What Are the Chances?

Explaining the Epsilon Parameter in Differential Privacy

Priyanka Nanayakkara
Northwestern University

Mary Anne Smart
University of California San Diego

Rachel Cummings*
Columbia University

Gabriel Kaptchuk*
Boston University

Elissa M. Redmiles*
Max Planck Institute for Software Systems

** equal advising*

USENIX Security '23 | Anaheim, California


Should I share
my data?

Desfontaines, D (2023). A list of real-world uses of differential privacy

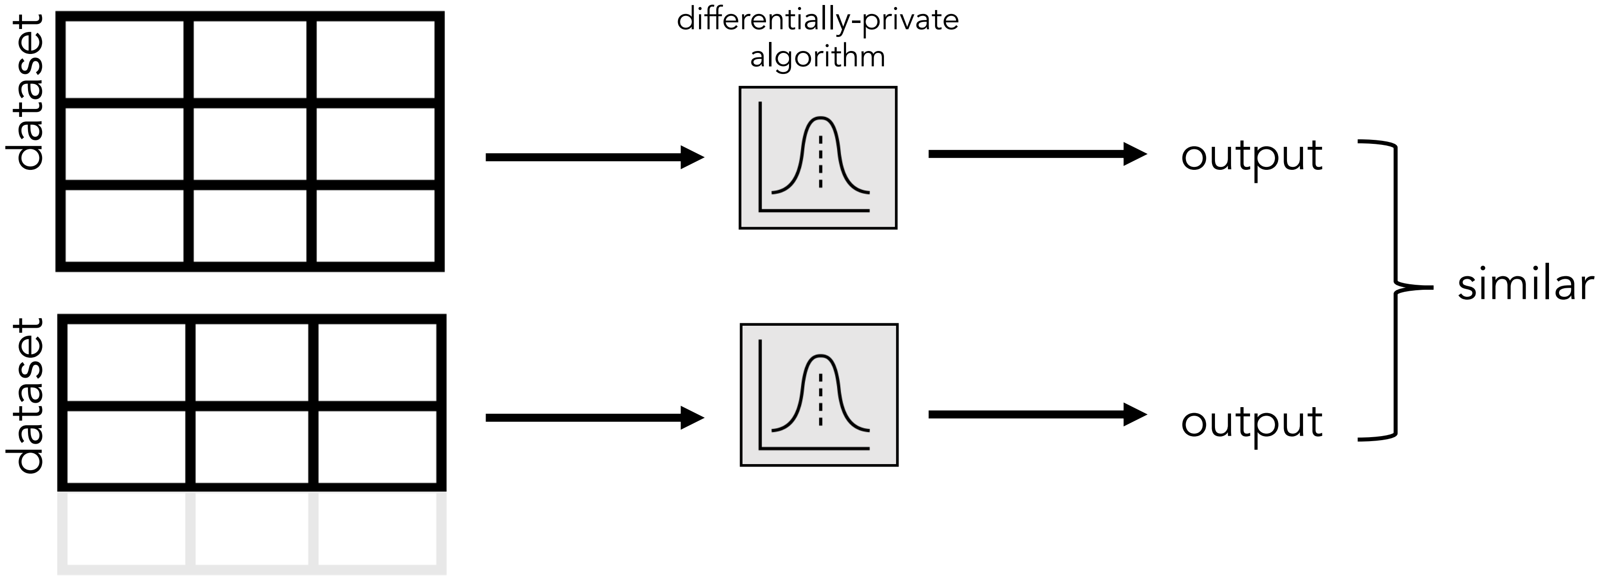

Diagram adapted from
Wood, Altman, Bembenek et al. (2020). Differential Privacy: A Primer for a Non-Technical Audience
Near, Darais, Boeckl (2020). Differential Privacy for Privacy-Preserving Data Analysis: An Introduction to our Blog Series

Dwork, McSherry, Nissim, Smith (2006). Calibrating noise to sensitivity in private data analysis *stronger privacy protections* privacy loss budget (ε) *weaker privacy protections*

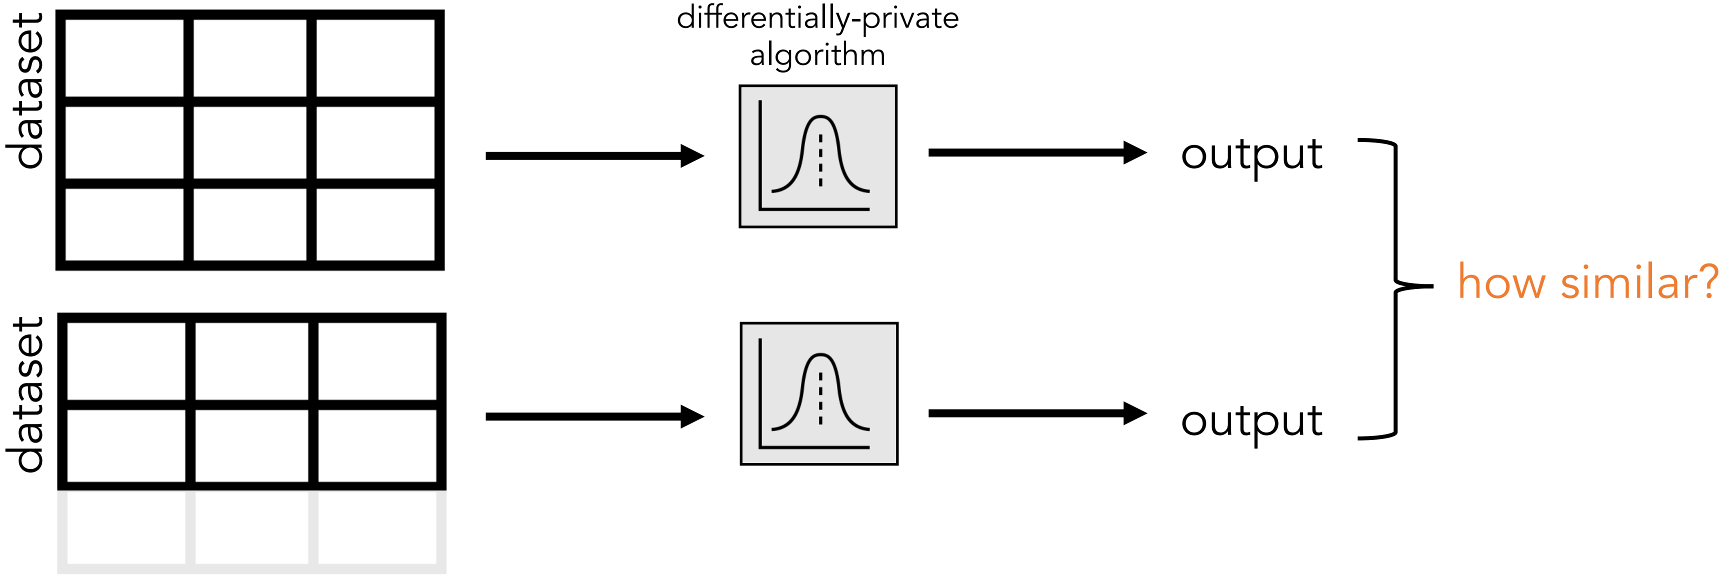

Diagram adapted from
Wood, Altman, Bembenek et al. (2020). Differential Privacy: A Primer for a Non-Technical Audience
Near, Darais, Boeckl (2020). Differential Privacy for Privacy-Preserving Data Analysis: An Introduction to our Blog Series

Dwork, McSherry, Nissim, Smith (2006). Calibrating noise to sensitivity in private data analysis

Epsilon Registry

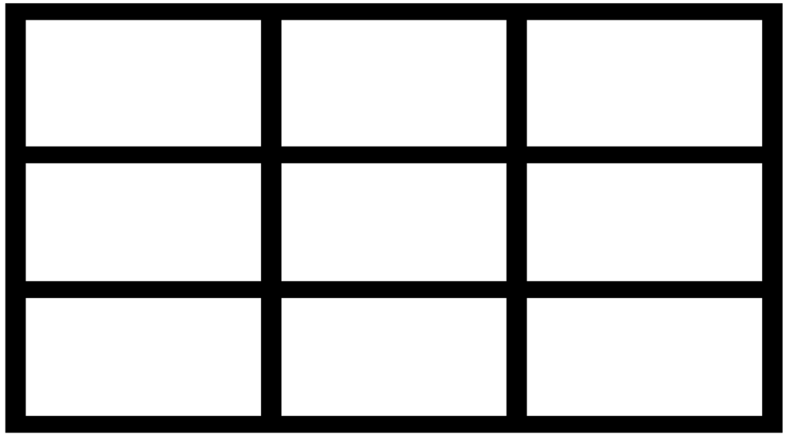

Dwork, Kohli, Mulligan (2019). Differential privacy in practice: Expose your epsilons!

Epsilon Registry

Need ε explanations!

Dwork, Kohli, Mulligan (2019). Differential privacy in practice: Expose your epsilons!

To reduce the intrusion into personal privacy, the company says they will use a technique called differential privacy. Differential privacy injects statistical noise into collected data in a way that protects privacy without significantly changing conclusions.

Adapted from Cummings et al.'s "Techniques" description

Cummings, Kaptchuk, Redmiles (2021). "I need a better description": An Investigation Into User Expectations For Differential Privacy

To reduce the intrusion into personal privacy, the company says they will use a technique called differential privacy. Differential privacy injects statistical noise into collected data in a way that protects privacy without significantly changing conclusions.

Adapted from Cummings et al.'s "Techniques" description

No ε information!

Cummings, Kaptchuk, Redmiles (2021). "I need a better description": An Investigation Into User Expectations For Differential Privacy

Challenges to explaining ε

unit-less & contextless

probabilistic guarantees

Kahneman, Slovic, Tversky (1982). Judgment under uncertainty: Heuristics and biases
Slovic (2000). The perception of risk

Explanation methods for *ε* that increase

objective risk comprehension

subjective privacy understanding

self-efficacy

confidence deciding   enough information

Odds-Based (Text)	Odds-Based (Visual)	Example-Based

Portable explanation methods for ε

If you do not participate, x out of 100 potential *DP outputs* will lead adversary A to believe you responded d_{true}.

If you participate, y out of 100 potential *DP outputs* will lead adversary A to believe you responded d_{true}.

If you do not participate, *x* out of 100 potential *DP outputs* will lead adversary *A* to believe you responded d_{true}.

If you participate, *y* out of 100 potential *DP outputs* will lead adversary *A* to believe you responded d_{true}.

If you do not participate, x out of 100 potential *DP outputs* will lead adversary A to believe you responded d_{true}.

If you participate, y out of 100 potential *DP outputs* will lead adversary A to believe you responded d_{true}.

If you $\underbrace{\textbf{do not participate}}^{\textbf{Probabilities reflect immediate decisions}}$, x out of 100 potential *DP outputs* will lead adversary A to believe you responded d_{true}.

If you **participate**, y out of 100 potential *DP outputs* will lead adversary A to believe you responded d_{true}.

Framing probabilities as frequencies vs. percentages supports statistical reasoning & has been applied in privacy contexts

If you do not participate, ***x* out of 100** potential *DP outputs* will lead adversary *A* to believe you responded d_{true}.

If you participate, ***y* out of 100** potential *DP outputs* will lead adversary *A* to believe you responded d_{true}.

Gigerenzer and Hoffrage (1995). How to improve Bayesian reasoning without instruction: Frequency formats
Hoffrage and Gigerenzer (1998). Using natural frequencies to improve diagnostic inferences
Slovic (2000). The perception of risk
Kaptchuk, Goldstein, Hargittai, Hofman, and Redmiles (2020). How good is good enough for COVID19 apps? …
Franzen, Nuñez von Voigt, Sörries, Tschorsch, Müller-Birn (2022). "Am I private and if so, how many?" …

If you **do not participate**,
x out of 100 potential *DP outputs* will lead adversary *A* to believe you responded d_{true}.

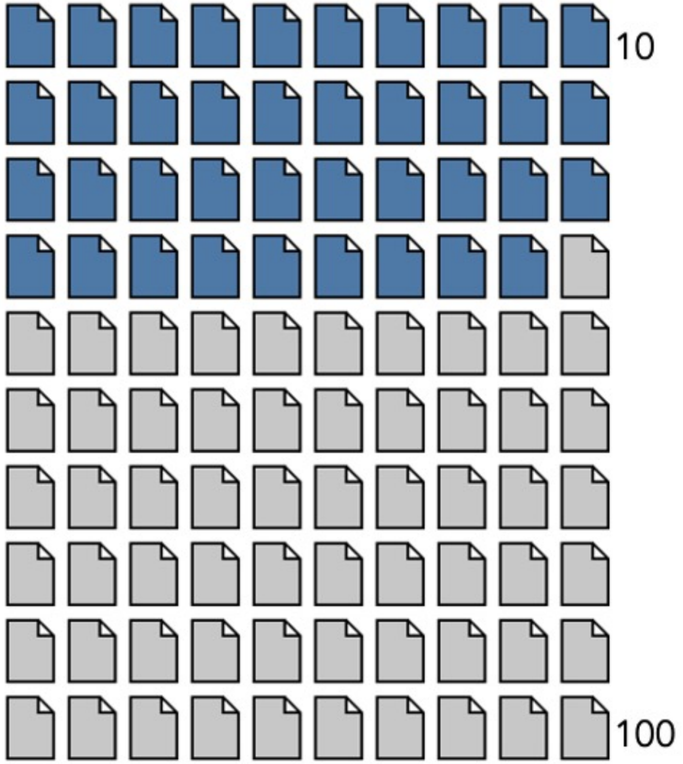

If you **participate**,
y out of 100 potential *DP outputs* will lead adversary *A* to believe you responded d_{true}.

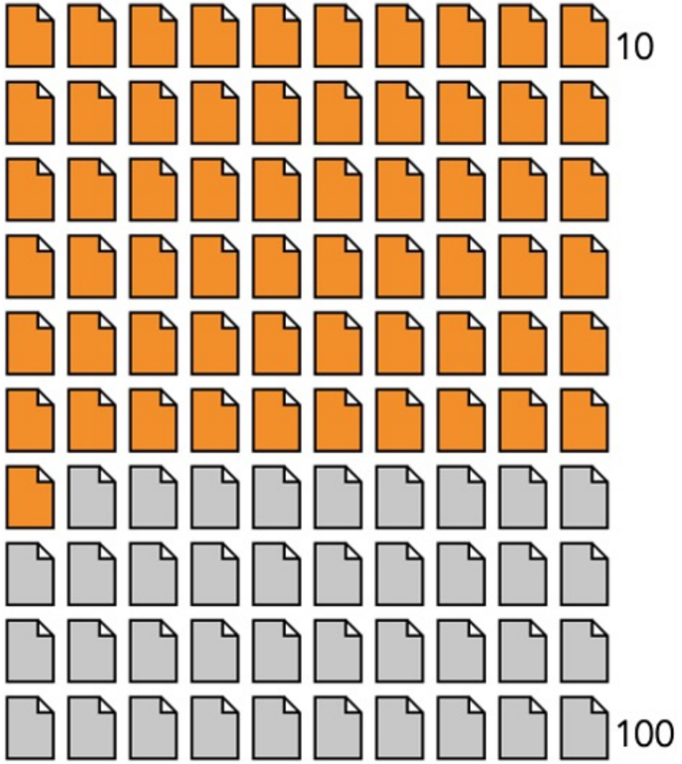

Icon arrays assume x = 39 and y = 61 for illustration purposes.

Galesic, Garcia-Retamero, Gigerenzer (2009). Using icon arrays to communicate medical risks: Overcoming low numeracy

If you **do not participate**, below are examples of potential DP outputs adversary *A* might receive:

Potential Output	x_1
Potential Output	x_2
Potential Output	x_3
Potential Output	x_4
Potential Output	x_5

If you **participate**, below are examples of potential DP outputs adversary *A* might receive:

Potential Output	y_1
Potential Output	y_2
Potential Output	y_3
Potential Output	y_4
Potential Output	y_5

Harbach, Hettig, Weber, Smith (2014). Using personal examples to improve risk communication for security & privacy decisions

If you **do not participate**, below are examples of potential DP outputs adversary *A* might receive:

Potential Output	x_1
Potential Output	x_2
Potential Output	x_3
Potential Output	x_4
Potential Output	x_5

Are these similar to

If you **participate**, below are examples of potential DP outputs adversary *A* might receive:

Potential Output	y_1
Potential Output	y_2
Potential Output	y_3
Potential Output	y_4
Potential Output	y_5

these?

Evaluation Criteria

objective risk comprehension

subjective privacy understanding

self-efficacy

confidence deciding   enough information

Evaluation Criteria

objective risk comprehension

subjective privacy understanding

self-efficacy

confidence deciding   enough information

willingness to share data

Between-subjects **vignette survey study** (n = 963)

Workplace scenario with a **data-sharing decision**

Hainmueller, Hangartner, Yamamoto (2015). Validating vignette and conjoint survey experiments against real-world behavior

Survey Flow

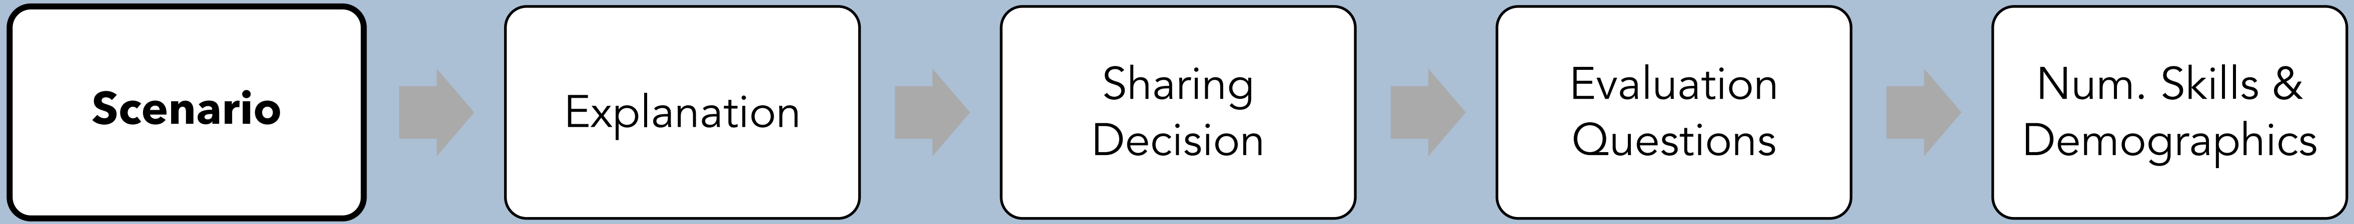

Between-subjects **vignette survey study** (n = 963)

Workplace scenario with a **data-sharing decision**

mandatory vs. optional

Survey Flow

Scenario → Explanation → Sharing Decision → Evaluation Questions → Num. Skills & Demographics

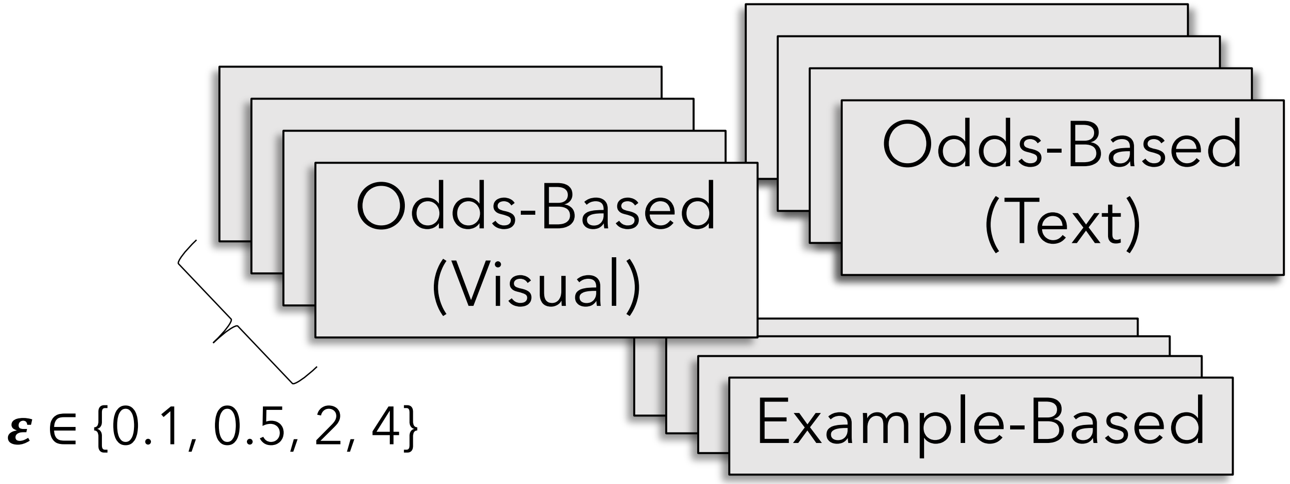

Survey Flow

Scenario → **Explanation** → Sharing Decision → Evaluation Questions → Num. Skills & Demographics

Odds-Based (Visual)

Odds-Based (Text)

Example-Based

$\boldsymbol{\varepsilon} \in \{0.1, 0.5, 2, 4\}$

Deterministic Control

Xiong et al. (2020) Control

Xiong, Wang, Li, Jha (2020). Towards effective differential privacy communication for users' data sharing decision and comprehension

Survey Flow

Scenario → **Explanation** → Sharing Decision → Evaluation Questions → Num. Skills & Demographics

Would you share your data?

Yes/No

Briefly explain your reasoning.

Survey Flow

Scenario → Explanation → **Sharing Decision** → Evaluation Questions → Num. Skills & Demographics

Objective risk comprehension T/F

Subjective privacy understanding Likert-style

Self-efficacy Likert-style

Survey Flow

Scenario → Explanation → Sharing Decision → **Evaluation Questions** → Num. Skills & Demographics

Results

Compared to our ***Example-Based Method****,*
Odds-Based Text and **Odds-Based Visual** improved:

Objective risk comprehension *(O.R. = 4.7; 7.6)*

Subjective privacy understanding *(O.R. = 1.7; 1.5)*

Self-efficacy (enough info) *(O.R. = 1.7; 1.6)*

Results

Over **2x** as likely to answer an additional **objective risk comprehension** question correctly with **Odds-Based Visual** vs. ***Deterministic Control***

Negative effect of our **Example-Based** Method (O.R. = 0.32)

No significant effect of Odds-Based Text

Results

*Compared to the **Xiong et al. Control**,*
Odds-Based Text and **Odds-Based Visual** improved
self-efficacy (enough info) *(O.R. = 1.8; 1.7)*

No significant effect of our Example-Based Method

Results: Willingness to Share Data

*Compared to the **Xiong et al. Control**,*

over 2x, **nearly 2x**, **over 4x**
as likely to **share data** when given
Odds-Based Text,
Odds-Based Visual,
& **Example-Based** respectively

Results: Willingness to Share Data

*Compared to the **Xiong et al. Control**,*

over 2x, **nearly 2x**, **over 4x**
as likely to **share data** when given
Odds-Based Text,
Odds-Based Visual,
& **Example-Based** respectively

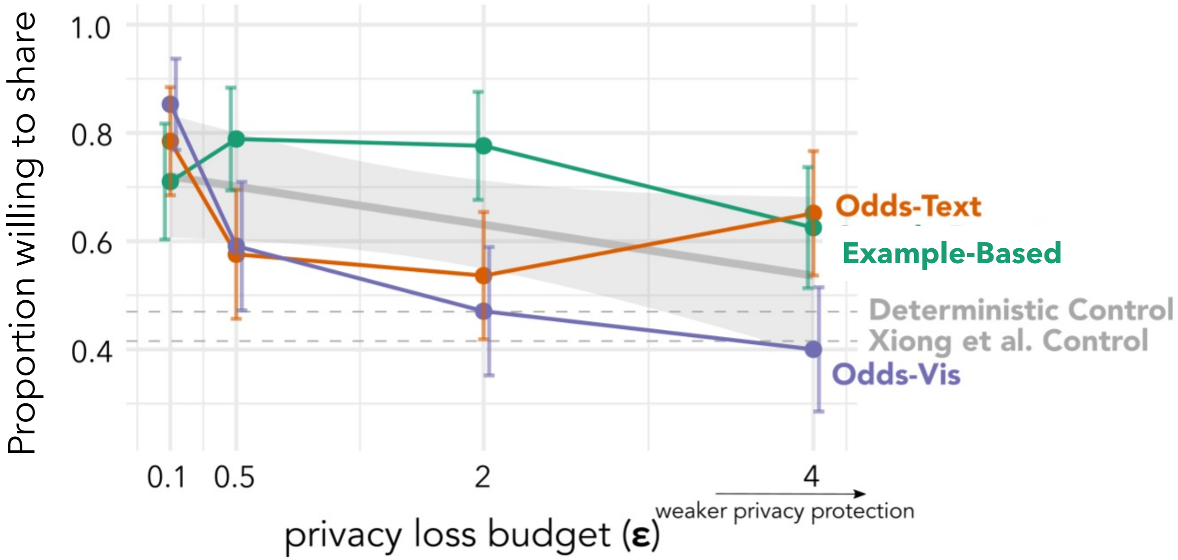

Decreased willingness to share as privacy strength decreases

Takeaways

- → Odds-based methods are promising for explaining ε to end users
- → Explanations should include ε information, since it supports self-efficacy
- → People's willingness to share data is sensitive to changes in ε
- → Explanation methods can support auditing & public deliberation over differential privacy deployments

Takeaways

- → Odds-based methods are promising for explaining ε to end users
- → Explanations should include ε information, since it supports self-efficacy
- → People's willingness to share data is sensitive to changes in ε
- → Explanation methods can support auditing & public deliberation over differential privacy deployments

Future

- → Explain impacts of ε on accuracy & utility
- → Port our methods into real-world settings & create developer tools

Takeaways

→ Odds-based methods are promising for explaining $\boldsymbol{\varepsilon}$ to end users

→ Explanations should include $\boldsymbol{\varepsilon}$ information, since it supports self-efficacy

→ People's willingness to share data is sensitive to changes in $\boldsymbol{\varepsilon}$

→ Explanation methods can support auditing & public deliberation over differential privacy deployments

Future

→ Explain impacts of $\boldsymbol{\varepsilon}$ on accuracy & utility

→ Port our methods into real-world settings & create developer tools

Thank you!

Priyanka Nanayakkara (priyankan@u.northwestern.edu | @priyakalot | @priyakalot@hci.social)
Coauthors: Mary Anne Smart, Rachel Cummings, Gabriel Kaptchuk, Elissa M. Redmiles