Can Swapping be Differentially Private?
A Refreshment Stirred, not Shaken

James Bailie,* Ruobin Gong† & Xiao-Li Meng*

*Harvard, †Rutgers

October 31, 2023

Statistics Canada

Motivation

- interchanges the values of sensitive variables in a randomly selected subset of records.
- It was used as the primary disclosure avoidance method in the 1990, 2000 and 2010 US Censuses.
- For the 2020 Census, disclosure avoidance was overhauled with the primary aim of satisfying (DP).

Can we also understand data swapping from the perspective of DP?

Motivation

- *Data swapping* interchanges the values of sensitive variables in a randomly selected subset of records (Dalenius and Reiss 1982; Fienberg and McIntyre 2004).
- It was used as the primary disclosure avoidance method in the 1990, 2000 and 2010 US Censuses.
- For the 2020 Census, disclosure avoidance was overhauled with the primary aim of satisfying *differential privacy* (DP) (Dwork et al. 2006b).

Can we also understand data swapping from the perspective of DP?

Motivation

- *Data swapping* interchanges the values of sensitive variables in a randomly selected subset of records (Dalenius and Reiss 1982; Fienberg and McIntyre 2004).
- It was used as the primary disclosure avoidance method in the 1990, 2000 and 2010 US Censuses.
- For the 2020 Census, disclosure avoidance was overhauled with the primary aim of satisfying *differential privacy* (DP) (Dwork et al. 2006b).

Can we also understand data swapping from the perspective of DP?

Motivation

- *Data swapping* interchanges the values of sensitive variables in a randomly selected subset of records (Dalenius and Reiss 1982; Fienberg and McIntyre 2004).
- It was used as the primary disclosure avoidance method in the 1990, 2000 and 2010 US Censuses.
- For the 2020 Census, disclosure avoidance was overhauled with the primary aim of satisfying *differential privacy* (DP) (Dwork et al. 2006b).

Can we also understand data swapping from the perspective of DP?

Motivation

- *Data swapping* interchanges the values of sensitive variables in a randomly selected subset of records (Dalenius and Reiss 1982; Fienberg and McIntyre 2004).
- It was used as the primary disclosure avoidance method in the 1990, 2000 and 2010 US Censuses.
- For the 2020 Census, disclosure avoidance was overhauled with the primary aim of satisfying *differential privacy* (DP) (Dwork et al. 2006b).

Can we also understand data swapping from the perspective of DP?

Differential privacy (Dwork et al. 2006b)

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

For example,

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

Differential privacy (Dwork et al. 2006b)

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

For example,

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

Differential privacy (Dwork et al. 2006b)

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

For example,

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

Differential privacy is Lipschitz continuity:

$$|T(\boldsymbol{x}) - T(\boldsymbol{x}')| \leq \epsilon |\boldsymbol{x} - \boldsymbol{x}'|,$$

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

"the *output* of T doesn't change much if the *input* doesn't change much" (*robustness*)

Differential privacy (Dwork et al. 2006b)

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

For example,

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

Differential privacy is Lipschitz continuity:

$$d_{\mathrm{Pr}}\left[T(\boldsymbol{x}), T(\boldsymbol{x}')\right] \leq \epsilon d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}'),$$

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

"the *output* of T doesn't change much if the *input* doesn't change much" (*robustness*)

Differential privacy (Dwork et al. 2006b)

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

For example,

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

Differential privacy is Lipschitz continuity:

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

for all possible data values $\boldsymbol{x}, \boldsymbol{x}'$, where $\mathsf{P}_{\boldsymbol{x}}$ is the distribution of T induced by the random noise Z.

"the *output* of T doesn't change much if the *input* doesn't change much" (*robustness*)

Output divergence d_{Pr} and data divergence $d_{\mathcal{X}}$

$d_{\mathcal{X}}$ is typically a graph distance on the data space $\mathcal{X}$.

Output divergence $d_{\Pr}$ and data divergence $d_{\mathcal{X}}$

$d_{\mathcal{X}}$ is typically a graph distance on the data space $\mathcal{X}$.

$$d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}') = k \quad \Leftrightarrow \quad \text{“}k\text{ units changed their responses”}$$

Output divergence d_{Pr} and data divergence $d_{\mathcal{X}}$

$d_{\mathcal{X}}$ is typically a graph distance on the data space $\mathcal{X}$.

$$d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}') = k \quad \Leftrightarrow \text{“}k \text{ units changed their responses”}$$

---

d_{Pr} can be the *multiplicative distance* (pure DP):

$$d_{\textsc{Mult}}(P, Q) = \sup_{E} \left| \ln \frac{P(E)}{Q(E)} \right|,$$

between probability distributions P, Q,

Output divergence $d_{\Pr}$ and data divergence $d_{\mathcal{X}}$

$d_{\mathcal{X}}$ is typically a graph distance on the data space $\mathcal{X}$.

$$d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}') = k \quad \Leftrightarrow \text{ “}k\text{ units changed their responses”}$$

$d_{\Pr}$ can be the *multiplicative distance* (pure DP):

$$d_{\textsc{Mult}}(P, Q) = \sup_{E} \left| \ln \frac{P(E)}{Q(E)} \right|,$$

between probability distributions P, Q, or the *normalised Rényi metric* D_{nor} (zero concentrated DP):

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

Does data swapping satisfy differential privacy?

- Not under the traditional formulation of DP...
- Because swapping has *invariants* 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 P_x and $\mathrm{P}_{x'}$ do not have common support, and so

$$d_{\text{MULT}}\left[\mathrm{P}_x, \mathrm{P}_{x'}\right] = D_{\text{nor}}\left[\mathrm{P}_x, \mathrm{P}_{x'}\right] = \infty.$$

Does data swapping satisfy differential privacy?

- Not under the traditional formulation of DP...
- Because swapping has *invariants* $\boldsymbol{c}_{\text{Swap}}$ – 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 P_x and $\text{P}_{x'}$ do not have common support, and so

$$d_{\text{MULT}}\left[\text{P}_x, \text{P}_{x'}\right] = D_{\text{nor}}\left[\text{P}_x, \text{P}_{x'}\right] = \infty.$$

Does data swapping satisfy differential privacy?

- Not under the traditional formulation of DP...
- Because swapping has *invariants* $\boldsymbol{c}_{\mathrm{Swap}}$ – 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...
- Because the TopDown Algorithm (TDA) has *invariants* c_{TDA}.

Does the 2020 US Census satisfy differential privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\mathrm{TDA}}$.

Does the 2020 US Census satisfy differential privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\mathrm{TDA}}$.

Modifying the definition of DP:

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

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.

Does the 2020 US Census satisfy differential privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\text{TDA}}$.

Modifying the definition of DP:

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

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.

Does the 2020 US Census satisfy differential privacy?

- Not under the traditional formulation of DP...
- Because the TopDown Algorithm (TDA) has *invariants* $\boldsymbol{c}_{\mathrm{TDA}}$.

Modifying the definition of DP:

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

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 V_{Match}.
Within each stratum:

1. Select each record independently with probability p (the swap rate).
2. Randomly derange swapping variable V_{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 $d_{\Pr} = d_{\text{MULT}}$ and budget

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

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

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 V_{Match}.
Within each stratum:

1. Select each record independently with probability p (the swap rate).
2. Randomly derange swapping variable V_{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

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

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

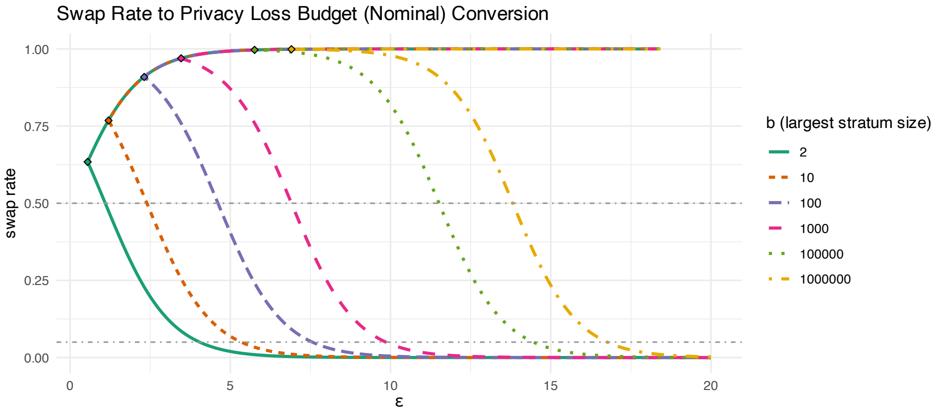

Conversion between the swap rate (p) and the nominal PLB (ϵ) at different levels of b. Note that:

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

Examples from the US Decennial Censuses

	d_{Pr}	$d_{\mathcal{X}}$ (Unit)	Invariants	Privacy Loss Budget
TopDown*	D_{nor}	d^{p}_{HAM} (person)	Population (state) Total housing units (block) Occupied group quarters (block) Structural zeros	PL & DHC: $\rho^2 = 15.29$ $\epsilon = 52.83$ ($\delta = 10^{-10}$)
SafeTab**	D_{nor}	d^{p}_{HAM} (person)	None	DDHC-A: $\rho^2 = 19.776$ DDHC-B & S-DHC: *TBD.*
Swapping	d_{MULT}	d^{h}_{HAM} (household)	Varies but greater than TDA	ϵ between 9.37-19.38

*(Abowd et al. 2022)

**(Tumult Labs 2022)

Four Components of a DP Flavour $(\mathscr{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{x}}\mathrm{P}(T(\boldsymbol{x}) \in \cdot)$ at every dataset $\boldsymbol{x}$ in the data universe $\mathcal{D}$.

Derivatives measure change in output per change in input. How do we measure change?

1. Data space $\mathscr{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{X}}$ on $\mathscr{X}$.
4. Divergence d_{Pr} on the space of (probability distributions over) the output.
2. Allow for multiple data universes $\mathcal{D} \subset \mathscr{X}$ from a data multiverse $\mathscr{D}$.

Four Components of a DP Flavour $(\mathscr{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{x}}\mathrm{P}(T(\boldsymbol{x}) \in \cdot)$ at every dataset $\boldsymbol{x}$ in the data universe $\mathcal{D}$.

Derivatives measure change in output per change in input. How do we measure change?

1. Data space $\mathcal{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{X}}$ on $\mathcal{X}$.
4. Divergence d_{Pr} on the space of (probability distributions over) the output.
2. Allow for multiple data universes $\mathcal{D} \subset \mathcal{X}$ from a data multiverse $\mathscr{D}$.

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\Pr})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{x}}\mathrm{P}(T(\boldsymbol{x}) \in \cdot)$ at every dataset $\boldsymbol{x}$ in the data universe $\mathcal{D}$.

Derivatives measure change in output per change in input. How do we measure change?

1. Data space $\mathcal{X}$ (the set of all theoretically-possible datasets).

3. Divergence $d_{\mathcal{X}}$ on $\mathcal{X}$.

4. Divergence $d_{\Pr}$ on the space of (probability distributions over) the output.

2. Allow for multiple data universes $\mathcal{D} \subset \mathcal{X}$ from a data multiverse $\mathscr{D}$.

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{x}}\mathsf{P}(T(\boldsymbol{x}) \in \cdot)$ at every dataset $\boldsymbol{x}$ in the data universe $\mathcal{D}$.

Derivatives measure change in output per change in input. How do we measure change?

1. Data space $\mathcal{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{X}}$ on $\mathcal{X}$.
4. Divergence d_{Pr} on the space of (probability distributions over) the output.
2. Allow for multiple data universes $\mathcal{D} \subset \mathcal{X}$ from a data multiverse $\mathscr{D}$.

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{x}}\mathrm{P}(T(\boldsymbol{x}) \in \cdot)$ at every dataset $\boldsymbol{x}$ in the data universe $\mathcal{D}$.

Derivatives measure change in output per change in input. How do we measure change?

1. Data space $\mathcal{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{X}}$ on $\mathcal{X}$.
4. Divergence d_{Pr} on the space of (probability distributions over) the output.
2. Allow for multiple data universes $\mathcal{D} \subset \mathcal{X}$ from a data multiverse $\mathscr{D}$.

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\Pr})$

Definition

A differential privacy flavour is a tuple $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\Pr})$.

A data release mechanism T *satisfies* $\mathsf{DP}(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\Pr})$ with budget ϵ if

$$d_{\Pr}\Big(\mathsf{P}_{\boldsymbol{x}}(T(\boldsymbol{x}) \in \cdot), \mathsf{P}_{\boldsymbol{x}'}(T(\boldsymbol{x}') \in \cdot)\Big) \leq \epsilon d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}'),$$

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

We aren't doing anything new here!

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

Definition

A differential privacy flavour is a tuple $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$.

A data release mechanism T *satisfies* $\mathrm{DP}(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$ with budget ϵ if

$$d_{\mathrm{Pr}}\Big(\mathsf{P}_{\boldsymbol{x}}(T(\boldsymbol{x}) \in \cdot), \mathsf{P}_{\boldsymbol{x}'}(T(\boldsymbol{x}') \in \cdot)\Big) \leq \epsilon d_{\mathcal{X}}(\boldsymbol{x}, \boldsymbol{x}'),$$

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

We aren't doing anything new here!

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

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

3. $d_{\mathcal{X}}$: $(\mathcal{R}, \epsilon)$-generic DP (Kifer and Machanavajjhala 2011) edge vs node privacy (Hay et al. 2009; McSherry and 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).

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

1. $\mathcal{X}$: DP for network data (Hay et al. 2009) for geospatial data (Andrés et al. 2013) Pufferfish DP (Kifer and 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).

Four Components of a DP Flavour $(\mathcal{X}, \mathcal{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

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

3. $d_{\mathcal{X}}$: $(\mathcal{R}, \epsilon)$-generic DP (Kifer and Machanavajjhala 2011) edge vs node privacy (Hay et al. 2009; McSherry and 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).

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

1. $\mathcal{X}$: DP for network data (Hay et al. 2009) for geospatial data (Andrés et al. 2013) Pufferfish DP (Kifer and 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).

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

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

3. $d_{\mathcal{X}}$: $(\mathcal{R}, \epsilon)$-generic DP (Kifer and Machanavajjhala 2011) edge vs node privacy (Hay et al. 2009; McSherry and 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).

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

1. $\mathcal{X}$: DP for network data (Hay et al. 2009) for geospatial data (Andrés et al. 2013) Pufferfish DP (Kifer and 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).

Four Components of a DP Flavour $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\Pr})$

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

3. $d_{\mathcal{X}}$: $(\mathcal{R}, \epsilon)$-generic DP (Kifer and Machanavajjhala 2011) edge vs node privacy (Hay et al. 2009; McSherry and 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).

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

1. $\mathcal{X}$: DP for network data (Hay et al. 2009) for geospatial data (Andrés et al. 2013) Pufferfish DP (Kifer and 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).

Five Building Blocks of DP $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}}, \epsilon_{\mathcal{D}})$

1. The protection domain (*what* can be protected?): as defined by the dataset space $\mathcal{X}$;
2. The scope of protection (*to where* does the protection extend?): as instantiated by the data multiverse $\mathscr{D}$, which is a collection of data universes $\mathcal{D} \subset \mathcal{X}$;
3. The protection unit (*who* are the units for data perturbation?): as conceptualized by the divergence $d_{\mathcal{X}}$ on the dataset space $\mathcal{X}$;
4. The standard of protection (*how* to measure the output variations?): as captured by the divergence d_{Pr} on the output probability distributions; and
5. The intensity of protection (*how much* protection is afforded?): as quantified by the privacy-loss budget $\epsilon_{\mathcal{D}}$.

Five Building Blocks of DP $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}}, \epsilon_{\mathcal{D}})$

1. The protection domain (*what* can be protected?): as defined by the dataset space $\mathcal{X}$;
2. The scope of protection (*to where* does the protection extend?): as instantiated by the data multiverse $\mathscr{D}$, which is a collection of data universes $\mathcal{D} \subset \mathcal{X}$;
3. The protection unit (*who* are the units for data perturbation?): as conceptualized by the divergence $d_{\mathcal{X}}$ on the dataset space $\mathcal{X}$;
4. The standard of protection (*how* to measure the output variations?): as captured by the divergence d_{Pr} on the output probability distributions; and
5. The intensity of protection (*how much* protection is afforded?): as quantified by the privacy-loss budget $\epsilon_{\mathcal{D}}$.

Five Building Blocks of DP $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}}, \epsilon_{\mathcal{D}})$

1. The protection domain (*what* can be protected?): as defined by the dataset space $\mathcal{X}$;
2. The scope of protection (*to where* does the protection extend?): as instantiated by the data multiverse $\mathscr{D}$, which is a collection of data universes $\mathcal{D} \subset \mathcal{X}$;
3. The protection unit (*who* are the units for data perturbation?): as conceptualized by the divergence $d_{\mathcal{X}}$ on the dataset space $\mathcal{X}$;
4. The standard of protection (*how* to measure the output variations?): as captured by the divergence d_{Pr} on the output probability distributions; and
5. The intensity of protection (*how much* protection is afforded?): as quantified by the privacy-loss budget $\epsilon_{\mathcal{D}}$.

Five Building Blocks of DP $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}}, \epsilon_{\mathcal{D}})$

1. The protection domain (*what* can be protected?): as defined by the dataset space $\mathcal{X}$;
2. The scope of protection (*to where* does the protection extend?): as instantiated by the data multiverse $\mathscr{D}$, which is a collection of data universes $\mathcal{D} \subset \mathcal{X}$;
3. The protection unit (*who* are the units for data perturbation?): as conceptualized by the divergence $d_{\mathcal{X}}$ on the dataset space $\mathcal{X}$;
4. The standard of protection (*how* to measure the output variations?): as captured by the divergence d_{Pr} on the output probability distributions; and
5. The intensity of protection (*how much* protection is afforded?): as quantified by the privacy-loss budget $\epsilon_{\mathcal{D}}$.

Five Building Blocks of DP $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}}, \epsilon_{\mathcal{D}})$

1. The protection domain (*what* can be protected?): as defined by the dataset space $\mathcal{X}$;
2. The scope of protection (*to where* does the protection extend?): as instantiated by the data multiverse $\mathscr{D}$, which is a collection of data universes $\mathcal{D} \subset \mathcal{X}$;
3. The protection unit (*who* are the units for data perturbation?): as conceptualized by the divergence $d_{\mathcal{X}}$ on the dataset space $\mathcal{X}$;
4. The standard of protection (*how* to measure the output variations?): as captured by the divergence d_{Pr} on the output probability distributions; and
5. The intensity of protection (*how much* protection is afforded?): as quantified by the privacy-loss budget $\epsilon_{\mathcal{D}}$.

The TopDown Algorithm (TDA) (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 $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, 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 $\boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{x})$.

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

The TopDown Algorithm (TDA) (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 $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, so that $\boldsymbol{T}$ satisfies $\text{DP}(\mathcal{X}_{\text{CEF}}, \{\mathcal{X}_{\text{CEF}}\}, d^p_{\text{HAM}}, D_{\text{nor}})$ 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 $\boldsymbol{c}_{\text{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\text{TDA}}(\boldsymbol{x})$.

TDA satisfies $\text{DP}(\mathcal{X}_{\text{CEF}}, \mathscr{D}_{\boldsymbol{c}_{\text{TDA}}}, d^p_{\text{HAM}}, D_{\text{nor}})$ with budget ρ_{TDA}.

The TopDown Algorithm (TDA) (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 $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, so that $\boldsymbol{T}$ satisfies $\text{DP}(\mathcal{X}_{\text{CEF}}, \{\mathcal{X}_{\text{CEF}}\}, d^p_{\text{HAM}}, D_{\text{nor}})$ 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 $\boldsymbol{c}_{\text{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\text{TDA}}(\boldsymbol{x})$.

TDA satisfies $\text{DP}(\mathcal{X}_{\text{CEF}}, \mathscr{D}_{\boldsymbol{c}_{\text{TDA}}}, d^p_{\text{HAM}}, D_{\text{nor}})$ with budget ρ_{TDA}.

The TopDown Algorithm (TDA) (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 $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, so that $\boldsymbol{T}$ satisfies $\text{DP}(\mathcal{X}_{\text{CEF}}, \{\mathcal{X}_{\text{CEF}}\}, d^p_{\text{HAM}}, D_{\text{nor}})$ 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 $\boldsymbol{c}_{\text{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\text{TDA}}(\boldsymbol{x})$.

TDA satisfies $\text{DP}(\mathcal{X}_{\text{CEF}}, \mathscr{D}_{\boldsymbol{c}_{\text{TDA}}}, d^p_{\text{HAM}}, D_{\text{nor}})$ with budget ρ_{TDA}.

The TopDown Algorithm (TDA) (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 $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, 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 $\boldsymbol{c}_{\text{TDA}}(\boldsymbol{z}) = \boldsymbol{c}_{\text{TDA}}(\boldsymbol{x})$.

TDA satisfies $\mathrm{DP}(\mathcal{X}_{\text{CEF}}, \mathscr{D}_{\boldsymbol{c}_{\text{TDA}}}, d^p_{\text{HAM}}, D_{\text{nor}})$ with budget ρ_{TDA}.

Theorem: TDA satisfies DP, subject to its Invariants

Denote the space of possible Census Edited Files by $\mathcal{X}_{\text{CEF}}$.

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:

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

Theorem: TDA satisfies DP, subject to its Invariants

Denote the space of possible Census Edited Files by $\mathcal{X}_{\text{CEF}}$.

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 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 DP$(\mathcal{X}_{\text{CEF}}, \mathscr{D}_{\boldsymbol{c}'}, d_{\mathcal{X}}, D_{\text{nor}})$ with any finite budget ρ.

Theorem: TDA satisfies DP, subject to its Invariants

Denote the space of possible Census Edited Files by $\mathcal{X}_{\text{CEF}}$.

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.

Contributions

- We supply a *framework for capturing and comparing* different flavours of DP which highlights often their overlooked components.
- We prove that *swapping satisfies DP, subject to its invariants*, putting its privacy guarantees on a comparable footing to the TopDown Algorithm.
- Our framework may help data custodians to systematically understand how traditional SDC methods can provide formal privacy protection.

Implications:

- What is the performance of reconstruction attacks on other formally-private mechanisms with invariants?
- Algorithmic and probabilistic transparency of swapping methods (for better data utility).

Contributions

- We supply a *framework for capturing and comparing* different flavours of DP which highlights often their overlooked components.
- We prove that *swapping satisfies DP, subject to its invariants*, putting its privacy guarantees on a comparable footing to the TopDown Algorithm.
- Our framework may help data custodians to systematically understand how traditional SDC methods can provide formal privacy protection.

Implications:

- What is the performance of reconstruction attacks on other formally-private mechanisms with invariants?
- Algorithmic and probabilistic transparency of swapping methods (for better data utility).

What if the 2020 Census used swapping?

The total nominal ϵ achievable by applying swapping to the 2020 Decennial Census for a variety of V_{Match}, V_{Swap}, and swap rate choices.

V_{Match}	V_{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 (V_{Match}, V_{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.

A Perverse Guide to Reducing the Privacy Loss ϵ (without adding more noise)

1. Add more invariants
2. Increase the granularity of the privacy units (inflate $d_{\mathcal{X}}$)
   - Persons instead of households
   - One day's worth of data, instead of all of an individual's data over time
3. Artificially shrink the output divergence d_{Pr}
   - Use (ϵ, δ)-DP instead of ϵ-DP.

References I

- Abowd, John et al. (June 2022). "The 2020 Census Disclosure Avoidance System TopDown Algorithm". In: *Harvard Data Science Review* Special Issue 2. DOI: 10.1162/99608f92.529e3cb9.
- Abowd, John M, Matthew J Schneider, and Lars Vilhuber (2013). "Differential privacy applications to Bayesian and linear mixed model estimation". In: *Journal of Privacy and Confidentiality* 5.1.
- Andrés, Miguel E., Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi (Nov. 2013). "Geo-Indistinguishability: Differential Privacy for Location-Based Systems". In: *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*. CCS '13. New York, NY, USA: Association for Computing Machinery, pp. 901–914. ISBN: 978-1-4503-2477-9. DOI: 10.1145/2508859.2516735.

References II

Ashmead, Robert, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, and William Sexton (2019). *EFFECTIVE PRIVACY AFTER ADJUSTING FOR INVARIANTS WITH APPLICATIONS TO THE 2020 Census*. Tech. rep.

Asi, Hilal, John C. Duchi, and O. Javidbakht (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, Borja, Gilles Barthe, and Marco Gaboardi (Jan. 2020). "Privacy Profiles and Amplification by Subsampling". In: *Journal of Privacy and Confidentiality* 10.1. ISSN: 2575-8527. DOI: `10.29012/jpc.726`.

Barber, Rina Foygel and John C. Duchi (Dec. 2014). *Privacy and Statistical Risk: Formalisms and Minimax Bounds*. `http://arxiv.org/abs/1412.4451`. DOI: `10.48550/arXiv.1412.4451`. arXiv: `1412.4451 [cs, math, stat]`.

References III

Barthe, Gilles and Federico Olmedo (2013). "Beyond Differential Privacy: Composition Theorems and Relational Logic for f-Divergences between Probabilistic Programs". In: *Automata, Languages, and Programming*. Ed. by Fedor V. Fomin, Rūsiņš Freivalds, Marta Kwiatkowska, and David Peleg. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 49–60. ISBN: 978-3-642-39212-2. DOI: 10.1007/978-3-642-39212-2_8.

Beimel, Amos, Shiva Prasad Kasiviswanathan, and Kobbi Nissim (Feb. 2010). "Bounds on the Sample Complexity for Private Learning and Private Data Release". In: *Proceedings of the 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland*. Ed. by Daniele Micciancio. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 437–454. DOI: 10.1007/978-3-642-11799-2_26.

References IV

Bhaskar, Raghav, Abhishek Bhowmick, Vipul Goyal, Srivatsan Laxman, and Abhradeep Thakurta (2011). "Noiseless Database Privacy". In: *Advances in Cryptology – ASIACRYPT 2011*. Ed. by Dong Hoon Lee and Xiaoyun Wang. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 215–232. ISBN: 978-3-642-25385-0. DOI: 10.1007/978-3-642-25385-0_12.

Bun, Mark, Jörg Drechsler, Marco Gaboardi, Audra McMillan, and Jayshree Sarathy (June 2022). "Controlling Privacy Loss in Sampling Schemes: An Analysis of Stratified and Cluster Sampling". In: *Foundations of Responsible Computing (FORC 2022)*, p. 24.

Bun, Mark and Thomas Steinke (2016). "Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds". In: *Theory of Cryptography*. Ed. by Martin Hirt and Adam Smith. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 635–658. ISBN: 978-3-662-53641-4. DOI: 10.1007/978-3-662-53641-4_24.

References V

Canonne, Clement, Gautam Kamath, and Thomas Steinke (July 2022). "The Discrete Gaussian for Differential Privacy". In: *Journal of Privacy and Confidentiality* 12.1. ISSN: 2575-8527. DOI: 10.29012/jpc.784. eprint: 2004.00010. (Visited on 04/28/2023).

Charest, Anne-Sophie and Yiwei Hou (2016). "On the meaning and limits of empirical differential privacy". In: *Journal of Privacy and Confidentiality* 7.3, pp. 53–66.

Chatzikokolakis, Konstantinos, Miguel E. Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi (2013). "Broadening the Scope of Differential Privacy Using Metrics". In: *Privacy Enhancing Technologies*. Ed. by Emiliano De Cristofaro and Matthew Wright. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 82–102. DOI: 10.1007/978-3-642-39077-7_5.

References VI

Dalenius, Tore and Steven P. Reiss (Jan. 1982). "Data-Swapping: A Technique for Disclosure Control". In: *Journal of Statistical Planning and Inference* 6.1, pp. 73–85. ISSN: 0378-3758. DOI: 10.1016/0378-3758(82)90058-1.

Dharangutte, Prathamesh, Jie Gao, Ruobin Gong, and Fang-Yi Yu (2023). "Integer Subspace Differential Privacy". In: *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-23).*

Dong, Jinshuo, Aaron Roth, and Weijie J. Su (2022). "Gaussian Differential Privacy". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 84.1, pp. 3–37. ISSN: 1467-9868. DOI: 10.1111/rssb.12454.

References VII

Dwork, Cynthia, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor (2006a). "Our Data, Ourselves: Privacy Via Distributed Noise Generation". In: *Advances in Cryptology - EUROCRYPT 2006*. Ed. by Serge Vaudenay. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 486–503. ISBN: 978-3-540-34547-3. DOI: 10.1007/11761679_29.

Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith (2006b). "Calibrating noise to sensitivity in private data analysis". In: *Theory of cryptography conference*. Springer, pp. 265–284.

References VIII

Dwork, Cynthia, Moni Naor, Toniann Pitassi, and Guy N. Rothblum (June 2010). "Differential Privacy under Continual Observation". In: *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*. STOC '10. https://dl.acm.org/doi/10.1145/1806689.1806787. New York, NY, USA: Association for Computing Machinery, pp. 715–724. ISBN: 978-1-4503-0050-6. DOI: 10.1145/1806689.1806787.

Ebadi, Hamid, David Sands, and Gerardo Schneider (Jan. 2015). "Differential Privacy: Now It's Getting Personal". In: *ACM SIGPLAN Notices* 50.1, pp. 69–81. ISSN: 0362-1340. DOI: 10.1145/2775051.2677005.

Feldman, Vitaly and Tijana Zrnic (Jan. 2022). *Individual Privacy Accounting via a Rényi Filter*. http://arxiv.org/abs/2008.11193. arXiv: 2008.11193 [cs, stat].

References IX

Fienberg, S. and J. McIntyre (2004). “Data Swapping: Variations on a Theme by Dalenius and Reiss”. In: *Privacy in Statistical Databases.* DOI: 10.1007/978-3-540-25955-8_2.

Gao, Jie, Ruobin Gong, and Fang-Yi Yu (June 2022). “Subspace Differential Privacy”. In: *Proceedings of the AAAI Conference on Artificial Intelligence.* Vol. 36. 4, pp. 3986–3995. DOI: 10.1609/aaai.v36i4.20315.

Gong, Ruobin and Xiao-Li Meng (2020). “Congenial differential privacy under mandated disclosure”. In: *Proceedings of the 2020 ACM-IMS on Foundations of Data Science Conference.* FODS ‘20, pp. 59–70.

Hay, Michael, Chao Li, Gerome Miklau, and David Jensen (Dec. 2009). “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.

References X

He, Xi, Ashwin Machanavajjhala, and Bolin Ding (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, Zach, Ting Yu, and Graham Cormode (Apr. 2015). "Conservative or Liberal? Personalized Differential Privacy". In: *2015 IEEE 31st International Conference on Data Engineering*. https://ieeexplore.ieee.org/document/7113353, pp. 1023–1034. DOI: 10.1109/ICDE.2015.7113353. (Visited on 09/30/2023).

Kifer, Daniel and Ashwin Machanavajjhala (2011). "No Free Lunch in Data Privacy". In: *Proceedings of the 2011 International Conference on Management of Data - SIGMOD '11*. Athens, Greece: ACM Press, pp. 193–204. ISBN: 978-1-4503-0661-4. DOI: 10.1145/1989323.1989345.

References XI

Kifer, Daniel and Ashwin Machanavajjhala (2014). "Pufferfish: A framework for mathematical privacy definitions". In: *ACM Transactions on Database Systems (TODS)* 39.1, pp. 1–36.

McSherry, Frank and Ratul Mahajan (Aug. 2010). "Differentially-Private Network Trace Analysis". In: *Proceedings of the ACM SIGCOMM 2010 Conference*. SIGCOMM '10. New York, NY, USA: Association for Computing Machinery, pp. 123–134. ISBN: 978-1-4503-0201-2. DOI: 10.1145/1851182.1851199.

Mironov, Ilya (Aug. 2017). "Rényi Differential Privacy". In: *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275. DOI: 10.1109/CSF.2017.11. eprint: 1702.07476. (Visited on 01/14/2020).

O'Keefe, Christine M and Anne-Sophie Charest (2019). "Bootstrap differential privacy". In: *Transactions on Data Privacy* 12, pp. 1–28.

References XII

Redberg, Rachel and Yu-Xiang Wang (2021). "Privately Publishable Per-Instance Privacy". In: *Advances in Neural Information Processing Systems*. Vol. 34. Curran Associates, Inc., pp. 17335–17346. (Visited on 03/29/2023).

Seeman, Jeremy, Matthew Reimherr, and Aleksandra Slavkovic (May 2022). *Formal Privacy for Partially Private Data*. http://arxiv.org/abs/2204.01102. arXiv: 2204.01102 [cs, stat].

Seeman, Jeremy, William Sexton, David Pujol, and Ashwin Machanavajjhala (2023+). "Per-Record Differential Privacy: Modeling Dependence between Individual Privacy Loss and Confidential Records". In.

References XIII

Soria-Comas, Jordi, Josep Domingo-Ferrer, David Sánchez, and David Megías (June 2017). "Individual Differential Privacy: A Utility-Preserving Formulation of Differential Privacy Guarantees". In: *IEEE Transactions on Information Forensics and Security* 12.6, pp. 1418–1429. ISSN: 1556-6013, 1556-6021. DOI: 10.1109/TIFS.2017.2663337. (Visited on 03/29/2023).

Tumult Labs (Mar. 2022). *SafeTab: DP Algorithms for 2020 Census Detailed DHC Race & Ethnicity*. Tech. rep.

Wang, Yu-Xiang (Nov. 2018). *Per-Instance Differential Privacy*. http://arxiv.org/abs/1707.07708. arXiv: 1707.07708 [cs, stat].

References XIV

Zhou, Shuheng, Katrina Ligett, and Larry Wasserman (June 2009). "Differential Privacy with Compression". In: *Proceedings of the 2009 IEEE International Conference on Symposium on Information Theory - Volume 4*. ISIT'09. Coex, Seoul, Korea: IEEE Press, pp. 2718–2722. ISBN: 978-1-4244-4312-3.

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	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

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	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

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	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

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	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

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

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

$\boldsymbol{V}_{\text{Hold}} - \boldsymbol{V}_{\text{Match}}$

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	$\cdots$
Boston					
Cambridge	-1			+1	
Brookline					
Somerville	+1			-1	
Watertown					
$\vdots$					

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 $\boldsymbol{V}_{\text{Hold}} - \boldsymbol{V}_{\text{Match}} \times \boldsymbol{V}_{\text{Swap}}$.

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					
⋮					

Changes: Interior cells of $\boldsymbol{V}_{\text{Hold}} - \boldsymbol{V}_{\text{Match}} \times \boldsymbol{V}_{\text{Swap}}$.

Invariants:

1. $\boldsymbol{V}_{\text{Hold}}$
2. $\boldsymbol{V}_{\text{Match}} \times \boldsymbol{V}_{\text{Swap}}$

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 $[n^{\boldsymbol{Z}}_{jkl}]$

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 $\mathcal{D}$,

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

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 $\Pr(\sigma)/Pr(\sigma')$ 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 (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 $\boldsymbol{c}_{\mathrm{TDA}}$ (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 $\boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{Z}_p, \boldsymbol{Z}_h) = \boldsymbol{c}_{\mathrm{TDA}}(\boldsymbol{X}_p, \boldsymbol{X}_h)$ 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 $\boldsymbol{c} : \mathcal{X} \to \mathbb{R}^l$.

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}

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

   - *Pure ϵ-DP* (Dwork et al. 2006b): d_{Pr} is the multiplicative distance

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

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

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

   - *Zero Concentrated DP* (Bun and Steinke 2016):

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

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

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

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 et al. 2006b): d_{Pr} is the multiplicative distance

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

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

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

   - *Zero Concentrated DP* (Bun and Steinke 2016):

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

   where D_α is the *Rényi divergence* of order α:

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

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 et al. 2006b): $d_{\Pr}$ 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 et al. 2006a):

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

   - *Zero Concentrated DP* (Bun and Steinke 2016):

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

where D_α is the *Rényi divergence* of order α:

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

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 et al. 2006b): $d_{\Pr}$ 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 et al. 2006a):

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

   - *Zero Concentrated DP* (Bun and Steinke 2016):

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

   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

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

The invariants $\boldsymbol{c}_{\text{Swap}}$ 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

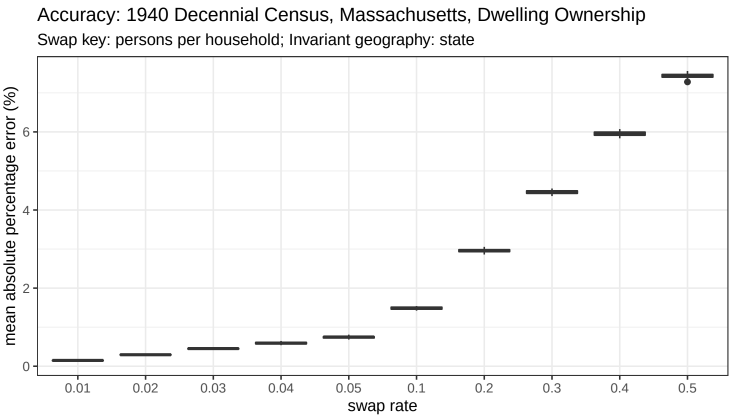

Mean absolute percentage error (MAPE) in the two-way tabulation of dwelling ownership by county induced by the Permutation Algorithm applied to the 1940 Census full count data of Massachusetts, at different swap rates from 1% to 50%. Each boxplot reflects 20 independent runs of the Algorithm at that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

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

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

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

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 $(\mathcal{X}, \mathscr{D}_{\boldsymbol{c}}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$ with *privacy budget* $\epsilon_{\mathcal{D}} = 0$.

2. Now suppose $d_{\mathrm{Pr}}(\mathrm{P}, \mathrm{Q}) = \infty$ 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 $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$ for any $\epsilon_{\mathcal{D}_0} < \infty$*.

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 $(\mathcal{X}, \mathscr{D}_{\boldsymbol{c}}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$ only if $\epsilon_{\mathcal{D}_0} > 0$*.