We thank all three reviewers for their comments and suggestions.

Reviewer 1:

Discrete Laplace comparison: We consider the comparison with prior related work to be important, as it highlights the value of our discrete Gaussian over the discrete Laplace. This comparison is crucial to help the reader/practicioner(/reviewer) weigh the pros and cons of each method. Unfortunately, the advantages of the Gaussian over the Laplace are not as widely known as they should be, so we consider it worthwhile to reiterate this. Nonetheless, we will consider reallocating space in the final version.

Runtime analysis: The running time of our algorithm can be shown to be subexponential (measured by the number of arithmetic operations, where the bit complexity of each operation is determined by the bit complexity of the input parameters). We will add a statement on this to the final version of the main document. (Currently there is a brief discussion in the supplement, which we will flesh out into a full proof.)

Sampling Algorithm: We like our algorithm and consider it to be a contribution, as it is simpler than the prior work by Karney. However, we chose not to emphasize it because there is prior work; instead we emphasized our other contributions which are more clearly novel. We also did not want to dedicate space to a comparison with Karney's algorithm.

Reviewer 2:

Technical content: The technical contribution of the paper is to solve the problem of practical implementation of differentially private noise addition. We combine a number of techniques to solve the problem. We agree that some of the techniques that we apply, such as the Poisson summation formula, are known within certain communities. However, to the best of our knowledge, these techniques are not well known to the NeurIPS and Privacy communities, and have not been applied before in this context. Furthermore, we believe that all of our lemmata and theorems are novel and fundamental statements. Additionally, the focus of the cryptography community is specifically on the high-dimensional discrete Gaussian, which is crucially believed to be *hard* to sample as the dimension grows, a desirable feature for cryptographic applications; while we rely on the efficient sampling for the univariate case. In that sense, the viewpoint and goals are fundamentally different.

Reviewer 3:

Parameterization of CDP: The parameterizations ρ-CDP and $\frac{\varepsilon^2}{2}$-CDP are interchangeable. We prefer the latter as it puts CDP on the same familiar "scale" as pure ε-DP and approximate (ε, δ)-DP (indeed, ε-DP implies $\frac{\varepsilon^2}{2}$-CDP which in turn implies $(O(\varepsilon\sqrt{\log(1/\delta)}), \delta)$-DP) – i.e., people are more accustomed to ε as a privacy parameter than ρ. We revert to ρ when we are comparing CDP with approx DP as otherwise the parameter ε would be overloaded. We will make this correspondence of $\rho = \frac{\varepsilon^2}{2}$ more explicit in the final version.

Numerical comparison of RDP to DP conversion with previous results: We did not include the comparison because the numerical instability issues created weird artifacts in our plots. The code for the plots is included in the supplementary material and this comparison can be added to the plot by simply changing `numeric=False` to `numeric=True` at the end of the file and re-running the code. The resulting plots are included below.

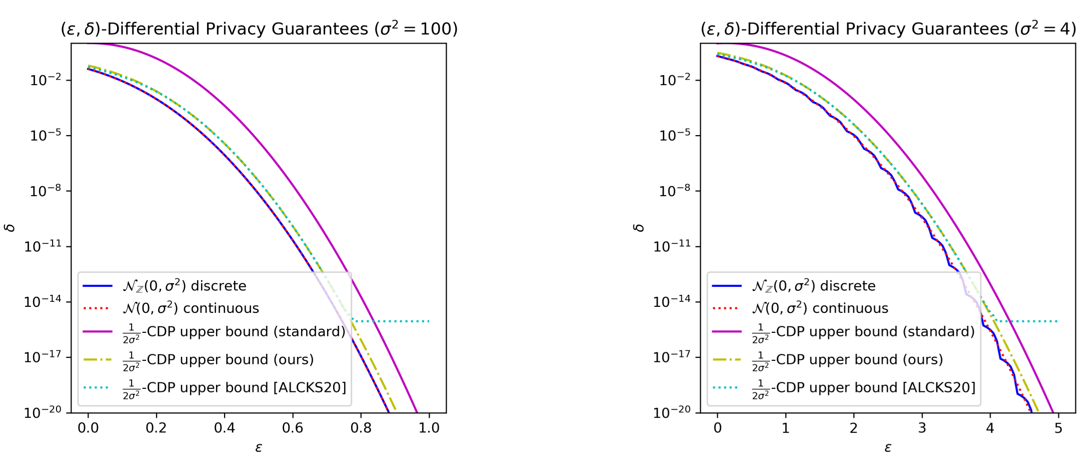

Figure 1: Numerical comparison of RDP to DP conversion with previous results