Can Swapping Be Differentially Privacy?

A Refreshment Stirred, not Shaken

James Bailie[†], Ruobin Gong[‡] and Xiao-Li Meng[†]

Statistics Department, [[†]Harvard/[‡]Rutgers] University

Privacy and Public Policy Conference

Georgetown University

September 14, 2024

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications)

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change:

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy *is* to limit this rate of change.

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy *is* to limit this rate of change.
- Different DP specifications correspond to different choices of how (and where) to measure *input* and *output changes*, in addition to how much to control this rate.

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy ***is*** to limit this rate of change.
- Different DP specifications correspond to different choices of how (and where) to measure *input* and *output changes*, in addition to how much to control this rate.
- These choice are the building blocks of a DP specification:

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy *is* to limit this rate of change.
- Different DP specifications correspond to different choices of how (and where) to measure *input* and *output changes*, in addition to how much to control this rate.
- These choice are the building blocks of a DP specification:
    1. The protection *domain* ($\mathcal{X}$)

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy *is* to limit this rate of change.
- Different DP specifications correspond to different choices of how (and where) to measure *input* and *output changes*, in addition to how much to control this rate.
- These choice are the building blocks of a DP specification:
    1. The protection *domain* ($\mathcal{X}$)
    2. The *scope* of protection ($\mathscr{D}$)

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy *is* to limit this rate of change.
- Different DP specifications correspond to different choices of how (and where) to measure *input* and *output changes*, in addition to how much to control this rate.
- These choice are the building blocks of a DP specification:
    1. The protection *domain* ($\mathcal{X}$)
    2. The *scope* of protection ($\mathscr{D}$)
    3. The protection *unit* ($d_{\mathcal{X}}$)

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy ***is*** to limit this rate of change.
- Different DP specifications correspond to different choices of how (and where) to measure *input* and *output changes*, in addition to how much to control this rate.
- These choice are the building blocks of a DP specification:
  1. The protection *domain* ($\mathcal{X}$)
  2. The *scope* of protection ($\mathscr{D}$)
  3. The protection *unit* ($d_{\mathcal{X}}$)
  4. The *standard* of protection (d_{Pr})

Differential Privacy (DP) (Dwork, McSherry, et al., 2006)

A large family of technical standards (i.e. mathematical specifications) that conceptualizes privacy loss as a rate of change: the *change* in the (distribution of) the *output* statistics per unit *change* in the *input* data.

- Key idea: To protect privacy *is* to limit this rate of change.
- Different DP specifications correspond to different choices of how (and where) to measure *input* and *output changes*, in addition to how much to control this rate.
- These choice are the building blocks of a DP specification:
  1. The protection *domain* ($\mathcal{X}$)
  2. The *scope* of protection ($\mathscr{D}$)
  3. The protection *unit* ($d_{\mathcal{X}}$)
  4. The *standard* of protection (d_{Pr})
  5. The *intensity* of protection (ε)

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \overset{\textit{Data release}}{\rightsquigarrow} \textit{Statistic } T(\boldsymbol{x}, Z)$$

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Object of interest: A statistic T – i.e. a function of the data $\boldsymbol{x} \in \mathcal{X}$

For example,

$$T(\boldsymbol{x}) = \frac{1}{n}\sum_{i=1}^{n} x_i$$

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Object of interest: A statistic T – i.e. a function of the data $\boldsymbol{x} \in \mathcal{X}$ and some auxiliary random noise Z.

For example,

$$T(\boldsymbol{x}) = \frac{1}{n}\sum_{i=1}^{n} x_i + Z.$$

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about T as a function of the dataset $\boldsymbol{x} \in \mathcal{X}$, its derivative is

$$\lim_{\boldsymbol{x}' \to \boldsymbol{x}} \frac{T(\boldsymbol{x}', Z) - T(\boldsymbol{x}, Z)}{\boldsymbol{x}' - \boldsymbol{x}}.$$

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathrm{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its derivative is

$$\lim_{\boldsymbol{x}' \to \boldsymbol{x}} \frac{\mathrm{P}_{\boldsymbol{x}'}(T) - \mathrm{P}_{\boldsymbol{x}}(T)}{\boldsymbol{x}' - \boldsymbol{x}}.$$

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathrm{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its derivative is

$$\lim_{\boldsymbol{x}' \to \boldsymbol{x}} \frac{d_{\mathrm{Pr}}(\mathrm{P}_{\boldsymbol{x}'}, \mathrm{P}_{\boldsymbol{x}})}{\boldsymbol{x}' - \boldsymbol{x}}.$$

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathrm{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its derivative is

$$\lim_{\boldsymbol{x}' \to \boldsymbol{x}} \frac{d_{\mathrm{Pr}}(\mathrm{P}_{\boldsymbol{x}'}, \mathrm{P}_{\boldsymbol{x}})}{d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x})}.$$

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathsf{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its derivative is

$$\cancel{\lim_{\boldsymbol{x}' \to \boldsymbol{x}}} \frac{d_{\mathsf{Pr}}(\mathsf{P}_{\boldsymbol{x}'}, \mathsf{P}_{\boldsymbol{x}})}{d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x})},$$

for all $\boldsymbol{x}, \boldsymbol{x}'$.

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathrm{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its ~~derivative~~ Lipschitz constant is the smallest ε such that

$$d_{\mathrm{Pr}}(\mathrm{P}_{\boldsymbol{x}'}, \mathrm{P}_{\boldsymbol{x}}) \le \varepsilon d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x}),$$

for all $\boldsymbol{x}, \boldsymbol{x}'$.

The *Derivative* of DP

$$\text{A population} \xrightarrow{\text{Data collection}} \text{Dataset } \boldsymbol{x} \rightsquigarrow^{\text{Data release}} \text{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathrm{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its ~~derivative~~ Lipschitz constant is the smallest ε such that

$$d_{\mathrm{Pr}}(\mathrm{P}_{\boldsymbol{x}'}, \mathrm{P}_{\boldsymbol{x}}) \leq \varepsilon d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x}),$$

for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}$ and all universes $\mathcal{D} \in \mathscr{D}$.

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathsf{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its ~~derivative~~ Lipschitz constant is the smallest ε such that

$$d_{\mathsf{Pr}}(\mathsf{P}_{\boldsymbol{x}'}, \mathsf{P}_{\boldsymbol{x}}) \leq \varepsilon d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x}),$$

for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}$ and all universes $\mathcal{D} \in \mathscr{D}$.

Definition: The statistic T is ε-differentially private if its Lipschitz constant is ε.

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathsf{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its ~~derivative~~ Lipschitz constant is the smallest ε such that

$$d_{\mathsf{Pr}}(\mathsf{P}_{\boldsymbol{x}'}, \mathsf{P}_{\boldsymbol{x}}) \leq \varepsilon d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x}),$$

for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}$ and all universes $\mathcal{D} \in \mathscr{D}$.

Definition: The statistic T is ε-differentially private if its Lipschitz constant is ε.

- Recall that Lipschitz continuity $\approx$ differentiability.

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathrm{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its ~~derivative~~ Lipschitz constant is the smallest ε such that

$$d_{\mathrm{Pr}}(\mathrm{P}_{\boldsymbol{x}'}, \mathrm{P}_{\boldsymbol{x}}) \leq \varepsilon d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x}),$$

for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}$ and all universes $\mathcal{D} \in \mathscr{D}$.

Definition: The statistic T is ε-differentially private if its Lipschitz constant is ε.

- Recall that Lipschitz continuity $\approx$ differentiability.
- Lipschitz constant is the supremum of the derivative.

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $P_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its ~~derivative~~ Lipschitz constant is the smallest ε such that

$$d_{\mathrm{Pr}}(P_{\boldsymbol{x}'}, P_{\boldsymbol{x}}) \leq \varepsilon d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x}),$$

for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}$ and all universes $\mathcal{D} \in \mathscr{D}$.

Definition: The statistic T is ε-differentially private if its Lipschitz constant is ε.

- Recall that Lipschitz continuity $\approx$ differentiability.
- Lipschitz constant is the supremum of the derivative.

Takeaway: Differential privacy is a "bound on the derivative" of T.

The *Derivative* of DP

$$\textit{A population} \xrightarrow{\textit{Data collection}} \textit{Dataset } \boldsymbol{x} \rightsquigarrow^{\textit{Data release}} \textit{Statistic } T(\boldsymbol{x}, Z)$$

Thinking about the distribution $\mathrm{P}_{\boldsymbol{x}}$ of T as a function of $\boldsymbol{x} \in \mathcal{X}$, its ~~derivative~~ Lipschitz constant is the smallest ε such that

$$d_{\mathrm{Pr}}(\mathrm{P}_{\boldsymbol{x}'}, \mathrm{P}_{\boldsymbol{x}}) \leq \varepsilon d_{\mathcal{X}}(\boldsymbol{x}', \boldsymbol{x}),$$

for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}$ and all universes $\mathcal{D} \in \mathscr{D}$.

Definition: The statistic T is ε-differentially private if its Lipschitz constant is ε.

- Recall that Lipschitz continuity $\approx$ differentiability.
- Lipschitz constant is the supremum of the derivative.

Takeaway: Differential privacy is a "bound on the derivative" of T.

- The choice of $\mathcal{X}$, $\mathscr{D}$, d_{Pr} and $d_{\mathcal{X}}$ determine the *flavour* of DP.

A DP Specification $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathsf{Pr}}, \varepsilon)$

The building blocks of DP:

- The protection domain
- The scope of protection
- The protection unit
- The standard of protection
- The intensity of protection

A DP Specification $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}}, \varepsilon)$

The building blocks of DP:

- The protection domain
  - *Who* is eligible for protection?
  - Defined by the set $\mathcal{X}$ of possible input datasets.
- The scope of protection
- The protection unit
- The standard of protection
- The intensity of protection

A DP Specification $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathsf{Pr}}, \varepsilon)$

The building blocks of DP:

- The protection domain
  - *Who* is eligible for protection?
  - Defined by the set $\mathcal{X}$ of possible input datasets.
- The scope of protection
  - *Where* does the protection extend to?
  - Instantiated by the multiverse $\mathscr{D}$, which is a collection of universes $\mathcal{D} \subset \mathcal{X}$.
- The protection unit
- The standard of protection
- The intensity of protection

A DP Specification $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathsf{Pr}}, \varepsilon)$

The building blocks of DP:

- The protection domain
  - *Who* is eligible for protection?
  - Defined by the set $\mathcal{X}$ of possible input datasets.
- The scope of protection
  - *Where* does the protection extend to?
  - Instantiated by the multiverse $\mathscr{D}$, which is a collection of universes $\mathcal{D} \subset \mathcal{X}$.
- The protection unit
  - *What* is the granularity of protection?
  - Conceptualized by the input divergence $d_{\mathcal{X}}$ on $\mathcal{X}$.
- The standard of protection
- The intensity of protection

A DP Specification $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}}, \varepsilon)$

The building blocks of DP:

- The protection domain
  - *Who* is eligible for protection?
  - Defined by the set $\mathcal{X}$ of possible input datasets.
- The scope of protection
  - *Where* does the protection extend to?
  - Instantiated by the multiverse $\mathscr{D}$, which is a collection of universes $\mathcal{D} \subset \mathcal{X}$.
- The protection unit
  - *What* is the granularity of protection?
  - Conceptualized by the input divergence $d_{\mathcal{X}}$ on $\mathcal{X}$.
- The standard of protection
  - *How* to measure change in the output variations?
  - Captured by the output divergence d_{Pr} on probability distributions.
- The intensity of protection

A DP Specification $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathsf{Pr}}, \varepsilon)$

The building blocks of DP:

- The protection domain
  - *Who* is eligible for protection?
  - Defined by the set $\mathcal{X}$ of possible input datasets.
- The scope of protection
  - *Where* does the protection extend to?
  - Instantiated by the multiverse $\mathscr{D}$, which is a collection of universes $\mathcal{D} \subset \mathcal{X}$.
- The protection unit
  - *What* is the granularity of protection?
  - Conceptualized by the input divergence $d_{\mathcal{X}}$ on $\mathcal{X}$.
- The standard of protection
  - *How* to measure change in the output variations?
  - Captured by the output divergence d_{Pr} on probability distributions.
- The intensity of protection
  - *How much* protection is afforded?
  - Quantified by the privacy-loss budget $\varepsilon_{\mathcal{D}}$.

Some Examples in the Literature

$\underline{\mathcal{X}}$: DP for network data (Hay et al., 2009) for geospatial data (Andrés et al., 2013) Pufferfish DP (Kifer & Machanavajjhala, 2014) noiseless privacy (Bhaskar et al., 2011) privacy under partial knowledge (Seeman et al., 2022) privacy amplification (Beimel et al., 2010; Balle et al., 2020; Bun et al., 2022)

$\underline{\mathscr{D}}$: privacy under invariants (Ashmead et al., 2019; Gong & Meng, 2020; Gao et al., 2022; Dharangutte et al., 2023) conditioned or empirical DP (J. M. Abowd et al., 2013; Charest & Hou, 2016) personalized DP (Ebadi et al., 2015; Jorgensen et al., 2015) individual DP (Soria-Comas et al., 2017; Feldman & Zrnic, 2022) bootstrap DP (O'Keefe & Charest, 2019) stratified DP (Bun et al., 2022) per-record DP (Seeman et al., 2023+) per-instance DP (Wang, 2018; Redberg & Wang, 2021)

$\underline{d_{\mathcal{X}}}$: $(\mathcal{R}, \varepsilon)$-generic DP (Kifer & Machanavajjhala, 2011) edge vs node privacy (Hay et al., 2009; McSherry & Mahajan, 2010) d-metric DP (Chatzikokolakis et al., 2013) Blowfish privacy (He et al., 2014) element level DP (Asi et al., 2022) distributional privacy (Zhou et al., 2009) event-level vs user-level DP (Dwork et al., 2010)

$\underline{d_{\mathrm{Pr}}}$: (ε, δ)-approximate DP (Dwork, Kenthapadi, et al., 2006) Rényi DP (Mironov, 2017) concentrated DP (Bun & Steinke, 2016a) f-divergence privacy (Barber & Duchi, 2014; Barthe & Olmedo, 2013) f-DP (including Gaussian DP) (Dong et al., 2022)

Comparisons: US Decennial Censuses

- We provide a data swapping algorithm which is reminiscent of the statistical disclosure method used in the 1990, 2000 and 2010 US Decennial Censuses.

Comparisons: US Decennial Censuses

- We provide a data swapping algorithm which is reminiscent of the statistical disclosure method used in the 1990, 2000 and 2010 US Decennial Censuses.
- This method does not satisfy the original DP specification – pure ε-DP (Dwork, McSherry, et al., 2006) – because it leaves *invariant* some statistics.

Comparisons: US Decennial Censuses

- We provide a data swapping algorithm which is reminiscent of the statistical disclosure method used in the 1990, 2000 and 2010 US Decennial Censuses.
- This method does not satisfy the original DP specification – pure ε-DP (Dwork, McSherry, et al., 2006) – because it leaves *invariant* some statistics.
- For the same reason, the principal method used in the 2020 Census, the TopDown Algorithm, does not satisfy any standard DP specification.

Comparisons: US Decennial Censuses

- We provide a data swapping algorithm which is reminiscent of the statistical disclosure method used in the 1990, 2000 and 2010 US Decennial Censuses.
- This method does not satisfy the original DP specification – pure ε-DP (Dwork, McSherry, et al., 2006) – because it leaves *invariant* some statistics.
- For the same reason, the principal method used in the 2020 Census, the TopDown Algorithm, does not satisfy any standard DP specification.
- Instead, we prove that

Comparisons: US Decennial Censuses

- We provide a data swapping algorithm which is reminiscent of the statistical disclosure method used in the 1990, 2000 and 2010 US Decennial Censuses.
- This method does not satisfy the original DP specification – pure ε-DP (Dwork, McSherry, et al., 2006) – because it leaves *invariant* some statistics.
- For the same reason, the principal method used in the 2020 Census, the TopDown Algorithm, does not satisfy any standard DP specification.
- Instead, we prove that
  - our swapping algorithm satisfies ε-DP, subject to the invariants it induces; and

Comparisons: US Decennial Censuses

- We provide a data swapping algorithm which is reminiscent of the statistical disclosure method used in the 1990, 2000 and 2010 US Decennial Censuses.
- This method does not satisfy the original DP specification – pure ε-DP (Dwork, McSherry, et al., 2006) – because it leaves *invariant* some statistics.
- For the same reason, the principal method used in the 2020 Census, the TopDown Algorithm, does not satisfy any standard DP specification.
- Instead, we prove that
  - our swapping algorithm satisfies ε-DP, subject to the invariants it induces; and
  - TopDown satisfies ρ-zCDP (Bun & Steinke, 2016b), subject to its invariants.

Comparisons: US Decennial Censuses

	d_{Pr}	$d_{\mathcal{X}}$ (Post-Imputation Unit)	Invariants ($\mathscr{D}$)	Privacy-Loss Budget
TopDown*	D_{nor}	d^p_{Ham} (person)	Population (state) Total housing units (block) Occupied group quarters (block) Structural zeros	PL & DHC: $\rho^2 = 15.29$ $\varepsilon = 52.83$ ($\delta = 10^{-10}$)
SafeTab**	D_{nor}	d^p_{Ham} (person)	None	DDHC-A: $\rho^2 = 19.776$ DDHC-B & S-DHC: *TBD.*
Swapping	$d_{\textsc{Mult}}$	d^h_{Ham} (household)	Varies but much greater than TDA	ε between 9.37-19.38

*(J. Abowd et al., 2022) **(Tumult Labs, 2022)

- $\mathcal{X}$ is always the space of possible Census Edited Files, $\mathcal{X}_{\mathrm{CEF}}$.
- $D_{\mathrm{nor}}(P, Q) = \sup_{\alpha > 1} \frac{1}{\sqrt{\alpha}} \max \left[\sqrt{D_\alpha(P||Q)}, \sqrt{D_\alpha(Q||P)} \right]$ is the normalised Rényi metric [zero concentrated DP] (with D_α the Rényi divergence of order);
- $d_{\textsc{Mult}}(P, Q) = \sup_{S \in \mathcal{F}} \left| \ln \frac{P(S)}{Q(S)} \right|$ is the multiplicative distance (pure DP); and
- d^u_{Ham} is the Hamming distance on units u (with p = post-imputation person, h = post-imputation household).
- $\mathscr{D}$ is the invariant-induced multiverse $\mathscr{D}_{\boldsymbol{c}} = \left\{ \{ \boldsymbol{x}' \in \mathcal{X}' : \boldsymbol{c}(\boldsymbol{x}) = \boldsymbol{c}(\boldsymbol{x}') \} : \boldsymbol{x} \in \mathcal{X} \right\}$.

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{\mathrm{P}_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{\mathrm{P}_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$
   - Conceives of data privacy as a limit on *probabilistic inference*

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{\mathrm{P}_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$
   - Conceives of data privacy as a limit on *probabilistic inference*
   - Focus on two aspects of forward-looking harms: the probability and strength of an *inferential, individual-based disclosure* (IIB)

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{P_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$
   - Conceives of data privacy as a limit on *probabilistic inference*
   - Focus on two aspects of forward-looking harms: the probability and strength of an *inferential, individual-based disclosure* (IIB)
   - Assumes a way to quantify IIBs (e.g. via the privacy loss random variable)

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{P_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$
   - Conceives of data privacy as a limit on *probabilistic inference*
   - Focus on two aspects of forward-looking harms: the probability and strength of an *inferential, individual-based disclosure* (IIB)
   - Assumes a way to quantify IIBs (e.g. via the privacy loss random variable)
3. DP is not a holistic framework for assessing privacy

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{P_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$
   - Conceives of data privacy as a limit on *probabilistic inference*
   - Focus on two aspects of forward-looking harms: the probability and strength of an *inferential, individual-based disclosure* (IIB)
   - Assumes a way to quantify IIBs (e.g. via the privacy loss random variable)
3. DP is not a holistic framework for assessing privacy
   - The theory of DP brackets other privacy concerns

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{P_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$
   - Conceives of data privacy as a limit on *probabilistic inference*
   - Focus on two aspects of forward-looking harms: the probability and strength of an *inferential, individual-based disclosure* (IIB)
   - Assumes a way to quantify IIBs (e.g. via the privacy loss random variable)
3. DP is not a holistic framework for assessing privacy
   - The theory of DP brackets other privacy concerns
   - The practice of DP is often left stranded

DP's Framing of Data Privacy (Seeman & Susser, 2023)

1. DP is a condition on the statistic T:
   - Conceives of data privacy as robustness
   - Focuses on *forward-looking, individual-based harms*
2. (More exactly) DP is a restriction on the *data-release model* $\{\mathrm{P}_{\boldsymbol{x}} : \boldsymbol{x} \in \mathcal{X}\}$
   - Conceives of data privacy as a limit on *probabilistic inference*
   - Focus on two aspects of forward-looking harms: the probability and strength of an *inferential, individual-based disclosure* (IIB)
   - Assumes a way to quantify IIBs (e.g. via the privacy loss random variable)
3. DP is not a holistic framework for assessing privacy
   - The theory of DP brackets other privacy concerns
   - The practice of DP is often left stranded
   - DP needs to be integrated into broader theories of privacy (Benthall & Cummings, 2024)

Data Swapping Visualisation

State	Location	Number of adults	Number of children	Age1	Race1	$\cdots$
MA	Cambridge	2	2	45	White	$\cdots$
TX	Houston	1	0	28	Hispanic	$\cdots$
WA	Tacoma	5	0	67	Asian	$\cdots$
MA	Somerville	2	2	50	Black	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

Data Swapping Visualisation

State	Location	Number of adults	Number of children	Age1	Race1	$\cdots$
MA	Cambridge	2	2	45	White	$\cdots$
TX	Houston	1	0	28	Hispanic	$\cdots$
WA	Tacoma	5	0	67	Asian	$\cdots$
MA	Somerville	2	2	50	Black	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

$\mathbf{V}_{\text{Stratify}}$

Data Swapping Visualisation

State	Location	Number of adults	Number of children	Age1	Race1	⋯
MA	Cambridge	2	2	45	White	⋯
TX	Houston	1	0	28	Hispanic	⋯
WA	Tacoma	5	0	67	Asian	⋯
MA	Somerville	2	2	50	Black	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

$\mathbf{V}_{\text{Stratify}}$

Data Swapping Visualisation

State	Location	Number of adults	Number of children	Age1	Race1	⋯
MA	Cambridge	2	2	45	White	⋯
TX	Houston	1	0	28	Hispanic	⋯
WA	Tacoma	5	0	67	Asian	⋯
MA	Somerville	2	2	50	Black	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

$\mathbf{V}_{\text{Stratify}}$
$\mathbf{V}_{\text{Swap}}$

Data Swapping Visualisation

State	Location	Number of adults	Number of children	Age1	Race1	⋯
MA	Somerville	2	2	45	White	⋯
TX	Houston	1	0	28	Hispanic	⋯
WA	Tacoma	5	0	67	Asian	⋯
MA	Cambridge	2	2	50	Black	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

$\mathbf{V}_{\text{Stratify}}$

$\mathbf{V}_{\text{Swap}}$

Data Swapping Visualisation

State	Location	Number of adults	Number of children	Age1	Race1	$\cdots$
MA	Somerville	2	2	45	White	$\cdots$
TX	Houston	1	0	28	Hispanic	$\cdots$
WA	Tacoma	5	0	67	Asian	$\cdots$
MA	Cambridge	2	2	50	Black	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

$\boldsymbol{V}_{\text{Stratify}}$

$\boldsymbol{V}_{\text{Swap}}$

$\boldsymbol{V}_{\text{Rest}}$

Data Swapping Visualisation

Massachusetts: Location by Race (head of household) Contingency Table

	White	Hispanic	Asian	Black	...
Boston					
Cambridge					
Brookline					
Somerville					
Watertown					
⋮					

Data Swapping Visualisation

Massachusetts: Location by Race (head of household) Contingency Table

	White	Hispanic	Asian	Black	⋯
Boston					
Cambridge	-1			+1	
Brookline					
Somerville	+1			-1	
Watertown					
⋮					

Data Swapping Visualisation

Massachusetts: Location by Race (head of household) Contingency Table

	White	Hispanic	Asian	Black	$\cdots$
Boston					
Cambridge	-1			+1	
Brookline					
Somerville	+1			-1	
Watertown					
$\vdots$					

Changes: Interior cells of $\mathbf{V}_{\text{Rest}} \times \mathbf{V}_{\text{Swap}}$.

Data Swapping Visualisation

Massachusetts: Location by Race (head of household) Contingency Table

	White	Hispanic	Asian	Black	$\cdots$
Boston					
Cambridge	-1			+1	
Brookline					
Somerville	+1			-1	
Watertown					
$\vdots$					

Changes: Interior cells of $\mathbf{V}_{\text{Rest}} \times \mathbf{V}_{\text{Swap}}$.
Invariants:

1. $\mathbf{V}_{\text{Stratify}} \times \mathbf{V}_{\text{Rest}}$
2. $\mathbf{V}_{\text{Stratify}} \times \mathbf{V}_{\text{Swap}}$

Does Swapping Satisfy Differential Privacy?

- Not under the traditional formulation of DP…

Does Swapping Satisfy Differential Privacy?

- Not under the traditional formulation of DP...
- Because swapping has *invariants* $\mathbf{c}_{Swap}$ – functions of the observed data which are released without noise.

Does Swapping Satisfy Differential Privacy?

- Not under the traditional formulation of DP...
- Because swapping has *invariants* $\mathbf{c}_{Swap}$ – functions of the observed data which are released without noise.

If a mechanism T contains an invariant (and x, x' have different values for this invariant), then $\mathrm{P}_{\boldsymbol{x}}$ and $\mathrm{P}_{\boldsymbol{x}'}$ do not have common support, and so

$$d_{\textsc{Mult}}\left[\mathrm{P}_{\boldsymbol{x}}, \mathrm{P}_{\boldsymbol{x}'}\right] = D_{\mathrm{nor}}\left[\mathrm{P}_{\boldsymbol{x}}, \mathrm{P}_{\boldsymbol{x}'}\right] = \infty.$$

Does the 2020 US Census Satisfy Differential Privacy?

- Not under the traditional formulation of DP…

Does the 2020 US Census Satisfy Differential Privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\text{TDA}}$.– functions of the observed data which are released without any noise added.

Does the 2020 US Census Satisfy Differential Privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\mathrm{TDA}}$.– functions of the observed data which are released without any noise added.

Modifying the definition of DP:

$$d_{\mathsf{Pr}}\left[\mathsf{P}_{\boldsymbol{x}}, \mathsf{P}_{\boldsymbol{x}'}\right] \leq \varepsilon \, d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}').$$

for all possible data values $\boldsymbol{x}, \boldsymbol{x}'$

Does the 2020 US Census Satisfy Differential Privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\mathrm{TDA}}$.– functions of the observed data which are released without any noise added.

Modifying the definition of DP:

$$d_{\mathsf{Pr}}\left[\mathsf{P}_{\boldsymbol{x}}, \mathsf{P}_{\boldsymbol{x}'}\right] \leq \varepsilon \, d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}').$$

for all possible data values $\boldsymbol{x}, \boldsymbol{x}'$ which agree on the invariants.

Does the 2020 US Census Satisfy Differential Privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\mathrm{TDA}}$.– functions of the observed data which are released without any noise added.

Modifying the definition of DP:

$$d_{\mathsf{Pr}}\left[\mathsf{P}_{\boldsymbol{x}}, \mathsf{P}_{\boldsymbol{x}'}\right] \leq \varepsilon \, d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}').$$

for all possible data values $\boldsymbol{x}, \boldsymbol{x}'$ which agree on the invariants.

- This is a necessary and sufficient modification for the release of invariants.

Swapping Satisfies DP, Subject to Its Invariants

Permutation swapping

Input: a dataset $\boldsymbol{x}$.
Define strata as groups of records which match on the swap key $\mathbf{V}_{\text{Stratify}}$.
Within each stratum:

1. Select each record independently with probability p (the swap rate).
2. Randomly permute swapping variable $\mathbf{V}_{\text{Swap}}$ of selected records.

Output: the *swapped* dataset $\boldsymbol{w}$.

Swapping Satisfies DP, Subject to Its Invariants

Permutation swapping

Input: a dataset $\boldsymbol{x}$.
Define strata as groups of records which match on the swap key $\mathbf{V}_{\text{Stratify}}$.
Within each stratum:

1. Select each record independently with probability p (the swap rate).
2. Randomly permute swapping variable $\mathbf{V}_{\text{Swap}}$ of selected records.

Output: the *swapped* dataset $\boldsymbol{w}$.

Permutation swapping is DP subject to its invariants, with input divergence $d_{\mathcal{X}} = d^{u}_{\text{Ham}}$, output divergence $d_{\text{Pr}} = d_{\text{MULT}}$ and budget

$$\varepsilon = \begin{cases} \ln(b+1) - \ln o & \text{if } 0 < p \leq 0.5, \\ \max\left\{\ln o, \ln(b+1) - \ln o\right\} & \text{if } 0.5 < p < 1, \end{cases}$$

where $o = p/(1-p)$ and b is the maximum stratum size.

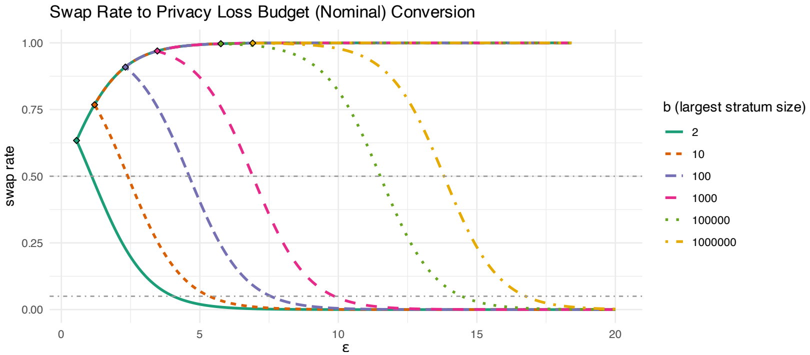

Conversion between the swap rate (p) and the nominal PLB (ε) at different levels of b. Note that:

1. For each b, there's a **smallest attainable** $\varepsilon_b > 0$;
2. For each b, every $\varepsilon > \varepsilon_b$ is satisfied by **two** different swap rates;
3. (counterintuitive) For the same swap rate, the larger the b, the **larger** the ε!

The TopDown Algorithm (TDA) (J. Abowd et al., 2022)

Two-step procedure:

0. Start with a Census edited file $\boldsymbol{x} \in \mathcal{X}_{\mathrm{CEF}}$.

The TopDown Algorithm (TDA) (J. Abowd et al., 2022)

Two-step procedure:

0. Start with a Census edited file $\boldsymbol{x} \in \mathcal{X}_{\mathrm{CEF}}$.
1. Add Gaussian noise to cells:

$$\boldsymbol{T}(\boldsymbol{x}) = \boldsymbol{q}(\boldsymbol{x}) + \boldsymbol{W},$$

where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \boldsymbol{\Sigma})$, so that $\boldsymbol{T}$ satisfies $\mathrm{DP}(\mathcal{X}_{\mathrm{CEF}}, \{\mathcal{X}_{\mathrm{CEF}}\}, d^{p}_{\mathrm{Ham}}, D_{\mathrm{nor}})$ with budget ρ_{TDA} (Canonne et al., 2022).

The TopDown Algorithm (TDA) (J. Abowd et al., 2022)

Two-step procedure:

0. Start with a Census edited file $\boldsymbol{x} \in \mathcal{X}_{\text{CEF}}$.
1. Add Gaussian noise to cells:

$$\boldsymbol{T}(\boldsymbol{x}) = \boldsymbol{q}(\boldsymbol{x}) + \boldsymbol{W},$$

   where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \boldsymbol{\Sigma})$, so that $\boldsymbol{T}$ satisfies $\mathrm{DP}(\mathcal{X}_{\text{CEF}}, \{\mathcal{X}_{\text{CEF}}\}, d^{p}_{\text{Ham}}, D_{\text{nor}})$ with budget ρ_{TDA} (Canonne et al., 2022).

2. "Post-process": find dataset $\boldsymbol{z}$ with $\boldsymbol{q}(\boldsymbol{z})$ close to $\boldsymbol{T}(\boldsymbol{x})$ such that $\boldsymbol{c}_{\text{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\text{TDA}}(\boldsymbol{x})$.

The TopDown Algorithm (TDA) (J. Abowd et al., 2022)

Two-step procedure:

0. Start with a Census edited file $\boldsymbol{x} \in \mathcal{X}_{\mathrm{CEF}}$.
1. Add Gaussian noise to cells:

$$\boldsymbol{T}(\boldsymbol{x}) = \boldsymbol{q}(\boldsymbol{x}) + \boldsymbol{W},$$

   where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \boldsymbol{\Sigma})$, so that $\boldsymbol{T}$ satisfies $\mathrm{DP}(\mathcal{X}_{\mathrm{CEF}}, \{\mathcal{X}_{\mathrm{CEF}}\}, d^{p}_{\mathrm{Ham}}, D_{\mathrm{nor}})$ with budget ρ_{TDA} (Canonne et al., 2022).
2. "Post-process": find dataset $\boldsymbol{z}$ with $\boldsymbol{q}(\boldsymbol{z})$ close to $\boldsymbol{T}(\boldsymbol{x})$ such that $\boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{x})$.

The TopDown Algorithm (TDA) (J. Abowd et al., 2022)

Two-step procedure:

0. Start with a Census edited file $\boldsymbol{x} \in \mathcal{X}_{\mathrm{CEF}}$.
1. Add Gaussian noise to cells:

$$\boldsymbol{T}(\boldsymbol{x}) = \boldsymbol{q}(\boldsymbol{x}) + \boldsymbol{W},$$

   where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \boldsymbol{\Sigma})$, so that $\boldsymbol{T}$ satisfies $\mathrm{DP}(\mathcal{X}_{\mathrm{CEF}}, \{\mathcal{X}_{\mathrm{CEF}}\}, d_{\mathrm{Ham}}^{p}, D_{\mathrm{nor}})$ with budget ρ_{TDA} (Canonne et al., 2022).
2. "Post-process": find dataset $\boldsymbol{z}$ with $\boldsymbol{q}(\boldsymbol{z})$ close to $\boldsymbol{T}(\boldsymbol{x})$ such that $\boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{x})$.

TDA satisfies $\mathrm{DP}(\mathcal{X}_{\mathrm{CEF}}, \mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}}, d_{\mathrm{Ham}}^{p}, D_{\mathrm{nor}})$ with budget ρ_{TDA}.

Theorem: TDA Satisfies DP, Subject to Its Invariants

Let $\boldsymbol{c}_{\mathrm{TDA}} : \mathcal{X}_{\mathrm{CEF}} \to \mathbb{R}^l$ be the invariants of TDA and let $\mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}}$ be the induced data multiverse:

$$\mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}} = \{\mathcal{D} \subset \mathcal{X}_{\mathrm{CEF}} \mid \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{x}) = \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{x}') \; \forall \boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}\}.$$

Theorem: TDA Satisfies DP, Subject to Its Invariants

Let $\boldsymbol{c}_{\mathrm{TDA}} : \mathcal{X}_{\mathrm{CEF}} \to \mathbb{R}^l$ be the invariants of TDA and let $\mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}}$ be the induced data multiverse:

$$\mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}} = \{\mathcal{D} \subset \mathcal{X}_{\mathrm{CEF}} \mid \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{x}) = \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{x}') \; \forall \boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}\}.$$

- TDA satisfies DP$(\mathcal{X}_{\mathrm{CEF}}, \mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}}, d^p_{\mathrm{Ham}}, D_{\mathrm{nor}})$ with privacy budget $\rho_{\mathrm{TDA}} = 2.63$ (for the PL Redistricting File) and $\rho_{\mathrm{TDA}} = 15.29$ (for the DHC).

Theorem: TDA Satisfies DP, Subject to Its Invariants

Let $\boldsymbol{c}_{\text{TDA}} : \mathcal{X}_{\text{CEF}} \to \mathbb{R}^l$ be the invariants of TDA and let $\mathscr{D}_{\boldsymbol{c}_{\text{TDA}}}$ be the induced data multiverse:

$$\mathscr{D}_{\boldsymbol{c}_{\text{TDA}}} = \{\mathcal{D} \subset \mathcal{X}_{\text{CEF}} \mid \boldsymbol{c}_{\text{TDA}}(\boldsymbol{x}) = \boldsymbol{c}_{\text{TDA}}(\boldsymbol{x}') \; \forall \boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}\}.$$

- TDA satisfies $\text{DP}(\mathcal{X}_{\text{CEF}}, \mathscr{D}_{\boldsymbol{c}_{\text{TDA}}}, d^p_{\text{Ham}}, D_{\text{nor}})$ with privacy budget $\rho_{\text{TDA}} = 2.63$ (for the PL Redistricting File) and $\rho_{\text{TDA}} = 15.29$ (for the DHC).
- Let $\boldsymbol{c}'$ be any proper subset of TDA's invariants. TDA does not satisfy $\text{DP}(\mathcal{X}_{\text{CEF}}, \mathscr{D}_{\boldsymbol{c}'}, d_{\mathcal{X}}, D_{\text{nor}})$ with any finite budget ρ.

What if the 2020 Census Used Swapping?

The total nominal ε achievable by applying swapping to the 2020 Decennial Census for a variety of $\mathbf{V}_{\text{Stratify}}$, $\mathbf{V}_{\text{Swap}}$, and swap rate choices.

$\mathbf{V}_{\text{Stratify}}$	$\mathbf{V}_{\text{Swap}}$	b	total ε $p = 5\%$	total ε $p = 50\%$	Largest stratum
state	county	13680081	19.38	16.43	California
state $\times$ household size	county	3653802	18.06	15.11	California, 3-household
county	tract	3445076	18.00	15.05	LA County
county $\times$ household size	tract	853003	16.60	13.66	LA County, 3-household
block group	block	21535	12.92	9.98	a FL block group
block group $\times$ household size	block	11691	12.31	9.37	a FL block group, 3-household

Note. For a fixed ($\mathbf{V}_{\text{Stratify}}$, $\mathbf{V}_{\text{Swap}}$, p) setting, the nominal ε would be the **total PLB** for all data products derived from the swapped dataset, including P.L. 94-171, DHC, Detailed DHC for both persons and household product types.

Permutation Swapping

Input: Dataset $\boldsymbol{X}$

1: **for** $j = 1, \ldots, \mathcal{J}$ **do**
2:     **if** $n_j = 0$ or $n_j = 1$ **then**
3:         **continue**
4:     **end if**
5:     **for** record i with category j **do**
6:         Select i with probability p
7:     **end for**
8:     **if** 0 records selected **then**
9:         **continue**
10:     **else if** exactly 1 record selected **then**
11:         **go to** line 5
12:     **end if**
13:     Sample uniformly at random a derangement σ of the selected records.
14:     /* *Permute the swapping variable of the selected records according to σ:* */
15:         Save copy $\boldsymbol{X}_0 \leftarrow \boldsymbol{X}$ before permutation
16:         Let $k^{\boldsymbol{X}}(i)$ be the value of the swapping variable of record i in dataset $\boldsymbol{X}$.
17:         **for all** selected records i **do**
18:             Set $k^{\boldsymbol{X}}(i) \leftarrow k^{\boldsymbol{X}_0}(\sigma(i))$
19:         **end for**
20: **end for**
21: Set $\boldsymbol{Z} \leftarrow \boldsymbol{X}$ to be the swapped dataset.
22: **return** contingency table $[n_{jkl}^{\boldsymbol{Z}}]$

Intuition of the Proof that Permutation Swapping Is DP

1. We need to show that, for fixed datasets $\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{w}$ in the same data universe $\mathcal{D}$,

$$\mathsf{Pr}(\sigma(\boldsymbol{x}) = \boldsymbol{w}) \leq \exp(d^u_{\text{Ham}}(\boldsymbol{x}, \boldsymbol{x}')\varepsilon)\, \mathsf{Pr}(\sigma'(\boldsymbol{x}') = \boldsymbol{w}),$$

2. We can show that there exists a derangement ρ of m records such that $\boldsymbol{x} = \rho(\boldsymbol{x}')$.
3. There is a bijection between the possible σ and σ' given by $\sigma' = \sigma \circ \rho$.
4. Hence, if m_σ is the number of records deranged by σ, we have

$$m_\sigma - m \leq m_{\sigma'} \leq m_\sigma + m.$$

5. This gives a bound on $\mathsf{Pr}(\sigma)/Pr(\sigma')$ in terms of $o^{m_\sigma - m_{\sigma'}}$ and the ratio between the number of derangements of $m_{\sigma'}$ and of m_σ.
6. For $o \leq 1$, this can be bounded by $o^{-m}(b+1)^m$ using the above inequality. The result for $0 < p \leq 0.5$ then follows with some algebraic simplification.

The TopDown Algorithm (J. Abowd et al., 2022)

Input:

Census Edited Files $\boldsymbol{X}_p, \boldsymbol{X}_h$ at the person and household levels
Person queries $\boldsymbol{Q}_p$
Household queries $\boldsymbol{Q}_h$
Privacy noise scales $\boldsymbol{D}_p$ and $\boldsymbol{D}_h$
Constraints $\boldsymbol{c}_{\text{TDA}}$ (including invariants, edit constraints and structural zeroes)
(Optional) previously released statistics $\boldsymbol{P}$, as aggregated from a microdata file (where the aggregation was achieved using a function $\boldsymbol{H}$)

```
Step 1: Noise Infusion
    Sample discrete Gaussian noise
        W_p ~ N_Z(0, D_p)
        W_h ~ N_Z(0, D_h)
    Compute Noisy Measurement Files:
        T_p(X_p) ← Q_p(X_p) + W_p
        T_h(X_h) ← Q_h(X_h) + W_h
Step 2: Post-Processing
    Compute Privacy-Protected Microdata Files Z_p, Z_h as a solution to the optimisation
    problem:
        Minimize loss l between [T_p(X_p), T_h(X_h)] and [Q_p(Z_p), Q_h(Z_h)]
            subject to constraints c_TDA(Z_p, Z_h) = c_TDA(X_p, X_h) and H(Z_p, Z_h) = P.
```

Step 1: Noise Infusion

- Sample discrete Gaussian noise
  - $\boldsymbol{W}_p \sim \mathcal{N}_{\mathbb{Z}}(\boldsymbol{0}, \boldsymbol{D}_p)$
  - $\boldsymbol{W}_h \sim \mathcal{N}_{\mathbb{Z}}(\boldsymbol{0}, \boldsymbol{D}_h)$
- Compute Noisy Measurement Files:
  - $\boldsymbol{T}_p(\boldsymbol{X}_p) \leftarrow \boldsymbol{Q}_p(\boldsymbol{X}_p) + \boldsymbol{W}_p$
  - $\boldsymbol{T}_h(\boldsymbol{X}_h) \leftarrow \boldsymbol{Q}_h(\boldsymbol{X}_h) + \boldsymbol{W}_h$

Step 2: Post-Processing

- Compute Privacy-Protected Microdata Files $\boldsymbol{Z}_p, \boldsymbol{Z}_h$ as a solution to the optimisation problem:
  - Minimize loss l between $[\boldsymbol{T}_p(\boldsymbol{X}_p), \boldsymbol{T}_h(\boldsymbol{X}_h)]$ and $[\boldsymbol{Q}_p(\boldsymbol{Z}_p), \boldsymbol{Q}_h(\boldsymbol{Z}_h)]$
    - subject to constraints $\boldsymbol{c}_{\text{TDA}}(\boldsymbol{Z}_p, \boldsymbol{Z}_h) = \boldsymbol{c}_{\text{TDA}}(\boldsymbol{X}_p, \boldsymbol{X}_h)$ and $\boldsymbol{H}(\boldsymbol{Z}_p, \boldsymbol{Z}_h) = \boldsymbol{P}$.

Output:

Privacy-Protected Microdata Files $\boldsymbol{Z}_p, \boldsymbol{Z}_h$, and
Noisy Measurement Files $\boldsymbol{T}_p(\boldsymbol{X}_p), \boldsymbol{T}_h(\boldsymbol{X}_h)$ at the person and household levels.

Examples of $\mathscr{D}$, $d_{\mathcal{X}}$ and d_{Pr}

1. An *invariant-compliant data universe*:

$$\mathscr{D}_{\boldsymbol{c}} = \left\{ \mathcal{D} \subset \mathcal{X} : \boldsymbol{c}(\boldsymbol{x}) = \boldsymbol{c}(\boldsymbol{x}') \ \forall \boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D} \right\},$$

for some invariants $\boldsymbol{c} : \mathcal{X} \to \mathbb{R}^l$.

Examples of $\mathscr{D}$, $d_{\mathcal{X}}$ and d_{Pr}

1. An *invariant-compliant data universe*:

$$\mathscr{D}_{\boldsymbol{c}} = \left\{ \mathcal{D} \subset \mathcal{X} : \boldsymbol{c}(\boldsymbol{x}) = \boldsymbol{c}(\boldsymbol{x}') \; \forall \boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D} \right\},$$

   for some invariants $\boldsymbol{c} : \mathcal{X} \to \mathbb{R}^l$.

2. *Data divergence* $d_{\mathcal{X}}$ induced by a “neighbour” relation:

$$d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}') = \begin{cases} 0 & \text{if } \boldsymbol{x} = \boldsymbol{x}', \\ 1 & \text{if } \boldsymbol{x} \text{ and } \boldsymbol{x}' \text{ are “neighbours”}, \\ \infty & \text{otherwise.} \end{cases}$$

Examples of $\mathscr{D}$, $d_{\mathcal{X}}$ and d_{Pr}

3. *Divergence* d_{Pr} on (the probability distributions over) the output space

Examples of $\mathscr{D}$, $d_{\mathcal{X}}$ and d_{Pr}

3. *Divergence* d_{Pr} on (the probability distributions over) the output space
   - *Pure ε-DP* (Dwork, McSherry, et al., 2006): d_{Pr} is the multiplicative distance

$$\textsc{Mult}(\mathsf{P}, \mathsf{Q}) = \sup \left\{ \left| \ln \frac{\mathsf{P}(S)}{\mathsf{Q}(S)} \right| : \text{event } S \right\}.$$

Examples of $\mathscr{D}$, $d_{\mathcal{X}}$ and d_{Pr}

3. *Divergence* d_{Pr} on (the probability distributions over) the output space
   - *Pure ε-DP* (Dwork, McSherry, et al., 2006): d_{Pr} is the multiplicative distance

$$\textsc{Mult}(\mathsf{P}, \mathsf{Q}) = \sup \left\{ \left| \ln \frac{\mathsf{P}(S)}{\mathsf{Q}(S)} \right| : \text{event } S \right\}.$$

   - *Approximate (ε, δ)-DP* (Dwork, Kenthapadi, et al., 2006):

$$\textsc{Mult}^{\delta}(\mathsf{P}, \mathsf{Q}) = \sup_{\text{event } S} \left\{ \ln \frac{[\mathsf{P}(S) - \delta]^+}{\mathsf{Q}(S)}, \ln \frac{[\mathsf{Q}(S) - \delta]^+}{\mathsf{P}(S)}, 0 \right\},$$

Examples of $\mathscr{D}$, $d_{\mathcal{X}}$ and d_{Pr}

3. *Divergence* d_{Pr} on (the probability distributions over) the output space
   - *Pure ε-DP* (Dwork, McSherry, et al., 2006): d_{Pr} is the multiplicative distance

     $$\textsc{Mult}(\mathsf{P}, \mathsf{Q}) = \sup\left\{\left|\ln \frac{\mathsf{P}(S)}{\mathsf{Q}(S)}\right| : \text{event } S\right\}.$$

   - *Approximate (ε, δ)-DP* (Dwork, Kenthapadi, et al., 2006):

     $$\textsc{Mult}^{\delta}(\mathsf{P}, \mathsf{Q}) = \sup_{\text{event } S}\left\{\ln \frac{[\mathsf{P}(S)-\delta]^{+}}{\mathsf{Q}(S)}, \ln \frac{[\mathsf{Q}(S)-\delta]^{+}}{\mathsf{P}(S)}, 0\right\},$$

   - *Zero Concentrated DP* (Bun & Steinke, 2016a):

     $$D_{\mathrm{nor}}(\mathsf{P}, \mathsf{Q}) = \sup_{\alpha > 1} \frac{1}{\sqrt{\alpha}} \max\left[\sqrt{D_{\alpha}(\mathsf{P}||\mathsf{Q})}, \sqrt{D_{\alpha}(\mathsf{Q}||\mathsf{P})}\right],$$

     where D_{α} is the *Rényi divergence* of order α:

     $$D_{\alpha}(\mathsf{P}||\mathsf{Q}) = \frac{1}{\alpha - 1} \ln \int \left[\frac{d\mathsf{P}}{d\mathsf{Q}}\right]^{\alpha} d\mathsf{Q},$$

Numerical demonstration: 1940 Census full count data

- $\mathbf{V}_{\text{Swap}}$: household's county;
- $\mathbf{V}_{\text{Stratify}}$ (swap key): the number of persons per household $\times$ household's state;
- $\mathbf{V}_{\text{Hold}} - \mathbf{V}_{\text{Stratify}}$: dwelling ownership.

The invariants $\mathbf{c}_{\text{Swap}}$ are

1. Total *number of owned vs rented dwellings* at each household size at the state level;
2. Total *number of dwellings* at each household size at the county level.

swap rate	0.01	0.05	0.10	0.50
ε	17.08	15.43	14.68	12.48

Table: Conversion of swap rate to ε (PLB). Under this swapping scheme, the largest stratum size is $b = 264,331$, the number of all two-person households of Massachusetts.

Numerical Demonstration: 1940 Census Full Count Data

Table: Two-way tabulations of dwelling ownership by county based on the 1940 Census full count for Massachusetts (left) and one instantiation of the Permutation Algorithm at $p = 50\%$ (right). Total dwellings per county, as well as total owned versus rented units per state, are invariant. All invariants induced by the Algorithm are not shown.

county	owned	rented	total	owned (swapped)	rented (swapped)	total (swapped)
Barnstable	7461	3825	11286	5907	5379	11286
Berkshire	14736	18417	33153	13770	19383	33153
Bristol	33747	63931	97678	35537	62141	97678
Dukes	1207	534	1741	946	795	1741
Essex	53936	81300	135236	52631	82605	135236
Franklin	7433	6442	13875	6337	7538	13875
Hampden	30597	58166	88763	32267	56496	88763
Hampshire	9427	8630	18057	8145	9912	18057
Middlesex	104144	147687	251831	100372	151459	251831
Nantucket	593	432	1025	471	554	1025
Norfolk	44885	40285	85170	38566	46604	85170
Plymouth	24857	23882	48739	21549	27190	48739
Suffolk	49656	176553	226209	67357	158852	226209
Worcester	53126	78535	131661	51950	79711	131661
total	435805	708619	1144424	435805	708619	1144424

Numerical Demonstration: 1940 Census Full Count Data

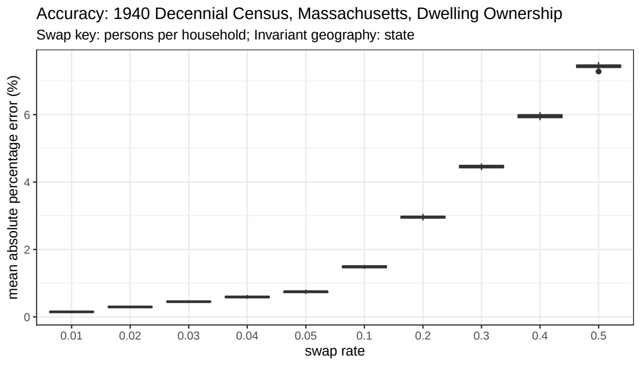

Figure: Mean absolute percentage error (MAPE) in the two-way tabulation of dwelling ownership by county induced by the Permutation Algorithm applied to the 1940 Census full count data of Massachusetts, at different swap rates from 1% to 50%. Each boxplot reflects 20 independent runs of the Algorithm at that swap rate.

Extending "Neighbour" Divergences to Metrics on $\mathcal{X}$

A divergence defined by neighbours:

$$d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}') = \begin{cases} 0 & \text{if } \boldsymbol{x} = \boldsymbol{x}', \\ 1 & \text{if } \boldsymbol{x} \text{ and } \boldsymbol{x}' \text{ are ``neighbours''}, \\ \infty & \text{otherwise}, \end{cases}$$

can always be sharpened to a metric $d^*_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}')$ defined as the length of a shortest path between $\boldsymbol{X}$ and $\boldsymbol{X}'$ in the graph on $\mathcal{X}$ with edges given by r. For example the extension of the bounded-neighbours is the Hamming distance on unordered datasets:

$$d^u_{\mathrm{Ham}}(\boldsymbol{x}, \boldsymbol{x}') = \begin{cases} \frac{1}{2}|\boldsymbol{x} \ominus \boldsymbol{x}'| & \text{if } |\boldsymbol{x}| = |\boldsymbol{x}|, \\ \infty & \text{otherwise} \end{cases}$$

and the extension of unbounded-neighbours is the symmetric difference distance:

$$d^u_{\mathrm{SymDiff}}(\boldsymbol{X}, \boldsymbol{X}') = \left|\boldsymbol{X} \ominus \boldsymbol{X}'\right|.$$

The superscript ^{u} emphasizes that these distances are defined with respect to a choice of the privacy unit u.

Sufficiency and Necessity of Restricting the Data Universe $\mathcal{D}$

1. For any $d_{\mathcal{X}}$ and d_{Pr}, the mechanism $T(\boldsymbol{x}) = \boldsymbol{c}(\boldsymbol{x})$ that *releases the invariants exactly* satisfies $(\mathcal{X}, \mathscr{D}_{\boldsymbol{c}}, d_{\mathcal{X}}, d_{\mathsf{Pr}})$ with *privacy budget* $\varepsilon_{\mathcal{D}} = 0$.

2. Now suppose $d_{\mathsf{Pr}}(\mathrm{P}, \mathrm{Q}) = \infty$ if $d_{\mathrm{TV}}(\mathrm{P}, \mathrm{Q}) = 1$. Let $\mathscr{D}$ be a data multiverse such that there exists datasets $\boldsymbol{x}_1, \boldsymbol{x}_2$ in some data universe $\mathcal{D}_0 \in \mathscr{D}$ with $d_{\mathcal{X}}(\boldsymbol{x}_1, \boldsymbol{x}_2) < \infty$ and $\boldsymbol{c}(\boldsymbol{x}_1) \neq \boldsymbol{c}(\boldsymbol{x}_2)$. Then *$T$ does not satisfy $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathsf{Pr}})$ for any $\varepsilon_{\mathcal{D}_0} < \infty$*.

3. Suppose that a mechanism T varies within some universe $\mathcal{D}_0 \in \mathscr{D}_{\boldsymbol{c}}$ in the sense that there exists $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}_0$ with $d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}') < \infty$ but $\mathrm{P}_{\boldsymbol{x}} \neq \mathrm{P}_{\boldsymbol{x}'}$.
When d_{Pr} is a metric, *T satisfies $(\mathcal{X}, \mathscr{D}_{\boldsymbol{c}}, d_{\mathcal{X}}, d_{\mathsf{Pr}})$ only if $\varepsilon_{\mathcal{D}_0} > 0$*.

References I

Abowd, J., Ashmead, R., Cumings-Menon, R., Garfinkel, S., Heineck, M., Heiss, C., … Zhuravlev, P. (2022, June). The 2020 Census disclosure avoidance system TopDown algorithm. *Harvard Data Science Review*(Special Issue 2). doi: 10.1162/99608f92.529e3cb9

Abowd, J. M., Schneider, M. J., & Vilhuber, L. (2013). Differential privacy applications to Bayesian and linear mixed model estimation. *Journal of Privacy and Confidentiality*, *5*(1).

Andrés, M. E., Bordenabe, N. E., Chatzikokolakis, K., & Palamidessi, C. (2013, November). Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security* (pp. 901–914). New York, NY, USA: Association for Computing Machinery. doi: 10.1145/2508859.2516735

Ashmead, R., Kifer, D., Leclerc, P., Machanavajjhala, A., & Sexton, W. (2019). *Effective privacy after adjusting for invariants with applications to the 2020 Census* (Tech. Rep.). https://github.com/uscensusbureau/census2020-das-e2e/blob/master/doc/20190711_0941_Effective_Privacy_after_Adjusting_for_Constraints__With_applications_to_the_2020_Census.pdf.

Asi, H., Duchi, J. C., & Javidbakht, O. (2022). Element level differential privacy: The right granularity of privacy. In *AAAI Workshop on Privacy-Preserving Artificial Intelligence.* Association for the Advancement of Artificial Intelligence.

Balle, B., Barthe, G., & Gaboardi, M. (2020, January). Privacy profiles and amplification by subsampling. *Journal of Privacy and Confidentiality*, *10*(1). doi: 10.29012/jpc.726

Barber, R. F., & Duchi, J. C. (2014, December). *Privacy and statistical risk: Formalisms and minimax bounds* (No. arXiv:1412.4451). http://arxiv.org/abs/1412.4451. arXiv. doi: 10.48550/arXiv.1412.4451

Barthe, G., & Olmedo, F. (2013). Beyond differential privacy: Composition theorems and relational logic for f-divergences between probabilistic programs. In F. V. Fomin, R. Freivalds, M. Kwiatkowska, & D. Peleg (Eds.), *Automata, languages, and programming* (pp. 49–60). Berlin, Heidelberg: Springer. doi: 10.1007/978-3-642-39212-2_8

Beimel, A., Kasiviswanathan, S. P., & Nissim, K. (2010, February). Bounds on the sample complexity for private learning and private data release. In D. Micciancio (Ed.), *Proceedings of the 7th theory of cryptography conference, TCC 2010, Zurich, Switzerland* (pp. 437–454). Berlin, Heidelberg: Springer. doi: 10.1007/978-3-642-11799-2_26

Benthall, S., & Cummings, R. (2024, January). *Integrating differential privacy and contextual integrity* (No. arXiv:2401.15774). http://arxiv.org/abs/2401.15774. arXiv. doi: 10.48550/arXiv.2401.15774

Bhaskar, R., Bhowmick, A., Goyal, V., Laxman, S., & Thakurta, A. (2011). Noiseless database privacy. In D. H. Lee & X. Wang (Eds.), *Advances in cryptology – ASIACRYPT 2011* (pp. 215–232). Berlin, Heidelberg: Springer. doi: 10.1007/978-3-642-25385-0_12

References II

Bun, M., Drechsler, J., Gaboardi, M., McMillan, A., & Sarathy, J. (2022, June). Controlling privacy loss in sampling schemes: An analysis of stratified and cluster sampling. In *Foundations of Responsible Computing (FORC 2022)* (p. 24).

Bun, M., & Steinke, T. (2016a). Concentrated differential privacy: Simplifications, extensions, and lower bounds. In M. Hirt & A. Smith (Eds.), *Theory of cryptography* (pp. 635–658). Berlin, Heidelberg: Springer. doi: 10.1007/978-3-662-53641-4_24

Bun, M., & Steinke, T. (2016b, May). *Concentrated differential privacy: Simplifications, extensions, and lower bounds* (No. arXiv:1605.02065). arXiv. doi: 10.48550/arXiv.1605.02065

Canonne, C., Kamath, G., & Steinke, T. (2022, July). The discrete Gaussian for differential privacy. *Journal of Privacy and Confidentiality*, *12*(1). doi: 10.29012/jpc.784

Charest, A.-S., & Hou, Y. (2016). On the meaning and limits of empirical differential privacy. *Journal of Privacy and Confidentiality*, *7*(3), 53–66.

Chatzikokolakis, K., Andrés, M. E., Bordenabe, N. E., & Palamidessi, C. (2013). Broadening the Scope of Differential Privacy Using Metrics. In E. De Cristofaro & M. Wright (Eds.), *Privacy Enhancing Technologies* (pp. 82–102). Berlin, Heidelberg: Springer. doi: 10.1007/978-3-642-39077-7_5

Dharangutte, P., Gao, J., Gong, R., & Yu, F.-Y. (2023). Integer subspace differential privacy. In *Proceedings of the aaai conference on artificial intelligence (aaai-23).*

Dong, J., Roth, A., & Su, W. J. (2022). Gaussian differential privacy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *84*(1), 3–37. doi: 10.1111/rssb.12454

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., & Naor, M. (2006). Our data, ourselves: Privacy via distributed noise generation. In S. Vaudenay (Ed.), *Advances in cryptology - EUROCRYPT 2006* (pp. 486–503). Berlin, Heidelberg: Springer. doi: 10.1007/11761679_29

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference* (pp. 265–284).

Dwork, C., Naor, M., Pitassi, T., & Rothblum, G. N. (2010, June). Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing* (pp. 715–724). New York, NY, USA: Association for Computing Machinery. (`https://dl.acm.org/doi/10.1145/1806689.1806787`) doi: 10.1145/1806689.1806787

Ebadi, H., Sands, D., & Schneider, G. (2015, January). Differential Privacy: Now it's Getting Personal. *ACM SIGPLAN Notices*, *50*(1), 69–81. doi: 10.1145/2775051.2677005

Feldman, V., & Zrnic, T. (2022, January). *Individual privacy accounting via a Rényi filter* (No. arXiv:2008.11193). `http://arxiv.org/abs/2008.11193`. arXiv.

Gao, J., Gong, R., & Yu, F.-Y. (2022, June). Subspace differential privacy. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 36, pp. 3986–3995). doi: 10.1609/aaai.v36i4.20315

References III

Gong, R., & Meng, X.-L. (2020). Congenial differential privacy under mandated disclosure. In *Proceedings of the 2020 acm-ims on foundations of data science conference* (pp. 59–70).

Hay, M., Li, C., Miklau, G., & Jensen, D. (2009, December). Accurate estimation of the degree distribution of private networks. In *2009 Ninth IEEE International Conference on Data Mining* (pp. 169–178). doi: 10.1109/ICDM.2009.11

He, X., Machanavajjhala, A., & Ding, B. (2014). Blowfish privacy: Tuning privacy-utility trade-offs using policies. In *Proceedings of the 2014 acm sigmod international conference on management of data* (pp. 1447–1458).

Jorgensen, Z., Yu, T., & Cormode, G. (2015, April). Conservative or liberal? Personalized differential privacy. In *2015 IEEE 31st International Conference on Data Engineering* (pp. 1023–1034). (`https://ieeexplore.ieee.org/document/7113353`) doi: 10.1109/ICDE.2015.7113353

Kifer, D., & Machanavajjhala, A. (2011). No free lunch in data privacy. In *Proceedings of the 2011 international conference on Management of data - SIGMOD '11* (pp. 193–204). Athens, Greece: ACM Press. doi: 10.1145/1989323.1989345

Kifer, D., & Machanavajjhala, A. (2014). Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems (TODS)*, *39*(1), 1–36.

McSherry, F., & Mahajan, R. (2010, August). Differentially-private network trace analysis. In *Proceedings of the ACM SIGCOMM 2010 conference* (pp. 123–134). New York, NY, USA: Association for Computing Machinery. doi: 10.1145/1851182.1851199

Mironov, I. (2017, August). Rényi differential privacy. *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, 263–275. doi: 10.1109/CSF.2017.11

O'Keefe, C. M., & Charest, A.-S. (2019). Bootstrap differential privacy. *Transactions on Data Privacy*, *12*, 1–28.

Redberg, R., & Wang, Y.-X. (2021). Privately publishable per-instance privacy. In *Advances in Neural Information Processing Systems* (Vol. 34, pp. 17335–17346). Curran Associates, Inc.

Seeman, J., Reimherr, M., & Slavkovic, A. (2022, May). *Formal privacy for partially private data* (No. arXiv:2204.01102). `http://arxiv.org/abs/2204.01102`. arXiv.

Seeman, J., Sexton, W., Pujol, D., & Machanavajjhala, A. (2023+). Per-record differential privacy: Modeling dependence between individual privacy loss and confidential records.

Seeman, J., & Susser, D. (2023, October). Between privacy and utility: On differential privacy in theory and practice. *ACM Journal on Responsible Computing*. (`https://dl.acm.org/doi/10.1145/3626494`) doi: 10.1145/3626494

References IV

Soria-Comas, J., Domingo-Ferrer, J., Sánchez, D., & Megías, D. (2017, June). Individual differential privacy: A utility-preserving formulation of differential privacy guarantees. *IEEE Transactions on Information Forensics and Security*, *12*(6), 1418–1429. doi: 10.1109/TIFS.2017.2663337

Tumult Labs. (2022, March). *SafeTab: DP algorithms for 2020 Census Detailed DHC Race & Ethnicity* (Tech. Rep.). https://www2.census.gov/about/partners/cac/sac/meetings/2022-03/dhc-attachment-1-safetab-dp-algorithms.pdf.

Wang, Y.-X. (2018, November). *Per-instance Differential Privacy* (No. arXiv:1707.07708). `http://arxiv.org/abs/1707.07708`. arXiv.

Zhou, S., Ligett, K., & Wasserman, L. (2009, June). Differential privacy with compression. In *Proceedings of the 2009 IEEE international conference on Symposium on Information Theory - Volume 4* (pp. 2718–2722). Coex, Seoul, Korea: IEEE Press.