Differential Privacy Surprises: Utility is not (always) monotonic on epsilon

Based on CCS ’23 paper by:

Mario Alvim[1], Natasha Fernandes[2], Annabelle McIver[2],
Carroll Morgan[3] and Gabriel Nunes[1]

[1] UFMG, Brazil [2] Macquarie University, Australia [3] UNSW, Australia

Differential Privacy

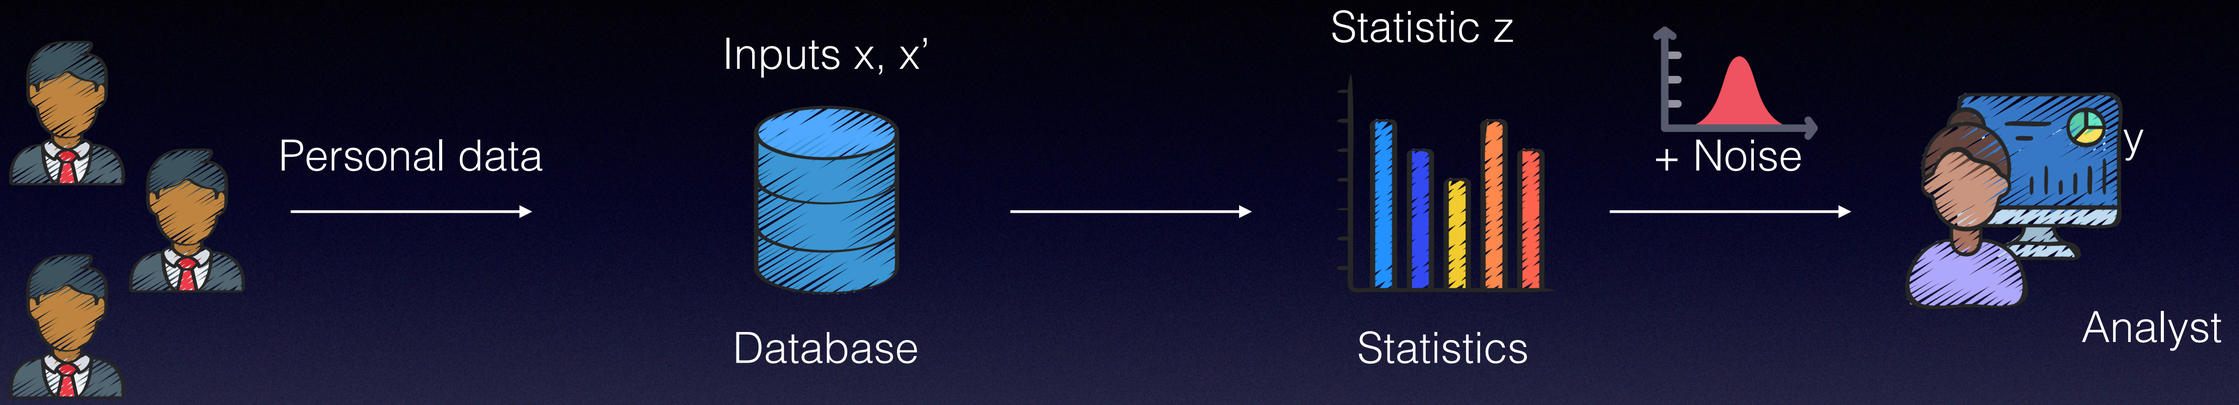

* Add statistical noise to prevent inferences about individuals' data.
* Noise tuning: $P(y \mid x) \leq e^{\epsilon} P(y \mid x')$

  where x, x' are datasets differing in one individual.
* Designed to protect against membership inference attacks.

* Also want to preserve useful information (inferences about z given y)

Local Differential Privacy

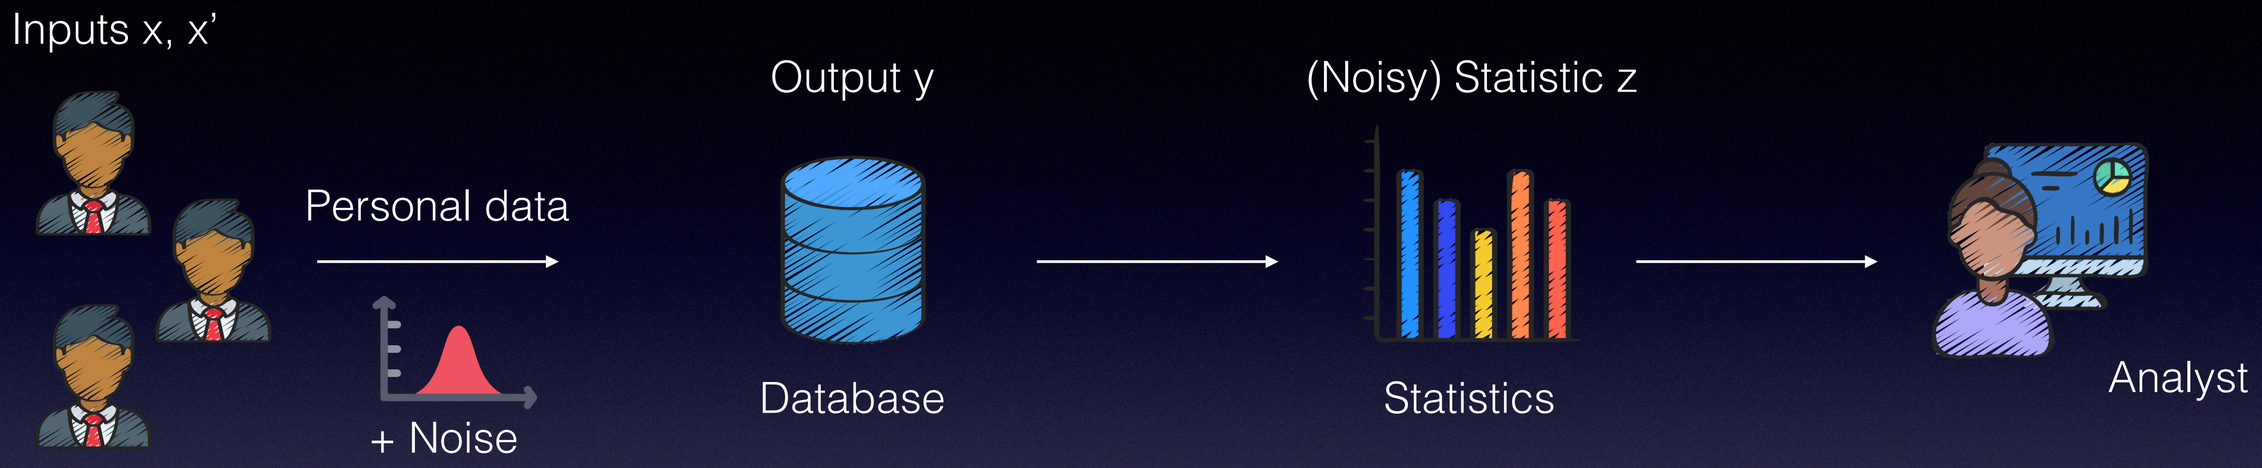

* Add statistical noise to prevent inferences about individuals' data.
* Noise tuning: $P(y \mid x) \leq e^{\epsilon} P(y \mid x')$

  where x, x' are different data values of an individual.
* Designed to protect against attribute inference attacks.

* Also want to preserve useful information (observation z vs true statistic.)

General Privacy Workflows

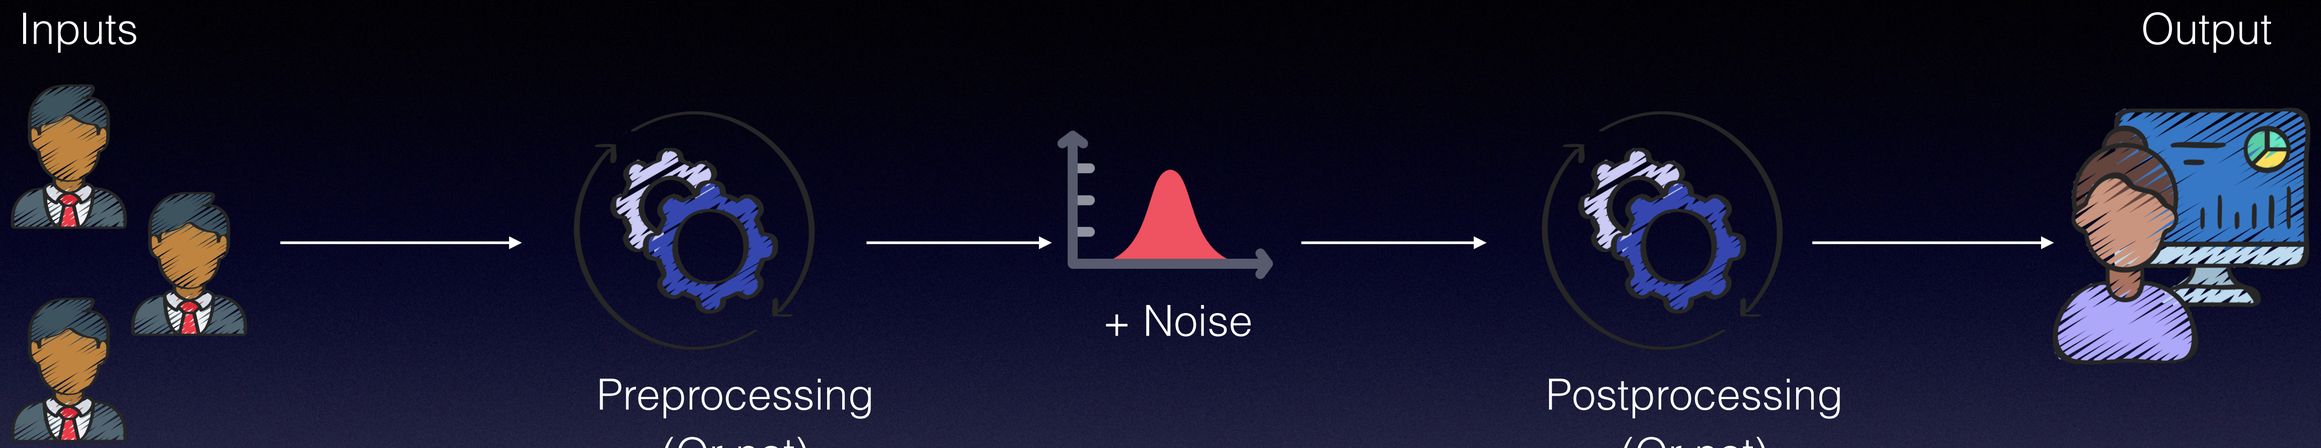

Example: DP-SGD

How to tune the noise to optimise the privacy-utility trade-off?

Assumption: Increasing epsilon (decreasing privacy) causes utility to increase — **monotonicity**

Our result: Utility is not always monotonic on epsilon in general privacy workflows

Quantitative Information Flow

Differential privacy: $C_{x,y} \leq e^{\epsilon} C_{x',y}$

$$C = \begin{bmatrix} 2/3 & 1/6 & 1/6 \\ 1/3 & 1/3 & 1/3 \\ 1/6 & 1/6 & 2/3 \end{bmatrix}$$ <- log4-differentially private

Information flow channel

Refinement: $A \sqsubseteq B$ means channel A is more useful than channel B
to a Bayesian analyst equipped with **any** prior and **any** loss function

Theorem (Coriaceous):

$A \sqsubseteq B$ iff there exists a channel W s.t. $A \cdot W = B$

where $\cdot$ is matrix multiplication (post-processing)

General Privacy Workflows

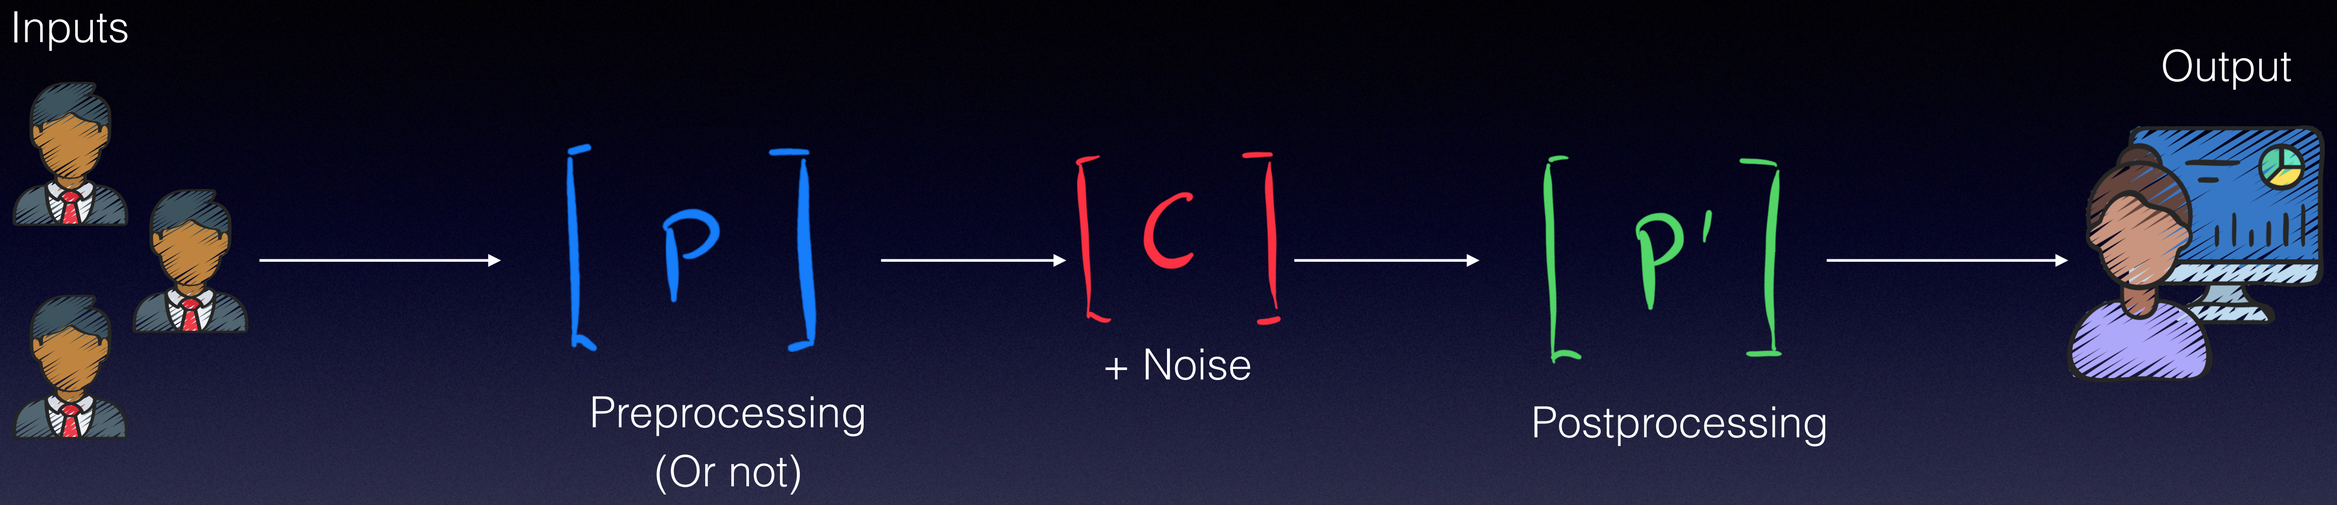

How to tune the noise to optimise the privacy-utility trade-off?

OR: If I change epsilon, how does utility change?

OR: What is the relationship between epsilon and refinement?

i.e. if $\epsilon \geq \epsilon'$ then is it true that $P \cdot C_{\epsilon} \cdot P' \sqsubseteq P \cdot C_{\epsilon'} \cdot P'$?

Properties of Refinement

Theorem*: Given channels C, C' it holds that

$$C \sqsubseteq C' \implies \epsilon(C) \geq \epsilon(C')$$

(i.e. Epsilon is always monotonic on utility)

Theorem*: If channels C, C' belong to the same "family" it holds that

$$\epsilon(C) \geq \epsilon(C') \implies C \sqsubseteq C'$$

(i.e. Utility is monotonic on epsilon).
Note: This was proven for KRR and Geometric "families" only.

* **Chatzikokolakis et al: Comparing mechanisms: max-case refinement orders and application to differential privacy, CSF 2019.**

Properties of Refinement

Theorem*: Given channels C, C' and a pre-processing step P it holds that

$$C \sqsubseteq C' \implies P \cdot C \sqsubseteq P \cdot C'$$

Corollary: If C, C' are both KRR or both Geometric then

$$\epsilon(C) \geq \epsilon(C') \implies P \cdot C \sqsubseteq P \cdot C'$$

i.e. Increasing epsilon also increases utility (monotonicity holds).

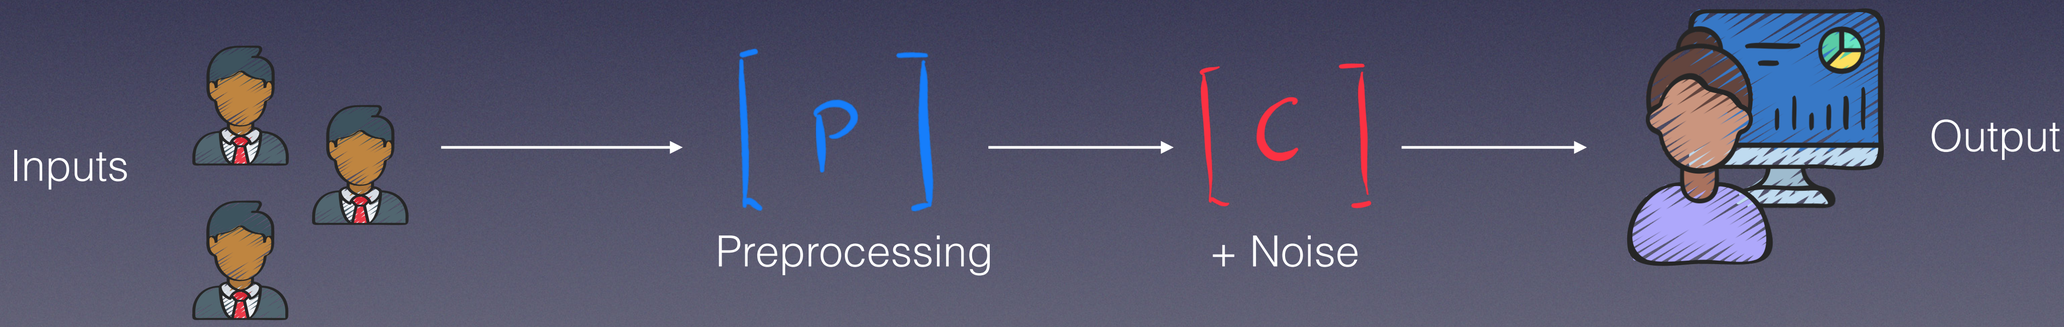

* **Alvim et al: The Science of Quantitative Information Flow, published by Springer 2020.**

Properties of Refinement

Theorem*: Given channels C, C’ and a post-processing step P then

$$C \sqsubseteq C' \implies C \cdot P \sqsubseteq C' \cdot P$$

Corollary: If C, C’ are both KRR or both Geometric then

$$\epsilon(C) \geq \epsilon(C') \implies C \cdot P \sqsubseteq C' \cdot P$$

BUT: Can’t apply this result directly to our local DP workflow…

* **Alvim et al: The Science of Quantitative Information Flow, published by Springer 2020.**

Privacy Workflows

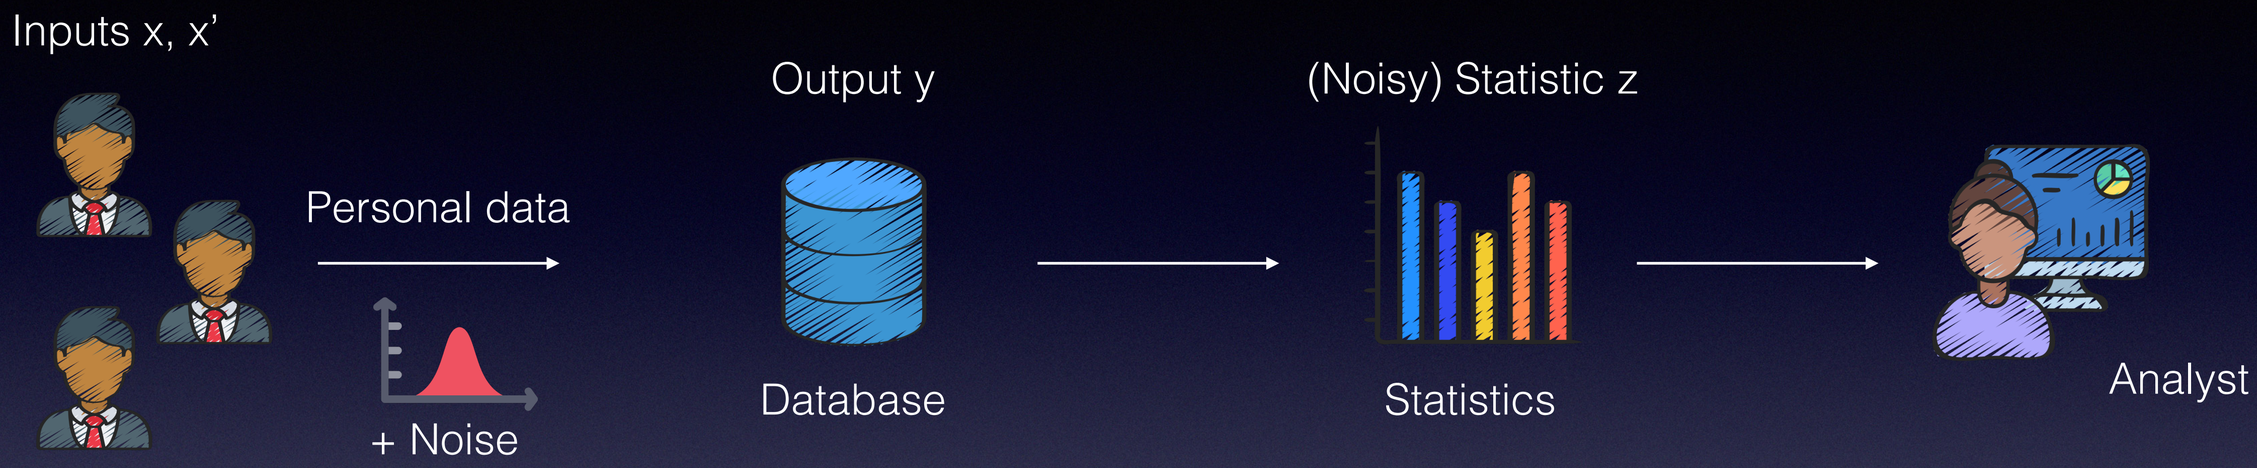

Issue: Noise is added to each individual but post-processing is done on the combination of individuals

Not $C_\epsilon \cdot P$ but n-fold composition of C_ϵ followed by P

Kronecker Composition

Personal data

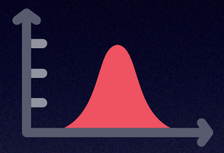

+ Noise

	0	1
0	3/4	1/4
1	1/4	3/4

Personal data

+ Noise

	(0,0)	(0,1)	(1,0)	(1,1)
(0,0)	9/16	3/16	3/16	1/16
(0,1)	3/16	9/16	1/16	3/16
(1,0)	3/16	1/16	9/16	3/16
(1,1)	1/16	3/16	3/16	9/16

“Kronecker product”

Kronecker Composition

Properties of Kronecker product:

Associativity: $(A \otimes B) \otimes C = A \otimes (B \otimes C)$

Bilinearity: $A \otimes (B + C) = A \otimes B + A \otimes C$ *and* $(B + C) \otimes A = B \otimes A + C \otimes A$

Product respecting: $(A \otimes B) \cdot (C \otimes D) = (A \cdot C) \otimes (B \cdot D)$ *if* $A \cdot C$ *and* $B \cdot D$ *are defined.*

Invertibility: If A, B *are invertible, then so is* $A \otimes B$*. The inverse is* $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$

But $\otimes$ is not commutative in general.

Kronecker Composition

Theorem: Given channels C, C' then

$$C \sqsubseteq C' \implies C^{\otimes N} \sqsubseteq C'^{\otimes N}$$

Corollary: Given channels C, C' then

$$\epsilon(C) \geq \epsilon(C') \implies \epsilon(C^{\otimes N}) \geq \epsilon(C'^{\otimes N})$$

Corollary: Given channels C, C' both KRR or both Geometric then

$$\epsilon(C) \geq \epsilon(C') \implies C^{\otimes N} \sqsubseteq C'^{\otimes N}$$

Non-Monotonicity in Theory

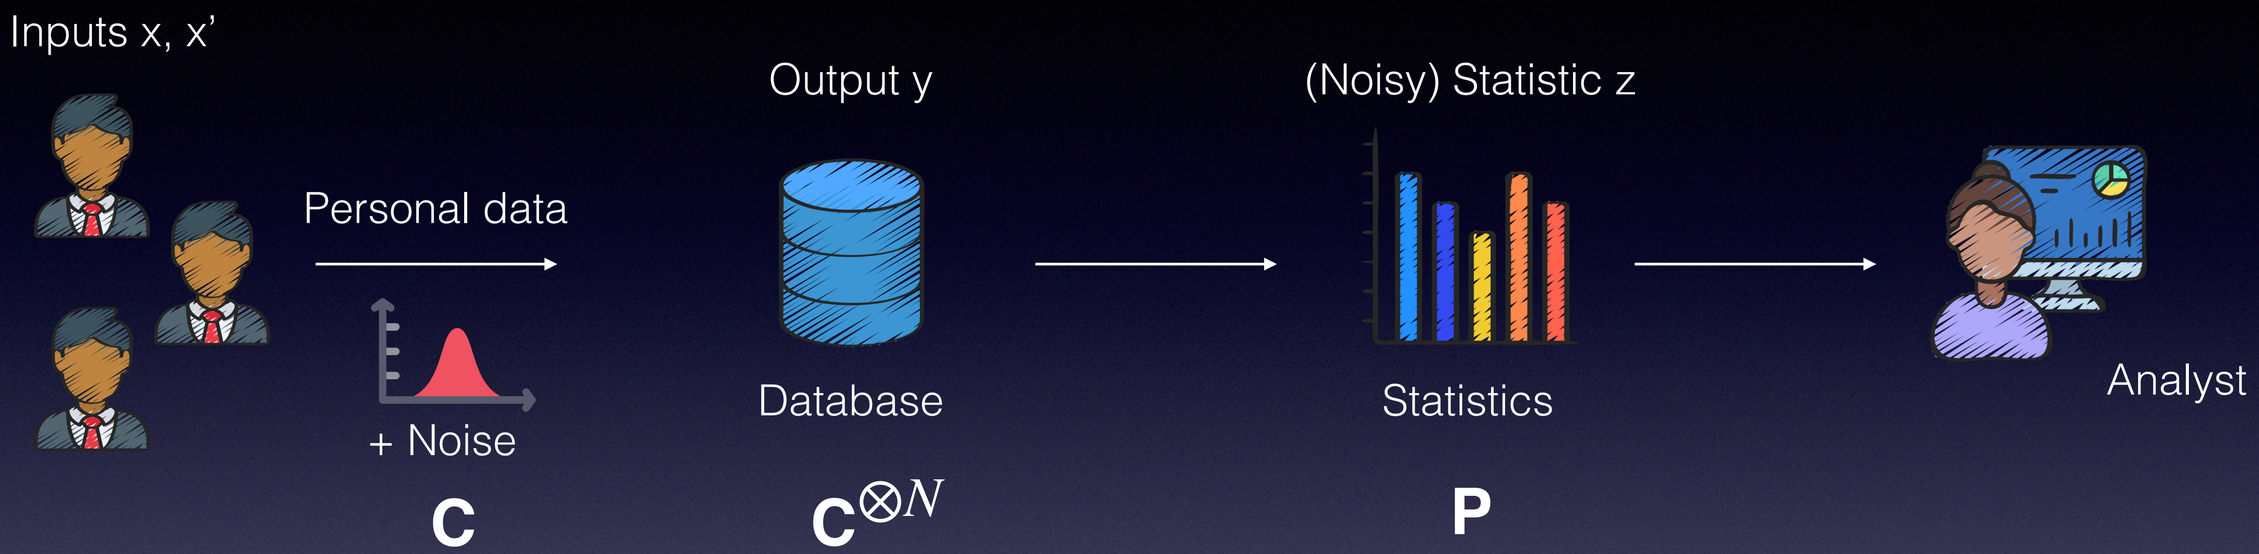

Theorem: Given noise-adding channels C, C' in the same family,

$$\epsilon(C) \geq \epsilon(C') \implies C \sqsubseteq C' \implies C^{\otimes N} \sqsubseteq C'^{\otimes N} \not\Longleftrightarrow C^{\otimes N} \cdot P \sqsubseteq C'^{\otimes N} \cdot P$$

i.e. There exist cases where we can get more privacy and more utility.

Non-Monotonicity in Practice

But… is this a problem for utility measures that are meaningful?

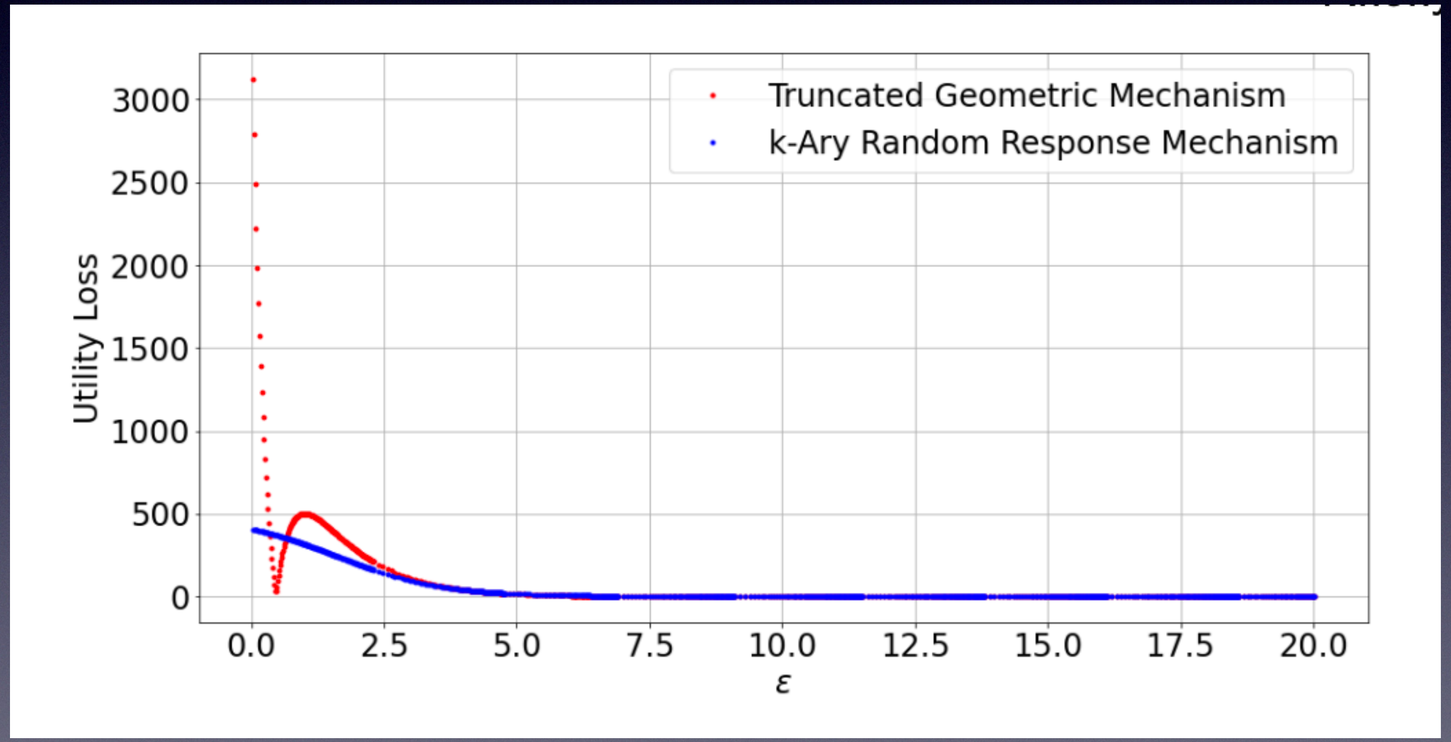

Yes!

Utility measured using mean absolute error

Loss(x, y) = |y - x|

Data source: https://github.com/propublica/compas-analysis

Conditions for Monotonicity

Under what conditions does the following hold?

$$C^{\otimes N} \sqsubseteq C'^{\otimes N} \implies C^{\otimes N} \cdot P \sqsubseteq C'^{\otimes N} \cdot P$$

Theorem: Given channels C, C' s.t. $C \sqsubseteq C'$ with witness W, and a post-processor P, the following are sufficient conditions for monotonicity of utility on epsilon:

1. P has a left inverse $^{-1}$P
2. P · $^{-1}$P · W · P = W · P

Monotonicity Results

Theorem: If R, R' are KRR mechanisms then

$$R^{\otimes N} \sqsubseteq R'^{\otimes N} \implies R^{\otimes N} \cdot P \sqsubseteq R'^{\otimes N} \cdot P$$

whenever P is a "counting" query.

Corollary: If R, R' are KRR mechanisms then

$$\epsilon(R) \geq \epsilon(R') \implies R^{\otimes N} \cdot P \sqsubseteq R'^{\otimes N} \cdot P$$

whenever P is a "counting" query.

Non-Monotonicity Results

Many negative results:

Theorem: If R, R' are "KRR" channels then

$$R^{\otimes N} \sqsubseteq R'^{\otimes N} \not\Longleftrightarrow R^{\otimes N} \cdot P \sqsubseteq R'^{\otimes N} \cdot P$$

whenever P is a "sum" query.

Also does not hold in general for Geometric channels for counting or sum P.

Statistical Postprocessors

1. “Counting” query - deterministic channel T

$R^{\otimes 2}$	(yes, yes)	(yes, no)	(no, yes)	(no, no)
(yes, yes)	$3/4 \times 3/4$	$3/4 \times 1/4$	$1/4 \times 3/4$	$1/4 \times 1/4$
(yes, no)	$3/4 \times 1/4$	$3/4 \times 3/4$	$1/4 \times 1/4$	$1/4 \times 3/4$
(no, yes)	$1/4 \times 3/4$	$1/4 \times 1/4$	$3/4 \times 3/4$	$3/4 \times 1/4$
(no, no)	$1/4 \times 1/4$	$1/4 \times 3/4$	$3/4 \times 1/4$	$3/4 \times 3/4$

T	0	1	2
(yes, yes)	0	0	1
(yes, no)	0	1	0
(no, yes)	0	1	0
(no, no)	1	0	0

$R^{\otimes 2}$ T	0	1	2
(yes, yes)	$1/4 \times 1/4$	$3/4 \times 1/4 + 1/4 \times 3/4$	$3/4 \times 3/4$
(yes, no)	$1/4 \times 3/4$	$3/4 \times 3/4 + 1/4 \times 1/4$	$3/4 \times 1/4$
(no, yes)	$3/4 \times 1/4$	$1/4 \times 1/4 + 3/4 \times 3/4$	$1/4 \times 3/4$
(no, no)	$3/4 \times 3/4$	$1/4 \times 3/4 + 3/4 \times 1/4$	$1/4 \times 1/4$

Statistical Postprocessors

2. “Sum” query - deterministic channel S

R	0	1	2
0	1/2	1/4	1/4
1	1/4	1/2	1/4
2	1/4	1/4	1/2

$\otimes$

R	0	1	2
0	1/2	1/4	1/4
1	1/4	1/2	1/4
2	1/4	1/4	1/2

S	0	1	2	3	4
00	1	0	0	0	0
01	0	1	0	0	0
02	0	0	1	0	0
10	0	1	0	0	0
11	0	0	1	0	0
12	0	0	0	1	0
20	0	0	1	0	0
21	0	0	0	1	0
22	0	0	0	0	1

Conclusion/Future Work

Main Takeaways:

- Refinement is a useful tool for reasoning about utility in differential privacy
- Kronecker products allow reasoning about utility in local differential privacy contexts
- Utility is not always monotonic on epsilon!

Future work:

- Study of monotonicity in machine learning contexts eg. DP-SGD.