


Article

On the Lift, Related Privacy Measures, and Applications to Privacy–Utility Trade-Offs[†]

Mohammad Amin Zarrabian^{1,*} , Ni Ding²  and Parastoo Sadeghi^{3,*} ¹ College of Engineering, Computing and Cybernetics, Australian National University, Canberra, ACT 2601, Australia² School of Computing and Information Systems, University of Melbourne, Parkville, VIC 3010, Australia; ni.ding@unimelb.edu.au³ School of Engineering and Information Technology, University of New South Wales, Canberra, ACT 2600, Australia

* Correspondence: mohammad.zarrabian@anu.edu.au (M.A.Z.); p.sadeghi@unsw.edu.au (P.S.)

[†] Preliminary results of this work have been published in part at the 2022 IEEE International Conference on Acoustics, Speech and Signal Processing, Singapore, 22–27 May 2022, and IEEE Information Theory Workshop, Mumbai, India, 1–2 November 2022.

Abstract: This paper investigates lift, the likelihood ratio between the posterior and prior belief about sensitive features in a dataset. Maximum and minimum lifts over sensitive features quantify the adversary’s knowledge gain and should be bounded to protect privacy. We demonstrate that max- and min-lifts have a distinct range of values and probability of appearance in the dataset, referred to as *lift asymmetry*. We propose asymmetric local information privacy (ALIP) as a compatible privacy notion with lift asymmetry, where different bounds can be applied to min- and max-lifts. We use ALIP in the watchdog and optimal random response (ORR) mechanisms, the main methods to achieve lift-based privacy. It is shown that ALIP enhances utility in these methods compared to existing local information privacy, which ensures the same (symmetric) bounds on both max- and min-lifts. We propose subset merging for the watchdog mechanism to improve data utility and subset random response for the ORR to reduce complexity. We then investigate the related lift-based measures, including ℓ_1 -norm, χ^2 -privacy criterion, and α -lift. We reveal that they can only restrict max-lift, resulting in significant min-lift leakage. To overcome this problem, we propose corresponding lift-inverse measures to restrict the min-lift. We apply these lift-based and lift-inverse measures in the watchdog mechanism. We show that they can be considered as relaxations of ALIP, where a higher utility can be achieved by bounding only average max- and min-lifts.

Keywords: local information privacy; local differential privacy; watchdog privacy mechanism; optimal random response



Citation: Zarrabian, M.A.; Ding, N.; Sadeghi, P. On the Lift, Related Privacy Measures, and Applications to Privacy–Utility Trade-Offs. *Entropy* **2023**, *25*, 679. <https://doi.org/10.3390/e25040679>

Academic Editors: Eirik Rosnes and Hsuan-Yin Lin

Received: 20 February 2023

Revised: 6 April 2023

Accepted: 14 April 2023

Published: 18 April 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

With the recent emergence of “Big-Data”, generating, sharing, and analysing data are proliferating via the advancement of communication systems and machine learning methods. While sharing datasets is essential to achieve social and economic benefits, it may lead to the leakage of private information, which has raised great concern about the privacy preservation of individuals. The main approach to protect privacy is perturbing the data via a privacy mechanism. Consider some raw data denoted by random variable X and some sensitive features denoted by S , which are correlated via a joint distribution $P_{SX} \neq P_S \times P_X$. A privacy mechanism (characterised by the transition probability $P_{Y|X}$) is applied to publish Y as a sanitised version of X to protect S .

The design of a privacy mechanism depends on the privacy measure. Differential privacy (DP) [1–3] is a widely used notion of privacy. DP restricts the chance of revealing the individual’s presence in a dataset from the outcome of analysis over that dataset [4]. It

ensures that neighbored sensitive features s and s' , which differ in only one entry, result in a similar output probability distribution, by restricting the ratio between posterior beliefs $P_{Y|S}(y|s)/P_{Y|S}(y|s')$ below a threshold e^ϵ . The neighbourhood assumption is relaxed in the local differential privacy (LDP) [5–9], where the ratio between posterior beliefs is restricted below e^ϵ for *any* two sensitive features s and s' , denoted by ϵ -LDP. The quantity of ϵ is known as the *privacy budget*. DP and LDP are considered context-free privacy notions, i.e., they do not take into account the prior distribution P_S . In contrast, in information-theoretic (IT) privacy, also known as context-aware privacy [10,11], it is assumed that the distribution of data or an estimation of them is available. Some of the dominant IT privacy measures are mutual information (MI) [10,12,13], maximal leakage [14–16], α -leakage [17], and local information privacy (LIP) [11,18–28]. A challenge is that while data perturbation restricts privacy leakage, it necessarily reduces data resolution and datasets' usefulness. Therefore, a privacy mechanism is desired to deliver a satisfactory level of data utility. Depending on the application, data utility is quantified either by measures of similarity between X and Y , such as f-divergence [7] and MI [7,12], or measures of dissimilarity and error, such as Hamming distortion [8,9] and mean square error [22], respectively. This tension between privacy and utility is known as the privacy–utility trade-off (PUT).

In this paper, we consider lift, a pivotal element in IT privacy measures, which is the likelihood ratio between the posterior belief $P_{S|Y}(s|y)$ and prior belief $P_S(s)$ about sensitive features in a dataset:

$$l(s, y) = \frac{P_{S|Y}(s|y)}{P_S(s)} = \frac{P_{SY}(s, y)}{P_S(s)P_Y(y)}. \quad (1)$$

The logarithm of the lift $i(s, y) = \log l(s, y)$, which we call *log-lift*, is the information density [24]. For each y , the more $P_{S|Y}(s|y)$ differs from $P_S(s)$, the more the adversary gains knowledge about s [29]. Consequently, both min-lift and max-lift, denoted by $\min_s l(s, y)$ and $\max_s l(s, y)$, respectively, quantify the highest privacy leakage for each y . In [29], the role of min-lift and max-lift in privacy breach was proposed, based on which information privacy was introduced in [18]. Accordingly, min-lift is associated with revealing what values are less probable for s after observing y , while max-lift is associated with the more probable values. In addition, recently, other operational meanings for max-lift have been revealed in guessing frameworks [30] and quantitative information flow [31]. In LIP, min-lift and max-lift are bounded below and above by thresholds $e^{-\epsilon}$ and e^ϵ , respectively, to restrict the adversary's knowledge gain, denoted by ϵ -LIP. The main privacy mechanisms to achieve ϵ -LIP are the watchdog mechanism [24,25] and optimal random response (ORR) [28]. Watchdog mechanism bipartitions the alphabet of X into low-risk and high-risk symbols, and only high-risk ones are randomised. It was proved in [25] that X -invariant randomisation (e.g., merging all high-risk symbols) minimises privacy leakage for the watchdog mechanism. ORR is an optimal mechanism for ϵ -LIP, which maximises MI as the utility measure.

Contributions

We investigate lift and its related privacy notions such as LIP. We demonstrate that min-lift and max-lift have distinct values and probability of appearance in the dataset. More specifically, min-lifts have a broader range of values than max-lifts, while max-lifts have a higher likelihood $P_{SY}(s, y)$ of appearing in the dataset. We call this property *lift asymmetry*. However, ϵ -LIP allocates symmetric privacy budgets to $\min_s i(s, y)$ and $\max_s i(s, y)$ ($-\epsilon$ and ϵ , respectively), which is incompatible with the lift asymmetry. Thus, we propose asymmetric-LIP (ALIP) as an amenable privacy notion to the lift properties, where asymmetric privacy budgets can be allocated to $\min_s i(s, y)$ and $\max_s i(s, y)$, denoted by $-\epsilon_l$ and ϵ_u , respectively. We demonstrate that ALIP implies ϵ -LDP and can result in better utility than LIP in the watchdog and ORR mechanisms. Utility increases by relaxing the bound on the min-lift, which has a lower probability of appearance in the dataset.

We propose two randomization methods to overcome the low utility of the watchdog mechanism and the high complexity of the ORR mechanism. In the watchdog mechanism,

X-invariant randomisation perturbs all high-risk symbols together and deteriorates data resolution and utility. On the other hand, ORR suffers from high complexity, which is exponential in the size of datasets. To overcome these problems, we propose *subset merging* and *subset random response* (SRR) perturbation methods that make finer subsets of high-risk symbols and privatise each subset separately. Subset merging enhances utility in the watchdog mechanism by applying X-invariant randomisation to disjoint subsets of high-risk symbols. In addition, SRR relaxes the complexity of ORR for large datasets by applying random response solutions on disjoint subsets of high-risk symbols, which results in near-optimal utility.

Besides LIP, we also consider some recently proposed privacy measures, which we call *lift-based* measures, including ℓ_1 -norm [32], χ^2 -strong privacy [33], and α -lift [34]. They have been proposed as the privacy notions stronger than their corresponding average leakages: the total variation distance [35], χ^2 -divergence [36], and Sibson MI [16,34], respectively. We clarify that they only bound max-lift leakage and can cause significant min-lift leakage. Therefore, we propose a corresponding modified version of these measures to restrict min-lift leakage, which we call *lift-inverse* measures. We apply lift-based and lift-inverse measures in the watchdog mechanism with subset merging randomisation to investigate their PUT. They result in higher utility than ALIP since they are functions of average lift over sensitive features and, thus, can be considered as relaxations of the max- and min-lift.

2. Preliminaries

2.1. Notation

We use the following notation throughout the paper. Capital letters denote discrete random variables, corresponding capital calligraphic letters denote their finite supports, and lowercase letters denote any of their realisations. For example, a random variable X has the support \mathcal{X} , and its realisation is $x \in \mathcal{X}$. For random variables S and X , we use P_{SX} to indicate their joint probability distribution, $P_{S|X}$ for the conditional distribution of S given X , and P_S and P_X for the marginal distributions. Bold capital and lowercase letters are used for matrices and vectors, respectively, and lowercase letters for the corresponding elements of the vectors, e.g., $\mathbf{v} = [v_1, v_2, \dots, v_n]^T$. We also use $|\cdot|$ for the cardinality of a set, e.g., $|\mathcal{X}|$. We denote the natural logarithm by \log and the set of integers $\{1, 2, \dots, n\}$ by $[n]$. The indicator function is shown by $\mathbf{1}_{\{f\}}$, which is 1 when f is true and zero otherwise.

2.2. System Model and Privacy Measures

Consider some useful data intended for sharing and denoted by random variable X with alphabet \mathcal{X} . It is correlated with some sensitive features S with the alphabet \mathcal{S} through a discrete joint distribution P_{SX} . To protect the sensitive features, a privacy mechanism is applied to generate a sanitised version of X , denoted by Y with the alphabet \mathcal{Y} . We assume P_S and P_X have full support, and $P_{Y|X,S}(y|x,s) = P_{Y|X}(y|x)$, which results in the Markov chain $S - X - Y$.

The main privacy measure is lift (since we assume P_S and P_Y have full supports, $l(s,y)$ is finite), given in (1). Lift and its logarithm, log-lift, quantify multiplicative information gain on each sensitive feature $s \in \mathcal{S}$ via accessing $y \in \mathcal{Y}$. There are two cases: $l(s,y) > 1 \Rightarrow P_{S|Y}(s|y) > P_S(s)$ indicates the increment of the belief about s after releasing y ; $l(s,y) \leq 1 \Rightarrow P_{S|Y}(s|y) \leq P_S(s)$ means that releasing y decreases the belief. The more the posterior belief deviates from the prior belief, the more an adversary gains knowledge about s . Thus, for each $y \in \mathcal{Y}$, the $\max_s l(s,y)$ and $\min_s l(s,y)$ determine the highest knowledge gain of sensitive features, and they should be restricted to protect privacy. We use the following notation for these quantities:

$$\Psi(y) \triangleq \min_{s \in \mathcal{S}} l(s,y) \quad \text{and} \quad \Lambda(y) \triangleq \max_{s \in \mathcal{S}} l(s,y). \quad (2)$$

In Appendix A, we explain the operational meaning of $\Psi(y)$ and $\Lambda(y)$ in privacy breach based on the work in [29].

The lift has been applied in local information privacy [24,25,28] to provide protection of sensitive features, and is defined as follows.

Definition 1. For $\varepsilon \in \mathbb{R}_+$, a privacy mechanism $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$ is ε -local information private or ε -LIP, with respect to S , if for all $y \in \mathcal{Y}$,

$$e^{-\varepsilon} \leq \Psi(y) \quad \text{and} \quad \Lambda(y) \leq e^{\varepsilon}. \quad (3)$$

Another instance-wise measure is local differential privacy [5,6,28],

Definition 2. For $\varepsilon \in \mathbb{R}_+$, a privacy mechanism $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$ is ε -local differential private or ε -LDP, with respect to S , if for all $s, s' \in \mathcal{S}$ and all $y \in \mathcal{Y}$,

$$\Gamma(y) = \sup_{s, s' \in \mathcal{S}} \frac{P_{Y|S}(y|s)}{P_{Y|S}(y|s')} = \frac{\Lambda(y)}{\Psi(y)} \leq e^{\varepsilon}. \quad (4)$$

3. Asymmetric Local Information Privacy

According to (3), LIP restricts the decrement of $\log \Psi(y)$ and increment of $\log \Lambda(y)$ by the symmetric bounds. However, we demonstrate that these metrics have a distinct range of values and probabilities of appearance in the dataset, $P_{SY}(s, y)$. We plot the histogram of $\log \Psi(y)$ and $\log \Lambda(y)$ for 10^3 randomly generated distributions in Figure 1, where $|\mathcal{X}| = 17$ and $|\mathcal{S}| = 5$. In this figure, the range of $\log \Psi(y)$ is $[-12, -0.06]$, much larger than the range of $\log \Lambda(y)$, $[0.02, 1.64]$. Moreover, the maximum probability of $\log \Psi(y)$ is much lower than the maximum probability of $\log \Lambda(y)$. We refer to these properties as *lift asymmetry*. Since high values of $|\log \Psi(y)|$ have a significantly lower probability (for example, in Figure 1, the probability of $|\log \Psi(y)| \geq 6$ is near zero) than the $\log \Lambda(y)$, we can relax the min-lift privacy by allocating a higher privacy budget to it while applying a stricter bound for the max-lift. Thus, we propose asymmetric local information privacy (ALIP), where we consider different privacy budgets ε_l and ε_u for $|\log \Psi(y)|$ and $\log \Lambda(y)$, respectively.




Figure 1. Histogram of $\log \Psi(y) = \min_s i(s, y)$ and $\log \Lambda(y) = \max_s i(s, y)$ for 10^3 randomly generated distributions, where $|\mathcal{X}| = 17$, $|\mathcal{S}| = 5$.

This will result in the following notion of privacy, which is more compatible with the lift asymmetry property.

Definition 3. For $\varepsilon_l, \varepsilon_u \in \mathbb{R}_+$, a privacy mechanism $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$ is $(\varepsilon_l, \varepsilon_u)$ -asymmetric local information private, or $(\varepsilon_l, \varepsilon_u)$ -ALIP, with respect to S , if for all $y \in \mathcal{Y}$,

$$e^{-\varepsilon_l} \leq \Psi(y) \quad \text{and} \quad \Lambda(y) \leq e^{\varepsilon_u}. \quad (5)$$

The following proposition indicates how (ϵ_l, ϵ_u) -ALIP restricts average privacy leakage measures and LDP.

Proposition 1. *If (ϵ_l, ϵ_u) -ALIP is satisfied, then*

1. $I(S; Y) \leq \epsilon_u$;
2. $T(S; Y) \leq \frac{1}{2}(e^{\epsilon_u} - 1)$ and $\chi^2(S; Y) \leq (e^{\epsilon_u} - 1)^2$;
3. $I_\alpha^S(S; Y) \leq \frac{\alpha}{\alpha-1}\epsilon_u$ and $I_\alpha^A(S; Y) \leq \frac{\alpha}{\alpha-1}\epsilon_u$;
4. ϵ -LDP is satisfied where $\epsilon = \epsilon_l + \epsilon_u$;

where $T(S; Y)$ is the total variation distance, $\chi^2(S; Y)$ is χ^2 -divergence, $I_\alpha^S(S; Y)$ is Sibson MI, and $I_\alpha^A(S; Y)$ is Arimoto MI.

Proof. The proof is given in Appendix B. \square

Proposition 1–3 demonstrate that average measures are bounded with the max-lift privacy budget. In Section 3.1, we show that ALIP can enhance utility via relaxing min-lift $\epsilon_l > \epsilon_u$, where a smaller upper bound is allocated to the max-lift and average measures in Proposition 1. Proposition 1–4 shows the relationship between (ϵ_l, ϵ_u) -ALIP and ϵ -LDP. We introduce a variable $\lambda \in (0, 1)$ to have a convenient representation of this relationship as follows: for an LDP privacy budget ϵ , if we set $\epsilon_l = \lambda\epsilon$ and $\epsilon_u = (1 - \lambda)\epsilon$, we have $\epsilon_l + \epsilon_u = \epsilon$. Thus, varying λ gives rise to different (ϵ_l, ϵ_u) -ALIP scenarios within the same budget for ϵ -LDP. If $\lambda < 0.5$, we have relaxation on the max-lift privacy; if $\lambda > 0.5$, it implies relaxation on the min-lift privacy. When $\lambda = 0.5$, we have the *symmetric* case of $\frac{\epsilon}{2}$ -LIP, where $\epsilon_l = \epsilon_u = \frac{\epsilon}{2}$.

3.1. ALIP Privacy–Utility Trade-Off

In this subsection, we propose a watchdog mechanism based on ALIP and LDP and an asymmetric ORR (AORR) mechanism for ALIP to perturb data and achieve privacy protection. We observe the PUT of ALIP and LDP, where the utility is measured by MI between X and Y , $I(X; Y)$.

3.1.1. Watchdog Mechanism

Watchdog privacy mechanism bipartitions \mathcal{X} into low-risk and high-risk subsets denoted by \mathcal{X}_L and \mathcal{X}_H , respectively, and only randomises high-risk symbols. In the existing LIP, \mathcal{X}_L and \mathcal{X}_H are determined by symmetric bounds. We propose to use ALIP to obtain \mathcal{X}_L and \mathcal{X}_H :

$$\mathcal{X}_L \triangleq \{x \in \mathcal{X} : e^{-\epsilon_l} \leq \Psi(x) \text{ and } \Lambda(x) \leq e^{\epsilon_u}\} \text{ and } \mathcal{X}_H = \mathcal{X} \setminus \mathcal{X}_L. \quad (6)$$

For LDP, \mathcal{X}_L and \mathcal{X}_H are given by

$$\mathcal{X}_L \triangleq \{x \in \mathcal{X} : \Gamma(x) \leq e^\epsilon\} \text{ and } \mathcal{X}_H = \mathcal{X} \setminus \mathcal{X}_L. \quad (7)$$

After obtaining \mathcal{X}_L and \mathcal{X}_H , the privacy mechanism will be

$$\mathcal{M} = \begin{cases} \mathbf{1}_{\{x=y\}}, & x, y \in \mathcal{X}_L = \mathcal{Y}_L, \\ r(y|x), & x \in \mathcal{X}_H, y \in \mathcal{Y}_H, \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where $\mathbf{1}_{\{x=y\}}$ indicates the publication of low-risk symbols without alteration, and $r(y|x)$ is the randomisation on high-risk symbols, where $\sum_{y \in \mathcal{Y}_H} r(y|x) = 1$.

An instance of $r(y|x)$ is the X -invariant randomisation, if $r(y|x) = \mathcal{R}(y)$ for $x \in \mathcal{X}_H, y \in \mathcal{Y}_H$, and $\sum_{y \in \mathcal{Y}_H} \mathcal{R}(y) = 1$. An example of $\mathcal{R}(y)$ is the uniform randomisation $\mathcal{R}(y) = \frac{1}{|\mathcal{Y}_H|}$ with the special case of *complete merging*, where $|\mathcal{Y}_H| = 1$, and all $x \in \mathcal{X}_H$ are mapped to one super symbol $y^* \in \mathcal{Y}_H$. It was proved in [25] for LIP that

X -invariant randomisation minimises privacy leakage in \mathcal{X}_H . Accordingly, if we apply ALIP in the watchdog mechanism, for $\mathcal{X}_H \neq \emptyset$, the minimum leakages over \mathcal{X}_H are

$$\bar{\epsilon}_u := \max_{s \in \mathcal{S}} i(s, \mathcal{X}_H) = \max_{s \in \mathcal{S}} \log l(s, \mathcal{X}_H) = \max_{s \in \mathcal{S}} \log \frac{P(\mathcal{X}_H|s)}{P(\mathcal{X}_H)}, \quad (9)$$

$$\bar{\epsilon}_l := \left| \min_{s \in \mathcal{S}} i(s, \mathcal{X}_H) \right| = \left| \min_{s \in \mathcal{S}} \log l(s, \mathcal{X}_H) \right| = \left| \min_{s \in \mathcal{S}} \log \frac{P(\mathcal{X}_H|s)}{P(\mathcal{X}_H)} \right|, \quad (10)$$

where $P(\mathcal{X}_H|s) = \sum_{x \in \mathcal{X}_H} P_{X|S}(x|s)$ and $P(\mathcal{X}_H) = \sum_{x \in \mathcal{X}_H} P_X(x)$.

X -invariant randomisation is also applicable for LDP, and the following theorem shows that it minimises LDP privacy leakage in \mathcal{X}_H .

Theorem 1. *In the LDP watchdog mechanism where \mathcal{X}_L and \mathcal{X}_H are determined according to (7), X -invariant randomisation minimises privacy leakage in \mathcal{X}_H measured by $\Gamma(y)$ in (4).*

Proof. The proof is given in Appendix C. \square

In the watchdog mechanism with X -invariant randomisation, the resulting utility measured by MI between X and Y is given by

$$I(X; Y) = H(X) - \sum_{x \in \mathcal{X}_H} P_X(x) \log \frac{P(\mathcal{X}_H)}{P_X(x)}. \quad (11)$$

In [25], it was verified that $I(X; Y)$ in (11) is monotonic in \mathcal{X}_H : if $\mathcal{X}'_H \subset \mathcal{X}_H$ then $I(X; Y) < I'(X; Y)$, where $I'(X; Y)$ is the resulting utility of \mathcal{X}'_H .

Proposition 2. *In the watchdog mechanism with X -invariant randomisation, for a given LDP privacy budget ϵ , $\lambda \in (0, 1)$, and ALIP privacy budgets $\epsilon_l = \lambda\epsilon, \epsilon_u = (1 - \lambda)\epsilon$, LDP results in higher utility and than ALIP.*

Proof. Denote the high-risk subset for LDP by \mathcal{X}'_H and for ALIP by \mathcal{X}_H . We have

$$\mathcal{X}'_H = \{x \in \mathcal{X} : \frac{\Lambda(x)}{\Psi(x)} > e^\epsilon\} \quad \text{and} \quad \mathcal{X}_H = \{x \in \mathcal{X} : \Lambda(x) > e^{(1-\lambda)\epsilon} \text{ or } \Psi(x) < e^{-\lambda\epsilon}\}.$$

Based on the remark following (11), it is enough to prove that $\mathcal{X}'_H \subseteq \mathcal{X}_H$. If $x \in \mathcal{X}'_H$, for any given $\lambda \in (0, 1)$, there are only two possible cases: either $\Lambda(x) > e^{(1-\lambda)\epsilon}$ or $\Lambda(x) \leq e^{(1-\lambda)\epsilon}$. If $\Lambda(x) > e^{(1-\lambda)\epsilon}$, then $x \in \mathcal{X}_H$, and our claim is true. If $\Lambda(x) \leq e^{(1-\lambda)\epsilon}$, we have $x \in \mathcal{X}'_H \Rightarrow \frac{\Lambda(x)}{\Psi(x)} > e^\epsilon \Rightarrow \Psi(x) < \Lambda(x)e^{-\epsilon}$. We assumed that $\Lambda(x) \leq e^{(1-\lambda)\epsilon}$; therefore, $\Psi(x) < \Lambda(x)e^{-\epsilon} \leq e^{(1-\lambda)\epsilon}e^{-\epsilon} = e^{-\lambda\epsilon}$. Since $\Psi(x) < e^{-\lambda\epsilon}$, we have $x \in \mathcal{X}_H$. \square

This proposition shows the result of applying LDP and ALIP in the watchdog mechanism in terms of the privacy–utility trade-off. While both ϵ -LDP and $(\lambda\epsilon, (1 - \lambda)\epsilon)$ -ALIP imply the same LDP privacy budget, LDP results in fewer high-risk symbols compared to ALIP. This needs to be considered when one applies the watchdog mechanism to achieve LDP or ALIP privacy. Although having fewer high-risk symbols provides better utility, it may compromise privacy. In other words, when $\mathcal{X}'_H \subseteq \mathcal{X}_H$, then the privacy leakage of the partition $\{\mathcal{X}'_L, \mathcal{X}'_H\}$ is greater than or equal to the privacy leakage of the partition $\{\mathcal{X}_L, \mathcal{X}_H\}$. As a result, it is possible that ϵ -LDP cannot achieve the desired $(\lambda\epsilon, (1 - \lambda)\epsilon)$ -ALIP privacy level for a given λ .

Watchdog mechanism with X -invariant randomisation is a powerful method with low complexity that can be easily applied to instance-wise measures. However, it significantly degrades the utility [25] because X -invariant randomisation obfuscates all high-risk symbols together to minimise privacy leakage, with the cost of deteriorating data resolution. In

Section 4, we propose subset merging randomisation to enhance the utility of the watchdog mechanism.

3.1.2. Asymmetric Optimal Random Response (AORR)

ORR was proposed in [28] as a localised instance-wise replacement of the privacy funnel [12]. It is the solution to the optimal utility problem subject to ε -LIP or ε -LDP constraints. For ALIP, we propose asymmetric optimal random response (AORR), which is defined as

$$\begin{aligned} & \max_{P_{X|Y}, P_Y} I(X; Y) \\ \text{s.t. } & S - X - Y \\ & e^{-\varepsilon_l} \leq \Psi(y) \quad \text{and} \quad \Lambda(y) \leq e^{\varepsilon_u}, \quad \forall y \in \mathcal{Y}. \end{aligned} \quad (12)$$

Privacy constraints in this optimisation problem form a closed, bounded, convex polytope [28]. It has been proved that vertices of this polytope are the feasible candidates that maximise MI and satisfy privacy constraints [7,28,37]. However, the number of vertices grows exponentially in the dimension of the polyhedron, which is $|\mathcal{X}|(|\mathcal{X}| - 1)$ for LDP and $(|\mathcal{X}| - 1)$ for LIP. This makes ORR computationally cumbersome for large $|\mathcal{X}|$. Accordingly, [28] suggests some approaches with lower complexity than ORR to avoid vertex enumeration for the larger sizes of \mathcal{X} , but this comes at the cost of lower utility.

3.1.3. Numerical Results

Here, we demonstrate the privacy leakage and utility of AORR and the watchdog mechanism under ALIP. For the utility, we use normalised MI (NMI)

$$\text{NMI} = \frac{I(X; Y)}{H(X)} \in [0, 1].$$

It is clear that the maximum possible utility is obtained when X is published without randomisation, where $Y = X$ and $I(X; Y) = H(X)$. Thus, $I(X; Y) \leq H(X)$ and $\text{NMI} \leq 1$.

We present numerical results for both synthetic and real datasets using MATLAB. For synthetic data, we randomly generated 10^3 distributions for the watchdog mechanism and 100 distributions for the AORR where $|\mathcal{X}| = 17$ and $|\mathcal{S}| = 5$. These distributions were generated by normalising the output of the *rand* function in MATLAB. For real datasets, we used the Adult dataset [38] and set $S = \{\text{relationship}\}$ and $X = \{\text{Occupation}\}$, where $|\mathcal{S}| = 5$ and $|\mathcal{X}| = 15$. In all scenarios, ε varies from 0.25 to 8, and we consider three cases for $(\varepsilon_u, \varepsilon_l)$ -ALIP, where $\lambda \in \{0.35, 0.5, 0.65\}$, $\varepsilon_l = \lambda\varepsilon$, and $\varepsilon_u = (1 - \lambda)\varepsilon$. The results of the watchdog mechanism are shown in Figures 2 and 3 for synthetic and real data, respectively; while the AORR results are presented in Figures 4 and 5. The figures display NMI, $\log(\max_y \Lambda(y))$ (max-lift leakage), and $|\log(\min_y \Psi(y))|$ (min-lift leakage) versus the LDP privacy budget ε for real data, and the mean values of the same quantities are shown for synthetic data.

In Figures 2a and 3a, we observe that in the watchdog mechanism, LDP provides higher utility and leakage than ALIP for all values of ε and λ , which confirms Proposition 2. Figures 2a–5a demonstrate that the min-lift relaxation, $\lambda = 0.65$, enhances utility in the watchdog and AORR mechanisms for $\varepsilon > 1$. Note that in all figures, $\lambda = 0.5$ refers to $\frac{\varepsilon}{2}$ -LIP. On the other hand, $\lambda = 0.35$ results in lower utility. Generally, any value of $\lambda < 0.5$ reduces utility since it strictly bounds the min-lift while relaxing the max-lift. As the min-lift has a wider range of values, achieving this strict bound enlarges the set \mathcal{X}_H and requires randomising more symbols, which reduces utility. Another observation here is that AORR incurs significantly higher utility than the watchdog mechanism. For instance, when $\lambda = 0.5$ and $\varepsilon = 2$, the watchdog mechanism results in a utility of 0.52 for synthetic data and 0.73 for the real data, while AORR has a utility of 0.94 and 0.96 for the synthetic and real data, respectively. AORR finds the optimal utility, which, due to PUT, necessarily results in

the highest leakage subject to privacy constraints. However, the watchdog mechanism is a nonoptimal solution that minimises leakage of high-risk symbols to provide strong privacy protection, which deteriorates utility. To solve this drawback of the watchdog mechanism, we propose a subset randomisation method in the following section.




Figure 2. Privacy–utility trade-off of the watchdog mechanism with complete merging randomisation for synthetic data under ϵ -LDP and (ϵ_l, ϵ_u) -ALIP, where $|\mathcal{X}| = 17$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.35, 0.5, 0.65\}$, $\epsilon_l = \lambda\epsilon$, and $\epsilon_u = (1 - \lambda)\epsilon$.




Figure 3. Privacy–utility trade-off of the watchdog mechanism with complete merging randomisation for Adult dataset under ϵ -LDP and (ϵ_l, ϵ_u) -ALIP, where $S = \{\text{relationship}\}$, $X = \{\text{occupation}\}$, $|\mathcal{X}| = 15$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.35, 0.5, 0.65\}$, $\epsilon_l = \lambda\epsilon$, and $\epsilon_u = (1 - \lambda)\epsilon$.




Figure 4. Privacy–utility trade-off of AORR for synthetic data where $|\mathcal{X}| = 17$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.35, 0.5, 0.65\}$, $\epsilon_l = \lambda\epsilon$, and $\epsilon_u = (1 - \lambda)\epsilon$.




Figure 5. Privacy–utility trade-off of AORR for Adult dataset where $S = \{\text{relationship}\}$, $X = \{\text{occupation}\}$, $|\mathcal{X}| = 15$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.35, 0.5, 0.65\}$, $\epsilon_l = \lambda\epsilon$, and $\epsilon_u = (1 - \lambda)\epsilon$.

4. Subset Merging in Watchdog Mechanism

The watchdog mechanism with X -invariant randomisation is a low-complexity method that can be easily applied when the privacy measures are symbol-wise. X -invariant randomisation is the optimal privacy protection for the high-risk symbols that minimises privacy leakage in \mathcal{X}_H and necessarily results in the worst data resolution. Thus, in this section, we propose the *subset merging* algorithm to improve data resolution by randomising disjoint subsets of high-risk symbols and enhancing utility in the watchdog mechanism. In the following, we show that applying X -invariant randomisation to disjoint subsets of \mathcal{X}_H increases the utility.

Let $\mathcal{G}_{\mathcal{X}_H} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_g\}$ be a partition of \mathcal{X}_H , where for every $i \in [g]$, $\mathcal{X}_i \subseteq \mathcal{X}_H$: $\mathcal{X}_i \cap \mathcal{X}_j = \emptyset, i \neq j$, and $\mathcal{X}_H = \bigcup_{i=1}^g \mathcal{X}_i$. We randomise each subset $\mathcal{X}_i \in \mathcal{G}_{\mathcal{X}_H}$ by X -invariant randomisation $\mathcal{R}_{\mathcal{Y}_i}(y)$ for $x \in \mathcal{X}_i$ and $y \in \mathcal{Y}_i$, where $\sum_{y \in \mathcal{Y}_i} \mathcal{R}_{\mathcal{Y}_i}(y) = 1$. The resulting MI between X and Y is

$$I(X; Y) = H(X) - \sum_{i=1}^g \sum_{x \in \mathcal{X}_i} P_X(x) \log \frac{P(\mathcal{X}_i)}{P_X(x)}. \quad (13)$$

Definition 4. Assume two partitions, $\mathcal{G}_{\mathcal{X}_H} = \{\mathcal{X}_1, \dots, \mathcal{X}_g\}$ and $\mathcal{G}'_{\mathcal{X}_H} = \{\mathcal{X}'_1, \dots, \mathcal{X}'_{g'}\}$. We say that $\mathcal{G}'_{\mathcal{X}_H}$ is a refinement of $\mathcal{G}_{\mathcal{X}_H}$, or $\mathcal{G}_{\mathcal{X}_H}$ is an aggregation of $\mathcal{G}'_{\mathcal{X}_H}$, if for every $i \in [g]$, $\mathcal{X}_i = \bigcup_{j \in J_i} \mathcal{X}'_j$ where $J_i \subseteq [g']$, and $P(\mathcal{X}_i) = \sum_{j \in J_i} P(\mathcal{X}'_j)$ (this definition is inspired from [39] (Definition 10)).

If $\mathcal{G}'_{\mathcal{X}_H}$ is a refinement of $\mathcal{G}_{\mathcal{X}_H}$, then $I_{\mathcal{G}_{\mathcal{X}_H}}(X; Y) \leq I_{\mathcal{G}'_{\mathcal{X}_H}}(X; Y)$.

Obtaining the optimal $\mathcal{G}_{\mathcal{X}_H}$ that maximises utility and satisfies privacy constraints is a combinatorial optimisation problem over all possible partitions of \mathcal{X}_H , which is cumbersome to solve. Therefore, we propose a heuristic method in the following.

4.1. Greedy Algorithm to Make Refined Subsets of High-Risk Symbols

In Algorithm 1, we propose a bottom-up algorithm that constitutes a partition of \mathcal{X}_H by merging high-risk symbols in disjoint subsets. It works based on a leakage risk metric for each $x \in \mathcal{X}_H$: $\omega(x) = \Lambda(x) + \Psi(x)$ for ALIP and $\omega(x) = \Gamma(x)$ for LDP. For LIP, $\omega(x) = \max\{\log \Lambda(x), |\log \Psi(x)|\}$. This metric is used to order the subsets by the privacy risk level. Accordingly, to constitute a subset $\mathcal{X}_i \subseteq \mathcal{X}_H$, Algorithm 1 bootstraps from the highest risk symbol $\mathcal{X}_i = \{\arg\max_{x \in \mathcal{X}_H} \omega(x)\}$ (line 5). Then, it merges a symbol x^* with \mathcal{X}_i that minimises $\omega(\mathcal{X}_i \cup x^*)$ (line 7), as long as the privacy constraints are satisfied in \mathcal{X}_i (line 6). The ALIP privacy constraints for a subset \mathcal{X}_i are given by

$$e^{-\epsilon_l} \leq \Psi(\mathcal{X}_i) \quad \text{and} \quad \Lambda(\mathcal{X}_i) \leq e^{\epsilon_u}, \quad (14)$$

where $\Psi(\mathcal{X}_i) = \min_{s \in \mathcal{S}} \frac{\sum_{x \in \mathcal{X}_i} P_{X|S}(x|s)}{\sum_{x \in \mathcal{X}_i} P_X(x)}$ and $\Lambda(\mathcal{X}_i) = \max_{s \in \mathcal{S}} \frac{\sum_{x \in \mathcal{X}_i} P_{X|S}(x|s)}{\sum_{x \in \mathcal{X}_i} P_X(x)}$. For LDP constraint, we have $\Gamma(\mathcal{X}_i) = \frac{\Lambda(\mathcal{X}_i)}{\Psi(\mathcal{X}_i)} \leq e^\epsilon$. In Algorithm 1, we used ALIP privacy constraints for the while loops condition in lines 4, 6, and 12. For LDP, the privacy constraint is changed to $\Gamma(\mathcal{X}_Q) > \epsilon$, and $\omega(x)$ for LDP is applied. After the constitution of the partition $\mathcal{G}_{\mathcal{X}_H}$, the last subset \mathcal{X}_g may not meet privacy constraints. Therefore, the leakage of \mathcal{X}_g is checked (line 12), and if there is a privacy breach, an agglomerate \mathcal{X}_g is made by merging other subsets to it that minimises subset risk, $\omega(\mathcal{X}_g) = \Lambda(\mathcal{X}_g) + \Psi(\mathcal{X}_g)$ (lines 13–14), until privacy constraints are satisfied.

Algorithm 1: Subset merging in the watchdog mechanism.

```

1 Input:  $\mathcal{X}, \varepsilon_l, \varepsilon_u, P_{SX}$ .
2 Output:  $\mathcal{G}_{\mathcal{X}_H} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_g\}$ .
3 Initialize: Obtain  $\{\mathcal{X}_L, \mathcal{X}_H\}$  in (6),  $\mathcal{X}_Q \leftarrow \mathcal{X}_H$ , and  $g = 1$ .
4 while  $(\Psi(\mathcal{X}_Q) < e^{-\varepsilon_l} \text{ or } \Lambda(\mathcal{X}_Q) > e^{\varepsilon_u}) \text{ and } |\mathcal{X}_Q| > 0$  do
5    $\mathcal{X}_g = \arg\max_{x \in \mathcal{X}_Q} \omega(x)$ , and  $\mathcal{X}_Q \leftarrow \mathcal{X}_Q \setminus \mathcal{X}_g$ ;
6   while  $(\Psi(\mathcal{X}_g) < e^{-\varepsilon_l} \text{ or } \Lambda(\mathcal{X}_g) > e^{\varepsilon_u}) \text{ and } |\mathcal{X}_Q| > 0$  do
7      $x^* = \arg\min_{x \in \mathcal{X}_Q} \omega(\mathcal{X}_g \cup \{x\})$ ;
8      $\mathcal{X}_g \leftarrow \mathcal{X}_g \cup \{x^*\}$ , and  $\mathcal{X}_Q \leftarrow \mathcal{X}_Q \setminus \{x^*\}$ ;
9   end
10   $\mathcal{G}_{\mathcal{X}_Q} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_g\}$ , and  $g \leftarrow g + 1$ ;
11 end
12 while  $(\Psi(\mathcal{X}_g) < e^{-\varepsilon_l} \text{ or } \Lambda(\mathcal{X}_g) > e^{\varepsilon_u}) \text{ and } |\mathcal{G}_{\mathcal{X}_Q}| > 1$  do
13    $i^* = \arg\min_{1 \leq i < g} \omega(\mathcal{X}_g \cup \mathcal{X}_i)$ , and  $\mathcal{X}_g \leftarrow \mathcal{X}_g \cup \mathcal{X}_{i^*}$ ;
14   For  $i^* + 1 \leq j \leq g$  update the indices of  $\mathcal{X}_j$ 's to  $\mathcal{X}_{j-1}$  and  $g \leftarrow g - 1$ ;
15    $\mathcal{G}_{\mathcal{X}_Q} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_g\}$ ;
16 end

```

4.2. Numerical Results

We show PUT for ALIP and LDP under subset merging randomisation in Figures 6 and 7 for synthetic and real data, respectively, with the same setup for the watchdog mechanism in Section 3.1.3. Compared with the complete merging (Figures 2 and 3), the utility was enhanced significantly for both LDP and ALIP in all scenarios under the same privacy constraint. For instance, consider the symmetric case $\lambda = 0.5$ when $\varepsilon = 1$, and compare PUT between the subset and complete merging. Figures 6a and 7a demonstrate a utility value of around 0.73 for the subset merging compared to the utilities of 0.17 and 0.28 for the complete merging in Figures 2a and 3a, which are almost 320% and 160% utility enhancement. Moreover, as Figures 6b,c and 7b,c illustrate, privacy constraints are satisfied in all cases.




Figure 6. Privacy–utility trade-off of subset merging randomisation under ε -LDP and $(\varepsilon_l, \varepsilon_u)$ -ALIP, where $|\mathcal{X}| = 17$, $|\mathcal{S}| = 5$, $\varepsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.35, 0.5, 0.65\}$, $\varepsilon_l = \lambda\varepsilon$, and $\varepsilon_u = (1 - \lambda)\varepsilon$.




Figure 7. Privacy–utility trade-off of subset merging randomisation for Adult dataset under ϵ -LDP and (ϵ_l, ϵ_u) -ALIP, where $S = \{\text{relationship}\}$, $X = \{\text{occupation}\}$, $|\mathcal{X}| = 15$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.35, 0.5, 0.65\}$, $\epsilon_l = \lambda\epsilon$, and $\epsilon_u = (1 - \lambda)\epsilon$.

5. Subset Random Response

In the previous section, we showed that subset merging enhances utility in the watchdog mechanism significantly. In this section, we propose a method to decrease the complexity of AORR for large datasets. We adopt AORR for subsets of \mathcal{X}_H to decrease the complexity of AORR for large sets such that random response becomes applicable for typically an order of magnitude larger \mathcal{X} .

The AORR optimisation problem in (12) is equivalent to the following problem:

$$H(x) - \min_{P_{X|Y}, P_Y} H(X|Y) \quad (15)$$

$$\text{s.t. } S - X - Y$$

$$e^{-\epsilon_l} \leq \Psi(y) \quad \text{and} \quad \Lambda(y) \leq e^{\epsilon_u}, \quad \forall y \in \mathcal{Y}.$$

To reduce the complexity of (15), we divide \mathcal{X} into \mathcal{X}_L and \mathcal{X}_H , similar to the watchdog mechanism, and make a partition $\mathcal{G}_{\mathcal{X}_H} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_g\}$ from \mathcal{X}_H . We randomise each subset $\mathcal{X}_i \in \mathcal{G}_{\mathcal{X}_H}$, $i \in [g]$, separately by a randomisation pair $(\mathbf{Q}_i, \mathbf{q}_i)$, where \mathbf{Q}_i is a matrix in $\mathbb{R}^{|\mathcal{X}_i| \times |\mathcal{Y}_i|}$ and \mathbf{q}_i is a vector in $\mathbb{R}^{|\mathcal{Y}_i|}$.

The elements of \mathbf{Q}_i and \mathbf{q}_i are given by

$$\mathbf{Q}_i(x|y) = \Pr[X = x|Y = y], \quad x \in \mathcal{X}_i, y \in \mathcal{Y}_i, \quad (16)$$

$$\mathbf{q}_i(y) = \Pr[Y = y], \quad y \in \mathcal{Y}_i. \quad (17)$$

For each $y \in \mathcal{Y}_i$, we have $\sum_{x \in \mathcal{X}_i} \mathbf{Q}_i(x|y) = 1$. Consequently, $H(X|Y) = \sum_{i \in [g]} H_i(X|Y)$, where

$$H_i(X|Y) = - \sum_{y \in \mathcal{Y}_i} \mathbf{q}_i(y) \sum_{x \in \mathcal{X}_i} \mathbf{Q}_i(x|y) \log \mathbf{Q}_i(x|y), \quad i \in [g]. \quad (18)$$

This setting turns (15) into g optimisation problems for each subset $\mathcal{X}_i \in \mathcal{G}_{\mathcal{X}_H}$, $i \in [g]$ as follows:

$$\min_{\mathbf{Q}_i, \mathbf{q}_i} H_i(X|Y) \quad (19)$$

$$\text{s.t. } 0 \leq \mathbf{q}_i(y), \quad \forall y \in \mathcal{Y}_i, \quad (20)$$

$$0 \leq \mathbf{Q}_i(x|y), \quad \forall x \in \mathcal{X}_i, \forall y \in \mathcal{Y}_i, \quad (21)$$

$$\sum_{x \in \mathcal{X}_i} \mathbf{Q}_i(x|y) = 1, \quad \forall y \in \mathcal{Y}_i, \quad (22)$$

$$\sum_{y \in \mathcal{Y}_i} \mathbf{Q}_i(x|y) \mathbf{q}_i(y) = P_X(x), \quad x \in \mathcal{X}_i, \quad (23)$$

$$e^{-\epsilon_l} P_S(s) \leq \sum_{x \in \mathcal{X}_i} P_{S|X}(s|x) \mathbf{Q}_i(x|y) \leq e^{\epsilon_u} P_S(s), \quad \forall s \in \mathcal{S}, y \in \mathcal{Y}_i. \quad (24)$$

The columns of the randomisation matrix \mathbf{Q}_i , $i \in [g]$ can be expressed as the members of a convex and bounded polytope Π_i , which is given by the following constraints:

$$\Pi_i = \left\{ \begin{array}{l} \mathbf{v} \in \mathbb{R}^{|\mathcal{X}_i|} : \\ 0 \leq v_k, \forall k \in [|\mathcal{X}_i|], \\ \sum_{k=1}^{|\mathcal{X}_i|} v_k = 1, \\ e^{-\varepsilon_l} P_S(s) \leq \sum_{x \in \mathcal{X}_i} P_{S|X}(s|x) v_k \leq e^{\varepsilon_u} P_S(s), \forall s \in \mathcal{S}, k \in [|\mathcal{X}_i|] \end{array} \right\}. \quad (25)$$

For each $\mathcal{X}_i \in \mathcal{G}_{\mathcal{X}_H}$, $i \in [g]$, let $\mathcal{V}_i = \{\mathbf{v}_1^i \cdots, \mathbf{v}_M^i\}$ be the vertices of Π_i in (25), $H(\mathbf{v}_k^i)$ be the entropy of each \mathbf{v}_k^i for $k \in [M]$, \mathbf{P}_X^i be the probability vector of $x \in \mathcal{X}_i$, and β^i be the solution to the following optimisation:

$$\begin{aligned} & \min_{\beta^i \in \mathbb{R}^M} \sum_{k=1}^M H(\mathbf{v}_k^i) \beta_k^i, \\ \text{s.t. } & \beta_k^i \geq 0, \quad \forall k \in [M], \\ & \sum_{k=1}^M \mathbf{v}_k^i \beta_k^i = \mathbf{P}_X^i. \end{aligned} \quad (26)$$

Then, \mathcal{V}_i and $(\mathbf{Q}_i, \mathbf{q}_i)$ are given by

$$\mathcal{V}_i = \{y : \beta_y^i \neq 0\}; \quad (27)$$

$$\mathbf{q}_i(y) = \beta_y^i \quad \text{and} \quad \mathbf{Q}_i(\cdot|y) = \mathbf{v}_y^i, \quad y \in \mathcal{V}_i. \quad (28)$$

The $(\varepsilon_l, \varepsilon_u)$ -ALIP protocol, $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$, is given by the pair of $(P_{X|Y}, P_Y)$ as follows:

$$P_{X|Y} = \begin{cases} \mathbf{1}_{\{x=y\}}, & x, y \in \mathcal{X}_L = \mathcal{V}_L, \\ \mathbf{Q}_i(x|y), & x \in \mathcal{X}_i, y \in \mathcal{V}_i, i \in [g], \\ 0, & \text{otherwise;} \end{cases} \quad (29)$$

$$P_Y = \begin{cases} P_X(y), & y \in \mathcal{V}_L, \\ \mathbf{q}_i(y), & y \in \mathcal{V}_i, i \in [g]. \end{cases} \quad (30)$$

5.1. Algorithm for Subset Random Response

We propose Algorithm 2 to implement AORR for subsets of \mathcal{X}_H , which we call *subset random response* (SRR). In this algorithm, first, we obtain a partition $\mathcal{G}_{\mathcal{X}_H} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_g\}$ of \mathcal{X}_H via Algorithm 1. Then, based on $\mathcal{G}_{\mathcal{X}_H}$, we make another partition $\mathcal{O}_{\mathcal{X}_H}$ and find the optimal random response for each subset $\mathcal{X}'_i \in \mathcal{O}_{\mathcal{X}_H}$ (line 5). By obtaining the optimal random responses for all subsets, we obtain a pair $(\mathbf{Q}_i, \mathbf{q}_i)$ for each subset \mathcal{X}'_i , and consequently $P_{X|Y}$ and P_Y , by (29) and (30) (lines 20–23). The while loop in lines 7–14 is for the particular cases when the polytope in (25) is empty for a subset \mathcal{X}'_i . It may occur for strict privacy conditions where the privacy budget is very small. Since we reduce the dimension of the original polytope in (15), it increases the possibility that no feasible random response exists in some cases. Therefore, in such cases, we make a union with other subsets until we have a nonempty polytope. If $|\mathcal{G}_{\mathcal{X}_H}| > 0$, then we make a union with another subset in $\mathcal{G}_{\mathcal{X}_H}$ (line 9). If $|\mathcal{G}_{\mathcal{X}_H}| = 0$, it means that \mathcal{X}'_g is the last subset in $\mathcal{O}_{\mathcal{X}_H}$; therefore, we make a union with previously made subsets in $\mathcal{O}_{\mathcal{X}_H}$ and update the index g' (line 12). Then, if the condition in lines 17–18 is for the cases where there is no feasible polytope after making a union of all subsets, this means that the problem in (15) cannot be solved. Whenever this occurs, we apply the subset merging mechanism.

The number of polytope vertices in AORR is $n \sim O(\exp(|\mathcal{X}| - 1))$, and the time complexity is $O(n)$. As $|\mathcal{X}|$ increases, the complexity of AORR increases exponentially in the

$|\mathcal{X}| - 1$. In SRR, for each $\mathcal{X}_i \in \mathcal{G}_{\mathcal{X}_H}$, $i \in [g]$, the number of vertices is $n_i \sim O(\exp(|\mathcal{X}_i| - 1))$, and the complexity of SRR is $O(\max_i n_i)$.

Algorithm 2: Subset random response.

```

1 Input:  $\mathcal{X}, \varepsilon_l, \varepsilon_u, P_{SX}$ .
2 Output:  $\mathcal{O}_{\mathcal{X}_H} = \{\mathcal{X}'_1, \mathcal{X}'_2, \dots, \mathcal{X}'_{g'}\}$ ,  $P_{X|Y}$ , and  $P_Y$ .
3 Initialize: Obtain  $\mathcal{G}_{\mathcal{X}_H} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_g\}$  by Algorithm 1 and  $g' = 1$ .
4 while  $|\mathcal{G}_{\mathcal{X}_H}| > 0$  do
5    $\mathcal{X}'_{g'} = \mathcal{X}_1$ , find  $\Pi_{g'}$  in (25), and  $\mathcal{G}_{\mathcal{X}_H} \leftarrow \mathcal{G}_{\mathcal{X}_H} \setminus \mathcal{X}_1$ ;
6   Update subset indices in  $\mathcal{G}_{\mathcal{X}_H}$ , such that  $\mathcal{X}_{i-1} \leftarrow \mathcal{X}_i$  for  $2 \leq i \leq g$  and
    $g \leftarrow g - 1$ ;
7   while  $\Pi_{g'} = \emptyset$  and  $(|\mathcal{G}_{\mathcal{X}_H}| > 0$  or  $|\mathcal{O}_{\mathcal{X}_H}| > 0)$  do
8     if  $|\mathcal{G}_{\mathcal{X}_H}| > 0$  then
9        $\mathcal{X}'_{g'} = \mathcal{X}'_{g'} \cup \mathcal{X}_1$ , find  $\Pi_{g'}$  in (25), and  $\mathcal{G}_{\mathcal{X}_H} \leftarrow \mathcal{G}_{\mathcal{X}_H} \setminus \mathcal{X}_1$ ;
10      Update subset indices in  $\mathcal{G}_{\mathcal{X}_H}$ , such that  $\mathcal{X}_{i-1} \leftarrow \mathcal{X}_i$  for  $2 \leq i \leq g$  and
       $g \leftarrow g - 1$ ;
11    else
12       $\mathcal{X}'_{g'} = \mathcal{X}'_{g'} \cup \mathcal{X}'_{g'-1}$ ,  $g' \leftarrow g' - 1$ , and find  $\Pi_{g'}$  in (25);
13    end
14  end
15  Set  $\mathcal{O}_{\mathcal{X}_H} = \{\mathcal{X}'_1, \mathcal{X}'_2, \dots, \mathcal{X}'_{g'}\}$  and  $g' \leftarrow g' + 1$ ;
16 end
17 if  $|\mathcal{O}_{\mathcal{X}_H}| = 1$  and  $\Pi_1 = \emptyset$  then
18   Apply subset merging mechanism for  $\mathcal{G}_{\mathcal{X}_H}$  in Algorithm 1.
19 else
20   for  $i \leftarrow 1$  to  $g'$  do
21     Solve optimisation in (26) for  $\mathcal{X}'_i$  and obtain  $\beta_i$ ,  $\mathbf{Q}_i$ , and  $\mathbf{q}_i$ ;
22   end
23   Obtain  $P_{X|Y}$  in (29) and  $P_Y$  in (30).
24 end

```

5.2. Numerical Results

Here, we compare the PUT of AORR with SRR in Algorithm 2 and subset merging in Algorithm 1. Figure 8 depicts mean values of utility, leakage, and time complexity for 100 randomly generated distributions where $\lambda = 0.65$ and simulation setup is the same as that in Section 3.1.3. The result of the Adult dataset is also shown in Figure 9.

Figures 8a–c and 9 demonstrate that SRR results in better utility and higher leakage than subset merging, and its PUT is very close to AORR. Figure 8d illustrates the processing time of each mechanism for synthetic data from which we observe that the complexity of AORR and SRR is much higher than the subset merging. Running SRR is less complex than AORR for strict privacy constraints ($\varepsilon < 1$) and for $\varepsilon > 2.5$. While SRR shows higher complexity for some privacy budgets, $1 \leq \varepsilon \leq 2.5$, it has the advantage in high-dimension systems. Figure 10 shows a PUT comparison between SRR and subset merging for synthetic data where $|\mathcal{X}| = 200$, $|\mathcal{S}| = 15$, $\varepsilon \in \{1, 1.25, \dots, 8\}$, and $\lambda = 0.5$. This experiment shows that both SRR and subset merging are applicable to large datasets. Obviously, SRR provides better utility (Figure 10a) and higher leakage (Figure 10b,c), which is still below the given budgets ε_l and ε_u .

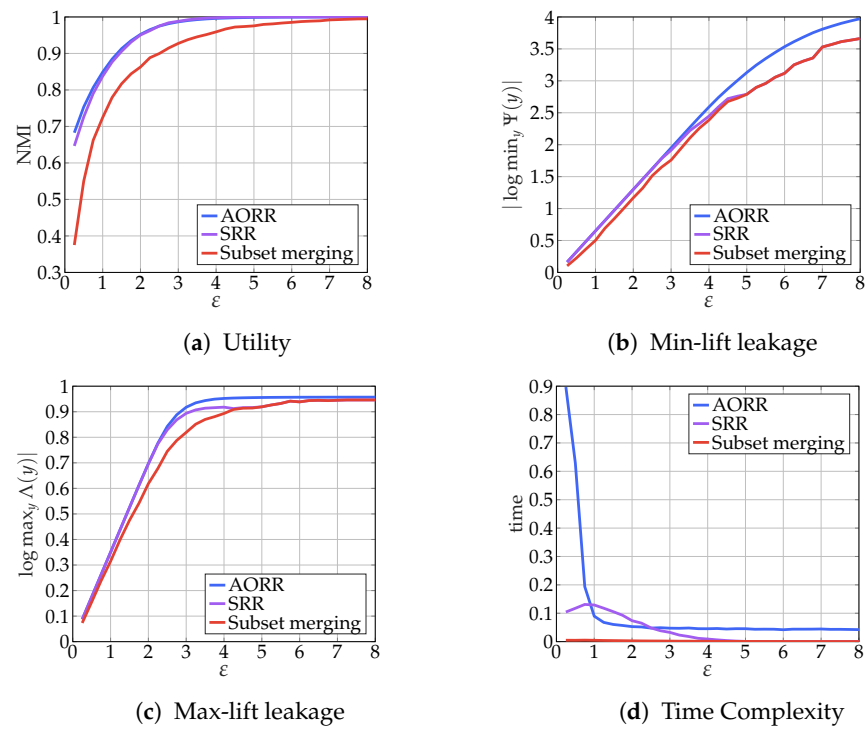


Figure 8. Comparison of privacy–utility trade-off and time complexity between AORR, SRR (Algorithm 2), and subset merging (Algorithm 1) for synthetic data, where $|\mathcal{X}| = 17$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda = 0.65$, $\epsilon_l = \lambda \epsilon_{\text{LDP}}$, and $\epsilon_u = (1 - \lambda) \epsilon_{\text{LDP}}$.




Figure 9. Comparison of privacy–utility trade-off between AORR, SRR (Algorithm 2), and subset merging (Algorithm 1) for Adult dataset, where $S = \{\text{relationship}\}$, $X = \{\text{occupation}\}$, $|\mathcal{X}| = 15$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda = 0.65$, $\epsilon_l = \lambda \epsilon_{\text{LDP}}$, and $\epsilon_u = (1 - \lambda) \epsilon_{\text{LDP}}$.




Figure 10. Comparison of privacy–utility trade-off and time complexity between SRR (Algorithm 2) and subset merging (Algorithm 1) for synthetic data, where $|\mathcal{X}| = 200$, $|\mathcal{S}| = 15$, $\epsilon_{\text{LDP}} \in \{1, 1.25, 1.5, 1.75, \dots, 8\}$, and $\epsilon_l = \epsilon_u = \frac{\epsilon_{\text{LDP}}}{2}$.

6. Lift-Based and Lift-Inverse Measures

In this section, we consider some recently proposed privacy measures that quantify the divergence between the posterior and prior belief on sensitive features, including ℓ_1 -norm [32], strong χ^2 -privacy criterion [33], and α -lift [34]. They have been proposed as a stronger version of their corresponding average measures, which are the total variation distance [35], χ^2 -divergence [40], and Sibson MI [16], respectively. We call them *lift-based* measures and define them in the following.

Definition 5. For each $y \in \mathcal{Y}$, lift-based privacy measures are defined as follows:

- The ℓ_1 -lift is given by

$$\Lambda_{\ell_1}(y) \triangleq \sum_{s \in \mathcal{S}} P_S(s) |l(s, y) - 1|, \quad (31)$$

then the total variation distance will be

$$T(S; Y) = \frac{1}{2} \mathbb{E}_Y[\Lambda_{\ell_1}(Y)].$$

- The χ^2 -lift is given by

$$\Lambda_{\chi^2}(y) \triangleq \sum_{s \in \mathcal{S}} P_S(s) (l(s, y) - 1)^2, \quad (32)$$

then χ^2 -divergence will be

$$\chi^2(S; Y) = \mathbb{E}_Y[\Lambda_{\chi^2}(Y)].$$

- The α -lift is given by

$$\Lambda_{\alpha}^S(y) \triangleq \left(\sum_{s \in \mathcal{S}} P_S(s) l(s, y)^{\alpha} \right)^{1/\alpha}, \quad (33)$$

then Sibson MI will be

$$I_{\alpha}^S(S; Y) = \frac{\alpha}{\alpha - 1} \log \mathbb{E}_Y[\Lambda_{\alpha}^S(Y)].$$

Here, we reveal their relationship with ALIP.

Proposition 3. If $(\varepsilon_l, \varepsilon_u)$ -ALIP is satisfied, then

1. $\max_{y \in \mathcal{Y}} \Lambda_{\ell_1}(y) \leq e^{\varepsilon_u} - 1$;
2. $\max_{y \in \mathcal{Y}} \Lambda_{\chi^2}(y) \leq (e^{\varepsilon_u} - 1)^2$;
3. $\max_{y \in \mathcal{Y}} \Lambda_{\alpha}^S(y) \leq e^{\varepsilon_u}$.

The proof is given in Appendix D.

Proposition 4. Let $s_y = \operatorname{argmax}_{s \in \mathcal{S}} [l(s, y)]$ and $\bar{y} = \operatorname{argmax}_{y \in \mathcal{Y}} \Lambda(y)$, then

1. If $\max_{y \in \mathcal{Y}} \Lambda_{\ell_1}(y) \leq \varepsilon \Rightarrow \max_{y \in \mathcal{Y}} \Lambda(y) \leq \varepsilon / P_S(s_{\bar{y}}) + 1$;
2. If $\max_{y \in \mathcal{Y}} \Lambda_{\chi^2}(y) \leq \varepsilon \Rightarrow \max_{y \in \mathcal{Y}} \Lambda(y) \leq \sqrt{\varepsilon / P_S(s_{\bar{y}})} + 1$;
3. If $\max_{y \in \mathcal{Y}} \Lambda_{\alpha}^S(y) \leq \varepsilon \Rightarrow \max_{y \in \mathcal{Y}} \Lambda(y) \leq \varepsilon / P_S(s_{\bar{y}})^{\frac{1}{\alpha}}$.

Proof. The proof is given in Appendix E. \square

Proposition 3 shows that lift-based measures, similar to their corresponding average leakages, are upper-bounded by the max-lift bound. Proposition 4 indicates that if we bound lift-based measures, they can only restrict the max-lift leakage. Accordingly, if one only applies a lift-based measure to protect privacy, such as in previous works [32–34],

it may cause significant leakage on the min-lift. Therefore, in the following, we propose lift-inverse measures to bound the min-lift leakage.

6.1. Lift-Inverse Measures

In Propositions 3 and 4, we showed that lift-based measures only bound the max-lift. In this subsection, we present lift-inverse measures to restrict the min-lift leakage, $\min_{y \in \mathcal{Y}} \Psi(y) = \min_{y \in \mathcal{Y}} [\min_{s \in \mathcal{S}} l(s, y)]$.

Definition 6. For lift-based measures in (31) to (33), we replace $l(s, y)$ with $\frac{1}{l(s, y)}$ and call the resulting quantities lift-inverse measures.

- The ℓ_1 -lift-inverse is given by

$$\Psi_{\ell_1}(y) = \sum_{s \in \mathcal{S}} P_S(s) \left| \frac{1}{\ell(s, y)} - 1 \right|. \quad (34)$$

- The χ^2 -lift-inverse is given by

$$\Psi_{\chi^2}(y) \triangleq \sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{\ell(s, y)} - 1 \right)^2. \quad (35)$$

- The α -lift-inverse is given by

$$\Psi_{\alpha}^S(y) \triangleq \left(\sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{\ell(s, y)} \right)^{\alpha} \right)^{1/\alpha}. \quad (36)$$

In the following propositions, we show the relationship between $(\varepsilon_l, \varepsilon_u)$ -ALIP.

Proposition 5. If $(\varepsilon_l, \varepsilon_u)$ -ALIP is achieved, we have

1. $\max_{y \in \mathcal{Y}} \Psi_{\ell_1}(y) \leq e^{\varepsilon_l} - 1$;
2. $\max_{y \in \mathcal{Y}} \Psi_{\chi^2}(y) \leq (e^{\varepsilon_l} - 1)^2$;
3. $\max_{y \in \mathcal{Y}} \Psi_{\alpha}^S(y) \leq e^{\varepsilon_l}$.

Proof. The proof is provided in Appendix F. \square

Proposition 6. Let $s_y = \operatorname{argmin}_s l(s, y)$ and $\underline{y} = \operatorname{argmin}_y [\Psi(y)]$, then

1. If $\max_{y \in \mathcal{Y}} \Psi_{\ell_1}(y) \leq \varepsilon \Rightarrow \min_{y \in \mathcal{Y}} \Psi(y) \geq \frac{P_S(s_y)}{\varepsilon + P_S(s_y)}$;
2. If $\max_{y \in \mathcal{Y}} \Psi_{\chi^2}(y) \leq \varepsilon \Rightarrow \min_{y \in \mathcal{Y}} \Psi(y) \geq \frac{\sqrt{P_S(s_y)}}{\sqrt{\varepsilon} + \sqrt{P_S(s_y)}}$;
3. If $\max_{y \in \mathcal{Y}} \Psi_{\alpha}^S(y) \leq \varepsilon \Rightarrow \min_{y \in \mathcal{Y}} \Psi(y) \geq \varepsilon^{-1} P_S(s_y)^{\frac{1}{\alpha}}$.

Proof. The proof is provided in Appendix G. \square

Propositions 5 and 6 demonstrate that the aforementioned lift-inverse measures are associated with the min-lift, and bounding them can restrict the min-lift leakage. Since lift-based and lift-inverse measures quantify privacy leakage by a function of lift averaged over sensitive features, they can be regarded as more relaxed measures than the min- and max-lifts.

6.2. PUT and Numerical Results

Optimal randomisations for ℓ_1 -lift and χ^2 -lift privacy, which maximise MI as the utility measure, were proposed in [32,33], respectively. Note that ORR is not applicable to these measures since their privacy constraints are not convex. However, here, we apply the watchdog mechanism with subset merging randomisation to investigate the PUT for lift-based and lift-inverse measures. This application shows that the watchdog mechanism with X-invariant randomisation is a low-complexity method that can be applied to all the aforementioned measures. Moreover, our subset merging algorithm significantly enhances the utility, which is comparable to the optimal solutions.

To apply lift-based and lift-inverse measures to the subset merging algorithm, we replace $\Lambda(y)$ and $\Psi(y)$ in (6) and Algorithm 1, with the corresponding lift-based and lift-inverse measures in Definitions 5 and 6, respectively. For example, \mathcal{X}_L and \mathcal{X}_H for the α -lift are obtained as

$$\mathcal{X}_L \triangleq \{x \in \mathcal{X} : \Psi_\alpha^S(x) \leq e^{\varepsilon_l} \text{ and } \Lambda_\alpha^S(x) \leq e^{\varepsilon_u}\} \text{ and } \mathcal{X}_H = \mathcal{X} \setminus \mathcal{X}_L. \quad (37)$$

The privacy risk measure in Algorithm 1 is also given by $\omega(x) = \Psi_\alpha^S(x) + \Lambda_\alpha^S(x)$. We compare the PUT of lift-based and lift-inverse privacy with $(\varepsilon_l, \varepsilon_u)$ -ALIP, where lift-based and lift-inverse measures are bounded as follows:

- ℓ_1 -privacy: $\Lambda_{\ell_1}(y) \leq e^{\varepsilon_u} - 1$ and $\Psi_{\ell_1}(y) \leq e^{\varepsilon_l} - 1, \quad \forall y \in \mathcal{Y}.$
- χ^2 -privacy: $\Lambda_{\chi^2}(y) \leq (e^{\varepsilon_u} - 1)^2$ and $\Psi_{\chi^2}(y) \leq (e^{\varepsilon_l} - 1)^2, \quad \forall y \in \mathcal{Y}.$
- α -lift-privacy: $\Lambda_\alpha^S(y) \leq e^{\varepsilon_u}$ and $\Psi_\alpha^S(y) \leq e^{\varepsilon_l}, \quad \forall y \in \mathcal{Y}.$

We apply the subset merging mechanism with the simulation setup in Section 3.1.3. Figures 11 and 12 demonstrate the PUT of ALIP for $\lambda = 0.5$ and ℓ_1 and χ^2 privacy for $\lambda = \{0.5, 0.65\}$ for synthetic and Adult dataset, respectively. When $\lambda = 0.5$, ℓ_1 and χ^2 privacy result in higher utility compared to ALIP for all values of ε since lift-based and lift-inverse measures are relaxations of max- and min-lift. To observe the effect of the asymmetric scenario, we depict ℓ_1 and χ^2 privacy for $\lambda = 0.65$. From Figures 11a and 12a, we observe that lift-inverse relaxation ($\lambda = 0.65$) enhances utility significantly for $\varepsilon > 1$, but worsens utility for $\varepsilon < 1$. The reason is that when $\varepsilon < 1$, lift-inverse privacy constraint in the asymmetric scenario is strict, which requires more symbols to be merged in each subset and causes larger subsets and utility degradation.

A comparison between α -lift privacy and ALIP for synthetic data is shown in Figure 13 for $\alpha \in \{2, 10, 100\}$. α -lift privacy is tunable such that when $\alpha = \infty$, it is equivalent to ALIP, and when $\alpha < \infty$, it results in a relaxation scenario. We observe tunable property in Figure 13 where $\alpha = 2$ has the highest utility. Moreover, when α increases, the PUT of α -lift privacy becomes closer to ALIP.




Figure 11. Comparison of privacy–utility trade-off between ALIP, ℓ_1 -privacy, and χ^2 -privacy for synthetic data, where $|\mathcal{X}| = 17$, $|\mathcal{S}| = 5$, $\varepsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.5, 0.65\}$, $\varepsilon_l = \lambda\varepsilon$, and $\varepsilon_u = (1 - \lambda)\varepsilon$.




Figure 12. Comparison of privacy–utility trade-off between ALIP, ℓ_1 -privacy, and χ^2 -privacy for Adult dataset where, $S = \{\text{relationship}\}$, $X = \{\text{occupation}\}$, $|\mathcal{X}| = 15$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\lambda \in \{0.5, 0.65\}$, $\epsilon_l = \lambda\epsilon$, and $\epsilon_u = (1 - \lambda)\epsilon$.




Figure 13. Comparison of privacy–utility trade-off between ALIP and α -lift-privacy, where $|\mathcal{X}| = 17$, $|\mathcal{S}| = 5$, $\epsilon_{\text{LDP}} \in \{0.25, 0.5, 0.75, \dots, 8\}$, $\epsilon_l = \epsilon_u = \frac{\epsilon_{\text{LDP}}}{2}$, and $\alpha \in \{2, 10, 100\}$.

7. Conclusions

In this paper, we studied lift, the likelihood ratio between posterior and prior belief about sensitive features in a dataset. We demonstrated the distinction between the min- and max-lifts in terms of data privacy concerns. We proposed ALIP as a generalised version of LIP to have a more compatible notion of privacy with lift asymmetry. ALIP can enhance utility in the watchdog and ORR mechanisms, two main approaches to achieve lift-based privacy. We proposed two subset randomisation methods to enhance the utility of the watchdog mechanism and reduce ORR complexity for large datasets. We also investigated the existing lift-based measures, showing that they could incur significant leakage on the min-lift. Thus, we proposed lift-inverse measures to restrict the min-lift leakage. Finally, we applied the watchdog mechanism to study the PUT of lift-based and lift-inverse measures. For future work, one can consider the applicable operational meaning of the min-lift and max-lift. Subset randomisation can be applied to decrease the complexity and enhance the utility of other privacy mechanisms. Moreover, optimal randomisation for α -lift is also unknown and could be considered.

Author Contributions: Conceptualization, N.D. and P.S.; Methodology, N.D. and P.S.; Validation, P.S.; Formal analysis, M.A.Z.; Writing—original draft, M.A.Z.; Writing—review & editing, N.D. and P.S.; Visualization, M.A.Z.; Supervision, N.D. and P.S.; Project administration, P.S.; Funding acquisition, P.S. All authors have read and agreed to the published version of the manuscript.

Funding: The work of P. Sadeghi and M.A. Zarrabian is supported by the ARC Future Fellowship FT190100429 and partly by the Data61 CRP: IT-PPUB.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: For real-world data simulations we have used public Adult dataset from UCI machine learning repository <https://archive-beta.ics.uci.edu/dataset/2/adult>.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Appendix A

In this appendix, we represent the privacy breach in terms of min-lift and max-lift based on the work of Evfimievsk et al. [29]. In this work, they proposed the following definition for privacy breach and provided a detailed example [29] (Example 1).

Definition A1. For $\rho_1, \rho_2 \in \mathbb{R}_+ : 0 < \rho_1 < \rho_2 < 1$, we say that there is a straight or upward ρ_1 -to- ρ_2 privacy breach with respect to a property $Q(S)$ if for some $y \in \mathcal{Y}$

$$\Pr[Q(S)] \leq \rho_1, \quad \Pr[Q(S)|Y = y] > \rho_2.$$

We say that there is a downward ρ_2 -to- ρ_1 privacy breach with respect to $Q(S)$ if for some $y \in \mathcal{Y}$

$$\Pr[Q(S)] \geq \rho_2, \quad \Pr[Q(S)|Y = y] < \rho_1.$$

Using the property $Q'(S) = \neg Q(S)$, we have

$$\Pr[Q'(S)] \leq 1 - \rho_2, \quad \Pr[Q'(S)|Y = y] \geq 1 - \rho_1.$$

Therefore, a downward ρ_2 -to- ρ_1 privacy breach implies an upward $(1 - \rho_2)$ -to- $(1 - \rho_1)$ for the complement event $Q'(S)$. In the original definition, the constraint for $\Pr[Q(S)|Y = y]$ includes equality. Here, we remove the equality constraint for consistency with the context. They proved that in the ϵ -LDP scenario, the sufficient condition to prevent both upward ρ_1 -to- ρ_2 and downward ρ_2 -to- ρ_1 privacy breach is

$$e^\epsilon \leq \frac{\rho_2}{\rho_1} \cdot \frac{1 - \rho_1}{1 - \rho_2}. \quad (\text{A1})$$

Here, we provide sufficient conditions to prevent upward ρ_1 -to- ρ_2 and downward ρ_2 -to- ρ_1 in terms of min- and max-lift as follows:

Proposition A1. To prevent upward and downward privacy breaches it is sufficient to have:

$$\frac{1 - \rho_2}{1 - \rho_1} \leq \Psi(y), \quad \Lambda(y) \leq \frac{\rho_2}{\rho_1}, \quad \forall y \in \mathcal{Y}.$$

Proof. We prove this for an upward privacy breach on a property $Q(S)$ and a downward privacy breach on its complement $\neg Q(S)$. For any property $Q(S)$ with $\Pr[Q(S)] \leq \rho_1$ we have

$$\Pr[Q(S)|Y = y] = \sum_{s \in Q(S)} P_{S|Y}(s|y) = \sum_{s \in Q(S)} \frac{P_S(s) \cdot P_{Y|S}(y|s)}{P_Y(y)} \leq \sum_{s \in Q(S)} P_S(s) \max_s \frac{P_{Y|S}(y|s)}{P_Y(y)} = \Lambda(y) \Pr[Q(S)].$$

Since $\Pr[Q(S)] \leq \rho_1$, if $\Lambda(y) \leq \frac{\rho_2}{\rho_1}$, then $\Pr[Q(S)|Y = y] \leq \rho_2$ and there is no upward privacy breach for $Q(S)$. Similarly, for downward privacy breach on $\neg Q(S)$, we have

$$\begin{aligned} \Pr[\neg Q(S)|Y = y] &= \sum_{s \in \neg Q(S)} P_{S|Y}(s|y) = \sum_{s \in \neg Q(S)} \frac{P_S(s) \cdot P_{Y|S}(y|s)}{P_Y(y)} \\ &\geq \sum_{s \in \neg Q(S)} P_S(s) \min_s \frac{P_{Y|S}(y|s)}{P_Y(y)} = \Psi(y) \Pr[\neg Q(S)]. \end{aligned}$$

Since $\Pr[\neg Q(S)] \geq 1 - \rho_1$, if $\Psi(y) \geq \frac{1-\rho_2}{1-\rho_1}$, then $\Pr[\neg Q(S)|Y = y] \geq 1 - \rho_2$ and there is no downward privacy breach for $\neg Q(S)$. \square

Appendix B

1. For MI, we have

$$I(S; Y) = \mathbb{E}_{P_{SY}}[i(S, Y)] \leq \mathbb{E}_{P_{SY}}[\varepsilon_u] = \varepsilon_u.$$

2. • For the total variation distance, we have

$$\begin{aligned} T(S; Y) &= \frac{1}{2} \sum_{y \in \mathcal{Y}} P_Y(y) \sum_{s \in \mathcal{S}} P_S(s) |l(s, y) - 1| \leq \frac{1}{2} \sum_{y \in \mathcal{Y}} P_Y(y) \sum_{s \in \mathcal{S}} P_S(s) |\Lambda(y) - 1| \\ &= \frac{1}{2} \sum_{y \in \mathcal{Y}} P_Y(y) |\Lambda(y) - 1| \leq \frac{1}{2} \sum_{y \in \mathcal{Y}} P_Y(y) |e^{\varepsilon_u} - 1| = \frac{1}{2} (e^{\varepsilon_u} - 1). \end{aligned}$$

• For χ^2 -divergence, we have

$$\begin{aligned} \chi^2(S; Y) &= \sum_{y \in \mathcal{Y}} P_Y(y) \sum_{s \in \mathcal{S}} P_S(s) (l(s, y) - 1)^2 \leq \sum_{y \in \mathcal{Y}} P_Y(y) \sum_{s \in \mathcal{S}} P_S(s) (\Lambda(y) - 1)^2 \\ &= \sum_{y \in \mathcal{Y}} P_Y(y) (\Lambda(y) - 1)^2 \leq \sum_{y \in \mathcal{Y}} P_Y(y) (e^{\varepsilon_u} - 1)^2 = (e^{\varepsilon_u} - 1)^2. \end{aligned}$$

3. • For Sibson MI, we have

$$\begin{aligned} I_\alpha^S(S; Y) &= \frac{\alpha}{\alpha - 1} \log \sum_{y \in \mathcal{Y}} P_Y(y) \left(\sum_{s \in \mathcal{S}} P_S(s) l(s, y)^\alpha \right)^{1/\alpha} \\ &\leq \frac{\alpha}{\alpha - 1} \log \sum_{y \in \mathcal{Y}} P_Y(y) \left(\sum_{s \in \mathcal{S}} P_S(s) \Lambda(y)^\alpha \right)^{1/\alpha} \\ &\leq \frac{\alpha}{\alpha - 1} \log \sum_{y \in \mathcal{Y}} P_Y(y) \left(\sum_{s \in \mathcal{S}} P_S(s) e^{\varepsilon_u \alpha} \right)^{1/\alpha} = \frac{\varepsilon_u \alpha}{\alpha - 1}. \end{aligned}$$

• For Arimoto MI, we have

$$\begin{aligned} I_\alpha^A(S; Y) &= \frac{\alpha}{\alpha - 1} \log \sum_{y \in \mathcal{Y}} P_Y(y) \left(\sum_{s \in \mathcal{S}} P_{S_\alpha}(s) l(s, y)^\alpha \right)^{1/\alpha} \\ &\leq \frac{\alpha}{\alpha - 1} \log \sum_{y \in \mathcal{Y}} P_Y(y) \left(\sum_{s \in \mathcal{S}} P_{S_\alpha}(s) \Lambda(y)^\alpha \right)^{1/\alpha} \\ &\leq \frac{\alpha}{\alpha - 1} \log \sum_{y \in \mathcal{Y}} P_Y(y) \left(\sum_{s \in \mathcal{S}} P_{S_\alpha}(s) e^{\varepsilon_u \alpha} \right)^{1/\alpha} = \frac{\varepsilon_u \alpha}{\alpha - 1}, \end{aligned}$$

$$\text{where } P_{S_\alpha}(s) = \frac{P_S(s)^\alpha}{\sum_{s \in \mathcal{S}} P_S(s)^\alpha}.$$

4. In LDP, for all $y \in \mathcal{Y}$, we have

$$\Gamma(y) = \sup_{s, s' \in \mathcal{S}} \frac{P_{Y|S}(y|s)}{P_{Y|S}(y|s')} = \frac{\max_{s \in \mathcal{S}} P_{Y|S}(y|s)}{\min_{s \in \mathcal{S}} P_{Y|S}(y|s)} = \frac{\max_{s \in \mathcal{S}} P_{Y|S}(y|s) / P_Y(y)}{\min_{s \in \mathcal{S}} P_{Y|S}(y|s) / P_Y(y)} = \frac{\Lambda(y)}{\Psi(y)} \leq \frac{e^{\varepsilon_u}}{e^{-\varepsilon_l}} = e^{\varepsilon_l + \varepsilon_u}.$$

Appendix C

Here, we prove that X -invariant randomisation minimises privacy leakage in \mathcal{X}_H for LDP.

Proposition A2. A randomisation $r(y|x)$, $x, y \in \mathcal{X}_H$ can attain $(\varepsilon, \mathcal{X}_H)$ -LDP if and only if:

$$\Gamma_{LDP}(\mathcal{X}_H) = \frac{\max_{s \in \mathcal{S}} P(\mathcal{X}_H|s)}{\min_{s \in \mathcal{S}} P(\mathcal{X}_H|s)} \leq e^\varepsilon. \quad (\text{A2})$$

Proof. Sufficient condition: Consider an X -invariant randomization where $r(y|x) = \mathcal{R}(y)$, $\forall x \in \mathcal{X}_H$, and $y \in \mathcal{Y}_H$. If (A2) holds, then for all $s, s' \in \mathcal{S}$, we have

$$\begin{aligned} \frac{P(\mathcal{X}_H|s)}{P(\mathcal{X}_H|s')} &\leq \frac{\max_{s \in \mathcal{S}} P(\mathcal{X}_H|s)}{\min_{s \in \mathcal{S}} P(\mathcal{X}_H|s)} \leq e^\varepsilon \Rightarrow \\ P(\mathcal{X}_H|s) &\leq P(\mathcal{X}_H|s')e^\varepsilon \Rightarrow \\ \mathcal{R}(y)P(\mathcal{X}_H|s) &\leq \mathcal{R}(y)P(\mathcal{X}_H|s')e^\varepsilon \Rightarrow \\ \sum_{x \in \mathcal{X}_H} r(y|x)P_{X|S}(x|s) &\leq \sum_{x \in \mathcal{X}_H} r(y|x)P_{X|S}(x|s')e^\varepsilon \Rightarrow \\ P_{Y|S}(y|s) &\leq P_{Y|S}(y|s')e^\varepsilon. \end{aligned}$$

For the necessary condition, note that for all $s, s' \in \mathcal{S}$ and $y \in \mathcal{Y}_H$, we have

$$P_{Y|S}(y|s) \leq e^\varepsilon P_{Y|S}(y|s') \Rightarrow \sum_{x \in \mathcal{X}_H} r(y|x)P_{X|S}(x|s) \leq e^\varepsilon \sum_{x \in \mathcal{X}_H} r(y|x)P_{X|S}(x|s'),$$

then by a summation over all $y \in \mathcal{Y}_H$ on both sides, we obtain

$$\begin{aligned} \sum_{x \in \mathcal{X}_H} P_{X|S}(x|s) \underbrace{\sum_{y \in \mathcal{Y}_H} r(y|x)}_{=1} &\leq e^\varepsilon \sum_{x \in \mathcal{X}_H} P_{X|S}(x|s') \underbrace{\sum_{y \in \mathcal{Y}_H} r(y|x)}_{=1} \\ \Rightarrow P(\mathcal{X}_H|s) &\leq P(\mathcal{X}_H|s')e^\varepsilon \Rightarrow \frac{P(\mathcal{X}_H|s)}{P(\mathcal{X}_H|s')} \leq e^\varepsilon. \end{aligned} \quad (\text{A3})$$

Because (A3) holds for all $s, s' \in \mathcal{S}$, we have

$$\max_{s, s' \in \mathcal{S}} \frac{P(\mathcal{X}_H|s)}{P(\mathcal{X}_H|s')} = \frac{\max_{s \in \mathcal{S}} P(\mathcal{X}_H|s)}{\min_{s \in \mathcal{S}} P(\mathcal{X}_H|s)} \leq e^\varepsilon.$$

□

Appendix D

1. For ℓ_1 -lift, we have

$$\Lambda_{\ell_1}(y) = \sum_{s \in \mathcal{S}} P_S(s) |l(s, y) - 1| \leq \sum_{s \in \mathcal{S}} P_S(s) |\Lambda(y) - 1| = \Lambda(y) - 1 \leq e^{\varepsilon_u} - 1.$$

2. For χ^2 -lift, we have

$$\Lambda_{\chi^2}(y) = \sum_{s \in \mathcal{S}} P_S(s) (l(s, y) - 1)^2 \leq \sum_{s \in \mathcal{S}} P_S(s) (\Lambda(y) - 1)^2 = (\Lambda(y) - 1)^2 \leq (e^{\varepsilon_u} - 1)^2.$$

3. For the α -lift, we have

$$\Lambda_\alpha^S(y) = \left(\sum_{s \in \mathcal{S}} P_S(s) l(s, y)^\alpha \right)^{1/\alpha} \leq \left(\sum_{s \in \mathcal{S}} P_S(s) \Lambda(y)^\alpha \right)^{1/\alpha} = \Lambda(y) \leq e^{\varepsilon_u}.$$

Appendix E

If $s_y = \operatorname{argmax}_{s \in \mathcal{S}} l(s, y)$, then we have $\Lambda(y) = l(s_y, y)$. Recall that $\bar{y} = \operatorname{argmax}_{y \in \mathcal{Y}} [\Lambda(y)]$.

1. When $\max_{y \in \mathcal{Y}} \Lambda_{\ell_1}(y) \leq \varepsilon$, for all $y \in \mathcal{Y}$, we have

$$\Lambda_{\ell_1}(y) = \sum_{s \in \mathcal{S}} P_S(s) |l(s, y) - 1| \leq \varepsilon \Rightarrow P_S(s_y) |l(s_y, y) - 1| = P_S(s_y) (\Lambda(y) - 1) \leq \varepsilon,$$

which results in

$$\max_{y \in \mathcal{Y}} \Lambda(y) \leq \frac{\varepsilon}{P_S(s_{\bar{y}})} + 1.$$

2. When $\max_{y \in \mathcal{Y}} \Lambda_{\chi^2}(y) \leq \varepsilon$, for all $y \in \mathcal{Y}$, we have

$$\Lambda_{\chi^2}(y) = \sum_{s \in \mathcal{S}} P_S(s) (l(s, y) - 1)^2 \leq \varepsilon \Rightarrow P_S(s_y) (l(s_y, y) - 1)^2 = P_S(s_y) (\Lambda(y) - 1)^2 \leq \varepsilon,$$

which results in

$$\max_{y \in \mathcal{Y}} \Lambda(y) \leq \sqrt{\frac{\varepsilon}{P_S(s_{\bar{y}})}} + 1.$$

3. When $\max_{y \in \mathcal{Y}} \Lambda_{\alpha}^S(y) \leq \varepsilon$, for all $y \in \mathcal{Y}$, we have

$$\Lambda_{\alpha}^S(y) = \left(\sum_{s \in \mathcal{S}} P_S(s) l(s, y)^{\alpha} \right)^{1/\alpha} \leq \varepsilon \Rightarrow P_S(s_y) l(s_y, y)^{\alpha} = P_S(s_y) \Lambda(y)^{\alpha} \leq \varepsilon^{\alpha},$$

which results in

$$\max_{y \in \mathcal{Y}} \Lambda(y) \leq \frac{\varepsilon}{P_S(s_{\bar{y}})^{\frac{1}{\alpha}}}.$$

Appendix F

Since $(\varepsilon_l, \varepsilon_u)$ -ALIP is satisfied, for all $y \in \mathcal{Y}$, we have $e^{-\varepsilon_u} \leq \frac{1}{l(s, y)} \leq e^{\varepsilon_l}$ and $\max_s \left(\frac{1}{l(s, y)} \right) = \frac{1}{\Psi(y)}$.

1. For ℓ_1 -lift-inverse, we have

$$\Psi_{\ell_1}(y) = \sum_{s \in \mathcal{S}} P_S(s) \left| \frac{1}{l(s, y)} - 1 \right| \leq \sum_{s \in \mathcal{S}} P_S(s) \left| \frac{1}{\Psi(y)} - 1 \right| = \frac{1}{\Psi(y)} - 1 \leq e^{\varepsilon_l} - 1.$$

2. For χ^2 -lift-inverse, we have

$$\Psi_{\chi^2}(y) = \sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{l(s, y)} - 1 \right)^2 \leq \sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{\Psi(y)} - 1 \right)^2 = \left(\frac{1}{\Psi(y)} - 1 \right)^2 \leq (e^{\varepsilon_l} - 1)^2.$$

3. For α -lift-inverse, we have

$$\Psi_{\alpha}^S(y) = \left(\sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{l(s, y)} \right)^{\alpha} \right)^{\frac{1}{\alpha}} \leq \left(\sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{\Psi(y)} \right)^{\alpha} \right)^{\frac{1}{\alpha}} = \frac{1}{\Psi(y)} \leq e^{\varepsilon_l}.$$

Appendix G

If $s_y = \operatorname{argmin}_s l(s, y)$, then we have $\Psi(y) = l(s_y, y)$. Recall that $\underline{y} = \operatorname{argmin}_y [\Psi(y)]$.

1. When $\max_{y \in \mathcal{Y}} \Psi_{\ell_1}(y) \leq \varepsilon$, for all $y \in \mathcal{Y}$, we have

$$\Psi_{\ell_1}(y) = \sum_{s \in \mathcal{S}} P_S(s) \left| \frac{1}{l(s, y)} - 1 \right| \leq \varepsilon \Rightarrow P_S(s_y) \left| \frac{1}{l(s_y, y)} - 1 \right| = P_S(s_y) \left(\frac{1}{\Psi(y)} - 1 \right) \leq \varepsilon,$$

which results in

$$\min_{y \in \mathcal{Y}} \Psi(y) \geq \frac{P_S(s_y)}{\varepsilon + P_S(s_y)}.$$

2. When $\max_{y \in \mathcal{Y}} \Psi_{\chi^2}(y) \leq \varepsilon$, for all $y \in \mathcal{Y}$, we have

$$\Psi_{\chi^2}(y) = \sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{l(s, y)} - 1 \right)^2 \leq \varepsilon \Rightarrow P_S(s_y) \left(\frac{1}{l(s_y, y)} - 1 \right)^2 = P_S(s_y) \left(\frac{1}{\Psi(y)} - 1 \right)^2 \leq \varepsilon,$$

which results in

$$\min_{y \in \mathcal{Y}} \Psi(y) \geq \frac{\sqrt{P_S(s_y)}}{\sqrt{\varepsilon} + \sqrt{P_S(s_y)}}.$$

3. When $\max_{y \in \mathcal{Y}} \Psi_{\alpha}^S(y) \leq \varepsilon$, for all $y \in \mathcal{Y}$, we have

$$\Psi_{\alpha}^S(y) = \left(\sum_{s \in \mathcal{S}} P_S(s) \left(\frac{1}{l(s, y)} \right)^{\alpha} \right)^{1/\alpha} \leq \varepsilon \Rightarrow P_S(s_y) \left(\frac{1}{l(s_y, y)} \right)^{\alpha} = P_S(s_y) \left(\frac{1}{\Psi(y)} \right)^{\alpha} \leq \varepsilon^{\alpha},$$

which results in

$$\min_{y \in \mathcal{Y}} \Psi(y) \geq \varepsilon^{-1} P_S(s_y)^{\frac{1}{\alpha}}.$$

References

1. Dwork, C.; McSherry, F.; Nissim, K.; Smith, A. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*; Halevi, S., Rabin, T., Eds.; Springer: Berlin/Heidelberg, Germany, 2006; pp. 265–284.
2. Dwork, C. Differential Privacy. In Proceedings of the 33rd International Colloquium on Automata, Languages and Programming, part II (ICALP 2006), Venice, Italy, 10–14 July 2006; Volume 4052, pp. 1–12.
3. Dwork, C. Differential privacy. In *Encyclopedia of Cryptography and Security*; Springer: Boston, MA, USA, 2011; pp. 338–340. [CrossRef]
4. Dwork, C.; Roth, A. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* **2014**, *9*, 211–407. [CrossRef]
5. Kasiviswanathan, S.P.; Lee, H.K.; Nissim, K.; Raskhodnikova, S.; Smith, A. What can we learn privately? *SIAM J. Comput.* **2011**, *40*, 793–826. [CrossRef]
6. Duchi, J.C.; Jordan, M.I.; Wainwright, M.J. Local Privacy and Statistical Minimax Rates. In Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, Berkeley, CA, USA, 26–29 October 2013; pp. 429–438. [CrossRef]
7. Kairouz, P.; Oh, S.; Viswanath, P. Extremal mechanisms for local differential privacy. *Adv. Neural Inf. Process. Syst.* **2014**, *4*, 2879–2887.
8. Sarwate, A.D.; Sankar, L. A rate-distortion perspective on local differential privacy. In Proceedings of the 2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton), Monticello, IL, USA, 30 September–3 October 2014; pp. 903–908. [CrossRef]
9. Kalantari, K.; Sankar, L.; Sarwate, A.D. Robust Privacy-Utility Tradeoffs Under Differential Privacy and Hamming Distortion. *IEEE Trans. Inf. Forensics Secur.* **2018**, *13*, 2816–2830. [CrossRef]
10. Sankar, L.; Rajagopalan, S.R.; Poor, H.V. Utility-privacy tradeoffs in databases: An information-theoretic approach. *IEEE Trans. Inf. Forensics Secur.* **2013**, *8*, 838–852. [CrossRef]
11. Jiang, B.; Li, M.; Tandon, R. Context-aware Data Aggregation with Localized Information Privacy. In Proceedings of the 2018 IEEE Conference on Communications and Network Security (CNS), Beijing, China, 30 May–1 June 2018; pp. 1–9. [CrossRef]
12. Makhdoumi, A.; Salamatian, S.; Fawaz, N.; Médard, M. From the Information Bottleneck to the Privacy Funnel. In Proceedings of the 2014 IEEE Information Theory Workshop (ITW 2014), Hobart, TAS, Australia, 2–5 November 2014; pp. 501–505. [CrossRef]
13. Salamatian, S.; du Pin Calmon, F.; Fawaz, N.; Makhdoumi, A.; Médard, M. Privacy-Utility Tradeoff and Privacy Funnel. 2020. Available online: http://www.mit.edu/~salmansa/files/privacy_TIFS.pdf (accessed on 16 January 2020).
14. Issa, I.; Kamath, S.; Wagner, A.B. An operational measure of information leakage. In Proceedings of the 2016 Annual Conference on Information Science and Systems (CISS), Princeton, NJ, USA, 16–18 March 2016; pp. 234–239.
15. Issa, I.; Kamath, S.; Wagner, A.B. Maximal leakage minimization for the Shannon cipher system. In Proceedings of the 2016 IEEE International Symposium on Information Theory (ISIT), Barcelona, Spain, 10–15 July 2016; pp. 520–524. [CrossRef]
16. Issa, I.; Wagner, A.B.; Kamath, S. An Operational Approach to Information Leakage. *IEEE Trans. Inf. Theory* **2020**, *66*, 1625–1657. [CrossRef]

17. Liao, J.; Kosut, O.; Sankar, L.; Calmon, F.P. A tunable measure for information leakage. In Proceedings of the IEEE International Symposium on Information Theory (ISIT), Vail, CO, USA, 17–22 June 2018; pp. 701–705.
18. du Pin Calmon, F.; Fawaz, N. Privacy against statistical inference. In Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton), Monticello, IL, USA, 1–5 October 2012; pp. 1401–1408.
19. Jiang, B.; Li, M.; Tandon, R. Local Information Privacy with Bounded Prior. In Proceedings of the 2019 IEEE International Conference on Communications (ICC), Shanghai, China, 20–24 May 2019; pp. 1–7. [\[CrossRef\]](#)
20. Seif, M.; Tandon, R.; Li, M. Context Aware Laplacian Mechanism for Local Information Privacy. In Proceedings of the 2019 IEEE Information Theory Workshop (ITW), Visby, Sweden, 25–28 August 2019; pp. 1–5. [\[CrossRef\]](#)
21. Jiang, B.; Li, M.; Tandon, R. Local Information Privacy and Its Application to Privacy-Preserving Data Aggregation. *IEEE Trans. Dependable Secur. Comput.* **2020**, *19*, 1918–1935. [\[CrossRef\]](#)
22. Jiang, B.; Seif, M.; Tandon, R.; Li, M. Context-Aware Local Information Privacy. *IEEE Trans. Inf. Forensics Secur.* **2021**, *16*, 3694–3708. [\[CrossRef\]](#)
23. Ding, N.; Liu, Y.; Farokhi, F. A Linear Reduction Method for Local Differential Privacy and Log-lift. In Proceedings of the 2021 IEEE International Symposium on Information Theory (ISIT), Melbourne, VIC, Australia, 12–20 July 2021; pp. 551–556. [\[CrossRef\]](#)
24. Hsu, H.; Asodeh, S.; Calmon, F.P. Information-Theoretic Privacy Watchdogs. In Proceedings of the 2019 IEEE International Symposium on Information Theory (ISIT), Paris, France, 7–12 July 2019; pp. 552–556.
25. Sadeghi, P.; Ding, N.; Rakotoarivelo, T. On Properties and Optimization of Information-theoretic Privacy Watchdog. In Proceedings of the 2020 IEEE Information Theory Workshop (ITW), Riva del Garda, Italy, 11–15 April 2020.
26. Zarrabian, M.A.; Ding, N.; Sadeghi, P.; Rakotoarivelo, T. Enhancing utility in the watchdog privacy mechanism. In Proceedings of the ICASSP 2022—2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Virtual, 7–13 May 2022; pp. 2979–2983.
27. Razeghi, B.; Calmon, F.; Gunduz, D.; Voloshynovskiy, S. On Perfect Obfuscation: Local Information Geometry Analysis. *arXiv* **2020**, arXiv:2009.04157.
28. Lopuhaä-Zwakenberg, M.; Tong, H.; Škorić, B. Data Sanitisation Protocols for the Privacy Funnel with Differential Privacy Guarantees. *Int. J. Adv. Secur.* **2021**, *13*, 162–174.
29. Evfimievski, A.; Gehrke, J.; Srikant, R. Limiting Privacy Breaches in Privacy Preserving Data Mining. In Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '03, San Diego, CA, USA, 9–12 June 2003; pp. 211–222. [\[CrossRef\]](#)
30. Saeidian, S.; Cervia, G.; Oechtering, T.J.; Skoglund, M. Pointwise Maximal Leakage. *arXiv* **2022**, arXiv:2205.04935.
31. Fernandes, N.; McIver, A.; Sadeghi, P. Explaining epsilon in differential privacy through the lens of information theory. *arXiv* **2022**, arXiv:2210.12916.
32. Zamani, A.; Oechtering, T.J.; Skoglund, M. Data Disclosure With Non-Zero Leakage and Non-Invertible Leakage Matrix. *IEEE Trans. Inf. Forensics Secur.* **2022**, *17*, 165–179. [\[CrossRef\]](#)
33. Zamani, A.; Oechtering, T.J.; Skoglund, M. A Design Framework for Strongly χ^2 -Private Data Disclosure. *IEEE Trans. Inf. Forensics Secur.* **2021**, *16*, 2312–2325. [\[CrossRef\]](#)
34. Ding, N.; Zarrabian, M.A.; Sadeghi, P. α -Information-theoretic Privacy Watchdog and Optimal Privatization Scheme. In Proceedings of the 2021 IEEE International Symposium on Information Theory (ISIT), Melbourne, VIC, Australia, 12–20 July 2021; pp. 2584–2589.
35. Rassouli, B.; Gunduz, D. Optimal utility-privacy trade-off with total variation distance as a privacy measure. *IEEE Trans. Inf. Forensics Secur.* **2019**, *15*, 594–603. [\[CrossRef\]](#)
36. Asodeh, S.; Diaz, M.; Alajaji, F.; Linder, T. Estimation efficiency under privacy constraints. *IEEE Trans. Inf. Theory* **2018**, *65*, 1512–1534. [\[CrossRef\]](#)
37. Rassouli, B.; Rosas, F.E.; Gunduz, D. Data Disclosure under Perfect Sample Privacy. *arXiv* **2019**, arXiv:1904.01711.
38. Becker, B.; Kohavi, R. Adult. UC Irvine Machine Learning Repository. Available online: <https://archive-beta.ics.uci.edu/dataset/2/adult> (accessed on 5 January 1996).
39. Liu, Y.; Sadeghi, P.; Arbabjolfaei, F.; Kim, Y.H. Capacity Theorems for Distributed Index Coding. *IEEE Trans. Inf. Theory* **2020**, *66*, 4653–4680. [\[CrossRef\]](#)
40. Wang, H.; Vo, L.; Calmon, F.P.; Médard, M.; Duffy, K.R.; Varia, M. Privacy With Estimation Guarantees. *IEEE Trans. Inf. Theory* **2019**, *65*, 8025–8042. [\[CrossRef\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.