



HAL
open science

Pointwise Maximal Leakage

Sara Saeidian, Giulia Cervia, Tobias J. Oechtering, Mikael Skoglund

► **To cite this version:**

Sara Saeidian, Giulia Cervia, Tobias J. Oechtering, Mikael Skoglund. Pointwise Maximal Leakage. 2022 IEEE International Symposium on Information Theory, Jun 2022, Aalto, Finland. 10.1109/ISIT50566.2022.9834814 . hal-04155244

HAL Id: hal-04155244

<https://hal.science/hal-04155244v1>

Submitted on 7 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pointwise Maximal Leakage

Sara Saeidian, *Member, IEEE*, Giulia Cervia, *Member, IEEE*,

Tobias J. Oechtering, *Senior Member, IEEE* Mikael Skoglund, *Fellow, IEEE*

Abstract

We introduce a privacy measure called *pointwise maximal leakage*, defined based on the pre-existing notion of *maximal leakage*, which quantifies the amount of information leaking about a secret X by disclosing a single outcome of a (randomized) function calculated on X . Pointwise maximal leakage is a robust and operationally meaningful privacy measure that captures the largest amount of information leaking about X to adversaries seeking to guess arbitrary (possibly randomized) functions of X , or equivalently, aiming to maximize arbitrary gain functions. We study several properties of pointwise maximal leakage, e.g., how it composes over multiple outcomes, how it is affected by pre- and post-processing, etc. Furthermore, we propose to view privacy leakage as a *random variable* which, in turn, allows us to regard privacy *guarantees* as requirements imposed on different statistical properties of the privacy leakage random variable. We define several privacy guarantees and study how they behave under pre-processing, post-processing and composition. Finally, we examine the relationship between pointwise maximal leakage and other privacy notions such as local differential privacy, local information privacy, f -information, and so on.

Index Terms

Privacy, information leakage, maximal leakage, g -leakage, pointwise leakage, privacy random variable, information density.

I. INTRODUCTION

Suppose X is a random variable representing some data containing sensitive information. As we aim to remain general, we intentionally keep X abstract. For example, X may be a single data entry collected from an individual (the *local* setting), X may represent an entire database containing sensitive information (the *centralized* setting), or X may be a secret such as a password that must be kept confidential (the *side-channel* setting).

Further, suppose Y is the output of a (randomized) function with input X . For example, in the local setting, Y denotes the perturbed version of a single user's data which is collected by a data curator. In the centralized setting, Y may be some aggregate statistic calculated on a database, and in the side-channel setting, Y is the output of a side-channel with input X , for instance, the inter-keystroke delays when typing in a password. In all of these scenarios, we are interested in answering the following question: How much information is Y leaking about X ?

The above question has been studied and answered in different contexts using a myriad of different privacy measures. For instance, differential privacy (DP) [1, 2] was introduced in a centralized setting within the context

of database privacy in order to ensure that no single individual’s participation can be revealed from the output of a function calculated on a database (i.e., ensuring *membership privacy*). Later on, borrowing ideas from survey privacy [3], the adaptation of differential privacy to the local setting led to the concept of local differential privacy [4–6], where, roughly speaking, the goal is to provide plausible deniability for all possible values of input X . Parallel to these developments, in the information theory literature, a wide range of privacy measures have been proposed and studied that aim to measure the dependence between two random variables X and Y . These include, for example, mutual information [7–10], and its generalizations [11, 12], divergence-based measures (e.g., metrics based on f -divergence [13, 14]), probability of correctly guessing [15], information privacy [16, 17] and indistinguishability [9]. The two recent surveys by Wagner and Eckhoff [18] and Bloch et al. [19] contain an extensive list of various privacy measures.

In much of the literature, the prevalent approach to tackling privacy problems has been to start with a particular definition of privacy, study the properties that follow from the definition, and design/optimize mechanisms that guarantee a certain level of privacy and utility. An alternative approach is to start from a *threat model* in which an adversary with explicitly-described capabilities is pursuing a specific objective, and study the system’s vulnerability as a result of this adversarial model. This approach has several advantages. First, it encourages us to make our assumptions about the capabilities of the adversary (e.g., in terms of computational power or prior knowledge of the system) and her objectives explicit. Second, the privacy definition obtained by studying a threat model is operationally meaningful and easier to interpret. Third, the discussions around the advantages and limitations of different privacy measures become more transparent and objective. Note that here we are making a subtle but essential distinction: Several notable privacy measures such as differential privacy and differential indistinguishability [20] can be *interpreted* using powerful adversarial models, but their definitions do not follow from any such model. In fact, not making this distinction may lead to misconceptions/disagreements about what a privacy definition does or does not promise. An example of this is the long ongoing debate about whether or not differential privacy (implicitly) requires assumptions about the data generating (prior) distribution, such as the assumption that the entries in a database are drawn independently [21, 22]. This disagreement may have been avoided had the assumptions about the adversarial model been made explicit in the definition.

The threat-model approach to privacy has been adopted in a line of work termed *quantitative information flow* [23–28], in which several notions of information leakage are motivated, defined and studied. One such notion is *min-entropy leakage* [23, 24] which is defined in a setup where a passive but computationally-unbounded adversary is trying to guess the value of the secret X in one try. Min-entropy leakage, then, quantifies the increase in the probability of correctly guessing X having observed the output Y , compared to guessing the secret with no observations. Naturally, min-entropy leakage depends both on the prior distribution of X , denoted by P_X , and the channel from X to Y , denoted by $P_{Y|X}$. Therefore, to obtain a privacy measure that depends only on the channel $P_{Y|X}$, Braun et al. [24] consider the problem of maximizing min-entropy leakage over all possible priors, which leads to a quantity (later) called *maximal leakage*, and defined as

$$\mathcal{L}(P_{Y|X}) = \log \sum_y \max_{x: P_X(x) > 0} P_{Y|X=x}(y). \quad (1)$$

Interestingly, the worst-case prior in this problem is the uniform distribution [24].

Subsequent works in this area have extended the above adversarial model [25–28], developed maximal leakage into a practical tool for quantifying privacy in learning applications [29], and considered maximal leakage as the privacy constraint in privacy-utility tradeoff problems [30–32]. We find two such works particularly interesting: the g -leakage framework introduced by Alvim et al. [26] and the maximal leakage definition of Issa et al. [33]. The g -leakage [26] formulation generalizes the setup by considering an adversary aiming to construct a guess of X that maximizes a certain *gain* function, and constitutes a useful tool for modeling a variety of different adversarial goals, such as guessing the secret X in $k \geq 1$ attempts, or approximately guessing the secret [26]. Moreover, it has been shown that for all prior distributions, maximizing g -leakage over all possible gain functions yields a quantity that is equal to maximal leakage [27]. The setup put forward by Issa et al. [33], on the other hand, develops the framework by considering adversaries interested in guessing a possibly-randomized discrete function of X , called U . Then, taking the supremum over all such U 's, the resulting quantity is once again equal to maximal leakage.

The above results lead us to believe that maximal leakage is a powerful privacy measure with a robust definition that is relevant in a multitude of different scenarios. However, one apparent limitation of the definition given in (1) is that maximal leakage is defined for the *average outcome* Y . Hence, a privacy guarantee given in terms of maximal leakage does not allow us to distinguish between individual outcomes based on how much information they leak. To see why this can be problematic, suppose X is a uniformly distributed ternary random variable defined over the set $\mathcal{X} = \{x_1, x_2, x_3\}$, and consider the following channels from X to Y :

$$P_{Y|X} = \begin{bmatrix} 1 & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}, \quad Q_{Y|X} = \begin{bmatrix} \frac{2}{3} & \frac{1}{6} & \frac{1}{6} \\ \frac{1}{6} & \frac{2}{3} & \frac{1}{6} \\ \frac{1}{6} & \frac{1}{6} & \frac{2}{3} \end{bmatrix},$$

where Y is a random variable defined over $\mathcal{Y} = \{y_1, y_2, y_3\}$ and $[P_{Y|X}]_{ij} = P_{Y|X=x_i}(y_j)$, $i, j \in \{1, 2, 3\}$. It is easy to see that both channels have equal maximal leakage $\mathcal{L}(P_{Y|X}) = \mathcal{L}(Q_{Y|X}) = \log 2$; however, they are qualitatively different. For example, the third outcome in $P_{Y|X}$, y_3 , can be considered far less private compared to the other two as it completely reveals the value of the secret (which is x_3). On the other hand, due to the symmetry in $Q_{Y|X}$, we may expect that all outcomes leak the same amount of information. One justification for this issue is that maximal leakage remains small as long as highly revealing outputs occur with small probability [33]. Nevertheless, the average-case guarantee provided by maximal leakage may be deemed insufficient in privacy-critical applications with strict requirements, which motivates the search for alternative privacy measures that capture the *distribution* of privacy leakage over the outcomes. Using the leakage distribution not only do we get a precise description of privacy in a given system, but we will also have the flexibility to adapt our definition of a private system to each specific application. For example, system designers will be able to decide whether or not highly-revealing but low-probability outcomes pose a privacy risk on a per-application basis.

We should point out that a distributional view of privacy also exists in the differential privacy literature, where the log likelihood ratio is referred to as the *privacy loss random variable* [34–36]. This approach is further formalized in [37] by introducing the notion of *privacy loss distribution*. The framework we develop in this paper differs from

these works mainly in that here we let X represent any type of data containing sensitive information, while the works on differential privacy target specifically the centralized setting, where the notion of neighboring databases (that is, databases that differ in a single entry) plays a central role in all definitions. On top of that, the privacy measure introduced in this paper has a clear and useful operational meaning, and is obtained by analyzing specific threat models (more on this below).

A. Overview

1) *Introducing pointwise maximal leakage*: In this paper, our main goal is to define a robust and operationally meaningful privacy measure that describes the amount of information leaking about the secret X due to disclosing a single outcome $Y = y$. We consider the same adversarial models that were used to obtain maximal leakage as a privacy measure (e.g., [33, 26]), but redirect our attention from the “average outcome” characterization of the previous works to individual outcomes, and obtain a new privacy definition which we will call *pointwise maximal leakage* (PML), denoted by $\ell_{P_{XY}}(X \rightarrow y)$, and expressed as

$$\ell_{P_{XY}}(X \rightarrow y) = \log \max_{x: P_X(x) > 0} \frac{P_{X|Y=y}(x)}{P_X(x)}. \quad (2)$$

In Section II-A, we start from the threat model of [33], in which an adversary is trying to guess the outcome of a randomized function of X , denoted by U . More concretely, we define a quantity $\ell_U(X \rightarrow y)$ as the logarithm of the ratio of the probability of correctly guessing U having observed an outcome y , and the probability of correctly guessing U with no observations. Then, we define pointwise maximal leakage as the supremum of $\ell_U(X \rightarrow y)$ over all U , that is

$$\ell_{P_{XY}}(X \rightarrow y) := \sup_{P_{U|X}} \ell_U(X \rightarrow y),$$

and obtain expression (2). We will call this approach the *randomized function view* of leakage. Afterwards, in Section II-B, we consider the threat model of the g -leakage framework [26], in which an adversary is trying to maximize a certain gain function g . More precisely, we define a quantity $\ell_g(X \rightarrow y)$ as the logarithm of the ratio of the expected gain having observed an outcome y , and the expected gain with no observations. We will call this approach the *gain function view* of leakage, and show that it is in fact equivalent to the randomized-function view of leakage. Specifically, we show that for all joint distributions over X and Y and for each randomized function of X , denoted by U , there exists a gain function g_U such that $\ell_U(X \rightarrow y) = \ell_{g_U}(X \rightarrow y)$. Conversely, we show that for each gain function g , there exists a randomized function of X , denoted by U_g , such that $\ell_g(X \rightarrow y) = \ell_{U_g}(X \rightarrow y)$. It follows that PML can alternatively be defined as

$$\ell_{P_{XY}}(X \rightarrow y) := \sup_g \ell_g(X \rightarrow y).$$

This result not only unifies two seemingly different ways of defining pointwise maximal leakage, but also signifies the *robustness* of pointwise maximal leakage as a privacy measure against a large class of adversaries with different objectives. Once we have established the definition of pointwise maximal leakage, in Section II-C, we study several properties following from the definition, e.g., how it composes due to observing several outcomes, how the leakage is affected by pre- and post-processing, and so on.

2) *Defining privacy guarantees based on PML*: Our second objective in this work is to argue in favor of viewing privacy leakage as a *random variable*. The idea behind this is simple: The amount of information leaked due to disclosing an outcome $Y = y$ is equal to $\ell_{P_{XY}}(X \rightarrow y)$ which is a function of y . Since Y is a random variable distributed according to P_Y (i.e., the output distribution induced by P_X and $P_{Y|X}$), this in turn allows us to define a random variable $\ell_{P_{XY}}(X \rightarrow Y)$ whose distribution is induced by P_Y . Adopting this view, a *privacy guarantee* is essentially a requirement we impose on some statistical property of $\ell_{P_{XY}}(X \rightarrow Y)$; thus, we have the flexibility to define different types of guarantees depending on how strict we may want to be. For example, we may require small privacy leakage *with probability one*. Less stringently, we can define guarantees that bound either the *tail* of the privacy leakage random variable, or its *expectation*.¹ These privacy guarantees, which are the subject of Section III, can be informally expressed as follows: Given an arbitrary but fixed prior P_X , we say that

- $P_{Y|X}$ satisfies ϵ -PML with $\epsilon \geq 0$ if $\ell_{P_{XY}}(X \rightarrow Y)$ is bounded by ϵ with probability one,
- $P_{Y|X}$ satisfies (ϵ, δ) -PML with $\epsilon \geq 0$ and $0 \leq \delta \leq 1$ if $\ell_{P_{XY}}(X \rightarrow Y)$ is bounded by ϵ with probability at least $1 - \delta$, and
- $P_{Y|X}$ satisfies $\mathcal{L}(P_{Y|X}) \leq \epsilon$ with $\epsilon \geq 0$, if the expectation of $\exp(\ell_{P_{XY}}(X \rightarrow Y))$ is bounded by $\exp(\epsilon)$, where $\mathcal{L}(\cdot)$ denotes maximal leakage.

In the rest of Section III, we study the *data-processing* and *composition* properties of the privacy guarantees introduced in this paper. Specifically, in Section III-A, we study how ϵ -PML and (ϵ, δ) -PML are affected by pre- and post-processing. Interestingly, it turns out that (ϵ, δ) -PML is not closed under post-processing, that is, given a privacy mechanism $P_{Y|X}$ that satisfies (ϵ, δ) -PML, we may be able to come up with a post-processing mechanism $P_{Z|Y}$ such that the overall mechanism $P_{Z|X}$ does not satisfy (ϵ, δ) -PML (which might come across as somewhat counter-intuitive). In response to this observation, we introduce another privacy guarantee, called (ϵ, δ) -EML (where EML stands for *event maximal leakage*) with $\epsilon \geq 0$ and $0 \leq \delta \leq 1$, which resembles (ϵ, δ) -PML but is closed under post-processing. Informally, a mechanism $P_{Y|X}$ satisfies (ϵ, δ) -EML if all post-processed outcomes of $P_{Y|X}$ with probability at least δ have their PML bounded by ϵ .

Next, in Section III-B, we study how different privacy guarantees change as a result of composing privacy mechanisms. More concretely, we are interested to find out what types of guarantees we can get for a mechanism $P_{YZ|X}$ which is obtained by adaptively composing two mechanisms $P_{Y|X}$ and $P_{Z|XY}$. Naturally, one can formulate different problems by making various assumptions about the involved mechanisms $P_{Y|X}$ and $P_{Z|XY}$. We present several such problem formulations and their corresponding results.

3) *Comparing PML with other privacy notions*: In Section IV, we study how pointwise maximal leakage relates to several other privacy/statistical notions, namely, max-information [38, 39], local differential privacy [17], local information privacy [16], local differential identifiability [9], mutual information, f -information [13] and total-variation privacy [14]. We derive bounds between the different notions and discuss their implications.

As a final note, in Section II-D we discuss a privacy framework called the *dynamic consumption of secrecy* [25] which, in the same spirit as our work, attempts to quantify the privacy leakage due to disclosing a single outcome of

¹By considering the expectation of privacy leakage, we retrieve the original definition of maximal leakage.

the random variable Y . Somewhat surprisingly, [25] argues that the privacy definition resulting from this dynamic view suffers from limitations that convince the authors against pursuing this line of research. In Section II-D, we discuss what these limitations are, and explain why they do not apply to pointwise maximal leakage.

B. Notation

In this work, we restrict our attention to finite random variables, therefore, all sets are assumed to be finite. We use uppercase letters to refer to random variables, e.g., X . Sets are represented by uppercase calligraphic letters, for example, the alphabet of X will be denoted by \mathcal{X} . Let $\mathcal{E} \subseteq \mathcal{X}$. We will use both $P_X(\mathcal{E})$ and $\mathbb{P}_{X \sim P_X}[\mathcal{E}]$ to describe the probability of an event \mathcal{E} according to distribution P_X . Similarly, we will use $\mathbb{E}_{X \sim P_X}[\cdot]$ to represent expectation with respect to P_X . The notation $\text{supp}(P_X) := \{x \in \mathcal{X} : P_X(x) > 0\}$ will be used to refer to the support set of distribution P_X . Given probability distributions P_X and Q_X over a set \mathcal{X} , we write $P_X \ll Q_X$ to imply that P_X is absolutely continuous with respect to Q_X .

Let n be a positive integer. We use $[n] := \{1, \dots, n\}$ to denote the set of all positive integers smaller than or equal to n . Suppose X is a random variable defined over an alphabet with cardinality $|\mathcal{X}| = n$, and Y is a random variable induced by a channel $P_{Y|X}$, such that $|\mathcal{Y}| = m$ is the cardinality of the alphabet of Y . Then, the channel $P_{Y|X} \in [0, 1]^{n \times m}$ is a row-stochastic matrix with elements $(P_{Y|X})_{ij} = P_{Y|X=x_i}(y_j)$ for $i \in [n]$ and $j \in [m]$. We say that a channel $P_{Y|X}$ is deterministic if it consists of only zeros and ones. Similarly, we say that an outcome y_j with $j \in [m]$ is deterministic if its corresponding column in the matrix $P_{Y|X}$ consists of only zeros and ones. Suppose the Markov chain $X - Y - Z$ holds. We write $P_{Z|X} = P_{Z|Y} \circ P_{Y|X}$ to denote marginalization over Y , that is, $P_{Z|X=x}(z) = \sum_{y \in \mathcal{Y}} P_{Z|Y=y}(z) P_{Y|X=x}(y)$ for $x \in \mathcal{X}$ and $z \in \mathcal{Z}$. Finally, we use $\log(\cdot)$ to denote the natural logarithm and $\mathbb{1}[\cdot]$ to denote the indicator function.

II. DEFINITION, INTERPRETATIONS AND PROPERTIES

A. Randomized Function View of Leakage

We begin by describing our threat model, which is a pointwise adaptation of the model described in [33]. Suppose X is a random variable defined over a finite alphabet \mathcal{X} . We will use X to represent some data containing sensitive information. Further, suppose Y is a random variable taking values in a finite alphabet \mathcal{Y} which is the output of a channel (i.e., kernel) $P_{Y|X}$ with input X . We will also refer to the channel $P_{Y|X}$ as a *privacy mechanism*. Consider an adversary who is interested in guessing the (realized) value of a possibly randomized function of X , called U , characterized by $P_{U|X}$. The adversary, who is computationally unbounded, observes an outcome $y \in \text{supp}(P_Y)$ (where P_Y is the output distribution induced by P_X and $P_{Y|X}$) and constructs a guess of U called \hat{U} according to a kernel $P_{\hat{U}|Y=y}$. The adversary is passive in the sense that she cannot affect the outcomes of the system, but can verify if her guess is correct. Furthermore, the adversary knows the joint distribution P_{UXY} , and therefore, can optimize her choice of guessing kernel $P_{\hat{U}|Y=y}$ to maximize her chances of correctly guessing U .

To measure the privacy leakage of a disclosed outcome y , the system designer considers the ratio of the probability of correctly guessing U having observed y , and the probability of correctly guessing U with no observations (in




Fig. 1. System model for the randomized function view of leakage: An adversary observes an outcome y of the channel $P_{Y|X}$, and tries to guess the value of a randomized function of X , denoted by U .

this case, the best guess is the most probable outcome according to P_U). Accordingly, we define the pointwise U -leakage of X as follows:

$$\ell_U(X \rightarrow y) := \log \frac{\sup_{P_{\hat{U}|Y=y}} \mathbb{P}[U = \hat{U} | Y = y]}{\max_{u \in \mathcal{U}} P_U(u)}, \quad (3)$$

where \mathcal{U} denotes the alphabet of the random variable U . As the system designer may not know what U the adversary is interested in, or different adversaries may be interested in guessing different U 's, we investigate the worst-case scenario by taking the supremum of (3) over all possible U 's. Considering this setup, we define pointwise maximal leakage (PML) denoted by $\ell_{P_{XY}}(X \rightarrow y)$ as follows.

Definition 1 (Pointwise maximal leakage): Let P_{XY} denote the joint distribution of X and Y , and suppose the Markov chain $U - X - Y - \hat{U}$ holds. The pointwise maximal leakage from X to $y \in \text{supp}(P_Y)$, $\ell_{P_{XY}}(X \rightarrow y)$, is defined as

$$\begin{aligned} \ell_{P_{XY}}(X \rightarrow y) &:= \sup_{P_{U|X}} \ell_U(X \rightarrow y) \\ &= \log \sup_{P_{U|X}} \frac{\sup_{P_{\hat{U}|Y=y}} \mathbb{P}[U = \hat{U} | Y = y]}{\max_{u \in \mathcal{U}} P_U(u)}. \end{aligned} \quad (4)$$

In the following result, we show that $\ell_{P_{XY}}(X \rightarrow y)$ can be written as a simple expression.

Theorem 1: Given a joint distribution P_{XY} over finite alphabets \mathcal{X} and \mathcal{Y} , the pointwise maximal leakage from X to $y \in \text{supp}(P_Y)$ is given by

$$\ell_{P_{XY}}(X \rightarrow y) = \log \max_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)}. \quad (5)$$

Proof: Fix an arbitrary random variable U . The numerator of (3) can be written as

$$\begin{aligned} \sup_{P_{\hat{U}|Y=y}} \mathbb{P}[U = \hat{U} | Y = y] &= \sup_{P_{\hat{U}|Y=y}} \sum_{u, \hat{u}} \mathbb{1}[u = \hat{u}] P_{U\hat{U}|Y=y}(u, \hat{u}) \\ &= \sup_{P_{\hat{U}|Y=y}} \sum_{u, \hat{u}} \mathbb{1}[u = \hat{u}] P_{U|Y=y}(u) P_{\hat{U}|Y=y}(\hat{u}) \\ &= \sup_{P_{\hat{U}|Y=y}} \sum_u P_{U|Y=y}(u) P_{\hat{U}|Y=y}(u) \\ &= \max_{u \in \mathcal{U}} P_{U|Y=y}(u), \end{aligned}$$

where the last equality follows from the fact that the optimal estimator $P_{\hat{U}|Y=y}^*$ in the above problem is the MAP estimator defined as

$$P_{\hat{U}|Y=y}^*(u) = \begin{cases} 1, & \text{for some } u \in \arg \max_{u \in \mathcal{U}} P_{U|Y=y}(u), \\ 0, & \text{otherwise.} \end{cases}$$

Thus, we can write

$$\begin{aligned} \exp(\ell_U(X \rightarrow y)) &= \frac{\sup_{P_{\hat{U}|Y=y}} \mathbb{P}[U = \hat{U} | Y = y]}{\max_{u' \in \mathcal{U}} P_U(u')} \\ &= \frac{\max_{u \in \mathcal{U}} P_{U|Y=y}(u)}{\max_{u' \in \mathcal{U}} P_U(u')} \\ &= \frac{\max_{u \in \mathcal{U}} \sum_{x \in \text{supp}(P_X)} P_{UX|Y=y}(u, x)}{\max_{u' \in \mathcal{U}} P_U(u')} \\ &= \frac{\max_{u \in \mathcal{U}} \sum_{x \in \text{supp}(P_X)} P_{X|Y=y}(x) P_{U|X=x}(u)}{\max_{u' \in \mathcal{U}} P_U(u')} \\ &= \max_{u \in \mathcal{U}} \sum_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)} P_{X|U=u}(x) \frac{P_U(u)}{\max_{u' \in \mathcal{U}} P_U(u')} \\ &\leq \max_{u \in \mathcal{U}} \sum_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)} P_{X|U=u}(x) \end{aligned} \tag{6a}$$

$$\leq \max_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)}. \tag{6b}$$

Taking the supremum over all U 's satisfying $U - X - Y$ we obtain

$$\ell_{P_{XY}}(X \rightarrow y) \leq \log \max_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)}. \tag{7}$$

To prove the reverse inequality, we construct a U achieving the bound in (7). Note that inequality (6b) holds with equality if there exists a $u^* \in \mathcal{U}$ such that

$$P_{X|U=u^*}(x) = \begin{cases} 1, & \text{for some } x \in \arg \max_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)}, \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

Furthermore, u^* will also satisfy (6a) with equality if it holds that

$$P_U(u^*) = \max_{u \in \mathcal{U}} P_U(u). \tag{9}$$

An example of U satisfying both of the above conditions can be obtained through the ‘‘shattering’’ channel $P_{U|X}$ defined in the proof of [33, Thm. 1]. Roughly speaking, the shattering channel breaks down each $x \in \text{supp}(P_X)$ with probability $P_X(x)$ into $k(x)$ corresponding elements with probability $\min_{x \in \text{supp}(P_X)} P_X(x)$, thus creating a random variable U with an (almost) uniform distribution. We recall the definition of the shattering channel $P_{U|X}$ for completeness.

Definition 2 (Shattering channel [33]): Let $p^* := \min_{x \in \text{supp}(P_X)} P_X(x)$. For each $x \in \text{supp}(P_X)$, let $k(x) := P_X(x)/p^*$, and let $\mathcal{U} = \bigcup_{x \in \text{supp}(P_X)} \{(x, 1), \dots, (x, \lceil k(x) \rceil)\}$, where $\lceil k(x) \rceil$ denotes the smallest integer greater than or equal to $k(x)$. For each $u = (i_u, j_u) \in \mathcal{U}$ and $x \in \text{supp}(P_X)$, the shattering channel $P_{U|X}$ is defined as

$$P_{U|X=x}((i_u, j_u)) = \begin{cases} \frac{p^*}{P_X(x)}, & \text{if } i_u = x, j_u = 1, \dots, \lfloor k(x) \rfloor, \\ 1 - \frac{\lfloor k(x) \rfloor p^*}{P_X(x)} & \text{if } i_u = x, j_u = \lceil k(x) \rceil, \\ 0, & \text{otherwise,} \end{cases}$$

where $\lfloor k(x) \rfloor$ denotes the largest integer smaller than or equal to $k(x)$.

The above channel induces the following joint distribution P_{UX} :

$$P_{UX}((i_u, j_u), x) = \begin{cases} p^*, & \text{if } i_u = x, j_u = 1, \dots, \lfloor k(x) \rfloor, \\ P_X(x) - \lfloor k(x) \rfloor p^* & \text{if } i_u = x, j_u = \lceil k(x) \rceil, \\ 0, & \text{otherwise,} \end{cases}$$

and P_U is obtained as

$$P_U(i_u, j_u) = \begin{cases} p^*, & \text{if } i_u = x, j_u = 1, \dots, \lfloor k(x) \rfloor, \\ P_X(x) - \lfloor k(x) \rfloor p^* & \text{if } i_u = x, j_u = \lceil k(x) \rceil, \end{cases}$$

for $(i_u, j_u) \in \mathcal{U}$. Clearly, each u is mapped to exactly one x , so condition (8) holds. Furthermore, each x corresponds to at least one u with probability $P_U(u) = \max_{u'} P_U(u') = p^*$, so condition (9) also holds. Consequently, the random variable U obtained through the shattering channel satisfies both (8) and (9), and attains the bound in (7). ■

Remark 1: Given $y \in \text{supp}(P_Y)$, pointwise maximal leakage can alternatively be written as

$$\begin{aligned} \ell_{P_{XY}}(X \rightarrow y) &= \max_{x \in \text{supp}(P_X)} \log \frac{P_{X|Y=y}(x)}{P_X(x)} \\ &= \max_{x \in \text{supp}(P_X)} \log \frac{P_{Y|X=x}(y)}{P_Y(y)} \\ &= D_\infty(P_{X|Y=y} \| P_X) \\ &= \max_{x \in \text{supp}(P_X)} i_{P_{XY}}(x; y), \end{aligned}$$

where $D_\infty(P_{X|Y=y} \| P_X)$ denotes the Rényi divergence of order infinity [40, 41] between the posterior $P_{X|Y=y}$ and the prior P_X , while $i_{P_{XY}}(x; y) := \log \frac{P_{XY}(x, y)}{P_X(x)P_Y(y)}$ denotes (the value of) the information density of the joint distribution P_{XY} at x and y . The above result is also related to a recent result by Kurri et al. [42], where a variational formula for Rényi divergence of order infinity is derived as the ratio of the expected gains in guessing a randomized function of X .

B. Gain Function View of Leakage

The threat model assumed in Theorem 1 considers a scenario in which an adversary is interested in guessing a possibly randomized function U of X . In this section, we argue that pointwise maximal leakage can be obtained




Fig. 2. System model for the gain function view of leakage: An adversary observes an outcome y of the channel $P_{Y|X}$, and tries to construct a guess \hat{X} of X in order to maximize a gain function g .

using an alternative threat model based on (a pointwise adaptation of) the g -leakage framework introduced in [26]. First, we describe this alternative threat model.

Suppose a passive and computationally unbounded adversary observes $y \in \text{supp}(P_Y)$, an outcome of the channel $P_{Y|X}$, and constructs a guess \hat{X} of X using a kernel $P_{\hat{X}|Y=y}$ in order to maximize her expected *gain*. The adversary selects her guess from a non-empty finite set $\hat{\mathcal{X}}$ (not necessarily equal to \mathcal{X}), and her gain is captured by a function g of the form $g : \mathcal{X} \times \hat{\mathcal{X}} \rightarrow \mathbb{R}^+$. In order to measure the amount of information leaking from y , the system designer considers the ratio of the expected adversarial gain having observed y , and the expected adversarial gain with no observations. Hence, we define the pointwise g -leakage of X as follows:

$$\ell_g(X \rightarrow y) := \log \frac{\sup_{P_{\hat{X}|Y=y}} \mathbb{E} [g(X, \hat{X}) | Y = y]}{\sup_{P_{\hat{X}}} \mathbb{E} [g(X, \hat{X})]}. \quad (10)$$

In Theorem 2, we will show that the randomized function view and the gain function view of leakage are equivalent in the sense that for every gain function g , there exists a corresponding randomized function of X , U_g , such that $\ell_g(X \rightarrow y) = \ell_{U_g}(X \rightarrow y)$, and conversely, for every randomized function of X , U , there exists a corresponding gain function g_U such that $\ell_U(X \rightarrow y) = \ell_{g_U}(X \rightarrow y)$. Before presenting this result, let us demonstrate through a few examples how gain functions can be used to model different problems.

Example 1 (The identity gain function [26, Def. 3.5]): The simplest type of gain function is the identity gain which models an adversary interested in guessing the secret X itself, who is only rewarded for correct guesses. Here, the guessing space of the adversary is $\hat{\mathcal{X}} = \mathcal{X}$, and her gain function is given by $g^{\text{identity}}(x, \hat{x}) = \mathbb{1}[x = \hat{x}]$. The g -leakage for the identity gain is

$$\ell_{g^{\text{identity}}}(X \rightarrow y) = \frac{\max_{x \in \mathcal{X}} P_{X|Y=y}(x)}{\max_{x \in \mathcal{X}} P_X(x)}, \quad (11)$$

which is equal to the dynamic min-entropy leakage defined in [25, Def. 3]. We will further discuss the identity gain and its associated g -leakage in Section II-D.

Example 2 (Membership/group privacy): Consider a centralized setting in which each $x \in \mathcal{X}$ represents a database whose rows constitute data collected from individuals, and the random variable X describes the random selection of a database according to some distribution P_X . Suppose an adversary is interested in guessing whether or not Alice's data is included in (the realization of) X , and is rewarded with a binary gain depending on whether or not

her guess is correct. We can model this problem as follows: Let $\mathcal{X}_1 = \{x \in \mathcal{X} : x \text{ contains Alice's data}\}$ denote the set of databases that contain Alice's data, and let $\mathcal{X}_0 = \mathcal{X} \setminus \mathcal{X}_1$ be the set of databases that do not contain Alice's data. An adversary who is interested in finding out Alice's membership makes a binary guess from $\hat{\mathcal{X}} = \{0, 1\}$, and is rewarded according to $g(x, x') = \mathbb{1}[\hat{x} = i]$ with $i \in \{0, 1\}$ and $x \in \mathcal{X}_i$.

More generally, suppose the adversary has a list of k individuals and is interested in guessing if any of their data is included in (the realization of) X . Further, suppose the adversary is rewarded based on how many correct guesses she makes. To model this problem, we (bi-)partition the set of all databases in k different ways, one for each individual on the list. Let $\mathcal{X}_{j1} = \{x \in \mathcal{X} : x \text{ contains the } j\text{-th individual's data}\}$ and $\mathcal{X}_{j0} = \mathcal{X} \setminus \mathcal{X}_{j1}$ for $j = 1, \dots, k$. Then, $\hat{\mathcal{X}} = \{0, 1\}^k$ is the guessing space of the adversary, and $g(x, \hat{x}) = \sum_{j=1}^k \mathbb{1}[\hat{x}_j = j_i]$ with $x \in \mathcal{X}_{j_i}$ and $i \in \{0, 1\}$ is her gain function. This example can be easily extended to model cases where different individuals signify different gains for the adversary.

As a side note, we should point out that in the membership privacy example above (or more generally, in our consideration of the centralized setting) we are *not* assuming that the adversary is *informed* [43] (an informed adversary knows all the entries in the database except for a single entry which may be Alice's). In our setup, we assume that the adversary knows the joint distribution P_{XY} (and the spaces \mathcal{X} and \mathcal{Y}), while any other side information should be explicitly modeled as such. The concept of an informed adversary was originally proposed as a model for a very powerful adversary. However, it has been argued that more side information does not necessarily make an adversary more effective [21]. For example, [21] provides three definitions of privacy against adversaries that either (i) know all the entries in a database except for a single entry, (ii) know all the attributes in a database except for a single attribute of a single entry, (iii) know all the bits in a database except for a single bit of a single entry. Then, it is shown that the privacy definition that seeks to limit the inference of the more knowledgeable adversary (i.e., option (iii)) may actually leak more sensitive information to the less knowledgeable adversaries.

In Section II-C, we will define a conditional form of pointwise maximal leakage that can be used to model an adversary who possesses some side information about the secret. There, we will see that side information can both increase and decrease the privacy leakage due to observing an outcome (more on this in Remark 5).

Example 3 (Multiple guesses (the k -tries gain function in [26])): Consider a side-channel setting in which X is a random variable representing a password and Y is a random variable denoting some information leaking about the password, for example, through the inter-keystroke delays. Suppose an adversary is allowed $k \geq 1$ attempts at guessing the password correctly before getting cut off from the system. Let \mathcal{X} be the set of all possible passwords and let $\hat{\mathcal{X}} = \{\hat{x} \subset \mathcal{X} : |\hat{x}| \leq k\}$ denote the collection of subsets of \mathcal{X} containing k or less passwords. Then, we can model the adversary's gain through the function $g(x, \hat{x}) = \mathbb{1}[x \in \hat{x}]$, where \hat{x} denotes the set of k or less attempts the adversary makes at guessing the true password x .

Example 4 (Metric spaces [26]): Suppose (\mathcal{X}, ρ) is a metric space, where \mathcal{X} is a finite set, and ρ is a metric on \mathcal{X} . Suppose the goal of the adversary is to construct a guess \hat{x} of x that minimizes $\rho(x, \hat{x})$. This scenario can be modeled by taking $\hat{\mathcal{X}} = \mathcal{X}$ and some non-negative gain function that is decreasing in $\rho(x, \hat{x})$, for example, $g(x, \hat{x}) = \exp(-\rho(x, \hat{x}))$. Many problems can be modeled as metric spaces. A simple example is in geo-location applications where the goal of an adversary may be to locate a user as accurately as possible based on partial or

noisy measurements.

We have now seen how a variety of adversarial objectives can be modeled using gain functions. In the following result, which is one of the main contributions of this paper, we show that the definition of g -leakage given in (10) is equivalent to the definition of U -leakage in (3). Thus, we unify two seemingly different ways of defining (pointwise) maximal leakage. The proof of the theorem is deferred to Appendix A.

Theorem 2: For all joint distributions P_{XY} over finite sets \mathcal{X} and \mathcal{Y} , the randomized function view and the gain function view of leakage are equivalent. That is, for every randomized function of X , described by the random variable U , there exists a space $\hat{\mathcal{X}}_U$ and a gain function $g_U : \mathcal{X} \times \hat{\mathcal{X}}_U \rightarrow \mathbb{R}^+$ such that $\ell_U(X \rightarrow y) = \ell_{g_U}(X \rightarrow y)$. Conversely, for every gain function $g : \mathcal{X} \times \hat{\mathcal{X}} \rightarrow \mathbb{R}^+$ there exists a randomized function of X , called U_g , such that $\ell_g(X \rightarrow y) = \ell_{U_g}(X \rightarrow y)$.

Note that while the above result establishes the equivalence of the gain-function view and the randomized-function view for pointwise leakages, it generalizes readily to the average-case leakages of [27] and [33]. Furthermore, we have the following corollary which provides an alternative definition of pointwise maximal leakage.

Corollary 1: Pointwise maximal leakage can be obtained as

$$\begin{aligned} \ell_{P_{XY}}(X \rightarrow y) &:= \sup_g \ell_g(X \rightarrow y) \\ &= \log \sup_g \frac{\sup_{P_{\hat{X}|Y=y}} \mathbb{E} \left[g(X, \hat{X}) \mid Y = y \right]}{\sup_{P_{\hat{X}}} \mathbb{E} \left[g(X, \hat{X}) \right]}, \end{aligned}$$

where the supremum is taken over all gain functions with non-negative and finite range.

Remark 2: Unlike privacy measures such as maximal leakage and (local) differential privacy that depend only on the mechanism $P_{Y|X}$, pointwise maximal leakage depends both on the mechanism $P_{Y|X}$ and the prior P_X , i.e., it is a property of the joint distribution P_{XY} . Thus, in the rest of this paper, we commonly assume that the prior P_X is arbitrary but fixed, and study pointwise maximal leakage as a function of the mechanism $P_{Y|X}$.

C. Properties

In this section, we recount several useful properties of $\ell_{P_{XY}}(X \rightarrow y)$. For instance, we discuss how pointwise maximal leakage composes over multiple outcomes, how it is affected by pre- and post-processing, and so on. Before we discuss these properties, let us first define a conditional form of pointwise maximal leakage which allows us to model adversaries who possess some side information about the secret X .

Definition 3 (Conditional pointwise maximal leakage): Let P_{XYZ} denote the joint distribution of random variables X , Y , and Z . Given $z \in \text{supp}(P_Z)$, where P_Z denotes the marginal distribution of Z , the conditional pointwise maximal leakage from X to $y \in \text{supp}(P_{Y|Z=z})$ is defined as

$$\ell_{P_{XY|Z}}(X \rightarrow y \mid z) := \log \sup_{P_{U|X}} \frac{\sup_{P_{\hat{U}|Y=y, Z=z}} \mathbb{P} \left[U = \hat{U} \mid Y = y, Z = z \right]}{\sup_{P_{\hat{U}|Z=z}} \mathbb{P} \left[U = \hat{U} \mid Z = z \right]}.$$

To obtain a simpler expression for $\ell_{P_{XY|Z}}(X \rightarrow y | z)$, we condition all the distributions in the proof of Theorem 1 on $Z = z$ and get

$$\begin{aligned} \ell_{P_{XY|Z}}(X \rightarrow y | z) &= \log \max_{x \in \text{supp}(P_{X|Z=z})} \frac{P_{X|Y=y, Z=z}(x)}{P_{X|Z=z}(x)} \\ &= \log \max_{x \in \text{supp}(P_{X|Z=z})} \frac{P_{Y|X=x, Z=z}(y)}{P_{Y|Z=z}(y)} \\ &= D_\infty(P_{X|Y=y, Z=z} \| P_{X|Z=z}) \\ &= \max_{x \in \text{supp}(P_{X|Z=z})} i_{P_{XY|Z}}(x; y | z), \end{aligned}$$

where

$$i_{P_{XY|Z}}(x; y | z) := \log \frac{P_{XY|Z=z}(x, y)}{P_{X|Z=z}(x)P_{Y|Z=z}(y)},$$

denotes the (value of the) *conditional information density*. In the remainder of this paper, when the joint distribution used to calculate pointwise maximal leakage or information density is clear from the context, we do not specify it as a subscript.

The following lemma provides several useful properties of pointwise maximal leakage. The proof of the lemma is provided in Appendix B.

Lemma 1: Pointwise maximal leakage satisfies the following properties:

- 1) (Upper/lower bounds). Given an arbitrary but fixed prior P_X , for all mechanisms $P_{Y|X}$ and all $y \in \text{supp}(P_Y)$ it holds that

$$0 \leq \ell(X \rightarrow y) \leq -\log \left(\min_{x \in \text{supp}(P_X)} P_X(x) \right),$$

where the left-hand side inequality holds with equality if and only if $P_{Y|X=x}(y) = P_{Y|X=x'}(y)$ for all $x, x' \in \text{supp}(P_X)$, and the right-hand side inequality holds with equality if and only if $P_{X|Y=y}(x^*) = 1$ for some $x^* \in \arg \min_{x \in \text{supp}(P_X)} P_X(x)$.

- 2) (Independence/deterministic mappings). If X and Y are independent random variables, then $\ell(X \rightarrow y) = 0$ for all $y \in \text{supp}(P_Y)$. If Y is the output of a deterministic mapping with input X , then $\ell(X \rightarrow y) = -\log P_Y(y)$.
- 3) (Pre-processing). Suppose the Markov chain $X - Y - Z$ holds, where Y represents some pre-processing of the secret X , and Z denotes the observable outcome of a channel $P_{Z|Y}$ with input Y . For all $z \in \text{supp}(P_Z)$ we have

$$\ell(X \rightarrow z) \leq \ell(Y \rightarrow z),$$

with equality if $i(y; z) = i(y'; z) = \max_{y \in \text{supp}(P_Y)} i(y; z)$ for all $y, y' \in \text{supp}(P_{Y|X=x^*})$, where $x^* \in \arg \max_{x \in \text{supp}(P_X)} i(x; z)$.

- 4) (Post-processing). Suppose the Markov chain $X - Y - Z$ holds, where Y represents the observable outcome of a channel with input X , and Z denotes some post-processing of Y . For all $z \in \text{supp}(P_Z)$ we have

$$\ell(X \rightarrow z) \leq \max_{y \in \text{supp}(P_Y)} \ell(X \rightarrow y),$$

where the inequality holds with equality if either of the following conditions are satisfied:

- a) X and Y are independent, or
 - b) there exists $y_z \in \text{supp}(P_Y)$ such that $P_{Y|Z=z}(y_z) = 1$ and $\ell(X \rightarrow y_z) = \max_{y \in \text{supp}(P_Y)} \ell(X \rightarrow y)$.
- 5) (Conditionally-independent side information). Suppose the Markov chain $Z - X - Y$ holds, where Z represents some side information about X . For all $z \in \text{supp}(P_Z)$ and $y \in \text{supp}(P_{Y|Z=z})$ we have

$$\ell(X \rightarrow y | z) = \max_{x \in \text{supp}(P_X)} i(x; y | z) = \ell(X \rightarrow y) - i(y; z).$$

- 6) (Composition). Given a prior P_X and a mechanism $P_{Y|Z|X}$, for all $(y, z) \in \text{supp}(P_{YZ})$ it holds that

$$\begin{aligned} \ell(X \rightarrow y, z) &= \max_{x \in \text{supp}(P_X)} i(x; y, z) \\ &\leq \ell(X \rightarrow z) + \ell(X \rightarrow y | z), \end{aligned}$$

with equality if and only if the two sets $\arg \max_{x \in \text{supp}(P_X)} i(x; y | z)$ and $\arg \max_{x \in \text{supp}(P_X)} i(x; z)$ have non-empty intersection.

Remark 3: The second property in the above lemma describes the privacy leakage of deterministic mechanisms. Surprisingly, not all deterministic outcomes leak the same amount of information, and outcomes with lower probabilities have higher leakage. This is because pointwise maximal leakage is a relative privacy measure in which, roughly speaking, the information leaked to an adversary scales depending on how consistent the observed outcome is with the adversary's prior beliefs (captured by the joint distribution P_{XY}). As such, deterministic outcomes with smaller probabilities leak more information since an adversary would be "more surprised" by observing them.

Remark 4: Concerning the post-processing property stated above, one may hope for the stronger statement $\ell(X \rightarrow z) \leq \ell(X \rightarrow y)$ for all $y \in \text{supp}(P_Y)$. To see why this statement is not valid, consider the situation where $Z = Y$ (or Z is a deterministic mapping of Y). In this case, the best bound we can have is indeed $\ell(X \rightarrow y) \leq \max_y \ell(X \rightarrow y)$.

Remark 5: In general, side information can both increase and decrease privacy leakage. As an example, suppose we have three binary random variables X, Y, Z where $\mathcal{X} = \mathcal{Y} = \mathcal{Z} = \{0, 1\}$. Assume X is uniformly distributed, and the joint distribution P_{XYZ} is described by the following channels:

$$P_{Z|X} = \begin{bmatrix} \frac{2}{5} & \frac{3}{5} \\ \frac{3}{5} & \frac{2}{5} \end{bmatrix}, \quad P_{Y|XZ} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{3} & \frac{2}{3} \\ \frac{2}{3} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix},$$

where the rows in the $P_{Y|XZ}$ matrix correspond to (x, z) equals $(0, 0)$, $(0, 1)$, $(1, 0)$, and $(1, 1)$ from top to bottom. Then, it can be verified through simple calculations that $\ell(X \rightarrow y = 0 | z = 0) = \log \frac{10}{9}$, $\ell(X \rightarrow y = 0 | z = 1) = \log \frac{5}{4}$ and $\ell(X \rightarrow y = 0) = \log \frac{6}{5}$. Therefore,

$$\ell(X \rightarrow y = 0 | z = 0) < \ell(X \rightarrow y = 0) < \ell(X \rightarrow y = 0 | z = 1).$$

D. Dynamic Consumption of Secrecy

In the last part of this section, we discuss a notion of privacy introduced in [25, Section 2.2] that, similar to our work, aims to measure the privacy leakage associated with individual observations. In [25], the *dynamic min-entropy leakage* of an outcome $y \in \text{supp}(P_Y)$ is defined as:

$$\ell^{\text{dynamic}}(X \rightarrow y) := \log \frac{\max_{x \in \mathcal{X}} P_{X|Y=y}(x)}{\max_{x \in \mathcal{X}} P_X(x)},$$

which is equal to the g -leakage associated with the identity gain given in (11). That is, the dynamic leakage is derived under the assumption that the adversary is trying to guess the secret X itself, but does not consider other gains that an adversary may be interested in.

The authors of [25] withdraw from further developing the idea of measuring the pointwise privacy leakage based on $\ell^{\text{dynamic}}(X \rightarrow y)$ as they believe the above privacy measure suffers from two drawbacks. First, they argue that the above definition cannot be axiomatically justified as it is shown that $\ell^{\text{dynamic}}(X \rightarrow y)$ may be negative (see [25, Example 4]), that is, the adversary's certainty about the secret may actually *decrease* by observing an outcome y . Second, they believe that dynamic policy enforcement based on individual outcomes (for example, discarding high-leakage outcomes) may reveal information about the secret X .

Note that the first issue mentioned above does not apply to $\ell(X \rightarrow y)$. It is easy to see that $\ell^{\text{dynamic}}(X \rightarrow y) \leq \ell(X \rightarrow y)$, and we have shown in Lemma 1 that $\ell(X \rightarrow y) \geq 0$, which implies that in our current setup, observations can never decrease certainty about a secret X . This is because $\ell(X \rightarrow y)$ is defined by considering all possible gain functions an adversary may be interested in, while $\ell^{\text{dynamic}}(X \rightarrow y)$ is defined only for the identity gain of Example 1.

Furthermore, while it is true that some policies defined based on individual outcomes, such as discarding high-leakage outcomes, may reveal information about the secret, we believe that this is not sufficient reason for abandoning the subject area altogether. In fact, contrary to [25], we believe that effective policy enforcement depends crucially on the ability to quantify the information leaking from individual outcomes as this allows us to treat privacy leakage as a *random variable*. Viewing privacy leakage as a random variable, we have the flexibility to define different types of privacy guarantees by specifying requirements on the statistical properties of privacy leakage. The resulting framework is then versatile enough to be applied to a wide range of problems. We develop this idea in the next section.

III. PRIVACY GUARANTEES

In Theorem 1, we showed that $\ell(X \rightarrow y)$ can be written as a function of y . Since Y is a random variable distributed according to P_Y , this in turn allows us to define a random variable $\ell(X \rightarrow Y)$ with a distribution induced by P_Y . From this point of view, a privacy guarantee is essentially a requirement we impose on some statistical property of $\ell(X \rightarrow Y)$; thus, we have the flexibility to define different types of privacy guarantees depending on how strict privacy requirements we need to meet. This section contains several examples of such guarantees: The *almost-sure* guarantee, which bounds privacy leakage with probability one, the *tail-bound* guarantee,

which bounds leakage with high probability, and the *average-case* guarantee, which bounds maximal leakage. We start by defining the almost-sure guarantee.

Definition 4 (Almost-sure guarantee): Given an arbitrary but fixed prior P_X , we say that a privacy mechanism $P_{Y|X}$ satisfies ϵ -PML with $\epsilon \geq 0$ if

$$\mathbb{P}_{Y \sim P_Y}[\ell(X \rightarrow Y) \leq \epsilon] = 1. \quad (12)$$

As we are assuming that the random variables X and Y are finite, the above condition can also be expressed as

$$\max_{x \in \text{supp}(P_X)} \max_{y \in \text{supp}(P_Y)} i_{P_{XY}}(x; y) \leq \epsilon. \quad (13)$$

The expression $\max_{x,y} i_{P_{XY}}(x; y)$ coincides with the definition of maximal realizable leakage [33, Def. 8] and also max-information [38]. Moreover, expression (13) is also related to the notion of information privacy leakage [17, 16]. We will discuss the relationship between our privacy definitions and pre-existing notions from the literature in the next section.

The following lemma establishes some basic facts about ϵ -PML guarantees.

Lemma 2: Given an arbitrary but fixed prior P_X , we have:

- 1) All privacy mechanisms $P_{Y|X}$ satisfy ϵ_{\max} -PML, where $\epsilon_{\max} := -\log \min_{x \in \text{supp}(P_X)} P_X(x)$ and $\epsilon_{\max} \geq \log 2$. Furthermore, we have

$$\inf\{\epsilon \geq 0: \mathbb{P}_{Y \sim P_Y}[\ell(X \rightarrow Y) \leq \epsilon] = 1\} = \epsilon_{\max},$$

if and only if there exists $y \in \text{supp}(P_Y)$ such that $P_{X|Y=y}(x^*) = 1$, or equivalently, $P_{Y|X=x}(y) = 0$ for all $x \neq x^*$, where $x^* \in \arg \min_x P_X(x)$.

- 2) A privacy mechanism $P_{Y|X}$ satisfies ϵ -PML with $\epsilon = 0$ if and only if X and Y are independent random variables.

Proof:

- 1) For all $P_{Y|X}$ and all $y \in \text{supp}(P_Y)$ we have

$$\begin{aligned} \log \max_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)} &\leq \log \max_{x \in \text{supp}(P_X)} \frac{1}{P_X(x)} \\ &= \log \frac{1}{\min_{x \in \text{supp}(P_X)} P_X(x)}. \end{aligned}$$

Note that $\min_{x \in \text{supp}(P_X)} P_X(x) \leq \frac{1}{2}$ which implies that $\epsilon_{\max} \geq \log 2$. The second half of the statement is clear from the above inequality.

- 2) If X and Y are independent, then $P_{X|Y=y}(x) = P_X(x)$ for all x, y , thus the mechanism $P_{Y|X}$ satisfies ϵ -PML with $\epsilon = 0$. Conversely, if $P_{Y|X}$ satisfies ϵ -PML with $\epsilon = 0$ this implies that $P_{X|Y=y}(x) = P_X(x)$ for all x, y which means that X and Y are independent. ■

In order for an ϵ -PML guarantee to hold, all $y \in \text{supp}(P_Y)$ must satisfy $\ell(X \rightarrow y) \leq \epsilon$. As this condition may prove to be too restrictive in practice, in what follows we define two possible relaxations of the almost-sure

guarantee: We either bound the privacy leakage by ϵ with high probability, or we bound the expected leakage (i.e., maximal leakage) by ϵ .

Definition 5 (Tail-bound guarantee): Given an arbitrary but fixed prior P_X , we say that a mechanism $P_{Y|X}$ satisfies (ϵ, δ) -PML with $\epsilon \geq 0$ and $0 \leq \delta \leq 1$ if

$$\mathbb{P}_{Y \sim P_Y}[\ell(X \rightarrow Y) \leq \epsilon] \geq 1 - \delta. \quad (14)$$

Clearly, ϵ -PML and $(\epsilon, 0)$ -PML are equivalent. Also, note that given an arbitrary but fixed prior P_X , if a channel $P_{Y|X}$ satisfies (ϵ, δ) -PML, then it also satisfies (ϵ', δ') -PML for all $\epsilon \leq \epsilon'$ and all $\delta \leq \delta' \leq 1$.

Definition 6 (Average-case guarantee): Given an arbitrary but fixed prior P_X , we say that the expected privacy leakage of a mechanism $P_{Y|X}$ is bounded by $\epsilon \geq 0$ if

$$\mathbb{E}_{Y \sim P_Y} \left[\exp(\ell(X \rightarrow Y)) \right] \leq e^\epsilon,$$

or equivalently,

$$\mathcal{L}(P_{Y|X}) \leq \epsilon,$$

where $\mathcal{L}(P_{Y|X})$ denotes maximal leakage as defined in [33, Thm. 1].

Note that here we denote maximal leakage by $\mathcal{L}(P_{Y|X})$ instead of $\mathcal{L}(X \rightarrow Y)$ used in [33] to emphasize that maximal leakage is a property of the channel $P_{Y|X}$ and does not depend on the prior P_X .²

The following lemma shows how almost-sure and tail-bound guarantees translate into maximal leakage guarantees.

Lemma 3: Given an arbitrary but fixed prior P_X , it holds that

- 1) ϵ -PML implies $\mathcal{L}(P_{Y|X}) \leq \epsilon$.
- 2) (ϵ, δ) -PML implies

$$\mathcal{L}(P_{Y|X}) \leq \log \left(\exp(\epsilon) + \delta \exp(\epsilon_{\max}) \right),$$

where $\epsilon_{\max} := -\log \min_{x \in \text{supp}(P_X)} P_X(x)$.

Proof:

- 1) By definition,

$$\begin{aligned} \mathcal{L}(P_{Y|X}) &= \log \sum_y P_Y(y) \exp(\ell(X \rightarrow y)) \\ &\leq \log \left(\exp(\epsilon) \sum_y P_Y(y) \right) = \epsilon. \end{aligned}$$

²Technically, maximal leakage depends on the support set of the prior P_X , but we can without loss of generality assume that P_X has full support.

2) We define the events $E_1(\epsilon) = \{y \in \text{supp}(P_Y) : \ell(X \rightarrow y) \leq \epsilon\}$ and $E_2(\epsilon) = \{y \in \text{supp}(P_Y) : \ell(X \rightarrow y) > \epsilon\}$. Then, $E_1(\epsilon)$ and $E_2(\epsilon)$ are disjoint and $\text{supp}(P_Y) = E_1(\epsilon) \cup E_2(\epsilon)$, for all $\epsilon \geq 0$. Clearly, (ϵ, δ) -PML implies that $\mathbb{P}[E_2(\epsilon)] \leq \delta$. Therefore, we have

$$\begin{aligned} \mathcal{L}(P_{Y|X}) &= \log \sum_y P_Y(y) \exp(\ell(X \rightarrow y)) \\ &= \log \left(\sum_{y \in E_1(\epsilon)} P_Y(y) \exp(\ell(X \rightarrow y)) + \sum_{y \in E_2(\epsilon)} P_Y(y) \exp(\ell(X \rightarrow y)) \right) \\ &\leq \log \left(\exp(\epsilon) \mathbb{P}[E_1(\epsilon)] + \exp(\epsilon_{\max}) \mathbb{P}[E_2(\epsilon)] \right) \\ &\leq \log \left(\exp(\epsilon) + \delta \exp(\epsilon_{\max}) \right). \end{aligned}$$

■

A. Data-processing Properties

Data-processing inequalities are often used while analyzing the end-to-end information leakage in larger systems. While the properties presented in Lemma 1 allow us to assess pointwise maximal leakage for the outcomes of a pre- or post-processed random variable, it is also of practical benefit to understand how different privacy guarantees are affected by pre- and post-processing. This type of characterization is useful when we do not have access to the distribution of the leakage over the outcomes, but know that a privacy mechanism satisfies a certain privacy guarantee.

What we are specifically interested in is to understand whether or not different privacy guarantees are *closed*³ under pre- and post-processing (in [44], a privacy guarantee that is closed under pre-processing is said to satisfy the *linkage* inequality). Suppose the channel $P_{Y|X}$ satisfies some privacy guarantee, say, ϵ -PML. If the ϵ -PML guarantee is closed under post-processing, then we can rest assured that for all post-processing channels $P_{Z|Y}$, the overall channel $P_{Z|X}$ also satisfies ϵ -PML, that is, there exists no channel $P_{Z|Y}$ that an adversary could use to undermine the original guarantee. Similarly, if the ϵ -PML guarantee is closed under pre-processing, then all (randomized) functions of X would be at least as well-protected as X .

The following lemma collects the data-processing properties satisfied by the privacy guarantees defined above. Part 4 of the result concerning maximal leakage was shown in [33, Lemma 1], and we re-state it here for completeness.

Proposition 1: Suppose the three random variables X , Y , and Z form the Markov chain $X - Y - Z$ and that the prior P_X is arbitrary but fixed. Given $\epsilon \geq 0$ and $0 \leq \delta \leq 1$, we have

- 1) If $P_{Z|Y}$ satisfies ϵ -PML, then $P_{Z|X}$ also satisfies ϵ -PML.
- 2) If $P_{Y|X}$ satisfies ϵ -PML, then $P_{Z|X}$ also satisfies ϵ -PML.
- 3) If $P_{Z|Y}$ satisfies (ϵ, δ) -PML, then $P_{Z|X}$ also satisfies (ϵ, δ) -PML.

³Suppose the channel $P_{Y|X}$ satisfies Property A . Given a Markov chain $X - Y - Z$, we say that Property A is *closed under post-processing* if the fact that $P_{Y|X}$ satisfies Property A implies that $P_{Z|X}$ also satisfies Property A . Closedness under pre-processing is defined similarly.

$$4) \mathcal{L}(P_{Z|X}) \leq \min\{\mathcal{L}(P_{Y|X}), \mathcal{L}(P_{Z|Y})\}.$$

Proof:

- 1) By the pre-processing property of Lemma 1, if $P_{Z|Y}$ satisfies ϵ -PML then for all $z \in \text{supp}(P_Z)$ we have $\ell(Y \rightarrow z) \leq \epsilon$. Hence,

$$\max_{z \in \text{supp}(P_Z)} \ell(X \rightarrow z) \leq \max_{z \in \text{supp}(P_Z)} \ell(Y \rightarrow z) \leq \epsilon,$$

and $P_{Z|X}$ satisfies ϵ -PML.

- 2) By the post-processing property of Lemma 1,

$$\max_{z \in \text{supp}(P_Z)} \ell(X \rightarrow z) \leq \max_{y \in \text{supp}(P_Y)} \ell(X \rightarrow y) \leq \epsilon,$$

so $P_{Z|X}$ satisfies ϵ -PML.

- 3) Similarly to the above, the pre-processing property of Lemma 1 yields

$$\mathbb{P}_{Z \sim P_Z} [\ell(X \rightarrow Z) > \epsilon] \leq \mathbb{P}_{Z \sim P_Z} [\ell(Y \rightarrow Z) > \epsilon] \leq \delta,$$

hence, $P_{Z|X}$ satisfies (ϵ, δ) -PML. ■

Conspicuous by its absence in the above result is the post-processing property for the (ϵ, δ) -PML privacy guarantee. It turns out that, in general, (ϵ, δ) -PML is not closed under post-processing. To understand why, let us consider the following example.

Example 5: Suppose X is a uniformly distributed random variable defined over an alphabet with four elements, and that the Markov chain $X - Y - Z$ holds. Suppose the channels $P_{Y|X}$ and $P_{Z|Y}$ are defined as

$$P_{Y|X} = \begin{bmatrix} 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} \end{bmatrix}, \quad P_{Z|Y} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

It can be easily verified that $\ell(X \rightarrow y_1) = \ell(X \rightarrow y_2) = \log 4$, and $\ell(X \rightarrow y_3) = \ell(X \rightarrow y_4) = \log \frac{6}{5}$. Since $P_Y(y_1) = P_Y(y_2) = \frac{1}{12}$, $P_{Y|X}$ satisfies (ϵ_1, δ_1) -PML with $\epsilon_1 = \log \frac{6}{5}$ and $\delta_1 = \frac{1}{6}$. On the other hand, one may also verify that $P_Z(z_1) = P_Z(z_2) = \frac{1}{2}$ and $\ell(X \rightarrow z_1) = \ell(X \rightarrow z_2) = \log \frac{4}{3}$. Hence, $P_{Z|X}$ does not satisfy (ϵ_1, δ_1) -PML; instead, it satisfies ϵ_2 -PML with $\epsilon_2 = \log \frac{4}{3} > \epsilon_1$ (and $\delta_2 = 0$). Note that the outcome z_1 is equivalent to the event $\{y_1, y_3\}$, z_2 is equivalent to the event $\{y_2, y_4\}$, and both outcomes have probability greater than δ_1 .

Informally speaking, when we say that a mechanism $P_{Y|X}$ satisfies (ϵ, δ) -PML this implies that $\text{supp}(P_Y)$ can be partitioned into two sets: a set of “good” y ’s with probability at least $1 - \delta$ whose members satisfy $\ell(X \rightarrow y) \leq \epsilon$, and a set of “bad” y ’s with probability at most δ with $\ell(X \rightarrow y) > \epsilon$ (see Definition 5). However, through a post-processing channel $P_{Z|Y}$, we may define new outcomes as a combination of the members of the good and bad sets of y (as in Example 5). As a result, the probability of the set whose members satisfy $\ell(X \rightarrow z) > \epsilon$ (that is, the set of “bad” z ’s) may no longer be bounded by δ . Also, note that while in Example 5 we have $\epsilon_2 > \epsilon_1$

and $\delta_2 < \delta_1$, this need not always be the case; one may come up with examples where both ϵ and δ increase by post-processing.

Remark 6: Interestingly, a similar behavior has been observed in differential privacy. Specifically, it has been shown that *probabilistic DP*, that is, a type of privacy guarantee where we require pure DP to hold with probability at least $1 - \delta$ [45, 46], is not closed under post-processing [47, 46]. Differential privacy resolves this problem by introducing *approximate DP* (i.e., (ϵ, δ) -DP defined based on an *additive* parameter δ [48]) which is closed under post-processing. Note that approximate DP is a strictly weaker guarantee compared to probabilistic DP in the sense that probabilistic DP implies approximate DP but the reverse direction does not necessarily hold [46]. In our current work with PML, we take a different approach to solving this issue and come up with a new privacy guarantee that maintains its probabilistic flavor.

Now, we define a new probabilistic privacy guarantee that is similar to (ϵ, δ) -PML, but is closed under post-processing. Drawing on Example 5, our new definition ensures that all post-processed outcomes with probability at least δ have their PML bounded by ϵ . We provide two alternative formulations of our new privacy guarantee: The first one in Definition 7 describes a somewhat technical condition, so we re-state it in a more intuitive form in Definition 11.

Definition 7: Given an arbitrary but fixed prior P_X , we say that a privacy mechanism $P_{Y|X}$ satisfies (ϵ, δ) -closedness with $\epsilon \geq 0$ and $0 \leq \delta \leq 1$, if for all post-processing channels $P_{Z|Y}$, $P_Z(z) \geq \delta$ implies $\ell_{P_{XZ}}(X \rightarrow z) \leq \epsilon$ where $z \in \text{supp}(P_Z)$.

Based on Definition 7, to check whether or not a certain mechanism $P_{Y|X}$ satisfies the desired closedness property, one needs to examine all possible post-processing channels $P_{Z|Y}$. This raises the question of whether it is possible to come up with a definition equivalent to Definition 7, which can be stated as a property of the channel $P_{Y|X}$ itself. In what follows, we show that this is indeed possible, but we need a few other ingredients before we are ready to state this alternative definition. First, we recall two concepts from [49].

Definition 8 (Similar outcomes [49]): Given a channel $P_{Y|X}$, we say that the outcomes $y, y' \in \text{supp}(P_Y)$ are similar if their corresponding columns in the matrix of $P_{Y|X}$ are scalar multiples of each other, or equivalently, if $P_{X|Y=y}(x) = P_{X|Y=y'}(x)$ for all $x \in \text{supp}(P_X)$.

Remark 7: Note that if the outcomes $y, y' \in \text{supp}(P_Y)$ are similar, then $i(x; y) = i(x; y')$ for all $x \in \text{supp}(P_X)$ and $\ell(X \rightarrow y) = \ell(X \rightarrow y')$.

Definition 9 (Reduced channel [49, Def. 3]): Given a channel $P_{Y|X}$, its reduced channel denoted by $P_{Y_r|X}$ is formed by removing all-zero columns from $P_{Y|X}$, and merging (i.e., adding) the columns corresponding to similar outcomes.

Let $P_{Y_r|X}$ denote the reduced channel of the mechanism $P_{Y|X}$. We can define an equivalence relation $P_{\bar{Y}|X} \sim P_{Y|X}$ if $P_{\bar{Y}|X}$ has $P_{Y_r|X}$ as its reduced channel. Then, the equivalence class of $P_{Y|X}$, denoted by $\mathcal{C}(P_{Y|X})$, is the collection of all mechanisms whose reduced channel is $P_{Y_r|X}$. Suppose $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$. We will use \bar{Y} to denote the (output) random variable induced by the channel $P_{\bar{Y}|X}$ whose alphabet is represented by $\bar{\mathcal{Y}}$, and whose marginal distribution is denoted by $P_{\bar{Y}}$.

Similar outcomes lead to the same posterior distribution, information density, and PML. Thus, the channels in a

class $\mathcal{C}(P_{Y|X})$ behave identically with respect to information measures that are defined based on the information density, such as mutual information and maximal leakage. In the following, we show that if $P_{Y|X}$ satisfies (ϵ, δ) -closedness, then all $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$ also satisfy (ϵ, δ) -closedness.

Proposition 2: Given an arbitrary but fixed prior P_X and an (ϵ, δ) pair with $\epsilon \geq 0$ and $0 \leq \delta \leq 1$, if a mechanism $P_{Y|X}$ satisfies (ϵ, δ) -closedness then all $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$ also satisfy (ϵ, δ) -closedness.

Proof: Let the function $f : \mathcal{C}(P_{Y|X}) \times [0, 1] \rightarrow \mathbb{R}^+$ be defined as

$$f(P_{\bar{Y}|X}, \delta) = \sup_{P_{Z|\bar{Y}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \ell_{P_{XZ}}(X \rightarrow z),$$

that is, f represents the largest PML over all outcomes z of all post-processing channels with probability at least δ . We argue that $f(\cdot, \delta)$ is constant on $\mathcal{C}(P_{Y|X})$ for all $0 \leq \delta \leq 1$. To see this, fix an arbitrary $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$ and note that the Markov chain $X - \bar{Y} - Y_r$ holds, where Y_r denotes the random variable induced by the reduced channel $P_{Y_r|X}$. By definition, $I(X; \bar{Y}) = I(X; Y_r)$, therefore Y_r is a sufficient statistic of \bar{Y} for X , and the Markov chain $X - Y_r - \bar{Y}$ also holds.⁴ Now, we write

$$\begin{aligned} f(P_{\bar{Y}|X}, \delta) &= \sup_{P_{Z|\bar{Y}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \ell_{P_{XZ}}(X \rightarrow z) \\ &= \sup_{\substack{P_{Z|Y_r}: \\ P_{Z|Y_r} = P_{Z|\bar{Y}} \circ P_{\bar{Y}|Y_r}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \ell_{P_{XZ}}(X \rightarrow z) \\ &\leq \sup_{P_{Z|Y_r}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \ell_{P_{XZ}}(X \rightarrow z) \\ &= f(P_{Y_r|X}, \delta). \end{aligned}$$

Reversing the role of \bar{Y} and Y_r , it can also be established that $f(P_{Y_r|X}, \delta) \leq f(P_{\bar{Y}|X}, \delta)$; hence, we obtain $f(P_{Y_r|X}, \delta) = f(P_{\bar{Y}|X}, \delta)$ for all $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$ and $0 \leq \delta \leq 1$. Finally, if $P_{Y|X}$ satisfies (ϵ, δ) -closedness then $f(P_{Y|X}, \delta) \leq \epsilon$, which implies that $\sup_{P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})} f(P_{\bar{Y}|X}, \delta) \leq \epsilon$. ■

One last concept that we need to introduce is the notion of the maximal leakage associated with arbitrary events (that is, subsets) of $\text{supp}(P_Y)$. We will call this new form of leakage *event maximal leakage* (EML), which is defined fairly similarly to PML. That said, the real benefit of EML is in that it allows us to come up with an alternative formulation of Definition 7.

Definition 10 (Event maximal leakage (EML)): Suppose the random variables X and Y are described by the joint distribution P_{XY} defined over the finite sets \mathcal{X} and \mathcal{Y} . Given an event $\mathcal{E} \subseteq \text{supp}(P_Y)$, the maximal leakage from X to \mathcal{E} is defined as

$$\begin{aligned} \ell_{P_{XY}}(X \rightarrow \mathcal{E}) &:= \log \max_{x \in \text{supp}(P_X)} \frac{P_{Y|X=x}(\mathcal{E})}{P_Y(\mathcal{E})} = \log \max_{x \in \text{supp}(P_X)} \frac{\sum_{y \in \mathcal{E}} P_{Y|X=x}(y)}{\sum_{y' \in \mathcal{E}} P_Y(y')} \\ &= \log \max_{x \in \text{supp}(P_X)} \sum_{y \in \mathcal{E}} \frac{P_Y(y)}{\sum_{y' \in \mathcal{E}} P_Y(y')} \exp(i_{P_{XY}}(x; y)). \end{aligned}$$

⁴We thank the anonymous reviewer of our paper for suggesting to use the concept of sufficient statistic to shorten the proof.

Having defined the maximal leakage associated with arbitrary subsets of $\text{supp}(P_Y)$, we are now ready to provide an alternative form of Definition 7.

Definition 11: Given an arbitrary but fixed prior P_X , we say that a privacy mechanism $P_{Y|X}$ satisfies (ϵ, δ) -EML with $\epsilon \geq 0$ and $0 \leq \delta \leq 1$ if for all $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$ and all events $\mathcal{E} \subseteq \text{supp}(P_{\bar{Y}})$, $P_{\bar{Y}}(\mathcal{E}) \geq \delta$ implies $\ell_{P_{X\bar{Y}}}(X \rightarrow \mathcal{E}) \leq \epsilon$.

Clearly, $(\epsilon, 0)$ -EML and ϵ -PML are equivalent. Furthermore, given an arbitrary but fixed prior P_X , if a channel $P_{Y|X}$ satisfies (ϵ, δ) -EML, then it also satisfies (ϵ', δ') -EML for all $\epsilon \leq \epsilon'$, and all $\delta \leq \delta' \leq 1$.

Next, we show that a privacy mechanism $P_{Y|X}$ satisfies (ϵ, δ) -closedness if and only if it satisfies (ϵ, δ) -EML. That is, Definitions 7 and 11 are equivalent.

Theorem 3: Given an arbitrary but fixed prior P_X , and a pair (ϵ, δ) with $\epsilon \geq 0$ and $0 \leq \delta \leq 1$, a privacy mechanism $P_{Y|X}$ satisfies (ϵ, δ) -closedness if and only if it satisfies (ϵ, δ) -EML.

Proof: Suppose without loss of generality that P_X has full support. We first show that if a privacy mechanism $P_{Y|X}$ satisfies (ϵ, δ) -closedness, then it satisfies (ϵ, δ) -EML. Informally, this result follows from the fact that for all $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$, optimizing over the events in $\text{supp}(P_{\bar{Y}})$ with probability at least δ is equivalent to optimizing over the outcomes of all deterministic mappings $P_{Z|\bar{Y}}$ with probability at least δ . More concretely, suppose $P_{Z|\bar{Y}}$ is a deterministic channel, that is, the matrix form of $P_{Z|\bar{Y}}$ consists only of zeros and ones. Then, each outcome $z \in \text{supp}(P_Z)$ corresponds to some event $\mathcal{E}_z \subseteq \bar{\mathcal{Y}}$ such that $P_{Z|\bar{Y}=y}(z) = 1$ for all $y \in \mathcal{E}_z$, and $P_{\bar{Y}}(\mathcal{E}_z) = P_Z(z)$. Let $\mathcal{D}_{\bar{\mathcal{Y}}}$ denote the set of all deterministic mappings with domain $\bar{\mathcal{Y}}$. We can write

$$\begin{aligned}
\epsilon &\geq \sup_{P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})} \sup_{P_{Z|\bar{Y}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \ell_{P_{XZ}}(X \rightarrow z) \\
&\geq \sup_{P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})} \sup_{P_{Z|\bar{Y}} \in \mathcal{D}_{\bar{\mathcal{Y}}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \ell_{P_{XZ}}(X \rightarrow z) \\
&= \sup_{P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})} \sup_{P_{Z|\bar{Y}} \in \mathcal{D}_{\bar{\mathcal{Y}}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \log \max_x \frac{P_{Z|X=x}(z)}{P_Z(z)} \\
&= \sup_{P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})} \sup_{P_{Z|\bar{Y}} \in \mathcal{D}_{\bar{\mathcal{Y}}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \log \max_x \frac{P_{\bar{Y}|X=x}(\mathcal{E}_z)}{P_{\bar{Y}}(\mathcal{E}_z)} \\
&= \sup_{P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})} \max_{\substack{\mathcal{E} \subseteq \text{supp}(P_{\bar{Y}}): \\ P_{\bar{Y}}(\mathcal{E}) \geq \delta}} \log \max_x \frac{P_{\bar{Y}|X=x}(\mathcal{E})}{P_{\bar{Y}}(\mathcal{E})} \\
&= \sup_{P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})} \max_{\substack{\mathcal{E} \subseteq \text{supp}(P_{\bar{Y}}): \\ P_{\bar{Y}}(\mathcal{E}) \geq \delta}} \ell_{P_{X\bar{Y}}}(X \rightarrow \mathcal{E}),
\end{aligned}$$

where the first inequality follows from Proposition 2. Thus, $P_{Y|X}$ satisfies (ϵ, δ) -EML.

Now, we show that if a mechanism $P_{Y|X}$ satisfies (ϵ, δ) -EML, then it satisfies (ϵ, δ) -closedness. Let the function $h : \mathcal{C}(P_{Y|X}) \times [0, 1] \rightarrow [1, \infty)$ be defined as

$$h(P_{\bar{Y}|X}, \delta) = \sup_{P_{Z|\bar{Y}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \exp\left(\ell_{P_{XZ}}(X \rightarrow z)\right). \quad (15)$$

We can write h as

$$h(P_{\bar{Y}|X}, \delta) = \max_x h_x(P_{\bar{Y}|X}, \delta),$$

where

$$\begin{aligned} h_x(P_{\bar{Y}|X}, \delta) &:= \sup_{P_{Z|\bar{Y}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \frac{P_{Z|X=x}(z)}{P_Z(z)} \\ &= \sup_{P_{Z|\bar{Y}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \frac{\sum_{y \in \text{supp}(P_{\bar{Y}})} P_{Z|\bar{Y}=y}(z) P_{\bar{Y}|X=x}(y)}{\sum_{y' \in \text{supp}(P_{\bar{Y}})} P_{Z|\bar{Y}=y'}(z) P_{\bar{Y}}(y')} \\ &= \sup_{P_{Z|\bar{Y}}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \sum_{y \in \text{supp}(P_{\bar{Y}})} \frac{P_{Z|\bar{Y}=y}(z) P_{\bar{Y}}(y)}{\sum_{y' \in \text{supp}(P_{\bar{Y}})} P_{Z|\bar{Y}=y'}(z) P_{\bar{Y}}(y')} \left(\frac{P_{\bar{Y}|X=x}(y)}{P_{\bar{Y}}(y)} \right). \end{aligned} \quad (16)$$

Using Proposition 2, it suffices to show that for each x , there exists $P_{\bar{Y}|X} \in \mathcal{C}(P_{Y|X})$ satisfying $h_x(P_{\bar{Y}|X}, \delta) \leq \exp(\epsilon)$. Hence, we solve the above optimization problem for the reduced channel associated with the class $\mathcal{C}(P_{Y|X})$, denoted by $P_{Y_r|X}$.

Fix some $x \in \mathcal{X}$. Let $n_r := |\mathcal{Y}_r|$ denote the cardinality of \mathcal{Y}_r (recall that P_{Y_r} has full support). We re-write (16) for $P_{Y_r|X}$ as

$$\begin{aligned} h_x(P_{Y_r|X}, \delta) &= \max_{a_1, \dots, a_{n_r}} \sum_{j=1}^{n_r} \frac{a_j P_{Y_r}(y_j)}{\sum_{j'=1}^{n_r} a_{j'} P_{Y_r}(y_{j'})} \exp(i_{P_{X Y_r}}(x; y_j)), \\ &\text{subject to } \sum_{j=1}^{n_r} a_j P_{Y_r}(y_j) \geq \delta, \\ &0 \leq a_j \leq 1, \quad \forall j \in [n_r], \end{aligned} \quad (17)$$

where $\{a_j\}$ specify $P_{Z|Y_r=y_j}(z)$ for the $z \in \text{supp}(P_Z)$ with the largest PML which also satisfies $P_Z(z) \geq \delta$.

Suppose the elements in \mathcal{Y}_r are labelled such that $i_{P_{X Y_r}}(x; y_1) \geq i_{P_{X Y_r}}(x; y_2) \geq \dots \geq i_{P_{X Y_r}}(x; y_{n_r})$. Given an integer $k \in [n_r]$, let $\mathcal{F}_k := \{y_1, \dots, y_k\}$ be the set containing k elements from \mathcal{Y}_r that have the largest information density with x . Let $k^* \in [n_r]$ be the smallest integer such that $P_{Y_r}(\mathcal{F}_{k^*}) \geq \delta$. The objective function in problem (17) is a linear-fractional function which is quasi-convex (in fact, quasi-linear) [50, Section 3.4], and the feasible region is a convex polytope. Therefore, the optimal solution is an extreme point of the feasible region given by

$$a_j^* = \begin{cases} 1, & \text{if } j = 1, \dots, k^* - 1, \\ \zeta, & \text{if } j = k^*, \\ 0, & \text{otherwise,} \end{cases}$$

where the parameter $0 < \zeta \leq 1$ can be calculated by

$$P_{Y_r}(\mathcal{F}_{k^*-1}) + \zeta P_{Y_r}(y_{k^*}) = \delta. \quad (18)$$

Hence, we obtain $h_x(P_{Y_r|X}, \delta)$ as

$$h_x(P_{Y_r|X}, \delta) = \frac{1}{\delta} \left(P_{Y_r|X=x}(\mathcal{F}_{k^*-1}) + \zeta P_{Y_r|X=x}(y_{k^*}) \right). \quad (19)$$

Since k^* is the smallest integer such that $P_{Y_r}(\mathcal{F}_{k^*}) \geq \delta$, we need to consider the following two possibilities: We either have $P_{Y_r}(\mathcal{F}_{k^*}) = \delta$ or $P_{Y_r}(\mathcal{F}_{k^*}) > \delta$. First, suppose $\mathcal{F}_{k^*} = \delta$. In this case, the optimal parameters become

$$a_j^* = \begin{cases} 1, & \text{if } j = 1, \dots, k^*, \\ 0, & \text{otherwise,} \end{cases}$$

that is, $\zeta = 1$ in (18). Since $\{a_j^*\}$ consist of only zeros and ones, it in fact specifies a deterministic outcome $z^* \in \text{supp}(P_{Z^*})$ for some channel $P_{Z^*|Y_r}$. This outcome corresponds to the event \mathcal{F}_{k^*} in the sense that $P_{Z|Y_r=y}(z^*) = 1$ for all $y \in \mathcal{F}_{k^*}$, and $P_{Y_r}(\mathcal{F}_{k^*}) = P_Z(z^*)$. Thus, we get

$$\begin{aligned} h_x(P_{Y_r|X}, \delta) &= \sup_{P_{Z|Y}} \max_{\substack{z \in \text{supp}(P_Z): \\ P_Z(z) \geq \delta}} \frac{P_{Z|X=x}(z)}{P_Z(z)} \\ &= \frac{P_{Z^*|X=x}(z^*)}{P_{Z^*}(z^*)} \\ &= \frac{P_{Y_r|X=x}(\mathcal{F}_{k^*})}{P_{Y_r}(\mathcal{F}_{k^*})} \\ &\leq \exp(\epsilon), \end{aligned} \tag{20}$$

where the last inequality follows from the fact that $P_{Y|X}$ satisfies (ϵ, δ) -EML.

Now, suppose $P_{Y_r}(\mathcal{F}_{k^*}) > \delta$ which implies that $0 < \zeta < 1$ in (18). In this case, we construct $P_{\hat{Y}|X} \in \mathcal{C}(P_{Y|X})$ whose columns are identical to the columns of $P_{Y_r|X}$, except that the k^* -th column of $P_{Y_r|X}$ is split into two corresponding columns in $P_{\hat{Y}|X}$ given by

$$P_{\hat{Y}|X=x}(y_{k^*_{(1)}}) = \zeta P_{Y_r|X=x}(y_{k^*}), \quad \text{and} \quad P_{\hat{Y}|X=x}(y_{k^*_{(2)}}) = (1 - \zeta) P_{Y_r|X=x}(y_{k^*}),$$

for all $x \in \mathcal{X}$. Note that the outcomes $y_{k^*_{(1)}}, y_{k^*_{(2)}} \in \text{supp}(P_{\hat{Y}})$ defined above are similar, and satisfy

$$i_{P_{X\hat{Y}}}(x; y_{k^*_{(1)}}) = i_{P_{X\hat{Y}}}(x; y_{k^*_{(2)}}) = i_{P_{XY_r}}(x; y_{k^*}).$$

Now, we find $h_x(P_{\hat{Y}|X}, \delta)$. Forming the optimization problem (16) for $P_{\hat{Y}|X}$, it is easy to see that the optimal parameters are

$$a_j^* = \begin{cases} 1, & \text{if } j = 1, \dots, k^* - 1, k^*_{(1)} \\ 0, & \text{otherwise,} \end{cases}$$

which, once again, specifies a deterministic outcome. Using arguments similar to (20), we get $h_x(P_{\hat{Y}|X}, \delta) \leq \exp(\epsilon)$.

Finally, as x was chosen arbitrarily, we conclude that $h(P_{Y|X}, \delta) \leq \exp(\epsilon)$, that is, $P_{Y|X}$ satisfies (ϵ, δ) -closedness. \blacksquare

Remark 8: The proof of Theorem 3 sheds light on the role of the class $\mathcal{C}(P_{Y|X})$: For each $0 \leq \delta \leq 1$, there exists $P_{\hat{Y}|X} \in \mathcal{C}(P_{Y|X})$ and a ‘‘least private’’ event $\mathcal{E}^* \subseteq \text{supp}(P_{\hat{Y}})$ satisfying $P_{\hat{Y}}(\mathcal{E}^*) = \delta$. As such, without loss of generality, we unify the channels in the equivalence class $\mathcal{C}(P_{Y|X})$ and assume that $\mathcal{E}^* \subseteq \text{supp}(P_Y)$. Then, to show that $P_{Y|X}$ satisfies (ϵ, δ) -EML, it suffices to show that $\ell_{P_{XY}}(X \rightarrow \mathcal{E}^*) \leq \epsilon$.

Now, recall that our motivation for introducing the notion of (ϵ, δ) -EML was to obtain a probabilistic privacy guarantee which is closed under both pre- and post-processing. The following result formally shows that this is indeed the case.

Proposition 3: Suppose the three random variables X , Y , and Z form the Markov chain $X - Y - Z$. Given $\epsilon \geq 0$ and $0 \leq \delta \leq 1$, it holds that:

- 1) (Pre-processing) If $P_{Z|Y}$ satisfies (ϵ, δ) -EML, then $P_{Z|X}$ satisfies (ϵ, δ) -EML.
- 2) (Post-processing) If $P_{Y|X}$ satisfies (ϵ, δ) -EML, then $P_{Z|X}$ satisfies (ϵ, δ) -EML.

Proof: In both cases, we use Theorem 3 and verify the conditions of Definition 7. Consider the Markov chain $X - Y - Z - T$.

- 1) Fix an arbitrary $P_{T|Z}$ and $t \in \text{supp}(P_T)$ satisfying $P_T(t) \geq \delta$. Then,

$$\ell(X \rightarrow t) \leq \ell(Y \rightarrow t) \leq \epsilon,$$

where the first inequality is due to Lemma 1 and the second inequality follows by the assumption that $P_{Z|Y}$ satisfies (ϵ, δ) -EML. Thus, $P_{Z|X}$ satisfies (ϵ, δ) -EML.

- 2) The result follows directly by noticing that T is a post-processing of Y through the channel $P_{T|Y} = P_{T|Z} \circ P_{Z|Y}$. Hence, $P_T(t) \geq \delta$ implies $\ell(X \rightarrow t) \leq \epsilon$ with $t \in \text{supp}(P_T)$ and $P_{Z|X}$ satisfies (ϵ, δ) -EML. ■

Now, let us re-visit Example 5 and analyze it through the lens of event maximal leakage.

Example 6: Suppose P_X , $P_{Y|X}$ and, $P_{Z|Y}$ are defined as in Example 5, and let $\delta = \frac{1}{6}$. Our goal is to find the smallest $\epsilon_1 \geq 0$ such that $P_{Y|X}$ satisfies (ϵ_1, δ) -EML, and the smallest $\epsilon_2 \geq 0$ such that $P_{Z|X}$ satisfies (ϵ_2, δ) -EML. First, note that the outcomes y_3 and y_4 are similar; hence, by merging them, we obtain the reduced channel $P_{Y_r|X}$ as

$$P_{Y_r|X} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & \frac{1}{3} & \frac{2}{3} \\ \frac{1}{3} & 0 & \frac{2}{3} \end{bmatrix}.$$

Now, for each x , we find $h_x(P_{Y_r|X}, \delta)$ defined in (19):

$$\begin{aligned} h_{x_1}(P_{Y_r|X}, \delta) &= h_{x_2}(P_{Y_r|X}, \delta) = \frac{6}{5}, \\ h_{x_3}(P_{Y_r|X}, \delta) &= h_{x_4}(P_{Y_r|X}, \delta) = \frac{12}{5}, \end{aligned}$$

which implies that $P_{Y|X}$ satisfies (ϵ_1, δ) -EML with

$$\epsilon_1 = \log \max_x h_{x_1}(P_{Y_r|X}, \delta) = \log \frac{12}{5}.$$

Furthermore, the channel $P_{Z|X}$ is given by

$$P_{Z|X} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \\ \frac{1}{3} & \frac{2}{3} \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}.$$

Since $P_Z(z_1) = P_Z(z_2) = \frac{1}{2}$ and $\ell(X \rightarrow z_1) = \ell(X \rightarrow z_2) = \log \frac{4}{3}$, it follows that $P_{Z|X}$ satisfies (ϵ_2, δ) -EML with $\epsilon_2 = \log \frac{4}{3}$. Note that since $\epsilon_2 < \epsilon_1$, $P_{Z|X}$ also satisfies (ϵ_1, δ) -EML which was expected from Proposition 3.

Remark 9: The above example also describes the computational complexity of finding the smallest ϵ associated with a given δ in an (ϵ, δ) -EML privacy guarantee (or alternatively, finding the least private event with probability δ). For each $x \in \mathcal{X}$, finding $h_x(P_{Y|X}, \delta)$ requires sorting the vector of information density $i(x; y_1), \dots, i(x; y_{|\mathcal{Y}|})$, which can be implemented with time complexity of $O(|\mathcal{Y}| \log |\mathcal{Y}|)$. Hence, the overall procedure can be implemented relatively efficiently.

As the final topic in this section, we discuss the relationship between the (ϵ, δ) -EML and (ϵ, δ) -PML privacy guarantees. By Examples 5 and 6, it is clear that (ϵ, δ) -PML does not imply (ϵ, δ) -EML. In general, (ϵ, δ) -EML does not imply (ϵ, δ) -PML either. For example, consider the binary symmetric channel

$$P_{Y|X} = \begin{bmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \end{bmatrix},$$

with uniform P_X . It is straightforward to verify that $P_{Y|X}$ satisfies (ϵ_1, δ_1) -EML with $\epsilon_1 = \log \frac{34}{30}$ and $\delta_1 = 0.6$. However, due to the symmetry of the channel, $P_{Y|X}$ satisfies (ϵ_2, δ_2) -PML with $\epsilon_2 = \log \frac{6}{5}$ and all $0 \leq \delta_2 < 1$. Note that for $\delta_1 = \delta_2 = 0.6$ we have $\epsilon_1 < \epsilon_2$.

While in general (ϵ, δ) -EML does not imply (ϵ, δ) -PML, there exists a condition under which (ϵ, δ) -EML does in fact imply (ϵ, δ) -PML. This is shown in our next result.

Proposition 4: Given an arbitrary but fixed prior P_X , suppose the privacy mechanism $P_{Y|X}$ satisfies (ϵ, δ) -EML. Let $\mathcal{A} = \{y \in \text{supp}(P_Y) : \ell(X \rightarrow y) > \epsilon\}$. If there exists $x^* \in \text{supp}(P_X)$ satisfying $x^* \in \arg \max_{x \in \text{supp}(P_X)} i(x; y)$ for all $y \in \mathcal{A}$, then $P_{Y|X}$ satisfies (ϵ, δ) -PML.

Proof: We need to show that $P_Y(\mathcal{A}) \leq \delta$. We can write

$$\begin{aligned} \ell(X \rightarrow \mathcal{A}) &= \log \max_{x \in \text{supp}(P_X)} \frac{\sum_{y \in \mathcal{A}} P_{Y|X=x}(y)}{\sum_{y \in \mathcal{A}} P_Y(y)} \\ &= \log \frac{\sum_{y \in \mathcal{A}} \max_{x \in \text{supp}(P_X)} P_{Y|X=x}(y)}{\sum_{y \in \mathcal{A}} P_Y(y)} \end{aligned} \tag{21a}$$

$$\begin{aligned} &\geq \log \min_{y \in \mathcal{A}} \max_{x \in \text{supp}(P_X)} \frac{P_{Y|X=x}(y)}{P_Y(y)} \\ &= \log \min_{y \in \mathcal{A}} \ell(X \rightarrow y) \\ &> \epsilon, \end{aligned} \tag{21b}$$

where (21a) follows from the assumption that all $P_{Y|X=x}(y)$ are maximized at the same x , and (21b) follows from the definition of the event \mathcal{A} . Since $P_{Y|X}$ satisfies (ϵ, δ) -EML, $\ell(X \rightarrow \mathcal{A}) > \epsilon$ implies that $P_Y(\mathcal{A}) < \delta$. ■

B. Composition Properties

A second group of properties that are helpful while assessing the privacy levels of more complicated systems are composition properties. Let X denote the sensitive data, and let Y be the output of a channel $P_{Y|X}$. Suppose

X and Y are fed into another channel $P_{Z|XY}$ inducing a random variable Z in its output. What we are interested in is to find out how much information the overall channel $P_{ZY|X}$ leaks about X .

The composition property stated in Lemma 1 describes an upper bound on the PML resulting from composing the two channels $P_{Y|X}$ and $P_{Z|XY}$. In this section, our goal is to understand how different privacy guarantees, namely, ϵ -PML, (ϵ, δ) -PML and (ϵ, δ) -EML are affected by adaptively composing two channels. We note that some related results have been derived in [51], where *non-adaptive* composition is studied asymptotically for maximal leakage and Sibson mutual information [52, 11]. The bounds we derive in this section differ from previous works in that here we study *adaptive* composition, that is, we allow Z to depend arbitrarily on both X and Y .

Naturally, one can formulate various problems by making different assumptions about the channels $P_{Y|X}$ and $P_{Z|XY}$. The following result contains several such problem formulations and results, and its proof is deferred to Appendix C.

Theorem 4: Consider three random variables X , Y and Z where X denotes the secret, Y is the output of a channel $P_{Y|X}$, and Z is the output of a channel $P_{Z|XY}$.

- 1) Suppose $P_{Y|X}$ satisfies ϵ_1 -PML, and for all $y \in \text{supp}(P_Y)$, the channel $P_{Z|X,Y=y}$ satisfies ϵ_2 -PML, where $\epsilon_1, \epsilon_2 \geq 0$. Then, the channel $P_{YZ|X}$ satisfies ϵ -PML with $\epsilon = \epsilon_1 + \epsilon_2$.
- 2) Suppose $P_{Y|X}$ satisfies (ϵ_1, δ_1) -PML, and for all $y \in \text{supp}(P_Y)$, $P_{Z|X,Y=y}$ satisfies (ϵ_2, δ_2) -PML. Then, the channel $P_{YZ|X}$ satisfies (ϵ, δ) -PML with $\epsilon = \epsilon_1 + \epsilon_2$ and $\delta = \delta_1 + \delta_2 - \delta_1\delta_2$.
- 3) Suppose $P_{Y|X}$ satisfies (ϵ_1, δ_1) -PML, and

$$\mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow Z | Y) \leq \epsilon_2 \right] \geq 1 - \delta_2.$$

Then, the channel $P_{YZ|X}$ satisfies (ϵ, δ) -PML with $\epsilon = \epsilon_1 + \epsilon_2$ and $\delta = \delta_1 + \delta_2$.

- 4) Suppose $P_{Y|X}$ satisfies (ϵ_1, δ_1) -EML, and for all $y \in \text{supp}(P_Y)$, $P_{Z|X,Y=y}$ satisfies (ϵ_2, δ_2) -EML. Given an event $\mathcal{E} \subseteq \text{supp}(P_{YZ})$, define the sets

$$\mathcal{E}_Y := \{y \in \text{supp}(P_Y) : (y, z) \in \mathcal{E}, z \in \text{supp}(P_Z)\}$$

$$\mathcal{E}_Z(y) := \{z \in \text{supp}(P_Z) : (y, z) \in \mathcal{E}\}.$$

Then, for all events \mathcal{E} such that $0 \leq \delta_2 \leq \min_{y \in \mathcal{E}_Y} P_{Z|Y=y}(\mathcal{E}_Z(y))$, $P_{YZ}(\mathcal{E}) \geq \delta_1$ implies $\ell(X \rightarrow \mathcal{E}) \leq \epsilon_1 + \epsilon_2$. Specifically, if $P_{Y|X}$ satisfies (ϵ_1, δ_1) -EML, and $P_{Z|X,Y=y}$ satisfies ϵ_2 -PML, then $P_{YZ|X}$ satisfies $(\epsilon_1 + \epsilon_2, \delta_1)$ -EML.

- 5) Suppose $P_{Y|X}$ satisfies (ϵ_1, δ_1) -EML, and for all $y \in \text{supp}(P_Y)$, $P_{Z|X,Y=y}$ satisfies (ϵ_2, δ_2) -EML. Then, $P_{YZ|X}$ satisfies (ϵ, δ) -EML with $\epsilon = \log \left(\frac{\delta_2}{\delta_1 + \delta_2} \cdot \exp(\epsilon_{\max}) + \exp(\epsilon_1 + \epsilon_2) \right)$ and $\delta = \delta_1 + \delta_2$, where $\epsilon_{\max} := -\log \min_{x \in \text{supp}(P_X)} P_X(x)$.

IV. RELATIONSHIP TO OTHER PRIVACY/STATISTICAL NOTIONS

In this section, we study how pointwise maximal leakage and the privacy guarantees defined in the previous section relate to several pre-existing privacy/statistical notions from the literature. More specifically, we discuss max-information [38, 39], local differential privacy [5, 6], local information privacy [17, 16], local differential

TABLE I
SUMMARY OF THE RESULTS OF SECTION IV

Privacy/Statistical Notion	Relation/Bound	Ref.
Max-information	$I_\infty(X; Y) = \max_y \ell(X \rightarrow y)$	Def. 13
Approximate max-information	(ϵ, δ) -PML $\implies I_\infty^\delta(X; Y) \leq \epsilon$	Prop. 5
Local information privacy	ϵ -LIP $\implies \epsilon$ -PML	Def. 15
Local differential privacy	ϵ -LDP $\implies \log \frac{1}{p_{\min} + e^{-\epsilon}(1-p_{\min})}$ -PML	Prop. 6
Local differential identifiability	ϵ -LDI $\implies \log \frac{1}{p_{\min}(1 + e^{-\epsilon}(\text{supp}(P_X) - 1))}$ -PML	Prop. 6
Mutual information	$I(X; Y) \leq \mathbb{E}_{Y \sim P_Y} [\ell(X \rightarrow Y)]$	Prop. 7
f -information	$I_f(P_{XY}) \leq \mathbb{E}_{Y \sim P_Y} [\max \{f(\exp(\ell(X \rightarrow Y))), f(0)\}]$	Prop. 8
Total variation privacy	$T(X; Y) \leq \min \left\{ \frac{1}{2} \mathbb{E}_{Y \sim P_Y} [\max \{\exp(\ell(X \rightarrow Y)) - 1, 1\}], \exp(\mathcal{L}(P_{Y X})) - 1 \right\}$	Prop. 8 Prop. 9

identifiability [20, 9], mutual information, f -information [13], and total variation privacy [14]. The results of this section are summarized in Table I.

A. Max-information

Max-information is a statistical quantity that was introduced as a tool for studying generalization in adaptive data analysis [38, 39]. Note that while max-information has not been developed as a notion of privacy, it is defined similarly to pointwise maximal leakage, and therefore, their comparison is appropriate. Before we give the definition of (approximate) max-information, we need the following definition of approximate max-divergence which is a weakening of max-divergence (that is, Rényi divergence of order infinity).

Definition 12 (Approximate max-divergence [38]): Let P and Q be two probability distributions over a finite set Ω and suppose $P \ll Q$. Given $0 \leq \delta \leq 1$, the δ -approximate max-divergence between P and Q is defined as

$$D_\infty^\delta(P||Q) = \log \max_{\mathcal{E} \subseteq \Omega, P(\mathcal{E}) \geq \delta} \frac{P(\mathcal{E}) - \delta}{Q(\mathcal{E})}.$$

Note that for $\delta = 0$ the above definition reduces to the max-divergence between P and Q , denoted by $D_\infty(P||Q)$.

Definition 13 ((Approximate) max-information [38]): Suppose V and W are two random variables supported over finite sets \mathcal{V} and \mathcal{W} , respectively, and let P_{VW} denote their joint distribution. The max-information between V and W is defined as

$$\begin{aligned} I_\infty(V; W) &:= D_\infty(P_{VW}||P_V P_W) \\ &= \log \max_{v \in \mathcal{V}, w \in \mathcal{W}} \frac{P_{VW}(v, w)}{P_V(v)P_W(w)}. \end{aligned}$$

Similarly, the δ -approximate max-information between V and W is defined as

$$\begin{aligned} I_\infty^\delta(V; W) &:= D_\infty^\delta(P_{VW} \| P_V P_W) \\ &= \log \max_{\mathcal{E} \subseteq \mathcal{V} \times \mathcal{W}: P_{VW}(\mathcal{E}) \geq \delta} \frac{P_{VW}(\mathcal{E}) - \delta}{P_V P_W(\mathcal{E})}, \end{aligned}$$

where

$$P_V P_W(\mathcal{E}) = \sum_{(v,w) \in \mathcal{E}} P_V(v) P_W(w).$$

It follows from the definition of max-information that, given a fixed prior P_X , a mechanism $P_{Y|X}$ satisfies ϵ -PML if and only if $I_\infty(X; Y) \leq \epsilon$. Therefore, in the following we examine how (ϵ, δ) -PML privacy compares with guarantees given in terms of the δ -approximate max-information. To do this, first we recall a lemma from [53]. As the proof of the lemma was omitted in [53], here we also provide a proof.

Lemma 4 ([53, Lemma 18]): Let P and Q be two probability distributions over a finite set Ω and suppose $P \ll Q$. Let the event $\mathcal{O} \subseteq \Omega$ be defined as

$$\mathcal{O} := \{\omega \in \Omega : \frac{P(\omega)}{Q(\omega)} > e^\epsilon\}.$$

If $P(\mathcal{O}) \leq \delta$, then $D_\infty^\delta(P \| Q) \leq \epsilon$.

Proof: Let $\tilde{\mathcal{O}} \subseteq \Omega$ be any event satisfying $P(\tilde{\mathcal{O}}) \geq \delta$. Then, we have

$$\begin{aligned} P(\tilde{\mathcal{O}}) &= P(\tilde{\mathcal{O}} \cap \mathcal{O}) + P(\tilde{\mathcal{O}} \cap \mathcal{O}^c) \\ &\leq P(\mathcal{O}) + Q(\tilde{\mathcal{O}} \cap \mathcal{O}^c) e^\epsilon \\ &\leq \delta + Q(\tilde{\mathcal{O}}) e^\epsilon, \end{aligned}$$

which gives

$$\frac{P(\tilde{\mathcal{O}}) - \delta}{Q(\tilde{\mathcal{O}})} \leq e^\epsilon, \quad \forall \tilde{\mathcal{O}} \subseteq \Omega, P(\tilde{\mathcal{O}}) \geq \delta. \quad \blacksquare$$

Now, we use Lemma 4 to relate $\mathbb{P}_{Y \sim P_Y}[\ell(X \rightarrow Y) \leq \epsilon]$ and the δ -approximate max-information.

Proposition 5: Given an arbitrary but fixed P_X , if the channel $P_{Y|X}$ satisfies (ϵ, δ) -PML, then $I_\infty^\delta(X; Y) \leq \epsilon$.

Proof: First, note that

$$\mathbb{P}_{(X,Y) \sim P_{XY}} \left[\frac{P_{XY}(X, Y)}{P_X(X) P_Y(Y)} \leq e^\epsilon \right] \geq \mathbb{P}_{Y \sim P_Y} \left[\max_{x \in \text{supp}(P_X)} \frac{P_{XY}(x, Y)}{P_X(x) P_Y(Y)} \leq e^\epsilon \right], \quad (22)$$

since the RHS of the above equality can be written as

$$\begin{aligned} \mathbb{P}_{Y \sim P_Y} \left[\max_{x \in \text{supp}(P_X)} \frac{P_{XY}(x, Y)}{P_X(x) P_Y(Y)} \leq e^\epsilon \right] &= \sum_{y \in \text{supp}(P_Y)} P_Y(y) \mathbb{1} \left[\max_{x \in \text{supp}(P_X)} \frac{P_{XY}(x, y)}{P_X(x) P_Y(y)} \leq e^\epsilon \right], \\ &= \sum_{(x,y) \in \text{supp}(P_{XY})} P_{XY}(x, y) \mathbb{1} \left[\max_{x \in \text{supp}(P_X)} \frac{P_{XY}(x, y)}{P_X(x) P_Y(y)} \leq e^\epsilon \right], \end{aligned}$$

and for all $(x, y) \in \text{supp}(P_{XY})$ we have

$$\frac{P_{XY}(x, y)}{P_X(x) P_Y(y)} \leq \max_{x \in \text{supp}(P_X)} \frac{P_{XY}(x, y)}{P_X(x) P_Y(y)}.$$

Now, since the privacy mechanism $P_{Y|X}$ satisfies (ϵ, δ) -PML, we have

$$\mathbb{P}_{Y \sim P_Y} \left[\max_{x \in \text{supp}(P_X)} \frac{P_{XY}(x, Y)}{P_X(x)P_Y(Y)} \leq e^\epsilon \right] = \mathbb{P}_{Y \sim P_Y} [\ell(X \rightarrow Y) \leq \epsilon] \geq 1 - \delta,$$

which combined with (22) yields

$$\mathbb{P}_{(X, Y) \sim P_{XY}} \left[\frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} \leq e^\epsilon \right] \geq 1 - \delta,$$

or equivalently,

$$\mathbb{P}_{(X, Y) \sim P_{XY}} \left[\frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} > e^\epsilon \right] \leq \delta.$$

Finally, using the above inequality and Lemma 4 we conclude that

$$I_\infty^\delta(X; Y) = D_\infty^\delta(P_{XY} \| P_X P_Y) \leq \epsilon. \quad \blacksquare$$

The previous result shows that (ϵ, δ) -PML is a stronger guarantee compared to $I_\infty^\delta(X; Y) \leq \epsilon$. Roughly speaking, this is because under a $I_\infty^\delta(X; Y) \leq \epsilon$ guarantee, the “good” y ’s are those that have small information density $i(x; y)$ *with high probability* over the x ’s. However, under an (ϵ, δ) -PML guarantee, the “good” y ’s need to have small $i(x; y)$ for *all* x ’s in $\text{supp}(P_X)$, that is, $i(x; y)$ must be small *with probability one* over the x ’s. In addition, note that $I_\infty^\delta(X; Y)$ treats random variables X and Y symmetrically and the probability of a good event is calculated according to P_{XY} . On the other hand, under (ϵ, δ) -PML, the probability of a good event (i.e., low leakage) is calculated according to P_Y over the y ’s.

B. LDP, LIP, and LDI

Now, we discuss the relationship between PML-based guarantees and three other notions of privacy, namely, local differential privacy (LDP) [5, 6], local information privacy (LIP) [17, 16] and local differential identifiability (LDI) [20, 9]. First, we recall their definitions.

Definition 14 (Local differential privacy [5, 6]): A privacy mechanism $P_{Y|X}$ satisfies ϵ -LDP with $\epsilon \geq 0$ if for all $y \in \text{supp}(P_Y)$, $x, x' \in \text{supp}(P_X)$ we have

$$\frac{P_{Y|X=x}(y)}{P_{Y|X=x'}(y)} \leq e^\epsilon.$$

Note that the above definition depends only on the channel $P_{Y|X}$, so an ϵ -LDP guarantee is valid for all priors P_X .

Definition 15 (Local information privacy [17, 16]): Given an arbitrary but fixed prior P_X , we say that a privacy mechanism $P_{Y|X}$ satisfies ϵ -LIP with $\epsilon \geq 0$ if for all $y \in \text{supp}(P_Y)$ and $x \in \text{supp}(P_X)$ we have

$$e^{-\epsilon} \leq \frac{P_{X|Y=y}(x)}{P_X(x)} \leq e^\epsilon.$$

While the definition of ϵ -LIP privacy is similar to ϵ -PML privacy, it differs from ϵ -PML in that it requires an additional lower bound on the information density, that is, we must also have

$$\frac{P_X(x)}{P_{X|Y=y}(x)} \leq e^\epsilon.$$

The above bound has no clear operational interpretation in our current framework, and may be superfluous. To see why, consider the following simple example. Suppose P_X is a binary random variable that is uniformly distributed over $\mathcal{X} = \{x_1, x_2\}$. Assume $P_{X|Y=y}(x_1) = p$ for some $y \in \text{supp}(P_Y)$, where $p > 0$ is a small positive number. Since we must have $P_{X|Y=y}(x_1) + P_{X|Y=y}(x_2) = 1$, this in turn implies that $P_{X|Y=y}(x_2) = 1 - p$ where $1 - p$ is close to one. Thus, the pointwise maximal leakage $\ell(X \rightarrow y)$ becomes

$$\ell(X \rightarrow y) = \log \max_{x \in \{x_1, x_2\}} \frac{P_{X|Y=y}(x)}{P_X(x)} = \log(2(1 - p)),$$

which is close to $\epsilon_{\max} = \log 2$. Hence, the outcome y has large privacy leakage.

The point of the above example is to show that small values of the ratio $\frac{P_{X|Y=y}(x)}{P_X(x)}$ can increase the pointwise maximal leakage simply because we must have $\sum_x P_{X|Y=y}(x) = 1$ for all $y \in \text{supp}(P_Y)$. As such, it may not be necessary to impose a lower bound on the information density as a separate constraint; a privacy guarantee defined based on an upper bound on information density will be automatically penalized for small values of the ratio $\frac{P_{X|Y=y}(x)}{P_X(x)}$.

Definition 16 (Local differential identifiability [20, 9]): Given an arbitrary but fixed prior P_X , we say that a privacy mechanism $P_{Y|X}$ satisfies ϵ -LDI with $\epsilon \geq 0$ if for all $y \in \text{supp}(P_Y)$, $x, x' \in \text{supp}(P_X)$ we have

$$\frac{P_{X|Y=y}(x)}{P_{X|Y=y}(x')} \leq e^\epsilon.$$

Note that the notion of identifiability has been previously considered in centralized settings [20, 9], where x and x' denote neighboring databases. Here, we have given a local version of the definition, where the ratio $\frac{P_{X|Y=y}(x)}{P_{X|Y=y}(x')}$ must be bounded by e^ϵ for all $x, x' \in \text{supp}(P_X)$.

What the above three notions of privacy have in common is that they are strictly intolerant of zero-probability assignments in the channel $P_{Y|X}$, that is, the existence of a single input-output pair (x, y) such that $P_{Y|X=x}(y) = 0$ (where $x \in \text{supp}(P_X)$ and $y \in \text{supp}(P_Y)$) immediately implies that the channel does not satisfy any of the above notions of privacy. On the other hand, $\ell(X \rightarrow y)$ is always bounded by ϵ_{\max} , so naturally a guarantee given in terms of pointwise maximal leakage may not satisfy any of the above notions of privacy.

To see why zero-probability assignments in the channel $P_{Y|X}$ do not necessarily imply “bad privacy”, consider the following simple example. Suppose X and Y are random variables defined over sets with cardinality n , and assume that X is uniformly distributed. Consider the following channel:

$$P_{Y|X=x_1}(y_i) = \begin{cases} 0, & \text{if } i = 1, \\ \frac{1}{n-1} & \text{otherwise,} \end{cases}$$

and $P_{Y|X=x_j}(y_i) = \frac{1}{n}$ with $i \in \{1, \dots, n\}$ and $j \in \{2, \dots, n\}$. Intuitively, for large n the above channel leaks very little information which also becomes apparent by calculating the pointwise maximal leakages:

$$\begin{aligned} \ell(X \rightarrow y_1) &= \log \frac{n}{n-1}, \\ \ell(X \rightarrow y_i) &= \log \frac{n^2}{n^2 - n + 1}, \quad i = 2, \dots, n. \end{aligned}$$

However, under LDP/LIP/LDI, no matter how large n is, the above channel is considered to be equally non-private as a deterministic mapping from X to Y . This is clearly an overly pessimistic privacy assessment.

That being said, it is straightforward to verify that ϵ -LDP/LIP/LDI guarantees imply a guarantee based on pointwise maximal leakage. According to Definition 15, ϵ -LIP implies ϵ -PML. We also have the following relations.

Proposition 6: Let $p_{\min} := \min_{x \in \text{supp}(P_X)} P_X(x)$. Given an arbitrary but fixed prior P_X , if a mechanism $P_{Y|X}$ satisfies ϵ -LDP, then $P_{Y|X}$ satisfies ϵ' -PML, where

$$\epsilon' = \log \frac{1}{p_{\min} + e^{-\epsilon}(1 - p_{\min})}.$$

Moreover, if $P_{Y|X}$ satisfies ϵ -LDI, then it satisfies $\tilde{\epsilon}$ -PML, where

$$\tilde{\epsilon} = \log \frac{1}{p_{\min}(1 + e^{-\epsilon}(|\text{supp}(P_X)| - 1))}.$$

Proof: The first statement regarding LDP is an intermediate step in the proof of [16, Thm. 3]. To see the second statement, note that $P_{Y|X}$ satisfies ϵ -LDI if

$$\frac{P_{Y|X=x}(y)P_X(x)}{P_{Y|X=x'}(y)P_X(x')} \leq e^\epsilon,$$

for all $y \in \text{supp}(P_Y)$, $x, x' \in \text{supp}(P_X)$. Fix $y \in \text{supp}(P_Y)$. We write

$$\begin{aligned} \max_{x \in \text{supp}(P_X)} \frac{P_{Y|X=x}(y)}{P_Y(y)} &= \max_{x \in \text{supp}(P_X)} \frac{P_{Y|X=x}(y)}{\sum_{x'} P_{Y|X=x'}(y)P_X(x')} \\ &= \max_{x \in \text{supp}(P_X)} \frac{P_{Y|X=x}(y)}{P_{Y|X=x}(y)P_X(x) + \sum_{x' \neq x} P_{Y|X=x'}(y)P_X(x')} \\ &\leq \max_{x \in \text{supp}(P_X)} \frac{P_{Y|X=x}(y)}{P_{Y|X=x}(y)P_X(x) + \sum_{x' \neq x} P_{Y|X=x'}(y)P_X(x)e^{-\epsilon}} \\ &= \max_{x \in \text{supp}(P_X)} \frac{1}{P_X(x)(1 + e^{-\epsilon}(|\text{supp}(P_X)| - 1))} \\ &= \frac{1}{p_{\min}(1 + e^{-\epsilon}(|\text{supp}(P_X)| - 1))}. \end{aligned}$$

■

C. Mutual Information

Mutual information has been studied as a privacy measure in a number of works, e.g., [7–10], although it has been argued that in some cases mutual information underestimates privacy leakage [33]. Previous works have shown that maximal leakage upper bounds mutual information, and in fact, no scalar multiple of mutual information can upper bound maximal leakage [26, 33]. Here, we show how mutual information can be bounded in terms of $\ell(X \rightarrow Y)$.

Proposition 7: Given a joint distribution P_{XY} , it holds that $I(X; Y) \leq \mathbb{E}_{Y \sim P_Y}[\ell(X \rightarrow Y)]$ with equality if and only if $P_{Y|X=x}(y) = P_{Y|X=x'}(y)$ for all $x, x' \in \mathcal{X}$ and $y \in \mathcal{Y}$ such that $P_{XY}(x, y) > 0$ and $P_{XY}(x', y) > 0$. Hence, given an arbitrary but fixed prior P_X ,

- 1) if $P_{Y|X}$ satisfies ϵ -PML, then $I(X; Y) \leq \epsilon$, and
- 2) if $P_{Y|X}$ satisfies (ϵ, δ) -PML, then $I(X; Y) \leq \epsilon + \delta \cdot \epsilon_{\max}$.

Proof:

$$\begin{aligned} I(X; Y) &= \mathbb{E}_{Y \sim P_Y} \left[\mathbb{E}_{X \sim P_{X|Y}} \left[\log \frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} \right] \right] \\ &\leq \mathbb{E}_{Y \sim P_Y} \left[\max_{x \in \text{supp}(P_X)} \log \frac{P_{XY}(x, Y)}{P_X(x)P_Y(Y)} \right] \\ &= \mathbb{E}_{Y \sim P_Y} [\ell(X \rightarrow Y)], \end{aligned}$$

where the inequality holds with equality if and only if for all x, x', y such that $P_{XY}(x, y) > 0$ and $P_{XY}(x', y) > 0$ we have $i(x; y) = i(x'; y)$, or equivalently, $P_{Y|X=x}(y) = P_{Y|X=x'}(y)$ (the condition for equality has also been noted in [33, Lemma 2]). The remaining statements then follow directly from the definitions of ϵ -PML and (ϵ, δ) -PML, and the above inequality. \blacksquare

D. f -information

f -information [13] refers to a class of information measures that are defined based on f -divergences [54]. First, we recall the definition of f -divergence.

Definition 17 (f -divergence): Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a convex function satisfying $f(1) = 0$. Let P and Q be two probability distributions defined over a finite set Ω , and suppose $P \ll Q$. The f -divergence between P and Q is defined as

$$D_f(P||Q) := \mathbb{E}_Q \left[f \left(\frac{P}{Q} \right) \right] = \sum_{\omega \in \Omega} Q(\omega) f \left(\frac{P(\omega)}{Q(\omega)} \right).$$

Common f -divergences include, for example, KL-divergence, total variation distance, and χ^2 -divergence. Now, we can define f -information as the f -divergence between the joint distribution and the product of the marginals of two random variables.

Definition 18 (f -information [13, Def. 7]): Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a convex function such that $f(1) = 0$. Given a joint distribution P_{VW} defined over the finite alphabets \mathcal{V} and \mathcal{W} , the f -information of P_{VW} is defined as

$$I_f(P_{VW}) := D_f(P_{VW}||P_V P_W) = \mathbb{E}_{(V,W) \sim P_V P_W} \left[f \left(\frac{P_{VW}(V, W)}{P_V(V)P_W(W)} \right) \right].$$

Diaz et al. [13] justify the definition of f -information as an information/privacy measure on account of its operational meaning for specific choices of the function f , such as mutual information (associated with KL-divergence) and χ^2 -information [55, 56] (associated with χ^2 -divergence). Next, we show how PML can be used to upper bound f -information.

Proposition 8: Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a convex function satisfying $f(1) = 0$, and suppose $\lim_{t \rightarrow 0^+} f(t) < \infty$. Then,

$$I_f(P_{XY}) \leq \mathbb{E}_{Y \sim P_Y} \left[\max \{ f(\exp(\ell(X \rightarrow Y))), f(0) \} \right],$$

where $f(0)$ is defined by continuity as $f(0) := f(0^+)$.

Proof:

$$\begin{aligned}
I_f(P_{XY}) &= \mathbb{E}_{Y \sim P_Y} \left[\mathbb{E}_{X \sim P_X} \left[f \left(\frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} \right) \right] \right] \\
&\leq \mathbb{E}_{Y \sim P_Y} \left[\max_{x \in \text{supp}(P_X)} f \left(\frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} \right) \right] \\
&= \mathbb{E}_{Y \sim P_Y} \left[\max \left\{ f \left(\max_{x \in \text{supp}(P_X)} \frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} \right), f \left(\min_{x \in \text{supp}(P_X)} \frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} \right) \right\} \right] \quad (23a) \\
&\leq \mathbb{E}_{Y \sim P_Y} \left[\max \left\{ f \left(\max_{x \in \text{supp}(P_X)} \frac{P_{XY}(X, Y)}{P_X(X)P_Y(Y)} \right), f(0) \right\} \right] \quad (23b) \\
&= \mathbb{E}_{Y \sim P_Y} \left[\max \{ f(\exp(\ell(X \rightarrow Y))), f(0) \} \right],
\end{aligned}$$

where (23a) follows from the fact that the maximum of a convex function is attained at an extreme point, and (23b) follows from $\min_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)} \geq 0$ for all $y \in \text{supp}(P_Y)$. ■

E. Total variation privacy

Let P and Q be two probability distributions defined over a finite set Ω . The total variation distance between P and Q is defined as

$$\text{TV}(P, Q) := \frac{1}{2} \mathbb{E}_Q \left[\left| \frac{P}{Q} - 1 \right| \right] = \frac{1}{2} \sum_{\omega \in \Omega} |P(\omega) - Q(\omega)|,$$

which is an f -divergence with $f(x) = \frac{1}{2}|x - 1|$. Total variation privacy [14] is a privacy measure that is defined as the expected total variation distance between the posterior distribution $P_{X|Y}$ and the prior P_X , given by

$$T(X; Y) := \mathbb{E}_{Y \sim P_Y} [\text{TV}(P_{X|Y}, P_X)] = \sum_{y \in \mathcal{Y}} P_Y(y) \text{TV}(P_{X|Y=y}, P_X). \quad (24)$$

Rassouli and Gündüz [14] motivate the use of total variation distance as a privacy measure by arguing that $T(X; Y)$ is closed under pre- and post-processing, and by showing that controlling $T(X; Y)$ restricts the inference quality of an adversary optimizing an additive gain function, described in [17]. Furthermore, the following bound between $T(X; Y)$ and maximal leakage is derived:

$$T(X; Y) \leq (|\mathcal{X}| - 1) \cdot \max_{x \in \mathcal{X}} P_X(x) \cdot (\exp(\mathcal{L}(P_{Y|X})) - 1), \quad (25)$$

which is a rather loose bound as it depends on the cardinality of \mathcal{X} (considering that $0 \leq T(X; Y) \leq 1$). Note that $T(X; Y)$ is actually the f -information associated with the total variation distance, so by applying Proposition 8 with $f(x) = \frac{1}{2}|x - 1|$ we get

$$T(X; Y) \leq \frac{1}{2} \mathbb{E}_{Y \sim P_Y} \left[\max \{ \exp(\ell(X \rightarrow Y)) - 1, 1 \} \right], \quad (26)$$

which is tighter than (25). In the following result, we derive another upper bound on $T(X; Y)$ in terms of maximal leakage $\mathcal{L}(P_{Y|X})$, and also show how (ϵ, δ) -PML privacy guarantee constrains $T(X; Y)$.

Proposition 9: The following relationship holds between $T(X; Y)$ and $\mathcal{L}(P_{Y|X})$:

$$T(X; Y) \leq \exp(\mathcal{L}(P_{Y|X})) - 1.$$

Furthermore, given an arbitrary but fixed prior P_X , if the channel $P_{Y|X}$ satisfies (ϵ, δ) -PML, then $T(X; Y)$ is bounded as follows:

1) if $\epsilon \leq \log \frac{3}{2}$, then

$$T(X; Y) \leq e^\epsilon - 1 + \frac{\delta}{2} (e^{\epsilon_{\max}} - 1),$$

2) if $\log \frac{3}{2} \leq \epsilon \leq \log 2$, then

$$T(X; Y) \leq \frac{1}{2} + \frac{\delta}{2} (e^{\epsilon_{\max}} - 1),$$

3) if $\epsilon \geq \log 2$, then

$$T(X; Y) \leq \frac{1}{2} (e^\epsilon - 1) + \frac{\delta}{2} (e^{\epsilon_{\max}} - 1).$$

Proof: Fix some $y \in \text{supp}(P_Y)$ and define the set $\mathcal{A}_y := \{x \in \text{supp}(P_X) : P_{X|Y=y}(x) \geq P_X(x)\}$. We can write

$$\begin{aligned} \text{TV}(P_{X|Y=y}, P_X) &= \frac{1}{2} \sum_{x \in \text{supp}(P_X)} |P_{X|Y=y}(x) - P_X(x)| \\ &= \sum_{x \in \mathcal{A}_y} P_{X|Y=y}(x) - P_X(x) \\ &= \sum_{x \in \mathcal{A}_y} \left(\frac{P_{X|Y=y}(x)}{P_X(x)} - 1 \right) P_X(x) \\ &\leq \max_{x \in \mathcal{A}_y} \left(\frac{P_{X|Y=y}(x)}{P_X(x)} - 1 \right) \sum_{x \in \mathcal{A}_y} P_X(x) \\ &\leq \exp(\ell(X \rightarrow y)) - 1. \end{aligned} \tag{27}$$

Taking the expectation of the above expression over y , we get

$$T(X; Y) \leq \exp(\mathcal{L}(P_{Y|X})) - 1.$$

Now, define the function $\eta(y) := \exp(\ell(X \rightarrow y)) - 1$, $y \in \text{supp}(P_Y)$. By (26) and (27), we obtain

$$T(X; Y) \leq \mathbb{E}_{Y \sim P_Y} \left[\min \left\{ \eta(Y), \max \left\{ \frac{1}{2} \eta(Y), \frac{1}{2} \right\} \right\} \right].$$

Suppose P_X is arbitrary but fixed, and the mechanism $P_{Y|X}$ satisfies (ϵ, δ) -PML. Define $\eta_{\max} := e^{\epsilon_{\max}} - 1$. Using the fact that $\epsilon_{\max} \geq \log 2$, we conclude that with probability smaller than δ over Y , we have

$$\text{TV}(P_{X|Y=y}, P_X) \leq \frac{1}{2} \eta_{\max} = \frac{1}{2} (e^{\epsilon_{\max}} - 1). \tag{28}$$

We need to consider the following three cases for ϵ :

1) $\epsilon \leq \log \frac{3}{2}$: With probability at least $1 - \delta$ we have $\text{TV}(P_{X|Y=y}, P_X) \leq \eta(y) \leq e^\epsilon - 1$, which implies that

$$T(X; Y) \leq e^\epsilon - 1 + \frac{\delta}{2} (e^{\epsilon_{\max}} - 1).$$

2) $\log \frac{3}{2} \leq \epsilon \leq \log 2$: With probability at least $1 - \delta$ we have $\text{TV}(P_{X|Y=y}, P_X) \leq \frac{1}{2}$, which implies that

$$T(X; Y) \leq \frac{1}{2} + \frac{\delta}{2} (e^{\epsilon_{\max}} - 1).$$

3) $\epsilon \geq \log 2$: With probability at least $1 - \delta$ we have $\text{TV}(P_{X|Y=y}, P_X) \leq \frac{1}{2} \eta(y) \leq \frac{1}{2} (e^\epsilon - 1)$, which implies that

$$T(X; Y) \leq \frac{1}{2} (e^\epsilon - 1) + \frac{\delta}{2} (e^{\epsilon_{\max}} - 1).$$

V. CONCLUSIONS

In this paper, we have introduced a new privacy measure called pointwise maximal leakage defined based on the pre-existing notion of maximal leakage, which quantifies the amount of information leaking about a secret X by disclosing a single outcome of a (randomized) function calculated on X . Our results demonstrate that a framework centered around PML can be used to reason about privacy in a wide range of problems: First, we make no assumptions about the nature of the sensitive data, e.g., X can represent an entire database or a single data point collected from an individual. Then, PML is operationally meaningful in the sense that it is obtained by analyzing threat models in which all assumptions about adversaries are made explicit. Next, PML is a robust measure of privacy since it is meaningful when considering any adversary whose objective can be described by a gain function. In addition, PML satisfies useful properties, for example, it behaves well under composition, enables us to model side information, and satisfies data-processing (i.e., pre- and post-processing) inequalities. Last but not least, pointwise maximal leakage allows us to view privacy leakage as a random variable; consequently, we have the freedom to define different types of privacy guarantees based on the particular requirements of different applications, both in terms of privacy and utility.

APPENDIX A

PROOF OF THEOREM 2

Suppose without loss of generality that P_X has full support. Given an arbitrary gain function g , the g -leakage of X can be written as

$$\begin{aligned} \ell_g(X \rightarrow y) &= \log \frac{\sup_{P_{\hat{X}|Y=y}} \mathbb{E} [g(X, \hat{X}) | Y = y]}{\sup_{P_{\hat{X}}} \mathbb{E} [g(X, \hat{X})]} \\ &= \log \frac{\sup_{P_{\hat{X}|Y=y}} \sum_{x \in \mathcal{X}} \sum_{\hat{x} \in \hat{\mathcal{X}}} g(x, \hat{x}) P_{X|Y=y}(x) P_{\hat{X}|Y=y}(\hat{x})}{\sup_{P_{\hat{X}}} \sum_{x \in \mathcal{X}} \sum_{\hat{x} \in \hat{\mathcal{X}}} g(x, \hat{x}) P_X(x) P_{\hat{X}}(\hat{x})} \\ &= \log \frac{\max_{\hat{x} \in \hat{\mathcal{X}}} \sum_{x \in \mathcal{X}} g(x, \hat{x}) P_{X|Y=y}(x)}{\max_{\hat{x} \in \hat{\mathcal{X}}} \sum_{x \in \mathcal{X}} g(x, \hat{x}) P_X(x)}, \end{aligned} \tag{29}$$

where the last equality follows by plugging in $P_{\hat{X}|Y=y}^*$ defined as

$$P_{\hat{X}|Y=y}^*(\hat{x}) := \begin{cases} 1, & \text{for some } \hat{x} \in \arg \max_{\hat{x}} \sum_{x \in \mathcal{X}} g(x, \hat{x}) P_{X|Y=y}(x) \\ 0, & \text{otherwise,} \end{cases}$$

in the numerator, and $P_{\hat{X}}^*$ defined as

$$P_{\hat{X}}^*(\hat{x}) := \begin{cases} 1, & \text{for some } \hat{x} \in \arg \max_{\hat{x}} \sum_{x \in \mathcal{X}} g(x, \hat{x}) P_X(x), \\ 0, & \text{otherwise,} \end{cases}$$

in the denominator. Furthermore, as shown in the proof of Theorem 1, given a randomized function of X denoted by U , the U -leakage of X can be expressed as

$$\begin{aligned} \ell_U(X \rightarrow y) &= \log \frac{\max_{u \in \mathcal{U}} P_{U|Y=y}(u)}{\max_u P_U(u)} \\ &= \log \frac{\max_{u \in \mathcal{U}} \sum_{x \in \mathcal{X}} P_{U|X=x}(u) P_{X|Y=y}(x)}{\max_{u \in \mathcal{U}} \sum_{x \in \mathcal{X}} P_{U|X=x}(u) P_X(x)}. \end{aligned} \quad (30)$$

To show the equivalence, first, we prove the simpler direction by showing that each U -leakage can be written as a g -leakage. Given an arbitrary randomized function of X denoted by U , define $\hat{\mathcal{X}}_U := \mathcal{U}$ such that each $u \in \mathcal{U}$ corresponds uniquely to some $\hat{x}_u \in \hat{\mathcal{X}}_U$, and let $g_U(x, \hat{x}_u) := P_{U|X=x}(u)$ for all $x \in \mathcal{X}$ and $u \in \mathcal{U}$. By computing expressions (29) and (30), it is easy to see that $\ell_U(X \rightarrow y) = \ell_{g_U}(X \rightarrow y)$. This construction implies that a randomized function of X is simply a gain function that satisfies $\sum_{\hat{x}} g_U(x, \hat{x}) = 1$, for all $x \in \mathcal{X}$, that is, the total gain associated with each secret $x \in \mathcal{X}$ is a constant.

Now, we show that each g -leakage can be written as a U -leakage. Fix an arbitrary gain function g . Without loss of generality, suppose $g(x, \hat{x}) \leq 1$ for all $x \in \mathcal{X}$ and $\hat{x} \in \hat{\mathcal{X}}$ (this can be achieved by normalizing the gain function by $\max_{x, \hat{x}} g(x, \hat{x})$). In what follows, we construct a randomized function of X using a channel that generalizes the shattering channel of Definition 2. We need to consider the following two cases:

Case 1: The same \hat{x} maximizes the numerator and the denominator in (29).

Here, we will construct a randomized function of X , denoted by V , which is described by the kernel $P_{V|X}$ and satisfies $\ell_V(X \rightarrow y) = \ell_g(X \rightarrow y)$. Let $\hat{x}_V \in \arg \max_{\hat{x}} \sum_x g(x, \hat{x}) P_{X|Y=y}(x)$ which (following from the definition of Case 1) also satisfies $\hat{x}_V \in \arg \max_{\hat{x}} \sum_x g(x, \hat{x}) P_X(x)$. Informally, \hat{x}_V denotes the adversary's best guess both after observing $Y = y$, and with no observations.

Define the set $\mathcal{X}_V := \{x \in \mathcal{X} : g(x, \hat{x}_V) > 0\}$. For now, let us assume that $\mathcal{X}_V = \mathcal{X}$ but later we will discuss how the proof can be adapted if \mathcal{X}_V is a proper subset of \mathcal{X} . For each $x \in \mathcal{X}_V$, let $k_V(x) := 1/g(x, \hat{x}_V)$, and define $k_V := \max_{x \in \mathcal{X}_V} k_V(x)$. Roughly speaking, $k_V(x)$ is the cardinality of the support of $P_{V|X=x}$ while k_V is the cardinality of the support of P_V .

Now, we construct a random variable V taking values in an alphabet \mathcal{V} such that $|\mathcal{V}| = \lceil k_V \rceil$. For all $x \in \mathcal{X}_V$, the kernel $P_{V|X=x}$ is defined by

$$P_{V|X=x}(v_i) := \begin{cases} g(x, \hat{x}_V), & \text{if } 1 \leq i \leq \lfloor k_V(x) \rfloor, \\ 1 - \lfloor k_V(x) \rfloor g(x, \hat{x}_V), & \text{if } i = \lceil k_V(x) \rceil, \\ 0, & \text{if } \lceil k_V(x) \rceil + 1 \leq i \leq \lceil k_V \rceil. \end{cases}$$

Informally, for each x , the above kernel allocates chunks of probability equal to $g(x, \hat{x}_V)$ to the first $\lfloor k_V(x) \rfloor$ letters, and the $\lfloor k_V(x) \rfloor$ -th letter is used to contain the remaining probability $1 - \lfloor k_V(x) \rfloor g(x, \hat{x}_V)$. Note that the above kernel indeed satisfies

$$\sum_{i=1}^{\lfloor k_V \rfloor} P_{V|X=x}(v_i) = 1,$$

for all $x \in \mathcal{X}_V$.

Renaming V to U_g , we verify that the random variable constructed above satisfies $\ell_{U_g}(X \rightarrow y) = \ell_g(X \rightarrow y)$:

$$\begin{aligned} \max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_{X|Y=y}(x) &= \sum_x g(x, \hat{x}_{U_g}) P_{X|Y=y}(x) \\ &= \max_{\hat{x}} \sum_x g(x, \hat{x}) P_{X|Y=y}(x), \end{aligned}$$

where the last equality follows from the definition of \hat{x}_{U_g} . Similarly, we have

$$\begin{aligned} \max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_X(x) &= \sum_x g(x, \hat{x}_{U_g}) P_X(x) \\ &= \max_{\hat{x}} \sum_x g(x, \hat{x}) P_X(x). \end{aligned}$$

Case 2: The maximizing \hat{x} 's in the numerator and the denominator of (29) are different.

We will construct two randomized functions of X denoted by V and W , one for the numerator of (29) and one for the denominator. Let

$$\begin{aligned} \hat{x}_V &\in \arg \max_{\hat{x}} \sum_x g(x, \hat{x}) P_{X|Y=y}(x), \\ \hat{x}_W &\in \arg \max_{\hat{x}} \sum_x g(x, \hat{x}) P_X(x), \end{aligned}$$

and define

$$\begin{aligned} \mathcal{X}_V &:= \{x \in \mathcal{X} : g(x, \hat{x}_V) > 0\}, \\ \mathcal{X}_W &:= \{x \in \mathcal{X} : g(x, \hat{x}_W) > 0\}, \end{aligned}$$

where \hat{x}_V denotes the adversary's best guess having observed $Y = y$, and \hat{x}_W denotes the adversary's best guess without an observation. We need to consider the following two cases:

Case 2.1: $\mathcal{X}_V = \mathcal{X}_W$.

Let $\mathcal{X}_{U_g} := \mathcal{X}_V = \mathcal{X}_W$. Once again, let us assume that $\mathcal{X}_{U_g} = \mathcal{X}$. Similarly to what we had in Case 1, for all $x \in \mathcal{X}_{U_g}$ we define

$$\begin{aligned} k_V(x) &:= 1/g(x, \hat{x}_V), & k_V &:= \max_{x \in \mathcal{X}_{U_g}} k_V(x), \\ k_W(x) &:= 1/g(x, \hat{x}_W), & k_W &:= \max_{x \in \mathcal{X}_{U_g}} k_W(x). \end{aligned}$$

Let \mathcal{V} denote the support set of random variable V , and \mathcal{W} denote the support set of random variable W , where $|\mathcal{V}| = \lceil k_V \rceil$ and $|\mathcal{W}| = \lceil k_W \rceil$. For all $x \in \mathcal{X}_{U_g}$, the kernels $P_{V|X=x}$ and $P_{W|X=x}$ are defined as

$$P_{V|X=x}(v_i) := \begin{cases} g(x, \hat{x}_V), & \text{if } 1 \leq i \leq \lfloor k_V(x) \rfloor, \\ 1 - \lfloor k_V(x) \rfloor g(x, \hat{x}_V), & \text{if } i = \lceil k_V(x) \rceil, \\ 0, & \text{if } \lceil k_V(x) \rceil + 1 \leq i \leq \lceil k_V \rceil, \end{cases}$$

and

$$P_{W|X=x}(w_j) := \begin{cases} g(x, \hat{x}_W), & \text{if } 1 \leq j \leq \lfloor k_W(x) \rfloor, \\ 1 - \lfloor k_W(x) \rfloor g(x, \hat{x}_W), & \text{if } j = \lceil k_W(x) \rceil, \\ 0, & \text{if } \lceil k_W(x) \rceil + 1 \leq j \leq \lceil k_W \rceil. \end{cases}$$

Finally, we define the random variable U_g as the Bernoulli mixture of V and W . Let $\mathcal{U}_g := \mathcal{V} \cup \mathcal{W}$ denote the alphabet of U_g . For all $x \in \mathcal{X}_{U_g}$, we define $P_{U_g|X=x}(u) := \frac{1}{2}P_{V|X=x}(u) + \frac{1}{2}P_{W|X=x}(u)$, where $P_{V|X=x}(u) = 0$ for $u \in \mathcal{W}$ and $P_{W|X=x}(u) = 0$ for $u \in \mathcal{V}$.⁵ Let us verify that U_g satisfies $\ell_{U_g}(X \rightarrow y) = \ell_g(X \rightarrow y)$:

$$\begin{aligned} \max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_{X|Y=y}(x) &= \sum_x \frac{1}{2} g(x, \hat{x}_V) P_{X|Y=y}(x) \\ &= \frac{1}{2} \max_{\hat{x}} \sum_x g(x, \hat{x}) P_{X|Y=y}(x), \end{aligned}$$

and also,

$$\begin{aligned} \max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_X(x) &= \sum_x \frac{1}{2} g(x, \hat{x}_W) P_X(x) \\ &= \frac{1}{2} \max_{\hat{x}} \sum_x g(x, \hat{x}) P_X(x). \end{aligned}$$

Thus, we have

$$\begin{aligned} \ell_{U_g}(X \rightarrow y) &= \log \frac{\max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_{X|Y=y}(x)}{\max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_X(x)} \\ &= \log \frac{\frac{1}{2} \max_{\hat{x}} \sum_x g(x, \hat{x}) P_{X|Y=y}(x)}{\frac{1}{2} \max_{\hat{x}} \sum_x g(x, \hat{x}) P_X(x)} \\ &= \ell_g(X \rightarrow y). \end{aligned}$$

Case 2.2: $\mathcal{X}_V \neq \mathcal{X}_W$.

Let n_V and n_W be positive integers. Here, the idea is that we increase the sizes of the sets \mathcal{V} and \mathcal{W} by n_V and n_W , respectively, where these extra letters are used to support those x 's for which we either have $g(x, \hat{x}_V) = 0$ or $g(x, \hat{x}_W) = 0$. We need to distinguish between three types of x 's:

⁵This is a slight abuse of notation. Strictly speaking, $P_{V|X=x}(u)$ is defined only for $u \in \mathcal{V}$ and $P_{W|X=x}(u)$ is defined only for $u \in \mathcal{W}$.

$$\begin{array}{c}
\begin{array}{cccccccccccc}
\overbrace{\hspace{2cm}}^{n_V} & & \overbrace{\hspace{4cm}}^{k_V} & & \overbrace{\hspace{4cm}}^{k_W} & & \overbrace{\hspace{2cm}}^{n_W} \\
\left[\begin{array}{cccccccccccc}
0 & \cdots & 0 & \frac{1}{2}g(x_1, \hat{x}_V) & \frac{1}{2}g(x_1, \hat{x}_V) & \cdots & \frac{1}{2}g(x_1, \hat{x}_V) & \frac{1 - \lfloor k_V(x_1) \rfloor g(x_1, \hat{x}_V)}{2} & \frac{1}{2}g(x_1, \hat{x}_W) & \frac{1}{2}g(x_1, \hat{x}_W) & \cdots & 0 & 0 & 0 & \cdots & 0 \\
0 & \cdots & 0 & \frac{1}{2}g(x_2, \hat{x}_V) & \frac{1}{2}g(x_2, \hat{x}_V) & \cdots & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \frac{1}{2n_W} & \cdots & \frac{1}{2n_W} \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
\frac{1}{2n_V} & \cdots & \frac{1}{2n_V} & 0 & 0 & \cdots & 0 & 0 & \frac{1}{2}g(x_l, \hat{x}_W) & \frac{1}{2}g(x_l, \hat{x}_W) & \cdots & 0 & 0 & 0 & \cdots & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
0 & \cdots & 0 & \frac{1}{2}g(x_m, \hat{x}_V) & \frac{1}{2}g(x_m, \hat{x}_V) & \cdots & \frac{1 - \lfloor k_V(x_m) \rfloor g(x_m, \hat{x}_V)}{2} & 0 & \frac{1}{2}g(x_m, \hat{x}_W) & \frac{1}{2}g(x_m, \hat{x}_W) & \cdots & \frac{1}{2}g(x_m, \hat{x}_W) & \frac{1 - \lfloor k_W(x_m) \rfloor g(x_m, \hat{x}_W)}{2} & 0 & \cdots & 0
\end{array} \right] \\
\hline
\mathcal{V} & & & & \mathcal{W} & & & & & & & & & & & &
\end{array}
\end{array}$$

Fig. 3. Illustration of the channel $P_{U_g|X}$ constructed in Case 2.2. We assume $\mathcal{X} = \mathcal{X}_V \cup \mathcal{X}_W$, where \mathcal{X} contains m elements. Each row of the above matrix corresponds to an x from \mathcal{X} , and the columns correspond to letters from \mathcal{U}_g . In this example, we have $k_V(x_1) = k_V$ and $k_W(x_m) = k_W$. Also, we are assuming $x_2 \in \mathcal{X}_V \setminus \mathcal{X}_W$ and $x_l \in \mathcal{X}_W \setminus \mathcal{X}_V$, where $2 < l < m$.

1) For $x \in \mathcal{X}_V \cap \mathcal{X}_W$ we define

$$P_{V|X=x}(v_i) := \begin{cases} g(x, \hat{x}_V), & \text{if } 1 \leq i \leq \lfloor k_V(x) \rfloor, \\ 1 - \lfloor k_V(x) \rfloor g(x, \hat{x}_V), & \text{if } i = \lceil k_V(x) \rceil, \\ 0, & \text{if } \lceil k_V(x) \rceil + 1 \leq i \leq \lceil k_V \rceil + n_V, \end{cases}$$

and

$$P_{W|X=x}(w_j) := \begin{cases} g(x, \hat{x}_W), & \text{if } 1 \leq j \leq \lfloor k_W(x) \rfloor, \\ 1 - \lfloor k_W(x) \rfloor g(x, \hat{x}_W), & \text{if } j = \lceil k_W(x) \rceil, \\ 0, & \text{if } \lceil k_W(x) \rceil + 1 \leq j \leq \lceil k_W \rceil + n_W. \end{cases}$$

2) For $x \in \mathcal{X}_V \setminus \mathcal{X}_W$ we let

$$P_{V|X=x}(v_i) := \begin{cases} g(x, \hat{x}_V), & \text{if } 1 \leq i \leq \lfloor k_V(x) \rfloor, \\ 1 - \lfloor k_V(x) \rfloor g(x, \hat{x}_V), & \text{if } i = \lceil k_V(x) \rceil, \\ 0, & \text{if } \lceil k_V(x) \rceil + 1 \leq i \leq \lceil k_V \rceil + n_V, \end{cases}$$

and

$$P_{W|X=x}(w_j) := \begin{cases} 0, & \text{if } 1 \leq j \leq \lceil k_W \rceil, \\ \frac{1}{n_W}, & \text{if } \lceil k_W \rceil + 1 \leq j \leq \lceil k_W \rceil + n_W. \end{cases}$$

3) For $x \in \mathcal{X}_W \setminus \mathcal{X}_V$ we let

$$P_{V|X=x}(v_i) := \begin{cases} 0, & \text{if } 1 \leq i \leq \lceil k_V \rceil, \\ \frac{1}{n_V}, & \text{if } \lceil k_V \rceil + 1 \leq i \leq \lceil k_V \rceil + n_V, \end{cases}$$

and

$$P_{W|X=x}(w_j) := \begin{cases} g(x, \hat{x}_W), & \text{if } 1 \leq j \leq \lfloor k_W(x) \rfloor, \\ 1 - \lfloor k_W(x) \rfloor g(x, \hat{x}_W), & \text{if } j = \lfloor k_W(x) \rfloor + 1, \\ 0, & \text{if } \lfloor k_W(x) \rfloor + 1 \leq j \leq \lfloor k_W(x) \rfloor + n_W. \end{cases}$$

Now, we define U_g as before. Suppose $\mathcal{U}_g = \mathcal{V} \cup \mathcal{W}$ is the alphabet of U_g . For $x \in \mathcal{X}_V \cup \mathcal{X}_W$ (where we are assuming $\mathcal{X} = \mathcal{X}_V \cup \mathcal{X}_W$), let $P_{U_g|X=x}(u) = \frac{1}{2}P_{V|X=x}(u) + \frac{1}{2}P_{W|X=x}(u)$, where $P_{V|X=x}(u) = 0$ for $u \in \mathcal{W}$ and $P_{W|X=x}(u) = 0$ for $u \in \mathcal{V}$. Then, we can write

$$\max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_{X|Y=y}(x) = \max \left\{ \sum_{x \in \mathcal{X}_V} \frac{g(x, \hat{x}_V)}{2} P_{X|Y=y}(x), \sum_{x \in \mathcal{X}_W \setminus \mathcal{X}_V} \frac{1}{2n_V} P_{X|Y=y}(x) \right\}.$$

By taking n_V to be large enough, we can ensure that

$$\sum_{x \in \mathcal{X}_W \setminus \mathcal{X}_V} \frac{1}{2n_V} P_{X|Y=y}(x) \leq \sum_{x \in \mathcal{X}_V} \frac{1}{2} g(x, \hat{x}_V) P_{X|Y=y}(x),$$

which yields

$$\begin{aligned} \max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_{X|Y=y}(x) &= \sum_{x \in \mathcal{X}_V} \frac{1}{2} g(x, \hat{x}_V) P_{X|Y=y}(x) \\ &= \frac{1}{2} \max_{\hat{x}} \sum_{x \in \mathcal{X}} g(x, \hat{x}) P_{X|Y=y}(x). \end{aligned}$$

Similarly, for n_W large enough we have

$$\begin{aligned} \max_{u \in \mathcal{U}_g} \sum_x P_{U_g|X=x}(u) P_X(x) &= \max \left\{ \sum_{x \in \mathcal{X}_W} \frac{1}{2} g(x, \hat{x}_W) P_X(x), \sum_{x \in \mathcal{X}_V \setminus \mathcal{X}_W} \frac{1}{2n_W} P_X(x) \right\} \\ &= \sum_{x \in \mathcal{X}_W} \frac{1}{2} g(x, \hat{x}_W) P_X(x) \\ &= \frac{1}{2} \max_{\hat{x}} \sum_{x \in \mathcal{X}} g(x, \hat{x}) P_X(x). \end{aligned}$$

Hence, we conclude that $\ell_{U_g}(X \rightarrow y) = \ell_g(X \rightarrow y)$.

The only point left to discuss is regarding the case where $\mathcal{X}_V \cup \mathcal{X}_W$ is a proper subset of \mathcal{X} . Let n_O be a positive integer. Once again, we increase the size of the alphabet \mathcal{U}_g by n_O letters, where these extra letters are used to support the x 's in $\mathcal{X} \setminus (\mathcal{X}_V \cup \mathcal{X}_W)$. Hence, we let $\mathcal{U}_g = \mathcal{V} \cup \mathcal{W} \cup \mathcal{O}$, where \mathcal{O} is a finite set containing n_O elements. For $x \in \mathcal{X} \setminus (\mathcal{X}_V \cup \mathcal{X}_W)$ we define the channel $P_{U_g|X=x}$ as

$$P_{U_g|X=x}(u) = \begin{cases} \frac{1}{n_O}, & \text{if } u \in \mathcal{O}, \\ 0, & \text{otherwise.} \end{cases}$$

For $x \in \mathcal{X}_V \cup \mathcal{X}_W$, we let $P_{U_g|X=x}(u) = 0$ when $u \in \mathcal{O}$; otherwise $P_{U_g|X=x}(u)$ is defined as in Case 2.2. It is straightforward to verify that for n_O large enough ($\frac{1}{n_O}$ small enough), the values of the numerator and the denominator in the expression of $\ell_{U_g}(X \rightarrow y)$ remain as before, from which we may conclude $\ell_{U_g}(X \rightarrow y) = \ell_g(X \rightarrow y)$. ■

APPENDIX B
PROOF OF LEMMA 1

1) Upper bound:

$$\begin{aligned}\ell(X \rightarrow y) &= \log \max_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)} \\ &\leq \log \max_{x \in \text{supp}(P_X)} \frac{1}{P_X(x)} = \log \frac{1}{\min_{x \in \text{supp}(P_X)} P_X(x)},\end{aligned}$$

where the inequality holds with equality if and only if $P_{X|Y=y}(x^*) = 1$ with $x^* \in \arg \min_{x \in \text{supp}(P_X)} P_X(x)$.

Lower bound: Here, the idea is that since both $P_{X|Y=y}$ and P_X are probability distributions over $\text{supp}(P_X)$, then for any fixed $y \in \text{supp}(P_Y)$, there exists at least one $x \in \text{supp}(P_X)$ such that $P_{X|Y=y}(x) \geq P_X(x)$. Suppose to the contrary that for all $x \in \mathcal{X}$, $P_{X|Y=y}(x) < P_X(x)$. Then, $1 = \sum_x P_{X|Y=y}(x) < \sum_x P_X(x) = 1$ which is a contradiction. Therefore,

$$\ell(X \rightarrow y) = \log \max_{x \in \text{supp}(P_X)} \frac{P_{X|Y=y}(x)}{P_X(x)} \geq \log 1 = 0.$$

The above inequality holds with equality if and only if $\max_x P_{Y|X=x}(y) = P_Y(y) = \sum_x P_{Y|X=x}(y)P_X(x)$ which holds whenever $P_{Y|X=x}(y) = P_{Y|X=x'}(y)$ for all $x, x' \in \text{supp}P_X$.

2) Both statements follow directly from the definition.

3) Let $x^* \in \arg \max_x P_{Z|X=x}(z)$. Then,

$$\begin{aligned}\ell(X \rightarrow z) &= \log \max_x \frac{P_{Z|X=x}(z)}{P_Z(z)} \\ &= \log \frac{\sum_{y \in \text{supp}(P_{Y|X=x^*})} P_{Z|Y=y}(z)P_{Y|X=x^*}(y)}{P_Z(z)} \\ &\leq \log \max_{y' \in \text{supp}(P_{Y|X=x^*})} \frac{P_{Z|Y=y'}(z)}{P_Z(z)} \sum_y P_{Y|X=x^*}(y) \\ &\leq \log \max_{y' \in \text{supp}(P_Y)} \frac{P_{Z|Y=y'}(z)}{P_Z(z)} = \ell(Y \rightarrow z),\end{aligned}$$

where the first inequality holds with equality if $P_{Z|Y=y}(z) = P_{Z|Y=y'}(z)$ for all $y, y' \in \text{supp}(P_{Y|X=x^*})$, and the second inequality holds with equality if $\max_{y \in \text{supp}(P_Y)} i(y; z)$ is attained at some $y \in \text{supp}(P_{Y|X=x^*})$.

4)

$$\begin{aligned}\ell(X \rightarrow z) &= \log \max_x \frac{P_{Z|X=x}(z)}{P_Z(z)} \\ &= \log \max_x \frac{\sum_{y \in \text{supp}(P_Y)} P_{Z|Y=y}(z)P_{Y|X=x}(y)}{\sum_{y \in \text{supp}(P_Y)} P_{Z|Y=y}(z)P_Y(y)} \\ &\leq \log \max_x \max_{y \in \text{supp}(P_Y)} \frac{P_{Y|X=x}(y)}{P_Y(y)} \\ &= \max_{y \in \text{supp}(P_Y)} \ell(X \rightarrow y).\end{aligned}$$

Now, if X and Y are independent then $\ell(X \rightarrow z) = \ell(X \rightarrow y) = 0$ for all y, z , and the inequality holds with equality. Furthermore, if $P_{Y|Z=z}$ is deterministic, then z is mapped uniquely to some $y_z \in \text{supp}(P_Y)$. This implies that $P_{Z|Y=y}(z) = 0$ for $y \neq y_z$, hence, we have

$$\begin{aligned}\ell(X \rightarrow z) &= \log \max_x \frac{P_{Z|X=x}(z)}{P_Z(z)} \\ &= \log \max_x \frac{P_{Y|X=x}(y_z)}{P_Y(y_z)} \\ &= \ell(X \rightarrow y_z) \\ &\leq \max_{y \in \text{supp}(P_Y)} \ell(X \rightarrow y),\end{aligned}$$

with equality if $\ell(X \rightarrow y_z) = \max_{y \in \text{supp}(P_Y)} \ell(X \rightarrow y)$.

5)

$$\begin{aligned}\ell(X \rightarrow y | z) &= \log \max_x \frac{P_{Y|X=x, Z=z}(y)}{P_{Y|Z=z}(y)} \\ &= \log \max_x \frac{P_{Y|X=x}(y) P_Y(y)}{P_Y(y) P_{Y|Z=z}(y)} \\ &= \log \max_x \frac{P_{Y|X=x}(y)}{P_Y(y)} + \log \frac{P_Y(y)}{P_{Y|Z=z}(y)} \\ &= \ell(X \rightarrow y) - i(y; z).\end{aligned}$$

6)

$$\begin{aligned}\ell(X \rightarrow y, z) &= \log \max_x \frac{P_{YZ|X=x}(y, z)}{P_{YZ}(y, z)} \\ &= \log \max_x \frac{P_{Y|X=x, Z=z}(y) P_{Z|X=x}(z)}{P_{Y|Z=z}(y) P_Z(z)} \\ &\leq \log \max_x \frac{P_{Y|X=x, Z=z}(y)}{P_{Y|Z=z}(y)} + \log \max_x \frac{P_{Z|X=x}(z)}{P_Z(z)} \\ &= \ell(X \rightarrow y | z) + \ell(X \rightarrow z),\end{aligned}$$

with equality if and only if there exists $x^* \in \text{supp}(P_X)$ maximizing both $i(x; y | z)$ and $i(x; z)$. ■

APPENDIX C

PROOF OF THEOREM 4

1) This result is an immediate consequence of the composition property given in Lemma 1. For all $y \in \text{supp}(P_Y)$ and all $z \in \text{supp}(P_Z)$ we have

$$\begin{aligned}\ell(X \rightarrow y, z) &\leq \ell(X \rightarrow y) + \ell(X \rightarrow z | y) \\ &\leq \max_{y \in \text{supp}(P_Y)} \ell(X \rightarrow y) + \max_{\substack{y \in \text{supp}(P_Y), \\ z \in \text{supp}(P_Z)}} \ell(X \rightarrow z | y) \\ &\leq \epsilon_1 + \epsilon_2.\end{aligned}$$

Therefore, $P_{Y|Z|X}$ satisfies $\epsilon_1 + \epsilon_2$ -PML.

2) Since $\ell(X \rightarrow y, z) \leq \ell(X \rightarrow y) + \ell(X \rightarrow z | y)$ for all $(y, z) \in \text{supp}(P_Y) \times \text{supp}(P_Z)$, we can write

$$\begin{aligned} \mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow Y, Z) > \epsilon_1 + \epsilon_2 \right] &\leq \mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow Y) + \ell(X \rightarrow Z | Y) > \epsilon_1 + \epsilon_2 \right] \\ &= 1 - \mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow Y) + \ell(X \rightarrow Z | Y) \leq \epsilon_1 + \epsilon_2 \right]. \end{aligned}$$

We define the following ‘‘good’’ events:

$$\mathcal{G} := \{(y, z) \in \text{supp}(P_Y) \times \text{supp}(P_Z) : \ell(X \rightarrow y) \leq \epsilon_1 \text{ and } \ell(X \rightarrow z | y) \leq \epsilon_2\},$$

$$\mathcal{G}_Y := \{y \in \text{supp}(P_Y) : (y, z) \in \mathcal{G} \text{ for some } z \in \text{supp}(P_Z)\},$$

$$\mathcal{G}_Z(y) := \{z \in \text{supp}(P_Z) : (y, z) \in \mathcal{G}\}.$$

Our goal is to lower bound the probability of event \mathcal{G} . We can write

$$\begin{aligned} P_{YZ}(\mathcal{G}) &= \sum_{(y,z) \in \mathcal{G}} P_Y(y) P_{Z|Y=y}(z) = \sum_{y \in \mathcal{G}_Y} P_Y(y) P_{Z|Y=y}(\mathcal{G}_Z(y)) \\ &\geq (1 - \delta_2) \sum_{y \in \mathcal{G}_Y} P_Y(y) \end{aligned} \tag{31a}$$

$$\geq (1 - \delta_2)(1 - \delta_1), \tag{31b}$$

where

- (31a) follows from the fact that for all $y \in \text{supp}(P_Y)$, $P_{Z|X,Y=y}$ satisfies (ϵ_2, δ_2) -PML which implies that

$$P_{Z|Y=y}(\mathcal{G}_Z(y)) = \mathbb{P}_{Z \sim P_{Z|Y=y}} \left[\ell(X \rightarrow Z | y) \leq \epsilon_2 \right] \geq 1 - \delta_2,$$

- and (31b) follows from the fact that $P_{Y|X}$ satisfies (ϵ_1, δ_1) -PML, that is,

$$P_Y(\mathcal{G}_Y) = \mathbb{P}_{Y \sim P_Y} \left[\ell(X \rightarrow Y) \leq \epsilon_1 \right] \geq 1 - \delta_1.$$

It follows that

$$\mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow y) + \ell(X \rightarrow z | y) \leq \epsilon_1 + \epsilon_2 \right] \geq P_{YZ}(\mathcal{G}) \geq (1 - \delta_2)(1 - \delta_1),$$

which yields

$$\mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow Y, Z) > \epsilon_1 + \epsilon_2 \right] \leq \delta_1 + \delta_2 - \delta_1 \delta_2.$$

3) Define the event

$$\mathcal{A}_Y = \{y \in \text{supp}(P_Y) : \ell(X \rightarrow y) \leq \epsilon_1\}.$$

As $P_{Y|X}$ satisfies (ϵ_1, δ_1) -PML, we have

$$1 - \delta_1 \leq P_Y(\mathcal{A}_Y) = \sum_{y \in \mathcal{A}_Y} P_Y(y) = \sum_{(y,z) \in \mathcal{A}_Y \times \text{supp}(P_Z)} P_{YZ}(y, z) = P_{YZ}(\mathcal{A}),$$

where $\mathcal{A} := \mathcal{A}_Y \times \text{supp}(P_Z)$. Moreover, we define the event

$$\mathcal{B} := \{(y, z) \in \text{supp}(P_Y) \times \text{supp}(P_Z) : \ell(X \rightarrow z | y) \leq \epsilon_2\}.$$

By assumption, $P_{YZ}(\mathcal{B}) \geq 1 - \delta_2$. Therefore,

$$\begin{aligned} \mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow Y, Z) \leq \epsilon_1 + \epsilon_2 \right] &\geq \mathbb{P}_{(Y,Z) \sim P_{YZ}} \left[\ell(X \rightarrow Y) + \ell(X \rightarrow Z | Y) \leq \epsilon_1 + \epsilon_2 \right] \\ &\geq P_{YZ}(\mathcal{A} \cap \mathcal{B}) \\ &= 1 - P_{YZ}(\mathcal{A}^c \cup \mathcal{B}^c) \\ &\geq 1 - \delta_1 - \delta_2. \end{aligned}$$

4) Let $\mathcal{E} \subseteq \text{supp}(P_{YZ})$ be an event satisfying $P_{YZ}(\mathcal{E}) \geq \delta_1$ and $0 \leq \delta_2 \leq \min_{y \in \mathcal{E}_Y} P_{Z|Y=y}(\mathcal{E}_Z(y))$. Since $P_{Z|X,Y=y}$ satisfies (ϵ_2, δ_2) -EML for all $y \in \mathcal{E}_Y$, we have

$$\max_{x \in \text{supp}(P_X)} \frac{P_{Z|Y=y, X=x}(\mathcal{E}_Z(y))}{P_{Z|Y=y}(\mathcal{E}_Z(y))} \leq \exp(\epsilon_2). \quad (32)$$

Now, we write

$$\begin{aligned} \exp(\ell(X \rightarrow \mathcal{E})) &= \max_{x \in \text{supp}(P_X)} \frac{P_{YZ|X=x}(\mathcal{E})}{P_{YZ}(\mathcal{E})} \\ &= \max_{x \in \text{supp}(P_X)} \frac{\sum_{y \in \mathcal{E}_Y} \sum_{z \in \mathcal{E}_Z(y)} P_{YZ|X=x}(y, z)}{\sum_{y \in \mathcal{E}_Y} \sum_{z \in \mathcal{E}_Z(y)} P_{YZ}(y, z)} \\ &= \max_{x \in \text{supp}(P_X)} \frac{\sum_{y \in \mathcal{E}_Y} P_{Y|X=x}(y) \sum_{z \in \mathcal{E}_Z(y)} P_{Z|Y=y, X=x}(z)}{\sum_{y \in \mathcal{E}_Y} P_Y(y) \sum_{z \in \mathcal{E}_Z(y)} P_{Z|Y=y}(z)} \\ &= \max_{x \in \text{supp}(P_X)} \sum_{y \in \mathcal{E}_Y} \frac{P_Y(y) P_{Z|Y=y}(\mathcal{E}_Z(y))}{\sum_{y' \in \mathcal{E}_Y} P_Y(y') P_{Z|Y=y'}(\mathcal{E}_Z(y'))} \left(\frac{P_{Y|X=x}(y)}{P_Y(y)} \right) \left(\frac{P_{Z|Y=y, X=x}(\mathcal{E}_Z(y))}{P_{Z|Y=y}(\mathcal{E}_Z(y))} \right) \\ &\leq \exp(\epsilon_2) \max_{x \in \text{supp}(P_X)} \sum_{y \in \mathcal{E}_Y} \frac{P_Y(y) P_{Z|Y=y}(\mathcal{E}_Z(y))}{\sum_{y' \in \mathcal{E}_Y} P_Y(y') P_{Z|Y=y'}(\mathcal{E}_Z(y'))} \left(\frac{P_{Y|X=x}(y)}{P_Y(y)} \right) \end{aligned} \quad (33a)$$

$$\leq \exp(\epsilon_2) \max_{x \in \text{supp}(P_X)} h_x(P_{Y|X=x}, \delta_1) \quad (33b)$$

$$\leq \exp(\epsilon_2 + \epsilon_1), \quad (33c)$$

where

- (33a) follows from inequality (32),
- the function h_x in (33b) is defined in (16),
- and (33c) follows from the fact that $P_{Y|X}$ satisfies (ϵ_1, δ_1) -EML.

5) Let $\mathcal{E} \subseteq \text{supp}(P_{YZ})$ be an event satisfying $P_{YZ}(\mathcal{E}) \geq \delta_1 + \delta_2$. We define the following “bad” sets

$$\mathcal{B}_Y := \{y \in \mathcal{E}_Y : P_Z(\mathcal{E}_Z(y)) < \delta_2\},$$

$$\mathcal{B} := \{(y, z) \in \mathcal{E} : y \in \mathcal{B}_Y\},$$

and the “good” sets $\mathcal{G}_Y = \mathcal{E}_Y \setminus \mathcal{B}_Y$ and $\mathcal{G} = \mathcal{E} \setminus \mathcal{B}$. Note that

$$P_{YZ}(\mathcal{B}) = \sum_{y \in \mathcal{B}_Y} P_Y(y) P_{Z|Y=y}(\mathcal{E}_Z(y)) < \delta_2,$$

which implies that $P_{YZ}(\mathcal{G}) = P_{YZ}(\mathcal{E}) - P_{YZ}(\mathcal{B}) > \delta_1$. Now, similarly to the previous part, we write

$$\begin{aligned} \exp(\ell(X \rightarrow \mathcal{E})) &= \max_{x \in \text{supp}(P_X)} \frac{\sum_{(y,z) \in \mathcal{E}} P_{YZ|X=x}(y,z)}{P_{YZ}(\mathcal{E})} \\ &\leq \max_{x \in \text{supp}(P_X)} \frac{\sum_{(y,z) \in \mathcal{B}} P_{YZ}(y,z) \left(\frac{P_{YZ|X=x}(y,z)}{P_{YZ}(y,z)} \right)}{P_{YZ}(\mathcal{E})} + \max_{x \in \text{supp}(P_X)} \frac{\sum_{(y,z) \in \mathcal{G}} P_{YZ|X=x}(y,z)}{P_{YZ}(\mathcal{E})} \\ &\leq \frac{\delta_2}{\delta_1 + \delta_2} \exp(\epsilon_{\max}) + \max_{x \in \text{supp}(P_X)} \frac{\sum_{(y,z) \in \mathcal{G}} P_{YZ|X=x}(y,z)}{P_{YZ}(\mathcal{G})} \end{aligned} \quad (34a)$$

$$\leq \frac{\delta_2}{\delta_1 + \delta_2} \exp(\epsilon_{\max}) + \exp(\epsilon_1 + \epsilon_2), \quad (34b)$$

where

- (34a) follows from the fact that $P_{YZ}(\mathcal{B}) < \delta_2$, $P_{YZ}(\mathcal{E}) \geq \delta_1 + \delta_2$, and for all $(y,z) \in \text{supp}(P_{YZ})$,

$$\max_{x \in \text{supp}(P_X)} \frac{P_{YZ|X=x}(y,z)}{P_{YZ}(y,z)} \leq \exp(\epsilon_{\max}),$$

- and (34b) follows from the definition of the set \mathcal{G} and the previous part. ■

REFERENCES

- [1] 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.
- [2] C. Dwork, “Differential privacy,” in *Automata, Languages and Programming*, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 1–12.
- [3] S. L. Warner, “Randomized response: A survey technique for eliminating evasive answer bias,” *Journal of the American Statistical Association*, vol. 60, no. 309, pp. 63–69, 1965.
- [4] A. Evfimievski, J. Gehrke, and R. Srikant, “Limiting privacy breaches in privacy preserving data mining,” in *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 2003, pp. 211–222.
- [5] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, “What can we learn privately?” *SIAM Journal on Computing*, vol. 40, no. 3, pp. 793–826, 2011.
- [6] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Local privacy and statistical minimax rates,” in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, 2013, pp. 429–438.
- [7] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, “Information extraction under privacy constraints,” *Information*, vol. 7, no. 1, p. 15, 2016.
- [8] S. Asoodeh, F. Alajaji, and T. Linder, “On maximal correlation, mutual information and data privacy,” in *2015 IEEE 14th Canadian workshop on information theory (CWIT)*. IEEE, 2015, pp. 27–31.
- [9] W. Wang, L. Ying, and J. Zhang, “On the relation between identifiability, differential privacy, and mutual-information privacy,” *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5018–5029, 2016.
- [10] A. Makhdoomi, S. Salamatian, N. Fawaz, and M. Médard, “From the information bottleneck to the privacy funnel,” in *2014 IEEE Information Theory Workshop (ITW 2014)*. IEEE, 2014, pp. 501–505.
- [11] S. Verdú, “ α -mutual information,” in *2015 Information Theory and Applications Workshop (ITA)*. IEEE, 2015, pp. 1–6.
- [12] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, “Tunable measures for information leakage and applications to privacy-utility tradeoffs,” *IEEE Transactions on Information Theory*, vol. 65, no. 12, pp. 8043–8066, 2019.
- [13] M. Diaz, H. Wang, F. P. Calmon, and L. Sankar, “On the robustness of information-theoretic privacy measures and mechanisms,” *IEEE Transactions on Information Theory*, vol. 66, no. 4, pp. 1949–1978, 2019.
- [14] B. Rassouli and D. Gündüz, “Optimal utility-privacy trade-off with total variation distance as a privacy measure,” *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 594–603, 2019.

- [15] 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.
- [16] B. Jiang, M. Seif, R. Tandon, and M. Li, "Context-aware local information privacy," *IEEE Transactions on Information Forensics and Security*, 2021.
- [17] F. P. Calmon and N. Fawaz, "Privacy against statistical inference," in *2012 50th annual Allerton conference on communication, control, and computing (Allerton)*. IEEE, 2012, pp. 1401–1408.
- [18] I. Wagner and D. Eckhoff, "Technical privacy metrics: a systematic survey," *ACM Computing Surveys (CSUR)*, vol. 51, no. 3, pp. 1–38, 2018.
- [19] M. Bloch, O. Günlü, A. Yener, F. Oggier, H. V. Poor, L. Sankar, and R. F. Schaefer, "An overview of information-theoretic security and privacy: Metrics, limits and applications," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 5–22, 2021.
- [20] J. Lee and C. Clifton, "Differential identifiability," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2012, pp. 1041–1049.
- [21] D. Kifer and A. Machanavajjhala, "No free lunch in data privacy," in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, 2011, pp. 193–204.
- [22] M. C. Tschantz, S. Sen, and A. Datta, "Sok: Differential privacy as a causal property," in *2020 IEEE Symposium on Security and Privacy (SP)*, 2020, pp. 354–371.
- [23] 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.
- [24] C. Braun, K. Chatzikokolakis, and C. Palamidessi, "Quantitative notions of leakage for one-try attacks," *Electronic Notes in Theoretical Computer Science*, vol. 249, pp. 75–91, 2009.
- [25] B. Espinoza and G. Smith, "Min-entropy as a resource," *Information and Computation*, vol. 226, pp. 57–75, 2013.
- [26] M. S. Alvim, K. Chatzikokolakis, C. Palamidessi, and G. Smith, "Measuring information leakage using generalized gain functions," in *2012 IEEE 25th Computer Security Foundations Symposium*, 2012, pp. 265–279.
- [27] M. S. Alvim, K. Chatzikokolakis, A. McIver, C. Morgan, C. Palamidessi, and G. Smith, "Additive and multiplicative notions of leakage, and their capacities," in *2014 IEEE 27th Computer Security Foundations Symposium*, 2014, pp. 308–322.
- [28] M. S. Alvim, K. Chatzikokolakis, A. McIver, C. Morgan, C. Palamidessi, and G. Smith, *The Science of Quantitative Information Flow*. Springer, 2020.
- [29] S. Saeidian, G. Cervia, T. J. Oechtering, and M. Skoglund, "Quantifying membership privacy via information leakage," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 3096–3108, 2021.
- [30] J. Liao, L. Sankar, F. P. Calmon, and V. Y. Tan, "Hypothesis testing under maximal leakage privacy constraints," in *2017 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2017, pp. 779–783.
- [31] B. Wu, A. B. Wagner, and G. E. Suh, "Optimal mechanisms under maximal leakage," in *2020 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 2020, pp. 1–6.
- [32] S. Saeidian, G. Cervia, T. J. Oechtering, and M. Skoglund, "Optimal maximal leakage-distortion tradeoff," in *2021 IEEE Information Theory Workshop (ITW)*. IEEE, 2021, pp. 1–6.
- [33] I. Issa, A. B. Wagner, and S. Kamath, "An operational approach to information leakage," *IEEE Transactions on Information Theory*, vol. 66, no. 3, pp. 1625–1657, 2019.
- [34] C. Dwork and G. N. Rothblum, "Concentrated differential privacy," *arXiv preprint arXiv:1603.01887*, 2016.
- [35] M. Bun and T. Steinke, "Concentrated differential privacy: Simplifications, extensions, and lower bounds," in *Theory of Cryptography Conference*. Springer, 2016, pp. 635–658.
- [36] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [37] D. Sommer, S. Meiser, and E. Mohammadi, "Privacy loss classes: The central limit theorem in differential privacy," *Cryptology ePrint Archive*, 2018.
- [38] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, "Generalization in adaptive data analysis and holdout reuse," *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [39] R. Rogers, A. Roth, A. Smith, and O. Thakkar, "Max-information, differential privacy, and post-selection hypothesis testing," in *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2016, pp. 487–494.
- [40] A. Rényi, "On measures of entropy and information," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and*

- Probability, Volume 1: Contributions to the Theory of Statistics*. University of California Press, 1961, pp. 547–561.
- [41] T. Van Erven and P. Harremoës, “Rényi divergence and Kullback-Leibler divergence,” *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 3797–3820, 2014.
- [42] G. R. Kurri, O. Kosut, and L. Sankar, “A variational formula for infinity-Rényi divergence with applications to information leakage,” *arXiv preprint arXiv:2202.06040*, 2022.
- [43] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.
- [44] Y. Wang, Y. O. Basciftci, and P. Ishwar, “Privacy-utility tradeoffs under constrained data release mechanisms,” *arXiv preprint arXiv:1710.09295*, 2017.
- [45] A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber, “Privacy: Theory meets practice on the map,” in *2008 IEEE 24th international conference on data engineering*. IEEE, 2008, pp. 277–286.
- [46] S. Meiser, “Approximate and probabilistic differential privacy definitions,” *Cryptology ePrint Archive*, 2018.
- [47] D. Kifer and B.-R. Lin, “An axiomatic view of statistical privacy and utility,” *Journal of Privacy and Confidentiality*, vol. 4, no. 1, 2012.
- [48] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, “Our data, ourselves: Privacy via distributed noise generation,” in *Annual international conference on the theory and applications of cryptographic techniques*. Springer, 2006, pp. 486–503.
- [49] A. McIver, C. Morgan, G. Smith, B. Espinoza, and L. Meinicke, “Abstract channels and their robust information-leakage ordering,” in *International Conference on Principles of Security and Trust*. Springer, 2014, pp. 83–102.
- [50] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [51] B. Wu, A. B. Wagner, G. E. Suh, and I. Issa, “Strong asymptotic composition theorems for sibson mutual information,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 2222–2227.
- [52] R. Sibson, “Information radius,” *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 14, no. 2, pp. 149–160, 1969.
- [53] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, “Generalization in adaptive data analysis and holdout reuse,” *arXiv preprint arXiv:1506.02629*, 2015.
- [54] I. Csiszár, “Information-type measures of difference of probability distributions and indirect observation,” *studia scientiarum Mathematicarum Hungarica*, vol. 2, pp. 229–318, 1967.
- [55] F. P. Calmon, A. Makhdoumi, M. Médard, M. Varia, M. Christiansen, and K. R. Duffy, “Principal inertia components and applications,” *IEEE Transactions on Information Theory*, vol. 63, no. 8, pp. 5011–5038, 2017.
- [56] H. Wang and F. P. Calmon, “An estimation-theoretic view of privacy,” in *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2017, pp. 886–893.