

P C Privacy: Automatic Privacy Measurement and Control of Data Processing

Ha she Xiao
hsxiao@mit.edu

Sri ivas Devadas
devadas@mit.edu

June 21, 2023

Abstract

We propose and study a new privacy definition, termed Probably ϵ -approximately Correct (P C) Privacy. P C Privacy characterizes the information-theoretic hardness to recover sensitive data given arbitrary information disclosure/leakage during/after any processing. Unlike the classic cryptographic definition and Differential Privacy (DP), which consider the *adversarial (i put-i depe de t) worst case*, P C Privacy is a simulatable metric that quantifies *the i sta ce-based* impossibility of inference. A fully automatic analysis and proof generation framework is proposed: security parameters can be produced with arbitrarily high confidence via Monte-Carlo simulation for any black-box data processing oracle. This appealing automation property enables analysis of complicated data processing, where the worst-case proof in the classic privacy regime could be loose or even intractable. Moreover, we show that the produced P C Privacy guarantees enjoy simple composition bounds and the automatic analysis framework can be implemented in an online fashion to analyze the composite P C Privacy loss even under correlated randomness. On the utility side, the magnitude of (necessary) perturbation required in P C Privacy is *not* lower bounded by $\Theta(\sqrt{d})$ for a d -dimensional release but could be $O(1)$ for many practical data processing tasks, which is in contrast to the input-independent worst-case information-theoretic lower bound. Example applications of P C Privacy are included with comparisons to existing works.

Keywords: P C Privacy; Automatic Security Proof; f -Divergence; Mutual Information; Instance-based Posterior Advantage; Reconstruction Hardness; Inference Hardness; Membership Attack.

1 Introduction

Privacy concerns with information leakage from data processing are receiving increasing attention. Obtaining usable security which provides high utility, low implementation overhead and meaningful interpretation, is challenging. Theoretically, the underlying problem can be described by the following generic model. Assume some sensitive data X is processed by some mechanism \mathcal{M} and the output $\mathcal{M}(X)$ is published. We want to characterize, given $\mathcal{M}(X)$, how much information an adversary can infer about X . In the applications of data analysis and machine learning, $\mathcal{M}(X)$ could be any statistics or model learned from X [1]; in encryption, $\mathcal{M}(X)$ could be ciphertexts of X [37]; in side channel attacks, $\mathcal{M}(X)$ could be any physical signal observed by the adversary, such as memory access pattern [59] or network traffic [20] when accessing or sending X .

To formalize and quantify the information leakage, a large number of security metrics have been proposed varying from the perfect secrecy [56] of Shannon, to computational indistinguishability [38], which forms the foundation of modern cryptography, to Differential Privacy (DP) [29] and its variants based on various statistical divergence metrics [49], and many other definitions based on information-theoretic quantities, such as Fisher-information privacy [32].

Most existing privacy or security definitions (including all above-mentioned examples) are rooted in indistinguishability likelihoods. They set out to quantify the difference between the likelihoods $\mathbb{P}(\mathcal{M}(X) = o|X = X_0)$ and $\mathbb{P}(\mathcal{M}(X) = o|X = X'_0)$ for different input selections from various perspectives. For instance,

perfect secrecy requires that for arbitrary input data X , the distribution of $\mathcal{M}(X)$ must be identical. Cryptographic security relaxes the statistical indistinguishability to that any polynomial-time adversary cannot distinguish the distributions of $\mathcal{M}(X_0)$ and $\mathcal{M}(X'_0)$ for any input candidates X_0 and X'_0 . Similarly, DP and its variants quantify the difference between $\mathcal{M}(X_0)$ and $\mathcal{M}(X'_0)$ for any two *adjacent datasets*¹ by some divergence function $\mathcal{D}(\mathbb{P}_{\mathcal{M}(X_0)}\|\mathbb{P}_{\mathcal{M}(X'_0)})$. For example, classic ϵ -DP [29] is based on maximal divergence $\mathcal{D}_\infty(\mathbb{P}_{\mathcal{M}(X_0)}\|\mathbb{P}_{\mathcal{M}(X'_0)}) = \sup_o \log(\mathbb{P}(\mathcal{M}(X_0) = o)/\mathbb{P}(\mathcal{M}(X'_0) = o))$.

The interested reader may wonder if there is an intuitive privacy interpretation of the limited likelihood difference? In particular, can likelihood indistinguishability tell us given $\mathcal{M}(X)$, how much of the sensitive data X an adversary can recover? Unfortunately, none of above-mentioned definitions can directly answer this question. This is *not* because those metrics are incomplete but because additional assumptions must be made. In general, to quantify reconstruction hardness, we have to specify the adversary’s prior knowledge and the (conditional) entropy of objective X . In the extreme case where the adversary already has the full knowledge of X , or where X is simply a public constant, then no matter what kind of mechanism \mathcal{M} is applied to encode X , the adversary can perfectly recover it, regardless of $\mathcal{M}(X)$. Thus, *the likelihood difference only partially captures the changes to the adversary’s belief or the additional information regarding X after observing $\mathcal{M}(X)$.*

Compared to the extensive studies along the line of (in)distinguishability, formal characterization of reconstruction hardness is a long-standing, challenging problem. The estimation of the priors regarding sensitive data X and resultant output distribution of $\mathcal{M}(X)$ could already be intractable – this is especially likely in the high-dimensional scenario. Thus, as mentioned before, most security/privacy definitions only focus on the likelihood divergence and avoid complicated, such as Bayesian, analysis to properly quantify the posterior inference hardness. As a consequence, to make a meaningful guarantee, most classic definitions treat the sensitive input *deterministically* and consider the adversarial worst-case, i.e., the maximal distribution divergence between $\mathcal{M}(X_0)$ and $\mathcal{M}(X'_0)$ for two arbitrary input candidates X_0 and X'_0 . Such worst-case indistinguishability is a very strong guarantee since it is independent of any particular input distribution, but is also expensive, especially for statistical data processing. At a high level, there are two key obstacles in the application of such indistinguishability measurement:

- (a) **Adversarial Worst-case Proof is Hard:** It is, in general, non-trivial to tightly or analytically characterize the worst case and privatize an (adaptively-iterative) algorithm to ensure satisfied output divergence. For example, in DP, calculation of the sensitivity, the maximal possible change on the output by replacing a single datapoint, is, in general, NP-hard [66]. Even in common applications such as to differentially privately train a neural network, the tight sensitivity bound of Stochastic Gradient Descent (SGD) remains open. The only known generic privatization approach, DP-SGD, restricts the sensitivity using clipped/normalized per-sample gradient, and adds noise step by step to derive an upper bound in a compositional fashion [1]. However, the tight composition bound calculation is also known to be #P hard [51]. Though many efforts have been devoted to obtaining efficient and tighter composition [49], the applications of DP-SGD in modern deep learning, and large-scale private database queries still encounter severe utility compromise. Therefore, an automatic and usable privacy analysis framework for arbitrary data processing is desirable.
- (b) **Data Entropy Matters in Reconstruction:** Reconstruction (or inference) of data is, in general, a harder problem for an adversary than simply distinguishing [39]. However, worst-case indistinguishability cannot capture the data entropy. We give a simple example. A dataset $X = \{x_1, x_2, \dots, x_n\}$ is generated as i.i.d. samples from some Gaussian distribution and we simply publish its empirical average, i.e., $\mathcal{M}(X) = 1/n \cdot (\sum_{i=1}^n x_i)$. Even without any further obfuscation on $\mathcal{M}(X)$ like perturbation, intuitively, we see some potential hardness to reconstruct X from its average $\mathcal{M}(X)$. For an adversary who only knows the data distribution, it is impossible to perfectly reconstruct X deterministically due to ambiguity. But the data leakage of this trivial aggregation cannot be properly analyzed from a distinguishability angle, since one can always determine the true input from two datasets whose means are different. However, in practical processing even for a given data pool, many kinds of randomness

¹Here, we say X_0 and X'_0 are adjacent if they only differ in one datapoint, i.e., their Hamming distance is 1.

can enforce data to be processed to enjoy entropy, for example, subsampling or data augmentation [58], where the participants are randomized. Therefore, a more generic security/privacy definition, which enables *instance-based analysis* in terms of both data priors and specific inference objective, is attractive.

In this paper, we tackle the fundamental problem of quantifying the relationship among data entropy, disclosure, and reconstruction hardness. Besides a more meaningful interpretation via inference impossibility, to address the above two mentioned challenges in classic security/privacy regimes, we want a more generic definition, which enables (a) automatic security proof, and (b) instance-based worst-case analysis. To this end, a new privacy definition, termed Probably Approximately Correct (PAC) Privacy, is proposed, which enables a user to keep track of the data leakage of an arbitrary (black-box) processing mechanism with confidence, via simulation. At a high level, (δ, ρ, ϵ) PAC Privacy characterizes the following reconstruction hardness:

For some sensitive data X distributed in \mathcal{X} and a mechanism \mathcal{M} , given $\mathcal{M}(X)$, there does not exist an adversary (possibly computationally-unbounded) who can successfully return \tilde{X} such that $\rho(\tilde{X}, X) = 1$ under measure function ρ with probability at least $(1 - \delta)$.

We use the notion “PAC” as we borrow the idea of well-known PAC learning theory [63] and describe the adversary’s reconstruction/inference task as a generic learning problem. The difference compared to classic learning theory is that we study impossibility results. The $\rho(\cdot, \cdot)$ captures an arbitrary inference task of interest and $\rho(\tilde{X}, X) = 1$ if and only if adversary’s inference \tilde{X} satisfies the reconstruction criterion. For example, if we set out to prevent the adversary from recovering a single digital bit of data X , then we may select ρ such that $\rho(\tilde{X}, X)$ returns 1 only when \tilde{X} and X collide in at least one coordinate. More discussions can be found below Definition 1.

Our presented PAC Privacy framework has the following appealing properties and important implications:

1. **Wide applicability:** The PAC Privacy framework presented can theoretically be used to analyze any data processing. Via f -divergence, we present a simple quantification of the posterior advantage that the adversary gains by observing the disclosure $\mathcal{M}(X)$ for arbitrary inference problems. Thus, PAC Privacy can characterize very generic reconstruction/inference hardness under any selection of the measure function ρ , which captures classic attacks such as membership inference and identification. Through instance-based (Gaussian) perturbation, we show how to efficiently control PAC Privacy parameters. In contrast to the adversarial worst-case, the noise magnitude in PAC Privacy does not necessarily scale with dimensionality but can be $O(1)$ in high-dimensional release.
2. **Automatic Privacy analysis:** To our knowledge, PAC Privacy is the first non-trivial automatic security definition. Theoretically, no worst-case proof is needed but one can determine a perturbation mechanism (if necessary) fully automatically in polynomial time to ensure desired security parameters with arbitrarily high confidence. More importantly, this universal framework only requires sufficient Monte-Carlo simulations while the underlying data processing mechanism could be a black-box oracle – no algorithmic analysis is required. This is very different from classic definitions such as cryptography and DP, where the key challenge is to derive the algorithmic worst-case proof.
3. **Connections to the Input-Independent Worst Case and Generalization Error:** We also connect PAC Privacy, DP and generalization error. Given assumptions on priors, we show DP can produce PAC Privacy, and PAC Privacy can then provide a generalization error guarantee, while the reverse direction, in general, does not hold. However, in practical data applications, PAC Privacy could bring significant (empirical) perturbation improvement, and can characterize the optimal instance-based perturbation required. We show concrete examples in private release of local samples, mean estimation and deep learning models. Though the tools we apply to develop the foundation of PAC Privacy include f -divergence and mutual information, we stress that PAC Privacy is *not* defined using any existing information-theoretical metric or statistical divergence, but by the impossibility of inference described above. Therefore, PAC Privacy enjoys a very straightforward interpretation. More discussion can be found in Section 7.

Paper Organization: At a high level, the paper has two parts, i.e., the definition of P C Privacy and the control of security parameters via automatic analysis. In Section 2, we formally present the definitions and interpretations of P C Privacy. In Section 3, we show generic bounds of P C Privacy parameters using f -divergence and mutual information in various setups. With these bounds, we move on to present automatic algorithms to determine proper perturbation (if necessary) to produce desired P C Privacy. In Section 4.1, we present the methods for deterministic data processing, and in Section 4.2, we study the case of randomized mechanisms. In Section 4.3, we discuss how to approximate the optimal perturbation in general. In Section 5, we study the composition of P C Privacy in various setups. In Section 6, we proceed to study and compare the adversarial and instance-based worst case. Practical applications of P C Privacy are included. We summarize related works in Section 7, where we also discuss the relationship between P C Privacy, Differential Privacy and generalization error. Finally, we conclude in Section 8. The generalization to local P C Privacy via secure implementation of automatic privacy analysis in a decentralized setup can be found in the full version.

2 P C Privacy

We assume each private datapoint x is defined over some domain \mathcal{X} , and the dataset $X \in \mathcal{X}^*$. Our goal is to formalize the privacy leakage from a data processing mechanism \mathcal{M} , where $\mathcal{M} : \mathcal{X}^* \rightarrow \mathcal{Y} \subset \mathbb{R}^d$. Both \mathcal{X} and \mathcal{Y} are measurable spaces. The data distribution \mathbb{P} is defined over \mathcal{X}^* . \mathcal{M} can be either deterministic or randomized. Different analyses will be provided later for these two cases. The formal definition of P C Privacy is as follows.

Definition 1 ((δ, ρ) P C Privacy). *For a data processing mechanism \mathcal{M} , given some data distribution \mathbb{P} , and a measure function $\rho(\cdot, \cdot)$, we say \mathcal{M} satisfies (δ, ρ) -P C Privacy if the following experiment is impossible:*

user generates data X from distribution \mathbb{P} and sends $\mathcal{M}(X)$ to an adversary. The adversary who knows \mathbb{P} and \mathcal{M} is asked to return an estimation $\tilde{X} \in \mathcal{X}^$ on X such that with probability at least $(1 - \delta)$, $\rho(\tilde{X}, X) = 1$.*

In Definition 1, we adopt a way to describe the reconstruction hardness as a game between the user and the adversary. The adversary sets out to learn something from the disclosure $\mathcal{M}(X)$. The measure function (on the adversary’s learning performance) $\rho(\cdot, \cdot)$ can be arbitrarily selected according to the adversarial inference task of interest. For example, $\mathcal{X} \subset \mathbb{R}^{d'}$ for some dimension d' and a safe reconstruction error is set to be ϵ in the l_2 norm. Then, we can define $\rho(\tilde{X}, X) = 1$ if and only if $\|\tilde{X} - X\|_2 \leq \epsilon$ and Definition 1 suggests it is impossible for an adversary to recover private data X within ϵ error with probability at least $(1 - \delta)$. Another example, \mathcal{X}^* is some finite set and $X \in \mathcal{X}^*$ contains n elements. Similarly, we define $\rho(\tilde{X}, X) = 1$ iff $\#\{\tilde{X} \cap X\} > (n - \epsilon)$. Then, Definition 1 implies that there does not exist an adversary who can identify more than $(n - \epsilon)$ elements/participants of X from $\mathcal{M}(X)$ with probability at least $(1 - \delta)$. As a remark, a neat expression of ρ is not important or even necessary in P C Privacy and we simply use ρ to indicate whether the adversary’s reconstruction satisfies a certain criterion.

It is worthwhile to mention that in Definition 1 and results presented later, we do not put any specific restrictions on \mathbb{P} or the adversary’s strategy. \mathbb{P} could be a generic joint distribution of datapoints generated, though with additional assumptions we may obtain tighter bounds (such as Theorem 2). Before proceeding, we augment the definition of P C Privacy to incorporate (instance-based) additional *posterior advantage*, commonly adopted in classic privacy definitions. That is, for any specific priors given, we set out to quantify how much the disclosure of $\mathcal{M}(X)$ can help the adversary to implement successful inference. To measure such advantage, we first formally introduce the definition of f -divergence.

Definition 2 (f -Divergence). *Let $f : (0, +\infty) \rightarrow \mathbb{R}$ be a convex function with $f(1) = 0$. Let \mathbb{P} and \mathbb{Q} be two distributions on some measurable space, and the f -divergence \mathcal{D}_f between \mathbb{P} and \mathbb{Q} is defined as*

$$\mathcal{D}_f(\mathbb{P} \parallel \mathbb{Q}) := \mathbb{E}_{\mathbb{Q}} \left[f\left(\frac{d\mathbb{P}}{d\mathbb{Q}}\right) \right].$$

Here, we use dP (dQ) to represent the probability density (or mass) function of the continuous (or discrete) probability distribution P (Q).

In Definition 2, when we select $f(x) = x \log(x)$, it becomes the Kullback-Leibler (KL) divergence; when we select $f(x) = \frac{1}{2}|x - 1|$, it becomes the total variation, where $\mathcal{D}_{TV}(P\|Q) = \frac{1}{2} \cdot \int |dP - dQ|$. In the following, we define $(1 - \delta_o^\rho)$ as the optimal *a priori* chance that an adversary who only knows the data distribution can complete the objective reconstruction task *before* observing the release $\mathcal{M}(X)$.

Definition 3 (Optimal Prior Success Rate). *For given measure ρ and data distribution \mathcal{D} , we define $(1 - \delta_o^\rho)$ as*

$$1 - \delta_o^\rho := \sup_{\tilde{X} \in \mathcal{X}} \Pr_{X \sim \mathcal{D}} (\rho(\tilde{X}, X) = 1),$$

which represents the optimal a priori success rate only based on adversary's prior knowledge on \mathcal{X} to recover X such that the reconstruction satisfies the criterion defined by ρ .

Definition 4 ($(\Delta_f \delta, \rho, \delta_o^\rho)$ P-C Advantage Privacy). *A mechanism \mathcal{M} is termed $(\Delta_f \delta, \rho, \delta_o^\rho)$ P-C-advantage private if it is $(\delta, \rho, \delta_o^\rho)$ P-C private and*

$$\Delta_f \delta := \mathcal{D}_f(\mathbf{1}_\delta \| \mathbf{1}_{\delta_o^\rho}) = \delta_o^\rho f\left(\frac{\delta}{\delta_o^\rho}\right) + (1 - \delta_o^\rho) f\left(\frac{1 - \delta}{1 - \delta_o^\rho}\right).$$

Here, $\mathbf{1}_\delta$ and $\mathbf{1}_{\delta_o^\rho}$ represent two Bernoulli distributions of parameters δ and δ_o^ρ , respectively.

In Definition 4, $\Delta_f \delta$ captures the additional (posterior) advantage gained by the adversary after observing $\mathcal{M}(X)$. When we select \mathcal{D}_f to be total variation, $\Delta_{TV} \delta$ simply becomes the difference $\delta_o^\rho - \delta$, as commonly used in classic cryptographic analysis. Accordingly, once we have an upper bound on $\delta_o^\rho - \delta$, plugging in δ_o^ρ provides a lower bound on the posterior failure rate δ .

The seemingly artificial aspect of Definition 4 in using f -divergence as the measure is motivated by Theorem 1, where we give generic upper bounds on $\Delta_f \delta$. In particular, when \mathcal{D}_f is selected to be KL-divergence, we can simply bound $\Delta_{KL} \delta$ via mutual information. In addition, it is not hard to see that when \mathcal{M} satisfies perfect secrecy, $\Delta_f \delta$ is 0 for any f . Before the end of this section, we have the following important remarks on P-C Privacy.

Reconstruction Metric Selection: Definitions 1 and 4 provide a unified framework to capture generic inference hardness. Membership and identification attacks are now special cases by selecting a specific data generation and a specific measure ρ .

Definition 5 ($(\delta, \rho, U, \delta_o^\rho)$ P-C Membership Privacy). *For a data processing mechanism \mathcal{M} , given some measure ρ and a data set $U = (u_1, u_2, \dots, u_N)$, we say \mathcal{M} satisfies $(\delta, \rho, U, \delta_o^\rho)$ -P-C Membership Privacy if the following experiment is impossible:*

user applies certain sampling (described by ρ) on U to generate a dataset $X \sim \mathcal{D}$, and sends $\mathcal{M}(X)$ to an adversary. The adversary who knows U , and \mathcal{M} is asked to return an N -dimensional binary vector $\tilde{\mathbf{1}}_U = (\tilde{\mathbf{1}}_{u_1}, \dots, \tilde{\mathbf{1}}_{u_N})$ to predict the participation of each u_i , denote by $\mathbf{1}_U = (\mathbf{1}_{u_1}, \dots, \mathbf{1}_{u_N})$, where $\mathbf{1}_{u_i}$ is an indicator to represent the participation of u_i , such that with probability at least $(1 - \delta)$, $\rho(\tilde{\mathbf{1}}_U, \mathbf{1}_U) = 1$.

One can imagine the distribution \mathcal{D} could be Poisson (i.i.d.) sampling or random sampling with(out) replacement from U . From a DP angle, such privacy amplification from sampling has been studied in [4], [65]. P-C Membership Privacy is a very important variant of P-C Privacy in practical applications, which can give a formal security proof against the well-known and fundamental *membership inference attack* problem [57].

We may consider the classic statistical query model, where there is some sensitive data set U , for example, medical records of N patients, held by a server. The server is asked to answer some query, expressed as disclosing the results of \mathcal{M} . The server randomly samples a subset X of elements from U ; for example, each datapoint in U will be independently selected with probability $1/2$, to participate in the computation of $\mathcal{M}(X)$. If we select $\rho(\tilde{\mathbf{1}}_U, \mathbf{1}_U) = 1$ only when $\tilde{\mathbf{1}}_U$ and $\mathbf{1}_U$ collide in at least one entry, then $(\delta, \rho, U, \delta_o^\rho)$ -P-C

Membership Privacy suggests that *even if the adversary already knows the whole data pool U , he still cannot successfully infer any participants in the processing mechanism \mathcal{M}* . In other words, P-C Membership Privacy quantifies the *risk* incurred by any group of persons who allow their data to be processed through \mathcal{M} .

If, on the other hand, our privacy concern pertains to individuals, we simply consider the reconstruction hardness of each single datapoint in the set X , and similarly define P-C Individual Privacy as follows.

Definition 6 ((δ, ρ, n, \cdot) P-C Individual Privacy). *For a data processing mechanism \mathcal{M} , given some data distribution \mathcal{D} , the total number n of datapoints, a measure $\rho(\cdot, \cdot)$, we say \mathcal{M} satisfies (δ, ρ, n, \cdot) -P-C Individual Privacy if for any $i \in [1 : n]$ the following experiment is impossible:*

user generates data $X = (x_1, x_2, \dots, x_n)$ from distribution \mathcal{D} and sends $\mathcal{M}(X)$ to an adversary. The adversary who knows \mathcal{D} and \mathcal{M} is asked to return an estimation \tilde{x}_i on x_i such that with probability at least $(1 - \delta)$, $\rho(\tilde{x}_i, x_i) = 1$.

If the input domain \mathcal{D} is some finite set and we want to measure how hard it is for an adversary to identify the true selection of x_i , we may set ρ to be $\rho(\tilde{x}_i, x_i) = 1$ iff $\tilde{x}_i = x_i$. Then, the reconstruction becomes the *identification problem*. We have to stress that a proper selection of measure ρ is important to make P-C Privacy meaningful in different scenarios. Different error metrics will capture different perspectives of data privacy. For biological or financial data, such as ages and incomes, l_p norm, such as l_1 and l_2 , is usually meaningful. But for sensitive image or natural language data, l_p norm may not be able to fully characterize the leakage of contextual information, and membership and identification hardness are better options. P-C Privacy provides a unified framework for arbitrary inference tasks and privacy concerns.

Lack of Knowledge: In all the above definitions, we assume the adversary has full knowledge of the *a priori* data generation distribution \mathcal{D} , or \mathcal{D} is public. However, our current P-C Privacy framework cannot properly quantify the inference hardness due to the *lack/absence* of prior knowledge. For example, suppose the data X is generated from some distribution \mathcal{D} , say a random picture selected from a pool of cat images. For an adversary who does not have full knowledge but only knows that the users' pictures could be selected from an image pool of cats or dogs, the inference would be harder compared to the full knowledge setup. Unfortunately, our techniques used in P-C Privacy rely on the entropy of data, which does not change conditional on such lack of knowledge. However, in the other direction, P-C Privacy *does* quantify the inference hardness in a more generic scenario when the adversary has more knowledge rather than simply knowing \mathcal{D} . We refer the reader to Remark 1 in Section 3.

lthough our current framework cannot exactly quantify the inference hardness amplification from the lack of knowledge, P-C Privacy enjoys strong and realistic privacy interpretation from this perspective. If a certain inference problem is already difficult for an adversary who has full knowledge, then it is even more difficult for an adversary who lacks knowledge. Suppose a user holds a set U of cat images and randomly samples, say, 1000 samples from them, to produce X and discloses $\mathcal{M}(X)$. If $\mathcal{M}(X)$ satisfies P-C Privacy such that even an adversary who knows U and the data generation \mathcal{D} cannot identify or recover the 1000 participants, then it is even more difficult for an adversary to identify or recover the 1000 selected samples from the universe of cat images.

3 Foundation of P-C Privacy

In this section, we present our generic results to enable concrete P-C Privacy analysis. Before we state the theorems, we need the following additional notations to express the results.

Definition 7 (Entropy). *For a random variable x defined on a discrete set \mathcal{X} , the (Shannon) entropy $\mathcal{H}(x)$ of x is defined as $\mathcal{H}(x) = \sum_{x_0 \in \mathcal{X}} -\Pr(x = x_0) \log \Pr(x = x_0)$. If x is in a continuous distribution on some domain \mathcal{X} , then the (differential) entropy $h(x)$ of x is defined as $h(x) = \int_{x_0 \in \mathcal{X}} -\mathbb{P}(x = x_0) \log \mathbb{P}(x = x_0) dx_0$, where \mathbb{P} represents the probability density function.*

With a slight abuse of notation, we simply apply $\mathcal{H}(x)$ to denote the (Shannon/differential) entropy of x depending on its distribution for simplicity. Accordingly, we can define the conditional entropy $\mathcal{H}(x|w)$ for

two random variables x and w in some joint distribution, where $\mathcal{H}(x|w) = \sum_{w_0 \in \mathcal{W}} \mathcal{H}(x|w = w_0) \Pr(w = w_0)$. When w is in some continuous distribution, we can similarly define the conditional entropy in an integral form as differential entropy. In the following, we will use \mathbb{P}_x to denote the distribution of a random variable x .

Definition 8 (Mutual Information). *For two random variables x and w in some joint distribution, the mutual information $\text{MI}(x; w)$ is defined as*

$$\text{MI}(x; w) := \mathcal{H}(x) - \mathcal{H}(x|w) = \mathcal{D}_{KL}(\mathbb{P}_{x,w} \| \mathbb{P}_x \otimes \mathbb{P}_w),$$

i.e., the KL-divergence between the joint distribution of (x, w) and the product of the marginal distributions of x and w , respectively.

Mutual information measures the dependence between x and w . $\text{MI}(x; w) = \text{MI}(w; x)$ is always non-negative and equals 0 if x is independent of w .

Theorem 1. *For any selected f -divergence \mathcal{D}_f , a mechanism $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies $(\Delta_f \delta, \rho)$ P C advantage Privacy if*

$$\Delta_f \delta = \mathcal{D}_f(\mathbf{1}_\delta \| \mathbf{1}_{\delta_o^p}) \leq \int_{\mathbb{P}_W} \mathcal{D}_f(\mathbb{P}_{X, \mathcal{M}(X)} \| \mathbb{P}_X \otimes \mathbb{P}_W). \quad (1)$$

Here, $\mathbf{1}_\delta$ and $\mathbf{1}_{\delta_o^p}$ are two Bernoulli distributions of parameters δ and δ_o^p , respectively; $\mathbb{P}_{X, \mathcal{M}(X)}$ and \mathbb{P}_X are the joint distribution and the marginal distribution of $(X, \mathcal{M}(X))$ and X , respectively; and \mathbb{P}_W represents the distribution of an arbitrary random variable $W \in \mathcal{Y}$.

In particular, when we select \mathcal{D}_f to be the KL-divergence and $\mathbb{P}_W = \mathcal{P}_{\mathcal{M}(X)}$, \mathcal{M} satisfies $(\Delta_{KL} \delta, \rho)$ P C advantage Privacy where

$$\Delta_{KL} \delta := \mathcal{D}_{KL}(\mathbf{1}_\delta \| \mathbf{1}_{\delta_o^p}) \leq \text{MI}(X; \mathcal{M}(X)). \quad (2)$$

Here, $\mathcal{D}_{KL}(\mathbf{1}_\delta \| \mathbf{1}_{\delta_o^p}) = \delta \log(\frac{\delta}{\delta_o^p}) + (1 - \delta) \log(\frac{1 - \delta}{1 - \delta_o^p})$.

Proof. See appendix . □

Remark 1. *In general, if the adversary has more prior knowledge (besides the data distribution) on X , denoted by a variable dv , we can obtain a corresponding P C Privacy bound using Theorem 1, where the only difference is that in (2), $\text{MI}(X; \mathcal{M}(X))$ and δ_o^p are changed to the conditional $\text{MI}(X; \mathcal{M}(X) | dv)$ and the corresponding failure probability of optimal a priori guessing on X conditional on dv , respectively.*

Theorem 1 is rooted in the well-known data processing inequality of f -divergence. Roughly speaking, this inequality says postprocessing cannot increase information. Given priors, the adversary's reconstruction or decision is determined by (a postprocessing of) his observation on $\mathcal{M}(X)$. Therefore, to information-theoretically bound the posterior advantage, it suffices to quantify the correlation between X and $\mathcal{M}(X)$. Theorem 1 serves as the foundation of P C Privacy analysis. From (1), we know that $\Delta_f \delta = \mathcal{D}_f(\mathbf{1}_\delta \| \mathbf{1}_{\delta_o^p})$, the change to the f -divergence between the optimal *a priori* and posterior reconstruction/inference performance, is bounded by the corresponding f -divergence between the distributions of $(X, \mathcal{M}(X))$ and (X, W) , where W is an arbitrary random variable independent of X . With the freedom of the selections of both f and W , if we apply KL-divergence and set W to be distributed the same as $\mathcal{M}(X)$, we have $\Delta_{KL} \delta \leq \text{MI}(X; \mathcal{M}(X))$, the mutual information between X and $\mathcal{M}(X)$.

It is noted that we are also free to select *arbitrary* ρ as the objective inference task and the success criterion in (1). Thus, another important implication from Theorem 1 is that once we have an upper bound on $\mathcal{D}_f(\mathbb{P}_{X, \mathcal{M}(X)} \| \mathbb{P}_X \otimes \mathbb{P}_{\mathcal{M}(X)})$, we indeed also control *the posterior advantage of all possible adversaries' inference tasks*. This is also the foundation of the automatic proof of P C Privacy, where the *entirety of data generation and processing \mathcal{M} could be a black-box oracle*; as shown in Section 4, we are able to control such f -divergence without any algorithmic analysis of \mathcal{M} . We can compute δ_o^p based on ρ and for any given inference task. Given δ_o^p , once we have the upper bound on $\Delta_f \delta$, we can consequently lower bound the adversary's failure probability δ for inference since \mathcal{D}_f is a deterministic function.

In the following, we will mainly focus on $\Delta_{KL}\delta$ to control P C Privacy since mutual information is a well-studied quantity in information theory and has many nice properties to simplify our analysis. Intuitively, $\text{MI}(X; \mathcal{M}(X))$ quantifies the dependence between X and $\mathcal{M}(X)$, and shows how much such instance-based information (in nats/bits by taking \log_e/\log_2 in KL-divergence) is provided by $\mathcal{M}(X)$ to support the adversary's inference. Especially, when we select \mathcal{D}_f to be total variation and still $\mathbb{P}_W = \mathbb{P}_{\mathcal{M}(X)}$,

$$\Delta_{TV}\delta = \delta_o^\rho - \delta \leq \mathcal{D}_{TV}(\mathbb{P}_{X, \mathcal{M}(X)} \| \mathbb{P}_X \otimes \mathbb{P}_{\mathcal{M}(X)}).$$

By Pinsker's inequality, which states that for two distributions \mathbb{P} and \mathbb{Q} on the same measurable space, $\Delta_{TV}(\mathbb{P} \| \mathbb{Q}) \leq \sqrt{\frac{1}{2} \cdot \Delta_{KL}(\mathbb{P} \| \mathbb{Q})}$, we have the following corollary to give a quantification on the more intuitive and classic posterior advantage.

Corollary 1. *With the same setup as Theorem 1, the posterior advantage $(\delta_o^\rho - \delta)$ for arbitrary ρ is upper bounded as*

$$\Delta_{TV}\delta = \delta_o^\rho - \delta \leq \sqrt{\frac{1}{2} \cdot \text{MI}(X; \mathcal{M}(X))}.$$

On the other hand, we can also apply Theorem 1 to interpret the classic worst-case security definitions. To derive the *data-independent* maximal leakage when exposing $\mathcal{M}(X)$, we may take a supremum on the right side of (2) and study $\sup_{X \sim \bar{\cdot}} \text{MI}(X; \mathcal{M}(X))$ instead. Thus, the adversarial worst-case analysis can be used to produce a (loose) upper bound of P C Privacy. For example, if we know \mathcal{M} is ϵ -DP, then $\sup_{X \sim \bar{\cdot}} \text{MI}(X; \mathcal{M}(X)) \leq 0.5\epsilon^2 n^2$ for an n -datapoint set X (Theorem 1.10 and Proposition 1.4 in [14]).

With Theorem 1, our goal for the remainder of this paper is clear, where we set out to determine or bound $\text{MI}(X; \mathcal{M}(X))$ with high confidence. We first move on to analyze an important special case where x_i is i.i.d. generated and $\mathcal{M}(\cdot)$ is a symmetric mechanism. Here, we call $\mathcal{M}(x_1, x_2, \dots, x_n)$ symmetric if for an arbitrary perturbation π on $[1 : n]$, the distribution of $\mathcal{M}(x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(n)})$ is identical, where each element x_i is treated equally. Many data processing tasks satisfy the symmetric property, for example mean/median estimation or empirical risk minimization (ERM) in machine learning, where the order of samples does not change the result. The main motivation to study this special case is that we can show tighter and simpler bounds to control the security parameter δ in many applications, especially when the mechanism \mathcal{M} is deterministic with respect to X . This will become clear in Section 4. For notional clarity, we will use $X \sim \bar{\cdot}$ to denote that x_i is i.i.d. generated from some distribution $\bar{\cdot}$ to differentiate from the generic joint distribution case.

Theorem 2. *mechanism \mathcal{M} satisfies $(\delta, \rho, n, \bar{\cdot})$ P C Individual Privacy if there exist δ_i for $i = 1, 2, \dots, n$, satisfying $\delta_i \geq \delta$ and*

$$\mathcal{D}_{KL}(\mathbf{1}_{\delta_i} \| \mathbf{1}_{\delta_o^\rho, i}) \leq \text{MI}(x_i; \mathcal{M}(X)). \quad (3)$$

Here, $1 - \delta_o^\rho = \sup_{\tilde{x}_i \in \bar{\cdot}} \Pr(\rho(\tilde{x}_i, x_i) = 1)$ represents the optimal prior success rate of recovering x_i .

In particular, when $X \sim \bar{\cdot}$ and \mathcal{M} is a symmetric mechanism, then

$$\delta \geq \sum_{j=1}^n \frac{\tilde{\delta}^j}{n},$$

for any selection of $\tilde{\delta}^j$ satisfying $\mathcal{D}_{KL}(\mathbf{1}_{\tilde{\delta}^j} \| \mathbf{1}_{\tilde{\delta}^j, \rho}) \leq \text{MI}(X; \mathcal{M}(X))$, where

$$1 - \tilde{\delta}_o^{j, \rho} = \sum_{l=j}^n \binom{n}{l} (\tilde{\delta}_o^\rho)^{-l} \cdot (1 - \tilde{\delta}_o^\rho)^l, \quad (4)$$

and $1 - \tilde{\delta}_o^\rho = \sup_{\tilde{x} \in \bar{D}} \Pr_{x \sim \bar{D}}(\rho(\tilde{x}, x) = 1)$.

Proof. See appendix B. □

In Theorem 2, we point out that if our objective is individual privacy, we may get tighter bounds by applying $\text{MI}(x_i; \mathcal{M}(X))$ rather than $\text{MI}(X; \mathcal{M}(X))$. By the chain rule of mutual information [24], we know that $\text{MI}(X; \mathcal{M}(X))$ is an upper bound for any $\text{MI}(x_i; \mathcal{M}(X))$ since for any i ,

$$\text{MI}(X; \mathcal{M}(X)) = \text{MI}(x_i; \mathcal{M}(X)) + \text{MI}(X \setminus \{x_i\}; \mathcal{M}(X)|x_i) \geq \text{MI}(x_i; \mathcal{M}(X)).$$

However, to analyze $\text{MI}(x_i; \mathcal{M}(X))$, generally we need to take the other elements $X \setminus \{x_i\}$ as part of randomness seeds and enforce \mathcal{M} to be randomized, which makes it more difficult to give tight bounds. Therefore, in the second part of Theorem 2, in the case of i.i.d. data with symmetric processing, we give a more fine-grained analysis but stick to $\text{MI}(X; \mathcal{M}(X))$. The idea behind this is that we consider the chance, denoted by $(1 - \delta^j)$, that the adversary can successfully recover at least j many elements out of an n -sample set X , and use the symmetric property to bound the success rate of each individual x_i .

Example 1 (Applications of Theorems 1 and 2). *We consider a set X of n datapoints where each datapoint x_i is i.i.d. uniformly selected from a finite set of size N and \mathcal{M} is some symmetric mechanism. Suppose $N = 100$ and for any given i , we know the optimal prior success rate $(1 - \delta_0^p)$ to identify the true selection of x_i is 0.01. When $\text{MI}(X; \mathcal{M}(X)) = 1$, via Theorem 1, by numerically solving $\mathcal{D}_{KL}(\mathbf{1}_{\delta_i} \| \mathbf{1}_{0.99}) \leq 1$, we have a global bound $(1 - \delta) \leq 0.36$ for any n .*

Applying Theorem 2, a larger n will produce a tighter bound of $(1 - \delta)$. Given $\text{MI}(X; \mathcal{M}(X)) = 1$, we numerically solve $\mathcal{D}_{KL}(\mathbf{1}_{\tilde{\delta}_j} \| \mathbf{1}_{\tilde{\delta}_j, \rho}) \leq 1$ to obtain the lower bound of $\tilde{\delta}_j$, the failure probability that an adversary can recover at least j elements for $j = 1, 2, \dots, n$, and consequently obtain the lower bound of individual inference failure probability δ . When $n = 10$, we have $(1 - \delta) \leq 0.17$, and when $n = 50$, we have $(1 - \delta) \leq 0.06$.

So far, we have presented the generic bounds of P C Privacy, and to ensure desired security parameters, the problem is reduced to controlling the mutual information $\text{MI}(X; \mathcal{M}(X))$ (or $\text{MI}(x_i; \mathcal{M}(X))$) with high confidence.

4 Automatic Control of Mutual Information

In this section, we present our main results on automatic privacy measurement and control. At a high level, we want an automatic protocol which takes security parameters as input and returns a privatization scheme on \mathcal{M} to ensure required privacy guarantees with high confidence. In particular, we want the least assumptions on mechanism \mathcal{M} , which ideally could be a black-box oracle and no specific algorithmic analysis is needed to produce the security proof.

One natural idea to achieve information leakage control is perturbation: when $\text{MI}(X; \mathcal{M}(X))$ is not small enough to produce satisfied P C Privacy, we may add additional noise, say Gaussian, to produce smaller $\text{MI}(X; \mathcal{M}(X) + \cdot)$. In P C Privacy, the role of noise is not simply perturbation but to also enforce the output of a black-box oracle into a class of parameterized distributions. As shown below for either deterministic or randomized algorithms, with Gaussian noise, the analysis is reduced to the study of divergences of Gaussian mixture models. The key question left is how to automatically determine the parameters of the noise. We give solutions in this section. In Section 4.1, we focus on the deterministic $\mathcal{M}(\cdot)$ w.r.t. X , while, in Section 4.2, we analyze generic randomized $\mathcal{M}(\cdot)$ w.r.t. either X or a single datapoint x_i . Finally, in Section 4.3, we discuss how to approximate the optimal perturbation.

4.1 Deterministic Mechanism

The estimation of mutual information is a long-standing open problem, though it enjoys a very simple expression via (conditional) entropy or KL-divergence. Practical mutual information estimations have been studied in many empirical works using deep learning [10], [21] or kernel methods [50] with different motivations such as information bottleneck [62] and causality [15]. However, one key obstacle to use those heuristic estimations is that we cannot derive a high confidence bound and design necessary perturbation for rigorous security guarantees.

In this section, we assume that \mathcal{M} is some deterministic function w.r.t. X , for example, \mathcal{M} returns the average/median of X or the global optimum of some loss function determined by X . Beyond the deterministic assumption, the only other restriction is that the output $\mathcal{M}(X)$ is bounded. Without loss of generality, we assume the l_2 norm $\|\mathcal{M}(x)\|_2 \leq r$. Since the distribution of $\mathcal{M}(X)$ could be arbitrary over \mathbb{R}^d , the d -dimensional l_2 ball of radius r centered at zero, in general, there do not exist non-trivial upper bounds of $\text{MI}(X; \mathcal{M}(X))$. Indeed, possibly counter-intuitively, we cannot even say that the aggregation of a larger sample set would incur less information leakage. For example, imagine a finite set \mathcal{U} where the empirical means of its subsets are distinct. X represents a subset of \mathcal{U} randomly selected and $\mathcal{M}(X)$ outputs the average of X . Since there is a bijection between X and $\mathcal{M}(X)$, $\text{MI}(X; \mathcal{M}(X)) = \mathcal{H}(X) = O(\log 2^{|\mathcal{U}|})$, increasing with $|\mathcal{U}|$.

Fortunately, adding *continuous* noise, especially Gaussian noise, enables us to derive generic control of $\text{MI}(X; \mathcal{M}(X) + \text{noise})$. Continuous noise enforces the distribution of output $\mathcal{M}(X) + \text{noise}$ to be parameterized and continuous regardless of \mathcal{M} and X . Below, we show how to apply covariance estimation to fully automatically determine a noise to control $\text{MI}(X; \mathcal{M}(X) + \text{noise})$ with confidence for any black-box oracle \mathcal{M} . In the following, we simply set noise to be some multivariate Gaussian in a form $\mathcal{N}(0, \Sigma)$, whose covariance is Σ . Similarly, $\Sigma_{\mathcal{M}(X)}$ represents the covariance matrix of $\mathcal{M}(X)$.

Theorem 3. *For an arbitrary deterministic mechanism \mathcal{M} and a Gaussian noise of the form $\text{noise} \sim \mathcal{N}(0, \Sigma)$,*

$$\text{MI}(X; \mathcal{M}(X) + \text{noise}) \leq \frac{1}{2} \cdot \log \det(\mathbf{I}_d + \Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1}). \quad (5)$$

In particular, let the eigenvalues of $\Sigma_{\mathcal{M}(X)}$ be $(\lambda_1, \dots, \lambda_d)$, then there exists some Σ such that $\mathbb{E}[\|\text{noise}\|_2^2] = (\sum_{j=1}^d \sqrt{\lambda_j})^2$, and $\text{MI}(X; \mathcal{M}(X) + \text{noise}) \leq 1/2$.

Proof. See appendix C. □

From Theorem 3, we have a simple upper bound on the mutual information after perturbation which only requires the knowledge of the covariance of $\mathcal{M}(X)$ when \mathcal{M} is deterministic. Another important and appealing property is that the noise calibrated to ensure the target mutual information bound is *not* explicitly dependent on the output dimensionality d but instead on the square root sum of eigenvalues of $\Sigma_{\mathcal{M}(X)}$. When $\mathcal{M}(X)$ is largely distributed in a p -rank subspace of \mathbb{R}^d , Theorem 3 suggests that a noise of scale $O(\sqrt{p})$ is needed. This is different from DP where the expected l_2 norm of noise is in a scale of $\Theta(\sqrt{d})$ given constant l_2 -norm sensitivity; we will present a comprehensive comparison in Section 6.

Based on Theorem 3, we now proceed to present an automatic protocol (Algorithm 1) to determine Σ and produce an upper bound such that $\text{MI}(X; \mathcal{M}(X) + \text{noise}) \leq (v + \beta)$ with high confidence, where v and β are positive parameters selected as explained below. After sufficiently many simulations, the following theorem ensures that we can obtain accurate enough estimation with arbitrarily high probability. The characterization on the relationship among security parameter c , simulation complexity m and confidence parameter γ is given in (6).

Theorem 4. *Assume that $\mathcal{M}(X) \in \mathbb{R}^d$ and $\|\mathcal{M}(X)\|_2 \leq r$ for some constant r uniformly for any X , and apply Algorithm 1 to obtain the Gaussian noise covariance Σ for a specified mutual information bound $v + \beta$. v and β can be chosen independently, and c is a security parameter. Then, there exists a fixed and universal constant κ such that one can ensure $\text{MI}(X; \mathcal{M}(X) + \text{noise}) \leq v + \beta$ with confidence at least $(1 - \gamma)$ once the selections of c , m and γ satisfy,*

$$c \geq \kappa r \left(\max\left\{ \sqrt{\frac{d + \log(4/\gamma)}{m}}, \frac{d + \log(4/\gamma)}{m} \right\} + \sqrt{\frac{d \log(4/\gamma)}{m}} \right). \quad (6)$$

Proof. See appendix D. □

Algorithm 1 $(1 - \gamma)$ -Confidence Noise Determination of Deterministic Mechanism \mathcal{M}

- 1: **Input:** deterministic mechanism $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}^d$, data distribution \mathcal{D} , sampling complexity m , security parameter c , and mutual information quantities v and β .
- 2: **for** $k = 1, 2, \dots, m$ **do**
- 3: Independently generate data $X^{(k)}$ from distribution \mathcal{D} .
- 4: Record $y^{(k)} = \mathcal{M}(X^{(k)})$.
- 5: **end for**
- 6: Calculate empirical mean $\hat{\mu} = \sum_{k=1}^m y^{(k)}/m$ and the empirical covariance estimation $\hat{\Sigma} = \sum_{k=1}^m (y^{(k)} - \hat{\mu})(y^{(k)} - \hat{\mu})^T/m$.
- 7: Apply singular value decomposition (SVD) on $\hat{\Sigma}$ and obtain the decomposition as $\hat{\Sigma} = \hat{U}\hat{\Lambda}\hat{U}^T$, where $\hat{\Lambda}$ is the diagonal matrix of eigenvalues $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$.
- 8: Determine the maximal index $j_0 = \arg \max_j \hat{\lambda}_j$, for those $\hat{\lambda}_j > c$.
- 9: **if** $\min_{1 \leq j \leq j_0, 1 \leq l \leq d} |\hat{\lambda}_j - \hat{\lambda}_l| > r\sqrt{dc} + 2c$ **then**
- 10: **for** $j = 1, 2, \dots, d$ **do**
- 11: Determine the j -th element of a diagonal matrix Λ_B as

$$\lambda_{B,j} = \frac{2v}{\sqrt{\hat{\lambda}_j + 10cv/\beta} \cdot (\sum_{j=1}^d \sqrt{\hat{\lambda}_j + 10cv/\beta})}.$$

- 12: **end for**
 - 13: Determine the Gaussian noise covariance as $\Sigma = \hat{U}\Lambda^{-1}\hat{U}^T$.
 - 14: **else**
 - 15: Determine the Gaussian noise covariance as $\Sigma = (\sum_{j=1}^d \hat{\lambda}_j + dc)/(2v) \cdot \mathbf{I}_d$.
 - 16: **end if**
 - 17: **Output:** Gaussian covariance matrix Σ .
-

At a high level, Algorithm 1 shows a way to estimate both the eigenvectors and the spectrum (eigenvalues) of the covariance matrix $\Sigma_{\mathcal{M}(X)}$ (lines 6-7), which afterwards determines the noise parameter Σ (lines 8-16) to ensure desired mutual information upper bound $(v + \beta)$. In the first part, we simply apply the eigenvalue and eigenvectors of the empirical covariance obtained by the results of m simulations as the estimations of those of the true covariance matrix $\Sigma_{\mathcal{M}(X)}$. In the second part to determine the noise Σ , depending on the eigenvalue gap (line 9), we present two different upper bounds of Σ as explained later.

Theorem 4 presents a generic tradeoff between noise, sampling complexity and confidence. The role of c in Algorithm 1 is a safe parameter that provides a lower bound on noise. From (6), in general, we may use a larger noise (larger c for looser noise estimation) to produce higher confidence (smaller γ) given limited computation power m for sampling and simulation; on the other hand, both c and β can also be arbitrarily small provided sufficiently many simulations m to produce accurate enough estimation.

The main reason for the need to estimate the eigenvectors of $\Sigma_{\mathcal{M}(X)}$ in Algorithm 1 is because we need the instance-based noise to fit the geometry of the eigenspace of $\Sigma_{\mathcal{M}(X)}$. This is different from the adversarial worst-case, such as DP, which only focuses on the magnitude of sensitivity and adds isometric noise; *in P-C Privacy, we add anisotropic noise, as much as needed in each direction*. The key intuition behind why we can ensure arbitrarily high confidence is because the output domain of $\mathcal{M}(X)$ is bounded and thus $\mathcal{M}(X)$ is in a certain sub-Gaussian distribution with concentration guarantees captured by bounded r . Therefore, provided sufficiently many simulations, we can approximate the true covariance matrix $\Sigma_{\mathcal{M}(X)}$ with arbitrary precision with a mild assumption on the eigenvalue gap.

It is noted that depending on the *eigenvalue gap* defined in line 9 of Algorithm 1, there are two cases of perturbation. When the *significant eigenvalues* of $\Sigma_{\mathcal{M}(X)}$ are distinct, then for arbitrary high confidence $(1 - \gamma)$, the determined noise covariance Σ matches that described in Theorem 3 given large enough m , which allows selection of small enough c . However, though the spectrum of $\Sigma_{\mathcal{M}(X)}$ can be estimated with arbitrary high precision, we cannot ensure a satisfied eigenvector space approximation in our current

analysis framework without the assumption on the eigenvalue gap. To this end, in the special case without the eigenvalue gap guarantee, we have to switch to a more pessimistic perturbation (described in line 15 in Algorithm 1), where the l_2 norm of the noise is then explicitly dimension dependent, $\Theta(\sqrt{d})$. In practice, such a case is very rare, and we leave the improvement of Algorithm 1 or a tighter analysis as an open question.

4.2 Randomized Mechanism

We proceed to consider the more complicated case where the mechanism \mathcal{M} is randomized, where to be specific we use the form $\mathcal{M}(X, \theta)$ and θ is the randomness seed. Given selection of θ , $\mathcal{M}(\cdot, \theta)$ becomes deterministic. As mentioned earlier, to analyze $\text{MI}(x_i; \mathcal{M}(X))$, we can simply combine $X \setminus \{x_i\}$, the other elements in X without x_i , with θ as randomness seeds. Therefore, without loss of generality, we only focus on the control of $\text{MI}(X; \mathcal{M}(X) + \text{noise}(\theta))$. The randomness seed θ of mechanism \mathcal{M} is randomly selected from a set $\Theta = \{\theta_1, \theta_2, \dots, \theta_{|\Theta|}\}$. As a simple generalization of Theorem 3, we have the following bound.

Corollary 2.

$$\text{MI}(X; \mathcal{M}(X, \theta) + \text{noise}(\theta)) \leq \mathbb{E}_\theta \frac{\log \det(\mathbf{I}_d + \Sigma_{\mathcal{M}(X, \theta)} \cdot \Sigma_{\text{noise}(\theta)}^{-1})}{2}. \quad (7)$$

Here, we assume a random-seed dependent Gaussian noise $\text{noise}(\theta)$, where when $\theta = \theta_0$, $\text{noise}(\theta) \sim \mathcal{N}(0, \Sigma_{\text{noise}(\theta_0)})$.

Corollary 2 utilizes the following properties of mutual information.

$$\text{MI}(\theta, \mathcal{M}(X, \theta); X) = \text{MI}(\theta; X) + \text{MI}(\mathcal{M}(X, \theta); X|\theta) = \text{MI}(\mathcal{M}(X, \theta); X|\theta);$$

$$\text{MI}(\theta, \mathcal{M}(X, \theta); X) = \text{MI}(\mathcal{M}(X, \theta); X) + \text{MI}(\theta; X|\mathcal{M}(X, \theta)) \geq \text{MI}(\mathcal{M}(X, \theta); X).$$

Therefore, $\text{MI}(X; \mathcal{M}(X, \theta) + \text{noise}(\theta))$ is upper bounded by $\text{MI}(X; \mathcal{M}(X, \theta) + \text{noise}(\theta)|\theta)$, where given θ , \mathcal{M} is deterministic. Thus, we can apply the results of Section 4.1 to control each subcase dependent on the selection of θ , and Corollary 2 follows straightforwardly. However, there are two potential limitations of Corollary 2. First, there could be exponentially many possible selections of θ and estimations on all $\Sigma_{\mathcal{M}(X, \theta)}$ could be intractable. The other issue is that (7) is tight only if for given θ , $\mathcal{M}(X, \theta)$ is concentrated. However, in some applications this may not be true. One example is non-convex optimization, e.g., we run SGD in deep learning. If the generalization error is small, the global distributions of $\mathcal{M}(X, \theta)$ and $\mathcal{M}(X', \theta)$ for any two representative datasets X and X' could be very close, but different X with different randomness may converge to very different local minima. Therefore, we propose an alternative approach as Algorithm 2 and its security proof is shown in Theorem 5.

In Algorithm 2, we need the following definition of *minimal-permutation distance*. The minimal-permutation distance between two k -block vectors \mathbf{a} and \mathbf{b} is defined as $\mathbf{d}_\pi(\mathbf{a}, \mathbf{b}) = \min_\pi \sum_j \|\mathbf{a}(j) - \mathbf{b}(\pi(j))\|^2/k$, where $\mathbf{a}(j)$ and $\mathbf{b}(j)$ represent their j -th blocks, respectively, and π is some permutation on the block index $[1 : k]$.

Theorem 5. Assume that $\mathcal{M}(X, \theta) \in \mathbb{R}^d$ and $\|\mathcal{M}(X, \theta)\|_2 \leq r$ for some constant r , and we apply Algorithm 2 to obtain the Gaussian noise covariance Σ_{noise} for a specified mutual information bound v . Then, if $\frac{1}{\tau}$ is an integer, where $|\Theta|$ is the total number of the randomness seeds of $\mathcal{M}(\cdot, \theta)$ and τ is the number of seeds selected in Algorithm 2, and c is a safe parameter in Algorithm 2, one can ensure that $\text{MI}(X; \mathcal{M}(X, \theta) + \text{noise}) \leq v$ with confidence at least $1 - \gamma$ once

$$m \geq \frac{8r^4 \log(1/\gamma)}{c^2}.$$

Proof. See Appendix E. □

In Algorithm 2, our key idea is to do sampling on the randomness seeds across Θ rather than accessing all possible selections of θ . Similar to Algorithm 1, we set a safe parameter c , which could be arbitrarily small, to lower bound the magnitude of noise added. Given $c > 0$, combining the assumption that $\|\mathcal{M}(X, \theta)\| \leq r$,

Algorithm 2 $(1 - \gamma)$ -Confidence Noise Determination of Randomized Mechanism $\mathcal{M}(\cdot, \theta)$

- 1: **Input:** randomized mechanism $\mathcal{M}(\cdot, \theta) : \mathcal{X} \rightarrow \mathcal{Y}^d$, $\theta \in \Theta$, data distribution \mathcal{D} , sampling complexity m , mutual information bound v , and security parameters c and τ .
 - 2: **for** $k = 1, 2, \dots, m$ **do**
 - 3: From distribution \mathcal{D} , independently sample $X_1^{(k)}$ and $X_2^{(k)}$.
 - 4: Randomly select τ seeds, denoted by $\mathcal{C}^{(k)} = \{\theta_1^{(k)}, \theta_2^{(k)}, \dots, \theta_\tau^{(k)}\}$, from Θ .
 - 5: Compute and record

$$\mathbf{y}^{(k,1)} = (\mathcal{M}(X_1^{(k)}, \theta_1^{(k)}), \dots, \mathcal{M}(X_1^{(k)}, \theta_\tau^{(k)})),$$

$$\mathbf{y}^{(k,2)} = (\mathcal{M}(X_2^{(k)}, \theta_1^{(k)}), \dots, \mathcal{M}(X_2^{(k)}, \theta_\tau^{(k)})).$$
 - 6: Compute minimal-permutation distance $\psi^{(k)} = \mathbf{d}_\pi(\mathbf{y}^{(k,1)}, \mathbf{y}^{(k,2)})$.
 - 7: **end for**
 - 8: Calculate empirical mean $\bar{\psi} = \sum_{k=1}^m \psi^{(k)} / m$.
 - 9: **Output:** Gaussian covariance $\Sigma = \frac{\bar{\psi} + c}{2v} \cdot \mathbf{I}_d$.
-

we have a very rough but uniform upper bound on $\text{MI}(X; \mathcal{M}(X, \theta))$ for any possible distributions of X and $\mathcal{M}(X, \theta)$. With this fact, in each iteration in Algorithm 2, we independently sample two data inputs $\bar{X}^k = \{X_1^{(k)}, X_2^{(k)}\}$, and $\mathcal{C}^{(k)} = \{\theta_1^{(k)}, \dots, \theta_\tau^{(k)}\}$ as a random τ -subset selection of Θ . It can be proved that

$$\text{MI}(X; \mathcal{M}(X, \theta) + \Sigma) \leq \mathbb{E}_{\bar{X}^{(k)}, \mathcal{C}^{(k)}, \theta \in \mathcal{C}^{(k)}} \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_1^{(k)}, \theta)_+} \parallel \mathbb{P}_{\mathcal{M}(X_2^{(k)}, \theta)_+}), \quad (8)$$

Thus, it suffices to consider a conditional local distribution when θ is restricted to the τ -subset selected. τ does not influence the lower bound of sampling complexity m required. However, a larger τ , which also requires higher computational complexity to evaluate more $\mathcal{M}(X, \theta)$, will produce a tighter upper bound. The reason behind this is that the conditional sub-sampled distribution with $\theta \in \mathcal{C}^{(k)}$ gets closer to the global distribution with $\theta \in \Theta$ as τ increases. We assume $\lfloor m/\tau \rfloor$ is an integer, i.e., we can split \mathcal{D} into multiple τ -subsets, to enable the application of minimal-perturbation distance \mathbf{d}_π (Lemma 6 in Appendix E in the full version) to give a tighter upper bound of the objective KL-divergence to estimate.

However, for any proper selection of τ , the upper bound (8) always holds. Thus, once we have a high-confidence estimation of the expectation on the right side of (8), we also obtain a high-confidence upper bound on the objective $\text{MI}(X; \mathcal{M}(X, \theta) + \Sigma)$. Further, recall that KL-divergence between $\mathbb{P}_{\mathcal{M}(X_1^{(k)}, \theta)_+}$ and $\mathbb{P}_{\mathcal{M}(X_2^{(k)}, \theta)_+}$ for $\theta \in \mathcal{C}^{(k)}$ is always bounded given $c > 0$ and bounded r . Therefore, through i.i.d. sampling, we can simply use the empirical mean to approximate the expectation with a high probability bound.

Before the end of this section, we have some comments on the computation complexity of Algorithms 1 and 2. As mentioned before, the deterministic \mathcal{M} is a special case of the randomized processing and Algorithm 2 also applies to the deterministic \mathcal{M} where we may view \mathcal{D} as only containing a single randomness seed. As a comparison, Algorithm 1 requires $O(\frac{d+1}{c^2} \frac{g(1/\gamma)}{\gamma})$ simulation trials, while Algorithm 2 only requires $O(\frac{1}{c^2} \frac{g(1/\gamma)}{\gamma})$ simulation trials. Here, we assume the parameters r and β are all constants. The reason behind the dependence on the dimensionality d in Algorithm 1 is because to characterize the optimal Gaussian noise, we need to estimate the power of output distribution $\mathcal{M}(X)$ in each direction in \mathbb{R}^d , captured by the spectrum of $\Sigma_{\mathcal{M}(X)}$. Though the estimation on l_2 -norm based (minimal-permutation) distance is more efficient, as a tradeoff, the noise upper bound produced by Algorithm 2 is looser, $\Theta(\sqrt{d})$, with a strict dependence on d . We will give a concrete example for this gap between the noise bound derived from Algorithm 1 and Algorithm 2 in Example 3. We leave a more generic tradeoff between the simulation efficiency and noise control to future work.

Algorithm 3 $(1 - \gamma)$ -Confidence Verification of Perturbation Proposal

- 1: **Input:** randomized mechanism $\mathcal{M}(\cdot, \theta) : \mathcal{X} \rightarrow \mathcal{Y}$, $\theta \in \Theta$, $\mathcal{X} = \{X_1, X_2, \dots, X_N\}$, a perturbation proposal $(X, \theta) \sim Q(X, \theta)$, sampling complexity m , security parameters c , τ_1, τ_2 and τ_3 .
- 2: **for** $k = 1, 2, \dots, m$ **do**
- 3: Randomly sample τ_1 selections of X from \mathcal{X} , denoted by $\bar{X}^{(k,1)} = \{X_1^{(k,1)}, \dots, X_{\tau_1}^{(k,1)}\}$, and independently sample another τ_2 selections of X , denoted by $\bar{X}^{(k,2)} = \{X_1^{(k,2)}, \dots, X_{\tau_2}^{(k,2)}\}$.
- 4: Randomly sample τ_3 selections of θ from Θ , denoted by $\mathcal{C}^{(k)} = \{\theta_1^{(k)}, \theta_2^{(k)}, \dots, \theta_{\tau_3}^{(k)}\}$.
- 5: Define two sets of distributions for independent $c \sim \mathcal{N}(0, c \cdot \mathbf{I}_d)$,

$$P_i^{(k,1)} = \sum_{j=1}^{\tau_3} \frac{1}{\tau_3} \cdot (P_{\mathcal{M}(X_i^{(k,1)}, \theta_j^{(k)})+c} \parallel P_{(X_i^{(k,1)}, \theta_j^{(k)})+c}), i = 1, 2, \dots, \tau_1,$$

$$P_l^{(k,2)} = \sum_{j=1}^{\tau_3} \frac{1}{\tau_3} \cdot (P_{\mathcal{M}(X_l^{(k,2)}, \theta_j^{(k)})+c} \parallel P_{(X_l^{(k,2)}, \theta_j^{(k)})+c}), l = 1, 2, \dots, \tau_2.$$

- 6: Compute the average of KL-divergences,

$$\psi^{(k)} = \sum_{i=1}^{\tau_1} \left(\frac{1}{\tau_1} \cdot \mathcal{D}_{KL}(P_i^{(k,1)} \parallel \sum_{l=1}^{\tau_2} \frac{1}{\tau_2} \cdot P_l^{(k,2)}) \right). \quad (10)$$

- 7: **end for**

- 8: Calculate empirical mean $\bar{\psi} = \sum_{k=1}^m \psi^{(k)} / m$.

- 9: **Output:** $\bar{\psi}$.
-

4.3 Towards Optimal Perturbation

The algorithms presented so far all produce conservative bounds of perturbation. Indeed, even for mean estimation, we will show (Example 5 in Section 6) that the gap between the proposed and the optimal perturbation could be asymptotically large. Indeed, in both algorithms 1 and 2, the proposal of perturbation distribution P of noise c is independent of the selections of input selection X and randomness seed θ , while the optimal perturbation could be dependent on both. Theoretically, we can define the optimal perturbation problem as follows.

Definition 9 (Optimal Perturbation). *Given a data generation distribution \mathcal{P} , a mechanism $\mathcal{M}(\cdot, \theta)$, an objective mutual information bound v , and a perturbation scheme Q in a form that $(X, \theta) \sim Q^*(X, \theta)$, i.e., \mathcal{P} is generated from distribution $Q^*(X, \theta)$ given selections of X and θ , we call the noise distribution $Q^*(X, \theta)$ optimal w.r.t. $(\mathcal{P}, \mathcal{M}, \mathcal{K})$ for some loss function \mathcal{K} , if*

$$Q^*(X, \theta) = \arg \min_Q \mathbb{E}_{X, \theta} \mathcal{K}(c \parallel (X, \theta)), \quad (9)$$

s.t. $\text{MI}(X; \mathcal{M}(X, \theta) + c \parallel (X, \theta)) \leq v$.

Optimal perturbation characterizes the least noise (measured by loss function \mathcal{K}) that we have to add to achieve the required mutual information bound v . For example, in many applications, we are mainly concerned with the expected l_2 norm of perturbation, where $\mathcal{K}(\mathcal{B}(X, \theta)) = \|\mathcal{B}(X, \theta)\|_2$. Though a generic efficient solution to find the optimal perturbation may not exist, we present a *propose-then-verify* framework to approximate the optimal solution. Without loss of generality, we consider the generic setup where $\mathcal{M}(\cdot, \theta)$, $\theta \in \Theta$, could be randomized and we consider an equivalent discrete data generation procedure to randomly select X from a set $\mathcal{X} = \{X_1, X_2, \dots, X_N\}$ of size N , where the X_i s are not necessarily all distinct.

Theorem 6. *Assume that $\mathcal{M}(X, \theta) \in \mathbb{R}^d$ and $\|\mathcal{M}(X, \theta)\|_2 \leq r$ for some constant r , and for any perturbation proposal $Q(X, \theta)$, $(X, \theta) \sim Q(X, \theta)$, we apply Algorithm 3 to obtain $\bar{\psi}$. Then, if $\frac{N}{\tau_2}$ and $\frac{N}{\tau_3}$ are integers,*

one can ensure that $\text{MI}(X; \mathcal{M}(X, \theta) + \mathcal{N}(X, \theta) + c) \leq \bar{\psi} + \beta$ for an independent Gaussian noise of the form $c \sim \mathcal{N}(0, c \cdot \mathbf{I}_d)$ with confidence at least $1 - \gamma$ when

$$m \geq \frac{2 \log(1/\gamma)}{\beta^2} \cdot ((2r^2/c)^2/\tau_1 + \beta/3 \cdot (2r^2/c)).$$

Proof. See appendix F. □

In Algorithm 3, we can give a high-confidence claim on the security parameters produced by any perturbation proposal $Q(X, \theta)$. In practice, when input set X is large and representative enough, and the mechanism \mathcal{M} can stably learn populational information from X , the distributions of $\mathcal{M}(X)$ would be close to each other for most selections of X . Therefore, an empirical approximation framework of optimal perturbation can be described as follows.

1. We follow Steps 3-5 in Algorithm 2 to sample and evaluate on some large enough sets with respect to the selections of X and θ , respectively, and optimize the noise distribution to minimize its \mathcal{K} loss in (9) such that the estimated divergence defined in (10) is smaller than $(1/O(1)) \cdot v$. Here, v is the desired mutual information bound.
2. Provided the locally-optimized noise proposal \hat{Q} , we can apply Algorithm 3 to test \hat{Q} . If $\mathcal{M} \sim \hat{Q}$ does not produce the desired mutual information bound v , we adjust $\mathcal{M} \sim \hat{Q}$ in a form $\mathcal{M} + \alpha$, where α is an independent Gaussian noise of the form $\mathcal{N}(0, \alpha \cdot \mathbf{I}_d)$. We then perform binary search on α via Algorithm 3 until we find a proper value.

5 Composition

In this section, we study the composition of P C Privacy. Composition is an important requirement of a practical security definition. Since the privacy loss of any release cannot be revoked, it is necessary to keep track of the accumulated leakage. Ideally, the respective privacy loss of multiple releases is expected to be analytically aggregated to produce an upper bound on the accumulated loss incurred. For example, it is well known that for two mechanisms \mathcal{M}_1 and \mathcal{M}_2 , which each satisfy (ϵ_0, δ_0) -DP and use *independent* randomness, the composite mechanism $\bar{\mathcal{M}}(X) = (\mathcal{M}_1(X), \mathcal{M}_2(X))$ then satisfies $(2\epsilon_0, 2\delta_0)$ -DP. With advanced composition, one may obtain a sharpened bound $\tilde{O}(\sqrt{T \log \delta'} \epsilon_0, T\delta_0 + \delta')$ -DP for a free parameter $\delta' > 0$ after a T -fold composition of mechanisms each satisfying (ϵ_0, δ_0) -DP.

Composition also plays an important role in a classic privacy regime, especially for DP, to analyze complicated algorithms. As mentioned before, since the worst-case sensitivity is hard to compute, in general, to enable non-trivial privacy analysis, a complex algorithm usually needs to be artificially decomposed into multiple, relatively simpler operations. Composition then provides a way to upper bound the entire leakage by aggregating per-operation privacy loss. DP-SGD, one of the most widely-applied DP optimization methods, is an example. The DP analysis is developed in a compositional fashion where it is pessimistically assumed that intermediate gradients across each iteration are released. The entire iterative optimization is thus decomposed into multiple per-sample gradient aggregations [1]. Accordingly, the DP composition result also shows a simple way to determine the scale of noise required, which increases as $\tilde{O}(\sqrt{T})$ to produce a fixed privacy guarantee.

However, we have to stress that in DP, the above described results only hold when the randomness in different mechanisms is independent. This limitation makes tight Differential Privacy in many applications intractable. For example, if $\mathcal{M}_{[1:T]}$ represent T sequential processings on subsampled data from the sensitive X set but the subsampling schemes in different $\mathcal{M}_{[1:T]}$ are correlated, then more involved compositional analysis is needed. In comparison, our P C Privacy analysis framework can indeed handle composition of arbitrary (possibly correlated) mechanisms. Besides, P C Privacy allows us to measure the end-to-end privacy leakage of the output from an arbitrary (black-box) \mathcal{M} . Therefore, there is no need to artificially

decompose an algorithm and control leakage step-by-step. This is one key difference compared to DP: we use composition in P C Privacy only when multiple releases are necessary.

We first consider the simplest case for T mechanisms $\mathcal{M}_{[1:T]}$ where the data generation and randomness are both independent. By the chain rule of mutual information, we then have the following equation,

$$\text{MI}(X_{[1:T]}; \mathcal{M}_1(X_1, \theta_1), \dots, \mathcal{M}_T(X_T, \theta_T)) = \sum_{t=1}^T \text{MI}(X_t; \mathcal{M}_t(X_t, \theta_t)), \quad (11)$$

which states that mutual information can be simply summed. Therefore, by plugging $\sum_{k=1}^T \text{MI}(X_k; \mathcal{M}_k(X_k))$ into (2) in Theorem 1, we can obtain an upper bound on the posterior advantage $\Delta_{KL}\delta$, and accordingly a lower bound of δ .

Now, we turn to consider the more generic scenario where $\mathcal{M}_{[1:T]}$ apply independent randomnesses but share the same input X generated from \mathcal{D} . Unfortunately, (11) does not hold anymore in this setup and mutual information does *not* enjoy a triangle inequality: $\text{MI}(X; \mathcal{M}_1(X), \mathcal{M}_2(X))$ is not upper bounded by $\text{MI}(X; \mathcal{M}_1(X)) + \text{MI}(X; \mathcal{M}_2(X))$ in general. This is also true for most other selections of f -divergence and thus P C Privacy, in general, cannot be simply summed up to produce an upper bound. Fortunately, we will show below that the upper bound on mutual information studied in Algorithm 2 and Theorem 5 can be simply summed up to form a new bound under the composition. In other words, the P C Privacy of privatized algorithms using the proposed automatic analysis framework *does* enjoy a simple summable composition bound.

Theorem 7. *For arbitrary T mechanisms $\mathcal{M}_t(X, \theta_t)$, $t = 1, 2, \dots, T$, whose randomness seeds θ_t are independently selected, and data generation $X \sim \mathcal{D}$, if via Algorithm 2, for each \mathcal{M}_t a noise v_t is determined and independently generated such that with confidence $(1 - \gamma_t)$, $\text{MI}(X; \mathcal{M}_t(X, \theta_t)) \leq v_t$, then with confidence $(1 - \sum_{t=1}^T \gamma_t)$, $\text{MI}(X; \mathcal{M}_1(X, \theta_1), \dots, \mathcal{M}_T(X, \theta_T)) \leq \sum_{t=1}^T v_t$.*

Proof. See Appendix G. □

From Theorem 7, when the mechanisms $\mathcal{M}_{[1:T]}$ are identical, we need to increase the noise by a factor at most \sqrt{T} to ensure the same bound of the T -composite mutual information compared to the case of a single mechanism. Before proceeding, we have several important remarks. Even if one does not know the mechanisms to be composed in advance, one can choose v_t for each \mathcal{M}_t in an online fashion and use Algorithm 2 to determine appropriate noise for each exposure. Though the composition bound in Theorem 7 is simple, tighter analysis exists using the proposed automatic framework. It is noted that one can always build a new joint mechanism $\bar{\mathcal{M}}$ which outputs $(\mathcal{M}_1(X, \theta_1), \dots, \mathcal{M}_T(X, \theta_T))$, and simply apply Algorithm 2 on $\bar{\mathcal{M}}$ (or Algorithm 1 in the deterministic case).

In the following, we move on to the even more generic case for adaptive composition with correlated randomness. We consider T mechanisms $\mathcal{M}_{[1:T]}$ which are implemented in a sequential manner with possibly correlated selection of randomness seeds. More formally, the output $\mathcal{M}_t(X, \theta_t, \mathcal{M}_j(X, \theta_j), j = 1, 2, \dots, t-1)$ is defined by X , the randomness θ_t , and the previous outputs $\mathcal{M}_j(X, \theta_j)$ for $j = 1, 2, \dots, t-1$. Without loss of generality, in the following, we use $\mathcal{M}_t(X, \bar{\theta}_t)$ to denote the t -th mechanism where $\bar{\theta}_t = (\theta_1, \dots, \theta_t)$ represents the joint randomness seed and θ_t may be selected dependent on $\bar{\theta}_{t-1}$. On the other hand, we are given a sequence of privacy budgets $v_0 = 0 < v_1 < v_2 < \dots < v_T$ and our goal is to ensure that for any $t = 1, 2, \dots, T$, the accumulated privacy loss measured as mutual information is upper bounded by v_t with high confidence.

To this end, we propose an online implementation of Algorithm 2 described as Algorithm 4. The key idea is still to incrementally aggregate existing mechanisms to form a joint one, that we can apply similar techniques in Algorithm 2 to analyze. Algorithm 4 has an initialization step (lines 2-6) where m input pairs $(X_1^{(k)}, X_2^{(k)})$, $k = 1, 2, \dots, m$, are randomly generated from \mathcal{D} , and for each pair we select τ randomness seeds $\bar{\theta}_{1,l}^{(k)}$, $l = 1, 2, \dots, \tau$, for \mathcal{M}_1 (lines 3-4). In the online phase, for steps $t > 1$, we generate new randomness seeds $\bar{\theta}_{t,l}^{(k)}$ conditional on $\bar{\theta}_{t-1,l}^{(k)}$ determined previously (line 9) and continue the simulations on newly-incoming \mathcal{M}_t over the same m input pairs selected (line 11). This saves computation, avoiding reevaluating $\mathcal{M}_{[1:t-1]}$ on

fresh samples and randomness. Given previous noise schemes $\{\Sigma_{[1:(t-1)]}\}_{t=1}^T$ determined by Algorithm 4 such that $\text{MI}(X; \mathcal{M}_1(X, \bar{\theta}_1) + \Sigma_1, \dots, \mathcal{M}_{t-1}(X, \bar{\theta}_{t-1}) + \Sigma_{t-1}) \leq v_{t-1}$, we can always determine some independent Gaussian noise Σ_t such that

$$\text{MI}(X; \mathcal{M}_1(X, \bar{\theta}_1) + \Sigma_1, \dots, \mathcal{M}_{t-1}(X, \bar{\theta}_{t-1}) + \Sigma_{t-1}, \mathcal{M}_t(X, \bar{\theta}_t) + \Sigma_t) \leq v_t,$$

for any $v_t > v_{t-1}$ with high confidence. The following theorem describes the tradeoff between the confidence level and sampling complexity.

Theorem 8. *Assume that $\mathcal{M}_t(X, \bar{\theta}_t) \in \mathcal{R}^d$ and $\|\mathcal{M}_t(X, \bar{\theta}_t)\|_2 \leq r$ uniformly for some constant r , for $t = 1, 2, \dots, T$, and we apply Algorithm 4 which returns a sequence of Gaussian noise covariances Σ_t in an online setup. Then, one can ensure that $\text{MI}(X; \mathcal{M}_1(X, \bar{\theta}_1) + \Sigma_1, \dots, \mathcal{M}_t(X, \bar{\theta}_t) + \Sigma_t) \leq v_t$ for any $t = 1, 2, \dots, T$, with confidence at least $1 - \gamma$ once $m \geq \frac{8r \cdot \log(T/\gamma)}{c^2}$. c as before is a security parameter.*

Proof. See Appendix H. □

As a final remark, though the deterministic algorithm is a special case of the randomized algorithm (with a single randomness seed as discussed earlier) and we can always use Algorithm 2 for privacy analysis, Algorithm 1 can usually produce tighter noise bounds. However, the composition of the P C Privacy result derived from Algorithm 1 is more involved, which depends on the eigenspace estimated. We leave the generalization of Theorem 7 for Algorithm 1 as an open problem. However, with simulations on incrementally aggregated mechanisms to be composed, one can similarly build an online version of Algorithm 1 to determine the perturbation sequentially given increasing privacy budget for composition.

Example 2 (Composition). *We take the Hospital dataset from Mathworks Sample² as an example. This set contains the records of 100 patients and we consider two mechanisms \mathcal{M}_1 and \mathcal{M}_2 which return the average of ages, weight, and the average of blood pressure range, respectively. We consider randomly selecting 50 individual samples to form the input X and via Algorithm 2 we need Gaussian noises $\mathbb{E}[\|\Sigma_1\|_2] = 6.85$ and $\mathbb{E}[\|\Sigma_2\|_2] = 2.60$ to ensure that $\text{MI}(X; \mathcal{M}_1(X) + \Sigma_1) \leq 0.5$ and $\text{MI}(X; \mathcal{M}_2(X) + \Sigma_2) \leq 0.5$. With the simple composition bound, we have that $\text{MI}(X; \mathcal{M}_1(X) + \Sigma_1, \mathcal{M}_2(X) + \Sigma_2) \leq 1$. As a comparison, by applying the automatic analysis on the joint mechanism $\bar{\mathcal{M}} = (\mathcal{M}_1, \mathcal{M}_2)$, we obtain a tighter noise bound where there exists certain Gaussian noise $\bar{\Sigma}$ with $\mathbb{E}[\|\bar{\Sigma}\|_2] = 4.57$ to ensure $\text{MI}(X; \bar{\mathcal{M}}(X) + \bar{\Sigma}) \leq 1$. Using the simple composition, it is noted that the noise magnitude $\mathbb{E}[\|(\Sigma_1, \Sigma_2)\|_2] = 7.32$, when we handle \mathcal{M}_1 and \mathcal{M}_2 separately with lower-complexity privacy analysis.*

6 Comparison between Adversarial and Instance-based Worst Case

We start by a lower bound of the Gaussian noise if we want to ensure an input-independent upper bound of the mutual information.

Theorem 9 (Lower Bound of Adversarial Gaussian Mechanism). *Given that $\|\mathcal{M}(X)\|_2 \leq r$, for an arbitrary Gaussian perturbation mechanism $Q(y) = \mathcal{N}(0, \Sigma(y))$, i.e., when $\mathcal{M}(X) = y \in \mathbb{R}^d$, we generate $\Sigma \sim Q(y)$ and publish $\mathcal{M}(X) + \Sigma$. Given $v = o(1)$, if*

$$\sup_{\mathcal{M}, \mathbb{P}} \text{MI}(X; \mathcal{M}(X) + \Sigma) \leq v,$$

then for any $y \in \mathbb{R}^d$ and $\Sigma \sim Q(y)$, $\mathbb{E}[\|\Sigma\|_2] = \Omega(r\sqrt{d}/\sqrt{v})$.

Proof. See Appendix I. □

²<https://www.mathworks.com/help/stats/sample-data-sets.html>

Algorithm 4 Online Noise Determination of ϵ -Adaptively-Composed Randomized Mechanisms $\mathcal{M}_{[1:T]}$

- 1: **Input:** set of randomized mechanisms $\mathcal{M}_t(\cdot, \bar{\theta}_t) : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{Y}$, $\bar{\theta}_t \in \mathcal{Y} \times \dots \times \mathcal{Y}$, where $\mathcal{M}_t(\cdot, \bar{\theta}_t)$ is provided at time slot t for $t = 1, 2, \dots, T$, data distribution \mathcal{D} , sampling complexity m , security parameters c, τ , and privacy budget $\{v_0 = 0, v_1, v_2, \dots, v_T\}$ such that $0 < v_1 < v_2 < \dots < v_T$.
 - 2: **for** $k = 1, 2, \dots, m$ **do**
 - 3: From distribution \mathcal{D} , independently sample $X_1^{(k)}$ and $X_2^{(k)}$.
 - 4: Initially randomly select τ seeds, denoted by $\bar{C}_1^{(k)} = \{\bar{\theta}_{1,1}^{(k)}, \bar{\theta}_{1,2}^{(k)}, \dots, \bar{\theta}_{1,\tau}^{(k)}\}$, $\bar{\theta}_{1,l}^{(k)} \in \mathcal{Y}$.
 - 5: **end for**
 - 6: **for** $t = 1, 2, \dots, T$ **do**
 - 7: **for** $k = 1, 2, \dots, m$ **do**
 - 8: **if** $t > 1$ **then**
 - 9: Conditional on $\bar{\theta}_{t-1,l}^{(k)}$, randomly select $\theta_{t,l}^{(k)} \in \mathcal{Y}$ and form a new joint seed $\bar{\theta}_{t,l}^{(k)} = (\bar{\theta}_{t-1,l}^{(k)}, \theta_{t,l}^{(k)})$, for $l = 1, 2, \dots, \tau$.
 - 10: **end if**
 - 11: Compute and record

$$\mathbf{y}_t^{(k,1)} = (\mathcal{M}_t(X_1^{(k)}, \bar{\theta}_{t,1}^{(k)}), \dots, \mathcal{M}_t(X_1^{(k)}, \bar{\theta}_{t,\tau}^{(k)})),$$

$$\mathbf{y}_t^{(k,2)} = (\mathcal{M}_t(X_2^{(k)}, \bar{\theta}_{t,1}^{(k)}), \dots, \mathcal{M}_t(X_2^{(k)}, \bar{\theta}_{t,\tau}^{(k)})).$$
 - 12: Compute averaged l_2 -norm pairwise distance $\psi_t^{(k)} = \|\mathbf{y}_t^{(k,1)} - \mathbf{y}_t^{(k,2)}\|_2^2 / \tau$.
 - 13: **end for**
 - 14: Calculate empirical mean $\bar{\psi}_t = \sum_{k=1}^m \psi_t^{(k)} / m$.
 - 15: **Output:** Gaussian covariance $\Sigma_t = \frac{\bar{\psi}_t + c}{2(v_t - v_{t-1})} \cdot \mathbf{I}_d$.
 - 16: **end for**
-

From Theorem 9, we see that if we want a data-independent Gaussian perturbation to ensure bounded mutual information in the adversarial worst case, we must add a noise of magnitude of $\Omega(r\sqrt{d}/\sqrt{v})$ to *any possible output* $\mathcal{M}(X)$ ³. On the other hand, from Theorem 3, we know this bound is also tight where a Gaussian noise of magnitude $O(r\sqrt{d}/\sqrt{v})$ can universally ensure the mutual information is bounded by v . This matches our intuition that without any other assumptions if we want to hide arbitrary information disclosure (bounded by r in l_2 norm), we need noise on each direction in \mathbb{R}^d in a scale of $\Omega(r)$, which finally produces the $\Omega(r\sqrt{d})$ magnitude. This, in general, makes the perturbed release useless. Therefore, for a particular processing mechanism \mathcal{M} , when the practical data is far away from the adversarial worst case or the data-independent worst-case proof cannot be tightly produced, there could be a huge utility compromise relative to the P-C Privacy measurement from an instance-based angle.

Example 3 (Membership Inference from Mean Estimation). *We suppose U to be CIFAR-10 [45], a canonical test dataset commonly used in computer vision, which consists of 60,000 32×32 color images. We normalize each pixel to within $[0, 1]$ and take each sample as a 3072-dimensional vector. We suppose each of the 60,000 3072-dimensional samples is i.i.d. randomly selected into the set X , where the expected cardinality of $|X|$ equals 30,000, to conduct the mean estimation analysis, where $\mathcal{M}(X)$ simply returns the sum of X divided by 30,000, as an unbiased mean estimator. Via Algorithm 1, we have that an independent Gaussian noise whose $\mathbb{E}[\|\cdot\|_2] = 0.28$ is sufficient to ensure $\text{MI}(X; \mathcal{M}(X) + \cdot) \leq 1$. On the other hand, in the adversarial worst-case setup, the data of normalized pixel suggests that the l_∞ -norm of the sensitivity is bounded by 1. The mutual information bound requires $\xi = 1/60000$ -zCDP [14], in turn requiring a noise $\mathbb{E}[\|\cdot\|_2] = 17.7$, which is $63\times$ larger than the instance-based analysis. As a comparison, if we apply Algorithm 2 to compute the noise bound, it is required that $\mathbb{E}[\|\cdot\|_2] = 3.94$ which is looser compared to 0.28 obtained from Algorithm 1*

³In Theorem 9, we restrict the noise distribution to be Gaussian. We leave a generic lower bound for arbitrary noise distribution as an open problem.

1, though it is more efficient and only requires the estimation on the (pair-wise) l_2 distance among simulation outputs.

In Example 3, we already utilize the property that each data point is independently sampled and the individual sensitivity of the averaging operator has a closed form $O(1/n)$ to get a tighter bound for the adversarial worst case [14]. However, in general, when the sampling is not independent or the sensitivity is intractable, one may need to adopt the loose noise bound shown in Theorem 9, where a noise in a scale of 2.2×10^3 is needed. Below, we provide such an example where producing a non-trivial adversarial worst-case proof remains open.

Example 4 (Private Machine Learning). *We consider measuring data leakage of deep learning on practical data. In general, in machine learning, our goal is to optimize the weights/parameters w of some model, viewed as some function $G(\cdot, w)$, to minimize the loss $\min_w \mathbb{E}_{(x,y)} \mathcal{L}(G(x, w), y)$. Here, \mathcal{L} represents some loss function, while x and y represent the feature and label of a sample, respectively. Ideally, we expect the trained out model $G(x, w)$ to predict the true label y well.*

In this example, we consider training a three-layer fully-connected neural network on the MNIST dataset [47]. MNIST contains 70,000 28×28 handwritten-digit images. We consider a data generation X by randomly sampling 35,000 samples out of the entire data set. The neural network has three layers, where the weights of the first layer are a 784×30 matrix, those of the second layer are a 30×30 matrix, and those of the last layer are a 30×10 matrix. We select the activation function between layers to be the Relu function and use cross-entropy as the loss function, as a common setup in deep learning [46]. The total number of parameters in this network is 24,790. Let the mechanism \mathcal{M} correspond to running full-batch gradient descent for 1,500 iterations with step size 0.05 and outputting the final weight obtained. From Algorithm 1, we have that an independent Gaussian noise $\mathbb{E}[\| \cdot \|_2] = 3.7$ is sufficient to ensure $\text{MI}(X; \mathcal{M}(X) + \cdot) \leq 1$. Non-privately, the trained-out neural network achieves 94.8% classification accuracy; under the perturbation to ensure ϵ -C Privacy, we achieve an accuracy of 93.5%, i.e., a slight compromise. Unfortunately, for the adversarial data-independent worst case, the sensitivity/stability of generic non-convex optimization remains open ⁵.

Before the end of this section, we want to point out that there could be an asymptotically large gap between our proposed conservative perturbation and the optimal noise required, in particular for the individual privacy case.

Example 5 (Gap to Optimal Perturbation). *We return to the problem of mean estimation of Gaussian distributions mentioned in Section 1, where we assume each $x_i \in \mathbb{R}^d \sim \mathcal{N}(\mu, \Sigma)$ and $\mathcal{M}(X) = 1/n \cdot \sum_{i=1}^n x_i$. We focus on the individual privacy $\text{MI}(x_i; \mathcal{M}(X))$ and set out to quantify the least noise needed. For any independent Gaussian noise, $\cdot' \sim \mathcal{N}(0, \Sigma \cdot')$, we have*

$$\text{MI}(x_i; \mathcal{M}(X) + \cdot') = \frac{1}{2} \cdot \log \left(\det \left(\left(\frac{1}{n} \cdot \Sigma + \Sigma \cdot' \right) \cdot \left(\frac{n-1}{n^2} \cdot \Sigma + \Sigma \cdot' \right)^{-1} \right) \right). \quad (12)$$

On the other hand, when we apply Algorithm 1 and measure $\text{MI}(X; \mathcal{M}(X))$ as a uniform upper bound,

$$\text{MI}(X; \mathcal{M}(X) + \cdot) = \frac{1}{2} \cdot \log \left(\det \left(\left(\frac{1}{n} \cdot \Sigma + \Sigma \cdot \right) \cdot \Sigma^{-1} \right) \right). \quad (13)$$

From (12) and (13), we can determine the corresponding Gaussian noises to ensure a mutual information bound v and their gap shows as a lower bound of the gap between \cdot' and the optimal perturbation \cdot , since

If each entry of X is independently generated, and \cdot satisfies ξ -Concentrated Differential Privacy (CDP), then $\text{MI}(X; \mathcal{M}(X)) \leq \xi$ [14].

⁵ n upper bound on individual sensitivity of Empirical Risk Minimization is only known for strongly-convex optimization [18] with an additional Lipschitz assumption and the loss function needs to be a sum of individual losses on each sample.

in this example we already restrict the noises to be Gaussian. With some calculation, we have

$$\mathbb{E}[\| \cdot \|_2^2] - \mathbb{E}[\| \cdot \|_2^2] \begin{cases} \mathbb{E}[\| \cdot \|_2^2] = \frac{1}{2v} \cdot \left(\sum_{j=1}^d \sqrt{\lambda_j}\right)^2, & v \geq \frac{d}{n-1} \\ \sum_{j=1}^d \min \left\{ \sqrt{\lambda_j} \left(\sum_{j=1}^d \sqrt{\lambda_j}\right) / (2v), (n-1)\lambda_j \right\}, & v < \frac{d}{n-1}. \end{cases} \quad (14)$$

Here, $\lambda_{[1:d]}$ are the eigenvalues of Σ . From (14), when the privacy budget $v \geq d/(n-1)$, there is even no need to add noise; in the high privacy regime, such a gap could be $\Omega(n \text{Tr}(\Sigma_{\mathcal{M}(X)}))$. Therefore, there is still much room to improve the current analysis.

7 Related Works

7.1 Differential Privacy

DP is the most successful and practical information-theoretic privacy notion over the last two decades. Lower bounds both in DP noise [34], [41] and Theorem 9 suggest that, in general, there is no free lunch for input-independent privacy guarantees – to produce meaningful security parameters, the perturbation could be much larger than the objective disclosure and destroy the released information. This is especially true in the high-dimensional case, known as the curse of dimensionality. DP and our P C Privacy propose two different solutions to overcome this challenge.

Differential Privacy (DP) can be most naturally interpreted in a binary hypothesis form. Consider two adjacent datasets S_0 and S_1 which only differ in one individual datapoint and two hypotheses: H_0 where the input set is S_0 , and H_1 where the input set is S_1 . Strong DP guarantee implies that the hypothesis testing error (a combination of Type I and II errors) is large [28]. In general, when the adversarial objective is specified to identify the secret from a class of candidates, such inference can be naturally characterized by a (multiple) hypothesis testing model, and one natural privacy definition is to measure the lower bound of the probability when the adversary outputs a wrong guess [17].

The beauty of DP in bridging meaningful privacy and tolerable utility compromise is to focus on an individual datapoint rather than the entire sensitive dataset. DP stays within the data-independent privacy regime. ϵ -DP for constant ϵ cannot make meaningful guarantees for the entire data, because the security parameter grows linearly with $n\epsilon$ for the privacy concern of n datapoints (group privacy). However, a noise of magnitude $O(\sqrt{d}/n)$ can hide any individual among the n -population if one can ensure an $O(1/n)$ sensitivity. Therefore, once the population size n is large enough and we can ensure the contribution from each individual to the output is limited enough, reasonably small perturbation can ensure good individual privacy, interpreted as some constant hypothesis test error [5]. However, as mentioned before, the worst-case sensitivity proof becomes the key step in DP, which is generally not easy beyond simple aggregation.

P C Privacy tackles this challenge by enforcing the privacy measurement to be instance-based to get rid of the restriction from data-independent impossibilities. In terms of the objective to protect, P C Privacy is stronger and more generic, where now the privacy of *the entire dataset* could be preserved once the processing mechanism $\mathcal{M}(X)$ learns or behaves stably based on some population property of the underlying distribution \mathcal{D} , from which the sensitive data X is generated. Individual Privacy is only a special case in P C Privacy. One can imagine that if the output of \mathcal{M} is only dependent on the distribution \mathcal{D} , X is then independent of $\mathcal{M}(X)$ and we achieve perfect secrecy with respect to the whole dataset. In particular, for applications in machine learning, P C Privacy would benefit from strong learning algorithms \mathcal{M} of good generalization. More details can be found in Section 7.2.

However, P C Privacy also has limitations since it is not applicable when we cannot simulate or we don't have access to the data generation of X . In contrast, data-independent privacy (if it can be proven) can always work with meaningful interpretation regardless of any assumptions on priors. But for most statistical data processing, where we are allowed to conduct the analysis on subsampled data, this limitation can be overcome. Even for a given deterministic dataset, through sub-sampling we can enforce that the selected

samples enjoy entropy and P C Privacy provides a strong security interpretation: even the adversary, who has full knowledge of the data pool, cannot infer too much about selected samples in the processing mechanism \mathcal{M} .

7.2 Other Privacy Metrics and Learning Theory Qualities

Mutual Information (MI) based Attack and Privacy Definition: MI is a well-studied measure on the dependence between two random variables quantified in information bits. Naturally, one may consider the MI between the secret and the release/leakage for security purposes. A large number of existing works apply MI to construct side-channel attacks [9, 35] or even simply adopt MI as a privacy definition [16, 48]. For example, given known plaintexts, [35] built a distinguisher based on empirical MI estimation where the adversary proposes a guessing of the encryption key with the highest correlation score regarding the side-channel leakage. [16] showed further asymptotic analysis of empirical MI estimations under discrete randomness. Though we use f -divergence, where MI based on KL-divergence is a special case, to develop impossibility results as the foundation of P C Privacy, our motivations and results are very different from those prior works. On one hand, P C Privacy is rigorously resistant to arbitrary adversary attacks⁶, without any restrictions on the adversary’s inference strategy or the computation power. In addition, P C Privacy does *not* use MI or other information metrics, such as Fisher information [32], discussed below, or anonymity set [55, 60], as the leakage measurement, which, in general, lack semantic security interpretation. Unfortunately, except DP, most of those metrics are not semantic. For example, recent work [23] argued that the concept of anonymity set does not resist singling-out attacks and does not produce desirable individual privacy guarantees. In contrast, the entire theory developed for P C Privacy is devoted to answering the fundamental question of how to quantify the hardness of an arbitrary adversarial inference task.

DP under Adversary Uncertainty: With a different motivation to add less noise, there is a line of works such as distributional DP [6] and noiseless privacy [12], which study the relaxation of DP and take the data entropy into account. As mentioned before, the original DP definition ensures indistinguishable likelihoods of any two adjacent datasets in the worst case. Instead, [6, 12] considered statistical adjacent datasets where both their common part and the differing datapoint are generated from some distributions. The goal of such relaxation is to exploit the data entropy to substitute (part of) the external perturbation for the release. However, since those works are still developed within the DP framework, [6, 12] encounter similar restrictions as those in generic DP analysis. First, analytical distributional DP [6] and noiseless privacy [12] bounds are only known for a limited number of applications with very specific assumptions on the underlying data distributions. Second, the privacy concern is still regarding the participation of an individual. In contrast, P C Privacy can automatically handle any black-box data generation and processing mechanism for arbitrary adversarial inference.

Empirical Membership Inference/Data Reconstruction: As a consequence of the loose/conservative worst-case privacy analysis such as DP [29] and maximal leakage [42], a large number of works are devoted to empirically measuring the actual information leakage [57, 68] or the actual privacy guarantees produced by existing privatization methods [52, 61] especially in machine learning. For example, [57] measured the influence of overfitting on the adversary’s advantage of membership or attribute inference. The results in [57] suggest that without proper privacy preservation, many popular machine learning algorithms could have severe privacy risks. On the other hand, [52] empirically studied the actual privacy guarantee provided by DP-SGD and showed that there is a substantial gap between the best theoretical bound we can claim so far and the practical distinguishing advantage of the adversary. [61] studied DP-SGD from a different angle. [61] showed that compared to distinguishing individual participation, DP-SGD may provide a much stronger privacy guarantee against extraction of rare features in training data. In other words, certain data reconstruction tasks could be much harder than simply identifying the enrollment of an individual sample. This matches our earlier discussion on data reconstruction and distinguishing attack in the introduction section. Empirical works on overfitting and generalization control are complementary, and can instruct the

⁶Besides MI, there are many other efficient side-channel attacks based on different statistical tests such as Pearson correlation [13].

design of more stable processing mechanisms with better privacy-utility tradeoff under P C Privacy.

Fisher Privacy: recently proposed privacy metric, which also measures the data reconstruction error, is rooted in the Fisher information and the Cramér–Rao bound, and is termed Fisher information leakage [32, 39, 40]. Given an upper bound of Fisher information, one can lower bound the expected l_2 norm of reconstruction error for an adversary with *given bias*. However, Fisher information leakage is still input-independent and based on likelihood difference. There are several essential differences compared to priors-based P C Privacy. First, to apply the Cramér–Rao bound in data reconstruction, the bias of the adversary’s reconstruction on each selection of X must be specified beforehand, while the optimal selection of such point-wise bias is, in general, unknown. In contrast, P C Privacy does not put any restrictions or assumptions on the adversary’s strategy. Second, Fisher information only measures the mean squared error whereas P C Privacy handles impossibility of arbitrary inference tasks and criteria.

Generalization Error: In learning theory, generalization error is defined as the gap between the loss (prediction accuracy) by applying the trained model on local data (seen) and on the data population (unseen) [67]. To be formal, in our context, let $\mathcal{M}(X)$ represent the model learned via the (possibly randomized) algorithm \mathcal{M} from a set of samples X , generated from some distribution \mathcal{D} . Given some loss \mathcal{L} , the generalization error is defined as

$$\text{Gen}(\mathcal{D}, \mathcal{M}) = \mathbb{E}_{\bar{X} \perp X \sim \mathcal{D}}(\mathcal{L}(\bar{X}, \mathcal{M}(X))) - \mathbb{E}_{X \sim \mathcal{D}}(\mathcal{L}(X, \mathcal{M}(X))).$$

$\bar{X} \perp X \sim \mathcal{D}$ represents that \bar{X} and X are i.i.d. in \mathcal{D} . $\mathbb{E}_{\bar{X} \perp X \sim \mathcal{D}}(\mathcal{L}(\bar{X}, \mathcal{M}(X)))$ captures the expected accuracy from the model $\mathcal{M}(X)$ on population data, while $\mathbb{E}_{X \sim \mathcal{D}}(\mathcal{L}(X, \mathcal{M}(X)))$ captures the expected training accuracy. It is well known that DP can control the generalization error [7]. Indeed, the relationship between DP, P C Privacy and generalization error can be described as follows. As mentioned earlier, DP guarantees of $\mathcal{M}(X)$ can upper bound the mutual information $\text{MI}(X; \mathcal{M}(X))$ [14], and from [67], $\text{MI}(X; \mathcal{M}(X))$ can then upper bound the generalization error. In the following, we show that P C Privacy can also control the generalization error. For simplicity, we assume that the distribution \mathcal{D} of X is uniform over a finite set $\mathcal{X} = \{X_1, X_2, \dots, X_N\}$ of N elements.

Theorem 10. *For any given $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$, and loss function $\mathcal{L} : (\mathcal{X}, \mathcal{Y}) \rightarrow (0, 1)$, suppose \mathcal{D} is a uniform distribution over \mathcal{X} . If \mathcal{M} satisfies $\Delta_{KL} \delta \leq v$, for the criterion $\rho(\bar{X}, X) = 1$ if $\bar{X} = X$, i.e., the identification problem, then we have $\text{Gen}(\mathcal{D}, \mathcal{M}) \leq \sqrt{v/2}$.*

Proof. See appendix J. □

However, we need to stress that the reverse direction does not hold in general. Recall the example of mutual information in mean estimation in Section 4.1: we can construct some discrete distribution where given sufficiently many samples one can learn the true mean arbitrarily accurately but the mean estimator \mathcal{M} could be a bijective function, where $\mathcal{M}(X)$ essentially leaks everything about X . However, small generalization error does help P C Privacy if the learning mechanism ensures the trained out model is close to the population optimum. Thus, even in a case that $\mathcal{M}(\cdot)$ is bijective, we only need small noise to satisfy a mutual information bound.

8 Conclusions and Prospects

In this paper, we propose and study a new instance-based information-theoretic privacy notion, termed P C Privacy. P C Privacy enjoys intuitive and meaningful interpretations: it enables concrete measurement on the adversary’s success rate or the posterior advantage for arbitrary data inference/reconstruction task with the observation of disclosures. Simple quantification via generic f -divergence is presented. More importantly, based on data priors, we show an automatic privacy analysis and proof generation framework, where theoretically no algorithmic worst-case proof is needed. Data leakage control of black-box oracles becomes possible via P C Privacy. Though the instance-based and adversarial setups are, in general, incomparable, P C Privacy does have strong connections to cryptography and DP. Indeed, many problems or concerns

in the research or applications of DP and cryptography have a corresponding version in P C Privacy. The underlying technical challenges are not necessarily the same but motivations are very similar. We list some interesting problems here.

Gap between the Local and Central: Collaboration of semi-honest users to amplify local DP via shuffling models has been recently proposed and studied in [22], [30]. The key idea is to securely aggregate the local noisy response from each user, where the underlying locally added noises also aggregate to produce a more powerful perturbation with amplified privacy. As analyzed in Appendix K, via MPC, P C Privacy analysis can be implemented in a decentralized setup and each user’s local P C Privacy would then also benefit from other people’s data entropy, similar to a centralized case. However, local P C Privacy puts more requirements on MPC, and cheap MPC implementation is a key challenge in making it practical in large-scale systems.

Optimal Perturbation: Minimal perturbation is another fundamental problem in both P C Privacy and DP. In DP, the theoretical study of (asymptotically) optimal utility loss in different data processing remains very active. Many tight results are known for mean estimation [3], convex ERM optimization [8] with specific solution restrictions like on l_1 norm [2], stochastic optimization [33], and principal component analysis (PCA) [19]. All these problems can also be systemically studied under P C Privacy. While in P C Privacy, privacy at a certain level could come for free (see Example 5), our current automatic framework needs noise to produce generic high-confidence security parameters. Therefore, besides a theoretically tight perturbation bound, P C Privacy also raises its own special research problem, i.e., how to efficiently implement privacy analysis protocols and produce perturbation schemes matching the optimal bound? For example, with assistance of public data, how can we properly truncate the output domain and reduce sampling/simulation complexity?

acknowledgements: We gratefully acknowledge the support of DST Singapore, Cisco Systems, Capital One, and a MathWorks fellowship. We also thank the anonymous reviewers for their helpful comments.

References

- [1] Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC conference on computer and communications security. pp. 308–318 (2016)
- [2] Shi, H., Feldman, V., Koren, T., Talwar, K.: Private stochastic convex optimization: Optimal rates in ℓ_1 geometry. In: International Conference on Machine Learning. pp. 393–403. PMLR (2021)
- [3] Shi, H., Feldman, V., Talwar, K.: Optimal algorithms for mean estimation under local differential privacy. arXiv preprint arXiv:2205.02466 (2022)
- [4] Balle, B., Barthe, G., Gaboardi, M.: Privacy amplification by subsampling: Tight analyses via couplings and divergences. *Advances in Neural Information Processing Systems* **31** (2018)
- [5] Balle, B., Barthe, G., Gaboardi, M., Hsu, J., Sato, T.: Hypothesis testing interpretations and renyi differential privacy. In: International Conference on Artificial Intelligence and Statistics. pp. 2496–2506. PMLR (2020)
- [6] Bassily, R., Groce, A., Katz, J., Smith, A.: Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science. pp. 439–448. IEEE (2013)
- [7] Bassily, R., Nissim, K., Smith, A., Steinke, T., Stemmer, U., Ullman, J.: Algorithmic stability for adaptive data analysis. In: Proceedings of the forty-eighth annual ACM symposium on Theory of Computing. pp. 1046–1059 (2016)
- [8] Bassily, R., Smith, A., Thakurta, A.: Private empirical risk minimization: Efficient algorithms and tight error bounds. In: 2014 IEEE 55th annual symposium on foundations of computer science. pp. 464–473. IEEE (2014)
- [9] Batina, L., Gierlichs, B., Prouff, E., Rivain, M., Standaert, F.X., Veyrat-Charvillon, N.: Mutual information analysis: a comprehensive study. *Journal of Cryptology* **24**(2), 269–291 (2011)
- [10] Belghazi, M.I., Baratin, A., Rajeswar, S., Ozair, S., Bengio, Y., Courville, A., Hjelm, R.D.: Mine: mutual information neural estimation. arXiv preprint arXiv:1801.04062 (2018)
- [11] Ben-Sasson, E., Fehr, S., Ostrovsky, R.: Near-linear unconditionally-secure multiparty computation with a dishonest minority. In: Annual Cryptology Conference. pp. 663–680. Springer (2012)
- [12] Bhaskar, R., Bhowmick, A., Goyal, V., Laxman, S., Thakurta, A.: Noiseless database privacy. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 215–232. Springer (2011)
- [13] Brier, E., Clavier, C., Olivier, F.: Correlation power analysis with a leakage model. In: International workshop on cryptographic hardware and embedded systems. pp. 16–29. Springer (2004)
- [14] Bun, M., Steinke, T.: Concentrated differential privacy: Simplifications, extensions, and lower bounds. In: Theory of Cryptography Conference. pp. 635–658. Springer (2016)
- [15] Butte, A.J., Kohane, I.S.: Mutual information relevance networks: functional genomic clustering using pairwise entropy measurements. In: Biocomputing 2000, pp. 418–429. World Scientific (1999)
- [16] Chatzikokolakis, K., Chothia, T., Guha, R.: Statistical measurement of information leakage. In: International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 390–404. Springer (2010)

- [17] Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: Probability of error in information-hiding protocols. In: 20th IEEE Computer Security Foundations Symposium (CSF'07). pp. 341–354. IEEE (2007)
- [18] Chaudhuri, K., Monteleoni, C., Sarwate, D.: Differentially private empirical risk minimization. *Journal of Machine Learning Research* **12**(3) (2011)
- [19] Chaudhuri, K., Sarwate, D., Sinha, K.: Near-optimal differentially private principal components. *Advances in neural information processing systems* **25** (2012)
- [20] Chen, S., Wang, R., Wang, X., Zhang, K.: Side-channel leaks in web applications: reality today, a challenge tomorrow. In: 2010 IEEE Symposium on Security and Privacy. pp. 191–206. IEEE (2010)
- [21] Cheng, P., Hao, W., Dai, S., Liu, J., Gan, Z., Carin, L.: Club: contrastive log-ratio upper bound of mutual information. In: International conference on machine learning. pp. 1779–1788. PMLR (2020)
- [22] Cheu, S., Smith, A., Ullman, J., Zeber, D., Zhilyaev, M.: Distributed differential privacy via shuffling. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 375–403. Springer (2019)
- [23] Cohen, E., Nissim, K.: Towards formalizing the gdpr’s notion of singling out. *Proceedings of the National Academy of Sciences* **117**(15), 8344–8352 (2020)
- [24] Cover, T.M.: Elements of information theory. John Wiley & Sons (1999)
- [25] Cramer, R., Damgård, I.B., et al.: Secure multiparty computation. Cambridge University Press (2015)
- [26] Davies, S., Mazumdar, A., Pal, S., Rashtchian, C.: Lower bounds on the total variation distance between mixtures of two Gaussians. In: International Conference on Algorithmic Learning Theory. pp. 319–341. PMLR (2022)
- [27] Devroye, L., Mehrabian, A., Reddad, T.: The total variation distance between high-dimensional gaussians. arXiv preprint arXiv:1810.08693 (2018)
- [28] Dong, J., Roth, A., Su, W.J.: Gaussian differential privacy. arXiv preprint arXiv:1905.02383 (2019)
- [29] Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Theory of cryptography conference. pp. 265–284. Springer (2006)
- [30] Erlingsson, Ú., Feldman, V., Mironov, I., Raghunathan, A., Talwar, K., Thakurta, R.: Amplification by shuffling: From local to central differential privacy via anonymity. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 2468–2479. SIAM (2019)
- [31] Erlingsson, Ú., Feldman, V., Mironov, I., Raghunathan, A., Talwar, K., Thakurta, R.: Amplification by shuffling: From local to central differential privacy via anonymity. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 2468–2479. SIAM (2019)
- [32] Farokhi, F., Sandberg, H.: Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries. *IEEE Transactions on Smart Grid* **9**(5), 4726–4734 (2017)
- [33] Feldman, V., Koren, T., Talwar, K.: Private stochastic convex optimization: optimal rates in linear time. In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. pp. 439–449 (2020)
- [34] Geng, Q., Viswanath, P.: Optimal noise adding mechanisms for approximate differential privacy. *IEEE Transactions on Information Theory* **62**(2), 952–969 (2015)
- [35] Gierlichs, B., Batina, L., Tuyls, P., Preneel, B.: Mutual information analysis. In: International Workshop on Cryptographic Hardware and Embedded Systems. pp. 426–442. Springer (2008)

- [36] Goldwasser, S., Klein, S., Mossel, E., Tamuz, O.: Publicly verifiable randomness (2018)
- [37] Goldwasser, S., Micali, S.: Probabilistic encryption & how to play mental poker keeping secret all partial information. In: Proceedings of the fourteenth annual ACM symposium on Theory of computing. pp. 365–377 (1982)
- [38] Goldwasser, S., Micali, S.: Probabilistic encryption. *Journal of computer and system sciences* **28**(2), 270–299 (1984)
- [39] Guo, C., Karrer, B., Chaudhuri, K., van der Maaten, L.: Bounding training data reconstruction in private (deep) learning. *arXiv preprint arXiv:2201.12383* (2022)
- [40] Hannun, A., Guo, C., van der Maaten, L.: Measuring data leakage in machine-learning models with fisher information. In: Uncertainty in Artificial Intelligence. pp. 760–770. PMLR (2021)
- [41] Hardt, M., Talwar, K.: On the geometry of differential privacy. In: Proceedings of the forty-second ACM symposium on Theory of computing. pp. 705–714 (2010)
- [42] Issa, I., Wagner, D.B., Kamath, S.: An operational approach to information leakage. *IEEE Transactions on Information Theory* **66**(3), 1625–1657 (2019)
- [43] Jin, C., Netrapalli, P., Ge, R., Kakade, S.M., Jordan, M.I.: A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736* (2019)
- [44] Kasiviswanathan, S.P., Lee, H.K., Nissim, K., Raskhodnikova, S., Smith, A.: What can we learn privately? *SIAM Journal on Computing* **40**(3), 793–826 (2011)
- [45] Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
- [46] LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
- [47] LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
- [48] Makhdoomi, A., Salamatian, S., Fawaz, N., Médard, M.: From the information bottleneck to the privacy funnel. In: 2014 IEEE Information Theory Workshop (ITW 2014). pp. 501–505. IEEE (2014)
- [49] Mironov, I.: Rényi differential privacy. In: 2017 IEEE 30th computer security foundations symposium (CSF). pp. 263–275. IEEE (2017)
- [50] Moon, Y.I., Rajagopalan, B., Lall, U.: Estimation of mutual information using kernel density estimators. *Physical Review E* **52**(3), 2318 (1995)
- [51] Murtagh, J., Vadhan, S.: The complexity of computing the optimal composition of differential privacy. In: Theory of Cryptography Conference. pp. 157–175. Springer (2016)
- [52] Nasr, M., Songi, S., Thakurta, A., Papemoti, N., Carlin, N.: Adversary instantiation: Lower bounds for differentially private machine learning. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 866–882. IEEE (2021)
- [53] Pinsker, M.S., Prelov, V.V., Verdu, S.: Sensitivity of channel capacity. *IEEE Transactions on Information Theory* **41**(6), 1877–1888 (1995)
- [54] Sason, I., Verdú, S.: f -divergence inequalities. *IEEE Transactions on Information Theory* **62**(11), 5973–6006 (2016)
- [55] Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Privacy Enhancing Technologies: Second International Workshop, PET 2002 San Francisco, CA, USA, April 14–15, 2002 Revised Papers 2. pp. 41–53. Springer (2003)

- [56] Shannon, C.E.: Communication theory of secrecy systems. *The Bell system technical journal* **28**(4), 656–715 (1949)
- [57] Shokri, R., Stronati, M., Song, C., Shmatikov, V.: Membership inference attacks against machine learning models. In: *2017 IEEE symposium on security and privacy (SP)*. pp. 3–18. IEEE (2017)
- [58] Shorten, C., Khoshgoftaar, T.M.: survey on image data augmentation for deep learning. *Journal of big data* **6**(1), 1–48 (2019)
- [59] Stefanov, E., Dijk, M.V., Shi, E., Chan, T.H.H., Fletcher, C., Ren, L., Yu, X., Devadas, S.: Path OR M: an extremely simple oblivious R M protocol. *Journal of the CM (J CM)* **65**(4), 1–26 (2018)
- [60] Steinbrecher, S., Köpsell, S.: Modelling unlinkability. In: *International workshop on privacy enhancing technologies*. pp. 32–47. Springer (2003)
- [61] Stock, P., Shilov, I., Mironov, I., Sablayrolles, .: Defending against reconstruction attacks with r\`enyi differential privacy. *arXiv preprint arXiv:2202.07623* (2022)
- [62] Tishby, N., Pereira, F.C., Bialek, W.: The information bottleneck method. *arXiv preprint physics/0004057* (2000)
- [63] Valiant, L.G.: theory of the learnable. *Communications of the CM* **27**(11), 1134–1142 (1984)
- [64] Vershynin, R.: *High-dimensional probability: n introduction with applications in data science*, vol. 47. Cambridge university press (2018)
- [65] Wang, Y.X., Balle, B., Kasiviswanathan, S.P.: Subsampled rényi differential privacy and analytical moments accountant. In: *The 22nd International Conference on rtificial Intelligence and Statistics*. pp. 1226–1235. PMLR (2019)
- [66] Xiao, X., Tao, Y.: Output perturbation with query relaxation. *Proceedings of the VLDB Endowment* **1**(1), 857–869 (2008)
- [67] Xu, ., Raginsky, M.: Information-theoretic analysis of generalization capability of learning algorithms. *dvances in Neural Information Processing Systems* **30** (2017)
- [68] Yeom, S., Giacomelli, I., Fredrikson, M., Jha, S.: Privacy risk in machine learning: nalyzing the connection to overfitting. In: *2018 IEEE 31st computer security foundations symposium (CSF)*. pp. 268–282. IEEE (2018)

ppendix

Proof of Theorem 1 and Corollary 1

To start, we need the following lemmas.

Lemma 1. *Given any f -divergence $\mathcal{D}_f(\cdot\|\cdot)$, and three Bernoulli distributions $\mathbf{1}_a$, $\mathbf{1}_b$ and $\mathbf{1}_c$ of parameters a , b and c , respectively, where $0 \leq a \leq b \leq c \leq 1$. Then, $\mathcal{D}_f(\mathbf{1}_a\|\mathbf{1}_b) \leq \mathcal{D}_f(\mathbf{1}_a\|\mathbf{1}_c)$.*

Proof. By the definition, $g(x) = \mathcal{D}_f(\mathbf{1}_a\|\mathbf{1}_x) = xf(\frac{a}{x}) + (1-x)f(\frac{1-a}{1-x})$ and we want to show $g(x)$ is non-decreasing for $x \geq a$. With some calculation, $g'(x) = (f(\frac{a}{x}) - \frac{a}{x}f'(\frac{a}{x})) - (f(\frac{1-a}{1-x}) - \frac{1-a}{1-x}f'(\frac{1-a}{1-x}))$. It is noted that $\frac{a}{x} \leq \frac{1-a}{1-x}$ for $x \geq a$. Thus, to show $g'(x) \geq 0$ for $x \geq a$, it suffices to show $t(y) = f(y) - yf'(y)$ is non-increasing with respect to $y \in [0, 1]$. On the other hand, $t'(y) = f'(y) - f'(y) - yf''(y) \leq 0$ due to the convex assumption of f . Therefore, the claim holds. \square

Lemma 2 (Data Processing Inequality [54]). *Consider a channel that produces Z given Y based on the law described as a conditional distribution $\mathbb{P}_{Z|Y}$. If \mathbb{P}_Z is the distribution of Z when Y is generated by \mathbb{P}_Y , and \mathbb{Q}_Z is the distribution of Z when Y is generated by \mathbb{Q}_Y , then for any f -divergence \mathcal{D}_f ,*

$$\mathcal{D}_f(\mathbb{P}_Z\|\mathbb{Q}_Z) \leq \mathcal{D}_f(\mathbb{P}_Y\|\mathbb{Q}_Y).$$

Now, we return to prove Theorem 1. First, we have the observation that for a random variable $X' \in \mathcal{X}$ in an arbitrary distribution but independent of X , $\delta_o^p \leq \Pr_{X' \perp X}(\rho(X', X) = 1)$, since δ_o^p is the minimum failure probability achieved by optimal *a priori* estimation. Here, $a \perp b$ represents that a is independent of b . Let the indicator be a function that for two random variables a and b , $\mathbf{1}(a, b) = 1$ if $\rho(a, b) = 1$, otherwise 0.

Apply Lemma 1 and Lemma 2, where we view $\mathbf{1}(\cdot, \cdot)$ as a post-processing on (X, \tilde{X}) and (X, X') , respectively, we have that

$$\mathcal{D}_f(\mathbf{1}_\delta\|\mathbf{1}_{\delta_o^p}) \leq \mathcal{D}_f(\mathbf{1}(X, \tilde{X})\|\mathbf{1}(X, X')) \leq \mathcal{D}_f(\mathbb{P}_{X, \tilde{X}}\|\mathbb{P}_{X, X'}) = \mathcal{D}_f(\mathbb{P}_{X, \tilde{X}}\|\mathbb{P}_X \otimes \mathbb{P}_{X'}).$$

On the other hand, we know $X \rightarrow \mathcal{M}(X) \rightarrow \tilde{X}$ forms a Markov chain, where the adversary's estimation \tilde{X} is dependent on observation $\mathcal{M}(X)$. Let the adversary's strategy be some operator g_{adv} where $\tilde{X} = g_{adv}(\mathcal{M}(X))$. Therefore, we can apply the data processing inequality again, where

$$\mathcal{D}_f(\mathbb{P}_{X, \tilde{X}}\|\mathbb{P}_{X, X'}) \leq \mathcal{D}_f(\mathbb{P}_{X, \mathcal{M}(X)}\|\mathbb{P}_{X, W}) = \mathcal{D}_f(\mathbb{P}_{X, \mathcal{M}(X)}\|\mathbb{P}_X \otimes \mathbb{P}_W).$$

Here, $X' = g_{adv}(W)$ and W is still independent of X . Since the above inequalities hold for arbitrarily distributed X' once it is independent of X , W could also be an arbitrary random variable on the same support domain as $\mathcal{M}(X)$ and independent of X . Therefore,

$$\mathcal{D}_f(\mathbf{1}_\delta\|\mathbf{1}_{\delta_o^p}) \leq \int_{\mathbb{P}_W} \mathcal{D}_f(\mathbb{P}_{X, \mathcal{M}(X)}\|\mathbb{P}_X \otimes \mathbb{P}_W) = \int_{\mathbb{P}_W} \mathcal{D}_f(\mathbb{P}_{\mathcal{M}(X)|X}\|\mathbb{P}_W|\mathbb{P}_X).$$

Here, we use $|\mathbb{P}_X$ to denote that it is conditional on X in a distribution \mathbb{P}_X . In particular, if we select \mathbb{P}_W to be the distribution of $\mathcal{M}(X)$, and take \mathcal{D}_f to be KL-divergence, we have $\mathcal{D}_{KL}(\mathbf{1}_\delta\|\mathbf{1}_{\delta_o^p}) \leq \text{MI}(X; \mathcal{M}(X))$.

B Proof of Theorem 2

The first part of Theorem 2 is a simple corollary of Theorem 1. When we separately consider the adversary's inference on each x_i , take x_i as the X in the proof of Theorem 1 and a similar reasoning will produce the upper bound on each posterior success rate $(1 - \delta_i)$, and the claim follows.

Consider the second part of Theorem 2, where we assume the x_i is i.i.d. generated and \mathcal{M} is some symmetric mechanism. We apply Theorem 1 to upper bound the success rate that the adversary can recover at least j many x_i s out of the total X , as some generic inference on the whole data X . Due to the i.i.d. assumption, we know the optimal *a priori* success rate to recover a single x_i equals $1 - \bar{\delta}_o^\rho = \sup_{\tilde{x}} \Pr(\rho(\tilde{x}, x) = 1)$. Consequently, the optimal *a priori* success rate to recover at least j out of n samples is $\bar{\delta}_o^{j, \rho} = \sum_{l=j}^n \binom{n}{l} (\bar{\delta}_o^\rho)^{n-l} \cdot (1 - \bar{\delta}_o^\rho)^l$. Thus, applying Theorem 1, we have that the optimal posterior success rate on recovering at least j samples $(1 - \tilde{\delta}^j)$ satisfies

$$\mathcal{D}_{KL}(\mathbf{1}_{\tilde{\delta}^j} \| \mathbf{1}_{\bar{\delta}_o^{j, \rho}}) \leq \text{MI}(X; \mathcal{M}(X)).$$

On the other hand, due to the i.i.d. and the symmetric assumption of \mathcal{M} , we know $\text{MI}(x_i; \mathcal{M}(X))$ is identical, for any $i \in [1 : n]$. The optimal posterior success rate to recover every single x_i is the same as $(1 - \delta)$. Therefore, let the indicator $\mathbf{1}_{x_i}$ represent whether x_i is successfully recovered. The expected number of samples that an adversary can recover satisfies

$$n(1 - \delta) = \mathbb{E}[\sum_{i=1}^n \mathbf{1}_{x_i}] \leq \sum_{j=1}^n (1 - \tilde{\delta}^j).$$

Here, we use Fubini's theorem on the expectation calculation, where for a random variable Y defined over \mathbb{Z}^+ ,

$$\mathbb{E}[Y] = \sum_{j=1}^{+\infty} \Pr(Y \geq j).$$

Thus, we obtain the desired lower bound of δ expressed by $\tilde{\delta}^j$ for $j = 1, 2, \dots, n$.

C Proof of Theorem 3

$\text{MI}(X; \mathcal{M}(X) + B)$

$$\begin{aligned} &= \int \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0)+B} \| \mathbb{P}_B) \mathbb{P}(X = X_0) dX_0 - \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X)+B} \| \mathbb{P}_B) \\ &= \int \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0)+B} \| \mathbb{P}_B) \mathbb{P}(X = X_0) dX_0 - (\mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X)+B} \| \mathbb{P}_{\text{Gau}(\mathcal{M}(X)+B)}) + \mathcal{D}_{KL}(\mathbb{P}_{\text{Gau}(\mathcal{M}(X)+B)} \| \mathbb{P}_B)) \\ &\leq \int \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0)+B} \| \mathbb{P}_B) \mathbb{P}(X = X_0) dX_0 - \mathcal{D}_{KL}(\mathbb{P}_{\text{Gau}(\mathcal{M}(X)+B)} \| \mathbb{P}_B). \end{aligned} \tag{15}$$

Given the definition of mutual information, we first apply the results of Gaussian approximation [53], where $\text{Gau}(\cdot)$ represents a (multivariate) Gaussian variable with the same mean and (co)variance as those of \cdot . Then, we drop a negative term to obtain the final inequality of (15). Next, focusing on the integral (first) term in (15),

$$\mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0)+B} \| \mathbb{P}_B) = \frac{1}{2} \cdot (\mathcal{M}(X_0))^T \Sigma^{-1} (\mathcal{M}(X_0)),$$

and therefore

$$\int \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0)+B} \| \mathbb{P}_B) \mathbb{P}(X = X_0) dX_0 = \frac{1}{2} \cdot \mathbb{E}_X [(\mathcal{M}(X))^T \Sigma^{-1} (\mathcal{M}(X))].$$

As for the second term in the last equation of (15), we have the covariance of $\mathcal{M}(X)$ equals $\Sigma_{\mathcal{M}(X)} = \mathbb{E}_X [(\mathcal{M}(X) - \mathbb{E}[\mathcal{M}(X)])(\mathcal{M}(X) - \mathbb{E}[\mathcal{M}(X)])^T]$, while the mean $\mu_{\mathcal{M}(X)} = \mathbb{E}[\mathcal{M}(X)]$. The KL divergence between two multivariate Gaussians has a closed form, where $\mathcal{D}_{KL}(\mathbb{P}_{\text{Gau}(\mathcal{M}(X)+B)} \| \mathbb{P}_B)$ equals

$$\begin{aligned} \mathcal{D}_{KL}(\mathbb{P}_{\text{Gau}(\mathcal{M}(X)+B)} \| \mathbb{P}_B) &= \frac{1}{2} \cdot (\text{Trace}(\Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1}) \\ &\quad + \mathbb{E}_X [(\mathcal{M}(X))^T \Sigma^{-1} \mathbb{E}_X [\mathcal{M}(X)] - \log \det(\mathbf{I}_d + \Sigma_{\mathcal{M}(X)} \Sigma^{-1})]). \end{aligned} \tag{16}$$

On the other hand, note that

$$\begin{aligned}
& \mathbb{E}_X [(\mathcal{M}(X))^T \Sigma^{-1} (\mathcal{M}(X))] - \mathbb{E}_X [\mathcal{M}(X)]^T \Sigma^{-1} \mathbb{E}_X [\mathcal{M}(X)] - \text{Trace}(\Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1}) \\
&= \text{Trace}(\mathbb{E}[\mathcal{M}(X) - \mathbb{E}[\mathcal{M}(X)]] \cdot \Sigma^{-1} \cdot \mathbb{E}[\mathcal{M}(X) - \mathbb{E}[\mathcal{M}(X)]]^T) - \text{Trace}(\Sigma_{\mathcal{M}(X)} \Sigma^{-1}) \\
&= \text{Trace}(\mathbb{E}[\mathcal{M}(X) - \mathbb{E}[\mathcal{M}(X)]] \cdot \mathbb{E}[\mathcal{M}(X) - \mathbb{E}[\mathcal{M}(X)]]^T \cdot \Sigma^{-1} - \Sigma_{\mathcal{M}(X)} \Sigma^{-1}) \\
&= \text{Trace}(\Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1} - \Sigma_{\mathcal{M}(X)} \Sigma^{-1}) = 0.
\end{aligned} \tag{17}$$

In (17), we use the following facts that for two arbitrary vectors $v_1, v_2 \in \mathbb{R}^d$, $(v_1)^T v_2 = \text{Trace}(v_1 \cdot (v_2)^T)$, and for two arbitrary matrices $\Sigma_1, \Sigma_2 \in \mathbb{R}^{d \times d}$, $\text{Trace}(\Sigma_1 \Sigma_2) = \text{Trace}(\Sigma_2 \Sigma_1)$. Therefore, putting it all together, we have a simplified form of the right hand of (15), where

$$\text{MI}(X; \mathcal{M}(X) + B) \leq \frac{\log \det(\mathbf{I}_d + \Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1})}{2}. \tag{18}$$

In the following, we turn to find a construction of Σ to control the right side of (18). Assume the SVD of $\Sigma_{\mathcal{M}(X)} = U \Lambda U^T$, where Λ represents a diagonal matrix whose diagonal elements are the eigenvalues $\lambda_1, \dots, \lambda_d$ of $\Sigma_{\mathcal{M}(X)}$. We take $\Sigma = U \Lambda^{-1} U^T$, where Λ^{-1} is a diagonal matrix of elements $\lambda^{-1}_1, \dots, \lambda^{-1}_d$. Using the fact that $\log(1+x) \leq x$, we can further upper bound the right side of (18) and obtain,

$$\text{MI}(X; \mathcal{M}(X) + B) \leq \frac{1}{2} \cdot \sum_{j=1}^d \frac{\lambda_j}{\lambda^{-1}_j}.$$

On the other hand, note that $\mathbb{E}[\|\cdot\|_2^2] = \sum_{j=1}^d \lambda^{-1}_j$. Therefore, to find the minimal noise, it is equivalent to solving the following optimization problem.

$$\text{mi} \sum_{j=1}^d \lambda^{-1}_j, \quad \text{s.t.} \sum_{j=1}^d \frac{\lambda_j}{\lambda^{-1}_j} \leq 2v, \lambda^{-1}_j \geq 0, j = 1, 2, \dots, d. \tag{19}$$

By Hölder's inequality on

$$\sum_{j=1}^d \lambda^{-1}_j \cdot \sum_{j=1}^d \frac{\lambda_j}{\lambda^{-1}_j} \leq 2v \cdot \sum_{j=1}^d \lambda^{-1}_j,$$

the optimal solution of (19) is $\lambda^{-1}_j = (2v)^{-1} \cdot \sqrt{\lambda_j} \cdot (\sum_{l=1}^d \sqrt{\lambda_l})$. As a consequence, the expected l_2 norm of the noise $\mathbb{E}[\|\cdot\|_2]$ is upper bounded by $\sqrt{(2v)^{-1} \cdot (\sum_{j=1}^d \sqrt{\lambda_j})}$.

D Proof of Theorem 4

We first consider two arbitrary positive semi-definite Σ and $\hat{\Sigma}$, where $\Delta \Sigma = \Sigma - \hat{\Sigma}$ denotes the difference. Let their SVD decompositions be $\Sigma = U \Lambda U^T = \sum_{j=1}^d \lambda_j u_j u_j^T$ and $\hat{\Sigma} = \hat{U} \hat{\Lambda} \hat{U}^T = \sum_{j=1}^d \hat{\lambda}_j \hat{u}_j \hat{u}_j^T$. Before the proof of Theorem 4, we need the following lemmas.

Lemma 3 (Weyl's theorem). *For two $d \times d$ positive semi-definite matrices Σ and $\hat{\Sigma}$, let their eigenvalues be $\lambda_1 \geq \dots \geq \lambda_d$ and $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_d$, both in a non-increasing order. Then, for any $1 \leq j \leq d$,*

$$|\lambda_j - \hat{\lambda}_j| \leq \|\Sigma - \hat{\Sigma}\|_2.$$

Lemma 4. *Assume that $Y = \{y_1, y_2, \dots, y_m\}$ is i.i.d. selected from some distribution \mathcal{P}_y whose support domain is \mathcal{B}_r^d and the covariance matrix is Σ . Let $\hat{\Sigma}$ be the empirical covariance of Y , then we have that with probability at least $(1 - \gamma)$, $\|\Delta \Sigma = \hat{\Sigma} - \Sigma\|$ is upper bounded by*

$$\kappa r \left(\max \left\{ \sqrt{\frac{d + \log(4/\gamma)}{m}}, \frac{d + \log(4/\gamma)}{m} \right\} + \sqrt{\frac{d \log(4/\gamma)}{m}} \right).$$

Proof. We will use some notions and the high-dimensional statistics results from [64]. First, we have the following w.r.t. the empirical covariance. Let $\hat{y} = \frac{1}{m} \cdot \sum_{i=1}^m y_i$, and the empirical covariance is defined as $\hat{\Sigma} = \frac{1}{m} \sum_{i=1}^m (y_i - \hat{y})(y_i - \hat{y})^T$. Then, for any $\mu \in \mathbb{R}^d$, we have

$$\frac{1}{m} \cdot \sum_{i=1}^m (y_i - \mu)(y_i - \mu)^T = \hat{\Sigma} + (\mu - \hat{y})(\mu - \hat{y})^T.$$

Second, since $y \sim y$ is uniformly bounded and we know y is some r -subGaussian random vector (see Def 3.4.1 in [64]). Now, we select $\mu = \mathbb{E}[y]$, and for zero-mean r -sub-Gaussian random vectors, from Thm 4.7.1 and Ex 4.7.3 in [64], we have that there exists some universal constant κ' such that with probability at least $(1 - \gamma')$,

$$\|\Sigma - \frac{1}{m} \cdot \sum_{i=1}^m (y_i - \mu)(y_i - \mu)^T\| \leq \kappa' r \cdot \max\left\{\sqrt{\frac{d \log(2/\gamma')}{m}}, \frac{d \log(2/\gamma')}{m}\right\}.$$

On the other hand, we provide a high-probability concentration on the mean estimation $\|\mu - \hat{y}\|$. We use the results on *norm-subGaussian* [43], where

$$\Pr(\|\mu - \hat{y}\| \geq t) \leq 2e^{-\frac{mt^2}{2dr^2}}.$$

Therefore, with probability at least $(1 - \gamma')$, $\|(\mathbb{E}[y] - \hat{y})(\mathbb{E}[y] - \hat{y})^T\| \leq \sqrt{\frac{-2dr^2 \log(\gamma'/2)}{m}}$. Now, let $\gamma' = \frac{\gamma}{2}$, and put things together, we have that there exists some universal constant κ ,

$$\|\hat{\Sigma} - \Sigma\| \leq \kappa r \left(\max\left\{\sqrt{\frac{d + \log(4/\gamma)}{m}}, \frac{d + \log(4/\gamma)}{m}\right\} + \sqrt{\frac{d \log(4/\gamma)}{m}} \right).$$

□

□

Lemma 5. *Let $\hat{\Sigma} = \Sigma + \Delta\Sigma = \hat{U}\hat{\Lambda}\hat{U}^T$ for $\hat{\Lambda} = \{\hat{\lambda}_1, \dots, \hat{\lambda}_d\}$. If $\|\Delta\Sigma\| \leq c$, then for some particular $j \in [1 : d]$, if the gap $\min_{l \neq j} |\lambda_j - \lambda_l| \geq \tau$, then*

$$\sum_{l \neq j} \lambda_l (u_l^T \cdot \hat{u}_j)^2 \leq 4c + 4\lambda_j \min\{1, (d-1)c^2/\tau^2\}.$$

Proof. For $\hat{\Sigma}$, we can rewrite its j -th eigenvector as the following

$$\hat{\Sigma} \hat{u}_j = (\Sigma + \Delta\Sigma)(u_j + \Delta u_j) = \hat{\lambda}_j \hat{u}_j = (\lambda_j + \Delta\lambda_j)(u_j + \Delta u_j). \quad (20)$$

From (20), since $\Sigma u_j = \lambda_j u_j$, we have that

$$\Sigma \Delta u_j = (\lambda_j + \Delta\lambda_j)(u_j + \Delta u_j) - \lambda_j u_j - \Delta\Sigma(u_j + \Delta u_j). \quad (21)$$

We multiply $(\Delta u_j)^T$ on both sides of (21) and have that

$$(\Delta u_j)^T \Sigma \Delta u_j = \lambda_j \|\Delta u_j\|^2 - (\Delta u_j)^T (\Delta\Sigma - \Delta\lambda_j \cdot \mathbf{I}_d) \cdot (u_j + \Delta u_j). \quad (22)$$

Now, we apply the equivalent expression $\Sigma = \sum_{l=1}^d \lambda_l u_l^T u_l$ in (22) and have

$$\begin{aligned} \lambda_j (u_j^T \cdot \Delta u_j)^2 + \sum_{l \neq j} \lambda_l (u_l^T \cdot \Delta u_j)^2 &\leq \lambda_j \|\Delta u_j\|^2 + 2\|\Delta u_j\| \|\Delta\Sigma\| \|\hat{u}_j\| \\ &\leq \lambda_j \|\Delta u_j\|^2 + 4\|\Delta\Sigma\|. \end{aligned} \quad (23)$$

Here, we use the following results. First, from Weyl's theorem on eigenvalues, $|\Delta\lambda_j| = |\lambda_j - \hat{\lambda}_j| \leq \|\Delta\Sigma\|$ for any j . Second, $\|\Delta u_j\| = \|u_j - \hat{u}_j\| \leq 2$.

Now, we consider multiplying u_l^T on both sides of 21 and we have the following,

$$\lambda_l(u_l^T \cdot \Delta u_j) = \lambda_j(u_l^T \cdot \Delta u_j) - u_l^T(\Delta \Sigma - \Delta \lambda_j \cdot \mathbf{I}_d) \cdot (\hat{u}_j).$$

Rearrange the above and take absolute value of both sides, then we have for any $l \neq j$,

$$|u_l^T \cdot \Delta u_j| \leq 2\|\Delta \Sigma\|/|\lambda_l - \lambda_j| \leq 2\|\Delta \Sigma\|/\tau.$$

With the above preparation, we can proceed to upper bound $\|\Delta u_j\|^2 - (u_j^T \cdot \Delta u_j)^2$ as follows. Since $u_{[1:d]}$ are a set of orthogonal basis of \mathbb{R}^d , for a given j ,

$$\|\Delta u_j\|^2 - (u_j^T \cdot \Delta u_j)^2 = \sum_{l \neq j} (u_l^T \cdot \Delta u_j)^2 \leq 4(d-1)\|\Delta \Sigma\|^2/\tau^2. \quad (24)$$

Now we put (24) back into (23) and obtain that

$$\sum_{l \neq j} \lambda_l(u_l^T \cdot \Delta u_j)^2 \leq 4(\lambda_j(d-1)\|\Delta \Sigma\|^2/\tau^2 + \|\Delta \Sigma\|).$$

It is noted that

$$\sum_{l \neq j} \lambda_l(u_l^T \cdot \hat{u}_j)^2 = \sum_{l \neq j} \lambda_l(u_l^T \cdot \Delta u_j)^2,$$

since $\langle u_l, u_j \rangle = 0$ for $l \neq j$, and on the other hand we know $\lambda_j\|\Delta u_j\|^2 \leq 4\lambda_j$, and therefore we obtain the upper bound claimed straightforwardly. \square

Back to the proof of Theorem 4. First, from Theorem 3, we have

$$\text{MI}(X; \mathcal{M}(X) + \Sigma) \leq \frac{1}{2} \cdot \log \det(\mathbf{I}_d + \Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1}) \leq \frac{1}{2} \cdot \text{Tr}(\Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1}). \quad (25)$$

Now, let $\Sigma_{\mathcal{M}(X)} = U^T \Lambda U$ and $\Sigma^{-1} = \hat{U}^T \hat{\Lambda} \hat{U}$. Similarly, we follow the notations used in Lemma 5 and let $\hat{u}_j = u_j + \Delta u_j$. Then, for the first scenario defined in Algorithm 1, let $j_0 = \arg \max_j \lambda_j \geq c_0$, $\min_{1 \leq j \leq j_0, 1 \leq l \leq d} |\hat{\lambda}_j - \lambda_l| > c_1$ and $\lambda_{j_0} = \frac{c_2}{\sqrt{\lambda_{j_0} + c_3} \cdot (\sum_{j=1}^d \sqrt{\lambda_j + c_3})}$, for parameters c_0, c_1, c_2 and c_3 to be determined, and we assume that $\|\Sigma - \hat{\Sigma}\| \leq c$. Then, we expand the term $\text{Tr}(\Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1})$ as follows,

$$\begin{aligned} \text{Tr}(\Sigma_{\mathcal{M}(X)} \cdot \Sigma^{-1}) &= \text{Tr} \left(\left(\sum_{j=1}^d \lambda_j u_j u_j^T \right) \cdot \left(\sum_{l=1}^d \lambda_{-l} \hat{u}_l \hat{u}_l^T \right) \right) \\ &= \text{Tr} \left(\sum_{j=1}^d \sum_{l=1}^d \lambda_j \lambda_{-l} \cdot (u_j^T \hat{u}_l)^2 \right) = \text{Tr} \left(\sum_{j=1}^d \sum_{l=1}^d \lambda_j \lambda_{-l} \cdot (u_j^T (u_l + \Delta u_l))^2 \right) \\ &\leq \sum_{j=1}^d \lambda_j \lambda_{-j} + \sum_{j=1}^d \sum_{l \neq j} \lambda_j \lambda_{-l} (u_j^T \Delta u_l)^2 \\ &\leq \sum_{j=1}^d \frac{c_2 \lambda_j}{\sqrt{\hat{\lambda}_j + c_3} \cdot (\sum_{j=1}^d \sqrt{\hat{\lambda}_j + c_3})} + \sum_{j=1}^d \frac{c_2}{dc_3} (4c + 4\lambda_j \min\{1, (d-1)c^2/\tau_j^2\}) \\ &\leq \sum_{j=1}^d \frac{c_2 \lambda_j}{\sqrt{\hat{\lambda}_j + c_3} \cdot (\sum_{j=1}^d \sqrt{\hat{\lambda}_j + c_3})} + \frac{j_0(4c + 16dc^2r^2/(c_1 - 2c)^2)c_2 + (d - j_0)(4c + 4c_0)c_2}{dc_3}. \end{aligned} \quad (26)$$

Here, in the second inequality, we use the fact that $\lambda_{\cdot,j} \leq \frac{c_2}{dc_3}$. In the third inequality, we use the fact that $\lambda_j \leq 4r^2$. For $\tau_j = \min_{1 \leq l \leq d} |\lambda_j - \lambda_l|$, via Weyl's Lemma (Lemma 3), $|\lambda_j - \hat{\lambda}_j| \leq c$ and thus for $j \in [1 : j_0]$, $\tau_j \geq c_1 - 2c$, and for any $c_3 \geq c$ we have

$$\sum_{j=1}^d \frac{c_2 \lambda_j}{\sqrt{\hat{\lambda}_j + c_3} \cdot (\sum_{j=1}^d \sqrt{\hat{\lambda}_j + c_3})} \leq c_2.$$

Now, we select $c_0 = c$, $c_1 = r\sqrt{dc} + 2c$, $c_2 = 2v$, $c_3 = 10cv/\beta$, and we can upper bound the right hand of (26) as $2(v + \beta)$. Thus, by (25), the objective $\text{MI}(X; \mathcal{M}(X) + B) \leq v + \beta$.

The second case in Algorithm 1 is straightforward to analyze, where

$$\text{Tr}(\Sigma \cdot \Sigma_B^{-1}) = \frac{c_2 \sum_{j=1}^d \lambda_j}{\sum_{j=1}^d \hat{\lambda}_j + dc} \leq c_2.$$

Here, we still apply Weyl's lemma where when $\|\hat{\Sigma} - \Sigma\|_2 \leq c$, then $|\text{Tr}(\hat{\Sigma}) - \text{Tr}(\Sigma)| \leq dc$.

E Proof of Theorem 5

First, regarding $\text{MI}(X; \mathcal{M}(X, \theta) + B)$ we have the following observation,

$$\begin{aligned} & \text{MI}(X; \mathcal{M}(X, \theta) + B) \\ &= \sum_{X_0} \Pr(X = X_0) \cdot \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0, \theta) + B} \| \mathbb{P}_{\mathcal{M}(X, \theta) + B}). \end{aligned} \quad (27)$$

Based on our assumptions, $\mathcal{M}(X_0, \theta) + B$ is distributed as a Gaussian mixture, where for some Gaussian $B \sim \mathcal{N}(\mathbf{0}, \sigma^2 \cdot \mathbf{I}_d)$, $\mathbb{P}_{\mathcal{M}(X_0, \theta) + B}$ can be equivalently written as

$$\mathbb{P}_{\mathcal{M}(X_0, \theta) + B} = \sum_{j=1}^{|J|} \frac{1}{|J|} \cdot \mathcal{N}(\mathcal{M}(X_0, \theta_j), \sigma^2 \cdot \mathbf{I}_d).$$

Similarly, we can write $\mathbb{P}_{\mathcal{M}(X, \theta) + B}$ in the following equivalent form,

$$\mathbb{P}_{\mathcal{M}(X, \theta) + B} = \sum_{X'_0} \sum_{j=1}^{|J|} \frac{1}{|J|} \cdot \mathbb{P}_{X'_0} \cdot \mathcal{N}(\mathcal{M}(X'_0, \theta_j), \sigma^2 \cdot \mathbf{I}_d).$$

Applying the convexity of KL-divergence, we have that

$$\begin{aligned} & \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0, \theta) + B} \| \mathbb{P}_{\mathcal{M}(X, \theta) + B}) \leq \sum_{X'_0} \mathbb{P}_{X'_0} \cdot \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0, \theta) + B} \| \mathbb{P}_{\mathcal{M}(X'_0, \theta) + B}) \\ & \leq \sum_{X'_0} \sum_{j=1}^{|J|} \frac{1}{|J|} \cdot \mathbb{P}_{X'_0} \cdot \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0, \theta_j) + B} \| \mathbb{P}_{\mathcal{M}(X'_0, \theta_{\pi(X'_0, j)}) + B}) \\ & = \sum_{X'_0} \sum_{j=1}^{|J|} \frac{1}{|J|} \cdot \mathbb{P}_{X'_0} \cdot \frac{\|\mathcal{M}(X_0, \theta_j) - \mathcal{M}(X'_0, \theta_{\pi(X'_0, j)})\|^2}{2\sigma^2}, \end{aligned} \quad (28)$$

for any permutation $\pi_{X'_0}$ on $[1 : |J|]$ for each X'_0 . Here, we use $\mathbb{P}_{X'_0} = \Pr(X = X'_0)$. Therefore, if we let $\psi(X_0, X'_0) = \mathbf{d}_{\pi}(\mathbf{y}(X_0), \mathbf{y}(X'_0))$, where

$$\mathbf{y}(X_0) = (\mathcal{M}(X_0, \theta_1), \dots, \mathcal{M}(X_0, \theta_{|J|})), \mathbf{y}(X'_0) = (\mathcal{M}(X'_0, \theta_1), \dots, \mathcal{M}(X'_0, \theta_{|J|})),$$

then we have a simplified upper bound on $\text{MI}(X; \mathcal{M}(X, \theta) + B)$ as follows,

$$\text{MI}(X; \mathcal{M}(X, \theta) + B) = \mathbb{E}_{X_0} \mathcal{D}_{KL}(\mathbb{P}_{\mathcal{M}(X_0, \theta) + B} \| \mathbb{P}_{\mathcal{M}(X, \theta) + B}) \leq \mathbb{E}_{X_0, X'_0} \frac{\psi(X_0, X'_0)}{2\sigma^2}. \quad (29)$$

Therefore, to derive the upper bound on the objective mutual information, it suffices to estimate the expectation $\mathbb{E}_{X_0, X'_0} \frac{\psi(X_0, X'_0)}{2\sigma^2}$, which is the mean of the permutation distance between the evaluations of $\mathcal{M}(\cdot, \theta)$ on two randomly selected input datasets X_0 and X'_0 from distribution \mathcal{P} . On one hand, we know $\psi(X_0, X'_0) \leq 4r^2$ due to the bounded output assumption for any X_0 and X'_0 . In the following, we prove the following lemma,

Lemma 6. *Let $\tau = \{\theta_{(1)}, \theta_{(2)}, \dots, \theta_{(\tau)}\}$ be a random τ -subset from \mathcal{P} for some τ that can divide $|\mathcal{P}|$. We define $\psi(X_0, X'_0, \tau)$ to be permutation distance between $(\mathcal{M}(X_0, \theta_{(1)}), \dots, \mathcal{M}(X_0, \theta_{(\tau)}))$ and $(\mathcal{M}(X'_0, \theta_{(1)}), \dots, \mathcal{M}(X'_0, \theta_{(\tau)}))$ restricted to the evaluations on τ . Then,*

$$\mathbb{E}_{\tau} \psi(X_0, X'_0, \tau) \geq \psi(X_0, X'_0).$$

Proof. We consider a random permutation π on $[1 : |\mathcal{P}|]$ and we split

$$(\mathcal{M}(X_0, \theta_{\pi(1)}), \dots, \mathcal{M}(X_0, \theta_{\pi(|\mathcal{P}|)}))$$

into $|\mathcal{P}|/\tau$ segments in order. For the j -th segment, we consider the minimal permutation distance between the two sets $(\mathcal{M}(X_0, \theta_{\pi((j-1)\tau+1)}), \dots, \mathcal{M}(X_0, \theta_{\pi(j\tau)}))$, and $(\mathcal{M}(X'_0, \theta_{\pi((j-1)\tau+1)}), \dots, \mathcal{M}(X'_0, \theta_{\pi(j\tau)}))$, denoted by $\psi_j(X_0, X'_0, \pi)$. Since $\psi(X_0, X'_0)$ is the global minimum, it is clear that

$$\psi(X_0, X'_0) \leq \sum_{j=1}^{|\mathcal{P}|/\tau} \frac{\psi_j(X_0, X'_0, \pi)}{|\mathcal{P}|/\tau}.$$

When we take the expectation on both sides, we have that

$$\psi(X_0, X'_0) \leq \mathbb{E}_{\tau} \psi(X_0, X'_0, \tau).$$

□

□

With Lemma 6, we can proceed to show a high-probability mean estimation of $\mathbb{E}_{X_0, X'_0} \mathbb{E}_q [\psi(X_0, X'_0, q)]$. By Hoeffding's inequality, since $\psi(X_0, X'_0, q) \leq 4r^2$, and let $\bar{\psi}$ be the empirical mean of m i.i.d. sampling on $\psi(X_0, X'_0, q)$, then for any $t > 0$,

$$\Pr(\mathbb{E}_{X_0, X'_0} \mathbb{E}_q [\psi(X_0, X'_0, q)] > \bar{\psi} + t) \leq e^{-\frac{2mt^2}{16r^2}}.$$

Thus, when $m \geq \frac{-8r^2 \ln \gamma}{c^2}$, with confidence $(1 - \gamma)$, one can ensure that

$$\mathbb{E}_{X_0, X'_0} \psi(X_0, X'_0) \leq \bar{\psi} + c.$$

As a consequence, by selecting $B \sim \mathcal{N}(0, \sigma^2 \cdot \mathbf{I}_d)$, for $\sigma^2 = (\bar{\psi} + c)/(2v)$, we have that $\text{MI}(X; \mathcal{M}(X) + B) \leq v$ with confidence $(1 - \gamma)$.

F Proof of Theorem 6

The proof of Theorem 6 is essentially a generalization of Lemma 6. We rewrite the expression of $\text{MI}(X; \mathcal{M}(X))$ as follows,

$$\frac{1}{N} \sum_{i=1}^N \mathcal{D}_{KL} \left(\sum_{l=1}^{|\mathcal{P}|} \frac{\mathbb{P}_{\mathcal{M}(X_i, \theta_l) + (X_i, \theta_l)}}{|\mathcal{P}|} \parallel \sum_{j=1}^N \sum_{l=1}^{|\mathcal{P}|} \frac{\mathbb{P}_{\mathcal{M}(X_j, \theta_l) + (X_j, \theta_l)}}{N|\mathcal{P}|} \right). \quad (30)$$

Now, we apply Lemma 6 w.r.t subsampling on θ . For a subset τ_3 of τ_3 elements randomly selected from τ_3 , once $\tau_3 \parallel \tau_3$, we have that (30) is upper bounded by,

$$\mathbb{E}_{\tau_3} \frac{1}{N} \sum_{i=1}^N \mathcal{D}_{KL} \left(\sum_{\theta \in \tau_3} \frac{\mathbb{P}_{\mathcal{M}(X_i, \theta) + c} (X_i, \theta)}{\tau_3} \parallel \sum_{j=1}^N \sum_{\theta \in \tau_3} \frac{\mathbb{P}_{\mathcal{M}(X_j, \theta) + c} (X_j, \theta)}{N \tau_3} \right). \quad (31)$$

With a similar idea, for a τ_2 -subset τ_2^* randomly selected from τ_2 , we have that (31) can be upper bounded by

$$\mathbb{E}_{\tau_3, \tau_2^*} \frac{1}{N} \sum_{i=1}^N \mathcal{D}_{KL} \left(\sum_{\theta \in \tau_3} \frac{\mathbb{P}_{\mathcal{M}(X_i, \theta) + c} (X_i, \theta)}{\tau_3} \parallel \sum_{\theta \in \tau_3, X_l \in \tau_2^*} \frac{\mathbb{P}_{\mathcal{M}(X_l, \theta) + c} (X_l, \theta)}{\tau_2 \tau_3} \right). \quad (32)$$

Finally, with another independent sampling on τ_2^* to randomly produce a τ_1 -set τ_1^* , we obtain an unbiased estimation on (32) as

$$\mathbb{E}_{\tau_3, \tau_2^*, \tau_1^*} \frac{1}{\tau_1} \sum_{X_j \in \tau_1^*} \mathcal{D}_{KL} \left(\sum_{\theta \in \tau_3} \frac{\mathbb{P}_{\mathcal{M}(X_j, \theta) + c} (X_j, \theta)}{\tau_3} \parallel \sum_{\theta \in \tau_3, X_l \in \tau_2^*} \frac{\mathbb{P}_{\mathcal{M}(X_l, \theta) + c} (X_l, \theta)}{\tau_2 \tau_3} \right). \quad (33)$$

Now, if we replace each $\mathbb{P}_{\mathcal{M}(X, \theta)}$ by $\mathbb{P}_{\mathcal{M}(X, \theta) + c}$, for an independent Gaussian noise $c \sim \mathcal{N}(0, c \cdot \mathbf{I}_d)$, we have that for arbitrary X_j and mechanism \mathcal{M} , $\mathcal{D}_{KL} \left(\sum_{\theta \in \tau_3} \frac{\mathbb{P}_{\mathcal{M}(X_j, \theta) + c} (X_j, \theta)}{\tau_3} \parallel \sum_{\theta \in \tau_3, X_l \in \tau_2^*} \frac{\mathbb{P}_{\mathcal{M}(X_l, \theta) + c} (X_l, \theta)}{\tau_2 \tau_3} \right)$ is uniformly upper bounded by $2r^2/c$. Therefore, we apply the Bernstein inequality to obtain the concentration control of such a mean estimation from i.i.d. samples, where let μ to be (33) and

$$\Pr(\bar{\phi} - \mu \geq \beta) \leq \exp\left(-\frac{0.5m\beta^2}{(2r^2/c)^2/\tau_1 + \beta/3 \cdot (2r^2/c)}\right),$$

and therefore for a confidence $(1 - \gamma)$, we may select

$$m \geq \frac{2 \log(1/\gamma)}{\beta^2} \cdot ((2r^2/c)^2/\tau_1 + \beta/3 \cdot (2r^2/c)).$$

G Proof of Theorem 7

We adopt the same notations as those used in the proof of Theorem 5. Without loss of generality, we consider two arbitrary randomized mechanisms $\mathcal{M}_1(X, \theta_1)$ and $\mathcal{M}_2(X, \theta_2)$ where θ_1 and θ_2 are independently selected. We construct a joint mechanism $\bar{\mathcal{M}}(X, \bar{\theta}) = (\mathcal{M}_1(X, \theta_1), \mathcal{M}_2(X, \theta_2))$, where $\bar{\theta} = (\theta_1, \theta_2)$, and a joint noise $\bar{B} = (B_1, B_2)$. Therefore, for any given X_0 , the distribution of $\bar{\mathcal{M}}(X_0, \bar{\theta})$ is uniform in $\{\mathcal{M}_1(X_0, \theta_1), \theta_1 \in \tau_1\} \times \{\mathcal{M}_2(X_0, \theta_2), \theta_2 \in \tau_2\}$. With a similar reasoning as (28), we have that

$$\begin{aligned} \text{MI}(X; \bar{\mathcal{M}}(X, \bar{\theta}) + \bar{B}) &= \mathbb{E}_{X_0} \mathcal{D}_{KL}(\mathbb{P}_{\bar{\mathcal{M}}(X_0, \bar{\theta}) + \bar{B}} \parallel \mathbb{P}_{\bar{\mathcal{M}}(X, \bar{\theta}) + \bar{B}}) \\ &\leq \sum_{X_0, X'_0} \mathbb{P}_{X_0} \cdot \mathbb{P}_{X'_0} \cdot \mathcal{D}_{KL}(\mathbb{P}_{\bar{\mathcal{M}}(X_0, \bar{\theta}) + \bar{B}} \parallel \mathbb{P}_{\bar{\mathcal{M}}(X'_0, \bar{\theta}) + \bar{B}}) \\ &\leq \sum_{X_0, X'_0} \sum_{\theta_{1j}, \theta_{2l}} \frac{\mathbb{P}_{X_0} \cdot \mathbb{P}_{X'_0}}{|\tau_1| \cdot |\tau_2|} \cdot \mathcal{D}_{KL}(\mathbb{P}_{\bar{\mathcal{M}}(X_0, \theta_{1j}, \theta_{2l}) + \bar{B}} \parallel \mathbb{P}_{\bar{\mathcal{M}}(X'_0, \theta_{1\pi_1(X'_0, j)}, \theta_{2\pi_2(X'_0, l)}) + \bar{B}}) \\ &\leq \sum_{X_0, X'_0} \sum_{\theta_{1j}, \theta_{2l}} \frac{\mathbb{P}_{X_0} \cdot \mathbb{P}_{X'_0}}{|\tau_1| \cdot |\tau_2|} \cdot \left(\frac{\|\mathcal{M}_1(X_0, \theta_{1j}) - \mathcal{M}_1(X'_0, \theta_{1\pi_1(X'_0, j)})\|^2}{2\sigma_1^2} + \frac{\|\mathcal{M}_2(X_0, \theta_{2l}) - \mathcal{M}_2(X'_0, \theta_{2\pi_2(X'_0, l)})\|^2}{2\sigma_2^2} \right), \end{aligned} \quad (34)$$

for arbitrary permutations $\pi_1(X'_0, \cdot)$ and $\pi_2(X'_0, \cdot)$ on $[1 : |\tau_1|]$ and $[1 : |\tau_2|]$. Here, we assume $B_1 \sim \mathcal{N}(0, \sigma_1^2 \cdot \mathbf{I})$ and $B_2 \sim \mathcal{N}(0, \sigma_2^2 \cdot \mathbf{I})$ and use the following fact of KL-divergence between two multivariate

Gaussians

$$\mathcal{D}_{KL}(\mathbb{P}_{(\mu_1, \mu_2) + (B_1, B_2)} \| \mathbb{P}_{(\mu'_1, \mu'_2) + (B_1, B_2)}) = \frac{1}{2} \left(\frac{\|\mu_1 - \mu'_1\|^2}{\sigma_1^2} + \frac{\|\mu_2 - \mu'_2\|^2}{\sigma_2^2} \right).$$

Therefore, if we let $\psi_1(X_0, X'_0)$ equal the minimum-permutation distance $d_{\pi_1}(\mathbf{y}_1(X_0), \mathbf{y}_1(X'_0))$, where

$$\mathbf{y}_1(X_0) = (\mathcal{M}_1(X_0, \theta_1), \dots, \mathcal{M}_1(X_0, \theta_{1|_{-1}})), \mathbf{y}_1(X'_0) = (\mathcal{M}_1(X'_0, \theta_1), \dots, \mathcal{M}_1(X'_0, \theta_{1|_{-1}})),$$

and we similarly define $\psi_2(X_0, X'_0)$ with π_2 , then we have a simplified upper bound on $\text{MI}(X; \bar{\mathcal{M}}(X, \bar{\theta}) + \bar{B})$ as follows,

$$\text{MI}(X; \bar{\mathcal{M}}(X, \bar{\theta}) + \bar{B}) \leq \mathbb{E}_{X_0, X'_0} \left(\frac{\psi_1(X_0, X'_0)}{2\sigma_1^2} + \frac{\psi_2(X_0, X'_0)}{2\sigma_2^2} \right). \quad (35)$$

Therefore, if we separately run Algorithm 2 on \mathcal{M}_1 and \mathcal{M}_2 and determine σ_1 and σ_2 such that with confidence $(1 - \gamma_1)$ and $(1 - \gamma_2)$,

$$\mathbb{E}_{X_0, X'_0} \frac{\psi_1(X_0, X'_0)}{2\sigma_1^2} \leq v_1, \mathbb{E}_{X_0, X'_0} \frac{\psi_2(X_0, X'_0)}{2\sigma_2^2} \leq v_2,$$

respectively, then by union bound we have $\text{MI}(X; \bar{\mathcal{M}}(X, \bar{\theta}) + \bar{B}) \leq v_1 + v_2$ with confidence at least $1 - (\gamma_1 + \gamma_2)$. The above analysis can be straightforwardly generalized to study any T -composition with independent randomness.

H Proof of Theorem 8

The proof is essentially a simple generalization of that of Theorem 5 and Theorem 7. With a similar reasoning, for any $t \in [1 : T]$, the mutual information $\text{MI}(X; \mathcal{M}_1(X; \bar{\theta}_1) + \dots, \mathcal{M}_t(X; \bar{\theta}_t) + \dots)$ can be upper bounded by

$$\text{MI}(X; \mathcal{M}_1(X; \bar{\theta}_1) + \dots, \mathcal{M}_t(X; \bar{\theta}_t) + \dots) \leq \mathbb{E}_{X_0, X'_0} \sum_{j=1}^t \frac{\psi_j(X_0, X'_0)}{2\sigma_j^2}. \quad (36)$$

Here, $\psi_j(X_0, X'_0) = \|\mathbf{y}_j(X_0) - \mathbf{y}_j(X'_0)\|_2^2$, where $\mathbf{y}_j(X_0) = (\mathcal{M}_j(X_0, \bar{\theta}_{j1}), \mathcal{M}_j(X_0, \bar{\theta}_{j2}), \dots)$, i.e., the concatenation of all possible outputs of \mathcal{M}_j with different selections of randomness seeds θ_j , and $\theta_j \sim \mathcal{N}(0, \sigma_j^2 \cdot \mathbf{I})$. Therefore, the goal of Algorithm 4 is essentially to determine σ_t in the t -th round such that the increment of (36), i.e., $\mathbb{E}_{X_0, X'_0} \frac{\psi_t(X_0, X'_0)}{2\sigma_t^2}$, is no larger than $v_t - v_{t-1}$ determined by the privacy budget prescribed.

The only difference is that in an online setup for adaptive composition, we need to ensure that the estimation in each step is accurate enough (without failure in any one of the steps with high probability). Therefore, for a T -composition, it suffices to ensure that the noise mechanisms determined in the t -th round can ensure current joint mechanism $\text{MI}(X; \mathcal{M}_t(X; \bar{\theta}_t) + \dots) \leq v_t$ with high confidence $(1 - \gamma/T)$ and by the union bound, we can claim that $\text{MI}(X; \mathcal{M}_t(X; \bar{\theta}_t) + \dots) \leq v_t$ for any $t \in [1 : T]$ with confidence at least $(1 - \gamma)$. Applying Theorem 5 for each step, we have the bound claimed.

I Proof of Theorem 9

By Pinsker's inequality, we know the KL-divergence can be lower bounded by the total variation. Thus, in the following, we simply consider the $\mathcal{D}_{TV}(\mathbb{P}_{X, \mathcal{M}(X)} \| \mathbb{P}_X \otimes \mathbb{P}_{\mathcal{M}(X)})$. We prove Theorem 9 by contradiction: if there exists some $y \in \mathcal{Y}$ such that $\mathbb{E}_{Q(y)} \|\cdot\|_2 = o(r\sqrt{d}/\sqrt{v})$, then there exist some \mathbb{P}_X and \mathcal{M} such that $\mathcal{D}_{TV}(\mathbb{P}_{X, \mathcal{M}(X)} \| \mathbb{P}_X \otimes \mathbb{P}_{\mathcal{M}(X)}) > \sqrt{v}$ for $v = o(1)$. If not, we consider the following construction of X

and $\mathcal{M}(X)$. X is randomly selected from a set $\{X_1, X_2\}$ and we consider a mechanism \mathcal{M} which maps X_1 to y and X_2 to y' . Thus, we have the following w.r.t. the total variation.

$$\begin{aligned}
\mathcal{D}_{TV}(\mathbb{P}_{X, \mathcal{M}(X)} \| \mathbb{P}_X \otimes \mathbb{P}_{\mathcal{M}(X)}) &= \mathbb{E}_{X_0 \sim \cdot} \mathcal{D}_{TV}(\mathbb{P}_{\mathcal{M}(X_0)} \| \mathbb{P}_{\mathcal{M}(X)}) \\
&= \frac{1}{4} \cdot \int_z (|\mathbb{P}_y(z) - \frac{1}{2} \cdot (\mathbb{P}_y(z) + \mathbb{P}_{y'}(z))| + |\mathbb{P}_{y'}(z) - \frac{1}{2} \cdot (\mathbb{P}_y(z) + \mathbb{P}_{y'}(z))|) dz \\
&= \frac{1}{4} \cdot \int_z (|\mathbb{P}_y(z) - \mathbb{P}_{y'}(z)| + |\mathbb{P}_{y'}(z) - \mathbb{P}_y(z)|) dz \\
&= \mathcal{D}_{TV}(\mathcal{N}(y, \Sigma(y)) \| \mathcal{N}(y', \Sigma(y'))).
\end{aligned} \tag{37}$$

Here, we simply use $\mathbb{P}_y(z)$ to denote the probability density of $Q(y)$ at point z . Now, we can apply the results about the lower bound of total variation for two multivariate Gaussians [26], [27], where

$$\mathcal{D}_{TV}(\mathcal{N}(y, \Sigma(y)) \| \mathcal{N}(y', \Sigma(y'))) = \Omega\left(1, \frac{\|y - y'\|^2}{\sqrt{(y - y')^T \Sigma_y (y - y')}}\right). \tag{38}$$

Since y' is free to select, we can set $y' - y$ to lie on the same direction of the j -th eigenvector of $\Sigma(y)$. On the other hand, as we assume $v = o(1)$, from the right side of (38), it implies $\frac{\|y - y'\|^2}{\sqrt{\lambda_j \|y - y'\|}} \leq \sqrt{v}$. Since we can select y' such that $\|y - y'\| = \Omega(r)$, therefore $\sqrt{\lambda_j} = \Omega(r/\sqrt{v})$ for any $j \in \{1, 2, \dots, d\}$. Therefore, $\mathbb{E}_{\sim Q(y)}[\| \cdot \|_2] = \Omega(r\sqrt{d}/\sqrt{v})$, a contradiction.

J Proof of Theorem 10

We prove the theorem by a contradiction. Without loss of generality we assume the distribution of $\mathcal{M}(X)$ is continuous and by definition,

$$\begin{aligned}
\text{Gen}(\cdot, \mathcal{M}) &= \mathbb{E}_{\tilde{X} \perp X \sim \cdot, \mathcal{M}}(\mathcal{L}(\tilde{X}, \mathcal{M}(X))) - \mathbb{E}_{X \sim \cdot, \mathcal{M}}(\mathcal{L}(X, \mathcal{M}(X))) \\
&= \frac{1}{N} \sum_{i=1}^N \left(\frac{\sum_{l=1}^N \int_y \mathcal{L}(X_i, y) \cdot \mathbb{P}_{\mathcal{M}(X_l)}(y) dy}{N} - \int_y \mathcal{L}(X_i, y) \cdot \mathbb{P}_{\mathcal{M}(X_i)}(y) dy \right) > \sqrt{v/2}.
\end{aligned} \tag{39}$$

With a rearrangement of the terms in (39), we have that

$$\int_y \sum_{i=1}^N \frac{\mathcal{L}(X_i, y)}{N} \cdot \left(\frac{\sum_{l=1}^N \mathbb{P}_{\mathcal{M}(X_l)}(y)}{N} - \mathbb{P}_{\mathcal{M}(X_i)}(y) \right) dy \geq \sqrt{v/2}. \tag{40}$$

Since the output of loss function \mathcal{L} is bounded and within $[0, 1]$, (40) suggests that

$$\frac{1}{N} \cdot \int_y \sum_{i=1}^N \left| \frac{\sum_{l=1}^N \mathbb{P}_{\mathcal{M}(X_l)}(y)}{N} - \mathbb{P}_{\mathcal{M}(X_i)}(y) \right| dy \geq \sqrt{v/2}. \tag{41}$$

Now, we consider the inference task that an adversary sets out to identify the true selection of X from \cdot with observations $\mathcal{M}(X)$, and we construct a strategy as follows. Given any observation $y = \mathcal{M}(X)$, the adversary will output $\tilde{X} = X_{i_0}$ such that $i_0 = \arg \max_i \mathbb{P}_{\mathcal{M}(X_i)}(y)$, with a *maximum a posteriori* estimation.

It is clear that the optimal *a priori* success rate $(1 - \delta_o^{\epsilon, \rho}) = 1/N$, while the posterior success $(1 - \delta)$ of the above strategy is at least

$$\begin{aligned}
&\frac{1}{N} \cdot \int_y \mathbb{P}_{\mathcal{M}(X_{i_0})}(y) dy = \frac{1}{N} \cdot \int_y \max_i \{\mathbb{P}_{\mathcal{M}(X_1)}(y), \dots, \mathbb{P}_{\mathcal{M}(X_N)}(y)\} dy \\
&= \frac{1}{N} \cdot \int_y \left(\frac{\sum_{i=1}^N \mathbb{P}_{\mathcal{M}(X_i)}(y)}{N} + \mathbb{P}_{\mathcal{M}(X_{i_0})}(y) - \frac{\sum_{i=1}^N \mathbb{P}_{\mathcal{M}(X_i)}(y)}{N} \right) dy \\
&\geq \frac{1}{N} + \frac{1}{N} \cdot \int_y \sum_{i=1}^N \left| \frac{\sum_{l=1}^N \mathbb{P}_{\mathcal{M}(X_l)}(y)}{N} - \mathbb{P}_{\mathcal{M}(X_i)}(y) \right| dy > \frac{1}{N} + \sqrt{v/2}.
\end{aligned} \tag{42}$$

In the last line of (42), we use the fact that given $\mathbb{P}_{\mathcal{M}(X_{i_0})}(y)$ is the maximal across $\mathbb{P}_{\mathcal{M}(X_i)}(y)$, $\mathbb{P}_{\mathcal{M}(X_{i_0})}(y) - \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathcal{M}(X_i)}(y)$ is also the maximal across all the possible $|\mathbb{P}_{\mathcal{M}(X_i)}(y) - \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathcal{M}(X_i)}(y)|$ for any $i = 1, 2, \dots, N$. Thus, based on the maximal is no less than the average of a set of numbers, we have the lower bound in (42).

Now, by Pinsker's inequality, if $\Delta_{KL}\delta \leq v$, then $\Delta_{TV}\delta \leq \sqrt{v/2}$. Therefore, $\delta_o^{\epsilon, \rho} - \delta \leq \sqrt{v/2}$, which leads to a contradiction.

K Local P C Privacy

In this section, we generalize P C Privacy to a local/decentralized setup. The first application is similar to local DP [31], [44], where a user does not trust any other parties and plans to publish her perturbed sample. The perturbed released individual data could then be collected and postprocessed, which does not incur additional privacy loss. In a scenario where a user randomly samples x from her local sample set \mathbf{U} and publishes, we may select \mathcal{M} to be the identity function, i.e., $\mathcal{M}(x) = x$ and measure the corresponding P C Privacy.

Example 6 (Data Publishing via Local DP and P C Privacy). *We suppose \mathbf{U} to be CIF R10 [45], a canonical test dataset commonly used in computer vision, which consists of 60,000 32×32 color images. We normalize each pixel to within $[0, 1]$ and take each sample as a 3072-dimensional vector. Let the local data publishing be $\mathcal{M}(x) = x$ where x is randomly sampled from CIF R10. With a target bound $\text{MI}(x; \mathcal{M}(x) + \text{noise}) \leq 1$, given l_∞ -norm sensitivity, the adversarial worst-case local DP requires (d) -large noise, where in this case $\mathbb{E}[\|\text{noise}\|_2] = d/\sqrt{2} = 2.2 \times 10^3$ via the Gaussian Mechanism [14]. On the other hand, via the upper bound in Theorem 3, for P C Privacy, we only need noise $\mathbb{E}[\|\text{noise}\|_2] \leq 1.6 \times 10^2$, $13\times$ smaller than that of local DP.*

The case where we consider a decentralized system of n semi-honest users is more complicated. Each of the users will provide some local data x_i . With $X = (x_1, x_2, \dots, x_n)$, they set out to collaboratively implement some mechanism $\mathcal{M}(x_1, \dots, x_n)$, and the output $\mathcal{M}(x_1, \dots, x_n)$ is the only *necessary* information they need to disclose. Meanwhile, they wish to minimize the leakage of x_i s from the exposure of $\mathcal{M}(x_1, \dots, x_n)$.

Recall the P C Privacy framework, which is basically three parts: (a) through simulation, we determine the parameters of noise distribution $Q(X, \theta)$; (b) we generate X and θ , and evaluate $\mathcal{M}(X, \theta)$; (c) we generate $(X, \theta) \sim Q(X, \theta)$ and publish perturbed $\mathcal{M}(X, \theta) + \text{noise}$. However, without a trusted server, a straightforward implementation of this protocol is problematic. First, the mechanism \mathcal{M}' we need to analyze now integrates the three procedures (a)-(c) mentioned above, where all intermediate simulations may cause additional privacy loss. Second and more important, to analyze the privacy of the new mechanism \mathcal{M}' , we then need \mathcal{M}'' to incorporate both \mathcal{M}' and its new privacy analysis simulation, which turns into an endless loop.

To this end, we need a secure computation for all three procedures (a)-(c) to fill the gap between local and centralized P C Privacy. Fortunately, if we slightly relax the threat model, Secure Multi-party Computation (MPC) [25] provides a generic solution. Theoretically, MPC enables a group of users to compute any generic function, and meanwhile it ensures that besides the output of the function and a (possibly colluding) user's own input, there is negligible additional information leakage to any polynomial-computation-bounded adversary⁷. Therefore, when the adversary is relaxed to be computationally-bounded, the group of semi-honest users⁸ may implement MPC for both privacy analysis simulation and the computation of $\mathcal{M}(X) + \text{noise}$, where the output finally disclosed is a single (noisy) evaluation $\mathcal{M}(X) + \text{noise}$ (if noise is necessary). We formally describe such *Locally-collaborative P C Individual Privacy* as follows:

Definition 10 ($(\delta, \rho, n, \epsilon)$ Locally-collaborative P C Individual Privacy). *For a data processing mechanism \mathcal{M} and n semi-honest computationally-bounded users, who (collaboratively) generate a dataset $X =$*

⁷Though MPC can ensure negligible information leakage in decentralized computing, there is no guarantee for the disclosure of output $\mathcal{M}(X)$. Therefore, we need a notion such as P C Privacy for generic privacy measurement.

⁸We need to assume users are semi-honest to ensure that the data generation follows the prescribed distribution Q , or additional verifiable computation on data generation.

(x_1, x_2, \dots, x_n) from some data distribution where x_i is from the i -th user, we say \mathcal{M} satisfies $(\delta, \rho, n, \epsilon)$ -locally-collaborative P - C Individual Privacy if for any $i \neq j \in [1 : n]$ the following experiment is impossible: dataset $X = (x_1, x_2, \dots, x_n)$ is generated from distribution \mathcal{D} and $\mathcal{M}(X)$ is published. The j -th user who knows x_j , and \mathcal{M} , $j \neq i$, is asked to return an estimation \tilde{x}_i on x_i such that with probability at least $(1 - \delta)$, $\rho(\tilde{x}_i, x_i) = 1$.

Compared to the centralized case, the only difference in the analysis is that we need to control the conditional $\text{MI}(x_i; \mathcal{M}(X)|x_j)$ for any $i \neq j \in [1 : n]$ rather than simply $\text{MI}(x_i; \mathcal{M}(X))$. We need to ensure that the j -th user who already knows her own data x_j still cannot infer too much about another i -th user's data from $\mathcal{M}(X)$. By the chain rule, we have that

$$\text{MI}(x_i; \mathcal{M}(X)|x_j) \leq \text{MI}(X; \mathcal{M}(X)) - \text{MI}(x_j; \mathcal{M}(X)) \leq \text{MI}(X; \mathcal{M}(X)),$$

for any i and j . Therefore, $\text{MI}(X; \mathcal{M}(X))$ is still an upper bound on the posterior advantage in the locally-collaborative version. As a final remark, the above semi-honest assumption can be further relaxed to the adversarial setup with additional applications of unconditionally secure MPC [11] and verifiable randomness [36]. The need of verifiable randomness is to ensure the distributions of data generation and randomness seed selection are consistent in both privacy analysis steps (a,b) and output generation steps (c) for each user.

Remark 2. Though $\text{MI}(X; \mathcal{M}(X))$ is an upper bound on both $\max_i \text{MI}(x_i; \mathcal{M}(X))$ and $\max_{i,j} \text{MI}(x_i; \mathcal{M}(X)|x_j)$, which can produce security parameters of centralized and local individual P - C Privacy in Definitions 6 and 10, respectively, we must stress that, using $\text{MI}(X; \mathcal{M}(X))$, we also obtain stronger and more meaningful P - C Privacy for the whole set X . Take membership inference as an example. If each user decides to randomly participate in the analysis with her local data x_i , $\text{MI}(X; \mathcal{M}(X))$ quantifies the maximal posterior advantage for any possible inference the adversary could implement on all participants rather than that of a single one.