

On the Robustness of Information-Theoretic Privacy Measures and Mechanisms

Mario Diaz, Hao Wang, *Student Member, IEEE*, Flavio P. Calmon, *Member, IEEE*, Lalitha Sankar, *Senior Member, IEEE*

Abstract—Consider a data publishing setting for a dataset composed by both private and non-private features. The publisher uses an empirical distribution, estimated from n i.i.d. samples, to design a privacy mechanism which is applied to new fresh samples afterward. In this paper, we study the discrepancy between the privacy-utility guarantees for the empirical distribution, used to design the privacy mechanism, and those for the true distribution, experienced by the privacy mechanism in practice. We first show that, for any privacy mechanism, these discrepancies vanish at speed $O(1/\sqrt{n})$ with high probability. These bounds follow from our main technical results regarding the Lipschitz continuity of the considered information leakage measures. Then we prove that the optimal privacy mechanisms for the empirical distribution approach the corresponding mechanisms for the true distribution as the sample size n increases, thereby establishing the statistical consistency of the optimal privacy mechanisms. Finally, we introduce and study *uniform privacy mechanisms* which, by construction, provide privacy to all the distributions within a neighborhood of the estimated distribution and, thereby, guarantee privacy for the true distribution with high probability.

Index Terms—Robustness, information leakage measures, privacy-utility trade-off, large deviations.

I. INTRODUCTION

PUBLISHING of individual-level data is increasingly pervasive [2]–[4]. In many settings, directly disclosing unmodified data may lead to a privacy risk through unwanted inferences, particularly if the data contains (or is correlated with) private information. For example, suppose that a hospital wishes to share patient-level data to further advance research on a certain disease. As shown by Sweeney [5], simply removing private features from the dataset (e.g., name) does not necessarily guarantee privacy as these features usually correlate with the remaining features in the dataset (e.g., ZIP code and gender). This modern data publishing issue, also present in politics [6], business [7], and welfare [8] to name

Manuscript received November 14, 2018; revised June 14, 2019; accepted August 17, 2019. This material is based upon work supported by the National Science Foundation under Grant No. CCF-1350914, CIF-1815361, and an ASU seed grant. Some parts of this paper were presented at the 2018 IEEE International Symposium on Information Theory (ISIT) [1]. *Mario Diaz and Hao Wang are co-first authors.*

M. Diaz is with Centro de Investigación en Matemáticas, A.C., Guanajuato, GTO 36023 Mexico (email: diaztorres@cimat.mx). A substantial part of this work was developed while being a postdoctoral scholar at Arizona State University and Harvard University.

H. Wang and F. P. Calmon are with Harvard University, Cambridge, MA 02138 USA (e-mail: hao_wang@g.harvard.edu; flavio@seas.harvard.edu).

L. Sankar is with Arizona State University, Tempe, AZ 85287 USA (e-mail: lsankar@asu.edu).

a few, requires data disclosure methodologies with provable privacy guarantees, whilst still enabling a desired level of utility to be derived from the disclosed data.

A common approach to ensure privacy consists of perturbing the data using a *privacy mechanism*: a (possibly) randomized mapping designed to control private information leakage, the nature of which depends on how privacy is quantified. Different notions of privacy leakage have been proposed to capture the ability of different adversaries to learn (estimate) private information, e.g., Shannon's mutual information [9], [10], k -anonymity [11], differential privacy [12], maximal leakage [13], total variation [14], among others. In addition, privacy mechanisms should also guarantee the statistical utility of the disclosed data, usually quantified by some measure of similarity between the original and the perturbed data, e.g., distortion [15], f -divergence [16], minimum mean-squared error [17]. In general, privacy and utility objectives compete with each other, making the design of privacy mechanisms a non-trivial task. When data privacy and statistical utility are measured using information-theoretic quantities (e.g., mutual information), most methods for the analysis and design of privacy mechanisms rely on the implicit assumption that the data distribution is, for the most part, known [10], [14], [18]–[25]. In practice, the designer has access only to a sample from the true distribution.

In this work we consider the following setup. We assume that data has both private and non-private features and, based on a sample of such pairs, the designer creates a randomized mapping called a privacy mechanism. As new samples arrive, the designed mechanism is applied to the non-private features in order to produce a sanitized version of them which is later disclosed. In this context, we assume that the adversary (data analyst) knows the true distribution of the data, the privacy mechanism being used, and the disclosed data set. The adversary's objective is then to illegitimately infer the private features associated to the disclosed data. We illustrate this procedure in Figure 1. Observe that this adversary is the strongest in terms of statistical knowledge and hence serves as a worst case benchmark. Apart from its statistical knowledge, we assume that the adversary has no side information regarding the disclosed data.

Since privacy mechanisms are designed using a sample, their privacy-utility guarantees might not generalize to the true distribution, thereby creating a privacy threat as the *de facto* guarantees might be considerably different from those

in the design phase. In this work, we study the effect of the discrepancy between the empirical and true distributions on the analysis and design of privacy mechanisms.

First, we consider the setting where the privacy mechanism is designed using an estimate of the data distribution to evaluate the privacy-utility guarantees. We derive bounds for the discrepancy between the privacy-utility guarantees for the empirical distribution (type) and the true (unknown) data distribution. In this context, these bounds can be used to assess the *de facto* guarantees of privacy mechanisms when deployed in practice.

Next, we investigate the *statistical consistency* of the optimal privacy mechanisms¹. More specifically, assume that the privacy mechanism designer constructs an optimal privacy mechanism for the empirical distribution. As the number of samples increases, the optimal privacy mechanism designed for the empirical distribution naturally changes. We show that if such a sequence of privacy mechanisms converges, its limit is an optimal privacy mechanism for the true data distribution.

Finally, we introduce the notion of *uniform* privacy mechanism. When privacy is a priority, the privacy mechanism designer may be required to guarantee a specific level of privacy for the true distribution despite having access only to an estimate of it. Motivated by this setting, we consider privacy mechanisms that, by design, assure privacy for *every* distribution within a specific neighborhood of the empirical distribution. In this case, large deviations results imply that privacy is guaranteed for the true distribution with a certain probability (depending on the neighborhood). Since privacy is guaranteed *uniformly* on a neighborhood of the empirical distribution, we name these privacy mechanisms as *uniform privacy mechanisms*.

The paper is organized as follows. In the remainder of this section we present our main contributions, review related work, and introduce notation used in this paper. In Section II we review preliminaries, including information leakage measures and large deviations results, and recall the framework of privacy-utility trade-offs. Our main results for the discrepancy of privacy-utility guarantees, convergence properties, and uniform privacy mechanisms are presented in Sections III, IV, and V, respectively. Finally, we illustrate some of the results derived in this paper through two numerical experiments in Section VI and provide concluding remarks in Section VII.

A. Contributions

In Section III, we provide probabilistic upper bounds for the difference between the privacy-utility guarantees for the empirical and the true distributions under five different information metrics: probability of correctly guessing, f -information with f locally Lipschitz, Arimoto's mutual information (α -leakage) of order $\alpha > 1$, Sibson's mutual information of order $\alpha > 1$, and maximal α -leakage of order $\alpha > 1$. These bounds, which scale as $O(1/\sqrt{n})$ with n being the sample size, hold uniformly across all privacy mechanisms and, hence, can be

¹Loosely speaking, a privacy mechanism is said to be optimal if it delivers the best possible utility constrained to a prescribed amount of privacy leakage. Precise definitions are given in Section II-C.

used even for privacy mechanisms that do not have explicit descriptions, e.g., the privacy mechanisms implemented using neural networks in [26]. Explicit constants depending on the information metrics under consideration are provided. The proofs of these bounds rely on known large deviations results [27] and Lipschitz continuity properties of information leakage measures established in this paper. These continuity properties can be combined with results for other estimation frameworks different from the large deviations one, e.g., the ℓ_1 -minimax setting in [28] and references therein. In those cases, the role of the empirical estimator is naturally replaced by other estimators, e.g., the add-constant estimator in [28].

We study the convergence properties of optimal privacy mechanisms in Section IV. Specifically, we consider the case when a privacy mechanism designer constructs a sequence of optimal privacy mechanisms for a sequence of joint distributions converging to the true distribution. In this context, we show that while the sequence of optimal privacy mechanisms do not necessarily form a Cauchy sequence, the distance between the privacy mechanisms in the sequence and the set of optimal privacy mechanisms with respect to the true distribution does go to zero. This convergence is analyzed for information measures which satisfy three technical conditions outlined in Section IV. These conditions are satisfied by probability of correctly guessing, f -information with f locally Lipschitz, and Arimoto's mutual information (α -leakage) of order $\alpha > 1$. As a by-product, we prove that, under these conditions, the privacy-utility function as a function of the joint distribution is continuous except possibly at the boundary of its domain.

In Section V, we introduce the notion of uniform privacy mechanism and prove the existence of *optimal uniform privacy mechanisms*, i.e., uniform privacy mechanisms that attain the best utility in a max-min sense. Optimal uniform privacy mechanisms are considerably harder to design than their non-uniform counterparts as privacy has to be guaranteed for any distribution within a neighborhood of the empirical distribution. Nonetheless, we show that they can be approximated by certain mechanisms designed to deliver privacy *only* for the empirical distribution. This approximation result circumvents the highly non-trivial task of designing optimal uniform privacy mechanisms.

B. Related Work

Differential privacy is a popular measure of privacy which aims at answering queries while simultaneously ensuring privacy of individual records in the database [12], [29]. It does not take into account the distribution of the entire dataset which results in its robustness with respect to the dataset distribution, a property that is not satisfied *a priori* by information-theoretic privacy mechanisms. In this work we analyze the degree of robustness to variations of the dataset distribution of certain information-theoretic measures of privacy.

The study of privacy from an information-theoretic point of view heavily depends on the chosen information metric. One commonly-used information leakage metric is Shannon's mutual information [30] and its generalizations [31]. In particular, both α -leakage and f -information — a special case of

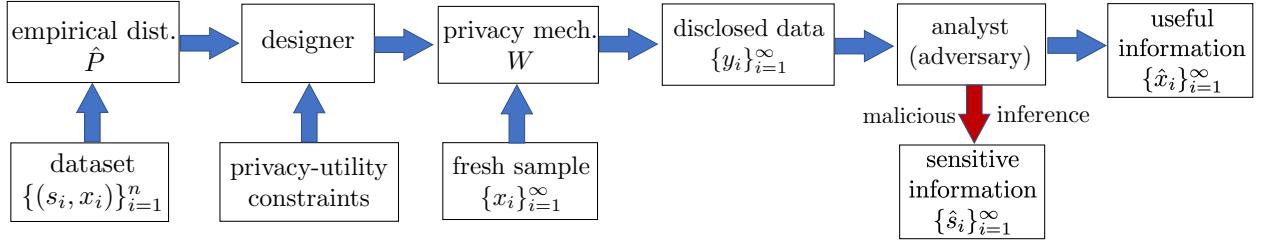


Fig. 1: Typical setting for the design and deployment of privacy mechanisms.

Csiszár's f -divergence [32] — have been used to formulate the privacy-utility trade-offs problem, see e.g., Liao *et al.* [33]. These generalizations seek to study privacy using information metrics that have appeared in the information theory literature and carry some operational (statistical) meaning. There are also information measures which have been introduced specifically in the context of privacy and information leakage. For example, Issa *et al.* [13], [34] introduced a metric called maximal leakage to measure privacy leakage, later extended by Liao *et al.* [35] in a tunable measure called α -leakage. Duchi *et al.* [36] presented locally differential privacy. We refer the reader to Wagner and Eckhoff [37] for a survey of privacy metrics and remark that information metrics also shed light on many secrecy problems, e.g., cryptosystems with short keys [38], side-channel attack [39], and entropic security [40], [41].

Many methods to design information-theoretic privacy mechanisms (see Figure 1) have been proposed in the past, see, e.g., [9], [10], [17], [18], [39], [42]. However, all these works assume that the designer has full knowledge of the data distribution. In practice, the designer may only have access to n i.i.d. samples, creating a mismatch between the distribution used to design the privacy mechanism and the true one. We significantly generalize the work of Wang and Calmon [22], which studies the robustness of privacy mechanisms when privacy and utility are measured using χ^2 -information, to encompass a wide range of information leakage measures. To be more specific, we show that for every information leakage measure considered (see Theorem 1), the discrepancy between the privacy-utility guarantees for the empirical and true distributions scales as $O(1/\sqrt{n})$. It is important to remark that, in the context of the information bottleneck, Shamir *et al.* [43] established similar upper bounds for mutual information that scale as $O(\log(n)/\sqrt{n})$. Hence, mutual information appears to fall outside the scope of the techniques presented in this paper. A related problem is that of estimating information-theoretic measures from a limited number of samples, as studied by, for example, Jiao *et al.* [44], Wu and Yang [45] and Issa *et al.* [39] in the context of privacy. Our work differs from those as we do not focus on the construction of estimators suited for a specific information leakage measure, instead we analyze the performance of the plug-in estimator under different leakage measures. Given the ease of implementation of this estimator, we believe our contributions might be of interest to practitioner.

The fundamental trade-offs between privacy guarantees and statistical utility in data disclosure have been studied in several papers from an information-theoretic view. For example, Yamamoto [15] developed the trade-off between rate, distortion, and equivocation for a specific source coding model. Rebollo-Monedero *et al.* [9], Calmon and Fawaz [18], and Sankar *et al.* [10] characterized privacy-utility trade-offs using tools from rate-distortion theory. Varodayan and Khisti [46] and Tan *et al.* [47] considered the privacy-energy efficiency trade-off in smart meter systems. Makhdoumi *et al.* [48] showed the connection between the privacy funnel and the information bottleneck proposed by Tishby *et al.* [49]. The privacy-utility function, as presented in Section II-C, has been analyzed to find the fundamental limits of privacy-utility trade-offs under different metrics, see, e.g., Calmon *et al.* [50] and Asoodeh *et al.* [20]. It has been shown that this function exhibits some common properties across different metrics. For example, it is continuous, concave, and strictly increasing with respect to the privacy parameter, see also Wang and Calmon [22] and Asoodeh *et al.* [17].

The study of robustness has appeared in many different areas, such as optimization [51], statistics [52], artificial intelligence [53], and machine learning [54]. The notion of robustness considered in this work is closely related to generalization in machine learning [55], which aim at analyzing different performances of the learner's output between training and testing phases. More specifically, a learner may output a classifier h which minimizes a given empirical risk over a sample $S = (z_1, \dots, z_m)$, namely,

$$L_S(h) \triangleq \frac{1}{m} \sum_{i=1}^m \ell(h, z_i), \quad (1)$$

where ℓ is a given loss function. Nonetheless, in practice, the same classifier h might exhibit a very different true risk, defined as

$$L(h) \triangleq \mathbb{E} [\ell(h, Z)], \quad (2)$$

where Z is a random variable distributed according to the dataset distribution. The goal of generalization is to analyze the difference between the empirical risk and the true risk and derive bounds for this discrepancy. This problem has been studied in computer science [56] and, recently, has been analyzed using tools from information theory by Xu and Raginsky [57] and Russo and Zou [58]. Since most information-theoretic metrics do not seem to be expressible

in terms of a loss function as in (2), classical generalization results cannot be directly applied for analyzing the robustness of information-theoretic privacy mechanisms. Recently, the robustness of privacy mechanisms has been investigated using tools from information theory in the works of Wang and Calmon [22] and Issa *et al.* [39]. It is important to mention that some techniques developed for privacy preservation, such as differential privacy, can be useful to analyze generalization guarantees, as shown by Dwork *et al.* [59]–[61].

C. Notation

Given random variables U and U' supported over a finite alphabet \mathcal{U} and another random variable V supported over a finite alphabet \mathcal{V} , we let $P_{U'} \cdot P_{V|U}$ be the joint distribution over $\mathcal{U} \times \mathcal{V}$ determined by

$$(P_{U'} \cdot P_{V|U})(u, v) = P_{U'}(u)P_{V|U}(v|u). \quad (3)$$

Observe that with this notation, $P_U \cdot P_{V|U} = P_{U,V}$. Given another random variable \tilde{V} supported over a finite alphabet $\tilde{\mathcal{V}}$, we let $P_{\tilde{V}|U}P_{V|\tilde{V}}$ be the transition probability matrix determined by

$$(P_{\tilde{V}|U}P_{V|\tilde{V}})(v|u) = \sum_{\tilde{v} \in \tilde{\mathcal{V}}} P_{\tilde{V}|U}(\tilde{v}|u)P_{V|\tilde{V}}(v|\tilde{v}). \quad (4)$$

We let S and X be two random variables with discrete supports \mathcal{S} and \mathcal{X} , respectively. We let P denote the joint distribution $P_{S,X}$ and \hat{P} be a generic estimate of P . The empirical distribution obtained from n i.i.d. samples drawn from P is denoted by \hat{P}_n . We let $(P_n)_{n=1}^\infty$ be any sequence of joint distributions converging to the joint distribution P . Any generic distribution over $\mathcal{S} \times \mathcal{X}$ is denoted by Q . Also, any sequence of distributions over $\mathcal{S} \times \mathcal{X}$ is denoted by $(Q_n)_{n=1}^\infty$.

Let $(\mathcal{A}, d_{\mathcal{A}})$ and $(\mathcal{B}, d_{\mathcal{B}})$ be two metric spaces. We say that $f : \mathcal{A} \rightarrow \mathcal{B}$ is Lipschitz over $\mathcal{A}' \subseteq \mathcal{A}$ if there exists $L \geq 0$ such that, for all $a_1, a_2 \in \mathcal{A}'$,

$$d_{\mathcal{B}}(f(a_1), f(a_2)) \leq L d_{\mathcal{A}}(a_1, a_2). \quad (5)$$

For convenience, we summarize some common notation in Table I.

II. PRELIMINARIES AND PROBLEM SETUP

In this section, we review preliminary material regarding information leakage measures and large deviations bounds for distribution estimation. In addition, we formally introduce the problems addressed in this paper.

A. Information Leakage Measures

Here we review the definition and some basic properties of information leakage measures which are commonly used in the literature. An information leakage measure $\mathcal{L}(U \rightarrow V)$ quantifies how much information V leaks about U . Specifically, U and V can represent raw and disclosed datasets, respectively, and, in this case, $\mathcal{L}(U \rightarrow V)$ measures the utility of the disclosed dataset. On the other hand, when U and V represent private information and information available to an adversary, then $\mathcal{L}(U \rightarrow V)$ quantifies the adversary's ability

of inferring U from the observation of V . In the privacy literature it is valuable to relate information leakage measures with a specific loss/gain function (implicitly) adopted by an adversary. In what follows, we review a range of information leakage measures along with their operational meanings in terms of loss/gain functions.

1) α -Leakage: In an effort to unify several information leakage measures within a single framework, Liao *et al.* recently introduced an information leakage measure called α -leakage [35].

Definition 1 (Definition 4, [35]). Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. For $\alpha \in (1, \infty)$, the α -leakage from U to V is defined as

$$\begin{aligned} \mathcal{L}_\alpha(U \rightarrow V) & \triangleq \frac{\alpha}{\alpha-1} \log \frac{\max_{\hat{U}: U \rightarrow V \rightarrow \hat{U}} \mathbb{E} \left[\Pr(\hat{U} = U|U, V)^{\frac{\alpha-1}{\alpha}} \right]}{\max_{\hat{U}: U \perp\!\!\!\perp \hat{U}} \mathbb{E} \left[\Pr(\hat{U} = U|U)^{\frac{\alpha-1}{\alpha}} \right]}. \end{aligned} \quad (6)$$

The value of $\mathcal{L}_\alpha(U \rightarrow V)$ is extended by continuity to $\alpha = 1$ and $\alpha = \infty$.

In terms of the adversarial setting described at the beginning of this section, α -leakage with $\alpha \in (1, \infty]$ can be interpreted as the multiplicative increase, upon observing V , of the expected gain of an adversary. In order to make this observation precise, note that the expected value in the denominator of the logarithmic term in (6) equals

$$\mathbb{E} \left[\Pr(\hat{U} = U|U)^{\frac{\alpha-1}{\alpha}} \right] = \sum_{u \in \mathcal{U}} P_U(u)P_{\hat{U}}(u)^{\frac{\alpha-1}{\alpha}}. \quad (7)$$

In particular, the RHS of (7) equals expected gain of $P_{\hat{U}}$ w.r.t. the gain function

$$\text{gain}(u, P_{\hat{U}}) = P_{\hat{U}}(u)^{\frac{\alpha-1}{\alpha}}. \quad (8)$$

Thus, the denominator of the logarithmic term in (6) is the largest expected gain of an adversary with no additional information apart from the distribution of U . Also, the expected value in the numerator of the logarithmic term in (6) equals

$$\begin{aligned} & \mathbb{E} \left[\Pr(\hat{U} = U|U, V)^{\frac{\alpha-1}{\alpha}} \right] \\ &= \sum_{u \in \mathcal{U}} \sum_{v \in \mathcal{V}} P_{U,V}(u, v)P_{\hat{U}|V}(u|v)^{\frac{\alpha-1}{\alpha}}. \end{aligned} \quad (9)$$

Thus, the numerator of the logarithmic term in (6) is the largest expected gain of an adversary that has access to V and the joint distribution of U and V . From these observations we conclude that $\mathcal{L}_\alpha(U \rightarrow V)$ measures the multiplicative increase of the best expected gain upon observing V . In a similar way, $\mathcal{L}_1(U \rightarrow V)$ can be interpreted as additive increase, upon observing V , of the expectation of the gain function

$$\text{gain}(u, P_{\hat{U}}) = \log P_{\hat{U}}(u). \quad (10)$$

In addition to its operational definition, α -leakage can be related to Arimoto's mutual information [62], [63] whose definition is provided below.

Category	Notation	Meaning
Norms	$\ a\ _\alpha$	α -norm with $\alpha \in [1, \infty]$: $(\sum_i a_i ^\alpha)^{1/\alpha}$
	$\ a\ _\infty$	∞ -norm: $\max_i a_i $
Joint Distributions	P	True joint distribution: $P_{S,X}$
	\hat{P}	Estimated distribution: $\hat{P}_{S,X}$
	\hat{P}_n	Empirical distribution obtained from n i.i.d. samples
	$(P_n)_{n=1}^\infty$	Any sequence of joint distributions converging to P
	Q	Any joint distribution over $\mathcal{S} \times \mathcal{X}$
	$(Q_n)_{n=1}^\infty$	Any sequence of joint distributions over $\mathcal{S} \times \mathcal{X}$
Channels	W	Any privacy mechanism (channel) from \mathcal{X} to \mathcal{Y}
Sets	\mathcal{W}	Set of all privacy mechanisms
	\mathcal{W}_N	Set of all privacy mechanisms from \mathcal{X} to $\{1 \dots, N\}$
	\mathcal{P}	Set of all joint distributions over $\mathcal{S} \times \mathcal{X}$
	\mathcal{Q}	Any closed subset of \mathcal{P}

TABLE I: List of notation used in this paper.

Definition 2. Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. Their Arimoto's mutual information of order $\alpha \in (1, \infty)$ is defined as

$$I_\alpha^A(P_{U,V}) \triangleq \frac{\alpha}{\alpha-1} \log \frac{\sum_v \|P_{U,V}(\cdot, v)\|_\alpha}{\|P_U(\cdot)\|_\alpha}. \quad (11)$$

Also, by continuous extension,

$$I_1^A(P_{U,V}) \triangleq I(P_{U,V}), \quad (12)$$

$$I_\infty^A(P_{U,V}) \triangleq \log \frac{\sum_v \max_u P_{U,V}(u, v)}{\max_u P_U(u)}, \quad (13)$$

where $I(P_{U,V})$ denotes Shannon's mutual, i.e.,

$$I(P_{U,V}) \triangleq \sum_{u \in \mathcal{U}} \sum_{v \in \mathcal{V}} P_{U,V}(u, v) \log \frac{P_{U,V}(u, v)}{P_U(u)P_V(v)}. \quad (14)$$

In [35], Liao *et al.* proved that α -leakage is indeed equal to Arimoto's mutual information of order α .

Proposition 1 (Theorem 1, [35]). *Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. For $\alpha \in [1, \infty]$, α -leakage satisfies*

$$\mathcal{L}_\alpha(U \rightarrow V) = I_\alpha^A(P_{U,V}). \quad (15)$$

The previous proposition shows that α -leakage recovers Shannon's mutual information in the extremal case $\alpha = 1$. It is important to point out that Shannon's mutual information has been extensively used in the context of privacy. Of particular relevance for the present papers are the works of Rebollo-Monedero *et al.* [9], Sankar *et al.* [10], Calmon and Fawaz [18], Asoodeh *et al.* [20], and references therein. On the other extreme, when $\alpha = \infty$, α -leakage is closely related to another information leakage measure: probability of correctly guessing. This information leakage measure, whose definition is recalled next, has been used recently in the context of

privacy-utility trade-offs by Asoodeh *et al.* [17] and in the broader privacy literature by Smith [64], Braun *et al.* [65], Barthe and Kopf [66], and references therein.

Definition 3. Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. The probability of correctly guessing U and the probability of correctly guessing U given V are given by

$$P_c(U) \triangleq \max_{u \in \mathcal{U}} P_U(u), \quad (16)$$

$$P_c(U|V) \triangleq \sum_{v \in \mathcal{V}} \max_{u \in \mathcal{U}} P_{U,V}(u, v). \quad (17)$$

As its name suggests, $P_c(U)$ is equal to the (largest) probability of correctly guessing U without any side information but the distribution of U . Similarly, $P_c(U|V)$ is equal to the (largest) probability of correctly guessing U given V and the joint distribution $P_{U,V}$. This interpretation can be made precise using the adversarial setting and the gain function defined in (8) with $\alpha = \infty$. Indeed, a simple manipulation shows that

$$P_c(U) = \max_{\hat{U}: U \perp\!\!\!\perp \hat{U}} \mathbb{E}(\mathbb{1}_{U=\hat{U}}), \quad (18)$$

$$P_c(U|V) = \max_{\hat{U}: U \rightarrow V \rightarrow \hat{U}} \mathbb{E}(\mathbb{1}_{U=\hat{U}}), \quad (19)$$

which corresponds to the adversarial setting under the 0-1 loss. Observe that under this loss, the optimal adversary strategy is the same as the maximum a posteriori (MAP) decoder in the communication literature. Finally, observe that

$$\mathcal{L}_\infty(U \rightarrow V) = \log \frac{P_c(U|V)}{P_c(U)}, \quad (20)$$

evidencing the intrinsic relation between ∞ -leakage and probability of correctly guessing.

Remark 1. Despite the close relation between ∞ -leakage and probability of correctly guessing, we always treat them

separately. On the one hand, probability of correctly guessing is of interest in its own right. On the other hand, the results for probability of correctly guessing are easier to derive than their ∞ -leakage counterparts. As a result, working explicitly with probability of correctly guessing allows us to present the key techniques used in this paper while keeping the technical difficulties at a minimum.

In [34], [39], Issa *et al.* introduced the notion of maximal leakage in order to quantify the adversary's capability to infer any (randomized) function of U from V . Motivated by this notion, Liao *et al.* introduced *maximal α -leakage*, a tunable information leakage measure that is equal to maximal leakage when $\alpha = \infty$.

Definition 4 (Definition 5, [35]). Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. For $\alpha \in [1, \infty]$, the maximal α -leakage from U to V is defined as

$$\mathcal{L}_\alpha^{\max}(U \rightarrow V) \triangleq \sup_{T: T \rightarrow U \rightarrow V} \mathcal{L}_\alpha(T \rightarrow V), \quad (21)$$

where T represents any (randomized) function of U that takes values on a finite but arbitrary alphabet.

Arimoto's mutual information can be regarded as a generalization of Shannon's mutual information, see, e.g., [31]. Another such generalization is Sibson's mutual information [67] which, as shown below, is also related to maximal α -leakage.

Definition 5 (Definition 4, [31]). Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. Their Sibson's mutual information of order $\alpha \in (1, \infty)$ is defined as

$$I_\alpha^S(P_{U,V}) \triangleq \frac{\alpha}{\alpha-1} \log \sum_{v \in \mathcal{V}} \|P_U(\cdot)^{1/\alpha} P_{V|U}(v|\cdot)\|_\alpha. \quad (22)$$

Also, $I_1^S(P_{U,V}) = I(P_{U,V})$ and

$$I_\infty^S(P_{U,V}) = \log \left(\sum_{v \in \mathcal{V}} \max_{u \in \mathcal{U}} P_{V|U}(v|u) \right). \quad (23)$$

The following proposition shows that maximal α -leakage is the Arimoto's capacity under an input support constraint which, in turn, is equal to Sibson's capacity under the same constraint.

Proposition 2 (Theorem 2, [35]). Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. For $\alpha \in (1, \infty]$, maximal α -leakage satisfies

$$\mathcal{L}_\alpha^{\max}(U \rightarrow V) = \sup_{P_{\tilde{U}}} I_\alpha^A(P_{\tilde{U}} \cdot P_{V|U}) \quad (24)$$

$$= \sup_{P_{\tilde{U}}} I_\alpha^S(P_{\tilde{U}} \cdot P_{V|U}), \quad (25)$$

where the support of $P_{\tilde{U}}$ is a subset of the support of P_U ². Also, $\mathcal{L}_1^{\max}(U \rightarrow V) = I(P_{U,V})$.

In particular, maximal α -leakage of order infinity (i.e., maximal leakage) is the Sibson's mutual information of order infinity.

²In [35, Theorem 2], $P_{\tilde{U}}$ ranges over the distributions with the same support as P_U . However, their proof readily implies Proposition 2 as stated in this work.

Proposition 3 (Corollary 1, [34]). Let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. Then,

$$\mathcal{L}_\infty^{\max}(U \rightarrow V) = I_\infty^S(P_{U,V}). \quad (26)$$

2) *f*-Information: We finish our review on information leakage measures by recalling the definition and some basic properties of *f*-information. This information leakage measure is defined in terms of Csiszár's *f*-divergence whose definition is recalled next for the reader's convenience.

Definition 6 (Definition 1.1, [32]). Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a convex function with $f(1) = 0$. Assume that Q_1 and Q_2 are two probability distributions over a finite set \mathcal{Z} and that $Q_1 \ll Q_2$. The *f*-divergence between Q_1 and Q_2 is given by

$$D_f(Q_1 \| Q_2) \triangleq \sum_{z \in \mathcal{Z}} Q_2(z) f \left(\frac{Q_1(z)}{Q_2(z)} \right). \quad (27)$$

Given a function f as in the previous definition, its corresponding *f*-information is defined as follows.

Definition 7. Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a convex function with $f(1) = 0$. Furthermore, let U and V be two random variables supported over finite sets \mathcal{U} and \mathcal{V} , respectively. Their *f*-information is defined by

$$I_f(P_{U,V}) \triangleq D_f(P_{U,V} \| P_U P_V) = \sum_{u \in \mathcal{U}} \sum_{v \in \mathcal{V}} P_U(u) P_V(v) f \left(\frac{P_{U,V}(u,v)}{P_U(u) P_V(v)} \right). \quad (28)$$

While an *f*-information may not have a straightforward interpretation in operational terms for a specific function f , many functions f do, e.g., mutual information and χ^2 -information discussed below. Hence, a general treatment allows us to simultaneously handle those *f*-information for which there is a concrete operational meaning and those whose operational interpretation is yet to be discovered. More recent developments about the properties of *f*-divergence can be found in Raginsky [68], Calmon *et al.* [69], and the references therein.

In the context of privacy, χ^2 -information is a fine example of an *f*-information with a tangible operational interpretation. Following the standard convention, the χ^2 -information between two random variables U and V is defined as $\chi^2(P_{U,V}) \triangleq I_f(P_{U,V})$ with $f(t) = (t-1)^2$. Relying on the so-called *principal inertia components* (PICs) [70], Calmon *et al.* [71] showed that if $\chi^2(U; V) < \epsilon$ for some $0 < \epsilon < 1$, then the *minimum mean-squared error* (MMSE) of reconstructing any zero-mean unit-variance function of U given V is lower bounded by $1 - \epsilon$, i.e., no function of U can be reconstructed with small MMSE given an observation of V . Thus, χ^2 -information measures an adversary's ability to reconstruct functions of U from V under an MMSE loss. Recently, χ^2 -information has been used in the context of privacy-utility trade-offs by Wang and Calmon in [22].

B. Large Deviations Bounds for Distribution Estimation

Now we review some preliminary results regarding the distance between the empirical and true distributions. Throughout

this work, we measure the distance between two probability distributions over \mathcal{Z} , say Q_1 and Q_2 , by their ℓ_1 -distance which is given by

$$\|Q_1 - Q_2\|_1 \triangleq \sum_{z \in \mathcal{Z}} |Q_1(z) - Q_2(z)|. \quad (29)$$

A result by Weissman *et al.* [27, Theorem 2.1] establishes that, for all $\epsilon > 0$,

$$\Pr\left(\|\hat{P}_n - P\|_1 \geq \epsilon\right) \leq (2^{|\mathcal{Z}|} - 2) \exp(-n\phi(\pi_P)\epsilon^2/4), \quad (30)$$

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

$$\phi(p) \triangleq \begin{cases} \frac{1}{1-2p} \log \frac{1-p}{p} & p \in [0, 1/2], \\ 2 & p = 1/2. \end{cases} \quad (31)$$

Note that $\phi(p) \geq 2$ for all $p \in [0, 1/2]$. Hence,

$$\Pr\left(\|\hat{P}_n - P\|_1 \geq \epsilon\right) \leq \exp(|\mathcal{Z}|) \exp(-n\epsilon^2/2). \quad (32)$$

By taking $\mathcal{Z} = \mathcal{S} \times \mathcal{X}$ and $\epsilon = \sqrt{\frac{2}{n}(|\mathcal{S}|\cdot|\mathcal{X}| - \log \beta)}$, inequality (32) implies that, with probability at least $1 - \beta$,

$$\|\hat{P}_n - P\|_1 \leq \sqrt{\frac{2}{n}(|\mathcal{S}|\cdot|\mathcal{X}| - \log \beta)}. \quad (33)$$

Even though in this paper we focus on large deviations results, it is worth pointing out that the order $O(\sqrt{|\mathcal{S}|\cdot|\mathcal{X}|/n})$ is present in other fundamental settings, e.g., the minimax expected loss framework in [28, Corollary 9].

C. Problem Setup

Suppose that S is a variable to be hidden (e.g., political preference) and X is an observed variable (e.g., movie ratings) that is correlated with S . In order to receive some utility (e.g., personalized recommendations), we would like to disclose as much information about X without compromising S . An approach with rigorous privacy guarantees is to release a new random variable Y produced by applying a randomized mapping to X , hence forming a Markov chain $S \rightarrow X \rightarrow Y$. This mapping, called the privacy mechanism, is designed to satisfy a specific privacy constraint.

In the sequel, we assume that S and X are discrete random variables with support sets \mathcal{S} and \mathcal{X} , respectively. Let \mathcal{W} be the set of all privacy mechanisms which take X as input and display a discrete random variable Y as output. More specifically, let

$$\mathcal{W} \triangleq \bigcup_{N \geq 1} \left\{ W \in [0, 1]^{|\mathcal{X}| \times N} : \forall x \in \mathcal{X}, \sum_{y=1}^N W(x, y) = 1 \right\}. \quad (34)$$

Through this work, we denote the joint distribution of S and X by P and the privacy mechanism producing Y from X by W , i.e., $W = P_{Y|X}$. The privacy leakage and the utility generated by a mapping $W \in \mathcal{W}$ for the true distribution P are denoted by $\mathcal{L}(P, W)$ and $\mathcal{U}(P, W)$, respectively. Throughout this paper, \mathcal{L} and \mathcal{U} are either probability of correctly guessing,

f -information with f locally Lipschitz, Arimoto's mutual information (α -leakage) of order $\alpha \in (1, \infty]$, Sibson's mutual information of order $\alpha \in (1, \infty]$, or maximal α -leakage of order $\alpha \in (1, \infty]$.

In this framework, a natural problem is to characterize the fundamental trade-off between privacy and utility as captured by the following definition. We denote by \mathcal{P} the set of all probability distributions over $\mathcal{S} \times \mathcal{X}$.

Definition 8. For a given $P \in \mathcal{P}$ and $\epsilon \geq \inf_{W \in \mathcal{W}} \mathcal{L}(P, W)$, the *privacy-utility function* is defined as

$$\mathsf{H}(P; \epsilon) \triangleq \sup_{W \in \mathcal{D}(P; \epsilon)} \mathcal{U}(P, W), \quad (35)$$

where $\mathcal{D}(P; \epsilon) \triangleq \{W \in \mathcal{W} : \mathcal{L}(P, W) \leq \epsilon\}$. Furthermore, the collection of all optimal privacy mechanisms for P at ϵ is defined as

$$\mathcal{W}^*(P; \epsilon) \triangleq \{W \in \mathcal{W} : \mathcal{L}(P, W) \leq \epsilon, \mathcal{U}(P, W) = \mathsf{H}(P; \epsilon)\}. \quad (36)$$

Observe that, by definition, $\mathcal{D}(P; \epsilon)$ is the set of all privacy mechanisms providing an ϵ -privacy guarantee for P . Hence, the privacy-utility function in Definition 8 quantifies the best utility achieved by any privacy mechanism providing an ϵ -privacy guarantee for P .

This specific type of privacy-utility trade-off (PUT), and the optimal privacy mechanisms associated to it, has been investigated for several measures of privacy and utility, see, for example, [14], [17], [20], [22]. These investigations often rely on the implicit assumption that the data distribution is, for the most part, known. However, in practice, the data distribution may only be accessed through a limited number of samples. In this work, we revisit this assumption and study its implications in the design and performance of privacy mechanisms. Next we present the problems addressed and the main contributions of this paper.

1) *Discrepancy of Privacy-Utility Guarantees:* In practice, the designer may not have access to the true distribution P , but only to samples drawn from this distribution. In this case, the privacy-utility guarantees for a distribution estimated from the samples, say \hat{P} , and the true distribution P might be different. For any given privacy mechanism W , these discrepancies are effectively quantified by

$$|\mathcal{L}(\hat{P}, W) - \mathcal{L}(P, W)| \quad \text{and} \quad |\mathcal{U}(\hat{P}, W) - \mathcal{U}(P, W)|. \quad (39)$$

In Section III we provide probabilistic upper bounds for the discrepancies in (39) when the estimated distribution \hat{P} is the empirical distribution \hat{P}_n of n i.i.d. samples drawn from P . These upper bounds depend on the sample size n , the alphabet sizes $|\mathcal{S}|$ and $|\mathcal{X}|$, and, in some cases, the probability of the least likely symbol of the marginals of \hat{P}_n . Their derivations rely on the large deviations results recalled in Section II-B and continuity properties of information leakage measures established in this work. We summarize these results in the following meta theorem (see Theorem 1). For ease of notation, we let

$$\mathcal{L}_c(Q, W) \triangleq P_c(S|Y) \quad \text{and} \quad \mathcal{U}_c(Q, W) \triangleq P_c(X|Y), \quad (40)$$

Theorem 1. Let \hat{P}_n be the empirical distribution obtained from n i.i.d. samples drawn from P . Then, with probability at least $1 - \beta$, for any $W \in \mathcal{W}$, we have

$$|\mathcal{L}(\hat{P}_n, W) - \mathcal{L}(P, W)| \leq \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)} \cdot \begin{cases} 1 & \text{when } \mathcal{L} = \mathcal{L}_c, \\ C_{f, \bar{m}_S} & \text{when } \mathcal{L} = \mathcal{L}_f \text{ with } f \text{ locally Lipschitz,} \\ \frac{2\alpha}{\alpha-1} |\mathcal{S}|^{1-1/\alpha} & \text{when } \mathcal{L} = \mathcal{L}_\alpha^A \text{ with } \alpha \in (1, \infty], \\ \frac{2\alpha+1}{(\alpha-1)\bar{m}_S^{1-1/\alpha}} & \text{when } \mathcal{L} = \mathcal{L}_\alpha^S \text{ with } \alpha \in (1, \infty), \\ \frac{2}{\min_s \sum_x \hat{P}_n(s, x)} & \text{when } \mathcal{L} = \mathcal{L}_\infty^S \text{ or } \mathcal{L} = \mathcal{L}_\infty^{\max}, \\ \frac{4\alpha |\mathcal{S}|^{1-1/\alpha}}{(\alpha-1) \min_s \sum_x \hat{P}_n(s, x)} & \text{when } \mathcal{L} = \mathcal{L}_\alpha^{\max} \text{ with } \alpha \in (1, \infty), \end{cases} \quad (37)$$

$$|\mathcal{U}(\hat{P}_n, W) - \mathcal{U}(P, W)| \leq \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)} \cdot \begin{cases} 1 & \text{when } \mathcal{U} = \mathcal{U}_c, \\ C_{f, \bar{m}_X} & \text{when } \mathcal{U} = \mathcal{U}_f \text{ with } f \text{ locally Lipschitz,} \\ \frac{2\alpha}{\alpha-1} |\mathcal{X}|^{1-1/\alpha} & \text{when } \mathcal{U} = \mathcal{U}_\alpha^A \text{ with } \alpha \in (1, \infty], \\ \frac{1}{(\alpha-1)\bar{m}_X^{1-1/\alpha}} & \text{when } \mathcal{U} = \mathcal{U}_\alpha^S \text{ with } \alpha \in (1, \infty), \\ 0 & \text{when } \mathcal{U} = \mathcal{U}_\infty^S \text{ or } \mathcal{U} = \mathcal{U}_\alpha^{\max} \text{ with } \alpha \in (1, \infty], \end{cases} \quad (38)$$

where $C_{f,u} = 2K_{f,u-1} + (2u^{-1} + 1)L_{f,u-1}$ with $K_{g,u}$, $L_{g,u}$, \bar{m}_S and \bar{m}_X as defined in (46), (47), (61), and (62), respectively.

where $S \rightarrow X \rightarrow Y$ is such that $P_{S,X} = Q$ and $P_{Y|X} = W$. With this notation, we also let $\mathcal{L}_f(Q, W) \triangleq I_f(P_{S,Y})$, $\mathcal{L}_\alpha^A(Q, W) \triangleq I_\alpha^A(P_{S,Y})$, $\mathcal{L}_\alpha^S(Q, W) \triangleq I_\alpha^S(P_{S,Y})$, and, by abuse of notation, $\mathcal{L}_\alpha^{\max}(Q, W) \triangleq \mathcal{L}_\alpha^{\max}(S \rightarrow Y)$. The analogues for utility are defined in a similar way.

We illustrate the discrepancies of privacy-utility guarantees in (37) and (38) along with the corresponding upper bounds through a synthetic and a real-world datasets in Section VI.

2) *Convergence of Optimal Privacy Mechanisms:* Assume there is a sequence of joint distributions $(P_n)_{n=1}^\infty$ converging to the true distribution P . The results provided in Section III establish that the privacy-utility guarantees for P_n converge to those of P as n goes to infinity. Nonetheless, they cannot be used to study the convergence properties of the optimal privacy mechanisms of P_n as defined in (36). In Section IV we address the convergence of optimal privacy mechanisms by further exploiting some of the Lipschitz continuity properties established for the information leakage measures under consideration. To be more specific, assume that the privacy mechanism designer constructs an optimal privacy mechanism W_n^* for each P_n , i.e., $W_n^* \in \mathcal{W}^*(P_n; \epsilon)$. We establish the convergence of the sequence of optimal privacy mechanisms $(W_n^*)_{n=1}^\infty$ to the set of optimal privacy mechanisms for P , i.e., $\mathcal{W}^*(P; \epsilon)$. Due to technical considerations, this result covers only probability of correctly guessing, f -information with f locally Lipschitz, and Arimoto's mutual information of order α with $\alpha \in (1, \infty]$.

3) *Uniform Privacy Mechanisms:* In applications where privacy is a priority, a specific privacy guarantee for the true distribution P may be required, even though the designer has only access to a distribution \hat{P} estimated from samples. We propose the following procedure to overcome this difficulty: (a) use large deviations results to find a probabilistic upper bound, say r , for the distance between \hat{P} and P ; (b) design privacy mechanisms that deliver the required privacy guarantee

for *all* distributions at distance less or equal than r from \hat{P} . Based on this procedure, we introduce the definition of uniform privacy mechanisms in Section V. We prove that *optimal* uniform privacy mechanisms exist and while their design might be challenging, they can be efficiently approximated by optimal privacy mechanisms for \hat{P} as defined in (36). We finish Section V establishing convergence properties of optimal uniform privacy mechanisms similar to those in Section IV.

III. DISCREPANCY OF PRIVACY-UTILITY GUARANTEES

We provide probabilistic upper bounds for the difference between the privacy-utility guarantees of the empirical and the true distributions. These upper bounds do not depend on the specific privacy mechanism. In order to simplify the exposition, we assume that privacy leakage and utility are measured using the same information metric. Nonetheless, the case when different information metrics are used can be handled in a straightforward manner. In this section, we study five different information metrics: probability of correctly guessing, f -information with f locally Lipschitz, Arimoto's mutual information (α -leakage) of order $\alpha > 1$, Sibson's mutual information of order $\alpha > 1$, and maximal α -leakage of order $\alpha > 1$.

A. Probability of Correctly Guessing

The main result of this subsection relies on the following lemma which establishes the Lipschitz continuity of the mappings $\mathcal{L}_c(\cdot, W)$ and $\mathcal{U}_c(\cdot, W)$. Recall that for $Q \in \mathcal{P}$ and $W \in \mathcal{W}$,

$$\mathcal{L}_c(Q, W) \triangleq P_c(S|Y) \quad \text{and} \quad \mathcal{U}_c(Q, W) \triangleq P_c(X|Y), \quad (41)$$

where $S \rightarrow X \rightarrow Y$ is such that $P_{S,X} = Q$ and $P_{Y|X} = W$.

Lemma 1. For any $Q_1, Q_2 \in \mathcal{P}$ and $W \in \mathcal{W}$, we have

$$|\mathcal{L}_c(Q_1, W) - \mathcal{L}_c(Q_2, W)| \leq \|Q_1 - Q_2\|_1, \quad (42)$$

$$|\mathcal{U}_c(Q_1, W) - \mathcal{U}_c(Q_2, W)| \leq \|Q_1 - Q_2\|_1. \quad (43)$$

Proof. See Appendix A-A. \square

In particular, when $Q_1 = \hat{P}$ is an estimate of $Q_2 = P$, the previous lemma implies that any upper bound on $\|\hat{P} - P\|_1$ translates into an upper bound for difference of the corresponding privacy-utility guarantees. The following theorem specializes this observation to the empirical distribution by means of the large deviations inequality in (33).

Theorem 2. Let \hat{P}_n be the empirical distribution obtained from n i.i.d. samples drawn from P . Then, with probability at least $1 - \beta$, for any $W \in \mathcal{W}$, we have

$$|\mathcal{L}_c(\hat{P}_n, W) - \mathcal{L}_c(P, W)| \leq \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}, \quad (44)$$

$$|\mathcal{U}_c(\hat{P}_n, W) - \mathcal{U}_c(P, W)| \leq \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}. \quad (45)$$

In Section VI-B we illustrate the fitness of the bounds in Theorem 2 when applied to ProPublica's COMPAS dataset [72].

Remark 2. Lemma 1 and Theorem 2 illustrate the general technique we use to derive (probabilistic) upper bounds for the difference between the privacy-utility guarantees of the empirical and the true distributions. Specifically, we relate the difference between the privacy-utility guarantees of two joint distributions with their ℓ_1 distance, and then we use large deviations results to provide upper bound for the ℓ_1 distance between the empirical and the true distributions.

B. f -Information with f Locally Lipschitz

We start establishing the following notation. For a given function $g : [0, \infty) \rightarrow \mathbb{R}$ and $u > 0$, we let

$$K_{g,u} \triangleq \sup\{|g(t)| : t \in [0, u]\}. \quad (46)$$

The constant $K_{g,u}$ is the so-called supremum norm of g on $[0, u]$. In addition, if g is Lipschitz on $[0, u]$, we let $L_{g,u}$ be its Lipschitz constant on $[0, u]$, i.e.,

$$L_{g,u} \triangleq \min\{L \geq 0 : |g(t_1) - g(t_2)| \leq L|t_1 - t_2|, \forall t_1, t_2 \in [0, u]\}. \quad (47)$$

A function $g : [0, \infty) \rightarrow \mathbb{R}$ is called locally Lipschitz if g is Lipschitz on $[0, u]$ for every $u > 0$. Note that a locally Lipschitz function is not necessarily Lipschitz on $[0, \infty)$. For example, the function $g(t) = t^2$ is locally Lipschitz with $L_{g,u} = 2u$ for all $u > 0$, but it is not Lipschitz on $[0, \infty)$. For any two distributions $Q_1, Q_2 \in \mathcal{P}$, we denote

$$m_S \triangleq \min \left\{ \sum_{x \in \mathcal{X}} Q_i(s, x) : s \in \mathcal{S}, i \in \{1, 2\} \right\}, \quad (48)$$

$$m_X \triangleq \min \left\{ \sum_{s \in \mathcal{S}} Q_i(s, x) : x \in \mathcal{X}, i \in \{1, 2\} \right\}. \quad (49)$$

Observe that m_S (resp. m_X) equals the probability of the least likely symbol among the marginal distributions over \mathcal{S} (resp. \mathcal{X}) of Q_1 and Q_2 . In other words, if (S_i, X_i) has joint distribution Q_i for each $i \in \{1, 2\}$, then

$$m_S = \min\{P_{S_i}(s) : s \in \mathcal{S}, i \in \{1, 2\}\}, \quad (50)$$

$$m_X = \min\{P_{X_i}(x) : x \in \mathcal{X}, i \in \{1, 2\}\}. \quad (51)$$

The following lemma serves as an analogue of Lemma 1 for an f -information with f locally Lipschitz. Recall that for $Q \in \mathcal{P}$ and $W \in \mathcal{W}$,

$$\mathcal{L}_f(Q, W) \triangleq I_f(P_{S,Y}) \quad \text{and} \quad \mathcal{U}_f(Q, W) \triangleq I_f(P_{X,Y}),$$

where $S \rightarrow X \rightarrow Y$ is such that $P_{S,X} = Q$ and $P_{Y|X} = W$.

Lemma 2. If $f : [0, \infty) \rightarrow \mathbb{R}$ is locally Lipschitz, then, for any $Q_1, Q_2 \in \mathcal{P}$ and $W \in \mathcal{W}$, we have

$$|\mathcal{L}_f(Q_1, W) - \mathcal{L}_f(Q_2, W)| \leq C_{f,m_S} \|Q_1 - Q_2\|_1, \quad (52)$$

$$|\mathcal{U}_f(Q_1, W) - \mathcal{U}_f(Q_2, W)| \leq C_{f,m_X} \|Q_1 - Q_2\|_1, \quad (53)$$

where, for $u \in \{m_S, m_X\}$,

$$C_{f,u} \triangleq 2K_{f,u-1} + (2u^{-1} + 1)L_{f,u-1}. \quad (54)$$

Proof. See Appendix A-B. \square

Remark 3. Despite the similarity between Lemmas 1 and 2, the latter does not imply that the mapping $\mathcal{L}_f(\cdot, W)$ is Lipschitz continuous. Indeed, the factor C_{f,m_S} depends on Q_1 and Q_2 through m_S . Nonetheless, Lemma 2 does show that the local Lipschitzianity of f is bequeathed to $\mathcal{L}_f(\cdot, W)$. Specifically, for every $\delta > 0$, the mapping $\mathcal{L}_f(\cdot, W)$ is Lipschitz continuous over

$$\left\{ Q \in \mathcal{P} : \delta \leq \min_{s \in \mathcal{S}} \sum_{x \in \mathcal{X}} Q(s, x) \right\}, \quad (55)$$

and its Lipschitz constant is less than or equal to $C_{f,\delta}$. Indeed, this assertion follows from the fact that $u \mapsto C_{f,u}$ is a non-increasing function, as exhibited in Appendix A-C. Hence, any lower bound for m_S translates into an upper bound for C_{f,m_S} in (52). As shown below, the *local Lipschitzianity* of $\mathcal{L}_f(\cdot, W)$ is enough to derive a bound similar to (44). A similar discussion can be held for $\mathcal{U}_f(\cdot, W)$.

We illustrate the value of $C_{f,u}$, defined in Lemma 2, through the following examples:

- *Total variation distance.* If $f(t) = |t - 1|/2$, then f is a (globally) Lipschitz function with $L_{f,u} = 1/2$ and $K_{f,u} = \max\{1, u - 1\}/2$ for all $u > 0$. Consequently,

$$C_{f,u} = u^{-1} \max\{1 + 3u/2, 2 - u/2\}; \quad (56)$$

- *χ^2 -divergence.* If $f(t) = (t - 1)^2$, then f is a locally Lipschitz function with $L_{f,u} = 2 \max\{1, u - 1\}$ and $K_{f,u} = \max\{1, (u - 1)^2\}$ for all $u > 0$. Consequently,

$$C_{f,u} = 2u^{-2} \max\{2u + 2u^2, 3 - 3u\}; \quad (57)$$

- *Hellinger divergence.* If $f_\alpha(t) = \frac{t^\alpha - 1}{\alpha - 1}$ with $\alpha > 1$, then f_α is locally Lipschitz with $L_{f_\alpha,u} = \frac{\alpha u^{\alpha-1}}{\alpha - 1}$ and $K_{f,u} = \frac{1}{\alpha-1} \max\{1, u^\alpha - 1\}$ for all $u > 0$. Consequently, $C_{f,u} = \frac{u^{-\alpha}}{\alpha - 1} \max\{2\alpha + \alpha u + 2u^\alpha, (2\alpha + 2) + \alpha u - 2u^\alpha\}$.

Note, however, that mutual information cannot be handled by Lemma 2 since the function $f(t) = t \log(t)$ is not locally Lipschitz. In fact, mutual information seems to have a different asymptotic behavior w.r.t. the sample size n when compared to the theorems in this paper, cf. [43, Theorem 3].

Relying on Lemma 2 and the large deviations inequality in (33), the following theorem provides a data dependent bound for the discrepancy between the guarantees provided for the empirical and the true distributions. For $x \in \mathbb{R}$, we define $(x)_+ \triangleq \max\{0, x\}$.

Theorem 3. *Let \hat{P}_n be the empirical distribution obtained from n i.i.d. samples drawn from P . If $f : [0, \infty) \rightarrow \mathbb{R}$ is locally Lipschitz, then, with probability at least $1 - \beta$, for any $W \in \mathcal{W}$, we have*

$$|\mathcal{L}_f(\hat{P}_n, W) - \mathcal{L}_f(P, W)| \leq C_{f, \bar{m}_S} \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}, \quad (58)$$

$$|\mathcal{U}_f(\hat{P}_n, W) - \mathcal{U}_f(P, W)| \leq C_{f, \bar{m}_X} \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}, \quad (59)$$

where, for $u \in \{\bar{m}_S, \bar{m}_X\}$,

$$C_{f,u} \triangleq 2K_{f,u^{-1}} + (2u^{-1} + 1)L_{f,u^{-1}}, \quad (60)$$

$$\begin{aligned} \bar{m}_S &\triangleq \left(\min \left\{ \sum_{x \in \mathcal{X}} \hat{P}_n(s, x) : s \in \mathcal{S} \right\} \right. \\ &\quad \left. - \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)} \right)_+, \end{aligned} \quad (61)$$

$$\begin{aligned} \bar{m}_X &\triangleq \left(\min \left\{ \sum_{s \in \mathcal{S}} \hat{P}_n(s, x) : x \in \mathcal{X} \right\} \right. \\ &\quad \left. - \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)} \right)_+. \end{aligned} \quad (62)$$

Proof. See Appendix A-C. \square

Remark 4. Observe that up to an additive factor of $(2(|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)/n)^{1/2}$, which is negligible in the large n regime, \bar{m}_S equals the probability of the least likely symbol of the marginal distribution over \mathcal{S} of \hat{P}_n . As a consequence, the bound in (58) is data-dependent, while the bound in (44) is data-independent. This discrepancy comes mainly from the fact that $\mathcal{L}_c(\cdot, W)$ is Lipschitz continuous, as established in Lemma 1, while $\mathcal{L}_f(\cdot, W)$ is *locally Lipschitz* in the sense of Remark 3. A similar discussion holds for utility.

In most practical scenarios, the alphabet of X is significantly larger than the alphabet of S , e.g., S might be political preference while X is movie ratings, S might be private

household information while X is smart meter data, S might be gender while X is a profile picture, etc. In particular, when the sample size is limited, the bounds in Theorem 3 might become ineffective due to the potentially small value of \bar{m}_X . In order to alleviate this issue, we propose the following pre-processing technique which combines the symbols of \mathcal{X} with less observations in the dataset.

Given $\gamma \geq 0$ and a symbol x_0 not belonging to \mathcal{X} , we let Π_γ be the mapping with input alphabet \mathcal{X} , output alphabet

$$\mathcal{X}_\gamma \triangleq \{x_0\} \cup \left\{ x \in \mathcal{X} : \sum_{s \in \mathcal{S}} \hat{P}(s, x) \geq \gamma \right\}, \quad (63)$$

and determined by

$$\Pi_\gamma(x) = \begin{cases} x & \text{if } \sum_s \hat{P}(s, x) \geq \gamma, \\ x_0 & \text{otherwise.} \end{cases} \quad (64)$$

The proposed pre-processing technique consists in applying this mapping to the samples $\{(s_i, x_i)\}_{i=1}^n$ to obtain the modified samples $\{(s_i, \Pi_\gamma(x_i))\}_{i=1}^n$. Clearly, this pre-processing technique improves the bounds in Theorem 3 by increasing \bar{m}_X . However, this improvement might come at a price in utility, as shown in the following proposition.

Proposition 4. *Let $\gamma \geq 0$ be given and let \hat{P}_n be the empirical distribution of n samples $\{(s_i, x_i)\}_{i=1}^n$. If \hat{P}_γ is the empirical distribution of the modified samples $\{(s_i, \Pi_\gamma(x_i))\}_{i=1}^n$, then, for any $\epsilon \in \mathbb{R}$ such that $\mathcal{W}^*(\hat{P}_\gamma; \epsilon) \neq \emptyset$,*

$$\mathsf{H}_f(\hat{P}_\gamma; \epsilon) \leq \mathsf{H}_f(\hat{P}_n; \epsilon), \quad (65)$$

where H_f denotes the privacy-utility function when both privacy leakage and utility are measured using f -information.

Proof. See Appendix A-D. \square

The proof of Proposition 4 relies on the following lemma.

Lemma 3. *Let $\gamma \geq 0$ be given. If $X \rightarrow X_0 \rightarrow Y_0$ is a Markov chain with $X_0 = \Pi_\gamma(X)$, then, for every f -information,*

$$I_f(P_{X,Y_0}) = I_f(P_{X_0,Y_0}). \quad (66)$$

Observe that Lemma 3 is an immediate consequence of the data processing inequality (DPI) for f -information [73, Remark 2.3] and the fact that X_0 is a deterministic function of X . Indeed, if $\Pi : \mathcal{X} \rightarrow \mathcal{X}_0$ is any deterministic function and $X_0 \triangleq \Pi(X)$, then any random variable Y_0 satisfies $X_0 \rightarrow X \rightarrow Y_0$. In particular, the DPI implies that $I_f(P_{X,Y_0}) \geq I_f(P_{X_0,Y_0})$. If, in addition, Y_0 is a randomized function of X_0 , i.e., $X \rightarrow X_0 \rightarrow Y_0$, then the DPI leads to $I_f(P_{X,Y_0}) \leq I_f(P_{X_0,Y_0})$.

C. Arimoto's Mutual Information, Sibson's Mutual Information, and Maximal α -Leakage

We now establish continuity properties of Arimoto's mutual information, Sibson's mutual information, and maximal α -leakage similar to those proved in Lemmas 1 and 2 for probability of correctly guessing and f -information, respectively. Upon these properties, we state two theorems regarding the

discrepancy of privacy-utility guarantees for the information measures at hand.

Recall that for $Q \in \mathcal{P}$ and $W \in \mathcal{W}$,

$$\mathcal{L}_\alpha^A(Q, W) \triangleq I_\alpha^A(P_{S,Y}) \quad \text{and} \quad \mathcal{U}_\alpha^A(Q, W) \triangleq I_\alpha^A(P_{X,Y}),$$

where $S \rightarrow X \rightarrow Y$ is such that $P_{S,X} = Q$ and $P_{Y|X} = W$. The next lemma shows the Lipschitz continuity of $\mathcal{L}_\alpha^A(\cdot, W)$ and $\mathcal{U}_\alpha^A(\cdot, W)$.

Lemma 4. *Let $\alpha \in (1, \infty]$. For any $Q_1, Q_2 \in \mathcal{P}$ and $W \in \mathcal{W}$,*

$$|\mathcal{L}_\alpha^A(Q_1, W) - \mathcal{L}_\alpha^A(Q_2, W)| \leq \frac{2\alpha}{\alpha-1} |\mathcal{S}|^{1-1/\alpha} \|Q_1 - Q_2\|_1, \quad (67)$$

$$|\mathcal{U}_\alpha^A(Q_1, W) - \mathcal{U}_\alpha^A(Q_2, W)| \leq \frac{2\alpha}{\alpha-1} |\mathcal{X}|^{1-1/\alpha} \|Q_1 - Q_2\|_1. \quad (68)$$

Proof. See Appendix A-E. \square

Remark 5. Recall that probability of correctly guessing and Arimoto's mutual information of order ∞ are related through the formula

$$I_\infty^A(U; V) = \log \frac{P_c(U|V)}{P_c(U)}. \quad (69)$$

Observe that, despite this relation, the bound for probability of correctly guessing in (42) and the corresponding bound for Arimoto's mutual information of order ∞ in (67) differ by a factor of $2|\mathcal{S}|$. As shown in the proof of Lemma 4, the $|\mathcal{S}|$ factor comes from the log in (69) via the minimum in the following inequality

$$\left| \log \frac{a}{b} \right| \leq \frac{|a-b|}{\min\{a,b\}}, \quad a, b > 0. \quad (70)$$

A similar argument explains the extra factor of $2|\mathcal{X}|$ appearing in (68) w.r.t. its counterpart in (43).

Now we consider Sibson's mutual information. Recall that for $Q \in \mathcal{P}$ and $W \in \mathcal{W}$,

$$\mathcal{L}_\alpha^S(Q, W) \triangleq I_\alpha^S(P_{S,Y}) \quad \text{and} \quad \mathcal{U}_\alpha^S(Q, W) \triangleq I_\alpha^S(P_{X,Y}),$$

where $S \rightarrow X \rightarrow Y$ is such that $P_{S,X} = Q$ and $P_{Y|X} = W$. The following lemma establishes that the mappings $\mathcal{L}_\alpha^S(\cdot, W)$ and $\mathcal{U}_\alpha^S(\cdot, W)$ have similar continuity properties to those of $\mathcal{L}_f(\cdot, W)$ and $\mathcal{U}_f(\cdot, W)$ with f locally Lipschitz.

Lemma 5. *For $Q_1, Q_2 \in \mathcal{P}$, we define m_S and m_X as in (48) and (49), respectively. If $\alpha \in (1, \infty)$, then, for any $W \in \mathcal{W}$,*

$$|\mathcal{L}_\alpha^S(Q_1, W) - \mathcal{L}_\alpha^S(Q_2, W)| \leq \frac{2\alpha+1}{\alpha-1} \frac{\|Q_1 - Q_2\|_1}{m_S^{1-1/\alpha}}, \quad (71)$$

$$|\mathcal{U}_\alpha^S(Q_1, W) - \mathcal{U}_\alpha^S(Q_2, W)| \leq \frac{1}{\alpha-1} \frac{\|Q_1 - Q_2\|_1}{m_X^{1-1/\alpha}}. \quad (72)$$

If $\alpha = \infty$, then, for any $W \in \mathcal{W}$,

$$|\mathcal{L}_\infty^S(Q_1, W) - \mathcal{L}_\infty^S(Q_2, W)| \leq \frac{2\|Q_1 - Q_2\|_1}{\min_s \sum_x Q_1(s, x)}, \quad (73)$$

$$|\mathcal{U}_\infty^S(Q_1, W) - \mathcal{U}_\infty^S(Q_2, W)| = 0. \quad (74)$$

Proof. See Appendix A-F. \square

Remark 6. A straightforward manipulation shows that

$$\exp(I_\infty^A(P_{U,V})) = \sum_{v \in \mathcal{V}} \max_{u \in \mathcal{U}} \frac{P_U(u)}{\max_{u'} P_U(u')} P_{V|U}(v|u), \quad (75)$$

$$\exp(I_\infty^S(P_{U,V})) = \sum_{v \in \mathcal{V}} \max_{u \in \mathcal{U}} P_{V|U}(v|u). \quad (76)$$

From (75) and (76), we observe that Sibson's mutual information of order infinity is independent of the input distribution P_U , while Arimoto's mutual information of the same order is not. By comparing (67) and (73), we also observe that the privacy leakage bound for Sibson's mutual information of order infinity *does* depend on the probability of the least likely symbol of the input³, while its Arimoto's counterpart does not. These seemingly contradicting facts have a rather intuitive cause: it is hard to estimate $P_{V|U}(\cdot|u)$ reliably if $P_U(u)$ is small. Observe that while $P_{V|U}(v|u)$ appears in both (75) and (76), the factor $P_U(u)/\max_{u'} P_U(u')$ in (75) makes Arimoto's mutual information of order infinity less dependent on symbols u with small probabilities $P_U(u)$. The difficulty in estimating Sibson's mutual information in comparison to its Arimoto counterpart is also natural from a privacy perspective. As proved by Issa *et al.* in [34], Sibson's mutual information guarantees privacy for *any* (possibly randomized) function of S , while Arimoto's mutual information guarantees privacy only for S itself.

We end up this section showing that maximal α -leakage ($\alpha > 1$) behaves in a similar way to Sibson's mutual information of order ∞ . Recall the notation in Section I-C and that, for $Q \in \mathcal{P}$ and $W \in \mathcal{W}$,

$$\mathcal{L}_\alpha^{\max}(Q, W) \triangleq \sup_{P_{\tilde{S}}} I_\alpha^A(P_{\tilde{S}} \cdot (P_{X|S}W)), \quad (77)$$

$$\mathcal{U}_\alpha^{\max}(Q, W) \triangleq \sup_{P_{\tilde{X}}} I_\alpha^A(P_{\tilde{X}} \cdot W), \quad (78)$$

where $S \rightarrow X \rightarrow Y$ is such that $P_{S,X} = Q$ and $P_{Y|X} = W$.

Lemma 6. *Let $\alpha \in (1, \infty)$. For any $Q_1, Q_2 \in \mathcal{P}$ and $W \in \mathcal{W}$,*

$$|\mathcal{L}_\alpha^{\max}(Q_1, W) - \mathcal{L}_\alpha^{\max}(Q_2, W)| \leq \frac{4\alpha}{\alpha-1} \frac{|\mathcal{S}|^{1-1/\alpha}}{\min_s \sum_x Q_1(s, x)} \|Q_1 - Q_2\|_1, \quad (79)$$

$$|\mathcal{U}_\alpha^{\max}(Q_1, W) - \mathcal{U}_\alpha^{\max}(Q_2, W)| = 0. \quad (80)$$

Proof. See Appendix A-G. \square

We summarize the probabilistic upper bounds derived in this subsection in the following two theorems. They follow immediately from Lemma 4, 5, 6 and the large deviations inequality in (33).

³In recent work [39], Issa *et al.* established a bound similar to (81) which also depends on the probability of the least likely symbol of the input. Thus, this dependency seems unavoidable at the moment.

Theorem 4. Let \hat{P}_n be the empirical distribution obtained from n i.i.d. samples drawn from P . Then, with probability at least $1 - \beta$, for any $W \in \mathcal{W}$, we have

$$|\mathcal{L}(\hat{P}_n, W) - \mathcal{L}(P, W)| \leq \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)} \\ \times \begin{cases} \frac{2\alpha}{\alpha-1} |\mathcal{S}|^{1-1/\alpha} & \text{when } \mathcal{L} = \mathcal{L}_\alpha^A \text{ with } \alpha \in (1, \infty], \\ \frac{2\alpha+1}{(\alpha-1)\bar{m}_S^{1-1/\alpha}} & \text{when } \mathcal{L} = \mathcal{L}_\alpha^S \text{ with } \alpha \in (1, \infty), \end{cases}$$

$$|\mathcal{U}(\hat{P}_n, W) - \mathcal{U}(P, W)| \leq \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)} \\ \times \begin{cases} \frac{2\alpha}{\alpha-1} |\mathcal{X}|^{1-1/\alpha} & \text{when } \mathcal{U} = \mathcal{U}_\alpha^A \text{ with } \alpha \in (1, \infty], \\ \frac{1}{(\alpha-1)\bar{m}_X^{1-1/\alpha}} & \text{when } \mathcal{U} = \mathcal{U}_\alpha^S \text{ with } \alpha \in (1, \infty), \end{cases}$$

where \bar{m}_S and \bar{m}_X are defined in (61) and (62), respectively.

Theorem 5. Let \hat{P}_n be the empirical distribution obtained from n i.i.d. samples drawn from P . Then, with probability at least $1 - \beta$, for any $W \in \mathcal{W}$, we have

$$|\mathcal{L}(\hat{P}_n, W) - \mathcal{L}(P, W)| \leq \frac{2\sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}}{\min_s \sum_x \hat{P}_n(s, x)} \\ \times \begin{cases} 1 & \text{when } \mathcal{L} = \mathcal{L}_\infty^S \text{ or } \mathcal{L} = \mathcal{L}_\infty^{\max}, \\ \frac{2\alpha|\mathcal{S}|^{1-1/\alpha}}{\alpha-1} & \text{when } \mathcal{L} = \mathcal{L}_\alpha^{\max} \text{ with } \alpha \in (1, \infty), \end{cases} \quad (81)$$

and $|\mathcal{U}(\hat{P}_n, W) - \mathcal{U}(P, W)| = 0$ when $\mathcal{U} = \mathcal{U}_\infty^S$ or $\mathcal{U} = \mathcal{U}_\infty^{\max}$ with $\alpha \in (1, \infty]$.

IV. CONVERGENCE OF OPTIMAL PRIVACY MECHANISMS

Now we study the consistency of the optimal privacy mechanisms using, among other results, the (local) Lipschitz continuity of the mappings $\mathcal{L}(\cdot, W)$ and $\mathcal{U}(\cdot, W)$ established in the previous section.

Assume that the (privacy) mechanism designer is given an $\epsilon \in \mathbb{R}$ and a sequence of joint distributions $(P_n)_{n=1}^\infty$ converging to the true distribution P . Furthermore, the designer constructs an optimal ϵ -private mechanism W_n^* for each P_n . In this section, we analyze some convergence properties of the sequence of optimal privacy mechanisms $(W_n^*)_{n=1}^\infty$.

In the sequel, we assume that prior knowledge about the true distribution P might be available. Accordingly, we let $\mathcal{Q} \subseteq \mathcal{P}$ be the set of all joint distributions compatible with such prior knowledge. Note that when no prior knowledge is available, we simply let $\mathcal{Q} = \mathcal{P}$. The main results of this and the following section are established under the following conditions.

There exist a closed set $\mathcal{Q} \subseteq \mathcal{P}$ and $N \in \mathbb{N}$ such that

(C.1) *Continuity.* For every $Q \in \mathcal{Q}$, the mappings $\mathcal{L}(Q, \cdot)$ and $\mathcal{U}(Q, \cdot)$ are continuous over

$$\mathcal{W}_N \triangleq \mathcal{W} \cap \mathbb{R}^{|\mathcal{X}| \times N}, \quad (82)$$

where \mathcal{W} , as defined in (34), denotes the set of all privacy mechanisms. Furthermore, the mapping $\mathsf{H}(Q, \cdot)$ is continuous over $[\epsilon_{\min}(Q), \infty)$, where

$$\epsilon_{\min}(Q) \triangleq \inf \{\mathcal{L}(Q, W) : W \in \mathcal{W}\}. \quad (83)$$

(C.2) *Lipschitz Continuity.* There exist positive constants C_L and C_U such that, for all $Q_1, Q_2 \in \mathcal{Q}$ and $W \in \mathcal{W}_N$,

$$|\mathcal{L}(Q_1, W) - \mathcal{L}(Q_2, W)| \leq C_L \|Q_1 - Q_2\|_1, \quad (84)$$

$$|\mathcal{U}(Q_1, W) - \mathcal{U}(Q_2, W)| \leq C_U \|Q_1 - Q_2\|_1. \quad (85)$$

(C.3) *Support.* For each $Q \in \mathcal{Q}$ and $\epsilon \geq \epsilon_{\min}(Q)$, the intersection $\mathcal{W}^*(Q; \epsilon) \cap \mathcal{W}_N$ is not empty.

These three conditions might seem restrictive at a first glance. Nonetheless, as shown in Section IV-A below, they are satisfied by probability of correctly guessing, f -information with f locally Lipschitz, and Arimoto's mutual information of order α with $\alpha \in (1, \infty]$.

Note that condition (C.1) is a continuity requirement. Specifically, it requires the privacy leakage and utility functions to be continuous w.r.t. the privacy mechanism W . Also, it requires the privacy-utility function to be continuous w.r.t. the privacy parameter ϵ . Similarly, condition (C.2) requires the privacy leakage and utility functions to be Lipschitz continuous w.r.t. the joint distribution Q . Observe that the constants C_L and C_U in (84) and (85), respectively, do not depend on the joint distributions Q_1 and Q_2 neither on the privacy mechanism W . Nonetheless, these constants may depend on \mathcal{Q} and N .

Remark 7. Observe that the set \mathcal{W}_N is in correspondence with the set of all privacy mechanisms from \mathcal{X} to $\{1, \dots, N\}$. In particular, condition (C.3) requires that, for each $Q \in \mathcal{Q}$ and $\epsilon \geq \epsilon_{\min}(Q)$, there exists an optimal ϵ -private mechanism for Q supported over $\{1, \dots, N\}$. As a consequence,

$$\mathsf{H}(Q; \epsilon) = \max_{W \in \mathcal{D}_N(Q; \epsilon)} \mathcal{U}(Q, W), \quad (86)$$

where $\mathcal{D}_N(Q; \epsilon) \triangleq \{W \in \mathcal{W}_N : \mathcal{L}(Q, W) \leq \epsilon\}$. Furthermore, the intersection of \mathcal{W}_N and all optimal privacy mechanisms is defined as

$$\begin{aligned} \mathcal{W}_N^*(P; \epsilon) &\triangleq \{W \in \mathcal{W}_N : \mathcal{L}(P, W) \leq \epsilon, \mathcal{U}(P, W) = \mathsf{H}(P; \epsilon)\}. \end{aligned} \quad (87)$$

By comparing (35) and (86), we can observe that condition (C.3) allows us to replace the unbounded set \mathcal{W} with the compact space \mathcal{W}_N in the optimization defining the privacy-utility function H . This technical difference is crucial in the proofs of the subsequent results. In a similar spirit, observe that condition (C.3) also implies that

$$\epsilon_{\min}(Q) = \min \{\mathcal{L}(Q, W) : W \in \mathcal{W}_N\}. \quad (88)$$

A two-variable function might be continuous in each argument without being jointly continuous, see, e.g., [74, Ex. 4.7]. The following lemma is an immediate, yet important, consequence of conditions (C.1–2), as it establishes the joint continuity of the privacy leakage and utility functions.

Lemma 7. Assume that conditions (C.1–2) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. The functions $\mathcal{L}(\cdot, \cdot)$ and $\mathcal{U}(\cdot, \cdot)$ are continuous over $\mathcal{Q} \times \mathcal{W}_N$.

Proof. See Appendix B-A. \square

Building upon Lemma 7, we establish the following proposition which plays a key role in the proofs of the main results of this section.

Proposition 5. *Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. For any given $\epsilon \in \mathbb{R}$, the mapping $H(\cdot; \epsilon)$ is continuous over the set $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$.*

Proof. See Appendix B-B. \square

Remark 8. The mapping $H(\cdot; \epsilon)$ might *not* be continuous over $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) \leq \epsilon\}$. In order to prove this claim, consider the following example.

Example 1. For each $\zeta \in [0, 1/2]$, let Q_ζ be the joint distribution of (S, X_ζ) where $S \sim \text{Ber}(1/2)$ and $P_{X_\zeta|S} = \text{BSC}(\zeta)$. By definition, H_c denotes the privacy-utility function when both privacy leakage and utility are measured using probability of correctly guessing. By Theorem 2 in [17], we have that, for every $\zeta \in [0, 1/2]$,

$$\epsilon_{\min}(Q_\zeta) = 1/2 \quad \text{and} \quad H_c(Q_\zeta; 1/2) = 1/2. \quad (89)$$

Since S and $X_{1/2}$ are independent, we also have that

$$\epsilon_{\min}(Q_{1/2}) = 1/2 \quad \text{and} \quad H_c(Q_{1/2}; 1/2) = 1. \quad (90)$$

Therefore, we have that $\lim_{\zeta \uparrow 1/2} H_c(Q_\zeta; 1/2) \neq H_c(Q_{1/2}; 1/2)$ although $\lim_{\zeta \uparrow 1/2} Q_\zeta = Q_{1/2}$.

Despite that $H(\cdot; \epsilon)$ might be discontinuous at the boundary $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) = \epsilon\}$, as presented in Example 1, in some situations such pathological behavior is not exhibited. In such cases, the following corollary provides conditions for which the pointwise convergence established in Proposition 5 upgrades into uniform convergence. This technical result will be useful in explaining the numerical experiments in Section VI.

Corollary 1. *Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. In addition, assume that there exists $\epsilon_0 \in \mathbb{R}$ such that $\epsilon_{\min}(Q) = \epsilon_0$ for all $Q \in \mathcal{Q}$. If $P_n \in \mathcal{Q}$ for each $n \in \mathbb{N}$, $\lim_n P_n = P$, and $\lim_n H(P_n; \epsilon_0) = H(P; \epsilon_0)$, then $H(P_n; \cdot)$ converges uniformly to $H(P; \cdot)$ over $[\epsilon_0, \infty)$, i.e.,*

$$\lim_{n \rightarrow \infty} \sup_{\epsilon \in [\epsilon_0, \infty)} |H(P_n; \epsilon) - H(P, \epsilon)| = 0. \quad (91)$$

Proof. See Appendix B-C. \square

The next theorem shows that if W_n^* is an optimal ϵ -private mechanism for P_n and $(P_n)_{n=1}^\infty$ converges to P , then the sequence of optimal ϵ -private mechanisms $(W_n^*)_{n=1}^\infty$ converges to the set of optimal ϵ -private mechanisms for P . For $W \in \mathcal{W}_N$ and a subset $\mathcal{W}' \subseteq \mathcal{W}_N$, the distance between W and \mathcal{W}' is defined as

$$\text{dist}(W, \mathcal{W}') \triangleq \inf\{\|W - W'\|_1 : W' \in \mathcal{W}'\}, \quad (92)$$

where $\|W - W'\|_1 = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} |W(x, y) - W'(x, y)|$.

Theorem 6. *Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. Let $\epsilon \in \mathbb{R}$*

be given. If $P_n \in \mathcal{Q}$ for each $n \in \mathbb{N}$, $\lim_n P_n = P$, and $\epsilon_{\min}(P) < \epsilon$, then, for any sequence $(W_n^)_{n=1}^\infty \subset \mathcal{W}_N$ such that $W_n^* \in \mathcal{W}_N^*(P_n; \epsilon)$,*

$$\lim_{n \rightarrow \infty} \text{dist}(W_n^*, \mathcal{W}_N^*(P; \epsilon)) = 0. \quad (93)$$

Furthermore, if $P_n = \hat{P}_n$ is the empirical distribution obtained from n i.i.d. samples drawn from P , then

$$\Pr\left(\lim_{n \rightarrow \infty} \text{dist}(W_n^*, \mathcal{W}_N^*(P; \epsilon)) = 0\right) = 1. \quad (94)$$

Proof. See Appendix B-D. \square

Remark 9. Under the assumptions in Theorem 6, it may be possible that $\mathcal{W}_N^*(P_n; \epsilon)$ is empty for some values of n . Nonetheless, under the same assumptions, $\mathcal{W}_N^*(P_n; \epsilon)$ is non-empty for n sufficiently large. This fact can be easily derived from the continuity of the mapping $\epsilon_{\min}(\cdot)$ which is the content of Lemma 15 in Appendix B-D.

The following corollary follows directly from Theorem 6 and the compactness of $\mathcal{W}_N^*(P; \epsilon)$ established in Lemma 16 in Appendix B-D. It shows that the limit of optimal ϵ -private mechanisms is also an optimal ϵ -private mechanism.

Corollary 2. *In addition to the assumptions in Theorem 6, also assume that $\lim_n W_n^* = W_0$ for some $W_0 \in \mathcal{W}_N$. Then, $W_0 \in \mathcal{W}_N^*(P; \epsilon)$.*

It is important to note that Theorem 6 does not imply that the sequence of optimal privacy mechanisms $(W_n^*)_{n=1}^\infty$ converges to a privacy mechanism. Indeed, it has been noticed in simulation that a small perturbation to the joint distribution might lead to a big change to the optimal privacy mechanism returned by some optimization algorithms. The following theorem, whose proof relies on standard results from *set-valued analysis* [75], shows that given a sequence $(P_n)_{n=1}^\infty$ with $\lim_n P_n = P$, it is possible to choose $W_n^* \in \mathcal{W}_N^*(P_n; \epsilon)$ such that $(W_n^*)_{n=1}^\infty$ is convergent. However, we prove this property only for a residual set which, by definition, is a countable intersection of dense open subsets of $\mathcal{Q} \subset \mathbb{R}^{|\mathcal{S}| \times |\mathcal{X}|}$.

Theorem 7. *Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. For any given $\epsilon \in \mathbb{R}$ and any $\delta > 0$, there exists a residual set*

$$\mathcal{Q}' \subseteq \{Q \in \mathcal{Q} : \epsilon_{\min}(Q) + \delta \leq \epsilon\} \quad (95)$$

which satisfies that:

For any joint distribution $Q_0 \in \mathcal{Q}'$, any sequence of joint distributions $(Q_n)_{n=1}^\infty \subset \{Q \in \mathcal{Q} : \epsilon_{\min}(Q) + \delta \leq \epsilon\}$ with $\lim_n Q_n = Q_0$, and any optimal privacy mechanism $W_0^ \in \mathcal{W}_N^*(Q_0; \epsilon)$, there exists a sequence of privacy mechanisms $(W_n^*)_{n=1}^\infty$ such that $W_n^* \in \mathcal{W}_N^*(Q_n; \epsilon)$ for each $n \in \mathbb{N}$ and $\lim_n W_n^* = W_0^*$.*

Proof. See Appendix B-E. \square

Remark 10. Baire's theorem states that any residual subset of a complete metric space is dense, see, e.g., [76, Thm. 5.6]. In this sense, residual sets are considered to be *big* in topological terms. However, residual sets might be negligible in measure theoretic terms, see, e.g., [77, Exercise 5.3.31].

A. Continuity and Compactness Conditions

In this section we show that probability of correctly guessing, f -information with f locally Lipschitz, Arimoto's mutual information of order α with $\alpha \in (1, \infty]$ satisfy conditions (C.1–3).

1) *Probability of Correctly Guessing:* We choose $\mathcal{Q} = \mathcal{P}$ and $\mathcal{Y} = \{1, \dots, N\}$ with $N \geq |\mathcal{X}|+1$. Note that $\mathcal{Q} = \mathcal{P}$ corresponds to the case where no prior information about the true distribution P is available. Recall that, by definition,

$$\mathcal{L}_c(Q, W) = \sum_{y \in \mathcal{Y}} \max_{s \in \mathcal{S}} \sum_{x \in \mathcal{X}} Q(s, x) W(x, y), \quad (96)$$

$$\mathcal{U}_c(Q, W) = \sum_{y \in \mathcal{Y}} \max_{x \in \mathcal{X}} \sum_{s \in \mathcal{S}} Q(s, x) W(x, y). \quad (97)$$

Since the maximum and the sum of continuous functions are continuous functions as well, we have that, for every $Q \in \mathcal{P}$, the mappings $\mathcal{L}_c(Q, \cdot)$ and $\mathcal{U}_c(Q, \cdot)$ are continuous over \mathcal{W}_N . In [17], Asoodeh *et al.* showed that $H_c(Q; \cdot)$, the privacy-utility function associated to \mathcal{L}_c and \mathcal{U}_c , is continuous and piecewise linear over⁴ $[\epsilon_{\min}(Q), \infty)$. Therefore condition (C.1) is satisfied. Furthermore, Lemma 1 shows that, for any $Q_1, Q_2 \in \mathcal{P}$ and $W \in \mathcal{W}$,

$$|\mathcal{L}_c(Q_1, W) - \mathcal{L}_c(Q_2, W)| \leq \|Q_1 - Q_2\|_1, \quad (98)$$

$$|\mathcal{U}_c(Q_1, W) - \mathcal{U}_c(Q_2, W)| \leq \|Q_1 - Q_2\|_1, \quad (99)$$

which implies that condition (C.2) is satisfied with $C_L = 1$ and $C_U = 1$. It was established in [17] that, for every $\epsilon \geq \epsilon_{\min}(Q)$, there is always an optimal ϵ -private mechanism using at most $|\mathcal{X}|+1$ symbols. As a consequence, condition (C.3) holds true as $N \geq |\mathcal{X}|+1$ by assumption.

2) *f -Information with f Locally Lipschitz:* Assume that the function $f : [0, \infty) \rightarrow \mathbb{R}$ is locally Lipschitz and convex with $f(1) = 0$. Given $\gamma > 0$, we choose $N \geq |\mathcal{X}|+1$ and

$$\mathcal{Q} = \left\{ Q \in \mathcal{P} : \gamma \leq \min_{s \in \mathcal{S}} \sum_{x \in \mathcal{X}} Q(s, x), \gamma \leq \min_{x \in \mathcal{X}} \sum_{s \in \mathcal{S}} Q(s, x) \right\}. \quad (100)$$

In this case, the prior information about the true distribution $P = P_{S,X}$ comes in the form of the assumption that the probability mass functions of S and X are bounded away from 0. Recall that, by definition,

$$\mathcal{L}_f(Q, W) = \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} P_S(s) P_Y(y) f\left(\frac{P_{S,Y}(s, y)}{P_S(s) P_Y(y)}\right), \quad (101)$$

$$\mathcal{U}_f(Q, W) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} P_X(x) P_Y(y) f\left(\frac{P_{X,Y}(x, y)}{P_X(x) P_Y(y)}\right), \quad (102)$$

where $S \rightarrow X \rightarrow Y$ is such that $P_{S,X} = Q$ and $P_{Y|X} = W$. Note that if $Q \in \mathcal{Q}$, then, for all $(s, x, y) \in \mathcal{S} \times \mathcal{X} \times \mathcal{Y}$,

$$\max \left\{ \frac{P_{S,Y}(s, y)}{P_S(s) P_Y(y)}, \frac{P_{X,Y}(x, y)}{P_X(x) P_Y(y)} \right\} \leq \gamma^{-1}. \quad (103)$$

Upon this fact, it is straightforward to verify that the mappings $\mathcal{L}_f(Q, \cdot)$ and $\mathcal{U}_f(Q, \cdot)$ are continuous over \mathcal{W}_N . In [78], Hsu *et al.*

⁴Indeed, in this setting $\epsilon_{\min}(Q) = \max_s \sum_x Q(s, x)$.

established the continuity of $H_f(Q; \cdot)$ over $[0, \infty)$. Hence, condition (C.1) is satisfied. Recall the definitions of $K_{g,u}$ and $L_{g,u}$ in (46) and (47), respectively. Lemma 2 shows that, for any $Q_1, Q_2 \in \mathcal{Q}$ and $W \in \mathcal{W}$,

$$\begin{aligned} & |\mathcal{L}_f(Q_1, W) - \mathcal{L}_f(Q_2, W)| \\ & \leq (2K_{f,\gamma^{-1}} + (2\gamma^{-1} + 1)L_{f,\gamma^{-1}}) \|Q_1 - Q_2\|_1, \end{aligned} \quad (104)$$

$$\begin{aligned} & |\mathcal{U}_f(Q_1, W) - \mathcal{U}_f(Q_2, W)| \\ & \leq (2K_{f,\gamma^{-1}} + (2\gamma^{-1} + 1)L_{f,\gamma^{-1}}) \|Q_1 - Q_2\|_1, \end{aligned} \quad (105)$$

which implies that condition (C.2) is satisfied with

$$C_L = C_U = 2K_{f,\gamma^{-1}} + (2\gamma^{-1} + 1)L_{f,\gamma^{-1}}. \quad (106)$$

For example, if $f(x) = |x - 1|$, then $C_L = C_U \leq 4\gamma^{-1} + 1$; and if $f(x) = x^2 - 1$, then $C_L = C_U \leq 8\gamma^{-2}$. As with probability of correctly guessing, there is always an optimal ϵ -private mechanism using at most $|\mathcal{X}|+1$ symbols [78]. From this fact, condition (C.3) follows immediately.

3) *Arimoto's Mutual Information:* Let $\alpha \in (1, \infty]$. We choose $N \geq |\mathcal{X}|+1$ and $\mathcal{Q} = \mathcal{P}$, i.e., no prior information is assumed. As with the previous information measures, the continuity of the mappings $\mathcal{L}_\alpha^A(Q, \cdot)$ and $\mathcal{U}_\alpha^A(Q, \cdot)$ is evident. We denote the privacy-utility function by H_α^A when both privacy leakage and utility are measured using Arimoto's mutual information of order α . Note that the graph of $H_\alpha^A(Q; \cdot)$ corresponds to the upper boundary of the set

$$\mathcal{A} \triangleq \{(\mathcal{L}_\alpha^A(Q, W), \mathcal{U}_\alpha^A(Q, W)) : W \in \mathcal{W}\}. \quad (107)$$

More specifically, $H_\alpha^A(Q; \epsilon) = \sup\{u : (p, u) \in \mathcal{A}, p \leq \epsilon\}$. Consider the transformation which maps (p, u) to

$$\left(\left\| \sum_{x \in \mathcal{X}} Q(\cdot, x) \right\|_\alpha e^{\alpha p / (\alpha - 1)}, \left\| \sum_{s \in \mathcal{S}} Q(s, \cdot) \right\|_\alpha e^{\alpha u / (\alpha - 1)} \right). \quad (108)$$

It is straightforward to verify that this transformation is one-to-one, continuous, and monotone coordinatewise. Using this transformation, it can be shown that (the upper boundary of) \mathcal{A} is homeomorphic to (the upper boundary of)

$$\mathcal{B} \triangleq \{(\phi(Q, W), \psi(Q, W)) : W \in \mathcal{W}\}, \quad (109)$$

where

$$\phi(Q, W) \triangleq \sum_{y \in \mathcal{Y}} \left\| \sum_{x \in \mathcal{X}} Q(\cdot, x) W(x, y) \right\|_\alpha, \quad (110)$$

$$\psi(Q, W) \triangleq \sum_{y \in \mathcal{Y}} \left\| \sum_{s \in \mathcal{S}} Q(s, \cdot) W(\cdot, y) \right\|_\alpha. \quad (111)$$

By the homogeneity of the α -norm, it can be verified that

$$\phi(Q, [\lambda_1 W_1, \dots, \lambda_K W_K]) = \sum_{k=1}^K \lambda_k \phi(Q, W_k), \quad (112)$$

whenever $K \in \mathbb{N}$, $W_1, \dots, W_K \in \mathcal{W}$, and $\lambda_1, \dots, \lambda_K \geq 0$ with $\sum_k \lambda_k = 1$. A similar equality holds for ψ . Therefore, \mathcal{B} is a convex set and, as a consequence, its upper boundary is the graph of a continuous function. Since the upper boundaries of \mathcal{A} and \mathcal{B} are homeomorphic, we conclude that the mapping $H_\alpha^A(Q; \cdot)$ is continuous. Therefore, condition (C.1) is satisfied.

Also, Lemma 4 implies that condition (C.2) is satisfied with $C_L = \frac{2\alpha}{\alpha-1}|\mathcal{S}|^{1-1/\alpha}$ and $C_U = \frac{2\alpha}{\alpha-1}|\mathcal{X}|^{1-1/\alpha}$. Observe that proving condition (C.3) with $N \geq |\mathcal{X}|+1$ is equivalent to show that any point in the upper boundary of \mathcal{A} can be written as $(\mathcal{L}_\alpha^A(Q, W), \mathcal{U}_\alpha^A(Q, W))$ for some $W \in \mathcal{W}_N$. Since \mathcal{A} and \mathcal{B} are homeomorphic, the latter property can be established from an analogous property for \mathcal{B} which in turn follows from a minor adaptation of the argument in [79].

Remark 11. Regarding condition (C.3), both [17] and [78] build upon the convex analysis argument employed by Witsenhausen and Wyner in [79]. More specifically, they rely on an extension of Carathéodory's theorem, the so-called Fenchel-Eggleston theorem [80], in order to find a bound for the size of the output alphabet of an optimal privacy mechanism.

V. UNIFORM PRIVACY MECHANISMS

In this section we consider the scenario in which delivering privacy is the top priority for the privacy mechanism designer. We introduce privacy mechanisms that, with a certain probability, guarantee privacy for the true distribution despite having access only to an estimate of it. These mechanisms are constructed in the following conceptual way. Recall that large deviations results show that, with a certain probability, the true distribution is within some ℓ_1 ball of the empirical distribution. Consequently, if a mechanism guarantees privacy *uniformly* for every distribution within such an ℓ_1 ball, it necessarily guarantees privacy for the true distribution with at least the same probability. In this section, we prove that there is a well-defined notion of optimality for *uniform* privacy mechanisms and that these optimal mechanisms can be approximated by appropriately chosen (non-uniform) optimal privacy mechanisms as previously defined in (36).

In order to introduce uniform privacy mechanisms precisely, recall that inequality (33) shows that, with probability at least $1 - \exp(-|\mathcal{S}|\cdot|\mathcal{X}| - nr^2/2)$, the true distribution P is within the ℓ_1 ball of radius r centered at the empirical distribution \hat{P}_n ,

$$\mathcal{Q}_r(\hat{P}_n) \triangleq \{Q \in \mathcal{Q} : \|Q - \hat{P}_n\|_1 \leq r\}. \quad (113)$$

Based on this observation, we consider uniform privacy mechanisms which guarantee privacy for every joint distribution within $\mathcal{Q}_r(\hat{P})$, where \hat{P} is any estimate of P . Despite the fact that $P \in \mathcal{Q}_r(\hat{P})$ with high probability, we do not know the true value of P . For this reason, we define optimal uniform privacy mechanisms as those that achieve the best worst-case utility within $\mathcal{Q}_r(\hat{P})$.

Definition 9. Let $\mathcal{Q} \subset \mathcal{P}$ and $N \in \mathbb{N}$. For $\hat{P} \in \mathcal{Q}$, $\epsilon \geq 0$, and $r \geq 0$, we define the set of uniform privacy mechanisms for $\mathcal{Q}_r(\hat{P})$ at ϵ as

$$\mathcal{D}_{\mathcal{Q},N}(\hat{P}; \epsilon, r) \triangleq \bigcap_{Q \in \mathcal{Q}_r(\hat{P})} \{W \in \mathcal{W}_N : \mathcal{L}(Q, W) \leq \epsilon\}. \quad (114)$$

Furthermore, we define the set of optimal uniform privacy mechanisms for $\mathcal{Q}_r(\hat{P})$ at ϵ as

$$\mathcal{W}_{\mathcal{Q},N}^\dagger(\hat{P}; \epsilon, r) \triangleq \arg \max_{W \in \mathcal{D}_{\mathcal{Q},N}(\hat{P}; \epsilon, r)} \mathcal{U}_r(\hat{P}, W), \quad (115)$$

where $\mathcal{U}_r(\hat{P}, W) \triangleq \min_{Q \in \mathcal{Q}_r(\hat{P})} \mathcal{U}(Q, W)$.

Recall that, as defined in Remark 7,

$$\mathcal{D}_N(Q; \epsilon) = \{W \in \mathcal{W}_N : \mathcal{L}(Q, W) \leq \epsilon\} \quad (116)$$

is the set of all privacy mechanisms in \mathcal{W}_N delivering an ϵ -privacy guarantee for Q . Thus,

$$\mathcal{D}_{\mathcal{Q},N}(\hat{P}; \epsilon, r) = \bigcap_{Q \in \mathcal{Q}_r(\hat{P})} \mathcal{D}_N(Q; \epsilon) \quad (117)$$

is the set of all privacy mechanisms in \mathcal{W}_N that deliver an ϵ -privacy guarantee uniformly for all the distributions in $\mathcal{Q}_r(\hat{P})$, i.e., all the distributions at a distance less than or equal to r from \hat{P} . For a given privacy mechanism W , $\mathcal{U}_r(\hat{P}, W)$ quantifies the least utility $\mathcal{U}(Q, W)$ attained by W over all the distributions in $\mathcal{Q}_r(\hat{P})$. Thus, by definition, $\mathcal{W}_{\mathcal{Q},N}^\dagger(\hat{P}; \epsilon, r)$ is the set of all uniform privacy mechanisms for $\mathcal{Q}_r(\hat{P})$ at ϵ with the best worst-case utility.

The following lemma shows that Definition 9 is well-defined under conditions (C.1–2). Specifically, it shows that the infimum defining $\mathcal{U}_r(\hat{P}, W)$ and the supremum defining $\mathcal{W}_{\mathcal{Q},N}^\dagger(\hat{P}; \epsilon, r)$ are attainable.

Lemma 8. Assume that conditions (C.1–2) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. Then

- (i) the infimum $\inf\{\mathcal{U}(Q, W) : Q \in \mathcal{Q}_r(\hat{P})\}$ is attainable for every $W \in \mathcal{W}_N$;
- (ii) the supremum $\sup\{\mathcal{U}_r(\hat{P}, W) : W \in \mathcal{D}_{\mathcal{Q},N}(\hat{P}; \epsilon, r)\}$ is attainable whenever $\mathcal{D}_{\mathcal{Q},N}(\hat{P}; \epsilon, r)$ is not empty.

Proof. See Appendix C-A. \square

Observe that (ii) shows that optimal uniform privacy mechanisms do exist as long as $\mathcal{D}_{\mathcal{Q},N}(\hat{P}; \epsilon, r)$ is not empty. Nonetheless, their construction might be challenging as the construction of optimal privacy mechanisms in the non-uniform sense, as defined in (36), is already non-trivial in most cases. The following theorem shows that optimal privacy mechanism can be modified to deliver a uniform privacy guarantee without incurring in a big cost in terms of utility. Throughout this section, we consistently use W^* to denote optimal privacy mechanisms, i.e., elements in \mathcal{W}_N^* as defined in (87), and W^\dagger to denote their uniform counterparts as defined in (115).

Theorem 8. Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. If $P \in \mathcal{Q}_r(\hat{P})$ and $\epsilon - C_L r \geq \epsilon_{\min}(\hat{P})$, then

$$\emptyset \neq \mathcal{D}_N(\hat{P}; \epsilon - C_L r) \subseteq \mathcal{D}_{\mathcal{Q},N}(\hat{P}; \epsilon, r). \quad (118)$$

Furthermore, for every $W^* \in \mathcal{W}_N^*(\hat{P}; \epsilon - C_L r)$ and every $W^\dagger \in \mathcal{W}_{\mathcal{Q},N}^\dagger(\hat{P}; \epsilon, r)$,

$$\begin{aligned} & \mathcal{U}(P, W^*) \\ & \geq \mathcal{U}(P, W^\dagger) - \left(\mathsf{H}(\hat{P}; \epsilon + C_L r) - \mathsf{H}(\hat{P}; \epsilon - C_L r) + 2C_U r \right). \end{aligned} \quad (119)$$

Proof. See Appendix C-B. \square

Observe that (118) establishes that privacy mechanisms for \hat{P} at $\epsilon - C_L r$ are, in fact, uniform privacy mechanisms at ϵ . In other words, (118) shows that by imposing a slightly stronger privacy requirement for \hat{P} , we obtain *uniform* privacy mechanisms for $\mathcal{Q}_r(\hat{P})$. Furthermore, under condition (C.1),

$$\lim_{r \downarrow 0} \left(H(\hat{P}; \epsilon + C_L r) - H(\hat{P}; \epsilon - C_L r) + 2C_U r \right) = 0. \quad (120)$$

Therefore, (119) shows that, when r is small, any optimal privacy mechanism for \hat{P} at $\epsilon - C_L r$ performs almost as well as any optimal uniform privacy mechanism for $\mathcal{Q}_r(\hat{P})$ at ϵ . Of course, all these conclusions hold as long as $P \in \mathcal{Q}_r(\hat{P})$ but, as pointed out before, this is the case (with high probability) in the large sample size regime.

Remark 12. Note that if $H(\hat{P}; \cdot)$ is Lipschitz continuous with Lipschitz constant L , then (119) becomes

$$\mathcal{U}(P, W^*) \geq \mathcal{U}(P, W^\dagger) - 2(C_U + LC_L)r. \quad (121)$$

Observe that if $H(\hat{P}; \cdot)$ is a convex function differentiable at $\epsilon_0 \triangleq \epsilon_{\min}(\hat{P})$, then $L = H'(\hat{P}; \epsilon_0)$. The value of $H'(\hat{P}; \epsilon_0)$ is closely related with the notion of *reverse* strong data processing inequality [19], see also [20].

Now we consider a specific example to illustrate the above results.

Example 2. Recall that H_c is the privacy-utility function associated to \mathcal{L}_c and \mathcal{U}_c as given in (41). For ease of notation, we define

$$p\#q \triangleq \begin{pmatrix} (1-p)(1-q) & (1-p)q \\ pq & p(1-q) \end{pmatrix}. \quad (122)$$

Let $\mathcal{Q} = \{p\#q : p \in [1/2, 1], q \in [0, 1-p]\}$ and $N \geq 2$. This selection of \mathcal{Q} captures the case where $S \sim \text{Ber}(p)$ with $p \in [1/2, 1]$ and $P_{X|S} = \text{BSC}(q)$ with $q \in [0, 1-p]$. By Theorem 2 in [17], for all $Q = p\#q \in \mathcal{Q}$,

$$H_c(Q; \epsilon) = 1 - \frac{1-q}{p-q}(p+q-2pq) + \epsilon \frac{p+q-2pq}{p-q}, \quad (123)$$

whenever $\epsilon \in [p, 1-q]$. In particular, $H_c(Q; \cdot)$ is Lipschitz continuous with Lipschitz constant $\frac{p+q-2pq}{p-q}$. Recall that, for probability of correctly guessing, $C_L = C_U = 1$ as established in Section IV-A1. Hence, under the assumptions of Theorem 8, (121) becomes

$$\mathcal{U}(P, W^*) \geq \mathcal{U}(P, W^\dagger) - \frac{2\hat{p}(1-\hat{q})}{\hat{p}-\hat{q}}r, \quad (124)$$

where $W^* \in \mathcal{W}_N^*(\hat{P}; \epsilon - r)$ and $W^\dagger \in \mathcal{W}_{\mathcal{Q}, N}^\dagger(\hat{P}; \epsilon, r)$ with $\hat{P} \triangleq \hat{p}\#\hat{q} \in \mathcal{Q}$. Furthermore, by taking $\hat{P} = \hat{P}_n$ and $r = (2(4 - \log \beta)/n)^{1/2}$, inequality (33) implies that, with probability at least $1 - \beta$,

$$\mathcal{U}(P, W^*) \geq \mathcal{U}(P, W^\dagger) - \frac{2\hat{p}(1-\hat{q})}{\hat{p}-\hat{q}} \sqrt{\frac{2}{n}(4 - \log \beta)}. \quad (125)$$

We finish this section by studying some convergence properties of uniform privacy mechanism, similar to those studied in Theorem 6. More specifically, the next theorem shows that although a sequence $(W_n^\dagger)_{n=1}^\infty$ with $W_n^\dagger \in \mathcal{W}_{\mathcal{Q}, N}^\dagger(P_n; \epsilon, r_n)$

may not be convergent, the distance between each W_n^\dagger and $\mathcal{W}_N^*(P; \epsilon)$ converges to zero as long as $\lim_n P_n = P$ and $\lim_n r_n = 0$. Furthermore, it also shows that, under certain conditions, this convergence can be guaranteed almost surely for the empirical distribution estimator.

Theorem 9. Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. Let $\epsilon \in \mathbb{R}$ and $P \in \mathcal{Q}$ be given. If $P_n \in \mathcal{Q}$ with $\|P_n - P\|_1 \leq r_n$ for each $n \in \mathbb{N}$, $\lim_n r_n = 0$, and $\epsilon > \epsilon_{\min}(P)$, then, for any sequence $(W_n^\dagger)_{n=1}^\infty \subset \mathcal{W}_N$ such that $W_n^\dagger \in \mathcal{W}_{\mathcal{Q}, N}^\dagger(P_n; \epsilon, r_n)$ for all $n \geq 1$,

$$\lim_{n \rightarrow \infty} \text{dist}(W_n^\dagger, \mathcal{W}_N^*(P; \epsilon)) = 0. \quad (126)$$

Furthermore, if $P_n = \hat{P}_n$ is the empirical estimator obtained from n i.i.d. samples drawn from P and, for some $p > 1$, $r_n \geq \sqrt{\frac{2p \log(n)}{n}}$ for all $n \geq 1$, then,

$$\Pr \left(\lim_{n \rightarrow \infty} \text{dist}(W_n^\dagger, \mathcal{W}_N^*(P; \epsilon)) = 0 \right) = 1. \quad (127)$$

Proof. See Appendix C-C. \square

The following corollary follows immediately from Theorem 9 by taking $P = P_n = \hat{P}$.

Corollary 3. Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. If $\hat{P} \in \mathcal{Q}$ and $\epsilon > \epsilon_{\min}(\hat{P})$, then

$$\limsup_{r \downarrow 0} \mathcal{W}_{\mathcal{Q}, N}^\dagger(\hat{P}; \epsilon, r) \subseteq \mathcal{W}_N^*(\hat{P}; \epsilon), \quad (128)$$

where

$$\begin{aligned} & \limsup_{r \downarrow 0} \mathcal{W}_{\mathcal{Q}, N}^\dagger(\hat{P}; \epsilon, r) \\ & \triangleq \left\{ W \in \mathcal{W}_N : \liminf_{r \downarrow 0} \text{dist}(W, \mathcal{W}_{\mathcal{Q}, N}^\dagger(\hat{P}; \epsilon, r)) = 0 \right\}. \end{aligned}$$

Theorem 8 shows that in terms of performance, optimal privacy mechanisms approximate optimal uniform privacy mechanisms in the small r regime. The previous corollary shows that, in the same regime, optimal uniform privacy mechanisms are *geometrically close* to the set of optimal privacy mechanisms. These two results evidence an intrinsic relation between optimal privacy mechanisms and their uniform counterparts.

VI. NUMERICAL EXPERIMENTS

We illustrate some of the results derived in the previous sections through two numerical experiments. The first experiment, conducted on a synthetic dataset, illustrates the convergence of the empirical privacy-utility function and optimal privacy mechanisms to their corresponding limits as the sample size increases. The second experiment, performed on a real-world dataset, displays the discrepancy between the privacy-utility guarantees during the design and testing of a privacy mechanism.

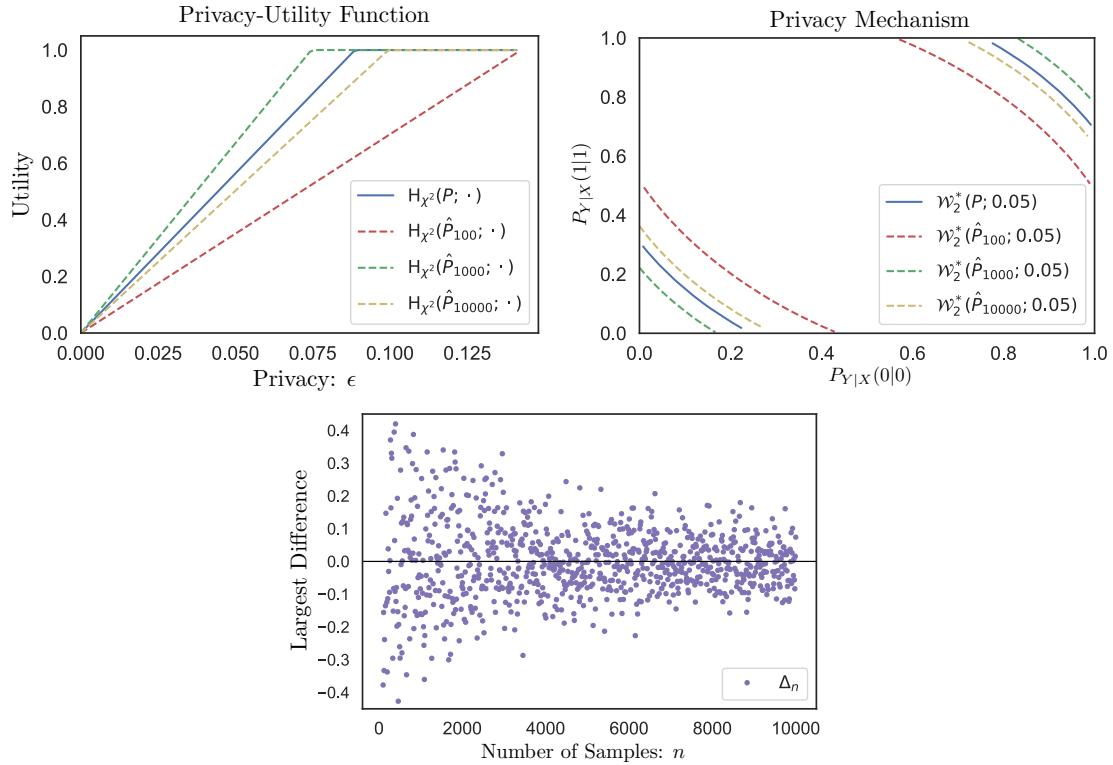


Fig. 2: Both privacy leakage and utility are measured using the f -information with $f(t) = (t - 1)^2$. Top left: privacy-utility function $H_{\chi^2}(P; \cdot)$ and empirical privacy-utility function $H_{\chi^2}(\hat{P}_n; \cdot)$. Top right: corresponding optimal privacy mechanisms for $\epsilon = 0.05$. Bottom: largest (signed) difference between $H_{\chi^2}(P_n; \cdot)$ and $H_{\chi^2}(P; \cdot)$.

A. Synthetic Dataset

Here we illustrate the convergence of the empirical privacy-utility function and optimal privacy mechanisms as the sample size increases. The joint distribution matrix \mathbf{P} , determined by $\mathbf{P}_{s,x} \triangleq P(s, x)$, is chosen to be

$$\mathbf{P} = \begin{pmatrix} 0.42 & 0.18 \\ 0.16 & 0.24 \end{pmatrix}, \quad (129)$$

and both privacy leakage and utility are measured using χ^2 -information, i.e., the f -information associated to the function $f(t) = (t - 1)^2$. For any given value of n , we randomly select n i.i.d. samples $\{(s_i, x_i)\}_{i=1}^n$ drawn from the joint distribution P and compute the empirical distribution \hat{P}_n .

In Figure 2 (top left), we depict the privacy-utility function $H_{\chi^2}(P; \cdot)$ and the empirical privacy-utility function $H_{\chi^2}(\hat{P}_n; \cdot)$ for three different values of n . Note that the privacy-utility function of the empirical distribution converges (pointwise) to the corresponding function for the true distribution, i.e., for any given $\epsilon \geq 0$ we have that $H_{\chi^2}(\hat{P}_n; \epsilon)$ approaches $H_{\chi^2}(P; \epsilon)$ as n grows. This corroborates the conclusion of Proposition 5 which establishes that, for any given $\epsilon \in \mathbb{R}$, $H(\cdot; \epsilon)$ is continuous over $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$. In Figure 2 (top right), we depict the corresponding optimal privacy mechanisms. In order to visualize a privacy mechanism $P_{Y|X}$ in the xy -plane, we let the x -axis and y -axis to be the values of $P_{Y|X}(0|0)$ and $P_{Y|X}(1|1)$, respectively. Observe that optimal privacy mechanisms are not unique. Hence, depending on the algorithm used to obtain such mechanisms, a sequence

of empirical optimal privacy mechanisms may not converge to a fixed mechanism. However, as observed from Figure 2 (top right), the distance between any such sequence and the set of optimal privacy mechanisms for the true distribution will necessarily converge to zero which echoes Theorem 6.

Finally we scatter plot the largest (signed) gap, denoted by Δ_n , between $H_{\chi^2}(P_n; \cdot)$ and $H_{\chi^2}(P; \cdot)$. In particular, the absolute value of Δ_n is equal to the uniform norm of the function $H_{\chi^2}(P_n; \cdot) - H_{\chi^2}(P; \cdot)$, i.e.,

$$|\Delta_n| = \sup_{\epsilon \in [0, \infty)} |H_{\chi^2}(P_n; \epsilon) - H_{\chi^2}(P; \epsilon)|. \quad (130)$$

We depict the value of Δ_n in Figure 2 (bottom). As shown, when the number of samples increases, $|\Delta_n|$ tends to decrease which illustrates the uniform convergence established in Corollary 1. Due to the mismatch between the empirical and true distributions, privacy and utility might be over or under-estimated, leading to the variable sign of Δ_n .

B. ProPublica's COMPAS Recidivism Dataset

We now illustrate Theorems 2 and 4 in Section III through ProPublica's COMPAS dataset [72], which contains the criminal history, jail and prison time, demographics and COMPAS risk scores for defendants in Broward County from 2013 and 2014. We process the original dataset by dropping records with missing information and quantizing some of the interest variables. Our final dataset contains 5278 records. We choose the private variable (S): $\text{Race} \in \{\text{Caucasian}, \text{African-American}\}$;

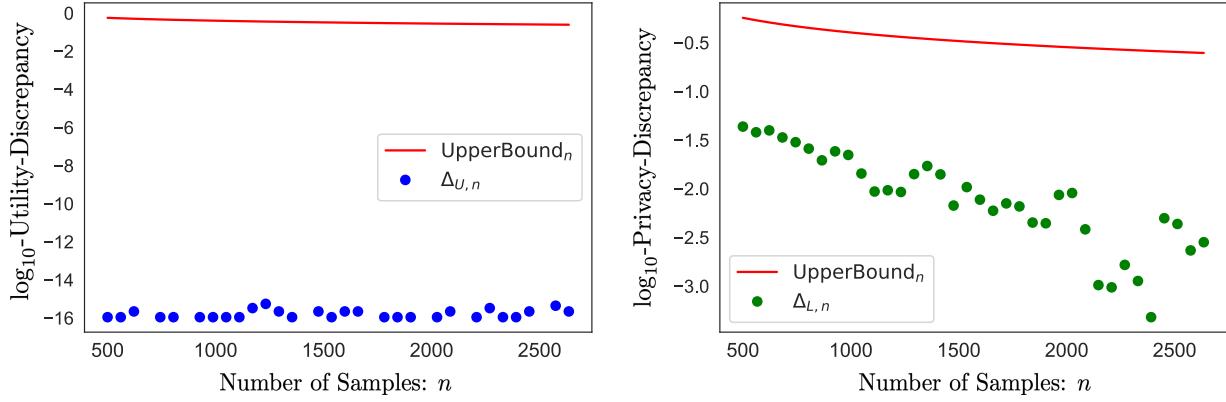


Fig. 3: Both privacy leakage and utility are measured using probability of correctly guessing. Left: discrepancy between utility guarantees for the training and testing sets. Right: discrepancy between privacy guarantees for the training and testing sets. In both pictures the theoretical bound in (136) is shown in red.

the useful variable (X): $PriorCounts \in \{0, 1 - 3, > 3\}$ and $AgeCategory \in \{< 25, 25 - 45, > 45\}$. Hence, the size of support sets are $|\mathcal{S}| = 2$ and $|\mathcal{X}| = 9$.

Given $n \in \mathbb{N}$, we choose n records from the dataset as our training set and another n different records as our testing set. We use probability of correctly guessing to measure both privacy and utility. We compute the empirical distribution $\hat{P}_{n,train}$ based on the training set and let W_n be the privacy mechanism that solves the following optimization problem:

$$\max_{W \in \mathcal{W}_R} \mathcal{U}_c(\hat{P}_{n,train}, W) \quad (131)$$

$$\text{s.t. } \mathcal{L}_c(\hat{P}_{n,train}, W) \leq 0.65, \quad (132)$$

where \mathcal{W}_R denotes the set of *randomized response mechanisms* [16], [81], i.e., the set of all privacy mechanisms W with output alphabet \mathcal{X} such that, for some $\rho \geq 0$,

$$W(i, j) = \begin{cases} \frac{e^\rho}{e^\rho + |\mathcal{X}| - 1} & \text{if } i = j, \\ \frac{1}{e^\rho + |\mathcal{X}| - 1} & \text{if } i \neq j. \end{cases} \quad (133)$$

Then, we let $\hat{P}_{n,test}$ be the empirical distribution of the testing set and compute the privacy-utility guarantees attained by W_n for this distribution. Finally, we evaluate the discrepancy between the privacy-utility guarantees provided for $\hat{P}_{n,train}$ and $\hat{P}_{n,test}$:

$$\Delta_{L,n} \triangleq |\mathcal{L}_c(\hat{P}_{n,test}, W_n) - \mathcal{L}_c(\hat{P}_{n,train}, W_n)|, \quad (134)$$

$$\Delta_{U,n} \triangleq |\mathcal{U}_c(\hat{P}_{n,test}, W_n) - \mathcal{U}_c(\hat{P}_{n,train}, W_n)|. \quad (135)$$

By the triangle inequality and Theorem 2, with probability at least $1 - \beta$, these two discrepancies are upper bounded by

$$\text{UpperBound}_n \triangleq 2\sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}. \quad (136)$$

Figure 3 depicts the discrepancies $\Delta_{L,n}$ and $\Delta_{U,n}$, as functions of n , along with the upper bound UpperBound_n for $\beta = 10^{-1}$. Observe that when number of samples increases, the discrepancy between the privacy-utility guarantees tends to decrease.

To further illustrate our results, we perform a similar experiment using Arimoto's mutual information of order 2.

Specifically, we compute the privacy mechanism W_n by solving the following optimization problem:

$$\max_{W \in \mathcal{W}_Z} \mathcal{U}_2^A(\hat{P}_{n,train}, W) \quad (137)$$

$$\text{s.t. } \mathcal{L}_2^A(\hat{P}_{n,train}, W) \leq 0.05, \quad (138)$$

where \mathcal{W}_Z denotes the set of privacy mechanisms W with output alphabet \mathcal{X} such that, for some $\bar{x} \in \mathcal{X}$ and $\zeta \geq 0$,

$$W(i, j) = \begin{cases} 1 & \text{if } i = j = \bar{x}, \\ 1 - \zeta & \text{if } i = j \neq \bar{x}, \\ \zeta & \text{if } i \neq \bar{x}, j = \bar{x}. \end{cases} \quad (139)$$

In the binary case, \mathcal{W}_Z is nothing but the collection of Z-channels. It has been proved [17] that these channels are capable to achieve optimal privacy-utility trade-offs in some cases. Let $\Delta_{L,n}$ (resp. $\Delta_{U,n}$) be the discrepancy between the privacy (resp. utility) guarantees for $\hat{P}_{n,train}$ and $\hat{P}_{n,test}$. By Theorem 4, with probability at least $1 - \beta$, $\Delta_{L,n}$ and $\Delta_{U,n}$ are upper bounded by

$$\text{Privacy:UpperBound}_n \triangleq 8\sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta) \cdot |\mathcal{S}|}, \quad (140)$$

$$\text{Utility:UpperBound}_n \triangleq 8\sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta) \cdot |\mathcal{X}|}, \quad (141)$$

respectively. Figure 4 depicts the discrepancies $\Delta_{L,n}$ and $\Delta_{U,n}$, as functions of n , along with their corresponding upper bounds for $\beta = 10^{-1}$.

VII. CONCLUDING REMARKS

In this work, we analyzed the effect of a limited sample size on the disclosure of data under privacy constraints. We considered a setting where data is released upon observing non-private features correlated with a set of private features. We evaluated privacy and utility using one out of five of the most commonly used information-theoretic privacy: probability of correctly guessing, f -information, Arimoto's mutual information, Sibson's mutual information, and maximal α -leakage.

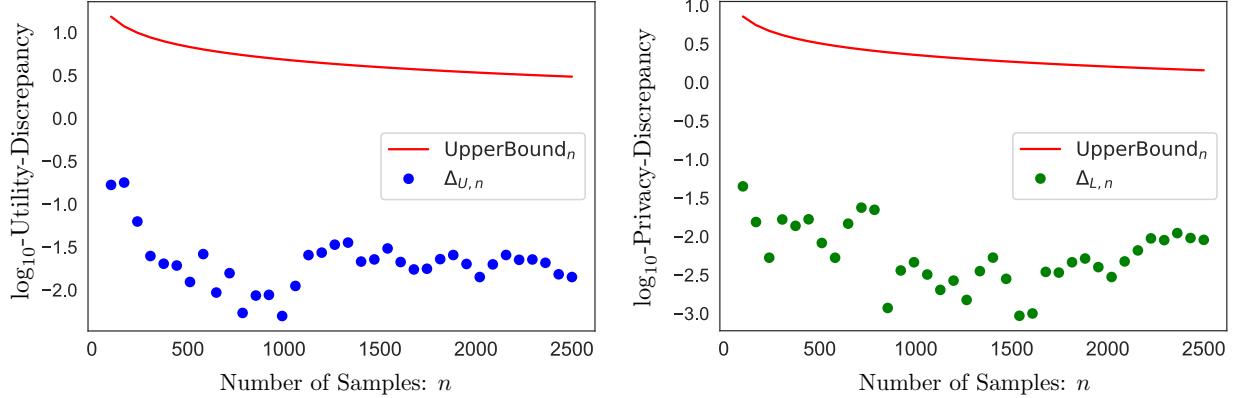


Fig. 4: Both privacy leakage and utility are measured using Arimoto’s mutual information of order 2. Left: discrepancy between utility guarantees for the training and testing sets. Right: discrepancy between privacy guarantees for the training and testing sets. The theoretical upper bounds, in (140) and (141) respectively, are shown in red.

The effect of a limited sample size was assessed via probabilistic upper bounds for the difference between the privacy-utility guarantees for the empirical and true distributions. An important feature of these bounds is that they are completely independent of the privacy mechanism at hand. On a technical level, the proofs of these bounds depend on large deviations results already available in the literature and continuity properties of information leakage measures established in this work. Furthermore, we have established new continuity properties of privacy-utility functions. Using these properties, we have shown that the limit of a convergent sequence of optimal privacy mechanisms is an optimal privacy mechanism itself.

In order to mitigate the effect of a limited sample size on the privacy guarantees delivered to the true distribution, we introduced the notion of *uniform* privacy mechanisms. By definition, these mechanisms provide a specific privacy guarantee for *every* distribution in a given subset of the probability simplex. In particular, when this subset is a neighborhood of the empirical distribution, large deviations results imply that privacy is guaranteed for the true distribution with high probability. While the construction of optimal uniform privacy mechanisms might be challenging, we proved that these mechanisms can be approximated in a natural way by optimal privacy mechanisms in the non-uniform sense. More specifically, we have proved that an optimal privacy mechanism for the empirical distribution delivers a slightly weaker privacy guarantee for a whole neighborhood of the empirical distribution and performs almost as well as any optimal uniform privacy mechanism. By establishing convergence results regarding optimal uniform privacy mechanisms, we have further exhibited the intrinsic relation between optimal privacy mechanisms and their uniform counterparts.

While this work predominantly focused on large deviations bounds, the continuity properties derived in this paper can be applied together with contemporary results to derive similar upper bounds in other estimation frameworks, e.g., ℓ_1 minimax estimation.

ACKNOWLEDGEMENTS

The authors would like to thank Dr. Shahab Asoodeh (Harvard University) for his keen comments regarding the proof of Lemma 3. The authors would also like to thank the anonymous reviewers and the associate editor for their valuable and constructive feedback.

APPENDIX A PROOFS FROM SECTION III

A. Lemma 1

Consider the following elementary lemma whose proof is included for the sake of completeness.

Lemma 9. *Let $Q_1, Q_2 \in \mathcal{P}$ and $W \in \mathcal{W}$ be given. For each $i \in \{1, 2\}$, let $S_i \rightarrow X_i \rightarrow Y_i$ be such that $P_{S_i, X_i} = Q_i$ and $P_{Y_i | X_i} = W$. Then,*

$$\|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1 \leq \|Q_1 - Q_2\|_1. \quad (142)$$

Proof. By the definition of $\|\cdot\|_1$, we have that

$$\|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1 = \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} |P_{S_1, Y_1}(s, y) - P_{S_2, Y_2}(s, y)|. \quad (143)$$

By assumption, $P_{S_i, Y_i}(s, y) = \sum_x Q_i(s, x)W(x, y)$ for each $i \in \{1, 2\}$. Hence, by the triangle inequality,

$$\begin{aligned} & \|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1 \\ & \leq \sum_{y \in \mathcal{Y}} \sum_{s \in \mathcal{S}} \sum_{x \in \mathcal{X}} |Q_1(s, x) - Q_2(s, x)|W(x, y) \end{aligned} \quad (144)$$

$$= \|Q_1 - Q_2\|_1, \quad (145)$$

where we used that $\sum_y W(x, y) = 1$ for all $x \in \mathcal{X}$. \square

The proof of Lemma 1 relies on the following elementary observation: Let $n \in \mathbb{N}$. If $a_i, b_i \in \mathbb{R}$ for all $i = 1, \dots, n$, then

$$\left| \max_{i=1, \dots, n} a_i - \max_{i=1, \dots, n} b_i \right| \leq \max_{i=1, \dots, n} |a_i - b_i| \quad (146)$$

$$\leq \sum_{i=1}^n |a_i - b_i|. \quad (147)$$

Proof. For ease of notation, let

$$\Delta_L \triangleq |\mathcal{L}_c(Q_1, W) - \mathcal{L}_c(Q_2, W)|. \quad (148)$$

For each $i \in \{1, 2\}$, let $S_i \rightarrow X_i \rightarrow Y_i$ be such that $P_{S_i, X_i} = Q_i$ and $P_{Y_i | X_i} = W$. With this notation, $\mathcal{L}_c(Q_i, W) = P_c(S_i | Y_i)$ where

$$P_c(S_i | Y_i) = \sum_{y \in \mathcal{Y}} \max_{s \in \mathcal{S}} P_{S_i, Y_i}(s, y). \quad (149)$$

By the triangle inequality, we have that

$$\Delta_L \leq \sum_{y \in \mathcal{Y}} \left| \max_{s \in \mathcal{S}} P_{S_1, Y_1}(s, y) - \max_{s \in \mathcal{S}} P_{S_2, Y_2}(s, y) \right|. \quad (150)$$

An immediate application of (147) leads to

$$\Delta_L \leq \sum_{y \in \mathcal{Y}} \sum_{s \in \mathcal{S}} |P_{S_1, Y_1}(s, y) - P_{S_2, Y_2}(s, y)| \quad (151)$$

$$= \|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1 \quad (152)$$

$$\leq \|Q_1 - Q_2\|_1, \quad (153)$$

where the last inequality follows from Lemma 9. Mutatis mutandis, it can be shown that

$$|\mathcal{U}_c(Q_1, W) - \mathcal{U}_c(Q_2, W)| \leq \|Q_1 - Q_2\|_1, \quad (154)$$

as required. \square

B. Lemma 2

The proof of the following lemma is similar to the one of Lemma 9. The details are left to the reader.

Lemma 10. *Let $Q_1, Q_2 \in \mathcal{P}$ and $W \in \mathcal{W}$ be given. For each $i \in \{1, 2\}$, let $S_i \rightarrow X_i \rightarrow Y_i$ be such that $P_{S_i, X_i} = Q_i$ and $P_{Y_i | X_i} = W$. Then,*

$$\begin{aligned} \|P_{Y_1} - P_{Y_2}\|_1 &\leq \max\{\|P_{S_1} - P_{S_2}\|_1, \|P_{X_1} - P_{X_2}\|_1\} \\ &\leq \|Q_1 - Q_2\|_1. \end{aligned} \quad (155)$$

Now we are in position to prove Lemma 2.

Proof. For ease of notation, let

$$\Delta_L \triangleq |\mathcal{L}_f(Q_1, W) - \mathcal{L}_f(Q_2, W)|. \quad (156)$$

For each $i \in \{1, 2\}$, let $S_i \rightarrow X_i \rightarrow Y_i$ be such that $P_{S_i, X_i} = Q_i$ and $P_{Y_i | X_i} = W$. With this notation, $\mathcal{L}_f(Q_i, W) = I_f(P_{S_i, Y_i})$. By the definition of f -information, we have that

$$\begin{aligned} \Delta_L &= \left| \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} \left(P_{S_1}(s) P_{Y_1}(y) f\left(\frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)}\right) \right. \right. \\ &\quad \left. \left. - P_{S_2}(s) P_{Y_2}(y) f\left(\frac{P_{S_2, Y_2}(s, y)}{P_{S_2}(s) P_{Y_2}(y)}\right) \right) \right|. \end{aligned} \quad (157)$$

An application of the triangle inequality leads to $\Delta_L \leq I + II$, where

$$\begin{aligned} I &= \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} |P_{S_1}(s) P_{Y_1}(y) - P_{S_2}(s) P_{Y_2}(y)| \\ &\quad \times \left| f\left(\frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)}\right) - f\left(\frac{P_{S_2, Y_2}(s, y)}{P_{S_2}(s) P_{Y_2}(y)}\right) \right|, \end{aligned} \quad (158)$$

$$\begin{aligned} II &= \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} P_{S_2}(s) P_{Y_2}(y) \\ &\quad \times \left| f\left(\frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)}\right) - f\left(\frac{P_{S_2, Y_2}(s, y)}{P_{S_2}(s) P_{Y_2}(y)}\right) \right|. \end{aligned} \quad (159)$$

First we provide an upper bound for I. Recall the definition of the supremum norm in (46). Observe that $P_{S_i, Y_i}(s, y) \leq P_{Y_i}(y)$ for all $i \in \{1, 2\}$, $s \in \mathcal{S}$, and $y \in \mathcal{Y}$. Hence, for all $s \in \mathcal{S}$ and $y \in \mathcal{Y}$,

$$\begin{aligned} &\max \left\{ \frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)}, \frac{P_{S_2, Y_2}(s, y)}{P_{S_2}(s) P_{Y_2}(y)} \right\} \\ &\leq \max \left\{ \frac{1}{P_{S_1}(s)}, \frac{1}{P_{S_2}(s)} \right\} \leq m_S^{-1}, \end{aligned} \quad (160)$$

and thus $\left| f\left(\frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)}\right) \right| \leq K_{f, m_S^{-1}}$, as $|f(t)| \leq K_{f, m_S^{-1}}$ for all $t \in [0, m_S^{-1}]$. As a consequence,

$$I \leq K_{f, m_S^{-1}} \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} |P_{S_1}(s) P_{Y_1}(y) - P_{S_2}(s) P_{Y_2}(y)| \quad (161)$$

$$\begin{aligned} &\leq K_{f, m_S^{-1}} \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} \left[P_{Y_1}(y) |P_{S_1}(s) - P_{S_2}(s)| \right. \\ &\quad \left. + P_{S_2}(s) |P_{Y_1}(y) - P_{Y_2}(y)| \right] \end{aligned} \quad (162)$$

$$= K_{f, m_S^{-1}} (\|P_{S_1} - P_{S_2}\|_1 + \|P_{Y_1} - P_{Y_2}\|_1). \quad (163)$$

By Lemma 10, we conclude that

$$I \leq 2K_{f, m_S^{-1}} \|Q_1 - Q_2\|_1. \quad (164)$$

Now we focus on II. Recall that f is locally Lipschitz by assumption, and thus it is Lipschitz on $[0, m_S^{-1}]$. As defined in (47), let $L_{f, m_S^{-1}}$ be the Lipschitz constant of f on the latter interval. Hence, (160) implies that

$$\begin{aligned} &\left| f\left(\frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)}\right) - f\left(\frac{P_{S_2, Y_2}(s, y)}{P_{S_2}(s) P_{Y_2}(y)}\right) \right| \\ &\leq L_{f, m_S^{-1}} \left| \frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)} - \frac{P_{S_2, Y_2}(s, y)}{P_{S_2}(s) P_{Y_2}(y)} \right|. \end{aligned} \quad (165)$$

As a result, we have that

$$\begin{aligned} II &\leq L_{f, m_S^{-1}} \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} \frac{1}{P_{S_1}(s) P_{Y_1}(y)} \\ &\quad \times |P_{S_2}(s) P_{Y_2}(y) P_{S_1, Y_1}(s, y) \\ &\quad - P_{S_1}(s) P_{Y_1}(y) P_{S_2, Y_2}(s, y)|. \end{aligned} \quad (166)$$

By the triangle inequality, the numerator of the quotient in (166) is upper bounded by

$$\begin{aligned} &P_{S_1, Y_1}(s, y) |P_{S_2}(s) P_{Y_2}(y) - P_{S_1}(s) P_{Y_1}(y)| \\ &+ P_{S_1}(s) P_{Y_1}(y) |P_{S_1, Y_1}(s, y) - P_{S_2, Y_2}(s, y)|. \end{aligned} \quad (167)$$

In particular,

$$\begin{aligned} II &\leq L_{f, m_S^{-1}} \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} \frac{P_{S_1, Y_1}(s, y)}{P_{S_1}(s) P_{Y_1}(y)} \\ &\quad \times |P_{S_2}(s) P_{Y_2}(y) - P_{S_1}(s) P_{Y_1}(y)| \\ &\quad + L_{f, m_S^{-1}} \sum_{s \in \mathcal{S}} \sum_{y \in \mathcal{Y}} |P_{S_1, Y_1}(s, y) - P_{S_2, Y_2}(s, y)| \\ &\leq m_S^{-1} L_{f, m_S^{-1}} (\|P_{Y_1} - P_{Y_2}\|_1 + \|P_{S_1} - P_{S_2}\|_1) \\ &\quad + L_{f, m_S^{-1}} \|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1, \end{aligned}$$

where the last inequality follows from (160) and the argument used in (163). Hence, Lemma 9 and Lemma 10 imply that

$$\text{II} \leq (2m_S^{-1} + 1) L_{f,m_S^{-1}} \|Q_1 - Q_2\|_1. \quad (168)$$

Since $\Delta_L \leq \text{I} + \text{II}$, we conclude that

$$\Delta_L \leq \left(2K_{f,m_S^{-1}} + (2m_S^{-1} + 1) L_{f,m_S^{-1}} \right) \|Q_1 - Q_2\|_1. \quad (169)$$

Mutatis mutandis, it can be shown that

$$\begin{aligned} & |\mathcal{U}_f(Q_1, W) - \mathcal{U}_f(Q_2, W)| \\ & \leq \left(2K_{f,m_X^{-1}} + (2m_X^{-1} + 1) L_{f,m_X^{-1}} \right) \|Q_1 - Q_2\|_1. \end{aligned} \quad (170)$$

The proof is complete. \square

C. Theorem 3

Proof. By choosing $Q_1 = \hat{P}_n$ and $Q_2 = P$ in Lemma 2,

$$|\mathcal{L}_f(\hat{P}_n, W) - \mathcal{L}_f(P, W)| \leq C_{f,m_S} \|\hat{P}_n - P\|_1, \quad (171)$$

where

$$\begin{aligned} m_S = \min \left\{ \left\{ \sum_{x \in \mathcal{X}} \hat{P}_n(s, x) : s \in \mathcal{S} \right\} \right. \\ \left. \cup \left\{ \sum_{x \in \mathcal{X}} P(s, x) : s \in \mathcal{S} \right\} \right\}. \end{aligned} \quad (172)$$

Observe that, for all $s \in \mathcal{S}$,

$$\begin{aligned} \sum_{x \in \mathcal{X}} P(s, x) & \geq \sum_{x \in \mathcal{X}} \hat{P}_n(s, x) - \sum_{x \in \mathcal{X}} |\hat{P}_n(s, x) - P(s, x)| \\ & \geq \sum_{x \in \mathcal{X}} \hat{P}_n(s, x) - \|\hat{P}_n - P\|_1. \end{aligned} \quad (173)$$

In particular, we have that

$$m_S \geq \left(\min \left\{ \sum_{x \in \mathcal{X}} \hat{P}_n(s, x) : s \in \mathcal{S} \right\} - \|\hat{P}_n - P\|_1 \right)^+. \quad (174)$$

Inequality (33) shows that, with probability at least $1 - \beta$,

$$\|\hat{P}_n - P\|_1 \leq \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}, \quad (175)$$

and, thus, $m_S \geq \bar{m}_S$ whenever (175) holds true. Recall the definitions of $K_{g,u}$ and $L_{g,u}$ in (46) and (47), respectively. It is straightforward to verify that the mappings $u \mapsto K_{f,u}$ and $u \mapsto L_{f,u}$ are non-decreasing. Since the mapping $u \mapsto u^{-1}$ is non-increasing, we conclude that

$$u \mapsto C_{f,u} = 2K_{f,u^{-1}} + (2u^{-1} + 1)L_{f,u^{-1}} \quad (176)$$

is non-increasing. Therefore, under (175), we have

$$C_{f,m_S} \leq C_{f,\bar{m}_S}. \quad (177)$$

Combining (171), (175) and (177), we have that, with probability at least $1 - \beta$,

$$|\mathcal{L}_f(\hat{P}_n, W) - \mathcal{L}_f(P, W)| \leq C_{f,\bar{m}_S} \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}. \quad (178)$$

Mutatis mutandis, it can be shown that

$$|\mathcal{U}_f(\hat{P}_n, W) - \mathcal{U}_f(P, W)| \leq C_{f,\bar{m}_S} \sqrt{\frac{2}{n} (|\mathcal{S}| \cdot |\mathcal{X}| - \log \beta)}, \quad (179)$$

whenever (175) holds true. \square

D. Proposition 4

Proof. Assume that \tilde{S} , \tilde{X} and \tilde{X}_0 satisfy that $P_{\tilde{S},\tilde{X}} = \hat{P}_n$ and $\tilde{X}_0 = \Pi_\gamma(\tilde{X})$. With this notation, it can be verified that $\tilde{S} \rightarrow \tilde{X} \rightarrow \tilde{X}_0$ and

$$P_{\tilde{S},\tilde{X}_0} = \hat{P}_n P_{\tilde{X}_0|\tilde{X}} = \hat{P}_\gamma, \quad (180)$$

where, by definition,

$$\left(\hat{P}_n P_{\tilde{X}_0|\tilde{X}} \right) (s, x_0) = \sum_{x \in \mathcal{X}} \hat{P}_n(s, x) P_{\tilde{X}_0|\tilde{X}}(x_0|x). \quad (181)$$

By assumption, $\mathcal{W}^*(\hat{P}_\gamma; \epsilon) \neq \emptyset$. Let $\tilde{W} \in \mathcal{W}^*(\hat{P}_\gamma; \epsilon)$, i.e.,

$$\mathcal{L}_f(\hat{P}_\gamma, \tilde{W}) \leq \epsilon \quad \text{and} \quad \mathcal{U}_f(\hat{P}_\gamma, \tilde{W}) = \mathsf{H}_f(\hat{P}_\gamma; \epsilon). \quad (182)$$

Let \tilde{Y}_0 be such that $\tilde{S} \rightarrow \tilde{X} \rightarrow \tilde{X}_0 \rightarrow \tilde{Y}_0$ and $P_{\tilde{Y}_0|\tilde{X}_0} = \tilde{W}$.

Observe that $P_{\tilde{Y}_0|\tilde{X}} = P_{\tilde{X}_0|\tilde{X}} \tilde{W}$. In particular,

$$P_{\tilde{S},\tilde{Y}_0} = \hat{P}_n P_{\tilde{Y}_0|\tilde{X}} = \hat{P}_n P_{\tilde{X}_0|\tilde{X}} \tilde{W} = \hat{P}_\gamma \tilde{W}, \quad (183)$$

where the last equality follows from (180). Therefore,

$$\mathcal{L}_f(\hat{P}_\gamma, \tilde{W}) = I_f(P_{\tilde{S},\tilde{Y}_0}) = \mathcal{L}_f(\hat{P}_n, P_{\tilde{Y}_0|\tilde{X}}). \quad (184)$$

Furthermore, Lemma 3 implies that

$$\mathcal{U}_f(\hat{P}_\gamma, \tilde{W}) = I_f(P_{\tilde{X}_0,\tilde{Y}_0}) = I_f(P_{\tilde{X},\tilde{Y}_0}) = \mathcal{U}_f(\hat{P}_n, P_{\tilde{Y}_0|\tilde{X}}). \quad (185)$$

By (182), (184) and (185), we have

$$\mathcal{L}_f(\hat{P}_n, P_{\tilde{Y}_0|\tilde{X}}) \leq \epsilon \quad \text{and} \quad \mathcal{U}_f(\hat{P}_n, P_{\tilde{Y}_0|\tilde{X}}) = \mathsf{H}_f(\hat{P}_\gamma; \epsilon).$$

Recall that, by definition, $\mathsf{H}_f(\hat{P}_n; \epsilon) = \sup_W \mathcal{U}_f(\hat{P}_n, W)$ where the supremum is taken over all $W \in \mathcal{W}$ such that $\mathcal{L}_f(\hat{P}_n, W) \leq \epsilon$. Thus,

$$\mathsf{H}_f(\hat{P}_n; \epsilon) \geq \mathcal{U}_f(\hat{P}_n, P_{\tilde{Y}_0|\tilde{X}}) = \mathsf{H}_f(\hat{P}_\gamma; \epsilon), \quad (186)$$

as we wanted to prove. \square

E. Lemma 4

Proof. For ease of notation, let

$$\Delta_L \triangleq |\mathcal{L}_\alpha^A(Q_1, W) - \mathcal{L}_\alpha^A(Q_2, W)|. \quad (187)$$

For each $i \in \{1, 2\}$, let $S_i \rightarrow X_i \rightarrow Y_i$ be such that $P_{S_i,X_i} = Q_i$ and $P_{Y_i|X_i} = W$. With this notation,

$$\Delta_L = |I_\alpha^A(P_{S_1,Y_1}) - I_\alpha^A(P_{S_2,Y_2})| \quad (188)$$

$$= \frac{\alpha}{\alpha - 1} \left| \log \frac{\sum_y \|P_{S_1,Y_1}(\cdot, y)\|_\alpha}{\sum_y \|P_{S_2,Y_2}(\cdot, y)\|_\alpha} - \log \frac{\|P_{S_1}(\cdot)\|_\alpha}{\|P_{S_2}(\cdot)\|_\alpha} \right|, \quad (189)$$

where the second equality follows directly from the definition of Arimoto's mutual information. By the triangle inequality, we obtain that $\frac{\alpha-1}{\alpha}\Delta_L \leq \text{I} + \text{II}$ where

$$\text{I} \triangleq \left| \log \frac{\sum_y \|P_{S_1, Y_1}(\cdot, y)\|_\alpha}{\sum_y \|P_{S_2, Y_2}(\cdot, y)\|_\alpha} \right|, \quad (190)$$

$$\text{II} \triangleq \left| \log \frac{\|P_{S_1}(\cdot)\|_\alpha}{\|P_{S_2}(\cdot)\|_\alpha} \right|. \quad (191)$$

Notice that if $a, b > 0$, then $\left| \log \frac{a}{b} \right| \leq \frac{|a-b|}{\min\{a, b\}}$. Hence,

$$\text{I} \leq \frac{\left| \sum_y \|P_{S_1, Y_1}(\cdot, y)\|_\alpha - \sum_y \|P_{S_2, Y_2}(\cdot, y)\|_\alpha \right|}{\min\{\sum_y \|P_{S_1, Y_1}(\cdot, y)\|_\alpha, \sum_y \|P_{S_2, Y_2}(\cdot, y)\|_\alpha\}}. \quad (192)$$

Observe that, by the triangle inequality,

$$\begin{aligned} & \left| \sum_{y \in \mathcal{Y}} \|P_{S_1, Y_1}(\cdot, y)\|_\alpha - \sum_{y \in \mathcal{Y}} \|P_{S_2, Y_2}(\cdot, y)\|_\alpha \right| \\ & \leq \sum_{y \in \mathcal{Y}} \left| \|P_{S_1, Y_1}(\cdot, y)\|_\alpha - \|P_{S_2, Y_2}(\cdot, y)\|_\alpha \right|. \end{aligned} \quad (193)$$

By Minkowski's inequality, $\|a\|_\alpha - \|b\|_\alpha \leq \|a-b\|_\alpha$ for every $a, b \in \mathbb{R}^n$ and $\alpha \geq 1$. Therefore,

$$\left| \sum_{y \in \mathcal{Y}} \|P_{S_1, Y_1}(\cdot, y)\|_\alpha - \sum_{y \in \mathcal{Y}} \|P_{S_2, Y_2}(\cdot, y)\|_\alpha \right| \quad (194)$$

$$\leq \sum_{y \in \mathcal{Y}} \|P_{S_1, Y_1}(\cdot, y) - P_{S_2, Y_2}(\cdot, y)\|_\alpha \quad (195)$$

$$\leq \sum_{y \in \mathcal{Y}} \|P_{S_1, Y_1}(\cdot, y) - P_{S_2, Y_2}(\cdot, y)\|_1, \quad (196)$$

$$= \|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1, \quad (197)$$

where the second inequality follows from the fact that $\|a\|_\alpha \leq \|a\|_1$ for all $a \in \mathbb{R}^n$ and $\alpha \geq 1$. By assumption, we have that $P_{Y_1|X_1} = W = P_{Y_2|X_2}$. Hence, Lemma 9 implies that

$$\|P_{S_1, Y_1} - P_{S_2, Y_2}\|_1 \leq \|Q_1 - Q_2\|_1. \quad (198)$$

This leads to

$$\left| \sum_{y \in \mathcal{Y}} \|P_{S_1, Y_1}(\cdot, y)\|_\alpha - \sum_{y \in \mathcal{Y}} \|P_{S_2, Y_2}(\cdot, y)\|_\alpha \right| \leq \|Q_1 - Q_2\|_1. \quad (199)$$

For $i \in \{1, 2\}$, Minkowski's inequality implies that

$$\sum_{y \in \mathcal{Y}} \|P_{S_i, Y_i}(\cdot, y)\|_\alpha \geq \left\| \sum_{y \in \mathcal{Y}} P_{S_i, Y_i}(\cdot, y) \right\|_\alpha = \|P_{S_i}(\cdot)\|_\alpha. \quad (200)$$

The generalized mean inequality implies that

$$\|P_{S_i}(\cdot)\|_\alpha \geq |\mathcal{S}|^{1/\alpha-1} \|P_{S_i}(\cdot)\|_1 = |\mathcal{S}|^{1/\alpha-1}, \quad (201)$$

as $\|P_{S_i}(\cdot)\|_1 = \sum_s P_{S_i}(s) = 1$. Hence,

$$\sum_{y \in \mathcal{Y}} \|P_{S_i, Y_i}(\cdot, y)\|_\alpha \geq |\mathcal{S}|^{1/\alpha-1}. \quad (202)$$

Therefore, by plugging (199) and (202) in (192),

$$\text{I} \leq |\mathcal{S}|^{1/\alpha-1} \|Q_1 - Q_2\|_1. \quad (203)$$

Similarly, we have that

$$\text{II} \leq \frac{|\|P_{S_2}(\cdot)\|_\alpha - \|P_{S_1}(\cdot)\|_\alpha|}{\min\{\|P_{S_2}(\cdot)\|_\alpha, \|P_{S_1}(\cdot)\|_\alpha\}}. \quad (204)$$

As in (197) and (198),

$$|\|P_{S_2}(\cdot)\|_\alpha - \|P_{S_1}(\cdot)\|_\alpha| \leq \|P_{S_1} - P_{S_2}\|_1 \leq \|Q_1 - Q_2\|_1. \quad (205)$$

As established in (202), for all $i \in \{1, 2\}$,

$$\|P_{S_i}(\cdot)\|_\alpha \geq |\mathcal{S}|^{1/\alpha-1}. \quad (206)$$

Thus, $\text{II} \leq |\mathcal{S}|^{1-1/\alpha} \|Q_1 - Q_2\|_1$. Since $\frac{\alpha-1}{\alpha} \Delta_L \leq \text{I} + \text{II}$, we conclude that

$$|\mathcal{L}_\alpha^A(Q_1, W) - \mathcal{L}_\alpha^A(Q_2, W)| \leq \frac{2\alpha}{\alpha-1} |\mathcal{S}|^{1-1/\alpha} \|Q_1 - Q_2\|_1. \quad (207)$$

Mutatis mutandis, one can obtain that

$$|\mathcal{U}_\alpha^A(Q_1, W) - \mathcal{U}_\alpha^A(Q_2, W)| \leq \frac{2\alpha}{\alpha-1} |\mathcal{X}|^{1-1/\alpha} \|Q_1 - Q_2\|_1, \quad (208)$$

as required \square

F. Lemma 5

Consider the following lemma.

Lemma 11. *If S_1, S_2 and X_1, X_2 are random variables supported over \mathcal{S} and \mathcal{X} respectively, then,*

$$\|P_{S_1} \cdot P_{X_1|S_1} - P_{S_1} \cdot P_{X_2|S_2}\|_1 \leq 2\|P_{S_1, X_1} - P_{S_2, X_2}\|_1. \quad (209)$$

Proof. For ease of notation, let

$$\Delta \triangleq \|P_{S_1} \cdot P_{X_1|S_1} - P_{S_1} \cdot P_{X_2|S_2}\|_1. \quad (210)$$

By the triangle inequality, we have that

$$\begin{aligned} \Delta & \leq \|P_{S_1} \cdot P_{X_1|S_1} - P_{S_2} \cdot P_{X_2|S_2}\|_1 \\ & \quad + \|P_{S_2} \cdot P_{X_2|S_2} - P_{S_1} \cdot P_{X_2|S_2}\|_1. \end{aligned} \quad (211)$$

By Lemma 10, and the fact that $P_{S_i} \cdot P_{X_i|S_i} = P_{S_i, X_i}$ for every $i \in \{1, 2\}$,

$$\Delta \leq \|P_{S_1, X_1} - P_{S_2, X_2}\|_1 + \|P_{S_1} - P_{S_2}\|_1 \quad (212)$$

$$\leq 2\|P_{S_1, X_1} - P_{S_2, X_2}\|_1, \quad (213)$$

as required \square

Now we are in position to prove Lemma 5.

Proof. For ease of notation, let

$$\Delta_L \triangleq |\mathcal{L}_\alpha^S(Q_1, W) - \mathcal{L}_\alpha^S(Q_2, W)|. \quad (214)$$

For each $i \in \{1, 2\}$, let $S_i \rightarrow X_i \rightarrow Y_i$ be such that $P_{S_i, X_i} = Q_i$ and $P_{Y_i|X_i} = W$. Assume that $\alpha \in (1, \infty)$. Then

$$\Delta_L = |I_\alpha^S(P_{S_1, Y_1}) - I_\alpha^S(P_{S_2, Y_2})| \quad (215)$$

$$= \frac{\alpha}{\alpha-1} \left| \log \frac{\sum_y \|P_{S_1}(\cdot)^{1/\alpha} P_{Y_1|S_1}(y|\cdot)\|_\alpha}{\sum_y \|P_{S_2}(\cdot)^{1/\alpha} P_{Y_2|S_2}(y|\cdot)\|_\alpha} \right|, \quad (216)$$

where the second equality follows directly from the definition of Sibson's mutual information. Notice that if $a, b > 0$, then $|\log \frac{a}{b}| \leq \frac{|a-b|}{\min\{a, b\}}$. Hence,

$$\Delta_L \leq \frac{\alpha}{\alpha-1} \frac{|a-b|}{\min\{a, b\}}, \quad (217)$$

where

$$a \triangleq \sum_y \|P_{S_1}(\cdot)^{1/\alpha} P_{Y_1|S_1}(y|\cdot)\|_\alpha, \quad (218)$$

$$b \triangleq \sum_y \|P_{S_2}(\cdot)^{1/\alpha} P_{Y_2|S_2}(y|\cdot)\|_\alpha. \quad (219)$$

By Minkowski's inequality, we have that

$$\sum_{y \in \mathcal{Y}} \|P_{S_i}(\cdot)^{1/\alpha} P_{Y_i|S_i}(y|\cdot)\|_\alpha \geq \|P_{S_i}(\cdot)^{1/\alpha}\|_\alpha = 1, \quad (220)$$

where we used the fact that $\sum_y P_{Y_i|S_i}(y|s) = 1$ for every $s \in \mathcal{S}$ and $i \in \{1, 2\}$. Hence, (217) becomes

$$\begin{aligned} \Delta_L &\leq \frac{\alpha}{\alpha-1} \sum_{y \in \mathcal{Y}} \left| \|P_{S_1}(\cdot)^{1/\alpha} P_{Y_1|S_1}(y|\cdot)\|_\alpha \right. \\ &\quad \left. - \|P_{S_2}(\cdot)^{1/\alpha} P_{Y_2|S_2}(y|\cdot)\|_\alpha \right| \\ &\leq \frac{\alpha}{\alpha-1} \sum_{y \in \mathcal{Y}} \|P_{S_1}(\cdot)^{1/\alpha} P_{Y_1|S_1}(y|\cdot) \\ &\quad - P_{S_2}(\cdot)^{1/\alpha} P_{Y_2|S_2}(y|\cdot)\|_\alpha, \end{aligned}$$

where the last inequality follows from another application of Minkowski's inequality. Since $\|a\|_\alpha \leq \|a\|_1$ for all $a \in \mathbb{R}^n$ and $\alpha > 1$, we obtain that

$$\begin{aligned} \Delta_L &\leq \frac{\alpha}{\alpha-1} \sum_{y \in \mathcal{Y}} \|P_{S_1}(\cdot)^{1/\alpha} P_{Y_1|S_1}(y|\cdot) \\ &\quad - P_{S_2}(\cdot)^{1/\alpha} P_{Y_2|S_2}(y|\cdot)\|_1. \end{aligned}$$

Observe that, for each $i \in \{1, 2\}$,

$$P_{S_i}(s)^{1/\alpha} P_{Y_i|S_i}(y|s) = \sum_{x \in \mathcal{X}} P_{S_i}(s)^{1/\alpha} P_{X_i|S_i}(x|s) W(x, y).$$

Thus, a straightforward manipulation leads to

$$\begin{aligned} \Delta_L &\leq \frac{\alpha}{\alpha-1} \sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} \|P_{S_1}(\cdot)^{1/\alpha} P_{X_1|S_1}(x|\cdot) \\ &\quad - P_{S_2}(\cdot)^{1/\alpha} P_{X_2|S_2}(x|\cdot)\|_1 W(x, y) \\ &= \frac{\alpha}{\alpha-1} \sum_{x \in \mathcal{X}} \|P_{S_1}(\cdot)^{1/\alpha} P_{X_1|S_1}(x|\cdot) \\ &\quad - P_{S_2}(\cdot)^{1/\alpha} P_{X_2|S_2}(x|\cdot)\|_1. \quad (221) \end{aligned}$$

By adding and subtracting the term $P_{S_1}(\cdot)^{1/\alpha} P_{X_2|S_2}(x|\cdot)$ inside the norm in (221), Minkowski's inequality implies that $\frac{\alpha-1}{\alpha} \Delta_L \leq \sum_x I_x + \sum_x II_x$ where, for each $x \in \mathcal{X}$,

$$I_x \triangleq \|P_{S_1}(\cdot)^{1/\alpha} (P_{X_1|S_1}(x|\cdot) - P_{X_2|S_2}(x|\cdot))\|_1, \quad (222)$$

$$II_x \triangleq \|(P_{S_1}(\cdot)^{1/\alpha} - P_{S_2}(\cdot)^{1/\alpha}) P_{X_2|S_2}(x|\cdot)\|_1. \quad (223)$$

Observe that, for each $x \in \mathcal{X}$,

$$I_x = \sum_{s \in \mathcal{S}} P_{S_1}(s)^{1/\alpha} |P_{X_1|S_1}(x|s) - P_{X_2|S_2}(x|s)| \quad (224)$$

$$= \sum_{s \in \mathcal{S}} \frac{P_{S_1}(s)}{P_{S_1}(s)^{1-1/\alpha}} |P_{X_1|S_1}(x|s) - P_{X_2|S_2}(x|s)|. \quad (225)$$

Since $P_{S_1}(s) \geq m_S$ for all $s \in \mathcal{S}$, we obtain that

$$\sum_{x \in \mathcal{X}} I_x \leq \sum_{x \in \mathcal{X}} \sum_{s \in \mathcal{S}} \frac{P_{S_1}(s) |P_{X_1|S_1}(x|s) - P_{X_2|S_2}(x|s)|}{m_S^{1-1/\alpha}} \quad (226)$$

$$= \frac{\|P_{S_1} \cdot P_{X_1|S_1} - P_{S_1} \cdot P_{X_2|S_2}\|_1}{m_S^{1-1/\alpha}}. \quad (227)$$

Therefore, Lemma 11 leads to

$$\sum_{x \in \mathcal{X}} I_x \leq \frac{2\|Q_1 - Q_2\|_1}{m_S^{1-1/\alpha}}. \quad (228)$$

By definition, $II_x = \sum_s |P_{S_1}(s)^{1/\alpha} - P_{S_2}(s)^{1/\alpha}| P_{X_2|S_2}(x|s)$ for every $x \in \mathcal{X}$. Hence,

$$\sum_{x \in \mathcal{X}} II_x = \sum_{s \in \mathcal{S}} |P_{S_1}(\cdot)^{1/\alpha} - P_{S_2}(\cdot)^{1/\alpha}|. \quad (229)$$

Since the function $t \mapsto t^{1/\alpha}$ is Lipschitz continuous on $[m_S, 1]$ with Lipschitz constant $(\alpha m_S^{1-1/\alpha})^{-1}$,

$$\sum_{x \in \mathcal{X}} II_x \leq \frac{\|P_{S_1} - P_{S_2}\|_1}{\alpha m_S^{1-1/\alpha}} \leq \frac{\|Q_1 - Q_2\|_1}{\alpha m_S^{1-1/\alpha}}, \quad (230)$$

where the last inequality follows from Lemma 10. By (228) and (230), we conclude that

$$\Delta_L \leq \frac{2\alpha+1}{\alpha-1} \frac{\|Q_1 - Q_2\|_1}{m_S^{1-1/\alpha}}. \quad (231)$$

Mutatis mutandis, it can be shown that

$$|\mathcal{U}_\alpha^S(Q_1, W) - \mathcal{U}_\alpha^S(Q_2, W)| \leq \frac{1}{\alpha-1} \frac{\|Q_1 - Q_2\|_1}{m_X^{1-1/\alpha}}. \quad (232)$$

Recall that $\lim_{\alpha \rightarrow \infty} I_\alpha^S(P_{X,Y}) = I_\infty^S(P_{X,Y})$ [31]. Therefore, (74) follows after taking the limit $\alpha \rightarrow \infty$ in both sides of (232). Using a similar argument as above, it can be shown that

$$\Delta_L \leq \sum_{s \in \mathcal{S}} \sum_{x \in \mathcal{X}} |P_{X_1|S_1}(x|s) - P_{X_2|S_2}(x|s)| \quad (233)$$

$$\leq \sum_{s \in \mathcal{S}} \sum_{x \in \mathcal{X}} \frac{P_{S_1}(s) |P_{X_1|S_1}(x|s) - P_{X_2|S_2}(x|s)|}{\min_{s'} P_{S_1}(s')} \quad (234)$$

$$= \frac{\|P_{S_1} \cdot P_{X_1|S_1} - P_{S_1} \cdot P_{X_2|S_2}\|_1}{\min_{s'} P_{S_1}(s')}. \quad (235)$$

By applying Lemma 11 and noting that

$$\min_{s' \in \mathcal{S}} P_{S_1}(s') = \min_{s' \in \mathcal{S}} \sum_{x \in \mathcal{X}} Q_1(s', x), \quad (236)$$

the result follows. \square

G. Lemma 6

Proof. For ease of notation, let

$$\Delta_L \triangleq |\mathcal{L}_\alpha^{\max}(Q_1, W) - \mathcal{L}_\alpha^{\max}(Q_2, W)|. \quad (237)$$

For each $i \in \{1, 2\}$, let $S_i \rightarrow X_i \rightarrow Y_i$ be such that $P_{S_i, X_i} = Q_i$ and $P_{Y_i|X_i} = W$. Observe that $P_{Y_i|S_i} = P_{X_i|S_i}W$ where, by definition,

$$(P_{X_i|S_i}W)(y|s) = \sum_{x \in \mathcal{X}} P_{X_i|S_i}(x|s)W(x, y). \quad (238)$$

In particular,

$$\mathcal{L}_\alpha^{\max}(Q_i, W) \triangleq \sup_{P_{\tilde{S}}} I_\alpha^A(P_{\tilde{S}} \cdot P_{Y_i|S_i}) \quad (239)$$

$$= \sup_{P_{\tilde{S}}} I_\alpha^A(P_{\tilde{S}} \cdot P_{X_i|S_i}W). \quad (240)$$

By the definition of $\mathcal{L}_\alpha^A(Q, W)$, we have that

$$\mathcal{L}_\alpha^{\max}(Q_i, W) = \sup_{P_{\tilde{S}}} \mathcal{L}_\alpha^A(P_{\tilde{S}} \cdot P_{X_i|S_i}, W). \quad (241)$$

It is easy to verify that if \mathcal{I} is an arbitrary index set and $a_\iota, b_\iota \in \mathbb{R}$ for all $\iota \in \mathcal{I}$, then

$$\left| \sup_{\iota \in \mathcal{I}} a_\iota - \sup_{\iota \in \mathcal{I}} b_\iota \right| \leq \sup_{\iota \in \mathcal{I}} |a_\iota - b_\iota|. \quad (242)$$

Therefore, (241) implies that

$$\begin{aligned} \Delta_L &= \left| \sup_{P_{\tilde{S}}} \mathcal{L}_\alpha^A(P_{\tilde{S}} \cdot P_{X_1|S_1}, W) - \sup_{P_{\tilde{S}}} \mathcal{L}_\alpha^A(P_{\tilde{S}} \cdot P_{X_2|S_2}, W) \right| \\ &\leq \sup_{P_{\tilde{S}}} |\mathcal{L}_\alpha^A(P_{\tilde{S}} \cdot P_{X_1|S_1}, W) - \mathcal{L}_\alpha^A(P_{\tilde{S}} \cdot P_{X_2|S_2}, W)|. \end{aligned}$$

By Lemma 4, we obtain that

$$\Delta_L \leq \frac{2\alpha}{\alpha-1} |\mathcal{S}|^{1-1/\alpha} \sup_{P_{\tilde{S}}} \|P_{\tilde{S}} \cdot P_{X_1|S_1} - P_{\tilde{S}} \cdot P_{X_2|S_2}\|_1. \quad (243)$$

A straightforward computation shows that, for any $P_{\tilde{S}}$,

$$\begin{aligned} &\|P_{\tilde{S}} \cdot P_{X_1|S_1} - P_{\tilde{S}} \cdot P_{X_2|S_2}\|_1 \\ &= \sum_{s \in \mathcal{S}} P_{\tilde{S}}(s) \|P_{X_1|S_1}(\cdot|s) - P_{X_2|S_2}(\cdot|s)\|_1 \\ &\leq \sum_{s \in \mathcal{S}} \|P_{X_1|S_1}(\cdot|s) - P_{X_2|S_2}(\cdot|s)\|_1 \\ &\leq \frac{1}{\min_{s'} P_{S_1}(s')} \sum_{s \in \mathcal{S}} P_{S_1}(s) \|P_{X_1|S_1}(\cdot|s) - P_{X_2|S_2}(\cdot|s)\|_1 \\ &= \frac{\|P_{S_1} \cdot P_{X_1|S_1} - P_{S_1} \cdot P_{X_2|S_2}\|_1}{\min_{s'} P_{S_1}(s')}. \end{aligned}$$

Therefore, Lemma 11 implies that

$$\sup_{P_{\tilde{S}}} \|P_{\tilde{S}} \cdot P_{X_1|S_1} - P_{\tilde{S}} \cdot P_{X_2|S_2}\|_1 \leq \frac{2\|Q_1 - Q_2\|_1}{\min_{s'} P_{S_1}(s')}. \quad (244)$$

By (243) and (244), we conclude that

$$\Delta_L \leq \frac{4\alpha}{\alpha-1} \frac{|\mathcal{S}|^{1-1/\alpha}}{\min_{s'} P_{S_1}(s')} \|Q_1 - Q_2\|_1. \quad (245)$$

By noting that $\min_{s'} P_{S_1}(s') = \min_{s'} \sum_x Q_1(s', x)$, (79) follows. Finally, note that $\mathcal{U}_\alpha^{\max}(Q, W)$ only depends on W . Therefore, $\mathcal{U}_\alpha^{\max}(Q_1, W) = \mathcal{U}_\alpha^{\max}(Q_2, W)$ which trivially leads to (80). \square

APPENDIX B PROOFS FROM SECTION IV

A. Lemma 7

Lemma 7 is an immediate consequence of the following standard result whose proof is included for the sake of completeness.

Lemma 12. *Let $(\mathcal{A}, d_{\mathcal{A}})$ and $(\mathcal{B}, d_{\mathcal{B}})$ be two metric spaces. If $f : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ satisfies that*

(i) $|f(a_1, b) - f(a_2, b)| \leq d_{\mathcal{A}}(a_1, a_2)$ for all $a_1, a_2 \in \mathcal{A}$ and $b \in \mathcal{B}$;

(ii) $f(a, \cdot)$ is continuous over \mathcal{B} for all $a \in \mathcal{A}$, then $f(\cdot, \cdot)$ is continuous over $\mathcal{A} \times \mathcal{B}$.

Proof. In order to prove the continuity of $f(\cdot, \cdot)$ over $\mathcal{A} \times \mathcal{B}$, we will show that for any sequence $\{(a_n, b_n)\}_{n=0}^\infty$ such that $\lim_n (a_n, b_n) = (a_0, b_0)$, it holds true that $\lim_n f(a_n, b_n) = f(a_0, b_0)$.

Let $\{(a_n, b_n)\}_{n=0}^\infty$ be such that $\lim_n (a_n, b_n) = (a_0, b_0)$. By the triangle inequality, for every $n \in \mathbb{N}$,

$$|f(a_n, b_n) - f(a_0, b_0)| \quad (246)$$

$$\leq |f(a_n, b_n) - f(a_0, b_n)| + |f(a_0, b_n) - f(a_0, b_0)| \quad (247)$$

$$\leq d_{\mathcal{A}}(a_n, a_0) + |f(a_0, b_n) - f(a_0, b_0)|, \quad (248)$$

where the last inequality follows from assumption (i). By assumption (ii), $f(a_0, \cdot)$ is continuous over \mathcal{B} . Since $\lim_n (a_n, b_n) = (a_0, b_0)$ is equivalent to $\lim_n a_n = a_0$ and $\lim_n b_n = b_0$, we conclude that

$$\lim_{n \rightarrow \infty} |f(a_n, b_n) - f(a_0, b_0)| = 0, \quad (249)$$

as required. \square

B. Proposition 5

We start recalling an elementary fact about convergent sequences in metric spaces, see, e.g., [82, Exercise 2.4.11].

Theorem 10. *Let \mathcal{M} be a metric space. A sequence $(a_n)_{n=1}^\infty \subset \mathcal{M}$ converges to $a \in \mathcal{M}$ if and only if every subsequence of $(a_n)_{n=1}^\infty$ has a further subsequence which converges to a .*

We proceed with the proof of Proposition 5.

Proof. In order to prove the continuity of $\mathsf{H}(\cdot; \epsilon)$ over the set $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$, we will show that for any sequence $(Q_n)_{n=0}^\infty \subset \{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$ such that

$$\lim_{n \rightarrow \infty} \|Q_n - Q_0\|_1 = 0, \quad (250)$$

it holds true that

$$\lim_{n \rightarrow \infty} \mathsf{H}(Q_n; \epsilon) = \mathsf{H}(Q_0; \epsilon). \quad (251)$$

In order to prove (251), Theorem 10 implies that it is enough to show that any subsequence $(\mathsf{H}(Q_{n_k}; \epsilon))_{k=1}^\infty$ has a further subsequence converging to $\mathsf{H}(Q_0; \epsilon)$.

Let $(n_k)_{k=1}^\infty \subset \mathbb{N}$ be given. By Remark 7, there exists $(W_{n_k}^*)_{k=1}^\infty \subset \mathcal{W}_N$ such that, for each $k \geq 1$, we have that $\mathcal{L}(Q_{n_k}, W_{n_k}^*) \leq \epsilon$ and

$$\mathsf{H}(Q_{n_k}; \epsilon) = \mathcal{U}(Q_{n_k}, W_{n_k}^*). \quad (252)$$

Since the space \mathcal{W}_N is compact, there exists a further subsequence $(n_{k_j})_{j=1}^\infty$ such that

$$\lim_{j \rightarrow \infty} W_{n_{k_j}}^* = W_0, \quad (253)$$

for some $W_0 \in \mathcal{W}_N$. Lemma 7 shows that $\mathcal{U}(\cdot, \cdot)$ is continuous over $\mathcal{Q} \times \mathcal{W}_N$. Then (250) and (253) imply that

$$\mathcal{U}(Q_0, W_0) = \lim_{j \rightarrow \infty} \mathcal{U}(Q_{n_{k_j}}, W_{n_{k_j}}^*) = \lim_{j \rightarrow \infty} \mathcal{H}(Q_{n_{k_j}}; \epsilon),$$

where the second equality follows from (252). A similar reasoning shows that

$$\mathcal{L}(Q_0, W_0) = \lim_{j \rightarrow \infty} \mathcal{L}(Q_{n_{k_j}}, W_{n_{k_j}}^*) \leq \epsilon. \quad (254)$$

Therefore, by the maximality of $\mathcal{H}(Q_0; \epsilon)$,

$$\lim_{j \rightarrow \infty} \mathcal{H}(Q_{n_{k_j}}; \epsilon) = \mathcal{U}(Q_0, W_0) \leq \mathcal{H}(Q_0; \epsilon). \quad (255)$$

Next, we show that $\lim_j \mathcal{H}(Q_{n_{k_j}}; \epsilon) \geq \mathcal{H}(Q_0; \epsilon)$. Recall that, by assumption, $\epsilon_{\min}(Q_0) < \epsilon$. Let $\delta > 0$ be such that $\epsilon_{\min}(Q_0) \leq \epsilon - \delta$. By condition (C.3), there exists $W'_0 \in \mathcal{W}_N$ such that $\mathcal{L}(Q_0, W'_0) \leq \epsilon - \delta$ and $\mathcal{U}(Q_0, W'_0) = \mathcal{H}(Q_0; \epsilon - \delta)$. By the Lipschitz continuity given in condition (C.2) and (250),

$$\lim_{j \rightarrow \infty} \mathcal{L}(Q_{n_{k_j}}, W'_0) = \mathcal{L}(Q_0, W'_0) \leq \epsilon - \delta, \quad (256)$$

$$\lim_{j \rightarrow \infty} \mathcal{U}(Q_{n_{k_j}}, W'_0) = \mathcal{U}(Q_0, W'_0) = \mathcal{H}(Q_0; \epsilon - \delta). \quad (257)$$

Hence, for j large enough we have that $\mathcal{L}(Q_{n_{k_j}}, W'_0) \leq \epsilon$ and, by the maximality of $\mathcal{H}(Q_{n_{k_j}}; \epsilon)$,

$$\begin{aligned} \lim_{j \rightarrow \infty} \mathcal{H}(Q_{n_{k_j}}; \epsilon) &\geq \lim_{j \rightarrow \infty} \mathcal{U}(Q_{n_{k_j}}, W'_0) \\ &= \mathcal{H}(Q_0; \epsilon - \delta), \end{aligned} \quad (258)$$

where the last equality follows from (257). Recall that, by condition (C.1), the mapping $\mathcal{H}(Q_0; \cdot)$ is continuous over $[\epsilon_{\min}(Q_0), \infty)$. Hence, by taking limits in (258),

$$\lim_{j \rightarrow \infty} \mathcal{H}(Q_{n_{k_j}}; \epsilon) \geq \lim_{\delta \downarrow 0} \mathcal{H}(Q_0; \epsilon - \delta) = \mathcal{H}(Q_0; \epsilon). \quad (259)$$

By combining (255) and (259), we conclude that

$$\lim_{j \rightarrow \infty} \mathcal{H}(Q_{n_{k_j}}; \epsilon) = \mathcal{H}(Q_0; \epsilon), \quad (260)$$

as required. \square

C. Corollary 1

Before we prove Corollary 1, let us recall two standard results from analysis on metric spaces, see, e.g., Chapter 0.1 in [83] and Lemma 14 in [84], respectively.

Lemma 13. *Let $a < b$ be given. For each $n \in \mathbb{N}$, assume that $f_n : [a, b] \rightarrow \mathbb{R}$ is a non-decreasing function. If there exists a continuous function $f : [a, b] \rightarrow \mathbb{R}$ such that, for all $x \in [a, b]$, $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$, then f_n converges uniformly to f .*

Lemma 14. *Let \mathcal{A} and \mathcal{B} be two metric spaces with \mathcal{B} compact. If $f : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ is continuous, then the function $g : \mathcal{A} \rightarrow \mathbb{R}$ defined by $g(a) \triangleq \inf_{b \in \mathcal{B}} f(a, b)$ is also continuous.*

Observe that the infimum in the previous lemma can be replaced by a supremum. We proceed with the proof of Corollary 1.

Proof. We first introduce

$$\epsilon_{\max}(Q) \triangleq \max\{\mathcal{L}(Q, W) : W \in \mathcal{W}_N\}. \quad (261)$$

Note that, by the maximality of $\epsilon_{\max}(Q)$, for all $\epsilon \geq \epsilon_{\max}(Q)$,

$$\mathcal{H}(Q; \epsilon) = \mathcal{H}(Q; \epsilon_{\max}(Q)). \quad (262)$$

By Lemma 7 the function $\mathcal{L}(\cdot, \cdot)$ is continuous over $\mathcal{Q} \times \mathcal{W}_N$. Since \mathcal{W}_N is compact, Lemma 14 implies that $\epsilon_{\max}(\cdot)$ is continuous over \mathcal{Q} . By hypothesis $\lim_n P_n = P$, thus $\lim_n \epsilon_{\max}(P_n) = \epsilon_{\max}(P)$ and, in particular, there exists $\epsilon_1 > \epsilon_0$ such that $\epsilon_{\max}(P_n) \leq \epsilon_1$ for all $n \in \mathbb{N}$. By (262), for all $\epsilon \geq \epsilon_1$ and all $n \in \mathbb{N}$,

$$\mathcal{H}(P_n; \epsilon) = \mathcal{H}(P_n; \epsilon_{\max}(P_n)) = \mathcal{H}(P_n; \epsilon_1). \quad (263)$$

Observe that an analogous equality holds for P . Hence, we obtain that, for all $n \in \mathbb{N}$,

$$\sup_{\epsilon \geq \epsilon_1} |\mathcal{H}(P_n; \epsilon) - \mathcal{H}(P; \epsilon)| = |\mathcal{H}(P_n; \epsilon_1) - \mathcal{H}(P; \epsilon_1)|. \quad (264)$$

As established by Proposition 5, the right hand side of (264) vanishes as n goes to infinity. Hence, we conclude that $\mathcal{H}(P_n; \cdot)$ converge uniformly to $\mathcal{H}(P; \cdot)$ over $[\epsilon_1, \infty)$. Finally, recall that the function $\mathcal{H}(P_n; \cdot)$ is non-decreasing over $[\epsilon_0, \epsilon_1]$ for each $n \in \mathbb{N}$. Furthermore, by condition (C.1), the function $\mathcal{H}(P; \cdot)$ is continuous over $[\epsilon_0, \epsilon_1]$. Therefore, Proposition 5 and Lemma 13 imply that $\mathcal{H}(P_n; \cdot)$ converges uniformly to $\mathcal{H}(P; \cdot)$ over $[\epsilon_0, \epsilon_1]$. This shows that uniform convergence holds over $[\epsilon_0, \infty)$. \square

D. Theorem 6

The following lemma follows directly from Lemma 14 by noticing that $\mathcal{L}(\cdot, \cdot)$ is continuous over $\mathcal{Q} \times \mathcal{W}_N$ and \mathcal{W}_N is compact.

Lemma 15. *Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. Then $\epsilon_{\min}(\cdot)$, as given in (88), is continuous over \mathcal{Q} .*

We prove the following lemma which will be used in the proof of Theorem 6.

Lemma 16. *Assume that conditions (C.1) and (C.3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. Then the set $\mathcal{W}_N^*(P; \epsilon)$, defined in (87), is compact.*

Proof. Observe that

$$\begin{aligned} \mathcal{W}_N^*(P; \epsilon) &= \{W \in \mathcal{W}_N : \mathcal{L}(P, W) \leq \epsilon\} \\ &\cap \{W \in \mathcal{W}_N : \mathcal{U}(P, W) = \mathcal{H}(P; \epsilon)\}. \end{aligned} \quad (265)$$

Recall that, by condition (C.1), the mappings $\mathcal{L}(P, \cdot)$ and $\mathcal{U}(P, \cdot)$ are continuous. Since both sets in the RHS of (265) are the preimage of a closed set under a continuous function, they are closed. Therefore, $\mathcal{W}_N^*(P; \epsilon)$ is closed as the finite intersection of closed sets is closed. Finally, we conclude that $\mathcal{W}_N^*(P; \epsilon)$ is compact since it is a closed subset of the compact set \mathcal{W}_N . \square

Observe that, by Lemma 15, $\lim_n \epsilon_{\min}(P_n) = \epsilon_{\min}(P)$ whenever $\lim_n P_n = P$. Hence, if $\epsilon_{\min}(P) < \epsilon$, then $\epsilon_{\min}(P_n) < \epsilon$ for n large enough. Therefore, under the

assumptions of Theorem 6, for n large enough we have that the set $\mathcal{W}_N^*(P_n; \epsilon)$ is non-empty and, by condition (C.3), $\mathcal{W}_N^*(P_n; \epsilon)$ is non-empty as well. Now we are in position to prove Theorem 6.

Proof. We assume, without loss of generality, that $\mathcal{W}_N^*(P_n; \epsilon)$ is non-empty for all $n \in \mathbb{N}$. In order to reach contradiction, assume that (93) does not hold true, i.e., there exists $W_n^* \in \mathcal{W}_N^*(P_n; \epsilon)$ such that

$$\limsup_{n \rightarrow \infty} \text{dist}(W_n, \mathcal{W}_N^*(P; \epsilon)) > 0. \quad (266)$$

Since \mathcal{W}_N and $\mathcal{W}_N^*(P; \epsilon)$ are compact from Lemma 16, a routine argument shows that there exist a subsequence $(W_{n_k}^*)_{k=1}^\infty$ and $W_0 \in \mathcal{W}_N$ such that $\lim_k W_{n_k}^* = W_0$ and $W_0 \notin \mathcal{W}_N^*(P; \epsilon)$. Lemma 7 shows that $\mathcal{L}(\cdot, \cdot)$ is continuous over $\mathcal{Q} \times \mathcal{W}_N$. Therefore, we have

$$\mathcal{L}(P, W_0) = \lim_{k \rightarrow \infty} \mathcal{L}(P_{n_k}, W_{n_k}^*) \leq \epsilon. \quad (267)$$

Hence, by the maximality of the privacy-utility function, $\mathcal{U}(P, W_0) \leq \mathcal{H}(P; \epsilon)$. Using a similar argument, one can show that

$$\mathcal{U}(P, W_0) = \lim_{k \rightarrow \infty} \mathcal{U}(P_{n_k}, W_{n_k}^*) = \lim_{k \rightarrow \infty} \mathcal{H}(P_{n_k}, \epsilon) = \mathcal{H}(P; \epsilon),$$

where the last equality follows from the fact that the mapping $\mathcal{H}(\cdot; \epsilon)$ is continuous, as established in Proposition 5. Altogether,

$$\mathcal{L}(P, W_0) \leq \epsilon \quad \text{and} \quad \mathcal{U}(P, W_0) = \mathcal{H}(P; \epsilon). \quad (268)$$

Therefore, $W_0 \in \mathcal{W}_N^*(P; \epsilon)$ which contradicts the fact that $W_0 \notin \mathcal{W}_N^*(P; \epsilon)$.

Furthermore, by the strong law of large numbers, $(\hat{P}_n)_{n=1}^\infty$ converges almost surely to P as n goes to infinity. Therefore, if $P_n = \hat{P}_n$, then

$$\Pr \left(\lim_{n \rightarrow \infty} \text{dist}(W_n^*, \mathcal{W}_N^*(P; \epsilon)) = 0 \right) = 1, \quad (269)$$

as claimed. \square

Note that in the previous proof we implicitly assume that, for each $n \in \mathbb{N}$, the (possibly random) privacy mechanism $W_n^* \in \mathcal{W}_N^*(\hat{P}_n; \epsilon)$ is a random variable, i.e., a measurable function. This is a minor subtlety given the rareness of non-measurable sets (functions) [85].

E. Theorem 7

Before we prove Theorem 7, let us recall some definitions and theorems from set-valued analysis [75].

Given two metric spaces \mathcal{A} and \mathcal{B} , a set-valued mapping F from \mathcal{A} to \mathcal{B} , denoted by $F : \mathcal{A} \rightsquigarrow \mathcal{B}$, is a function from \mathcal{A} to the subsets of \mathcal{B} . The domain of a set-valued mapping $F : \mathcal{A} \rightsquigarrow \mathcal{B}$ is defined as

$$\text{Dom}(F) \triangleq \{a \in \mathcal{A} : F(a) \neq \emptyset\}. \quad (270)$$

For $r > 0$, $\mathcal{A}_r(a) \triangleq \{a' \in \mathcal{A} : d_{\mathcal{A}}(a, a') < r\}$ where $d_{\mathcal{A}}$ denotes the metric associated to \mathcal{A} .

Definition 10 (Definition 1.4.1, [75]). A set-valued mapping $F : \mathcal{A} \rightsquigarrow \mathcal{B}$ is called upper semicontinuous at $a \in \text{Dom}(F)$

if and only if for every neighborhood \mathcal{N} of $F(a)$ there exists $r > 0$ such that $F(a') \subset \mathcal{N}$ for all $a' \in \mathcal{A}_r(a)$. A set-valued mapping is said to be upper semicontinuous if and only if it is upper semicontinuous at every point of its domain.

Definition 11 (Definition 1.4.2, [75]). A set-valued mapping $F : \mathcal{A} \rightsquigarrow \mathcal{B}$ is called lower semicontinuous at $a \in \text{Dom}(F)$ if and only if for every $b \in F(a)$ and every sequence $(a_n)_{n=1}^\infty \subset \text{Dom}(F)$ with $\lim_n a_n = a$, there exists a sequence $(b_n)_{n=1}^\infty$ such that $b_n \in F(a_n)$ for each $n \in \mathbb{N}$ and $\lim_n b_n = b$. A set-valued mapping is said to be lower semicontinuous if and only if it is lower semicontinuous at every point of its domain.

By definition, a the set-valued mapping $F : \mathcal{A} \rightsquigarrow \mathcal{B}$ is continuous at $a \in \mathcal{A}$ if and only if it is both upper and lower semicontinuous at a .

Theorem 11 (Theorem 1.4.13, [75]). *Let F be a set-valued map from a complete metric space \mathcal{A} to a complete separable metric space \mathcal{B} . If F is upper semicontinuous, then it is continuous on a residual set of \mathcal{A} .*

The proof of Theorem 7 relies on the following lemma.

Lemma 17. *Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. For any $\epsilon \in \mathbb{R}$, the set-valued mapping $\mathcal{W}_N^*(\cdot; \epsilon)$ is upper semicontinuous over $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$.*

Proof. For ease of notation, let $F : \mathcal{Q} \rightsquigarrow \mathcal{W}_N$ be the set-valued mapping defined by $F(Q) \triangleq \mathcal{W}_N^*(Q; \epsilon)$. Observe that $\text{Dom}(F) = \{Q \in \mathcal{Q} : \epsilon_{\min}(Q) \leq \epsilon\}$. In order to reach contradiction, assume that the set-valued mapping F is not upper semicontinuous over $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$. In this case, there exist a joint distribution $Q \in \mathcal{Q}$ with $\epsilon_{\min}(Q) < \epsilon$ and a neighborhood \mathcal{N} of $F(Q)$ such that for every $r > 0$ there exists $Q_r \in \mathcal{Q}_r(Q)$ with $F(Q_r) \not\subset \mathcal{N}$. For each $n \in \mathbb{N}$, let W_n^* be such that $W_n^* \in F(Q_{1/n}) \cap \mathcal{N}^c$, where \mathcal{N}^c denotes the complement of \mathcal{N} w.r.t. \mathcal{W}_N . Since \mathcal{N}^c is closed and \mathcal{W}_N is compact, we have that \mathcal{N}^c is also compact. Hence, there exists a subsequence $(W_{n_k}^*)_{k=1}^\infty$ converging to some $W \in \mathcal{N}^c$. Since $F(Q) \subset \mathcal{N}$, we have that $W \notin F(Q)$. Recall that $\lim_k Q_{1/n_k} = Q$ by construction, and thus Corollary 2 implies that $W \in \mathcal{W}_N^*(Q; \epsilon)$. Since this contradicts the fact that $W \notin \mathcal{W}_N^*(Q; \epsilon)$, we conclude that F is necessarily upper semicontinuous over $\{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$. \square

Now we are in position to prove Theorem 7.

Proof. For ease of notation, let

$$\mathcal{Q}^{(\delta)} \triangleq \{Q \in \mathcal{Q} : \epsilon_{\min}(Q) + \delta \leq \epsilon\}. \quad (271)$$

Observe that $\mathcal{Q}^{(\delta)} \subset \{Q \in \mathcal{Q} : \epsilon_{\min}(Q) < \epsilon\}$. Thus, Lemma 17 implies that the mapping $\mathcal{W}_N^*(\cdot; \epsilon)$ is upper semicontinuous over $\mathcal{Q}^{(\delta)}$. Observe that both $\mathcal{Q}^{(\delta)}$ and \mathcal{W}_N are compact metric spaces. Therefore, Theorem 11 establishes the lower semicontinuity of the mapping $\mathcal{W}_N^*(\cdot; \epsilon)$ on a residual set of $\mathcal{Q}^{(\delta)}$. Theorem 7 follows immediately from this fact and Definition 11. \square

APPENDIX C PROOFS FROM SECTION V

A. Lemma 8

Proof.

(i) Let $W \in \mathcal{W}_N$ be fixed. By condition (C.2), the mapping $\mathcal{U}(\cdot, W)$ is (Lipschitz) continuous over $\mathcal{Q}_r(\hat{P})$. Since $\mathcal{Q}_r(\hat{P})$ is compact, the mapping $\mathcal{U}(\cdot, W)$ attains its infimum over $\mathcal{Q}_r(\hat{P})$.

(ii) Assume that $\mathcal{D}_{\mathcal{Q}, N}(\hat{P}; \epsilon, r)$ is non-empty. By (i), we have that

$$\mathcal{U}_r(\hat{P}, W) = \min_{Q \in \mathcal{Q}_r(\hat{P})} \mathcal{U}(Q, W) \quad (272)$$

is well defined. As established in Lemma 7, the mapping $\mathcal{U}(\cdot, \cdot)$ is continuous over $\mathcal{Q} \times \mathcal{W}_N$ and in particular over $\mathcal{Q}_r(\hat{P}) \times \mathcal{W}_N$. Hence, Lemma 14 implies that $\mathcal{U}_r(\hat{P}, \cdot)$ is continuous over \mathcal{W}_N .

Note that, by definition,

$$\mathcal{D}_{\mathcal{Q}, N}(\hat{P}; \epsilon, r) = \bigcap_{Q \in \mathcal{Q}_r(\hat{P})} \mathcal{D}_N(Q; \epsilon), \quad (273)$$

where $\mathcal{D}_N(Q; \epsilon) = \{W \in \mathcal{W}_N : \mathcal{L}(Q, W) \leq \epsilon\}$. Condition (C.1) implies that, for every $Q \in \mathcal{Q}$, the mapping $\mathcal{L}(Q, \cdot)$ is continuous over \mathcal{W}_N . Thus, we have that $\mathcal{D}_N(Q; \epsilon)$ is compact as \mathcal{W}_N is compact. Since the arbitrary intersection of compact sets is also compact, (273) readily shows that $\mathcal{D}_{\mathcal{Q}, N}(\hat{P}; \epsilon, r)$ is compact. Therefore, the (continuous) mapping $\mathcal{U}_r(\hat{P}, \cdot)$ attains its supremum over $\mathcal{D}_{\mathcal{Q}, N}(\hat{P}; \epsilon, r)$, as required. \square

B. Theorem 8

Consider the following lemma.

Lemma 18. *Assume that conditions (C.1–3) hold true for a given closed set $\mathcal{Q} \subseteq \mathcal{P}$ and a given $N \in \mathbb{N}$. Let $\epsilon \in \mathbb{R}$ and $r \geq 0$ be given. If $P \in \mathcal{Q}_r(\hat{P})$ and $\epsilon - C_L r \geq \epsilon_{\min}(\hat{P})$, then*

$$\mathcal{H}(P; \epsilon) \leq \mathcal{H}(\hat{P}; \epsilon + C_L r) + C_U r. \quad (274)$$

Proof. We start proving that $\epsilon \geq \epsilon_{\min}(P)$. Remark 7 implies that there exists $W' \in \mathcal{W}_N$ such that $\mathcal{L}(\hat{P}, W') = \epsilon_{\min}(\hat{P})$. By the Lipschitz continuity given in condition (C.2), we have that

$$\mathcal{L}(P, W') \leq \mathcal{L}(\hat{P}, W') + |\mathcal{L}(\hat{P}, W') - \mathcal{L}(P, W')| \quad (275)$$

$$\leq \epsilon_{\min}(\hat{P}) + C_L \|\hat{P} - P\|_1 \quad (276)$$

$$\leq \epsilon_{\min}(\hat{P}) + C_L r \leq \epsilon, \quad (277)$$

where the last inequality follows from the assumption $\epsilon - C_L r \geq \epsilon_{\min}(\hat{P})$. By the minimality of $\epsilon_{\min}(P)$, we conclude that $\epsilon_{\min}(P) \leq \epsilon$. Then, Remark 7 implies that there exists $W \in \mathcal{W}_N$ such that $\mathcal{L}(P, W) \leq \epsilon$ and

$$\mathcal{H}(P; \epsilon) = \mathcal{U}(P, W). \quad (278)$$

By the Lipschitz continuity given in condition (C.2), we have that

$$|\mathcal{U}(\hat{P}, W) - \mathcal{U}(P, W)| \leq C_U \|\hat{P} - P\|_1 \leq C_U r. \quad (279)$$

In particular,

$$\mathcal{H}(P; \epsilon) \leq \mathcal{U}(\hat{P}, W) + C_U r. \quad (280)$$

Similarly, since

$$|\mathcal{L}(\hat{P}, W) - \mathcal{L}(P, W)| \leq C_L \|\hat{P} - P\|_1 \leq C_L r, \quad (281)$$

we have $\mathcal{L}(\hat{P}, W) \leq \mathcal{L}(P, W) + C_L r \leq \epsilon + C_L r$. Therefore, from inequality (280) and the maximality of the privacy-utility function, we conclude that

$$\mathcal{H}(P; \epsilon) \leq \mathcal{H}(\hat{P}; \epsilon + C_L r) + C_U r, \quad (282)$$

as we wanted to prove. \square

Now we are in position to prove Theorem 8.

Proof. For any $W \in \mathcal{D}_N(\hat{P}; \epsilon - C_L r)$, we have $\mathcal{L}(\hat{P}, W) + C_L r \leq \epsilon$. By the Lipschitz continuity given in condition (C.2), for every $Q \in \mathcal{Q}_r(\hat{P})$,

$$|\mathcal{L}(\hat{P}, W) - \mathcal{L}(Q, W)| \leq C_L \|\hat{P} - Q\|_1 \leq C_L r. \quad (283)$$

Hence, $\mathcal{L}(Q, W) \leq \epsilon$ and $W \in \mathcal{D}_N(Q; \epsilon)$ for every $Q \in \mathcal{Q}_r(\hat{P})$. By (117), it follows that $W \in \mathcal{D}_{\mathcal{Q}, N}(\hat{P}; \epsilon, r)$.

Let $W^* \in \mathcal{W}_N^*(\hat{P}; \epsilon - C_L r)$ and $W^\dagger \in \mathcal{W}_{\mathcal{Q}, N}^\dagger(\hat{P}; \epsilon, r)$. By definition,

$$\mathcal{U}(\hat{P}, W^*) = \mathcal{H}(\hat{P}; \epsilon - C_L r). \quad (284)$$

By the Lipschitz continuity given in condition (C.2), we have that

$$|\mathcal{U}(\hat{P}, W^*) - \mathcal{U}(P, W^*)| \leq C_U \|\hat{P} - P\|_1 \leq C_U r. \quad (285)$$

Combining (284) and (285) together, we have

$$\mathcal{U}(P, W^*) \geq \mathcal{H}(\hat{P}; \epsilon - C_L r) - C_U r. \quad (286)$$

Since $W^\dagger \in \mathcal{W}_{\mathcal{Q}, N}^\dagger(\hat{P}; \epsilon, r) \subset \mathcal{D}_{\mathcal{Q}, N}(\hat{P}; \epsilon, r)$ and $P \in \mathcal{Q}_r(\hat{P})$, (117) implies that $W^\dagger \in \mathcal{D}_N(P; \epsilon)$. Thus, by the maximality of $\mathcal{H}(P; \epsilon)$, we have

$$\mathcal{U}(P, W^\dagger) \leq \mathcal{H}(P; \epsilon). \quad (287)$$

An immediate application of Lemma 18 shows that

$$\mathcal{U}(P, W^\dagger) \leq \mathcal{H}(P; \epsilon) \leq \mathcal{H}(\hat{P}; \epsilon + C_L r) + C_U r. \quad (288)$$

Therefore, combining (286) and (288) together,

$$\mathcal{U}(P, W^*) - \mathcal{U}(P, W^\dagger) \quad (289)$$

$$\geq - \left(\mathcal{H}(\hat{P}; \epsilon + C_L r) - \mathcal{H}(\hat{P}; \epsilon - C_L r) + 2C_U r \right), \quad (290)$$

which implied the desired conclusion. \square

C. Theorem 9

Proof. Recall that, by Lemma 15, the mapping $\epsilon_{\min}(\cdot)$ is continuous over \mathcal{Q} . In particular, we have that $\lim_n \epsilon_{\min}(P_n) = \epsilon_{\min}(P)$ whenever $\lim_n P_n = P$. Since $\epsilon_{\min}(P) < \epsilon$, $\lim_n P_n = P$, and $\lim_n r_n = 0$, we have that $\epsilon_{\min}(P_n) + C_L r_n < \epsilon$ for n large enough. Therefore, for n large enough, the first part of Theorem 8 establishes that the set $\mathcal{D}_{\mathcal{Q}, N}(P_n; \epsilon, r_n)$ is non-empty. Furthermore, part (ii) of Lemma 8 implies that $\mathcal{W}_{\mathcal{Q}, N}^\dagger(P_n; \epsilon, r_n)$ is non-empty as well. Without loss of generality, we assume that $\mathcal{W}_{\mathcal{Q}, N}^\dagger(P_n; \epsilon, r_n) \neq \emptyset$ for all $n \geq 1$.

In order to reach contradiction, we assume that the conclusion does not hold true, i.e., that there exists a sequence $(W_n^\dagger)_{n=1}^\infty$ such that $W_n^\dagger \in \mathcal{W}_{\mathcal{Q},N}^\dagger(P_n; \epsilon, r_n)$ for all $n \geq 1$ and

$$\limsup_{n \rightarrow \infty} \text{dist}(W_n^\dagger, \mathcal{W}_N^*(P; \epsilon)) > 0. \quad (291)$$

Since both \mathcal{W}_N and $\mathcal{W}_N^*(P; \epsilon)$ are compact (Lemma 16), there exists a subsequence of $(W_n^\dagger)_{n=1}^\infty$ converging to some $W_0 \in \mathcal{W}_N$ with $W_0 \notin \mathcal{W}_N^*(P; \epsilon)$. Without loss of generality, we assume that $\lim_n W_n^\dagger = W_0$. By the continuity of $\mathcal{L}(\cdot, \cdot)$ established in Lemma 7,

$$\mathcal{L}(P, W_0) = \lim_{n \rightarrow \infty} \mathcal{L}(P_n, W_n^\dagger) \leq \epsilon. \quad (292)$$

The last inequality implies that $W_0 \in \mathcal{D}_N(P; \epsilon)$ and hence $\mathcal{U}(P, W_0) \leq \mathsf{H}(P; \epsilon)$. Assume that we have already proved that $\mathcal{U}(P, W_0) \geq \mathsf{H}(P; \epsilon)$. In particular, we would have that $\mathcal{U}(P, W_0) = \mathsf{H}(P; \epsilon)$ and thus $W_0 \in \mathcal{W}_N^*(P; \epsilon)$. Since this conclusion contradicts the fact that $W_0 \notin \mathcal{W}_N^*(P; \epsilon)$, the result would follow. Now we focus on proving that $\mathcal{U}(P, W_0) \geq \mathsf{H}(P; \epsilon)$.

Since $\epsilon > \epsilon_{\min}(P)$ and $\lim_n r_n = 0$, we have that $\epsilon - 2C_L r_n > \epsilon_{\min}(P)$ for n large enough. Without loss of generality, we assume that $\epsilon - 2C_L r_n > \epsilon_{\min}(P)$ for all $n \geq 1$. For a given $n \geq 1$, let $W_0^* \in \mathcal{W}_N$ be such that $\mathcal{L}(P, W_0^*) \leq \epsilon - 2C_L r_n$ and $\mathcal{U}(P, W_0^*) = \mathsf{H}(P; \epsilon - 2C_L r_n)$. By condition (C.2), for any $Q \in \mathcal{Q}_{r_n}(P_n)$,

$$|\mathcal{L}(P, W_0^*) - \mathcal{L}(Q, W_0^*)| \leq C_L \|P - Q\|_1 \leq 2C_L r_n, \quad (293)$$

where the last inequality follows from the triangle inequality and the fact that $P, Q \in \mathcal{Q}_{r_n}(P_n)$. Thus,

$$\mathcal{L}(Q, W_0^*) \leq \mathcal{L}(P, W_0^*) + 2C_L r_n \leq \epsilon, \quad (294)$$

for any $Q \in \mathcal{Q}_{r_n}(P_n)$. Hence, $W_0^* \in \mathcal{D}_{\mathcal{Q},N}(P_n; \epsilon, r_n)$ and, by the definition of $\mathcal{W}_{\mathcal{Q},N}^\dagger(P_n; \epsilon, r_n)$,

$$\min_{Q \in \mathcal{Q}_{r_n}(P_n)} \mathcal{U}(Q, W_0^*) \leq \min_{Q \in \mathcal{Q}_{r_n}(P_n)} \mathcal{U}(Q, W_n^\dagger) \leq \mathcal{U}(P, W_n^\dagger). \quad (295)$$

As in (293), we have that

$$|\mathcal{U}(P, W_0^*) - \mathcal{U}(Q, W_0^*)| \leq 2C_U r_n. \quad (296)$$

Therefore, for any $Q \in \mathcal{Q}_{r_n}(P_n)$,

$$\mathcal{U}(Q, W_0^*) \geq \mathcal{U}(P, W_0^*) - 2C_U r_n \quad (297)$$

$$= \mathsf{H}(P; \epsilon - 2C_L r_n) - 2C_U r_n. \quad (298)$$

In particular, this implies that

$$\min_{Q \in \mathcal{Q}_{r_n}(P_n)} \mathcal{U}(Q, W_0^*) \geq \mathsf{H}(P; \epsilon - 2C_L r_n) - 2C_U r_n. \quad (299)$$

Combining (295) and (299) together, we have

$$\mathcal{U}(P, W_n^\dagger) \geq \min_{Q \in \mathcal{Q}_{r_n}(P_n)} \mathcal{U}(Q, W_0^*) \quad (300)$$

$$\geq \mathsf{H}(P; \epsilon - 2C_L r_n) - 2C_U r_n. \quad (301)$$

Therefore, by the continuity of $\mathcal{U}(P, \cdot)$ and $\mathsf{H}(P; \cdot)$ assumed in condition (C.1),

$$\mathcal{U}(P, W_0) = \lim_{n \rightarrow \infty} \mathcal{U}(P, W_n^\dagger) \quad (302)$$

$$\geq \lim_{n \rightarrow \infty} \mathsf{H}(P; \epsilon - 2C_L r_n) - 2C_U r_n \quad (303)$$

$$= \mathsf{H}(P; \epsilon). \quad (304)$$

This concludes the proof of the first part of the theorem.

Now we prove the second part of the theorem. Recall that in this case, for all $n \geq 1$, $P_n = \hat{P}_n$ and $r_n \geq (2p \log(n)/n)^{1/2}$ for some $p > 1$. Thus,

$$\sum_{n=1}^{\infty} \Pr(\|\hat{P}_n - P\|_1 \geq r_n) \quad (305)$$

$$\leq \sum_{n=1}^{\infty} \Pr\left(\|\hat{P}_n - P\|_1 \geq \sqrt{\frac{2p \log(n)}{n}}\right) \quad (306)$$

$$\leq \exp(|\mathcal{S}| \cdot |\mathcal{X}|) \sum_{n=1}^{\infty} \frac{1}{n^p}, \quad (307)$$

where (307) follows directly from (32). Since $p > 1$ by assumption, we have that $\sum_n 1/n^p$ is finite. Thus, by a routine application of the Borel-Cantelli Lemma, see, e.g., [86, Sec. 2.18],

$$\Pr\left(\left\{\omega : \|\hat{P}_n(\omega) - P\|_1 < r_n \text{ for all } n \geq N = N(\omega)\right\}\right) = 1.$$

Observe that, by the last equality, the hypotheses of the first part of this theorem are satisfied almost surely. A direct application of the former part of this theorem leads to

$$\Pr\left(\lim_{n \rightarrow \infty} \text{dist}(W_n^\dagger, \mathcal{W}_N^*(P; \epsilon)) = 0\right) = 1, \quad (308)$$

as we wanted to prove. \square

REFERENCES

- [1] 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, pp. 706–710.
- [2] M. Baker, "Is there a reproducibility crisis? A Nature survey lifts the lid on how researchers view the 'crisis rocking science and what they think will help,'" *Nature*, vol. 533, no. 7604, pp. 452–455, 2016.
- [3] J. M. Wicherts and M. Bakker, "Publish (your data) or (let the data) perish! Why not publish your data too??" *Intelligence*, vol. 40, no. 2, pp. 73–76, 2012.
- [4] D. Donoho, "50 years of data science," *J. Comput. Graph. Statist.*, vol. 26, no. 4, pp. 745–766, 2017.
- [5] L. Sweeney, "Only you, your doctor, and many others may know," *Technology Science*, vol. 2015092903, no. 9, p. 29, 2015.
- [6] K. Granville, "Facebook and Cambridge Analytica: What you need to know as fallout widens," *The New York Times*, vol. 19, 2018.
- [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] J. F. Handler and M. K. Rosenheim, "Privacy in welfare: Public assistance and juvenile justice," *Law & Contemp. Probs.*, vol. 31, pp. 377–412, 1966.
- [9] D. Rebolledo-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.
- [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] L. Sweeney, "k-anonymity: A model for protecting privacy," *Int. J. Uncertainty Fuzziness Knowl.-Based Syst.*, vol. 10, no. 05, pp. 557–570, 2002.

- [12] C. Dwork, "Differential privacy," in *Encyclopedia of Cryptography and Security*. Springer, 2011, pp. 338–340.
- [13] I. Issa, S. Kamath, and A. B. Wagner, "Maximal leakage minimization for the Shannon cipher system," in *Proc. 2016 IEEE Int. Symp. on Inf. Theory*, 2016, pp. 520–524.
- [14] B. Rassouli and D. Gunduz, "Optimal utility-privacy trade-off with total variation distance as a privacy measure," *IEEE Trans. Inf. Forensics Security*, 2019.
- [15] H. Yamamoto, "A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers," *IEEE Trans. Inf. Theory*, vol. 29, no. 6, pp. 918–923, 1983.
- [16] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," in *Adv. Neural Inf. Process. Syst.*, 2014, pp. 2879–2887.
- [17] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Estimation efficiency under privacy constraints," *IEEE Transactions on Information Theory*, vol. 65, no. 3, pp. 1512–1534, 2018.
- [18] F. P. Calmon and N. Fawaz, "Privacy against statistical inference," in *Proc. 50th Annu. Allerton Conf. Commun. Control Comput.*, 2012, pp. 1401–1408.
- [19] F. P. Calmon, A. Makhdoomi, and M. Médard, "Fundamental limits of perfect privacy," in *Proc. 2015 IEEE Int. Symp. on Inf. Theory*, 2015, pp. 1796–1800.
- [20] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Information extraction under privacy constraints," *Information*, vol. 7, no. 1, pp. 1–15, 2016.
- [21] Y. O. Basciftci, Y. Wang, and P. Ishwar, "On privacy-utility tradeoffs for constrained data release mechanisms," in *2016 IEEE Information Theory and Applications Workshop (ITA)*, 2016, pp. 1–6.
- [22] 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.
- [23] S. A. Osia, B. Rassouli, H. Haddadi, H. R. Rabiee, and D. Gündüz, "Privacy against brute-force inference attacks," *arXiv preprint arXiv:1902.00329*, 2019.
- [24] B. Rassouli, F. E. Rosas, and D. Gunduz, "Data disclosure under perfect sample privacy," *arXiv preprint arXiv:1904.01711*, 2019.
- [25] A. Nageswaran and P. Narayan, "Data privacy for a ρ -recoverable function," *IEEE Trans. Inf. Theory*, 2019.
- [26] C. Huang, P. Kairouz, X. Chen, L. Sankar, and R. Rajagopal, "Context-aware generative adversarial privacy," *Entropy*, vol. 19, no. 12, pp. 1–35, 2017, Art. ID 656.
- [27] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdú, and M. J. Weinberger, "Inequalities for the l_1 deviation of the empirical distribution," *Hewlett-Packard Labs, Tech. Rep.*, 2003.
- [28] S. Kamath, A. Orlitsky, D. Pichapati, and A. T. Suresh, "On learning distributions from their samples," in *Proc. 28th Conf. on Learn. Theory*, 2015, pp. 1066–1100.
- [29] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [30] C. E. Shannon, "Communication theory of secrecy systems," *Bell Syst. Tech. J.*, vol. 28, no. 4, pp. 656–715, 1949.
- [31] S. Verdú, " α -mutual information," in *2015 IEEE Information Theory and Applications Workshop (ITA)*, 2015, pp. 1–6.
- [32] I. Csiszár, "Information-type measures of difference of probability distributions and indirect observations," *Studia Sci. Math. Hungar.*, vol. 2, pp. 299–318, 1967.
- [33] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "Privacy under hard distortion constraints," *arXiv preprint arXiv:1806.00063*, 2018.
- [34] I. Issa, S. Kamath, and A. B. Wagner, "An operational measure of information leakage," in *Proc. IEEE Annu. Conf. on Inf. Sci. and Syst. (CISS)*, 2016, pp. 234–239.
- [35] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "A tunable measure for information leakage," in *Proc. 2018 IEEE Int. Symp. on Inf. Theory*, 2018, pp. 701–705.
- [36] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Privacy aware learning," *J. ACM*, vol. 61, no. 6, pp. 38:1–38:57, Dec. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2666468>
- [37] I. Wagner and D. Eckhoff, "Technical privacy metrics: A systematic survey," *ACM Comput. Surv.*, vol. 51, no. 3, pp. 57:1–57:38, Jun. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3168389>
- [38] C. T. Li and A. El Gamal, "Maximal correlation secrecy," *IEEE Trans. Inf. Theory*, vol. 64, no. 5, pp. 3916–3926, 2018.
- [39] I. Issa, A. B. Wagner, and S. Kamath, "An operational approach to information leakage," *arXiv preprint arXiv:1807.07878*, 2018.
- [40] A. Russell and H. Wang, "How to fool an unbounded adversary with a short key," in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2002, pp. 133–148.
- [41] Y. Dodis and A. Smith, "Entropic security and the encryption of high entropy messages," in *Theory of Cryptography Conference*. Springer, 2005, pp. 556–577.
- [42] 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.
- [43] O. Shamir, S. Sabato, and N. Tishby, "Learning and generalization with the information bottleneck," *Theor. Comput. Sci.*, vol. 411, no. 29–30, pp. 2696–2711, 2010.
- [44] J. Jiao, K. Venkat, Y. Han, and T. Weissman, "Minimax estimation of functionals of discrete distributions," *IEEE Trans. Inf. Theory*, vol. 61, no. 5, pp. 2835–2885, 2015.
- [45] Y. Wu and P. Yang, "Minimax rates of entropy estimation on large alphabets via best polynomial approximation," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3702–3720, 2016.
- [46] D. Varodayan and A. Khisti, "Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 1932–1935.
- [47] O. Tan, D. Gunduz, and H. V. Poor, "Increasing smart meter privacy through energy harvesting and storage devices," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 7, pp. 1331–1341, 2013.
- [48] A. Makhdoomi, 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.
- [49] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annu. Allerton Conf. Commun. Control Comput.*, 1999, pp. 368–377.
- [50] 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.
- [51] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton University Press, 2009.
- [52] P. J. Huber, "Robust statistics," in *International Encyclopedia of Statistical Science*. Springer, 2011, pp. 1248–1251.
- [53] L. G. Valiant, "Robust logics," *Artif. Intell. J.*, vol. 117, no. 2, pp. 231–253, 2000.
- [54] M. Charikar, J. Steinhardt, and G. Valiant, "Learning from untrusted data," in *Proc. 49th ACM SIGACT Symp. on Theory of Computing*, 2017, pp. 47–60.
- [55] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University press, 2014.
- [56] M. J. Kearns, U. V. Vazirani, and U. Vazirani, *An Introduction to Computational Learning Theory*. MIT press, 1994.
- [57] A. Xu and M. Raginsky, "Information-theoretic analysis of generalization capability of learning algorithms," in *Adv. Neural Inf. Process. Syst.*, 2017, pp. 2524–2533.
- [58] D. Russo and J. Zou, "How much does your data exploration overfit? Controlling bias via information usage," *arXiv preprint arXiv:1511.05219*, 2015.
- [59] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, "Preserving statistical validity in adaptive data analysis," in *Proc. 47th ACM Symp. on Theory of Computing*, 2015, pp. 117–126.
- [60] ———, "Generalization in adaptive data analysis and holdout reuse," in *Adv. Neural Inf. Process. Syst.*, 2015, pp. 2350–2358.
- [61] ———, "The reusable holdout: Preserving validity in adaptive data analysis," *Science*, vol. 349, no. 6248, pp. 636–638, 2015.
- [62] S. Arimoto, "Information measures and capacity of order α for discrete memoryless channels," *Topics in Information Theory, Coll. Math. Soc. J. Bolyai (I. Csiszár and P. Elias Eds.)*, vol. 16, pp. 41–52, 1977.
- [63] S. Fehr and S. Berens, "On the conditional Rényi entropy," *IEEE Trans. Inf. Theory*, vol. 60, no. 11, pp. 6801–6810, 2014.
- [64] G. Smith, "On the foundations of quantitative information flow," in *International Conference on Foundations of Software Science and Computational Structures*. Springer, 2009, pp. 288–302.
- [65] C. Braun, K. Chatzikokolakis, and C. Palamidessi, "Quantitative notions of leakage for one-try attacks," *Electron. Notes Theor. Comput. Sci.*, vol. 249, pp. 75–91, 2009.
- [66] G. Barthe and B. Kopf, "Information-theoretic bounds for differentially private mechanisms," in *2011 IEEE 24th Computer Security Foundations Symposium*. IEEE, 2011, pp. 191–204.
- [67] R. Sibson, "Information radius," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 14, no. 2, pp. 149–160, 1969.

- [68] 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.
- [69] 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.
- [70] M. J. Greenacre, *Theory and Applications of Correspondence Analysis*. San Diego, CA, USA: Academic, Mar. 1984.
- [71] 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.
- [72] J. Angwin, J. Larson, S. Mattu, and L. Kirchner, "Machine bias," *ProPublica*, 2016.
- [73] Y. Polyanskiy and Y. Wu, "Lecture notes on information theory," *Lecture Notes for ECE563 (UIUC) and 6.441 (MIT)*, 2012–2017. [Online]. Available: http://people.lids.mit.edu/yp/homepage/data/itlectures_v5.pdf
- [74] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed. New York: McGraw-hill, 1976.
- [75] J.-P. Aubin and H. Frankowska, *Set-valued Analysis*. Springer Science & Business Media, 2009.
- [76] W. Rudin, *Real and Complex Analysis*, 2nd ed. McGraw-Hill, 1987.
- [77] G. B. Folland, *Real Analysis: Modern Techniques and Their Applications*. Wiley-Interscience, New York, 1984.
- [78] 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.
- [79] 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.
- [80] H. G. Eggleston, *Convexity*. Cambridge University Press, 1958.
- [81] 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.
- [82] H. L. Royden, *Real Analysis*, 2nd ed. New York: Macmillan, 1968.
- [83] S. I. Resnick, *Extreme Values, Regular Variation and Point Processes*. Springer, 2013.
- [84] V. Y. Tan, A. Anandkumar, L. Tong, and A. S. Willsky, "A large-deviation analysis of the maximum-likelihood learning of Markov tree structures," *IEEE Trans. Inf. Theory*, vol. 57, no. 3, pp. 1714–1735, 2011.
- [85] J. C. Oxtoby, *Measure and Category: A Survey of the Analogies between Topological and Measure Spaces*. Springer Science & Business Media, 2013, vol. 2.
- [86] A. Gut, *Probability: A Graduate Course*. Springer Science & Business Media, 2013, vol. 75.

Mario Diaz currently holds a postdoctoral research appointment at Centro de Investigación en Matemáticas (CIMAT), Guanajuato, Mexico. He received his B.Eng. degree in Electrical Engineering in 2011 from the University of Guadalajara, Mexico, his M.Sc. degree in Probability and Statistics in 2013 from CIMAT, Mexico, and his Ph.D. degree in Mathematics and Statistics in 2017 from Queen's University, Canada. Before joining CIMAT in 2018, he was a joint postdoctoral scholar at Arizona State University and Harvard University. His research interests include the mathematical and statistical foundations of information privacy, information theory, theoretical machine learning, random matrix theory, and free probability theory. During his graduate studies, Dr. Diaz was a recipient of the Ontario Trillium Scholarship.

Hao Wang is currently a Ph.D. candidate at Harvard University. He received his M.S. degree in applied mathematics from Harvard University in 2019. In 2016, he received his B.S. degree in mathematics from the University of Science and Technology of China (USTC). His research interests are in information theory and statistics, with applications to fairness and privacy in machine learning.

Flavio P. Calmon is an Assistant Professor of Electrical Engineering at Harvard's John A. Paulson School of Engineering and Applied Sciences. Before joining Harvard, he was the inaugural data science for social good post-doctoral fellow at IBM Research in Yorktown Heights, New York. He received his Ph.D. in Electrical Engineering and Computer Science at MIT. His main research interests are information theory, inference, and statistics, with applications to fairness, privacy, machine learning, and communications engineering. Prof. Calmon has received the NSF CAREER Award, the Google Research Faculty Award, the IBM Open Collaborative Research Award, and Harvard's Lemann Brazil Research Fund Award.

Lalitha Sankar (S'02-M'07-SM'15) received the B.Tech. degree from the Indian Institute of Technology, Bombay, the M.S. degree from the University of Maryland, and the Ph.D. degree from Rutgers University. She is currently an Associate Professor in the School of Electrical, Computer, and Energy Engineering at Arizona State University. Prior to this, she was an Associate Research Scholar at Princeton University and a recipient of a three year Science and Technology Teaching postdoctoral fellowship from the Council on Science and Technology at Princeton University. Sankar has also worked as a Senior Member of Technical Staff at AT&T Shannon Labs and Polaroid Engineering R&D Labs. Her research interests include applying information sciences to study data privacy as well as cybersecurity and resilience in critical infrastructure networks. For her doctoral work, she received the 2007-2008 Electrical Engineering Academic Achievement Award from Rutgers University. She received the IEEE Globecom 2011 Best Paper Award for her work on privacy of side-information in multi-user data systems, and the National Science Foundation CAREER Award in 2014.