TRANSACTIONS ON DATA PRIVACY 4 (2011) 1 —17

Evaluating Laplace Noise Addition to Satisfy Differential Privacy for Numeric Data

Rathindra Sarathy* and Krishnamurty Muralidhar**
*Spears School of Business
Oklahoma State University
Stillwater OK 74078, USA
E-mail: rathin.sarathy@okstate.edu
** Gatton College of Business & Economics
University of Kentucky
Lexington KY 40506, USA
E-mail: krishm@uky.edu

Abstract. Laplace noise addition is often advanced as an approach for satisfying differential privacy. There have been several illustrations of the application of Laplace noise addition for count data, but no evaluation of its performance for numeric data. In this study we evaluate the privacy and utility performance of Laplace noise addition for numeric data. Our results indicate that Laplace noise addition delivers the promised level of privacy only by adding a large quantity of noise for even relatively large subsets. Because of this, even for simple mean queries, the responses for a masking mechanism that uses Laplace noise addition is of little value. We also show that Laplace noise addition may be vulnerable to a tracker attack. In order to avoid this, it may be necessary to increase the variance of the noise added as a function of the number of queries issued. This implies that the utility of the responses would be further reduced. These results raise serious questions regarding the viability of Laplace based noise addition for masking numeric data.

1 Introduction

Differential privacy [8][9] is a new privacy standard that has garnered considerable attention in the literature in recent years. Differential privacy attempts to ensure that even if an intruder has knowledge of all but one observation in a data set, an intruder, using the response to any query from the dataset should not be able to distinguish between the presence or absence of an individual. In this sense, differential privacy not only attempts to protect individuals who are

present in a particular data set, but all individuals in the universe of possible values for that data set. In addition, since differential privacy applies to any query, it is applicable to both count data as well as numeric data.

Laplace noise addition is the primary method that has been advanced for satisfying the differential privacy. There have been several applications of using Laplace noise addition for count data. However, there have been few illustrations, and no evaluation, of the application of Laplace noise addition to numeric data. The objective of this study is to provide some such an evaluation.

This paper is organized as follows. The second section provides a brief introduction into the concept of differential privacy and Laplace noise addition. In the third section, we evaluate the utility of the responses resulting from a Laplace noise addition mechanism to numeric data. In the fourth section, we provide an illustration of the disclosure risk resulting from a Laplace noise addition mechanism used to protect numeric data. In the final section, we provide the conclusions.

2 Differential Privacy

Differential privacy is simply a privacy standard whereby the response to any query including or excluding a particular observation is indistinguishable in a probabilistic sense. Note that differential privacy operates in the realm of output perturbation where it is assumed that the responses will be computed from the original data and the resulting output will then be perturbed. The alternative approach is source data perturbation where the original data is perturbed and all responses are computed from the perturbed data.

In general terms, the requirement for differential privacy can be described as follows [8][9]. Consider *any two possible datasets D_1 and D_2, from a population of values $\mathcal{D}$, that differ by exactly one record*. Let $\kappa_f(D_1)$ and $\kappa_f(D_2)$ be the responses from datasets D_1 and D_2, respectively, where $\kappa_f()$ represents a mechanism used to respond to an arbitrary query f(). For $\kappa_f()$ to satisfy differential privacy, it is necessary that

$$e^{-\varepsilon} \leq \frac{P[\kappa_f(D_1) = R]}{P[\kappa_f(D_2) = R]} \leq e^{\varepsilon}. \quad (1)$$

where R represents the response.[1] Abowd and Vilhuber [1] interpret the the ratio in the middle as the "knowledge gain ratio" for an intruder from one version

[1] Note that for continuous data, the ratio would be expressed as inequalities.

of the data set (D_1) over the other (D_2) assuming, without loss of generality, that the numerator is always larger. Differential privacy requires that the knowledge gain ratio be limited to e^{ε}. The ratio is better interpreted as an indistinguishability ratio. The larger the ratio, the greater the probability that the response was obtained from one data set than the other. As a decision tool, an intruder would clearly favor the data set with the larger probability of providing the specific perturbed response, even for moderately large values of the ratio such as 2. Note that differential privacy is predicated on a query/response framework. Since it would be very difficult (if not impossible) to show that differential privacy will be satisfied for every type of analyses that an intruder might employ when microdata is released, adoption of differential privacy precludes the release of masked microdata.

We would also like to briefly address the use of the term "ε differential privacy" when referring to a procedure that satisfies equation (1). In reality, a procedure that satisfies the ratio in equation (1) actually provides "e^{ε} differential privacy" since the intruder's knowledge gain is e^{ε}. In our opinion, a procedure that satisfies "ε differential privacy" should satisfy the requirement specified by the following [3]:

$$\frac{P[\kappa_f(D_1) = R]}{P[\kappa_f(D_2) = R]} \leq 1 + \varepsilon. \tag{2}$$

It is true that when ε is zero or very small, $e^{\varepsilon} \approx (1 + \varepsilon)$. However, the two measures diverge quite quickly. Even when ε = 0.20, e^{ε} exceeds (1 + ε) by at least a non-negligible 10%. It is easy to see that, as ε increases, the two values diverge considerably. Considering that we see specifications of ε much larger than 0.20, special care should be exercised to ensure that users do not overestimate the privacy provided by equation (1). For example when "ε = 2" as in [1], the true ratio equals $e^2 = 7.389$. Similarly, if the "overall ε for the procedure was 8.6" [11], the true ratio is $e^{8.6} = 5432$. In order to avoid any confusion, we suggest the use of the term "e^{ε} differential privacy" when privacy is evaluated as shown in equation (1) and the term "ε differential privacy" when privacy is evaluated using equation (2). Since Dwork [8] defines privacy as shown in equation (1), for the remainder of the paper, we will be using the term "e^{ε}-differential privacy."

In order to satisfy differential privacy, Dwork [8] suggests the use of Laplace based noise addition, again assuming a query/response situation. Assume that the intruder issues the query f(X) on a data set X for which the true response is *a*. Let a differential privacy satisfying mechanism $\kappa_f()$ be implemented for this data set and that the response from the system is R. Dwork [8] suggests that the masked response $\kappa_f(X) = R = a + y$ where y represents a noise term from a Laplace distribution with mean 0 and scale parameter $b = \Delta f/\varepsilon$ where *Δf represents*

the maximum difference in the value of f(X) when exactly one input to X is changed. This accounts for the situation, for example, when the intruder's data differs from that of the data set X by exactly one record.

The value of Δf represents the global sensitivity for the query f(X). To determine Δf one must consider all possible values for D_1 and D_2 in the population of values $\mathcal{D}$ and not just the specific values that may exist in the current dataset X that is being protected. Hence, in order to implement the globally sensitive version of the Laplace noise addition procedure, it is necessary to determine the value of Δf.

In binary databases where the queries are always assumed to be "count" queries, it is easy to see that the value of $\Delta f = 1$ since the maximum difference in the count between D_1 and D_2 is always 1. For numeric data, it is not even certain that we can determine the global sensitivity for an arbitrary dataset $\mathcal{D}$ and an arbitrary query f(X). Consider for instance, a numerical variable such as insurance claim that was used by Sarathy and Muralidhar (2009). Let us also assume that a simple sum query was issued. Since we have to protect the universe of possible income values, it would be necessary to determine the global sensitivity of the insurance claim variable for the sum query. But there is no simple approach to do this. Even assuming that insurance claim is a positive variable, we now have to answer the question "What is the largest possible insurance claim that could exist in the universe of insurance claim values?" Without answering this question, it is simply impossible to implement any procedure that satisfies differential privacy (not even Laplace based noise addition).)[2]. Wasserman and Zhou [14] also note this issue when they state that "In particular, it is difficult to extend differential privacy to unbounded domains."

In summary, in order to satisfy differential privacy, *it is necessary that the upper and lower bounds on the values in the database exist and are known. Without this knowledge, it is impossible to compute Δf for an arbitrary function f() and impossible to implement the Laplace based noise addition procedure to satisfy differential privacy.* Arbitrary approaches such as bottom and top coding may be a convenient solution, but in our opinion, defeats the very purpose of differential privacy whose primary objective is to protect even extreme values while simultaneously providing meaningful responses. Even in cases where there are clearly defined upper and lower bounds, there remains the question of the utility of the re-

[2] In evaluating whether differential privacy is satisfied, we only consider the original definition of differential privacy [8]. We do not consider any relaxations such as found in [12] as satisfying true e^{ε} differential privacy.

sponses resulting from a Laplace based masking mechanism. In the following section, we investigate this issue.

3 Utility of Responses from a Masking Mechanism using Laplace Noise Addition

The concept of differential privacy was motivated based on potential privacy breaches through auxiliary information using the example of Terry Gross' height [8] or Turing's height [9] as follows:

> Suppose one's exact height were considered a highly sensitive piece of information, and that revealing the exact height of an individual were a privacy breach. Assume that the database yields the average heights of women of different nationalities. An adversary who has access to the statistical database and the auxiliary information "Terry Gross is two inches shorter than the average Lithuanian woman" learns Terry Gross' height, while anyone learning only the auxiliary information, without access to the average heights, learns relatively little.

Dwork [8] then goes on to provide describe differential privacy and the Laplace based noise addition method to achieve the same. Although never explicitly stated, this illustration leaves the impression that the Laplace based noise addition would protect Terry Gross. But we never actually see the implications of using Laplace based noise addition and the level of protection it offers Terry Gross. In this study, we illustrate this issue by using a similar definition in a slightly different context.

For this illustration, rather than use height of women, we use a variable (income) that is usually considered sensitive. Consider the situation where the following auxiliary information is available: "Mr. Gold Sack's income is $5 million more than the average American." It is well known that the income of some hedge fund managers exceed $1 billion [13]. In order to protect such individuals, it is necessary that Δf must be at least 1 billion. Note that, in order to satisfy differential privacy, it is better to be conservative in estimating Δf. For the purposes of this illustration, let us assume that for the sum query, Δf = 1,000,000,000. For this illustration, all information was gathered from the 2006-2008 American Community Survey at the U.S. Census Bureau web site (http://factfinder.census.gov) [2]. We assume that the responses from this web site represent the true values although it is likely that these values have themselves been masked prior to release. These "true values" are masked using La-

place noise addition prior to release. In evaluating the responses from Laplace noise addition, we consider three differential levels of privacy by specifying three levels of $\varepsilon = 0.10, 0.25, 1.00$ corresponding to a maximum knowledge gain of approximately 10%, 28%, and 172%, respectively.

Our evaluation of the Laplace noise addition is restricted to the simple mean query. The response may be obtained either as (Response to the Sum query/n) or as a response to a direct Mean query with noise generated from Laplace(0, $\Delta f/\varepsilon n$). n represents the number of records in the query and Δf represents the global sensitivity for the Sum query. Either approach would result in exactly the same response distribution for the Mean query. We conducted the analysis at two different levels: National (entire USA), State and County (Whitley County, Kentucky).

3.1 Protecting Mr. Gold Sack's Information

The primary objective of this analysis is to prevent disclosure of Mr. Sack's income. Consistent with the Terry Gross example [8], we have the auxiliary information that Mr. Sack earns $5 million more than the "average American." From the US Census Bureau web site, the average per capita income for 2008 is provided as $27,466 (n = 143,195,793 and standard deviation ≈ 275000). Thus, if the auxiliary information is correct, we know that Mr.Sack's income is $5,027,466.

Now let us consider the impact of Laplace noise addition on this data. Note that the scale parameter b of the Laplace distribution for the Mean query is given by $\Delta f/\varepsilon n$ and corresponding to $\varepsilon = 0.10, 0.25$, and 1.00, the values of b = 69.83, 27.93, and 6.98, respectively. The variance of the noise added corresponding to $\varepsilon = 0.10, 0.25$, and 1.00, equals 9750, 1560, and 97, respectively. The actual variance of the income is approximately 275000^2 or (7.56×10^{10}). In other words, considering the variance of the data set, *the Laplace noise added is negligible*. The impact of adding so little noise is evident when we consider the protection afforded to Mr. Sack.

Even when $\varepsilon = 0.10$ (the highest privacy level), in 99% of the cases, the noise added is less than $275. In other words, 99% of the responses would provide the intruder with an estimate of Mr. Sack's income that is within $275. Given that Mr. Gold Sack's income is over $5 million, the privacy afforded to Mr. Sack is very small. The situation gets worse when we consider $\varepsilon = 0.25$ in which case the intruder is able to estimate the income to within $110 in 99% of the cases. When $\varepsilon = 1.00$, the intruder is able to estimate Mr. Sack's income to within *$30 of the true value*. The distribution of Mr. Sack's income based on the responses from Laplace noise addition for all three levels of ε are provided in Figure 1 which shows that an intruder gets an accurate estimate of Mr. Sack's income even after Laplace noise is added. Thus, when the intruder has the auxiliary information

that Mr. Gold Sack earns $5 million than the average American, adding Laplace noise provides little protection against disclosure of the true income.

Figure 1. Distribution of Responses for Mean Income (Entire US)

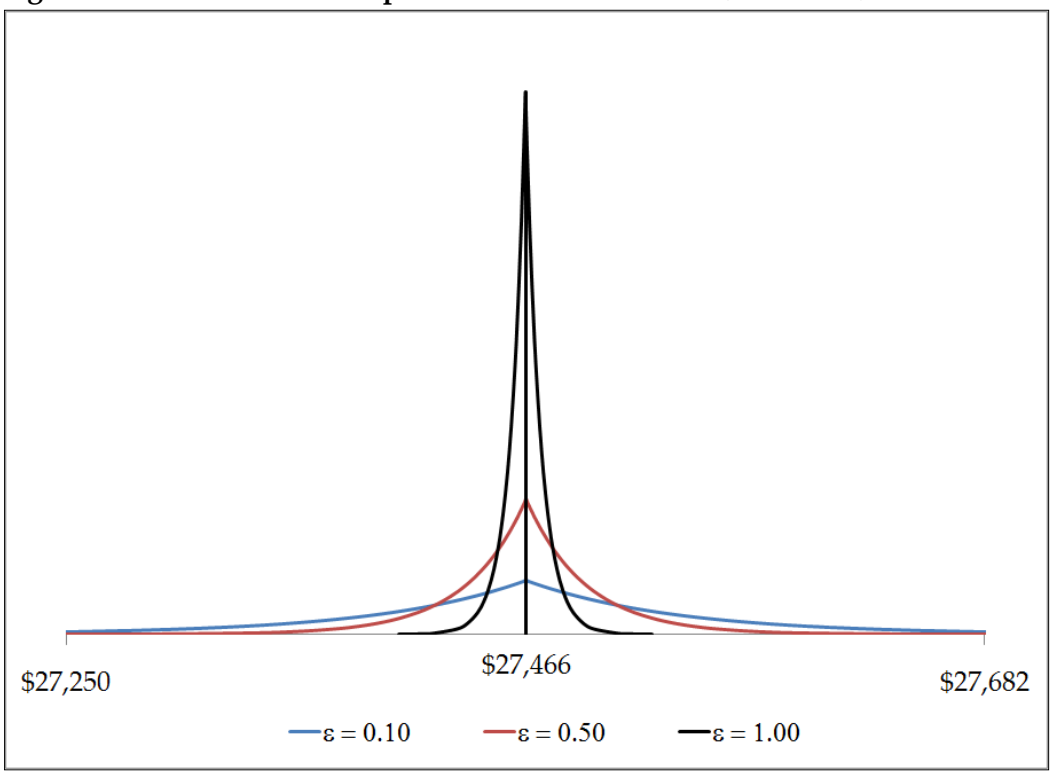

3.2 Analysis of Data Subsets

Data such as income are often analyzed not just at the entire data set level, but also for subsets of data. One simple illustration of this is the consideration of the average income of a county (Whitley County) in Kentucky. From the Census Bureau web site, we find that the average income of Whitley county is $16,708 (n = 12,685 and standard deviation ≈ 134,000). Given that this subset has 12,685 observations, one would expect that the response to the average income of this county would be very accurate. With Laplace noise addition, this is not the case.

When ε = 0.10, the variance of the noise added is over 7000% of the variance of the original data. For ε = 0.50 and 1.00, the variance of the noise added as a proportion of the variance of the data set are over 1000% and 72%, respectively. This implies that even for very low privacy levels (ε = 1.00), the probability of observing a response that is within $\pm$ $10,000 is less than 12%. In 88% of the cas-

es, the Laplace noise added response would be outside the $10,000 range. When ε = 0.25 and 0.10, the probability of observing within ±$10,000 reduces to approximately 3% and 1%, respectively. This implies that when ε = 0.10, 99% of the responses to the query "What is the mean income of individuals in Whitley county?" would be outside ±$10,000 of the true income. In other words, in most cases, the responses are of little value for analytical purposes since they are so different from the true income.

Figure 2 provides the response distribution for the mean income of Whitley county for all three ε levels. The wide range of the response distribution means that the Laplace noise added responses will be of little use to the legitimate user. This response distribution is particularly striking when compared with the response distribution in Figure 1. At the subset level, the distribution varies between approximately ±$3 million while at the overall dataset level (making the response useless for a legitimate user), the responses vary only by a meager ±$250 (providing little protection for Mr. Gold Sack).

Figure 2. Distribution of Responses for Mean Income (Whitley County)

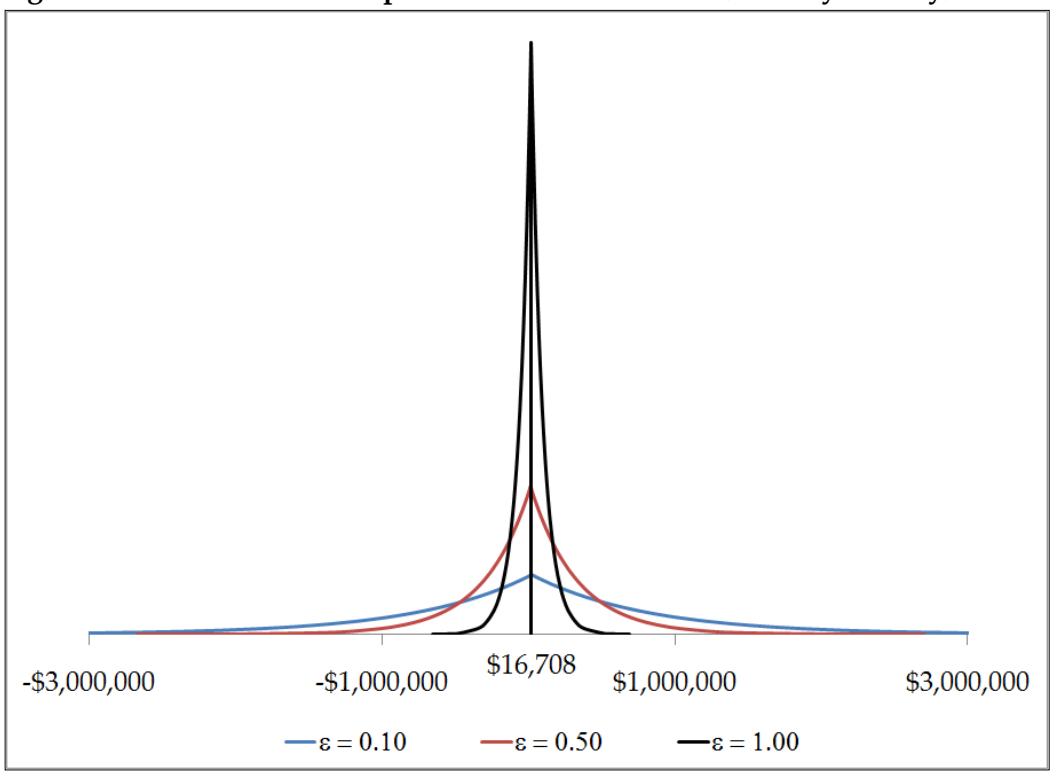

3.3 Summary of Analysis

The analysis in this section addresses two important aspect of the Laplace noise addition procedure. It shows that when n is large (such as the case when the

entire US population is considered), the privacy that is offered is very small. Consequently, when an intruder has the auxiliary information "Mr. Gold Sack earns $5 million more than the average American," Laplace noise addition offers very little protection against disclosure of confidential information. However, using the same parameters, when we consider smaller subsets, the level of noise added is so large as to make practically all responses from the system useless for the legitimate user. These weaknesses arise from the fact that that unlike traditional noise addition where the noise added is proportional to the variance in the data set, the noise added to satisfy differential privacy is independent of the actual data set and is based only on the value of Δf. In particular, this reflects an inherent characteristic of differential privacy in that it seeks to protect even extreme values. When the data set is skewed resulting in very large Δf, the noise variance can be very large compared to the variance of the data set. Consequently, the level of noise added to satisfy differential privacy can be of orders of magnitude greater than the variance in the data set, reducing the utility of the responses. The result is little or no protection for large data sets, while for small data sets the noise added is so large that the responses are practically useless. In contrast, traditional noise addition methods would only add noise that is proportional to the actual variance of the data set that is likely to provide more meaningful responses. The trade-off is that they would not satisfy the requirements of differential privacy.

4 Evaluating Laplace Noise Addition for Multiple Queries

For any procedure that satisfies differential privacy, it is assumed that since the intruder's knowledge gain is limited to the probability ratio in equation (1), disclosure is avoided. However, this evaluation relates to a single response and does not consider the case of a tracker who issues multiple queries and uses the responses from these queries to disclose one or more observations [7]. In this section, we show that the Laplace noise addition mechanism to satisfy differential privacy may be vulnerable to a tracker attack.

For the purpose of this illustration, let us assume that the data set consists of real numbers between the limits 0 and 1. For simplicity and without loss of generality, we will limit our discussion to the Sum query and $\varepsilon = 1$. For this data, set the global sensitivity $\Delta f = 1$. In order to satisfy e^{ε}-differential privacy, the noise term for the Sum queries must be generated from a Laplace distribution with mean 0, scale parameter $b = \Delta f/\varepsilon = 1$ (and resulting noise variance = $2b^2 = 2$). Consistent with the description in [8], we assume that the details of the perturbation (ε, Δf, etc.) are provided to the users. We generated a dataset $\mathbf{X}$ with 50

observations and we assume that the intruder has 49 observations (2, 3, …, 50) in the dataset (Table 1). The intruder's objective is to estimate the unknown observation x_1.

Note that the variance of the entire data set provided in Table 1 is 0.09339. However, in order to satisfy differential privacy, it is necessary that the variance of the noise distribution be 2. As noted earlier, unlike traditional noise addition where the noise added is proportional to the variance of the data set, the noise added to satisfy differential privacy can be of orders of magnitude greater than the variance in the data set, especially for skewed data sets, reducing the utility of the responses. In addition, since the range of the variable **X** is between (0, 1), the user knows that the sum of $(x_i + x_j)$ must be in the range (0, 2). However, due to the large noise variance and the fact that the Laplace distribution is unbounded, a large proportion of the responses (21 out of 49) are outside the meaningful range of (0, 2). This is a problem for most datasets.

The intruder issues the following series of queries $(x_1 + x_2)$, $(x_1 + x_3)$, …, $(x_1 + x_{50})$ resulting in a total of 49 queries. Let $a_2, a_3, \ldots, a_{50}$ represent the true response to the queries, respectively. Let $y_2, y_3, \ldots, y_{50}$ represent the noise terms generated from a Laplace(0,b) to the queries. Let $R_2 = (x_1 + x_2) + y_2$, $R_3 = (x_1 + x_3) + y_3$, …, $R_{50} = (x_1 + x_{50}) + y_{50}$ represent the responses from the system to the queries. These values are provided in Table 1 as well. Since the intruder knows the true values of $x_2, x_3, \ldots, x_{50}$, the intruder can simply subtract the respective value from the response to result in an estimate of x_1 as follows:

$$\hat{x}_1^i = R_i - x_i, \quad i=2, 3, \ldots, 50. \tag{3}$$

For example, $R_2 = 1.88222$, $x_2 = 0.85490$, and hence $\hat{x}_1^2 = 1.02732$. Table 1 provides the results of all 49 queries. The resulting variable from this process $\hat{x}_1$ is an iid Laplace(x_1,b) random variable.

Now consider $\bar{x}_1$the mean of $(\hat{x}_1)$. Let q represent the number of queries issued. In this illustration q = 49. From the central limit theorem, when q is large, we know that $\bar{x}_1 \sim$ Normal(x_1, $2b^2/q$). Even when q is small, since $\hat{x}_1$ is IID Laplace(x_1, b), we know that the mean and variance of $\bar{x}_1$ are x_1 and $2b^2/q$, respectively, although we do not know the exact distribution of $\bar{x}_1$. Thus, even when q is small, the variance of the resulting estimate is smaller by a factor of q compared to the variance of the original Laplace distribution.

The intruder now simply estimates the true value of x_1 as the mean of $\hat{x}_1^i, i = 2, 3, \ldots, 50$. From the data in Table 1, we know that $\bar{x}_1^{Est} = \sum \frac{\hat{x}_1^i}{49} = 0.97262$. Now consider the following probabilities,

$$P\left[\bar{x}_1^{Est} \geq 0.97262 | x_1 = 1\right] = P\left[\bar{x}_1^{Est} \geq 0.97262 | \text{Normal}\left(1, \frac{2}{49}\right)\right] = 0.576 \tag{4}$$

$$P\left[\bar{x}_1^{Est} \geq 0.97262 | x_1 = 0\right] = P\left[\bar{x}_1^{Est} \geq 0.97262 | \text{Normal}\left(0, \frac{2}{49}\right)\right] = (5 \times 10^{-12}). \tag{5}$$

Table 1. Data and computations for the example

Individual	x	$(x_1 + x_i)$	Random #	y_i	$R_i = (x_1 + x_i) + y_i$	Estimate of x_1
1	0.97032					
2	0.85490	1.82522	0.52770	0.05700	1.88222	1.02732
3	0.72936	1.69968	0.07910	-1.84387	-0.14419	-0.87355
4	0.06435	1.03467	0.43122	-0.14799	0.88668	0.82233
5	0.42397	1.39429	0.79758	0.90427	2.29856	1.87459
6	0.75934	1.72966	0.89064	1.52000	3.24966	2.49032
7	0.67422	1.64454	0.52139	0.04372	1.68826	1.01404
8	0.62075	1.59107	0.93254	2.00312	3.59419	2.97344
9	0.66039	1.63071	0.26885	-0.62044	1.01027	0.34988
10	0.54600	1.51632	0.70928	0.54225	2.05857	1.51257
11	0.22039	1.19071	0.50533	0.01072	1.20143	0.98104
12	0.98132	1.95164	0.56455	0.13822	2.08986	1.10854
13	0.22174	1.19206	0.76173	0.74119	1.93325	1.71151
14	0.88548	1.85580	0.92787	1.93619	3.79199	2.90651
15	0.95191	1.92223	0.18962	-0.96957	0.95266	0.00075
16	0.65780	1.62812	0.80940	0.96445	2.59257	1.93477
17	0.88826	1.85858	0.63536	0.31569	2.17427	1.28601
18	0.74429	1.71461	0.73987	0.65344	2.36805	1.62376
19	0.12368	1.09400	0.35452	-0.34386	0.75014	0.62646
20	0.59708	1.56740	0.10678	-1.54385	0.02355	-0.57353
21	0.03746	1.00778	0.76936	0.77376	1.78154	1.74408
22	0.82311	1.79343	0.14070	-1.26801	0.52542	-0.29769
23	0.03147	1.00179	0.81132	0.97456	1.97635	1.94488
24	0.32822	1.29854	0.40197	-0.21823	1.08031	0.75209
25	0.20763	1.17795	0.87744	1.40601	2.58396	2.37633
26	0.57210	1.54242	0.50602	0.01211	1.55453	0.98243
27	0.66724	1.63756	0.07235	-1.93313	-0.29557	-0.96281
28	0.36904	1.33936	0.26919	-0.61921	0.72015	0.35111
29	0.83805	1.80837	0.79157	0.87498	2.68335	1.84530
30	0.72112	1.69144	0.85190	1.21672	2.90816	2.18704
31	0.98357	1.95389	0.83221	1.09189	3.04578	2.06221
32	0.23028	1.20060	0.56917	0.14890	1.34950	1.11922
33	0.09613	1.06645	0.04275	-2.45934	-1.39289	-1.48902
34	0.00538	0.97570	0.13981	-1.27435	-0.29865	-0.30403
35	0.46984	1.44016	0.56482	0.13886	1.57902	1.10918
36	0.96043	1.93075	0.31687	-0.45610	1.47465	0.51422
37	0.20283	1.17315	0.22134	-0.81491	0.35824	0.15541
38	0.60845	1.57877	0.30995	-0.47820	1.10057	0.49212
39	0.45104	1.42136	0.13421	-1.31520	0.10616	-0.34488
40	0.63254	1.60286	0.95061	2.31488	3.91774	3.28520
41	0.49287	1.46319	0.62664	0.29205	1.75524	1.26237
42	0.47326	1.44358	0.87987	1.42603	2.86961	2.39635
43	0.87437	1.84469	0.17758	-1.03516	0.80953	-0.06484
44	0.01190	0.98222	0.86866	1.33681	2.31903	2.30713
45	0.89823	1.86855	0.53223	0.06663	1.93518	1.03695
46	0.54942	1.51974	0.42995	-0.15095	1.36879	0.81937
47	0.03274	1.00306	0.03394	-2.68988	-1.68682	-1.71956
48	0.61882	1.58914	0.36316	-0.31976	1.26938	0.65056
49	0.30366	1.27398	0.17065	-1.07496	0.19902	-0.10464
50	0.33108	1.30140	0.40338	-0.21474	1.08666	0.75558
					Average Of Estimates	0.97262

If we now consider the ratio of the two probabilities above, we get $\frac{0.575992308}{(5 \times 10^{-12})} = (1.17 \times 10^{11}) \gg e^1 = 2.7182$. Thus, the ratio of the two probabilities does not satisfy the requirements of e^{ε}-differential privacy (but does satisfy $e^{q\varepsilon}$-differential privacy as we discuss later).

Note that 0.97262 is an excellent point estimate of x_1 (0.97032) for such a relatively small sample size of 49. For more realistic situations such estimates are expected to be very close to the true value and would result in extremely sharp interval estimates with high confidence. By most standards of statistical disclosure control, this would be considered an unacceptable breach of confidentiality and privacy. Even with only 50 observations, if we assume that the intruder has 49 observations, the intruder is not limited to 49 queries. The intruder can also issue all possible combinations of queries involving x_1 and the remaining known observations. One such possible query is $(x_1 + x_2 + x_3)$. From the response to this query, the intruder can get the estimate of x_1 simply as the Response $- (x_2 + x_3)$. Even when n is relatively small, the intruder can issue a very large number of queries to the system in this manner, resulting in very large q resulting in $(2b^2/q) \to 0$ and $\bar{x}_1^{Est} \to x_1$. Thus, with increasing q, the intruder gets a very accurate estimate of the true value of x_1. Since this result is true for x_1, similarly we can show it to be true for any value x_1.

In summary, when the intruder has information regarding the (n – 1) observations, they can use this information to issue a series of (tracker) queries in order to estimate the value of the missing observation with a great deal of accuracy. This type of phenomena has been addressed previously in the statistical disclosure limitation literature [4][5][7], and others. It is interesting that this result is also consistent with the observations of Dinur and Nissim [6] who showed that an intruder, *with no prior information*, given an unlimited number of queries, can reconstruct the value of the entire database. What these results indicate is that, when we assume that the intruder has (n – 1) of the n observations, then *only the first query will provide the desired level of privacy. All subsequent queries will result in a reduction in privacy*.

It should be noted that Dwork [8] and other researchers have recognized the tracker problem. In a recent paper, Dwork and Smith [10, page 139] acknowledge the issue with multiple queries when they observe that:

> Differential privacy applies equally well to an interactive process, in which an adversary adaptively questions the curator about the data. The probability K(S) then depends on the adversary's strategy, so the definition becomes more delicate. However, one can prove that if the algorithm used to answer each question is ε-differentially private, and the adversary asks q questions, then

> the resulting process is qε-differentially private, no matter what the adversary's strategy is.

This is precisely the result that was illustrated in this section. The implications of this statement are far reaching than what it seems at first glance. Assume that differential privacy based Laplace noise addition has been implemented on a data set and some ε has been specified. The above statement implies that the intruder's knowledge gain for the very first query is e^{ε}; for the second query, it is $e^{2\varepsilon}$; … for the q^{th} query, it is $e^{q\varepsilon}$. In other words, the intruder's knowledge gain increases exponentially with the number of queries. *Consequently, after just a few queries, the intruder's knowledge gain is so large that differential privacy based Laplace noise addition procedure offers no privacy at all.*

One potential solution to eliminate the tracker problem is to increase the variance of the Laplace noise to compensate for the intruder's increase in knowledge. Since the privacy provided for the q^{th} query is $e^{q\varepsilon}$, in order to achieve the same privacy level for the q^{th} query as for the first query, it would be necessary that the scale parameter of the Laplace distribution for the q^{th} query equal $(q \times b)$, with resulting variance equal to $(q^2 \times 2b^2)$. In other words, in our current example, for the 10^{th} query, the variance of the noise added would be 200 units; and for the 50^{th} query, the noise variance would be 5000. Adding noise with variance of 5000 (or even 200) when the variance of the actual data set is only 0.09339 makes the query responses practically meaningless. Table 2 shows the impact of increasing the noise variance to account for the reduction in privacy. In generating this data, we have used exactly the same random numbers to generate the noise in this table as we did in Table 1. For instance, for the query sum of $(x_1 + x_{47})$ where the true sum is 1.00306, we get a masked response of -122.7312. Similarly, for the sum of $(x_1 + x_{40})$ where the true sum is 1.60286, the response from the system is 91.8833. As observed earlier, since $\mathbf{X}$ is in the range (0, 1), the sum of $(x_i + x_j)$ must be in the range (0, 2). However, we observe in Table 2 that as the number of queries increases practically none of the responses fall in the meaningful range of (0, 2). Out of the total of 49 queries, only 5 fall in the meaningful range. For any intelligent user who knows that the sum of two observations must be in the range (0, 2), practically all the responses from the system after the first few queries are useless. Hence, as shown in Table 2, increasing the variance as the number queries increases may maintain privacy, but makes the responses practically useless, and hence is simply not a feasible approach.

Table 2. Responses from the system with increasing noise variance

Individual	x	$(x_1 + x_i)$	Random #	y_i	R = $(x_1 + x_i) + y_i$
1	0.97032				
2	0.85490	1.82522	0.52770	0.0570	1.8822
3	0.72936	1.69968	0.07910	-3.6877	-1.9881
4	0.06435	1.03467	0.43122	-0.4440	0.5907
5	0.42397	1.39429	0.79758	3.6171	5.0114
6	0.75934	1.72966	0.89064	7.6000	9.3297
7	0.67422	1.64454	0.52139	0.2623	1.9069
8	0.62075	1.59107	0.93254	14.0218	15.6129
9	0.66039	1.63071	0.26885	-4.9635	-3.3328
10	0.54600	1.51632	0.70928	4.8802	6.3966
11	0.22039	1.19071	0.50533	0.1072	1.2979
12	0.98132	1.95164	0.56455	1.5204	3.4721
13	0.22174	1.19206	0.76173	8.8943	10.0863
14	0.88548	1.85580	0.92787	25.1705	27.0263
15	0.95191	1.92223	0.18962	-13.5740	-11.6517
16	0.65780	1.62812	0.80940	14.4668	16.0949
17	0.88826	1.85858	0.63536	5.0510	6.9096
18	0.74429	1.71461	0.73987	11.1084	12.8230
19	0.12368	1.09400	0.35452	-6.1894	-5.0954
20	0.59708	1.56740	0.10678	-29.3332	-27.7658
21	0.03746	1.00778	0.76936	15.4752	16.4830
22	0.82311	1.79343	0.14070	-26.6283	-24.8348
23	0.03147	1.00179	0.81132	21.4403	22.4421
24	0.32822	1.29854	0.40197	-5.0193	-3.7208
25	0.20763	1.17795	0.87744	33.7443	34.9223
26	0.57210	1.54242	0.50602	0.3027	1.8451
27	0.66724	1.63756	0.07235	-50.2613	-48.6237
28	0.36904	1.33936	0.26919	-16.7185	-15.3792
29	0.83805	1.80837	0.79157	24.4995	26.3078
30	0.72112	1.69144	0.85190	35.2849	36.9763
31	0.98357	1.95389	0.83221	32.7567	34.7106
32	0.23028	1.20060	0.56917	4.6158	5.8164
33	0.09613	1.06645	0.04275	-78.6988	-77.6323
34	0.00538	0.97570	0.13981	-42.0535	-41.0778
35	0.46984	1.44016	0.56482	4.7211	6.1613
36	0.96043	1.93075	0.31687	-15.9635	-14.0328
37	0.20283	1.17315	0.22134	-29.3368	-28.1636
38	0.60845	1.57877	0.30995	-17.6934	-16.1147
39	0.45104	1.42136	0.13421	-49.9775	-48.5562
40	0.63254	1.60286	0.95061	90.2804	91.8833
41	0.49287	1.46319	0.62664	11.6822	13.1454
42	0.47326	1.44358	0.87987	58.4673	59.9109
43	0.87437	1.84469	0.17758	-43.4769	-41.6322
44	0.01190	0.98222	0.86866	57.4829	58.4651
45	0.89823	1.86855	0.53223	2.9319	4.8004
46	0.54942	1.51974	0.42995	-6.7928	-5.2730
47	0.03274	1.00306	0.03394	-123.7343	-122.7312
48	0.61882	1.58914	0.36316	-15.0285	-13.4394
49	0.30366	1.27398	0.17065	-51.5983	-50.3243
50	0.33108	1.30140	0.40338	-10.5222	-9.2208

In summary, with the original specifications, the data administrator can be certain that e^{ε}-differential privacy is satisfied only for the first query. For all subsequent queries, the value of ε increases and the privacy level decreases. If we attempt to increase the noise variance to compensate for the reduction in privacy, the resulting responses to queries are practically useless. Finally, while we have illustrated this approach for a single data set, the intruder can adopt the tracker approach for any data set of any size. Hence, the results in this section can be generalized to any data set. The only solution to alleviate the above problem is to increase the variance of the noise added with the number of queries. Unfortunately, as we have shown, this has the consequence of making the responses useless after just a few queries. Thus, after just a few queries, Laplace noise addition results in either no privacy or no utility.

5 Conclusions

Differential privacy is often characterized by its proponents along the following lines [10, page 137]: "Differential privacy is therefore an *ad omnia* guarantee, as opposed to an *ad hoc* definition that provides guarantees only against a specific set of attacks or concerns." This characterization ignores the following very practical issues highlighted in this paper:

(1) In many situations, it may not even be possible to implement differential privacy since the numerical variable in question may not have known natural lower and upper bounds.
(2) Even when upper and lower bounds are known, because of global sensitivity,
    a. For large subsets, the level of noise added may be so small that it may not provide the desired level of protection for observations of large magnitude.
    b. For small subsets, the level of noise added may be so large as to make responses from such a system meaningless for many queries.
(3) Even if the above two requirements are satisfied, the intruder's knowledge gain will be limited to e^{ε} only for the first query. For subsequent queries, the intruder's knowledge gain increases exponentially, resulting in practically no privacy after just a few queries.
(4) If we attempt to address the issue in (3) above by increasing the noise variance, after just a few queries, the resulting noise variance is so large as to make all responses to all queries meaningless.

In conclusion, differential privacy and the associated Laplace noise addition procedure may sound like a good idea in theory. However, when we actually examine the applicability of this approach to numerical data as we do in this paper, we find that it has very limited applicability offering either very little privacy or very little utility, or neither. It is important to note that our criticism is not necessarily of differential privacy as a privacy standard, but of Laplace noise addition as the appropriate method to satisfy differential privacy. Our results indicate that while Laplace may satisfy differential privacy in theory, it is of little value in practice for numerical data.

References

[1] Abowd, J. M. and Vilhuber, L. (2008) How Protective Are Synthetic Data? In J. Domingo-Ferrer and Y. Saygın (Eds.): PSD 2008, LNCS 5262, Springer-Verlag, Berlin, 239-246

[2] American Factfinder, U.S. Census Bureau, http://factfinder.census.gov/home/

[3] Chaudhuri, K. and Monteleoni, C. (2008) Privacy-preserving logistic regression. In Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems (NIPS), 289-296

[4] Dalenius T. (1977) Towards a Methodology for Statistical Disclosure Control. *Statistisktidskrift* 5, 429-444

[5] Denning, D. E., Denning, P. J., and Schwartz, M. D. (1979) The tracker: A threat to statistical database security. ACM Transactions on Database Systems 4, 76-96

[6] Dinur, I. and Nissim, K. (2003) Revealing Information while Preserving Privacy. PODS 2003, San Diego, CA, 202-210

[7] Duncan, G.T. and Mukherjee, S. (2000) Optimal Disclosure Limitation Strategy in Statistical Databases: Deterring Tracker Attacks Through Additive Noise. *Journal of the American Statistical Association*, 95, 720-729

[8] Dwork, C. (2006) Differential privacy. In M. Bugliesi, B. Preneel, V. Sassone, and I.Wegener, eds., ICALP (2), Volume 4052, Lecture Notes in Computer Science, Springer, 1–12

[9] Dwork, C. (2011) A firm foundation for private data analysis. *Communications of the ACM*, 54(1), 86-95.

[10] Dwork, C. and Smith, A. (2009) Differential Privacy for Statistics: What we Know and What we Want to Learn. *Journal of Privacy and Confidentiality*, 1, 135-154

[11] Machanavajjhala, A., Kifer, D., Abowd, J., Gehrke, J., and Vilhuber, L. (2008) Privacy: From Theory to Practice on the Map. In ICDE. IEEE Computer Society, 277–286

[12] Nissim, K., Raskhodnokova, S. and Smith, A. (2007) Smooth Sensitivity and Sampling in Private Data Analysis. Proceedings of the thirty-ninth annual ACM symposium on Theory of computing. 75-84

[13] Story, L. (2009) Top Hedge Fund Managers Do Well in a Down Year. New York Times, March 24, http://www.nytimes.com/2009/03/25/business/25hedge.html

[14] Wasserman, L and Zhou, S. (2010) A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105, 375-389.