

WP. 9
ENGLISH ONLY

**UNITED NATIONS STATISTICAL COMMISSION and
ECONOMIC COMMISSION FOR EUROPE
CONFERENCE OF EUROPEAN STATISTICIANS**

**EUROPEAN COMMISSION
STATISTICAL OFFICE OF THE
EUROPEAN COMMUNITIES (EUROSTAT)**

Joint UNECE/Eurostat work session on statistical data confidentiality
(Bilbao, Spain, 2-4 December 2009)

Topic (ii): Synthetic and hybrid data

DIFFERENTIAL PRIVACY FOR NUMERIC DATA

Invited Paper

Prepared by Rathindra Sarathy, Oklahoma State University and
Krish Muralidhar, University of Kentucky

Differential Privacy for Numeric Data

Rathindra Sarathy* and Krish Muralidhar**

*Spears School of Business, Oklahoma State University, Stillwater, OK 74078, USA

**Gatton College of Business & Economics, University of Kentucky, Lexington, KY 40506, USA

Abstract: The concept of differential privacy has received considerable attention in the literature recently. In this paper we evaluate the masking mechanism based on Laplace noise addition to satisfy differential privacy. The results of this study indicate that the Laplace based noise addition procedure does not satisfy the requirements of differential privacy.

Introduction

The concept of differential privacy is outlined in detail in Dwork (2006), Dwork and Smith (2008) and in many other papers. The primary thrust of differential privacy “captures the increased risk to one’s privacy incurred by participating in a database” or that “any given disclosure will be, within a small multiplicative factor, just as likely whether or not the individual participates in the database.” (Dwork 2006) A perturbation approach to satisfy differential privacy using Laplace noise was also proposed in Dwork (2006). In this paper, for numeric data, we show that satisfying differential privacy in this manner does not assure the intruder’s knowledge gain will be within “a small multiplicative factor”.

An Output Perturbation Method to Satisfy Differential Privacy for Numeric Data

For the purpose of illustration, we will consider the following example although all the results that we have in this section can be extended to any data set (as we show in a later section). Consider a database consisting of 30 observations of the insurance claims made by individuals. The distribution of claims made by 29 of these individuals is normal with a mean of \$10,000 and standard deviation of \$2,000. In contrast to the 29 individuals, the 30th observation is Terry Gross whose house burnt down and made a claim of \$100,000 to rebuild the house.

Assume that an insurance company has the first 29 observations but not Terry Gross. The insurance company is liable to reject Terry Gross’s application if it was revealed that she has made a large insurance claim. Differential privacy is intended to specifically protect Terry Gross whose value is significantly higher than the others in the database. With differential privacy, any query response from this dataset with and without Terry Gross’s information being included in the database should be within ϵ . Thus, inclusion of Terry Gross’s information should not “substantially” affect the results for differential privacy to be satisfied. The data set is provided in Table 1.

Dwork (2006) describes a procedure for satisfying differential privacy for output perturbation of confidential numeric data using Laplace noise addition. Let $\mathbf{X} = (x_1, x_2, \dots, x^*, \dots, x_n)$ represent a set of continuous numeric data. Let $f_n(\mathbf{X})$ represent the response to an arbitrary query on the entire data set \mathbf{X} (consisting of all n observations). One key aspect of the procedure to satisfy differential privacy is to identify the maximum difference between the query response using all n observations ($f_n(\mathbf{X})$) and that using only $(n - 1)$ observations ($f_{n-1}(\mathbf{X})$). Inherently, this captures

the worst-case situation where an intruder has all but one observation in the data set and wishes to compromise the missing observation. Let $\Delta f_n(\mathbf{X})$ represent this difference. That is

$$\Delta f_n(\mathbf{X}) = \text{Max} ||f_n(\mathbf{X}) - f_{n-1}(\mathbf{X})|| = ||f_n(\mathbf{X}) - f_{n-1}(\mathbf{X}^*)||$$

(equation 2, Dwork 2006) where $\mathbf{X}^* = (x_1, x_2, \dots, \cancel{x^*}, \dots, x_n)$ represents $(n - 1)$ observations from the original data from which x^* , the most influential observation for the arbitrary query $f_n(\mathbf{X})$, has been removed. Note that x^* could be different for different queries. Based on the value of $\Delta f_n(\mathbf{X})$, the scale parameter (b_n) for the Laplace distribution is computed as $\Delta f_n(\mathbf{X})/\epsilon$, where ϵ represents the privacy level.

ID	Claims ('000\$)	ID	Claims ('000\$)	ID	Claims ('000\$)
1	9.91	11	13.37	21	7.09
2	9.16	12	8.55	22	10.68
3	10.59	13	11.18	23	9.96
4	11.27	14	9.45	24	9.69
5	10.50	15	12.65	25	9.33
6	11.89	16	9.49	26	10.08
7	10.62	17	11.24	27	10.38
8	12.54	18	8.69	28	12.32
9	9.29	19	9.97	29	9.92
10	8.92	20	10.64	30	100.00

Table 1. Example Data Set

Note that the selection of ϵ is completely independent of all other aspects of the procedure. The selection of ϵ is the equivalent of specifying the level of noise to be added in the perturbation method. Smaller values of ϵ represent greater levels of privacy with $\epsilon = 0$ provides absolute privacy (that is, no knowledge gain from the released output). Although many illustrations of differential privacy have used $\epsilon = 2$, there is no reason that the value of ϵ cannot be lower. In essence the value $\text{Exp}(\epsilon)$ represents the knowledge gain. If $\epsilon = 2$, then this implies that the potential knowledge gain could be a factor of $\text{Exp}(2) \approx 7.3891$. Thus, for example, if the probability based on knowledge of \mathbf{X}^* was 0.10, specifying $\epsilon = 2$ indicates that even if the probability based on knowledge of \mathbf{X}^* and the released response is as high as 0.70, this would still satisfy differential privacy since the improvement is less than a factor of 7.3891. Dwork and Smith (2009, page 3) suggest “We tend to think of ϵ as, say, 0.01, 0.1, or in some cases, $\ln 2$ or $\ln 3$.” Thus, although $\epsilon = 2$ is often used in illustrations, ϵ values greater than 1 may not be advisable in practice.

Based on this information, when the query $f_n(\mathbf{X})$ is issued, privacy mechanism $\kappa(f_n(\mathbf{X}))$ responds with

$$\kappa(f_n(\mathbf{X})) = f_n(\mathbf{X}) + \text{Laplace}^{-1}(u, 0, b_n)$$

where $Laplace^{-1}$ represents the inverse of the Laplace distribution, u is Uniform(0,1), 0 and b_n represent the mean and scale parameter of the Laplace distribution. Let $\kappa(f_n(\mathbf{X}))$ and $\kappa(f_{n-1}(\mathbf{X}))$ represent the masked responses to the same query from $f_n(\mathbf{X})$ and $f_{n-1}(\mathbf{X})$, respectively. The above approach gives differential privacy since,

$$\frac{P[\kappa(f_n(\mathbf{X})) = \tau]/P[\kappa(f_{n-1}(\mathbf{X})) = \tau]}{P[f_n(\mathbf{X}) + Laplace^{-1}(u, 0, b_n)]/P[f_{n-1}(\mathbf{X}) + Laplace^{-1}(u, 0, b_n)]} \leq e^\epsilon.$$

Thus, any response provided by the system using the entire database \mathbf{X} and a subset of the database consisting of $(n - 1)$ observations will be “indistinguishable” within a factor of e^ϵ .

For the purposes of this illustration, we will be using the mean query and $\epsilon = 2$. From the dataset, we can compute the following:

$$f_n(\mathbf{X}) = 13.312, f_{n-1}(\mathbf{X}^*) = 10.323, \Delta f_n(\mathbf{X}) = 2.989, \text{ and } b_n = 1.495.$$

It is easy to see that the “most influential” observation in this dataset is the last observation. For the purposes of illustration and without loss of generality, let us assume that

$$\kappa(f_n(\mathbf{X})) = f_n(\mathbf{X}) + Laplace^{-1}(0.75, 0, 1.495) = 13.312 + 1.036 = 14.348.$$

If we evaluate the following probabilities $p_1 = P[\kappa(f_n(\mathbf{X})) \geq 14.348] = 0.25$ and $p_2 = P[\kappa(f_{n-1}(\mathbf{X})) \geq 14.348] = 0.0338$, then the ratio of $p_1/p_2 \leq \text{Exp}(2)$. We can also generalize these results to any $(n - 1)$ observations and any query and any value of u , thereby satisfying differential privacy.

At first glance, this seems like a very good approach for output perturbation; one that protects an extreme outlier even from the user who has all but the most influential observation. When the probability of the system response is evaluated using the Laplace noise added to the mean of the entire observation set, and subsequently with one observation missing, the ratio of the two probabilities is within the specified limits $\text{Exp}[\epsilon]$. This is the spirit of differential privacy. An intruder who finds that the response is “equally likely within the limits specified” has gained no information about the missing observation. *Conversely, any ratio not within the specified limit $\text{Exp}[\epsilon]$ is interpreted to mean that the intruder has gained some information in violation of differential privacy.* For example, based on the system response, if the intruder can conclude that the missing value is much higher (lower) than the largest (smallest) observation in her dataset with high probability, it would be a violation of differential privacy.

We now proceed to show that the Laplace based output perturbation approach of Dwork (2006) is indeed vulnerable to this kind of differential privacy violation when numeric data is considered. As with the previous examples, this violation is the result of the manner in which the response is evaluated when deciding whether differential privacy has been satisfied. The current evaluation (equation (4), Dwork 2006) is *being performed from the perspective that the entire data set \mathbf{X} is available (which is the perspective of the data provider)*. That is, the system bases its calculation of $\Delta f_n(\mathbf{X})$ on the availability of the entire data set \mathbf{X} and assumes that the intruder would calculate it similarly. However, the information available to the intruder is \mathbf{X}^* . As we

show in the following section, *when we perform the very same evaluation from the information available in X^* (the perspective of the intruder), differential privacy is not satisfied.*

Evaluation of Differential Privacy – The Correct Perspective

Consider an intruder's data set X^* (that is all observations except the last observation) where the missing observation is the most influential observation for a particular query. We use this particular scenario, since it cannot be ruled out. The intruder will attempt to learn information about the missing observation based on the response provided by the system. Using X^* , we can compute $f_{n-1}(X^*) = 10.323$, $f_{n-2}(X^*) = 10.439$, $\Delta f_{n-1}(X^*) = 0.115$, and $b_{n-1} = 0.058$. Note that this information is computed directly from the knowledge of the masking procedure and X^* . No queries are issued to the system. Also note that when the last observation is removed from the data, the next most influential observation is the 21st observation with a value of 7.09.

Now a query is issued to the system regarding the mean of the entire data set. The response from the system as described in the previous section would be $r = 14.348$. The probability (p_3) of observing $r = 14.348$ given X^* can be computed as:

$$p_3 = P[\kappa(f_{n-1}(X^*)) \geq 14.348 | X^*] = P[\kappa(f_{n-1}(X^*)) \geq 14.348 | f_{n-1}(X^*), b_{n-1}] \approx 0.$$

This intruder behavior is perfectly rational. The intruder has knowledge of how the system constructs the response. Based on what the intruder does know and knowledge of b_{n-1} , the intruder can conclude with high probability that the missing observation is much higher than the remaining 29 observations. That is, when the response of 14.348 is provided by the system, *the intruder immediately knows that the remaining observation in the database must be much higher than the remaining 29 observations in X^* since the probability of observing this response given X^* is practically zero.* As discussed earlier, this is exactly the kind of information gain that differential privacy tries to prevent. Yet, the proposed Laplace output perturbation procedure does not prevent this information gain. This information gain becomes obvious when we note that the ratio of $p_1/p_3 > \text{Exp}(2)$. In other words, based on the knowledge of X^* , the appropriate evaluation is p_1/p_3 and not p_1/p_2 . Table 2 provides the ratio of p_1/p_3 for all possible combinations of $f_{n-1}(X)$. Note that *differential privacy is satisfied in every case except when the most influential (30th) observation is missing.*

The problem with the procedure suggested by Dwork (2006) is that the evaluation of differential privacy is performed using b_n and not using b_{n-1} . Note that when the intruder has X^* , knowledge of b_n results in complete disclosure of the missing data value. The primary objective of differential privacy is to protect that one influential observation (Terry Gross) that affects the response to the query. Yet it is precisely when information regarding Terry Gross is missing that the knowledge gain exceeds $\text{Exp}(\epsilon)$. *Assuming that the insurance company has the 29 observations other than Terry Gross, the response from the system immediately identifies that Terry Gross's claim is significantly higher than those of the other 29 individuals in the database. This directly defeats the primary purpose of differential privacy.* The difference between ratios p_1/p_2 (knowledge of b_n) and p_1/p_3 (knowledge of b_{n-1}) are provided in Figures 1a and 1b.

Missing Observation	Ratio of Probabilities	Missing Observation	Ratio of Probabilities	Missing Observation	Ratio of Probabilities
1	0.91	11	0.98	21	0.85
2	0.89	12	0.88	22	0.92
3	0.92	13	0.93	23	0.91
4	0.93	14	0.90	24	0.90
5	0.92	15	0.96	25	0.89
6	0.95	16	0.90	26	0.91
7	0.92	17	0.93	27	0.92
8	0.96	18	0.88	28	0.96
9	0.89	19	0.91	29	0.91
10	0.89	20	0.92	30	> 99999999

Table 2. Ratio of $P[\kappa(f_n(\mathbf{X})) \geq 14.348]$ to $P[\kappa(f_{n-1}(\mathbf{X})) \geq 14.348]$




Figure 1a. Evaluation based on b_n

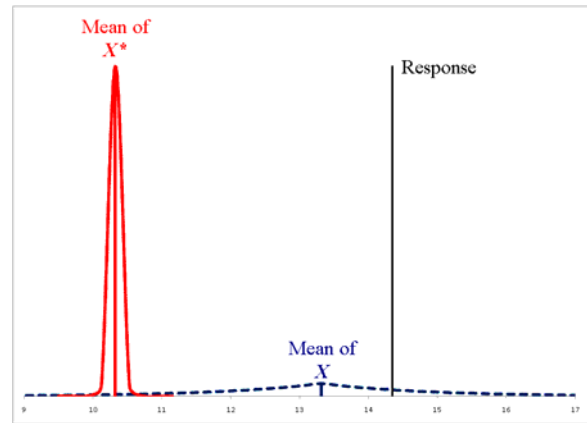


Figure 1b. Evaluation based on b_{n-1}

It is also easy to determine a lower bound for the value of the missing observation by using the following. Given X^* , the 99.99th percentile from the Laplace distribution can be computed as 10.81483 from which we can estimate the missing observation of being at least 25.075. Finally, we can easily show that the results obtained above can also be obtained in a similar fashion when we express the values in the database as 0-1 strings rather than numerical values.

In our illustration, we used a very large value (100.00) for the last observation. However, we can easily show that as long as the last observation is the most influential observation, the resulting ratio of $p_1/p_3 > \text{Exp}(\epsilon) = 7.389$ and hence fails to satisfy differential privacy. Table 3 provides the evaluation of the ratio of p_1/p_3 for different values of the last observation. We assume that the value of $u = 0.75$. It is easy to see that *in every case, the ratio exceeds $\text{Exp}(\epsilon)$* . The interesting aspect is that, *for any data set, when the mean query is issued, there exists a value of x_i and u for which the response from the masking mechanism fails to satisfy differential privacy.*

Value of the Last Observation	Ratio of p_1/p_3	Value of the Last Observation	Ratio of p_1/p_3
14	9	20	925
15	19	25	45140
16	41	30	2201886
17	90	40	5239083064
18	195	50	> 9999999999
19	425	100	> 9999999999

Table 3. The ratio of p_1/p_3 for different values of x_{30}

In fact, we can prove that for any numeric data set \mathbf{X} , there exists a data set consisting of $(n - 1)$ observations \mathbf{X}^ , some query $f_n(\mathbf{X})$, and some response r from the Laplace output perturbation procedure for which we can prove that $P[\kappa(f_n(\mathbf{X})) = r] / P[\kappa(f_{n-1}(\mathbf{X}^*)) = r] \geq \text{Exp}(\epsilon)$. Thus, we can show that for numeric data sets, the Laplace based noise addition approach does not satisfy differential privacy.*

Differential Privacy – A Statistical Perspective

The discussion in the previous section should be adequate to disprove the claims of differential privacy for the Laplace output perturbation procedure. However, in this section, we show that any security mechanism that attempts to achieve security without addressing the underlying statistical properties of the data set will fail.

Consider an insurance company which has 29 observations from the database and is missing only Terry Gross's information. A query is issued to the systems regarding the mean of the entire database resulting as before, with a response of 14.348 as suggested in Dwork (2006) to satisfy differential privacy. The insurance company completely ignores all information regarding the masking procedure and focuses on the data that is available on hand and the statistical properties of the mean. From the available 29 observations, we can compute mean = 10.32 with a standard deviation of 1.377. Using the central limit theorem, we know that the sampling distribution of the sample mean is approximately normal with standard error = $1.377/29^{0.5} = 0.256$. The insurance company simply evaluates the probability of observing the response from the system given this underlying sampling distribution. This evaluation can be performed as follows:

$$P[\kappa(f_n(\mathbf{X})) \geq 14.348 | \text{Normal}(10.32, 0.256)] \approx 0.$$

The insurance company is immediately aware $f_n(\mathbf{X})$ is significantly higher than $f_{n-1}(\mathbf{X}^*)$ and consequently, the last remaining observation must be significantly higher than the other 29 observations. The evaluation of privacy from the provider's perspective and the intruder's perspective are provided in Figures 3a and 3b.

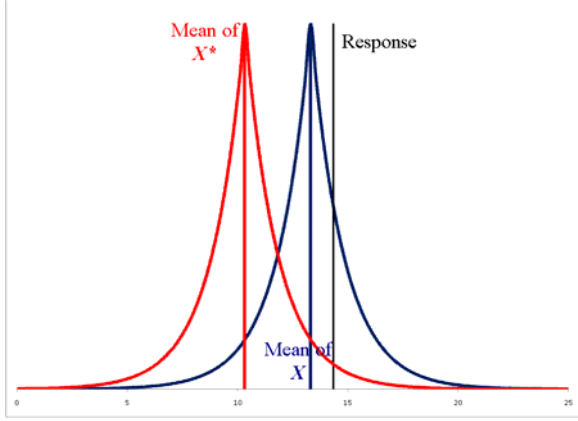


Figure 2a. Evaluation based on Laplace distribution

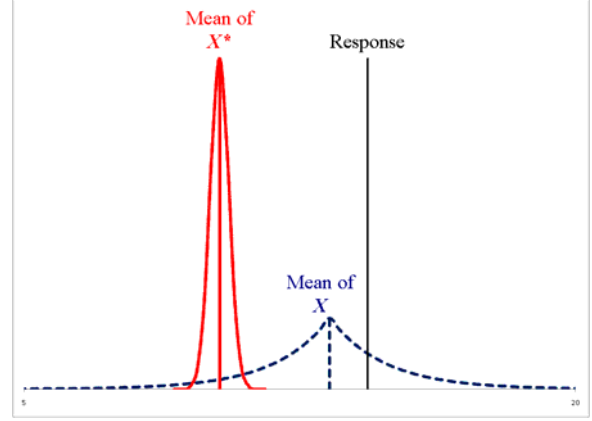


Figure 2b. Evaluation based on the sampling distribution of \bar{x}

The procedure adopted by the intruder here is very similar to basic hypothesis test for the mean that is used in inferential statistics. In essence, the intruder is performing a simple hypothesis test of the mean (using the information in \mathbf{X}^*) to assess whether the response is significantly different from the observed. If the response from the masking mechanism is within the reasonable limits (say $\alpha/2\%$ and $(1 - \alpha/2)\%$), the intruder would conclude that the $f_n(\mathbf{X})$ and $f_{n-1}(\mathbf{X}^*)$ are not significantly different. However, when the response from the masking mechanism is greater (lower) than the specified upper (lower) limit, then the intruder would conclude, with $(1 - \alpha)\%$ confidence, that the value of $f_n(\mathbf{X})$ is significantly higher (lower) than $f_{n-1}(\mathbf{X}^*)$.

A more important aspect is that such inferences are inevitable when the value of the remaining observation is very different from the other observations, which is the very situation that differential privacy purports to protect. No simple noise addition procedure can avoid these inferences, unless the level of noise is so large as to make response to any query completely meaningless. These inferences also cannot be prevented since the intruder is using the distribution of the statistic (in this case \bar{x}) to gain knowledge and does not care about the underlying masking mechanism.

An Even More Basic Problem with the Current Masking Mechanism

There is an even more basic problem with the Laplace noise addition procedure. Consider for the sake of example a data consisting of 10 observations $\mathbf{X} = (1, 2, \dots, 10)$. For the sum query, $f_n(\text{Sum of } \mathbf{X}) = 55$, $f_{n-1}(\text{Sum of } \mathbf{X}^*) = 45$, and $\Delta f_n(\text{Sum of } \mathbf{X}) = 10$. Assuming $\epsilon = 2$, this implies that the scale parameter for the Laplace distribution b_{Sum} must equal 5. The variance of the noise term from this Laplace distribution would equal $(2 \times b_{\text{Sum}}^2) = 50$. The (0.01%, 99.99%) percentile values of the noise term for sum query would range between $(-42.59, 42.59)$.

Now consider the same data for the mean query. For the mean query, $f_n(\text{Mean of } \mathbf{X}) = 5.5$, $f_{n-1}(\text{Mean of } \mathbf{X}^*) = 5$, $\Delta f_n(\text{Mean of } \mathbf{X}) = 0.50$, and $b_{\text{Mean}} = 0.25$. The variance of this Laplace distribution would equal $(2 \times b_{\text{Mean}}^2) = 0.125$. The (0.01%, 99.99%) percentile values of the noise term from this Laplace distribution would range between $(-2.13 \text{ to } 2.13)$. However, an intruder can easily estimate the sum of the data set using the simple transformation $\text{Sum} = n \times \kappa(f_n(\text{Mean of } \mathbf{X}))$. For a data set with $n = 10$ observations, the variance corresponding to this transformation

equal $(2 \times n^2 \times b_{Mean}^2) = 12.5$ with noise range between $(-21.29, 21.29)$. Hence, an intruder using the estimate $Sum = n \times \kappa(f_n(\text{Mean of } \mathbf{X}))$ would estimate the true sum with a far greater level of accuracy than the direct response to the Sum query. In effect, it is as if the intruder is getting responses to queries regarding the Sum from a Laplace distribution with *effective* $b_{Sum} = (n \times b_{Mean})$ which is much smaller than the actual b_{Sum} . Thus, we are left with *the contradiction that an intruder is able to estimate the true value of the sum with far greater accuracy using the mean, but each of the individual queries satisfy the requirements of differential privacy.*

The problem gets much worse if we consider other cases where the data set is shifted by some constant. Say for the sake of argument, that the values of \mathbf{X} have been shifted by a constant 100 to $(101, 102, 103, \dots, 110)$. In this case, $b_{Sum} = 55$ and $b_{Mean} = 0.25$, resulting in *effective* $b_{Sum} = (n \times b_{Mean}) = 12.5$. For the sum query, the (0.01%, 99.99%) percentile values for the noise term from the Laplace distribution is $(-468.45, 468.45)$. Using the mean query, we can estimate the value of the sum within $(-21.29, 21.29)$. Note that shifting the values in \mathbf{X} by a constant 100 did not change the value of b_{Mean} but dramatically altered the value of b_{Sum} . The disparity between b_{Sum} and the *effective* b_{Sum} (and correspondingly the ability of the intruder to estimate the sum with a much higher than intended level of accuracy) increases significantly when values in \mathbf{X} have been shifted by a large constant.

It is possible to view this information from the perspective of the specified ϵ level. From the definition of differential privacy, $b = \Delta f_n(\mathbf{X})/\epsilon$. Since in the above cases, the *effective* b_{Sum} is much smaller than the actual b_{Sum} , we can compute the *effective* ϵ for the sum query as $\Delta f_n(\text{Sum of } \mathbf{X})/(\text{effective } b_{Sum})$. Table 4 provides the results of the analysis for several data sets. The results indicate that the *effective* ϵ for the sum query is far less than the *specified level of* ϵ . As discussed earlier, higher values of ϵ result in lower level of security. Thus, using the mean query to estimate the sum query leads to a much lower level of security than the specified level of security while satisfying the requirements of differential privacy. Note that it is practically impossible to identify all such possible transformations for a given data set.

	$\mathbf{X} =$ (1, 2, ..., 10)		$\mathbf{X} =$ (101, 102, ..., 110)		$\mathbf{X} =$ (50001, 50002, ..., 50010)	
	Sum	Mean	Sum	Mean	Sum	Mean
$f_n(\mathbf{X})$	55	5.5	1055	105.5	500055	50005.5
$f_{n-1}(\mathbf{X}^*)$	45	5	945	105	450045	50005
$\Delta f_n(\mathbf{X})$	10	0.5	110	0.5	50010	0.5
ϵ	2	2	2	2	2	2
b	5	0.25	55	0.25	25005	0.25
<i>Effective</i> b	2.5		2.5		2.5	
<i>Effective</i> ϵ	4		44		20004	

Table 4. Comparison of specified ϵ and *effective* ϵ for the Sum query

There are only two possible conclusions that we could reach based on this obvious contradiction. *The first conclusion is that even though the intruder is allowed to estimate responses to arbitrary queries with a greater than desired level of accuracy, this procedure satisfies differential privacy*

since it satisfies the probability ratios as specified in Dwork (2006). If this is true, then it implies as long as $P[\kappa(f_n(\mathbf{X})) = r]/P[\kappa(f_{n-1}(\mathbf{X})) = r] \leq \text{Exp}(\varepsilon)$, differential privacy is really not concerned with how close the actual response is to the true response. This may really pose serious problems in many practical situations where numerical data needs protection. For example, a perfectly legitimate procedure would be to generate the noise terms from a truncated Laplace distribution such that the noise values are generated as $\text{Laplace}^{-1}(u^*, 0, b_n)$, where $u^* = (0.49999, 0.50001)$. The resulting $\kappa(f_n(\mathbf{X}))$ would be very close to $f_n(\mathbf{X})$ for all queries and all \mathbf{X} , but would still satisfy the requirements of differential privacy. From a practical perspective however, this would not be acceptable since the true values would be released with little or no noise. More importantly, *it may lead to the conclusion that differential privacy is not a viable security standard.*

The only other possible conclusion is that the masking mechanism based on Laplace noise addition does not satisfy differential privacy. In this case, it makes it necessary for us to consider other masking mechanisms that could possibly satisfy differential privacy. In our opinion, this is a far more palatable conclusion since it leaves open the possibility that differential privacy is a viable security standard that could possibly be satisfied by some as yet unknown masking mechanism.

Differential Privacy for Microdata Release

Adopting differential privacy requirements makes it practically impossible to release masked microdata. For the purposes of this discussion, let us assume that Laplace based noise addition does satisfy differential privacy. To release masked microdata that would satisfy differential privacy requirements, it would then be necessary to ensure that *for any query* the response from the masked microdata would satisfy differential privacy. Once released, it is impossible to prevent the user from analyzing the microdata in any way that they choose and *it is impossible to prove that a query for which differential privacy is violated does not exist.* Consequently, *the adoption of differential privacy as the security standard precludes the release of masked microdata.*

Conclusions

The objective of this study was to investigate the application of a masking mechanism based on Laplace noise addition in responding to queries as suggested in Dwork (2006) which is supposed to satisfy the requirements of differential privacy. We are able to show that the intruder's knowledge gain is far higher than the differential privacy limit of $\text{Exp}(\varepsilon)$ even though this mechanism is supposed to satisfy differential privacy. This contradiction can be easily explained by the fact that that evaluation of whether differential privacy is satisfied is performed from the perspective of the data provider and not the intruder. When we consider the true knowledge gain from the intruder's perspective, we find that the resulting knowledge gain is far higher than $\text{Exp}(\varepsilon)$ in violation of the requirements of differential privacy. Such knowledge gain also occurs when the intruder uses basic concepts from inferential statistics to infer the missing observation. We also identified a far more serious issue relating to the masking mechanism based on Laplace noise addition. These results show that an intruder will be able to gain very accurate estimates of the true value based on simple, commonly used numerical transformations. We illustrated this issue by showing that an intruder would be able to gain extremely accurate estimates of the sum of the data set using the mean query.

In summary, the results of this paper lead us to one of only two possible conclusions:

- (1) The masking mechanism based on Laplace noise addition satisfies differential privacy in spite of these issues. If this is the case, then we have to conclude that, the whole concept of differential privacy fails as a security standard since satisfying differential privacy does not guarantee the promised level of security.
- (2) The masking mechanism based on Laplace noise addition does not satisfy differential privacy. If this is the case, then we have to search for alternative masking mechanism that would satisfy differential privacy. At this point then, differential privacy is a security standard which like Dalenius' definition of privacy is unachievable in practice.

References

1. Dwork, C. (2006). Differential Privacy. M. Bugliesi et al. (Eds.): ICALP 2006, Part II, LNCS 4052, Springer-Verlag Berlin Heidelberg, 1-12.
2. Dwork, C. and A. Smith (2009). Differential Privacy for Statistics: What we Know and What we Want to Learn. NISS/NCHS workshop on Data Confidentiality: The Next Five Years, Hyattsville, MD.