



HAL
open science

Differential Privacy versus Quantitative Information Flow

Mário S. Alvim, Konstantinos Chatzikokolakis, Pierpaolo Degano, Catuscia Palamidessi

► **To cite this version:**

Mário S. Alvim, Konstantinos Chatzikokolakis, Pierpaolo Degano, Catuscia Palamidessi. Differential Privacy versus Quantitative Information Flow. 2010. hal-00548214

HAL Id: hal-00548214

<https://hal.science/hal-00548214v1>

Preprint submitted on 20 Dec 2010

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

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

Differential Privacy versus Quantitative Information Flow^{*}

Mário S. Alvim¹, Konstantinos Chatzikokolakis²,
Pierpaolo Degano³, and Catuscia Palamidessi¹

¹ INRIA and LIX, Ecole Polytechnique, France.

² Technical University of Eindhoven, The Netherlands.

³ Dipartimento di Informatica, Università di Pisa, Italy.

Abstract. Differential privacy is a notion of privacy that has become very popular in the database community. Roughly, the idea is that a randomized query mechanism provides sufficient privacy protection if the ratio between the probabilities of two different entries to originate a certain answer is bound by e^ϵ . In the fields of anonymity and information flow there is a similar concern for controlling information leakage, i.e. limiting the possibility of inferring the secret information from the observables. In recent years, researchers have proposed to quantify the leakage in terms of the information-theoretic notion of mutual information. There are two main approaches that fall in this category: One based on Shannon entropy, and one based on Rényi's min entropy. The latter has connection with the so-called Bayes risk, which expresses the probability of guessing the secret.

In this paper, we show how to model the query system in terms of an information-theoretic channel, and we compare the notion of differential privacy with that of mutual information. We show that the notion of differential privacy is strictly stronger, in the sense that it implies a bound on the mutual information, but not viceversa.

1 Introduction

The growth of information technology raises significant concerns about the vulnerability of sensitive information. The possibility of collecting and storing data in large amount and the availability of powerful data processing techniques open the way to the threat of inferring private and secret information, to such an extent that fully justifies the users' worries.

1.1 Differential privacy

The area of statistical databases has been, naturally, one of the first communities to consider the issues related to the protection of information. Already some

^{*} This work has been partially supported by the project ANR-09-BLAN-0169-01 PANDA and by the INRIA DRI Equipe Associée PRINTEMPS.

decades ago, Dalenius [10] proposed a famous “ad omnia” privacy desideratum: nothing about an individual should be learnable from the database that cannot be learned without access to the database.

Dalenius’ property, however, is too strong to be useful in practice: it has been shown by Dwork [11] that no useful database can provide it. In replacement Dwork has proposed the notion of *differential privacy*, which has had an extraordinary impact in the community. Intuitively, such notion is based on the idea that the presence or the absence of an item in the database should not change in a significant way the probability of obtaining a certain answer for a given query [11,12,13,14].

In order to explain the concept more precisely, let us consider the typical scenario: we have databases whose entries are values (possibly tuples) taken from a given universe. A database can be queried by users which have honest purposes, but also by attackers trying to infer secret or private data. In order to control the leakage of secret information, the curator uses some randomized mechanism, which causes a certain lack of precision in the answers. Clearly, there is a trade off between the need of obtaining answers as precise as possible for legitimate use, and the need to introduce some fuzziness to the purpose of confusing the attacker.

Let \mathcal{K} be the randomized function that provides the answers to the queries. We say that \mathcal{K} provides ϵ -differential privacy if for all databases D and D' , such that one is a subset of the other and the larger contains a single additional entry, and for all $S \subseteq \text{range}(\mathcal{K})$, the ratio between the probability that the result of $\mathcal{K}(D)$ is in S , and the probability that the result of $\mathcal{K}(D')$ is in S , is at most e^ϵ .

Dwork has also studied sufficient conditions for a randomized function \mathcal{K} to implement a mechanism satisfying ϵ -differential privacy. It suffices to consider a Laplacian distribution with variance depending on ϵ , and mean equal to the correct answer [13]. This is a technique quite diffused in practice.

1.2 Quantitative information flow and anonymity

The problem of preventing the leakage of secret information has been a pressing concern also in the area of software systems, and has motivated a very active line of research called *secure information flow*. Similarly to the case of privacy, also in this field, at the beginning, the goal was ambitious: to ensure *non-interference*, which means complete lack of leakage. But, as for Dalenius’ notion of privacy, no-interference is too strong for being obtainable in practice, and the community has started exploring weaker notions. Some of the most popular approaches are the quantitative ones, based on information theory. See for instance [6,7,8,15,16,17,21].

Independently the field of anonymity, which is concerned with the protection of the identity of agents performing certain tasks, has evolved towards similar approaches. In the case of anonymity it is even more important to consider a quantitative formulation, because *anonymity protocols typically use randomization* to obfuscate the link between the *culprit* (i.e. the agent which performs the task) and the observable effects of the task. The first notion of anonymity, due

to Chaum [5], required that the observation would not change the probability of an individual to be the culprit. In other words, the protocol should guarantee that the observation does not increase the chances of learning the identity of the culprit. This is very similar to Dalenius' notion of privacy, and equally unattainable in practice (at least, in the majority of real situations). Also in this case, researchers in the area have started considering weaker notions based on information theory, see for instance [3,18,22].

If we abstract from the kind of secrets and observables, anonymity and of information flow are similar problems: there is some information that we want to keep secret, there is a system that produces some kind of observable information depending on the secret one, and we want to prevent as much as possible that an attacker may infer the secrets from the observables. It is therefore not surprising that the foundations of the two fields have converged towards the same information theoretical approaches. The majority of these approaches are based on the idea of representing the system (or protocol) as an information-theoretic channel taking the secrets in input (X) and producing the observables in output (Y). The *entropy* of X , $H(X)$, represents the converse of the *a priori vulnerability*, i.e. the chance of the attacker to find out the secret. Similarly, the conditional entropy of X given Y , $H(X | Y)$, represents the converse of the *a posteriori vulnerability*, i.e. the chance of the attacker to find out the secret after having observed the output. The *mutual information* between X and Y , $I(X; Y) = H(X) - H(X | Y)$, represents the gain for the adversary provided by the observation, and is taken as definition of the *information leakage* of the system. Sometimes we may want to abstract from the distribution of X , in which case we can use the *capacity* of the channel, defined as the maximum of $I(X; Y)$ over all possible distributions on X . This represents the worst case for leakage.

The various approaches in literature differ, mainly, for the notion of entropy. Such notion is related to the kind of attackers we want to model, and to how we measure their success (see [15] for an illuminating discussion of such relation). Shannon entropy [20], on which most of the approaches are based, represents an adversary which tries to find out the secret x by asking questions of the form "does x belong to set S ?". Shannon entropy is precisely the average number of questions necessary to find out the exact value of x with an optimal strategy (i.e. an optimal choice of the S 's). The other most popular notion of entropy (in this area) is Rényi's min entropy [19]. The corresponding notion of attack is a *single try* of the form "is x equal to v ?". Rényi's min entropy is precisely the log of the probability of guessing the true value with the optimal strategy, which consists, of course, in selecting the v with the highest probability. Approaches based on this notion include [21] and [2].

It is worth noting that, while the Rényi's min entropy of X , $H_\infty(X)$, represents the a priori probability of success (of the single-try attack), the Rényi's min conditional entropy of X given Y , $H_\infty(X | Y)$, represents the a posteriori

probability of success¹. This a posteriori probability is the converse of the Bayes risk [9] , which has also been used as a measure of leakage [1,4].

1.3 Goal of the paper

From a mathematical point of view, privacy presents many similarities with information flow and anonymity. The private data of the entry constitute the secret, the answer to the query gives the observation, and the goal is to prevent as much as possible the inference of the secret from the observable. Differential privacy can be seen as a quantitative definition of the degree of leakage. The main goal of this paper is to explore the relation with the alternative definitions based on information theory, with the purpose of getting a better understanding of the notion of differential privacy, of the specific problems related to privacy, and of the models of attack used to formalize the notion of privacy, in relation to those used for anonymity and information flow.

1.4 Contribution

The contribution of this paper is as follows:

- We show how the problem of privacy can be formulated in an information-theoretic setting. More precisely, we show how the answer function \mathcal{K} can be associated to an information-theoretic channel.
- We prove that ϵ -differential privacy implies a bound on the Shannon mutual information of the channel, and that this bound approach 0 as ϵ approaches 0. Same for Rényi min mutual information.
- We show that the viceversa of the above point does not hold, i.e. that Shannon and Rényi min mutual information (and also the corresponding capacities) can approach 0 while the ϵ parameter of differential privacy approaches infinity.

1.5 Plan of the paper

Next section introduces some necessary background notions. Section 3 proposes an information-theoretic view of the database query systems. Section 4 show the main results of the paper, namely that differential privacy implies a bound on Shannon and Rényi min mutual information, but not viceversa. Section 5 concludes and presents some ideas for future work.

The proofs of the results are in the appendix. Such appendix will not be included in the proceeding version (for reasons of space), but the proofs will be made available on line.

¹ We should mention that Rényi did not define the conditional version of the min entropy, and that there have been various different proposals in literature for this notion. We use here the one proposed by Smith in [21].

2 Preliminaries

2.1 Differential privacy

We assume a fixed finite universe U in which the entries of databases may range. The concept of differential privacy is tightly connected to the concept of *adjacent* (or *neighbor*) databases.

Definition 1 ([13]). *A pair of databases (D', D'') is considered adjacent (or neighbors) if one is a proper subset of the other and the larger database contains just one additional entry.*

Dwork's definition of differential privacy is the following:

Definition 2 ([11]). *A randomized function \mathcal{K} satisfies ϵ -differential privacy if for all pairs of adjacent databases D' and D'' , and all $S \subseteq \text{Range}(\mathcal{K})$,*

$$\Pr[\mathcal{K}(D') \in S] \leq e^\epsilon \times \Pr[\mathcal{K}(D'') \in S] \quad (1)$$

2.2 Information theory and interpretation in terms of attacks

In the following, X, Y denote two discrete random variables with carriers $\mathcal{X} = \{x_1, \dots, x_n\}$, $\mathcal{Y} = \{y_1, \dots, y_m\}$, and probability distributions $p_X(\cdot)$, $p_Y(\cdot)$, respectively. An information-theoretic channel is constituted by an input X , an output Y , and the matrix of conditional probabilities $p_{Y|X}(\cdot | \cdot)$, where $p_{Y|X}(y | x)$ represent the probability that Y is y given that X is x . We will use $X \wedge Y$ to represent the random variable with carrier $\mathcal{X} \times \mathcal{Y}$ and joint probability distribution $p_{X \wedge Y}(x, y) = p_X(x) \cdot p_{Y|X}(y | x)$. We shall omit the subscripts on the probabilities when they are clear from the context.

2.3 Shannon entropy

The Shannon entropy of X is defined as

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x)$$

The minimum value $H(X) = 0$ is obtained when $p(\cdot)$ is concentrated on a single value (i.e. when $p(\cdot)$ is a delta of Dirac). The maximum value $H(X) = \log |\mathcal{X}|$ is obtained when $p(\cdot)$ is the uniform distribution. Usually the base of the logarithm is set to be 2 and, correspondingly, the entropy is measured in *bits*.

The *conditional entropy* of X given Y is

$$H(X | Y) = \sum_{y \in \mathcal{Y}} p(y) H(X | Y = y)$$

where

$$H(X | Y = y) = - \sum_{x \in \mathcal{X}} p(x | y) \log p(x | y)$$

We can prove that $0 \leq H(X | Y) \leq H(X)$. The minimum value, 0, is obtained when X is completely determined by Y . The maximum value $H(X)$ is obtained when Y reveals no information about X , i.e. when X and Y are independent.

The *mutual information* between X and Y is defined as

$$I(X; Y) = H(X) - H(X | Y) \quad (2)$$

and it measures the amount of information about X that we gain by observing Y . It can be shown that $I(X; Y) = I(Y; X)$ and $0 \leq I(X; Y) \leq H(X)$.

Shannon capacity is defined as the maximum mutual information over all possible input distributions:

$$C = \max_{p_X(\cdot)} I(X; Y)$$

2.4 Rényi min-entropy

In [19], Rényi introduced an one-parameter family of entropy measures, intended as a generalization of Shannon entropy. The Rényi entropy of order α ($\alpha > 0$, $\alpha \neq 1$) of a random variable X is defined as

$$H_\alpha(X) = \frac{1}{1-\alpha} \log \sum_{x \in \mathcal{X}} p(x)^\alpha$$

We are particularly interested in the limit of H_α as α approaches ∞ . This is called *min-entropy*. It can be proven that

$$H_\infty(X) \stackrel{\text{def}}{=} \lim_{\alpha \rightarrow \infty} H_\alpha(X) = -\log \max_{x \in \mathcal{X}} p(x)$$

Rényi defined also the α -generalization of other information-theoretic notions, like the Kullback-Leibler divergence. However, he did not define the α -generalization of the conditional entropy, and there is no agreement on what it should be. For the case $\alpha = \infty$, we adopt here the definition of conditional entropy proposed by Smith in [21]:

$$H_\infty(X | Y) = -\log \sum_{y \in \mathcal{Y}} p(y) \max_{x \in \mathcal{X}} p(x | y) \quad (3)$$

Analogously to (2), we can define the mutual information I_∞ as $H_\infty(X) - H_\infty(X | Y)$, and the capacity C_∞ as $\max_{p_X(\cdot)} I_\infty(X; Y)$. It has been proven in [2] that C_∞ is obtained at the uniform distribution, and that it is equal to the sum of the maxima of each column in the channel matrix:

$$C_\infty = \sum_{y \in \mathcal{Y}} \max_{x \in \mathcal{X}} p(y | x).$$

3 An information theoretic model of privacy

In this section we show how to represent a database query system (of the kind considered in differential privacy) in terms of an information-theoretic channel.

According to [11] and [13], differential privacy can be implemented by adding some appropriately chosen random noise to the answer $x = f(D)$, where f is the *query function* and D is the database. The function can operate in the entire database at once, and even though the query may be composed by a chain of sub-queries, we assume that subsequent sub-queries depend only on the *true answer* to previous sub-queries. Under this constraint, no matter how complex the query is, it is still a function f of the database D . The scenario where subsequent sub-queries can depend on the *reported answer* to previous queries corresponds to adaptive adversaries [11], and is not considered in this paper.

After the true answer x to the query is obtained from D , some noise is introduced in order to produce a reported answer y . The reported answer can be seen as a random variable Y dependent on the random variable X corresponding to the real answer, and the two random variables are related by a conditional probability distribution $p_{Y|X}(\cdot|\cdot)$. The conditional probabilities $p_{Y|X}(y|x)$ constitute the matrix of an information theoretic channel from X to Y .

Figure 1 shows the scheme of implementation of a differential privacy scheme.




Fig. 1. The channel corresponding to a differential privacy scheme.

In [11] it has been proved that a way to define the values of $p_{Y|X}(\cdot|\cdot)$ so to ensure ϵ -differential privacy, is by using the Laplace distribution:

$$P((Y = y)|(X = x), \Delta f/\epsilon) = \frac{\Delta f}{2\epsilon} e^{-|y-x|/\epsilon/\Delta f} \quad (4)$$

where Δf is the L1-sensitivity of f , defined as²

$$\Delta f = \max_{D', D'' \text{ adjacent}} |f(D') - f(D'')|.$$

² We give here the definition for the case in which the range of f is \mathbb{R} . In the more general case in which the range is \mathbb{R}^n we should replace $|f(D') - f(D'')|$ by the 1-norm of the vector $f(D') - f(D'')$.

4 Relation between differential privacy and mutual information

In this section we investigate the relation between differential privacy and information-theoretic notions. We start by considering an equivalent definition of differential privacy, easier to handle for our purposes.

4.1 Testing single elements

Definition 2 considers tests which check whether the result of $\mathcal{K}(D)$ belongs to a certain set or not. We prefer to simplify this definition by considering only tests over single elements:

Definition 3. *A randomized function \mathcal{K} gives δ -differential privacy if for all pairs adjacent datasets D' and D'' , and all $k \in \text{Range}(\mathcal{K})$,*

$$\Pr[\mathcal{K}(D') = k] \leq e^\delta \times \Pr[\mathcal{K}(D'') = k] \quad (5)$$

The following result shows that our definition of differential privacy is equivalent to the classical one.

Theorem 1. *A function \mathcal{K} gives ϵ -differential privacy iff it gives δ -differential privacy, with $\epsilon = \delta$.*

4.2 Databases with the same number of entries and differing in at most one entry

Consider two databases D' and D'' that have the same number of entries and differ in at most one entry as in Figure 2. Let D be the common part shared by both databases, and let r' and r'' be the rows in which they differ, namely $D' = D \cup \{r'\}$ and $D'' = D \cup \{r''\}$.




Fig. 2. Two databases differing in exactly one entry

We prove that δ -differential privacy imposes also a bound on the comparison between databases with the same number of entries, and which differ in the values of only one entry.

Lemma 1. *Let \mathcal{K} be a function that gives δ -differential privacy for all pairs of adjacent databases. Given two databases D' and D'' that have the same number of entries and differ in the value of at most one entry, then:*

$$\Pr[\mathcal{K}(D') = k] \leq e^{2\delta} \times \Pr[\mathcal{K}(D'') = k]$$

4.3 Shannon mutual information

We prove now that δ -differential privacy imposes a bound on Shannon mutual information, and that this bound approaches 0 as the parameter δ approaches 0.

Theorem 2. *If a randomized function \mathcal{K} gives δ -differential privacy according to Definition 3, then for every result x^* of the function f the Shannon mutual information between the true answers X (i.e. the results of f) and the reported answers Y (i.e. the results of \mathcal{K}) is bounded by:*

$$I(X; Y) \leq (e^{2\delta} + e^{-2\delta})\delta \log(e) + (e^{2\delta} - e^{-2\delta}) \sum_y p(y|x^*) \log(p(y|x^*))$$

It is easy to see that the expression which bounds I from above, $(e^{2\delta} + e^{-2\delta})\delta \log(e) + (e^{2\delta} - e^{-2\delta}) \sum_y p(y|x^*) \log(p(y|x^*))$, converges to 0 when δ approaches 0.

The converse of Theorem 2 does not hold. One reason is that mutual information is sensitive to the values of the input distribution, while differential privacy is not. Next example illustrates this point.

Example 1. Let n be the number of elements of the universe, and m the cardinality of the set of possible answers of f . Assume that $p(x_1) = \alpha$ and $p(x_i) = \frac{1-\alpha}{n-1}$ for $2 \leq i \leq n$. Let $p(y_1 | x_1) = \beta$, $p(y_j | x_1) = \frac{1-\beta}{m-1}$ for $2 \leq j \leq m$, and $p(y_j | x_i) = \frac{1}{m}$, otherwise. This channel is represented in Figure 0(a). It is easy to see that the Shannon mutual information approaches 0 as α approaches 0, independently of the value of β . Differential privacy, however, depends only on the value of β , more precisely, the parameter of differential privacy is $\max\{\log_e \frac{1}{m\beta}, \log_e m\beta, \log_e \frac{m-1}{m(1-\beta)}, \log_e \frac{m(1-\beta)}{m-1}\}$, and it is easy to see that such parameter is unbound and goes to infinity as β approaches 0.

The reasoning in the counterexample above is not valid anymore if we consider capacity instead than mutual information. However, there is another reason why the converse of Theorem 2 does not hold, and this remains the case also if we consider capacity. The situation is illustrated by the following example.

Example 2. Let n be the number of elements of the universe, and m the cardinality of the set of possible answers of f . Assume that $p(y_i | x_i) = \beta$ and $p(y_i | x_j) = \frac{1-\beta}{m-1}$ for $i \neq j$. This channel is represented in Figure 0(b). It is easy to see that the Shannon capacity is $C = \log m - (1-\beta) \log(m-1) + \beta \log \beta + (1-\beta) \log(1-\beta)$, and that C approaches 0 as β approaches 0 and m becomes large. Differential privacy, however, goes in the other direction when β approaches 0, and it is not very sensitive to the value of m . More precisely, the parameter of differential privacy is $\max\{\log_e \frac{1-\beta}{\beta(1-m)}, \frac{\beta(1-m)}{1-\beta}\}$, and it is easy to see that such parameter is unbound and goes to infinity as β approaches 0, independently of the value of m .

	(a) Example 1				(b) Example 2					
$p_X(\cdot)$	y_1	y_2	\dots	y_m	y_1	y_2	\dots	y_m		
α	x_1	β	$\frac{1-\beta}{m-1}$	\dots	$\frac{1-\beta}{m-1}$	x_1	β	$\frac{1-\beta}{m-1}$	\dots	$\frac{1-\beta}{m-1}$
$\frac{1-\alpha}{m-1}$	x_2	$\frac{1}{m}$	$\frac{1}{m}$	\dots	$\frac{1}{m}$	x_2	$\frac{1-\beta}{m-1}$	β	\dots	$\frac{1-\beta}{m-1}$
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
$\frac{1-\alpha}{m-1}$	x_n	$\frac{1}{m}$	$\frac{1}{m}$	\dots	$\frac{1}{m}$	x_n	$\frac{1-\beta}{m-1}$	$\frac{1-\beta}{m-1}$	\dots	β

Table 1. The channels of Examples 1 and 2

4.4 Rényi min mutual information

We show now that a result analogous to that of Section 4.3 holds also in the case of Rényi min entropy.

Theorem 3. *If a randomized function \mathcal{K} gives δ -differential privacy according to Definition 3, then the Rényi min mutual information between the true answer of the function X and the reported answer Y is bounded by*

$$I_\infty(X; Y) \leq 2\delta \log e.$$

The converse of Theorem 3 does not hold, not even if we consider capacity instead than mutual information. It is easy to prove, in fact, that Examples 1 and 2 lead to counterexamples also in the case of Rényi min mutual information and capacity.

5 Conclusion and future work

In this paper we have shown that the problem of privacy in statistical databases can be formulated in information-theoretic terms, in a way analogous to what has been done for information flow and anonymity: the database query system can be seen as a noisy channel, in the information-theoretic sense. Then we have considered Dwork's notion of differential privacy, and we have shown that it is strictly stronger than requiring the channel to have low capacity, both for the cases of Shannon and Rényi min entropy. It is natural to consider, then, whether a weaker notion would give enough privacy guarantees. As future work, we intend to investigate this question.

We first need to understand, of course, what are the constraints that could be relaxed in the notion of differential privacy. To this aim, Example 2 is quite interesting: whenever we get an answer y , there are $n - 1$ possible inputs (entries) which are equally likely to have generated that answer, and one input x that is much less likely than the others ($p(x|y) = \alpha$, where α is a very small value). The existence of the latter seems quite harmless, yet it is exactly that entry that causes differential privacy to fail (in the sense that its parameter is unbound).

The notion of Rényi min capacity seems a plausible candidate for the notion of privacy: its relation with the Bayes risk ensures that a bound C_∞ can be seen as a bound on the probability of guessing the right value of x (given the observable). In some scenario, this may be exactly what we want.

Acknowledgement

We wish to thank Daniel Le Métayer for having pointed out to us the notion of differential privacy, and brought to our attention the possible relation with quantitative information flow.

References

1. Christelle Braun, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Compositional methods for information-hiding. In *Proc. of FOSSACS*, volume 4962 of *LNCS*, pages 443–457. Springer, 2008.
2. Christelle Braun, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Quantitative notions of leakage for one-try attacks. In *Proc. of MFPS*, volume 249 of *ENTCS*, pages 75–91. Elsevier, 2009.
3. Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panangaden. Anonymity protocols as noisy channels. *Inf. and Comp.*, 206(2–4):378–401, 2008.
4. Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panangaden. On the Bayes risk in information-hiding protocols. *J. of Comp. Security*, 16(5):531–571, 2008.
5. David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
6. David Clark, Sebastian Hunt, and Pasquale Malacaria. Quantitative analysis of the leakage of confidential data. In *Proc. of QAPL*, volume 59 (3) of *Electr. Notes Theor. Comput. Sci.*, pages 238–251. Elsevier, 2001.
7. David Clark, Sebastian Hunt, and Pasquale Malacaria. Quantitative information flow, relations and polymorphic types. *J. of Logic and Computation*, 18(2):181–199, 2005.
8. Michael R. Clarkson, Andrew C. Myers, and Fred B. Schneider. Belief in information flow. *J. of Comp. Security*, 17(5):655–701, 2009.
9. Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. J. Wiley & Sons, Inc., second edition, 2006.
10. Tore Dalenius. Towards a methodology for statistical disclosure control. *Statistik Tidskrift*, 15:429 — 444, 1977.
11. Cynthia Dwork. Differential privacy. In *Automata, Languages and Programming, 33rd Int. Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proc., Part II*, volume 4052 of *LNCS*, pages 1–12. Springer, 2006.
12. Cynthia Dwork. Differential privacy in new settings. In *Proc. of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 174–183. SIAM, 2010.
13. Cynthia Dwork. A firm foundation for private data analysis. *Communications of the ACM*, 2010. To appear.
14. Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proc. of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 371–380. ACM, 2009.

15. Boris Köpf and David A. Basin. An information-theoretic model for adaptive side-channel attacks. In *Proc. of CCS*, pages 286–296. ACM, 2007.
16. Pasquale Malacaria. Assessing security threats of looping constructs. In *Proc. of POPL*, pages 225–235. ACM, 2007.
17. Pasquale Malacaria and Han Chen. Lagrange multipliers and maximum information leakage in different observational models. In *Proc. of PLAS*, pages 135–146. ACM, 2008.
18. Ira S. Moskowitz, Richard E. Newman, Daniel P. Crepeau, and Allen R. Miller. Covert channels and anonymizing networks. In *Proc. of PES*, pages 79–88. ACM, 2003.
19. Alfréd Rényi. On Measures of Entropy and Information. In *Proc. of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability*, pages 547–561, 1961.
20. Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 625–56, 1948.
21. Geoffrey Smith. On the foundations of quantitative information flow. In *Proc. of FOSSACS*, volume 5504 of *LNCS*, pages 288–302. Springer, 2009.
22. Ye Zhu and Riccardo Bettati. Anonymity vs. information leakage in anonymity systems. In *Proc. of ICDCS*, pages 514–524. IEEE, 2005.

Appendix

Theorem 4 (Theorem 1 in the paper). *A function \mathcal{K} gives ϵ -differential privacy iff it gives δ -differential privacy, with $\epsilon = \delta$.*

Proof.

\Rightarrow Let $k \in \text{Range}(\mathcal{K})$. Then for all pair of adjacent databases D', D'' we have

$$\begin{aligned} \Pr[\mathcal{K}(D') = k] &= \Pr[\mathcal{K}(D') \in \{k\}] && \text{(taking } S \text{ to be a singleton set)} \\ &\leq e^\epsilon \Pr[\mathcal{K}(D'') \in \{k\}] && \text{(by Definition 2)} \\ &= e^\epsilon \Pr[\mathcal{K}(D'') = k] \end{aligned}$$

\Leftarrow Let $S \subseteq \text{Range}(\mathcal{K})$

$$\begin{aligned} \Pr[\mathcal{K}(D') \in S] &= \sum_{k \in S} \Pr[\mathcal{K}(D') = k] && \text{(by union of elements)} \\ &\leq \sum_{k \in S} e^\delta \Pr[\mathcal{K}(D'') = k] && \text{(by Definition 3)} \\ &= e^\delta \sum_{k \in S} \Pr[\mathcal{K}(D'') = k] && \text{(by distributivity)} \\ &= e^\delta \Pr[\mathcal{K}(D'') \in S] && \text{(by union of elements)} \end{aligned}$$

□

Lemma 2 (Lemma 1 in the paper). *Let \mathcal{K} be a function that gives δ -differential privacy for all pairs of adjacent databases. Given two databases D' and D'' that have the same number of entries and differ in the value of at most one entry, then:*

$$\Pr[\mathcal{K}(D') = k] \leq e^{2\delta} \times \Pr[\mathcal{K}(D'') = k]$$

Proof. Let us call D the common part that D' and D'' share, and let us call r' and r'' the entries in which they differ, in such a way that $D' = D \cup \{r'\}$ and

$$\begin{aligned} \Pr[\mathcal{K}(D \cup \{r'\}) = k] &\leq e^\delta \times \Pr[\mathcal{K}(D) = k] && \text{(by Definition 3)} \\ &\leq e^\delta \times e^\delta \times \Pr[\mathcal{K}(D \cup \{r''\}) = k] && \text{(by Definition 3)} \\ &\leq e^{2\delta} \times \Pr[\mathcal{K}(D'') = k] \end{aligned}$$

□

Theorem 5 (Theorem 2 in the paper). *If a randomized function \mathcal{K} gives δ -differential privacy according to Definition 3, then for every result x^* of the function f the Shannon mutual information between the true answers X (i.e. the results of f) and the reported answers Y (i.e. the results of \mathcal{K}) is bounded by:*

$$I(X; Y) \leq (e^{2\delta} + e^{-2\delta})\delta \log(e) + (e^{2\delta} - e^{-2\delta}) \sum_y p(y|x^*) \log(p(y|x^*))$$

Proof. Let us calculate the Shannon mutual information using the formula $I(X; Y) = H(Y) - X(Y|X)$.

$$\begin{aligned}
H(Y) &= - \sum_y p(y) \log p(y) && \text{(by definition)} \\
&= - \sum_y \left(\sum_x p(x, y) \right) \log \left(\sum_x p(x, y) \right) && \text{(by probability laws)} \\
&= - \sum_y \left(\sum_x p(x) p(y|x) \right) \log \left(\sum_x p(x) p(y|x) \right) && \text{(by probability laws)} \\
&\leq - \sum_y \left(\sum_x p(x) e^{-2\delta} p(y|x^*) \right) \log \left(\sum_x p(x) e^{-2\delta} p(y|x^*) \right) && \text{(by Definition 3 and Lemma 1)} \\
&= - \sum_y e^{-2\delta} p(y|x^*) \left(\sum_x p(x) \right) \log \left(e^{-2\delta} p(y|x^*) \sum_x p(x) \right) \\
&= - \sum_y e^{-2\delta} p(y|x^*) \log(e^{-2\delta} p(y|x^*)) && \text{(by probability laws)} \\
&= - \sum_y (e^{-2\delta} p(y|x^*) \log e^{-2\delta}) - \sum_y (e^{-2\delta} p(y|x^*) \log p(y|x^*)) && \text{(by distributivity)} \\
&= -e^{-2\delta} \log e^{-2\delta} \left(\sum_y p(y|x^*) \right) - \sum_y (e^{-2\delta} p(y|x^*) \log p(y|x^*)) \\
&= \delta e^{-2\delta} \log e - e^{-2\delta} \sum_y p(y|x^*) \log p(y|x^*) && \text{(by probability laws)}
\end{aligned} \tag{6}$$

$$\begin{aligned}
H(Y|X) &= - \sum_x p(x) \sum_y p(y|x) \log p(y|x) && \text{(by definition)} \\
&\geq - \sum_x p(x) \sum_y e^{2\delta} p(y|x^*) \log(e^{2\delta} p(y|x^*)) && \text{(by Definition 3 and Lemma 1)} \\
&= - \left(\sum_y e^{2\delta} p(y|x^*) \log(e^{2\delta} p(y|x^*)) \right) \sum_x p(x) && \text{(by distributivity)} \\
&= - \sum_y e^{2\delta} p(y|x^*) \log(e^{2\delta} p(y|x^*)) && \text{(by probability laws)} \\
&= - \sum_y (e^{2\delta} p(y|x^*) \log(e^{2\delta})) - \sum_y (e^{2\delta} p(y|x^*) \log p(y|x^*)) \\
&= -e^\delta \log e^{2\delta} \left(\sum_y p(y|x^*) \right) - e^{2\delta} \sum_y p(y|x^*) \log p(y|x^*) \\
&= -\delta e^{2\delta} \log e - e^{2\delta} \sum_y p(y|x^*) \log p(y|x^*) && \text{(by probability laws)}
\end{aligned} \tag{7}$$

$$\begin{aligned}
I(X; Y) &= H(Y) - X(Y|X) && \text{(by definition)} \\
&\leq \delta e^{-2\delta} \log e - e^{-2\delta} \sum_x p(y|x^*) \log p(y|x^*) + \\
&\quad 2\delta e^{2\delta} \log e + e^{2\delta} \sum_y p(y|x^*) \log p(y|x^*) && \text{(by Equations 6 and 7)} \\
&= (e^{2\delta} + e^{-2\delta})\delta \log(e) + (e^{2\delta} - e^{-2\delta}) \sum_y p(y|x^*) \log(p(y|x^*)) \text{ (by distributivity)}
\end{aligned}$$

□

Theorem 6 (Theorem 3 in the paper). *If a randomized function \mathcal{K} gives δ -differential privacy according to Definition 3, then the Rényi min mutual information between the true answer of the function X and the reported answer Y is bounded by:*

$$I_\infty(X; Y) \leq 2\delta \log e.$$

Proof. Let us calculate the Rényi mutual information using the formula $I_\infty(X; Y) = H_\infty(X) - X_\infty(X|Y)$.

$$H_\infty(X) = -\log \max_x p(x) \quad \text{(by definition)} \quad (8)$$

$$\begin{aligned}
H_\infty(X|Y) &= -\log \sum_y p(y) \max_x p(x|y) && \text{(by definition)} \\
&= -\log \sum_y \max_x p(y)p(x|y) \\
&= -\log \sum_y \max_x p(x)p(y|x) && \text{(by probability laws)} \\
&\geq -\log \sum_y \max_x p(x)e^{2\delta} p(y|x^*) && \text{(by Definition 3 and Lemma 1)} \\
&= -\log \sum_y e^{2\delta} p(y|x^*) \max_x p(x) \\
&= -\log \left(e^{2\delta} \max_x p(x) \sum_y p(y|x^*) \right) \\
&= -\log \left(e^{2\delta} \max_x p(x) \right) && \text{(by probability laws)} \\
&= -2\delta \log e - \log \max_x p(x)
\end{aligned} \quad (9)$$

$$\begin{aligned}
I_\infty(X; Y) &= H_\infty(X) - H_\infty(X|Y) && \text{(by definition)} \\
&\leq -\log \max_x p(x) + 2\delta \log e + \log \max_x p(x) \text{ (by Equations 8 and 9)} \\
&= 2\delta \log e
\end{aligned}$$

□