Journals
The Institution of
Engineering and Technology

Non-Stochastic Hypothesis Testing for Privacy

ISSN 1751-8644
doi: 0000000000
www.ietdl.org

Farhad Farokhi[1,2]

[1] The University of Melbourne, Parkville, VIC 3010, Australia
** E-mail: farhad.farokhi@unimelb.edu.au*

Abstract: In this paper, we consider privacy against hypothesis testing adversaries within a non-stochastic framework. We develop a theory of non-stochastic hypothesis testing by borrowing the notion of uncertain variables from non-stochastic information theory. We define tests as binary-valued mappings on uncertain variables and prove a fundamental bound on the best performance of tests in non-stochastic hypothesis testing. We provide parallels between stochastic and non-stochastic hypothesis-testing frameworks. We use the performance bound in non-stochastic hypothesis testing to develop a measure of privacy. We then construct reporting policies with prescribed privacy and utility guarantees. The utility of a reporting policy is measured by the distance between the reported and original values. Finally, we present the notion of indistinguishability as a measure of privacy by extending identifiability from the privacy literature to the non-stochastic framework. We prove that linear quantizers can indeed achieve identifiability for responding to linear queries on private datasets.

1 Introduction

Advances in computation and communication capabilities have paved the way for using big data to solve important societal challenges. However, new tools developed for collection and analysis of data has caused the erosion of privacy, motivating investigation of methodologies for privacy analysis and preservation. For decades, stochastic or randomized policies have been used for privacy protection [1]. Provable privacy guarantees have been presented for stochastic policies within the frameworks of differential privacy [2–5], identifiability [6–8], and information-theoretic privacy [9–15].

Differential privacy uses randomization to ensure that the statistics of the reported outputs, i.e., query responses to datasets, do not change noticeably (extent of which is captured by the so-called privacy budget) by variations in an individual entry of the dataset. For real-valued datasets, differential privacy can be ensured by additive Laplace or Gaussian noise whose scales must be selected proportional to the sensitivity of query responses with respect to the individual entries of the private dataset. There are however difficulties associated with these mechanisms. In fact, differentially-private additive noises create undesirable computational properties [16] and lead to generation of unreasonable/unrealistic outputs [17–20].

Information-theoretic privacy, dating back to the secrecy problem [21], emphasizes on masking or equivocating of information from the intended primary receiver or a secondary receiver with as much information as the primary receiver (e.g., an eavesdropper) [15, 22–24] while providing guarantees on utility by bounding distortion, i.e., the distance between obfuscated and original reports. A common factor among all the information-theoretic privacy results (see [9–15] and the references therein) is that the private dataset must be randomly distributed (often with independence assumption among the entries of the dataset for the sake of analysis). This might not be the case in practice. Even if randomly distributed, the knowledge of the probability distributions of datasets might not be available at the time of design or might change over time. Furthermore, privacy can only be achieved with the aid of randomizations at the output. This is because the metrics used for capturing information leakage in those studies are based on probability theory and its application in traditional information theory. The guarantees of information-theoretic privacy policies are also presented in the form of averages, i.e., they bound the average amount of leaked information or the information leaked about a population rather than a specific individual.

Although the above-mentioned stochastic policies provide provable privacy guarantees, many organizations still use deterministic heuristic-based privacy-preserving methods, such as k-anonymity [25, 26] and ℓ-diversity [27]. For instance, anonymization by removing identifiable features, such as name or address, is frequently used by governments[*] or companies alike for releasing private data[†] to the broader public for analysis even though it is proved to be insufficient for privacy preservation [28–30]. Other policies, such as k-anonymity, are also shown to be vulnerable to attacks [27].

Stochastic privacy-preserving policies also cause problems, e.g., un-truthfulness [17], that are not desirable in practice [31]. Stochastic policies have been criticized in financial auditing [19, 20] and medical, health, or social sciences [32, 33]. In these situations, non-stochastic privacy-preserving policies can be better suited.

Motivated by these observations, a non-stochastic measure of privacy was introduced in [34] to prove that k-anonymity is in fact *not* privacy preserving. Although this was shown using numerical analysis in [27], the non-stochastic measure of privacy in [34] did not require extensive simulations and numerical studies. This is good because of two reasons. First, experimental analysis only provides sufficient results, i.e., lack of existence of an experiment that can unravel a privacy-preserving methodology does not imply that it is provably private (as there might exist an experiment for which the-said methodology leaks private information of individuals while it is not found or formalized yet). Furthermore, a measure of privacy allows for investigation of general methodologies and comparing their guarantees, e.g., to rank their promises or to convert their parameter settings to each other. The non-stochastic measure of privacy in [34] was introduced based on the theory of non-stochastic information theory [35–42] and was also successfully used to show that binning, a popular deterministic policy for privacy preservation, provides tunable privacy guarantees. The privacy measure in [34] is perfect for providing protection against generic adversaries; however, in some instances, more might be known about the privacy-intrusive adversaries, hence the privacy measure can be further refined.

[*] *See* `https://data.gov.au` *for anonymized government data.*
[†] *See* `https://www.kaggle.com` *for data of companies and individuals.*

A category of adversaries studied in privacy literature is the hypothesis-testing adversaries [4, 13, 43]. In this case, the adversary is interested in examining the validity of a hypothesis, e.g., if a house is occupied or if an individual has a certain disease. For this setup, the privacy risk can be measured by the error probability of the adversary when performing the hypothesis testing, e.g., [13]. In this paper, we expand privacy against hypothesis-testing adversaries to a non-stochastic framework.

We particularly develop a theory of non-stochastic hypothesis testing by borrowing the concept of uncertain variables from non-stochastic information theory [34, 40]. In order to establish non-stochastic hypothesis testing, we introduce uncertain variables as non-stochastic counterparts of random variables. Uncertain variables only consider support sets, referred to as ranges, and do not assign distributions/measures to variables. In non-stochastic hypothesis theory, we define tests as binary-valued functions on uncertain variables. We measure the performance of a test by combination of the size of true positive and negative sets. We prove a fundamental bound for the best performance of tests. This bound is used to develop a measure of privacy. We then construct reporting functions for given privacy and utility guarantees. The privacy-preserving reporting function is based on blurring the transition from validity of null to alternative hypothesis for a hypothesis-testing privacy-intrusive adversary, making it harder for the adversary to infer which hypothesis is valid. The wider the region in which the transition is blurred, the lower the accuracy and the higher the privacy become. This intuitively establishes a trade-off between privacy and accuracy. We provide parallels between stochastic and non-stochastic hypothesis-testing frameworks under some mild conditions. This provide an avenue for extending the guarantees of non-stochastic privacy to random variables. The non-stochastic hypothesis-testing framework allows us to extend the notion of identifiability from the privacy literature (see, e.g., [6–8]) to a non-stochastic setup. We refer to this notion of privacy as indistinguishability because of the similarity of its definition to the concept of semantic security, also know as indistinguishability under chosen plaintext attack in the encryption literature; see, e.g., [44]. We prove that concatenation of linear functions with quantizers (also known as binning) can guarantee indistinguishability if the quantizer has "few enough" levels. The upper bound on the number of the levels is a function of the sensitivity of the mapping, the range of the data, and indistinguishability budget (which is inversely proportional to the privacy level).

Finally, an early version of these results has been presented [45]. In this paper, the results are expanded by investigating the parallels between stochastic and non-stochastic hypothesis-testing frameworks, introducing indistinguishability as a notion of privacy, and investigating the non-stochastic information leakage associated with the new notions of privacy.

The remainder is organized as follows. We present background material on non-stochastic information theory in Section 2. We present non-stochastic hypothesis testing in Section 3. In Section 4, we provide parallels between stochastic and non-stochastic hypothesis testing. In Section 5, we investigate privacy against hypothesis-testing adversaries. Finally, we present indistinguishability as a notion of privacy in Section 6 and conclude the paper in Section 7.

2 Uncertain Variables and Non-Stochastic Information Theory

In this section, we review a few necessary concepts from non-stochastic information theory. We start with the notion of uncertain variables, non-stochastic counterparts of random variables.

Consider uncertainty set Ω whose elements $\omega \in \Omega$ are samples. The elements or samples $\omega \in \Omega$ are the source of uncertainty. An uncertain variable X is a mapping on Ω. For uncertain variable $X : \Omega \to \mathbb{X}$, $X(\omega)$ denotes a realization of the uncertain variable. Sometimes, in short, u.v. refers to uncertain variable. When the dependence of a realization of an uncertain variable to the sample ω is evident from the context, $X(\omega)$ is replaced by X. In this paper, we restrict ourselves to real-valued uncertain variables, e.g., $\mathbb{X} \subseteq \mathbb{R}^{n_x}$ for some integer $n_x \geq 0$.

Range of any uncertain variable X is $[\![X]\!] := \{X(\omega) : \omega \in \Omega\} \subseteq \mathbb{X}$. Range of an uncertain variable is essentially the same as the support set of the probability density of random variables. As uncertain variables do not possess any probability distributions, the range is enough to describe their behavior. The range determines the amount of uncertainty that can be caused by an uncertain variable. Hence, as shown subsequently, the size of the range is intimately related to *entropy* of uncertain variables, a measure of their internal disorder or unpredictability. *Joint range* of any two uncertain variables $X : \Omega \to \mathbb{X}$ and $Y : \Omega \to \mathbb{Y}$ is given by $[\![X, Y]\!] := \{(X(\omega), Y(\omega)) : \omega \in \Omega\} \subseteq \mathbb{X} \times \mathbb{Y}$. Joint range is similar to the support set of the joint probability density, capturing how two uncertain variable vary together. *Conditional range* of an uncertain variable X, conditioned on the realization of uncertain variable $Y(\omega) \in \mathcal{Y} \subseteq [\![Y]\!]$, is $[\![X|\mathcal{Y}]\!] := \{X(\omega) : \exists \omega \in \Omega \text{ such that } Y(\omega) \in \mathcal{Y}\} \subseteq [\![X]\!]$. When the set $\mathcal{Y}$ is a singleton $\{y\}$, $[\![X|y]\!]$ can be used instead of $[\![X|\mathcal{Y}]\!]$ to denote the conditional range. Conditional range is the same as the support set of conditional probability density, capturing how observing one uncertain variable informs the realizations of another uncertain variable.

Example 1 (Modeling a Completely Unfair Dice)**.** *Imagine playing a dice game against a player who can cheat arbitrarily. When the cheating player "rolls" the dice, there is no guarantee that the numbers on the dice appear with fair probability of* $1/6$. *The number can be anything that the cheating player desires. If we do not know the player's cheating strategy, this situation can be best modeled using an uncertain variable* $X : \Omega \to \mathbb{N}$. *Here, elements of uncertainty set* Ω *model the thinking or the desire of the cheating player. The range of* X *is* $[\![X]\!] = \{1, 2, \dots, 6\}$. *Now, imagine another uncertain variable* $Y : \Omega \to \{\mathit{yes}, \mathit{no}\}$ *determining if the number on the dice is even. The joint range of* X *and* Y *is* $[\![X, Y]\!] = \{(1, \mathit{no}), (2, \mathit{yes}), \dots, (6, \mathit{yes})\}$. *Also, the conditional range of* X *given* Y *is* $[\![X|\mathit{yes}]\!] = \{2, 4, 6\}$ *and* $[\![X|\mathit{no}]\!] = \{1, 3, 5\}$. *This is a trivial illustrative example; however, uncertain variables provide a powerful mechanism for modeling behaviors that are not random or, if random, their distributions are not known or are time varying.*

Figure 1 shows an illustrative example of uncertainty set Ω and two uncertain variables X and Y. The dotted arrow, connecting $\omega \in \Omega$ and $X(\omega) \in [\![X]\!]$, shows the relationship between samples in the uncertainty set and realizations of the uncertain variable. The dashed lines show the range of the mapping X for the entire uncertainty set Ω, illustrating the range $[\![X]\!]$. Note that $[\![X]\!] \subseteq \mathbb{X}$; however, equality $[\![X]\!] = \mathbb{X}$ might not always occur. Finally, the dash-dotted arrows illustrate the conditional range $[\![X|y]\!]$ as the image of the uncertain variable X for the inverse image set $Y^{-1}(y) := \{\omega \in \Omega : Y(\omega) = y\} \subseteq \Omega$.

Finally, if the range $[\![X]\!]$ is uncountably infinite for an uncertain variable X, we call it a *continuous* uncertain variable, similar to a continuous random variable. If the range $[\![X]\!]$ is countable for an uncertain variable X, we call it a *discrete* uncertain variable.

Non-stochastic entropy of a continuous uncertain variable X can be defined as

$$h_0(X) := \log_e(\mu([\![X]\!])) \in \mathbb{R} \cup \{\pm\infty\}, \tag{1}$$

where $\mu(\cdot)$ is the Lebesgue measure. The logarithm can be taken in any basis; however, in this paper for continuous uncertain variables, the logarithm is in the natural basis in line with the literature on differential entropy of continuous random variables. The non-stochastic entropy in (1) is sometimes referred to as Rényi differential 0-entropy [38]. *Non-stochastic entropy of a discrete uncertain variable* X can be defined as

$$H_0(X) := \log_2(|[\![X]\!]|) \in \mathbb{R}, \tag{2}$$

where $|\cdot|$ is the cardinality of a set. In this paper, for discrete uncertain variables, in line with the literature on entropy of discrete random variables, the logarithm is in the basis of two. *Entropy* of an uncertain variable measures the size of its range and, noting that the

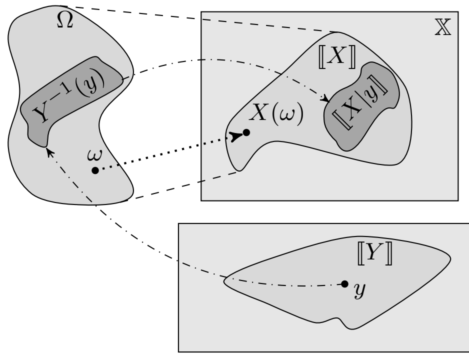

Fig. 1: An illustrative example of uncertainty set Ω and two uncertain variables X and Y. The dotted arrow shows the relationship between samples of the uncertain space and realizations of the uncertain variable. The dashed lines show limits of range of uncertain variable X. The dash-dotted arrows illustrate the relationship between realization $y = Y(\omega)$ and conditional range $[\![X|y]\!]$.

range of an uncertain variable models all possibilities that can occur, the entropy is a measure of internal disorder or unpredictability of the uncertain variable.

We can use a similar approach to [38, 46] to define *non-stochastic relative or conditional entropy* of uncertain variable X conditioned on uncertain variable Y as

$$h_0(X|Y) := \operatorname*{ess\,sup}_{y\in[\![Y]\!]} \log_e(\mu([\![X|y]\!])), \tag{3}$$

where essential supremum is $\operatorname{ess\,sup}_{x\in\mathcal{X}} f(x) := \inf\{b \in \mathbb{R} : \mu(\{x \in \mathcal{X} : f(x) > b\}) = 0\}$ for any $f : \mathcal{X} \to \mathbb{R}$. If Y is a discrete uncertain variable, we replace essential supremum with supremum, $\mu(\cdot)$ with $|\cdot|$, $\log_e$ with $\log_2$, and h_0 with H_0. Based on the definition of entropy and relative entropy, *non-stochastic information* between two uncertain variables X and Y is defined as

$$\begin{aligned} I_0(X;Y) :=& h_0(X) - h_0(X|Y) \\ =& \operatorname*{ess\,inf}_{y\in[\![Y]\!]} \log_e\left(\frac{\mu([\![X]\!])}{\mu([\![X|y]\!])}\right), \end{aligned} \tag{4}$$

where $\operatorname{ess\,inf}_{x\in\mathcal{X}} f(x) := -\operatorname{ess\,sup}_{x\in\mathcal{X}} -f(x)$, for $f : \mathcal{X} \to \mathbb{R}$, defines essential infimum. Again, if Y is a discrete uncertain variable, we replace essential infimum with infimum, $\mu(\cdot)$ with $|\cdot|$, and $\log_e$ with $\log_2$.

Kolmogorov's definition of *combinatorial conditional entropy* is $\log_e(\mu([\![X|y]\!]))$ based on which the *information gain* is $\log_e(\mu([\![X]\!])/\mu([\![X|y]\!]))$ [36]. Note that these quantities depend on the observed realization of uncertain variable $Y = y$ while the non-stochastic information in (4) relies on the worst-case ratio.

In [34], it was shown that the non-stochastic information in (4) is not a good measure of privacy and an alternative notion of non-stochastic information leakage was introduced. *Non-stochastic information leakage* is

$$L_0(X;Y) := \operatorname*{ess\,sup}_{y\in[\![Y]\!]} \log_e\left(\frac{\mu([\![X]\!])}{\mu([\![X|y]\!])}\right). \tag{5}$$

Non-stochastic information leakage $L_0(X;Y)$ captures the reduction in the complexity of brute-force guessing X after observing Y [47]. Now, we are ready to present our framework and results regarding non-stochastic hypothesis testing.

3 Non-Stochastic Hypothesis Testing

In this section, we introduce non-stochastic hypothesis testing. We define tests as functions that map observed uncertain variables to hypothesis. We define performance of a test by measuring the size of outputs that result in true positives and negatives. We prove that a specific category of tests, known as consistent tests, maximize this performance. This provides a fundamental bound on performance of non-stochastic hypothesis tests.

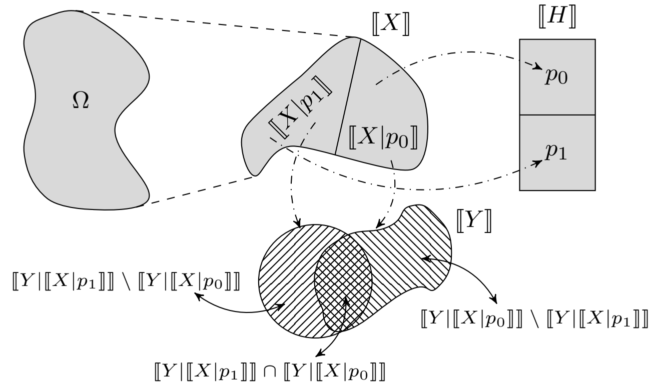
Ω
$[\![X]\!]$
$[\![H]\!]$
p_0
p_1
$[\![X|p_1]\!]$
$[\![X|p_0]\!]$
$[\![Y]\!]$
$[\![Y|[\![X|p_1]\!]]\!] \setminus [\![Y|[\![X|p_0]\!]]\!]$
$[\![Y|[\![X|p_0]\!]]\!] \setminus [\![Y|[\![X|p_1]\!]]\!]$
$[\![Y|[\![X|p_1]\!]]\!] \cap [\![Y|[\![X|p_0]\!]]\!]$

Fig. 2: Relationship between uncertain variables in non-stochastic hypothesis testing based on uncertain measurements. If the realization of uncertain measurement Y belongs to $[\![Y|p_0]\!] \cap [\![Y|p_1]\!]$, there is not enough evidence to accept or reject the null hypothesis p_0 or the alternative hypothesis p_1. However, if the realization of uncertain measurement Y belongs to $[\![Y|p_0]\!] \setminus [\![Y|p_1]\!]$ ($[\![Y|p_1]\!] \setminus [\![Y|p_0]\!]$), we can confidently accept (reject) the null hypothesis p_0 and reject (accept) the alternative hypothesis p_1.

Consider uncertain variable X denoting the original uncertain variable and we are interested in testing the validity of a hypothesis for its realizations. We clearly do not have access to realizations of this uncertain variable as otherwise hypothesis testing was trivial. We have access to an *uncertain measurement* of this variable denoted by Y. This is captured by that $Y = g_Y(X)$ for a mapping $g_Y : [\![X]\!] \to [\![Y]\!]$. Recalling that uncertain variables are mappings from the uncertainty set, it must be that $Y = g_Y \circ X$, where $\circ$ denotes composition of mappings. Similarly, we may define the *hypothesis* as an uncertain variable H with binary range $[\![H]\!] = \{p_0, p_1\}$, where p_0 denotes the *null hypothesis* and p_1 denotes the *alternative hypothesis*. We assume that there exists a mapping $g_H : [\![X]\!] \to [\![H]\!]$ such that $H = g_H \circ X$; the hypothesis is constructed based on the uncertain variable X as $H = g_H(X)$. This setup and the relationship between all uncertain variables is summarized in Figure 2. When the realization of uncertain measurement Y belongs to $[\![Y|p_0]\!] \setminus [\![Y|p_1]\!]$, we can confidently accept the null hypothesis p_0 and reject the alternative hypothesis p_1. Therefore, the size (or volume) of $[\![Y|p_0]\!] \setminus [\![Y|p_1]\!]$ corresponds to the highest "true negative" rate in the traditional hypothesis testing framework. Alternatively, when the realization of uncertain measurement Y belongs to $[\![Y|p_1]\!] \setminus [\![Y|p_0]\!]$, we can confidently reject the null hypothesis p_0 and accept the alternative hypothesis p_1. Hence, the size of $[\![Y|p_1]\!] \setminus [\![Y|p_0]\!]$ corresponds to the highest "true positive" rate.

A *test* is a function $T : [\![Y]\!] \to [\![H]\!] = \{p_0, p_1\}$. If $T(Y) = p_1$, the test rejects the null hypothesis in favour of the alternative hypothesis; however, if $T(Y) = p_0$, the test accepts the null hypothesis (and rejects the alternative hypothesis). The set of all tests is given by $[\![H]\!]^{[\![Y]\!]}$, which captures the set of all functions from $[\![Y]\!]$ to $[\![H]\!]$.

Let us, as a thought experiment, consider $y \in [\![Y]\!]$ for which $T(y) = p_0$ (if exists, otherwise the same thought experiment can be conducted $y \in [\![Y]\!]$ such that $T(y) = p_1$); hence, the null hypothesis p_0 is accepted by the test at this instance. The realization $Y(\omega) = y$ may correspond to many realizations of uncertain variable X (noting the uncertainty surrounding the measurement), i.e., all the elements of the set $[\![X|y]\!]$. Evidently, $T(y) = p_0$ is correct, or the test is correct for the output realization $Y(\omega) = y$, if $g_H(x) = p_0$ for all $x \in [\![X|y]\!]$, i.e., all realizations of uncertain variable X compatible with y are also compatible with the null hypothesis. The same also holds for the alternative hypothesis. In the following definition, recalling from the previous section, we use the notation $\llbracket H|\llbracket X|y\rrbracket\rrbracket = \{H(\omega) : \exists\omega\in\Omega$ such that $X(\omega)\in\llbracket X|y\rrbracket\} = \{h\in\llbracket H|x\rrbracket : x\in\llbracket X|y\rrbracket\}$.

Definition 1 (Correctness). *A test $T\in\llbracket H\rrbracket^{\llbracket Y\rrbracket}$ is correct at $y\in\llbracket Y\rrbracket$ if $\llbracket H|\llbracket X|y\rrbracket\rrbracket = \{T(y)\}$. The set of all outputs at which test T is correct is given by $\aleph(T) := \{y\in\llbracket Y\rrbracket : \llbracket H|\llbracket X|y\rrbracket\rrbracket = \{T(y)\}\}$.*

Based on the definition of correctness, we can define a performance measure for tests. If Y is a continuous uncertain variable, the performance is given by

$$\mathcal{P}(T) := \log_e(\mu(\aleph(T))). \tag{6}$$

Alternatively, if Y is a discrete uncertain variable, the performance can be captured according to

$$\mathcal{P}(T) := \log_2(|\aleph(T)|). \tag{7}$$

The performance of a test $\mathcal{P}(T)$ captures both true positives and negatives. We can break down positive and negative cases separately as in the traditional hypothesis testing; however, in what follows, we are interested in both true positives and negatives when measuring the success of hypothesis-testing privacy-intrusive adversaries: Any information gained in a privacy attack is important for privacy analysis. The problem of finding an optimal hypothesis test can be cast as an optimization problem:

$$T^*\in\operatorname*{arg\,max}_{T\in\llbracket H\rrbracket^{\llbracket Y\rrbracket}} \mathcal{P}(T). \tag{8}$$

Instead of dealing with the size of true positive and negatives sets, we could work with the combined size of type I and II errors $\log_e(\mu(\{y\in\llbracket Y\rrbracket : \llbracket H|\llbracket X|y\rrbracket\rrbracket \neq \{T(y)\}\}))$ or $\log_2(|\{y\in\llbracket Y\rrbracket : \llbracket H|\llbracket X|y\rrbracket\rrbracket \neq \{T(y)\}\}|)$. This combined measure of error corresponds to the probability of error in stochastic hypothesis testing [48]. Following this approach, the problem (8) should be reformulated as a minimization instead of a maximization (as we would want to minimize the measure of error). A family of tests play an important role in the optimization problem (8). The underlying property of these tests is captured in the following definition.

Definition 2 (Consistency). *A test $T:\llbracket Y\rrbracket\to\llbracket H\rrbracket$ is called consistent if (i) $T(y)=p_0$ only if $y\in\llbracket Y|p_0\rrbracket$ and (ii) $T(y)=p_1$ only if $y\in\llbracket Y|p_1\rrbracket$.*

In the following theorem, we prove that consistent tests are in fact optimal in the sense of $\mathcal{P}$ in (6) and (7).

Theorem 1 (Optimal Tests). *Any consistent test is a solution of* (8).

Proof: See Appendix 9. □

Note that, for any realization of uncertain measurement Y belonging to $\llbracket Y|p_0\rrbracket\cap\llbracket Y|p_1\rrbracket$, there is not enough evidence to accept or reject either the null hypothesis or the alternative hypothesis. This is because such realizations of uncertain measurement Y can be caused by the realizations of X that are consistent with the null hypothesis p_0 and the realizations of X that are consistent with the alternative hypothesis p_1. However, if the realization of the measurement Y belongs to $(\llbracket Y|p_0\rrbracket\setminus\llbracket Y|p_1\rrbracket)\cup(\llbracket Y|p_1\rrbracket\setminus\llbracket Y|p_0\rrbracket) = \llbracket Y|p_0\rrbracket\Delta\llbracket Y|p_1\rrbracket$, with Δ denoting the symmetric difference operator on the sets, we can confidently reject or accept the null hypothesis or the alternative hypothesis. This fact is used by the consistent tests to achieve the highest performance.

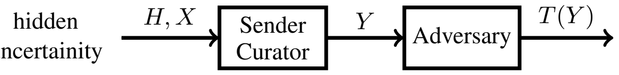

Fig. 3: Communication structure between a sender and a hypothesis-testing adversary.

Theorem 2. (Fundumental Performance Bound) *The performance of any test $T\in\llbracket H\rrbracket^{\llbracket Y\rrbracket}$ is upper bounded by*

$$\mathcal{P}(T)\le\log_e(\mu(\llbracket Y|p_0\rrbracket\Delta\llbracket Y|p_1\rrbracket)),$$

if Y is a continuous uncertain variable or

$$\mathcal{P}(T)\le\log_2(|\llbracket Y|p_0\rrbracket\Delta\llbracket Y|p_1\rrbracket|),$$

if Y is a discrete uncertain variable.

Proof: See Appendix 10. □

Theorem 2 presents a non-stochastic alternative to classical information-theoretic results on hypothesis testing [49]. In Theorem 2, the upper bounds $\log_e(\mu(\llbracket Y|p_0\rrbracket\Delta\llbracket Y|p_1\rrbracket))$ and $\log_2(|\llbracket Y|p_0\rrbracket\Delta\llbracket Y|p_1\rrbracket|)$ essentially capture the difference between the conditional ranges $\llbracket Y|p_0\rrbracket$ and $\llbracket Y|p_1\rrbracket$, resembling the *total variation distance* in a non-stochastic framework (incorporating support sets of uncertain variables instead of density functions of random variables).

Example 2 (Hypothesis testing using uncertain measurements). *Consider an uncertain variable $X=(X_1,X_2)\in[100,250]\times[-10,10]$, where X_1 denotes the height of an individual in centimetres and X_2 denotes a measurement error in centimetres. Let the uncertain measurement be $Y=g_Y(X)=X_1+X_2$. The uncertain variable capturing the hypothesis $H:\Omega\to\llbracket H\rrbracket=\{p_0,p_1\}$ is*

$$H(\omega)=g_H(X(\omega))=\begin{cases}p_0, & X_1(\omega)\le 150,\\ p_1, & X_1(\omega)>150.\end{cases}$$

The null hypothesis p_0 states that the individual whose height is captured by a realization of the uncertain variable X is short (in fact shorter than or equal to 150 centimetres) and the alternative hypothesis p_1 states that the individual is tall (in fact taller than 150 centimetres). Note that $\llbracket Y|p_0\rrbracket=\{X_1+X_2 : 100\le X_1\le 150, X_2\in[-10,10]\}=[90,160]$ and $\llbracket Y|p_1\rrbracket=\{X_1+X_2 : 150\le X_1\le 250, X_2\in[-10,10]\}=[140,260]$. Therefore, $\llbracket Y|p_0\rrbracket\cap\llbracket Y|p_1\rrbracket=[140,160]$. Let T be a test such that $T(Y)=p_0$ if $Y\in[90,150]$ and $T(Y)=p_1$ if $Y\in(150,260]$. Evidently, T is a consistent test. We get $\mathcal{P}(T)=\log_e(\mu(\llbracket Y|p_0\rrbracket\Delta\llbracket Y|p_1\rrbracket))=\log_e(\mu([90,140)\cup(160,260]))=\log_e(150)$. If we scale the performance by $h_0(Y)=\log_e(170)$, we get

$$\mathcal{P}(T)-h_0(Y)=\log_e(150)-\log_e(170)\approx -0.1251.$$

Let us imagine another example in which $X=(X_1,X_2)\in[100,250]\times[-20,20]$ with the interpretation that the magnitude of the additive measurement uncertainty is twice larger. In this case, we get

$$\mathcal{P}(T)-h_0(Y)=\log_e(150)-\log_e(190)\approx -0.2364.$$

Hence, by increasing the magnitude of the measurement uncertainty, the confidence of the test (its performance relative to the total uncertainty) is reduced, which aligns with our expectation. △

4 Stochastic versus Non-Stochastic Hypothesis Testing

In this section, we investigate the similarities between stochastic and non-stochastic hypothesis testing by endowing all the uncertain variables with a measure, therefore making them random variables. Doing so, we can extend the guarantees from non-stochastic privacy, subsequently defined based our results from non-stochastic hypothesis testing, to random variables.

Particularly, we assume that X, Y, H are jointly distributed random variables. We also assume that the test is no longer a determinstic function but a conditional distribution that maps Y to $\widehat{H} \in \{p_0, p_1\}$ with $\widehat{H} = p_0$ signifying the fact that the null hypothesis is accepted and $\widehat{H} = p_1$ capturing the case where the null hypothesis is rejected in favour of the alternative hypothesis. We assume that all probability distributions are absolutely continuous with respect to the Lebesgue measure, i.e., probability density functions exist (following the Radon-Nikodym theorem [50, pp. 419–427]). In what follows, we use the notation $\mathrm{supp}(\cdot)$ to denote the support set of the probability density functions. We use the notation $p_{Y|p_0}$ and $p_{Y|p_1}$, respectively, to denote the conditional probability density function for Y given p_0 and p_1. Also, in this section, $[\![Y|p_0]\!] = \mathrm{supp}(p_{Y|p_0})$ and $[\![Y|p_1]\!] = \mathrm{supp}(p_{Y|p_1})$.

Theorem 3. *Assume that there exists $\varepsilon > 0$ such that $p_{Y|p_0}(y) \geq \varepsilon$ and $p_{Y|p_1}(y) \geq \varepsilon$ for all $y \in [\![Y|p_0]\!]$ and $y \in [\![Y|p_1]\!]$, respectively. Then,*

$$\sup_{p_{\widehat{H}|Y}} \mathbb{P}\{\widehat{H} = H\} \geq \frac{\varepsilon}{2}\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!]),$$

where $p_{\widehat{H}|Y}$ is the conditional probability density of $\widehat{H}$ given Y.

Proof: See Appendix 11. □

Theorem 3 shows that the lower bound on the best achievable true positive rate in the stochastic hypothesis testing scales with $\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!])$. From Theorem 2, we know that $\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!])$ captures the fundamental bound on the performance of non-stochastic hypothesis tests. This provides a parallel between the stochastic and non-stochastic hypothesis testing frameworks. Therefore, when using the non-stochastic hypothesis testing for privacy preservation in the next section, we get parallels about privacy against stochastic hypothesis-testing adversaries.

5 Privacy Against Hypothesis-Testing Adversary

In this section, we define a non-stochastic notion of privacy against hypothesis-testing privacy-intrusive adversaries. Noting that sometimes we must perturb reporting functions to satisfy this new notion of privacy, we also define a notion of accuracy (or utility) for the reports to balance privacy and accuracy. We introduce a family of privacy-preserving reporting functions based on blurring the transition from validity of one hypothesis to other, making it harder for the adversary to infer which hypothesis is valid. The wider the region in which the transition is blurred, the lower the accuracy and the higher the privacy become. Following this, we establish a trade-off between privacy and accuracy.

Consider the communication diagram in Figure 3 between a sender or a curator, and an adversary. The adversary's ultimate aim is to accurately test a hypothesis H based on the communicated information from the sender Y. The sender wants to provide a message Y that is as close as possible to X while making the adversary's task in testing the validity of hypothesis H hard; therefore, there is a need for strategic manipulation of realizations of X during the communication. The policy of the sender is captured by the mapping from X to Y, denoted by $g_Y : [\![X]\!] \to [\![Y]\!]$.

We use the performance of the adversary in testing the privacy-intrusive hypothesis based on the communicated uncertain variable Y to define a measure of privacy. If Y is a continuous uncertain variable, the measure of privacy is

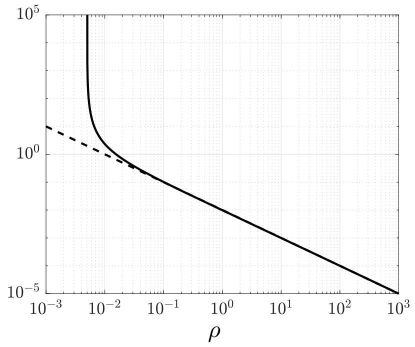

Fig. 4: Privacy guarantee, ϵ, versus accuracy level, ρ. The dashed line shows the asymptotic $\mathcal{O}(\rho^{-1})$.

$$\mathrm{Priv}(g_Y) := h_0(Y) - \log_e(\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!])), \quad (9)$$

while, if Y is a discrete uncertain variable, the measure of privacy is

$$\mathrm{Priv}(g_Y) := H_0(Y) - \log_2(|[\![Y|p_0]\!]\Delta[\![Y|p_1]\!]|). \quad (10)$$

Increasing the privacy measure $\mathrm{Priv}(g_Y)$ implies that the size of the set $[\![Y|p_0]\!]\Delta[\![Y|p_1]\!]$ measured by its volume or cardinality is decreased, thus degrading the performance of any test employed by the adversary in light of Theorem 2.

Definition 3 (ϵ-privacy). *Policy g_Y is ϵ-private for $\epsilon \in (0, \infty)$ if $\mathrm{Priv}(g_Y) \geq \log(1 + \epsilon)$.*

Evidently, we must balance the need for privacy against accuracy (or utility) of the report; otherwise the best policy in terms of protecting privacy is to report nothing. Therefore, we need to define a measure of accuracy.

Definition 4 (ρ-accuracy). *Policy g_Y is ρ-accurate for $\rho \in (0, +\infty)$ if $\sup_{x\in[\![X]\!]} \|x - g_Y(x)\| \leq 1/\rho$.*

Increasing ρ in ρ-accuracy implies that the worst-case distance between the realization of uncertain variable X and the message Y is decreased; thus improving the quality of the reported output Y by requiring it to stay consistently closer to X. In what follows, for any i, x_i denotes the i-th component of x and x_{-i} denotes all entries of x except the i-th component.

Theorem 4. *Assume that $[\![X]\!] \subseteq \mathbb{R}^{n_x}$, index $i \in \{1, \ldots, n_x\}$, and $g : \mathbb{R}^{n_x-1} \to \mathbb{R}$ exist such that*

$$g_H(x) = \begin{cases} p_0, & x_i - g(x_{-i}) \geq 0, \\ p_1, & x_i - g(x_{-i}) < 0, \end{cases}$$

where $x_{-i} = (x_j)_{j\neq i}$. Let

$$g_Y(x) = \begin{cases} (g(x_{-i}), x_{-i}), & g(x_{-i}) - \dfrac{1}{\rho} \leq x_i \leq g(x_{-i}) + \dfrac{1}{\rho}, \\ x, & \textit{otherwise}, \end{cases}$$

and

$$\epsilon = \left[\frac{\exp(h_0(X))}{\mu\left([\![X]\!] \cap \left\{x : g(x_{-i}) - \dfrac{1}{\rho} \leq x_i \leq g(x_{-i}) + \dfrac{1}{\rho}\right\}\right)} - 1\right]^{-1}.$$

Then, g_Y is ρ-accurate and ϵ-private.

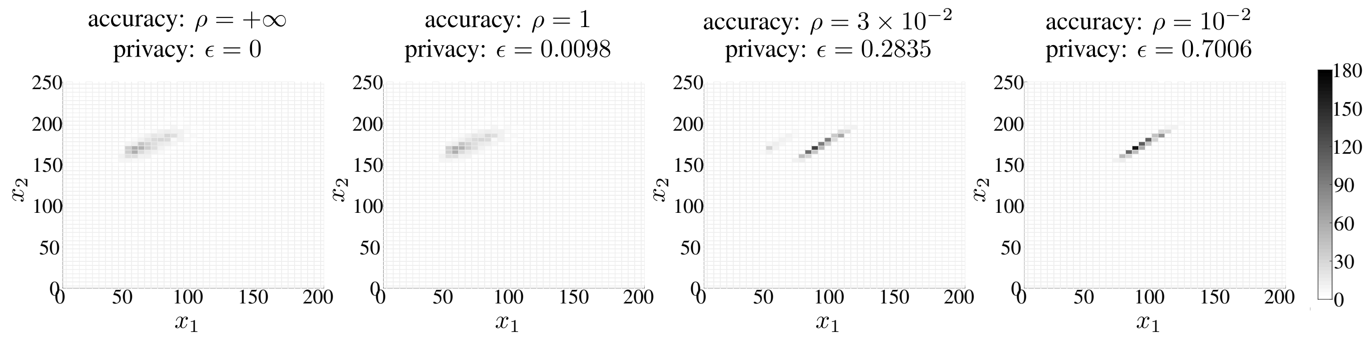

Fig. 5: The histogram of the reported weight and height of individuals for various levels of accuracy ρ. The darker colors show higher frequencies.

Proof: See Appendix 12. □

Theorem 4 provides a method for constructing ρ-accurate policies and computing the privacy budget ϵ of these reports. The reporting mechanism is based on blurring the transition area from one hypothesis to the other, making it harder to infer which hypothesis is valid. The width of the region in which the data is manipulated is given by $1/\rho$. The wider the region, the lower the accuracy becomes.

Remark 1 (Privacy×Accuracy=Constant). *For large enough ρ, it can be seen that $[\![X]\!] \cap \{x : g(x_{-i}) - \rho^{-1} \leq x_i \leq g(x_{-i}) + \rho^{-1}\} \approx \{x : g(x_{-i}) - \rho^{-1} \leq x_i \leq g(x_{-i}) + \rho^{-1}\}$, and, as a result, $\epsilon = \mathcal{O}(\rho^{-1})$. This implies that, for the policy in Theorem* 4, *we have $\epsilon\rho = \mathcal{O}(1)$.*

With the theoretical results in hand, we can demonstrate the effects of using the policy in Theorem 4 on a practical dataset.

Example 3 (Body Mass Index Privacy). *Let us consider the design of a privacy-preserving policy for reporting individuals' height (in centimetres) and weight (in kilograms) publicly. We consider an adversary who is interested in identifying individuals passing the obesity threshold in terms of body mass index (BMI), e.g., an insurance agency that may use publicly available data to increase premiums of obese people or deny them insurance. Therefore, there is a duty of care for releasing demographic data of individuals publicly. By the definition of the U.S. Department of Health & Human Services, a person, be it female or male, is considered obese if their BMI is greater than or equal to 30.*

Let uncertain variable X contain the weight and height of individuals, i.e., $X = [X_1\ X_2]^\top$ with $X_1 \in [0, 200]$ denoting the weight in kilograms and $X_2 \in [0, 250]$ denoting the height in centimetres. The hypothesis is given by

$$g_H(x) = \begin{cases} p_0, & \dfrac{x_1}{(x_2/100)^2} \geq 30, \\ p_1, & \dfrac{x_1}{(x_2/100)^2} < 30. \end{cases}$$

Following the notation of Theorem 4, *we can redefine the hypothesis using the sign of $x_1 - g(x_2)$ with $g : x_2 \mapsto 30(x_2/100)^2$. Define*

$$g_Y(x) = \begin{cases} \left(\dfrac{30x_2^2}{10^4}, x_2\right), & \dfrac{3x_2^2}{1000} - \dfrac{1}{\rho} \leq x_1 \leq \dfrac{3x_2^2}{1000} + \dfrac{1}{\rho}, \\ (x_1, x_2), & \textit{otherwise}. \end{cases} \tag{11}$$

Following Theorem 4, *$g_Y(\cdot)$ is ρ-accurate. Using Theorem* 4, *we can compute the level of privacy guarantee. The solid black line in Figure* 4 *illustrates privacy guarantee ϵ versus accuracy level ρ. The dashed line shows the asymptotic $\mathcal{O}(\rho^{-1})$. As expected from Remark* 1, *by increasing accuracy, the privacy guarantee can only be reduced and vice versa.*

Now, we use a real dataset to investigate the effects of the privacy-preserving policy in (11). *We use a dataset of preferences, interests, and demographics of young people, aged between 15-30, of Slovakian nationality [51]. The data was gathered in 2013 by students of an statistics class at FSEV UK through friends and families. The dataset consists of 1,010 records with 150 features (139 integer and 11 categorical) including height, weight, music preferences, eating habits, etc. This dataset is popular for analysis on Kaggle (an online platform for sharing data) with 99.4k views and 21.8k downloads on all continents within the last three years. Noting that the preferences of the individuals can be matched with publicly available datasets, such as IMDb (Internet Movie Database), to potentially identify the individuals, there is a need for obfuscating the data in order to avoid privacy breaches related to age, weight, and height.*

Assume that we use the policy (11) *for reporting weight and height of individuals so that potential future insurance agencies cannot test for obesity levels. Figure* 5 *illustrates the histogram of the reported weight and hight of individuals for various levels of accuracy ρ with darker colors showing higher frequencies. As expected, by decreasing ρ, the accuracy gets worse; the histogram changes more drastically.* △

We can relate the notion of privacy in Definition 3 with non-stochastic information leakage used in [34]. This is investigated in the following proposition.

Proposition 1 (ϵ-Privacy versus Non-Stochastic Information Leakage). $L_0(Y; H) \leq -\log_e(1 - \exp(-\epsilon))$.

Proof: See Appendix 13. □

Figure 6 illustrates the upper bound, developed in Proposition 1, for the non-stochastic information leakage $L_0(Y; H)$ versus ϵ for ϵ-privacy. As intuitively expected, by increasing privacy ϵ, the information leakage reduces. Note that the upper bound approaches zero as ϵ grows and the non-stochastic information leakage is also lower bounded by zero; thus the non-stochastic information leakage $L_0(Y; H)$ must also approach zero as ϵ increases.

6 Indistinguishability

In this section, we use the non-stochastic hypothesis-testing framework to adapt identifiability from the privacy literature to a non-stochastic setup. We prove that concatenation of linear functions with quantizers (also known as binning) can guarantee indistinguishability if the quantizer has "few enough" levels. The upper bound on the number of the levels is a function of the sensitivity of the mapping, the range of the data, and the privacy level.

Consider a private dataset, modelled by uncertain variable $X : \Omega \to \mathbb{R}^n$, where n denotes the number of individuals whose data is available in the dataset. The realisations of the dataset are essentially

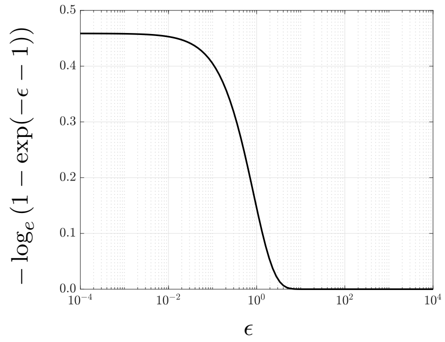

Fig. 6: Upper bound of the non-stochastic information leakage $L_0(Y;H)$ versus ϵ for ϵ-privacy.

vectors of the form

$$X(\omega)=\begin{bmatrix}X_1(\omega)\\X_2(\omega)\\\vdots\\X_n(\omega)\end{bmatrix},$$

where $X_i:\Omega\to\mathbb{R}$, $1\leq i\leq n$, is the uncertain variable modelling the data of i-th individual. The sender or curator, as in the previous section, must return an answer to a query, which is essentially a function $f:[\![X]\!]\to\mathbb{R}$ that must be evaluated for the realization of the uncertain variable X available to the curator. The curator must compute and provide the response $Y(\omega)=f(X(\omega))$. Note that $Y=f\circ X$ is an uncertain variable. Motivated by the definition of semantic security, or indistinguishability under chosen plaintext attack [44], we define the notation of indistinguishability for privacy. Assume that an adversary selects $x_i,x_i'\in[\![X_i]\!]$ and provides this information to the curator. The curator inserts uncertain variable $X_i:\Omega\to[\![X]\!]:=\{x_i,x_i'\}$ into the vector-valued uncertain variable X instead of the original data of individual i, generates a realization $X(\omega)$, computes $Y(\omega)=f(X(\omega))$, and provides $Y(\omega)$ to the adversary. The adversary then tests if the realization of the data of individual i is equal to x_i or x_i' (knowing that it is bound to be one of those values). To this aim, we can define $g_H:X(\omega)\mapsto H(\omega)$ as

$$H(\omega)=g_H(X(\omega))=\begin{cases}p_0, & X_i(\omega)=x_i,\\p_1, & X_i(\omega)=x_i'.\end{cases}$$

The uncertain variable $H=g_H\circ X$ models the hypothesis uncertain variable. If Y is a continuous uncertain variable, for any test T, Theorem 2 states that the performance of the adversary is bounded from the above by

$$\begin{aligned}\mathcal{P}(T)&\leq\log_e(\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!]))\\&=\log_e(\mu([\![Y|X_i(\omega)=x_i]\!]\Delta[\![Y|X_i(\omega)=x_i']\!])).\end{aligned}$$

Similarly, if Y is a discrete uncertain variable, we have

$$\mathcal{P}(T)\leq\log_2(|[\![Y|X_i(\omega)=x_i]\!]\Delta[\![Y|X_i(\omega)=x_i']\!]|).$$

Therefore, if $\log_e(\mu([\![Y|X_i(\omega)=x_i]\!]\Delta[\![Y|X_i(\omega)=x_i']\!]))$ or $\log_2(|[\![Y|X_i(\omega)=x_i]\!]\Delta[\![Y|X_i(\omega)=x_i']\!]|)$ is large, the adversary's performance can be good. For privacy preservation, we must ensure that the upper bound of $\mathcal{P}(T)$ is small. This allows for the introduction the following notion of privacy.

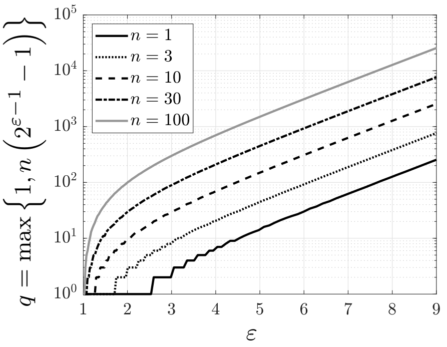

Fig. 7: Number of the quantization levels q versus ε for ε-indistinguishable.

Definition 5 (ε-Indistinguishability). *For continuous uncertain variable Y, the query f is ε-indistinguishable, for $\varepsilon>0$, if*

$$\mu([\![Y|X_i(\omega)=x_i]\!]\Delta[\![Y|X_i(\omega)=x_i']\!])\leq\exp(\varepsilon),\quad\forall x_i,x_i'\in[\![X_i]\!],\forall i.\quad(12)$$

Similarly, for discrete uncertain variable Y, the query f is ε-indistinguishable, for $\varepsilon\in\mathbb{N}$, if

$$|[\![Y|X_i(\omega)=x_i]\!]\Delta[\![Y|X_i(\omega)=x_i']\!]|\leq2^{\varepsilon},\quad\forall x_i,x_i'\in[\![X_i]\!],\forall i.\quad(13)$$

This notion of privacy is in essence close to identifiability [6, 8] for which privacy preservation relates to the potential of an adversary identifying the private data of individuals based on the received outputs. However, in identifiability, additive noise is used to complicate the adversary's task while, in indistinguishability, non-stochastic approaches, such as binning or quantization, are utilized.

Definition 6 (Quantizer). *A q-level quantizer $\mathcal{Q}:[x_{\min},x_{\max}]\to\{b_1,\dots,b_q\}$ is a piecewise constant function defined as*

$$\mathcal{Q}(x)=\begin{cases}b_1, & x\in[x_1,x_2),\\b_2, & x\in[x_2,x_3),\\\vdots & \vdots\\b_{q-1}, & x\in[x_{q-1},x_q),\\b_q, & x\in[x_q,x_{q+1}],\end{cases}$$

where $(b_i)_{i=1}^q$ are distinct symbols and $x_1\leq x_2\leq\dots\leq x_q$ are real numbers such that $x_1=x_{\min}$, $x_{q+1}=x_{\max}$, $x_{i+1}-x_i=(x_{\max}-x_{\min})/q$ for all $1\leq i\leq q$.

We can show that linear quantizers can achieve indistinguishability for linear queries on private datasets. This is proved in the next theorem.

Theorem 5. *Assume $f(x)=c^\top x$ and $[\![X_i]\!]=[x_{\min},x_{\max}]$, $\forall i$. Then $\mathcal{Q}\circ f$ with q-level quantizer $\mathcal{Q}$ is ε-indistinguishable if*

$$q\leq\max\left\{1,\frac{\varpi}{\|c\|_\infty(x_{\max}-x_{\min})}\left(2^{\varepsilon-1}-1\right)\right\},$$

with

$$\varpi:=\max_{x\in[x_{\min},x_{\max}]^n}c^\top x-\min_{x\in[x_{\min},x_{\max}]^n}c^\top x.$$

Proof: See Appendix 14. □

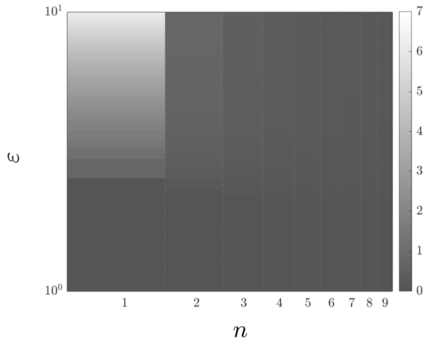

Fig. 8: Non-stochastic information leakage $L_0(X;Y)$ versus n and ε for the averaging query. The lighter colors show higher leakage.

Theorem 5 shows that concatenation of linear mappings with quantizers can ensure ϵ-indistinguishability if the quantizer has "few enough" levels. The upper bound on the number of the levels is a function of the sensitivity of the linear mapping ϖ, the range of the data $x_{\max} - x_{\min}$, and indistinguishability budget ε.

Example 4 (Averaging Query)**.** *The number of bins or the resolution of quantizer can be simplified for averaging queries. This allows us to illustrate the relationship between indistinguishability and non-stochastic information leakage. For the averaging query $f(x) = (x_1 + \cdots + x_n)/n$, $\mathcal{Q} \circ f$ with q-level quantizer $\mathcal{Q}$ is ε-indistinguishable if $q \leq \max\{1, n(2^{\varepsilon-1} - 1)\}$. Figure 7 illustrates the number of the quantization levels versus ε. For larger n, the quantizer resolution can be significantly higher for the same level of privacy. This is intuitive as privacy can be more easily preserved when hiding the data of individuals in aggregate reports of many people. Figure 8 shows the relationship between indistinguishability and non-stochastic information leakage, which was used as a measure of privacy in [34]. Clearly, as expected, the information leakage reduces significantly by decreasing ε or by increasing n. This is in line with the observations in [34] indicating that the quantization or binning are optimal privacy-preserving policies in the non-stochastic framework.* △

7 Conclusions and Future Work

We considered privacy against hypothesis testing adversaries using the theory of non-stochastic hypothesis testing. We constructed reporting policies with prescribed privacy and utility guarantees. Future work can focus on multiple directions:

- **Optimal trade-off**: We should develop optimal policies for balancing between privacy and accuracy. Note that the results in Theorems 4 and 5 are only sufficient for ensuring ϵ-privacy and ρ-accuracy and may not capture the optimal trade-off between privacy and accuracy.
- **Application**: We should investigate the application of the proposed notions of privacy, both ϵ-privacy and ε-indistinguishability, in practical scenarios, such as machine learning based on sensitive data (e.g., health or finance data), smart meter privacy, and traffic estimation based on crowd-sensed data. In these scenarios, there are many established adversaries for privacy attacks. It is important to investigate the effect of the proposed notions of privacy on the success of those privacy attacks, establishing practical implications and guarantees of the non-stochastic privacy-preserving policies.
- **Properties**: Many stochastic privacy-preserving policies, such as differential privacy, enjoy useful properties, such as composition. To be able to use the proposed non-stochastic notions of privacy in dynamic environments, such as when training machine learning models using gradient descent algorithm, we need to establish those results for non-stochastic privacy notions as well. This allows us to split the privacy budget across multiple reports.

8 References

1 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.

2 C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

3 J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pp. 429–438, IEEE, 2013.

4 P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," in *Advances in Neural Information Processing Systems*, pp. 2879–2887, 2014.

5 A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber, "Privacy: Theory meets practice on the map," in *Proceedings of the 2008 IEEE 24th International Conference on Data Engineering*, pp. 277–286, IEEE Computer Society, 2008.

6 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.

7 J. Lee and C. Clifton, "Differential identifiability," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1041–1049, ACM, 2012.

8 A. Bkakria, N. Cuppens-Boulahia, and F. Cuppens, "Linking differential identifiability with differential privacy," in *International Conference on Information and Communications Security*, pp. 232–247, Springer, 2018.

9 F. Farokhi and G. Nair, "Privacy-constrained communication," *IFAC-PapersOnLine*, vol. 49, no. 22, pp. 43–48, 2016.

10 M. J. Wainwright, M. I. Jordan, and J. C. Duchi, "Privacy aware learning," in *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pp. 1430–1438, 2012.

11 Y. Liang, H. V. Poor, and S. Shamai, "Information theoretic security," *Foundations and Trends in Communications and Information Theory*, vol. 5, no. 4–5, pp. 355–580, 2009.

12 L. Lai, S.-W. Ho, and H. V. Poor, "Privacy–security trade-offs in biometric security systems–Part I: Single use case," *IEEE Transactions on Information Forensics and Security*, vol. 6, no. 1, pp. 122–139, 2011.

13 Z. Li and T. Oechtering, "Privacy on hypothesis testing in smart grids," in *IEEE Information Theory Workshop (ITW) 2015, Jeju, Korea, Oct. 11-15, 2015*, pp. 337–341, IEEE, 2015.

14 F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4726–4734, 2018.

15 L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoffs in databases: An information-theoretic approach," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 6, pp. 838–852, 2013.

16 F. Farokhi, J. Milosevic, and H. Sandberg, "Optimal state estimation with measurements corrupted by Laplace noise," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pp. 302–307, IEEE, 2016.

17 R. Bild, K. A. Kuhn, and F. Prasser, "SafePub: A truthful data anonymization algorithm with strong privacy guarantees," *Proceedings on Privacy Enhancing Technologies*, vol. 2018, no. 1, pp. 67–87, 2018.

18 J. Bambauer, K. Muralidhar, and R. Sarathy, "Fool's gold: an illustrated critique of differential privacy," *Vanderbilt Journal of Entertainment & Technology Law*, vol. 16, p. 701, 2013.

19 R. Bhaskar, A. Bhowmick, V. Goyal, S. Laxman, and A. Thakurta, "Noiseless database privacy," in *International Conference on the Theory and Application of Cryptology and Information Security*, pp. 215–232, Springer, 2011.

20 S. U. Nabar, B. Marthi, K. Kenthapadi, N. Mishra, and R. Motwani, "Towards robustness in query auditing," in *Proceedings of the 32nd international conference on Very large data bases*, pp. 151–162, VLDB Endowment, 2006.

21 A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, no. 8, pp. 1355–1387, 1975.

22 T. Courtade, "Information masking and amplification: The source coding setting," in *Proceedings of the IEEE International Symposium on Information Theory Proceedings (ISIT)*, pp. 189–193, 2012.

23 H. Yamamoto, "A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers," *IEEE Transactions on Information Theory*, vol. 29, no. 6, pp. 918–923, 1983.

24 H. Yamamoto, "A rate-distortion problem for a communication system with a secondary decoder to be hindered," *IEEE Transactions on Information Theory*, vol. 34, no. 4, pp. 835–842, 1988.

25 P. Samarati, "Protecting respondents identities in microdata release," *IEEE transactions on Knowledge and Data Engineering*, vol. 13, no. 6, pp. 1010–1027, 2001.

26 L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.

27 A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramaniam, "ℓ-diversity: privacy beyond k-anonymity," in *22nd International Conference on Data Engineering (ICDE'06)*, pp. 24–24, 2006.

28 A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, pp. 111–125, IEEE, 2008.

29 J. Su, A. Shukla, S. Goel, and A. Narayanan, "De-anonymizing web browsing data with social networks," in *Proceedings of the 26th International Conference on World Wide Web*, pp. 1261–1269, 2017.

30 Y.-A. De Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," *Scientific reports*, vol. 3, p. 1376, 2013.

31 G. Poulis, A. Gkoulalas-Divanis, G. Loukides, S. Skiadopoulos, and C. Tryfonopoulos, "Secreta: A tool for anonymizing relational, transaction and rt-datasets," in *Medical Data Privacy Handbook* (A. Gkoulalas-Divanis and G. Loukides, eds.), pp. 83–109, Springer International Publishing, 2015.

32 F. K. Dankar and K. El Emam, "Practicing differential privacy in health care: A review," *Transactions on Data Privacy*, vol. 6, no. 1, pp. 35–67, 2013.

33 J. Mervis, "Researchers object to census privacy measure," *Science*, vol. 363, no. 6423, pp. 114–114, 2019.

34 F. Farokhi, "Development and analysis of deterministic privacy-preserving policies using non-stochastic information theory," *IEEE Transactions on Information Forensics and Security*, 2019. In Press.

35 R. V. L. Hartley, "Transmission of information," *Bell System Technical Journal*, vol. 7, no. 3, pp. 535–563, 1928.

36 A. N. Kolmogorov and V. M. Tikhomirov, "ε-entropy and ε-capacity of sets in function spaces," *Uspekhi Matematicheskikh Nauk*, vol. 14, no. 2, pp. 3–86, 1959. English translation American Mathematical Society Translations, series 2, vol. 17, pp. 277–364.

37 A. Renyi, "On measures of entropy and information," in *Proc. of the Fourth Berkeley Symp. on Math. Statist. and Prob.*, vol. 1, pp. 547–561, 1961.

38 G. N. Nair, "A nonstochastic information theory for communication and state estimation," *IEEE Transactions on Automatic Control*, vol. 58, no. 6, pp. 1497–1510, 2013.

39 D. Jagerman, "ε-entropy and approximation of bandlimited functions," *SIAM Journal on Applied Mathematics*, vol. 17, no. 2, pp. 362–377, 1969.

40 G. N. Nair, "A nonstochastic information theory for feedback," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pp. 1343–1348, IEEE, 2012.

41 P. Duan, F. Yang, S. L. Shah, and T. Chen, "Transfer zero-entropy and its application for capturing cause and effect relationship between variables," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 3, pp. 855–867, 2015.

42 M. Wiese, K. H. Johansson, T. J. Oechtering, P. Papadimitratos, H. Sandberg, and M. Skoglund, "Uncertain wiretap channels and secure estimation," in *Information Theory (ISIT), 2016 IEEE International Symposium on*, pp. 2004–2008, IEEE, 2016.

43 R. F. Barber and J. Duchi, "Privacy: A few definitional aspects and consequences for minimax mean-squared error," in *53rd IEEE Conference on Decision and Control*, pp. 1365–1369, IEEE, 2014.

44 J. Katz and Y. Lindell, *Introduction to Modern Cryptography, Second Edition*. Chapman & Hall/CRC Cryptography and Network Security Series, Taylor & Francis, 2 ed., 2014.

45 F. Farokhi, "Non-stochastic hypothesis testing with application to privacy against hypothesis-testing adversary," in *Proceedings of the Decision and Control (CDC), 2019 IEEE 58th Conference on*, 2019.

46 H. Shingin and Y. Ohta, "Disturbance rejection with information constraints: Performance limitations of a scalar system for bounded and gaussian disturbances," *IFAC Proceedings Volumes*, vol. 42, no. 20, pp. 304–309, 2009.

47 F. Farokhi and N. Ding, "Measuring information leakage in non-stochastic brute-force guessing," *arXiv preprint arXiv:2004.10911*, 2020.

48 I. Sason and S. Verdú, "Arimoto–Rényi conditional entropy and Bayesian m-ary hypothesis testing," *IEEE Transactions on Information theory*, vol. 64, no. 1, pp. 4–25, 2017.

49 B. Yu, "Assouad, Fano, and Le Cam," in *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics* (D. Pollard, E. Torgersen, and G. L. Yang, eds.), pp. 423–435, New York, NY: Springer New York, 1997.

50 P. Billingsley, *Probability and Measure*. John Wiley & Sons, 3 ed., 1995.

51 M. Sabo, "Young people survey: Explore the preferences, interests, habits, opinions, and fears of young people." available online at Kaggle.com, last visit: 8-Mar-2019. https://www.kaggle.com/miroslavsabo/young-people-survey.

9 Proof of Theorem 1

First, we proved three important claims.

Claim 1: $[\![H|[\![X|y]\!]]\!] = \{p_0, p_1\}, \forall y \in [\![Y|p_1]\!] \cap [\![Y|p_0]\!]$.

The proof for this claim is as follows. For any $y \in [\![Y|p_i]\!]$, there exists $x \in g_Y^{-1}(y) = [\![X|y]\!]$ such that $g_H(x) = p_i$. Note that, for any mapping $g : x \mapsto y$, we define the inverse image $g^{-1}(y) := \{x : g(x) = y\}$. Therefore, $\{p_i\} \subseteq g_H([\![X|y]\!]) = [\![H|[\![X|y]\!]]\!]$. This implies that, for any $y \in [\![Y|p_1]\!] \cap [\![Y|p_0]\!]$, $\{p_0, p_1\} \subseteq [\![H|[\![X|y]\!]]\!] \subseteq [\![H]\!] = \{p_0, p_1\}$.

Claim 2: $[\![H|[\![X|y]\!]]\!] = \{p_0\}, \forall y \in [\![Y|p_0]\!] \setminus [\![Y|p_1]\!]$.

The proof for this claim is as follows. If $y \notin [\![Y|p_1]\!]$, there must not exist $x \in [\![X|y]\!]$ such that $g_H(x) = p_1$. Therefore, $p_1 \notin [\![H|[\![X|y]\!]]\!]$. Therefore, for any $y \in [\![Y|p_0]\!] \setminus [\![Y|p_1]\!]$, it must be that $p_0 \in [\![H|[\![X|y]\!]]\!]$ and $p_1 \notin [\![H|[\![X|y]\!]]\!]$. This is only possible if $[\![H|[\![X|y]\!]]\!] = \{p_0\}$.

Claim 3: $[\![H|[\![X|y]\!]]\!] = \{p_1\}, \forall y \in [\![Y|p_1]\!] \setminus [\![Y|p_0]\!]$.

The proof for this claim is the same as *Claim 2*.

With these claims in hand, we are ready to prove the lemma. For any test T, we have

$$\begin{aligned}\aleph(T) =&\{y \in [\![Y]\!] : [\![H|[\![X|y]\!]]\!] = \{T(y)\}\}\\ =&\{y \in [\![Y|p_0]\!] \setminus [\![Y|p_1]\!] : [\![H|[\![X|y]\!]]\!] = \{T(y)\}\}\\ &\cup \{y \in [\![Y|p_0]\!] \cap [\![Y|p_1]\!] : [\![H|[\![X|y]\!]]\!] = \{T(y)\}\}\\ &\cup \{y \in [\![Y|p_1]\!] \setminus [\![Y|p_0]\!] : [\![H|[\![X|y]\!]]\!] = \{T(y)\}\}\\ =&\{y \in [\![Y|p_0]\!] \setminus [\![Y|p_1]\!] : [\![H|[\![X|y]\!]]\!] = \{T(y)\}\}\\ &\cup \{y \in [\![Y|p_1]\!] \setminus [\![Y|p_0]\!] : [\![H|[\![X|y]\!]]\!] = \{T(y)\}\},\end{aligned} \tag{14}$$

where the last equality follows from that, by *Claim 1*, $[\![H|[\![X|y]\!]]\!] = \{p_0, p_1\}$ while $\{T(y)\}$ can be either $\{p_0\}$ or $\{p_1\}$. From (14), we get

$$\begin{aligned}\aleph(T) \subseteq&([\![Y|p_0]\!] \setminus [\![Y|p_1]\!]) \cup ([\![Y|p_1]\!] \setminus [\![Y|p_0]\!])\\ =&[\![Y|p_0]\!]\Delta[\![Y|p_1]\!],\end{aligned}$$

where Δ denotes the symmetric difference operator on the operands. As a result, if Y is a continuous uncertain variable, for any test T, we get

$$\mathcal{P}(T) = \log_e(\mu(\aleph(T))) \leq \log(\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!])). \tag{15}$$

Similarly, if Y is a discrete uncertain variable, we get

$$\mathcal{P}(T) = \log_2(|\aleph(T)|) \leq \log(|[\![Y|p_0]\!]\Delta[\![Y|p_1]\!]|). \tag{16}$$

For any consistent test T, by application of the definition of consistency in (14), we have

$$\begin{aligned}\aleph(T) =&\{y \in [\![Y|p_0]\!] \setminus [\![Y|p_1]\!] : [\![H|[\![X|y]\!]]\!] = \{p_0\}\}\\ &\cup \{y \in [\![Y|p_1]\!] \setminus [\![Y|p_0]\!] : [\![H|[\![X|y]\!]]\!] = \{p_1\}\},\end{aligned}$$

which, by the virtue of *Claims 2* and *3*, results in $\aleph(T) = ([\![Y|p_0]\!] \setminus [\![Y|p_1]\!]) \cup ([\![Y|p_1]\!] \setminus [\![Y|p_0]\!])$. This shows that consistent tests attain the upper bound on the performance in (15) and (16).

10 Proof of Theorem 2

The upper bound in the statement of the theorem follows from the proof of Theorem 1; see (15) and (16).

11 Proof of Theorem 3

Using Le Cam's inequality [49], we get

$$\inf_{p_{\widehat{H}|Y}} \mathbb{P}\{\widehat{H} \neq H\} = 1 - \nu(p_{Y|p_0}, p_{Y|p_1}),$$

where ν is the total variation distance defined as

$$\nu(\xi_1, \xi_2) := \frac{1}{2}\int_{\mathrm{supp}(\xi_1)\cup\mathrm{supp}(\xi_2)} |\xi_1(x) - \xi_2(x)|\mathrm{d}x.$$

Hence,

$$\begin{aligned}\sup_{p_{\widehat{H}|Y}} \mathbb{P}\{\widehat{H} = H\} =& \sup_{p_{\widehat{H}|Y}} [1 - \mathbb{P}\{\widehat{H} \neq H\}]\\ =& 1 - \inf_{p_{\widehat{H}|Y}} \mathbb{P}\{\widehat{H} \neq H\} = \nu(p_{Y|p_0}, p_{Y|p_1}).\end{aligned}$$

Evidently, $\operatorname{supp}(p_{Y|p_0}) = [\![Y|p_0]\!]$, $\operatorname{supp}(p_{Y|p_1}) = [\![Y|p_1]\!]$, and thus $\operatorname{supp}(p_{Y|p_0}) \cup \operatorname{supp}(p_{Y|p_1}) = [\![Y]\!]$. We can see that

$$\begin{aligned}
\nu(p_{Y|p_0}&,p_{Y|p_1})\\
=&\frac{1}{2}\int_{[\![Y|p_0]\!]\setminus[\![Y|p_1]\!]} |p_{Y|p_0}(y) - p_{Y|p_1}(y)|\mathrm{d}y\\
&+\frac{1}{2}\int_{[\![Y|p_1]\!]\setminus[\![Y|p_0]\!]} |p_{Y|p_0}(y) - p_{Y|p_1}(y)|\mathrm{d}y\\
&+\frac{1}{2}\int_{[\![Y|p_0]\!]\cap[\![Y|p_1]\!]} |p_{Y|p_0}(y) - p_{Y|p_1}(y)|\mathrm{d}y\\
\geq&\frac{1}{2}\int_{[\![Y|p_0]\!]\setminus[\![Y|p_1]\!]} p_{Y|p_0}(y)\mathrm{d}y\\
&+\frac{1}{2}\int_{[\![Y|p_1]\!]\setminus[\![Y|p_0]\!]} p_{Y|p_1}(y)\mathrm{d}y,
\end{aligned}$$

where the inequality follows from that $p_{Y|p_1}(y) = 0$ for $y \in [\![Y|p_0]\!] \setminus [\![Y|p_1]\!]$, $p_{Y|p_0}(y) = 0$ for $y \in [\![Y|p_1]\!] \setminus [\![Y|p_0]\!]$, and $\int_{[\![Y|p_0]\!]\cap[\![Y|p_1]\!]} |\xi_1(x) - \xi_2(x)|\mathrm{d}x \geq 0$. Finally, noting that $p_{Y|p_0}(y), p_{Y|p_1}(y) \geq \varepsilon$, we get

$$\begin{aligned}
\nu(p_{Y|p_0}, p_{Y|p_1}) \geq&\frac{\varepsilon}{2}\left(\int_{[\![Y|p_0]\!]\setminus[\![Y|p_1]\!]}\mathrm{d}y + \int_{[\![Y|p_1]\!]\setminus[\![Y|p_0]\!]}\mathrm{d}y\right)\\
=&\frac{\varepsilon}{2}\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!]).
\end{aligned}$$

This concludes the proof.

12 Proof of Theorem 4

The proof of ρ-accuracy follows from that $\|X - g_Y(X)\| \leq |x_i - g(x_{-i})| \leq 1/\rho$. The proof for ϵ-privacy follows from that, if $[\![Y]\!]$, $([\![Y|p_0]\!]\Delta[\![Y|p_1]\!])$, $([\![Y|p_0]\!] \cap [\![Y|p_1]\!])$ are Lebesgue measurable, we get $\mu([\![Y]\!]) = \mu(([\![Y|p_0]\!]\Delta[\![Y|p_1]\!]) \cup ([\![Y|p_0]\!] \cap [\![Y|p_1]\!])) \geq \mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!]) + \mu([\![Y|p_0]\!] \cap [\![Y|p_1]\!])$ because $([\![Y|p_0]\!]\Delta[\![Y|p_1]\!]) \cup ([\![Y|p_0]\!] \cap [\![Y|p_1]\!]) = [\![Y|p_0]\!] \cup [\![Y|p_1]\!] = [\![Y]\!]$.

13 Proof of Proposition 1

Note that

$$\begin{aligned}
L_0(Y;H) =& \max_{p\in\{p_0,p_1\}} \log_e\left(\frac{\mu([\![Y]\!])}{\mu([\![Y|p]\!])}\right)\\
=&\log_e\left(\frac{\mu([\![Y]\!])}{\min\limits_{p\in\{p_0,p_1\}}\mu([\![Y|p]\!])}\right)\\
\leq&\log_e\left(\frac{\mu([\![Y]\!])}{\mu([\![Y|p_0]\!] \cap [\![Y|p_1]\!])}\right),
\end{aligned}$$

where the inequality follows from that $\mu([\![Y|p_0]\!] \cap [\![Y|p_1]\!]) \leq \min_{p\in\{p_0,p_1\}} \mu([\![Y|p]\!])$, because $[\![Y|p_0]\!] \cap [\![Y|p_1]\!] \subseteq \mu([\![Y|p_0]\!]$ and $[\![Y|p_0]\!] \cap [\![Y|p_1]\!] \subseteq \mu([\![Y|p_1]\!]$. Furthermore,

$$\begin{aligned}
\frac{\mu([\![Y|p_0]\!] \cap [\![Y|p_1]\!])}{\mu([\![Y]\!])}+&\frac{\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!])}{\mu([\![Y]\!])}\\
&\geq \frac{\mu([\![Y|p_0]\!] \cup [\![Y|p_1]\!])}{\mu([\![Y]\!])}\\
&= \frac{\mu([\![Y]\!])}{\mu([\![Y]\!])} = 1,
\end{aligned}$$

and as a result

$$\begin{aligned}
\frac{\mu([\![Y|p_0]\!] \cap [\![Y|p_1]\!])}{\mu([\![Y]\!])} &\geq 1 - \frac{\mu([\![Y|p_0]\!]\Delta[\![Y|p_1]\!])}{\mu([\![Y]\!])}\\
&\geq 1 - \exp(-\epsilon - 1).
\end{aligned}$$

Therefore, $L_0(Y;H) \leq -\log_e\left(1 - \exp(-\epsilon - 1)\right)$.

14 Proof of Theorem 5

By definition, $\mu([\![c^\top x]\!]) = \varpi$. Hence, ϖ/q is the resolution of the quantizer. Since $f(x) = c^\top x$ is continuous, there exists $\underline{x} \leq \overline{x}$ such that $[\![f(X)|X_i = x_i]\!] = [\underline{x}, \overline{x}]$. In fact,

$$\begin{aligned}
\underline{x} =& c_i x_i + \min_{x_j\in[\![X_j]\!],\forall j\neq i} \sum_{j\neq i} c_j x_j,\\
\overline{x} =& c_i x_i + \max_{x_j\in[\![X_j]\!],\forall j\neq i} \sum_{j\neq i} c_j x_j.
\end{aligned}$$

The connectivity of $[\![f(X)|X_i = x_i]\!]$ follows from the linearity of $f(x) = c^\top x$ and the convexity of $[\![X|X_i = x_i]\!]$. Furthermore, $[\![f(X)|X = x']\!] \subseteq [\underline{x} + \delta, \overline{x} + \delta]$ where $\delta := c_i(x_i' - x_i)$ with $|\delta| \leq c_i(x_{\max} - x_{\min})$. Then,

$$\begin{aligned}
|\mathcal{Q}([\![f(X)|X_i = x_i]\!])\Delta\mathcal{Q}([\![f(X)|X_i = x_i']\!])|&\\
\leq \begin{cases} 2\lceil q|\delta|/\varpi\rceil, & q > 1,\\ 0, & q = 1.\end{cases}&
\end{aligned}$$

The upper bound comes from that, in the most informative case, $[\underline{x} + \delta, \overline{x} + \delta]\Delta[\underline{x}, \overline{x}]$ is fully quantized. Selecting $q \leq (2^{\varepsilon-1} - 1)\varpi/(\max_i |c_i|(x_{\max} - x_{\min}))$ guarantees ε-indistinguishability.