

Article

Statistic Maximal Leakage [†]

Shuaiqi Wang ^{1,*}, Zinan Lin ²  and Giulia Fanti ¹ 
¹ Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA

² Microsoft Research, Redmond, WA 98052, USA; zinanlin@microsoft.com

* Correspondence: shuaiqi@andrew.cmu.edu

[†] This article is a revised and expanded version of a paper entitled Statistic maximal leakage, which was presented at ISIT 2024, Athens, Greece, 7–12 July 2024.

Abstract

We introduce a privacy measure called *statistic maximal leakage* that quantifies how much a privacy mechanism leaks about a specific secret random variable, relative to the adversary's prior information about that secret, in the worst case over all possible priors. Statistic maximal leakage is an extension of the well-known *maximal leakage* framework. Unlike maximal leakage, which protects an arbitrary, unknown secret random variable, statistic maximal leakage is designed to protect a known function of a public random variable. We show that statistic maximal leakage satisfies composition and post-processing properties. Additionally, we show how to efficiently compute it in the special case of deterministic data release mechanisms. We analyze two important mechanisms under statistic maximal leakage: the quantization mechanism and randomized response. We show theoretically and empirically that the quantization mechanism achieves better privacy–utility tradeoffs in the settings we study. This framework may benefit data holders and privacy practitioners who release data containing known secrets by enabling them to assess leakage without specifying an exact prior and to better preserve data utility while protecting those secrets.

Keywords: privacy; maximal leakage; data release; information leakage; privacy–utility tradeoffs

1. Introduction

Information disclosure while preventing privacy leakage is a central problem in the privacy and information theory literature. That is, how can we release a (realization of a) random variable without leaking correlated secret information? Since the seminal work of Yamamoto [1], many papers have studied this problem [2–15]. This article is a revised and expanded version of the paper titled Statistic Maximal Leakage, presented at ISIT 2024 [16]. We model a data holder with data drawn from a distribution f_{Θ} parameterized by a random variable Θ . The data holder aims to protect a secret, represented by a discrete random variable G , which is defined as a function of Θ ; that is, $G = g(\Theta)$. The data holder's goal is to release Θ' , a perturbed version of Θ , while limiting *leakage* about G and maximizing the *utility* of Θ' , in the sense of minimizing its total variation (TV) distance from Θ .

We seek a measure of privacy leakage that satisfies three properties:

1. *Prior independence:* The measure should not depend on any party's prior over the input data, up to determining a superset of the distribution's support; that is, it may depend on which input values are considered possible, but not on the probability mass assigned to those values. This arises because priors may be difficult to obtain



Academic Editors: Amirreza Zamani and Mina Sheikhalishahi

Received: 1 June 2026

Revised: 29 June 2026

Accepted: 7 July 2026

Published: 18 July 2026

Copyright: © 2026 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 \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

in practice; measures that require such knowledge may admit mechanisms that are fragile to prior misspecification [17–20].

2. *Secret-awareness*: We assume the secret function g is known, which describes how to obtain the secret G from the input data distribution. We want a measure that depends explicitly on g . This is for utility reasons; by exploiting known information, we hope to add less noise or perturbation to our data than a measure that is secret-agnostic.
3. *Composition and post-processing*: Informally, composition provides a bound on how a privacy measure degrades when one or more mechanisms are applied sequentially to the same data. For example, in differential privacy (DP), the privacy parameter ϵ degrades additively when a mechanism is applied multiple times to the same dataset [21]. Post-processing guarantees that if one applies an arbitrary (possibly random) function to the output of a privacy mechanism, as long as the function does not depend on the original data, the privacy measure in question does not degrade. Composition and post-processing are useful properties exhibited by DP, which have contributed to its widespread usage in practical settings such as machine learning pipelines [22].

To the best of our knowledge, existing privacy leakage measures that satisfy all three of these properties are all inspired by (generalizations of) differential privacy. Examples include attribute privacy [23], distribution privacy [24], and distribution inference [25]. Due to their assumptions and formulation, they require large amounts of noise in practice [17,26]. We discuss this tradeoff further in Sections 2, 3 and 6.

In this work, we study an information-theoretic privacy measure that satisfies the above three properties. For a special class of parametric data distributions whose (random) parameter vector is drawn from a finite set, we propose a privacy measure inspired by maximal leakage [7,8], which we call *statistic maximal leakage*. Our contributions in this work are as follows:

- *Formulation and properties*. We formulate the statistic maximal leakage privacy measure, and show that it trivially satisfies the first two properties. We prove that it also satisfies composition and post-processing. Note that although composition and post-processing have been previously proved for an extension of maximal leakage called *pointwise maximal leakage* [27], their result and proofs do not imply that statistic maximal leakage also satisfies these properties (details in Section 4.3).
- *Computation*. We next study how to compute statistic maximal leakage. We show that in general, computing statistic maximal leakage is NP-hard. However, for the class of deterministic data release mechanisms, we show that statistic maximal leakage can be computed in polynomial time by solving a minimum-cost flow problem over a graph whose construction depends on the mechanism parameters.
- *Mechanisms*. We next analyze two natural mechanisms that have been studied widely in the privacy literature: the quantization mechanism [17,28] and randomized response [29–31]. Quantization-based mechanisms have been shown to achieve (near)-optimal privacy–utility tradeoffs in several privacy frameworks, including summary statistic privacy [17] and non-randomized privacy [28]. Randomized response is a widely adopted mechanism [29,30] that achieves optimal privacy–utility tradeoffs for differentially private data collection [32]. We show that both mechanisms satisfy non-trivial statistic maximal leakage guarantees. Further, for a special case of tabular datasets with categorical entries, we analyze the privacy–utility tradeoffs of these two mechanisms; for this analysis, we define utility as a worst-case total variation distance between the original and the released data distributions. Our results show that for a tabular data release problem that aims to hide one of the marginal values, the

quantization mechanism achieves a better privacy–utility tradeoff than randomized response for practical privacy parameters.

- *Empirical evaluation.* Finally, we apply the quantization mechanism to a real-world tabular dataset. We show that when instantiated with an appropriate statistic maximal leakage parameter, the quantization mechanism effectively protects the secret while still ensuring the high utility of the released data. Moreover, we illustrate empirically that the quantization mechanism we propose for satisfying statistic maximal leakage incurs lower utility cost than mechanisms designed for other privacy frameworks: attribute privacy [23] and maximal leakage [8].

2. Related Work

There have been many privacy measures proposed for capturing leakage of sensitive information during information disclosure. We divide them into *information-theoretic measures* and *indistinguishability-based measures*.

2.1. Information-Theoretic Measures

Maximal leakage-based measures. While there are various metrics designed to measure the privacy leakage, *maximal leakage* is the one most related to our statistic maximal leakage. Maximal leakage, introduced by Issa et al. [8], is an information-theoretic measure designed to capture the worst-case leakage of sensitive information in a data release mechanism. It is defined using the Markov chain $G - \Theta - \Theta' - \hat{G}$, where $\hat{G}$ represents the adversary's guess of secret G . The measure is expressed as

$$\mathcal{L}_{ML} = \sup_{G - \Theta - \Theta' - \hat{G}} \log \frac{\mathbb{P}(G = \hat{G})}{\max_g \mathbb{P}_G(g)}, \quad (1)$$

where the sup is over G and $\hat{G}$, i.e., consider the worst-case secret and the strongest attack strategy. This formulation evaluates the ratio of probabilities of correctly guessing the secret with and without observing the released data, thereby quantifying the increase in guessing probability facilitated by the data release.

Several variants and extensions of maximal leakage have been proposed to address different privacy scenarios via gain functions [27,33–35], which penalize different values of $\mathbb{P}(G = \hat{G})$ differently. This framework can adapt to different threat models by modifying the evaluation of adversarial success, thus offering a tunable balance between privacy and utility. For example, maximal (α, β) -leakage [34] generalizes maximal leakage by introducing parameters α and β , which adjust the sensitivity of leakage to different probabilities of adversarial success. By tuning α and β , the metric can be converted to several known leakage measures, such as vanilla maximal leakage and local differential privacy. Pointwise maximal leakage, introduced by Saeidian et al. [27], examines the leakage given certain released data rather than averaging over the entire distribution, preventing disproportionate exposure of sensitive information. Binary maximal leakage [36] focuses on the scenario where only binary secret functions are assumed to be of interest to the adversary, and demonstrates that restricting the set of secrets to be hidden allows for a better utility.

Maximal leakage and its variants [8,27,33–36] assume the exact secret is unknown a priori, and hence are worst-case metrics over all secrets (or secrets of certain types). Their guarantees typically depend on the input prior through its support; the set of input values with nonzero probability mass must be specified. Furthermore, although maximal leakage and some of its variant [34,36] analyze the composition property, it holds only when successive outputs from the mechanism(s) are conditionally independent, conditioned on the input data. The only exception is pointwise maximal leakage [27], which has the property of adaptive composition in additive form.

Besides maximal leakage, several metrics are proposed to quantify the information leakage from Θ to Θ' based on mutual information [5,6,37–43], f -divergences [4,44], or min-entropy [2,3,45–51], which are detailed as follows.

Mutual information-based measures. Privacy Funnel [1,5] measures the privacy leakage by the mutual information between the secret G and released data. It minimizes the leakage of secret while ensuring the released data retains a certain level of utility from the original data, which is captured by the mutual information between Θ and Θ' . Similarly, rate-distortion theory, originally from source coding [52], models privacy problems by minimizing mutual information between Θ and Θ' subject to a distortion constraint, trading off privacy for utility [39,40]. However, in addition to requiring knowledge of a prior distribution of input data, as discussed in [8,17], mutual information does not align with practical notions of privacy and utility, as it does not account for the probability of correct guesses by an adversary.

Min-entropy-based measures. Several works measure information leakage with min-entropy [3,45–47], the probability of guessing a secret correctly, which provides a more direct measure of privacy risk. For example, summary statistic privacy [17,26] quantifies the privacy loss by assessing the worst-case probability the adversary successfully guesses the secret within a tolerance range. However, those measures require the knowledge of the prior distribution of Θ . Although [7] extends the measure to be prior-independent by considering the probability of guessing the secret under a worst-case prior over Θ , it considers on the privacy loss of the entire random variable Θ , rather than a specific secret; moreover, no composition and post-processing properties are provided.

Quantitative Information Flow. Quantitative Information Flow (QIF) measures how much information about a secret is leaked through observable outputs. Originating in the works of Denning [53] and Gray [54], QIF has evolved to encompass a range of privacy metrics. Early QIF frameworks used mutual information to quantify leakage [41–43], while Smith [2] adopted min-entropy as the metric after recognizing the shortcomings of mutual information as a privacy metric. Further developments introduced generalized metrics like g -leakage, which incorporate gain functions to model scenarios involving partial or multiple guesses [3]. Worst-case min-entropy leakage [7] can be viewed as a special case of g -leakage by using the 0–1 gain function for exactly guessing the input. In Property 1, we make its connection to our statistic maximal leakage precise by showing that statistic maximal leakage is upper- and lower-bounded by worst-case min-entropy leakage by up to an additive factor. QIF framework has been applied to various domains, such as local DP and cyberattack defenses [49–51]. However, these studies treat the whole dataset as the secret rather than hiding correlated random variables, and few analyze post-processing and composition properties.

Non-stochastic information theoretic methods. Non-stochastic information theoretic methods have also been adopted to assess privacy leakage [28,55–57], and most works in this area focus on the exploration of noiseless privacy-preserving mechanisms, e.g., quantization-based mechanisms. For example, Farokhi et al. [28,56] evaluate utility as the worst-case deviation between input and output data and measure the privacy by maximin information [58], which is defined through a unique taxicab partitioning of feasible input–output pairs. Ref. [28] identifies quantization as an optimal privacy mechanism within deterministic piecewise differentiable policies, i.e., policies quantizing the input dataset into bins and the output is differentiable for each bin, that maximize privacy under a utility constraint. However, similar to Quantitative Information Flow, these works consider the whole data as the secret and do not prove post-processing or composition properties.

2.2. Indistinguishability-Based Measures

Orthogonal to information-theoretic approaches, several studies measure the privacy leakage of statistical properties of the data based on a measure of indistinguishability over candidate inputs [23–25,59]; these techniques draw inspiration from *differential privacy* (DP) [21], one of the most widely adopted privacy metrics. Roughly, differential privacy quantifies the privacy of a data release mechanism by measuring the influence of the participation of an individual record (or a group of records) on the final output [21]. However, as discussed in [17,26], DP is designed to measure individual-level privacy and does not directly quantify information leakage of statistical (distribution-level) secrets. As an example, if we aim to protect the mean of a dataset of scalar values and apply a typical local DP mechanism [60] that adds zero-mean Gaussian noise to each record, this prevents the attacker from inferring the inclusion of individual records in the dataset but still allows an unbiased estimate of the dataset’s mean. As the number of records increases, the mean of released dataset converges to the original mean value [17]. Simple variants of this mechanism that add different noise scales to different features are also shown to fail in [17].

Distribution privacy [24] and distribution inference [25,59] tackle this challenge by ensuring that for any two input distributions with different parameters, the output distributions remain indistinguishable (up to multiplicative and additive factors), thus protecting sensitive distributional properties. For instance, attribute privacy [23] adopts an indistinguishability definition based on the pufferfish framework, while restricting the class of distributions that should be indistinguishable; it is particularly useful for scenarios where only a statistic of the dataset is shared. Both attribute and distribution privacy mechanisms can be very noisy, as they must obfuscate between entire distributions. This is challenging because two random variables may have similar secret values (e.g., means), while having arbitrarily different underlying distributions; hence, the amount of noise required to make these distributions indistinguishable can be substantial. We show an example of this privacy–utility tradeoff in Section 6.

3. Notation and Problem Setting

We use uppercase Greek and Roman letters (e.g., Θ, G) to denote random variables, and lowercase Greek letters (e.g., θ) to denote their realization, unless specified otherwise. Sets are denoted by **boldface** uppercase Greek and Roman letters.

A *data holder* has data drawn from a parametric distribution f_θ , for some $\theta \in \Theta$, where Θ is a finite set (we extend the formulation to infinite, continuous parameter sets in Appendix L). The parameter θ is itself a realization of a random variable $\Theta \sim \mathbb{P}_\Theta$. $\mathbb{P}_\Theta$ represents the prior distribution of the parameter random variable Θ , and can equivalently be viewed as the prior over the input data.

The data holder aims to release information about θ while limiting disclosure about a secret $g = g(\theta)$ to an *adversary*, where $g : \Theta \rightarrow \mathbf{G}$ is a function that is fixed and known, and $\mathbf{G} \triangleq \{g_1, g_2, \dots, g_s\}$ is the set of s possible secret values. Note that $g \in \mathbf{G}$ is a realization of random variable G , and we will also write $G = g(\Theta)$ to denote the random variable induced by g . In the following sections, we define precise notions of leakage about G and utility with respect to Θ . We use Θ_g to represent the set of original parameters whose secret values are g , i.e., $\Theta_g = \{\theta \in \Theta | g(\theta) = g\}$. In Section 5, we use Θ_i to represent $\Theta_{g_i}, \forall i \in [s]$, for convenience.

The data holder releases data via a randomized *data release mechanism* $\mathcal{M} = \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta)$, which is a channel mapping a realization θ of the random variable input parameter Θ to a conditional output distribution over the random output parameter Θ' taking values in Θ' (in general, $\Theta \neq \Theta'$). With some abuse of notation, we also write $\Theta' = \mathcal{M}(\theta)$.

Given a realization θ' of Θ' , the attacker outputs a (possibly random) estimate of the secret, $\hat{G}$. We assume the attacker knows the prior distribution of the data and the data release mechanism $\mathcal{M}$, and is computationally unbounded. The overall data sharing and attacker guessing process can be formulated as a Markov chain $G - \Theta - \Theta' - \hat{G}$.

Utility

To analyze the privacy–utility tradeoff for a mechanism $\mathcal{M}$, we define our utility measure as a distortion: the expected total variation (TV) distance between the original and released data distributions, represented by f_Θ and $f_{\Theta'}$ respectively, under the worst-case prior:

$$\Delta_{\mathcal{M}} = \sup_{\mathbb{P}_\Theta \in \mathcal{P}(\Theta)} \mathbb{E}_{\Theta, \Theta' = \mathcal{M}(\Theta)} [D_{\text{TV}}(f_\Theta \| f_{\Theta'})], \quad (2)$$

where $\mathcal{P}(\Theta)$ denotes the probability simplex over Θ and D_{TV} is the total variation distance.

Note that Equation (2) is computed for a worst-case prior, but it is average-case over the realizations of the data and the output of the data release mechanism. We have chosen to model utility in this way because we want the utility measure to apply to a data release mechanism, regardless of input. Hence, utility should not depend on $\mathbb{P}_\Theta$. Two natural options are to either consider the worst-case distribution $\mathbb{P}_\Theta$, or the average case. As we do not model a prior over $\mathbb{P}_\Theta$, we use the former. We subsequently take an expectation over Θ and Θ' , as is common in many works studying the utility of privacy mechanisms, both in information theory [40,46] and in differential privacy [61,62].

We next provide Lemma 1, which allows us to analyze the distortion of $\mathcal{M}$ as the expected TV distance between the original and released datasets under a worst-case input.

Lemma 1. *The distortion measure $\Delta_{\mathcal{M}}$ can be rewritten as*

$$\Delta_{\mathcal{M}} = \sup_{\theta \in \Theta} \mathbb{E}_{\Theta' = \mathcal{M}(\theta)} [D_{\text{TV}}(f_\theta \| f_{\Theta'})]. \quad (3)$$

(Proof in Appendix A.1.) Since our utility measure considers a worst-case prior distribution, the distortion of mechanisms proposed for attribute privacy [23], distribution privacy [24], or distribution inference [25] can reach the trivial upper bound 1 on distortion (we illustrate the case for attribute privacy in Section 6), i.e., the worst possible utility. Our goal is to understand whether there exist data release mechanisms that achieve meaningful utility while also satisfying a privacy guarantee with the desired properties from Section 1.

4. Statistic Maximal Leakage

We next present our proposed privacy leakage measure, discuss its properties, and explain how to compute it. Statistic maximal leakage (SML) measures the largest increase an adversary can gain in their guess of G ; it is a property of a data release mechanism $\mathcal{M}$ and a secret mapping g . We define it as follows:

$$\Pi_{\mathcal{M}, g} = \sup_{\mathbb{P}_\Theta, \mathbb{P}_{\hat{G}|\Theta'}} \log \frac{\mathbb{P}(\hat{G} = G)}{\sup_{g \in G} \mathbb{P}_G(g)}, \quad (4)$$

where the supremum is taken over all prior distributions $\mathbb{P}_\Theta$ over the distribution parameter $\theta \in \Theta$ and attack strategies $\mathbb{P}_{\hat{G}|\Theta'}$. The probability in the numerator is over the attacker's randomized estimator, the mechanism, and the secret. Note that the secret function g and the data release mechanism $\mathcal{M}$ are fixed and assumed to be known in this optimization.

4.1. Comparison to Other Leakage Metrics

Maximal Leakage. SML bears some similarities with maximal leakage [8] and worst-case min-entropy leakage [7], though these prior metrics do not simultaneously satisfy all three desirable properties we propose. Adopting our notation, maximal leakage $\mathcal{L}_{ML}$ is defined as follows:

$$\mathcal{L}_{ML} = \sup_{\mathbb{P}_{\Theta}: \text{Supp}(\mathbb{P}_{\Theta})=S, \mathbb{P}_{G|\Theta}, \mathbb{P}_{\hat{G}|\Theta'}} \log \frac{\mathbb{P}(\hat{G} = G)}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)}, \quad (5)$$

where $S \subseteq \Theta$ is a fixed and known support set; that is, maximal leakage optimizes over prior distributions that assign nonzero mass to every $S \in S$. We write it in this form because [8] shows that maximal leakage depends on the support of $\mathbb{P}_{\Theta}$. A stronger condition, which is presented in [8], requires knowledge of the prior $\mathbb{P}_{\Theta}$ over Θ . Maximal leakage assumes the secret function $g(\cdot)$ is unknown and, therefore, it optimizes over all $\mathbb{P}_{G|\Theta}$.

Min-entropy Leakage. Worst-case min-entropy leakage $\mathcal{L}_{MEL}$ can be viewed as a special case of g -leakage [3] by using the 0–1 gain function for exactly guessing the input; it treats the entire input data distribution as the secret to protect. For the Markov chain $\Theta - \Theta' - \hat{\Theta}$, $\mathcal{L}_{MEL}$ is defined as

$$\mathcal{L}_{MEL} = \sup_{\mathbb{P}_{\Theta}, \mathbb{P}_{\hat{\Theta}|\Theta'}} \log \frac{\mathbb{P}(\hat{\Theta} = \Theta)}{\sup_{\theta \in \Theta} \mathbb{P}_{\Theta}(\theta)}. \quad (6)$$

Under a fixed prior, maximal leakage can have a smaller value than SML since SML considers the worst-case prior. However, one can construct $\mathcal{M}$ and g for which maximal leakage, with a worst-case prior, achieves its largest possible value (i.e., $\min\{\log |\Theta'|, \log |\Theta|\}$) while SML is 0. The following property shows that SML is upper- and lower-bounded by both maximal leakage with a worst-case prior and worst-case min-entropy leakage, by up to an additive factor that depends on the secret and the parametric family (Proof in Appendix A.2.).

Property 1 (Relation to Maximal Leakage and Min-Entropy Leakage). *Statistic maximal leakage $\Pi_{\mathcal{M},g}$ satisfies*

$$\mathcal{L}_{MEL} - \sup_{g \in \mathbf{G}} \log |\Theta_g| \leq \Pi_{\mathcal{M},g} \leq \mathcal{L}_{MEL}, \quad (7)$$

$$\sup_{\mathbb{P}_{\Theta}} \mathcal{L}_{ML} - \sup_{g \in \mathbf{G}} \log |\Theta_g| \leq \Pi_{\mathcal{M},g} \leq \sup_{\mathbb{P}_{\Theta}} \mathcal{L}_{ML}. \quad (8)$$

4.2. Computation of Statistic Maximal Leakage

We next discuss how to compute statistic maximal leakage. Computing SML is more convenient under an alternative form, which shows that we only need to search over a restricted class of priors that assign nonzero probability mass to at most one parameter $\theta \in \Theta_g$, for each secret $g \in \mathbf{G}$. Under such a prior, for a fixed $g \in \mathbf{G}$, there is at most one $\theta \in \Theta$ such that $\mathbb{P}_{\Theta|G}(\theta|g) > 0$; we use θ_g to denote this value.

Proposition 1. *Statistic maximal leakage satisfies*

$$\Pi_{\mathcal{M},g} = \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g). \quad (9)$$

(Proof in Appendix A.3.) Based on Proposition 1, we explain how to compute SML with an accompanying example in Figure 1 (left) that has $\Theta = \{\theta_1, \theta_2, \theta_3\}$, $\Theta' = \{\theta'_1, \theta'_2\}$, and $\Theta_{g_1} = \{\theta_1, \theta_2\}$, $\Theta_{g_2} = \{\theta_3\}$. Given a mechanism $\mathcal{M} = \mathbb{P}_{\Theta'|\Theta}$, we can construct a policy matrix

where the value in the i -th row and j -th column is $\mathbb{P}_{\Theta'|\Theta}(\theta'_j|\theta_i)$. First, fix a prior $\mathbb{P}_{\Theta}$ such that $\mathbb{P}_{\Theta|G} \in \{0, 1\}$. Figure 1 illustrates a case where the prior satisfies $\theta_{g_1} = \theta_1$ and $\theta_{g_2} = \theta_3$ (θ_g is defined above Proposition 1). Next, fix a column θ'_j in the policy matrix. We can now find a secret value $\tilde{g} \in \mathbf{G}$ that maximizes $\mathbb{P}_{\Theta'|\Theta}(\theta'_j|\theta_{\tilde{g}})$ —i.e., $\tilde{g}$ is the maximum likelihood secret for an observed output θ'_j . For each column, the red outline denotes the input parameters in $\Theta_{\tilde{g}}$. Our example mechanism satisfies $\arg \sup_{\tilde{g}} \mathbb{P}_{\Theta'|\Theta}(\theta'_1|\theta_{\tilde{g}}) = g_1$, $\arg \sup_{\tilde{g}} \mathbb{P}_{\Theta'|\Theta}(\theta'_2|\theta_{\tilde{g}}) = g_2$. For each column, the blue square is the intersection of the red region with the row of $\theta_{\tilde{g}}$. We finally sum the likelihoods of all the blue squares. Our goal is to find the worst-case prior and calculate the maximum value of $\log \sum_{\theta' \in \Theta'} \sup_{\tilde{g} \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{\tilde{g}})$. In the worst case, this can be done in time exponential in the number of input parameters $|\Theta|$ by enumerating all feasible $P_{\Theta|G}$.




Figure 1. Given a mechanism $\mathcal{M} = \mathbb{P}_{\Theta'|\Theta}$, the left subfigure shows a policy matrix. For each column j , the red outlined region indicates rows of parameters with secret g maximizing $\mathbb{P}_{\Theta'|\Theta}(\theta'_j|\theta_g)$. The blue cell lies in the row of $\theta_{\tilde{g}}$. When the mechanism $\mathcal{M}$ is deterministic, SML calculation can be converted to a min-cost flow problem (right). The constructed directed graph contains three columns of nodes (representing $\mathbf{G}, \Theta, \Theta'$, respectively) between the source and sink nodes. The capacity of all edges is 1, and only the edges between nodes in Θ and Θ' columns have non-zero cost ($-\mathbb{P}_{\Theta'|\Theta}(\theta'_k|\theta_j)$ between θ_j and θ'_k). Edges are annotated as: Capacity (Cost).

4.2.1. Computation for Deterministic Mechanisms: Min-Cost Flow

When the mechanism $\mathcal{M}$ is *deterministic*, i.e., $\mathbb{P}_{\Theta'|\Theta} \in \{0, 1\}$, SML calculation process can be converted to a min-cost flow problem [63]. Given a directed graph where each edge is assigned a capacity and a cost, the min-cost flow problem aims to design a network flow satisfying the capacity constraint of each edge, while achieving the minimum cost. The final cost of the min-cost flow has a one-to-one correspondence to the SML of the underlying problem.

To construct the network, we start with a source and a sink node, and create three columns of nodes. The first $\mathbf{G}$ column contains all potential secret values (g_1, g_2 in Figure 1). The capacity of the edge between the source and each node in the $\mathbf{G}$ column is 1 and the cost is 0. The second Θ column contains all possible input parameter values ($\theta_1, \theta_2, \theta_3$ in Figure 1). There is an edge between node g_i and θ_j iff $\theta_j \in \Theta_{g_i}$. The capacity of each edge is 1 and the cost is 0. The third Θ' column contains all possible released parameter values (θ'_1, θ'_2 in Figure 1). This column is fully connected with the second column. The capacity of the edge between θ_j and θ'_k is 1 and the cost is $-\mathbb{P}_{\Theta'|\Theta}(\theta'_k|\theta_j)$. All nodes in the Θ' column are also connected with the sink node with edge capacity 1 and cost 0. Concretely, the network construction steps are detailed in Algorithm 1.

From Proposition 1, we know that among all the distributions we are optimizing over, there is only one θ_g satisfying $\mathbb{P}_{\Theta|G}(\theta_g|g) = 1 > 0, \forall g \in \mathbf{G}$. For any *deterministic* mechanism, there is only one $\theta' \in \Theta'$ satisfying $\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) = 1 > 0, \forall \theta \in \Theta$. Finally, for each $\theta' \in \Theta'$, we can only select one $\tilde{g} \in \mathbf{G}$ to calculate $\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{\tilde{g}})$ for the SML calculation, based on Proposition 1. Hence, we set the capacity of all edges as 1. It is known that there

exists a min-cost network flow such that the flow of each edge is either 1 or 0 [63]. In that case, for all nodes $\theta \in \Theta_g$ in the Θ column, only one can accept one unit of flow from g , and this node is θ_g . For each node in the Θ' column, it can also only accept one unit of flow from $\theta_g, \forall g \in \mathbf{G}$. Therefore, the min-cost flow problem under our constructed network in Algorithm 1 shares the same objective as Proposition 1. Importantly, the min-cost flow problem can be solved efficiently in polynomial time in $|\Theta| \cdot |\Theta'|$ [63].

Algorithm 1: Network construction under deterministic mechanism.

Input: parameter value set Θ , secret function g , secret value set $\mathbf{G}$, released parameter value set Θ' , data release mechanism $\mathcal{M} = \mathbb{P}_{\Theta'|\Theta}$

- 1 Construct a source node N_{src} and a sink node N_{sink} ;
 - 2 **for** each $g \in \mathbf{G}$: construct a node N_g connected to N_{src} with edge capacity 1 and cost 0;
 - 3 **for** each $g \in \mathbf{G}$ and $\theta \in \Theta_g$: construct a node N_θ connected to N_g with edge capacity 1 and cost 0;
 - 4 **for** each $\theta' \in \Theta'$ and $\theta \in \Theta$: construct a node $N_{\theta'}$ connected to N_θ with edge capacity 1 and cost $-\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta)$, and connected to N_{sink} with edge capacity 1 and cost 0.
-

The following proposition shows that under the network constructed through Algorithm 1, the final cost of the min-cost flow has a one-to-one correspondence to the SML of the underlying problem.

Proposition 2 (SML computation, deterministic mechanism). *Given a deterministic mechanism $\mathcal{M}$ and a secret mapping g , $\Pi_{\mathcal{M},g} = \log(-C_{\mathcal{M},g})$, where $C_{\mathcal{M},g}$ is the total cost of the min-cost flow under the network constructed through Algorithm 1.*

(Proof in Appendix B.) Intuitively, we can refer to our example, where we allocate 1 unit of flow from the source to each of g_1 and g_2 . For g_1 , the full flow goes either to θ_1 or θ_2 ; the selected node is dubbed θ_{g_1} under $\mathbb{P}_{\Theta'|\Theta}$. The flows from θ_{g_1} and θ_3 then go to θ'_1 and θ'_2 , respectively. This is because θ'_1 has g_1 as its ML secret, and θ'_2 has g_2 as its ML secret under $\mathbb{P}_{\Theta'|\Theta}$. Finally, the flows merge to the sink. The log of the negative total cost of the flow is the SML.

4.2.2. Computational Hardness for General Mechanisms

Although the computation of SML is efficient when the mechanism is deterministic, the computation is NP-hard in general, although there exist efficient approximation algorithms, as shown in Theorem 1.

Theorem 1 (Hardness of SML computation). *Given a mechanism $\mathcal{M}$ and a secret mapping g , the computation of SML $\Pi_{\mathcal{M},g}$ is NP-hard. There exists an approximation algorithm for SML calculation on $|\Theta'| - \exp(\Pi_{\mathcal{M},g})$ with approximation ratio $1 + \rho$ ($\rho > 0$) and running time polynomial in $|\Theta| \cdot |\Theta'|$ and $1/\rho$.*

(Proof in Appendix C.) We prove the NP-hardness of SML computation by reducing it to the three-set cover decision problem, for which NP completeness is known [64]. To approximately compute SML, we convert the computation to an edge cost flow problem, for which there exist practical (poly-time) approximation algorithms [65]. However, we do not use these approximation methods in the remainder of the paper; although our proposed mechanisms are not deterministic, they possess a special structure that enables us to determine their privacy guarantees in closed form.

4.3. Properties of Statistic Maximal Leakage

We next show that SML satisfies two natural desired properties: adaptive composition and post-processing. Adaptive composition bounds the total leakage of releasing multiple results from one or more possibly adaptive mechanisms applied sequentially.

Theorem 2 (Adaptive Composition). *Suppose a data holder sequentially applies m mechanisms $\mathcal{M}_1, \dots, \mathcal{M}_m$, where $\forall i \in [m]$, the i -th mechanism is a function of the input data distribution θ and all of the previous outputs, which we denote as $\theta^{(1)}, \dots, \theta^{(i-1)}$. That is, $\mathcal{M}_i(\theta, \theta^{(1)}, \dots, \theta^{(i-1)}) = \theta^{(i)}$. Suppose $\forall i \in [m]$, mechanism $\mathcal{M}_i$ satisfies a statistic maximal leakage guarantee with respect to $\mathfrak{g}$ of $\Pi_{\mathcal{M}_i, \mathfrak{g}}$, where the guarantee is taken in the worst case over all possible previous outputs. Let $\mathcal{M} = \mathcal{M}_1 \circ \mathcal{M}_2 \circ \dots \circ \mathcal{M}_m$ denote the composition of these adaptively chosen mechanisms. The SML with respect to an adversary that can see all intermediate outputs $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(m)}$ can be bounded as $\Pi_{\mathcal{M}, \mathfrak{g}} \leq \sum_{i \in [m]} \Pi_{\mathcal{M}_i, \mathfrak{g}}$.*

(Proof in Appendix D.) Theorem 2 shows that statistic maximal leakage degrades additively when one or more possibly adaptive mechanisms are applied multiple times sequentially to the original dataset. This additive result is similar in form to analogous composition results for other privacy metrics, including pointwise maximal leakage [27] (Thm. 13) and differential privacy [66] (Thm. III.1). In particular, we note that the result for pointwise maximal leakage does not imply our result, nor vice versa. Pointwise maximal leakage requires knowledge of the prior distribution of the input data and assumes the secret function is unknown; the composition property for the worst-case secret function or prior does not imply the composition property for an arbitrary secret or prior.

Theorem 3 (Post-Processing). *Let $\mathcal{M}$ be a data release mechanism whose SML is $\Pi_{\mathcal{M}, \mathfrak{g}}$. Let $\tilde{\mathcal{M}}$ be an arbitrary (possibly randomized) mechanism defined by $\mathbb{P}_{\Theta''|\Theta'}$. Then the SML of $\tilde{\mathcal{M}} \circ \mathcal{M}$ is $\Pi_{\tilde{\mathcal{M}} \circ \mathcal{M}, \mathfrak{g}} \leq \Pi_{\mathcal{M}, \mathfrak{g}}$.*

(Proof in Appendix E) Theorem 3 shows that applying an arbitrary (possibly randomized) mechanism to the output of a mechanism that satisfies statistic maximal leakage will not degrade statistic maximal leakage.

5. Mechanism Design for Tabular, Categorical Data

We next study two natural mechanisms for releasing *tabular* data with attributes that take values from a finite set under SML: randomized response and the quantization mechanism. Both have been widely used in practice [17,28,30,31]. Our goal is to understand (a) if each of these satisfies a SML guarantee, and (b) if so, which one has a better privacy–utility tradeoff?

5.1. Setup

The data holder holds a tabular dataset $\mathbf{D}$ with n rows and c columns, i.e., n samples with c attributes for each sample. Let Γ be the set of combinations of attributes for samples existing in the original dataset $\mathbf{D}$, and $d \triangleq |\Gamma|$. For example, suppose our dataset has binary columns “Above age 18?” and “Registered to vote in the U.S?” and only includes samples with attribute values “(Yes, Yes)” and “(Yes, No)”; then $d = 2$ and the number of columns $c = 2$.

We assume there is some unknown true *feasible* set of attribute combinations Γ^* , where $d^* \triangleq |\Gamma^*|$ and $\Gamma \subseteq \Gamma^*$. In our example, in the U.S., $d^* = 3$ because voters must be at least 18 years of age, so the only feasible combinations of attributes are (Above 18, Registered), (Above 18, Not registered), (Under 18, Not registered). We use $\hat{\Gamma}^*$ to denote the data holder’s estimate of the true support Γ^* (e.g., this could be obtained from public data), where $\hat{d}^* \triangleq |\hat{\Gamma}^*|$. In our voting example, Γ^* can be accurately estimated based on public information, i.e., $\hat{\Gamma}^* = \Gamma^*$. We start by analyzing this special case in Section 5.3, but

in general, they need not be the same set, and we subsequently analyze robustness to misspecification of the feasible attribute set in Section 5.4.

The data release mechanism is designed such that samples with attribute combinations in $\Gamma \cup \hat{\Gamma}^*$ may exist in the released dataset $\mathbf{D}'$. Let $\hat{\Gamma}_0^*$ and $\hat{\Gamma}_1^*$ be the feasible and infeasible attribute combinations within the estimated set $\hat{\Gamma}^*$ respectively, i.e., $\hat{\Gamma}_0^* = \hat{\Gamma}^* \cap \Gamma^*$ and $\hat{\Gamma}_1^* = \hat{\Gamma}^* \setminus \Gamma^*$. Let $\hat{d}_0^* = |\hat{\Gamma}_0^*|$ and $\hat{d}_1^* = |\hat{\Gamma}_1^*|$, so we have $\hat{d}^* = \hat{d}_0^* + \hat{d}_1^*$. We illustrate the relation between Γ , Γ^* , $\hat{\Gamma}^*$, $\hat{\Gamma}_0^*$ and $\hat{\Gamma}_1^*$ in Figure 2.




Figure 2. Relation between Γ , Γ^* , $\hat{\Gamma}^*$, $\hat{\Gamma}_0^*$ and $\hat{\Gamma}_1^*$. Γ^* is the set depicted in black, containing all feasible attribute combinations. Γ in red is the set of attribute combinations existing in the dataset $\mathbf{D}$, and $\Gamma \subseteq \Gamma^*$. $\hat{\Gamma}^*$ in dark yellow is the data holder's estimate of Γ^* ; it may contain both feasible attribute combinations, as shown in the blue sub-region with the notation $\hat{\Gamma}_0^*$, and infeasible attribute combinations, as shown in the green sub-region with the notation $\hat{\Gamma}_1^*$.

In the remainder of this paper, we model a tabular dataset as a one-dimensional histogram (which has a one-to-one correspondence with an underlying parameter vector θ). In particular, suppose the released dataset $\mathbf{D}'$ has the same size as the original dataset $\mathbf{D}$. $\mathbf{D}$ and $\mathbf{D}'$ can be represented by categorical distribution parameters θ and θ' respectively, where each category corresponds to the fraction of records (out of n) with a given combination of attribute values, e.g., (Yes, Yes) in our voting example. So for example, suppose our input dataset contained four records, as listed in Figure 3 below. Then we would generate from this dataset a histogram with $|\Gamma \cup \hat{\Gamma}^*|$ bins, where each bin contains the fraction of records with a specific combination of attribute values. This representation of tabular, categorical data captures correlations across columns by modeling the empirical joint distribution. Finally, let the data holder's precision level of the categorical distribution be τ , i.e., the probability mass of any category is a value within set $\{0, 1/\tau, 2/\tau, \dots, 1\}$. We assume $\tau \leq n$, where n is the number of samples in the dataset.




Figure 3. For a tabular dataset, we can generate a categorical distribution with each category corresponding to the fraction of records with a given combination of attribute values. θ is the parameter of the constructed categorical distribution.

Suppose the set of secret values admits a total ordering, and without loss of generality, suppose that $g_1 < g_2 < \dots < g_s$. For convenience, we use Θ_i to represent Θ_{g_i} , the parameter set with secret value g_i , $\forall i \in [s]$.

5.2. SML Mechanisms

In the following, we present two mechanisms that satisfy an SML guarantee: randomized response and quantization mechanism. Note that among these, only randomized

response also satisfies a DP guarantee. Finally, in Section 5.3, we compare their privacy–utility tradeoffs.

5.2.1. Randomized Response (RR)

RR is a widely-used mechanism in the DP literature for discrete distributions [31,67]. We first consider a RR mechanism $\mathcal{M}_{RR}$ that outputs the original distribution parameter with some probability, and otherwise releases a different distribution’s parameters uniformly at random over the parameter set. Specifically, $\mathcal{M}_{RR}$ can be written as follows, $\forall \theta \in \Theta$:

$$\mathbb{P}(\mathcal{M}_{RR}(\theta) = \theta') = \begin{cases} \frac{e^\epsilon}{|\Theta'| + e^\epsilon - 1}, & \theta' = \theta, \\ \frac{1}{|\Theta'| + e^\epsilon - 1}, & \theta' \in \Theta' \setminus \{\theta\}, \end{cases}$$

where ϵ is a mechanism hyperparameter controlling the probability of outputting the original distribution parameter; larger ϵ increases this probability.

Theorem 4 (Privacy and Distortion of Randomized Response). *For any secret function $\mathbf{g}$, the SML and distortion of RR are:*

$$\Pi_{\mathcal{M}_{RR}, \mathbf{g}} = \log \frac{1 + sr}{1 + r}, \quad \Delta_{\mathcal{M}_{RR}} = \frac{\hat{d}_0^* - 1}{\hat{d}_0^*(1 + r)}. \quad (10)$$

$$\text{where } r \triangleq \frac{e^\epsilon - 1}{|\Theta'|} = \frac{e^\epsilon - 1}{\binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^* - 1}}.$$

(Proof in Appendix F).

5.2.2. Quantization Mechanism (QM)

QM has also been adopted in privacy-preserving mechanism design [8,17,28]. This mechanism first partitions the secret values in $\mathbf{G}$ into bins of size I , so that each bin contains I elements. It then releases, uniformly at random, a distribution parameter from the parameter set associated with the median-index secret of the bin containing the original secret. Specifically, $\mathcal{M}_{QM}$ can be written as $\forall k \in \{0, 1, \dots, \lceil \frac{s}{I} \rceil - 1\}, j \in [I], \theta \in \Theta_{kI+j}$:

$$\mathcal{M}_{QM}(\theta) \sim \text{Unif}(\Theta'_{R(k)}),$$

where $R(k) = \lfloor (k + \frac{1}{2})I \rfloor + 1$ represents the median index of the k -th bin, and $\Theta'_{R(k)}$ represents the released parameter set with secret value $g_{R(k)}$. More generally, if there does not exist a total ordering over the set of secrets $\mathbf{G}$, QM releases uniformly at random a distribution parameter whose secret is contained in the corresponding bin. In the following analysis, we assume that the secret is the fraction of a category (i.e., one of the PMF values).

Theorem 5 (Privacy and Distortion of Quantization Mechanism). *For any secret function $\mathbf{g}$, the privacy of QM is*

$$\Pi_{\mathcal{M}_{QM}, \mathbf{g}} = \log \left\lceil \frac{s}{I} \right\rceil. \quad (11)$$

When the secret $\mathbf{g}$ is the fraction of the records with a given attribute combination ψ , i.e.,

$$\mathbf{g}(\theta) = \frac{n_\psi}{n}, \quad (12)$$

where n_ψ denotes the number of records whose attribute combination is ψ , the distortion of QM is

$$\Delta_{\mathcal{M}_{\text{QM}}} = \frac{1}{2} + \frac{\hat{d}_0^* \left\lfloor \frac{I}{2} \right\rfloor - \tau}{2\tau(\hat{d}_0^* - 1)}. \quad (13)$$

(Proof in Appendix F.)

5.3. Privacy–Utility Tradeoffs: A Special Case

When $\hat{\Gamma}^* = \Gamma^*$, i.e., the data holder correctly estimates the feasible attribute combination set, we analyze and compare the privacy–distortion tradeoffs between RR and QM as follows. Note that in our analysis of QM, we constrain the secret function $\mathbf{g}$ to be the PMF value for a specific category (e.g., “the fraction of people above age 18 registered to vote in the U.S.”) to simplify the analysis.

Corollary 1 (Mechanism Comparison). *When the secret $\mathbf{g}$ is the fraction of records with a given attribute combination ψ , i.e.,*

$$\mathbf{g}(\theta) = \frac{n_\psi}{n},$$

for any non-trivial privacy budget $T < \log s$, if $\Pi_{\mathcal{M}_{\text{QM}}, \mathbf{g}} = \Pi_{\mathcal{M}_{\text{RR}}, \mathbf{g}} \leq T$, we have

$$\lim_{\tau \rightarrow \infty} \frac{\Delta_{\mathcal{M}_{\text{RR}}}}{\Delta_{\mathcal{M}_{\text{QM}}}} \geq 1.$$

(Proof in Appendix F.) From Section 1, we know that when the precision level τ (Section 5.1) is high enough, for the same SML guarantee, QM performs at least as well as RR as long as SML does not achieve its upper bound $\log s$. Intuitively, this is because the output space of RR covers the full support of Θ' , while the output space of QM is significantly reduced, making the released distribution closer to the input distribution. We conjecture that a similar result may hold when comparing QM to general DP mechanisms that also satisfy an SML guarantee, since DP mechanisms must have nonzero probability mass on their entire output support, whereas this is not true of SML. Since our utility metric is worst-case over data distributions, DP mechanisms thus incur a high utility penalty.

5.4. Robustness to Misspecification of Feasible Attribute Combinations

When $\hat{\Gamma}^* \neq \Gamma^*$, i.e., the data holder only partially knows the feasible attribute combination set, we provide a robustness result for the privacy of the mechanisms. We define robustness to support mismatch below, using γ to quantify the divergence in the SML value, $\Pi_{\mathcal{M}, \mathbf{g}} - \Pi_{\mathcal{M}, \mathbf{g}}^*$, under a given level of support mismatch, $d^* - \hat{d}_0^*$.

Definition 1 (Robustness to support mismatch). *Consider a tabular dataset $\mathbf{D}$ with attribute combination set Γ . For any mechanism $\mathcal{M}$, let $\Pi_{\mathcal{M}, \mathbf{g}}$ be the SML of $\mathcal{M}$ if its released dataset contains samples with attribute combinations in $\hat{\Gamma}^* \cup \Gamma$, and $\Pi_{\mathcal{M}, \mathbf{g}}^*$ be the SML of $\mathcal{M}$ if its released dataset only contains samples with attribute combinations in Γ^* . The mechanism $\mathcal{M}$ is γ -robust if for any Γ ,*

$$\Pi_{\mathcal{M}, \mathbf{g}} - \Pi_{\mathcal{M}, \mathbf{g}}^* \leq \gamma(d^* - \hat{d}_0^*). \quad (14)$$

Proposition 3. *Consider a dataset $\mathbf{D}$ with τ samples. Regardless of the value of $\hat{d}_1^*$, RR is $\log 3$ -robust if its hyperparameter ϵ satisfies $e^\epsilon - 1 \leq (\tau + \hat{d}_1^* - 1)/s$, and QM with any interval length I is 1-robust when the secret is the fraction of the records with a certain arbitrary attribute combination, as specified in Equation (12).*

(Proof in Appendix G.) Proposition 3 indicates that RR is robust to support mismatch when it satisfies certain privacy constraints, and QM is robust under certain secret types.

We also analyze the upper and lower bounds on privacy and distortion for both mechanisms when $\hat{\Gamma}^* \neq \Gamma^*$, as well as the tightness of the privacy bounds. These results are deferred to Appendix H.

6. Empirical Evaluation

In this section, we conduct empirical experiments on a real-world dataset to analyze and compare the privacy–distortion tradeoffs of our proposed randomized response (RR) and quantization mechanism (QM), and evaluate the performance of the quantization mechanism under downstream tasks. We first introduce the experimental settings.

6.1. Experimental Setting

Dataset. We conduct our empirical evaluation on the Census Income dataset [68], which collects information from 48,842 individuals ($n = 48,842$) about their income, education level, age, gender, and more. We set the distribution precision level as $\tau = n$. The dataset contains 22,381 unique attribute combinations ($d = 22,381$), all of which we assume to be contained in the estimated attribute combination set $\hat{\Gamma}^*$, i.e., $\Gamma \subseteq \hat{\Gamma}^*$. As described in Section 5.1, the category distribution parameter θ can be constructed with each category corresponding to a certain attribute combination, e.g., (white, male, age 32, master’s degree, etc.), with the value being set as the fraction of samples with this attribute combination.

Baselines. Besides the quantization mechanism and randomized response proposed, we also include and compare with the mechanism of Wu et al. [69], which satisfies a maximal leakage guarantee. For ease of reference, we call this mechanism MaxL. It is known to exhibit a privacy–utility tradeoff with bounded sub-optimality under a class of utility functions that does not include distance measures like total variation distance. However, we do not know of other mechanisms that satisfy a maximal leakage guarantee and can be applied to our dataset.

In short, for a given distribution parameter set Θ and a released distribution parameter set Θ' , let $\mathbf{A}$ be a subset of Θ' , u be the cost function (total variation distance in our experiment), f be a mapping function that satisfies $f(\theta) = \arg \min_{\theta' \in \mathbf{A}} u(\theta, \theta')$, and $\mathcal{U}_{\mathbb{P}_{\Theta}}(\mathbf{A}) = \mathbb{E}_{\Theta}[u(\theta, f(\theta))]$ be the expected cost under the prior distribution $\mathbb{P}_{\Theta}$. MaxL works by initializing $\mathbf{A} = \emptyset$, and iteratively adding an element $\theta' \in \Theta' \setminus \mathbf{A}$ to $\mathbf{A}$, where θ' satisfies $\theta' = \arg \min_{\theta' \in \Theta' \setminus \mathbf{A}} \mathcal{U}_{\mathbb{P}_{\Theta}}(\mathbf{A} \cup \theta')$. MaxL terminates when $\mathbf{A} = \Theta'$ or $\min_{\theta' \in \Theta' \setminus \mathbf{A}} \mathcal{U}_{\mathbb{P}_{\Theta}}(\mathbf{A} \cup \theta') = \mathcal{U}_{\mathbb{P}_{\Theta}}(\mathbf{A})$, and design the data release mechanism $\mathcal{M}_{\text{MaxL}}$ according to the mapping function f .

Since we assume the prior distribution $\mathbb{P}_{\Theta}$ is unknown in our setting, we replace the expected cost $\mathcal{U}_{\mathbb{P}_{\Theta}}(\mathbf{A})$ in MaxL by $\max_{\theta \in \mathbf{A}} \mathcal{U}_{\mathbb{P}_{\Theta}}(\mathbf{A})$. When the cost function is set as total variation distance, the performance of MaxL is sensitive to the choice of the released parameter set Θ' . For example, if $\Theta' = \Theta$, we can easily check that MaxL will always release the input distribution parameter itself, i.e., $\mathcal{M}_{\text{MaxL}}(\theta) = \theta$, resulting in no privacy protection. To ensure a fair comparison, we design the released parameter set Θ' based on a similar idea to the Quantization Mechanism, that is, we partition the secret values into subsets of size I , and for each subset, we uniformly select a distribution parameter with the secret within the subset and add it to Θ' . We vary the interval length I to change the privacy guarantee of MaxL.

Evaluation Metrics. Due to the above results showing that attribute privacy has a poor privacy–utility tradeoff under our settings, we use SML to measure the privacy performance in the remainder of this paper. To measure the utility of the data release mechanism, we adopt the worst-case expected total variation distance from Equation (2). We quantify

downstream utility by conducting a classification task on the downstream data; namely, we learn a classifier predicting whether an individual has high income or not based on a random forest trained on the released (perturbed) data and compute ROC curves.

6.2. Qualitative Comparison to Attribute Privacy

We compare the privacy (and utility) of a mechanism (RR with hyperparameter $\epsilon = 4$) under the privacy measures SML and a differential privacy (DP)-based measure; namely, we consider attribute privacy (AP) [23]. The comparison is not meant to rank SML and AP as universally comparable privacy notions. Instead, it tests whether SML and AP are informative under our chosen utility metric and secret-statistic settings. We chose RR because it satisfies both AP and SML, whereas QM does not satisfy a finite AP guarantee for any interval size less than the full range of secrets. Additionally, the mechanism proposed in [23] applies only to continuous query results from the dataset, rather than to the entire dataset.

A fair comparison between SML and AP is challenging because these two measures have different definitions, and hence the numeric values of SML and attribute/inferential privacy cannot be directly compared. We thus run two experiments: (1) We fix the dataset and mechanism, and study how the values of SML and AP *change* with the number of possible secret values. (2) We evaluate the privacy–distortion tradeoff for the RR mechanism under the AP privacy measure. First, in Figure 4a, we vary s , the number of possible secret values the original dataset can take, by quantizing the secret quantity to different granularities (i.e., the smaller the quantization interval, the greater the number of possible secret values). As the number of possible secret values increases, the value of SML increases, i.e., the privacy measure degrades. This aligns with our intuition: for a given dataset, under the same randomized response mechanism, a larger number of possible secret values results in more uncertainty regarding the true secret value prior to data release. However, under randomized response, the information the adversary obtains after data release remains the same—specifically, the probability of RR outputting the original data remains the same regardless of the number of secrets. Therefore, we intuitively expect that RR with more secrets should leak more information, and we observe that SML has this property.




Figure 4. Under attribute privacy, the privacy level of RR does not vary with the number of possible secrets, indicating its inability to capture the influence of the secret value space on the information leakage (a); the distortion of RR always achieve its trivial upper bound 1 until the attribute privacy value is larger than 44,320, indicating that attribute privacy is a much more conservative metric (b). (a) Privacy value of RR with $\epsilon = 4$ vs. Number of possible secret values; (b) Distortion vs. Attribute Privacy for RR with $\epsilon = 4$.

In contrast, attribute privacy only measures whether the mechanism outputs are different under different secrets, and only depends on the mechanism itself; hence, the attribute privacy level of RR does not change with the number of possible secrets.

We additionally evaluate the privacy–distortion tradeoff of RR under attribute privacy in Figure 4b, where the distortion metric is the worst-case total variation distance from Equation (2). Figure 4b shows that the distortion of RR achieves its trivial upper bound of 1 until the attribute privacy value is larger than 44,320, which is impractically large. This suggests that for our utility metric (which is a worst-case metric, and hence may be conservative) attribute privacy incurs a poor privacy–utility tradeoff for reasonable privacy parameters.

6.3. Privacy–Distortion Tradeoffs

In this section, we compare the privacy–distortion tradeoffs between QM, RR, and MaxL with the assumption that $\Gamma = \Gamma^* = \hat{\Gamma}^*$, i.e., all feasible attribute combinations are contained in the dataset, and there is no estimate error for the attribute combinations. We defer the privacy–distortion tradeoff analysis under the case where $\Gamma \subseteq \Gamma^* \neq \hat{\Gamma}^*$ in Appendix J.

We consider the secret as the fraction of an arbitrary category—namely, (white, male, age 32, Master’s degree)—and compare the privacy–utility tradeoffs between RR, QM, and MaxL in Figure 5. To obtain different privacy guarantees, we vary the hyperparameter ϵ in RR and vary the interval length I in QM and MaxL. The SML for this setting takes values in $[0, 16]$, where the upper bound arises from the limited size of the secret value space. From Figure 5, we observe that for the same privacy guarantee, the distortion of RR is almost twice as large as that of QM under most SML levels, indicating better performance of QM, which is in line with the theoretical insight in Corollary 1. For MaxL, which can be regarded as a variant of QM in our setting, it also has a lower distortion than RR at most SML levels, while showing a worse privacy–distortion performance than QM. This intuitively makes sense, because MaxL is designed to prevent leakage of *any* secret, whereas QM is only protecting the specified secret.




Figure 5. Privacy–utility trade-offs of RR, QM, and MaxL when the secret is the fraction of an arbitrary category.

6.4. Evaluation of the Quantization Mechanism Under Downstream Tasks

We next analyze the downstream utility of QM by conducting a classification task to predict whether an individual has high income or not based on a random forest model trained on the released data. We set the secret as the difference between the proportion of white and non-white high-income people (>\$50 k/yr) among their own race groups. The SML for this setting takes values in $[0, 4.09]$.

6.4.1. Perfectly Calibrated Attribute Sets

We first consider the case where $\Gamma = \Gamma^* = \hat{\Gamma}^*$, and vary the quantization set size I to achieve different levels of SML from 0 to 2. For each SML level, we conduct the experiment 20 times with independent mechanism outputs, show the averaged ROC curve of the random forests trained on corresponding released datasets. We then compare this with the performance of the random forest trained on the original dataset in Figure 6. We observe that as SML increases (weaker privacy), AUC (area under the ROC curve) increases, indicating the improvement of the downstream task utility. When SML is as little as 2, the AUC is close to its upper bound (0.89 in raw dataset). When SML is 0 (perfect privacy), the utility drops mildly (0.14 AUC) compared to the utility of the original dataset. Note that a perfect privacy-preserving mechanism can still achieve reasonable utility on this task since we only aim to protect a secret of the dataset, rather than the whole data.




Figure 6. Comparison of ROC curves of random forests under QM with different SML.

6.4.2. Mis-Calibrated Attribute Sets

We then focus on the case where $\Gamma \subseteq \Gamma^* \neq \hat{\Gamma}^*$ in Figure 7. Under each setting, we conduct 20 independent experiments, show the average AUC with standard deviation of the random forests trained on the released datasets.

In Figure 7a, we compare the AUC of random forests with different SML levels when the data holder misses 5 feasible attribute combinations in their pre-held attribute set, i.e., $d^* - \hat{d}_0^* = 5$. We consider the case where the data holder does not introduce invalid attribute combinations, i.e., $\hat{d}_1^* = 0$, and vary the SML from 2.8 to 4. We observe that compared to Figure 6, the performance of the random forests is worse even under more relaxed privacy constraints (AUC = 0.82 when SML = 4 vs. AUC = 0.82 when SML = 1 in Figure 6).

In Figure 7b, we fix the privacy budget as SML = 4, and compare the utility when the data holder has different prior knowledge of feasible attribute combinations, i.e., different number of missing feasible combinations $d^* - \hat{d}_0^*$. We vary $d^* - \hat{d}_0^*$ from 0 to 100, and, similar to the settings of Figure 7a, set $\hat{d}_1^* = 0$. As we can observe, with the increasing number of missing feasible attribute combinations, the utility first drops sharply—AUC drops from 0.87 to 0.76 as $d^* - \hat{d}_0^*$ increases from 0 to 10, and then drops mildly—AUC drops from 0.76 to 0.74 as $d^* - \hat{d}_0^*$ increases from 10 to 100.

In Figure 7c–e, we fix the data holder's pre-held attribute combination sets with $d^* - \hat{d}_0^* = 5, 20$ and 100, respectively, and compare the utility under different numbers of introduced invalid attribute combinations, i.e., $\hat{d}_1^*$. We observe in Figure 7c that as $\hat{d}_1^*$ increases, the privacy upper bounds drops linearly while the utility only degrades slightly when $\hat{d}_1^*$ increases from 0 to 5 k, or from 10 k to 20 k. When the number of feasible attribute combinations the data holder misses increases to 20 in Figure 7d, the privacy upper bounds drops from 4.09 to 2.88 in a concave curve as $\hat{d}_1^*$ increases, from 50 k to 200 k, while the

utility only suffers a slight degradation as $\hat{d}_1^*$ increases from 50 k to 200 k. Those results indicate that introducing more invalid attribute combinations results in better (lower) SML while not severely harming utility. We also observe that as the number of missing feasible attribute combinations increases, more invalid attribute combinations are needed to achieve a given improvement in SML. However, as we can observe in Figure 7e, when $d^* - \hat{d}_0^* = 100$, the SML of the mechanism always achieves its trivial upper bound whatever $\hat{d}_1^*$ is, which indicates that with more missing feasible attribute combinations, adding invalid attribute combinations eventually cannot help improve the privacy performance of the mechanism. In Appendix I, we show through an approximation argument why the privacy expressions decay in the patterns observed in Figure 7c–e. Roughly, in the approximation regime where $\hat{d}_1^* \ll \hat{d}_0^*$ and $d^* - \hat{d}_0^* \ll \tau + \hat{d}_0^*$, there exists a positive constant $\alpha = \frac{\log s - \log \left[\frac{s}{1} \right]}{\log [1 + \tau / (\tau + \hat{d}_0^*)]}$ such that with increasing $\hat{d}_1^*$, the privacy value drops linearly when $0 \leq d^* - \hat{d}_0^* < \alpha$, and always achieves its trivial upper bound $\log s$ when $d^* - \hat{d}_0^* \geq \alpha$.




Figure 7. Comparisons of ROC curves of random forests trained on released data by quantization mechanism with different SML, missed feasible attribute combinations ($d^* - \hat{d}_0^*$), or invalid attribute combinations inserted ($\hat{d}_1^*$).

7. Future Work

In this section, we outline several promising directions for future research, including developing alternative leakage metrics for different threat models, extending statistic maximal leakage to continuous parameter spaces and broader classes of data, and characterizing how the choice of secret affects downstream utility.

Design of Alternative Metrics. The operational meaning of statistic maximal leakage is to measure the information gain in an attacker's knowledge about the most likely secret value after data release, under worst-case prior and attack strategy. However, in some scenarios, especially those related to social justice and fairness, protecting the most unlikely posterior secret value can also be crucial. Additionally, while statistic maximal leakage assesses the attack success probability over the randomness of the released data, a more conservative approach would measure this probability under the worst-case output. In Appendix K,

we take a first step towards alternative metrics under those scenarios and analyze their connections with the indistinguishability-based measure inspired by differential privacy.

Extension to Continuous Parameter Space. The definition and analysis of statistic maximal leakage in Section 4 are under the discrete parameter space. In Appendix L, we take an initial step toward extending statistic maximal leakage to the continuous distribution parameter and secret spaces. A complete extension, however, requires a rigorous measure-theoretic treatment to establish that the proposed formulation is well-defined and that the corresponding characterizations remain valid. Furthermore, efficiently calculating or estimating statistic maximal leakage over a continuous parameter space is an open problem, as is designing data release mechanisms with desirable privacy–utility tradeoffs.

For instance, we study tabular datasets with categorical values in this work. This is a very limited class of data. Designing mechanisms that satisfy an SML guarantee for more general classes of data is an important and interesting question.

Relation Between Secret and Utility. In our experiments, we (perhaps surprisingly) observed nonzero utility even at a perfect SML guarantee of zero. This is likely due to our choice of secret, which was correlated with the selected downstream task very loosely. An interesting question is to characterize conditions on the secret and the downstream task to explain how SML will affect utility.

8. Conclusions

In this work, we propose *statistic maximal leakage*, a prior-independent and secret-specific privacy measure quantifying the information leakage of a given, known secret. We show that SML satisfies post-processing and composition properties. Although the calculation of statistic maximal leakage is NP-hard in general, under deterministic mechanisms, we show an efficient statistic maximal leakage calculation process. We analyze and compare the privacy–utility tradeoffs of randomized response and the quantization mechanism. Both theoretical and empirical results suggest a better privacy–utility tradeoff for the quantization mechanism.

Author Contributions: Conceptualization, S.W., Z.L. and G.F.; methodology, S.W.; formal analysis, S.W.; investigation, S.W.; resources, G.F.; data curation, S.W.; writing—original draft preparation, S.W.; writing—review and editing, Z.L. and G.F.; visualization, S.W.; funding acquisition, G.F. All authors have read and agreed to the published version of the manuscript.

Funding: The authors gratefully acknowledge the support of NSF grants CIF-1705007, CCF-2338772, and RINGS-2148359, as well as support from the Sloan Foundation, Intel, J.P. Morgan Chase, Siemens, Bosch, and Cisco. This material is based upon work supported in part by the U.S. Army Research Office and the U.S. Army Futures Command under Contract No. W911NF20D0002.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting this study is available at: <https://doi.org/10.24432/C5XW20>.

Conflicts of Interest: Author Zinan Lin was employed by the company Microsoft Research. The authors declare that the research was conducted in the absence of any commercial or financial relationships, and they declare no conflict of interest.

Appendix A. Proofs of Lemma 1, Property 1, and Proposition 1

Appendix A.1. Proof of Lemma 1

Proof. We show that the proposed expression is both an upper bound and a lower bound on the original distortion measure.

Recall that the original distortion measure is defined as

$$\Delta_{\mathcal{M}} = \sup_{\mathbb{P}_{\Theta} \in \mathcal{P}(\Theta)} \mathbb{E}_{\Theta, \Theta' = \mathcal{M}(\Theta)} [D_{\text{TV}}(f_{\Theta} \| f_{\Theta'})].$$

Let $\mathbb{P}_{\Theta}^* = \arg \sup_{\mathbb{P}_{\Theta}} \mathbb{E}_{\Theta, \Theta' = \mathcal{M}(\Theta)} [D_{\text{TV}}(f_{\Theta} \| f_{\Theta'})]$, and for convenience, we define $P_{\theta} \triangleq \mathbb{P}_{\Theta}^*(\theta)$. We can get that $\Delta_{\mathcal{M}}$ can be upper bounded by

$$\begin{aligned} \Delta_{\mathcal{M}} &= \sup_{\mathbb{P}_{\Theta} \in \mathcal{P}(\Theta)} \mathbb{E}_{\Theta, \Theta' = \mathcal{M}(\Theta)} [D_{\text{TV}}(f_{\Theta} \| f_{\Theta'})] \\ &= \sum_{\theta \in \Theta} P_{\theta} \cdot \mathbb{E}_{\Theta' = \mathcal{M}(\theta)} [D_{\text{TV}}(f_{\theta} \| f_{\Theta'})] \\ &\leq \sup_{\theta \in \Theta} \mathbb{E}_{\Theta' = \mathcal{M}(\theta)} [D_{\text{TV}}(f_{\theta} \| f_{\Theta'})]. \end{aligned}$$

To get the lower bound of $\Delta_{\mathcal{M}}$, we construct a prior such that $\mathbb{P}_{\Theta}(\theta) = 1$ when $\theta = \arg \sup_{\theta \in \Theta} \mathbb{E}_{\Theta' = \mathcal{M}(\theta)} [D_{\text{TV}}(f_{\theta} \| f_{\Theta'})]$, and we have

$$\begin{aligned} \Delta_{\mathcal{M}} &= \sup_{\mathbb{P}_{\Theta} \in \mathcal{P}(\Theta)} \mathbb{E}_{\Theta, \Theta' = \mathcal{M}(\Theta)} [D_{\text{TV}}(f_{\Theta} \| f_{\Theta'})] \\ &\geq \sup_{\theta \in \Theta} \mathbb{E}_{\Theta' = \mathcal{M}(\theta)} [D_{\text{TV}}(f_{\theta} \| f_{\Theta'})]. \end{aligned}$$

Therefore, $\Delta_{\mathcal{M}} = \sup_{\theta \in \Theta} \mathbb{E}_{\Theta' = \mathcal{M}(\theta)} [D_{\text{TV}}(f_{\theta} \| f_{\Theta'})]$. $\square$

Appendix A.2. Proof of Property 1

Proof. For the worst-case min-entropy leakage $\mathcal{L}_{\text{MEL}}$, it can be derived as

$$\begin{aligned} \mathcal{L}_{\text{MEL}} &= \sup_{\mathbb{P}_{\Theta}} \log \frac{\sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta \Theta'}(\theta, \theta')}{\sup_{\theta \in \Theta} \mathbb{P}_{\Theta}(\theta)} \\ &= \sup_{\mathbb{P}_{\Theta}} \log \frac{\sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \mathbb{P}_{\Theta}(\theta)}{\sup_{\theta \in \Theta} \mathbb{P}_{\Theta}(\theta)} \\ &\leq \sup_{\mathbb{P}_{\Theta}} \log \frac{\sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \cdot \sup_{\theta \in \Theta} \mathbb{P}_{\Theta}(\theta)}{\sup_{\theta \in \Theta} \mathbb{P}_{\Theta}(\theta)} \\ &= \log \sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta). \end{aligned}$$

This upper bound can be achieved when $\mathbb{P}_{\Theta}$ satisfies

$$\forall \theta' \in \Theta' : \arg \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \cap \arg \sup_{\theta \in \Theta} \mathbb{P}_{\Theta}(\theta) \neq \emptyset.$$

Thus, $\mathcal{L}_{\text{MEL}} = \log \sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta)$.

For statistic maximal leakage $\Pi_{\mathcal{M}, g}$, we defer the derivations to Appendix A.3, from which we have

$$\begin{aligned}\Pi_{\mathcal{M},g} &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \cdot \mathbb{P}_{\Theta}(\theta)}{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta}(\theta)} \\ &\geq \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta)}{|\Theta_g|}\end{aligned}\quad (\text{A1})$$

$$\begin{aligned}&\geq \log \frac{\sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \sup_{\theta \in \Theta_g} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta)}{\sup_{g \in \mathbf{G}} |\Theta_g|} \\ &= \log \frac{\sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta)}{\sup_{g \in \mathbf{G}} |\Theta_g|} \\ &= \mathcal{L}_{MEL} - \sup_{g \in \mathbf{G}} \log |\Theta_g|.\end{aligned}\quad (\text{A2})$$

Equality is achieved in (A1) when $\mathbb{P}_{\Theta}$ satisfies $\forall g \in \mathbf{G}, \forall \theta_1, \theta_2 \in \Theta_g : \mathbb{P}_{\Theta}(\theta_1) = \mathbb{P}_{\Theta}(\theta_2)$, and (A2) holds because $\cup_{g \in \mathbf{G}} \Theta_g = \Theta$. Additionally, we have

$$\begin{aligned}\Pi_{\mathcal{M},g} &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \cdot \mathbb{P}_{\Theta}(\theta)}{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta}(\theta)} \\ &\leq \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\sup_{\theta \in \Theta_g} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \cdot \sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta}(\theta)}{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta}(\theta)} \\ &= \log \sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \\ &= \mathcal{L}_{MEL}.\end{aligned}$$

This gives

$$\mathcal{L}_{MEL} - \sup_{g \in \mathbf{G}} \log |\Theta_g| \leq \Pi_{\mathcal{M},g} \leq \mathcal{L}_{MEL}.\quad (\text{A3})$$

For maximal leakage $\mathcal{L}_{ML}$, from [8] (Thm. 1), we know that

$$\mathcal{L}_{ML} = \log \sum_{\theta' \in \Theta'} \sup_{\substack{\theta \in \Theta: \\ \mathbb{P}_{\Theta}(\theta) > 0}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta).$$

Therefore, we have

$$\begin{aligned}\sup_{\mathbb{P}_{\Theta}} \mathcal{L}_{ML} &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta' \in \Theta'} \sup_{\substack{\theta \in \Theta: \\ \mathbb{P}_{\Theta}(\theta) > 0}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \\ &= \log \sum_{\theta' \in \Theta'} \sup_{\theta \in \Theta} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \\ &= \mathcal{L}_{MEL}.\end{aligned}$$

Based on Equation (A3), we can get that

$$\sup_{\mathbb{P}_{\Theta}} \mathcal{L}_{ML} - \sup_{g \in \mathbf{G}} \log |\Theta_g| \leq \Pi_{\mathcal{M},g} \leq \sup_{\mathbb{P}_{\Theta}} \mathcal{L}_{ML}.$$

□

Appendix A.3. Proof of Proposition 1

Proof. We first analyze the upper bound of statistic maximal leakage under a fixed prior. Let Θ_g be the set of distribution parameters whose secret value is g , i.e.,

$\Theta_g = \{\theta \in \Theta | g(\theta) = g\}$. For statistic maximal leakage $\Pi_{\mathcal{M},g}$, we have $\Pi_{\mathcal{M},g} = \sup_{\mathbb{P}_\Theta} \log \frac{\sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{G\Theta'}(g, \theta')}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)}$. Under a fixed prior distribution $\mathbb{P}_\Theta$, we can get that

$$\begin{aligned}
 & \log \frac{\sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{G\Theta'}(g, \theta')}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)} \\
 &= \log \frac{\sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g) \mathbb{P}_G(g)}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)} \\
 &\leq \log \frac{\sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g) \cdot \sup_{g \in \mathbf{G}} \mathbb{P}_G(g)}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)} \\
 &= \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g) \\
 &= \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \sum_{\theta \in \Theta} \mathbb{P}_{\Theta'|G\Theta}(\theta'|g, \theta) \cdot \mathbb{P}_{\Theta|G}(\theta|g) \tag{A4} \\
 &= \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \sum_{\theta \in \Theta} \mathbb{P}_{\Theta'|G\Theta}(\theta'|g, \theta) \frac{\mathbb{P}_{G|\Theta}(g|\theta) \cdot \mathbb{P}_\Theta(\theta)}{\sum_{\theta \in \Theta} \mathbb{P}_{G|\Theta}(g|\theta) \cdot \mathbb{P}_\Theta(\theta)} \\
 &= \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \cdot \mathbb{P}_\Theta(\theta)}{\sum_{\theta \in \Theta_g} \mathbb{P}_\Theta(\theta)} \\
 &= \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{E}_\Theta \left[\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g \right] \\
 &\triangleq V_{\mathbb{P}_\Theta}.
 \end{aligned}$$

This upper bound can be achieved when $\mathbb{P}_\Theta$ satisfies

$$\forall \theta' \in \Theta', \exists g \in \mathbf{G} : g \in \arg \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g) \text{ and } g \in \arg \sup_{g \in \mathbf{G}} \mathbb{P}_G(g), \tag{A5}$$

where $\mathbb{P}_{\Theta'|G}(\theta'|g) = \mathbb{E}[\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g]$, $\mathbb{P}_G(g) = \sum_{\theta \in \Theta_g} \mathbb{P}_\Theta(\theta)$.

Let $\mathbb{P}_\Theta^+ \in \arg \sup_{\mathbb{P}_\Theta} V_{\mathbb{P}_\Theta}$ and $\Theta'(g)$ be the set of released parameters satisfying $g = \arg \sup_{g \in \mathbf{G}} \mathbb{E}_\Theta[\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g]$ under $\mathbb{P}_\Theta^+$. We can construct a prior $\mathbb{P}_\Theta^*$ as follows: for any $g \in \mathbf{G}$, select a distribution parameter $\theta_g \in \Theta$ such that $\theta_g \in \arg \sup_{\theta \in \Theta_g} \sum_{\theta' \in \Theta'(g)} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta)$ (Let $\arg \sup_{\theta \in \Theta_g} \sum_{\theta' \in \Theta'(g)} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) = \Theta_g$ when $\Theta'(g) = \emptyset$), and let $\mathbb{P}_\Theta(\theta_g) = \frac{1}{|\mathbf{G}|}$ and $\mathbb{P}_\Theta(\theta) = 0, \forall \theta \in \Theta_g, \theta \neq \theta_g$.

We then show that replacing $\mathbb{P}_\Theta^+$ by the constructed prior $\mathbb{P}_\Theta^*$ which places all conditional mass on one representative θ_g per secret value, cannot decrease the objective.

$$\begin{aligned}
 V_{\mathbb{P}_\Theta^+} &= \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{E}_\Theta \left[\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g \right] \\
 &= \log \sum_{g \in \mathbf{G}} \sum_{\theta' \in \Theta'(g)} \mathbb{E}_\Theta \left[\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g \right] \\
 &= \log \sum_{g \in \mathbf{G}} \mathbb{E}_\Theta \left[\sum_{\theta' \in \Theta'(g)} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \middle| \theta \in \Theta_g \right] \\
 &\leq \log \sum_{g \in \mathbf{G}} \sup_{\theta \in \Theta_g} \sum_{\theta' \in \Theta'(g)} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \\
 &= V_{\mathbb{P}_\Theta^*}.
 \end{aligned}$$

Since $\mathbb{P}_{\Theta}^+ \in \arg \sup_{\mathbb{P}_{\Theta}} V_{\mathbb{P}_{\Theta}}$, we can get that $\mathbb{P}_{\Theta}^* \in \arg \sup_{\mathbb{P}_{\Theta}} V_{\mathbb{P}_{\Theta}}$. Since $\mathbb{P}_{\Theta}^*$ satisfies Condition (A5) and $\mathbb{P}_{\Theta|G}(\theta|g) \in \{0, 1\}, \forall g \in \mathbf{G}, \theta \in \Theta$, we have

$$\begin{aligned} \Pi_{\mathcal{M}, \mathbf{g}} &= \sup_{\mathbb{P}_{\Theta}} \log \frac{\sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{G\Theta'}(g, \theta')}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)} \\ &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g) \\ &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) \cdot \mathbb{P}_{\Theta}(\theta)}{\sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta}(\theta)} \\ &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{E}_{\Theta} [\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g] \\ &= \sup_{\mathbb{P}_{\Theta}: \mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\ &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g), \end{aligned} \quad (\text{A6})$$

where θ_g satisfies $\mathbb{P}_{\Theta|G}(\theta_g|g) = 1, \forall g \in \mathbf{G}$, under a prior with $\mathbb{P}_{\Theta|G} \in \{0, 1\}$. $\square$

Appendix B. Proof of Proposition 2

Proof. Given a deterministic mechanism $\mathcal{M}$ and a secret mapping $\mathbf{g}$, let $\mathcal{C}_{\mathcal{M}, \mathbf{g}}$ be the total cost of the min-cost flow under the network we construct.

We first prove that $\log(-\mathcal{C}_{\mathcal{M}, \mathbf{g}}) \geq \Pi_{\mathcal{M}, \mathbf{g}}$ as follows. Let

$$\mathbb{P}_{\Theta|G}^* = \arg \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g),$$

and denote θ_g^* such that $\mathbb{P}_{\Theta|G}^*(\theta_g^*|g) = 1$. For any deterministic mechanism, we have $\mathbb{P}_{\Theta'|\Theta} \in \{0, 1\}$, and we denote θ'_g such that $\mathbb{P}_{\Theta'|\Theta}(\theta'_g|\theta_g^*) = 1$. From Proposition 1, we can get that $\Pi_{\mathcal{M}, \mathbf{g}} = \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g^*) = \log \sum_{\theta' \in \Theta'} \mathbb{1}\{\theta' = \theta'_g, \exists g \in \mathbf{G}\}$. Denote $\Theta'^{(*)} = \{\theta' | \theta' = \theta'_g, \exists g \in \mathbf{G}\}$, and then we have $\Pi_{\mathcal{M}, \mathbf{g}} = \log |\Theta'^{(*)}|$. For any $\theta' \in \Theta'^{(*)}$, there is at least one secret, denoted as $g_{\theta'}$, such that $\theta' = \theta'_{g_{\theta'}}$. We can construct a feasible network flow where for any $\theta' \in \Theta'^{(*)}$, there is a 1-unit flow in the path $\text{src} - g_{\theta'} - \theta'_{g_{\theta'}} - \theta' - \text{sink}$. Since we have $\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{g_{\theta'}}^*) = \mathbb{P}_{\Theta'|\Theta}(\theta'_{g_{\theta'}}|\theta_{g_{\theta'}}^*) = 1$, we can easily get that the total cost of the constructed flow is $-\log |\Theta'^{(*)}| \geq \mathcal{C}_{\mathcal{M}, \mathbf{g}}$. Therefore, we have $\log(-\mathcal{C}_{\mathcal{M}, \mathbf{g}}) \geq \Pi_{\mathcal{M}, \mathbf{g}}$.

We then prove that $\log(-\mathcal{C}_{\mathcal{M}, \mathbf{g}}) \leq \Pi_{\mathcal{M}, \mathbf{g}}$ as follows. We can construct a min-cost network flow such that the flow of each edge is either 1 or 0 [63]. Denote $\Theta'^{(+)}$ as the set of nodes in Θ' column that this min-cost flow goes through. Therefore, for any $\theta' \in \Theta'^{(+)}$, there is one 1-unit flow goes through it, and we denote this flow path as $\text{src} - g_{\langle \theta' \rangle} - \theta_{g_{\langle \theta' \rangle}}^{g_{\langle \theta' \rangle}} - \theta' - \text{sink}$. Since the capacity of each edge is 1, $\theta_{g_{\langle \theta' \rangle}}^{g_{\langle \theta' \rangle}}$ has a one-on-one correspondence to $g_{\langle \theta' \rangle}$ and θ' . We can construct a $\mathbb{P}_{\Theta|G}^+$ such that $\mathbb{P}_{\Theta|G}^+(\theta_{g_{\langle \theta' \rangle}}^{g_{\langle \theta' \rangle}} | g_{\langle \theta' \rangle}) = 1$, i.e., $\theta_{g_{\langle \theta' \rangle}}^{g_{\langle \theta' \rangle}} = \theta_{g_{\langle \theta' \rangle}}^+$ under $\mathbb{P}_{\Theta|G}^+$. Then we can get that $\log(-\mathcal{C}_{\mathcal{M}, \mathbf{g}}) = \log \sum_{\theta' \in \Theta'^{(*)}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{g_{\langle \theta' \rangle}}^+) = \log \sum_{\theta' \in \Theta'^{(*)}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{g_{\langle \theta' \rangle}}^+) \leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) = \Pi_{\mathcal{M}, \mathbf{g}}$.

Hence, we have $\log(-\mathcal{C}_{\mathcal{M}, \mathbf{g}}) = \Pi_{\mathcal{M}, \mathbf{g}}$. $\square$

Appendix C. Proof of Theorem 1

Proof. We first prove the NP-hardness of the SML computation. Consider a three-set cover decision problem where given a set $\mathcal{U}$, a collection $\mathcal{T}$ of m three size subsets whose union is $\mathcal{U}$, and an integer k , decide whether there are k subsets in $\mathcal{T}$ whose union equals to $\mathcal{U}$. It is well-known that the three-set cover decision problem is NP-complete [64]. We will prove the NP-hardness of SML computation by reducing the three-set cover decision problem to it.

Let $\mathcal{T} = \{T_1, \dots, T_m\}$, where $T_i = \{t_1^{(i)}, t_3^{(i)}, t_3^{(i)}\}, \forall i \in [m]$, and let $\Theta_g = \{\theta_1^{(g)}, \dots, \theta_{|\Theta_g|}^{(g)}\}, \forall g \in \mathbf{G}$. For any $\mathcal{U}, \mathcal{T}$ and k , we can construct the support of distribution parameter Θ , the secret function g , and the mechanism $\mathcal{M}$ such that (i) $\forall g \in \mathbf{G} : |\Theta_g| = m$; (ii) $|\mathbf{G}| = k$, i.e., $\mathbf{G} = \{g_1, \dots, g_k\}$; (iii) $\Theta' = \mathcal{U}$; (iv) $\forall g \in \mathbf{G}, i \in [m] : \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_i^{(g)}) = \frac{1}{3}$ if $\theta' \in T_i$, otherwise, 0.

We show that solving the three-set cover decision problem with sets $\mathcal{U}, \mathcal{T}$ and integer k is equivalent to deciding whether the SML of the mechanism constructed above is equal to $\log(|\mathcal{U}|/3)$ as follows.

If there is a sub-collection $\tilde{\mathcal{T}} = \{T_{j_1}, \dots, T_{j_k}\}$ whose union is $\mathcal{U}$, we have $\theta' \in \bigcup_{i \in [k]} T_{j_i}, \forall \theta' \in \Theta'$. Let $\theta_{g_i} = \theta_{j_i}^{(g_i)}, \forall i \in [k]$ (θ_g is defined above Proposition 1), and we can get that

$$\begin{aligned} \Pi_{\mathcal{M},g} &= \sup_{\mathbb{P}_{\Theta|\mathbf{G}} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\ &\geq \log \sum_{\theta' \in \Theta'} \sup_{i \in [k]} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{j_i}^{(g_i)}) \\ &= \log \sum_{\theta' \in \Theta'} \frac{1}{3} \\ &= \log \frac{|\mathcal{U}|}{3}. \end{aligned}$$

Additionally, we have

$$\begin{aligned} \Pi_{\mathcal{M},g} &= \sup_{\mathbb{P}_{\Theta|\mathbf{G}} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\ &\leq \sup_{\mathbb{P}_{\Theta|\mathbf{G}} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \frac{1}{3} \\ &= \log \frac{|\mathcal{U}|}{3}. \end{aligned}$$

Therefore, we have $\Pi_{\mathcal{M},g} = \log \frac{|\mathcal{U}|}{3}$.

If there is no sub-collection of $\mathcal{T}$ with size k that covers $\mathcal{U}$, then $\Pi_{\mathcal{M},g} < \log \frac{|\mathcal{U}|}{3}$, which can be proved by contradiction as follows. Suppose $\Pi_{\mathcal{M},g} = \log \frac{|\mathcal{U}|}{3}$, then we can construct a prior distribution of parameter with $\mathbb{P}_{\Theta|\mathbf{G}} \in \{0,1\}$, such that $\forall \theta' \in \Theta', \exists i \in [k] : \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{g_i}) = \frac{1}{3}$. For any $i \in [k]$, there exists $j_i \in [m]$ such that $\theta_{g_i} = \theta_{j_i}^{(g_i)}$, and we construct a sub-collection of $\mathcal{T}$ as $\tilde{\mathcal{T}} = \{T_{j_1}, \dots, T_{j_k}\}$. Since $\Theta' = \mathcal{U}$ and $\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_{g_i}) = \frac{1}{3}$ if and only if $\theta' \in T_{j_i} (i \in [k])$, we can get that $\forall u \in \mathcal{U}, \exists i \in [k] : u \in T_{j_i}$. Therefore, $\bigcup_{i \in [k]} T_{j_i} = \mathcal{U}$, which contradicts with the assumption that there is no sub-collection of $\mathcal{T}$ with size k that covers $\mathcal{U}$. Since $\Pi_{\mathcal{M},g} \leq \log \frac{|\mathcal{U}|}{3}$, we have $\Pi_{\mathcal{M},g} < \log \frac{|\mathcal{U}|}{3}$.

Hence, we can reduce three-set cover decision problem to SML computation, and therefore, SML computation is NP-hard.

Next, we convert the computation of SML to network flow problems, for which approximation or exact algorithms have been proposed.

Edge cost flow problem. Edge cost flow problem is a type of min cost–max flow problem, where the cost c of an edge satisfies

$$c = \begin{cases} b & f > 0 \\ 0 & f = 0 \end{cases}.$$

f is the flow through the edge and b is a constant. This cost function is different from the most common min cost–max flow problem, where $c = bf$.

Edge cost flow problem is shown to be an NP-hard problem [70] (Problem ND32), and calculating SML can be converted as solving an edge cost flow problem. We illustrate the conversion process in Figure A1. To construct the network, we introduce the source and sink nodes, and create three columns of nodes between them.

The first column is the g column, which contains all potential secret values (g_1 and g_2). The capacity of the edge between the source and the g column node is $|\Theta'|$ ($|\Theta'| = 2$ in Figure A1), and the cost is 0. Setting the each edge capacity as $|\Theta'|$ enables us to freely select the secret g for each released parameter value $\theta' \in \Theta'$.

The second column is θ column, which contains all potential parameter value (θ_1 , θ_2 and θ_3 in Figure A1). There will be an edge between node N_{g_i} and N_{θ_j} if $\theta_j \in \Theta_{g_i}$. The capacity of each edge is $|\Theta'|$ and the cost is 1. This ensures that to achieve the min-cost goal, all flows from a certain secret g go to a same node N_{θ} , which is in line with the SML calculation, i.e., $\mathbb{P}_{\Theta|G} \in \{0, 1\}$.

The third column is the θ' column, which contains all potential released parameter value (θ'_1 and θ'_2 in Figure A1). This column is fully connected with the second column as well as the sink node. To be in line with the privacy calculation, the capacity of the edge between the N_{θ_i} and $N_{\theta'_j}$ columns is set as 1 and the cost is set as $-\mathbb{P}(\theta'_j|\theta_i)$. The capacity of the edge between the sink and the third column is 1 and the cost is 0, which ensures that the flow that goes through a node in θ' column is from a single node in g column.




Figure A1. Given a mechanism $\mathcal{M} = \mathbb{P}_{\Theta'|G}$, the left subfigure shows a policy matrix. For each column j , the red outlined region indicates rows of parameters with secret g maximizing $\mathbb{P}_{\Theta'|G}(\theta'_j|\theta_g)$. The blue cell lies in the row of θ_g . SML calculation can be converted to an edge cost flow problem (**top right**) or a min-cost flow problem with conflict constraints (**bottom right**). Both constructed directed graphs contain three columns of nodes (representing G, Θ, Θ' , respectively) between the source and sink nodes.

Additionally, we fully connect the θ column and the sink, and set the capacity and cost as $|\Theta'|$ and 0, respectively, which are shown in red edges in Figure A1.

The max flow of the network is $|\mathbf{G}| \cdot |\Theta'|$. To achieve the min cost-max flow goal, every node in θ' column has nonzero flow going through it, since only the edge cost between the θ and θ' columns can be negative. Recall that the flow going through a node in the θ' column is from a single node in the g column, and to achieve the min-cost goal, all flows from a node in the g column will go through only one node in the θ column, denoted as $N_{\theta(g)}$. Let $\mathcal{C}_{\mathcal{M},g}$ be the total cost of the min cost-max flow under the network we construct. Similar to the proof of Proposition 2, we can get that

$$\begin{aligned}\mathcal{C}_{\mathcal{M},g} &= \inf_{N_{\theta(g)}: \theta(g) \in \Theta, g \in \mathbf{G}} \sum_{\theta' \in \Theta'} \inf_{g \in \mathbf{G}} \left(1 - \mathbb{P}(\theta' | \theta(g))\right) = |\Theta'| - \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \\ &= |\Theta'| - \exp(\Pi_{\mathcal{M},g}),\end{aligned}$$

i.e., $\log(|\Theta'| - \mathcal{C}_{\mathcal{M},g}) = \Pi_{\mathcal{M},g}$. Therefore, we can convert the calculation of SML as the edge cost flow problem.

Ref. [65] provides an approximation algorithm for the edge cost flow problem, with approximation ratio $1 + \rho$ ($\rho > 0$) and running time polynomial in m and $1/\rho$, where m represents the number of edges in the network. Therefore, there exists an approximation algorithm for SML calculation on $|\Theta'| - \exp(\Pi_{\mathcal{M},g})$ with an approximation ratio $1 + \rho$ ($\rho > 0$) and running time polynomial in $|\Theta| \cdot |\Theta'|$ and $1/\rho$.

Min-cost flow problem with conflict constraints. We can also convert the calculation process as the min-cost flow problem with conflict constraints, which is also proven to be NP-hard [71]. The difference of the constructed network with the previous one is that (1) the edge cost between the g column and the θ column is 0. (2) A conflict constraint applies to edges between the g column and the θ column: each g -node can only send flow to one θ -node. (3) There are no red edges between the θ column and the sink. Similarly, we can also conclude that the calculation of SML can be converted as the min-cost flow problem with conflict constraints. Ref. [71] provides an exact solution (with running time exponential to the number of edges) for the min-cost flow problem with conflict constraints. $\square$

Appendix D. Proof of Theorem 2

We first provide Lemma A1, which demonstrates how to compute the statistic maximal leakage of an adaptive composition of mechanisms.

Lemma A1. Suppose there are m adaptive data release mechanisms, and each mechanism $\mathcal{M}_i$ is defined by $\mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \Theta^{(k)}\}_{k \in [i-1]}}$, which is adaptively designed given the original data distribution, previous mechanisms $\{\mathcal{M}_k\}_{k \in [i-1]}$, and previous outputs $\{\Theta^{(k)}\}_{k \in [i-1]}$. Let θ' be the sequence of the released distribution parameter, i.e., $\theta' \in \Theta^{(1)} \times \Theta^{(2)} \times \dots \times \Theta^{(m)}$. The statistic maximal leakage of the adaptively composed mechanism $\mathcal{M} = \mathcal{M}_1 \circ \mathcal{M}_2 \circ \dots \circ \mathcal{M}_m$ is

$$\Pi_{\mathcal{M},g} = \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta'} \sup_{g \in \mathbf{G}} \prod_{i \in [m]} \mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \Theta^{(k)}\}_{k \in [i-1]}} \left(\theta^{(i)} | \theta_g, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]} \right),$$

where θ_g satisfies $\mathbb{P}_{\Theta|G}(\theta_g | g) = 1, \forall g \in \mathbf{G}$, under a prior with $\mathbb{P}_{\Theta|G} \in \{0,1\}$.

Proof. Based on Equation (A11), we have

$$\begin{aligned}\Pi_{\mathcal{M},g} &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\theta^{(1)}, \dots, \theta^{(m)}|G}(\theta'|g) \\ &= \sup_{\mathbb{P}_{\Theta}} \log \sum_{\theta'} \sup_{g \in \mathbf{G}} \sum_{\theta \in \Theta} \mathbb{P}_{\theta^{(1)}, \dots, \theta^{(m)}|G\Theta}(\theta'|g, \theta) \cdot \mathbb{P}_{\Theta|G}(\theta|g) \\ &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta'} \sup_{g \in \mathbf{G}} \prod_{i \in [m]} \mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]}} \left(\theta^{(i)} | \theta_g, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]} \right).\end{aligned}$$

□

We then prove Theorem 2 as follows.

Proof. From Lemma A1, we know that the calculation of the statistic maximal leakage of adaptively composed mechanism is similar to the process illustrated in Figure 1, except that the policy matrix is changed as $\prod_{i \in [m]} \mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]}}$ and the released set of parameters is changed as $\Theta^{(1)} \times \Theta^{(2)} \times \dots \times \Theta^{(m)}$. Using Lemma A1, we get that

$$\begin{aligned}\Pi_{\mathcal{M},g} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta'} \sup_{g \in \mathbf{G}} \prod_{i \in [m]} \mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]}} \left(\theta^{(i)} | \theta_g, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]} \right) \\ &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta'} \prod_{i \in [m]} \sup_{g \in \mathbf{G}} \mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]}} \left(\theta^{(i)} | \theta_g, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]} \right) \\ &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \prod_{i \in [m]} \sum_{\theta^{(i)} \in \Theta^{(i)}} \sup_{g \in \mathbf{G}, [\theta^{(1)}, \dots, \theta^{(i-1)}] \in \Theta^{(1)} \times \dots \times \Theta^{(i-1)}} \mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]}} \left(\theta^{(i)} | \theta_g, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]} \right) \\ &\leq \sum_{i \in [m]} \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta^{(i)} \in \Theta^{(i)}} \sup_{g \in \mathbf{G}, [\theta^{(1)}, \dots, \theta^{(i-1)}] \in \Theta^{(1)} \times \dots \times \Theta^{(i-1)}} \mathbb{P}_{\theta^{(i)}|\Theta, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]}} \left(\theta^{(i)} | \theta_g, \{\mathcal{M}_k, \theta^{(k)}\}_{k \in [i-1]} \right) \\ &= \sum_{i \in [m]} \Pi_{\mathcal{M}_i, g}.\end{aligned}$$

□

Appendix E. Proof of Theorem 3

Proof. Based on Proposition 1, we can get that

$$\begin{aligned}
\Pi_{\widetilde{\mathcal{M}} \circ \mathcal{M}, \mathfrak{g}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \sum_{\theta'' \in \Theta''} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta''|\Theta}(\theta''|\theta_g) \\
&= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \sum_{\theta'' \in \Theta''} \sup_{g \in \mathbf{G}} \sum_{\theta' \in \Theta'} \mathbb{P}_{\Theta''|\Theta'}(\theta''|\theta') \cdot \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\
&\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \sum_{\theta'' \in \Theta''} \sum_{\theta' \in \Theta'} \mathbb{P}_{\Theta''|\Theta'}(\theta''|\theta') \cdot \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\
&= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \sum_{\theta' \in \Theta'} \sum_{\theta'' \in \Theta''} \mathbb{P}_{\Theta''|\Theta'}(\theta''|\theta') \cdot \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\
&= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\
&= \Pi_{\mathcal{M}, \mathfrak{g}}.
\end{aligned}$$

□

Appendix F. Proofs of Theorem 4, Theorem 5, and Corollary 1

Proof. (Privacy and Distortion of Randomized Response.) We first analyze the statistic maximal leakage of the Randomized Response with hyperparameter ϵ . Based on Proposition 1, we can get that

$$\begin{aligned}
\Pi_{\mathcal{M}_{\text{RR}}, \mathfrak{g}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\
&= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\sum_{\theta' \in \{\theta_g\}_{g \in \mathbf{G}}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) + \sum_{\theta' \in \Theta' \setminus \{\theta_g\}_{g \in \mathbf{G}}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \right) \\
&= \log \left(|\mathbf{G}| \cdot \frac{e^\epsilon}{|\Theta'| + e^\epsilon - 1} + (|\Theta'| - |\mathbf{G}|) \cdot \frac{1}{|\Theta'| + e^\epsilon - 1} \right).
\end{aligned}$$

Denote $r = \frac{e^\epsilon - 1}{|\Theta'|}$. Since $|\mathbf{G}| = s$ and $|\Theta'| = |\Theta| = \binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^* - 1}$, we have

$$\Pi_{\mathcal{M}_{\text{RR}}, \mathfrak{g}} = \log \left(\frac{\binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^* - 1} + s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^* - 1} + e^\epsilon - 1} \right) = \log \left(\frac{1 + sr}{1 + r} \right).$$

We then analyze the distortion of randomized response. For the categorical distribution with parameter θ , denote m_{θ_i} as the probability mass of the i -th category, where $i \in [\hat{d}_0^*]$. From Lemma 1, we have $\Delta = \sup_{\theta} \mathbb{E}_{\Theta'}[D_{\text{TV}}(f_{\theta} \| f_{\Theta'})]$. Let $\theta^* = \arg \sup_{\theta} \mathbb{E}_{\Theta'}[D_{\text{TV}}(f_{\theta} \| f_{\Theta'})]$, we first prove by contradiction that θ^* satisfies that $\exists i \in [\hat{d}_0^*] : m_{\theta_i^*} = 1$ as follows.

Suppose $\max_{i \in [\hat{d}_0^*]} m_{\theta_i^*} = m_{\theta_j^*} < 1$, then we have $m_{\theta_j^*} \leq 1 - \frac{1}{\tau}$ and there exists another category $k \neq j$ such that $m_{\theta_j^*} \geq m_{\theta_k^*} \geq \frac{1}{\tau}$. We can construct another categorical distribution parameter $\tilde{\theta}$ where $m_{\tilde{\theta}_j} = m_{\theta_j^*} + \frac{1}{\tau}$, $m_{\tilde{\theta}_k} = m_{\theta_k^*} - \frac{1}{\tau}$, and $m_{\tilde{\theta}_i} = m_{\theta_i^*}, \forall i \in [\tau] \setminus \{j, k\}$. We can get that

$$D_{\text{TV}}(f_{\theta^*} \| f_{\Theta'}) = \begin{cases} D_{\text{TV}}(X_{\tilde{\theta}} \| f_{\Theta'}) + \frac{2}{\tau}, & m_{\theta_j^*} \geq m_{\tilde{\theta}_j}, m_{\theta_k^*} \leq m_{\tilde{\theta}_k}, \\ D_{\text{TV}}(f_{\tilde{\theta}} \| f_{\Theta'}) - \frac{2}{\tau}, & m_{\theta_j^*} < m_{\tilde{\theta}_j}, m_{\theta_k^*} > m_{\tilde{\theta}_k}, \\ D_{\text{TV}}(f_{\tilde{\theta}} \| f_{\Theta'}), & \text{otherwise.} \end{cases}$$

For every $\theta' \in \Theta'$ satisfying $m_{\theta'_j} \geq m_{\tilde{\theta}_j}, m_{\theta'_k} \leq m_{\tilde{\theta}_k}$, we can construct a unique θ'' as $m_{\theta''_j} = m_{\theta'_j}, m_{\theta''_k} = m_{\theta'_k}$, and $m_{\theta''_i} = m_{\theta'_i}, \forall i \in [\tau] \setminus \{j, k\}$. Since $m_{\theta'_j} = m_{\theta'_k} \leq m_{\tilde{\theta}_k} < m_{\theta'_j}$ and $m_{\theta'_k} = m_{\theta'_j} \geq m_{\tilde{\theta}_j} > m_{\theta'_k}$, we have $D_{TV}(f_{\theta^*} \| f_{\theta''}) = D_{TV}(f_{\tilde{\theta}} \| f_{\theta''}) - \frac{2}{\tau}$. Therefore, we can get that $\mathbb{E}_{\Theta'}[D_{TV}(f_{\theta^*} \| f_{\Theta'})] \leq \mathbb{E}_{\Theta'}[D_{TV}(f_{\tilde{\theta}} \| f_{\Theta'})]$, which contradicts with the definition of $\theta^* : \theta^* = \arg \sup_{\theta} \mathbb{E}_{\Theta'}[D_{TV}(f_{\theta} \| f_{\Theta'})]$. Therefore, we can get that $\exists i \in [\hat{d}_0^*] : m_{\theta_i^*} = 1$. Without loss of generality, let $m_{\theta_1^*} = 1$.

We can calculate the distortion as

$$\begin{aligned} \Delta_{\mathcal{M}_{RR}} &= \mathbb{E}_{\theta' \sim \mathcal{M}_{RR}(\theta^*)}[D_{TV}(v_{X_{\theta^*}} \| f_{\Theta'})] \\ &= \frac{|\Theta'|}{|\Theta'| + e^{\epsilon} - 1} \cdot \sum_{i \in [\tau]} \frac{i}{\tau} \cdot \mathbb{P}\left(m_{\theta'_1} = 1 - \frac{i}{\tau}\right) \\ &= \frac{|\Theta'|}{|\Theta'| + e^{\epsilon} - 1} \cdot \sum_{i \in [\tau]} \frac{i}{\tau} \cdot \frac{\binom{i + \hat{d}_0^* - 2}{\hat{d}_0^* - 2}}{|\Theta'|} \\ &= \frac{1}{\tau \binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^* - 1} + \tau(e^{\epsilon} - 1)} \cdot \sum_{i \in [\tau]} i \cdot \binom{i + \hat{d}_0^* - 2}{\hat{d}_0^* - 2} \\ &= \frac{1}{\tau \binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^* - 1} + \tau(e^{\epsilon} - 1)} \cdot \sum_{i \in [\tau]} (\hat{d}_0^* - 1) \cdot \binom{i + \hat{d}_0^* - 2}{\hat{d}_0^* - 1} \\ &= \frac{(\hat{d}_0^* - 1) \cdot \binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^*}}{\tau \binom{\tau + \hat{d}_0^* - 1}{\hat{d}_0^* - 1} + \tau(e^{\epsilon} - 1)} \\ &= \frac{\hat{d}_0^* - 1}{\hat{d}_0^*(1 + r)}. \end{aligned}$$

□

Proof. (Privacy and Distortion of Quantization Mechanism.) For the privacy of the quantization mechanism, we can get that

$$\begin{aligned} \Pi_{\mathcal{M}_{QM}, \mathbf{g}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta' | \theta_g) \\ &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1, \dots, \lceil \frac{s}{I} \rceil - 1\}} \sum_{\theta' \in \Theta'_{R(k)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta' | \theta_g) \\ &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1, \dots, \lceil \frac{s}{I} \rceil - 1\}} \sum_{\theta' \in \Theta'_{R(k)}} \frac{1}{|\Theta'_{R(k)}|} \\ &= \log \left\lceil \frac{s}{I} \right\rceil. \end{aligned}$$

For the distortion analysis, we consider the case where the secret is the fraction of a category, i.e., $\mathbf{g}(\theta) = n_{\psi_i}/n = m_{\theta_i}, i \in [\hat{d}_0^*]$. Without loss of generality, let $\mathbf{g}(\theta) = m_{\theta_1}$. The number of possible secret values is $s = \tau + 1$ as $s \in \{0, 1/\tau, 2/\tau, \dots, 1\}$, and let $g_l = \frac{l-1}{\tau}, \forall l \in [\tau + 1]$. Let $\theta^* = \arg \sup_{\theta} \mathbb{E}_{\Theta'}[D_{TV}(f_{\theta} \| f_{\Theta'})]$, with the same distortion analysis of randomized response, we can get that θ^* satisfies that $\exists j \in [\hat{d}_0^*] \setminus \{1\} : m_{\theta_j^*} = 1 - m_{\theta_1^*}$. Without loss of generality, let $m_{\theta_2^*} = 1 - m_{\theta_1^*}$. Let $\tilde{\theta}$ be the distribution parameter satisfying $m_{\tilde{\theta}_1} = g_{kl+j}$, where $k \in \{0, 1, \dots, \lceil \frac{s}{I} \rceil - 1\}, j \in [I]$, and $m_{\tilde{\theta}_2} = 1 - m_{\tilde{\theta}_1}$. We can get that

$$\begin{aligned}
& \mathbb{E}_{\Theta' = \mathcal{M}_{\text{QM}}(\theta)} [D_{\text{TV}}(f_{\theta} \| f_{\Theta'})] \\
&= \frac{1}{2} \left[\frac{|\mathbf{R}(k) - kI - j|}{\tau} + \sum_{i \in [\tau+1] \setminus [\mathbf{R}(k)]} \frac{|i - kI - j| + i - \mathbf{R}(k) - 1}{\tau} \mathbb{P}\left(m_{\theta'_2} = 1 - \frac{i-1}{\tau}\right) \right] \\
&\leq \frac{1}{2} \left[\frac{|\mathbf{R}(k) - kI - 1|}{\tau} + \sum_{i \in [\tau+1] \setminus [\mathbf{R}(k)]} \frac{|i - kI - 1| + i - \mathbf{R}(k) - 1}{\tau} \mathbb{P}\left(m_{\theta'_2} = 1 - \frac{i-1}{\tau}\right) \right] \\
&= \frac{|\mathbf{R}(k) - kI - 1|}{2\tau} + \sum_{i \in [\tau+1] \setminus [\mathbf{R}(k)]} \frac{|i - kI - 1| + i - \mathbf{R}(k) - 1}{2\tau} \frac{\binom{i - \mathbf{R}(k) + \hat{d}_0^* - 3}{\hat{d}_0^* - 3}}{\binom{\tau + 1 - \mathbf{R}(k) + \hat{d}_0^* - 2}{\hat{d}_0^* - 2}} \\
&\leq \frac{|\mathbf{R}(0) - 1|}{2\tau} + \sum_{i \in [\tau+1] \setminus [\mathbf{R}(0)]} \frac{|i - 1| + i - \mathbf{R}(0) - 1}{2\tau} \frac{\binom{i - \mathbf{R}(0) + \hat{d}_0^* - 3}{\hat{d}_0^* - 3}}{\binom{\tau + 1 - \mathbf{R}(0) + \hat{d}_0^* - 2}{\hat{d}_0^* - 2}} \\
&= \frac{1}{2\tau} \left\lfloor \frac{I}{2} \right\rfloor + \frac{1}{2\tau \binom{\tau - \lfloor \frac{I}{2} \rfloor + \hat{d}_0^* - 2}} \sum_{l \in [n - \lfloor \frac{I}{2} \rfloor] \cup \{0\}} \left(\left\lfloor \frac{I}{2} \right\rfloor + 2l - 1 \right) \binom{l + \hat{d}_0^* - 3}{\hat{d}_0^* - 3} \\
&= \frac{1}{2\tau} \left\lfloor \frac{I}{2} \right\rfloor + \frac{1}{2\tau \binom{\tau - \lfloor \frac{I}{2} \rfloor + \hat{d}_0^* - 2}} \cdot \left[\left\lfloor \frac{I}{2} \right\rfloor \cdot \binom{\tau - \lfloor \frac{I}{2} \rfloor + \hat{d}_0^* - 2}{\hat{d}_0^* - 2} + \sum_{l \in [n - \lfloor \frac{I}{2} \rfloor]} \binom{\hat{d}_0^* - 2}{\hat{d}_0^* - 2} \binom{l + \hat{d}_0^* - 3}{\hat{d}_0^* - 2} \right] \\
&= \frac{1}{\tau} \cdot \left\lfloor \frac{I}{2} \right\rfloor + \frac{(\hat{d}_0^* - 2) \binom{\tau - \lfloor \frac{I}{2} \rfloor + \hat{d}_0^* - 2}{\hat{d}_0^* - 1}}{2\tau \binom{\tau - \lfloor \frac{I}{2} \rfloor + \hat{d}_0^* - 2}} \\
&= \frac{1}{2} + \frac{\hat{d}_0^* \lfloor \frac{I}{2} \rfloor - \tau}{2\tau (\hat{d}_0^* - 1)},
\end{aligned}$$

where $\leq$ achieves '=' when $j = 1$, and $\leq$ achieves '=' when $k = 0$. Therefore, we can get that

$$\Delta_{\mathcal{M}_{\text{QM}}} = \mathbb{E}_{\Theta' = \mathcal{M}_{\text{QM}}(\theta^*)} [D_{\text{TV}}(f_{\theta^*} \| f_{\Theta'})] = \frac{1}{2} + \frac{\hat{d}_0^* \lfloor \frac{I}{2} \rfloor - \tau}{2\tau (\hat{d}_0^* - 1)}.$$

□

Proof. (Mechanism Comparison.)

Under the case where the secret is the fraction of a category, we compare the distortion of quantization and randomized response when they achieve the same non-trivial statistic maximal leakage, i.e., $\Pi_{\mathcal{M}_{\text{QM}}, \mathfrak{g}} = \Pi_{\mathcal{M}_{\text{RR}}, \mathfrak{g}} < \log s$. Under the same statistic maximal leakage, we have $\log \frac{1+s}{1+r} = \log \lceil \frac{s}{I} \rceil$, and therefore,

$$r = \frac{s-1}{s - \lceil \frac{s}{I} \rceil} - 1 \leq \frac{s-1}{s - \frac{s}{I} - 1} - 1.$$

For the distortion of randomized response, we have

$$\Delta_{\mathcal{M}_{\text{RR}}} = \frac{\hat{d}_0^* - 1}{\hat{d}_0^* (1+r)} \geq \frac{(Is - s - I)(\hat{d}_0^* - 1)}{I\hat{d}_0^* (s-1)}.$$

Since $s = \tau + 1$, we can get that

$$\begin{aligned}\frac{\Delta_{\mathcal{M}_{\text{RR}}}}{\Delta_{\mathcal{M}_{\text{QM}}}} &\geq \frac{(Is - s - I)(\hat{d}_0^* - 1)}{I\hat{d}_0^*(s - 1)} / \left(\frac{1}{2} + \frac{\hat{d}_0^*I - 2\tau}{4\tau(\hat{d}_0^* - 1)} \right) \\ &= \frac{4(\hat{d}_0^* - 1)^2(I\tau - \tau - 1)}{\left[\hat{d}_0^*I + 2\tau(\hat{d}_0^* - 2) \right] I\hat{d}_0^*}.\end{aligned}$$

When $\lim_{\tau \rightarrow \infty} I = \infty$, since $I \leq s = \tau + 1$, we have

$$\begin{aligned}\lim_{\tau \rightarrow \infty} \frac{\Delta_{\mathcal{M}_{\text{RR}}}}{\Delta_{\mathcal{M}_{\text{QM}}}} &\geq \frac{4\tau(\hat{d}_0^* - 1)^2}{\left[\hat{d}_0^*I + 2\tau(\hat{d}_0^* - 2) \right] \hat{d}_0^*} \\ &\geq \frac{4(\hat{d}_0^* - 1)^2}{(3\hat{d}_0^* - 4)\hat{d}_0^*} \\ &= 1 + \frac{(\hat{d}_0^* - 2)^2}{3(\hat{d}_0^*)^2 - 4\hat{d}_0^*} \\ &\geq 1.\end{aligned}$$

When $\lim_{\tau \rightarrow \infty} I < \infty$, to achieve non-trivial statistic maximal leakage, I should satisfy $\log \left\lceil \frac{s}{I} \right\rceil < \log s$, i.e., $I > 1$. We have

$$\begin{aligned}\lim_{\tau \rightarrow \infty} \frac{\Delta_{\mathcal{M}_{\text{RR}}}}{\Delta_{\mathcal{M}_{\text{QM}}}} &\geq \frac{4\tau(\hat{d}_0^* - 1)^2}{2\tau(\hat{d}_0^* - 2)\hat{d}_0^*} \cdot \frac{I - 1}{I} \\ &\geq \frac{(\hat{d}_0^* - 1)^2}{\hat{d}_0^*(\hat{d}_0^* - 2)} \\ &= 1 + \frac{1}{\hat{d}_0^*(\hat{d}_0^* - 2)} \\ &> 1.\end{aligned}$$

Hence, we can get that when secret is the fraction of a category, for any non-trivial privacy budget $T < \log s = \log(\tau + 1)$, when $\Pi_{\mathcal{M}_{\text{QM}}, \mathfrak{g}} = \Pi_{\mathcal{M}_{\text{RR}}, \mathfrak{g}} \leq T$, we have

$$\lim_{\tau \rightarrow \infty} \frac{\Delta_{\mathcal{M}_{\text{RR}}}}{\Delta_{\mathcal{M}_{\text{QM}}}} \geq 1.$$

□

Appendix G. Proof of Proposition 3

Proof. Denote $\Theta'^{(*)}$ as the set that contains all the distribution parameters whose support is or is within $\hat{\Gamma}^* \cup \Gamma^*$, and $\Theta'^{(0)}$ as the set of parameters whose support is or is within $\hat{\Gamma}^*$. Let $\beta(\theta')$ be the number of categories in $\Gamma^* \setminus \hat{\Gamma}^*$ that have nonzero probability mass in θ' , i.e., $\beta(\theta') = |\text{Supp}(\theta') \setminus \hat{\Gamma}^*|$. We analyze the robustness of RR and QM to support mismatch correspondingly as follows.

(1) *Randomized Response.*

For randomized response with hyperparameter ϵ , from Corollary 1, we know that

$$\Pi_{\mathcal{M}_{RR},g}^* = \log \frac{\binom{\tau+d^*-1}{d^*-1} + s(e^\epsilon - 1)}{\binom{\tau+d^*-1}{d^*-1} + e^\epsilon - 1} \geq \log \frac{\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1} + s(e^\epsilon - 1)}{\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1} + e^\epsilon - 1}.$$

For $\Pi_{\mathcal{M}_{RR},g}$, we can get that

$$\begin{aligned} \Pi_{\mathcal{M}_{RR},g} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'^{(*)}} \sup_{g \in G} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\ &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\sum_{\theta' \in \{\theta_g\}} \sup_{g \in G} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) + \sum_{\theta' \in \Theta'^{(*)} \setminus \{\theta_g\}} \sup_{g \in G} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \right) \\ &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\sum_{\theta' \in \{\theta_g\}} \frac{e^\epsilon}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*+\beta(\theta')-1}{\hat{d}_0^*+\hat{d}_1^*+\beta(\theta')-1} + e^\epsilon - 1} \right. \\ &\quad \left. + \sum_{\theta' \in \Theta'^{(*)} \setminus \{\theta_g\}} \frac{1}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*+\beta(\theta')-1}{\hat{d}_0^*+\hat{d}_1^*+\beta(\theta')-1} + e^\epsilon - 1} \right) \\ &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*-1}{\hat{d}_0^*+\hat{d}_1^*-1} + e^\epsilon - 1} + \sum_{\theta' \in \Theta'^{(*)}} \frac{1}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*+\beta(\theta')-1}{\hat{d}_0^*+\hat{d}_1^*+\beta(\theta')-1} + e^\epsilon - 1} \right) \\ &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*-1}{\hat{d}_0^*+\hat{d}_1^*-1} + e^\epsilon - 1} + \sum_{i \in \{0\} \cup [d^*-\hat{d}_0^*]} \frac{\binom{d^*-\hat{d}_0^*}{i} \binom{\tau-i+\hat{d}_0^*+\hat{d}_1^*+i-1}{\hat{d}_0^*+\hat{d}_1^*+i-1}}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*+i-1}{\hat{d}_0^*+\hat{d}_1^*+i-1} + e^\epsilon - 1} \right) \\ &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*-1}{\hat{d}_0^*+\hat{d}_1^*-1} + e^\epsilon - 1} \right. \\ &\quad \left. + \frac{\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1}}{\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1} + e^\epsilon - 1} \sum_{i \in \{0\} \cup [d^*-\hat{d}_0^*]} \frac{\binom{d^*-\hat{d}_0^*}{i} \binom{\tau-i+\hat{d}_0^*+\hat{d}_1^*+i-1}{\hat{d}_0^*+\hat{d}_1^*+i-1}}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*+i-1}{\hat{d}_0^*+\hat{d}_1^*+i-1}} \right) \end{aligned}$$

$$\begin{aligned}
 &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \right. \\
 &\quad \left. + \frac{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \sum_{i \in \{0\} \cup [d^* - \hat{d}_0^*]} \frac{\tau! (\tau + \hat{d}_0^* + \hat{d}_1^* - 1)! \binom{d^* - \hat{d}_0^*}{i}}{(\tau - i)! (\tau + i + \hat{d}_0^* + \hat{d}_1^* - 1)!} \right) \\
 &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \right. \\
 &\quad \left. + \frac{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \sum_{i \in \{0\} \cup [d^* - \hat{d}_0^*]} \left(\frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^i \binom{d^* - \hat{d}_0^*}{i} \right) \\
 &= \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right).
 \end{aligned}$$

When $e^\epsilon - 1 \leq \binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} / s$, we can get that

$$\begin{aligned}
 &\Pi_{\mathcal{M}_{RR}, \mathfrak{g}} - \Pi_{\mathcal{M}_{RR}, \mathfrak{g}}^* \\
 &= \log \left(\frac{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1}{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + s(e^\epsilon - 1)} \frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \right. \\
 &\quad \left. + \frac{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1}{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + s(e^\epsilon - 1)} \frac{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right) \\
 &\leq \log \left(\frac{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1}{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + s(e^\epsilon - 1)} \frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right) \\
 &\leq \log \left(1 + \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right) \\
 &< \log 3 \cdot (d^* - \hat{d}_0^*),
 \end{aligned}$$

which indicates that randomized response is log3-robust when ϵ satisfies $e^\epsilon - 1 \leq \binom{\tau + \hat{d}_1^* - 1}{\hat{d}_1^* - 1} / s$.

(2) Quantization Mechanism.

For quantization mechanism with interval length I , from Corollary 1, we know that

$$\Pi_{\mathcal{M}_{QM}, \mathfrak{g}}^* = \log \left\lceil \frac{s}{I} \right\rceil.$$

For convenience, let $g_{k,I}$ represent $g_{(k+\frac{1}{2})I}$, the output secret value of the k -th interval. For $\Pi_{\mathcal{M}_{QM},g}$, we have

$$\begin{aligned}
 \Pi_{\mathcal{M}_{QM},g} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta^{(*)}} \sup_{g \in G} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \\
 &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \sum_{\theta' \in \Theta_{R(k)}^{(*)}} \sup_{g \in G} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \\
 &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \sum_{\theta' \in \Theta_{R(k)}^{(*)}} \frac{1}{\binom{\tau(1-g_{k,I}) + \hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 2}{\hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 2}} \\
 &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \sum_{i \in \{0,1,\dots, d^* - \hat{d}_0^*\}} \frac{\binom{d^* - \hat{d}_0^*}{i} \binom{\tau(1-g_{k,I}) - i + \hat{d}_0^* + \hat{d}_1^* + i - 2}{\hat{d}_0^* + \hat{d}_1^* + i - 2}}{\binom{\tau(1-g_{k,I}) + \hat{d}_0^* + \hat{d}_1^* + i - 2}{\hat{d}_0^* + \hat{d}_1^* + i - 2}} \\
 &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \sum_{i \in \{0,1,\dots, d^* - \hat{d}_0^*\}} \frac{(\tau(1-g_{k,I}))! (\tau(1-g_{k,I}) + \hat{d}_0^* + \hat{d}_1^* - 2)!}{(\tau(1-g_{k,I}) - i)! (\tau(1-g_{k,I}) + i + \hat{d}_0^* + \hat{d}_1^* - 2)!} \binom{d^* - \hat{d}_0^*}{i} \\
 &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \sum_{i \in \{0,1,\dots, d^* - \hat{d}_0^*\}} \left(\frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^i \binom{d^* - \hat{d}_0^*}{i} \\
 &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \\
 &= \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right].
 \end{aligned}$$

Therefore, we have

$$\begin{aligned}
 \Pi_{\mathcal{M}_{QM},g} - \Pi_{\mathcal{M}_{QM},g}^* &\leq \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right] - \log \left\lceil \frac{s}{I} \right\rceil \\
 &= (d^* - \hat{d}_0^*) \log \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right) \\
 &< d^* - \hat{d}_0^*,
 \end{aligned}$$

which indicates that a quantization mechanism with any interval length I is 1-robust when the secret is the fraction of a category. $\square$

Appendix H. Privacy–Utility Tradeoff Analysis of RR and QM with Misspecified of Feasible Attribute Combinations

In this section, we focus on the case where there is misspecification of feasible attribute combinations, i.e., $\hat{\Gamma}^* \neq \Gamma^*$, and analyze the upper and lower bounds of privacy and distortion for both RR and QM in Appendices H.1, H.2 and H.4, as well as the tightness of the privacy bounds in Appendix H.3.

Similar to Appendix G, we denote $\Theta^{(*)}$ as the set that contains all the distribution parameters whose support is or is not within $\hat{\Gamma}^* \cup \Gamma^*$, and $\Theta^{(0)}$ as the set of parameters whose support is or is not within $\hat{\Gamma}^*$. Let $\beta(\theta')$ be the number of categories in $\Gamma^* \setminus \hat{\Gamma}^*$ that have nonzero probability mass in θ' , i.e., $\beta(\theta') = |\text{Supp}(\theta') \setminus \hat{\Gamma}^*|$.

Appendix H.1. Privacy Upper Bounds

Proposition A1. For any secret function $\mathbf{g}$, the SML upper bound of RR is

$$\Pi_{\mathcal{M}_{RR}, \mathbf{g}} \leq \min \left\{ \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}^* + \hat{d}_1^* - 1}{\hat{d}^* + \hat{d}_1^* - 1} \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*}}{\binom{\tau + \hat{d}^* + \hat{d}_1^* - 1}{\hat{d}^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \right), \right. \\ \left. \log \left(1 + s - s \cdot \frac{\left(\frac{\hat{d}_0^* + \hat{d}_1^*}{\tau + \hat{d}^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*}}{1 + (e^\epsilon - 1) / \binom{\tau + \hat{d}^* + \hat{d}_1^* - 1}{\hat{d}^* + \hat{d}_1^* - 1}} \right) \right\}.$$

Specifically, when $e^\epsilon - 1 \leq \binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} / s$, the upper bound can be written as

$$\Pi_{\mathcal{M}_{RR}, \mathbf{g}} < \min \left\{ \log \left(\frac{s}{1 + s} + \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right), \right. \\ \left. \log \left(1 + s - s \cdot \frac{\left(\frac{\hat{d}_0^* + \hat{d}_1^*}{\tau + \hat{d}^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*}}{1 + 1/s} \right) \right\}.$$

For any secret function $\mathbf{g}$, the SML upper bound of QM is

$$\Pi_{\mathcal{M}_{QM}, \mathbf{g}} \leq \min \left\{ \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right], \right. \\ \left. \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + \hat{d}^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right) \right\}.$$

Appendix H.1.1. Randomized Response

Proof. $\Pi_{\mathcal{M}_{RR}, \mathbf{g}}$ can be upper bounded by

$$\Pi_{\mathcal{M}_{RR}, \mathbf{g}} = \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta^{(*)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\ \leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\sup_{g \in \mathbf{G}} \sum_{\theta' \in \Theta^{(0)}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) + \sum_{g \in \mathbf{G}} \sum_{\theta' \in \Theta^{(*)} \setminus \Theta^{(0)}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \right) \\ \leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(1 + \sum_{g \in \mathbf{G}} \left(1 - \frac{|\Theta^{(0)}|}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* + \beta(\theta_g) - 1}{\hat{d}_0^* + \hat{d}_1^* + \beta(\theta_g) - 1} + e^\epsilon - 1 \right) \right) \\ \leq \log \left(1 + s \cdot \frac{\binom{\tau + \hat{d}^* + \hat{d}_1^* - 1}{\hat{d}^* + \hat{d}_1^* - 1} - \binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1}{\binom{\tau + \hat{d}^* + \hat{d}_1^* - 1}{\hat{d}^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \right).$$

Since we have

$$\frac{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*-1}{\hat{d}_0^*+\hat{d}_1^*-1}}{\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1}} = \frac{(\tau+\hat{d}_0^*+\hat{d}_1^*-1)!(d^*+\hat{d}_1^*-1)!}{(\tau+d^*+\hat{d}_1^*-1)!(\hat{d}_0^*+\hat{d}_1^*-1)!} \geq \left(\frac{\hat{d}_0^*+\hat{d}_1^*}{\tau+d^*+\hat{d}_1^*-1}\right)^{d^*-\hat{d}_0^*},$$

we can get that

$$\Pi_{\mathcal{M}_{RR}, \mathfrak{g}} \leq \log \left(1 + s - s \cdot \frac{\left(\frac{\hat{d}_0^*+\hat{d}_1^*}{\tau+d^*+\hat{d}_1^*-1}\right)^{d^*-\hat{d}_0^*}}{1 + (e^\epsilon - 1)/\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1}} \right).$$

Combining with the privacy upper bound provided in Appendix G, we have

$$\begin{aligned} \Pi_{\mathcal{M}_{RR}, \mathfrak{g}} \leq \min & \left\{ \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau+\hat{d}_0^*+\hat{d}_1^*-1}{\hat{d}_0^*+\hat{d}_1^*-1} + e^\epsilon - 1} + \frac{\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1} \left(1 + \frac{\tau}{\tau+\hat{d}_0^*+\hat{d}_1^*-1}\right)^{d^*-\hat{d}_0^*}}{\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1} + e^\epsilon - 1} \right), \right. \\ & \left. \log \left(1 + s - s \cdot \frac{\left(\frac{\hat{d}_0^*+\hat{d}_1^*}{\tau+d^*+\hat{d}_1^*-1}\right)^{d^*-\hat{d}_0^*}}{1 + (e^\epsilon - 1)/\binom{\tau+d^*+\hat{d}_1^*-1}{d^*+\hat{d}_1^*-1}} \right) \right\}. \end{aligned}$$

Specifically, when $e^\epsilon - 1 \leq \binom{\tau+\hat{d}_0^*+\hat{d}_1^*-1}{\hat{d}_0^*+\hat{d}_1^*-1}/s$, we have

$$\begin{aligned} \Pi_{\mathcal{M}_{RR}, \mathfrak{g}} & < \min \left\{ \log \left(\frac{s}{1+s} + \left(1 + \frac{\tau}{\tau+\hat{d}_0^*+\hat{d}_1^*-1}\right)^{d^*-\hat{d}_0^*} \right), \right. \\ & \left. \log \left(1 + s - s \cdot \frac{\left(\frac{\hat{d}_0^*+\hat{d}_1^*}{\tau+d^*+\hat{d}_1^*-1}\right)^{d^*-\hat{d}_0^*}}{1 + 1/s} \right) \right\}. \end{aligned}$$

□

Appendix H.1.2. Quantization Mechanism

Proof. $\Pi_{\mathcal{M}_{QM}, \mathfrak{g}}$ can be upper bounded by

$$\begin{aligned} \Pi_{\mathcal{M}_{QM}, \mathfrak{g}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta^{(*)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \\ &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{k \in \{0,1, \dots, \lceil \frac{s}{I} \rceil - 1\}} \left(\sup_{g \in \mathbf{G}} \sum_{\theta' \in \Theta_{R(k)}^{(0)}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \right. \\ &\quad \left. + \sum_{g \in \mathbf{G}} \sum_{\theta' \in \Theta_{R(k)}^{(*)} \setminus \Theta_{R(k)}^{(0)}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \right) \\ &\leq \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \left(\left\lceil \frac{s}{I} \right\rceil + \sum_{g \in \mathbf{G}} \sum_{\theta' \in \Theta^{(*)} \setminus \Theta^{(0)}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \right) \end{aligned}$$

$$\begin{aligned}
 &\leq \log \left(\left\lceil \frac{s}{I} \right\rceil + s \cdot \sup_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \left(1 - \frac{|\Theta'_{R(k)}(0)|}{\binom{\tau(1-g_{k,I}) + d^* + \hat{d}_1^* - 2}{d^* + \hat{d}_1^* - 2}} \right) \right) \\
 &= \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \sup_{k \in \{0,1,\dots, \lceil \frac{s}{I} \rceil - 1\}} \frac{\binom{\tau(1-g_{k,I}) + \hat{d}_0^* + \hat{d}_1^* - 2}{\hat{d}_0^* + \hat{d}_1^* - 2}}{\binom{\tau(1-g_{k,I}) + d^* + \hat{d}_1^* - 2}{d^* + \hat{d}_1^* - 2}} \right) \\
 &\leq \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right).
 \end{aligned}$$

Combining with the privacy upper bound provided in Appendix G, we have

$$\begin{aligned}
 \Pi_{\mathcal{M}_{QM}, \mathbf{g}} \leq \min \left\{ \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right], \right. \\
 \left. \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right) \right\}.
 \end{aligned}$$

□

Appendix H.2. Lower Bounds for SML

Proposition A2. For any secret function $\mathbf{g}$, the SML lower bound of RR is

$$\Pi_{\mathcal{M}_{RR}, \mathbf{g}} \geq \begin{cases} \log \left(\frac{s(e^e - 1)}{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^e - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} \left(1 + \frac{\tau - (d^* - \hat{d}_0^*)}{\tau + d^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^e - 1} \right), & d^* - \hat{d}_0^* \leq \log s, \\ \log \left(\frac{s(e^e - 1)}{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^e - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} \left(1 + \frac{\tau - \log s}{\tau + d^* + \hat{d}_1^* - 1} \right)^{\log s}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^e - 1} \right), & \log s < d^* - \hat{d}_0^* < s, \\ \log \left(s - (s - 1) \left(\frac{\hat{d}_0^* + \hat{d}_1^* + \zeta - 1}{\tau + \hat{d}_0^* + \hat{d}_1^*} \right)^\zeta \right), & d^* - \hat{d}_0^* \geq s. \end{cases}$$

For any secret function $\mathbf{g}$, the SML lower bound of QM is

$$\Pi_{\mathcal{M}_{QM}, \mathbf{g}} \geq \begin{cases} \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau I - 2s(d^* - \hat{d}_0^*)}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^{d^* - \hat{d}_0^*} \right], & d^* - \hat{d}_0^* \leq \log I, \\ \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau I - 2s \cdot \log I}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^{\log I} \right], & \log I < d^* - \hat{d}_0^* < I, \\ \log \left(s - (s - \lceil \frac{s}{I} \rceil) \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* + \zeta - 2}{\tau I / 2s + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^\zeta \right), & d^* - \hat{d}_0^* \geq I. \end{cases}$$

Appendix H.2.1. Randomized Response

Proof. When $d^* - \hat{d}_0^* \leq \log s$, we can find a prior distribution $\mathbb{P}_\Theta$ of the distribution parameter such that $\mathbb{P}_{\Theta|G} \in \{0, 1\}$ and $\forall \tilde{\Gamma} \subseteq \Gamma^* \setminus \hat{\Gamma}_0^*, \exists g \in \mathbf{G} : \text{Supp}(\theta_g) = \tilde{\Gamma} \cup \hat{\Gamma}_0^*$. Therefore, we can get that

$$\begin{aligned}
 \Pi_{\mathcal{M}_{RR}, \mathbf{g}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0, 1\}} \log \sum_{\theta' \in \Theta^{(*)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \\
 &= \sup_{\mathbb{P}_{\Theta|G} \in \{0, 1\}} \log \left(\sum_{\theta' \in \{\theta_g\}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) + \sum_{\theta' \in \Theta^{(*)} \setminus \{\theta_g\}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \right) \\
 &\geq \log \left(\sum_{\theta' \in \{\theta_g\}} \sup_{g \in \mathbf{G}} \frac{e^\epsilon}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 1}{\hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 1} + e^\epsilon - 1} + \sum_{\theta' \in \Theta^{(*)} \setminus \{\theta_g\}} \sup_{g \in \mathbf{G}} \frac{1}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 1}{\hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 1} + e^\epsilon - 1} \right) \\
 &\geq \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \sum_{\theta' \in \Theta^{(*)}} \frac{1}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 1}{\hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 1} + e^\epsilon - 1} \right) \\
 &= \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \sum_{i \in \{0\} \cup [d^* - \hat{d}_0^*]} \frac{\binom{d^* - \hat{d}_0^*}{i} \binom{\tau - i + \hat{d}_0^* + \hat{d}_1^* + i - 1}{\hat{d}_0^* + \hat{d}_1^* + i - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* + i - 1} + e^\epsilon - 1} \right) \\
 &\geq \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \sum_{i \in \{0\} \cup [d^* - \hat{d}_0^*]} \frac{\binom{d^* - \hat{d}_0^*}{i} \binom{\tau - i + \hat{d}_0^* + \hat{d}_1^* + i - 1}{\hat{d}_0^* + \hat{d}_1^* + i - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* + i - 1}} \right) \\
 &= \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} \sum_{i \in \{0\} \cup [d^* - \hat{d}_0^*]} \frac{\tau! (\tau + \hat{d}_0^* + \hat{d}_1^* - 1)! \binom{d^* - \hat{d}_0^*}{i}}{(\tau - i)! (\tau + i + \hat{d}_0^* + \hat{d}_1^* - 1)!}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \right) \\
 &\geq \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} \sum_{i \in \{0\} \cup [d^* - \hat{d}_0^*]} \left(\frac{\tau - (d^* - \hat{d}_0^*)}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^i \binom{d^* - \hat{d}_0^*}{i}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \right) \\
 &= \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \left(1 + \frac{\tau - (d^* - \hat{d}_0^*)}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right).
 \end{aligned}$$

When $\log s < d^* - \hat{d}_0^* < s$, we can easily get that

$$\Pi_{\mathcal{M}_{RR}, \mathbf{g}} \geq \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \left(1 + \frac{\tau - \log s}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{\log s} \right)$$

When $d^* - \hat{d}_0^* \geq s$, denote $\tilde{\Gamma}_\theta \triangleq \text{Supp}(\theta) \setminus \hat{\Gamma}_0^*$. We can find a prior distribution $\mathbb{P}_\Theta$ of the distribution parameter such that $\mathbb{P}_{\Theta|G} \in \{0, 1\}$ and $\forall g_1, g_2 \in \mathbf{G}, g_1 \neq g_2 : \tilde{\Gamma}_{\theta_{g_1}} \cap \tilde{\Gamma}_{\theta_{g_2}} = \emptyset$ and $|\tilde{\Gamma}_{\theta_{g_1}}| = |\tilde{\Gamma}_{\theta_{g_2}}| = \left\lfloor \frac{d^* - \hat{d}_0^*}{s} \right\rfloor$. Denote $\zeta \triangleq \left\lfloor \frac{d^* - \hat{d}_0^*}{s} \right\rfloor$. Therefore, we can get that

$$\begin{aligned}
 \Pi_{\mathcal{M}_{RR}, \mathbf{g}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0, 1\}} \log \sum_{\theta' \in \Theta^{(*)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \\
 &\geq \log \left(\frac{|\Theta^{(0)}|}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* + \zeta - 1}{\hat{d}_0^* + \hat{d}_1^* + \zeta - 1} + e^\epsilon - 1} + \sum_{g \in \mathbf{G}} \left(1 - \frac{|\Theta^{(0)}|}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* + \zeta - 1}{\hat{d}_0^* + \hat{d}_1^* + \zeta - 1} + e^\epsilon - 1} \right) \right)
 \end{aligned}$$

$$\begin{aligned} &\geq \log \left(s - (s-1) \cdot \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* + \zeta - 1}{\hat{d}_0^* + \hat{d}_1^* + \zeta - 1} + e^\epsilon - 1} \right) \\ &\geq \log \left(s - (s-1) \left(\frac{\hat{d}_0^* + \hat{d}_1^* + \zeta - 1}{\tau + \hat{d}_0^* + \hat{d}_1^*} \right)^\zeta \right). \end{aligned}$$

□

Appendix H.2.2. Quantization Mechanism

Proof. Denote $\mathbf{G}_k$ as the set of the secret values within the k -th interval, i.e., $\mathbf{G}_k = \{g_{kI+j}\}_{j \in [I]}$. When $d^* - \hat{d}_0^* \leq \log I$, we can find a prior distribution $\mathbb{P}_\Theta$ of the distribution parameter such that $\mathbb{P}_{\Theta|G} \in \{0, 1\}$ and $\forall \tilde{\Gamma} \subseteq \Gamma^* \setminus \hat{\Gamma}_0^*, k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}, \exists g \in \mathbf{G}_k : \text{Supp}(\theta_g) = \tilde{\Gamma} \cup \hat{\Gamma}_0^*$. Therefore, we can get that

$$\begin{aligned} \Pi_{\mathcal{M}_{\text{QM}, \mathbf{g}}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0, 1\}} \log \sum_{\theta' \in \Theta^{(*)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \\ &\geq \sup_{\mathbb{P}_{\Theta|G} \in \{0, 1\}} \log \sum_{k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}} \sum_{\theta' \in \Theta_{R(k)}^{(*)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'| \Theta}(\theta' | \theta_g) \\ &\geq \log \sum_{k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}} \sum_{\theta' \in \Theta_{R(k)}^{(*)}} \frac{1}{\binom{\tau(1-g_{k,I}) + \hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 2}{\hat{d}_0^* + \hat{d}_1^* + \beta(\theta') - 2}} \\ &= \log \sum_{k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}} \sum_{i \in \{0, 1, \dots, d^* - \hat{d}_0^*\}} \frac{\binom{d^* - \hat{d}_0^*}{i} \binom{\tau(1-g_{k,I}) - i + \hat{d}_0^* + \hat{d}_1^* + i - 2}{\hat{d}_0^* + \hat{d}_1^* + i - 2}}{\binom{\tau(1-g_{k,I}) + \hat{d}_0^* + \hat{d}_1^* + i - 2}{\hat{d}_0^* + \hat{d}_1^* + i - 2}} \\ &= \log \sum_{k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}} \sum_{i \in \{0, 1, \dots, d^* - \hat{d}_0^*\}} \frac{(\tau(1-g_{k,I}))! (\tau(1-g_{k,I}) + \hat{d}_0^* + \hat{d}_1^* - 2)! \binom{d^* - \hat{d}_0^*}{i}}{(\tau(1-g_{k,I}) - i)! (\tau(1-g_{k,I}) + i + \hat{d}_0^* + \hat{d}_1^* - 2)!} \\ &\geq \log \sum_{k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}} \sum_{i \in \{0, 1, \dots, d^* - \hat{d}_0^*\}} \left(\frac{\tau I - 2s(d^* - \hat{d}_0^*)}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^i \binom{d^* - \hat{d}_0^*}{i} \\ &= \log \sum_{k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}} \left(1 + \frac{\tau I - 2s(d^* - \hat{d}_0^*)}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^{d^* - \hat{d}_0^*} \\ &= \log \left[\left\lfloor \frac{s}{I} \right\rfloor \left(1 + \frac{\tau I - 2s(d^* - \hat{d}_0^*)}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^{d^* - \hat{d}_0^*} \right]. \end{aligned}$$

When $\log I < d^* - \hat{d}_0^* < I$, we can easily get that

$$\Pi_{\mathcal{M}_{\text{QM}, \mathbf{g}}} \geq \log \left[\left\lfloor \frac{s}{I} \right\rfloor \left(1 + \frac{\tau I - 2s \cdot \log I}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^{\log I} \right].$$

When $d^* - \hat{d}_0^* \geq I$, denote $\tilde{\Gamma}_\theta \triangleq \text{Supp}(\theta) \setminus \hat{\Gamma}_0^*$. We can find a prior distribution $\mathbb{P}_\Theta$ of the distribution parameter such that $\mathbb{P}_{\Theta|G} \in \{0, 1\}$ and $\forall k \in \{0, 1, \dots, \lfloor \frac{s}{I} \rfloor - 1\}, g_1, g_2 \in$

$\mathbf{G}_k, g_1 \neq g_2 : \tilde{\mathbf{r}}_{\theta_{g_1}} \cap \tilde{\mathbf{r}}_{\theta_{g_2}} = \emptyset$ and $|\tilde{\mathbf{r}}_{\theta_{g_1}}| = |\tilde{\mathbf{r}}_{\theta_{g_2}}| = \left\lfloor \frac{d^* - \hat{d}_0^*}{I} \right\rfloor$. Denote $\zeta \triangleq \left\lfloor \frac{d^* - \hat{d}_0^*}{I} \right\rfloor$. Therefore, we can get that

$$\begin{aligned} \Pi_{\mathcal{M}_{QM}, \mathbf{g}} &= \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta^{(*)}} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta' | \theta_g) \\ &\geq \log \left[\sum_{k \in \{0,1, \dots, \lceil \frac{s}{I} \rceil - 1\}} \left(\frac{|\Theta_{R(k)}^{(0)}|}{\binom{\tau(1-s_{k,I}) + \hat{d}_0^* + \zeta + \hat{d}_1^* - 2}{\hat{d}_0^* + \zeta + \hat{d}_1^* - 2}} + \min\{I, s - kI\} \cdot \left(1 - \frac{|\Theta_{R(k)}^{(0)}|}{\binom{\tau(1-s_{k,I}) + \hat{d}_0^* + \zeta + \hat{d}_1^* - 2}{\hat{d}_0^* + \zeta + \hat{d}_1^* - 2}} \right) \right) \right] \\ &\geq \log \left(s - (s - \lceil \frac{s}{I} \rceil) \cdot \sup_{k \in \{0,1, \dots, \lceil \frac{s}{I} \rceil - 1\}} \frac{\binom{\tau(1-s_{k,I}) + \hat{d}_0^* + \hat{d}_1^* - 2}{\hat{d}_0^* + \hat{d}_1^* - 2}}{\binom{\tau(1-s_{k,I}) + \hat{d}_0^* + \zeta + \hat{d}_1^* - 2}{\hat{d}_0^* + \zeta + \hat{d}_1^* - 2}} \right) \\ &\geq \log \left(s - (s - \lceil \frac{s}{I} \rceil) \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* + \zeta - 2}{\tau I / 2s + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^\zeta \right). \end{aligned}$$

□

Appendix H.3. Tightness of Lower and Upper Bounds for SML

Proposition A3. Let $\Delta_{\mathcal{M}_{RR}, \tau$ and $\Delta_{\mathcal{M}_{QM}, \tau$ be the gaps between SML lower and upper bounds of RR and QM under precision level τ , respectively.

For RR, we have

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau &\leq 1, \quad \text{when } d^* - \hat{d}_0^* < s, \quad e^\epsilon - 1 \leq \left(\frac{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} \right) / s, \\ \lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau &= 0, \quad \text{when } d^* - \hat{d}_0^* \geq s. \end{aligned}$$

For QM, we have

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{QM}, \tau &\leq 1, \quad \text{when } d^* - \hat{d}_0^* < I, \\ \lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{QM}, \tau &= 0, \quad \text{when } d^* - \hat{d}_0^* \geq I. \end{aligned}$$

Specifically, when s can be divided by I , $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{QM}, \tau = 0$ whatever $d^* - \hat{d}_0^*$ is.

Appendix H.3.1. Randomized Response

Proof. When $d^* - \hat{d}_0^* \leq \log s$ and $e^\epsilon - 1 \leq \left(\frac{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} \right) / s$, we have

$$\begin{aligned} &\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau \\ &\leq \lim_{\tau \rightarrow \infty} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right) \\ &\quad - \lim_{\tau \rightarrow \infty} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \left(1 + \frac{\tau - (d^* - \hat{d}_0^*)}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{d^* - \hat{d}_0^*} \right) \\ &= \lim_{\tau \rightarrow \infty} \log \left(\frac{\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \cdot 2^{d^* - \hat{d}_0^*}}{\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \cdot 2^{d^* - \hat{d}_0^*}} \right) \end{aligned}$$

$$\begin{aligned} &\leq \log \left(\frac{\frac{1}{1+1/s} + 2^{d^* - \hat{d}_0^*}}{\frac{1}{1+1/s} \cdot 2^{d^* - \hat{d}_0^*}} \right) \\ &= \log \left(1 + \frac{1}{s} + 2^{-(d^* - \hat{d}_0^*)} \right). \end{aligned}$$

Since we focus on the nontrivial case where $s > 1$ and $d^* - \hat{d}_0^* > 0$, we have $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau \leq 1$.

When $\log s < d^* - \hat{d}_0^* < s$ and $e^\epsilon - 1 \leq \binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} / s$, we have

$$\begin{aligned} &\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau \\ &\leq \log s - \lim_{\tau \rightarrow \infty} \log \left(\frac{s(e^\epsilon - 1)}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} + \frac{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} + e^\epsilon - 1} \left(1 + \frac{\tau - \log s}{\tau + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^{\log s} \right) \\ &\leq \log \left(\frac{s}{1 + 1/s} \cdot 2^{\log s} \right) \\ &= \log \left(1 + \frac{1}{s} \right) \\ &< 1. \end{aligned}$$

When $d^* - \hat{d}_0^* \geq s$, denote $\zeta \triangleq \left\lfloor \frac{d^* - \hat{d}_0^*}{s} \right\rfloor$ and we have

$$\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau \leq \log s - \lim_{\tau \rightarrow \infty} \log \left(s - (s-1) \left(\frac{\hat{d}_0^* + \hat{d}_1^* + \zeta - 1}{\tau + \hat{d}_0^* + \hat{d}_1^*} \right)^\zeta \right) = \log s - \log s = 0.$$

Hence, we can get that $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau \leq 1$ when $d^* - \hat{d}_0^* \leq s$ and $e^\epsilon - 1 \leq \binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1} / s$, and $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{RR}, \tau = 0$ when $d^* - \hat{d}_0^* > s$. $\square$

Appendix H.3.2. Quantization Mechanism

Proof. Let $\Delta_{\mathcal{M}_{QM}, \tau$ be the gap of lower and upper bounds of the privacy of randomized response under a dataset with τ samples. When $d^* - \hat{d}_0^* \leq \log I$, we have

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{QM}, \tau &\leq \lim_{\tau \rightarrow \infty} \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right] \\ &\quad - \lim_{\tau \rightarrow \infty} \log \left[\left\lfloor \frac{s}{I} \right\rfloor \left(1 + \frac{\tau I - 2s(d^* - \hat{d}_0^*)}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^{d^* - \hat{d}_0^*} \right] \\ &= \log \left\lceil \frac{s}{I} \right\rceil - \log \left\lfloor \frac{s}{I} \right\rfloor \\ &\leq 1. \end{aligned}$$

Specifically, when s can be divided by I , $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{QM}, \tau = 0$.

When $\log I < d^* - \hat{d}_0^* < I$, we have

$$\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{QM}, \tau \leq \log s - \lim_{\tau \rightarrow \infty} \log \left[\left\lfloor \frac{s}{I} \right\rfloor \left(1 + \frac{\tau I - 2s \log I}{\tau I + 2s(d^* + \hat{d}_1^* - 2)} \right)^{\log I} \right]$$

$$= \log s - \log \left(I \cdot \left\lfloor \frac{s}{I} \right\rfloor \right) < 1.$$

When $d^* - \hat{d}_0^* \geq I$, denote $\tilde{\Gamma}_\theta \triangleq \text{Supp}(\theta) \setminus \hat{\Gamma}_0^*$ and we have

$$\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{\text{QM}, \tau} \leq \log s - \lim_{\tau \rightarrow \infty} \log \left(s - (s - \left\lfloor \frac{s}{I} \right\rfloor) \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* + \zeta - 2}{\tau I / 2s + \hat{d}_0^* + \hat{d}_1^* - 1} \right)^\zeta \right) = \log s - \log s = 0.$$

Hence, we can get that $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{\text{QM}, \tau} \leq 1$ when $d^* - \hat{d}_0^* \leq I$, and $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{\text{QM}, \tau} = 0$ otherwise. Specifically, when s can be divided by I , $\lim_{\tau \rightarrow \infty} \Delta_{\mathcal{M}_{\text{QM}, \tau} = 0$ whatever $d^* - \hat{d}_0^*$ is. $\square$

Appendix H.4. Lower and Upper Bounds for Distortion

Appendix H.4.1. Randomized Response

Similar to the proof of Corollary 1, we can bound the distortion of randomized response as

$$\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{(\hat{d}_0^* + \hat{d}_1^*)(1 + r_1)} \leq \Delta_{\mathcal{M}_{\text{RR}}} \leq \frac{d^* + \hat{d}_1^* - 1}{(d^* + \hat{d}_1^*)(1 + r_2)},$$

$$\text{where } r_1 \triangleq \frac{e^\epsilon - 1}{\binom{\tau + \hat{d}_0^* + \hat{d}_1^* - 1}{\hat{d}_0^* + \hat{d}_1^* - 1}} \text{ and } r_2 \triangleq \frac{e^\epsilon - 1}{\binom{\tau + d^* + \hat{d}_1^* - 1}{d^* + \hat{d}_1^* - 1}}.$$

Appendix H.4.2. Quantization Mechanism

When secret is the fraction of a category, similar to the proof of Corollary 1, we can bound the distortion of the quantization mechanism as

$$\frac{1}{2} + \frac{(\hat{d}_0^* + \hat{d}_1^*) \left\lfloor \frac{I}{2} \right\rfloor - \tau}{2\tau(\hat{d}_0^* + \hat{d}_1^* - 1)} \leq \Delta_{\mathcal{M}_{\text{QM}}} \leq \frac{1}{2} + \frac{(d^* + \hat{d}_1^*) \left\lfloor \frac{I}{2} \right\rfloor - \tau}{2\tau(d^* + \hat{d}_1^* - 1)}.$$

Appendix I. Analysis of Privacy Performance with the Increase of $\hat{d}_1^*$

In this section, we provide theoretical analysis of the patterns of decrease in privacy value under different numbers of missing feasible attribute combinations $d^* - \hat{d}_0^*$; that is, with increasing $\hat{d}_1^*$, the privacy value drops linearly when $d^* - \hat{d}_0^* = 5$, drops in a concave curve when $d^* - \hat{d}_0^* = 20$, and always achieves its trivial upper bound $\log s$ when $d^* - \hat{d}_0^* = 100$.

From Proposition A1 we know that the privacy of Quantization mechanism can be upper bounded by

$$\Pi_{\mathcal{M}_{\text{RR}, \mathcal{G}}} \leq \min \left\{ \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right], \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right) \right\}.$$

For the Census Income Dataset we adopt in our experiment, we have $\tau = 48,842$, $\hat{d}_0^* = 22,381$, and $\log s = 4.09$. When $d^* - \hat{d}_0^* = 5$ and $\hat{d}_1^* \ll \tau + \hat{d}_0^*$, based on simple calculation, we can get that the privacy is upper bounded by $\Pi_{\mathcal{M}_{RR}, \mathfrak{g}} \leq \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^5 \right]$.

Then we have

$$\begin{aligned} \Pi_{\mathcal{M}_{RR}, \mathfrak{g}} &\leq \log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^5 \right] \\ &= \log \left\lceil \frac{s}{I} \right\rceil + 5 \left[\log(2\tau + \hat{d}_0^* + \hat{d}_1^* - 2) - \log(\tau + \hat{d}_0^* + \hat{d}_1^* - 2) \right] \\ &= \log \left\lceil \frac{s}{I} \right\rceil + 5 \left[\log \frac{2\tau + \hat{d}_0^* - 2}{\tau + \hat{d}_0^* - 2} + \log \left(1 + \frac{\hat{d}_1^*}{2\tau + \hat{d}_0^* - 2} \right) - \log \left(1 + \frac{\hat{d}_1^*}{\tau + \hat{d}_0^* - 2} \right) \right] \\ &= \log \left\lceil \frac{s}{I} \right\rceil + 5 \left[\log \frac{2\tau + \hat{d}_0^* - 2}{\tau + \hat{d}_0^* - 2} - \frac{\tau}{(2\tau + \hat{d}_0^* - 2)(\tau + \hat{d}_0^* - 2)} \cdot \hat{d}_1^* \right] + o \left(\frac{\hat{d}_1^*}{\tau + \hat{d}_0^*} \right), \end{aligned}$$

which indicates a linear function of $\hat{d}_1^*$ dominate the privacy upper bound. Therefore, when $d^* - \hat{d}_0^* = 5$ and $\hat{d}_1^* \ll \tau + \hat{d}_0^*$, the privacy upper bound of the quantization mechanism decrease linearly with the increase of $\hat{d}_1^*$.

When $d^* - \hat{d}_0^* = 20$ and $\tau \leq \hat{d}_1^* < 5\tau$, we can get that the privacy is upper bounded by $\Pi_{\mathcal{M}_{RR}, \mathfrak{g}} \leq \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{20} \right)$. Let $f(\hat{d}_1^*) \triangleq \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{20}$ and we have

$$\frac{d^2 f(\hat{d}_1^*)}{d(\hat{d}_1^*)^2} = \frac{20(\tau - 1) \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{18}}{(\tau + d^* + \hat{d}_1^* - 2)^4} \cdot \left[19(\tau - 1) - 2(\hat{d}_0^* + \hat{d}_1^* - 1) \right] > 0.$$

Therefore, $g(\hat{d}_1^*) = \left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{20} = \left\lceil \frac{s}{I} \right\rceil + s - s \cdot f(\hat{d}_1^*)$ is a concave function. Since $l(x) = \log(x)$ is a concave and monotonically increasing function, we have $h(\hat{d}_1^*) = \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{20} \right) = l \circ g(\hat{d}_1^*)$ is a concave function. Therefore, when $d^* - \hat{d}_0^* = 20$ and $\hat{d}_1^* < 5\tau$, with the increase of $\hat{d}_1^*$, the privacy upper bound of the quantization mechanism decreases in a concave curve.

When $d^* - \hat{d}_0^* = 100$ and $\hat{d}_1^* < 5\tau$, based on simple calculation, we can get that the privacy upper bound is always $\log s$.

More generally, we analyze the patterns of decrease in privacy value with increasing $\hat{d}_1^*$ under the case where $\hat{d}_1^* \ll \hat{d}_0^*$ and $d^* - \hat{d}_0^* \ll \tau + \hat{d}_0^*$. Under this setting, we can get that $\frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \approx \frac{\tau}{\tau + \hat{d}_0^*} \triangleq r$ and $\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \approx \frac{\hat{d}_0^*}{\tau + \hat{d}_0^*} = 1 - r$. Therefore, when $\log \left[\left\lceil \frac{s}{I} \right\rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right] \leq \log \left(\left\lceil \frac{s}{I} \right\rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + d^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right)$, we have

$$\left\lceil \frac{s}{I} \right\rceil (1 + r)^{d^* - \hat{d}_0^*} + s(1 - r)^{d^* - \hat{d}_0^*} \leq \left\lceil \frac{s}{I} \right\rceil + s.$$

Let $y(d^* - \hat{d}_0^*) = \lceil \frac{s}{I} \rceil (1+r)^{d^* - \hat{d}_0^*} + s(1-r)^{d^* - \hat{d}_0^*}$, and we can get that $\frac{d^2 y(d^* - \hat{d}_0^*)}{d(d^* - \hat{d}_0^*)^2} \geq 0$ when $d^* - \hat{d}_0^* \geq 0$. Let $\alpha = \frac{\log s - \log \lceil \frac{s}{I} \rceil}{\log(1+r)}$, since $(1-r)^\alpha \leq 1/(1+r)^\alpha$, we have

$$\lceil \frac{s}{I} \rceil (1+r)^\alpha + s(1-r)^\alpha \leq \lceil \frac{s}{I} \rceil + s.$$

Since $y(d^* - \hat{d}_0^*)$ is a convex function and $y(0) = \lceil \frac{s}{I} \rceil + s$, we can get that $\lceil \frac{s}{I} \rceil (1+r)^{d^* - \hat{d}_0^*} + s(1-r)^{d^* - \hat{d}_0^*} \leq \lceil \frac{s}{I} \rceil + s$ when $0 \leq d^* - \hat{d}_0^* \leq \alpha$. Additionally, we can easily get that when $d^* - \hat{d}_0^* > \alpha$,

$$\min \left\{ \log \left[\lceil \frac{s}{I} \rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right], \log \left(\lceil \frac{s}{I} \rceil + s - s \cdot \left(\frac{\hat{d}_0^* + \hat{d}_1^* - 1}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right) \right\} > \log s.$$

Therefore, the privacy value is upper bounded by

$$\Pi_{\mathcal{M}_{RR}, \mathfrak{g}} \leq \begin{cases} \log \left[\lceil \frac{s}{I} \rceil \left(1 + \frac{\tau}{\tau + \hat{d}_0^* + \hat{d}_1^* - 2} \right)^{d^* - \hat{d}_0^*} \right], & 0 \leq d^* - \hat{d}_0^* < \alpha, \\ \log s, & \beta \leq d^* - \hat{d}_0^*. \end{cases}$$

Similar to the analyses above, we can get that with increasing $\hat{d}_1^*$, the privacy value drops linearly when $0 \leq d^* - \hat{d}_0^* < \alpha$, and always achieves its trivial upper bound $\log s$ when $d^* - \hat{d}_0^* \geq \alpha$.

Appendix J. Privacy–Distortion Tradeoff Analysis When $\Gamma \subseteq \Gamma^* \neq \hat{\Gamma}^*$

We conduct numerical experiments to illustrate the lower and upper bounds of the privacy–distortion performance of our proposed quantization mechanism and randomized response. We consider a dataset with 5M samples and 100 feasible attribute combinations in total, i.e., $d^* = 100$. We vary the number of actually feasible attribute combinations that the data holder does not consider to be valid, i.e., $d^* - \hat{d}_0^*$, from 5 to 20. We assume the data holder introduces 40 invalid attribute combinations during the data release process, i.e., $\hat{d}_1^* = 40$. We can observe that with the increase of the number of missed feasible attribute combinations, the minimal achievable statistic maximal leakage of both mechanisms increase. The reason is that correlation may exists between the secret value and the appearance of those missed feasible attribute combinations in the original dataset. The existence of certain missed feasible attribute combinations in the released dataset attacker may serve as the indication of certain secret values. We can also observe that with smaller number of missed feasible attribute combinations, the gaps between the lower and upper bounds of the privacy–distortion performance for both mechanism are smaller. With different number of missed feasible attribute combination, quantization mechanism consistently outperforms randomized response under most leakage regimes (if statistic maximal leakage is smaller than 21).




Figure A2. The lower and upper bounds of the privacy–distortion performance of quantization mechanism and randomized response with different number of feasible attribute combinations the data holder misses, i.e., $d^* - \hat{d}_0^*$.

Appendix K. Alternative Privacy Measures

Recall that our proposed statistic maximal leakage can be written as

$$\Pi_{\mathcal{M},g} = \sup_{\mathbb{P}_{\Theta}, \mathbb{P}_{\hat{G}|\Theta'}} \log \frac{\mathbb{P}(\hat{G} = G)}{\sup_{g \in \mathcal{G}} \mathbb{P}_G(g)},$$

which aims to measure the information gain of the attacker’s knowledge about what the secret value is most likely to be after data release under the worst-case prior $\mathbb{P}_{\Theta}$ and attack strategy $\mathbb{P}_{\hat{G}|\Theta'}$.

In scenarios related to social justice and fairness, where knowing what the secret is most unlikely to be also reveals sensitive information, we propose an alternative privacy metric to measure the information gain of the attacker’s knowledge about what the secret value is most unlikely to be

$$\Pi_{\mathcal{M},g}^{\text{inf}} = \begin{cases} \infty, & \exists \mathbb{P}_{\Theta}, \mathbb{P}_{\hat{G}|\Theta'} : \mathbb{P}_G > 0 \text{ \& } \mathbb{P}(\hat{G} = G) = 0, \\ \sup_{\mathbb{P}_{\Theta}: \mathbb{P}_G > 0; \mathbb{P}_{\hat{G}|\Theta'}} \log \frac{\inf_{g \in \mathcal{G}} \mathbb{P}_G(g)}{\mathbb{P}(\hat{G} = G)}, & \text{otherwise.} \end{cases}$$

Similarly to statistic maximal leakage, $\Pi_{\mathcal{M},g}^{\text{inf}}$ is also defined under the worst-case prior $\mathbb{P}_{\Theta}$ and the attack strategy $\mathbb{P}_{\hat{G}|\Theta'}$, but consider the lowest prior probability of the secret value in the numerator of the ratio. A higher value of $\Pi_{\mathcal{M},g}^{\text{inf}}$ indicates that the attacker gains more information about the least likely secret. We set $\Pi_{\mathcal{M},g}^{\text{inf}}$ to infinity when the adversary can successfully infer what the secret value is not after the data release.

Furthermore, statistic maximal leakage measures the attack success probability over the randomness of the released data. In a more conservative scenario, however, one may

prefer to measure the attack success probability under the worst-case output. For this scenario, we propose alternative privacy measures, $\tilde{\Pi}_{\mathcal{M},g}$ and $\tilde{\Pi}_{\mathcal{M},g}^{\text{inf}}$, which quantify the information leakage of the most likely and most unlikely secret values, respectively:

$$\tilde{\Pi}_{\mathcal{M},g} = \sup_{\mathbb{P}_{\Theta}, \theta' \in \Theta', \mathbb{P}_{\hat{G}|\theta'}} \log \frac{\mathbb{P}(\hat{G} = G)}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)},$$

$$\tilde{\Pi}_{\mathcal{M},g}^{\text{inf}} = \begin{cases} \infty, & \exists \mathbb{P}_{\Theta}, \theta' \in \Theta', \mathbb{P}_{\hat{G}|\theta'} : \mathbb{P}_G > 0 \ \& \ \mathbb{P}(\hat{G} = G) = 0, \\ \sup_{\mathbb{P}_{\Theta} : \mathbb{P}_G > 0; \theta' \in \Theta'; \mathbb{P}_{\hat{G}|\theta'}} \log \frac{\inf_{g \in \mathbf{G}} \mathbb{P}_G(g)}{\mathbb{P}(\hat{G} = G)}, & \text{otherwise.} \end{cases}$$

In addition to the worst-case prior and attack strategy, $\tilde{\Pi}_{\mathcal{M},g}$ and $\tilde{\Pi}_{\mathcal{M},g}^{\text{inf}}$ are also defined under the worst-case data output θ' ; thus, the attack strategy is represented as $\mathbb{P}_{\hat{G}|\theta'}$.

Additionally, inspired by differential privacy, we design the indistinguishability-based privacy measure $\mathcal{L}_{DP}$, and we say that a data release mechanism $\mathcal{M}$ is μ - $\mathcal{L}_{DP}$ if

$$\mathbb{P}(\mathcal{M}(\theta_1) \in \Theta'_0) \leq e^\mu \cdot \mathbb{P}(\mathcal{M}(\theta_2) \in \Theta'_0), \quad \forall \theta_1, \theta_2 \in \Theta, g(\theta_1) \neq g(\theta_2); \forall \Theta'_0 \subseteq \Theta'.$$

We analyze the connections between those privacy measures in Proposition A4.

Proposition A4. *If a data release mechanism $\mathcal{M}$ is μ - $\mathcal{L}_{DP}$, then it satisfies $\Pi_{\mathcal{M},g} \leq \mu$ and $\Pi_{\mathcal{M},g}^{\text{inf}} \leq \mu$.*

If a data release mechanism $\mathcal{M}$ satisfies $\tilde{\Pi}_{\mathcal{M},g} \leq \mu_1$ and $\tilde{\Pi}_{\mathcal{M},g}^{\text{inf}} \leq \mu_2$, then it is $(\mu_1 + \mu_2)$ - $\mathcal{L}_{DP}$.

Proof. We first prove that $\Pi_{\mathcal{M},g} \leq \mu$ and $\Pi_{\mathcal{M},g}^{\text{inf}} \leq \mu$ if a data release mechanism $\mathcal{M}$ is μ - $\mathcal{L}_{DP}$ as follows.

When a mechanism $\mathcal{M}$ is μ - $\mathcal{L}_{DP}$, i.e.,

$$\mathbb{P}(\mathcal{M}(\theta_1) \in \Theta'_0) \leq e^\mu \cdot \mathbb{P}(\mathcal{M}(\theta_2) \in \Theta'_0), \quad \forall \theta_1, \theta_2 \in \Theta, g(\theta_1) \neq g(\theta_2); \forall \Theta'_0 \subseteq \Theta',$$

we have

$$\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_1) \leq e^\mu \cdot \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_2), \quad \forall \theta_1, \theta_2 \in \Theta, g(\theta_1) \neq g(\theta_2); \forall \theta' \in \Theta'. \quad (\text{A7})$$

From Proposition 1, we know that

$$\Pi_{\mathcal{M},g} = \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g).$$

We can find a $\theta_0 \in \Theta$ that satisfies $\exists g \in \mathbf{G}, \mathbb{P}_{\Theta|G}^* \in \arg \sup_{\mathbb{P}_{\Theta|G}} \Pi_{\mathcal{M},g} : \theta_0 = \theta_g$. Under the prior distribution $\mathbb{P}_{\Theta|G}^*$, we can get that

$$\Pi_{\mathcal{M},g} = \log \sum_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_g) \leq \log \sum_{\theta' \in \Theta'} e^\mu \cdot \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_0) = \mu.$$

From Equation (A7), we know that $\forall \theta \in \Theta, \theta' \in \Theta' : \mathbb{P}_{\Theta|\Theta}(\theta'|\theta) > 0$. For the value of $\tilde{\Pi}_{\mathcal{M},g}$, since $\forall \Theta : \mathbb{P}_G > 0, \forall \mathbb{P}_{\hat{G}|\Theta'}$,

$$\begin{aligned}\mathbb{P}(\hat{G} = G) &\geq \sum_{\theta' \in \Theta'} \inf_{g \in \mathbf{G}} \mathbb{P}_{G|\Theta'}(g|\theta') \mathbb{P}_{\Theta'}(\theta') \\ &= \sum_{\theta' \in \Theta'} \inf_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g) \mathbb{P}_G(g) \\ &= \sum_{\theta' \in \Theta'} \inf_{g \in \mathbf{G}} \sum_{\theta \in \Theta_g} \mathbb{P}_{\Theta|\Theta}(\theta'|\theta) \mathbb{P}_{\Theta|G}(\theta|g) \mathbb{P}_G(g) \\ &> 0,\end{aligned}$$

we can get that $\tilde{\Pi}_{\mathcal{M},g} < \infty$. Similar to the proof of Proposition 1, we can get that when $\tilde{\Pi}_{\mathcal{M},g} < \infty$,

$$\tilde{\Pi}_{\mathcal{M},g} = \sup_{\mathbb{P}_{\Theta|G} \in \{0,1\}} \log \frac{1}{\sum_{\theta' \in \Theta'} \inf_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g)}.$$

We can find a $\theta_1 \in \Theta$ that satisfies $\exists g \in \mathbf{G}, \mathbb{P}_{\Theta|G}^* \in \arg \sup_{\mathbb{P}_{\Theta|G}} \tilde{\Pi}_{\mathcal{M},g} : \theta_1 = \theta_g$. Under the prior distribution $\mathbb{P}_{\Theta|G}^*$, we can get that

$$\tilde{\Pi}_{\mathcal{M},g} = \log \frac{1}{\sum_{\theta' \in \Theta'} \inf_{g \in \mathbf{G}} \mathbb{P}_{\Theta'|G}(\theta'|g)} \leq \log \frac{1}{\sum_{\theta' \in \Theta'} e^{-\mu} \cdot \mathbb{P}_{\Theta'|G}(\theta'|\theta_0)} = \mu.$$

Above all, we have $\Pi_{\mathcal{M},g} \leq \mu$ and $\Pi_{\mathcal{M},g}^{\inf} \leq \mu$ when a data release mechanism $\mathcal{M}$ is μ - $\mathcal{L}_{DP}$.

If $\mathcal{M}$ satisfies $\tilde{\Pi}_{\mathcal{M},g} \leq \mu_1$ and $\tilde{\Pi}_{\mathcal{M},g}^{\inf} \leq \mu_2$, we have

$$\begin{aligned}\mu_1 + \mu_2 &\geq \tilde{\Pi}_{\mathcal{M},g} + \tilde{\Pi}_{\mathcal{M},g}^{\inf} \\ &= \sup_{\mathbb{P}_{\Theta}, \theta' \in \Theta', \mathbb{P}_{\hat{G}|\theta'}} \log \frac{\mathbb{P}(\hat{G} = G)}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)} + \sup_{\mathbb{P}_{\Theta}: \mathbb{P}_G > 0; \theta' \in \Theta'; \mathbb{P}_{\hat{G}|\theta'}} \log \frac{\inf_{g \in \mathbf{G}} \mathbb{P}_G(g)}{\mathbb{P}(\hat{G} = G)} \\ &\geq \sup_{\substack{\mathbb{P}_{\Theta}: \mathbb{P}_G = 1/s, \\ \theta' \in \Theta'}} \left(\log \frac{\sup_{\mathbb{P}_{\hat{G}|\theta'}} \mathbb{P}(\hat{G} = G)}{\sup_{g \in \mathbf{G}} \mathbb{P}_G(g)} + \log \frac{\inf_{g \in \mathbf{G}} \mathbb{P}_G(g)}{\inf_{\mathbb{P}_{\hat{G}|\theta'}} \mathbb{P}(\hat{G} = G)} \right) \\ &= \sup_{\substack{\mathbb{P}_{\Theta}: \mathbb{P}_G = 1/s, \\ \theta' \in \Theta'}} \log \frac{\sup_{\mathbb{P}_{\hat{G}|\theta'}} \mathbb{P}(\hat{G} = G)}{\inf_{\mathbb{P}_{\hat{G}|\theta'}} \mathbb{P}(\hat{G} = G)} \\ &= \sup_{\substack{\mathbb{P}_{\Theta}: \mathbb{P}_G = 1/s, \\ \theta' \in \Theta'}} \log \frac{\sup_{g \in \mathbf{G}} \mathbb{P}(g|\theta')}{\inf_{g \in \mathbf{G}} \mathbb{P}(g|\theta')} \\ &= \sup_{\substack{\mathbb{P}_{\Theta}: \mathbb{P}_G = 1/s, \\ \theta' \in \Theta'}} \log \frac{\sup_{g \in \mathbf{G}} \mathbb{P}(\theta'|g) \cdot \mathbb{P}_G(g)}{\inf_{g \in \mathbf{G}} \mathbb{P}(\theta'|g) \cdot \mathbb{P}_G(g)} \\ &\geq \sup_{\substack{\mathbb{P}_{\Theta}: \mathbb{P}_G = 1/s, \mathbb{P}_{\Theta|G} \in \{0,1\}, \\ \theta' \in \Theta'}} \log \frac{\sup_{g \in \mathbf{G}} \mathbb{P}(\theta'|g)}{\inf_{g \in \mathbf{G}} \mathbb{P}(\theta'|g)},\end{aligned}$$

where θ_g satisfies $\mathbb{P}_{\Theta|G}(\theta_g|g) = 1$. Therefore, we can get that

$$\mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_1) \leq e^{\mu_1 + \mu_2} \cdot \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta_2), \quad \forall \theta_1, \theta_2 \in \Theta, g(\theta_1) \neq g(\theta_2); \forall \theta' \in \Theta',$$

i.e., $\mathcal{M}$ is $(\mu_1 + \mu_2)$ - $\mathcal{L}_{DP}$. $\square$

Appendix L. Statistic Maximal Leakage Under Continuous Parameter Space

When extending to the continuous distribution parameter and secret spaces, statistic maximal leakage can be written as

$$\Pi_{\mathcal{M},g} = \sup_{\mathbf{f}_{\Theta}, \mathbf{f}_{\hat{G}|\Theta'}} \log \frac{\mathbf{f}(\hat{G} = G)}{\sup_{g \in \mathbf{G}} \mathbf{f}_G(g)}, \quad (\text{A8})$$

where $\mathbf{f}$ denotes the probability density function. Similar to Proposition 1, we provide an equivalent expression of Equation (A8) under continuous parameter space as follows.

Proposition A5. *Under continuous parameter space, statistic maximal leakage satisfies*

$$\Pi_{\mathcal{M},g} = \sup_{\mathbf{f}_{\Theta}: \mathbf{f}_{\Theta}(\theta) \in \{0, \mathbf{f}_G(g)\}, \forall \theta \in \Theta_g, g \in \mathbf{G}} \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{\Theta'|\Theta}(\theta'|\theta_g) d\theta',$$

where θ_g satisfies $\mathbf{f}_{\Theta}(\theta_g) = \mathbf{f}_G(g)$.

Proof. Let Θ_g be the set of distribution parameters whose secret value is g , i.e., $\Theta_g = \{\theta \in \Theta | g(\theta) = g\}$. For statistic maximal leakage $\Pi_{\mathcal{M},g}$, we have

$$\Pi_{\mathcal{M},g} = \sup_{\mathbf{f}_{\Theta}} \log \frac{\int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{G\Theta'}(g, \theta') d\theta'}{\sup_{g \in \mathbf{G}} \mathbf{f}_G(g)}.$$

Under a fixed distribution $\mathbf{f}_{\Theta}$, we can get that

$$\begin{aligned} & \log \frac{\int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{G\Theta'}(g, \theta') d\theta'}{\sup_{g \in \mathbf{G}} \mathbf{f}_G(g)} \\ &= \log \frac{\int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{\Theta'|G}(\theta'|g) \mathbf{f}_G(g) d\theta'}{\sup_{g \in \mathbf{G}} \mathbf{f}_G(g)} \\ &\leq \log \frac{\int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{\Theta'|G}(\theta'|g) \cdot \sup_{g \in \mathbf{G}} \mathbf{f}_G(g) d\theta'}{\sup_{g \in \mathbf{G}} \mathbf{f}_G(g)} \\ &= \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{\Theta'|G}(\theta'|g) d\theta' \\ &= \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \int_{\theta \in \Theta} \mathbf{f}_{\Theta'|G\Theta}(\theta'|g, \theta) \cdot \mathbf{f}_{\Theta|G}(\theta|g) d\theta d\theta' \\ &= \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \int_{\theta \in \Theta} \mathbf{f}_{\Theta'|G\Theta}(\theta'|g, \theta) \frac{\mathbf{f}_{G|\Theta}(g|\theta) \cdot \mathbf{f}_{\Theta}(\theta)}{\int_{\theta \in \Theta} \mathbf{f}_{G|\Theta}(g|\theta) \cdot \mathbf{f}_{\Theta}(\theta) d\theta} d\theta d\theta' \\ &= \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\int_{\theta \in \Theta_g} \mathbf{f}_{\Theta'|G\Theta}(\theta'|g, \theta) \cdot \mathbf{f}_{\Theta}(\theta) d\theta}{\int_{\theta \in \Theta_g} \mathbf{f}_{\Theta}(\theta) d\theta} d\theta' \\ &= \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{E}_{\Theta} [\mathbf{f}_{\Theta'|G\Theta}(\theta'|g, \theta) | \theta \in \Theta_g] d\theta' \\ &\triangleq V_{\mathbf{f}_{\Theta}}. \end{aligned} \quad (\text{A9})$$

This upper bound can be achieved when $\mathbf{f}_{\Theta}$ satisfies

$$\forall \theta' \in \Theta' : \arg \sup_{g \in \mathbf{G}} \mathbf{f}_{\Theta'|G}(\theta'|g) \cap \arg \sup_{g \in \mathbf{G}} \mathbf{f}_G(g) \neq \emptyset, \quad (\text{A10})$$

where $\mathbf{f}_{\Theta'|G}(\theta'|g) = \mathbb{E}_{\Theta} [\mathbf{f}_{\Theta'|G\Theta}(\theta'|g, \theta) | \theta \in \Theta_g]$, $\mathbf{f}_G(g) = \int_{\theta \in \Theta_g} \mathbf{f}_{\Theta}(\theta) d\theta$.

Let $\mathbf{f}_{\Theta}^+ \in \arg \sup_{\mathbf{f}_{\Theta}} V_{\mathbf{f}_{\Theta}}$ and $\Theta'(g)$ be the set of released parameters satisfying $g = \arg \sup_{g \in \mathbf{G}} \mathbb{E}_{\Theta} [\mathbf{f}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g]$ under $\mathbf{f}_{\Theta}^+$. We can construct a prior $\mathbf{f}_{\Theta}^*$ that satisfies

- $\mathbf{f}_G(g_1) = \mathbf{f}_G(g_2), \forall g_1, g_2 \in \mathbf{G}$;
- $\mathbf{f}_{\Theta}(\theta) \in \{0, \mathbf{f}_G(g)\}, \forall g \in \mathbf{G}, \theta \in \Theta_g$;
- $\mathbf{f}_{\Theta}(\theta) = 0$ if $\theta \notin \arg \sup_{\theta \in \Theta_g} \sum_{\theta' \in \Theta'(g)} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta), \forall g \in \mathbf{G}$.

We can get that

$$\begin{aligned} V_{\mathbf{f}_{\Theta}^+} &= \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{E}_{\Theta} [\mathbf{f}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g] d\theta' \\ &= \log \int_{g \in \mathbf{G}} \int_{\theta' \in \Theta'(g)} \mathbb{E}_{\Theta} [\mathbf{f}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g] d\theta' dg \\ &= \log \int_{g \in \mathbf{G}} \mathbb{E}_{\Theta} \left[\int_{\theta' \in \Theta'(g)} \mathbf{f}_{\Theta'|\Theta}(\theta'|\theta) d\theta' \middle| \theta \in \Theta_g \right] dg \\ &\leq \log \int_{g \in \mathbf{G}} \sup_{\theta \in \Theta_g} \int_{\theta' \in \Theta'(g)} \mathbb{P}_{\Theta'|\Theta}(\theta'|\theta) d\theta' dg \\ &= V_{\mathbf{f}_{\Theta}^*}. \end{aligned}$$

Since $\mathbf{f}_{\Theta}^+ \in \arg \sup_{\mathbf{f}_{\Theta}} V_{\mathbf{f}_{\Theta}}$, we can get that $\mathbf{f}_{\Theta}^* \in \arg \sup_{\mathbf{f}_{\Theta}} V_{\mathbf{f}_{\Theta}}$. Since $\mathbf{f}_{\Theta}^*$ satisfies condition (A10) and $\mathbf{f}_{\Theta}(\theta) \in \{0, \mathbf{f}_G(g)\}, \forall g \in \mathbf{G}, \theta \in \Theta_g$, we have

$$\begin{aligned} \Pi_{\mathcal{M}, \mathbf{g}} &= \sup_{\mathbf{f}_{\Theta}} \log \frac{\int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{G\Theta'}(g, \theta') d\theta'}{\sup_{g \in \mathbf{G}} \mathbf{f}_G(g)} \\ &= \sup_{\mathbf{f}_{\Theta}} \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{\Theta'|\Theta}(\theta'|\theta) d\theta' \\ &= \sup_{\mathbf{f}_{\Theta}} \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \frac{\int_{\theta \in \Theta_g} \mathbf{f}_{\Theta'|\Theta}(\theta'|\theta) \cdot \mathbf{f}_{\Theta}(\theta) d\theta}{\int_{\theta \in \Theta_g} \mathbf{f}_{\Theta}(\theta) d\theta} d\theta' \quad (\text{A11}) \\ &= \sup_{\mathbf{f}_{\Theta}} \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbb{E}_{\Theta} [\mathbf{f}_{\Theta'|\Theta}(\theta'|\theta) | \theta \in \Theta_g] d\theta' \\ &= \sup_{\mathbf{f}_{\Theta}: \mathbf{f}_{\Theta}(\theta) \in \{0, \mathbf{f}_G(g)\}, \forall g \in \mathbf{G}, \theta \in \Theta_g} \log \int_{\theta' \in \Theta'} \sup_{g \in \mathbf{G}} \mathbf{f}_{\Theta'|\Theta}(\theta'|\theta_g) d\theta', \end{aligned}$$

where θ_g satisfies $\mathbf{f}_{\Theta}(\theta_g) = \mathbf{f}_G(g)$. $\square$

References

1. Yamamoto, H. A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers (corresp.). *IEEE Trans. Inf. Theory* **1983**, *29*, 918–923.
2. Smith, G. On the foundations of quantitative information flow. In *Proceedings of the FoSSaCS*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 288–302.
3. Mario, S.A.; Chatzikokolakis, K.; Palamidessi, C.; Smith, G. Measuring information leakage using generalized gain functions. In *Proceedings of the CSF*; IEEE: New York, NY, USA, 2012.
4. 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.
5. Makhdoumi, A.; Salamatian, S.; Fawaz, N.; Médard, M. From the information bottleneck to the privacy funnel. In *Proceedings of the IEEE ITW 2014*; IEEE: New York, NY, USA, 2014; pp. 501–505.
6. Zamani, A.; Oechtering, T.J.; Skoglund, M. Bounds for privacy-utility trade-off with non-zero leakage. In *Proceedings of the IEEE ISIT 2022*; IEEE: New York, NY, USA, 2022; pp. 620–625.
7. Braun, C.; Chatzikokolakis, K.; Palamidessi, C. Quantitative notions of leakage for one-try attacks. *Electron. Notes Theor. Comput. Sci.* **2009**, *249*, 75–91.
8. Issa, I.; Wagner, A.B.; Kamath, S. An operational approach to information leakage. *IEEE Trans. Inf. Theory* **2019**, *66*, 1625–1657.

9. du Pin Calmon, F.; Makhdoumi, A.; Médard, M.; Varia, M.; Christiansen, M.; Duffy, K.R. Principal inertia components and applications. *IEEE Trans. Inf. Theory* **2017**, *63*, 5011–5038.
10. Shkel, Y.Y.; Blum, R.S.; Poor, H.V. Secrecy by design with applications to privacy and compression. *IEEE Trans. Inf. Theory* **2020**, *67*, 824–843.
11. Zamani, A.; Oechtering, T.J.; Skoglund, M. On the privacy-utility trade-off with and without direct access to the private data. *IEEE Trans. Inf. Theory* **2023**, *70*, 2177–2200.
12. Huang, T.H.; El Gamal, H. An efficient difference-of-convex solver for privacy funnel. In *Proceedings of the 2024 IEEE International Symposium on Information Theory Workshops (ISIT-W)*; IEEE: New York, NY, USA, 2024; pp. 1–6.
13. Shkel, Y.Y.; Poor, H.V. A compression perspective on secrecy measures. *IEEE J. Sel. Areas Inf. Theory* **2021**, *2*, 163–176.
14. Zamani, A.; Oechtering, T.J.; Gündüz, D.; Skoglund, M. Cache-aided private variable-length coding with zero and non-zero leakage. In *Proceedings of the 2023 21st International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*; IEEE: New York, NY, USA, 2023; pp. 247–254.
15. Zamani, A.; Oechtering, T.J.; Skoglund, M. Private Variable-Length Coding with Non-Zero Leakage. In *Proceedings of the 2023 IEEE International Workshop on Information Forensics and Security (WIFS)*; IEEE: New York, NY, USA, 2023; pp. 1–6.
16. Wang, S.; Lin, Z.; Fanti, G. Statistic Maximal Leakage. In *Proceedings of the 2024 IEEE International Symposium on Information Theory (ISIT)*; IEEE: New York, NY, USA, 2024; pp. 2742–2747.
17. Lin, Z.; Wang, S.; Sekar, V.; Fanti, G. Summary Statistic Privacy in Data Sharing. *IEEE J. Sel. Areas Inf. Theory* **2024**, *5*, 369–384. Special Issue: Information-Theoretic Methods for Trustworthy and Reliable Machine Learning.
18. Golden, R.M.; Henley, S.S.; White, H.; Kashner, T.M. Consequences of model misspecification for maximum likelihood estimation with missing data. *Econometrics* **2019**, *7*, 37.
19. Richardson, S.; Leblond, L. Some comments on misspecification of priors in Bayesian modelling of measurement error problems. *Stat. Med.* **1997**, *16*, 203–213.
20. Kleijn, B.J.; van der Vaart, A.W. Misspecification in infinite-dimensional Bayesian statistics. *Ann. Statist.* **2006**, *34*, 837–877.
21. Dwork, C.; McSherry, F.; Nissim, K.; Smith, A. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the TCC 2006, New York, NY, USA, 4–7 March 2006, Proceedings 3*; Springer: Berlin/Heidelberg, Germany, 2006; pp. 265–284.
22. Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H.B.; Mironov, I.; Talwar, K.; Zhang, L. Deep learning with differential privacy. In *Proceedings of the ACM CCS 2016*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 308–318.
23. Zhang, W.; Ohrimenko, O.; Cummings, R. Attribute privacy: Framework and mechanisms. In *Proceedings of the FACCT*; ACM: New York, NY, USA, 2022.
24. Kawamoto, Y.; Murakami, T. Local obfuscation mechanisms for hiding probability distributions. In *Proceedings of the Computer Security—ESORICS 2019, Luxembourg, 23–27 September 2019, Proceedings, Part I 24*; Springer: Berlin/Heidelberg, Germany, 2019; pp. 128–148.
25. Suri, A.; Evans, D. Formalizing and estimating distribution inference risks. *arXiv* **2021**, arXiv:2109.06024.
26. Wang, S.; Wei, R.; Ghassemi, M.; Kreacic, E.; Potluru, V.K. Guarding Multiple Secrets: Enhanced Summary Statistic Privacy for Data Sharing. *arXiv* **2024**, arXiv:2405.13804.
27. Saeidian, S.; Cervia, G.; Oechtering, T.J.; Skoglund, M. Pointwise maximal leakage. In *Proceedings of the IEEE ISIT 2022*; IEEE: New York, NY, USA, 2022; pp. 626–631.
28. Farokhi, F. Development and analysis of deterministic privacy-preserving policies using non-stochastic information theory. *IEEE Trans. Inf. Forensics Secur.* **2019**, *14*, 2567–2576.
29. Warner, S.L. Randomized response: A survey technique for eliminating evasive answer bias. *J. Am. Stat. Assoc.* **1965**, *60*, 63–69.
30. Chaudhuri, A.; Mukerjee, R. *Randomized Response: Theory and Techniques*; Routledge: Abingdon, UK, 2020.
31. Kairouz, P.; Oh, S.; Viswanath, P. Extremal mechanisms for local differential privacy. *J. Mach. Learn. Res.* **2016**, *17*, 492–542.
32. Wang, Y.; Wu, X.; Hu, D. Using Randomized Response for Differential Privacy Preserving Data Collection. In *Proceedings of the EDBT/ICDT Workshops*; CEUR: Bonn, Germany, 2016; Volume 1558, pp. 0090–6778.
33. Liao, J.; Kosut, O.; Sankar, L.; du Pin Calmon, F. Tunable measures for information leakage and applications to privacy-utility tradeoffs. *IEEE Trans. Inf. Theory* **2019**, *65*, 8043–8066.
34. Gilani, A.; Kurri, G.R.; Kosut, O.; Sankar, L. (α, β) -Leakage: A Unified Privacy Leakage Measure *arXiv* **2023**, arXiv:2304.07456.
35. Kurri, G.R.; Sankar, L.; Kosut, O. An Operational Approach to Information Leakage via Generalized Gain Functions **2022**, *70*, 1349–1375.
36. Cung, R.D.; Shkel, Y.Y.; Issa, I. Binary Maximal Leakage. In *Proceedings of the 2024 IEEE International Symposium on Information Theory (ISIT)*; IEEE: New York, NY, USA, 2024; pp. 2748–2753.
37. Calmon, F.P.; Makhdoumi, A.; Médard, M. Fundamental limits of perfect privacy. In *Proceedings of the ISIT*; IEEE: New York, NY, USA, 2015.
38. Rassouli, B.; Gündüz, D. On perfect privacy. *IEEE J. Sel. Areas Inf. Theory* **2021**, *2*, 177–191.

39. Biswas, S.; Palamidessi, C. PRIVIC: A privacy-preserving method for incremental collection of location data *arXiv* **2022**, arXiv:2206.10525.
40. Oya, S.; Troncoso, C.; Pérez-González, F. Back to the drawing board: Revisiting the design of optimal location privacy-preserving mechanisms. In *Proceedings of the 2017 ACM CCS*; ACM: New York, NY, USA, 2017; pp. 1959–1972.
41. Clark, D.; Hunt, S.; Malacaria, P. Quantitative analysis of the leakage of confidential data. *Electron. Notes Theor. Comput. Sci.* **2002**, *59*, 238–251.
42. Clark, D.; Hunt, S.; Malacaria, P. A static analysis for quantifying information flow in a simple imperative language. *J. Comput. Secur.* **2007**, *15*, 321–371.
43. Malacaria, P. Assessing security threats of looping constructs. In *Proceedings of the 2007 ACM POPL*; ACM: New York, NY, USA, 2007; pp. 225–235.
44. Rassouli, B.; Gündüz, D. Optimal utility-privacy trade-off with total variation distance as a privacy measure. *IEEE Trans. Inf. Forensics Secur.* **2019**, *15*, 594–603.
45. Alvim, M.S.; Chatzikokolakis, K.; McIver, A.; Morgan, C.; Palamidessi, C.; Smith, G. Additive and multiplicative notions of leakage, and their capacities. In *Proceedings of the IEEE CSF 2014*; IEEE: New York, NY, USA, 2014; pp. 308–322.
46. Asodeh, S.; Diaz, M.; Alajaji, F.; Linder, T. Privacy-aware guessing efficiency. In *Proceedings of the ISIT*; IEEE: New York, NY, USA, 2017.
47. Asodeh, S.; Diaz, M.; Alajaji, F.; Linder, T. Estimation efficiency under privacy constraints. *IEEE Trans. Inf. Theory* **2018**, *65*, 1512–1534.
48. Bordenabe, N.E.; Smith, G. Correlated secrets in quantitative information flow. In *Proceedings of the CSF*; IEEE: New York, NY, USA, 2016.
49. Jurado, M.; Gonze, R.G.; Alvim, M.S.; Palamidessi, C. Analyzing the Shuffle Model through the Lens of Quantitative Information Flow. *arXiv* **2023**, arXiv:2305.13075.
50. Jin, R.; He, X.; Dai, H. On the security-privacy tradeoff in collaborative security: A quantitative information flow game perspective. *IEEE Trans. Inf. Forensics Secur.* **2019**, *14*, 3273–3286.
51. Alvim, M.S.; Fernandes, N.; McIver, A.; Morgan, C.; Nunes, G.H. A novel analysis of utility in privacy pipelines, using Kronecker products and quantitative information flow. In *Proceedings of the 2023 ACM CCS*; ACM: New York, NY, USA, 2023; pp. 1718–1731.
52. Shannon, C.E. Coding theorems for a discrete source with a fidelity criterion. *Ire Int. Conv. Rec.* **1959**, *4*, 1.
53. Denning, D.E.R. *Cryptography and Data Security*; Addison-Wesley Reading: Reading, MA, USA, 1982; Volume 112.
54. Gray, J.W., III. Toward a mathematical foundation for information flow security. *J. Comput. Secur.* **1992**, *1*, 255–294.
55. Bhaskar, R.; Bhowmick, A.; Goyal, V.; Laxman, S.; Thakurta, A. Noiseless database privacy. In *Proceedings of the ASIACRYPT 2011, Seoul, South Korea, 4–8 December 2011, Proceedings 17*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 215–232.
56. Farokhi, F. Noiseless privacy: Definition, guarantees, and applications. *IEEE Trans. Big Data* **2021**, *9*, 51–62.
57. Farokhi, F.; Nair, G. Non-stochastic private function evaluation. In *Proceedings of the 2020 IEEE ITW*; IEEE: New York, NY, USA, 2021; pp. 1–5.
58. Nair, G.N. A nonstochastic information theory for communication and state estimation. *IEEE Trans. Autom. Control.* **2013**, *58*, 1497–1510.
59. Suri, A.; Lu, Y.; Chen, Y.; Evans, D. Dissecting Distribution Inference. In *Proceedings of the First IEEE Conference on Secure and Trustworthy Machine Learning*, Raleigh, NC, USA, 8–10 February 2023.
60. Bebensee, B. Local differential privacy: A tutorial. *arXiv* **2019**, arXiv:1907.11908.
61. Alvim, M.S.; Andrés, M.E.; Chatzikokolakis, K.; Degano, P.; Palamidessi, C. Differential privacy: on the trade-off between utility and information leakage. In *Proceedings of the Formal Aspects of Security and Trust: 8th International Workshop, FAST 2011, Leuven, Belgium, 12–14 September 2011, Revised Selected Papers 8*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 39–54.
62. Murakami, T.; Kawamoto, Y. {Utility-Optimized} local differential privacy mechanisms for distribution estimation. In *Proceedings of the 28th USENIX Security Symposium (USENIX Security 19)*; USENIX Association: Berkeley, CA, USA, 2019; pp. 1877–1894.
63. Ford, L.R. *Flows in Networks*; Torrossa: Fiesole, Italy, 2015.
64. Karp, R.M. *Reducibility Among Combinatorial Problems*; Springer: Berlin/Heidelberg, Germany, 2010.
65. Krumke, S.O.; Noltemeier, H.; Schwarz, S.; Wirth, H.C.; Ravi, R. Flow improvement and network flows with fixed costs. In *Proceedings of the Operations Research Proceedings 1998: Selected Papers of the International Conference on Operations Research Zurich, 31 August–3 September 1998*; Springer: Berlin/Heidelberg, Germany, 1999; pp. 158–167.
66. Dwork, C.; Rothblum, G.N.; Vadhan, S. Boosting and differential privacy. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*; IEEE: New York, NY, USA, 2010; pp. 51–60.
67. Dwork, C.; Roth, A. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* **2014**, *9*, 211–487.
68. Becker, B.; Kohavi, R. Adult. UCI Machine Learning Repository, 1996. Available online: <https://archive.ics.uci.edu/dataset/2/adult> (accessed on 31 May 2026).

69. Wu, B.; Wagner, A.B.; Suh, G.E. Optimal mechanisms under maximal leakage. In *Proceedings of the 2020 IEEE Conference on Communications and Network Security (CNS)*; IEEE: New York, NY, USA, 2020; pp. 1–6.
70. Lewis, H.R.; Garey, M.R.; Johnson, D.S. Computers and intractability. A guide to the theory of NP-completeness. *J. Symb. Log.* **1983**, *48*, 498–500.
71. Şuvak, Z.; Altinel, İ.K.; Aras, N. Minimum cost flow problem with conflicts. *Networks* **2021**, *78*, 421–442.

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.