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

James Bailie, Ruobin Gong & Xiao-Li Meng

4 May, 2023

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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 (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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	⋯
⋮	⋮	⋮	⋮	⋮	⋮	⋱

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

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

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

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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

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

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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

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

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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

Data Swapping (Dalenius and Reiss 1982; Fienberg and McIntyre 2004)

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

The Permutation Algorithm

Input: a dataset $\boldsymbol{X}$.
Define strata as groups of records which match on the swap key $\boldsymbol{V}_{\text{Match}}$.

The Permutation Algorithm

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

1. Select each record independently with probability p (the swap rate).

The Permutation Algorithm

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

1. Select each record independently with probability p (the swap rate).
2. Derange swapping variable $\boldsymbol{V}_{\text{Swap}}$ of selected records, uniformly at random.

The Permutation Algorithm

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

1. Select each record independently with probability p (the swap rate).
2. Derange swapping variable $\boldsymbol{V}_{\text{Swap}}$ of selected records, uniformly at random.

The Permutation Algorithm

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

1. Select each record independently with probability p (the swap rate).
2. Derange swapping variable $\boldsymbol{V}_{\text{Swap}}$ of selected records, uniformly at random.

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

The Permutation Algorithm

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

1. Select each record independently with probability p (the swap rate).
2. Derange swapping variable $\boldsymbol{V}_{\text{Swap}}$ of selected records, uniformly at random.

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

Theorem

The Permutation Algorithm satisfies pure differential privacy with privacy loss budget

$$\epsilon = \ln(b+1) - \ln(o), \quad \text{for } 0 < p \leq 0.5,$$

conditioning on the invariants it induces, where $o = p/(1-p)$ *and* b *is the largest stratum size.*

The Permutation Algorithm

Theorem

The Permutation Algorithm satisfies pure differential privacy with privacy loss budget

$$\epsilon = \ln(b+1) - \ln(o), \quad \textit{for } 0 < p \leq 0.5,$$

conditioning on the invariants it induces, where $o = p/(1-p)$ *and* b *is the largest stratum size.*

The Permutation Algorithm

Theorem

The Permutation Algorithm satisfies pure differential privacy with privacy loss budget

$$\epsilon = \ln(b+1) - \ln(o), \quad \text{for } 0 < p \leq 0.5,$$

conditioning on the invariants it induces, where $o = p/(1-p)$ and b is the largest stratum size.

Theorem (formal)

The Permutation Algorithm satisfies $(\mathscr{D}_{\boldsymbol{c}_{\mathrm{Swap}}}, d_{\mathrm{HamS}}, \mathrm{MULT})$ differential privacy with privacy loss budget

$$\epsilon = \ln(b+1) - \ln(o), \quad \text{for } 0 < p \leq 0.5,$$

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

The Three Components of Differential Privacy $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{X}} \Pr(T(\boldsymbol{X}) \in \cdot)$ 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?

The Three Components of Differential Privacy $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{X}} \Pr(T(\boldsymbol{X}) \in \cdot)$ 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?

2. Divergence $d_{\mathcal{X}}$ on the data input space $\mathcal{X}$ (the set of all theoretically-possible datasets).

The Three Components of Differential Privacy $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{X}} \Pr(T(\boldsymbol{X}) \in \cdot)$ 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?

2. Divergence $d_{\mathcal{X}}$ on the data input space $\mathcal{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{T}}$ on the space of (probability distributions over) the output.

The Three Components of Differential Privacy $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{X}} \Pr(T(\boldsymbol{X}) \in \cdot)$ 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?

2. Divergence $d_{\mathcal{X}}$ on the data input space $\mathcal{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{T}}$ on the space of (probability distributions over) the output.
1. Allow the data universe $\mathcal{D} - \mathscr{D}(\boldsymbol{X}^*)$ to be data-dependent.

The Three Components of Differential Privacy $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{X}} \Pr(T(\boldsymbol{X}) \in \cdot)$ 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?

2. Divergence $d_{\mathcal{X}}$ on the data input space $\mathcal{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{T}}$ on the space of (probability distributions over) the output.
1. Allow the data universe $\mathcal{D} - \mathscr{D}(\boldsymbol{X}^*)$ to be data-dependent.

Definition

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

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

$$d_{\mathcal{T}}\Big(P_{\boldsymbol{X}}(T(\boldsymbol{X}) \in \cdot), P_{\boldsymbol{X}'}(T(\boldsymbol{X}') \in \cdot)\Big) \leq \epsilon d_{\mathcal{X}}(\boldsymbol{X}, \boldsymbol{X}'),$$

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

The Three Components of Differential Privacy $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$

Intuition: DP is a bound on the *derivative* of a data-release mechanism $\frac{d}{d\boldsymbol{X}} \Pr(T(\boldsymbol{X}) \in \cdot)$ 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?

2. Divergence $d_{\mathcal{X}}$ on the data input space $\mathcal{X}$ (the set of all theoretically-possible datasets).
3. Divergence $d_{\mathcal{T}}$ on the space of (probability distributions over) the output.
1. Allow the data universe $\mathcal{D} - \mathscr{D}(\boldsymbol{X}^*)$ to be data-dependent.

Definition

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

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

$$d_{\mathcal{T}}\Big(P_{\boldsymbol{X}}(T(\boldsymbol{X}) \in \cdot), P_{\boldsymbol{X}'}(T(\boldsymbol{X}') \in \cdot)\Big) \leq \epsilon d_{\mathcal{X}}(\boldsymbol{X}, \boldsymbol{X}'),$$

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

We aren't doing anything new here!

Examples of $\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}}$ in the Decennial Censuses

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

* (Abowd et al. 2022)

** (Tumult Labs 2022)

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

1. An invariant-compliant data universe:

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

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

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

1. An invariant-compliant data universe:

$$\mathscr{D}_{\boldsymbol{c}}(\boldsymbol{X}) = \Big\{ \boldsymbol{X}' \in \mathcal{X} : \boldsymbol{c}(\boldsymbol{X}') = \boldsymbol{c}(\boldsymbol{X}) \Big\},$$

   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 $d_{\mathcal{T}}$

3. Divergence $d_{\mathcal{T}}$ on (the probability distributions over) the output space

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

3. Divergence $d_{\mathcal{T}}$ on (the probability distributions over) the output space
   - Pure ϵ-DP (Dwork et al. 2006b): $d_{\mathcal{T}}$ is the multiplicative distance

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

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

3. Divergence $d_{\mathcal{T}}$ on (the probability distributions over) the output space
   - Pure ϵ-DP (Dwork et al. 2006b): $d_{\mathcal{T}}$ is the multiplicative distance

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

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

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

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

3. Divergence $d_{\mathcal{T}}$ on (the probability distributions over) the output space
   - Pure ϵ-DP (Dwork et al. 2006b): $d_{\mathcal{T}}$ is the multiplicative distance

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

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

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

   - Zero Concentrated DP (Bun and Steinke 2016):

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

     where D_α is the Rényi divergence of order α:

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

Swapping Satisfies Differential Privacy, Conditioning on its Invariants

Theorem

The Permutation Algorithm satisfies $(\mathscr{D}_{\boldsymbol{c}_{\mathrm{Swap}}}, d^{u}_{\mathrm{HamS}}, \mathrm{MULT})$ *differential privacy with privacy loss budget*

$$\epsilon = \ln(b+1) - \ln(o), \quad \text{for } 0 < p \leq 0.5,$$

with $o = p/(1-p)$ *and* b *is the largest stratum size.*

Swapping Satisfies Differential Privacy, Conditioning on its Invariants

Theorem

The Permutation Algorithm satisfies $(\mathscr{D}_{\boldsymbol{c}_{\mathrm{Swap}}}, d^u_{\mathrm{HamS}}, \mathrm{MULT})$ *differential privacy with privacy loss budget*

$$\epsilon = \ln(b+1) - \ln(o), \quad \text{for } 0 < p \leq 0.5,$$

with $o = p/(1-p)$ *and* b *is the largest stratum size.*

For all $\boldsymbol{X}, \boldsymbol{X}'$ which share the same invariants – $\boldsymbol{c}_{\mathrm{Swap}}(\boldsymbol{X}) = \boldsymbol{c}_{\mathrm{Swap}}(\boldsymbol{X}')$ – and all possible output datasets $\boldsymbol{Z}$,

$$\Pr(T(\boldsymbol{X}) = \boldsymbol{Z}) \leq \exp(d^u_{\mathrm{HamS}}(\boldsymbol{X}, \boldsymbol{X}')\epsilon) \Pr(T(\boldsymbol{X}') = \boldsymbol{Z}),$$

where T is the Permutation Algorithm.

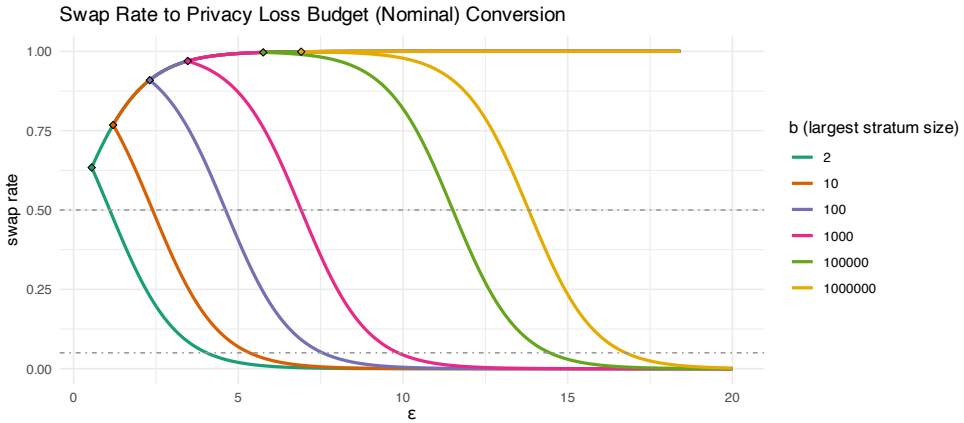

Conversion between the swap rate (p) and the nominal PLB (ϵ) at different levels of b: size of the largest stratum delineated by $\boldsymbol{V}_{\text{Match}}$. Note that:

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!

The TopDown Algorithm (Abowd et al. 2022)

Two-step procedure:

1. Add noise to cells independently:

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

where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, so that $\boldsymbol{T}$ satisfies $(\mathcal{X}, d^{p}_{\mathrm{HamS}}, D_{\mathrm{nor}})$-differential privacy with budget ρ_{TDA}.

The TopDown Algorithm (Abowd et al. 2022)

Two-step procedure:

1. Add noise to cells independently:

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

   where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, so that $\boldsymbol{T}$ satisfies $(\mathcal{X}, d^{p}_{\mathrm{HamS}}, D_{\mathrm{nor}})$-differential privacy with budget ρ_{TDA}.

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 (Abowd et al. 2022)

Two-step procedure:

1. Add noise to cells independently:

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

   where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, so that $\boldsymbol{T}$ satisfies $(\mathcal{X}, d^{p}_{\mathrm{HamS}}, D_{\mathrm{nor}})$-differential privacy with budget ρ_{TDA}.

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 (Abowd et al. 2022)

Two-step procedure:

1. Add noise to cells independently:

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

   where $\boldsymbol{W} \sim \mathcal{N}_{\mathbb{Z}}(0, \Sigma)$, so that $\boldsymbol{T}$ satisfies $(\mathcal{X}, d^p_{\mathrm{HamS}}, D_{\mathrm{nor}})$-differential privacy with budget ρ_{TDA}.
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 $(\mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}}, d^p_{\mathrm{HamS}}, D_{\mathrm{nor}})$-differential privacy with budget ρ_{TDA}.

Comparisons with 2020 Census

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

* (Abowd et al. 2022)

** (Tumult Labs 2022)

What if the 2020 Census Used Swapping?

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

$\boldsymbol{V}_{\text{Match}}$	$\boldsymbol{V}_{\text{Swap}}$	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 ($\boldsymbol{V}_{\text{Match}}$, $\boldsymbol{V}_{\text{Swap}}$, 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.

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

1. Add more invariants (decrease the size of the data universes $\mathcal{D}$)
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 $d_{\mathcal{T}}$
   - Use (ϵ, δ)-DP instead of ϵ-DP.

Contributions

- We supply a framework $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$ for capturing and comparing different types of differential privacy which highlights often overlooked components of DP.
- We prove that swapping satisfies DP, when conditioning on its invariants, putting its privacy guarantees on the same footing as the TopDown algorithm.
- Our framework may help data custodians to systematically understand how traditional SDC methods can afford formal privacy protection.

Contributions

- We supply a framework $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$ for capturing and comparing different types of differential privacy which highlights often overlooked components of DP.
- We prove that swapping satisfies DP, when conditioning on its invariants, putting its privacy guarantees on the same footing as the TopDown algorithm.
- Our framework may help data custodians to systematically understand how traditional SDC methods can afford formal privacy protection.

Implications:

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

Contributions

- We supply a framework $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$ for capturing and comparing different types of differential privacy which highlights often overlooked components of DP.
- We prove that swapping satisfies DP, when conditioning on its invariants, putting its privacy guarantees on the same footing as the TopDown algorithm.
- Our framework may help data custodians to systematically understand how traditional SDC methods can afford formal privacy protection.

Implications:

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

Extensions:

- Incorporating disclosure risk: Variable swap rate.
- Allowing flexible invariants: Probabilistic matching & Pre-swap noise infusion (Hawes and Rodriguez 2021).

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.

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.

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.

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.

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.

References II

Hawes, Michael and Rolando Rodriguez (2021). *Determining the Privacy-loss Budget Research into Alternatives to Differential Privacy*.

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

US Census Bureau (Apr. 2023). *2023-04-03 Privacy-loss Budget Allocations*. Tech. rep. https://www2.census.gov/programs-surveys/decennial/2020/program-management/data-product-planning/2010-demonstration-data-products/04-Demonstration_Data_Products_Suite/2023-04-03/2023-04-03_Privacy-Loss_Budget_Allocations.pdf [Accessed: 04-25-2023], p. 10. (Visited on 04/25/2023).

The Permutation Algorithm

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

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 $\boldsymbol{c}_{\text{TDA}}$ (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
    $\boldsymbol{W}_p \sim \mathcal{N}_{\mathbb{Z}}(\boldsymbol{0}, \boldsymbol{D}_p)$
    $\boldsymbol{W}_h \sim \mathcal{N}_{\mathbb{Z}}(\boldsymbol{0}, \boldsymbol{D}_h)$
  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.

Theorem: Swapping Satisfies DP, Conditioning on its Invariants

Let

$$b = \max\{0, n_j \mid \text{there are at least two different records in stratum } j\}.$$

Then the Permutation Algorithm is $(\boldsymbol{c}_{\mathrm{Swap}}, d^u_{\mathrm{HamS}}, \epsilon_{\mathcal{D}})$-DP where d^u_{HamS} is the symmetric Hamming distance

$$d^u_{\mathrm{HamS}}(\boldsymbol{X}, \boldsymbol{X}') = \frac{1}{2}\left|\boldsymbol{X} \ominus \boldsymbol{X}'\right|,$$

and $\epsilon_{\mathcal{D}} = 0$ if $b = 0$, otherwise

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

with $o = p/(1-p)$. On the other hand, for $p \in \{0, 1\}$ and for some $\mathcal{D}$ with $b > 0$, the Permutation Algorithm does not satisfy $(\boldsymbol{c}_{\mathrm{Swap}}, d^u_{\mathrm{HamS}}, \epsilon_{\mathcal{D}})$-DP for any finite $\epsilon_{\mathcal{D}}$.

Proof Intuition

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

$$\Pr(\sigma(\boldsymbol{X}) = \boldsymbol{Z}) \leq \exp(d^{u}_{\mathrm{HamS}}(\boldsymbol{X}, \boldsymbol{X}')\epsilon) \Pr(\sigma'(\boldsymbol{X}') = \boldsymbol{Z}),$$

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.

Theorem: TDA satisfies DP, Conditioning on its Invariants

Let $\boldsymbol{c}_{\mathrm{TDA}}$ be the invariants of TDA and let $\mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}}$ be the induced data universe function. Then TDA satisfies the differential privacy definition $(\mathscr{D}_{\boldsymbol{c}_{\mathrm{TDA}}}, d^{p}_{\mathrm{HamS}}, D_{\mathrm{nor}})$ with privacy budget $\rho_{\mathrm{TDA}} = 2.63$ (for the Census Redistricting Summary File) and $\rho_{\mathrm{TDA}} = 15.29$ (for the DHC). In the opposite direction, let $\boldsymbol{c}'$ be any proper subset of TDA's invariants. Then TDA does not satisfy $(\mathscr{D}_{\boldsymbol{c}'}, d_{\mathcal{X}}, D_{\mathrm{nor}})$ with any finite budget ρ.

		ρ	ϵ (with $\delta = 10^{-10}$)
PL	Household	0.07	2.70
	Person	2.56	17.90
DHC	Household	7.70	34.33
	Person	4.96	26.34
Total		15.29	52.83

Source: (US Census Bureau 2023).

Numerical demonstration: 1940 Census full count data

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

The invariants $\boldsymbol{c}_{\mathrm{Swap}}$ 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 1: 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 2: 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

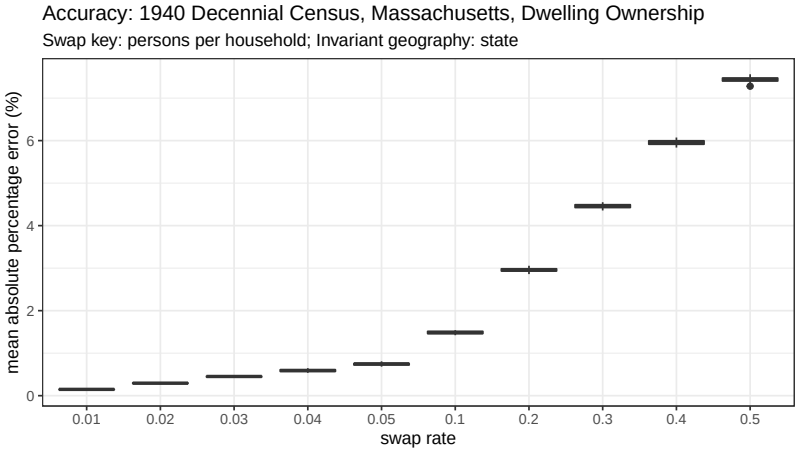

Mean absolute percentage error (MAPE) in the two-way tabulation of dwelling ownership by county induced by the Permutation Algorithm applied to the 1940 Census full count data of Massachusetts, at different swap rates from 1% to 50%. Each boxplot reflects 20 independent runs of the Algorithm at that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

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

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 $d_{\mathcal{T}}$, the mechanism $T(\boldsymbol{X}) = \boldsymbol{c}(\boldsymbol{X})$ that releases the invariants exactly satisfies $(\mathscr{D}_{\boldsymbol{c}}, d_{\mathcal{X}}, d_{\mathcal{T}})$ with privacy budget $\epsilon_{\mathcal{D}} = 0$.

2. Now suppose $d_{\mathcal{T}}(P, Q) = \infty$ if $d_{\mathrm{TV}}(P, Q) = 1$. Let $\mathscr{D}$ be a data universe function such that there exists datasets $\boldsymbol{X}_1, \boldsymbol{X}_2$ in some data universe $\mathcal{D}_0 \in \operatorname{Im} \mathscr{D}$ 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 $(\mathscr{D}, d_{\mathcal{X}}, d_{\mathcal{T}})$ for any $\epsilon_{\mathcal{D}_0} < \infty$.

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

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