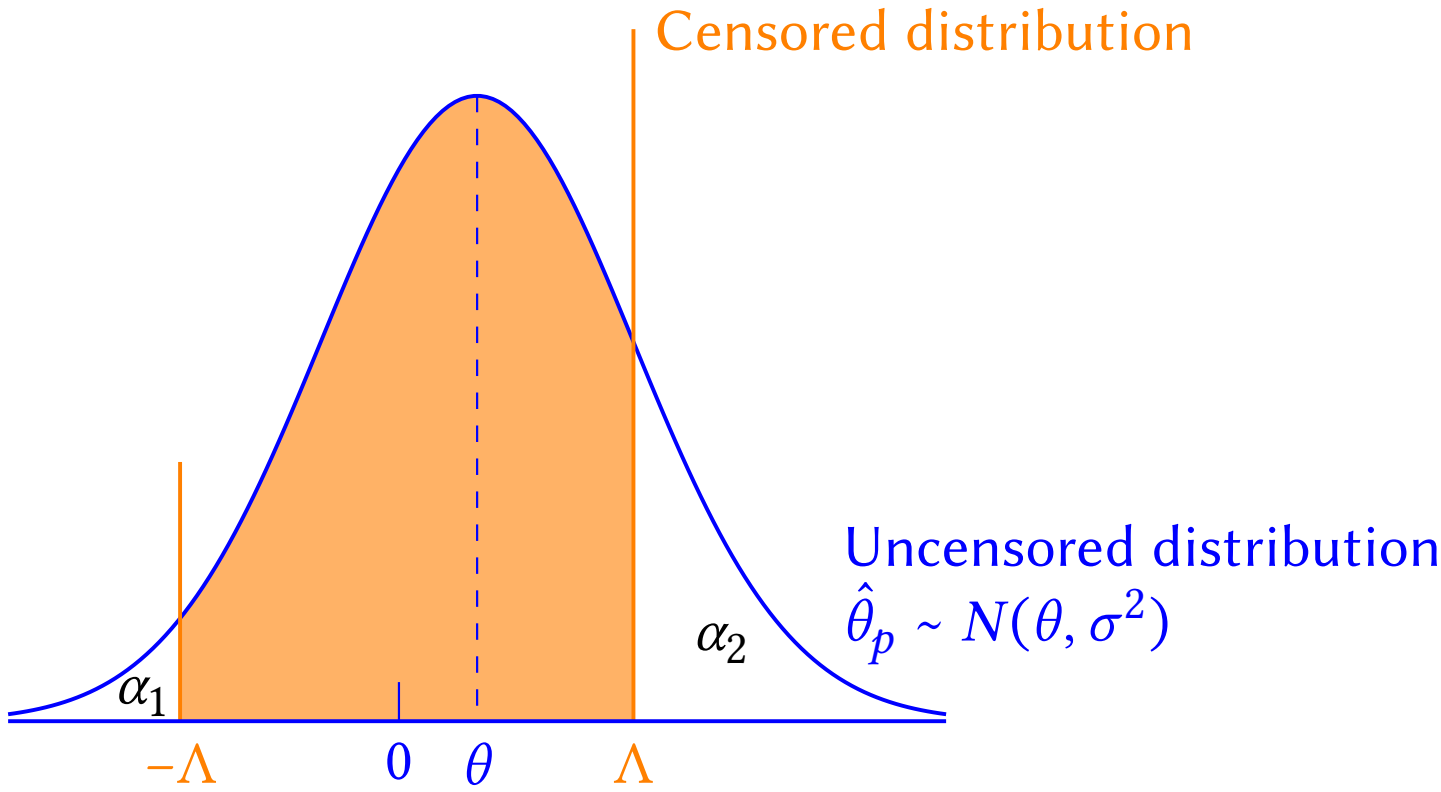

$$\int_{-\infty}^{-\Lambda} N(t \mid \theta, \sigma^2)dt \qquad \int_{\Lambda}^{\infty} N(t \mid \theta, \sigma^2)dt$$

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

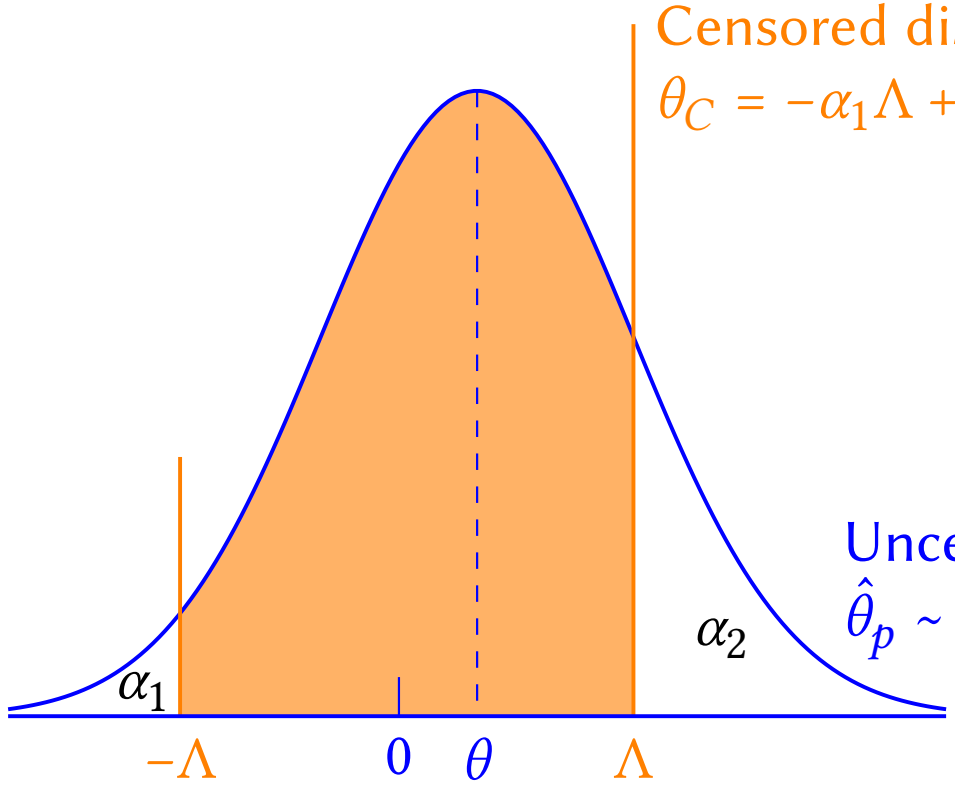

$$\int_{-\infty}^{-\Lambda} N(t \mid \theta, \sigma^2)dt \qquad \int_{\Lambda}^{\infty} N(t \mid \theta, \sigma^2)dt$$

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

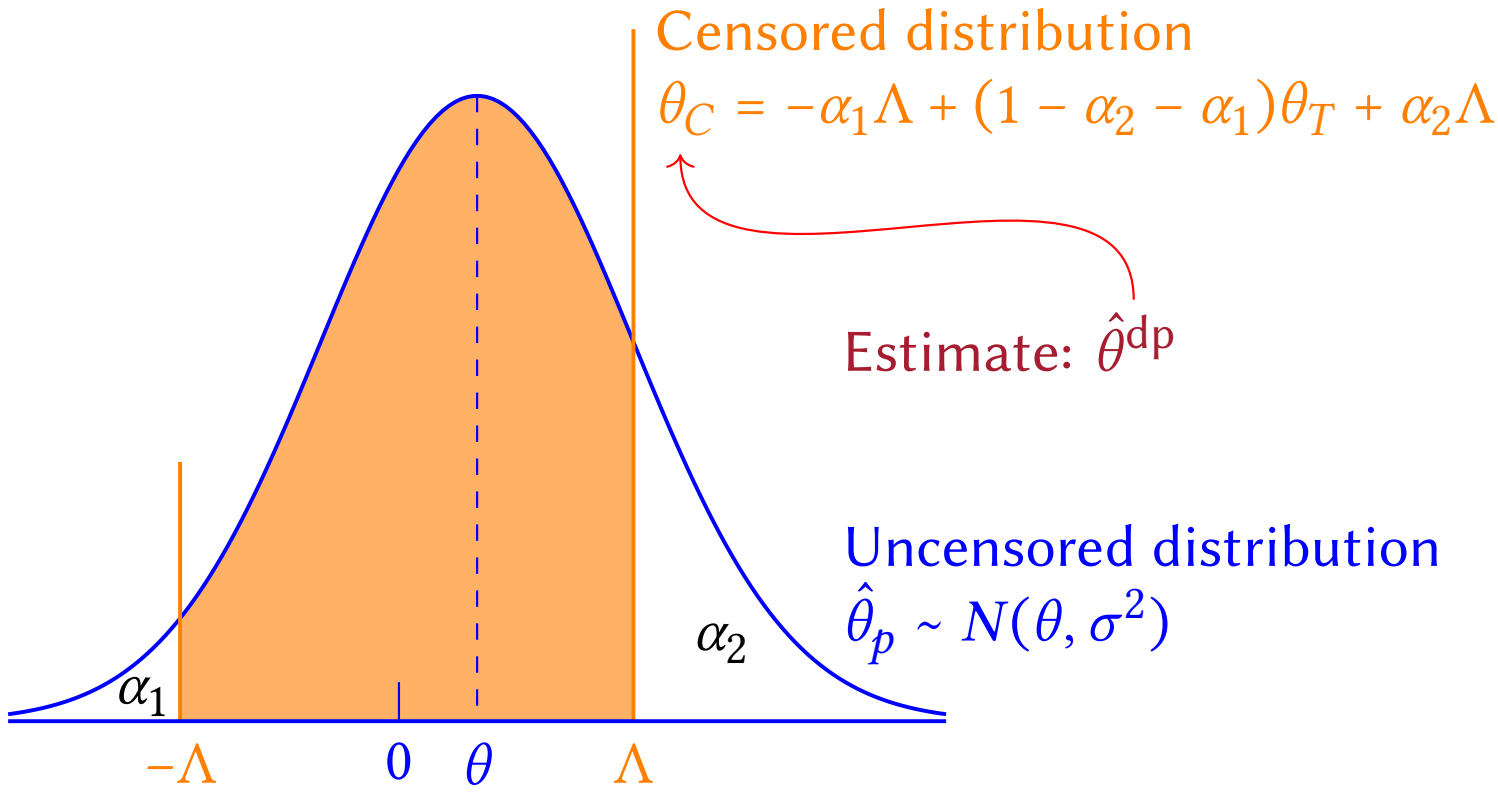

$$\int_{-\infty}^{-\Lambda} N(t \mid \theta, \sigma^2)dt \qquad \int_{\Lambda}^{\infty} N(t \mid \theta, \sigma^2)dt$$

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

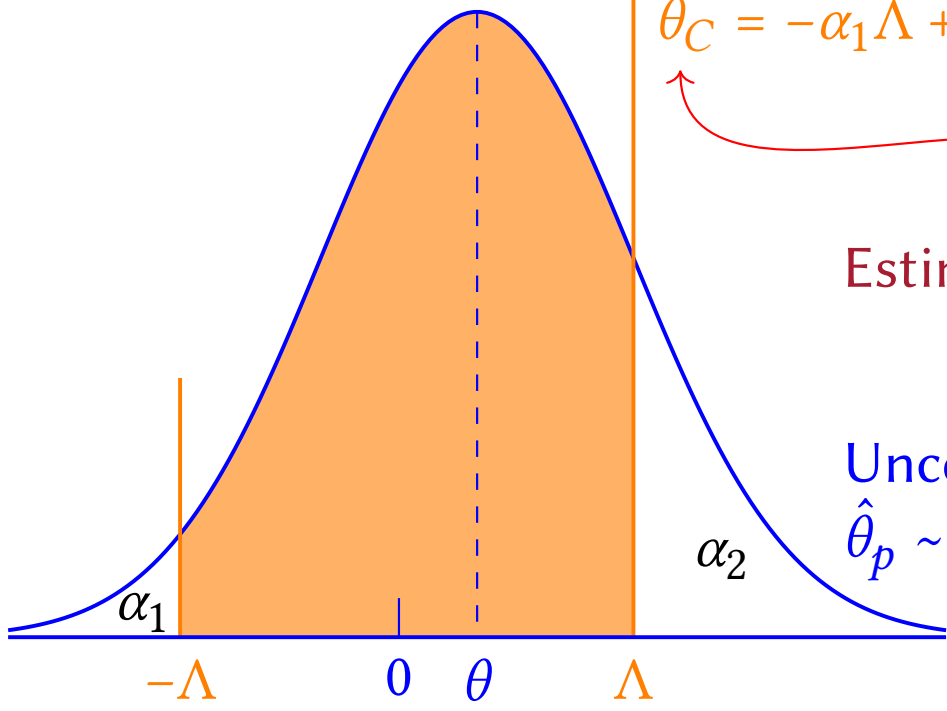

$$\int_{-\infty}^{-\Lambda} N(t \mid \theta, \sigma^2)dt \qquad \int_{\Lambda}^{\infty} N(t \mid \theta, \sigma^2)dt$$

3 eqns, 4 unknowns $\theta, \sigma^2, \alpha_1, \alpha_2$

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

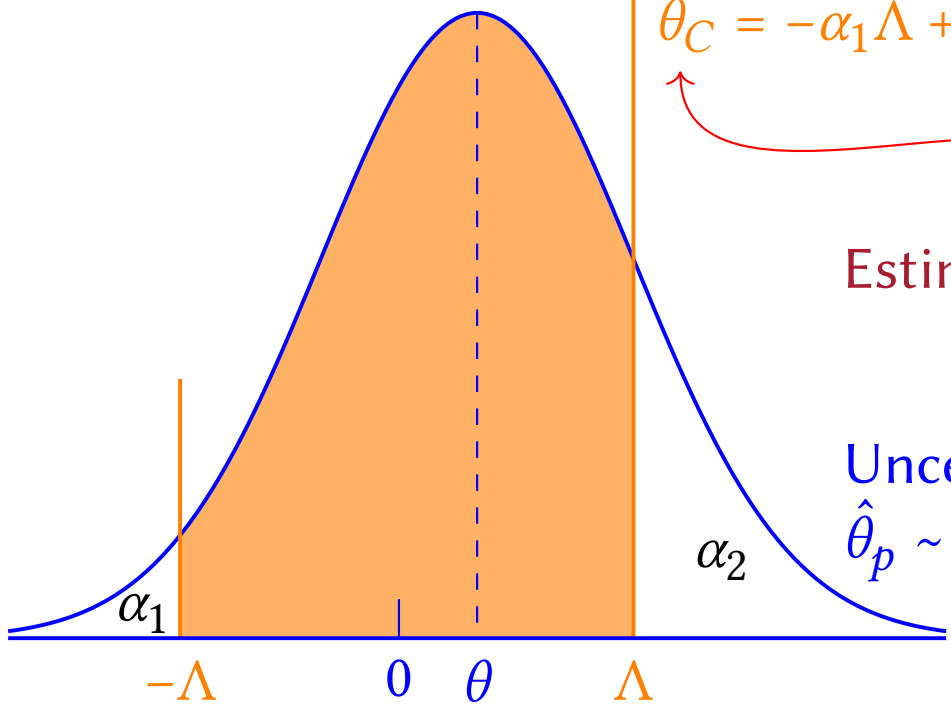

Censored distribution

$\theta_C = -\alpha_1\Lambda + (1 - \alpha_2 - \alpha_1)\theta_T + \alpha_2\Lambda$

Estimate: $\hat{\theta}^{\text{dp}}$, $\hat{\alpha}_2^{\text{dp}}$

Uncensored distribution

$\hat{\theta}_p \sim N(\theta, \sigma^2)$

$$\int_{-\infty}^{-\Lambda} N(t \mid \theta, \sigma^2)dt \qquad \int_{\Lambda}^{\infty} N(t \mid \theta, \sigma^2)dt$$

3 eqns, 4 unknowns $\theta, \sigma^2, \alpha_1, \alpha_2$

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

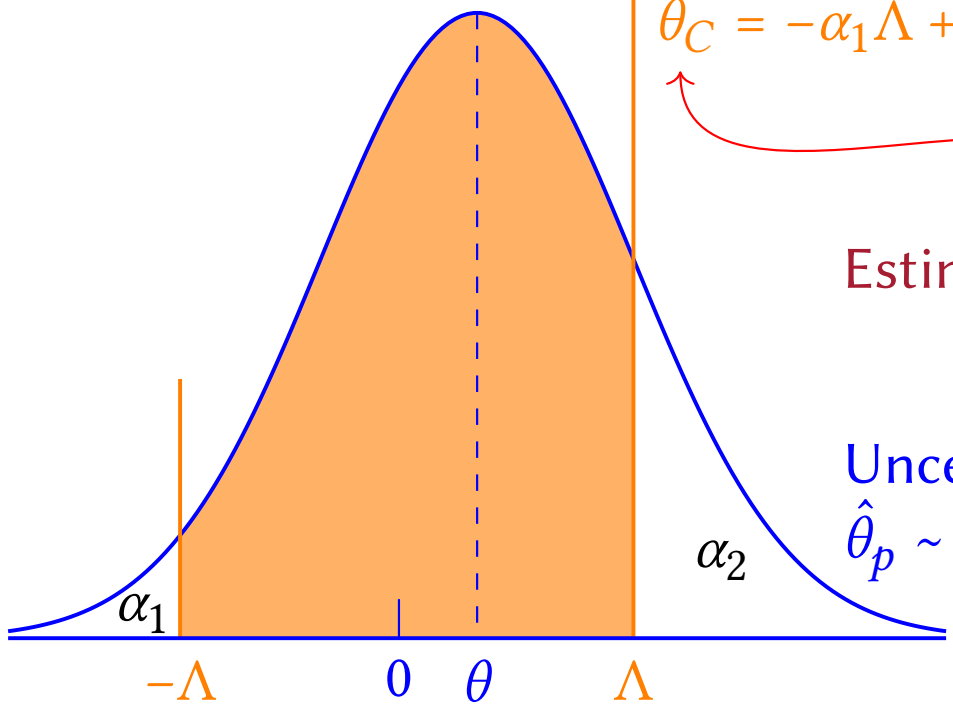

$$\int_{-\infty}^{-\Lambda} N(t \mid \theta, \sigma^2)dt \qquad \int_{\Lambda}^{\infty} N(t \mid \theta, \sigma^2)dt$$

~~3 eqns, 4 unknowns $\theta, \sigma^2, \alpha_1, \alpha_2$~~

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{4\Lambda}{\epsilon P}\right)$

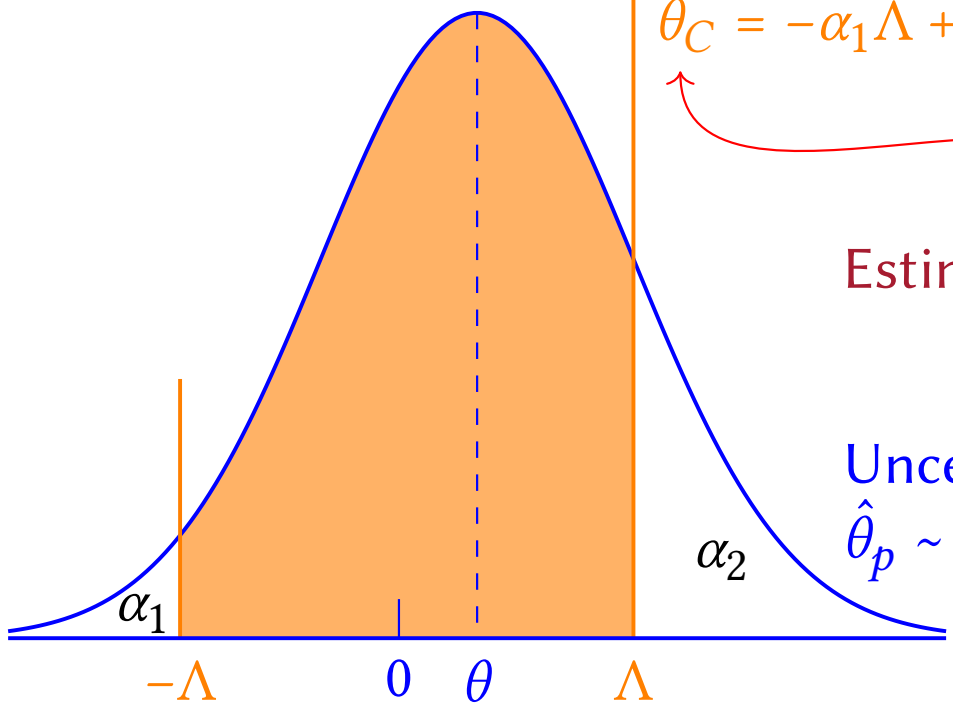

$$\int_{-\infty}^{-\Lambda} N(t \mid \theta, \sigma^2)dt \qquad \int_{\Lambda}^{\infty} N(t \mid \theta, \sigma^2)dt$$

~~3 eqns, 4 unknowns $\theta, \sigma^2, \alpha_1, \alpha_2$~~

Solve for θ (and σ^2, α_1)

Variance Estimation

Variance Estimation

- DP Variance is unhelpful: $V(\hat{\theta})^{\text{dp}} \neq V(\hat{\theta}^{\text{dp}})$

Variance Estimation

- DP Variance is unhelpful: $V(\hat{\theta})^{\text{dp}} \neq V(\hat{\theta}^{\text{dp}})$
- Simulate estimates via standard (Clarify) procedures:

$$\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i) \sim N\left(\begin{bmatrix} \hat{\theta}^{\text{dp}} \\ \hat{\alpha}_2^{\text{dp}} \end{bmatrix}, \begin{bmatrix} \hat{V}(\hat{\theta}^{\text{dp}}) & \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) \\ \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) & \hat{V}(\hat{\alpha}_2^{\text{dp}}) \end{bmatrix}\right)$$

Variance Estimation

- DP Variance is unhelpful: $V(\hat{\theta})^{\text{dp}} \neq V(\hat{\theta}^{\text{dp}})$
- Simulate estimates via standard (Clarify) procedures:

$$\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i) \sim N\left(\begin{bmatrix} \hat{\theta}^{\text{dp}} \\ \hat{\alpha}_2^{\text{dp}} \end{bmatrix}, \begin{bmatrix} \hat{V}(\hat{\theta}^{\text{dp}}) & \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) \\ \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) & \hat{V}(\hat{\alpha}_2^{\text{dp}}) \end{bmatrix}\right)$$

Functions of disclosed params

Variance Estimation

- DP Variance is unhelpful: $V(\hat{\theta})^{\text{dp}} \neq V(\hat{\theta}^{\text{dp}})$
- Simulate estimates via standard (Clarify) procedures:

$$\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i) \sim N\left(\begin{bmatrix} \hat{\theta}^{\text{dp}} \\ \hat{\alpha}_2^{\text{dp}} \end{bmatrix}, \begin{bmatrix} \hat{V}(\hat{\theta}^{\text{dp}}) & \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) \\ \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) & \hat{V}(\hat{\alpha}_2^{\text{dp}}) \end{bmatrix}\right)$$

Functions of disclosed params

- Bias correct simulated params:

$$\{\tilde{\theta}^{\text{dp}}(i), \hat{\alpha}_1^{\text{dp}}(i), \hat{\sigma}^2_{\text{dp}}(i)\} = \text{BiasCorrect}\left[\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i)\right]$$

Variance Estimation

- DP Variance is unhelpful: $V(\hat{\theta})^{\text{dp}} \neq V(\hat{\theta}^{\text{dp}})$
- Simulate estimates via standard (Clarify) procedures:

$$\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i) \sim N\left(\begin{bmatrix} \hat{\theta}^{\text{dp}} \\ \hat{\alpha}_2^{\text{dp}} \end{bmatrix}, \begin{bmatrix} \hat{V}(\hat{\theta}^{\text{dp}}) & \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) \\ \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) & \hat{V}(\hat{\alpha}_2^{\text{dp}}) \end{bmatrix}\right)$$

Functions of disclosed params

- Bias correct simulated params:

$$\{\tilde{\theta}^{\text{dp}}(i), \hat{\alpha}_1^{\text{dp}}(i), \hat{\sigma}^2_{\text{dp}}(i)\} = \text{BiasCorrect}\left[\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i)\right]$$

- Standard error, $\text{SE}(\tilde{\theta}^{\text{dp}})$: Standard deviation of $\tilde{\theta}^{\text{dp}}(i)$ over i

Variance Estimation

- DP Variance is unhelpful: $V(\hat{\theta})^{\text{dp}} \neq V(\hat{\theta}^{\text{dp}})$
- Simulate estimates via standard (Clarify) procedures:

$$\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i) \sim N\left(\begin{bmatrix} \hat{\theta}^{\text{dp}} \\ \hat{\alpha}_2^{\text{dp}} \end{bmatrix}, \begin{bmatrix} \hat{V}(\hat{\theta}^{\text{dp}}) & \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) \\ \widehat{\text{Cov}}(\hat{\alpha}_2^{\text{dp}}, \hat{\theta}^{\text{dp}}) & \hat{V}(\hat{\alpha}_2^{\text{dp}}) \end{bmatrix}\right)$$

Functions of disclosed params

- Bias correct simulated params:

$$\{\tilde{\theta}^{\text{dp}}(i), \hat{\alpha}_1^{\text{dp}}(i), \hat{\sigma}^2_{\text{dp}}(i)\} = \text{BiasCorrect}\left[\hat{\theta}^{\text{dp}}(i), \hat{\alpha}_2^{\text{dp}}(i)\right]$$

- Standard error, SE($\tilde{\theta}^{\text{dp}}$): Standard deviation of $\tilde{\theta}^{\text{dp}}(i)$ over i
- Bias correction (usually) reduces bias *and* variance:

$$V(\tilde{\theta}^{\text{dp}}) < V(\hat{\theta}^{\text{dp}})$$

Solving Political Problems Technologically

Differential Privacy & Inferential Validity

A General Purpose, Statistically Valid DP Algorithm

The Algorithm in Practice

Simulations: Finite Sample Evaluation

Simulations: Finite Sample Evaluation

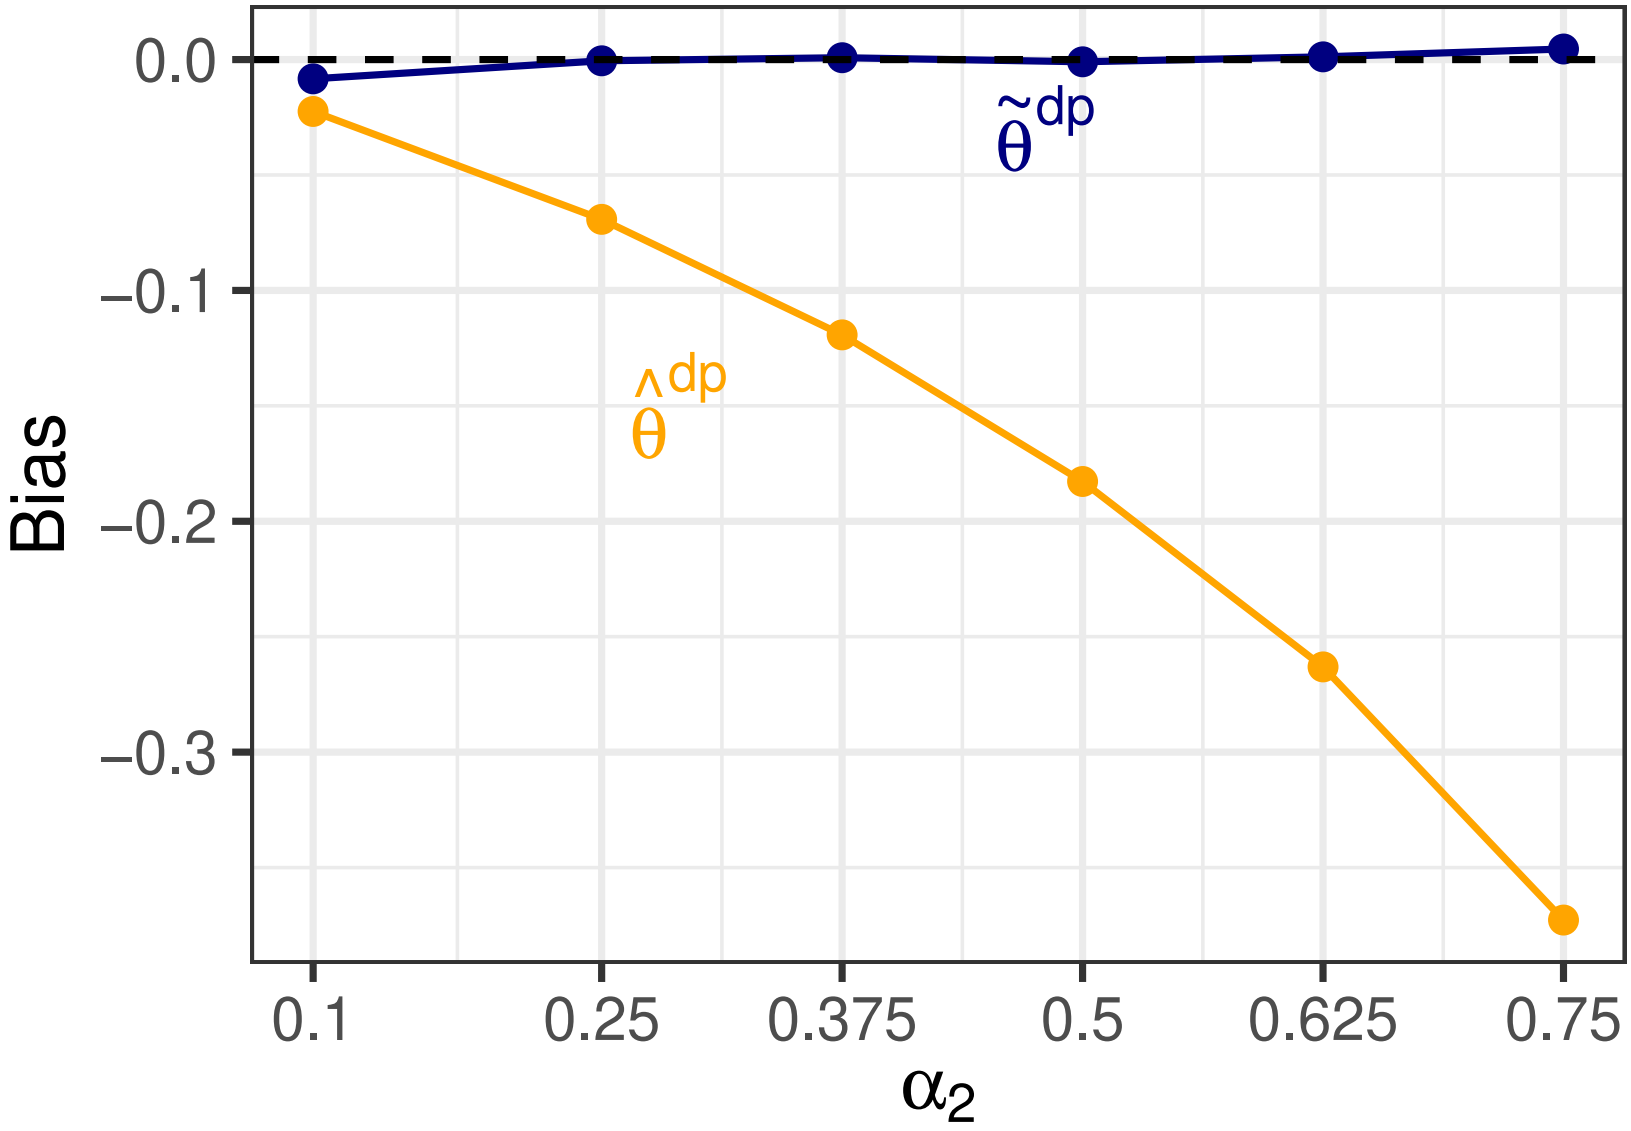

Simulations: Finite Sample Evaluation

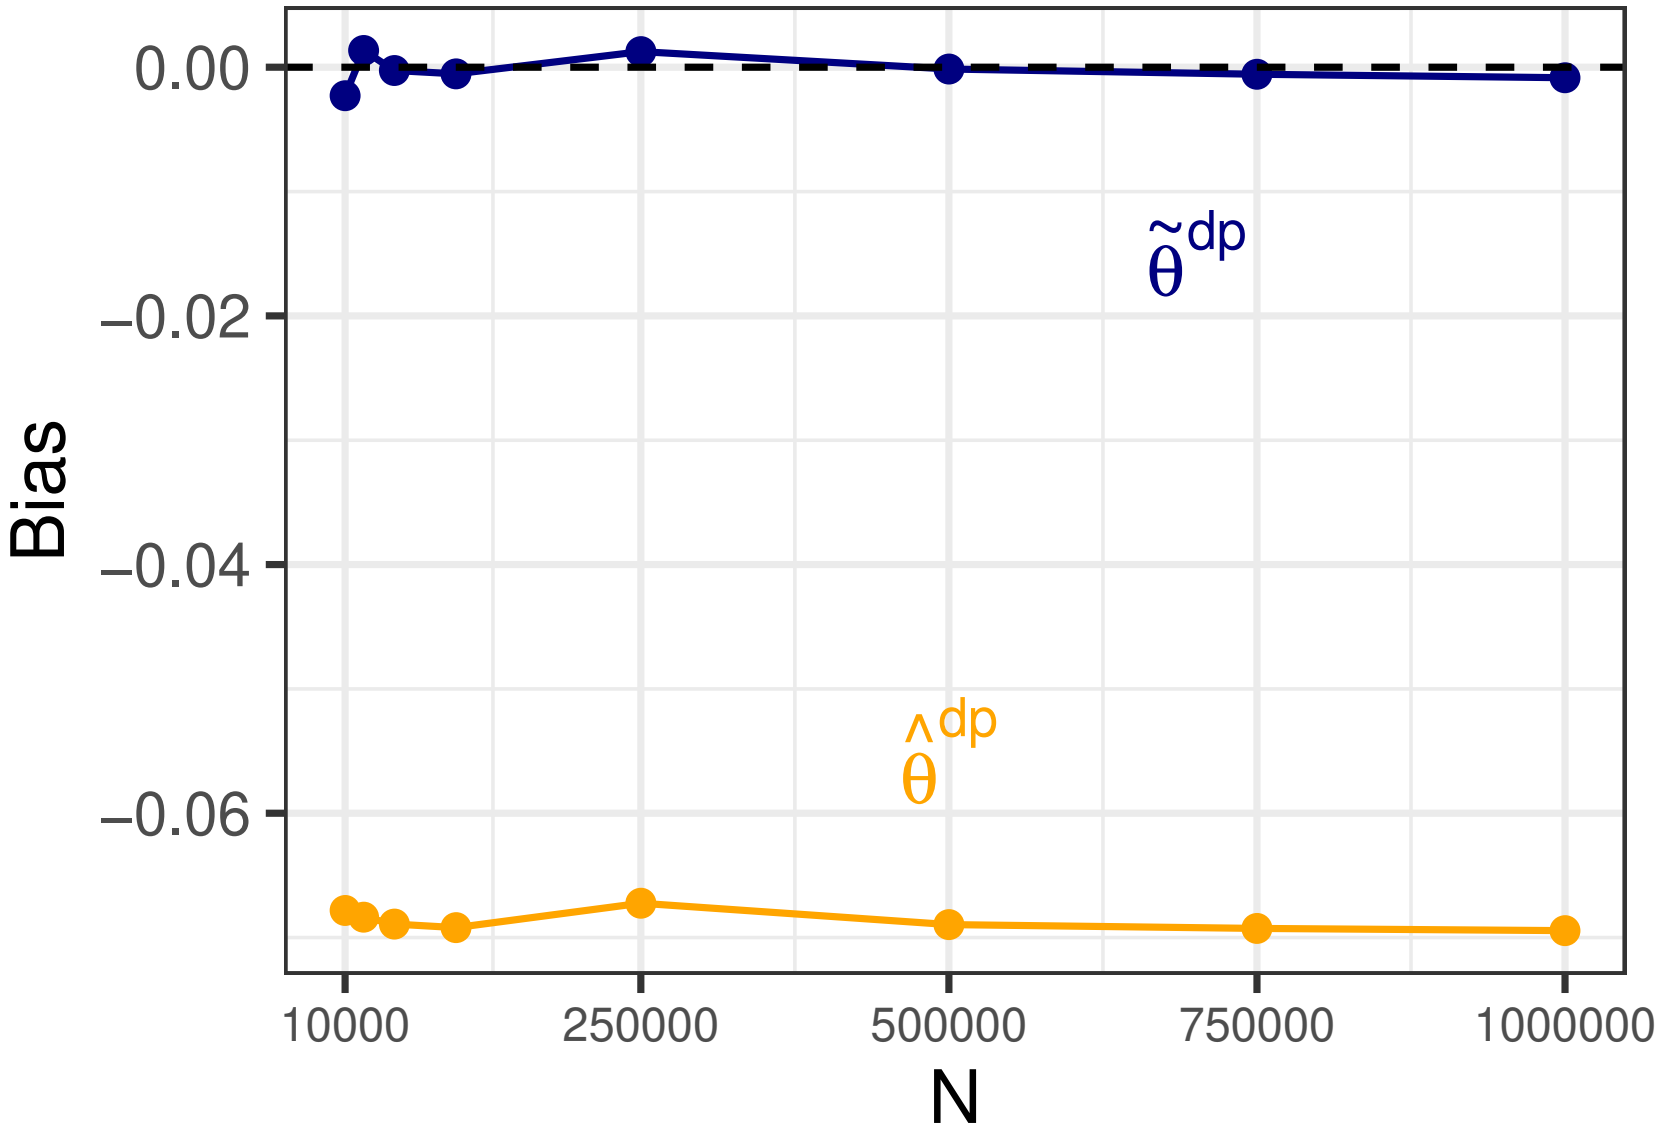

Simulations: Finite Sample Evaluation

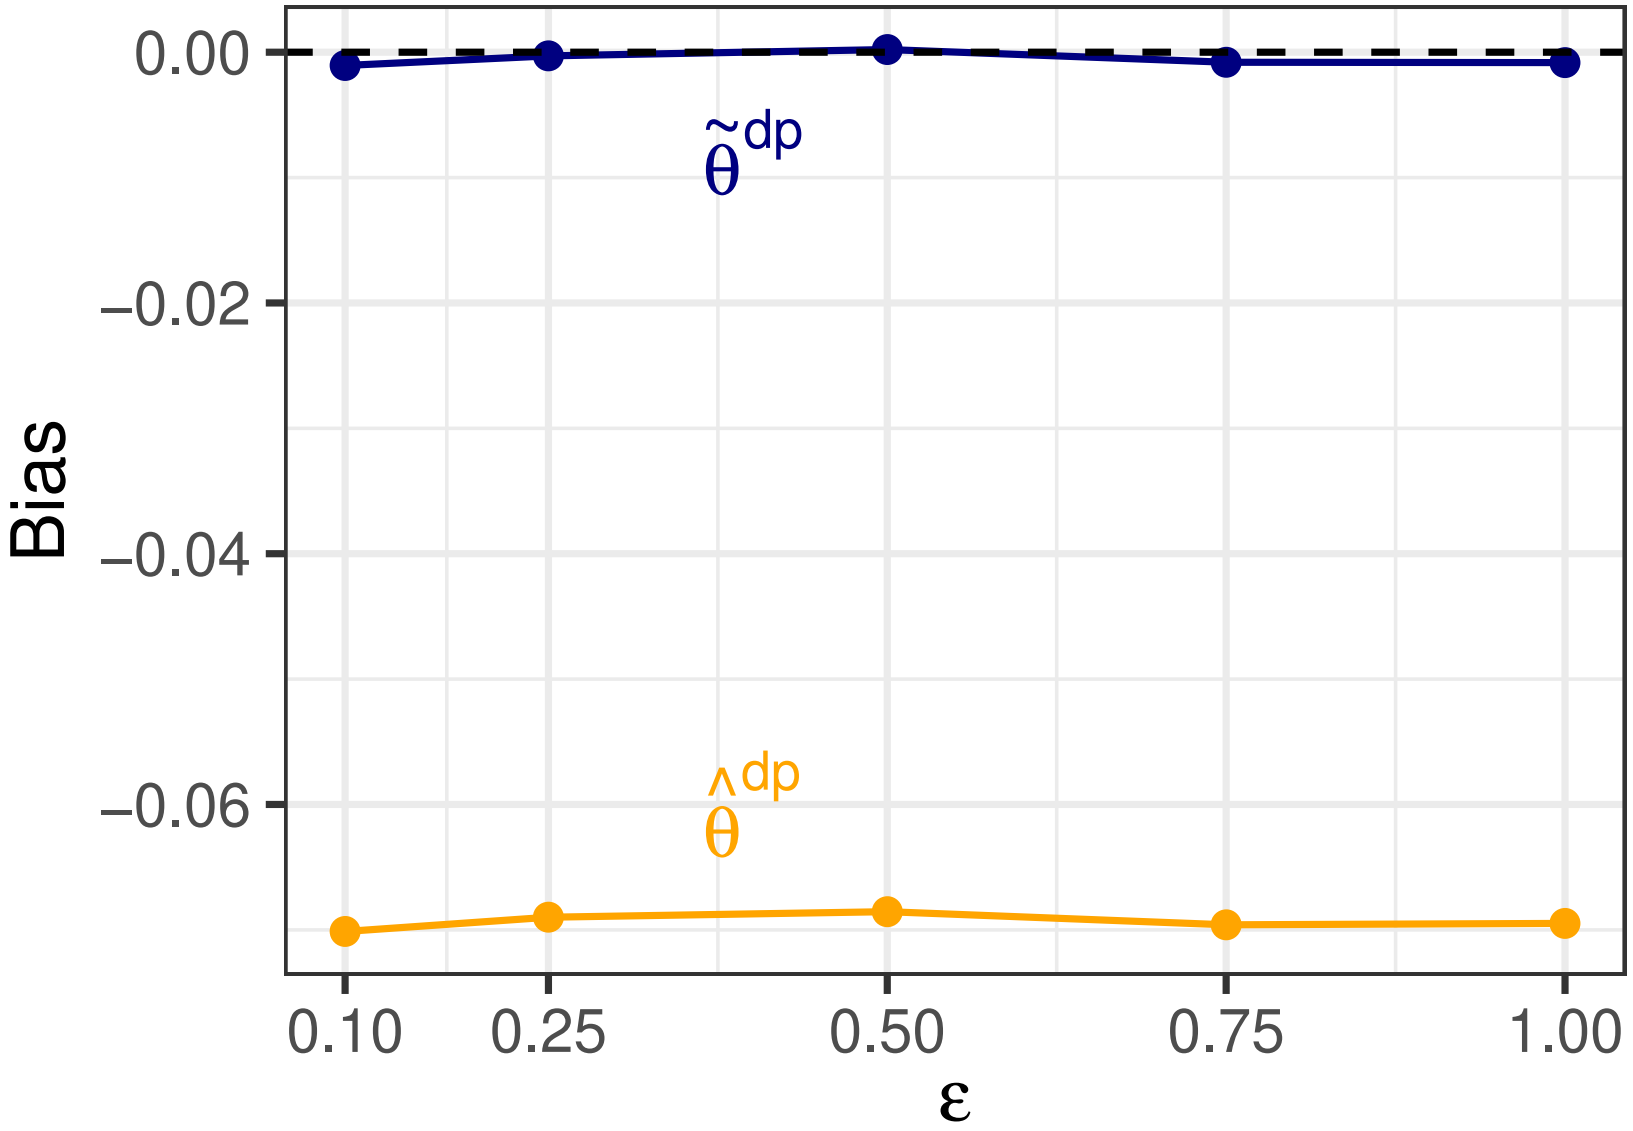

Simulations: Finite Sample Evaluation

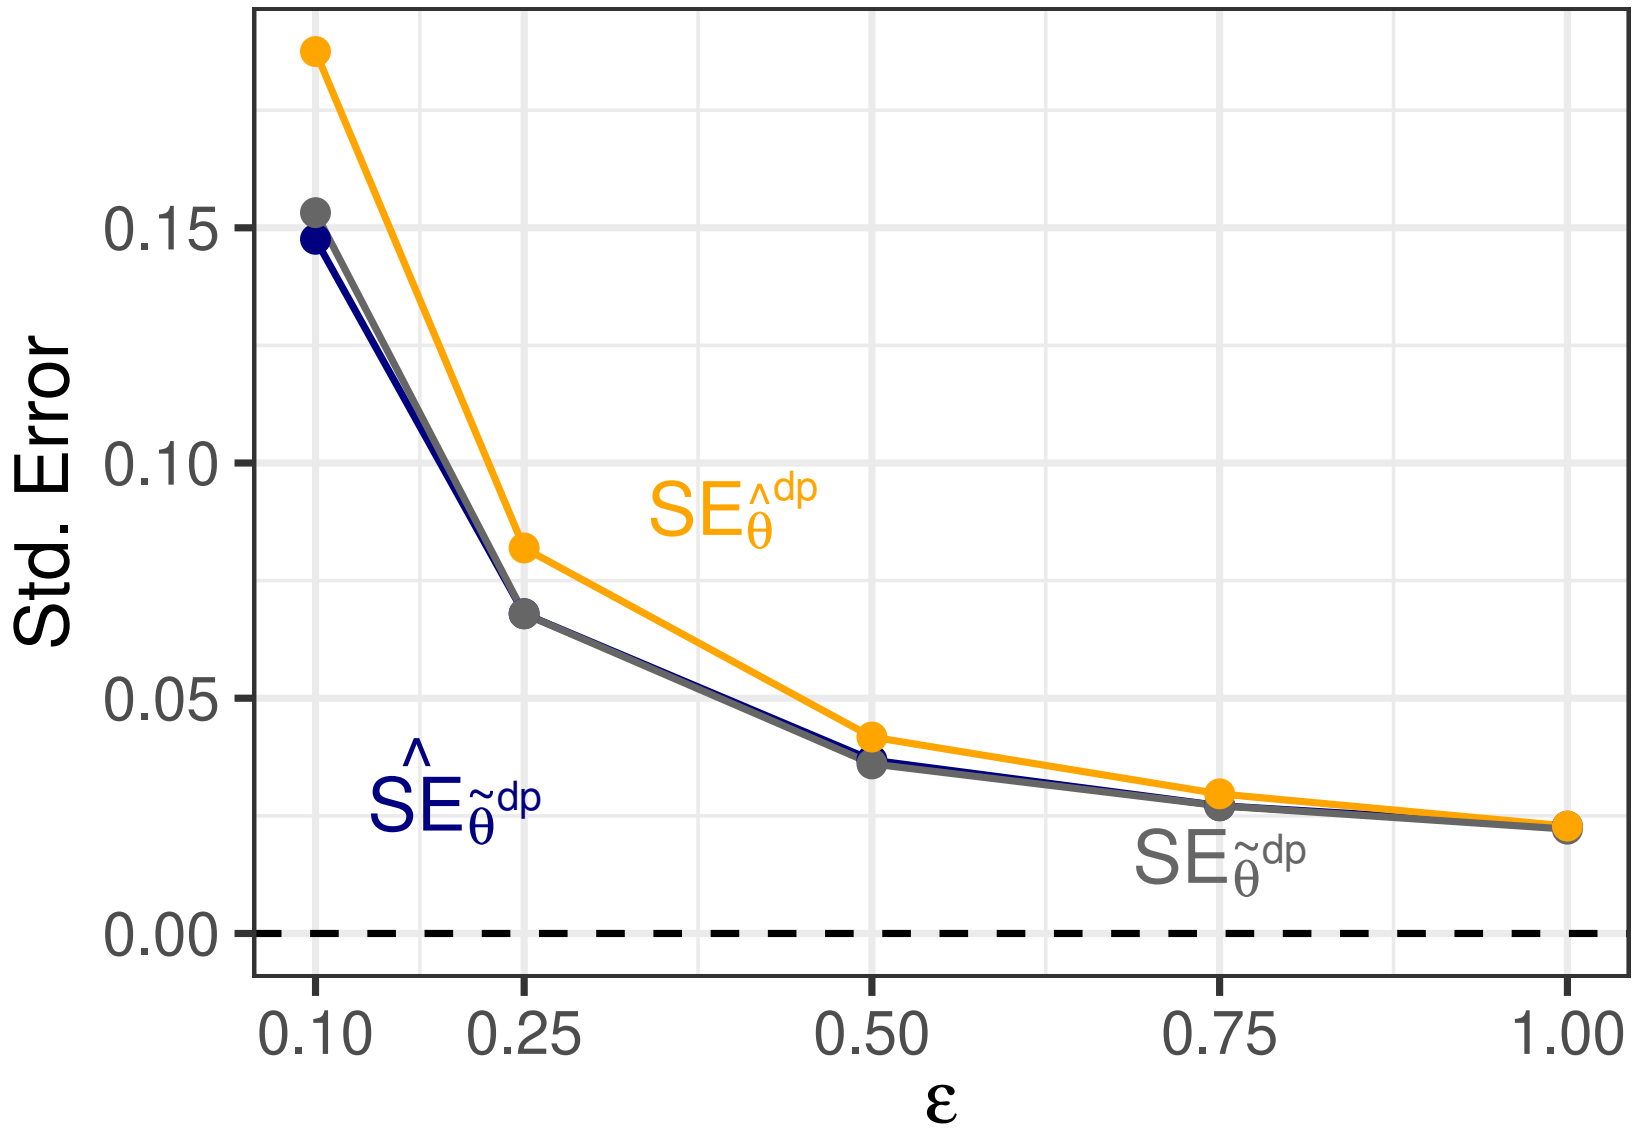

Simulations: Finite Sample Evaluation

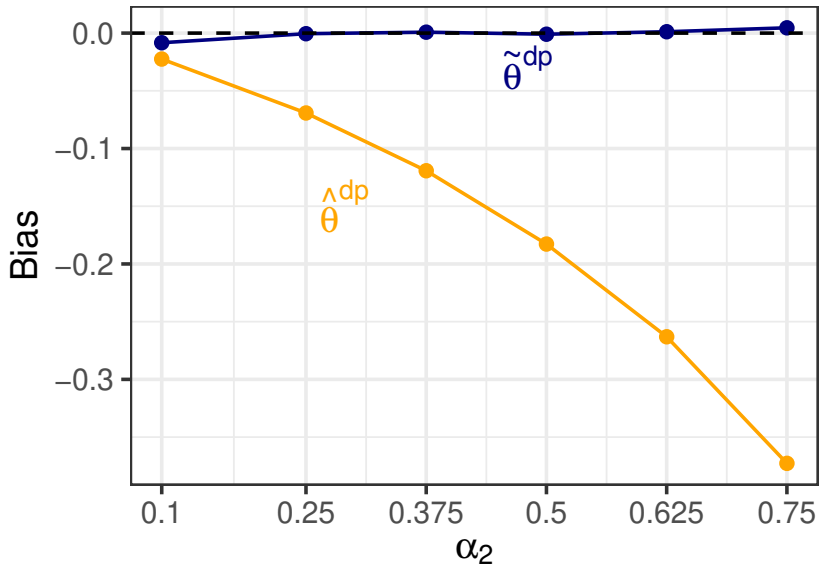

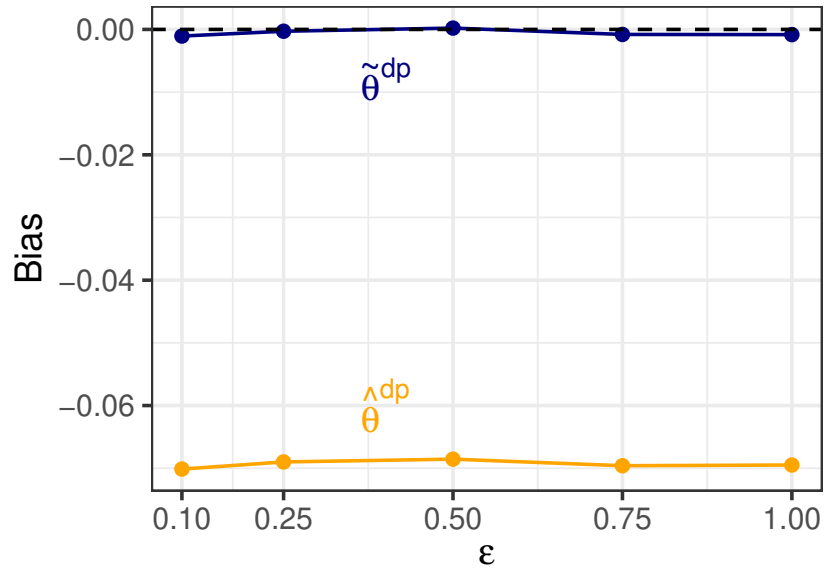

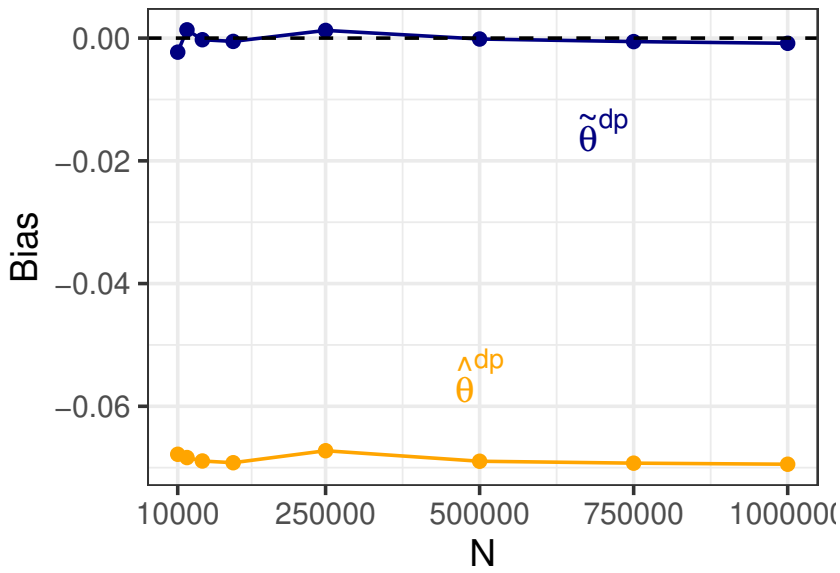

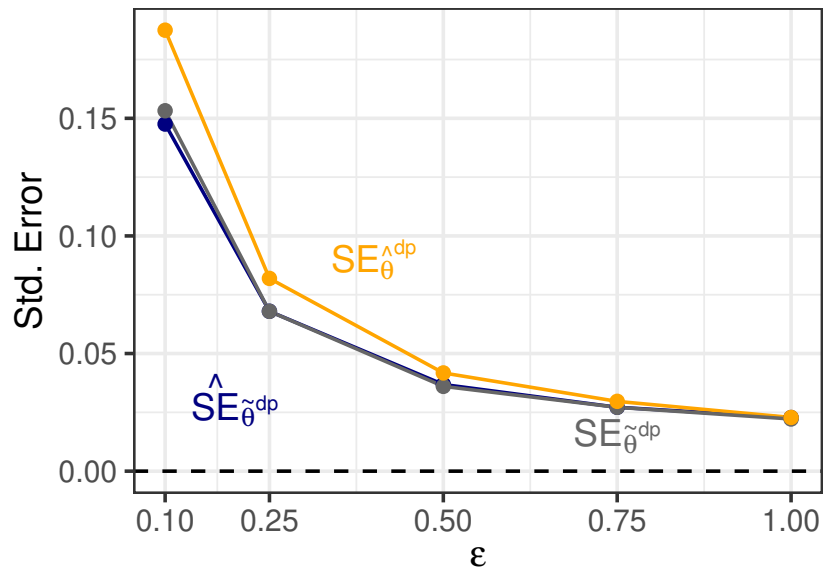

Theory and Practice

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ
  - Without bias correction: choose more censoring or more noise!

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ
  - Without bias correction: choose more censoring or more noise!
  - With bias correction: Keep $\max(\alpha_1, \alpha_2) < 0.6$

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ
  - Without bias correction: choose more censoring or more noise!
  - With bias correction: Keep $\max(\alpha_1, \alpha_2) < 0.6$
- Privacy Policies:

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ
  - Without bias correction: choose more censoring or more noise!
  - With bias correction: Keep $\max(\alpha_1, \alpha_2) < 0.6$
- Privacy Policies:
  - Science informs, but does not determine, policy

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ
  - Without bias correction: choose more censoring or more noise!
  - With bias correction: Keep $\max(\alpha_1, \alpha_2) < 0.6$
- Privacy Policies:
  - Science informs, but does not determine, policy
  - Few if any implementations exactly meet DP standards

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ
  - Without bias correction: choose more censoring or more noise!
  - With bias correction: Keep $\max(\alpha_1, \alpha_2) < 0.6$
- Privacy Policies:
  - Science informs, but does not determine, policy
  - Few if any implementations exactly meet DP standards
  - Most use larger ϵ and no budget, but with other protections

Theory and Practice

- Reducing DP's Societal Risks. Report:

$$\text{Effective reduction in } N \quad = \quad 1 - \frac{\hat{\sigma}^2_{\text{dp}}/P}{\text{SE}(\tilde{\theta}^{\text{dp}})}$$

- Choosing ϵ (like a power calculation):

$$\text{SE}(\tilde{\theta}^{\text{dp}})^2 < V(\hat{\theta}^{\text{dp}}) + \left(\frac{4\Lambda}{\epsilon P}\right)^2$$

- Choosing Λ
  - Without bias correction: choose more censoring or more noise!
  - With bias correction: Keep $\max(\alpha_1, \alpha_2) < 0.6$
- Privacy Policies:
  - Science informs, but does not determine, policy
  - Few if any implementations exactly meet DP standards
  - Most use larger ϵ and no budget, but with other protections
  - It's safer: de-identification + noise and censoring

Concluding Remarks

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty
  - Utility requires known statistical properties and valid uncertainty estimates

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty
  - Utility requires known statistical properties and valid uncertainty estimates
- Proposed algorithm

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty
  - Utility requires known statistical properties and valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty
  - Utility requires known statistical properties and valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased (if estimator is), lower variance

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty
  - Utility requires known statistical properties and valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased (if estimator is), lower variance
  - Valid uncertainty estimates

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty
  - Utility requires known statistical properties and valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased (if estimator is), lower variance
  - Valid uncertainty estimates
  - Computationally efficient

Concluding Remarks

- Data sharing $\rightsquigarrow$ data access
  - DP protects individual privacy
  - Enables inference to private database, not population
  - Usually biased, no uncertainty estimates
  - Fails to protect society from fallacious scientific conclusions
- Inferential validity
  - A scientific statement is not one that is correct; it is one that comes with an appropriate degree of uncertainty
  - Utility requires known statistical properties and valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased (if estimator is), lower variance
  - Valid uncertainty estimates
  - Computationally efficient
  - Easy to implement

For more information

Georgina-Evans.com

GaryKing.org

MegSchwenzfeier.com

bit.ly/AbhradeepThakurta

Paper, software, slides: GaryKing.org/dp