SUBSAMPLING IS NOT MAGIC: WHY LARGE BATCH SIZES WORK FOR DIFFERENTIALLY PRIVATE STOCHASTIC OPTIMISATION

Ossi Räisä[*], Joonas Jälkö[*] & Antti Honkela
University of Helsinki
Department of Computer Science
{ossi.raisa,joonas.jalko,antti.honkela}@helsinki.fi

ABSTRACT

We study the effect of the batch size to the total gradient variance in differentially private stochastic gradient descent (DP-SGD), seeking a theoretical explanation for the usefulness of large batch sizes. As DP-SGD is the basis of modern DP deep learning, its properties have been widely studied, and recent works have empirically found large batch sizes to be beneficial. However, theoretical explanations of this benefit are currently heuristic at best. We first observe that the total gradient variance in DP-SGD can be decomposed into subsampling-induced and noise-induced variances. We then prove that in the limit of an infinite number of iterations, the effective noise-induced variance is invariant to the batch size. The remaining subsampling-induced variance decreases with larger batch sizes, so large batches reduce the effective total gradient variance. We confirm numerically that the asymptotic regime is relevant in practical settings when the batch size is not small, and find that outside the asymptotic regime, the total gradient variance decreases even more with large batch sizes. We also find a sufficient condition that implies that large batch sizes similarly reduce effective DP noise variance for one iteration of DP-SGD.

1 INTRODUCTION

As deep learning models are being trained on ever larger datasets, the privacy of the subjects of these training datasets is a growing concern. *Differential privacy* (DP) (Dwork et al., 2006) is a property of an algorithm that formally quantifies the privacy leakage that can result from releasing the output of the algorithm. Due to the formal guarantee provided by DP, there is a great deal of interest in training deep learning models with a DP variant of stochastic gradient descent (DP-SGD) (Song et al., 2013; Bassily et al., 2014; Abadi et al., 2016).

One of the key properties of DP is so called *subsampling amplification* (Li et al., 2012; Beimel et al., 2014). Broadly speaking, subsampling the data before applying a DP mechanism adds an additional layer of protection to the data samples, leading to stronger privacy guarantees for a fixed amount of added noise. Quantifying the gains from subsampling amplification (Abadi et al., 2016; Zhu & Wang, 2019; Koskela et al., 2020; Zhu et al., 2022) has been a crucial component in making algorithms such as DP-SGD work under strict privacy guarantees.

On the other hand, many works have shown that the large batch sizes actually provide better performance in DP-SGD compared to heavy subsampling of the training data McMahan et al. (2018); De et al. (2022); Mehta et al. (2023). This would suggest, that maybe the gains from the subsampling amplification are outweighted by the stronger signal of gradients from larger batches.

In Poisson subsampled DP-SGD, each datapoint is included with probability q, called the *subsampling rate*, which is proportional to the expected batch size. The total gradient variance in DP-SGD can be decomposed into two parts: the subsampling variance and the Gaussian noise variance σ^2. A larger q reduces the subsampling variance, but the effect on the noise variance is not as clear. On one hand, a

[*]Equal contribution

larger q reduces the privacy amplification effect, necessitating a larger σ^2, but on the other hand, an unbiased gradient estimate must be divided by q, so the effective noise variance is $\frac{\sigma^2}{q^2}$.

We study how the effective noise variance scales with the subsampling rate, making the following contributions:

1. In Section 5, we prove that in the limit of an infinite number of iterations, there is a linear relationship $\sigma = cq$ between q and σ, meaning that the two effects q has on the effective noise variance cancel each other, leaving the effective noise variance invariant to the subsampling rate. This means that a larger subsampling rate always reduces the effective total gradient variance, since the subsampling-induced variance decreases with a larger subsampling rate.
2. In Section 6, we consider the case of a single iteration of the subsampled Gaussian mechanism. We find a sufficient condition which implies that large subsampling rates always reduce the effective injected DP noise variance, hence also reducing the effective total gradient variance. We check this condition numerically for a wide grid of hyperparameter values, and find that the condition holds amongst these hyperparameters.
3. We look at the relationship between the subsampling rate and noise standard deviation empirically in Section 5.1, and find that the asymptotic regime from our theory is reached quickly with small privacy parameters. Moreover, we find that when we are not in the asymptotic regime, the effective injected DP noise variance decreases even more with a large subsampling rate.

1.1 Related Work

The linear relationship between the subsampling rate and noise standard deviation we study has been suggested as a heuristic rule-of-thumb in previous work (Li et al., 2021; Sander et al., 2023) to explain why large batch sizes perform better in DP-SGD. The linear relationship also appears in several works that study Rényi DP (Mironov, 2017) accounting of the subsampled Gaussian mechanism (Abadi et al., 2016; Bun et al., 2018; Mironov et al., 2019), though these works do not make the connection with large subsampling rates, and instead assume a small subsampling rate. Rényi DP-based accounting of the subsampled Gaussian mechanism does not provide tight privacy bounds anyway, so these results would not imply that the linear relationship holds even asymptotically with tight accounting.

Sommer et al. (2019) and Dong et al. (2022) prove central limit theorems for privacy accounting, which essentially show that the privacy loss of any sufficiently well-behaved mechanism after many compositions is asymptotically like the privacy loss of the Gaussian mechanism. As we study the asymptotic behaviour of the subsampled Gaussian mechanism after many compositions, it is possible that these theorems could be used to prove our result. However, we opted for another route in our proof. Instead of showing that privacy accounting for the subsampled Gaussian mechanism is asymptotically similar to accounting for the Gaussian mechanism, we show that the subsampled mechanism itself is asymptotically similar to the Gaussian mechanism.

2 Background

In this section, we go through some background material on differential privacy. We start with the definition and basics in Section 2.1 and introduce DP-SGD in Section 2.2.

2.1 Differential privacy

Differential privacy (DP) (Dwork et al., 2006; Dwork & Roth, 2014) is a property of an algorithm that quantifies the privacy loss resulting from releasing the output of the algorithm. In the specific definition we use, the privacy loss is quantified by two numbers: $\epsilon \geq 0$ and $\delta \in [0, 1]$.

Definition 2.1. Let $\mathcal{M}$ be a randomised algorithm. $\mathcal{M}$ is (ϵ, δ)-DP if, for all measurable A and all neighbouring $x, x' \in \mathcal{X}$,

$$\Pr(\mathcal{M}(x) \in A) \leq e^{\epsilon} \Pr(\mathcal{M}(x') \in A) + \delta. \tag{1}$$

We focus on the add/remove neighbourhood in this work.

2.2 DIFFERENTIALLY PRIVATE SGD

Differentially private SGD (DP-SGD) (Song et al., 2013; Bassily et al., 2014; Abadi et al., 2016), uses

$$G_{DP} = \sum_{i \in \mathcal{B}} \text{clip}_C(g_i) + \mathcal{N}(0, \sigma^2 I_d) \tag{2}$$

in place of the non-private sum of gradients. G_{DP} can also be used in adaptive versions of SGD like Adam.

Privacy bounds for DP-SGD can be computed using numerical privacy accountants (Koskela et al., 2020; 2021; Gopi et al., 2021; Doroshenko et al., 2022; Alghamdi et al., 2023). For more technical background on privacy accounting and the composition of DP mechanisms, see Appendix A.

3 ACCOUNTING ORACLES

We use the concept of *accounting oracles* (Tajeddine et al., 2020) to make formalising properties of privacy accounting easier. The accounting oracle is the ideal privacy accountant that numerical accountants aim to approximate.

Definition 3.1. For a mechanism $\mathcal{M}$, the accounting oracle $\text{AO}_{\mathcal{M}}(\epsilon)$ returns the smallest δ such that $\mathcal{M}$ is (ϵ, δ)-DP.

In case $\mathcal{M}$ has hyperparameters that affect its privacy bounds, these hyperparameters will also be arguments of $\text{AO}_{\mathcal{M}}$. We write the accounting oracle for the T-fold composition of the Poisson subsampled Gaussian mechanism with sampling rate q and sensitivity Δ as $\text{AO}_S(\sigma, \Delta, q, T, \epsilon)$, and the accounting oracle of a composition of the Gaussian mechanism as $\text{AO}_G(\sigma, \Delta, T, \epsilon)$.

See Appendix B for properties of accounting oracles.

4 DP-SGD NOISE DECOMPOSITION

For now, let us denote the sum on the right in Equation (2), the sum of clipped gradients with subsampling rate q, as $G_q := \sum_{i \in \mathcal{B}} \text{clip}_C(g_i)$. In each step of DP-SGD, we are releasing this sum using Gaussian perturbation. However, due to the subsampling, our gradient sum can comprise of any number of terms between 0 and N. Therefore, before we do the step, we want to upscale the summed gradient to approximate the magnitude of the full data gradient G_1. We will use $1/q$ scaling for G_q, which also gives an unbiased estimator of G_1 (see Appendix C.1). We can decouple the noise-scale and clipping threshold C and write the DP gradient $\tilde{G}$ used to update the parameters as

$$\tilde{G} = \frac{1}{q} \left(G_q + C\sigma\eta \right), \tag{3}$$

where $\eta \sim \mathcal{N}(0, I_d)$. The clipping threshold C affects the update as a constant scale independent of q, and therefore for notational simplicity we set $C = 1$. Since the subsampled gradient G_q and the DP noise are independent, we can decompose the total gradient variance of the jth element of $\tilde{G}$ as

$$\underbrace{\text{Var}(\tilde{G}_j)}_{\text{Total}} = \underbrace{\frac{1}{q^2} \text{Var}(G_{q,j})}_{\text{Subsampling}} + \underbrace{\left(\frac{\sigma}{q}\right)^2}_{\text{Effective DP Noise}} . \tag{4}$$

The first component on the right in Equation (4), the subsampling variance, arises from the subsampling, so it is easy to see that it is decreasing w.r.t. q (see Appendix C.2) For the rest of the paper, we will use $\sigma^2_{\text{eff}} := \sigma^2/q^2$ to denote the second component of the variance decomposition in Equation (4), the effective noise variance.

In order to guarantee (ϵ, δ)-DP, the σ term in σ_{eff} needs to be adjusted based on the number of iterations T and the subsampling rate q. Hence we will treat the σ as a function of q and T, and denote the smallest σ that provides (ϵ, δ)-DP as $\sigma(q, T)$.

In the rest of the paper, we will study how $\sigma(q, T)$ behaves w.r.t. q in two settings: (i) when $T \to \infty$ we show that $\sigma(q, T)$ becomes a linear function of q (Section 5), and (ii) when $T = 1$ we derive a sufficient condition showing that $\sigma(q, 1)/q$ is decreasing w.r.t q (Section 6).

5 SUBSAMPLED GAUSSIAN MECHANISM IN THE LIMIT OF MANY COMPOSITIONS

Using the accounting oracle $\mathrm{AO}_S(\sigma, \Delta, q, T, \epsilon)$ to study the privacy accounting of DP-SGD, simplifies the analysis. In the setting we use, there is a single datapoint $x \in \{0, 1\}$, which is released T times with the Poisson subsampled Gaussian mechanism:

$$\mathcal{M}_i(x) \sim xB_q + \mathcal{N}(0, \sigma_T^2) \tag{5}$$

for $1 \leq t \leq T$, where B_q is a Bernoulli random variable. Since $\mathcal{M}$ is a Poisson subsampled Gaussian mechanism, its accounting oracle is $\mathrm{AO}_S(\sigma, 1, q, T, \epsilon)$.[1] We have $\mathbb{E}(\mathcal{M}_i(x)) = qx$ and $\mathrm{Var}(\mathcal{M}_i(x)) = x^2q(1-q) + \sigma_T^2$. As $\sigma_T^2 \to \infty$, the variance and σ_T^2 approach each other. As a result, we can approximate $\mathcal{M}_i$ with

$$\mathcal{M}_i'(x) \sim \mathcal{N}(qx, \sigma_T^2). \tag{6}$$

To prove this, we first need to find a lower bound on how quickly σ_T^2 must grow as $T \to \infty$. All proofs for this section are in Appendix D.

Theorem 5.1. *Let σ_T be such that $\mathrm{AO}_S(\sigma_T, \Delta, q, T, \epsilon) \leq \delta$ for all T, with $\delta < 1$. Then $\sigma_T^2 = \Omega(T)$.*

Now we can prove that the approximation is sound.

Theorem 5.2. *For $1 \leq i \leq T$, let*

$$\mathcal{M}_i(x) \sim xB_q + \mathcal{N}(0, \sigma_T^2), \tag{7}$$

$$\mathcal{M}_i'(x) \sim \mathcal{N}(qx, \sigma_T^2). \tag{8}$$

be independent for each i. Let $\mathcal{M}_{1:T}$ be the composition of $\mathcal{M}_i$, and let $\mathcal{M}'_{1:T}$ be the composition of $\mathcal{M}_i'$. Then

$$\sup_x \mathrm{TV}(\mathcal{M}_{1:T}(x), \mathcal{M}'_{1:T}(x)) \to 0. \tag{9}$$

The mechanism $\mathcal{M}_i'$ is nearly a Gaussian mechanism, since

$$\frac{1}{q}\mathcal{M}_i'(x) \sim x + \mathcal{N}\left(0, \frac{\sigma_T^2}{q^2}\right). \tag{10}$$

On the right is a Gaussian mechanism with noise standard deviation $\frac{\sigma_T}{q}$, which has the linear relationship between σ_T and q we are looking for. In the next theorem, we use this to prove our main result.

Theorem 5.3. *For any σ, q_1, q_2, Δ and ϵ*

$$|\mathrm{AO}_S(\sigma, \Delta, q_1, T, \epsilon) - \mathrm{AO}_S(\sigma \cdot q_2/q_1, \Delta, q_2, T, \epsilon)| \to 0 \tag{11}$$

as $T \to \infty$.

By setting $q_2 = 1$, we see that $\mathrm{AO}_S(\sigma, \Delta, q, T, \epsilon)$ behaves like $\mathrm{AO}_S(\sigma/q, \Delta, 1, T, \epsilon)$ in the $T \to \infty$ limit, so $\frac{\sigma}{q}$ must be a constant that does not depend on q to reach a given (ϵ, δ)-bound. We formalise this in the following corollary.

Corollary 5.4. *Let $\sigma(q, T)$ be the smallest σ such that $\mathrm{AO}_S(\sigma, \Delta, q, T, \epsilon) \leq \delta$. Then*

$$\lim_{T \to \infty} \frac{\sigma(q, T)}{q\sigma(1, T)} = 1. \tag{12}$$

5.1 EMPIRICAL RESULTS

In order to demonstrate how the σ_{eff} converges to $q\sigma(1, T)$, we use Google's open source implementation[2] of the PLD accounting Doroshenko et al. (2022) to compute the $\sigma(q, T)$.

[1] Considering $\Delta = 1$ suffices due to Lemma B.2.

[2] Available at https://github.com/google/differential-privacy/tree/main/python

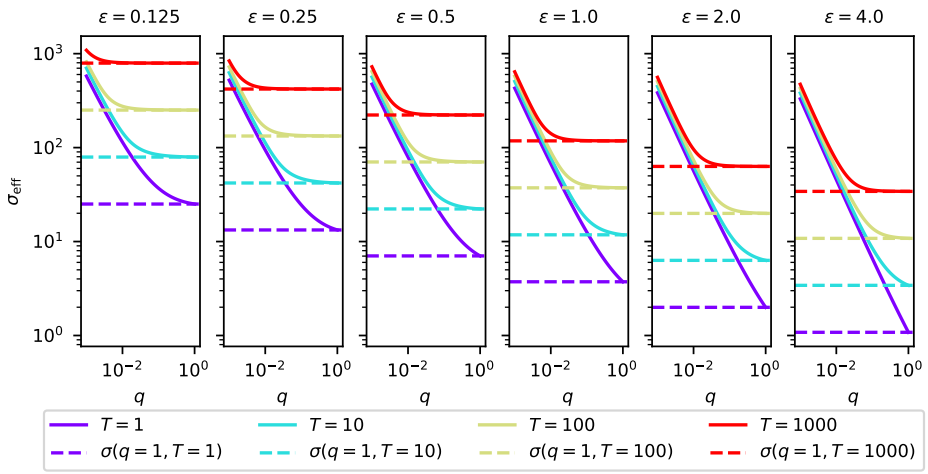

Figure 1: The $\sigma_{\text{eff}} := \sigma(q,T)/q$ decreases as q grows for all the ϵ and T values. As T grows, σ_{eff} approaches the $\sigma(1,T)$, as the asymptotic theory predicts. The privacy parameter δ was set to 10^{-5} when computing the $\sigma(q,T)$.

Figure 1 shows that as the the number of iterations grows, the σ_{eff} approaches the $\sigma(1,T)$ line. We can also see, that for smaller values ϵ, the convergence happens faster. This behaviour can be explained by the larger level of perturbation needed for smaller ϵ values, which will make the components of the Gaussian mixture in Equation (7) less distinguishable from each other.

We can also see, that for all the settings, the σ_{eff} stays above the $\sigma(1,T)$ line, and it is the furthest away when q is the smallest. This would suggest, that our observation in Section 6 hold also for $T > 1$, and that smaller values of q incur a disproportionally large DP-induced variance component in Equation (4).

6 Subsampled Gaussian Mechanism without Compositions

For now, let us consider $T = 1$, and denote $\sigma(q) := \sigma(q,1)$. We can express the δ as (Koskela et al., 2020; Zhu et al., 2022)

$$\delta(q) = q \Pr\left(Z \geq \sigma(q) \log\left(\frac{h(q)}{q}\right) - \frac{1}{2\sigma(q)}\right) - h(q) \Pr\left(Z \geq \sigma(q) \log\left(\frac{h(q)}{q}\right) + \frac{1}{2\sigma(q)}\right), \tag{13}$$

where $h(q) := e^{\epsilon} - (1-q)$. Recall that we have assumed $\sigma(q,T)$ to be a function that returns a noise-level matching to a particular (ϵ,δ) privacy level for given q and T. Therefore, the derivative of $\delta(q)$ in Equation (13) w.r.t. q is 0, and we can solve the derivative of the RHS for $\sigma'(q)$. See Appendix E for missing derivations and proofs in this section.

Let us denote

$$a := \frac{1}{2\sqrt{2}\sigma(q)}, \; b := \frac{\sigma(q)}{\sqrt{2}} \log\left(\frac{e^{\epsilon} - (1-q)}{q}\right). \tag{14}$$

We have the following Lemma

Lemma 6.1. *For the smallest $\sigma(q)$ that provides (ϵ,δ)-DP for the Poisson subsampled Gaussian mechanism, we have*

$$\sigma'(q) = \frac{\sigma(q)}{q} \frac{1}{2a} \frac{1}{\operatorname{erf}'(a-b)} (\operatorname{erf}(a-b) - \operatorname{erf}(-a-b)). \tag{15}$$

Now, Lemma 6.1 allows us to establish following result.

Theorem 6.2. *If $a < b$ for a and b defined in Equation* (14)*, then $\frac{\mathrm{d}}{\mathrm{d}q} \frac{\sigma(q)}{q} < 0$.*

Now, Theorem 6.2 implies that σ_{eff} is a decreasing function w.r.t. q, and therefore larger subsampling rates should be preferred when $a < b$. So now the remaining question is, when is a smaller than b.

Unfortunately, analytically solving the region where $a < b$ is intractable as we do not have a closed form expression for $\sigma(q)$. Therefore we make the following Conjecture, which we study numerically.

Conjecture 6.3. *For $\epsilon, q \geq 4\delta$, we have $a - b < 0$.*

Verifying the Conjecture 6.3 numerically requires computing the $\sigma(q)$ values for a range of q and ϵ values, which would be computationally inefficient. However, computing a, b and δ can be easily parallelized over multiple values of q and σ. We set $\delta_{\text{target}} = 10^{-5}$ and compute the a, b and δ for $q \in [4\delta_{\text{target}}, 1.0]$, $\sigma \in [\min(q)\sigma(1,\epsilon), \sigma(1,\epsilon)]$ and $\epsilon \in [4\delta_{\text{target}}, 4.0]$ which would be a reasonable range of ϵ values for practical use. The $\sigma(1,\epsilon)$ is a noise-level matching $(\epsilon, \delta_{\text{target}})$-DP guarantee for $q = 1$, and the lower bound $\min(q)\sigma(1,\epsilon)$ for σ was selected based on hypothesis that within $[\min(q)\sigma(1,\epsilon), \sigma(1,\epsilon)]$ we can find a noise-level closely matching the δ_{target}.

Figure 2 shows the largest $a - b$ value for our target ϵ values among the q and σ pairs that resulted into $\delta \leq \delta_{\text{target}}$. We can see that among these values there are no cases of $a > b$. While our evaluation covers a range of ϵ values, $a - b$ seems to be monotonically decreasing w.r.t. ϵ, which suggests that the conjecture should hold even for larger values of ϵ. Based on our numerical evaluation, the σ range resulted into δ values differing from $\delta_{\text{target}} = 10^{-5}$ at most $\approx 2 \times 10^{-7}$ in absolute difference. As a final remark, the constant 4 in the Conjecture 6.3 was found numerically. With smaller values of this constant we obtained cases for which $a - b > 0$. Furthermore, values of $q \approx \delta$ empirically produce results that even fail to satisfy the claim of Theorem 6.2.

Limitations Our analysis for the $T = 1$ case reduced the monotonicity of effective DP noise standard deviation σ_{eff} to a sufficient condition $a < b$. We were unable to provide a formal proof of when this condition is satisfied, but verified numerically that it appears valid for a broad range of practically relevant parameters with $\delta = 10^{-5}$, a standard value suggested by NIST.

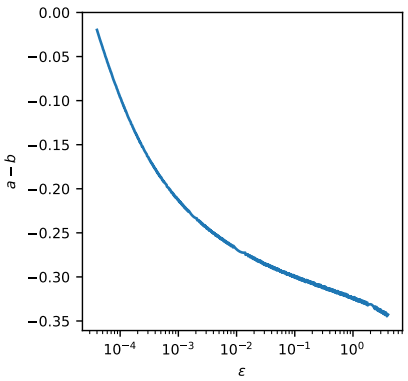

Figure 2: The largest $a - b$ computed for multiple (q, σ) pairs stays negative for a broad range of ϵ values. The a and b were selected so that the corresponding (q, σ) pair satisfies (ϵ, δ)-DP with $\delta \leq 10^{-5}$.

7 Conclusion

We studied the relationship between the effective noise variance and the subsampling rate in Poisson subsampled DP-SGD, and proved that as the number of iterations approaches infinity, the relationship becomes linear, which cancels the effect of the subsampling rate in the effective noise variance. This means that a large subsampling rate always reduces the effective total gradient variance. Furthermore, we demonstrated that under a wide range of ϵ values, a single application of the Poisson subsampled Gaussian mechanism actually incurs a monotonically decreasing effective noise variance w.r.t. subsampling rate. Our numerical experiments show that the asymptotic regime is relevant in practice, and that outside the asymptotic regime, smaller subsampling rates lead to increasingly large effective total gradient variances. This explains the observed benefits of large batch sizes in DP-SGD, which has so far had only empirical and heuristic explanations, furthering the theoretical understanding of DP-SGD.

For future work, it would be important to theoretically study how to interpolate our results between the $T = 1$ and the asymptotic case. Based on our numerical evaluations however, we expect our main conclusion, that the large batches provide smaller effective noise, to hold even for finite $T > 1$.

In all cases we have studied, less subsampling (larger q) always leads to better privacy–utility trade-off, at the cost of more compute. Thus the magic of subsampling amplification lies in saving compute, not in achieving higher utility than without subsampling.

ACKNOWLEDGEMENTS

This work was supported by the Research Council of Finland (Flagship programme: Finnish Center for Artificial Intelligence, FCAI as well as Grants 356499 and 359111), the Strategic Research Council at the Research Council of Finland (Grant 358247) as well as the European Union (Project 101070617). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 308–318. ACM, 2016.

Wael Alghamdi, Juan Felipe Gomez, Shahab Asoodeh, Flavio Calmon, Oliver Kosut, and Lalitha Sankar. The Saddle-Point Method in Differential Privacy. In *International Conference on Machine Learning*, 2023.

Borja Balle and Yu-Xiang Wang. Improving the Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 394–403. PMLR, 2018.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Proceedings of the 2014 IEEE 55th annual symposium on foundations of computer science*, FOCS '14, pp. 464–473. IEEE Computer Society, 2014.

Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. *Machine Learning*, 94(3):401–437, 2014.

Mark Bun, Cynthia Dwork, Guy N. Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated CDP. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pp. 74–86. ACM, 2018.

Soham De, Leonard Berrada, Jamie Hayes, Samuel L. Smith, and Borja Balle. Unlocking high-accuracy differentially private image classification through scale. arXiv: 2204.13650, 2022. URL `https://arxiv.org/abs/2204.13650`.

Jinshuo Dong, Aaron Roth, and Weijie J. Su. Gaussian Differential Privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 02 2022.

Vadym Doroshenko, Badih Ghazi, Pritish Kamath, Ravi Kumar, and Pasin Manurangsi. Connect the dots: Tighter discrete approximations of privacy loss distributions. In *Proceedings on Privacy Enhancing Technologies*, volume 4, pp. 552–570, 2022.

Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *Third Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pp. 265–284. Springer, 2006.

Sivakanth Gopi, Yin Tat Lee, and Lukas Wutschitz. Numerical composition of differential privacy. *Advances in Neural Information Processing Systems*, 34, 2021.

Mark Kelbert. Survey of Distances between the Most Popular Distributions. *Analytics*, 2(1):225–245, 2023.

Antti Koskela, Joonas Jälkö, and Antti Honkela. Computing tight differential privacy guarantees using FFT. In *The 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2560–2569. PMLR, 2020.

Antti Koskela, Joonas Jälkö, Lukas Prediger, and Antti Honkela. Tight Differential Privacy for Discrete-Valued Mechanisms and for the Subsampled Gaussian Mechanism Using FFT. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, 2021.

S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

Ninghui Li, Wahbeh H. Qardaji, and Dong Su. On sampling, anonymization, and differential privacy or, *k*-anonymization meets differential privacy. In *7th ACM symposium on information, compuer and communications security*, pp. 32–33. ACM, 2012.

Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. Large Language Models Can Be Strong Differentially Private Learners. In *International Conference on Learning Representations*, 2021.

H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *International conference on learning representations*, 2018.

Harsh Mehta, Abhradeep Guha Thakurta, Alexey Kurakin, and Ashok Cutkosky. Towards Large Scale Transfer Learning for Differentially Private Image Classification. *Transactions on Machine Learning Research*, 2023.

Ilya Mironov. Rényi Differential Privacy. In *30th IEEE Computer Security Foundations Symposium*, pp. 263–275, 2017.

Ilya Mironov, Kunal Talwar, and Li Zhang. Rényi Differential Privacy of the Sampled Gaussian Mechanism. arXiv:1908.10530, 2019. URL https://arxiv.org/abs/1908.10530.

Tom Sander, Pierre Stock, and Alexandre Sablayrolles. TAN Without a Burn: Scaling Laws of DP-SGD. In *International Conference on Machine Learning*, 2023.

David M. Sommer, Sebastian Meiser, and Esfandiar Mohammadi. Privacy Loss Classes: The Central Limit Theorem in Differential Privacy. *PoPETs*, 2019(2):245–269, 2019.

Shuang Song, Kamalika Chaudhuri, and Anand D. Sarwate. Stochastic gradient descent with differentially private updates. In *IEEE Global Conference on Signal and Information Processing*, pp. 245–248. IEEE, 2013.

Razane Tajeddine, Joonas Jälkö, Samuel Kaski, and Antti Honkela. Privacy-preserving Data Sharing on Vertically Partitioned Data. arXiv:2010.09293, 2020. URL https://arxiv.org/abs/2010.09293.

Ashkan Yousefpour, Igor Shilov, Alexandre Sablayrolles, Davide Testuggine, Karthik Prasad, Mani Malek, John Nguyen, Sayan Ghosh, Akash Bharadwaj, Jessica Zhao, Graham Cormode, and Ilya Mironov. Opacus: User-Friendly Differential Privacy Library in PyTorch. arXiv:2109.12298, 2021. URL https://arxiv.org/abs/2109.12298v4.

Yuqing Zhu and Yu-Xiang Wang. Poisson subsampled Rényi differential privacy. In *International Conference on Machine Learning*, pp. 7634–7642, 2019.

Yuqing Zhu, Jinshuo Dong, and Yu-Xiang Wang. Optimal accounting of differential privacy via characteristic function. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 4782–4817. PMLR, 2022.

A Composition of Differential Privacy

Another useful property of DP is composition, which means that running multiple DP algorithms in succession degrades the privacy bounds in a predictable way. The tight composition theorem we are interested in is most easily expressed through privacy loss random variables (Sommer et al., 2019) and dominating pairs (Zhu et al., 2022), which we introduce next. Differential privacy can also be expressed through *hockey-stick divergence*.

Definition A.1. For $\alpha \geq 0$ and random variables P, Q, the hockey-stick divergence is

$$H_\alpha(P, Q) = \mathbb{E}_{t\sim Q}\left(\left(\frac{\mathrm{d}P}{\mathrm{d}Q}(t) - \alpha\right)_+\right), \tag{16}$$

where $(x)_+ = \max\{x, 0\}$ and $\frac{\mathrm{d}P}{\mathrm{d}Q}$ is the Radon-Nikodym derivative, which simplifies to the density ratio if P and Q are continuous.

Zhu et al. (2022) showed that $\mathcal{M}$ is (ϵ, δ)-DP if and only if $\sup_{x\sim x'} H_{e^\epsilon}(\mathcal{M}(x), \mathcal{M}(x')) \leq \delta$.

Computing the hockey-stick directly for a mechanism can be challenging. Dominating pairs allow computing privacy bounds for a mechanism from simpler distributions than the mechanism itself.

Definition A.2 (Zhu et al. 2022)**.** A pair of random variables P, Q is called a dominating pair for mechanism $\mathcal{M}$ if for all $\alpha \geq 0$,

$$\sup_{x\sim x'} H_\alpha(\mathcal{M}(x), \mathcal{M}(x')) \leq H_\alpha(P, Q). \tag{17}$$

Definition A.3 (Sommer et al. 2019)**.** For distributions P, Q, the privacy loss function is $\mathcal{L}(t) = \ln \frac{\mathrm{d}P}{\mathrm{d}Q}(t)$ and the privacy loss random variable (PLRV) is $L = \mathcal{L}(T)$ where $T \sim P$.

The PLRV allows computing the privacy bounds of the mechanism with a simple expectation (Sommer et al., 2019):

$$\delta(\epsilon) = \mathbb{E}_{s\sim L}((1 - e^{\epsilon - s})_+). \tag{18}$$

PLRVs also allow expressing the tight composition theorem in a simple way.

Theorem A.4 (Sommer et al. 2019)**.** *If $L_1, \ldots, L_T$ are the PLRVs for mechanisms $\mathcal{M}_1, \ldots, \mathcal{M}_T$, the PRLV of the adaptive composition of $\mathcal{M}_1, \ldots, \mathcal{M}_T$ is the convolution of $L_1, \ldots, L_T$, which is the distribution of $L_1 + \cdots + L_T$.*

A.1 DP-SGD Privacy Accounting

To compute the privacy bounds for DP-SGD, we need to account for the *subsampling amplification* that comes from the subsampling in SGD. This requires fixing the subsampling scheme. We consider Poisson subsampling, where each datapoint in $\mathcal{B}$ is included in the subsample with probability q, independently of any other datapoints.

When the neighbourhood relation is add/remove, privacy accounting for the Poisson subsampled Gaussian mechanism, of which DP-SGD is an instance of, can be done by analysing the following *dominating pair* of random variables (Koskela et al., 2020; Zhu et al., 2022)

$$P = q\mathcal{N}(\Delta, \sigma^2) + (1 - q)\mathcal{N}(0, \sigma^2), \qquad Q = \mathcal{N}(0, \sigma^2). \tag{19}$$

With this dominating pair, we can form the PLRV for the mechanism, take a T-fold convolution for T iterations, and compute the privacy bounds from the expectation in (18). The computation is not trivial, but can be done using numerical privacy accountants (Koskela et al., 2020; 2021; Gopi et al., 2021; Doroshenko et al., 2022; Alghamdi et al., 2023). These accountants are used by libraries implementing DP-SGD like Opacus (Yousefpour et al., 2021) to compute privacy bounds in practice.

B Properties of accounting oracles

Accounting oracles make it easy to formally express symmetries of privacy accounting. For example, privacy bounds are invariant to post-processing with a bijection.

Lemma B.1. *Let $\mathcal{M}$ be a mechanism and f be a bijection. Then* $\mathrm{AO}_{\mathcal{M}}(\epsilon) = \mathrm{AO}_{f\circ\mathcal{M}}(\epsilon)$.

Proof. This follows by using post-processing immunity for f to show that $\mathrm{AO}_{f\circ\mathcal{M}}(\epsilon) \leq \mathrm{AO}_{\mathcal{M}}(\epsilon)$ and for f^{-1} to show that $\mathrm{AO}_{f\circ\mathcal{M}}(\epsilon) \geq \mathrm{AO}_{\mathcal{M}}(\epsilon)$. □

We can also formalise the lemma that considering $\Delta = 1$ is sufficient when analysing the (subsampled) Gaussian mechanism.

Lemma B.2.

$$\mathrm{AO}_G(\sigma, \Delta, T, \epsilon) = \mathrm{AO}_G(\sigma/\Delta, 1, T, \epsilon), \tag{20}$$
$$\mathrm{AO}_S(\sigma, \Delta, q, T, \epsilon) = \mathrm{AO}_S(\sigma/\Delta, 1, q, T, \epsilon). \tag{21}$$

Proof. Let $\mathcal{M}$ be the (subsampled) Gaussian mechanism. Then

$$\mathcal{M}(x) = f(x) + \eta = \Delta\left(\frac{1}{\Delta}f(x) + \frac{1}{\Delta}\eta\right). \tag{22}$$

The sensitivity of $\frac{1}{\Delta}f(x)$ is 1, so the part inside parenthesis in the last expression is a (subsampled) Gaussian mechanism with sensitivity 1 and noise standard deviation σ/Δ. Multiplying by Δ is bijective post-processing, so that mechanism must have the same privacy bounds as the original mechanism. In a composition, this transformation can be done separately for each individual mechanism of the composition. □

Since considering $\Delta = 1$ when analysing the subsampled Gaussian mechanism is enough by Lemma B.2, we occasionally shorten $\mathrm{AO}_S(\sigma, 1, q, T, \epsilon)$ to $\mathrm{AO}_S(\sigma, q, T, \epsilon)$.

The next lemma and corollary show that mechanisms close in total variation distance also have similar privacy bounds.

Lemma B.3. *Let $\mathcal{M}$ be an (ϵ, δ)-DP mechanism, and let $\mathcal{M}'$ be a mechanism with*

$$\sup_x \mathrm{TV}(\mathcal{M}(x), \mathcal{M}'(x)) \leq d, \tag{23}$$

for some $d \geq 0$. Then $\mathcal{M}'$ is $(\epsilon, \delta + (1 + e^\epsilon)d)$-DP.

Proof. For any measurable set A,

$$\begin{aligned}
\Pr(\mathcal{M}'(x) \in A) &\leq \Pr(\mathcal{M}(x) \in A) + d \\
&\leq e^\epsilon \Pr(\mathcal{M}(x') \in A) + \delta + d \\
&\leq e^\epsilon(\Pr(\mathcal{M}'(x') \in A) + d) + \delta + d \\
&= e^\epsilon \Pr(\mathcal{M}'(x') \in A) + \delta + (1 + e^\epsilon)d.
\end{aligned} \tag{24}$$

□

Lemma B.3 can also be expressed with accounting oracles.

Corollary B.4. *Let $\mathrm{AO}_{\mathcal{M}}$ and $\mathrm{AO}_{\mathcal{M}'}$ be the accounting oracles for mechanisms $\mathcal{M}$ and $\mathcal{M}'$, respectively. If*

$$\sup_x \mathrm{TV}(\mathcal{M}(x), \mathcal{M}'(x)) \leq d. \tag{25}$$

then

$$|\mathrm{AO}_{\mathcal{M}}(\epsilon) - \mathrm{AO}_{\mathcal{M}'}(\epsilon)| \leq (1 + e^\epsilon)d. \tag{26}$$

Proof. Let $\delta = \mathrm{AO}_{\mathcal{M}}(\epsilon)$. By Lemma B.3, $\mathcal{M}'$ is $(\epsilon, (1+e^\epsilon)d+\delta)$-DP, so $\mathrm{AO}_{\mathcal{M}'}(\epsilon) \leq \delta+(1+e^\epsilon)d$. If $\mathrm{AO}_{\mathcal{M}'}(\epsilon) < \delta - (1 + e^\epsilon)d$,

$$\mathrm{AO}_{\mathcal{M}}(\epsilon) < \delta - (1 + e^\epsilon)d + (1 + e^\epsilon)d = \delta \tag{27}$$

by Lemma B.3, which is a contradiction, so $\mathrm{AO}_{\mathcal{M}'}(\epsilon) \geq \delta - (1 + e^\epsilon)d$. Putting these together,

$$|\mathrm{AO}_{\mathcal{M}}(\epsilon) - \mathrm{AO}_{\mathcal{M}'}(\epsilon)| \leq (1 + e^\epsilon)d. \tag{28}$$

□

C GRADIENT VARIANCE AND BIAS IN DP-SGD

C.1 SCALING G_q WITH $1/q$ GIVES AN UNBIASED ESTIMATOR OF G_1

Recall $G_q = \sum_{i\in\mathcal{B}} \text{clip}_C(g_i)$. We have

$$\mathbb{E}\left[\frac{1}{q}G_q - G_1\right] = \mathbb{E}\left[\frac{1}{q}\sum_{i\in[N]} b_i\tilde{g}_i - \sum_{i\in[N]}\tilde{g}_i\right] = \sum_{i\in[N]}\tilde{g}_i\,\mathbb{E}\left[\frac{b_i}{q} - 1\right] = 0, \tag{29}$$

where $b_i \sim \text{Bernoulli}(q)$ and $\tilde{g}_i$ denote the clipped per-example gradients. Thus, G_q/q is an unbiased estimator of the G_1

C.2 THE SUBSAMPLING INDUCED VARIANCE DECREASES W.R.T. q

Recall from Section 4, that we denote the sum of clipped subsampled gradients with G_q:

$$G_q = \sum_{i\in\mathcal{B}} \text{clip}_C(g_i). \tag{30}$$

Now for the subsampling induced variance in noise decomposition of Equation (4) we have

$$\text{Var}\left(\frac{1}{q}G_{q,j}\right) = \frac{1}{q^2}\,\text{Var}\left(\sum_{i\in[N]} b_i g_{i,j}\right) \tag{31}$$

$$= \frac{1}{q^2}\sum_{i\in[N]} g_{i,j}^2\,\text{Var}\left(b_i\right) \tag{32}$$

$$= \frac{q(1-q)}{q^2}\sum_{i\in[N]} g_{i,j}^2 \tag{33}$$

$$= \frac{1-q}{q}\sum_{i\in[N]} g_{i,j}^2. \tag{34}$$

Now, it is easy to see that the sum in Equation (34) is a constant w.r.t. q, and the term $(1-q)/q$ is decreasing w.r.t. q. Thus the subsampling induced variance is decreasing w.r.t. q.

D MISSING PROOFS IN SECTION 5

D.1 KULLBACK-LEIBLER DIVERGENCE AND TOTAL VARIATION DISTANCE

Our proofs use two notions of distance between random variables, *total variation distance*, and *Kullback-Leibler (KL) divergence* (Kullback & Leibler, 1951).

Definition D.1. Let P and Q be random variables.

1. The total variation distance between P and Q is
$$\text{TV}(P,Q) = \sup_A |\Pr(P\in A) - \Pr(Q\in A)|. \tag{35}$$
The supremum is over all measurable sets A.
2. The KL divergence between P and Q with densities $p(t)$ and $q(t)$ is
$$\text{KL}(P,Q) = \mathbb{E}_{t\sim P}\left(\ln\frac{p(t)}{q(t)}\right). \tag{36}$$

The two notions of distance are related by Pinsker's inequality.

Lemma D.2 (Kelbert, 2023)**.**

$$\text{TV}(P,Q) \le \sqrt{\frac{1}{2}\text{KL}(P,Q)}. \tag{37}$$

D.2 USEFUL LEMMAS

Lemma D.3 (Kullback & Leibler, 1951)**.** *Properties of KL divergence:*

1. *If P and Q are joint distributions over independent random variables $P_1, \ldots, P_T$ and $Q_1, \ldots, Q_T$,*
$$\mathrm{KL}(P, Q) = \sum_{i=1}^{T} \mathrm{KL}(P_i, Q_i). \tag{38}$$
2. *If f is a bijection,*
$$\mathrm{KL}(f(P), f(Q)) = \mathrm{KL}(P, Q). \tag{39}$$

We will need to analyse the following function:

$$f_x(u) = \ln \frac{\mathcal{N}(x; 0, 1)}{q\mathcal{N}(x; u - qu, 1) + (1 - q)\mathcal{N}(x; -qu, 1)}. \tag{40}$$

In particular, we need the fourth-order Taylor approximation of $\mathbb{E}_x(f_x(u))$ for $x \sim \mathcal{N}(0, 1)$ at $u = 0$. We begin by looking at the Taylor approximation of $f_x(u)$ without the expectation, and then show that we can differentiate under the expectation.

Lemma D.4. *The fourth-order Taylor approximation of f_x at $u = 0$ is*

$$\begin{aligned} f_x(u) = &\frac{1}{2}(q-1)q(x^2-1)u^2 \\ &- \frac{1}{6}(q-1)q(2q-1)x(x^2-3)u^3 \\ &+ \frac{1}{24}q\left(-3 + 6x^2 - x^4 - 12q^2(2 - 4x^2 + x^4) + 6q^3(2 - 4x^2 + x^4) + q(15 - 30x^2 + 7x^4)\right)u^4 \\ &+ r_x(u)u^4, \end{aligned} \tag{41}$$

with $\lim_{u \to 0} r_x(u) = 0$.

Proof. The claim follows from Taylor's theorem after computing the first four derivatives of f_x at $u = 0$. We computed the derivatives with Mathematica. □

Lemma D.5. *When $x \sim \mathcal{N}(0, 1)$,*

$$\mathbb{E}_x\left(\frac{1}{2}(q-1)q(x^2-1)u^2\right) = 0 \tag{42}$$

$$\mathbb{E}_x\left(-\frac{1}{6}(q-1)q(2q-1)x(x^2-3)u^3\right) = 0 \tag{43}$$

$$\mathbb{E}_x\left(\frac{f_x^{(4)}(0)}{4!}u^4\right) = \frac{1}{4}(q-1)^2q^2u^4. \tag{44}$$

Proof. It is well-known that $\mathbb{E}(x) = 0, \mathbb{E}(x^2) = 1, \mathbb{E}(x^3) = 0$ and $\mathbb{E}(x^4) = 3$. The first expectation:

$$\begin{aligned} \mathbb{E}_x\left(\frac{1}{2}(q-1)q(x^2-1)u^2\right) &= \frac{1}{2}(q-1)qu^2\,\mathbb{E}_x\left(x^2-1\right) \\ &= \frac{1}{2}(q-1)qu^2\left(\mathbb{E}_x(x^2)-1\right) \\ &= 0. \end{aligned} \tag{45}$$

The second expectation:

$$\begin{aligned} \mathbb{E}_x\left(-\frac{1}{6}(q-1)q(2q-1)x(x^2-3)u^3\right) &= -\frac{1}{6}(q-1)q(2q-1)u^3\,\mathbb{E}_x\left(x(x^2-3)\right) \\ &= -\frac{1}{6}(q-1)q(2q-1)u^3\left(\mathbb{E}_x(x^3) - 3\,\mathbb{E}_x(x)\right) \\ &= 0. \end{aligned} \tag{46}$$

The third expectation:

$$\begin{aligned}
&\mathbb{E}_x\left(\frac{f_x^{(4)}(0)}{4!}u^4\right)\\
&=\mathbb{E}_x\left(\frac{1}{24}q\left(-3+6x^2-x^4-12q^2(2-4x^2+x^4)+6q^3(2-4x^2+x^4)+q(15-30x^2+7x^4)\right)u^4\right)\\
&=\frac{1}{24}q\left(-3+\mathbb{E}_x(6x^2)-\mathbb{E}_x(x^4)-12q^2\,\mathbb{E}_x(2-4x^2+x^4)+6q^3\,\mathbb{E}_x(2-4x^2+x^4)+q\,\mathbb{E}_x(15-30x^2+7x^4)\right)u^4\\
&=\frac{1}{24}q\left(-3+6-3-12q^2(2-4+3)+6q^3(2-4+3)+q(15-30+21)\right)u^4\\
&=\frac{1}{24}q\left(-12q^2+6q^3+6q\right)u^4\\
&=\frac{1}{4}q^2\left(-2q+q^2+1\right)u^4\\
&=\frac{1}{4}(q-1)^2q^2u^4.
\end{aligned}\tag{47}$$

□

Lemmas D.4 and D.5 show that the Taylor approximation of $\mathbb{E}_x(f_x(u))$ is

$$\mathbb{E}_x(f_x(u))=\frac{1}{4}(q-1)^2q^2u^4+r(u)u^4 \tag{48}$$

if we can differentiate under the expectation. Next, we show that this is possible in Lemma D.9, which requires several preliminaries.

Definition D.6. A function $g(x,u)$ is a polynomial-exponentiated simple polynomial (PESP) if

$$g(x,u)=\sum_{i=1}^{n}P_i(x,u)e^{Q_i(x,u)} \tag{49}$$

for some $n\in\mathbb{N}$ and polynomials $P_i(x,u)$ and $Q_i(x,u)$, $1\le i\le n$, with $Q_i(x,u)$ being first-degree in x.

Lemma D.7. *If g_1 and g_2 are PESPs,*

1. *g_1+g_2 is a PESP,*
2. *$g_1\cdot g_2$ is a PESP,*
3. *$\frac{\partial}{\partial u}g_1$ is a PESP.*

Proof. Let

$$g_j(x,u)=\sum_{i=1}^{n_j}P_{i,j}(x,u)e^{Q_{i,j}(x,u)} \tag{50}$$

for $j\in\{1,2\}$. (1) is clear by just writing the sums in g_1 and g_2 as a single sum. For (2),

$$\begin{aligned}
g_1(x,u)\cdot g_2(x,u)&=\sum_{i=1}^{n_1}\sum_{j=1}^{n_2}P_{i,1}(x,u)P_{j,2}(x,u)e^{Q_{i,1}(x,u)}e^{Q_{j,2}(x,u)}\\
&=\sum_{i=1}^{n_1}\sum_{j=1}^{n_2}P_{i,1}(x,u)P_{j,2}(x,u)e^{Q_{i,1}(x,u)+Q_{j,2}(x,u)}\\
&=\sum_{i=1}^{n_3}P_{i,3}(x,u)e^{Q_{i,3}(x,u)}
\end{aligned}\tag{51}$$

since the product of two polynomials is a polynomial, and the sum of two polynomials is a polynomial of the same degree.

For (3)

$$\begin{aligned}\frac{\partial}{\partial u}g_1(x,u) &= \sum_{i=1}^{n_1}\left(\frac{\partial}{\partial u}P_{i,1}(x,u)\right)e^{Q_{i,1}(x,u)} + \sum_{i=1}^{n_1}P_{i,1}(x,u)\left(\frac{\partial}{\partial u}Q_{i,1}(x,u)\right)e^{Q_{i,1}(x,u)} \\ &= \sum_{i=1}^{n_1}P_{i,4}(x,u)e^{Q_{i,1}(x,u)}\end{aligned} \tag{52}$$

since the partial derivatives, products and sums of polynomials are polynomials. □

Lemma D.8. *When $x \sim \mathcal{N}(0,1)$, for $a > 0, b \in \mathbb{R}$ and $k \in \mathbb{N}$,*

$$\mathbb{E}_x(a|x|^k e^{b|x|}) < \infty. \tag{53}$$

Proof.

$$\begin{aligned}\mathbb{E}_x(a|x|^k e^{b|x|}) &= \int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi}}a|x|^k e^{b|x|}e^{-\frac{1}{2}x^2}\mathrm{d}x \\ &\propto \int_{-\infty}^{\infty}|x|^k e^{b|x|}e^{-\frac{1}{2}x^2}\mathrm{d}x \\ &= 2\int_0^{\infty}x^k e^{bx}e^{-\frac{1}{2}x^2}\mathrm{d}x \\ &= 2\int_0^{\infty}x^k e^{-\frac{1}{2}(x^2-2bx)}\mathrm{d}x \\ &\propto \int_0^{\infty}x^k e^{-\frac{1}{2}(x^2-2bx+b^2)}\mathrm{d}x \\ &= \int_0^{\infty}x^k e^{-\frac{1}{2}(x-b)^2}\mathrm{d}x \\ &\leq \int_{-\infty}^{\infty}|x|^k e^{-\frac{1}{2}(x-b)^2}\mathrm{d}x \\ &< \infty\end{aligned} \tag{54}$$

since all absolute moments of Gaussian distributions are finite. □

Lemma D.9. *For any $k \in \mathbb{N}$, $k \geq 1$, there is a function $g_k(x)$ such that $|f_x^{(k)}(u)| \leq g_k(x)$ for all $u \in [-1,1]$ and $x \in \mathbb{R}$, and $\mathbb{E}_x(g_k(x)) < \infty$.*

Proof. We start by computing the first derivative of $f_x(u)$

$$\begin{aligned}f_x(u) &= \ln\frac{e^{-\frac{1}{2}x^2}}{qe^{-\frac{1}{2}((q-1)u+x)^2} + (1-q)e^{-\frac{1}{2}(qu+x)^2}} \\ &= -\ln\left(qe^{-\frac{1}{2}((q-1)u+x)^2} + (1-q)e^{-\frac{1}{2}(qu+x)^2}\right) - \frac{x^2}{2}\end{aligned} \tag{55}$$

$$f_x'(u) = -\frac{-q(q-1)((q-1)u+x)e^{-\frac{1}{2}((q-1)u+x)^2} - (1-q)q(qu+x)e^{-\frac{1}{2}(qu+x)^2}}{qe^{-\frac{1}{2}((q-1)u+x)^2} + (1-q)e^{-\frac{1}{2}(qu+x)^2}} \tag{56}$$

Since

$$e^{-\frac{1}{2}((q-1)u+x)^2} = e^{-\frac{1}{2}((q-1)^2u^2+2(q-1)xu+x^2)} = e^{-\frac{1}{2}x^2}e^{-\frac{1}{2}((q-1)^2u^2+2(q-1)xu)} \tag{57}$$

and

$$e^{-\frac{1}{2}(qu+x)^2} = e^{-\frac{1}{2}(q^2u^2+2qux+x^2)} = e^{-\frac{1}{2}x^2}e^{-\frac{1}{2}(q^2u^2+2qux)} \tag{58}$$

we can write the first derivative in the following form:

$$f'_x(u) = \frac{P_1(x,u)e^{Q_1(x,u)} + P_2(x,u)e^{Q_2(x,u)}}{qe^{Q_1(x,u)} + (1-q)e^{Q_2(x,u)}} \tag{59}$$

for polynomials $P_1(x,u)$, $P_2(x,u)$, $Q_1(x,u)$ and $Q_2(x,u)$, with $Q_1(x,u)$ and $Q_2(x,u)$ being first-degree in x.

We show by induction that further derivatives have a similar form:

$$f_x^{(k)}(u) = \frac{\sum_{i=1}^{n_k} P_{i,k}(x,u)e^{Q_{i,k}(x,u)}}{(qe^{Q_1(x,u)} + (1-q)e^{Q_2(x,u)})^{2^{k-1}}} \tag{60}$$

We also require the Q-polynomials to still be first degree in x, so the numerator must be a PESP. The claim is clearly true for $k = 1$. If the claim is true for k, then

$$f_x^{(k+1)}(u) = \frac{G(x,u)\frac{\partial}{\partial u}F(x,u) - F(x,u)\frac{\partial}{\partial u}G(x,u)}{(qe^{Q_1(x,u)} + (1-q)e^{Q_2(x,u)})^{2^k}} \tag{61}$$

where

$$F(x,u) = \sum_{i=1}^{n_k} P_{i,k}(x,u)e^{Q_{i,k}(x,u)} \tag{62}$$

$$G(x,u) = (qe^{Q_1(x,u)} + (1-q)e^{Q_2(x,u)})^{2^{k-1}} \tag{63}$$

by the quotient differentiation rule. The denominator has the correct form, so it remains to show that the numerator is a PESP. F is clearly a PESP, and so is G due to Lemma D.7. The numerator is a sum of products of F, G and their derivatives, so it is a PESP by Lemma D.7, which concludes the induction proof.

We have

$$qe^{Q_1(x,u)} + (1-q)e^{Q_2(x,u)} \geq \min\left(e^{Q_1(x,u)}, e^{Q_2(x,u)}\right). \tag{64}$$

We can split $\mathbb{R}$ into measurable subsets A_1 and A_2 such that

$$\min\left(e^{Q_1(x,u)}, e^{Q_2(x,u)}\right) = \begin{cases} e^{Q_1(x,u)} & x \in A_1 \\ e^{Q_2(x,u)} & x \in A_2 \end{cases} \tag{65}$$

Now

$$|f_x^{(k)}(u)| \leq \left|\frac{\sum_{i=1}^{n_k} P_{i,k}(x,u)e^{Q_{i,k}(x,u)}}{\min\left(e^{Q_1(x,u)}, e^{Q_2(x,u)}\right)^{2^{k-1}}}\right| \tag{66}$$

$$= \left|I_{A_1}\frac{\sum_{i=1}^{n_k} P_{i,k}(x,u)e^{Q_{i,k}(x,u)}}{(e^{Q_1(x,u)})^{2^{k-1}}}\right| + \left|I_{A_2}\frac{\sum_{i=1}^{n_k} P_{i,k}(x,u)e^{Q_{i,k}(x,u)}}{(e^{Q_2(x,u)})^{2^{k-1}}}\right| \tag{67}$$

$$\leq \left|\frac{\sum_{i=1}^{n_k} P_{i,k}(x,u)e^{Q_{i,k}(x,u)}}{e^{2^{k-1}Q_1(x,u)}}\right| + \left|\frac{\sum_{i=1}^{n_k} P_{i,k}(x,u)e^{Q_{i,k}(x,u)}}{e^{2^{k-1}Q_2(x,u)}}\right| \tag{68}$$

$$\leq \sum_{i=1}^{m_k} |R_{i,k}(x,u)|e^{S_{i,k}(x,u)} \tag{69}$$

where $R_{i,k}(x,u)$ and $S_{i,k}(x,u)$ are further polynomials, with the S-polynomials being of first degree in x.

Since $u \in [-1,1]$, for a monomial $ax^{k_1}u^{k_2}$

$$ax^{k_1}u^{k_2} \leq |ax^{k_1}u^{k_2}| \leq |ax^{k_1}|. \tag{70}$$

Using this inequality on each monomial of $R_{i,k}(x,u)$ and $S_{i,k}(x,u)$ gives upper bound polynomials of $|x|$ $\hat{R}_{i,k}(x)$ and $\hat{S}_{i,k}(x)$ such that

$$|f_x^{(k)}(u)| \leq \sum_{i=1}^{m_k} |R_{i,k}(x,u)|e^{S_{i,k}(x,u)} \leq \sum_{i=1}^{m_k} |\hat{R}_{i,k}(x)|e^{\hat{S}_{i,k}(x)}, \tag{71}$$

with the $\hat{S}$-polynomials being first degree.

Let

$$g_k(x) = \sum_{i=1}^{m_k} |\hat{R}_{i,k}(x)| e^{\hat{S}_{i,k}(x)}. \tag{72}$$

We have shown that $|f_x(u)| \leq g_k(x)$ for $u \in [-1, 1]$ and $x \in \mathbb{R}$. The integrability of $g_k(x)$ against a standard Gaussian follows from Lemma D.8, as we can first push the absolute value around $\hat{R}_{i,k}(x)$ to be around each monomial of $\hat{R}_{i,k}(x)$ with the triangle inequality, and then write the resulting upper bound as a sum with each term of the form $a|x|^k e^{b|x|}$, with $a > 0$, $b \in \mathbb{R}$ and $k \in \mathbb{N}$. □

Now we can put the preliminaries together to use the Taylor approximation of of $\mathbb{E}_x(f_x(u))$ to find its order of convergence.

Lemma D.10. *When $x \sim \mathcal{N}(0, 1)$, $\mathbb{E}_x(f_x(u)) = O(u^4)$ as $u \to 0$.*

Proof. Since $u \to 0$ in the limit, it suffices to consider $u \in [-1, 1]$. First, we find the fourth-order Taylor approximation of $\mathbb{E}_x(f_x(u))$. Lemma D.9 allows us to differentiate under the expectation four times. Then Taylor's theorem, and Lemmas D.4 and D.5 give

$$\mathbb{E}_x(f_x(u)) = \frac{1}{4}(q-1)^2 q^2 u^4 + r(u)u^4 \tag{73}$$

where $\lim_{u\to 0} r(u) = 0$. Now

$$\begin{aligned} \lim_{u\to 0} \frac{1}{u^4} \mathbb{E}(f_x(u)) &= \lim_{u\to 0} \frac{1}{u^4} \left(\frac{1}{4}(q-1)^2 q^2 u^4 + r(u)u^4 \right) \\ &= \frac{1}{4}(q-1)^2 q^2 + \lim_{u\to 0} r(u) \\ &= \frac{1}{4}(q-1)^2 q^2 \\ &< \infty, \end{aligned} \tag{74}$$

which implies the claim. □

D.3 Proof of Theorem 5.2

Theorem 5.2. *For $1 \leq i \leq T$, let*

$$\mathcal{M}_i(x) \sim xB_q + \mathcal{N}(0, \sigma_T^2), \tag{7}$$

$$\mathcal{M}_i'(x) \sim \mathcal{N}(qx, \sigma_T^2). \tag{8}$$

be independent for each i. Let $\mathcal{M}_{1:T}$ be the composition of $\mathcal{M}_i$, and let $\mathcal{M}'_{1:T}$ be the composition of $\mathcal{M}_i'$. Then

$$\sup_x \mathrm{TV}(\mathcal{M}_{1:T}(x), \mathcal{M}'_{1:T}(x)) \to 0. \tag{9}$$

Proof. It suffices to show

$$\sup_x T \cdot \mathrm{KL}(\mathcal{M}_i'(x), \mathcal{M}_i(x)) \to 0 \tag{75}$$

due to Pinsker's inequality and the additivity of KL divergence for products of independent random variables (Lemma D.3). When $x = 0$, the two mechanism are the same, so it suffices to look at $x = 1$.

KL-divergence is invariant to bijections, so

$$\begin{aligned} \mathrm{KL}(\mathcal{M}_i'(1), \mathcal{M}_i(1)) &= \mathrm{KL}\big(\mathcal{N}(q, \sigma_T^2), q\mathcal{N}(1, \sigma_T^2) + (1-q)\mathcal{N}(0, \sigma_T^2)\big) \\ &= \mathrm{KL}\left(\mathcal{N}\left(q\frac{1}{\sigma_T}, 1\right), q\mathcal{N}\left(\frac{1}{\sigma_T}, 1\right) + (1-q)\mathcal{N}\left(0, 1\right)\right) \\ &= \mathrm{KL}\left(\mathcal{N}\left(qu, 1\right), q\mathcal{N}\left(u, 1\right) + (1-q)\mathcal{N}\left(0, 1\right)\right) \\ &= \mathrm{KL}\left(\mathcal{N}\left(0, 1\right), q\mathcal{N}\left(u - qu, 1\right) + (1-q)\mathcal{N}\left(-qu, 1\right)\right) \end{aligned} \tag{76}$$

where we first divide both distributions by σ_T, then set $u = \frac{1}{\sigma_T}$, and finally subtract qu. As $\sigma_T^2 = \Omega(T)$, $u = O\left(\frac{1}{\sqrt{T}}\right)$.

From the definition of KL-divergence, $u^4 = O(\frac{1}{T^2})$ and Lemma D.10, when $x \sim \mathcal{N}(0,1)$ we have

$$\mathrm{KL}(\mathcal{M}_i'(1), \mathcal{M}_i(1)) = \mathbb{E}_x(f_x(u)) = O(u^4) = O\left(\frac{1}{T^2}\right). \tag{77}$$

This implies

$$\lim_{T\to\infty} T \cdot \mathrm{KL}(\mathcal{M}_i'(1), \mathcal{M}_i(1)) = 0, \tag{78}$$

which implies the claim. □

D.4 Proof for Theorem 5.3

Theorem 5.3. *For any σ, q_1, q_2, Δ and ϵ*

$$|\mathrm{AO}_S(\sigma, \Delta, q_1, T, \epsilon) - \mathrm{AO}_S(\sigma \cdot q_2/q_1, \Delta, q_2, T, \epsilon)| \to 0 \tag{11}$$

as $T \to \infty$.

Proof. It suffices to look at $\Delta = 1$ by Lemma B.2. Let $\mathrm{AO}'(\sigma, q, T, \epsilon)$ be the accounting oracle for $\mathcal{M}_{1:T}'$. By Lemma B.1 and (10),

$$\mathrm{AO}'(\sigma, q, T, \epsilon) = \mathrm{AO}_G\left(\frac{\sigma}{q}, T, \epsilon\right). \tag{79}$$

Let $\sigma_2 = \sigma \cdot q_2/q_1$,

$$d_T^{(1)} = \sup_x \mathrm{TV}(\mathcal{M}_{1:T}(x, q_1, \sigma), \mathcal{M}_{1:T}'(x, q_1, \sigma)), \tag{80}$$

and

$$d_T^{(2)} = \sup_x \mathrm{TV}(\mathcal{M}_{1:T}(x, q_2, \sigma_2), \mathcal{M}_{1:T}'(x, q_2, \sigma_2)). \tag{81}$$

Now

$$|\mathrm{AO}_S(\sigma, q_1, T, \epsilon) - \mathrm{AO}'(\sigma, q_1, T, \epsilon)| \le (1 + e^\epsilon) d_T^{(1)} \tag{82}$$

and

$$|\mathrm{AO}_S(\sigma_2, q_2, T, \epsilon) - \mathrm{AO}'(\sigma_2, q_2, T, \epsilon)| \le (1 + e^\epsilon) d_T^{(2)} \tag{83}$$

by Corollary B.4. Since $\frac{\sigma}{q_1} = \frac{\sigma_2}{q_2}$,

$$\begin{aligned} \mathrm{AO}'(\sigma, q_1, T, \epsilon) &= \mathrm{AO}_G\left(\frac{\sigma}{q_1}, T, \epsilon\right) \\ &= \mathrm{AO}_G\left(\frac{\sigma_2}{q_2}, T, \epsilon\right) \\ &= \mathrm{AO}'(\sigma_2, q_2, T, \epsilon). \end{aligned} \tag{84}$$

Now

$$\begin{aligned} &|\mathrm{AO}_S(\sigma, q_1, T, \epsilon) - \mathrm{AO}_S(\sigma \cdot q_2/q_1, q_2, T, \epsilon)| \\ &\le |\mathrm{AO}_S(\sigma, q_1, T, \epsilon) - \mathrm{AO}'(\sigma, q_1, T, \epsilon)| \\ &\quad + |\mathrm{AO}_S(\sigma_2, q_2, T, \epsilon) - \mathrm{AO}'(\sigma_2, q_2, T, \epsilon)| \\ &\le (1 + e^\epsilon)(d_T^{(1)} + d_T^{(2)}) \to 0 \end{aligned} \tag{85}$$

by Theorem 5.2. □

D.5 PROOF OF COROLLARY 5.4

With fixed Δ, T, ϵ, the function $\sigma \mapsto \mathrm{AO}_G(\sigma, \Delta, T, \epsilon)$ is strictly decreasing (Balle & Wang, 2018, Lemma 7) and continuous, so it has a continuous inverse $\delta \mapsto \mathrm{AO}_G^{-1}(\delta, \Delta, T, \epsilon)$. To declutter the notation, we omit the ϵ and Δ arguments from AO_S, AO_G and AO_G^{-1} in the rest of this section.

Lemma D.11. *For AO_G and its inverse,*

1. $\mathrm{AO}_G(\sigma, T) = \mathrm{AO}_G\left(\frac{\sigma}{\sqrt{T}}, 1\right)$,
2. $\mathrm{AO}_G^{-1}(\delta, T) = \mathrm{AO}_G^{-1}(\delta, 1)\sqrt{T}$.

Proof. Recall that the PLRV of the Gaussian mechanism is $\mathcal{N}(\mu_1, 2\mu_1)$ with $\mu_1 = \frac{\Delta^2}{2\sigma^2}$ (Sommer et al., 2019). By Theorem A.4, the PLRV of T compositions of the Gaussian mechanism is $\mathcal{N}(T\mu_1, 2T\mu_1)$. Denoting $\mu_T = T\mu_1 = \frac{T\Delta^2}{2\sigma^2}$, we see that the T-fold composition of the Gaussian mechanism has the same PLRV as a single composition of the Gaussian mechanism with standard deviation $\frac{\sigma}{\sqrt{T}}$, which proves (1).

To prove (2), first we have

$$\mathrm{AO}_G(\mathrm{AO}_G^{-1}(\delta, 1)\sqrt{T}, T) = \mathrm{AO}_G(\mathrm{AO}_G^{-1}(\delta, 1), 1) = \delta \tag{86}$$

by applying (1) to the outer AO_G. Applying $\mathrm{AO}_G^{-1}(\cdot, T)$ to both sides gives

$$\mathrm{AO}_G^{-1}(\delta, 1)\sqrt{T} = \mathrm{AO}_G^{-1}(\delta, T) \tag{87}$$

□

Corollary 5.4. *Let $\sigma(q, T)$ be the smallest σ such that $\mathrm{AO}_S(\sigma, \Delta, q, T, \epsilon) \leq \delta$. Then*

$$\lim_{T\to\infty} \frac{\sigma(q, T)}{q\sigma(1, T)} = 1. \tag{12}$$

Proof. By definition, $\mathrm{AO}_S(\sigma(q, T), q, T) = \delta$, so Theorem 5.3 implies

$$\lim_{T\to\infty} \mathrm{AO}_S\left(\frac{\sigma(q, T)}{q}, 1, T\right) = \delta. \tag{88}$$

Since $\mathrm{AO}_S(\sigma, 1, T) = \mathrm{AO}_G(\sigma, T)$ for any σ, we get from Lemma D.11

$$\left|\mathrm{AO}_G\left(\frac{\sigma(q, T)}{q\sqrt{T}}, 1\right) - \delta\right| = \left|\mathrm{AO}_G\left(\frac{\sigma(q, T)}{q}, T\right) - \delta\right| \to 0 \tag{89}$$

as $T \to \infty$. We have

$$\mathrm{AO}_G^{-1}\left(\mathrm{AO}_G\left(\frac{\sigma(q, T)}{q\sqrt{T}}, 1\right), 1\right) = \frac{\sigma(q, T)}{q\sqrt{T}} \tag{90}$$

and by Lemma D.11,

$$\mathrm{AO}_G^{-1}(\delta, 1) = \frac{1}{\sqrt{T}}\mathrm{AO}_G^{-1}(\delta, T) = \frac{\sigma(1, T)}{\sqrt{T}}. \tag{91}$$

Now, by the continuity of $\mathrm{AO}_G^{-1}(\cdot, 1)$,

$$\frac{1}{\sqrt{T}}\left|\frac{\sigma(q, T)}{q} - \sigma(1, T)\right| \to 0 \tag{92}$$

as $T \to \infty$.

Since $\sigma(1, T) = \Omega(\sqrt{T})$ by Theorem 5.1,

$$\left|\frac{\sigma(q, T)}{q\sigma(1, T)} - 1\right| = \frac{|\sigma(q, T)/q - \sigma(1, T)|}{\sigma(1, T)} \to 0 \tag{93}$$

as $T \to \infty$, which implies the claim. □

D.6 $\sigma_T^2 = \Omega(T)$

Lemma D.12. *Let L be the PLRV of a single iteration of the Poisson subsampled Gaussian mechanism. Then $\mathbb{E}(L) \geq \frac{q^2}{2\sigma^2}$.*

Proof. Recall, that for Poisson subsampled Gaussian mechanism we have the dominating pair $P = q\mathcal{N}(1, \sigma^2) + (1-q)\mathcal{N}(0, \sigma^2)$ and $Q = \mathcal{N}(0, \sigma^2)$. Let f_P and f_Q be their densities. Now the mean of the PLRV can be written as

$$\begin{aligned}
\mathbb{E}(L) &= \mathbb{E}_{t\sim P}\left[\log \frac{f_P(t)}{f_Q(t)}\right] \\
&= \mathbb{E}_{t\sim P}[\log f_P(t)] - \mathbb{E}_{t\sim P}[\log f_Q(t)] \\
&= -H(P) - \left(-\frac{1}{2}\log 2\pi\sigma^2 - \frac{1}{2\sigma^2}\,\mathbb{E}_{t\sim P}[t^2]\right) \\
&= -H(P) + \frac{1}{2}\left(\log 2\pi + \log\sigma^2 + \frac{1}{\sigma^2}(\sigma^2 + q)\right) \\
&= -H(P) + \frac{1}{2}\left(\log 2\pi + \log\sigma^2 + \frac{q}{\sigma^2} + 1\right),
\end{aligned} \tag{94}$$

where H denotes the differential entropy. The entropy term is analytically intractable for the mixture of Gaussians (the P). However, we can upper bound it with the entropy of a Gaussian with the same variance, as the Gaussian distribution maximises entropy among distributions with given mean and variance. The variance of P is $\sigma^2 + q - q^2 = \sigma^2 + q(1-q)$, and therefore

$$\begin{aligned}
H(P) &\leq H(\mathcal{N}(0, \sigma^2 + q(1-q))) \\
&= \frac{1}{2}\left(\log 2\pi(\sigma^2 + q(1-q)) + 1\right) \\
&= \frac{1}{2}\left(\log 2\pi + \log\big(\sigma^2 + q(1-q)\big) + 1\right)
\end{aligned} \tag{95}$$

Now, substituting this into (94) we get the following lower bound

$$\mathbb{E}(L) \geq \frac{1}{2}\left(\log\frac{\sigma^2}{\sigma^2 + q(1-q)} + \frac{q}{\sigma^2}\right) = \frac{1}{2}\left(-\log\left(1 + \frac{q - q^2}{\sigma^2}\right) + \frac{q}{\sigma^2}\right). \tag{96}$$

Since for all $x \geq -1$,

$$\log(1+x) \leq x \Leftrightarrow -\log(1+x) \geq -x, \tag{97}$$

we have

$$-\log\left(1 + \frac{q - q^2}{\sigma^2}\right) \geq -\frac{q - q^2}{\sigma^2} \tag{98}$$

which gives

$$\mathbb{E}(L) \geq \frac{1}{2}\left(-\frac{q - q^2}{\sigma^2} + \frac{q}{\sigma^2}\right) = \frac{q^2}{2\sigma^2}. \tag{99}$$

□

Lemma D.13. *Let L be the PLRV of a single iteration of the Poisson subsampled Gaussian mechanism. Then $\mathrm{Var}(L) \leq \frac{1}{\sigma^2} + \frac{1}{4\sigma^4}$*

Proof.

$$\begin{aligned}
\ln \frac{f_P(t)}{f_Q(t)} &= \ln \frac{q \exp\left(-\frac{(t-1)^2}{2\sigma^2}\right) + (1-q)\exp\left(-\frac{t}{2\sigma^2}\right)}{\exp\left(-\frac{t^2}{2\sigma^2}\right)} \\
&= \ln\left(q \exp\left(\frac{t^2-(t-1)^2}{2\sigma^2}\right) + (1-q)\right) \\
&= \ln\left(q \exp\left(\frac{2t-1}{2\sigma^2}\right) + (1-q)\right) \\
&\leq \ln \max\left\{\exp\left(\frac{2t-1}{2\sigma^2}\right), 1\right\} \\
&= \max\left\{\frac{2t-1}{2\sigma^2}, 0\right\}
\end{aligned} \tag{100}$$

Similarly, we also have

$$\ln \frac{f_P(t)}{f_Q(t)} \geq \ln \min\left\{\exp\left(\frac{2t-1}{2\sigma^2}\right), 1\right\} = \min\left\{\frac{2t-1}{2\sigma^2}, 0\right\} \tag{101}$$

so

$$\begin{aligned}
\left|\ln \frac{f_P(t)}{f_Q(t)}\right| &\leq \max\left\{\max\left\{\frac{2t-1}{2\sigma^2}, 0\right\}, -\min\left\{\frac{2t-1}{2\sigma^2}, 0\right\}\right\} \\
&= \max\left\{\max\left\{\frac{2t-1}{2\sigma^2}, 0\right\}, \max\left\{-\frac{2t-1}{2\sigma^2}, 0\right\}\right\} \\
&\leq \left|\frac{2t-1}{2\sigma^2}\right|
\end{aligned} \tag{102}$$

Since $\mathbb{E}_{t\sim P}(t) = q$ and $\mathbb{E}_{t\sim P}(t^2) = \sigma^2 + q$,

$$\begin{aligned}
\mathrm{Var}(L^{(T)}) &\leq \mathbb{E}((L^{(T)})^2) \\
&= \mathbb{E}_{t\sim P}\left(\left(\ln \frac{f_P(t)}{f_Q(t)}\right)^2\right) \\
&\leq \mathbb{E}_{t\sim P}\left(\left|\frac{2t-1}{2\sigma^2}\right|^2\right) \\
&= \mathbb{E}_{t\sim P}\left(\left(\frac{2t-1}{2\sigma^2}\right)^2\right) \\
&= \frac{1}{4\sigma^4}\,\mathbb{E}_{t\sim P}\left(4t^2 - 4t + 1\right) \\
&= \frac{\sigma^2+q}{\sigma^4} - \frac{q}{\sigma^4} + \frac{1}{4\sigma^4} \\
&= \frac{1}{\sigma^2} + \frac{1}{4\sigma^4}.
\end{aligned} \tag{103}$$

□

Lemma D.14. *Let L_T be the PLRV of T iterations of the Poisson subsampled Gaussian mechanism, and let $K \in \mathbb{N}$. If $\sigma_T^2 = \Omega(T)$ is not true, for any $\alpha_i > 0$, $b_i > 0$ with $1 \leq i \leq K$, it is possible to find a T such that,*

$$\Pr_{s\sim L_T}(s \leq b_i) \leq \alpha_i \tag{104}$$

holds simultaneously for all i.

Proof. If $\sigma_T^2 = \Omega(T)$ is not true,

$$\liminf_{T\to\infty} \frac{\sigma_T^2}{T} = 0. \tag{105}$$

By Lemma D.12 and the composition theorem, $\mathbb{E}(L_T) \geq T\frac{q^2}{2\sigma_T^2}$. By Lemma D.13, $\mathrm{Var}(L_T) \leq \frac{T}{\sigma_T^2} + \frac{T}{4\sigma_T^4}$.

Let $k_i = \frac{1}{\sqrt{\alpha_i}}$. Choose T such that

$$\frac{Tq^2}{\sigma_T^2} - k_i\sqrt{\frac{T}{\sigma_T^2} + \frac{T}{4\sigma_T^4}} \geq b_i. \tag{106}$$

for all i. This is possible by (105) by choosing $\frac{T}{\sigma_T^2}$ to be large enough to satisfy all the inequalities. Now

$$\mathbb{E}(L_T) - k_i\sqrt{\mathrm{Var}(L_T)} \geq \frac{Tq^2}{\sigma_T^2} - k_i\sqrt{\frac{T}{\sigma_T^2} + \frac{T}{4\sigma_T^4}} \geq b_i \tag{107}$$

for all i, so

$$\Pr_{s\sim L_T}(s \leq b_i) \leq \Pr_{s\sim L_T}\left(|s - \mathbb{E}(L_T)| \geq k_i\sqrt{\mathrm{Var}(L_T)}\right) \leq \frac{1}{k_i^2} = \alpha_i \tag{108}$$

for all i by Chebyshev's inequality. □

This means that it is possible to make $\Pr_{s\sim L_T}(s \leq b)$ arbitrarily small for any b by choosing an appropriate T, and to satisfy a finite number of these constraints simultaneously with a single T.

Theorem 5.1. *Let σ_T be such that $\mathrm{AO}_S(\sigma_T, \Delta, q, T, \epsilon) \leq \delta$ for all T, with $\delta < 1$. Then $\sigma_T^2 = \Omega(T)$.*

Proof. By Lemma B.2, it suffices to consider $\Delta = 1$. To obtain a contradiction, assume that $\sigma_T^2 = \Omega(T)$ is not true. Let L_T be the PLRV for T iterations of the Poisson subsampled Gaussian mechanism.

From (18),

$$\begin{aligned}
\mathrm{AO}_S(\sigma_T, q, T, \epsilon) &= \mathbb{E}_{s\sim L_T}((1 - e^{\epsilon - s})_+) \\
&= \mathbb{E}_{s\sim L_T}(I(s > \epsilon)(1 - e^{\epsilon - s})) \\
&= \mathbb{E}_{s\sim L_T}(I(s > \epsilon)) - \mathbb{E}_{s\sim L}(I(s > \epsilon)e^{\epsilon - s}) \\
&= \Pr_{s\sim L_T}(s > \epsilon) - \mathbb{E}_{s\sim L}(I(s > \epsilon)e^{\epsilon - s})
\end{aligned} \tag{109}$$

By choosing $b_1 = \epsilon$ and $\alpha_1 = 1 - \frac{1}{2}(\delta + 1)$ in Lemma D.14, we get $\Pr_{s\sim L_T}(s > \epsilon) \geq \frac{1}{2}(\delta + 1)$.

To bound the remaining term,

$$\begin{aligned}
\mathbb{E}_{s\sim L_T}(I(s>\epsilon)e^{\epsilon-s}) &= e^{\epsilon}\,\mathbb{E}_{s\sim L_T}(I(s>\epsilon)e^{-s}) \\
&= e^{\epsilon}\sum_{i=0}^{\infty}\mathbb{E}_{s\sim L}(I(\epsilon+i<s\le\epsilon+i+1)e^{-s}) \\
&\le e^{\epsilon}\sum_{i=0}^{\infty}\mathbb{E}_{s\sim L}(I(\epsilon+i<s\le\epsilon+i+1)e^{-\epsilon-i}) \\
&\le \sum_{i=0}^{\infty}e^{-i}\,\mathbb{E}_{s\sim L}(I(\epsilon+i<s\le\epsilon+i+1)) \\
&\le \sum_{i=0}^{\infty}e^{-i}\Pr_{s\sim L}(\epsilon+i<s\le\epsilon+i+1) \\
&= \sum_{i=0}^{K}e^{-i}\Pr_{s\sim L}(\epsilon+i<s\le\epsilon+i+1)+\sum_{i=K+1}^{\infty}e^{-i}\Pr_{s\sim L}(\epsilon+i<s\le\epsilon+i+1) \\
&\le \sum_{i=0}^{K}e^{-i}\Pr_{s\sim L}(s\le\epsilon+i+1)+\sum_{i=K+1}^{\infty}e^{-i}.
\end{aligned} \tag{110}$$

The series $\sum_{i=0}^{\infty}e^{-i}$ converges, so it is possible to make $\sum_{i=K+1}^{\infty}e^{-i}$ arbitrarily small by choosing an appropriate K, which does not depend on T.

If we choose K such that $\sum_{i=K+1}^{\infty}e^{-i} < \frac{1}{4}(1-\delta)$ and then choose T such that $e^{-i}\Pr(s\le\epsilon+i+1)<\frac{1}{4K}(1-\delta)$ for all $1\le i\le K$, we have

$$\mathbb{E}_{s\sim L_T}(I(s>\epsilon)e^{\epsilon-s}) < K\cdot\frac{1}{4K}(1-\delta)+\frac{1}{4}(1-\delta)=\frac{1}{2}(1-\delta). \tag{111}$$

Lemma D.14 allows multiple inequalities for a single T, so we can find a T that satisfies all of the $K+1$ inequalities we have required it to satisfy. With this T,

$$\mathrm{AO}_S(\sigma_T,q,T,\epsilon)=\Pr_{s\sim L_T}(s>\epsilon)-\mathbb{E}_{s\sim L}(I(s>\epsilon)e^{\epsilon-s})>\frac{1}{2}(\delta+1)-\frac{1}{2}(1-\delta)=\delta \tag{112}$$

which is a contradiction. □

E Missing Proofs in Section 6

E.1 Solving $\sigma'(q)$ for $T=1$ case

Lemma 6.1. *For the smallest $\sigma(q)$ that provides (ϵ,δ)-DP for the Poisson subsampled Gaussian mechanism, we have*

$$\sigma'(q)=\frac{\sigma(q)}{q}\frac{1}{2a}\frac{1}{\mathrm{erf}'(a-b)}(\mathrm{erf}(a-b)-\mathrm{erf}(-a-b)). \tag{15}$$

Proof. Recall Equation (13): for a single iteration Poisson subsampled Gaussian mechanism, we have

$$\begin{aligned}
\delta(q) =& q\Pr\left(Z\ge\sigma(q)\log\left(\frac{h(q)}{q}\right)-\frac{1}{2\sigma(q)}\right) \\
&-h(q)\Pr\left(Z\ge\sigma(q)\log\left(\frac{h(q)}{q}\right)+\frac{1}{2\sigma(q)}\right),
\end{aligned} \tag{113}$$

where $h(q)=e^{\epsilon}-(1-q)$. Since $\sigma(q)$ is a function that returns a noise-level matching any (ϵ,δ)-DP requirement for subsampling rate q, $\delta'(q)=0$. Using Mathematica, we can solve the derivative of

the RHS in Equation (113) for $\sigma'(q)$ and we get

$$\sigma'(q) = \frac{\sqrt{\frac{\pi}{2}}\sigma(q)^2 e^{\frac{1}{2}\sigma(q)^2\log^2\left(\frac{q+e^\epsilon-1}{q}\right)+\frac{1}{8\sigma(q)^2}}\left(\text{erf}\left(\frac{1-2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)}{2\sqrt{2}\sigma(q)}\right)-\text{erf}\left(-\frac{2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)+1}{2\sqrt{2}\sigma(q)}\right)\right)}{q\sqrt{\frac{q+e^\epsilon-1}{q}}}. \tag{114}$$

Note that

$$\exp\left(\frac{1}{2}\sigma(q)^2\log^2\left(\frac{q+e^\epsilon-1}{q}\right)+\frac{1}{8\sigma(q)^2}\right) \tag{115}$$

$$= \exp\left(\frac{1}{2}\left(\sigma(q)^2\log^2\left(\frac{q+e^\epsilon-1}{q}\right)+\frac{1}{4\sigma(q)^2}\right)\right) \tag{116}$$

$$= \exp\left(\frac{1}{2}\left(\left(\sigma(q)\log\left(\frac{q+e^\epsilon-1}{q}\right)-\frac{1}{2\sigma(q)}\right)^2+\log\left(\frac{q+e^\epsilon-1}{q}\right)\right)\right) \tag{117}$$

$$= \exp\left(\left(\frac{1}{\sqrt{2}}\sigma(q)\log\left(\frac{q+e^\epsilon-1}{q}\right)-\frac{1}{2\sqrt{2}\sigma(q)}\right)^2\right)\sqrt{\frac{q+e^\epsilon-1}{q}} \tag{118}$$

$$= \exp\left(\left(\frac{1-2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)}{2\sqrt{2}\sigma(q)}\right)^2\right)\sqrt{\frac{q+e^\epsilon-1}{q}} \tag{119}$$

and therefore the derivative becomes

$$\sigma'(q) = \frac{\sqrt{\frac{\pi}{2}}\sigma(q)^2 e^{\frac{1}{2}\left(\sigma(q)\log\left(\frac{q+e^\epsilon-1}{q}\right)-\frac{1}{2\sigma(q)}\right)^2}\left(\text{erf}\left(\frac{1-2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)}{2\sqrt{2}\sigma(q)}\right)-\text{erf}\left(-\frac{2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)+1}{2\sqrt{2}\sigma(q)}\right)\right)}{q} \tag{120}$$

$$= \frac{\sqrt{\frac{\pi}{2}}\sigma(q)^2 e^{\left(\frac{1-2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)}{2\sqrt{2}\sigma(q)}\right)^2}\left(\text{erf}\left(\frac{1-2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)}{2\sqrt{2}\sigma(q)}\right)-\text{erf}\left(-\frac{2\sigma(q)^2\log\left(\frac{q+e^\epsilon-1}{q}\right)+1}{2\sqrt{2}\sigma(q)}\right)\right)}{q} \tag{121}$$

Lets denote

$$a := \frac{1}{2\sqrt{2}\sigma(q)} \tag{122}$$

$$b := \frac{\sigma(q)}{\sqrt{2}}\log\left(\frac{e^\epsilon-(1-q)}{q}\right). \tag{123}$$

Using this notation we can write

$$q\sigma'(q) = \sqrt{\frac{\pi}{2}}\sigma(q)^2\exp\big((a-b)^2\big)(\text{erf}(a-b)-\text{erf}(-a-b)) \tag{124}$$

$$= \sqrt{\frac{\pi}{2}}\sigma(q)^2\exp\big((a-b)^2\big)(\text{erf}(a+b)-\text{erf}(b-a)). \tag{125}$$

Note that we have

$$\text{erf}'(a-b) = \frac{2}{\sqrt{\pi}}\exp\big(-(a-b)^2\big) \tag{126}$$

$$\Leftrightarrow \exp\big((a-b)^2\big) = \frac{2}{\sqrt{\pi}\text{erf}'(a-b)} \tag{127}$$

Hence

$$q\sigma'(q) = \sigma(q)^2 \frac{\sqrt{2}}{\text{erf}'(a-b)}(\text{erf}(a+b) - \text{erf}(b-a)) \tag{128}$$

$$= \sigma(q)^2 \frac{\sqrt{2}}{\text{erf}'(a-b)}(\text{erf}(a+b) + \text{erf}(a-b)) \tag{129}$$

$$= \sigma(q) \frac{1}{2a} \frac{1}{\text{erf}'(a-b)}(\text{erf}(a+b) + \text{erf}(a-b)) \tag{130}$$

$$= \sigma(q) \frac{1}{2a} \frac{1}{\text{erf}'(a-b)}(\text{erf}(a-b) - \text{erf}(-a-b)). \tag{131}$$

□

E.2 Proof for Theorem 6.2

Theorem 6.2. *If $a < b$ for a and b defined in Equation* (14)*, then $\frac{\mathrm{d}}{\mathrm{d}q}\frac{\sigma(q)}{q} < 0$.*

Proof. The $\text{erf}(x)$ is a convex function for $x \in \mathbb{R}_{<0}$. Since $a, b \geq 0$ we have $-a - b \leq 0$ and if $a - b < 0$ we get from the convexity that

$$\text{erf}(a-b) - \text{erf}(-a-b) < 2a\,\text{erf}'(a-b). \tag{132}$$

Substituting this upper bound into Equation (15) gives

$$\sigma'(q) < \frac{\sigma(q)}{q} \Leftrightarrow \frac{q\sigma'(q) - \sigma(q)}{q^2} = \frac{\mathrm{d}}{\mathrm{d}q}\frac{\sigma(q)}{q} < 0. \tag{133}$$

□