But Can You Use It?

Design Recommendations for Differentially Private Validation Servers

Joshua Snoke
RAND

FCSM 2024
October 24, 2024

In collaboration with Liudas Panavas (Northeastern), Claire Bowen, Erika Tyagi, and Aaron Williams (Urban)

What is a validation server?

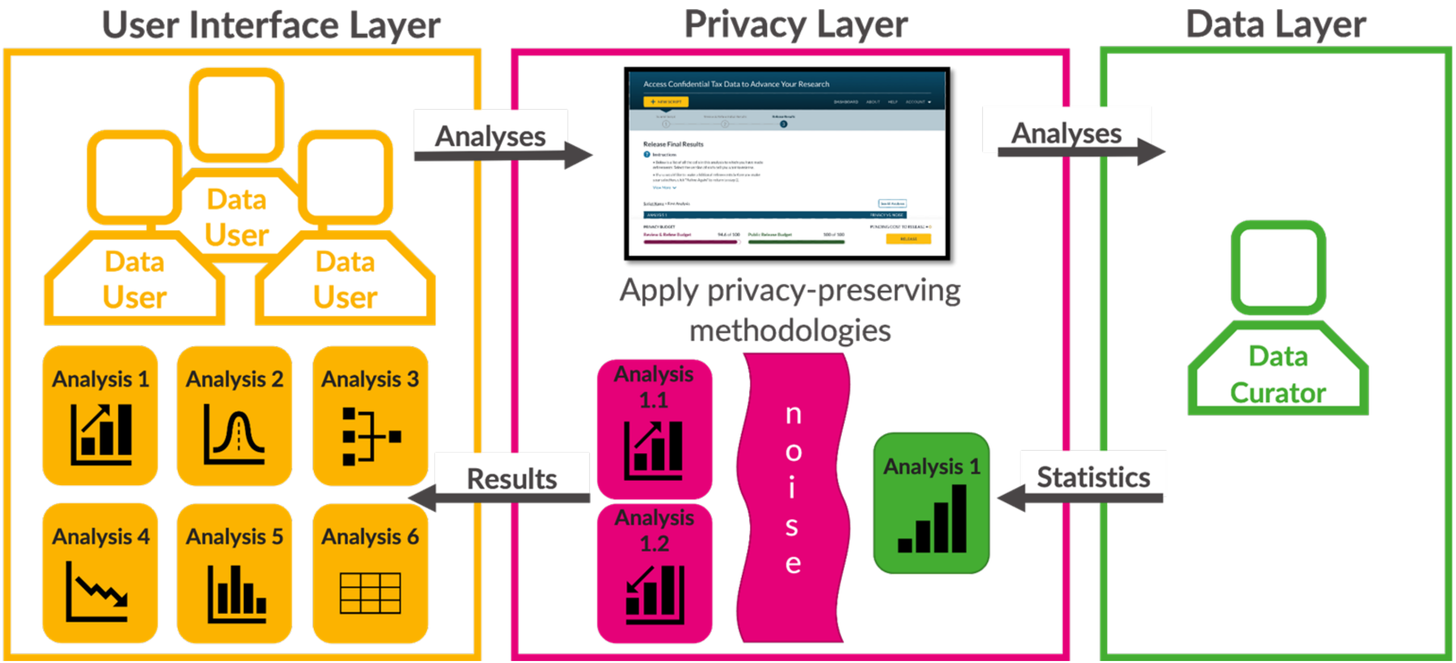

Focusing on a particular type of validation server

- Assuming the framework of differential privacy as a starting point
- Our target audience:
  - User of federal statistical data products
  - Uses traditional statistical methods
  - Wants to inform public policy
- Goal is to assess the *practical application* of DP validation servers
- Inspired by the Safe Data Technologies work with broader implications

Why do we want to use validation servers?

- More flexible than fixed releases
  - Fixed releases require determining *a priori* what statistics to preserve
- For social science/public policy researchers:
  - Significant limitations and skepticism of public data*
  - Interactive setting offers the opportunity for targeted analyses
  - But crucially, the results need to enable valid statistical inference to provide value
- Federal statistical systems are investing in a tiered approach [NASEM24]

*When it is transparent how the data are noisy

So what is the issue?

In practice, validation servers are *hard to use*

Why validation servers are hard in practice

- Automation (even partial) requires strong privacy protections
  - Differential privacy is a natural solution
- The framework of differential privacy presents issues in practice
  - Misalignment between mechanism design and statistical methods
  - Unrealistic assumptions about users

Issues: mechanisms unaligned with statistical methods

- Mechanisms designed for prediction problems or simple (e.g., count) queries
  - Methods have lagged for regression-based inferential methods
- Mechanisms assume well-behaved data generating processes
  - I.e., symmetric and gaussian
  - Theoretical guarantees do not hold under common issues such as skewness
- Mechanisms do not provide uncertainty estimates
  - Often assumed without practical means of achieving

Misalignment results in poor empirical performance

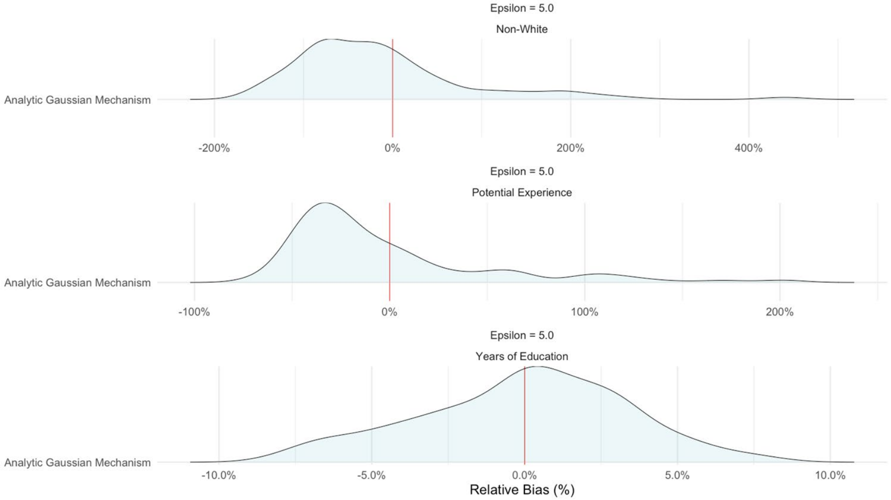

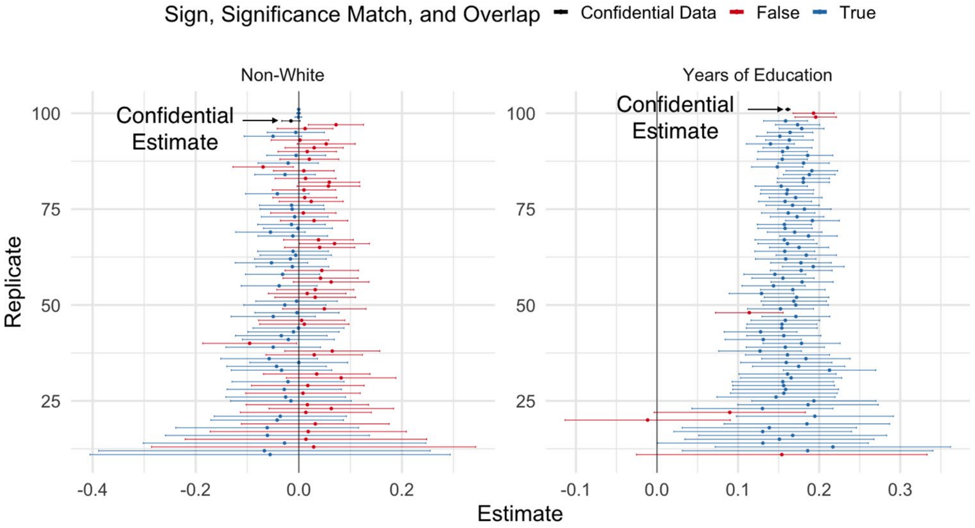

Source: [BWSB24]

Issues: limited ability to perform exploratory data analysis

- DP requires specifying the function and sensitivity without observing the data
- But…
    - Domain of the data and range of the outputs is often not known
    - Lack of desirable means of error handling
    - Very little work exists on applying DP to common EDA tools
    - Induces undesirable tradeoff between correctly specifying the function and the amount of noise (or privacy loss) [SBWB2024]

Issues: setting the privacy parameters

- DP assumes privacy parameters can be set a priori
- But…
  - Parameters lack consensus interpretation [WZ10, Kea22, Nea23]
  - Privacy parameters do not have absolute interpretations [SS23]
  - Users will have a threshold for finding the data useful [WSBB24]

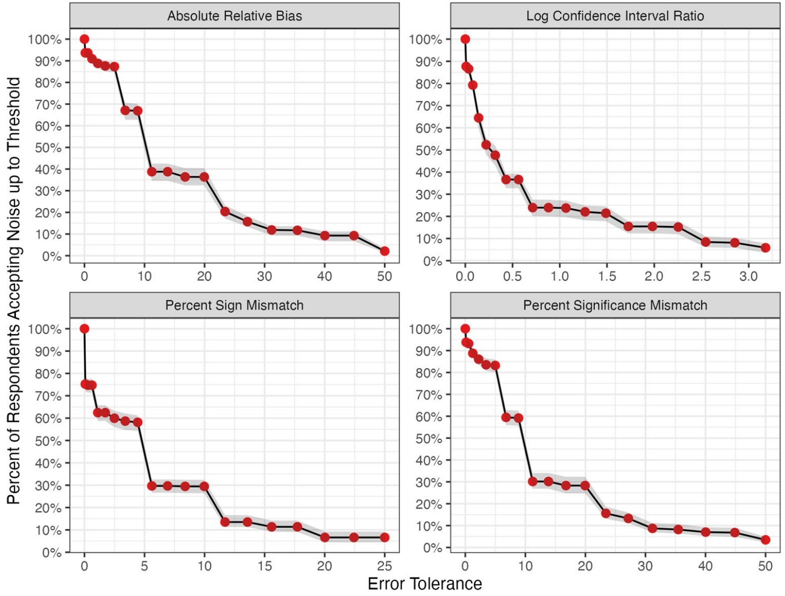

Issues: finite privacy budget constraints

- DP requires a finite privacy budget
- But…
  - What happens to the system when the budget runs out? [D23]
  - Efficiently allocating the privacy budget assumes knowledge of all queries a priori

IMMA LET YOU FINISH YOUR ANALYSIS
BUT YOUR PRIVACY BUDGET IS OUT
imgflip.com

Is a DP validation server possible?

- Pointing out incompatibilities, not making a value statement
    - Perhaps a DP system could be build as theorized
    - But it requires a different research environment than what currently exists
- In the real-world, all systems require some compromises:
    - Explore ideas for increasing practical usability
    - Determine how privacy relaxations can be applied
    - See also [CS24, SS22]

Let’s be clear about our design principles

- A validation server should incorporate the following principles:
  - Privacy assurance
  - Statistical utility
  - System usability

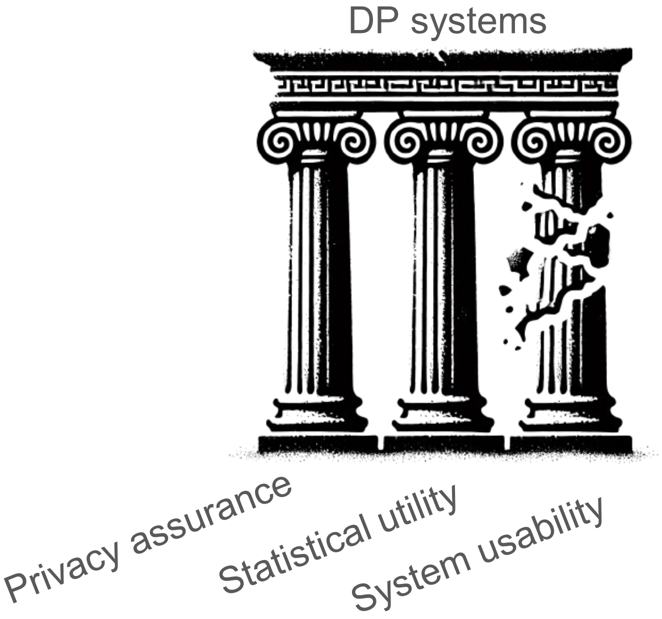

Design principle: privacy assurance

- Accounting
  - Quantify and track cumulative privacy loss
- Transparency
  - Articulate what is and *isn't* covered by our privacy mechanisms
- Threat modeling
  - Meaningfully interpret the privacy risks

Design principle: statistical utility

- Capacity
  - Relevance of the possible queries
- Coverage
  - Ability to make valid statistical inferences
- Power
  - Minimizing the loss in effective sample size

Design principle: system usability

- Design
  - How easy and efficient it is to interact with the system
- Knowledge
  - Required expertise of users
- Applicability
  - How well the system's outputs meet the specified user tasks

How do we move towards a practical validation server?

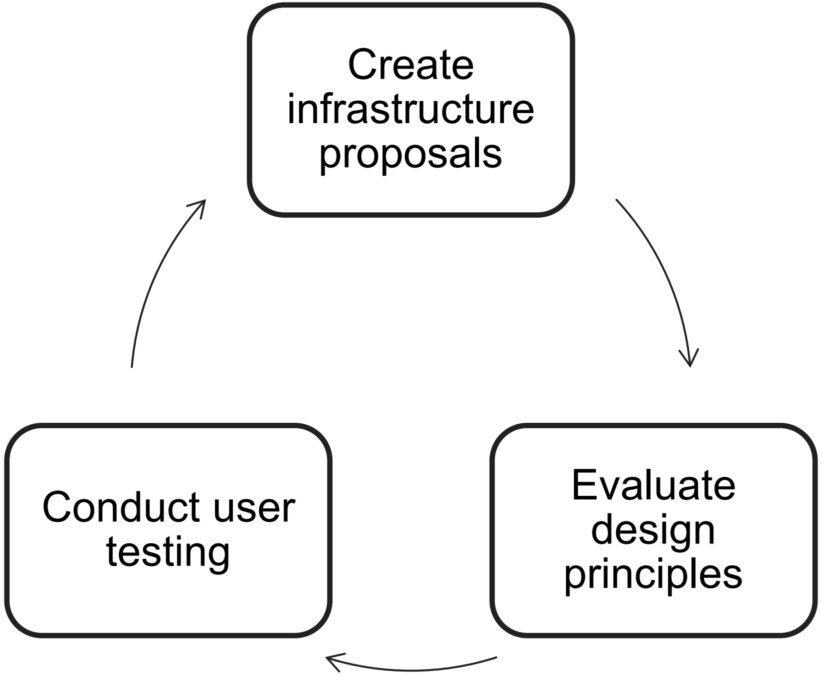

Create infrastructure proposals

Evaluate design principles

Conduct user testing

Example of putting this idea into practice

- **Proposal:** *provide synthetic data alongside validation server*
- Evaluate impact on design principles
  - Privacy: additional privacy loss
  - Statistical Utility: understanding impact of additional noise
  - Usability: enable EDA, help budget setting
- User testing
  - Does this improve users’ ability to correctly specify their queries?
  - Does this improve users’ ability to correctly specify their privacy budget?
  - What characteristics do the synthetic data need to have?

Closing thoughts

- Design recommendations help ensure we build systems that can be used
- Theory has far outpaced practice
  - We need to work out the barriers to practical use
- Collaboration is key
  - Privacy engineers, statisticians, and user-focused researchers all have a role

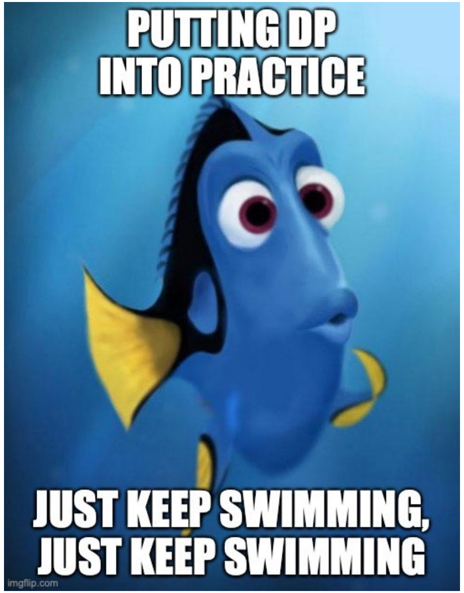

Closing thoughts

- Design recommendations help ensure we build systems that can be used
- Theory has far outpaced practice
  - We need to work out the barriers to practical use
- Collaboration is key
  - Privacy engineers, statisticians, and user-focused researchers all have a role

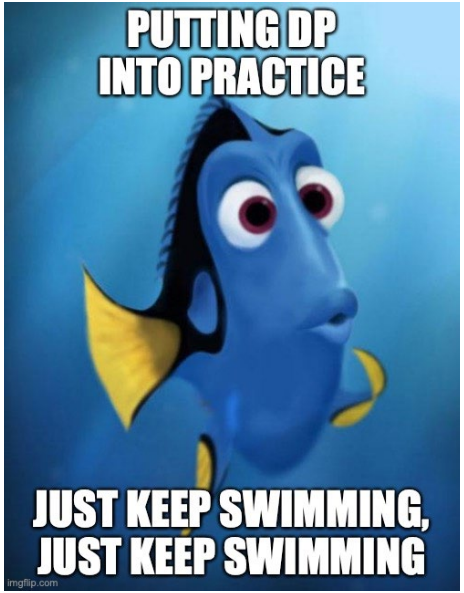

Thank you!
Comments/complaints/criticisms: jsnoke@rand.org