Statistically Valid Inferences from Privacy Protected Data[1]

Gary King[2]

Institute for Quantitative Social Science
Harvard University

Google, 9/18/2020

[1] Joint 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!)

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!)
- 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!)
- 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!)
- 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!)
- 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!)
- 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!)
- 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!)
- 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!)
- 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!)
- 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!)
- 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!)
- 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
⋮
Lindsay
Salil
Georgie
Gary
Meg
Abhradeep
Joshua
Annie
Bob
Ellen
Mean income: $48

Quantity of Interest

Theories of Inference: Statistics vs. CS

Population	Sample
⋮	✗
Lindsay	✓
Salil	✓
Georgie	✓
Gary	✓
Meg	✓
Abhradeep	✓
Joshua	✓
Annie	✓
Bob	✓
Ellen	✓
Mean income:	$48

Quantity of Interest

Theories of Inference: Statistics vs. CS

Population	Sample	$
⋮	✗	?
Lindsay	✓	122
Salil	✓	76
Georgie	✓	145
Gary	✓	96
Meg	✓	86
Abhradeep	✓	127
Joshua	✓	72
Annie	✓	132
Bob	✓	95
Ellen	✓	134
Mean income:	$48	$108

Quantity of Interest

Usually no direct relevance

Theories of Inference: Statistics vs. CS

Population	Sample	$
⋮	✗	?
Lindsay	✓	122
Salil	✓	76
Georgie	✓	145
Gary	✓	96
Meg	✓	86
Abhradeep	✓	127
Joshua	✓	72
Annie	✓	132
Bob	✓	95
Ellen	✓	134
Mean income:	$48	$108

Quantity of Interest

Usually no direct relevance

Theories of Inference: Statistics vs. CS

Population	Sample	$	+Privacy
⋮	✗	?	
Lindsay	✓	122	
Salil	✓	76	
Georgie	✓	145	Noise & Censoring
Gary	✓	96	
Meg	✓	86	
Abhradeep	✓	127	
Joshua	✓	72	
Annie	✓	132	
Bob	✓	95	
Ellen	✓	134	
Mean income:	$48		$108

Quantity of Interest

Usually no direct relevance

Theories of Inference: Statistics vs. CS

Population	Sample	$	+Privacy	=dp$
⋮	✗	?		
Lindsay	✓	122		85
Salil	✓	76		103
Georgie	✓	145	Noise & Censoring	75
Gary	✓	96		113
Meg	✓	86		125
Abhradeep	✓	127		97
Joshua	✓	72		101
Annie	✓	132		128
Bob	✓	95		83
Ellen	✓	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$
⋮	✗	?		
Lindsay	✓	122		85
Salil	✓	76		103
Georgie	✓	145	Noise & Censoring	75
Gary	✓	96		113
Meg	✓	86		125
Abhradeep	✓	127		97
Joshua	✓	72		101
Annie	✓	132		128
Bob	✓	95		83
Ellen	✓	134		201
Mean income:	$48	$108		$111

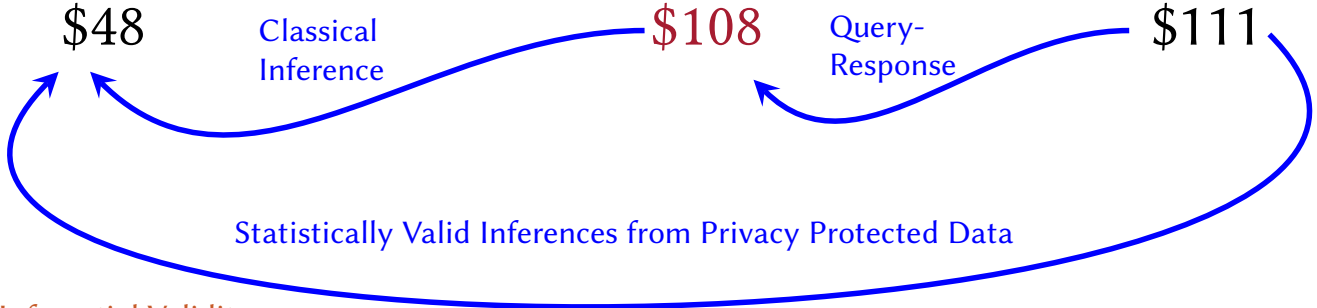

Differential Privacy and its Inferential Challenges

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 (simplifying)

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 (simplifying)
  - 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 (simplifying)
  - 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 (simplifying)
  - 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
  - $M(\text{mean}, D) = \frac{1}{n} \sum_{i=1}^{n} c(y_i, \Lambda) + N\left(0, \frac{8\Lambda}{n\epsilon}\right)$ (Λ, n, ϵ known)

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 (simplifying)
  - 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
  - $M(\text{mean}, D) = \frac{1}{n} \sum_{i=1}^{n} c(y_i, \Lambda) + N\left(0, \frac{8\Lambda}{n\epsilon}\right)$ (Λ, n, ϵ known)
  - 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 (simplifying)
  - 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
  - $M(\text{mean}, D) = \frac{1}{n} \sum_{i=1}^{n} c(y_i, \Lambda) + N\left(0, \frac{8\Lambda}{n\epsilon}\right)$ (Λ, n, ϵ known)
  - 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
- 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
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real 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
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - 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
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum 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
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:
    - P-hacking

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:
    - P-hacking
    - Threats to inference

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:
    - P-hacking $\rightsquigarrow$ pre-registration (e.g., clinical trials, Mars lander)
    - Threats to inference

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:
    - P-hacking $\rightsquigarrow$ pre-registration (e.g., clinical trials, Mars lander)
    - Threats to inference $\rightsquigarrow$ diagnostics, exploration, serendipity (e.g., observational data)

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:
    - P-hacking $\rightsquigarrow$ pre-registration (e.g., clinical trials, Mars lander)
    - Threats to inference $\rightsquigarrow$ diagnostics, exploration, serendipity (e.g., observational data)
  - With DP:

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:
    - P-hacking $\rightsquigarrow$ pre-registration (e.g., clinical trials, Mars lander)
    - Threats to inference $\rightsquigarrow$ diagnostics, exploration, serendipity (e.g., observational data)
  - With DP: ~~P-hacking~~,

Properties of Differential Privacy

- Post-processing: if $M(s, D)$ is DP, so is $f[M(s, D)]$
  - Useful for bias corrections
- Privacy risk quantified (ϵ), instead of 0/1 for re-ID
  - Helpful mathematically; insufficient in applications
- Real privacy loss $\ll$ maximum privacy loss
  - OK for worst case scenerio; unhelpful in practice
- Privacy Budget
  - Composition: ϵ_1-DP and ϵ_2-DP is $(\epsilon_1 + \epsilon_2)$-DP
  - Can limit maximum risks across analyses & researchers
  - When the budget is used, no new analyses can ever be run
- Completely changes statistical best practices
  - Without DP, we balance worries:
    - P-hacking $\rightsquigarrow$ pre-registration (e.g., clinical trials, Mars lander)
    - Threats to inference $\rightsquigarrow$ diagnostics, exploration, serendipity (e.g., observational data)
  - With DP: ~~P-hacking~~, surveys treated like the Mars lander

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

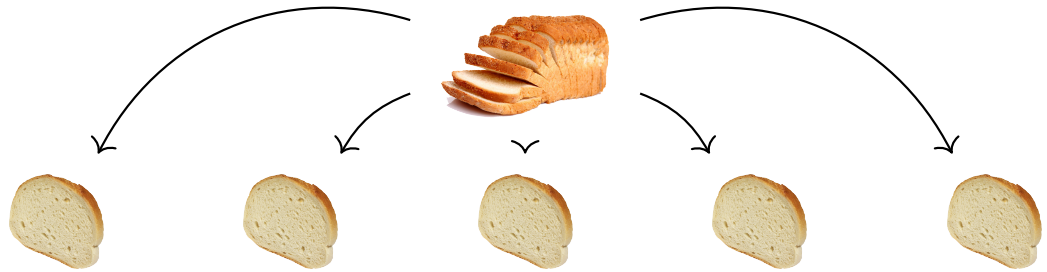

Private data

Partition

A Differentially Private Estimator

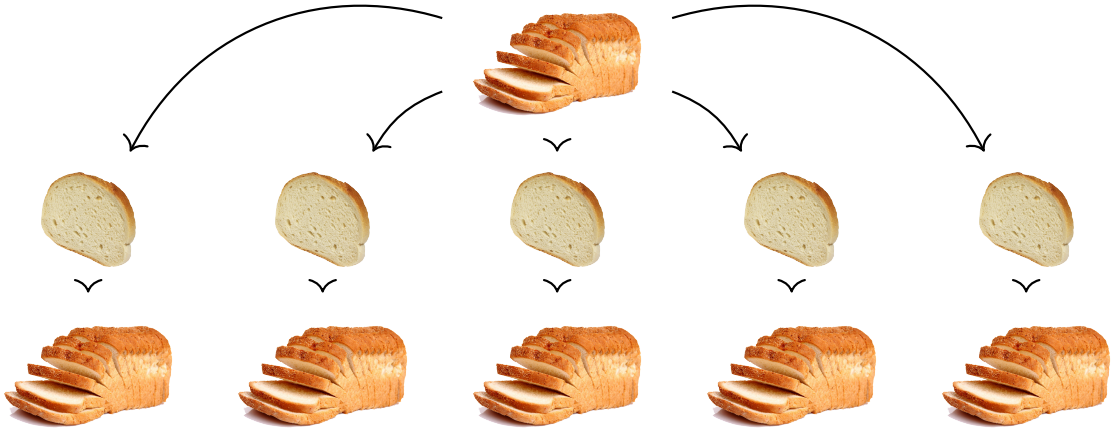

Private data

Partition

Bag of little bootstraps

A Differentially Private Estimator

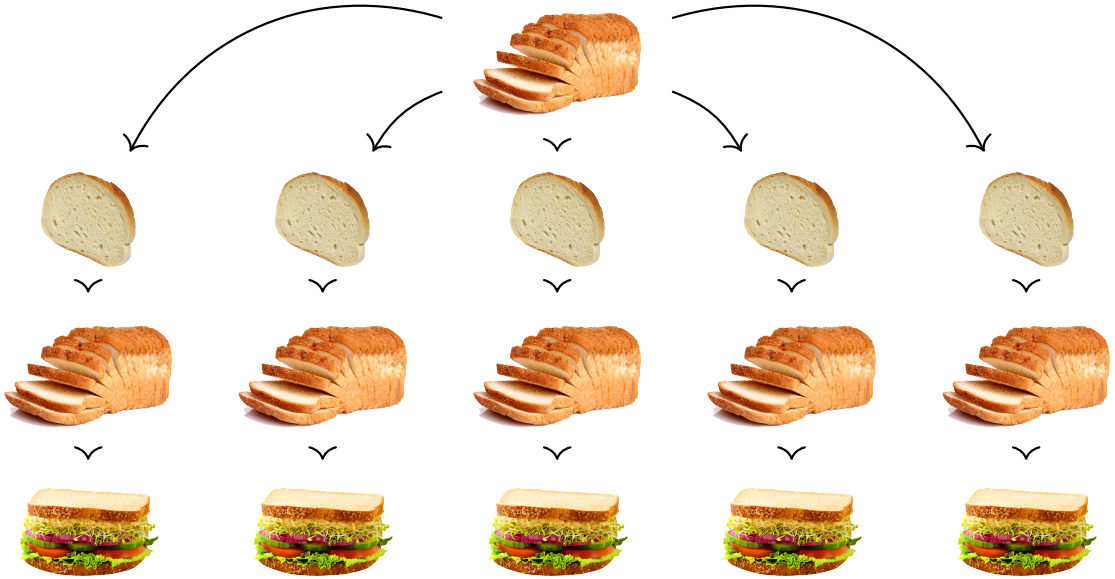

Private data

Partition

Bag of little bootstraps

Estimator

A Differentially Private Estimator

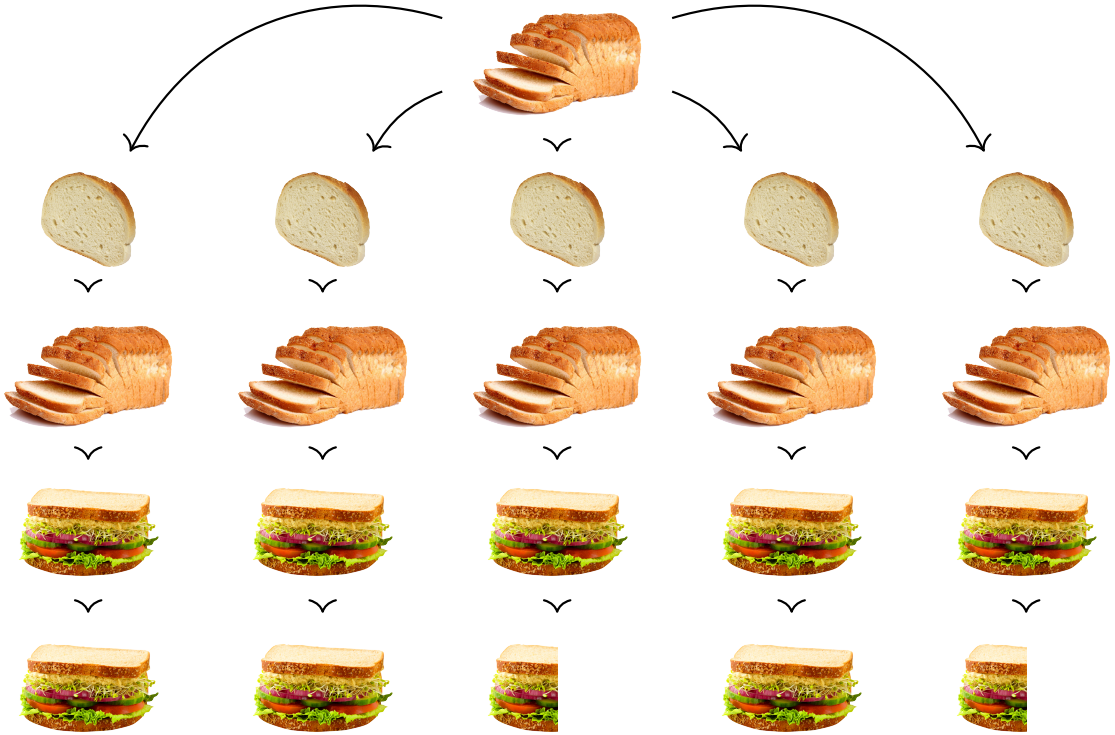

Private data

Partition

Bag of little bootstraps

Estimator

Censor

A Differentially Private Estimator

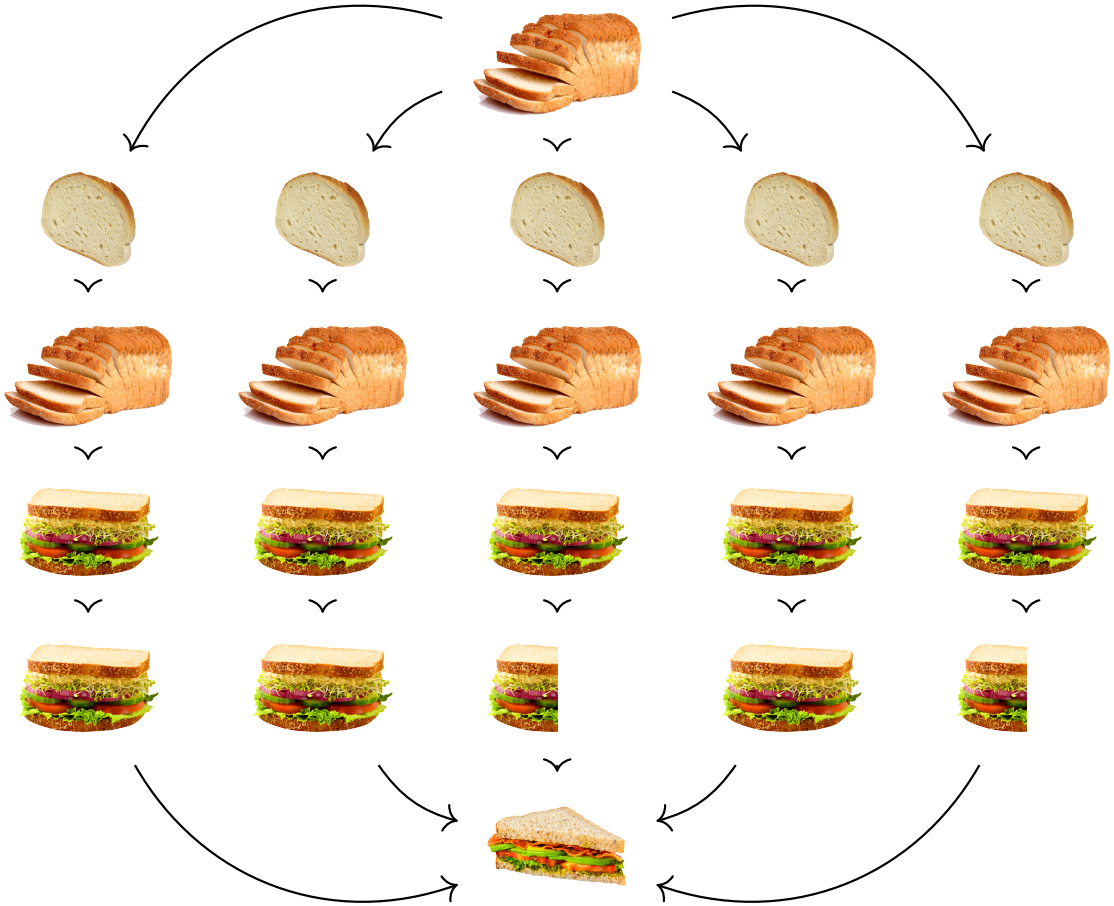

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

A Differentially Private Estimator

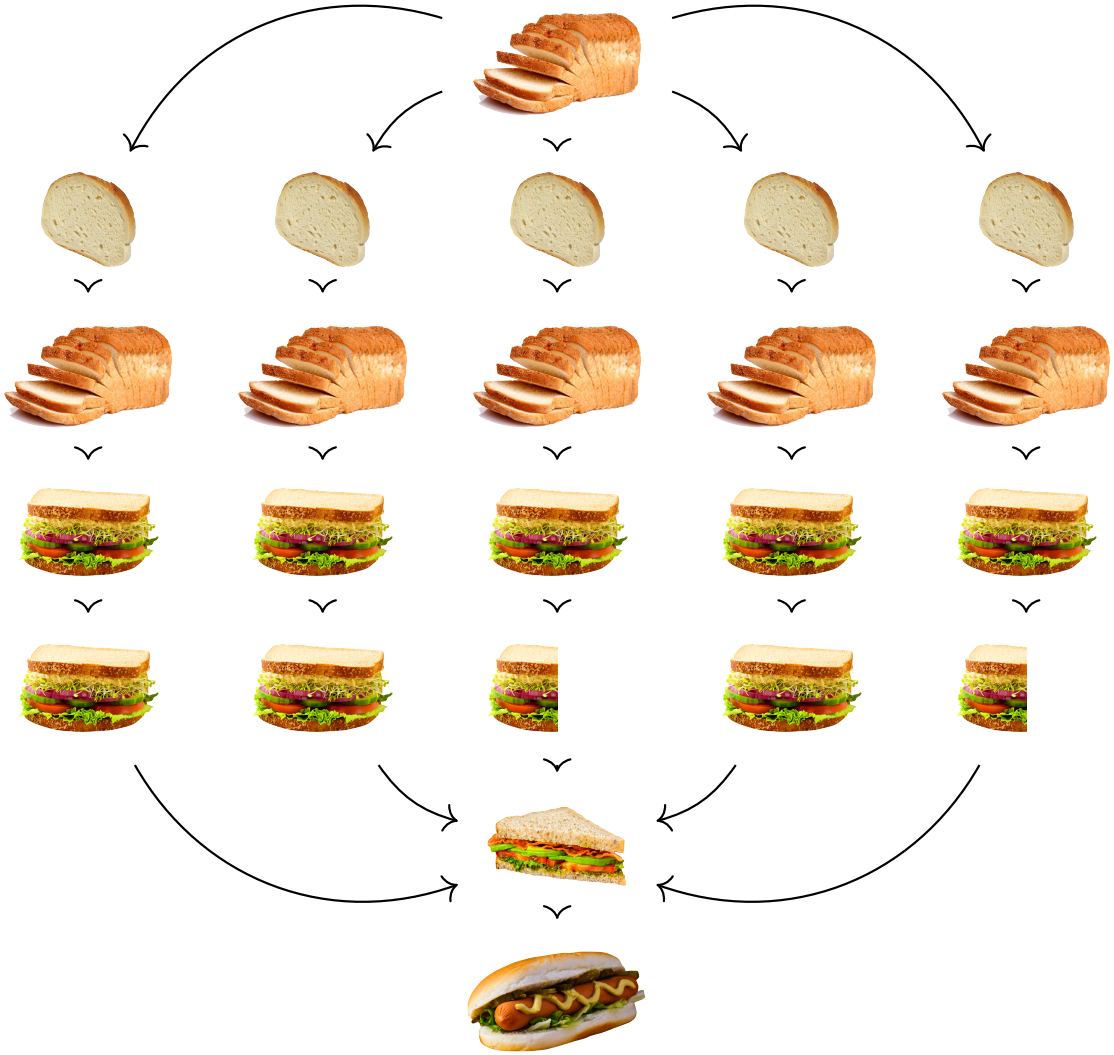

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

A Differentially Private Estimator

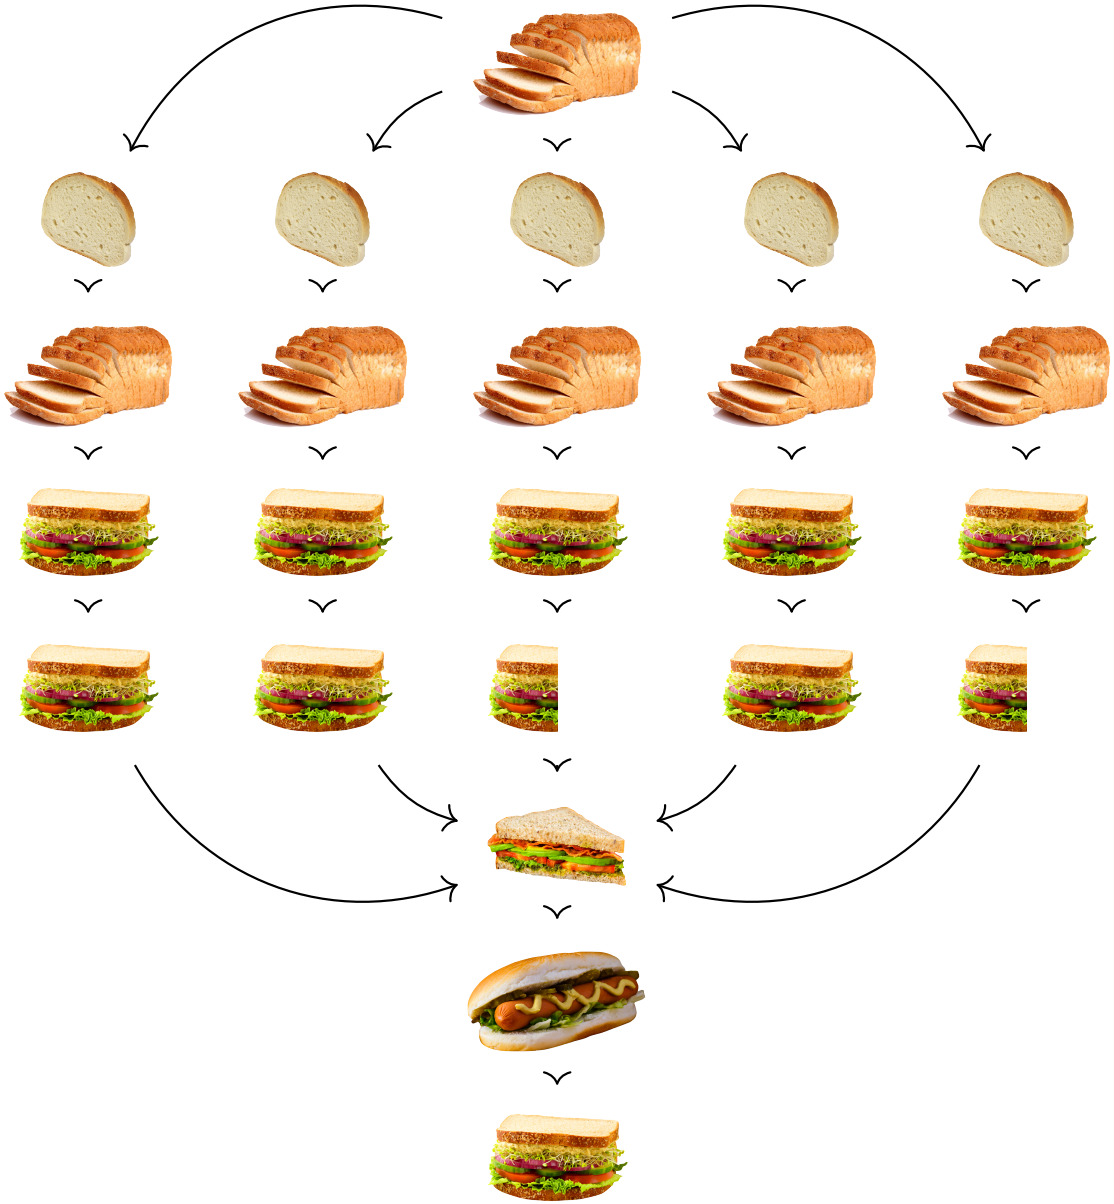

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction

A Differentially Private Estimator

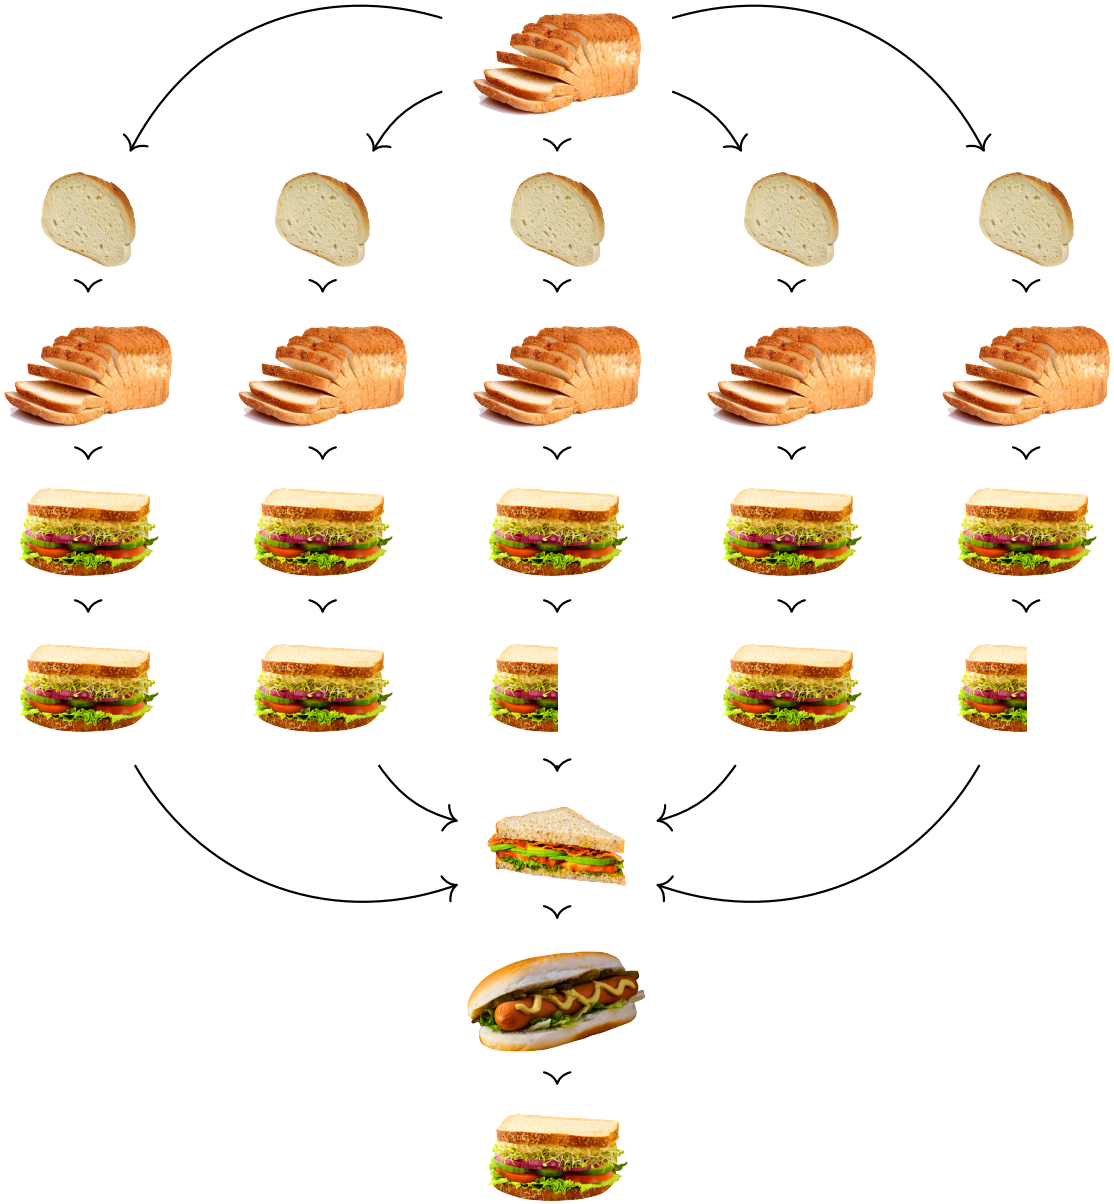

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

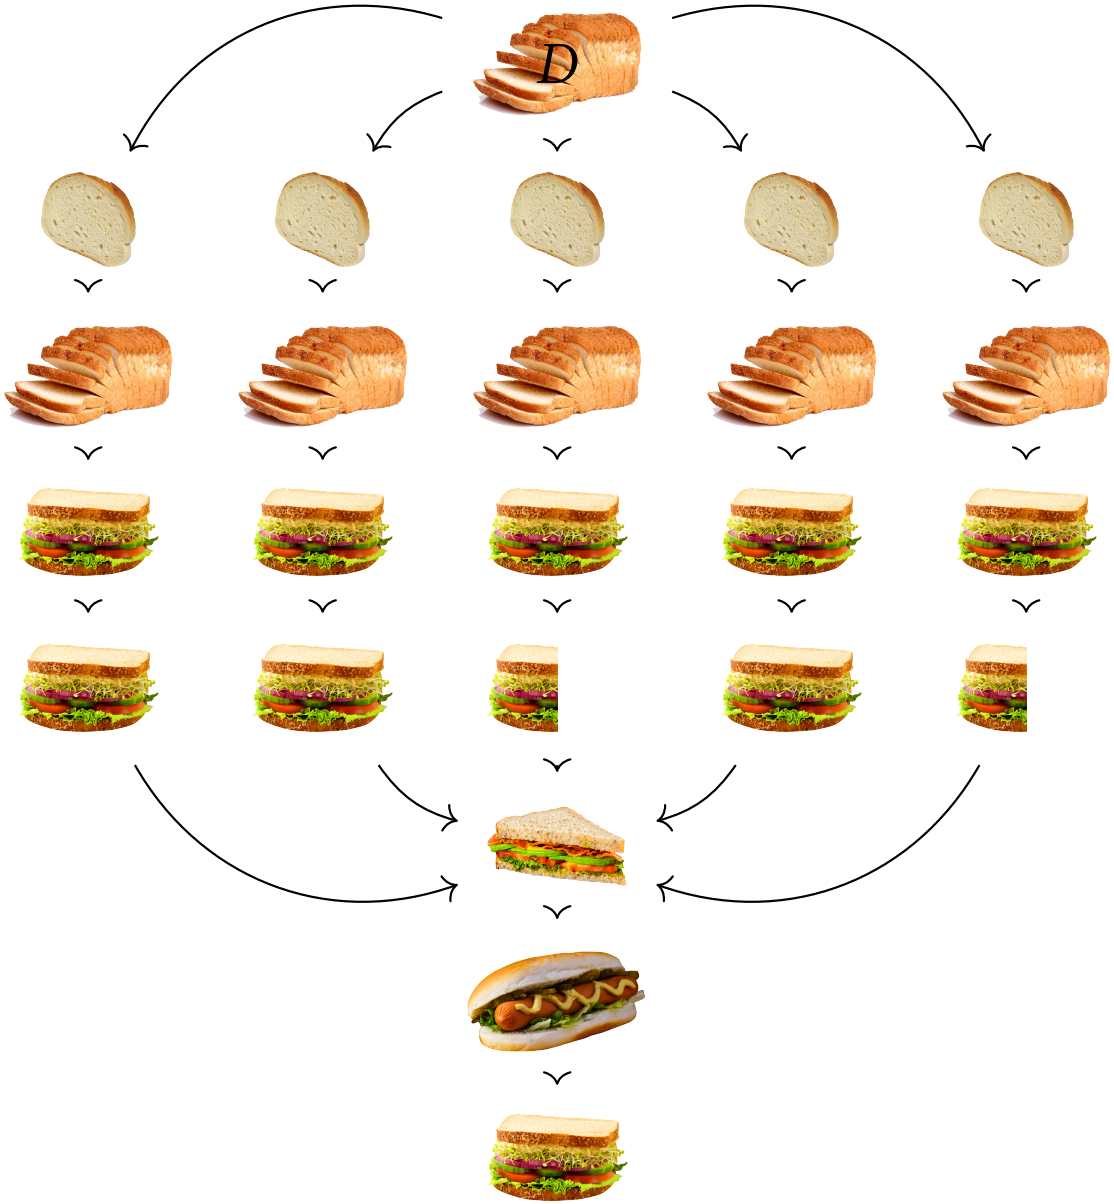

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

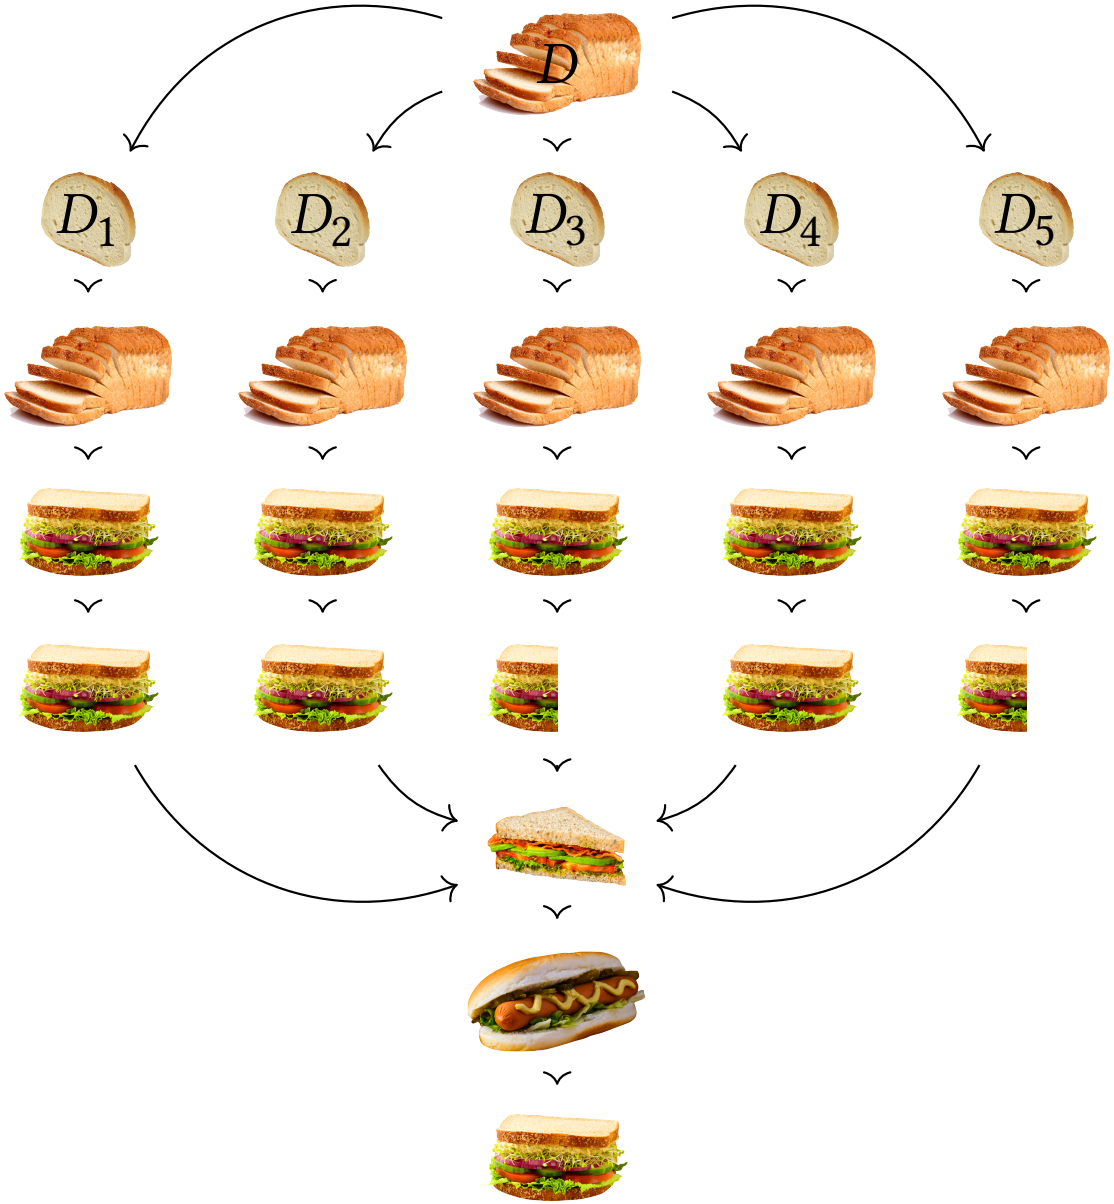

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

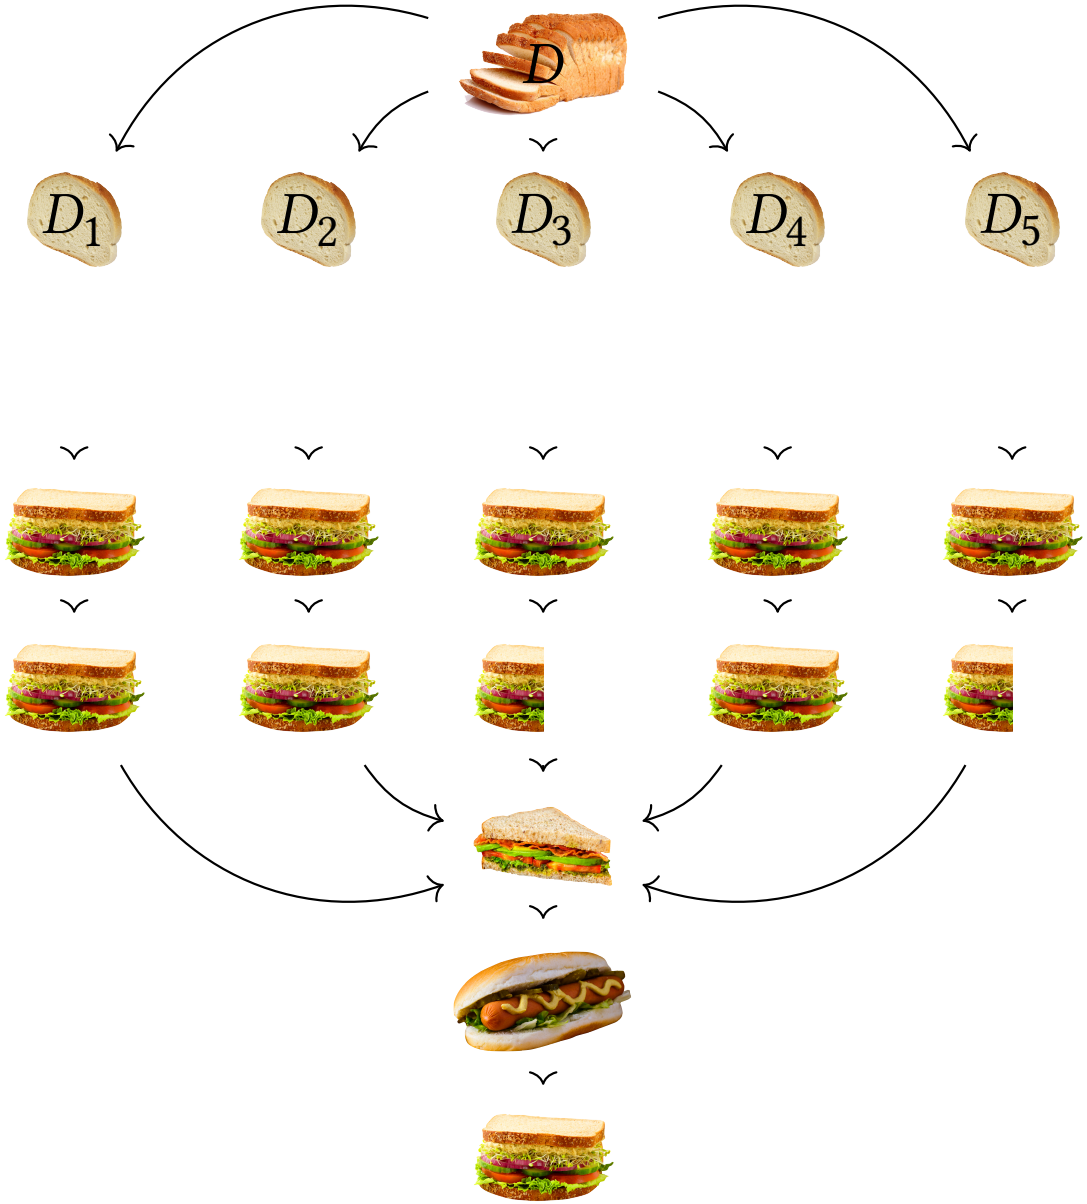

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

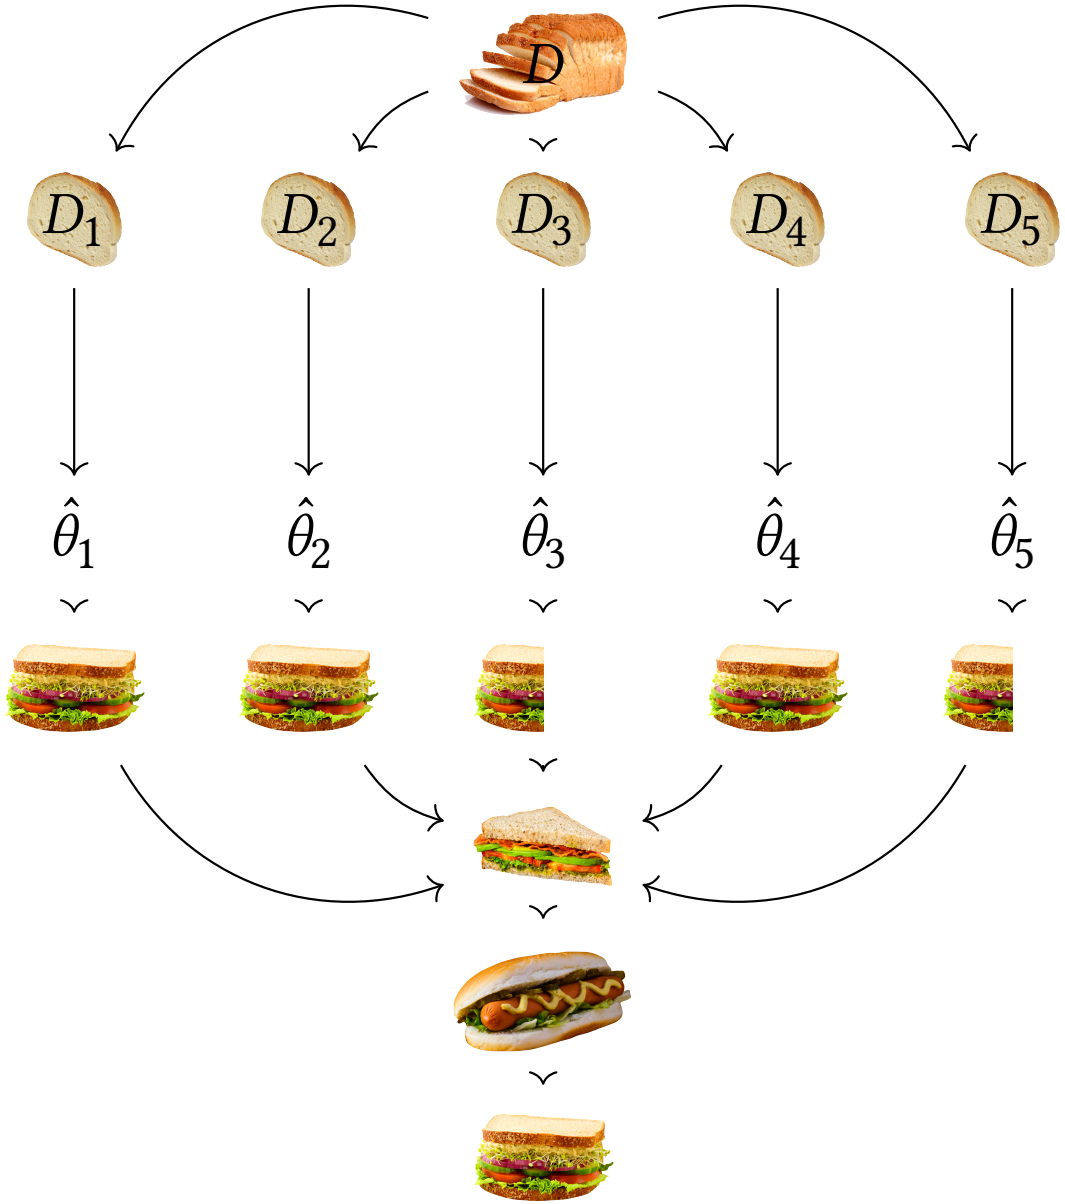

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

A Differentially Private Estimator

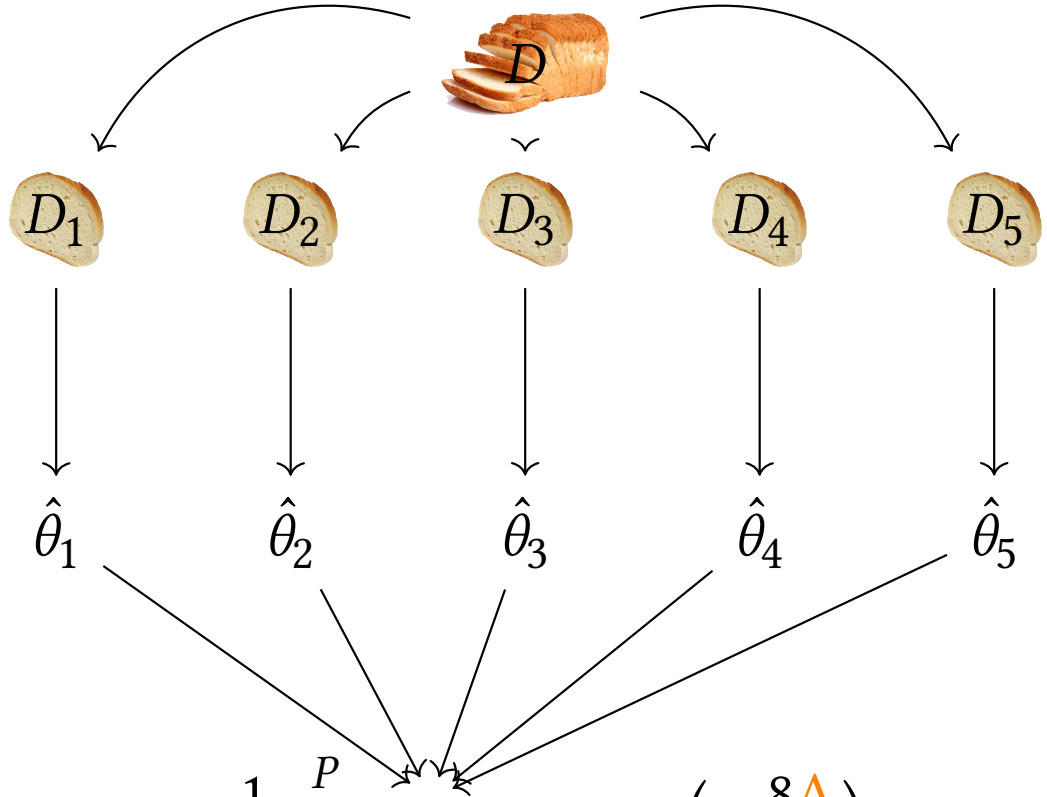

Private data

Partition

Bag of little bootstraps

Estimator

Censor

Average

Noise

Bias Correction
(& variance estimation)

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

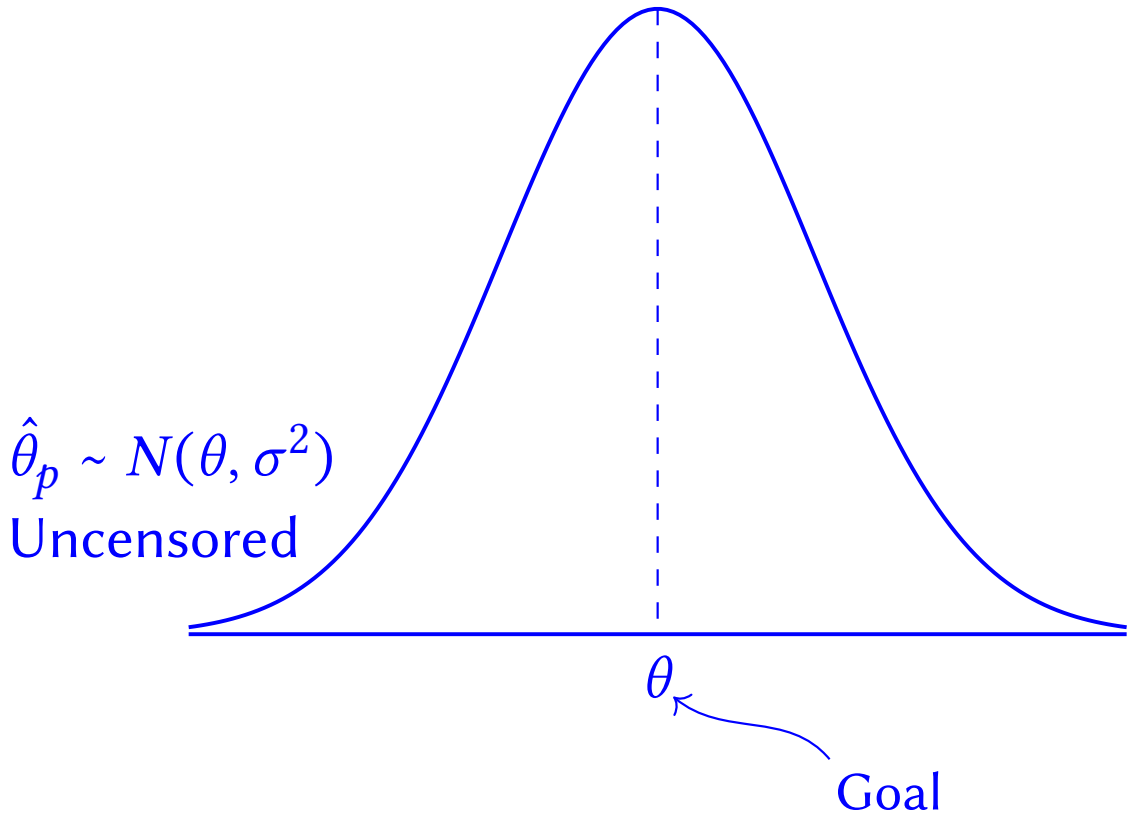

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

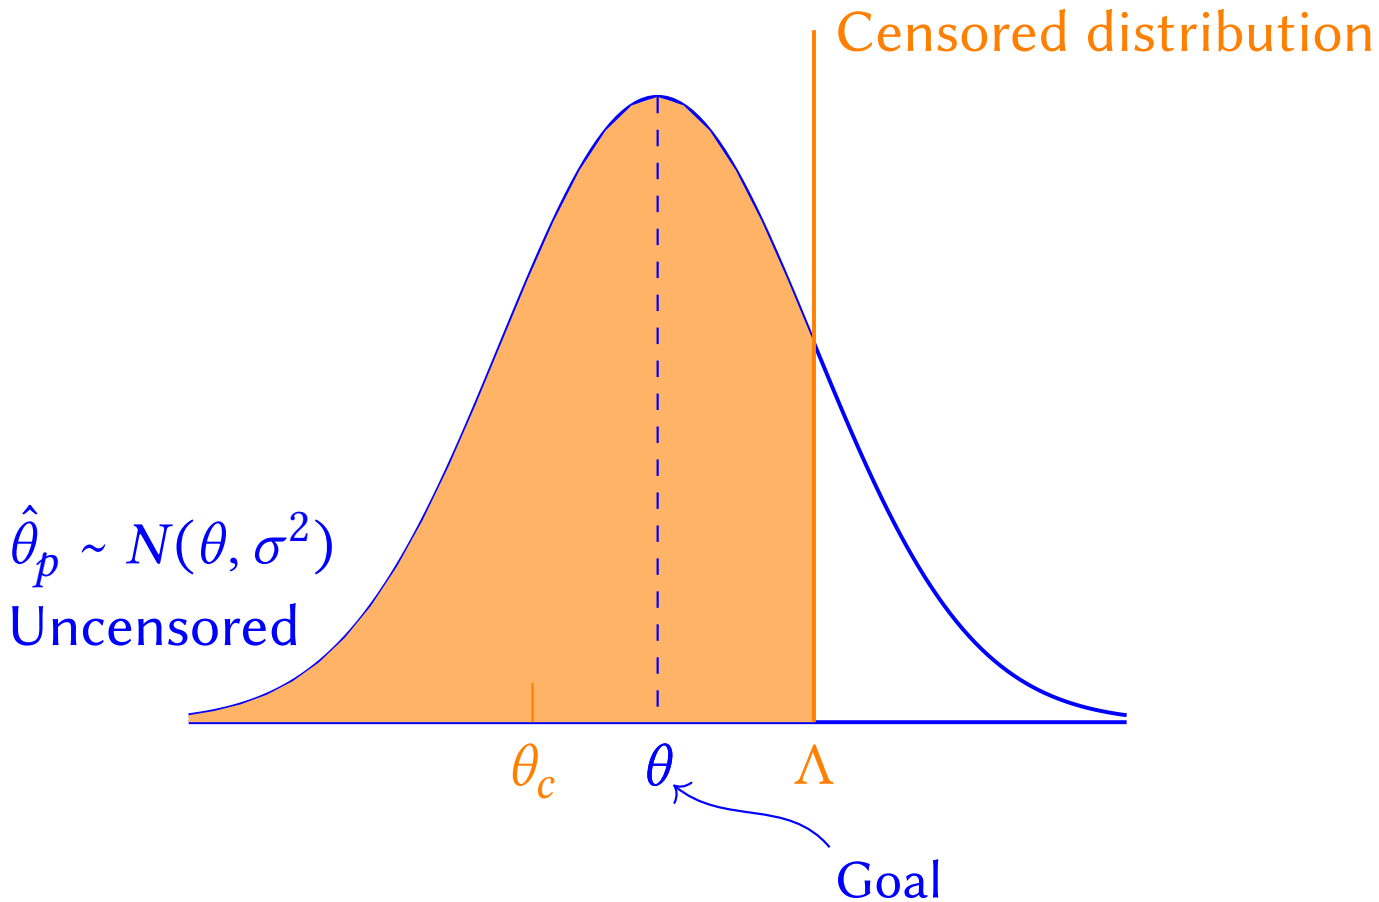

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

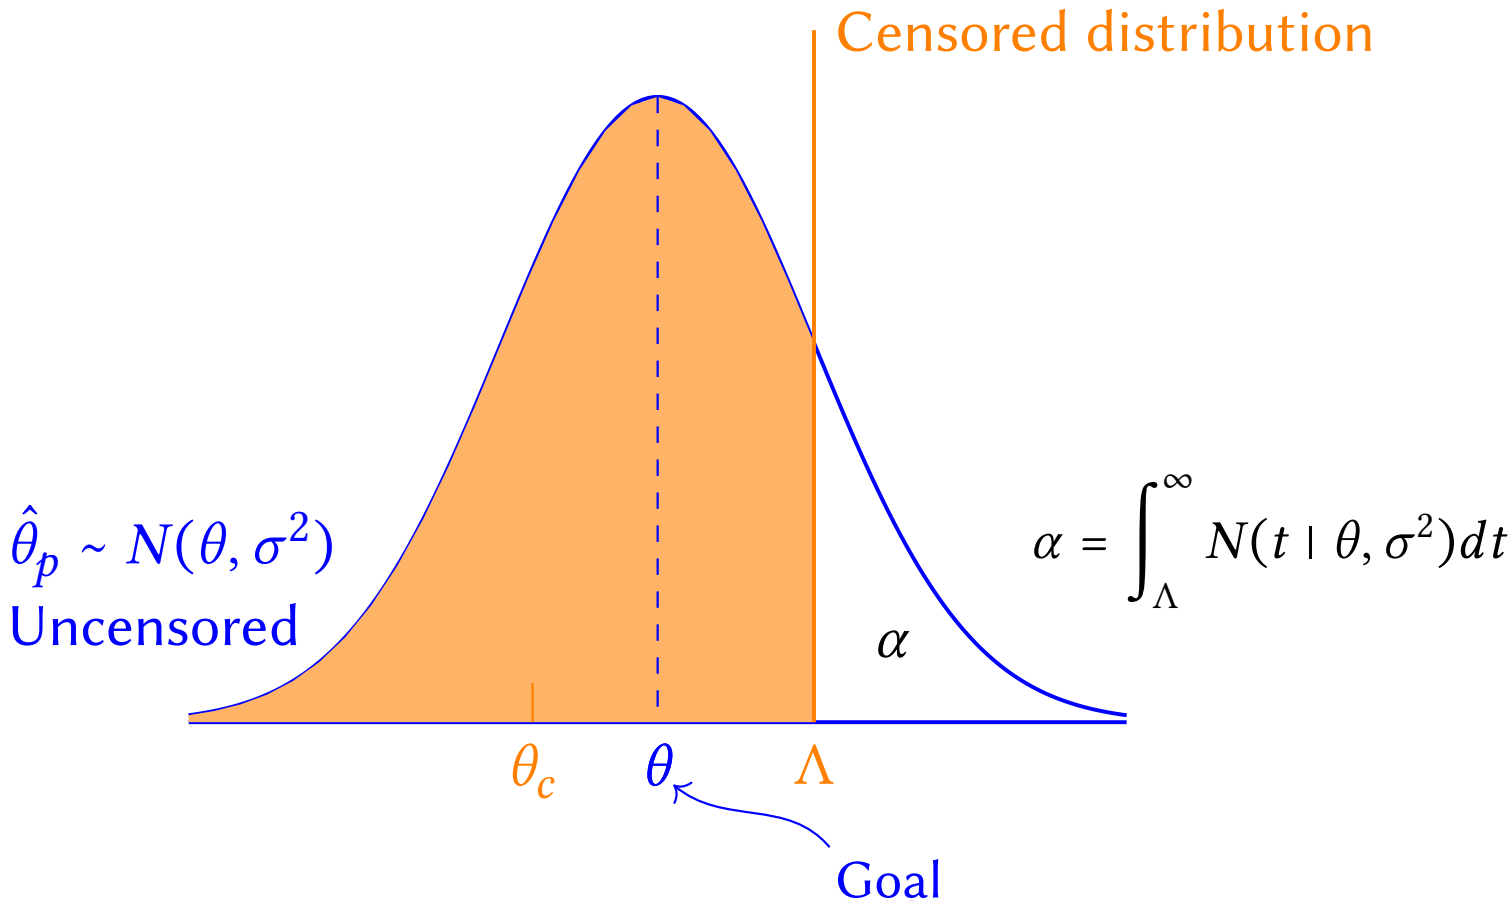

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

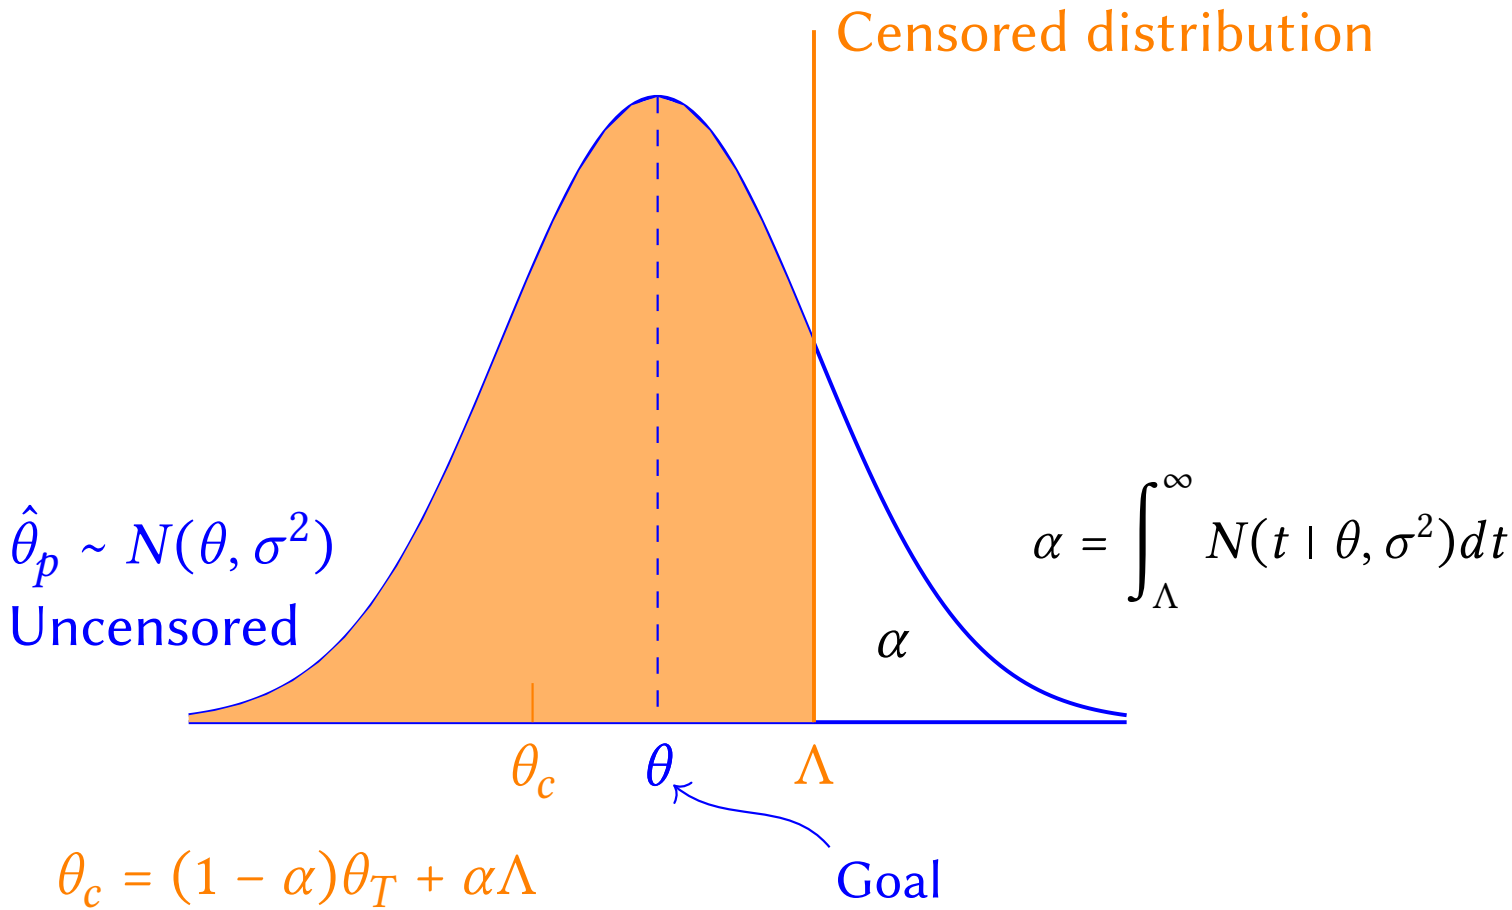

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

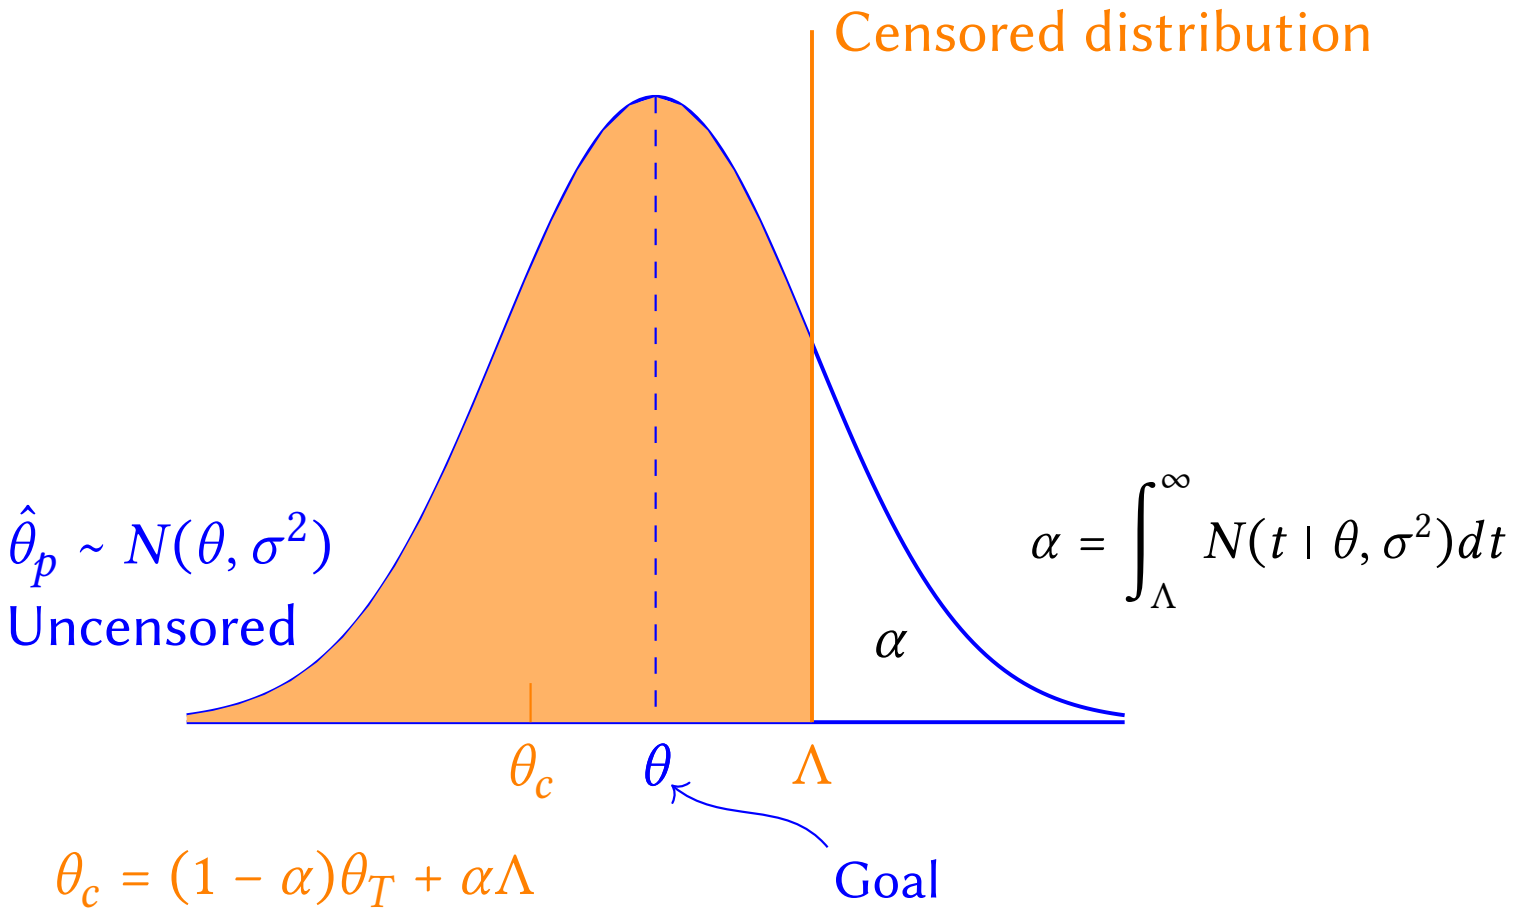

Equations: 2

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

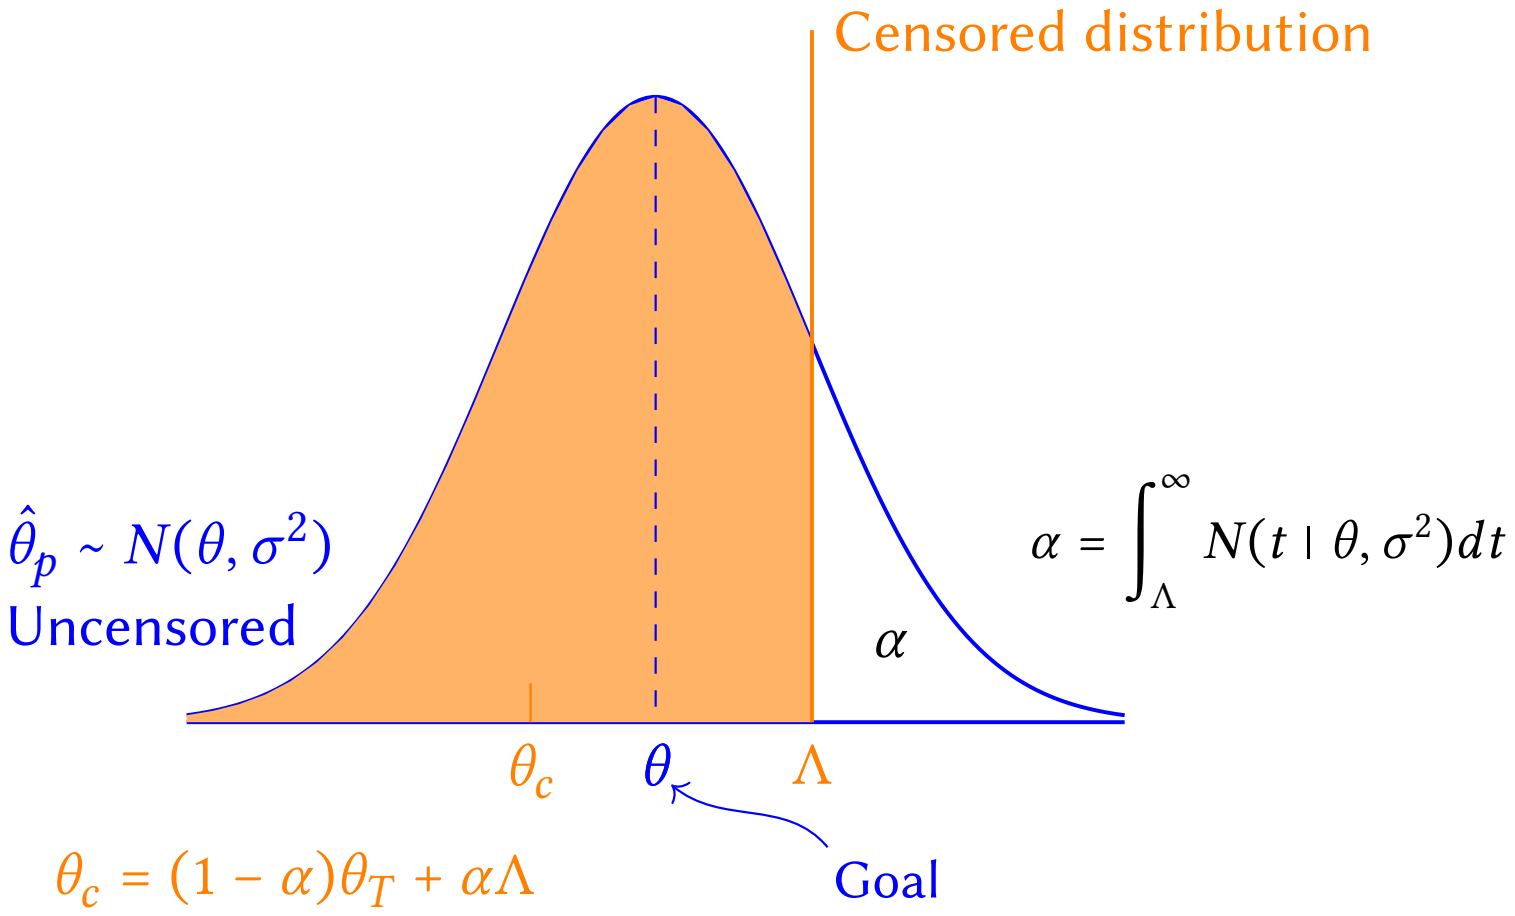

Equations: 2

Unknowns: θ, σ^2, α, θ_c

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

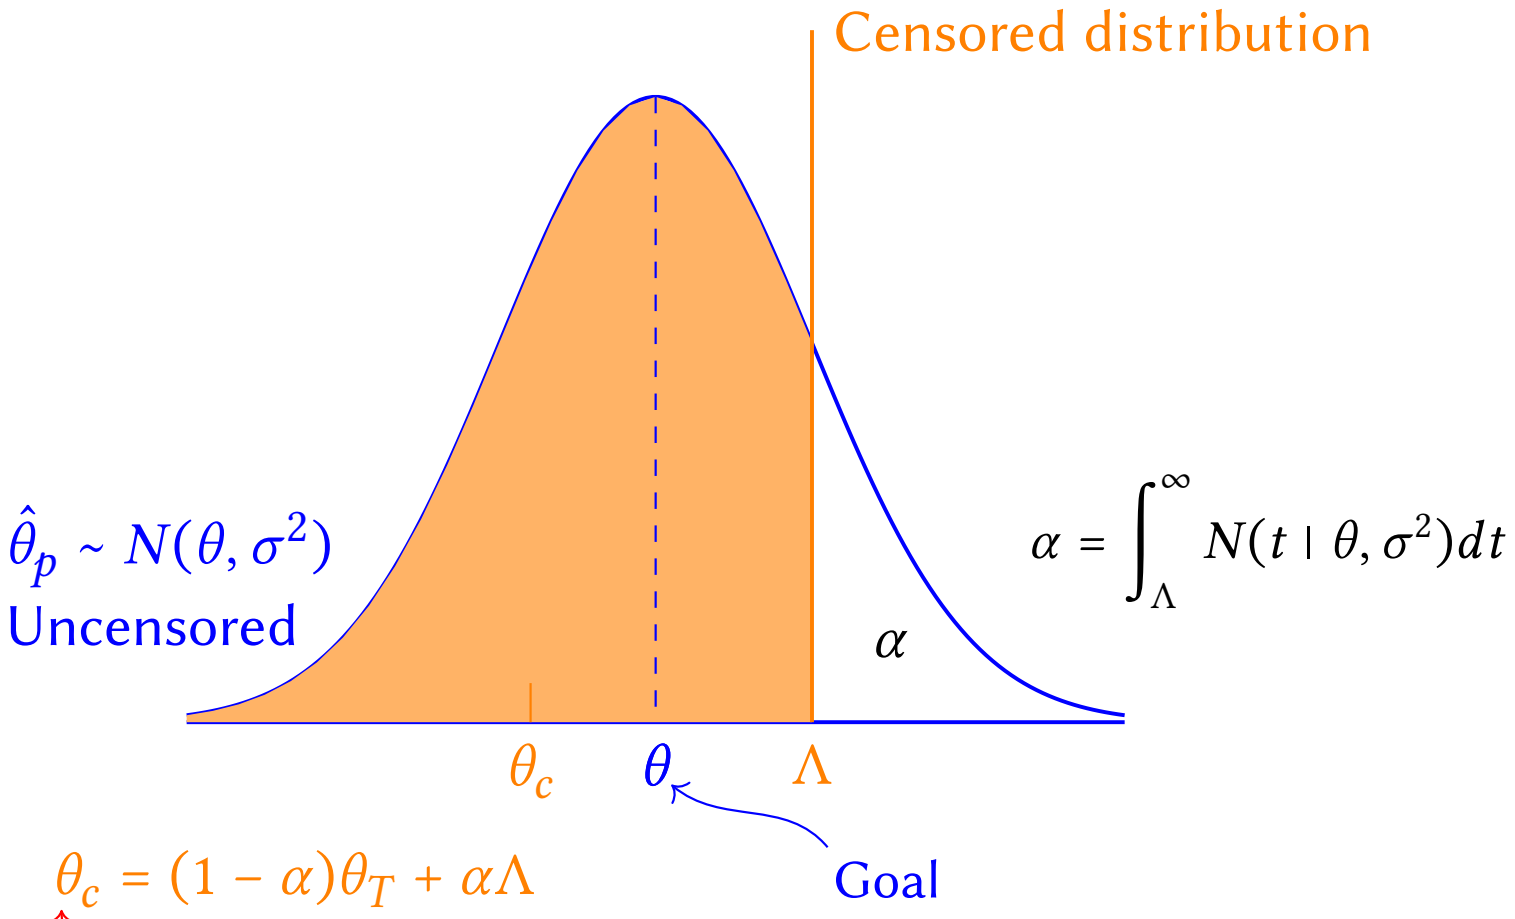

Equations: 2

Unknowns: θ, σ^2, α, ~~θ_c~~ ✗

Bias Correction of: $\hat{\theta}^{\text{dp}} = \frac{1}{P}\sum_{p=1}^{P} c(\hat{\theta}_p, \Lambda) + N\left(0, \frac{8\Lambda}{P\epsilon}\right)$ (Λ, P, ϵ known)

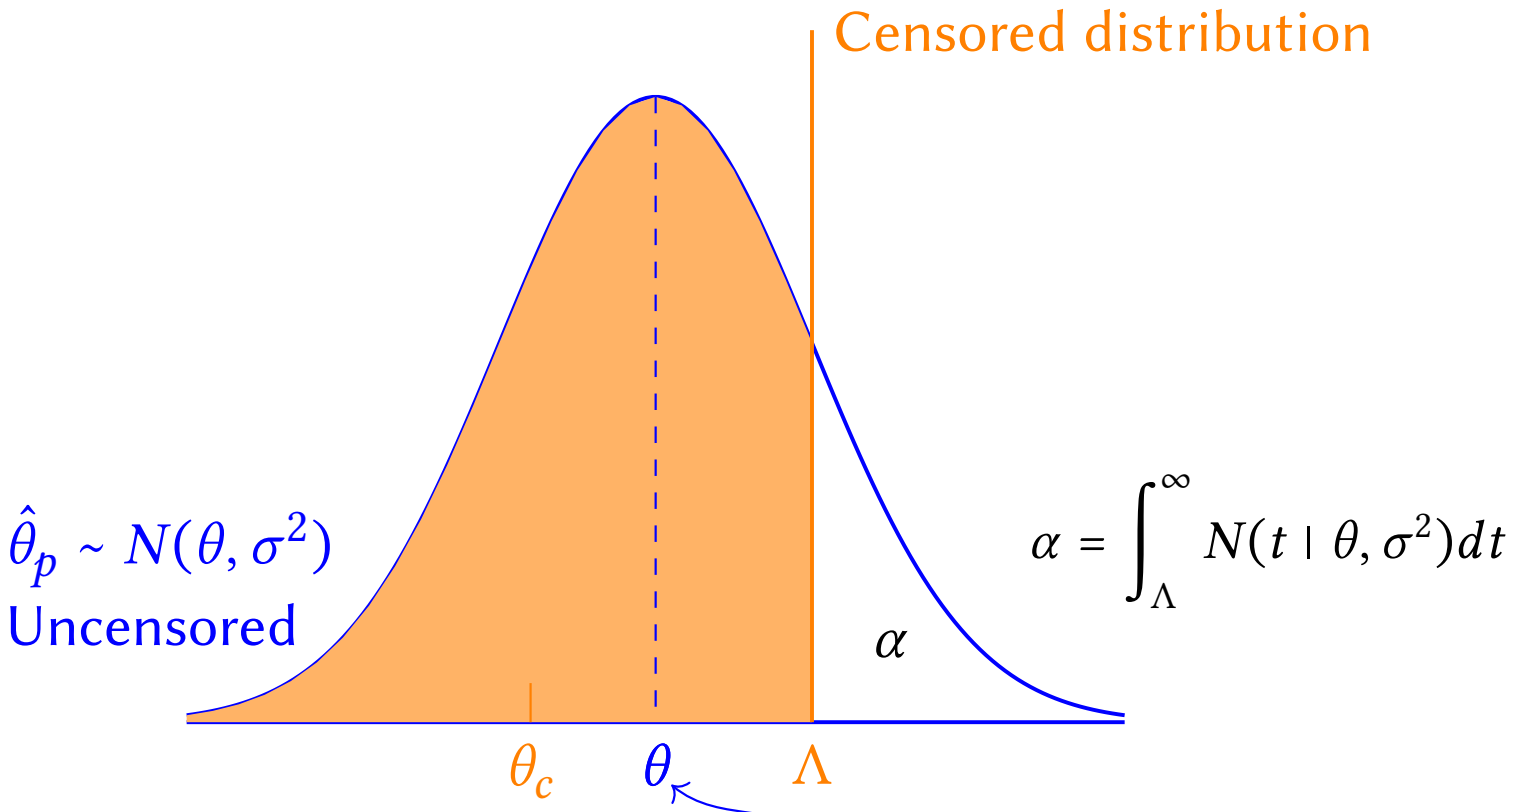

Equations: 2

Unknowns: θ, σ^2, ~~α~~, ~~θ_c~~

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:

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

Variance Estimation

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

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

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:

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

Functions of disclosed params

- Bias correct simulated params:

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

Variance Estimation

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

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

Functions of disclosed params

- Bias correct simulated params:

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

- Standard error: Standard deviation of $\tilde{\theta}^{\text{dp}}$ over simulations

Variance Estimation

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

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

Functions of disclosed params

- Bias correct simulated params:

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

- Standard error: Standard deviation of $\tilde{\theta}^{\text{dp}}$ over simulations
- Bias correction: reduces bias *and* variance:

$$E(\tilde{\theta}^{\text{dp}}) \approx \theta, \qquad V(\tilde{\theta}^{\text{dp}}) \lesssim V(\hat{\theta}^{\text{dp}})$$

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

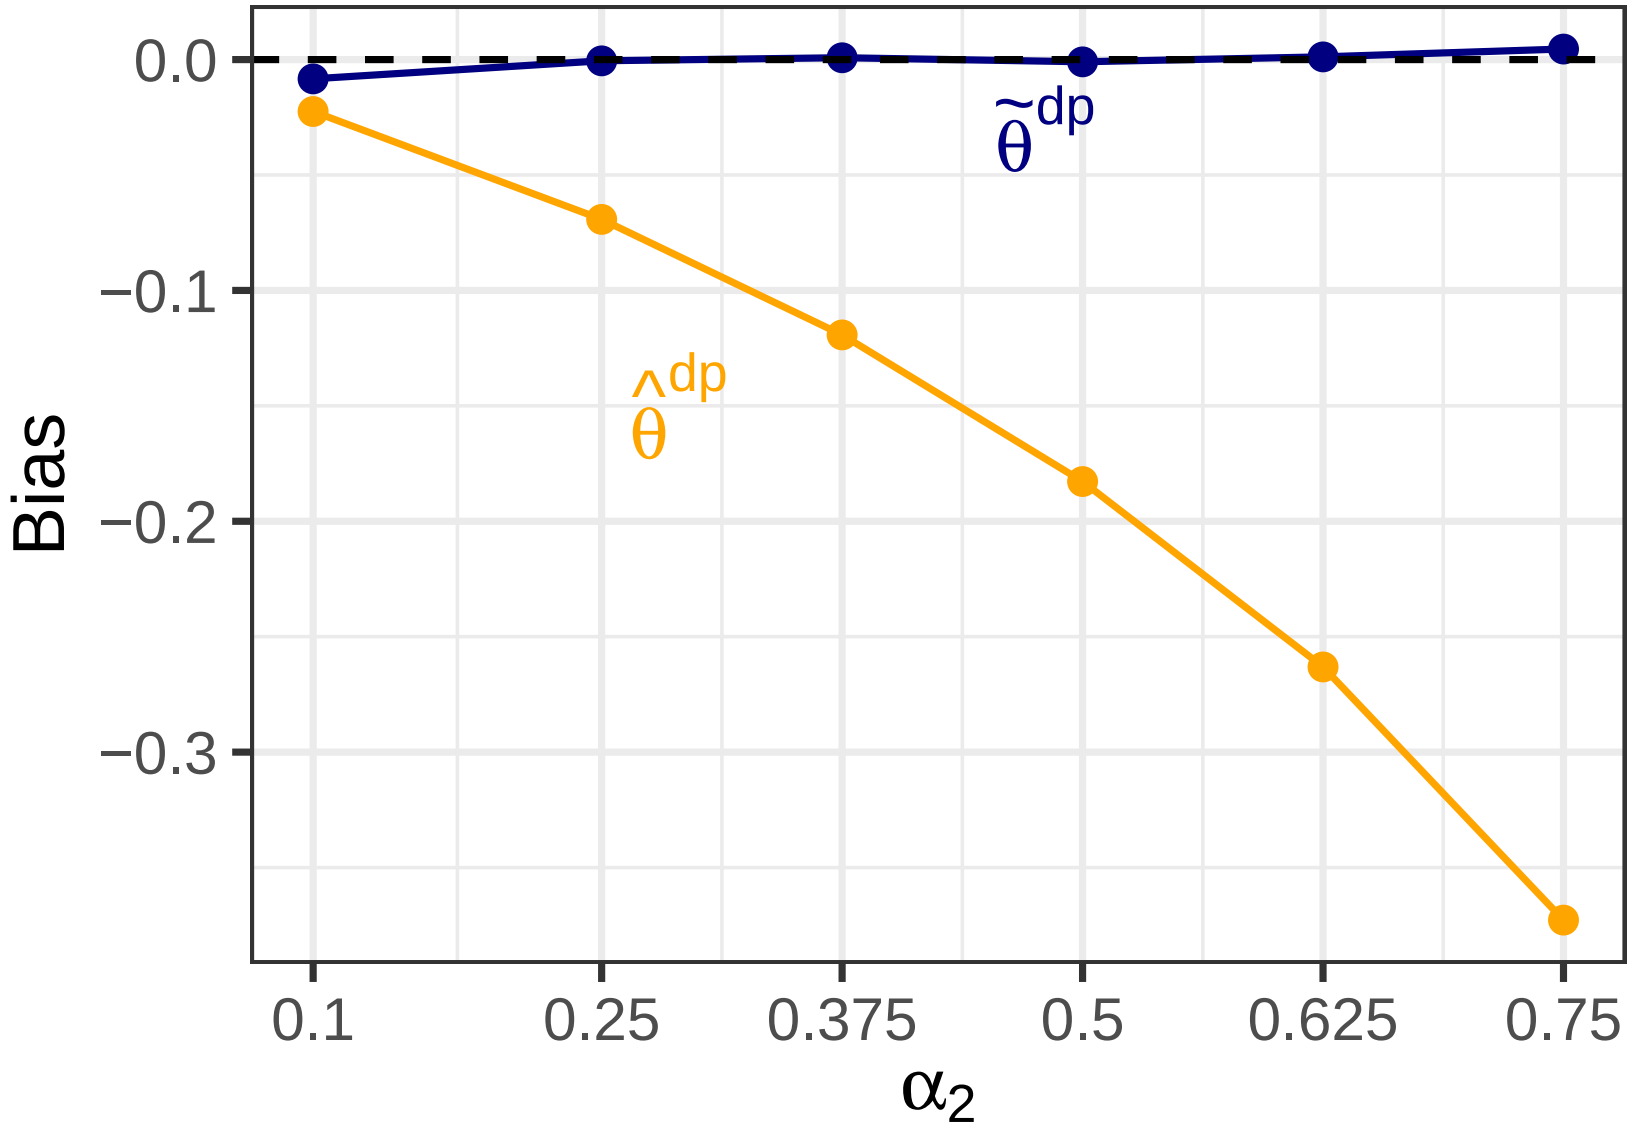

Simulations: Finite Sample Evaluation

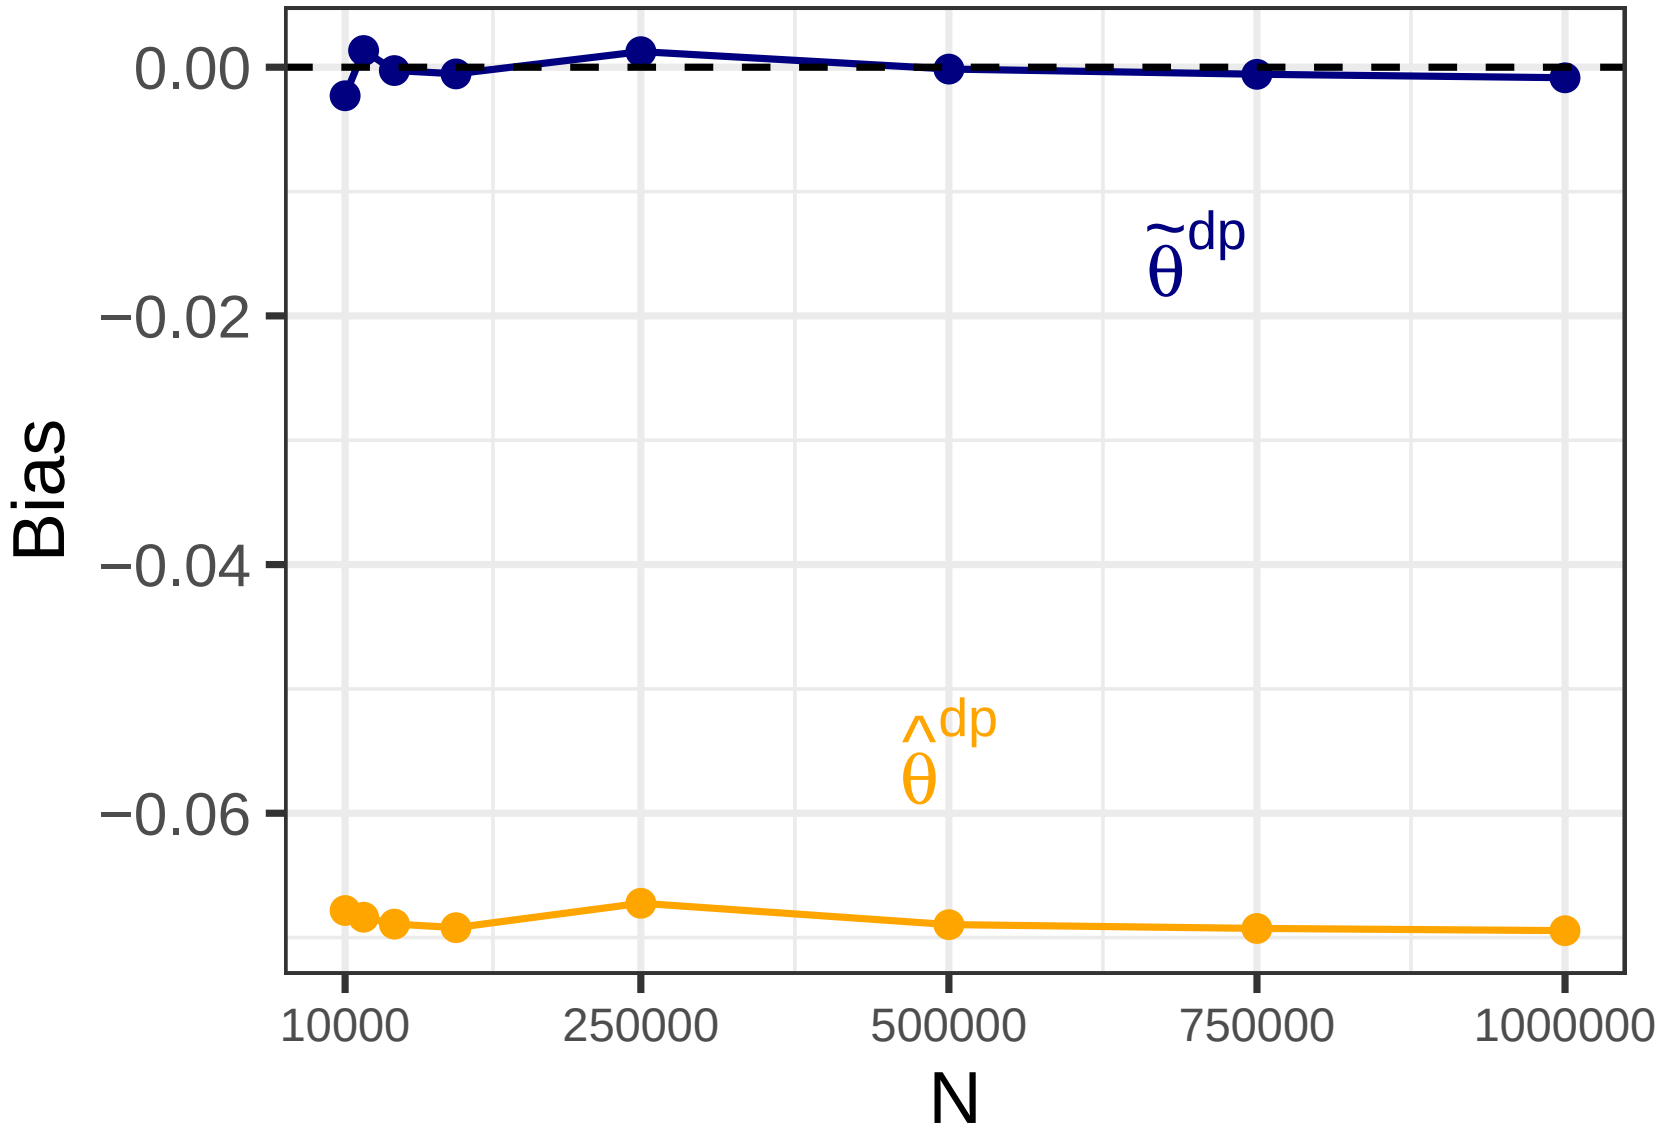

Simulations: Finite Sample Evaluation

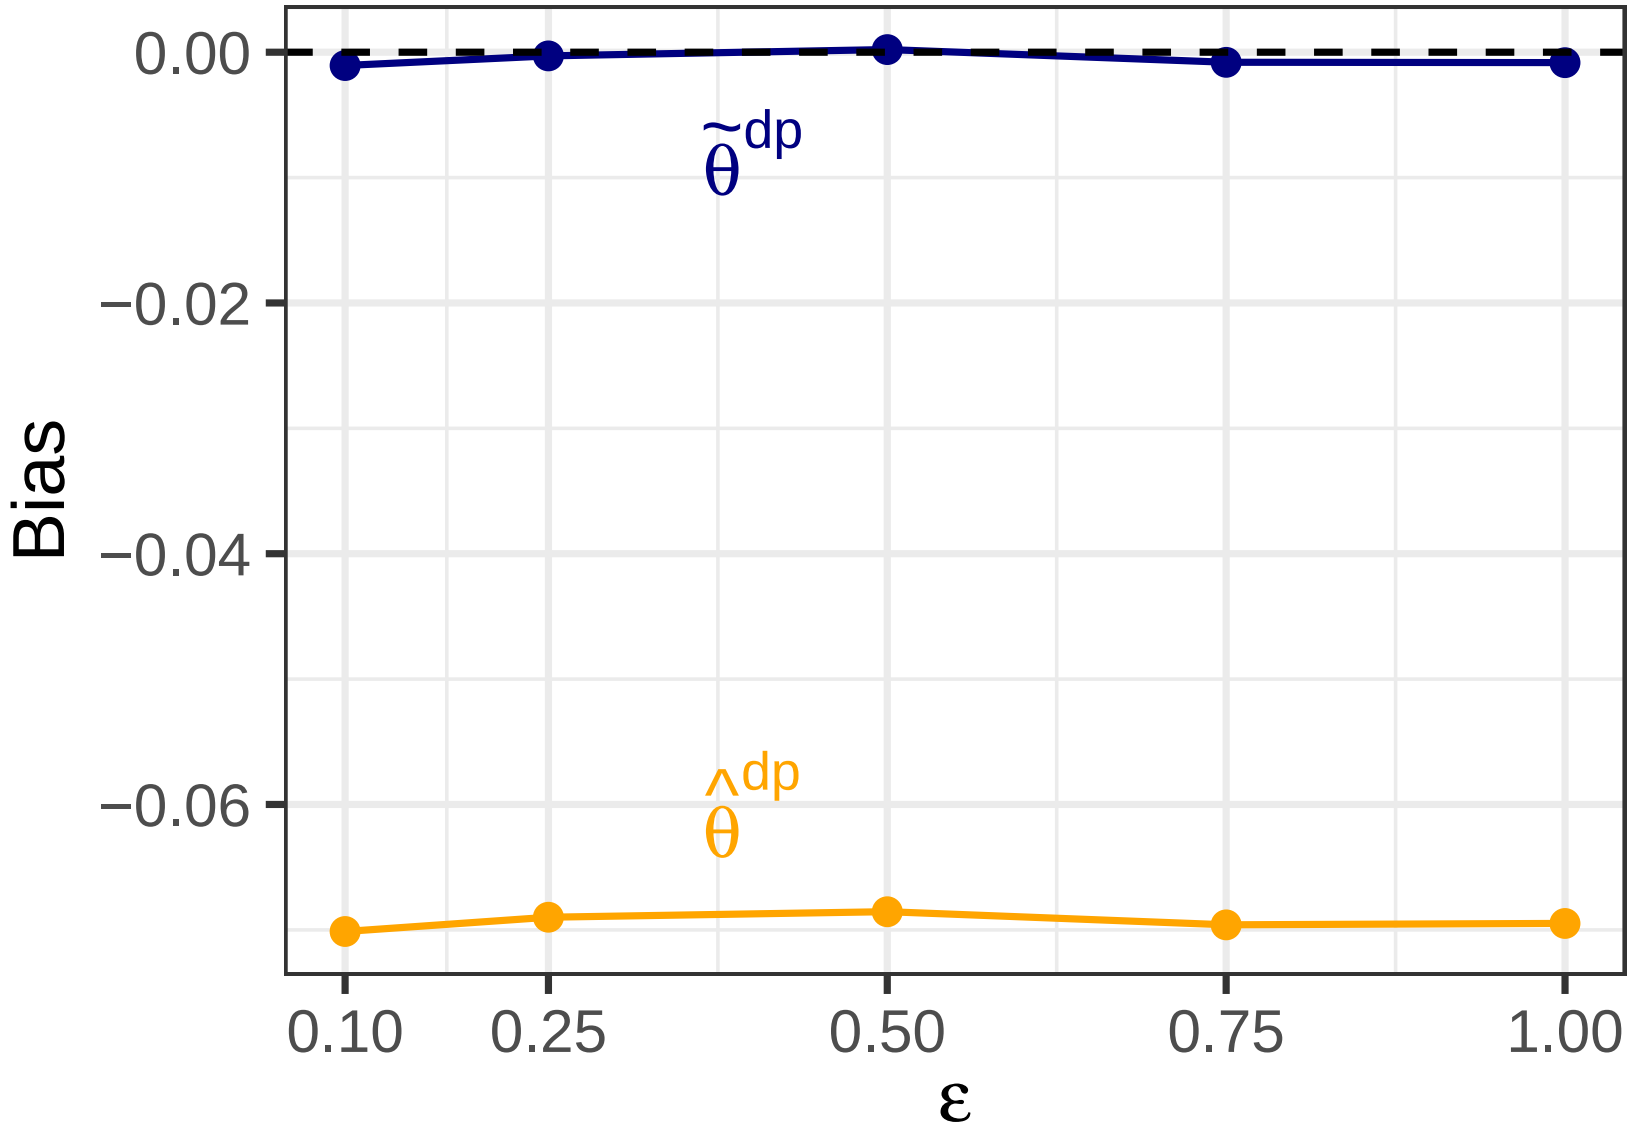

Simulations: Finite Sample Evaluation

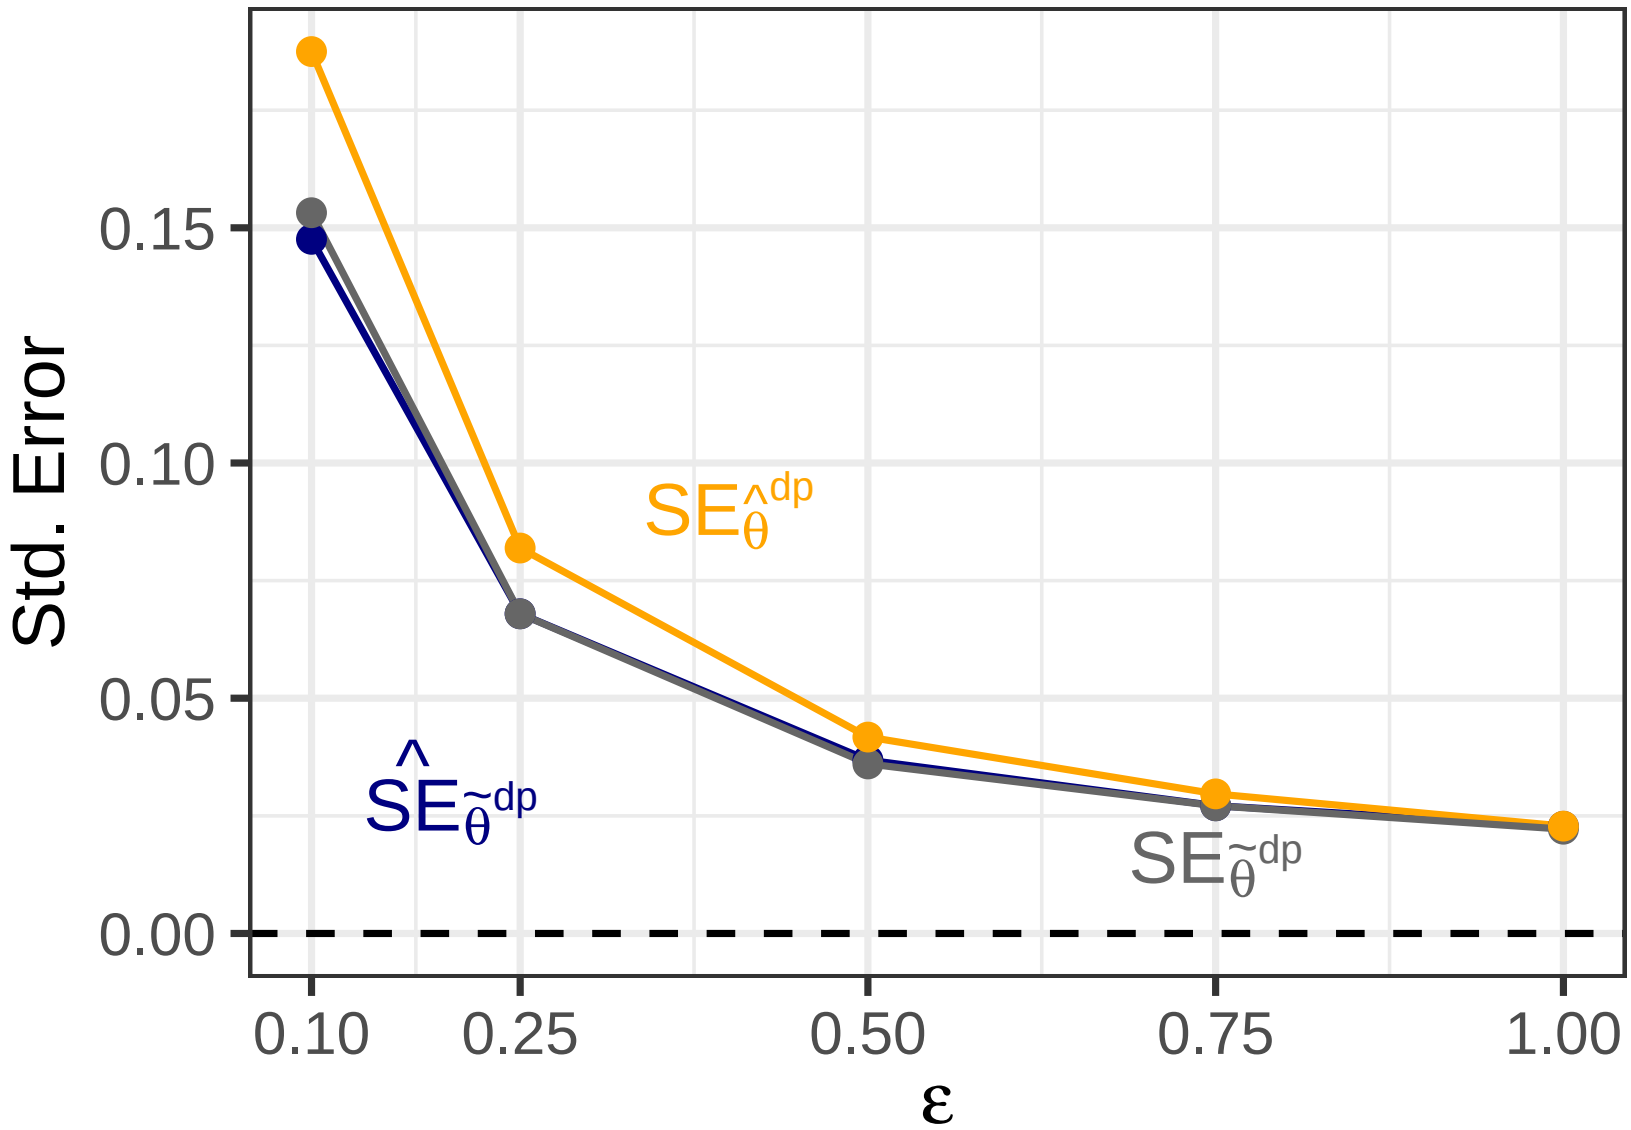

Simulations: Finite Sample Evaluation

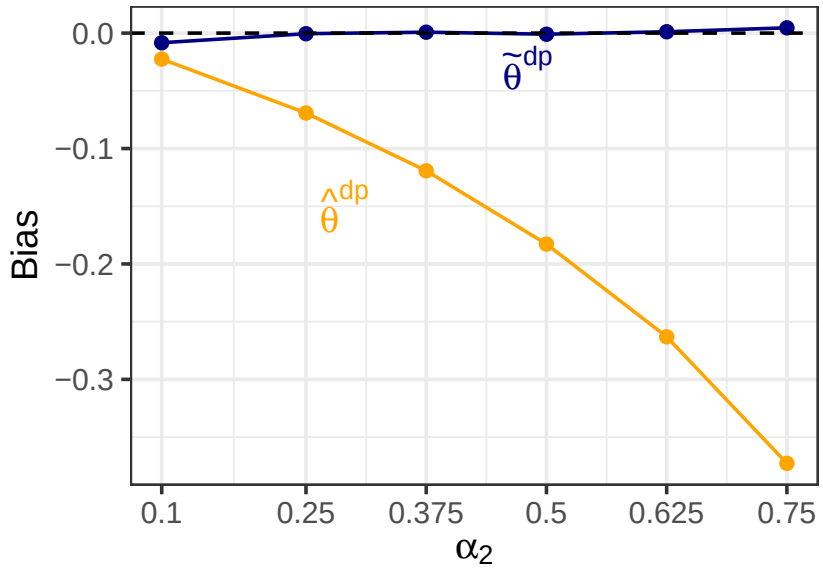

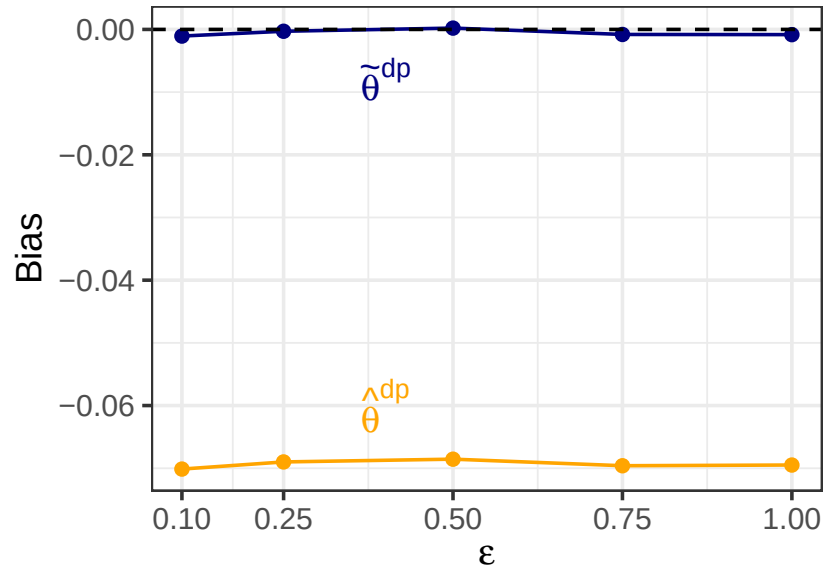

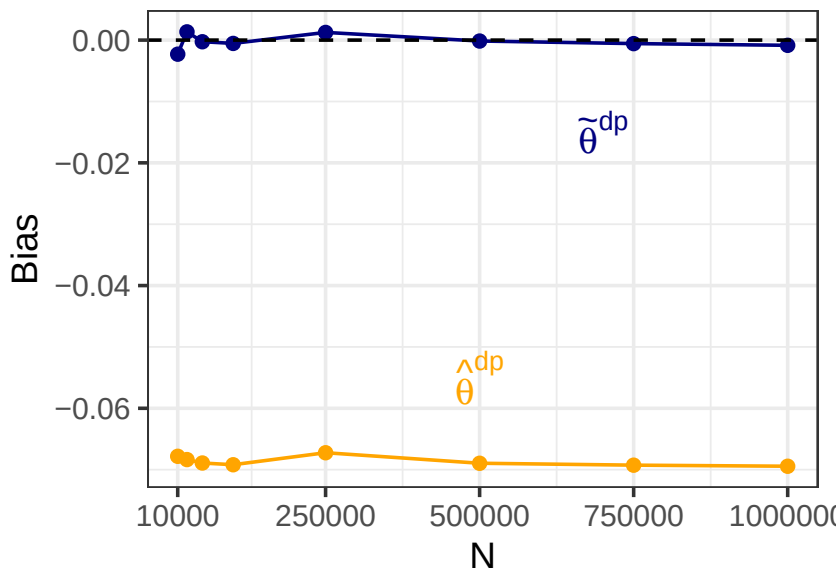

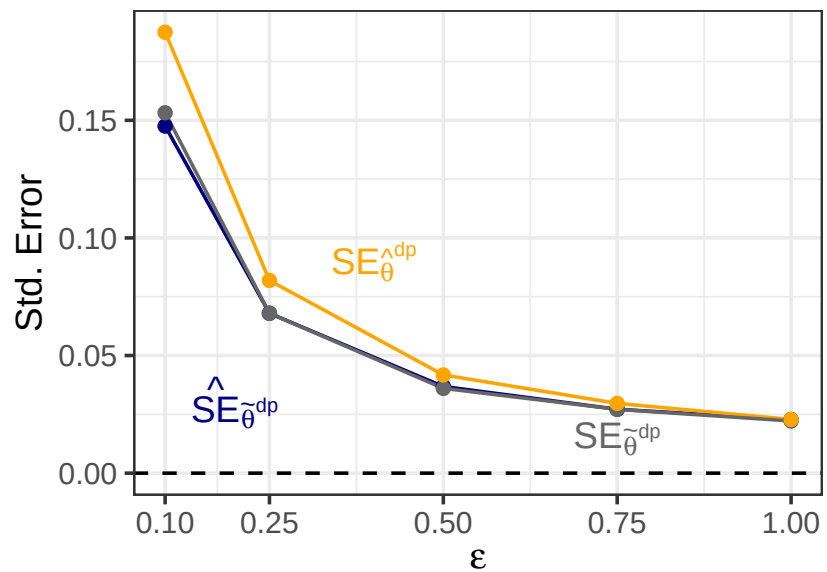

Concluding Remarks

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: not necessarily correct, but must have:

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: not necessarily correct, but must have:
  - known statistical properties & 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: not necessarily correct, but must have:
  - known statistical properties & 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: not necessarily correct, but must have:
  - known statistical properties & 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: not necessarily correct, but must have:
  - known statistical properties & valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased, 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: not necessarily correct, but must have:
  - known statistical properties & valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased, 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: not necessarily correct, but must have:
  - known statistical properties & valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased, 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: not necessarily correct, but must have:
  - known statistical properties & valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased, lower variance
  - Valid uncertainty estimates
  - Computationally efficient
  - Solves political problems technologically

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: not necessarily correct, but must have:
  - known statistical properties & valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased, lower variance
  - Valid uncertainty estimates
  - Computationally efficient
  - Solves political problems technologically
  - Implementations in progress:

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: not necessarily correct, but must have:
  - known statistical properties & valid uncertainty estimates
- Proposed algorithm
  - Generic: almost any statistical method or quantity of interest
  - Statistically unbiased, lower variance
  - Valid uncertainty estimates
  - Computationally efficient
  - Solves political problems technologically
  - Implementations in progress:
    - Facebook, Microsoft+Harvard/IQSS, OpenDP

For more information

Georgina-Evans.com

GaryKing.org

MegSchwenzfeier.com

bit.ly/AbhradeepThakurta

Paper, software, slides, video: GaryKing.org/dp