Finite Sample Differentially Private Confidence Intervals (Extended Abstract)*

Vishesh Karwa[1] and Salil Vadhan[2]

1 **Department of Statistics, The Ohio State University, Columbus, Ohio, USA**
karwa.8@osu.edu
2 **Center for Research on Computation & Society and School of Engineering & Applied Sciences, Harvard University, Cambridge, Massachussets, USA**
salil_vadhan@harvard.edu

Abstract

We study the problem of estimating finite sample confidence intervals of the mean of a normal population under the constraint of differential privacy. We consider both the known and unknown variance cases and construct differentially private algorithms to estimate confidence intervals. Crucially, our algorithms guarantee a finite sample coverage, as opposed to an asymptotic coverage. Unlike most previous differentially private algorithms, we do not require the domain of the samples to be bounded. We also prove lower bounds on the expected size of any differentially private confidence set showing that our the parameters are optimal up to polylogarithmic factors.

1998 ACM Subject Classification G.3 PROBABILITY AND STATISTICS

Keywords and phrases Differential Privacy, Confidence Intervals, Lower bounds, Finite Sample

Digital Object Identifier 10.4230/LIPIcs.ITCS.2018.<44>

1 Overview

Differential privacy [7] is a strong and by now widely accepted definition of privacy for statistical analysis of datasets with sensitive information about individuals. While there is now a rich and flourishing body of research on differential privacy, extending well beyond theoretical computer science, the following three basic goals for research in the area have not been studied in combination with each other:

Differentially private statistical inference: The vast majority of work in differential privacy has studied how well one can approximate statistical properties of the dataset itself, i.e. empirical quantities, rather than inferring statistics of an underlying *population* from which a dataset is drawn. Since the latter is the ultimate goal of most data analysis, it should also be a more prominent object of study in the differential privacy literature.

Conservative statistical inference: An important purpose of statistical inference is to limit the chance that data analysts draw incorrect conclusions because their dataset may not accurately reflect the underlying population, for example due to the sample size being too small. For this reason, classical statistical inference also offers measures of statistical significance such as p-values and confidence intervals. Constructing such measures for differentially private algorithms is more complex, as one must also take into account the additional noise that is introduced for the purpose of privacy protection. For this reason, we advocate that differentially private inference procedures should be *conservative*, and err

* A full version of the paper is available at [17], https://arxiv.org/abs/1711.03908

on the side of underestimating statistical significance, even at small sample sizes and for all settings of other parameters.

Rigorous analysis of the inherent price of privacy: As has been done extensively in the differential privacy literature for empirical statistics, we should also investigate the fundamental "privacy–utility tradeoffs" for (conservative) differentially private statistical inference. This involves both designing and analyzing differentially private statistical inference procedures, as well as proving negative results about the performance that can be achieved, using the best non-private procedures as a benchmark.

In this paper, we pursue all of these goals, using as a case study the problem of constructing a confidence interval for the mean of normal data. The latter is one of the most basic problems in statistical inference, yet already turns out to be nontrivial to fully understand under the constraints of differential privacy. We expect that most of our modeling and methods will find analogues for other inferential problems (e.g. hypothesis testing, Bayesian credible intervals, non-normal data, and estimating statistics other than the mean).

2 Confidence Intervals for a Normal Mean

We begin by recalling the problem of constructing a $(1-\alpha)$-level confidence interval for a normal mean without privacy. Let $X_1, \ldots, X_n$ be an independent and identically distributed (*iid*) random sample from a normal distribution with an unknown mean μ and variance σ^2. The goal is to design an estimator I that given $X_1, \ldots, X_n$, outputs an interval $I(X_1, \ldots, X_n) \subseteq \mathbb{R}$ such that

$$\mathbb{P}\left(I(X_1, \ldots, X_n) \ni \mu\right) \geq 1 - \alpha,$$

for all μ and σ. Here $1-\alpha$ is called the *coverage probability*. Given a desired coverage probability, the goal is minimize the *expected length* of the interval, namely $\mathbb{E}[|I(X_1, \ldots, X_n)|]$.

Known Variance. In the case that variance σ^2 is known (so only μ is unknown), the classic confidence interval for a normal mean is:

$$I(X_1, \ldots, X_n) = \bar{X} \pm \frac{\sigma}{\sqrt{n}} \cdot z_{1-\alpha/2},$$

where $\bar{X}$ is the sample mean and z_a represents the a^{th} quantile of a standard normal distribution.[1] It is known that this interval has the smallest expected size among all $1-\alpha$ level confidence sets for a normal mean, see for example, [20]. In this case, the length of the confidence interval is fixed and equal to

$$|I(X_1, \ldots, X_n)| = (2\sigma z_{1-\alpha/2})/\sqrt{n} = \Theta\left(\sigma\sqrt{\log(1/\alpha)/n}\right).$$

Unknown Variance. In the case that the variance σ^2 is unknown, the variance must be estimated from the data itself, and the classic confidence interval is:

$$I(X_1, \ldots, X_n) = \bar{X} \pm \frac{s}{\sqrt{n}} \cdot t_{n-1,1-\alpha/2},$$

where s^2 is the *sample variance* defined by

$$s^2 = \frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{X})^2,$$

[1] The proof that this is in fact a $(1-\alpha)$-confidence interval follows by observing that $\sqrt{n}\cdot(\bar{X}-\mu)$ has a standard normal distribution, and $[-z_{1-\alpha/2}, z_{1-\alpha/2}]$ covers a $1-\alpha$ fraction of the mass of this distribution.

and $t_{n-1,a}$ is the a^{th} quantile of a t-distribution with $n-1$ degrees of freedom (see the appendix for definitions).[2] Now the length of the interval is a random variable with expectation

$$\mathbb{E}[|I(X_1,\ldots,X_n)|] = \frac{2\sigma}{\sqrt{n}} \cdot k_n \cdot t_{n-1,1-\alpha/2} = \Theta\left(\sigma\sqrt{\log(1/\alpha)/n}\right),$$

for an appropriate constant $k_n = 1 - O(1/n)$. (See [20].)

Relation to Hypothesis Tests. In general, including both cases above, a confidence interval for a population parameter also gives rise to hypothesis tests, which is often how the confidence intervals are used in applied statistics. For example, if our null hypothesis is that the mean μ is nonnegative, then we could reject the null hypothesis if the interval $I(X_1,\ldots,X_n)$ does not intersect the positive real line. The significance level of this hypothesis test is thus at least $1-\alpha$. Minimizing the length of the confidence interval corresponds to being able to reject the alternate hypotheses that are closer to the null hypothesis; that is, when the confidence interval is of length at most β and μ is distance greater than β from the null hypothesis, then the test will reject with probability at least $1-\alpha$.

3 Statistical Inference with Differential Privacy

The Laplace mechanism [7], one of the most basic differentially private algorithm, is used for estimating a function $f(\underline{x})$ of the dataset $\underline{x}$, rather than the population from which $\underline{x}$ is drawn, and much of the differential privacy literature is about estimating such empirical statistics. There are several important exceptions, the earliest being the work on differentially private PAC learning ([2, 18]), but still many basic statistical inference questions have not been addressed.

However, a natural approach for inference was already suggested in early works on differential privacy. In many cases, we know that population statistics are well-approximated by empirical statistics, and thus we can try to estimate these empirical statistics with differential privacy. For example, the population mean μ for a normal population is well-approximated by the sample mean $\overline{X}$, which we can estimate using the Laplace mechanism:

$$M(X_1,\ldots,X_n) = \overline{X} + Z, \text{where } Z \sim \text{Lap}(2B/\epsilon n).$$

On the positive side, observe that the noise being introduced for privacy vanishes linearly in $1/n$, whereas $\overline{X}$ converges to the population mean at a rate of $1/\sqrt{n}$, so asymptotically we obtain privacy "for free" compared to the (optimal) non-private estimator $\overline{X}$.

However, this rough analysis hides some important issues. First, it is misleading to look only at the dependence on n. The other parameters, such as σ, ϵ, and B can be quite significant and should not be treated as constants. Indeed $\sigma/\sqrt{n} \gg B/\epsilon n$ only when $n \gg (B/\epsilon\sigma)^2$, which means that the asymptotics only kick in at a very large value of n. Thus it is important to determine whether the dependence on these parameters is necessary or can be improved. Second, the parameter B is supposed to be a (worst-case) bound on the range of the data, which is incompatible with a modeling the population as following a normal distribution (which is supported on the entire real line). Thus, there have been several works seeking the best asymptotic approximations we can obtain for population statistics under differential privacy, such as [6, 22, 27, 26, 12, 4, 5, 1].

[2] Again the proof follows by observing that $s \cdot (\bar{X} - \mu)$ follows a t distribution, with no dependence on the unknown parameters.

4 Conservative Statistical Inference with DP

The works discussed in the previous section focus on providing point estimates for population quantities, but as mentioned earlier, it is also important to be able to provide measures of statistical significance, to prevent analysts from drawing incorrect conclusions from the results. These measures of statistical significance need to take into account the uncertainty coming both from the sampling of the data and from the noise introduced for privacy. Ignoring the noise introduced for privacy can result in wildly incorrect results at finite sample sizes, as demonstrated empirically many times (e.g. [8, 14, 15]) and this can have severe consequences. For example, [9] found that naive use of differential privacy in calculating warfarin dosage would lead to unsafe levels of medication, but of course one should never use any sort of statistics for life-or-death decisions without some analysis of statistical significance.

Since calculating the exact statistical significance of differentially private computations seems difficult in general, we advocate *conservative* estimates of significance. That is, we require $\mathbb{P}(I(X_1, \ldots, X_n) \ni \mu) \geq 1-\alpha$, for *all* values of n, values of the population parameters, and values of the privacy parameter.

For sample sizes that are too small or privacy parameters that are too aggressive, we may achieve this property by allowing the algorithm to sometimes produce an extremely large confidence interval, but that is preferable to producing a small interval that does not actually contain the true parameter which may violate the desired coverage property. Note that what constitutes a sample size that is "too small" can depend on the unknown parameters of the population (e.g. the unknown variance σ^2) and their interplay with other parameters (such as the privacy parameter ϵ).

Returning to our example of estimating a normal mean with known variance under differential privacy, if we use the Laplace Mechanism to approximate the empirical mean (as discussed above), we can obtain a conservative confidence interval for the population mean by increasing the length of classical, non-private confidence interval to account for the likely magnitude of the Laplace noise. More precisely, starting with the differentially private mechanism

$$M(X_1, \ldots, X_n) = \overline{X} + Z, \text{where } Z \sim \text{Lap}(2B/\epsilon n),$$

the following is a $(1 - O(\alpha))$-level confidence interval for the population mean μ:

$$I(X_1, \ldots, X_n) = M(X_1, \ldots, X_n) \pm \left(\frac{\sigma}{\sqrt{n}} \cdot z_{1-\alpha/2} + \frac{B}{\epsilon n} \cdot \log(1/\alpha) \right).$$

The point is that with probability $1 - O(\alpha)$, the Laplace noise Z has magnitude at most $(B/\epsilon n) \cdot \log(1/\alpha)$, so increasing the interval by this amount will preserve coverage (up to an $O(\alpha)$ change in the probability). Again, the privacy guarantees of the Laplace mechanism relies on the data points being guaranteed to lie in $[-B, B]$; otherwise, points need to be clamped to lie in the range, which can bias the empirical mean and compromise the coverage guarantee. Thus, to be safe, a user may choose a very large value of B, but then this makes for a much larger (and less useful) interval, as the length of the interval grows linearly with B. Thus, a natural research question (which we investigate) is whether such a choice and corresponding cost is necessary.

Conservative hypothesis testing with differential privacy, where we require that the significance level is at least $1 - \alpha$, was advocated by [11]. Methods aimed at calculating the combined uncertainty due to sampling and privacy (for various differentially private algorithms) were given in [24, 28, 14, 13, 16, 15, 11, 23, 25, 19], but generally the utility of these methods (e.g. the expected length of a confidence interval or power of a hypothesis test)

is only evaluated empirically or the conservativeness only holds in a particular asymptotic regime. Rigorous, finite-sample analyses of conservative inference were given in [21] for confidence intervals on the coefficients from ordinary least-squares regression (which can be seen as a generalization of the problem we study to multivariate Gaussians) and in [3] for hypothesis testing of discrete distributions. However, neither paper provides matching lower bounds, and in particular, the algorithms of [21] only apply for bounded data (similar to the basic Laplace mechanism). In our work, we provide a comprehensive theoretical analysis of conservative differentially private confidence intervals for a normal mean, with both algorithms and lower bounds, without any bounded data assumption.

5 Our Results

As discussed above, in this paper we develop conservative differentially private estimators of confidence intervals for the mean μ of a normal distribution with known and unknown variance σ^2. Our algorithms are designed to be differentially private for all input datasets and they provide $(1-\alpha)$-level coverage whenever the data is generated from a normal distribution. Unlike the Laplace mechanism described above and many other differentially private algorithms, we do not make any assumptions on the boundedness of the data. Our pure DP (i.e. $(\epsilon, 0)$-DP) algorithms assume that the mean μ and variance σ^2 lie in a bounded (but possibly very large) interval, and we show (using lower bounds) that such an assumption is necessary. Our approximate (i.e. (ϵ, δ)) differentially private algorithms do not make any such assumptions, i.e. both the data and the parameters (μ, σ^2) can remain unbounded. We also show that the differentially private estimators that we construct have nearly optimal expected length, up to logarithmic factors. This is done by proving lower bounds on the length of differentially private confidence intervals. A key aspect of the confidence intervals that we construct is their conservativeness — the coverage guarantee holds in finite samples, as opposed to only holding asymptotically. We also show that as $n \to \infty$, the length of our differentially private confidence intervals is at most $1 + o(1)$ factor larger than length of their non-private counterparts.

Let $X_1, \ldots, X_n$ be an independent and identically distributed (*iid*) random sample from a normal distribution with an unknown mean μ and variance σ^2, where $\mu \in (-R, R)$ and $\sigma \in (\sigma_{\min}, \sigma_{\max})$. Our goal is to construct (ϵ, δ)-differentially private $(1-\alpha)$-level confidence sets for μ in both the known and the unknown variance case, i.e. we seek a set $I = I(X_1, \ldots, X_n)$ such that

1. $I(X_1, \ldots, X_n)$ is a $(1-\alpha)$-level confidence interval, and
2. $I(x_1, \ldots, x_n)$ is (ϵ, δ)-differentially private.
3. $\mathop{\mathbb{E}}_{X_1,\ldots,X_n,I}[I(X_1, \ldots, X_n)]$ is as small as possible.

Known Variance: For the known variance case, we construct differentially private algorithms that output a fixed width $(1-\alpha)$-level confidence interval for any n. Moreover, when n is large enough, the algorithm outputs a confidence interval of length β which is non-trivial in the sense that $\beta \ll R$. Specifically, β is a maximum of two terms: The first term is $\mathcal{O}\left(\sigma\sqrt{\log(1/\alpha)/n}\right)$ which is the same as the length of the non-private confidence interval discussed in Section 2 up to constant factors. The second term is $\mathcal{O}\left(\sigma/(\epsilon n)\right)$ up to polylogarithmic factors − it goes to 0 at the rate of $\tilde{\mathcal{O}}\left(1/n\right)$ which is faster than the rate at which the first term goes to 0. Thus for large n the increase in the length of the confidence interval due to privacy is mild. Note that, unlike the basic approach based on the Laplace mechanism discussed in Section 4, the length of the confidence interval has no dependence on the range of the data, or even the range $(-R, R)$ of the mean μ.

The sample complexity required for obtaining a non-trivial confidence interval is the minimum of two terms: $\mathcal{O}\left((1/\epsilon)\log(R/\alpha\sigma)\right)$ and $\mathcal{O}\left((1/\epsilon)\log(1/\alpha\delta)\right)$. The dependence of sample complexity on R/σ is only logarithmic. Thus one can choose a very large value of R. Moreover, when $\delta > 0$, we can set $R = \infty$ and hence there is no dependence of the sample complexity on R.

Unknown Variance: As in the known variance case, we construct (ϵ, δ) differentially private algorithms that output an $(1-\alpha)$ confidence interval of μ for all n. If n is large enough, the length of the confidence interval is a maximum of two terms, where the first term is same as the length of the non-private confidence interval and the second term goes to 0 at a faster rate.

As before, the dependence of sample complexity on $R/\sigma_{\min}$ and $\sigma_{\max}/\sigma_{\min}$ is logarithmic, as opposed to linear. Hence we can set these parameters to a large number. Moreover, when $\delta > 0$, we can set R and $\sigma_{\max}$ to be ∞ and $\sigma_{\min}$ to be 0. Thus when $\delta > 0$, there are no assumptions on the boundedness of the parameters.

Lower Bounds: We also prove lower bounds on the length of any $(1-\alpha)$-level (ϵ, δ)-differentially private confidence set of expected size β. Our lower bounds show that one must pay $\Omega\left(\sigma/(\epsilon n) \cdot \log(1/\alpha)\right)$ in the length of the confidence interval when R is very large. Our algorithms come quite close to this lower bound with an extra factor of $\text{polylog}(n/\alpha)$. We also show that the sample complexity required by our algorithms is necessary to obtain a confidence interval that saves more than a factor of 2 over the trivial interval $(-R, R)$.

6 Directions for Future Work

The most immediate direction for future work is to close the (small) gaps between our upper and lower bounds. We came to the problem of constructing confidence intervals for a normal mean as part of an effort to bring differential privacy to practice in the sharing of social science research data through the design of the software tool PSI [10], as confidence intervals are a workhorse of data analysis in the social sciences. However, our algorithms are not optimized for practical performance, but rather for asymptotic analysis of the confidence interval length. Initial experiments indicate that alternative approaches (not just tuning of parameters) may be needed to reasonably sized confidence intervals (e.g. length at most twice that of the non-private length) handle modest sample sizes (e.g. in the 1000's). Thus designing practical differentially private algorithms for confidence intervals remains an important open problem, whose solution could have wide applicability.

As mentioned earlier, we expect that much of the modelling and techniques we develop should also be applicable more widely. In particular, it would be natural to study the estimation of other population statistics, and families of distributions, such as other continuous random variables, Bernoulli random variables, and multivariate families. In particular, a natural generalization of the problem we consider is to construct confidence intervals for the parameters of a (possibly degenerate) multivariate Gaussian, which is closely related to the problem of ordinary least-squares regression (cf. [21]).

Finally, while we have advocated for conservative inference at finite sample size, to avoid spurious conclusions coming from the introduction of privacy, many practical, non-private inference methods rely on asymptotics also for measuring statistical significance. In particular, the standard confidence interval for a normal mean with unknown variance and its corresponding hypothesis test (see Section 2) is often applied on non-normal data, and heuristically justified using the Central Limit Theorem. (This is heuristic since the rate of convergence depends on the data distribution, which is unknown.) Is there a criterion

to indicate what asymptotics are "safe"? In particular, can we formalize the idea of only using the "same" asymptotics that are used without privacy? [19] analyze their hypothesis tests using asymptotics that constrain the setting of the privacy parameter in terms of the sample size n (e.g. $\epsilon \geq \Omega(1/\sqrt{n})$), but it's not clear that this relationship is safe to assume in general.

Acknowledgments. This research was supported by NSF grant CNS-1237235, a Simons Investigator Award to Salil Vadhan, a grant from the Sloan Foundation, and Cooperative Agreement CB16ADR0160001 with the Census Bureau. We are grateful to the Harvard Privacy Tools differential privacy research group for illuminating and motivating discussions, particularly James Honaker, Gary King, Kobbi Nissim and Uri Stemmer. We thank members of the Center for Disclosure Avoidance Research at the US Census Bureau, in particular Philip Leclerc, for carefully reading our paper and giving helpful feedback.

References

1 Rina Foygel Barber and John C Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.

2 Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 128–138. ACM, 2005.

3 Bryan Cai, Constantinos Daskalakis, and Gautam Kamath. Priv'it: Private and sample efficient identity testing. *arXiv preprint arXiv:1703.10127*, 2017.

4 John Duchi, Martin J Wainwright, and Michael I Jordan. Local privacy and minimax bounds: Sharp rates for probability estimation. In *Advances in Neural Information Processing Systems*, pages 1529–1537, 2013.

5 John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 429–438. IEEE, 2013.

6 Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 371–380. ACM, 2009.

7 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284. Springer, 2006.

8 Stephen E. Fienberg, Alessandro Rinaldo, and Xiaolin Yang. Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *Proceedings of the 2010 international conference on Privacy in statistical databases*, PSD'10, pages 187–199, Berlin, Heidelberg, 2010. Springer-Verlag. URL: `http://dl.acm.org/citation.cfm?id=1888848.1888869`.

9 Matthew Fredrikson, Eric Lantz, Somesh Jha, Simon Lin, David Page, and Thomas Ristenpart. Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. In *USENIX Security Symposium*, pages 17–32, 2014.

10 Marco Gaboardi, James Honaker, Gary King, Kobbi Nissim, Jonathan Ullman, and Salil Vadhan. Psi ({\Psi}): a private data sharing interface. *arXiv preprint arXiv:1609.04340*, 2016.

11 Marco Gaboardi, Ryan Rogers, and Salil Vadhan. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. 2016.

12 Rob Hall, Alessandro Rinaldo, and Larry Wasserman. Differential privacy for functions and functional data. *Journal of Machine Learning Research*, 14(Feb):703–727, 2013.

13 Vishesh Karwa, Dan Kifer, and Aleksandra B Slavković. Private posterior distributions from variational approximations. *arXiv preprint arXiv:1511.07896*, 2015.

14 Vishesh Karwa and Aleksandra Slavković. Differentially private graphical degree sequences and synthetic graphs. In *Privacy in Statistical Databases*, pages 273–285. Springer, 2012.

15 Vishesh Karwa and Aleksandra Slavković. Inference using noisy degrees: Differentially private β-model and synthetic graphs. *The Annals of Statistics*, 44(1):87–112, 2016.

16 Vishesh Karwa, Aleksandra B Slavković, and Pavel Krivitsky. Differentially private exponential random graphs. In *International Conference on Privacy in Statistical Databases*, pages 143–155. Springer, 2014.

17 Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. *arXiv preprint arXiv:1711.03908*, 2017.

18 Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

19 Daniel Kifer and Ryan Rogers. A new class of private chi-square tests. *arXiv preprint arXiv:1610.07662*, 2016.

20 Erich L Lehmann and Joseph P Romano. *Testing statistical hypotheses*. Springer Science & Business Media, 2006.

21 Or Sheffet. Differentially private ordinary least squares. In *International Conference on Machine Learning*, pages 3105–3114, 2017.

22 Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 813–822. ACM, 2011.

23 Eftychia Solea. Differentially private hypothesis testing for normal random variables. 2014.

24 Duy Vu and Aleksandra Slavkovic. Differential privacy for clinical trial data: Preliminary evaluations. In *Data Mining Workshops, 2009. ICDMW'09. IEEE International Conference on*, pages 138–143. IEEE, 2009.

25 Yue Wang, Jaewoo Lee, and Daniel Kifer. Differentially private hypothesis testing, revisited. *arXiv preprint arXiv:1511.03376*, 2015.

26 Larry Wasserman. Minimaxity, statistical thinking and differential privacy. *Journal of Privacy and Confidentiality*, 4(1):3, 2012.

27 Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *J. Amer. Statist. Assoc.*, 105(489):375–389, 2010. URL: `http://dx.doi.org/10.1198/jasa.2009.tm08651`, `doi:10.1198/jasa.2009.tm08651`.

28 Oliver Williams and Frank McSherry. Probabilistic inference and differential privacy. In *Advances in Neural Information Processing Systems*, pages 2451–2459, 2010.