

Privacy with Estimation Guarantees

Hao Wang, Lisa Vo, Flavio P. Calmon, Muriel Médard, Ken R. Duffy, Mayank Varia *

Abstract

We study the central problem in data privacy: how to share data with an analyst while providing both privacy and utility guarantees to the user that owns the data. In this setting, we present an estimation-theoretic analysis of the privacy-utility trade-off (PUT). Here, an analyst is allowed to reconstruct (in a mean-squared error sense) certain functions of the data (utility), while other private functions should not be reconstructed with distortion below a certain threshold (privacy). We demonstrate how chi-square information captures the fundamental PUT in this case and provide bounds for the best PUT. We propose a convex program to compute privacy-assuring mappings when the functions to be disclosed and hidden are known a priori and the data distribution is known. We derive lower bounds on the minimum mean-squared error of estimating a target function from the disclosed data and evaluate the robustness of our approach when an empirical distribution is used to compute the privacy-assuring mappings instead of the true data distribution. We illustrate the proposed approach through two numerical experiments.

Index Terms— Estimation, privacy-utility trade-off, minimum mean-squared error.

*This material is based upon work supported by the National Science Foundation under Grant No. CCF-1845852. This paper was presented in part at the 2014 52nd Annual Allerton Conference on Communication, Control, and Computing [1], and the 2017 55th Annual Allerton Conference on Communication, Control, and Computing [2]. H. Wang, L. Vo, and F. P. Calmon are with Harvard University (e-mail: hao_wang@g.harvard.edu; lisavo@college.harvard.edu; flavio@seas.harvard.edu). M. Médard is with the Massachusetts Institute of Technology (e-mail: medard@mit.edu). K. R. Duffy is with the Hamilton Institute, Maynooth University (e-mail: ken.duffy@nuim.ie). M. Varia is with Boston University (e-mail: varia@bu.edu).

1 Introduction

Data sharing and publishing is increasingly common within scientific communities [3], businesses [4], government operations [5], medical fields [6], and beyond. Data is usually shared with an application in mind, from which the data provider receives some utility. For example, when a user shares her movie ratings with a streaming service, she receives utility in the form of suggestions of new, interesting movie recommendations that fit her taste. As a second example, when a medical research group shares patient data, their aim is to enable a wider community of researchers and statisticians to learn patterns from that data. Utility is then gained through new scientific discoveries.

The disclosure of non-encrypted data incurs a privacy risk through unwanted inferences. In our previous examples, the streaming service may infer the user’s political preference (potentially deemed private by the user) from her movie ratings [7], or an insurance company may determine the identity of a patient within a medical dataset [6, 8, 9]. If privacy is a concern but the data has no immediate utility, then cryptographic methods suffice.

The dichotomy between privacy and utility has been widely studied by computer scientists, statisticians, and information theorists alike. While specific metrics and models vary among these communities, their desideratum is the same: to design mechanisms that perturb the data (or functions thereof) while achieving an acceptable privacy-utility trade-off (PUT). The feasibility of this goal depends on several factors including the chosen privacy and utility metric, as well as the topology and distribution of the data. The information-theoretic approach to privacy, and notably the results by Sankar *et al.* [10, 11], Issa *et al.* [12, 13], Asoodeh *et al.* [14, 15], Calmon *et al.* [16, 17], among others, seek to quantify the best possible PUT for *any* privacy mechanism. In those works, information-theoretic quantities, such as mutual information and maximal leakage [12, 13], have been used to characterize privacy, and bounds on the fundamental PUT were derived under assumptions on the distribution of the data. It is within this information-theoretic approach that the present work is inscribed.

Our aim is to characterize the fundamental limits of PUT from an estimation-theoretic perspective, and to design privacy-assuring mechanisms that provide estimation-theoretic guarantees. We use the *principal inertia components* (PICs) [16, 18–26] to formalize the privacy and utility constraints. The PICs quantify the minimum mean-squared error (MMSE) achievable for reconstructing both private and useful information from the disclosed data. We do not seek to claim that the estimation-based approach subsumes other privacy metrics, such as differential privacy [27]. Rather, our goal is to show that the MMSE viewpoint reveals an interesting facet of data disclosure which, in turn, can drive the design of privacy mechanisms used in practice.

In the remainder of this section, we present an overview of the paper and our main results, discuss related work, and introduce the notation adopted in the paper.

1.1 Overview and Main Contributions

Throughout this paper, we assume all random variables are discrete with finite support sets. We let S denote a private variable to be hidden (e.g., political preference) and X be a useful variable that depends on S (e.g., movie ratings). Our goal is to disclose a realization of a random variable Y , produced from X through a randomized mapping $P_{Y|X}$ called the *privacy-assuring mapping*. Here, S , X , and Y satisfy the Markov condition $S \rightarrow X \rightarrow Y$. We assume that an analyst will provide some utility based on an observation of Y (e.g., movie recommendations), while potentially trying to estimate S from Y . Denoting $[n] \triangleq \{1, \dots, n\}$, the support sets of S , X , and Y are $\mathcal{S} = [|\mathcal{S}|]$, $\mathcal{X} = [|\mathcal{X}|]$, and $\mathcal{Y} = [|\mathcal{Y}|]$, respectively.

In the sequel, we derive PUTs when both privacy and utility are measured in terms of the mean-squared error of reconstructing functions of S and X from an observation of Y . We analyze three related scenarios: (i) an *aggregate* setting, where certain functions of X can be, on average, reconstructed from the disclosed variable while controlling the MMSE of estimating functions of S and $P_{S,X}$ is known to the privacy mechanism designer, (ii) a *composite* setting, where specific functions of S and X have different privacy/utility reconstruction requirements and $P_{S,X}$ is known to the privacy mechanism designer, and (iii) a *restricted-knowledge* setting, where $P_{S,X}$ is unknown, but the correlation between a target function to be hidden and a set of functions which are known to be hard to infer from the disclosed variable is given. For the first two thrusts, we also analyze the robustness of privacy-assuring mappings designed using an empirical

estimate of $P_{S,X}$ computed from a finite number of samples. Next, we present the outline of the paper and a summary of our main contributions.

Aggregate PUTs

We start by studying the problem of limiting an untrusted party’s ability to estimate functions of S given an observation of Y , while controlling for the MMSE of reconstructing functions of X given Y . Here, privacy and utility are measured in terms of the χ^2 -information between S and Y and the χ^2 -information between X and Y , denoted by $\chi^2(S;Y)$ and $\chi^2(X;Y)$ (cf. (1)), respectively. We introduce the χ^2 -privacy-utility function in Section 3. Bounds of this function are presented in Theorem 2. In particular, the upper bound is cast in terms of the PICs of $P_{S,X}$ and provides an interpretation of the trade-off between privacy and utility that goes beyond simply using maximal correlation. We also prove that the upper bound is achievable in the high-privacy regime in Theorem 3.

Composite PUTs

χ^2 -based metrics guarantee privacy and utility in a uniform sense, capturing the aggregate mean-squared error of estimating *any* functions of the private and the useful variables. However, in many applications, specific functions of S and X that should be hidden/revealed are known *a priori*. This knowledge enables a more refined design of privacy-assuring mechanisms that specifically target these functions. We explore this finer-grained approach in Section 4, and propose a PIC-based convex program for computing privacy-assuring mappings within this setting. We demonstrate the practical feasibility of the convex programs through two numerical experiments in Section 7, deriving privacy-assuring mappings for a synthetic dataset and a real-world dataset. In the latter case, we approximate $P_{S,X}$ using its empirical distribution.

Restricted Knowledge of the Distribution

The aforementioned aggregate and composite PUTs require knowledge of the joint distribution $P_{S,X}$. In Section 5, we forgo this assumption, and study a simpler setting where $S = \phi(X)$ (i.e., the private variable is a function of the data) and the correlation between $\phi(X)$ and a set of functions (composed with the data) $\{\phi_j(X)\}_{j=1}^m$ is given. In practice, $\phi(X)$ may be a sensitive feature of the data X , and $\{\phi_j(X)\}_{j=1}^m$ is a collection of other features from which $\mathbb{E}[\phi(X)\phi_j(X)]$ can be accurately estimated.

Our goal here is to derive lower bounds on the MMSE of estimating a real-valued function of X , namely $\phi(X)$, from Y for any privacy-assuring mapping $P_{Y|X}$. These bounds are cast in terms of the MMSE of estimating $\phi_j(X)$ from Y and the correlation between $\phi(X)$ and $\{\phi_j(X)\}_{j=1}^m$. This leads to a converse result in Theorem 5: if the MMSE of estimating $\phi_j(X)$ from Y is large and $\phi(X)$ is strongly correlated with $\phi_j(X)$, then the MMSE of estimating $\phi(X)$ from Y will also be large and privacy is assured in an estimation-theoretic sense. The inverse result is straightforward: if $\phi(X)$ and $\phi_j(X)$ are strongly correlated and $\phi_j(X)$ can be reliably reconstructed from Y , then $\phi(X)$ can also be reliably estimated from Y . This intuitive trade-off is at the heart of the estimation-theoretic view of privacy, and demonstrates that no function of X can remain private whilst other strongly correlated functions are revealed through Y . The results in Section 5 make this intuition mathematically precise.

Finally, in Section 6 we investigate the resilience of privacy-assuring mappings when designed using an estimate of the distribution $P_{\hat{S},\hat{X}}$ computed as the empirical frequencies of S, X obtained from n i.i.d. samples. Here, the value of the privacy and utility guarantees estimated using $P_{\hat{S},\hat{X}}$ will not match the true values $\chi^2(S;Y)$ and $\chi^2(X;Y)$ obtained when the privacy-assuring mechanism is applied to fresh samples drawn from the true distribution $P_{S,X}$. We bound this performance gap in Theorem 6 and show that this gap scales as $O(\sqrt{1/n})$, while also depending on the alphabet size of the variables and the probability of the least likely symbols.

1.2 Related Work

Currently, the most adopted definition of privacy is differential privacy [27, 28], which enables queries to be computed over a database while simultaneously ensuring privacy of individual entries of the database.

Information-theoretic quantities, such as Rényi divergence, can be used to relax the definition of differential privacy [29]. Fundamental bounds on composition of differentially private mechanisms were given by Kairouz *et al.* [30]. Recently, a new privacy framework called Pufferfish [31] was developed for creating customized privacy definitions.

Several papers, such as Sankar *et al.* [10], Calmon and Fawaz [17], Asoodeh *et al.* [32], and Makhdoumi *et al.* [33], have studied information disclosure with privacy guarantees through an information-theoretic lens. For example, Sankar *et al.* [10] characterized PUTs in large databases using tools from rate-distortion theory. Calmon and Fawaz [17] used expected distortion and mutual information to measure utility and privacy, respectively, and characterized the PUT as an optimization problem. Makhdoumi *et al.* [33] introduced the privacy funnel, where both privacy and utility are measured in terms of mutual information, and showed its connection with the information bottleneck [34]. The PUT was also explored in [35] and [36] using mutual information as a privacy metric.

Other quantities from the information-theoretic literature have been used to quantify privacy and utility. For example, Asoodeh *et al.* [14] and Calmon *et al.* [16] used estimation-theoretic tools to characterize fundamental limits of privacy. Liao *et al.* [37, 38] explored the PUT within a hypothesis testing framework. Issa *et al.* [12, 39] introduced maximal leakage as an information leakage metric. There is also significant recent work in information-theoretic privacy in the context of network secrecy. For example, Li and Oechtering [40] proposed a new privacy metric based on distributed Bayesian detection which can inform privacy-aware system design. Recently, Tripathy *et al.* [41] and Huang *et al.* [42] used adversarial networks for designing privacy-assuring mappings that navigate the PUT. Takbiri *et al.* [43] considered obfuscation and anonymization techniques and characterized the conditions required to obtain perfect privacy.

MMSE-based analysis and maximal correlation have been investigated in the context of log-Sobolev inequalities and hypercontractivity, such as in the work of Raginsky [44], Anantharam *et al.* [45], and Polyanskiy and Wu [46]. The metric used in this paper, namely χ^2 -information, relates with χ^2 -divergence, which is a special case of f -divergence [47]. Also of note, the study of robustness of estimated distributions with finite sample size has appeared in [48–51].

1.3 Notation

Matrices are denoted in bold capital letters (e.g., \mathbf{P}) and vectors in bold lower-case letters (e.g., \mathbf{p}). For a vector \mathbf{p} , $\text{diag}(\mathbf{p})$ is defined as the matrix with diagonal entries equal to \mathbf{p} and all other entries equal to 0. The span of a set \mathcal{V} of vectors is

$$\text{span}(\mathcal{V}) \triangleq \left\{ \sum_{i=1}^k \lambda_i \mathbf{v}_i \mid k \in \mathbb{N}, \mathbf{v}_i \in \mathcal{V}, \lambda_i \in \mathbb{R} \right\}.$$

The dimension of a linear span is denoted by $\text{dim}(\text{span}(\mathcal{V}))$.

We denote independence of random variables U and V by $U \perp\!\!\!\perp V$, and write $U \sim V$ to indicate that U and V have the same distribution. When U , V , and W form a Markov chain, we write $U \rightarrow V \rightarrow W$. For a random variable U with probability distribution P_U , we denote

$$P_{U \min} \triangleq \inf\{P_U(u) \mid u \in \mathcal{U}\},$$

where \mathcal{U} is the support set of U . The MMSE of estimating U given V is

$$\text{mmse}(U|V) \triangleq \min_{U \rightarrow V \rightarrow \hat{U}} \mathbb{E} \left[(U - \hat{U})^2 \right] = \mathbb{E} \left[(U - \mathbb{E}[U|V])^2 \right].$$

The χ^2 -information between two random variables U and V is defined as

$$\chi^2(U; V) \triangleq \mathbb{E} \left[\left(\frac{P_{U,V}(U, V)}{P_U(U)P_V(V)} \right) \right] - 1. \quad (1)$$

Let P_U and Q_U be two probability distributions taking values in the same discrete and finite set \mathcal{U} . We denote $\|P_U - Q_U\|_1 \triangleq \sum_{u \in \mathcal{U}} |P_U(u) - Q_U(u)|$. For any real-valued random variable U , we denote the \mathcal{L}_p -norm of U as

$$\|U\|_p \triangleq (\mathbb{E}[|U|^p])^{1/p}.$$

The set of all functions that applied to a random variable U with distribution P_U result in an \mathcal{L}_2 -norm less than or equal to 1 is given by

$$\mathcal{L}_2(P_U) \triangleq \{f : \mathcal{U} \rightarrow \mathbb{R} \mid \|f(U)\|_2 \leq 1\}. \quad (2)$$

The conditional expectation operators $T_{V|U} : \mathcal{L}_2(P_V) \rightarrow \mathcal{L}_2(P_U)$ and $T_{U|V} : \mathcal{L}_2(P_U) \rightarrow \mathcal{L}_2(P_V)$ are given by $(T_{V|U}g)(u) \triangleq \mathbb{E}[g(V)|U = u]$ and $(T_{U|V}f)(v) \triangleq \mathbb{E}[f(U)|V = v]$, respectively.

2 Principal Inertia Components

We present next the properties of the PICs that will be used in this paper. For a more detailed overview, we refer the reader to [16] and the references therein. We use the definition of PICs presented in [16], but note that the PICs predate [16] by many decades (e.g., [18–24]). Recently, Huang *et al.* [52] considered the PICs by analyzing the “divergence transition matrix” [52, Eq. 2]. Specifically, there are different directions of local perturbation [53] of input distribution and the direction which leads to the greatest influence of the output distribution of a noisy channel can be identified [52] by specifying the singular vector decomposition of the divergence transition matrix. In follow-on work, Huang *et al.* [54] used the divergence transition matrix in the context of feature selection. The singular values of the divergence transition matrix are exactly the square root of the PICs considered here, and are also related to the singular values of the conditional expectation operator, as also noted by Makur and Zheng [26] and originally by Witsenhausen [21] and others [24]. We build on these prior works by using the PICs for quantifying privacy-utility trade-offs.

Definition 1 ([16, Definition 1]). Let U and V be random variables with support sets \mathcal{U} and \mathcal{V} , respectively, and joint distribution $P_{U,V}$. In addition, let $f_0 : \mathcal{U} \rightarrow \mathbb{R}$ and $g_0 : \mathcal{V} \rightarrow \mathbb{R}$ be the constant functions $f_0(u) = 1$ and $g_0(v) = 1$. For $k \in \mathbb{Z}_+$, we (recursively) define

$$\lambda_k(U; V) \triangleq \mathbb{E}[f_k(U)g_k(V)]^2, \quad (3)$$

where

$$(f_k, g_k) \triangleq \operatorname{argmax} \left\{ \mathbb{E}[f(U)g(V)]^2 \mid \begin{array}{l} f \in \mathcal{L}_2(P_U), g \in \mathcal{L}_2(P_V), \mathbb{E}[f(U)f_j(U)] = 0, \\ \mathbb{E}[g(V)g_j(V)] = 0, j \in \{0, \dots, k-1\} \end{array} \right\}. \quad (4)$$

The values $\lambda_k(U; V)$ are called the *principal inertia components* (PICs) of $P_{U,V}$. The functions f_k and g_k are called the *principal functions* of $P_{U,V}$.

Observe that the PICs satisfy $\lambda_k(U; V) \leq 1$, since $f_k \in \mathcal{L}_2(P_U)$, $g_k \in \mathcal{L}_2(P_V)$, and

$$|\mathbb{E}[f(U)g(V)]| \leq \|f(U)\|_2 \|g(V)\|_2 \leq 1.$$

Thus, from Definition 1, $0 \leq \lambda_{k+1}(U; V) \leq \lambda_k(U; V) \leq 1$.

The largest PIC satisfies $\lambda_1(U; V) = \rho_m(U; V)^2$ where $\rho_m(U; V)$ is the maximal correlation [23], defined as

$$\rho_m(U; V) \triangleq \max_{\substack{\mathbb{E}[f(U)] = \mathbb{E}[g(V)] = 0 \\ \mathbb{E}[f(U)^2] = \mathbb{E}[g(V)^2] = 1}} \mathbb{E}[f(U)g(V)]. \quad (5)$$

Definition 2 ([16, Definition 2]). For $\mathcal{U} = [m]$ and $\mathcal{V} = [n]$, let $\mathbf{P}_{U,V} \in \mathbb{R}^{m \times n}$ be a matrix with entries $[\mathbf{P}_{U,V}]_{i,j} = P_{U,V}(i, j)$, and $\mathbf{D}_U \in \mathbb{R}^{m \times m}$ and $\mathbf{D}_V \in \mathbb{R}^{n \times n}$ be diagonal matrices with diagonal entries $[\mathbf{D}_U]_{i,i} = P_U(i)$ and $[\mathbf{D}_V]_{j,j} = P_V(j)$, respectively, where $i \in [m]$ and $j \in [n]$. We define

$$\mathbf{Q}_{U,V} \triangleq \mathbf{D}_U^{-1/2} \mathbf{P}_{U,V} \mathbf{D}_V^{-1/2}. \quad (6)$$

We denote the singular value decomposition of $\mathbf{Q}_{U,V}$ by $\mathbf{Q}_{U,V} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$.

Definition 3 ([16, Definition 14]). Let $d \triangleq \min\{|\mathcal{U}|, |\mathcal{V}|\} - 1$, and $\lambda_d(U; V)$ the d -th PIC of $P_{U,V}$. We define

$$\delta(P_{U,V}) \triangleq \begin{cases} \lambda_d(U; V) & \text{if } |\mathcal{V}| \leq |\mathcal{U}|, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

We also denote $\lambda_d(U; V)$ and the corresponding principal functions f_d, g_d as $\lambda_{\min}(U; V)$ and f_{\min}, g_{\min} , respectively, when the alphabet size is clear from the context.

The next theorem illustrates the different characterizations of the PICs used in this paper.

Theorem 1 ([16, Theorem 1]). *The following characterizations of the PICs are equivalent:*

1. *The characterization given in Definition 1, where, for f_k and g_k given in (4), $g_k(V) = \frac{\mathbb{E}[f_k(U)|V]}{\|\mathbb{E}[f_k(U)|V]\|_2}$ and $f_k(U) = \frac{\mathbb{E}[g_k(V)|U]}{\|\mathbb{E}[g_k(V)|U]\|_2}$.*
2. *For any $k \in \mathbb{Z}_+$,*

$$1 - \lambda_k(U; V) = \text{mmse}(h_k(U)|V), \quad (8)$$

where

$$h_k \triangleq \text{argmin} \left\{ \text{mmse}(h(U)|V) \mid \|h(U)\|_2 = 1, \mathbb{E}[h(U)h_j(U)] = 0, j \in \{0, \dots, k-1\} \right\}. \quad (9)$$

If $\lambda_k(U; V)$ is unique, then $h_k = f_k$, given in (4).

3. *$\sqrt{\lambda_k(U; V)}$ is the $(k+1)$ -st largest singular value of $\mathbf{Q}_{U,V}$. The principal functions f_k and g_k in (4) correspond to the columns of the matrices $\mathbf{D}_U^{-1/2}\mathbf{U}$ and $\mathbf{D}_V^{-1/2}\mathbf{V}$, respectively, where $\mathbf{Q}_{U,V} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$.*

The equivalent characterizations of the PICs in the above theorem have the following intuitive interpretation: the principal functions can be viewed as a basis that decompose the mean-squared error of estimating functions of a hidden variable U given an observation V . In particular, for any zero-mean finite-variance function $f: \mathcal{U} \rightarrow \mathbb{R}$,

$$\text{mmse}(f(U)|V) = \sum_{i=1}^{|\mathcal{U}|-1} \mathbb{E}[f(U)f_i(U)]^2 (1 - \lambda_i(U; V)).$$

We remark that the χ^2 -information between U and V is the sum of all PICs. Specifically, it has been shown (e.g., [16, 21]) that $\chi^2(U; V) = \sum_{i=1}^d \lambda_i(U; V)$, where $d = \min\{|\mathcal{U}|, |\mathcal{V}|\} - 1$.

3 Aggregate PUTs: The Chi-Square-Privacy-Utility Function

We start our analysis by adopting χ^2 -information as a measure of both privacy and utility. As seen in the previous section, $\chi^2(S; Y) = \sum_{i=1}^d \lambda_i(S; Y)$, where $d = \min\{|\mathcal{S}|, |\mathcal{Y}|\} - 1$. If $\chi^2(S; Y) < 1$, then, from characterization 2 in Theorem 1, the MMSE of reconstructing *any* zero-mean, unit-variance function of S given Y is lower bounded by $1 - \chi^2(S; Y)$, i.e., all functions of S cannot be reconstructed with small MMSE given an observation of Y . Note that this argument also holds true when we replace χ^2 -information with the maximal correlation. In fact, in the high privacy regime, the PUT under χ^2 -information is essentially equivalent to the PUT when both privacy and utility are measured using maximal correlation. We make this intuition precise at the end of this section. When $1 \leq \chi^2(S; Y)$, certain private functions, on average, may be estimated from Y but, in general, most private functions are still kept in secret. Analogously, when $\chi^2(X; Y)$ is large, certain functions of X can be, on average, reconstructed (i.e., estimated) with small MMSE from Y . We demonstrate next that the PICs play a central role in bounding the PUT in this regime.

We first introduce the χ^2 -privacy-utility function. This function captures how well an analyst can reconstruct functions of the useful variable X while restricting the analyst's ability to estimate functions of the private variable S .

Definition 4. For a given joint distribution $P_{S,X}$ and $0 \leq \epsilon \leq \chi^2(S; X)$, we define the χ^2 -privacy-utility (trade-off) function as

$$F_{\chi^2}(\epsilon; P_{S,X}) \triangleq \sup_{P_{Y|X} \in \mathcal{D}(\epsilon; P_{S,X})} \chi^2(X; Y),$$

where $\mathcal{D}(\epsilon; P_{S,X}) \triangleq \{P_{Y|X} \mid S \rightarrow X \rightarrow Y, \chi^2(S; Y) \leq \epsilon\}$.

It has been proved in [55, 56] that there is always a privacy-assuring mapping $P_{Y|X}$ which achieves the supremum in $F_{\chi^2}(\epsilon; P_{S,X})$ using at most $|\mathcal{X}|+1$ symbols (i.e., $|\mathcal{Y}| \leq |\mathcal{X}|+1$). The following lemma gives an alternative way to compute the χ^2 -information, in the discrete, finite setting.

Lemma 1. *Suppose $S \rightarrow X \rightarrow Y$. Then*

$$\chi^2(X; Y) = \text{tr}(\mathbf{A}) - 1, \quad (10)$$

$$\chi^2(S; Y) = \text{tr}(\mathbf{B}\mathbf{A}) - 1, \quad (11)$$

where, using (6),

$$\mathbf{A} \triangleq \mathbf{Q}_{X,Y} \mathbf{Q}_{X,Y}^T, \quad \mathbf{B} \triangleq \mathbf{Q}_{S,X}^T \mathbf{Q}_{S,X}.$$

Proof. See Appendix A.1. □

The following lemma characterizes some properties of the χ^2 -privacy-utility function.

Lemma 2. *For a given joint distribution $P_{S,X}$, the χ^2 -privacy-utility function $F_{\chi^2}(\epsilon; P_{S,X})$ is a concave function in ϵ . Furthermore, $\epsilon \rightarrow \frac{1}{\epsilon} F_{\chi^2}(\epsilon; P_{S,X})$ is a non-increasing mapping.*

Proof. See Appendix A.2. □

The χ^2 -privacy-utility function has a simple upper bound,

$$F_{\chi^2}(\epsilon; P_{S,X}) \leq \epsilon + |\mathcal{X}| - 1 - \chi^2(S; X), \quad (12)$$

which follows immediately from the data-processing inequality:

$$\chi^2(S; X) + \chi^2(X; Y) \leq \chi^2(S; Y) + \chi^2(X; X). \quad (13)$$

We derive an upper bound for the χ^2 -privacy-utility function that significantly improves (12) by using properties of the PICs. The bound is piecewise linear, where each piece has a slope given in terms of a PIC of $P_{S,X}$. Intuitively, this bound corresponds to the privacy-assuring mapping $P_{Y|X}$ that achieves the best PUT if $P_{Y|X}$ was not constrained to be non-negative. We also provide a lower bound that follows directly from the concavity of the χ^2 -privacy-utility function. These bounds are illustrated in Fig. 1.

Definition 5. For $t_i \in [0, 1]$ ($i \in [n]$), $0 \leq \epsilon \leq \sum_{i \in [n]} t_i$, and $n \leq m$, $G_\epsilon^m(t_1, \dots, t_n)$ is defined as

$$G_\epsilon^m(t_1, \dots, t_n) \triangleq \max \left\{ \sum_{i=1}^m x_i \mid (x_1, \dots, x_m) \in \mathcal{D}_\epsilon^m(t_1, \dots, t_n) \right\},$$

where

$$\mathcal{D}_\epsilon^m(t_1, \dots, t_n) \triangleq \left\{ (x_1, \dots, x_m) \mid \sum_{i=1}^n t_i x_i \leq \epsilon, x_i \in [0, 1], i \in [m] \right\}.$$

For fixed m and t_i ($i \in [n]$), $G_\epsilon^m(t_1, \dots, t_n)$ is a piecewise linear function with respect to ϵ and can be expressed in closed-form (cf. Appendix A.3).




Figure 1: Piecewise linear upper bound and lower bound for the χ^2 -privacy-utility function when $\delta(P_{S,X})$, defined in (7), is positive.

Theorem 2. For the χ^2 -privacy-utility function $F_{\chi^2}(\epsilon; P_{S,X})$ introduced in Definition 4 and $\epsilon \in [0, \chi^2(S; X)]$,

$$\frac{|\mathcal{X}|-1}{\chi^2(S; X)}\epsilon \leq F_{\chi^2}(\epsilon; P_{S,X}) \leq G_{\epsilon}^{|\mathcal{X}|-1}(\lambda_1(S; X), \dots, \lambda_d(S; X)),$$

where $d \triangleq \min\{|S|, |\mathcal{X}|\} - 1$ and $\lambda_1(S; X), \dots, \lambda_d(S; X)$ are the PICs of $P_{S,X}$.

Proof. See Appendix A.4. □

Remark 1. The upper bound for the χ^2 -privacy-utility function given in Theorem 2 can also be proved by, for example, combining Theorem 4 in [57] with properties of the PICs.

We now illustrate the piecewise linear upper bound. Recall that the PIC decomposition of $P_{S,X}$ results in a set of basis functions $\mathcal{P} \triangleq \{f_1(S), \dots, f_d(S)\}$, with corresponding MMSE estimators $\mathcal{U} \triangleq \{g_1(X), \dots, g_d(X)\}$. Consider the following intuition for designing a sequence of privacy-assuring mappings. The first mapping enables the function $g_d(X)$ to be reliably estimated from Y while keeping all other functions in \mathcal{U} secret. In this case, the utility is one, since exactly one zero-mean, unit-variance function of X can be recovered from Y . The privacy leakage is $\lambda_d(S; X)$, since using $g_d(X)$ to estimate the private function $f_d(S)$ has mean-squared error $1 - \lambda_d(S; X)$. Following the same procedure, the second privacy-assuring mapping allows only $g_d(X)$ and $g_{d-1}(X)$ to be recovered from the disclosed variable and so on. This sequence of privacy-assuring mappings corresponds to the breakpoints of the upper bound. Note that such privacy-assuring mappings may not be feasible — hence the upper bound.

Note that $F_{\chi^2}(0; P_{S,X})$ characterizes the maximal aggregate MMSE of estimating useful functions while guaranteeing perfect privacy. Here perfect privacy means that no zero-mean, unit-variance function of S can be reconstructed from Y . If the value of $F_{\chi^2}(0; P_{S,X})$ is known, a better lower bound can be obtained from the concavity of $F_{\chi^2}(\epsilon; P_{S,X})$ as

$$\frac{|\mathcal{X}|-1 - F_{\chi^2}(0; P_{S,X})}{\chi^2(S; X)}\epsilon + F_{\chi^2}(0; P_{S,X}) \leq F_{\chi^2}(\epsilon; P_{S,X}). \quad (14)$$

When $S = X$, then $\chi^2(S; X) = |\mathcal{X}|-1$ and $F_{\chi^2}(\epsilon; P_{S,X}) = \epsilon$. Following from Definition 5 and noticing that all PICs of $P_{S,X}$ are 1, the upper bound and the lower bound for the χ^2 -privacy-utility function in

Theorem 2 are both ϵ , which is equal to $F_{\chi^2}(\epsilon; P_{S,X})$. In this sense, the upper bound and lower bound given in Theorem 2 are sharp. We investigate the tightness of the upper bound through numerical example in Section 7.1.

The following corollary of Lemma 2 and Theorem 2 shows that the χ^2 -privacy-utility function is strictly increasing with respect to ϵ .

Corollary 1. *For a given joint distribution $P_{S,X}$, the mapping $\epsilon \rightarrow F_{\chi^2}(\epsilon; P_{S,X})$ is strictly increasing for $\epsilon \in [0, \chi^2(S; X)]$.*

Proof. See Appendix A.5. □

We denote

$$\partial\mathcal{D}(\epsilon; P_{S,X}) \triangleq \{P_{Y|X} \mid S \rightarrow X \rightarrow Y, \chi^2(S; Y) = \epsilon\}.$$

By Corollary 1, $F_{\chi^2}(\epsilon; P_{S,X})$ is strictly increasing. Therefore,

$$F_{\chi^2}(\epsilon; P_{S,X}) = \max_{P_{Y|X} \in \partial\mathcal{D}(\epsilon; P_{S,X})} \chi^2(X; Y). \quad (15)$$

By Corollary 7 in [16], when $\delta(P_{S,X}) = 0$, defined in (7), then $F_{\chi^2}(0; P_{S,X}) > 0$ (i.e., there exists a privacy-assuring mapping that allows the disclosure of a non-trivial amount of useful functions while guaranteeing perfect privacy). On the other hand, when $\delta(P_{S,X}) > 0$, then $F_{\chi^2}(0; P_{S,X}) = 0$. The following theorem shows that when $\delta(P_{S,X}) > 0$, the upper bound of $F_{\chi^2}(\epsilon; P_{S,X})$ in Theorem 2 is achievable around zero, implying that the upper bound is tight around zero. The proof of this theorem also provides a specific way to construct an optimal privacy-assuring mapping (i.e., achieves the upper bound in Theorem 2).

Theorem 3. *Suppose $\delta(P_{S,X}) > 0$ and $P_{X \min} > 0$. Then there exists Y such that $S \rightarrow X \rightarrow Y$, $\chi^2(X; Y) = P_{X \min}$ and $\chi^2(S; Y) = P_{X \min} \lambda_{\min}(S; X)$.*

Proof. See Appendix A.6. □

When $\delta(P_{S,X}) > 0$ and $P_{X \min} > 0$, then $F_{\chi^2}(\hat{\epsilon}; P_{S,X}) = P_{X \min}$ where $\hat{\epsilon} = P_{X \min} \lambda_{\min}(S; X)$. Since $(\hat{\epsilon}, P_{X \min})$ is a point on the upper bound of the χ^2 -privacy-utility function given in Theorem 2, Theorem 3 shows that, in this case, the upper bound is achievable in the high-privacy region. We remark that the local behavior of privacy-utility functions in high-privacy region and high-utility region has been studied in the context of strong data processing inequalities (e.g., [58, 59] and the references therein).

Connections with Maximal Correlation

Maximal correlation has previously been considered as a privacy measure in [14, 15, 60, 61]. In particular, it has been proved [61] that when $\rho_m(S; Y)$ is small, then $\Pr(S \neq \hat{S})$ can be lower bounded for any $\hat{S} = h(Y)$. We show in Corollary 2 that, in the high privacy regime, the privacy-utility function under maximal correlation possesses similar properties to $F_{\chi^2}(\epsilon; P_{S,X})$. However, when $\rho_m(S; Y)$ is large, say $\rho_m(S; Y) = 1$, it is unclear whether one private function or several private functions can be recovered from the disclosed variable. In contrast, χ^2 -information can distinguish between these two cases and quantifies how many private functions, on average, can be reconstructed from the disclosed variable. For example, a user might be comfortable revealing that his/her age is above a certain threshold, but not the age itself. In this case, the privacy leakage measured by maximal correlation is one since there is a function of age which can be recovered from the disclosed variable. Thus, maximal correlation cannot distinguish between the cases where only one function of S and S itself can be estimated from the disclosed data. We will revisit this example in the next section and show how to design privacy-assuring mappings using PICs which target specific private functions and useful functions. Finally, we provide an example showing the limitation of maximal correlation as a utility measure.

Definition 6. For a given joint distribution $P_{S,X}$ and $0 \leq \epsilon \leq \rho_m(S; X)$, we define the maximal-correlation-privacy-utility (trade-off) function as

$$F_{\rho_m}(\epsilon; P_{S,X}) \triangleq \sup_{P_{Y|X} \in \mathcal{D}_{\rho_m}(\epsilon; P_{S,X})} \rho_m(X; Y),$$

where $\mathcal{D}_{\rho_m}(\epsilon; P_{S,X}) \triangleq \{P_{Y|X} \mid S \rightarrow X \rightarrow Y, \rho_m(S; Y) \leq \epsilon\}$.

The next corollary follows from the same proof techniques used in Theorem 2 and Theorem 3.

Corollary 2. For a given joint distribution $P_{S,X}$ and $\epsilon \in [0, \rho_m(S; X)]$, if $\delta(P_{S,X}) > 0$, then $F_{\rho_m}(\epsilon; P_{S,X}) \leq \epsilon / \sqrt{\lambda_{\min}(S; X)}$. Furthermore, if $P_{X \min} > 0$, then there exists Y such that $S \rightarrow X \rightarrow Y$, $\rho_m(X; Y) = \sqrt{P_{X \min}}$ and $\rho_m(S; Y) = \sqrt{P_{X \min} \lambda_{\min}(S; X)}$.

We illustrate the limitation of the maximal correlation as a utility measure through the following example.

Example 1. Let $\mathcal{S} = \{-1, 1\}^n$ and $\mathcal{X} = \{-1, 1\}^n$, and X^n be the result of passing S^n through a memoryless binary symmetric channel with crossover probability $\epsilon < 1/2$. We assume that S^n is composed of n uniform and i.i.d. bits. For $\mathcal{A} \subseteq [n]$, let $Y = \prod_{i \in \mathcal{A}} X_i$. In this case, one can show that $\rho_m(S^n; Y) = (1 - 2\epsilon)^{|\mathcal{A}|}$ and $\rho_m(X^n; Y) = 1$. If $|\mathcal{A}|$ is an increasing function of n , then $\rho_m(S^n; Y) \rightarrow 0$ as $n \rightarrow \infty$. In other words, we can disclose a function of X^n achieving nearly perfect privacy and utility as measured by $\rho_m(S^n; Y)$ and $\rho_m(X^n; Y)$, respectively, with large $|\mathcal{A}|$ and n . However, as n increases, the basis of functions in $\mathcal{L}_2(P_{X^n})$ will increase exponentially, and revealing only one function may not be enough for achieving utility. The crux of the limitation is that maximal correlation only takes into account the most reliably estimated function. The χ^2 -information overcomes this limitation by capturing all possible real-valued functions of X^n that can be recovered from Y . In particular, if $\chi^2(X^n; Y) = |\mathcal{X}| - 1$, then all zero-mean finite-variance functions of X^n can be reconstructed from Y . We will revisit this example again in Section 5 and Section 7.

4 Composite PUTs:

A Convex Program for Computing Privacy-Assuring Mappings

In the previous section, we studied χ^2 -based metrics for both privacy and utility. The optimization problem in the definition of χ^2 -privacy-utility function (Definition 4) is non-convex. Next, we provide a convex program for designing privacy-assuring mappings by adding more stringent constraints on privacy and utility.

More specifically, we explore an alternative, finer-grained approach for measuring both privacy and utility based on PICs (recall that χ^2 -information is the sum of all PICs). This approach has a practical motivation, since oftentimes there are specific well-defined features (functions) of the data (realizations of a random variable) that should be hidden or disclosed. For example, a user may be willing to disclose that they prefer documentaries over action movies, but not exactly which documentary they like. More abstractly, we consider the case where certain known functions should be disclosed (utility), whereas others should be hidden (privacy). This is a finer-grained setting than the one used in the last section, since χ^2 -information captures the aggregate reconstruction error across all zero-mean, unit-variance functions.

We denote the set of functions to be disclosed as

$$\mathcal{U}(X) \triangleq \{u_i : \mathcal{X} \rightarrow \mathbb{R} \mid \mathbb{E}[u_i(X)] = 0, \|u_i(X)\|_2 = 1, i \in [n]\},$$

and the set of functions to be hidden as

$$\mathcal{P}(S) \triangleq \{s_i : \mathcal{S} \rightarrow \mathbb{R} \mid \mathbb{E}[s_i(S)] = 0, \|s_i(S)\|_2 = 1, i \in [m]\}.$$

Our goal is to find the privacy-assuring mapping $P_{Y|X}$ such that $S \rightarrow X \rightarrow Y$ and Y satisfies the following privacy-utility constraints:

1. **Utility constraints:** $\max\{\text{mmse}(u_i(X)|Y)\}_{i \in [n]} \leq \Delta$ and $X \sim Y$.
2. **Privacy constraints:** $\text{mmse}(s_i(S)|Y) \geq \theta_i, i \in [m]$.

Note that the utility constraint $X \sim Y$ implies that the disclosed variable follows the same distribution as the useful variable. The practical motivation for adding this constraint is to enable Y to preserve overall population statistics about X , while hiding information about individual samples. This assumption also enables the problem of finding the optimal privacy-assuring mapping to be formulated as a convex program, described next.

We follow two steps – projection¹ and optimization – to find the privacy-assuring mapping. Private functions are projected to a new set of functions based on the useful variable in the first step. Then a PIC-based convex program is proposed in order to find the privacy-assuring mapping.

4.1 Projection

As a first step, we project (i.e., compute the conditional expectation) all private functions to the useful variable and obtain a new set of functions:

$$\mathcal{P}(X) \triangleq \left\{ \hat{s}_i(x) \triangleq \frac{\mathbb{E}[s_i(S)|X=x]}{\|\mathbb{E}[s_i(S)|X]\|_2} \mid i \in [m] \right\}.$$

It is worth noting that, after the projection, the obtained privacy-assuring mapping may not be an optimal solution to the original problem since the privacy constraints become stricter (see Lemma 3). Nonetheless, the advantage of this projection is twofold. First, it can significantly simplify the optimization program, since after the projection all functions are cast in terms of the useful variable alone. Second, the private variable is not needed as an input to the optimization after the projection. Therefore, the party that solves the optimization does not need access to the private data directly, further guaranteeing the safety of the sensitive information. The following lemma proves that privacy guarantees cast in terms of the projected functions still hold for the original functions.

Lemma 3. *Assume $S \rightarrow X \rightarrow Y$. For any function $f : \mathcal{S} \rightarrow \mathbb{R}$, if $\mathbb{E}[f(S)] = 0$ and $\|\mathbb{E}[f(S)|X]\|_2 \neq 0$, we have $\mathbb{E}[\mathbb{E}[f(S)|X]] = 0$ and*

$$\text{mmse} \left(\frac{f(S)}{\|f(S)\|_2} \middle| Y \right) \geq \text{mmse} \left(\frac{\mathbb{E}[f(S)|X]}{\|\mathbb{E}[f(S)|X]\|_2} \middle| Y \right).$$

Proof. See Appendix B.1. □

By Lemma 3, $\text{mmse}(s_i(S)|Y) \geq \text{mmse}(\hat{s}_i(X)|Y)$. Therefore, if the new set of functions satisfies the privacy constraints (i.e., $\text{mmse}(\hat{s}_i(X)|Y) \geq \theta_i$), the original set of functions also satisfies the privacy constraints (i.e., $\text{mmse}(s_i(S)|Y) \geq \theta_i$).

4.2 Optimization

We introduce next a PIC-based convex program to find the privacy-assuring mapping $P_{Y|X}$. First, we construct a matrix \mathbf{F} given by $(\mathbf{f}_0, \mathbf{f}_1, \dots, \mathbf{f}_{|\mathcal{X}|-1})$ such that

$$\mathbf{F}^T \mathbf{D}_X \mathbf{F} = \mathbf{I}, \tag{16}$$

$$\text{span}(\{\mathbf{f}_0, \dots, \mathbf{f}_{n'}\}) = \text{span}(\{\mathbf{f}_0, \mathbf{u}_1, \dots, \mathbf{u}_n\}), \tag{17}$$

where $\mathbf{f}_0 \triangleq (1, \dots, 1)^T$, $\mathbf{f}_i \triangleq (f_i(1), \dots, f_i(|\mathcal{X}|))^T$, $\mathbf{u}_i \triangleq (u_i(1), \dots, u_i(|\mathcal{X}|))^T$, and

$$n' \triangleq \dim(\text{span}(\{\mathbf{f}_0, \mathbf{u}_1, \dots, \mathbf{u}_n\})) - 1.$$

Following from (16), $\{f_k(x) \mid k = 0, \dots, |\mathcal{X}|-1\}$ is a basis of $\mathcal{L}_2(P_X)$ and, consequently, the functions $\hat{s}_i(x)$ can be decomposed as

$$\hat{s}_i(x) = \sum_{k=0}^{|\mathcal{X}|-1} \alpha_{i,k} f_k(x). \tag{18}$$

¹We call this step as projection because of the geometric interpretation of conditional expectation (see, e.g., [62]).

$$\max \text{obj}(\sigma_1, \dots, \sigma_{n'}) \quad (20)$$

$$\text{s.t.} \quad \sum_{k=1}^{|\mathcal{X}|-1} \alpha_{i,k}^2 \sigma_k^2 \leq 1 - \theta_i \quad (i = 1, \dots, m), \quad (21)$$

$$0 \leq \sigma_i \leq 1 \quad (i = 1, \dots, |\mathcal{X}|-1), \quad (22)$$

$$\mathbf{\Sigma} = \text{diag}(1, \sigma_1, \dots, \sigma_{|\mathcal{X}|-1}), \quad (23)$$

$$\mathbf{P}_{X,Y} = \mathbf{D}_X \mathbf{F} \mathbf{\Sigma} \mathbf{F}^T \mathbf{D}_X, \quad (24)$$

$$\mathbf{P}_{X,Y} \text{ has non-negative entries.} \quad (25)$$

Formulation 1: PIC-based convex program. Here σ_i ($i = 1, \dots, |\mathcal{X}|-1$) and θ_i ($i = 1, \dots, m$) are variables and privacy parameters, respectively. The objective function $\text{obj}(\sigma_1, \dots, \sigma_{n'})$ is chosen as a concave function and measures utility.

Since $\mathbb{E}[\hat{s}_i(X)] = 0$, then $\alpha_{i,0} = 0$. Similarly, since $\mathbf{u}_i \in \text{span}(\{\mathbf{f}_0, \dots, \mathbf{f}_{n'}\})$ and $\mathbb{E}[u_i(X)] = 0$, we have

$$u_i(x) = \sum_{k=1}^{n'} \beta_{i,k} f_k(x). \quad (19)$$

If $\mathbf{P}_{X,Y} = \mathbf{D}_X \mathbf{F} \mathbf{\Sigma} \mathbf{F}^T \mathbf{D}_X$ with $\mathbf{\Sigma} = \text{diag}(1, \sigma_1, \dots, \sigma_{|\mathcal{X}|-1})$ is a feasible joint distribution matrix (i.e., non-negative entries and all entries add to 1), then, following from Theorem 1,

$$\begin{aligned} \text{mmse}(\hat{s}_i(X)|Y) &= 1 - \sum_{k=1}^{|\mathcal{X}|-1} \alpha_{i,k}^2 \sigma_k^2, \\ \text{mmse}(u_i(X)|Y) &= \sum_{k=1}^{n'} \beta_{i,k}^2 (1 - \lambda_k(X; Y)) \leq 1 - \min_{k \in [n']} \lambda_k(X; Y) = 1 - \left(\min_{k \in [n']} \sigma_k \right)^2. \end{aligned}$$

Therefore, the design of the privacy-assuring mapping $P_{Y|X}$ with privacy-utility constraints is equivalent to solving the PIC-based convex program in Formulation 1. In this case, the objective function is chosen as $\text{obj}(\sigma_1, \dots, \sigma_{n'}) = \min\{\sigma_1, \dots, \sigma_{n'}\}^2$.

The objective function $\min\{\sigma_1, \dots, \sigma_{n'}\}$ maximizes the worst-case utility over all useful functions. On the other hand, we can choose the objective function to be a weighted sum $\sum_{i=1}^{n'} a_i \sigma_i$. Although maximizing the weighted sum is not equivalent to the desired utility constraints, this new formulation allows more flexibility in the optimization. In particular, this enables useful functions which do not highly correlate with private functions to achieve better utility, in terms of mean-squared error, under the same privacy constraints. Furthermore, the weights can be used to prioritize the reconstruction of certain useful functions.

The previous convex programs can be numerically solved by standard methods (e.g., CVXPY [63]). Note that when all useful functions and private functions are based on the same random variable, we can use optimization without projection. We defer the numerical results to Section 7, where we derive privacy-assuring mappings for a synthetic dataset and a real-world dataset using tools introduced in this section.

5 Lower Bounds for MMSE with Restricted Knowledge of the Data Distribution

So far we have assumed the information-theoretic setting where the probability distribution $P_{S,X}$ is known to the privacy mechanism designer beforehand. In this section, we forgo this assumption and consider a

²This is a convex program since one can add a constraint $\sigma_i \geq \sigma$ ($i \in [n']$) and maximize σ .

setting where $S = \phi(X)$ and the correlation between $\phi(X)$ and a set of functions (composed with the data) $\{\phi_j(X)\}_{j=1}^m$ is given. We derive lower bounds for the MMSE of estimating $\phi(X)$ given Y in terms of the MMSE of estimating $\phi_j(X)$ given Y . In privacy systems, X may be a user's data and Y a distorted version of X generated by a privacy-assuring mapping $P_{Y|X}$. The set $\{\phi_j(X)\}_{j=1}^m$ could then represent a set of functions that are known to be hard to infer from Y due to inherent privacy constraints of the setup. For example, when the mapping $P_{Y|X}$ is designed by the PIC-based convex programs in Formulations 1 and $\{\phi_j(X)\}_{j=1}^m$ is the set of private functions, $\text{mmse}(\phi_j(X)|Y)$ is lower bounded due to the privacy constraints.

The following lemma will be used to derive the lower bounds for the MMSE of $\phi(X)$ given Y .

Lemma 4. Let $L_n : (0, \infty)^n \times [0, 1]^n \rightarrow \mathbb{R}$ be given by

$$L_n(\mathbf{a}, \mathbf{b}) \triangleq \max \{ \mathbf{a}^T \mathbf{y} \mid \mathbf{y} \in \mathbb{R}^n, \|\mathbf{y}\|_2 \leq 1, \mathbf{y} \leq \mathbf{b} \}. \quad (26)$$

Let π be a permutation of $[n]$ such that $b_{\pi(1)}/a_{\pi(1)} \leq \dots \leq b_{\pi(n)}/a_{\pi(n)}$. If $b_{\pi(1)}/a_{\pi(1)} \geq 1$, $L_n(\mathbf{a}, \mathbf{b}) = \|\mathbf{a}\|_2$. Otherwise,

$$L_n(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^{k^*} a_{\pi(i)} b_{\pi(i)} + \sqrt{\left(\|\mathbf{a}\|_2^2 - \sum_{i=1}^{k^*} a_{\pi(i)}^2 \right) \left(1 - \sum_{i=1}^{k^*} b_{\pi(i)}^2 \right)}$$

where

$$k^* \triangleq \max \left\{ k \in [n] \mid \frac{b_{\pi(k)}}{a_{\pi(k)}} \leq \sqrt{\frac{\left(1 - \sum_{i=1}^{k-1} b_{\pi(i)}^2 \right)^+}{\|\mathbf{a}\|_2^2 - \sum_{i=1}^{k-1} a_{\pi(i)}^2}} \right\}. \quad (27)$$

Proof. See Appendix C.1. □

Throughout this section we assume $\|\phi_i(X)\|_2 = 1$ ($i \in [m]$) and $\mathbb{E}[\phi_i(X)\phi_j(X)] = 0$ ($i \neq j$). For a given ϕ_i , the inequality

$$\max_{\psi \in \mathcal{L}_2(P_Y)} \mathbb{E}[\phi_i(X)\psi(Y)] = \|\mathbb{E}[\phi_i(X)|Y]\|_2 \leq \nu_i \quad (28)$$

is satisfied, where $0 \leq \nu_i \leq 1$. This is equivalent to $\text{mmse}(\phi_i(X)|Y) \geq 1 - \nu_i^2$.

Theorem 4. Let $\|\phi(X)\|_2 = 1$ and $\mathbb{E}[\phi(X)\phi_i(X)] = \rho_i > 0$. Denoting $\boldsymbol{\rho} \triangleq (|\rho_1|, \dots, |\rho_m|)$, $\boldsymbol{\nu} \triangleq (\nu_1, \dots, \nu_m)$, $\rho_0 \triangleq \sqrt{1 - \sum_{i=1}^m \rho_i^2}$, $\boldsymbol{\rho}_0 \triangleq (\rho_0, \boldsymbol{\rho})$ and $\boldsymbol{\nu}_0 \triangleq (1, \boldsymbol{\nu})$, then

$$\|\mathbb{E}[\phi(X)|Y]\|_2 \leq B_m(\boldsymbol{\rho}_0, \boldsymbol{\nu}_0), \quad (29)$$

where

$$B_m(\boldsymbol{\rho}_0, \boldsymbol{\nu}_0) \triangleq \begin{cases} L_{m+1}(\boldsymbol{\rho}_0, \boldsymbol{\nu}_0) & \text{if } \rho_0 > 0, \\ L_m(\boldsymbol{\rho}, \boldsymbol{\nu}) & \text{otherwise,} \end{cases} \quad (30)$$

and L_n is given in (26). Consequently,

$$\text{mmse}(\phi(X)|Y) \geq 1 - B_m(\boldsymbol{\rho}_0, \boldsymbol{\nu}_0)^2. \quad (31)$$

Proof. See Appendix C.2. □

Denote $\psi_i(Y) \triangleq (T_{X|Y}\phi_i)(Y)/\|(T_{X|Y}\phi_i)(Y)\|_2$ ($i \in [m]$) and $\phi_0(X) \triangleq \rho_0^{-1}(\phi(X) - \sum_{i=1}^m \rho_i \phi_i(X))$ if $\rho_0 > 0$, otherwise $\phi_0(X) \triangleq 0$. The previous bounds, (29) and (31), can be further improved when $\mathbb{E}[\psi_i(Y)\phi_j(X)] = 0$ for $i \neq j$, $j \in \{0, \dots, m\}$.

Theorem 5. Let $\|\phi(X)\|_2 = 1$ and $|\mathbb{E}[\phi(X)\phi_i(X)]| = \rho_i > 0$ for $i \in [m]$. In addition, assume $\mathbb{E}[\psi_i(Y)\phi_j(X)] = 0$ for $i \neq j$, $i \in [t]$ and $j \in \{0, \dots, m\}$, where $0 \leq t \leq m$. Then

$$\|\mathbb{E}[\phi(X)|Y]\|_2 \leq \sqrt{\sum_{k=1}^t \nu_k^2 \rho_k^2 + B_{m-t}(\tilde{\boldsymbol{\rho}}, \tilde{\boldsymbol{\nu}})^2}, \quad (32)$$

where $\tilde{\boldsymbol{\rho}} = (\rho_0, \rho_{t+1}, \dots, \rho_m)$, $\tilde{\boldsymbol{\nu}} = (1, \nu_{t+1}, \dots, \nu_m)$, and B_m is defined in (30) (considering $B_0 = 0$). In particular, if $t = m$,

$$\|\mathbb{E}[\phi(X)|Y]\|_2 \leq \sqrt{\rho_0^2 + \sum_{k=1}^m \nu_k^2 \rho_k^2}, \quad (33)$$

and (33) is an equality when $\rho_0 = 0$. Furthermore,

$$\text{mmse}(\phi(X)|Y) \geq 1 - \sum_{k=1}^t \nu_k^2 \rho_k^2 - B_{m-t}(\tilde{\boldsymbol{\rho}}, \tilde{\boldsymbol{\nu}})^2. \quad (34)$$

Proof. See Appendix C.3. □

In what follows, we use three examples to illustrate different use cases of Theorem 4 and 5. Example 2 illustrates how Theorem 5 can be applied to the q -ary symmetric channel which could be perceived as a model of randomized response [64,65], and demonstrates that bound (33) is sharp. Example 3 illustrates Theorem 5 for the binary symmetric channel. Here the useful variable is composed by n uniform and independent bits. In this case, the basis can be expressed as the parity bits of the input to the channel. Finally, Example 4 illustrates Theorem 4 for one-bit functions. The same method used in the proof of Theorem 4 is applied to bound the probability of correctly guessing a one-bit function from an observation of the disclosed data.

Example 2 (q -ary symmetric channel). Let $\mathcal{X} = \mathcal{Y} = [q]$, and Y be the result of passing X through an (ϵ, q) -ary symmetric channel, which is defined by the transition probability

$$P_{Y|X}(y|x) = (1 - \epsilon)\mathbb{1}_{y=x} + \epsilon/q \text{ for all } x \in \mathcal{X}, y \in \mathcal{Y}. \quad (35)$$

We assume that X has a uniform distribution, which implies Y also has a uniform distribution. Any function $\phi \in \mathcal{L}_2(P_X)$ such that $\mathbb{E}[\phi(X)] = 0$ and $\|\phi(X)\|_2 = 1$ satisfies

$$\psi(Y) = (T_{X|Y}\phi)(Y) = (1 - \epsilon)\phi(Y),$$

and, consequently, $\|(T_{X|Y}\phi)(Y)\|_2 = (1 - \epsilon)$. We will use this fact to show that the bound (33) is sharp in this case.

Observe that for $\phi_i, \phi_j \in \mathcal{L}_2(P_X)$, if $\mathbb{E}[\phi_i(X)\phi_j(X)] = 0$ then $\mathbb{E}[\psi_i(Y)\psi_j(Y)] = 0$. Now let $\phi \in \mathcal{L}_2(P_X)$ satisfy $\mathbb{E}[\phi(X)] = 0$ and $\|\phi(X)\|_2 = 1$, and let $\mathbb{E}[\phi(X)\phi_i(X)] = \rho_i$ for $i \in [m]$, where $\{\phi_i\}$ satisfies the conditions in Theorem 5 and $\sum_{i=1}^m \rho_i^2 = 1$. In addition, $\|\psi_i(Y)\|_2 = (1 - \epsilon) = \nu_i$. Then, from (33) and noting that $\rho_0 = 0$, $t = m$, we have

$$\|(T_{X|Y}\phi)(Y)\|_2 \leq \sqrt{\sum_{i=1}^m \nu_i^2 \rho_i^2} = (1 - \epsilon) \sqrt{\sum_{i=1}^m \rho_i^2} = 1 - \epsilon,$$

which matches $\|(T_{X|Y}\phi)(Y)\|_2$, and the bound is tight in this case.

Example 3 (Binary channels with additive noise). Let $\mathcal{X} = \{-1, 1\}^n$ and $\mathcal{Y} = \{-1, 1\}^n$, and Y^n be the result of passing X^n through a memoryless binary symmetric channel with crossover probability $\epsilon < 1/2$. We assume that X^n is composed by n uniform and i.i.d. bits. For $\mathcal{S} \subseteq [n]$, let

$$\chi_{\mathcal{S}}(X^n) \triangleq \prod_{i \in \mathcal{S}} X_i.$$

Any function $\phi : \mathcal{X} \rightarrow \mathbb{R}$ can then be decomposed in terms of the basis $\chi_{\mathcal{S}}(X^n)$ as [66]

$$\phi(X^n) = \sum_{\mathcal{S} \subseteq [n]} c_{\mathcal{S}} \chi_{\mathcal{S}}(X^n),$$

where $c_{\mathcal{S}} = \mathbb{E}[\phi(X^n)\chi_{\mathcal{S}}(X^n)]$. Furthermore, since $\mathbb{E}[\chi_{\mathcal{S}}(X^n)|Y^n] = (1 - 2\epsilon)^{|\mathcal{S}|} \chi_{\mathcal{S}}(Y^n)$, it follows from Theorem 5 that

$$\text{mmse}(\phi(X^n)|Y^n) = 1 - \sum_{\mathcal{S} \subseteq [n]} c_{\mathcal{S}}^2 (1 - 2\epsilon)^{2|\mathcal{S}|}. \quad (36)$$

This result can be generalized for the case $X^n = Y^n \otimes Z^n$, where the operation \otimes denotes bit-wise multiplication, Z^n is drawn from $\{-1, 1\}^n$ and X^n is uniformly distributed. In this case

$$\text{mmse}(\phi(X^n)|Y^n) = 1 - \sum_{S \subseteq [n]} c_S^2 \mathbb{E}[\chi_S(Z^n)]^2. \quad (37)$$

Example 4 (One-Bit Functions). Let X be a hidden random variable with support \mathcal{X} , and let Y be a noisy observation of X . We denote by B_1, \dots, B_m a collection of m predicates of X , where $B_i = \phi_i(X)$, $\phi_i : \mathcal{X} \rightarrow \{-1, 1\}$ for $i \in [m]$ and, without loss of generality, $\mathbb{E}[B_i] = b_i \geq 0$.

We denote by \hat{B}_i an estimate of B_i given an observation of Y , where $B_i \rightarrow X \rightarrow Y \rightarrow \hat{B}_i$. We assume that for any \hat{B}_i

$$\left| \mathbb{E}[B_i \hat{B}_i] \right| \leq 1 - 2\alpha_i$$

for some $0 \leq \alpha_i \leq (1 - b_i)/2 \leq 1/2$. This condition is equivalent to imposing that $\Pr(B_i \neq \hat{B}_i) \geq \alpha_i$, since

$$\begin{aligned} \mathbb{E}[B_i \hat{B}_i] &= \Pr(B_i = \hat{B}_i) - \Pr(B_i \neq \hat{B}_i) \\ &= 1 - 2\Pr(B_i \neq \hat{B}_i). \end{aligned}$$

In particular, this captures the ‘‘hardness’’ of guessing B_i based solely on an observation of Y .

Now assume there is a bit B such that $\mathbb{E}[BB_i] = \rho_i$ for $i \in [m]$ and $\mathbb{E}[B_i B_j] = 0$ for $i \neq j$. We can apply the same method used in the proof of Theorem 4 to bound the probability of B being guessed correctly from an observation of Y :

$$\Pr(B \neq \hat{B}) \geq \frac{1}{2} (1 - B_m(\boldsymbol{\rho}, \boldsymbol{\nu})), \quad (38)$$

where $\nu_i = 1 - 2\alpha_i$.

6 Robustness of the PUTs

In this section we investigate the pipeline in Fig. 2 for designing privacy-assuring mappings in practice. During the training time, a reference dataset with n samples is drawn from $P_{S,X}$. The distribution of the source is estimated by computing the empirical distribution (type) $P_{\hat{S},\hat{X}}$ of the reference dataset. $P_{\hat{S},\hat{X}}$ and the privacy-utility constraints are then used as inputs to a convex program solver that returns the corresponding privacy-assuring mapping $W_{Y|\hat{X}}$ (if feasible). We denote by \hat{Y} the random variable produced by randomizing \hat{X} according to $W_{Y|\hat{X}}$, i.e., by applying the privacy-assuring mapping to a source with distribution $P_{\hat{S},\hat{X}}$. During the testing time, new i.i.d. samples from the source $P_{S,X}$ are randomized using the privacy-assuring mapping $W_{Y|\hat{X}}$ computed during the training time, resulting in the disclosed variable Y .

The privacy and utility constraints used for computing the privacy-assuring mapping hold for a data source with distribution $P_{\hat{S},\hat{X}}$, since this is the distribution used as an input to the optimization program. However, during the testing time, $W_{Y|\hat{X}}$ is applied to new samples from the source $P_{S,X}$. *Do the privacy and utility guarantees still hold during the testing time?* Since as n increases $P_{\hat{S},\hat{X}}$ converges to $P_{S,X}$, it is natural to expect that the privacy and utility guarantees during the testing time will not be far from the ones selected during the training time.

In what follows, we analyze the robustness of the PUT optimization using χ^2 -information, and characterize the gap between privacy and utility guarantees of the training and testing time in terms of the number of samples in the reference dataset and the probability of less likely symbols. The following lemma will be used to prove the main result in this section.

Lemma 5. *Suppose that $S_i \rightarrow X_i \rightarrow Y_i$ for $i = 1, 2$ and $P_{Y_1|X_1} = P_{Y_2|X_2}$. Let $m_S \triangleq \min\{P_{S_i}(s) \mid s \in \mathcal{S}, i = 1, 2\}$ and $m_X \triangleq \min\{P_{X_i}(x) \mid x \in \mathcal{X}, i = 1, 2\}$. Then*

$$\begin{aligned} |\chi^2(S_1; Y_1) - \chi^2(S_2; Y_2)| &\leq \frac{4}{m_S} \|P_{S_1, X_1} - P_{S_2, X_2}\|_1, \\ |\chi^2(X_1; Y_1) - \chi^2(X_2; Y_2)| &\leq \frac{4}{m_X} \|P_{S_1, X_1} - P_{S_2, X_2}\|_1. \end{aligned}$$




Figure 2: Flowchart for training time (red) and testing time (blue).

Proof. See Appendix D.1. □

Next, we illustrate the sharpness of the upper bounds in Lemma 5 through the following example.

Example 5. Let $\mathcal{X} = \mathcal{Y} = \{0, 1\}$, $S_i = X_i$ for $i \in \{1, 2\}$. Assume that $P_{X_1}(1) = p$ and $P_{X_2}(1) = p + \epsilon/2$ are such that $p, \epsilon \in (0, 1/2)$. Let $P_{Y|X}$ denote the Z-channel determined by

$$P_{Y|X}(y|x) = \begin{cases} a & \text{if } y = 1, x = 1, \\ 1 - a & \text{if } y = 0, x = 1, \\ 1 & \text{if } y = 0, x = 0, \end{cases} \quad (39)$$

where $a \in (0, 1)$. In this case, we have $\epsilon = \|P_{S_1, X_1} - P_{S_2, X_2}\|_1$ which does not depend on a and p , and

$$\chi^2(S_i; Y_i) = \chi^2(X_i; Y_i) = 1 - \frac{1 - a}{1 - aP_{X_i}(1)}, \text{ for } i \in \{1, 2\}.$$

We denote $\Delta \triangleq |\chi^2(S_1; Y_1) - \chi^2(S_2; Y_2)| = |\chi^2(X_1; Y_1) - \chi^2(X_2; Y_2)|$. By a simple manipulation and Lemma 5, for sufficiently small ϵ ,

$$\frac{(1 - a)a}{2(1 - ap)^2} \epsilon \leq \Delta \leq \frac{4}{p} \epsilon. \quad (40)$$

In particular, if we let $a = 2/3$ and $p = 1/3$, then $9\epsilon/49 \leq \Delta \leq 12\epsilon$. This example shows that the first-order dependence on the \mathcal{L}_1 -norm in the bounds of Lemma 5 cannot be improved in general. In what follows, this \mathcal{L}_1 -norm will be translated into the deviation between the underlying and empirical distributions which vanishes with order $1/\sqrt{n}$ where n is the number of samples.

The next theorem follows from Lemma 5 and large deviation results [67]. It answers the question raised at the beginning of this section and provides the bounds for the difference between the training and testing privacy-utility guarantees. Note that the bounds provided in the theorem hold for any channel $W_{Y|\hat{X}}$, not just the ones that optimize the PUT. In other words, the theorem holds for any privacy-assuring mapping returned by the solver in Fig. 2, even if this mapping is not a globally optimal solution.

Theorem 6. Let $P_{\hat{S}, \hat{X}}$ be the empirical distribution obtained from n i.i.d. samples drawn from the true distribution $P_{S, X}$. In addition, denote by Y and \hat{Y} the random variables obtained by passing X and \hat{X} through a given channel $W_{Y|\hat{X}}$, respectively. Let

$$m_S \triangleq \left(\min\{P_S(s) \mid s \in \mathcal{S}\} - \sqrt{\frac{2}{n} (M - \ln \beta)} \right)_+, \quad (41)$$




Figure 3: We depict the bounds of the χ^2 -privacy-utility function (see Theorem 2) and the privacy-utility values of the privacy-assuring mappings designed by the optimization methods in Section 4.

$$m_X \triangleq \left(\min\{P_X(x) \mid x \in \mathcal{X}\} - \sqrt{\frac{2}{n} (M - \ln \beta)} \right)_+ . \quad (42)$$

Then, with probability at least $1 - \beta$,

$$|\chi^2(S; Y) - \chi^2(\hat{S}; \hat{Y})| \leq \frac{4}{m_S} \sqrt{\frac{2}{n} (M - \ln \beta)}, \quad (43)$$

$$|\chi^2(X; Y) - \chi^2(\hat{X}; \hat{Y})| \leq \frac{4}{m_X} \sqrt{\frac{2}{n} (M - \ln \beta)}, \quad (44)$$

where $M = |\mathcal{S}| |\mathcal{X}|$.

Proof. See Appendix D.2. □

Note that the bounds provided in Theorem 6 depend on the probability of the least likely symbols. The bounds become weaker as m_X and m_S become smaller. We refer the reader to Wang *et al.* [51] for an alternate approach if m_X and m_S are near zero.

7 Numerical Results

We illustrate some of the results derived in this paper through two experiments. The first experiment, conducted on a synthetic dataset, verifies the tightness of the upper bound for the χ^2 -privacy-utility function. The second experiment, run on a real-world dataset, demonstrates the performance of the optimization methods proposed in Section 4.

7.1 Example 1: Parity Bits

We choose private variable $S = (S_1, S_2) \in \{-1, 1\}^2$, where S is composed by two independent bits with $\Pr(S_1 = 1) = 0.45$ and $\Pr(S_2 = 1) = 0.4$. The useful variable $X = (X_1, X_2) \in \{-1, 1\}^2$ is generated by passing S_1 and S_2 through BSC(0.2) and BSC(0.15), respectively.

We use the optimization methods proposed in Section 4 to design privacy-assuring mappings. The private and useful functions are selected as $s_1(S) = S_1$ and $u_1(X) = X_1 X_2$, $u_2(X) = X_2$, respectively. We first project the private function to the useful variable. Then we apply Formulation 1 with $\text{obj}(\sigma_1, \dots, \sigma_{n'}) = \sum_{i=1}^{n'} \sigma_i$ to find the privacy-assuring mappings.

In Fig. 3, we depict the privacy and utility, measured by χ^2 -information, of the privacy-assuring mappings. We also draw the upper bound and lower bound of the χ^2 -privacy-utility function. As shown, the




Figure 4: MMSE of estimating each function given the disclosed variable, where darker means harder to estimate. Here $(Education\ Years, Income)$ and $(Gender, Race)$ are useful variable and private variable, respectively. The privacy parameters θ_i are selected as the same for all i and increase from 0 to 1 (i.e., the privacy constraints are increasing from the top down). The privacy-assuring mappings are designed by Formulation 1 with $\text{obj}(\sigma_1, \dots, \sigma_{n'}) = \min\{\sigma_1, \dots, \sigma_{n'}\}$ (left) and with $\text{obj}(\sigma_1, \dots, \sigma_{n'}) = \sum_{i=1}^{n'} \sigma_i$ (right), respectively.

privacy-utility values of the designed mappings are very close to the upper bound. In particular, since the χ^2 -privacy-utility function is a concave function (see Lemma 2), the curve of this function is between its upper bound (red line) and the linear interpolation of the achievable privacy-utility values (dashed line).

7.2 Example 2: UCI Adult Dataset

We apply our formulations to the UCI Adult Dataset [68]. A natural selection for the private and useful variables are $S = (Gender, Race)$ and $X = (Education\ Years, Income)$, respectively. This allows us to interpret the results of our formulations in an intuitive way, as one would expect there to exist correlations between the chosen private and useful variables. Private functions and useful functions are represented by indicator functions. Furthermore, functions which are linear combinations of others are removed. Following the same procedure proposed in Section 4, we first project all private functions to the useful variable. We use QR decomposition [69] to construct the basis $\{f_k(x)\}$. Note that other decomposition methods can also be used for constructing basis and, in fact, different bases affect the behavior of the PIC-based convex program (e.g., the joint distribution matrix $\mathbf{P}_{X,Y}$ returned by the optimization may be different). Consequently, the solution produced from the optimization program may not be optimal. Finally, Formulation 1 is used to compute the privacy-assuring mappings.

In Fig. 4, we show the MMSE of estimating useful functions and private functions given the disclosed variable. As shown, when we use Formulation 1 with $\text{obj}(\sigma_1, \dots, \sigma_{n'}) = \min\{\sigma_1, \dots, \sigma_{n'}\}$ to compute privacy-assuring mappings, the estimation errors behave uniformly among all functions. This is because we aim at maximizing the worst-case utility over all useful functions. On the other hand, the privacy-assuring mappings designed by Formulation 1 with $\text{obj}(\sigma_1, \dots, \sigma_{n'}) = \sum_{i=1}^{n'} \sigma_i$ reveal more interpretable relationships between the private functions and useful functions. We see that *Income*, *Gender*, and *Race* are highly correlated, and it is not possible to reveal *Income* while maintaining privacy for *Gender* and *Race*. Of particular interest are

the subtle correlations between the three aforementioned functions and *Education Years*. There is a marked correlation between *Education Years* < 6 and, to a lesser degree, *Education Years* > 12 , with *Gender* and *Race*. This may be due to the fact that most members of the dataset do not end their education midway. That is, most individuals will either never have begun schooling in the first place or will not continue their education after the 12-year benchmark, which marks graduation from high school. Therefore, we observe that the relationship between *Education Years* and *Race* is manifested the most in the two extremities of *Education Years* (> 12 and < 6). Also of note is the correlation between the private functions and *Education Years* : 8. Though not as obvious, this relationship can, too, be explained by the fact that 8 years of education marks another benchmark: the beginning of high school, also a time when people are prone to terminating their education.

8 Conclusion

In this paper, we studied a fundamental PUT in data disclosure, where an analyst is allowed to reconstruct certain functions of the data, while other private functions should not be estimated with distortion below a certain threshold. First, χ^2 -information was used to measure both privacy and utility. Bounds on the best PUT were provided and the upper bound, in particular, was shown to be achievable in the high-privacy region. Moreover, a PIC-based convex program was proposed to design privacy-assuring mappings when the useful functions and private functions were known beforehand. We also derived lower bounds on the MMSE of estimating a target function from the disclosed data and analyzed the robustness of our method when the designer used empirical distribution to compute the privacy-assuring mappings. Finally, we performed two experiments and analyzed the numerical results. Our hope is that the methods presented here can inspire new, information-theoretically grounded and interpretable privacy mechanisms.

Acknowledgment

The authors would like to thank the anonymous reviewers and the Associate Editor for their careful reading of our manuscript and their many insightful comments and suggestions.

Appendix A Proofs from Section 3

A.1 Lemma 1

Proof. By the definition of χ^2 -information,

$$\chi^2(X; Y) + 1 = \sum_{x=1}^{|\mathcal{X}|} \sum_{y=1}^{|\mathcal{Y}|} \frac{P_{X,Y}(x, y)}{P_X(x)P_Y(y)} P_{X,Y}(x, y).$$

Note that $\mathbf{Q}_{X,Y} = \mathbf{D}_X^{-\frac{1}{2}} \mathbf{P}_{X,Y} \mathbf{D}_Y^{-\frac{1}{2}}$ which implies

$$\begin{aligned} \text{tr}(\mathbf{Q}_{X,Y} \mathbf{Q}_{X,Y}^T) &= \text{tr}(\mathbf{D}_X^{-\frac{1}{2}} \mathbf{P}_{X,Y} \mathbf{D}_Y^{-1} \mathbf{P}_{X,Y}^T \mathbf{D}_X^{-\frac{1}{2}}) = \text{tr}(\mathbf{D}_X^{-1} \mathbf{P}_{X,Y} \mathbf{D}_Y^{-1} \mathbf{P}_{X,Y}^T) \\ &= \sum_{x=1}^{|\mathcal{X}|} \sum_{y=1}^{|\mathcal{Y}|} \frac{P_{X,Y}(x, y)}{P_X(x)} \frac{P_{X,Y}(x, y)}{P_Y(y)}. \end{aligned}$$

Therefore,

$$\chi^2(X; Y) = \text{tr}(\mathbf{Q}_{X,Y} \mathbf{Q}_{X,Y}^T) - 1 = \text{tr}(\mathbf{A}) - 1.$$

Since

$$\mathbf{Q}_{S,Y} = \mathbf{D}_S^{-\frac{1}{2}} \mathbf{P}_{S,Y} \mathbf{D}_Y^{-\frac{1}{2}} = \mathbf{D}_S^{-\frac{1}{2}} \mathbf{P}_{S,X} \mathbf{D}_X^{-\frac{1}{2}} \mathbf{D}_X^{-\frac{1}{2}} \mathbf{P}_{X,Y} \mathbf{D}_Y^{-\frac{1}{2}} = \mathbf{Q}_{S,X} \mathbf{Q}_{X,Y},$$

then

$$\chi^2(S; Y) = \text{tr}(\mathbf{Q}_{S,Y} \mathbf{Q}_{S,Y}^T) - 1 = \text{tr}(\mathbf{Q}_{S,X} \mathbf{Q}_{X,Y} \mathbf{Q}_{X,Y}^T \mathbf{Q}_{S,X}^T) - 1 = \text{tr}(\mathbf{B}\mathbf{A}) - 1.$$

□

A.2 Lemma 2

Proof. For $0 \leq \epsilon_1 < \epsilon_2 < \epsilon_3 \leq \chi^2(S; X)$, it suffices to show that

$$\frac{F_{\chi^2}(\epsilon_3; P_{S,X}) - F_{\chi^2}(\epsilon_1; P_{S,X})}{\epsilon_3 - \epsilon_1} \leq \frac{F_{\chi^2}(\epsilon_2; P_{S,X}) - F_{\chi^2}(\epsilon_1; P_{S,X})}{\epsilon_2 - \epsilon_1},$$

which is equivalent to

$$\frac{\epsilon_2 - \epsilon_1}{\epsilon_3 - \epsilon_1} F_{\chi^2}(\epsilon_3; P_{S,X}) + \frac{\epsilon_3 - \epsilon_2}{\epsilon_3 - \epsilon_1} F_{\chi^2}(\epsilon_1; P_{S,X}) \leq F_{\chi^2}(\epsilon_2; P_{S,X}). \quad (45)$$

Let $P_{Y_1|X}$ and $P_{Y_3|X}$ be two optimal solutions in $\mathcal{D}(\epsilon_1; P_{S,X})$ and $\mathcal{D}(\epsilon_3; P_{S,X})$, respectively. Assume that Y_1 and Y_3 take values in $[m_1]$ and $[m_3]$, respectively. Furthermore, we denote $\lambda \triangleq \frac{\epsilon_2 - \epsilon_1}{\epsilon_3 - \epsilon_1}$. Next, we introduce a new privacy-assuring mapping defined as

$$P_{Y_\lambda|X}(y|x) \triangleq \begin{cases} \lambda P_{Y_3|X}(y|x) & \text{if } y \in [m_3], \\ (1 - \lambda) P_{Y_1|X}(y - m_3|x) & \text{if } y - m_3 \in [m_1]. \end{cases} \quad (46)$$

Consequently, we have

$$P_{Y_\lambda}(y) = \begin{cases} \lambda P_{Y_3}(y) & \text{if } y \in [m_3], \\ (1 - \lambda) P_{Y_1}(y - m_3) & \text{if } y - m_3 \in [m_1]. \end{cases}$$

Then

$$\begin{aligned} \chi^2(X; Y_\lambda) &= \mathbb{E} \left[\frac{P_{X,Y_\lambda}(X, Y_\lambda)}{P_X(X)P_{Y_\lambda}(Y_\lambda)} \right] - 1 \\ &= \sum_{y \in [m_3]} \sum_{x=1}^{|\mathcal{X}|} \frac{P_{X,Y_\lambda}(x, y)^2}{P_X(x)P_{Y_\lambda}(y)} + \sum_{y - m_3 \in [m_1]} \sum_{x=1}^{|\mathcal{X}|} \frac{P_{X,Y_\lambda}(x, y)^2}{P_X(x)P_{Y_\lambda}(y)} - 1 \\ &= \sum_{y \in [m_3]} \sum_{x=1}^{|\mathcal{X}|} \frac{P_{Y_\lambda|X}(y|x)^2 P_X(x)}{P_{Y_\lambda}(y)} + \sum_{y - m_3 \in [m_1]} \sum_{x=1}^{|\mathcal{X}|} \frac{P_{Y_\lambda|X}(y|x)^2 P_X(x)}{P_{Y_\lambda}(y)} - 1 \\ &= \sum_{y \in [m_3]} \sum_{x=1}^{|\mathcal{X}|} \frac{\lambda^2 P_{Y_3|X}(y|x)^2 P_X(x)}{\lambda P_{Y_3}(y)} + \sum_{y \in [m_1]} \sum_{x=1}^{|\mathcal{X}|} \frac{(1 - \lambda)^2 P_{Y_1|X}(y|x)^2 P_X(x)}{(1 - \lambda) P_{Y_1}(y)} - 1 \\ &= \lambda \chi^2(X; Y_3) + (1 - \lambda) \chi^2(X; Y_1). \end{aligned}$$

Similarly, we have

$$\chi^2(S; Y_\lambda) = \lambda \chi^2(S; Y_3) + (1 - \lambda) \chi^2(S; Y_1) \leq \epsilon_2, \quad (47)$$

which implies that $P_{Y_\lambda|X} \in \mathcal{D}(\epsilon_2; P_{S,X})$. Therefore,

$$\begin{aligned} F_{\chi^2}(\epsilon_2; P_{S,X}) &\geq \chi^2(X; Y_\lambda) \\ &= \lambda \chi^2(X; Y_3) + (1 - \lambda) \chi^2(X; Y_1) \\ &= \frac{\epsilon_2 - \epsilon_1}{\epsilon_3 - \epsilon_1} F_{\chi^2}(\epsilon_3; P_{S,X}) + \frac{\epsilon_3 - \epsilon_2}{\epsilon_3 - \epsilon_1} F_{\chi^2}(\epsilon_1; P_{S,X}), \end{aligned} \quad (48)$$

which implies that (45) is true, so $F_{\chi^2}(\epsilon; P_{S,X})$ is a concave function. Furthermore, $\epsilon \rightarrow \frac{1}{\epsilon} F_{\chi^2}(\epsilon; P_{S,X})$ is non-increasing since $F_{\chi^2}(\epsilon; P_{S,X})$ is non-negative and concave. \square

A.3 Closed-Form Expression of $G_\epsilon^m(t_1, \dots, t_n)$

Recall that, for $t_i \in [0, 1]$ ($i \in [n]$), $0 \leq \epsilon \leq \sum_{i \in [n]} t_i$, and $n \leq m$, $G_\epsilon^m(t_1, \dots, t_n)$ is defined as:

$$G_\epsilon^m(t_1, \dots, t_n) \triangleq \max \left\{ \sum_{i=1}^m x_i \mid (x_1, \dots, x_m) \in \mathcal{D}_\epsilon^m(t_1, \dots, t_n) \right\},$$

where

$$\mathcal{D}_\epsilon^m(t_1, \dots, t_n) \triangleq \left\{ (x_1, \dots, x_m) \mid \sum_{i=1}^n t_i x_i \leq \epsilon, x_i \in [0, 1], i \in [m] \right\}.$$

We assume $1 \geq t_1 \geq \dots \geq t_{n-s} > t_{n-s+1} = \dots = t_n = 0$ without loss of generality. Then we divide $[0, \sum_{i=1}^n t_i]$ into $n-s$ intervals:

$$\left[0, \sum_{i=1}^n t_i \right] = \bigcup_{j=0}^{n-1-s} \left[\sum_{i=n-s-j+1}^{n-s} t_i, \sum_{i=n-s-j}^{n-s} t_i \right].$$

If $\epsilon \in \left[\sum_{i=n-s-j+1}^{n-s} t_i, \sum_{i=n-s-j}^{n-s} t_i \right]$, then

$$G_\epsilon^m(t_1, \dots, t_n) = s + (m-n) + j + \frac{\epsilon - \sum_{i=n-s-j+1}^{n-s} t_i}{t_{n-s-j}}, \quad (49)$$

and it can be achieved by setting

$$\begin{aligned} x_i &= 1, \text{ for } i = n-s-j+1, \dots, m, \\ x_{n-s-j} &= \frac{\epsilon - \sum_{i=n-s-j+1}^{n-s} t_i}{t_{n-s-j}}, \\ x_i &= 0, \text{ for } i = 1, \dots, n-s-j-1. \end{aligned}$$

A.4 Theorem 2

Proof. The lower bound for $F_{\chi^2}(\epsilon; P_{S,X})$ follows immediately from the concavity of $F_{\chi^2}(\epsilon; P_{S,X})$ and

$$\begin{aligned} F_{\chi^2}(0; P_{S,X}) &\geq 0, \\ F_{\chi^2}(\chi^2(S; X); P_{S,X}) &= \chi^2(X; X) = |\mathcal{X}| - 1. \end{aligned}$$

Using Lemma 1, the χ^2 -privacy-utility function can be simplified as

$$\begin{aligned} F_{\chi^2}(\epsilon; P_{S,X}) &= \max_{P_{Y|X} \in \mathcal{D}(\epsilon; P_{S,X})} \text{tr}(\mathbf{A}) - 1, \\ \mathcal{D}(\epsilon; P_{S,X}) &= \{P_{Y|X} \mid S \rightarrow X \rightarrow Y, \text{tr}(\mathbf{B}\mathbf{A}) - 1 \leq \epsilon\}, \end{aligned}$$

where

$$\mathbf{A} = \mathbf{Q}_{X,Y} \mathbf{Q}_{X,Y}^T, \mathbf{B} = \mathbf{Q}_{S,X}^T \mathbf{Q}_{S,X}.$$

We denote the singular value decomposition of $\mathbf{Q}_{S,X}$ and $\mathbf{Q}_{X,Y}$ by $\mathbf{Q}_{S,X} = \mathbf{W}\mathbf{\Sigma}_1\mathbf{U}^T$ and $\mathbf{Q}_{X,Y} = \mathbf{V}\mathbf{\Sigma}_2\mathbf{M}^T$, respectively. Then $\mathbf{B} = \mathbf{U}\mathbf{\Sigma}_1^T\mathbf{\Sigma}_1\mathbf{U}^T = \mathbf{U}\mathbf{\Sigma}_B\mathbf{U}^T$, $\mathbf{A} = \mathbf{V}\mathbf{\Sigma}_2\mathbf{\Sigma}_2^T\mathbf{V}^T = \mathbf{V}\mathbf{\Sigma}_A\mathbf{V}^T$ where $\mathbf{\Sigma}_B \triangleq \mathbf{\Sigma}_1^T\mathbf{\Sigma}_1$ and $\mathbf{\Sigma}_A \triangleq \mathbf{\Sigma}_2\mathbf{\Sigma}_2^T$.

Let $\mathbf{A}_1 = \mathbf{U}^T\mathbf{A}\mathbf{U} = \mathbf{L}\mathbf{\Sigma}_A\mathbf{L}^T$ where $\mathbf{L} \triangleq \mathbf{U}^T\mathbf{V}$. Suppose the diagonal elements of \mathbf{A}_1 are $a_1, \dots, a_{|\mathcal{X}|}$. Then, from characterization 3 in Theorem 1, we have

$$\text{tr}(\mathbf{B}\mathbf{A}) - 1 = a_1 - 1 + \sum_{i=2}^{d+1} \lambda_{i-1}(S; X) a_i. \quad (50)$$

Suppose the i -th row of \mathbf{L} is $\mathbf{l}_i = (l_{i,1}, \dots, l_{i,|\mathcal{X}|})$, the i -th column of \mathbf{U} is \mathbf{u}_i^T and $\mathbf{\Sigma}_A = \text{diag}(\sigma_1, \dots, \sigma_{|\mathcal{X}|})$. By characterization 3 in Theorem 1, $\sigma_1 = 1$, $\sigma_{j+1} = \lambda_j(X; Y)$ for $j = 1, \dots, d$ and $\sigma_{j+1} = 0$ for $j = d+1, \dots, |\mathcal{X}|-1$. Then, for $\forall i \in [|\mathcal{X}|]$,

$$0 \leq a_i = \sum_{j=1}^{|\mathcal{X}|} \sigma_j l_{i,j}^2 \leq \sum_{j=1}^{|\mathcal{X}|} l_{i,j}^2 = 1.$$

Since, following from characterization 3 in Theorem 1, the first column of \mathbf{U} and that of \mathbf{V} are both $(\sqrt{P_X(1)}, \dots, \sqrt{P_X(|\mathcal{X}|)})^T$, then $\mathbf{l}_1 = \mathbf{u}_1 \mathbf{V} = (1, 0, \dots, 0)$. Therefore, $a_1 = \sigma_1 = 1$. If $P_{Y|X} \in \mathcal{D}(\epsilon; P_{S,X})$, then (50) shows

$$\text{tr}(\mathbf{BA}) - 1 = \sum_{i=2}^{d+1} \lambda_{i-1}(S; X) a_i \leq \epsilon,$$

which implies that

$$(a_2, \dots, a_{|\mathcal{X}|}) \in \mathcal{D}_\epsilon^{|\mathcal{X}|-1}(\lambda_1(S; X), \dots, \lambda_d(S; X)).$$

Thus,

$$\begin{aligned} F_{\chi^2}(\epsilon; P_{S,X}) &\leq \max_{(a_2, \dots, a_{|\mathcal{X}|}) \in \mathcal{D}_\epsilon^{|\mathcal{X}|-1}(\lambda_1(S; X), \dots, \lambda_d(S; X))} \sum_{i=2}^{|\mathcal{X}|} a_i \\ &= G_\epsilon^{|\mathcal{X}|-1}(\lambda_1(S; X), \dots, \lambda_d(S; X)), \end{aligned}$$

as required. \square

A.5 Corollary 1

Proof. First, $F_{\chi^2}(\epsilon; P_{S,X})$ is non-decreasing since, for any $0 \leq \epsilon_1 < \epsilon_2 \leq \chi^2(S; X)$, we have $\mathcal{D}(\epsilon_1; P_{S,X}) \subseteq \mathcal{D}(\epsilon_2; P_{S,X})$. Now suppose there exist ϵ_1 and ϵ_2 , such that $F_{\chi^2}(\epsilon_1; P_{S,X}) = F_{\chi^2}(\epsilon_2; P_{S,X})$. We denote $\chi^2(S; X)$ by ϵ_0 . Since $F_{\chi^2}(\epsilon; P_{S,X})$ is a concave and non-decreasing function, then for any $\epsilon > \epsilon_1$, $F_{\chi^2}(\epsilon; P_{S,X}) = F_{\chi^2}(\epsilon_1; P_{S,X})$. In particular, $F_{\chi^2}(\epsilon_1; P_{S,X}) = F_{\chi^2}(\epsilon_0; P_{S,X}) = |\mathcal{X}|-1$. This contradicts the upper bound of χ^2 -privacy-utility function in Theorem 2 since the upper bound implies that $F_{\chi^2}(\epsilon; P_{S,X}) < |\mathcal{X}|-1$ when $\epsilon < \epsilon_0$. \square

A.6 Theorem 3

Proof. Following from characterization 2 in Theorem 1, there exists $f \in \mathcal{L}_2(P_X)$ such that $\|f(X)\|_2 = 1$, $\mathbb{E}[f(X)] = 0$ and $\|\mathbb{E}[f(X)|S]\|_2^2 = \lambda_{\min}(S; X)$.

Fix $\mathcal{Y} = \{1, 2\}$ and the privacy-assuring mapping is defined as

$$P_{Y|X}(y|x) = \frac{1}{2} + (-1)^y \frac{\sqrt{P_X \min} f(x)}{2}. \quad (51)$$

Since

$$1 = \|f(X)\|_2^2 = \sum_{x=1}^{|\mathcal{X}|} f(x)^2 P_X(x) \geq f(x)^2 P_X(x),$$

for any $x \in [|\mathcal{X}|]$

$$|f(x)| \leq \frac{1}{\sqrt{P_X(x)}} \leq \frac{1}{\sqrt{P_X \min}}.$$

Therefore, $\left| \frac{\sqrt{P_X \min} f(x)}{2} \right| \leq \frac{1}{2}$, which implies that $P_{Y|X}(y|x)$ is feasible. Furthermore, $P_Y(y) = \frac{1}{2}$ because of $\mathbb{E}[f(X)] = 0$.

$$\chi^2(X; Y) = \sum_{x=1}^{|\mathcal{X}|} \sum_{y=1}^{|\mathcal{Y}|} \frac{P_{Y|X}(y|x)^2 P_X(x)}{P_Y(y)} - 1 = \sum_{x=1}^{|\mathcal{X}|} (P_X(x) + P_X \min f(x)^2 P_X(x)) - 1 = P_X \min.$$

Since

$$\begin{aligned} P_{Y|S}(y|s) &= \sum_{x=1}^{|\mathcal{X}|} P_{Y|X}(y|x) P_{X|S}(x|s) = \sum_{x=1}^{|\mathcal{X}|} \left(\frac{1}{2} + (-1)^y \frac{\sqrt{P_X \min} f(x)}{2} \right) P_{X|S}(x|s) \\ &= \frac{1}{2} + (-1)^y \frac{\sqrt{P_X \min}}{2} \mathbb{E}[f(X)|S=s], \end{aligned}$$

then

$$\begin{aligned}\chi^2(S; Y) &= \sum_{s=1}^{|S|} \sum_{y=1}^{|Y|} \frac{P_{Y|S}(y|s)^2 P_S(s)}{P_Y(y)} - 1 = \sum_{s=1}^{|S|} \left(P_S(s) + P_{X \min} \mathbb{E}[f(X)|S=s]^2 P_S(s) \right) - 1 \\ &= P_{X \min} \lambda_{\min}(S; X).\end{aligned}$$

Hence, this Y satisfies $\chi^2(X; Y) = P_{X \min}$ and $\chi^2(S; Y) = P_{X \min} \lambda_{\min}(S; X)$. \square

Appendix B Proofs from Section 4

B.1 Lemma 3

Proof. Suppose $\|f(S)\|_2 = 1$ without loss of generality. Observe that

$$\mathbb{E}[\mathbb{E}[f(S)|X]] = \mathbb{E}[f(S)] = 0.$$

Since $f(S) \rightarrow X \rightarrow Y$, then $\mathbb{E}[f(S)|X] = \mathbb{E}[f(S)|X, Y]$. Therefore,

$$\begin{aligned}\text{mmse} \left(\frac{\mathbb{E}[f(S)|X]}{\|\mathbb{E}[f(S)|X]\|_2} \middle| Y \right) &= \frac{\mathbb{E}[\mathbb{E}[f(S)|X]^2] - \mathbb{E}[\mathbb{E}[\mathbb{E}[f(S)|X]|Y]^2]}{\|\mathbb{E}[f(S)|X]\|_2^2} \\ &= \frac{\mathbb{E}[\mathbb{E}[f(S)|X]^2] - \mathbb{E}[\mathbb{E}[\mathbb{E}[f(S)|X, Y]|Y]^2]}{\|\mathbb{E}[f(S)|X]\|_2^2} \\ &= 1 - \frac{\mathbb{E}[\mathbb{E}[f(S)|Y]^2]}{\|\mathbb{E}[f(S)|X]\|_2^2} \\ &= \mathbb{E}[f(S)^2] - \frac{\mathbb{E}[\mathbb{E}[f(S)|Y]^2]}{\|\mathbb{E}[f(S)|X]\|_2^2} \\ &\leq \mathbb{E}[f(S)^2] - \mathbb{E}[\mathbb{E}[f(S)|Y]^2] \\ &= \text{mmse}(f(S)|Y),\end{aligned}$$

where the last inequality follows from Jensen's inequality:

$$1 = \mathbb{E}[f(S)^2] = \mathbb{E}[\mathbb{E}[f(S)^2|X]] \geq \mathbb{E}[\mathbb{E}[f(S)|X]^2] = \|\mathbb{E}[f(S)|X]\|_2^2.$$

\square

Appendix C Proofs from Section 5

C.1 Lemma 4

Proof. For fixed $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ where $a_i > 0$ and $b_i \geq 0$, let $L_P : \mathbb{R}^n \rightarrow \mathbb{R}$ and $L_D : \mathbb{R}^n \rightarrow \mathbb{R}$ be given by

$$\begin{aligned}L_P(\mathbf{y}) &\triangleq \mathbf{a}^T \mathbf{y}, \\ L_D(\mathbf{u}) &\triangleq \mathbf{a}^T \mathbf{b} + \mathbf{u}^T \mathbf{b} + \|\mathbf{u}\|_2.\end{aligned}$$

Furthermore, we define $\mathcal{A}(\mathbf{a}) \triangleq \{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \geq -\mathbf{a}\}$ and $\mathcal{B}(\mathbf{b}) \triangleq \{\mathbf{y} \in \mathbb{R}^n \mid \|\mathbf{y}\|_2 \leq 1, \mathbf{y} \leq \mathbf{b}\}$.

Assume, without loss of generality, that $b_1/a_1 \leq b_2/a_2 \leq \dots \leq b_n/a_n$, and let k^* be defined in (27). Note that $b_1 \leq 1$ and for $k \in [k^*]$

$$\sum_{i=1}^k b_i^2 \leq \frac{a_k^2}{\sum_{i=k}^n a_i^2} \left(1 - \sum_{i=1}^{k-1} b_i^2 \right)^+ + \sum_{i=1}^{k-1} b_i^2,$$

so $\sum_{i=1}^k b_i^2 \leq 1$. Especially, $\sum_{i=1}^{k^*} b_i^2 \leq 1$. For $c_j \triangleq \sqrt{\frac{(1-\sum_{i=1}^j b_i^2)}{\|\mathbf{a}\|_2^2 - \sum_{i=1}^j a_i^2}}$, let

$$\mathbf{y}^* = (b_1, \dots, b_{k^*}, a_{k^*+1}c_{k^*}, \dots, a_n c_{k^*})$$

and

$$\mathbf{u}^* = (-b_1/c_{k^*}, \dots, -b_{k^*}/c_{k^*}, -a_{k^*+1}, \dots, -a_n).$$

From the definition of k^* , $\mathbf{y}^* \in \mathcal{B}(\mathbf{b})$ and $\mathbf{u}^* \in \mathcal{A}(\mathbf{a})$. Furthermore,

$$\begin{aligned} L_P(\mathbf{y}^*) &= \mathbf{a}^T \mathbf{y}^* \\ &= \sum_{i=1}^{k^*} a_i b_i + \sum_{i=k^*+1}^n c_{k^*} a_i^2 \\ &= \sum_{i=1}^{k^*} a_i b_i + \sqrt{\left(\|\mathbf{a}\|_2^2 - \sum_{i=1}^{k^*} a_i^2 \right) \left(1 - \sum_{i=1}^{k^*} b_i^2 \right)}, \end{aligned} \tag{52}$$

and

$$\begin{aligned} L_D(\mathbf{u}^*) &= \mathbf{a}^T \mathbf{b} + \mathbf{u}^{*T} \mathbf{b} + \|\mathbf{u}^*\|_2 \\ &= \sum_{i=1}^{k^*} \left(a_i b_i - \frac{b_i^2}{c_{k^*}} \right) + c_{k^*}^{-1} \sqrt{\sum_{i=1}^{k^*} b_i^2 + c_{k^*}^2 \left(\|\mathbf{a}\|_2^2 - \sum_{i=1}^{k^*} a_i^2 \right)} \\ &= \sum_{i=1}^{k^*} a_i b_i + c_{k^*}^{-1} \left(1 - \sum_{i=1}^{k^*} b_i^2 \right) \\ &= \sum_{i=1}^{k^*} a_i b_i + \sqrt{\left(\|\mathbf{a}\|_2^2 - \sum_{i=1}^{k^*} a_i^2 \right) \left(1 - \sum_{i=1}^{k^*} b_i^2 \right)} = L_P(\mathbf{y}^*). \end{aligned}$$

Since both the primal and the dual achieve the same value at \mathbf{y}^* and \mathbf{u}^* , respectively, it follows that the value $L_P(\mathbf{y}^*)$ given in (52) is optimal. \square

C.2 Theorem 4

Proof. Let

$$h(x) \triangleq \begin{cases} \rho_0^{-1} (\phi(x) - \sum_{i=1}^m \rho_i \phi_i(x)) & \text{if } \rho_0 > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{53}$$

Note that when $\rho_0 > 0$, we have $\|h(X)\|_2 = 1$. Then for any $\psi \in \mathcal{L}_2(P_Y)$ and $\|\psi(Y)\|_2 = 1$,

$$\begin{aligned} |\mathbb{E}[\phi(X)\psi(Y)]| &= \left| \rho_0 \mathbb{E}[h(X)\psi(Y)] + \sum_{i=1}^m \rho_i \mathbb{E}[\phi_i(X)\psi(Y)] \right| \\ &\leq \rho_0 |\mathbb{E}[h(X)\psi(Y)]| + \sum_{i=1}^m |\rho_i \mathbb{E}[\phi_i(X)\psi(Y)]| \\ &= \rho_0 |\mathbb{E}[h(X)(T_{Y|X}\psi)(X)]| + \sum_{i=1}^m |\rho_i \mathbb{E}[\phi_i(X)(T_{Y|X}\psi)(X)]|, \end{aligned}$$

where $T_{Y|X}$ is defined in Section 1.3. Denoting $|\mathbb{E}[h(X)(T_{Y|X}\psi)(X)]| \triangleq x_0$, $|\mathbb{E}[\phi_i(X)(T_{Y|X}\psi)(X)]| \triangleq x_i$, $\mathbf{x} \triangleq (x_0, x_1, \dots, x_m)$, the last inequality can be rewritten as

$$|\mathbb{E}[\phi(X)\psi(Y)]| \leq \boldsymbol{\rho}_0^T \mathbf{x}. \tag{54}$$

Observe that $\|\mathbf{x}\|_2 \leq 1$ and $x_i \leq \nu_i$ for $i \in [m]$, and the right hand side of (54) can be maximized over all values of \mathbf{x} that satisfy these constraints. We assume, without loss of generality, that $\rho_0 > 0$ (otherwise set $x_0 = 0$). The left-hand side of (54) can be further bounded by

$$|\mathbb{E}[\phi(X)\psi(Y)]| \leq L_{m+1}(\boldsymbol{\rho}_0, \boldsymbol{\nu}_0), \quad (55)$$

where $\boldsymbol{\nu}_0 = (1, \nu_1, \dots, \nu_m)$ and L_{m+1} is defined in (26). The result follows directly from Lemma 4 and noting that $\max_{\psi \in \mathcal{L}_2(P_Y)} \mathbb{E}[\phi(X)\psi(Y)] = \|\mathbb{E}[\phi(X)|Y]\|_2$. \square

C.3 Theorem 5

Proof. For any $\psi \in \mathcal{L}_2(P_Y)$ with $\|\psi(Y)\|_2 = 1$, let $\alpha_i \triangleq \mathbb{E}[\psi(Y)\psi_i(Y)]$, $\alpha_0 \triangleq \sqrt{1 - \sum_{i=1}^t \alpha_i^2}$ and $\psi_0(Y) \triangleq \alpha_0^{-1}(\psi(Y) - \sum_{i=1}^t \alpha_i \psi_i(Y))$ if $\alpha_0 > 0$, otherwise $\psi_0(Y) \triangleq 0$. Observe that $\|\psi_0(Y)\|_2 = 1$ when $\alpha_0 > 0$. Also, $\mathbb{E}[\phi_i(X)\psi_j(Y)] = 0$ for $i \neq j$, $i \in \{0, \dots, m\}$, $j \in [t]$. Consequently,

$$\begin{aligned} \mathbb{E}[\phi(X)\psi(Y)] &= \mathbb{E}\left[\left(\sum_{i=0}^m \rho_i \phi_i(X)\right) \left(\sum_{j=0}^t \alpha_j \psi_j(Y)\right)\right] \\ &= \sum_{i=0}^m \sum_{j=0}^t \rho_i \alpha_j \mathbb{E}[\phi_i(X)\psi_j(Y)] \\ &\leq \left| \alpha_0 \sum_{i=0, i \notin [t]}^m \rho_i \mathbb{E}[\phi_i(X)\psi_0(Y)] \right| + \sum_{i=1}^t |\nu_i \rho_i \alpha_i| \\ &\leq |\alpha_0| B_{m-t}(\tilde{\boldsymbol{\rho}}, \tilde{\boldsymbol{\nu}}) + \sum_{i=1}^t |\nu_i \rho_i \alpha_i| \end{aligned} \quad (56)$$

$$\leq \sqrt{\sum_{k=1}^t \nu_k^2 \rho_k^2 + B_{m-t}(\tilde{\boldsymbol{\rho}}, \tilde{\boldsymbol{\nu}})^2}. \quad (57)$$

Inequality (56) follows from the similar proof of Theorem 4, and (57) follows by observing that $\sum_{i=0}^t \alpha_i^2 = 1$ and applying Cauchy-Schwarz inequality.

Finally, when $\rho_0 = 0$ and $t = m$, (57) can be achieved with equality by taking

$$\psi(Y) = \frac{\sum_{i=1}^m \nu_i \rho_i \psi_i(Y)}{\sqrt{\sum_{i=1}^m \nu_i^2 \rho_i^2}}.$$

\square

Appendix D Proofs from Section 6

D.1 Lemma 5

Proof. By the definition of χ^2 -information, we have

$$|\chi^2(X_1; Y_1) - \chi^2(X_2; Y_2)| \leq \sum_{x=1}^{|\mathcal{X}|} \sum_{y=1}^{|\mathcal{Y}|} \left| \frac{P_{X_1, Y_1}(x, y)^2}{P_{X_1}(x)P_{Y_1}(y)} - \frac{P_{X_2, Y_2}(x, y)^2}{P_{X_2}(x)P_{Y_2}(y)} \right|. \quad (58)$$

By the triangle inequality,

$$\begin{aligned}
\left| \frac{P_{X_1, Y_1}(x, y)^2}{P_{X_1}(x)P_{Y_1}(y)} - \frac{P_{X_2, Y_2}(x, y)^2}{P_{X_2}(x)P_{Y_2}(y)} \right| &\leq \frac{P_{X_1, Y_1}(x, y)}{P_{X_1}(x)P_{Y_1}(y)} |P_{X_1, Y_1}(x, y) - P_{X_2, Y_2}(x, y)| \\
&+ \frac{P_{X_1, Y_1}(x, y)}{P_{X_1}(x)P_{Y_1}(y)} \frac{P_{X_2, Y_2}(x, y)}{P_{X_2}(x)} |P_{X_1}(x) - P_{X_2}(x)| \\
&+ \frac{P_{X_1, Y_1}(x, y)}{P_{Y_1}(y)} \frac{P_{X_2, Y_2}(x, y)}{P_{X_2}(x)P_{Y_2}(y)} |P_{Y_1}(y) - P_{Y_2}(y)| \\
&+ \frac{P_{X_2, Y_2}(x, y)}{P_{X_2}(x)P_{Y_2}(y)} |P_{X_1, Y_1}(x, y) - P_{X_2, Y_2}(x, y)|. \tag{59}
\end{aligned}$$

Note that for $i = 1, 2$

$$\frac{P_{X_i, Y_i}(x, y)}{P_{X_i}(x)P_{Y_i}(y)} \leq \frac{1}{P_{X_i}(x)} \leq \frac{1}{m_X}.$$

Therefore, we have

$$\begin{aligned}
|\chi^2(X_1; Y_1) - \chi^2(X_2; Y_2)| &\leq \frac{2}{m_X} \|P_{X_1, Y_1} - P_{X_2, Y_2}\|_1 + \frac{1}{m_X} (\|P_{X_1} - P_{X_2}\|_1 + \|P_{Y_1} - P_{Y_2}\|_1) \\
&\leq \frac{4}{m_X} \|P_{X_1, Y_1} - P_{X_2, Y_2}\|_1. \tag{60}
\end{aligned}$$

Similarly,

$$|\chi^2(S_1; Y_1) - \chi^2(S_2; Y_2)| \leq \frac{4}{m_S} \|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1. \tag{61}$$

By the data processing inequality and the assumption $P_{Y_1|X_1} = P_{Y_2|X_2}$, we have

$$\begin{aligned}
\|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1 &\leq \|P_{S_1, X_1} - P_{S_2, X_2}\|_1, \\
\|P_{X_1, Y_1} - P_{X_2, Y_2}\|_1 &\leq \|P_{S_1, X_1} - P_{S_2, X_2}\|_1.
\end{aligned}$$

Therefore, we get the desired conclusion. \square

D.2 Theorem 6

Recall the following results by Weissman *et al.* [67, Theorem 2.1] for the \mathcal{L}_1 deviation of the empirical distribution.

For all $\epsilon > 0$,

$$\Pr\left(\|\hat{P}_n - P\|_1 \geq \epsilon\right) \leq (2^M - 2) \exp(-n\bar{\phi}(\pi_P)\epsilon^2/4), \tag{62}$$

where P is a probability distribution on the set $[M]$, \hat{P}_n is the empirical distribution obtained from n i.i.d. samples, $\pi_P \triangleq \max_{\mathcal{M} \subseteq [M]} \min(P(\mathcal{M}), 1 - P(\mathcal{M}))$, and

$$\bar{\phi}(p) \triangleq \begin{cases} \frac{1}{1-2p} \ln \frac{1-p}{p} & p \in [0, 1/2), \\ 2 & p = 1/2. \end{cases}$$

Note that $\bar{\phi}(\pi_P) \geq 2$, which implies that

$$\Pr\left(\|\hat{P}_n - P\|_1 \geq \epsilon\right) \leq \exp(M) \exp(-n\epsilon^2/2). \tag{63}$$

Therefore, by choosing $P = P_{S, X}$, $M = |\mathcal{S}||\mathcal{X}|$, and $\epsilon = \sqrt{\frac{2}{n}(M - \ln \beta)}$, (63) implies that, with probability at least $1 - \beta$,

$$\|P_{\hat{S}, \hat{X}} - P_{S, X}\|_1 \leq \sqrt{\frac{2}{n}(M - \ln \beta)}, \tag{64}$$

where $P_{\hat{S},\hat{X}}$ is the empirical distribution obtained from n i.i.d. samples drawn from $P_{S,X}$. Also of note,

$$\begin{aligned} \min\{P_{\hat{S}}(s) \mid s \in \mathcal{S}\} &\geq \min\{P_S(s) \mid s \in \mathcal{S}\} - \|P_{\hat{S}} - P_S\|_1 \\ &\geq \min\{P_S(s) \mid s \in \mathcal{S}\} - \|P_{\hat{S},\hat{X}} - P_{S,X}\|_1. \end{aligned} \quad (65)$$

Similarly,

$$\min\{P_{\hat{X}}(x) \mid x \in \mathcal{X}\} \geq \min\{P_X(x) \mid x \in \mathcal{X}\} - \|P_{\hat{S},\hat{X}} - P_{S,X}\|_1. \quad (66)$$

The proof of Theorem 6 then follows from Lemma 5 and (64), (65), (66).

References

- [1] F. P. Calmon, M. Varia, and M. Médard, “On information-theoretic metrics for symmetric-key encryption and privacy,” in *Proc. 52nd Annu. Allerton Conf. Commun. Control Comput.*, 2014, pp. 889–894.
- [2] H. Wang and F. P. Calmon, “An estimation-theoretic view of privacy,” in *Proc. 55th Annu. Allerton Conf. Commun. Control Comput.*, 2017, pp. 886–893.
- [3] C. Tenopir, S. Allard, K. Douglass, A. U. Aydinoglu, L. Wu, E. Read, M. Manoff, and M. Frame, “Data sharing by scientists: practices and perceptions,” *PLoS ONE*, vol. 6, no. 6, p. e21101, 2011.
- [4] G. Stefansson, “Business-to-business data sharing: A source for integration of supply chains,” *Int. J. Prod. Econ.*, vol. 75, no. 1, pp. 135–146, 2002.
- [5] B. Otjacques, P. Hitzelberger, and F. Feltz, “Interoperability of e-government information systems: Issues of identification and data sharing,” *J. Manag. Inf. Syst.*, vol. 23, no. 4, pp. 29–51, 2007.
- [6] L. Sweeney, “Guaranteeing anonymity when sharing medical data, the datafly system.” in *Proc. the AMIA Annu. Symp.*, 1997, p. 51.
- [7] A. Narayanan and V. Shmatikov, “Robust de-anonymization of large sparse datasets,” in *Proc. 2008 IEEE Symp. on Security and Privacy*, 2008, pp. 111–125.
- [8] L. Sweeney, “Only you, your doctor, and many others may know,” *Technology Science*, vol. 2015092903, no. 9, p. 29, 2015.
- [9] C. Safran, M. Bloomrosen, W. E. Hammond, S. Labkoff, S. Markel-Fox, P. C. Tang, and D. E. Detmer, “Toward a national framework for the secondary use of health data: an American medical informatics association white paper,” *J. Am. Med. Inform. Assoc.*, vol. 14, no. 1, pp. 1–9, 2007.
- [10] L. Sankar, S. R. Rajagopalan, and H. V. Poor, “Utility-privacy tradeoffs in databases: An information-theoretic approach,” *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 6, pp. 838–852, 2013.
- [11] —, “A theory of utility and privacy of data sources,” in *Proc. 2010 IEEE Int. Symp. on Inf. Theory*, 2010, pp. 2642–2646.
- [12] I. Issa, S. Kamath, and A. B. Wagner, “An operational measure of information leakage,” in *Proc. Annu. Conf. on Inf. Sci. and Syst. (CISS)*. IEEE, 2016, pp. 234–239.
- [13] —, “Maximal leakage minimization for the Shannon cipher system,” in *Proc. 2016 IEEE Int. Symp. on Inf. Theory*, 2016, pp. 520–524.
- [14] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, “Information extraction under privacy constraints,” *Information*, vol. 7, no. 1, p. 15, 2016.
- [15] S. Asoodeh, F. Alajaji, and T. Linder, “Privacy-aware mmse estimation,” in *Proc. 2016 IEEE Int. Symp. on Inf. Theory*, 2016, pp. 1989–1993.
- [16] F. P. Calmon, A. Makhdoumi, M. Médard, M. Varia, M. Christiansen, and K. R. Duffy, “Principal inertia components and applications,” *IEEE Trans. Inf. Theory*, vol. 63, no. 8, pp. 5011–5038, 2017.

- [17] F. P. Calmon and N. Fawaz, “Privacy against statistical inference,” in *Proc. 50th Annu. Allerton Conf. Commun. Control Comput.*, 2012, pp. 1401–1408.
- [18] H. O. Hirschfeld, “A connection between correlation and contingency,” in *Proc. Cambridge Philos. Soc.*, vol. 31, no. 4, 1935, pp. 520–524.
- [19] H. Gebelein, “Das statistische problem der korrelation als variations-und eigenwertproblem und sein zusammenhang mit der ausgleichsrechnung,” *ZAMM-Z. Angew. Math. Me.*, vol. 21, no. 6, pp. 364–379, 1941.
- [20] O. Sarmanov, “Maximum correlation coefficient (nonsymmetric case),” *Select. Transl. Math. Statist. Probab.*, vol. 2, pp. 207–210, 1962.
- [21] H. S. Witsenhausen, “On sequences of pairs of dependent random variables,” *SIAM J. Appl. Math.*, vol. 28, no. 1, pp. 100–113, 1975.
- [22] M. Greenacre and T. Hastie, “The geometric interpretation of correspondence analysis,” *J. Amer. Statist. Assoc.*, vol. 82, no. 398, pp. 437–447, 1987.
- [23] A. Rényi, “On measures of dependence,” *Acta Math. Hungarica*, vol. 10, no. 3-4, pp. 441–451, 1959.
- [24] A. Buja, “Remarks on functional canonical variates, alternating least squares methods and ace,” *Ann. Statist.*, vol. 18, no. 3, pp. 1032–1069, 1990.
- [25] A. Makur, F. Kozynski, S.-L. Huang, and L. Zheng, “An efficient algorithm for information decomposition and extraction,” in *Proc. 53rd Annu. Allerton Conf. Commun. Control Comput.*, 2015, pp. 972–979.
- [26] A. Makur and L. Zheng, “Polynomial singular value decompositions of a family of source-channel models,” *IEEE Trans. Inf. Theory*, vol. 63, no. 12, pp. 7716–7728, 2017.
- [27] C. Dwork, “Differential privacy,” in *Encyclopedia of Cryptography and Security*. Springer, 2011, pp. 338–340.
- [28] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Proc. TCC*, 2006, pp. 265–284.
- [29] I. Mironov, “Rényi differential privacy,” in *Proc. 2017 IEEE Computer Security Foundations Symposium (CSF)*, 2017, pp. 263–275.
- [30] P. Kairouz, S. Oh, and P. Viswanath, “The composition theorem for differential privacy,” *IEEE Trans. Inf. Theory*, vol. 63, no. 6, pp. 4037–4049, 2017.
- [31] D. Kifer and A. Machanavajjhala, “Pufferfish: A framework for mathematical privacy definitions,” *ACM Trans. Database Syst.*, vol. 39, no. 1, p. 3, 2014.
- [32] S. Asodeh, F. Alajaji, and T. Linder, “Notes on information-theoretic privacy,” in *Proc. 52nd Annu. Allerton Conf. Commun. Control Comput.*, 2014, pp. 1272–1278.
- [33] A. Makhdoumi, S. Salamatian, N. Fawaz, and M. Médard, “From the information bottleneck to the privacy funnel,” in *IEEE Inf. Theory Workshop (ITW)*, 2014, pp. 501–505.
- [34] N. Tishby, F. C. Pereira, and W. Bialek, “The information bottleneck method,” *arXiv preprint physics/0004057*, 2000.
- [35] A. Makhdoumi and N. Fawaz, “Privacy-utility tradeoff under statistical uncertainty,” in *Proc. 51st Annu. Allerton Conf. Commun. Control Comput.*, 2013, pp. 1627–1634.
- [36] D. Rebollo-Monedero, J. Forne, and J. Domingo-Ferrer, “From t-closeness-like privacy to postrandomization via information theory,” *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 11, pp. 1623–1636, 2010.
- [37] J. Liao, L. Sankar, V. Y. Tan, and F. P. Calmon, “Hypothesis testing in the high privacy limit,” in *Proc. 54th Annu. Allerton Conf. Commun. Control Comput.*, 2016, pp. 649–656.

- [38] J. Liao, L. Sankar, F. P. Calmon, and V. Y. Tan, “Hypothesis testing under maximal leakage privacy constraints,” in *Proc. 2017 IEEE Int. Symp. on Inf. Theory*, 2017, pp. 779–783.
- [39] I. Issa and A. B. Wagner, “Operational definitions for some common information leakage metrics,” in *Proc. 2017 IEEE Int. Symp. on Inf. Theory*, 2017, pp. 769–773.
- [40] Z. Li and T. J. Oechtering, “Privacy-aware distributed bayesian detection,” *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1345–1357, 2015.
- [41] A. Tripathy, Y. Wang, and P. Ishwar, “Privacy-preserving adversarial networks,” *arXiv preprint arXiv:1712.07008*, 2017.
- [42] C. Huang, P. Kairouz, X. Chen, L. Sankar, and R. Rajagopal, “Context-aware generative adversarial privacy,” *Entropy*, vol. 19, no. 12, p. 656, 2017.
- [43] N. Takbiri, A. Houmansadr, D. L. Goeckel, and H. Pishro-Nik, “Matching anonymized and obfuscated time series to users’ profiles,” *IEEE Trans. Inf. Theory*, vol. 65, no. 2, pp. 724–741, 2019.
- [44] M. Raginsky, “Strong data processing inequalities and Φ -Sobolev inequalities for discrete channels,” *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3355–3389, 2016.
- [45] V. Anantharam, A. A. Gohari, S. Kamath, and C. Nair, “On hypercontractivity and the mutual information between boolean functions,” in *Proc. 51st Annu. Allerton Conf. Commun. Control Comput.*, 2013, pp. 13–19.
- [46] Y. Polyanskiy and Y. Wu, “Dissipation of information in channels with input constraints,” *IEEE Trans. Inf. Theory*, vol. 62, no. 1, pp. 35–55, 2016.
- [47] F. Liese and I. Vajda, “On divergences and informations in statistics and information theory,” *IEEE Trans. Inf. Theory*, vol. 52, no. 10, pp. 4394–4412, 2006.
- [48] L. Paninski, “Estimation of entropy and mutual information,” *Neural computation*, vol. 15, no. 6, pp. 1191–1253, 2003.
- [49] O. Shamir, S. Sabato, and N. Tishby, “Learning and generalization with the information bottleneck,” in *Int. Conf. Alg. Learn. Theory*. Springer, 2008, pp. 92–107.
- [50] F. P. Calmon, D. Wei, B. Vinzamuri, K. N. Ramamurthy, and K. R. Varshney, “Optimized pre-processing for discrimination prevention,” in *Adv. Neural Inf. Process. Syst.*, 2017, pp. 3995–4004.
- [51] H. Wang, M. Diaz, F. P. Calmon, and L. Sankar, “The utility cost of robust privacy guarantees,” in *Proc. 2018 IEEE Int. Symp. on Inf. Theory*, 2018.
- [52] S.-L. Huang, A. Makur, F. Kozynski, and L. Zheng, “Efficient statistics: Extracting information from iid observations,” in *Proc. 52nd Annu. Allerton Conf. Commun. Control Comput.*, 2014, pp. 699–706.
- [53] S. Borade and L. Zheng, “Euclidean information theory,” in *Proc. 45th Annu. Allerton Conf. Commun. Control Comput.*, 2007, pp. 633–640.
- [54] S.-L. Huang, L. Zhang, and L. Zheng, “An information-theoretic approach to unsupervised feature selection for high-dimensional data,” in *IEEE Inf. Theory Workshop (ITW)*, 2017, pp. 434–438.
- [55] H. Witsenhausen and A. Wyner, “A conditional entropy bound for a pair of discrete random variables,” *IEEE Trans. Inf. Theory*, vol. 21, no. 5, pp. 493–501, 1975.
- [56] H. Hsu, S. Asoodeh, S. Salamatian, and F. P. Calmon, “Generalizing bottleneck problems,” in *Proc. 2018 IEEE Int. Symp. on Inf. Theory*, 2018, pp. 531–535.
- [57] B. Wang and F. Zhang, “Some inequalities for the eigenvalues of the product of positive semidefinite hermitian matrices,” *Linear Algebra Appl.*, vol. 160, pp. 113–118, 1992.
- [58] F. P. Calmon, Y. Polyanskiy, and Y. Wu, “Strong data processing inequalities for input constrained additive noise channels,” *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1879–1892, 2018.

- [59] A. Makur and L. Zheng, “Linear bounds between contraction coefficients for f -divergences,” *arXiv preprint arXiv: 1510.01844 v3*, 2017.
- [60] C. T. Li and A. El Gamal, “Maximal correlation secrecy,” *IEEE Trans. Inf. Theory*, vol. 64, no. 5, pp. 3916–3926, 2018.
- [61] F. P. Calmon, M. Varia, M. Médard, M. M. Christiansen, K. R. Duffy, and S. Tessaro, “Bounds on inference,” in *Proc. 51st Annu. Allerton Conf. Commun. Control Comput.*, 2013, pp. 567–574.
- [62] R. Durrett, *Probability: theory and examples*. Cambridge university press, 2010.
- [63] S. Diamond and S. Boyd, “CVXPY: A Python-embedded modeling language for convex optimization,” *J. Mach. Learn. Res.*, vol. 17, no. 83, pp. 1–5, 2016.
- [64] S. L. Warner, “Randomized response: A survey technique for eliminating evasive answer bias,” *J. Amer. Statist. Assoc.*, vol. 60, no. 309, pp. 63–69, 1965.
- [65] P. Kairouz, S. Oh, and P. Viswanath, “Extremal mechanisms for local differential privacy,” in *Adv. Neural Inf. Process. Syst.*, 2014, pp. 2879–2887.
- [66] R. O’Donnell, “Some topics in analysis of Boolean functions,” in *Proc. 40th ACM Symp. on Theory of Computing*. ACM, 2008, pp. 569–578.
- [67] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdu, and M. J. Weinberger, “Inequalities for the L_1 deviation of the empirical distribution,” *Hewlett-Packard Labs, Tech. Rep*, 2003.
- [68] M. Lichman, “UCI machine learning repository,” 2013. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [69] C. F. Van Loan, “Matrix computations (Johns Hopkins studies in mathematical sciences),” 1996.