Differential Privacy in Statistical Agencies
Challenges and Opportunities

James Bailie
Statistics Department, Harvard University

February 5, 2025

Les Houches

Background & motivation

- The U.S. Census Bureau has committed to adopting "*formal privacy*" for all their data products (U.S. Census Bureau 2022).
- Yet the *"science ... does not yet exist"* for a formally private solution to the American Community Survey (for example).
- In implementing differential privacy (DP), statistical agencies' data products come with their own set of *unique challenges and opportunities.*
- See also: "*Differential privacy for government agencies–Are we there yet?*" (Drechsler 2023) (Spoiler alert: No!)

Background & motivation

- Differential privacy (and its variants) are Lipschitz continuity conditions:

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

  for all $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{D}$ and all $\mathcal{D} \in \mathscr{D}$.
- Choices of $\mathcal{X}$, $\mathscr{D}$, $d_{\mathcal{X}}$ and d_{Pr} determine the flavour of DP.
- ε is the "intensity of protection" *in units dependent on* $\mathcal{X}$, $\mathscr{D}$, $d_{\mathcal{X}}$ *and* d_{Pr}.

Challenge: End-to-end pipelines for DP

- DP is a property of data processing (broadly construed, including e.g. sampling)
- Statistical agencies use complex data processing procedures (imputation, non-response adjustments, weighting, etc.)
- Efficient and true implementation of DP require translating existing processes into "DP-friendly" versions (i.e. ground-up re-building of agencies' infrastructure).
- Alternative: where can we safely cut corners?
    - What can be safely ignored? E.g. data-dependent sample design?
    - What constitutes principled corner cutting?
- Some preliminary work: "*Provable privacy with non-private pre-processing*" (Hu et al. 2024).
- Pufferfish is one framework for considering multi-phase pipelines (but does not compose).

Survey data pipeline

$$\cdots \longrightarrow \underset{\mathfrak{p}}{\textbf{Population}} \longrightarrow \underset{\mathfrak{f}}{\textbf{Frame}} \longrightarrow \underset{\mathfrak{s}}{\textbf{Sample}} \longrightarrow \textbf{Responding sample } \mathfrak{r} \longrightarrow \cdots \longrightarrow \textbf{Published statistics } t$$

Register for sampling

The collected data

What we want to study

Random set of units

Survey data pipeline

Register for sampling

The collected data

$\cdots \longrightarrow$ **Population** $\mathfrak{p}$ $\longrightarrow$ **Frame** $\mathfrak{f}$ $\longrightarrow$ **Sample** $\mathfrak{s}$ $\longrightarrow$ **Responding sample** $\mathfrak{r}$ $\longrightarrow \cdots \longrightarrow$ **Published statistics** t

What we want to study

Random set of units

Survey data pipeline

Register for sampling

The collected data

$\cdots \longrightarrow$ **Population** $\mathfrak{p}$ $\longrightarrow$ **Frame** $\mathfrak{f}$ $\longrightarrow$ **Sample** $\mathfrak{s}$ $\longrightarrow$ **Responding sample** $\mathfrak{r}$ $\longrightarrow \cdots \longrightarrow$ **Published statistics** t

What we want to study

Random set of units

Survey data pipeline

Register for sampling

The collected data

$\cdots \longrightarrow$ **Population** $\mathfrak{p}$ $\longrightarrow$ **Frame** $\mathfrak{f}$ $\longrightarrow$ **Sample** $\mathfrak{s}$ $\longrightarrow$ **Responding sample** $\mathfrak{r}$ $\longrightarrow \cdots \longrightarrow$ **Published statistics** t

What we want to study

Random set of units

Survey data pipeline

Register for sampling

The collected data

$\cdots \longrightarrow$ **Population** $\mathfrak{p}$ $\longrightarrow$ **Frame** $\mathfrak{f}$ $\longrightarrow$ **Sample** $\mathfrak{s}$ $\longrightarrow$ **Responding sample** $\mathfrak{r}$ $\longrightarrow \cdots \longrightarrow$ **Published statistics** t

What we want to study

Random set of units

DP settings for surveys (Bailie and Drechsler 2024)

$$\cdots \longrightarrow \mathfrak{p} \longrightarrow \mathfrak{f} \longrightarrow \mathfrak{s} \longrightarrow \mathfrak{r} \longrightarrow \cdots \longrightarrow t$$

Two considerations

- Where does the DP mechanism *start* in the data pipeline?
- Which of the previous steps in the pipeline are kept *invariant*?

For example,

1)

$$\begin{tikzcd}
 & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\
\cdots \arrow[ur] \arrow[dr] & & & & & & \\
 & \mathfrak{p}' \arrow[r] & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t'
\end{tikzcd}$$

2)

$$\begin{tikzcd}
 & & & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\
\cdots \arrow[r] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[ur] \arrow[dr] & & & & \\
 & & & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t'
\end{tikzcd}$$

Ten possible settings

$$\cdots \longrightarrow \mathfrak{p} \longrightarrow \mathfrak{f} \longrightarrow \mathfrak{s} \longrightarrow \mathfrak{r} \longrightarrow \cdots \longrightarrow t$$

---

1)

$$\begin{tikzcd}[row sep=small] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[ur] \arrow[dr] & & & & & & \\ & \mathfrak{p}' \arrow[r] & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

2)

$$\begin{tikzcd}[row sep=small] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[ur] \arrow[dr] & & & & & & \\ & \mathfrak{p}' \arrow[r] & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

3)

$$\begin{tikzcd}[row sep=small] & & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[r] & \mathfrak{p} \arrow[ur] \arrow[dr] & & & & & \\ & & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

4)

$$\begin{tikzcd}[row sep=small] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[ur] \arrow[dr] & & & & & & \\ & \mathfrak{p}' \arrow[r] & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

5)

$$\begin{tikzcd}[row sep=small] & & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[r] & \mathfrak{p} \arrow[ur] \arrow[dr] & & & & & \\ & & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

6)

$$\begin{tikzcd}[row sep=small] & & & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[r] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[ur] \arrow[dr] & & & & \\ & & & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

7)

$$\begin{tikzcd}[row sep=small] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[ur] \arrow[dr] & & & & & & \\ & \mathfrak{p}' \arrow[r] & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

8)

$$\begin{tikzcd}[row sep=small] & & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[r] & \mathfrak{p} \arrow[ur] \arrow[dr] & & & & & \\ & & \mathfrak{f}' \arrow[r] & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

9)

$$\begin{tikzcd}[row sep=small] & & & \mathfrak{s} \arrow[r] & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[r] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[ur] \arrow[dr] & & & & \\ & & & \mathfrak{s}' \arrow[r] & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

10)

$$\begin{tikzcd}[row sep=small] & & & & \mathfrak{r} \arrow[r] & \cdots \arrow[r] & t \\ \cdots \arrow[r] & \mathfrak{p} \arrow[r] & \mathfrak{f} \arrow[r] & \mathfrak{s} \arrow[ur] \arrow[dr] & & & \\ & & & & \mathfrak{r}' \arrow[r] & \cdots \arrow[r] & t' \end{tikzcd}$$

Opportunity: Building connections between DP and SDC

- Traditional statistical disclosure control (SDC) has a long and established literature starting in the 1970s.
- Almost all privacy methods currently implemented in statistical agencies are non-DP. (And those that are cut corners.)
- Can we provide some mathematical guarantees for these methods? Can DP inspire analogues of these methods which have guarantees?
  - Example – data swapping: A close analogue to the 2010 US Census satisfies DP subject to its invariants (Bailie et al. 2025).
- Can SDC inform the DP literature?
  - Example: 'Nosy neighbour' and 'journalist' attackers in sample data (Bailie and Drechsler 2024)

Challenge: DP's framing of data privacy (Seeman and Susser 2023)

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

Opportunity: Integrating DP into agencies' broader toolbox

- We need to fit DP into the bigger picture (Contextual Integrity, Five Safes).
- Agencies have many other tools for disclosure control (e.g. access controls).
- How can these tools work in tandem with statistical protections (such as DP)? (Bailie and Gong 2023)
- Are there other statistical theories of "formal privacy" beyond DP? (Why has a Lipschitz condition been so theoretically successful?)

Utility Considerations (I)

Privacy amplification by sampling

If $T(\mathfrak{s})$ is ε-DP and $\mathcal{S}(\mathfrak{f})$ randomly samples f fraction of the frame $\mathfrak{f}$, then $T' = T \circ \mathcal{S}$ is ε'-DP where $\varepsilon' \approx f\varepsilon$. (Balle et al. 2020)

- *Take-away:* If the sampling procedure is included, less noise is required to achieve the same privacy budget.
- *But* there is little privacy amplification when $\mathcal{S}$ is a complex sampling design. (Bun et al. 2022)

Utility Considerations (I)

Privacy amplification by sampling

If $T(\mathfrak{s})$ is ε-DP and $\mathcal{S}(\mathfrak{f})$ randomly samples f fraction of the frame $\mathfrak{f}$, then $T' = T \circ \mathcal{S}$ is ε'-DP where $\varepsilon' \approx f\varepsilon$. (Balle et al. 2020)

- *Take-away:* If the sampling procedure is included, less noise is required to achieve the same privacy budget.
- *But* there is little privacy amplification when $\mathcal{S}$ is a complex sampling design. (Bun et al. 2022)

Utility Considerations (I)

Privacy amplification by sampling

If $T(\mathfrak{s})$ is ε-DP and $\mathcal{S}(\mathfrak{f})$ randomly samples f fraction of the frame $\mathfrak{f}$, then $T' = T \circ \mathcal{S}$ is ε'-DP where $\varepsilon' \approx f\varepsilon$. (Balle et al. 2020)

- *Take-away:* If the sampling procedure is included, less noise is required to achieve the same privacy budget.
- *But* there is little privacy amplification when $\mathcal{S}$ is a complex sampling design. (Bun et al. 2022)

Utility Considerations (II)

- Surveys use weighted estimators $\sum_{i=1}^{n} w_i x_i$, which have increased sensitivity.
- Unweighted sums $\sum_{i=1}^{n} x_i$ have sensitivity $|\max x_i - \min x_i|$, where the max, min are over all possible values of x_i.
- Weighted estimators can have sensitivity

$$|\max w_i x_i - \min w_i x_i| + (n-1)(\max w_i - \min w_i)(|\max x_i| \vee |\min x_i|),$$

  because changing a record can change the weights of other records.
- Hence, weighted estimators require more noise to achieve the same privacy loss.
- Taking the frame as invariant means that the weights do not change.

Utility Considerations (II)

- Surveys use weighted estimators $\sum_{i=1}^{n} w_i x_i$, which have increased sensitivity.
- Unweighted sums $\sum_{i=1}^{n} x_i$ have sensitivity $|\max x_i - \min x_i|$, where the $\max, \min$ are over all possible values of x_i.
- Weighted estimators can have sensitivity

$$|\max w_i x_i - \min w_i x_i| + (n-1)(\max w_i - \min w_i)(|\max x_i| \vee |\min x_i|),$$

because changing a record can change the weights of other records.
- Hence, weighted estimators require more noise to achieve the same privacy loss.
- Taking the frame as invariant means that the weights do not change.

Utility Considerations (II)

- Surveys use weighted estimators $\sum_{i=1}^{n} w_i x_i$, which have increased sensitivity.
- Unweighted sums $\sum_{i=1}^{n} x_i$ have sensitivity $|\max x_i - \min x_i|$, where the $\max, \min$ are over all possible values of x_i.
- Weighted estimators can have sensitivity

$$|\max w_i x_i - \min w_i x_i| + (n-1)(\max w_i - \min w_i)(|\max x_i| \vee |\min x_i|),$$

  because changing a record can change the weights of other records.
- Hence, weighted estimators require more noise to achieve the same privacy loss.
- Taking the frame as invariant means that the weights do not change.

Utility Considerations (II)

- Surveys use weighted estimators $\sum_{i=1}^n w_i x_i$, which have increased sensitivity.
- Unweighted sums $\sum_{i=1}^n x_i$ have sensitivity $|\max x_i - \min x_i|$, where the $\max, \min$ are over all possible values of x_i.
- Weighted estimators can have sensitivity

$$|\max w_i x_i - \min w_i x_i| + (n-1)(\max w_i - \min w_i)(|\max x_i| \vee |\min x_i|),$$

because changing a record can change the weights of other records.
- Hence, weighted estimators require more noise to achieve the same privacy loss.
- Taking the frame as invariant means that the weights do not change.

Utility Considerations (II)

- Surveys use weighted estimators $\sum_{i=1}^{n} w_i x_i$, which have increased sensitivity.
- Unweighted sums $\sum_{i=1}^{n} x_i$ have sensitivity $|\max x_i - \min x_i|$, where the $\max, \min$ are over all possible values of x_i.
- Weighted estimators can have sensitivity

$$|\max w_i x_i - \min w_i x_i| + (n-1)(\max w_i - \min w_i)(|\max x_i| \vee |\min x_i|),$$

  because changing a record can change the weights of other records.
- Hence, weighted estimators require more noise to achieve the same privacy loss.
- Taking the frame as invariant means that the weights do not change.

Privacy Considerations (I)

Posterior-to-posterior privacy semantics

What would an attacker learn about a single record if it is included in the input dataset, relative to a counterfactual world in which it is not included?

- If T is ε-DP, then the posterior-to-posterior ratio is in $[e^{-\varepsilon}, e^{\varepsilon}]$. (Kifer et al. 2022)
- What record (in what input dataset) is being protected depends on where T *starts* in the data pipeline; and what counterfactual worlds are possible depends on what steps are *invariant*.

Privacy Considerations (I)

Posterior-to-posterior privacy semantics

What would an attacker learn about a single record if it is included in the input dataset, relative to a counterfactual world in which it is not included?

- If T is ε-DP, then the posterior-to-posterior ratio is in $[e^{-\varepsilon}, e^{\varepsilon}]$. (Kifer et al. 2022)
- What record (in what input dataset) is being protected depends on where T *starts* in the data pipeline; and what counterfactual worlds are possible depends on what steps are *invariant*.

Privacy Considerations (I)

Posterior-to-posterior privacy semantics

What would an attacker learn about a single record if it is included in the input dataset, relative to a counterfactual world in which it is not included?

- If T is ε-DP, then the posterior-to-posterior ratio is in $[e^{-\varepsilon}, e^{\varepsilon}]$. (Kifer et al. 2022)
- What record (in what input dataset) is being protected depends on where T *starts* in the data pipeline; and what counterfactual worlds are possible depends on what steps are *invariant*.

Privacy Considerations (I)

- Suppose $T(\mathfrak{s})$ is ε-DP and $\mathcal{S}(\mathfrak{f})$ randomly samples f fraction of $\mathfrak{f}$.
- $T' = T \circ \mathcal{S}$ is ε'-DP with $\varepsilon' \approx f\varepsilon < \varepsilon$.
- So the posterior-to-posterior ratio of T' should be in the interval $[e^{-\varepsilon'}, e^{\varepsilon'}]$.

Traditional statistical disclosure control attacker models

- The *nosy neighbor:* Knows that a record is in the sample.
- The *journalist:* Wants to learn about *any* record, so picks one in the sample.

For these attackers, the posterior-to-posterior ratio of T' is in the interval $[e^{-\varepsilon}, e^{\varepsilon}]$, *not* the interval $[e^{-\varepsilon'}, e^{\varepsilon'}]$.

Privacy Considerations (I)

- Suppose $T(\mathfrak{s})$ is ε-DP and $\mathcal{S}(\mathfrak{f})$ randomly samples f fraction of $\mathfrak{f}$.
- $T' = T \circ \mathcal{S}$ is ε'-DP with $\varepsilon' \approx f\varepsilon < \varepsilon$.
- So the posterior-to-posterior ratio of T' should be in the interval $[e^{-\varepsilon'}, e^{\varepsilon'}]$.

Traditional statistical disclosure control attacker models

- The *nosy neighbor:* Knows that a record is in the sample.
- The *journalist:* Wants to learn about *any* record, so picks one in the sample.

For these attackers, the posterior-to-posterior ratio of T' is in the interval $[e^{-\varepsilon}, e^{\varepsilon}]$, *not* the interval $[e^{-\varepsilon'}, e^{\varepsilon'}]$.

Privacy Considerations (I)

- Suppose $T(\mathfrak{s})$ is ε-DP and $\mathcal{S}(\mathfrak{f})$ randomly samples f fraction of $\mathfrak{f}$.
- $T' = T \circ \mathcal{S}$ is ε'-DP with $\varepsilon' \approx f\varepsilon < \varepsilon$.
- So the posterior-to-posterior ratio of T' should be in the interval $[e^{-\varepsilon'}, e^{\varepsilon'}]$.

Traditional statistical disclosure control attacker models

- The *nosy neighbor:* Knows that a record is in the sample.
- The *journalist:* Wants to learn about *any* record, so picks one in the sample.

For these attackers, the posterior-to-posterior ratio of T' is in the interval $[e^{-\varepsilon}, e^{\varepsilon}]$, *not* the interval $[e^{-\varepsilon'}, e^{\varepsilon'}]$.

Privacy Considerations (II)

- The composition theorem does not hold when there is dependency between the sample designs.
- For $i \in \{1, 2\}$, suppose $T_i(\mathfrak{s})$ is ε-DP, and $T_i' = T_i \circ \mathcal{S}$.
- Privacy loss of the composition (T_1', T_2') is not the sum of T_1' and T_2''s privacy losses.
- This will complicate global privacy loss calculations.

Privacy Considerations (II)

- The composition theorem does not hold when there is dependency between the sample designs.
- For $i \in \{1, 2\}$, suppose $T_i(\mathfrak{s})$ is ε-DP, and $T_i' = T_i \circ \mathcal{S}$.
- Privacy loss of the composition (T_1', T_2') is not the sum of T_1' and T_2''s privacy losses.
- This will complicate global privacy loss calculations.

Privacy Considerations (II)

- The composition theorem does not hold when there is dependency between the sample designs.
- For $i \in \{1, 2\}$, suppose $T_i(\mathfrak{s})$ is ε-DP, and $T_i' = T_i \circ \mathcal{S}$.
- Privacy loss of the composition (T_1', T_2') is not the sum of T_1' and T_2''s privacy losses.
- This will complicate global privacy loss calculations.

Privacy Considerations (II)

- The composition theorem does not hold when there is dependency between the sample designs.
- For $i \in \{1, 2\}$, suppose $T_i(\mathfrak{s})$ is ε-DP, and $T_i' = T_i \circ \mathcal{S}$.
- Privacy loss of the composition (T_1', T_2') is not the sum of T_1' and T_2''s privacy losses.
- This will complicate global privacy loss calculations.

Privacy Considerations (II)

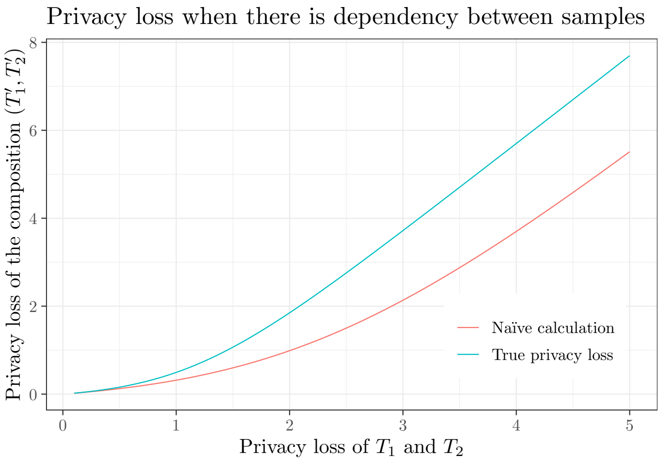

References I

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.

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.

References II

Bailie, James and Jörg Drechsler (May 13, 2024). "Whose Data Is It Anyway? Towards a Formal Treatment of Differential Privacy for Surveys". In: Data Privacy Protection and the Conduct of Applied Research: Methods, Approaches and Their Consequences. Washington D.C.: National Bureau of Economic Research, p. 33. URL: https://conference.nber.org/conf_papers/f194306.pdf.

Bailie, James and Ruobin Gong (2023). "The Five Safes as a Privacy Context". In: *The 5th Annual Symposium on Applications of Contextual Integrity*. The 5th Annual Symposium on Applications of Contextual Integrity. Toronto, Canada.

Bailie, James, Ruobin Gong, and Xiao-Li Meng (Jan. 14, 2025). *A Refreshment Stirred, Not Shaken (II): Invariant-preserving Deployments of Differential Privacy for the US Decennial Census*. DOI: 10.48550/arXiv.2501.08449. arXiv: 2501.08449 [cs]. URL: http://arxiv.org/abs/2501.08449 (visited on 01/16/2025). Pre-published.

References III

- 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].
- 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.

References IV

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.

Benthall, Sebastian and Rachel Cummings (Jan. 28, 2024). *Integrating Differential Privacy and Contextual Integrity.* DOI: 10.48550/arXiv.2401.15774. arXiv: 2401.15774 [cs]. URL: http://arxiv.org/abs/2401.15774 (visited on 06/04/2024).

References V

- 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 VI

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.

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)*.

References VII

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.

Drechsler, Jörg (Jan. 2023). "Differential Privacy for Government Agencies—Are We There Yet?" In: *Journal of the American Statistical Association* 118.541, pp. 761–773. ISSN: 0162-1459. DOI: 10.1080/01621459.2022.2161385. eprint: 2102.08847.

Dwork, Cynthia, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor (2006). "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.

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

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.

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.

References X

- Hu, Yaxi, Amartya Sanyal, and Bernhard Schölkopf (June 21, 2024). *Provable Privacy with Non-Private Pre-Processing*. arXiv: 2403.13041 [cs, stat]. URL: http://arxiv.org/abs/2403.13041 (visited on 08/21/2024).
- 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, John M. Abowd, Robert Ashmead, Ryan Cumings-Menon, Philip Leclerc, Ashwin Machanavajjhala, William Sexton, and Pavel Zhuravlev (Sept. 2022). *Bayesian and Frequentist Semantics for Common Variations of Differential Privacy: Applications to the 2020 Census*. Tech. rep. arXiv:2209.03310. DOI: 10.48550/arXiv.2209.03310. eprint: 2209.03310 (cs, stat). (Visited on 10/23/2022).

References XI

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.

— (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.

References XII

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.

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].

References XIII

- Seeman, Jeremy, William Sexton, David Pujol, and Ashwin Machanavajjhala (2023+). "Per-Record Differential Privacy: Modeling Dependence between Individual Privacy Loss and Confidential Records". In.
- Seeman, Jeremy and Daniel Susser (Oct. 2023). "Between Privacy and Utility: On Differential Privacy in Theory and Practice". In: *ACM Journal on Responsible Computing*. DOI: 10.1145/3626494.
- 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).

References XIV

U.S. Census Bureau (Dec. 2022). *Disclosure Avoidance Protections for the American Community Survey*. https://www.census.gov/newsroom/blogs/random-samplings/2022/12/disclosure-avoidance-protections-acs.html. (Visited on 12/17/2023).

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

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.

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 $\frac{d}{d\boldsymbol{x}}\mathrm{P}_{\boldsymbol{x}}(T \in \cdot)$ at every dataset $\boldsymbol{x}$ in every data universe $\mathcal{D} \in \mathscr{D}$.

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 $\frac{d}{d\boldsymbol{x}}\mathrm{P}_{\boldsymbol{x}}(T \in \cdot)$ at every dataset $\boldsymbol{x}$ in every data universe $\mathcal{D} \in \mathscr{D}$.

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 $\frac{d}{d\boldsymbol{x}}\mathrm{P}_{\boldsymbol{x}}(T \in \cdot)$ at every dataset $\boldsymbol{x}$ in every data universe $\mathcal{D} \in \mathscr{D}$.

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 $(\mathscr{X}, \mathscr{D}, d_{\mathscr{X}}, d_{\mathrm{Pr}})$

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

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 $d_{\mathscr{X}}$ 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 $(\mathcal{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

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

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 $(\mathcal{D}_0, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

Definition

A differential privacy flavour is a tuple $(\mathcal{D}_0, \mathscr{D}, d_{\mathcal{D}_0}, d_{\mathrm{Pr}})$.
A data release mechanism T *satisfies* $\mathrm{DP}(\mathcal{D}_0, \mathscr{D}, d_{\mathcal{D}_0}, d_{\mathrm{Pr}})$ with budget ϵ if

$$d_{\mathrm{Pr}}\Big(\mathrm{P}_{\mathfrak{d}}(T \in \cdot), \mathrm{P}_{\mathfrak{d}'}(T \in \cdot)\Big) \leq \epsilon d_{\mathcal{D}_0}(\mathfrak{d}, \mathfrak{d}'),$$

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

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. 2006) 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. $\mathcal{X}$: 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 $(\mathscr{X}, \mathscr{D}, d_{\mathcal{X}}, d_{\mathrm{Pr}})$

4. d_{Pr}: (ϵ, δ)-approximate DP (Dwork et al. 2006) 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. $\mathscr{X}$: 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. 2006) 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. $\mathcal{X}$: 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. 2006) 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. $\mathcal{X}$: 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 $(\mathcal{X}, \mathscr{D}, d_{\mathcal{D}_0}, d_{\Pr}, \epsilon_{\mathcal{D}})$

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 $(\mathcal{X}, \mathscr{D}, d_{\mathcal{D}_0}, d_{\mathrm{Pr}}, \epsilon_{\mathcal{D}})$

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 $(\mathcal{X}, \mathscr{D}, d_{\mathcal{D}_0}, d_{\Pr}, \epsilon_{\mathcal{D}})$

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 $(\mathcal{X}, \mathscr{D}, d_{\mathcal{D}_0}, d_{\Pr}, \epsilon_{\mathcal{D}})$

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 $(\mathcal{X}, \mathscr{D}, d_{\mathcal{D}_0}, d_{\Pr}, \epsilon_{\mathcal{D}})$

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 $\epsilon_{\mathcal{D}}$.