Journal of Privacy and Confidentiality
Vol. 14 (2) 2024

Submitted Jan 2023
Published June 2024

ON THE CONNECTION BETWEEN THE ABS PERTURBATION METHODOLOGY AND DIFFERENTIAL PRIVACY

PARASTOO SADEGHI AND CHIEN-HUNG CHIEN

School of Engineering & Technology, UNSW, Canberra
e-mail address: p.sadeghi@unsw.edu.au

Methodology Division, Australian Bureau of Statistics
e-mail address: joseph.chien@abs.gov.au

Abstract. This paper explores analytical connections between the perturbation methodology of the Australian Bureau of Statistics (ABS) and the differential privacy (DP) framework. We consider a single static counting query function and find the analytical form of the perturbation distribution with symmetric support for the ABS perturbation methodology. We then analytically measure the DP parameters, namely the (ε, δ) pair, for the ABS perturbation methodology under this setting. The results and insights obtained about the behaviour of (ε, δ) with respect to the perturbation support and variance are used to judiciously select the variance of the perturbation distribution to give a good δ in the DP framework for a given desired ε and perturbation support. Finally, we propose a simple sampling scheme to implement the perturbation probability matrix in the ABS Cellkey method. The post sampling (ε, δ) pair is numerically analysed as a function of the Cellkey size. It is shown that the best results are obtained for a larger Cellkey size, because the (ε, δ) pair post-sampling measures remain almost identical when we compare sampling and theoretical results.

INTRODUCTION

The Australian Bureau of Statistics (ABS) is committed to improving access to ABS statistics, while continuing to ensure privacy and confidentiality are maintained [ABS22]. The emergence of differential privacy (DP) methods provides opportunities to better quantify the trade-off between statistical utility and confidentiality protection in statistical outputs. As a result, the ABS is continuing to explore the opportunities offered by DP. This research builds on [FW05, ML11] and [BC19], and seeks to enhance the perturbation methodology in the ABS TableBuilder through the lens of DP. ABS perturbation methodology has two components—an entropy maximisation method for generating the perturbation probability transition matrix (or the perturbation table) and a cell key method to ensure consistent protections for statistical outputs [FW05, ML11]. At a high level, this work improves both components by first proposing an approach to incorporate the DP framework while

Notes Views expressed in this paper are those of the authors and do not necessarily represent those of the ABS. Where quoted or used, they should be attributed clearly to the authors. The work of P. Sadeghi was supported by the Australian Research Council Future Fellowship FT19100429.

JPrivacy Confidentiality

www.journalprivacyconfidentiality.org
DOI:10.29012/jpc.859

creating the perturbation table and then developing a sampling scheme to make full use of the perturbation table using a memory-efficient lookup table. Overall, this work offers tools and insights for analytical quantification of DP measures for the ABS perturbation methodology and improves its implementation efficiency. To the best of our knowledge, this is the first attempt to quantify the connection between ABS perturbation methodology design parameters and DP metrics in an analytical and provable format. In summary, the design-focused approach in this paper makes it different from [BC19], which empirically quantified the DP measures of the existing ABS perturbation methodology. We do not propose to change the entropy maximisation framework of [FW05, ML11]. Instead, we make the choice of its parameters relevant to DP through the analytical framework that we introduce for the design of its probability transition matrix.

More specifically, our contributions include: (1) introducing a method to analytically quantify the ϵ and δ DP parameters in the ABS perturbation methodology for a single counting query with symmetric perturbation support; (2) developing an approach to incorporate the ϵ DP parameter and the symmetric support of the distribution into the entropy maximisation process; (3) showing the importance of carefully choosing the variance parameter in the method proposed by [FW05, ML11] with respect to the DP parameters; and (4) proposing a sampling scheme to ensure the proposed method can be efficiently integrated with the cell key approach to improve ABS perturbation methodology and quantifying the (ϵ,δ)-DP parameters post sampling.

While we consider a specific case with a single counting query and a symmetric perturbation support, the methodology and insights have the potential to be extended to more advanced and complex cases. Addressing the following issues remain as ongoing challenges.

(1) How can national statistical organisations (NSO) use the DP framework to design disclosure protection mechanisms to better balance confidentiality and utility?
(2) How can we extend the proposed methodology to consider asymmetric perturbation support?
(3) How can we characterise composition of overall DP parameters if each count query has its own perturbation distribution?
(4) How can we study and apply DP when there is a significantly large query space?
(5) How can we study and apply scope-based perturbation, e.g., [TBE13] from the lens of DP (to protect against the range of count query being leaked from the perturbed output)?

The paper is structured as follows. Section 1 provides the key notation and describes the entropy maximisation proposed by [FW05, ML11]. Section 2 discusses the proposed analytical entropy maximisation approach to incorporate (ϵ,δ)-DP parameters for noise distributions with symmetric supports. We propose an approach to quantise and sample the probability mass function (pmf) with a simple lookup table in Section 3. We show the importance of increasing the size of the row index look up in Section 4. Finally, we provide a conclusion and propose future research directions in Section 5.

1. SYSTEM MODEL AND PRELIMINARIES

This section provides the notation conventions used throughout the paper. The set $\{a, \cdots, b\}$ for some $a, b \in \mathbb{Z}$, $a \leq b$ is compactly represented as $[a, b]$.

We consider a single counting query function q from a dataset $x \in \mathcal{X}$. Assume the true count is $q(x) = n$. In order to enhance the privacy of individuals in the dataset, a discrete-valued independent random variable Z with alphabet $\mathcal{Z}$ and probability mass function (pmf) p_Z is added to the true count to give the random query response

$$M(x) = q(x) + Z. \tag{1.1}$$

For brevity, we may simply refer to Z as noise. The parameters of the noise pmf are assumed to be independent of the dataset x. The probability mass of noise at $z \in \mathcal{Z}$ is denoted by $p_Z(Z = z)$. We may use the short-hand notation $p(z)$ where the context is clear.

References [FW05, ML11] show that given the above model and assumptions, the ABS TableBuilder aims to maximise statistical confusion induced by noise, measured by the Shannon entropy. It performs the following constrained optimisation to derive the noise parameters

$$\max_{p_Z} H(Z) = \max_{p_Z} \sum_{z \in \mathcal{Z}} p(z) \log \frac{1}{p(z)}, \tag{1.2}$$

$$\text{s.t.} \begin{cases} \mathbb{E}[Z] = 0, & \text{zero bias,} \\ \mathbb{E}[Z^2] \leq V, & \text{variance constraint,} \\ \sum_{z \in \mathcal{Z}} p(z) = 1, & \text{valid pmf,} \\ p(z) \geq 0, \quad \forall z \in \mathcal{Z}, & \text{valid pmf.} \end{cases} \tag{1.3}$$

We use the natural logarithm in this paper.

[FW05, ML11] solved the above optimisation problem numerically using standard solvers. However, it turns out that in the case where the noise support $\mathcal{Z}$ is symmetric, the solution to the problem becomes a discrete truncated Gaussian. In Section 2.1, we study this special case and analytically derive its DP measure in Section 2.2. In Sections 2.3 and 2.4, we take the derivations one step further and propose a design technique to keep the δ measure of the perturbation noise under control through judicious selection of its variance based on noise support and desired ε. We use the definition of differential privacy from [DMNS06, DR$^+$14]. Throughout this paper, we assume that $\varepsilon \in \mathbb{R}^+$.

Definition 1. (Approximate Differential Privacy) A randomised mechanism $M : \mathcal{X} \to \mathcal{Y}$ is said to satisfy (ε, δ)-differential privacy, or (ε, δ)-DP for short, if for all datasets $x, x' \in \mathcal{X}$ differing on a single element and all events $E \subset \mathcal{Y}$,

$$\mathbb{P}[M(x) \in E] \leq e^{\varepsilon} \mathbb{P}[M(x') \in E] + \delta.$$

If $\delta = 0$, we obtain pure or just ε-DP. If $0 < \delta \leq 1$, we obtain approximate (ε, δ)-DP.

2. MAIN RESULTS

2.1. **Analytical Distribution of the Symmetric TableBuilder Noise.** The noise range $\mathcal{Z}$ in the TableBuilder method is general and can be any subset of the integers $\mathbb{Z}$. However, to analytically characterise and optimise the differential privacy performance of the TableBuilder, we focus on the special symmetric case where $\mathcal{Z} = [-D, D]$ for some $D \in \mathbb{N}$. In order for the random query output $M(x)$ to remain nonnegative, it is required that the true count satisfy $q(x) = n \geq D$.[1] We specialise the TableBuilder optimisation problem in (1.2) as

$$\max_{p_Z} H(Z) = \max_{p_Z} \sum_{z=-D}^{D} p(z) \log \frac{1}{p(z)}, \tag{2.1}$$

$$\text{s.t.} \begin{cases} \sum_{z=-D}^{D} zp(z) = 0, & \text{zero bias,} \\ \sum_{z=-D}^{D} z^2 p(z) \leq V, & \text{variance constraint,} \\ \sum_{z=-D}^{D} p(z) = 1, & \text{valid pmf,} \\ p(z) \geq 0, \quad z \in [-D, D] & \text{valid pmf.} \end{cases} \tag{2.2}$$

Taking the derivative of the Lagrangian function for this problem and after some manipulations See Appendix A,, the optimal distribution $p(z)$ is of the form

$$p(z) = Ce^{-\gamma z^2}, \qquad z \in [-D, D], \tag{2.3}$$

where C is the normalisation constant satisfying

$$\sum_{z=-D}^{D} Ce^{-\gamma z^2} = 1,$$

so that

$$C = \frac{1}{2\sum_{z=1}^{D} e^{-\gamma z^2} + 1}. \tag{2.4}$$

Note that the optimal distribution is indeed a discrete zero-mean truncated Gaussian. The parameter γ is chosen to satisfy the variance constraint

$$\sum_{z=-D}^{D} z^2 Ce^{-\gamma z^2} = 2C \sum_{z=1}^{D} z^2 e^{-\gamma z^2} = V. \tag{2.5}$$

Combining (2.4) and (2.5),

$$\sum_{z=1}^{D} z^2 e^{-\gamma z^2} = V \left(2 \sum_{z=1}^{D} e^{-\gamma z^2} + 1 \right),$$

which implies that

$$\sum_{z=1}^{D} (2z^2 - 2V) e^{-\gamma z^2} - V = 0. \tag{2.6}$$

Let us denote $x := e^{-\gamma} > 0$. To find the pmf of noise, we need to solve the following polynomial equation of degree D^2 in x:

$$f(x) := \sum_{z=1}^{D} (2z^2 - 2V) x^{z^2} - V = 0. \tag{2.7}$$

[1] We remark that the ABS perturbation methodology has various ways to deal with small counts to protect confidentiality. This includes, but is not limited to zeroing out small counts or using asymmetric noise support. The details is beyond the scope of this paper and their study from the lens of DP is suggested as future work. Thus, the analysis in this paper applies to tables that are not sparse and all of whose entries exceed D. The DP consequences of this data dependence are not examined.

This equation has sparse nonzero coefficients at square degrees $D^2, (D-1)^2, \cdots, 9, 4, 1, 0$.

It is desirable for p_Z to have its highest probability at $Z = 0$, corresponding to the truthful count $M(x) = q(x) = n$ having the highest likelihood in the response. That is, we wish to have $0 < e^{-\gamma} < 1$ if and only if $\gamma > 0$. This means the polynomial $f(x)$ must have a root between 0 and 1. Note that $f(0) = -V < 0$, and also that

$$\begin{aligned} f(1) &= \sum_{z=1}^{D}(2z^2 - 2V) - V \\ &= \frac{D(D+1)(2D+1)}{3} - (2D+1)V \\ &= (2D+1)\left(\frac{D(D+1)}{3} - V\right). \end{aligned}$$

Therefore, if $f(1) > 0$, then $f(x)$ is guaranteed to have a root between 0 and 1. For $f(1) > 0$, we have the following proposition.

Proposition 1. For the TableBuilder pmf with symmetric support $\mathcal{Z} = [-D, D]$ to be a decreasing function of $|z|$, its variance V must satisfy

$$0 < V < \frac{D(D+1)}{3}. \tag{2.8}$$

This bound on variance V is consistent with the fact that among all probability mass functions over the support $[-D, D]$, the uniform distribution has the maximum entropy $H(Z) = \log(2D+1)$, zero bias, and variance $D(D+1)/3$. In the rest of Section 2, we will impose the constraint in Proposition 1.

2.2. **Differential Privacy Parameters of the ABS TableBuilder Method.**

We take a first-principles approach to computing the (ε, δ)-DP parameters of the TableBuilder mechanism. Our approach is similar in spirit to the one introduced in [BW18] for the continuous Gaussian mechanism and [CKS22] for the discrete Gaussian mechanism. However, the derivation of δ and optimisation of the TableBuilder noise pmf are very different and a main novelty of this paper.

Throughout this subsection, we assume the TableBuilder noise support $[-D, D]$ and noise variance V are given. Recall that the variance V determines γ in (2.3), which is found via solving (2.7). In summary, the TableBuilder noise pmf p_Z in (2.3) is parameterised by D and γ.

In this subsection, we characterise δ as a function of ε for given TableBuilder noise parameters D and γ. To make these dependencies clear, we denote it as $\delta_{\gamma,D}(\varepsilon)$. In the next subsections, we take the analysis one step further, where we study and optimise the effect of the TableBuilder noise parameters, γ and D, on δ.

First, the definition of (ε, δ)-DP (see Definition 1 on page 3) specifies that for every $E \subset \mathcal{Y}$, δ must satisfy

$$\delta \geq \mathbb{P}[M(x) \in E] - e^{\varepsilon}\mathbb{P}[M(x') \in E].$$

Therefore, the tightest lower bound on δ is

$$\delta \geq \sup_{E \subset \mathcal{Y}} \left\{\mathbb{P}[M(x) \in E] - e^{\varepsilon}\mathbb{P}[M(x') \in E]\right\}. \tag{2.9}$$

To characterise the worst-case event set E^* achieving the supremum, we need to define the privacy loss random variable. Using the results in [CKS22], detailed in Appendix B, we get

$$E^* := E^*_{\gamma,D}(\varepsilon) \triangleq \left\{ z : z \in [-D, D+1], \frac{p(z)}{p(z-1)} > e^{\varepsilon} \right\}, \tag{2.10}$$

leading to the computation of the lowest achievable δ as follows

$$\delta_{\gamma,D}(\varepsilon) := \mathbb{P}[M(x) \in E^*] - e^{\varepsilon}\mathbb{P}[M(x') \in E^*] = \sum_{z \in E^*} p(z) - e^{\varepsilon} p(z-1). \tag{2.11}$$

We first note that $z = -D$ always belongs to E^* regardless of ε and γ. This is because $p(Z = -D) \neq 0$ and $p(Z = -D-1) = 0$, resulting in the privacy loss ratio becoming infinite. Therefore, a simple lower bound on δ is the noise pmf value at $Z = -D$. That is, for all $\varepsilon > 0$,

$$\delta_{\gamma,D}(\varepsilon) \geq p_Z(Z = -D) = Ce^{-\gamma D^2}. \tag{2.12}$$

In order to fully characterise the set E^*, let us expand and simplify it as

$$\begin{aligned} E^*_{\gamma,D}(\varepsilon) &\triangleq \left\{ z : z \in [-D, D+1], \quad \frac{p(z)}{p(z-1)} = \frac{e^{-\gamma z^2}}{e^{-\gamma(z-1)^2}} > e^{\varepsilon} \right\} && (2.13)\\ &= \{ z : z \in [-D, D+1], \quad e^{-2\gamma z + \gamma} > e^{\varepsilon} \} && (2.14)\\ &= \{ z : z \in [-D, D+1], \quad -2\gamma z + \gamma > \varepsilon \} && (2.15)\\ &= \left\{ z : z \in [-D, D+1], \quad z < 0.5 - \frac{\varepsilon}{2\gamma} \right\}. && (2.16) \end{aligned}$$

Under the constraint detailed in Section 2.1 that $V < D(D+1)/3$, we will have that $\gamma > 0$ and hence, $z \geq 1$ cannot belong to E^*. Therefore, it suffices to determine whether each $z \in [-D+1 : 0]$ belongs to E^*. Let us define

$$F^* := F^*_{\gamma,D}(\varepsilon) \triangleq \left\{ z : z \in [-D+1 : 0], \quad z < 0.5 - \frac{\varepsilon}{2\gamma} \right\}. \tag{2.17}$$

Denote $z^* := \lfloor 0.5 - \varepsilon/2\gamma \rfloor$. We consider two cases:

(1) $0 < \varepsilon < \gamma$, which is the same as $0 < 0.5 - \varepsilon/2\gamma < 0.5$. Therefore, $z^* = 0$ and $F^* = \{-D+1, \cdots, 0\}$

(2) $\varepsilon > \gamma > 0$, which means that $0.5 - \varepsilon/2\gamma < 0$, and hence $z^* < 0$. Within this case, there are two sub-cases:
   (a) If $z^* \leq -D$, then $F^* = \emptyset$. This means $E^* = \{-D\}$.
   (b) If $-D+1 \leq z^* < 0$, then $F^* = [-D+1, z^*] \neq \emptyset$.

Therefore, the set E^* can be compactly written as

$$E^*_{\gamma,D}(\varepsilon) = \{-D\} \cup F^*_{\gamma,D}(\varepsilon) = [-D, \max\{-D, z^*\}]. \tag{2.18}$$

In summary, we analytically characterise $\delta_{\gamma,D}(\varepsilon)$ in the following proposition.

Proposition 2. Consider the mechanism in (1.1) for the single counting query q. The TableBuilder mechanism with noise pmf given in (2.3)-(2.5) and variance V satisfying

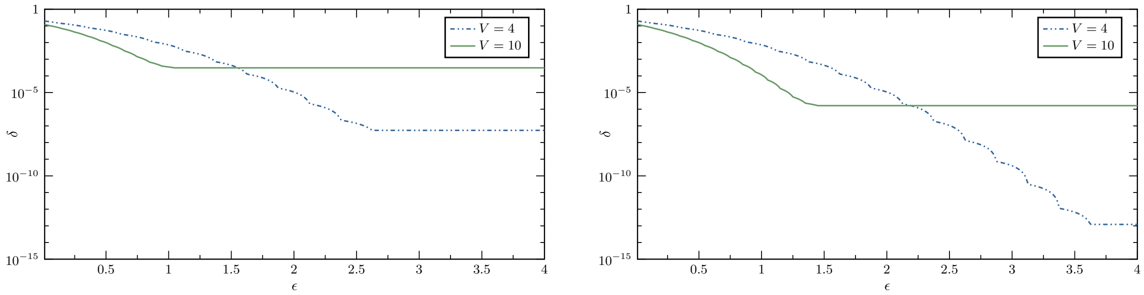

(A) Plots of $\delta_{\gamma,D}(\varepsilon)$ using (2.19) for $D = 11$ and two values of $\gamma = 0.125$ and $\gamma \approx 0.0498$ corresponding to two variances $V = 4$ and $V = 10$, respectively.

(B) Plots of $\delta_{\gamma,D}(\varepsilon)$ using (2.19) for $D = 15$ and the same variances $V = 4$ and $V = 10$.

FIGURE 1. Plots of $\delta_{\gamma,D}(\varepsilon)$ for various TableBuilder noise parameters D and γ.

Proposition 1 achieves (ε, δ)-DP such that

$$\delta = \delta_{\gamma,D}(\varepsilon) = \begin{cases} Ce^{-\gamma D^2}, & \lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor \le -D, \\ Ce^{-\gamma D^2} + C \sum_{z=-D+1}^{\lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor} (e^{-\gamma z^2} - e^{\varepsilon} e^{-\gamma (z-1)^2}), & -D < \lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor \le 0, \end{cases} \tag{2.19}$$

where γ is determined by V via solving the polynomial equation in (2.7).

In Appendix C, we verify that for $D = \infty$, the above derivation coincides with δ given in [CKS22] for the discrete Gaussian mechanism over the integers.

In Figure 1, we present evaluation of $\delta_{\gamma,D}(\varepsilon)$ according to (2.19) for four possible combinations of $D = 11$ and $D = 15$ with $\gamma = 0.125$ and $\gamma \approx 0.0498$ (corresponding to two variances $V = 4$ and $V = 10$, respectively). They are divided into Figure 1a and 1b for different values of D. There are a number of important observations that can be made from the two figures. Broadly speaking, when γ is fixed, increasing the noise support span D will decrease δ. However, the impact of γ on δ and its interactions with ε is complex. [FW05] and [ML11] rely on variance V and support D to design the noise and there are no specific relationships between V (or γ), δ and ε. Therefore, when we attempt introducing ε parameter to calculate δ using (2.19), we can observe increasing ε will hit a point where δ is not decreasing. This is because if the relation of ε with D and γ is such that $\lfloor 0.5 - \varepsilon/2\gamma \rfloor \le -D$, we will have a fixed $\delta = Ce^{-\gamma D^2}$ regardless of how much larger ε gets. Figure 1a shows that δ hits a plateau after reaching a certain point in ε. [FW05, ML11] preserve better confidentiality-utility trade off before hitting this plateau where δ is lower for a higher variance V. However, the plateau in δ prevents it from continuing this trend. Overall, a careful choice of parameters for the TableBuilder noise is needed to ensure a desired outcome. We will discuss this topic in greater detail in Subsections 2.3 and 2.4.

2.3. **Selection of TableBuilder Parameters under (ε, δ)-DP Framework.** In the previous subsection, we derived the (ε, δ)-DP parameters of the Tablebuilder mechanism. The derivation technique takes the TableBuilder γ and D as input parameters and determines what δ is achievable as a function of ε. We observed that for a fixed γ, there comes a threshold in ε beyond which increasing ε does not decrease δ. We attributed this plateauing phenomenon to the existence of the first case for $\delta_{\gamma,D}(\varepsilon)$ in (2.19) in Proposition 2. Even as we increase ε, we observed that $\delta_{\gamma,D}(\varepsilon)$ is bounded away from zero by $Ce^{-\gamma D^2}$. In this subsection, the core idea is to judiciously select γ (or variance) as ε increases to avoid a plateau in δ.

As we know from the first case in (2.19), making z^* smaller than $-D$ by increasing ε does not result in a reduction of δ. Therefore, we propose to choose γ such that $z^* = \lfloor 0.5 - \varepsilon/2\gamma \rfloor = -D$, *always*. This effectively means that in (2.17), $F^*_{\gamma,D}(\varepsilon) = \emptyset$ and the only element in E^* in (2.18) is $Z = -D$. Setting $z^* = \lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor = -D$ prevents it from unnecessarily becoming too small, thereby avoiding a plateau. That is, we propose to choose γ such that the first case in (2.19) always holds with equality $\lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor = -D$.[2] This will give $\delta = Ce^{-\gamma D^2}$.

Note that $\lfloor u \rfloor \le u < \lfloor u \rfloor + 1$. To have $z^* = \lfloor 0.5 - \varepsilon/2\gamma \rfloor = -D$, we need to ensure the following is satisfied

$$D \le 0.5 - \frac{\varepsilon}{2\gamma} < -D + 1,$$

i.e., that

$$\frac{\varepsilon}{2D+1} \le \gamma < \frac{\varepsilon}{2D-1}, \tag{2.20}$$

which is the proposed range for γ as a function of ε. We now find what range for the variance of the TableBuilder is required to ensure the desired γ. It turns out that we can find the corresponding range for V in analytical closed-form. Recall (2.7), which is polynomial in $x = e^{-\gamma}$, but is affine in V. We can solve (2.7) for V in terms of $x = e^{-\gamma}$:

$$V = \frac{\sum_{z=1}^{D} 2z^2 e^{-\gamma z^2}}{2\sum_{z=1}^{D} e^{-\gamma z^2} + 1}. \tag{2.21}$$

It can be verified that V in (2.21) is an increasing function of $x = e^{-\gamma}$ or a decreasing function of γ. Therefore, based on (**??**) and (2.21), the proposed range for V is

$$\frac{\sum_{z=1}^{D} 2z^2 e^{-\frac{\varepsilon}{2D-1} z^2}}{2\sum_{z=1}^{D} e^{-\frac{\varepsilon}{2D-1} z^2} + 1} < V \le \frac{\sum_{z=1}^{D} 2z^2 e^{-\frac{\varepsilon}{2D+1} z^2}}{2\sum_{z=1}^{D} e^{-\frac{\varepsilon}{2D+1} z^2} + 1}. \tag{2.22}$$

And from (2.4), the desired range for C (which is decreasing in $x = e^{-\gamma}$ or increasing in γ) is

$$\frac{1}{2\sum_{z=1}^{D} e^{-\frac{\varepsilon}{2D+1} z^2} + 1} \le C < \frac{1}{2\sum_{z=1}^{D} e^{-\frac{\varepsilon}{2D-1} z^2} + 1}. \tag{2.23}$$

Finally, we detail δ, which also has a range. It can be verified that

$$\delta = Ce^{-\gamma D^2} = \frac{e^{-\gamma D^2}}{2\sum_{z=1}^{D} e^{-\gamma z^2} + 1},$$

[2] Note that we are not claiming this choice for γ will minimise δ overall. This is because for simplicity of analysis, we are not considering both cases of (2.19) jointly to select the best γ for a given ε and D. Our proposed method is a heuristic technique, which focuses on optimising the first case in (2.19) and obtains an analytical achievable expression for δ in terms of ε and D. It is intuitive that focusing on the first case of (2.19) should be a good choice, as it does not suffer from additional terms for δ. See Figure 2 for a numerical corroboration.

is an increasing function of $x = e^{-\gamma}$ or a decreasing function of γ. The obtained range for δ is

$$\frac{e^{-\frac{\varepsilon}{2D-1}D^2}}{2\sum_{z=1}^{D} e^{-\frac{\varepsilon}{2D-1}z^2}+1} < \delta \le \frac{e^{-\frac{\varepsilon}{2D+1}D^2}}{2\sum_{z=1}^{D} e^{-\frac{\varepsilon}{2D+1}z^2}+1}. \tag{2.24}$$

Remark 1. We stress the importance of choosing an appropriate value for γ for a corresponding V (see (2.22)) and δ (see (2.24)). The choice of value for γ should be within the range $[\varepsilon/(2D+1), \varepsilon/(2D-1))$ according to (2.20).

It is important to note that the interval in (2.20) is open on the right hand side. It is not desirable for γ to be equal to $\varepsilon/(2D-1)$, because when this happens the derivations in Section 2.2 show that $E^*_{\gamma,D}$ will become strictly larger than $\{-D\}$ and this will contribute to a larger δ. This is shown in the second case in (2.19) in Proposition 2. To avoid this, γ has to be strictly smaller than $\varepsilon/(2D-1)$. We also remind the reader that an unnecessarily small γ is not desirable either. First, a value for γ that is strictly smaller than $\varepsilon/(2D+1)$ can result in a plateau in δ, as we saw before, because $E^*_{\gamma,D} = \{-D\}$ regardless of how small γ gets. Second, if $\gamma = \varepsilon/(2D+1)$, the upper bounds on the V and δ are reached according to (2.22) and (2.24) (Note that V and δ increase in (2.22) and (2.24) as γ decreases.)

We summarise the results of this subsection in the following proposition.

Proposition 3. Consider the mechanism in (1.1) for the single counting query q. For any given $\varepsilon > 0$, $D \in \mathbb{N}$ and γ in the range $[\varepsilon/(2D+1), \varepsilon/(2D-1))$, the TableBuilder mechanism with the following noise pmf

$$p_Z(Z=z) = \frac{e^{-\gamma z^2}}{2\sum_{z=1}^{D} e^{-\gamma z^2}+1}, \quad z \in [-D, D], \tag{2.25}$$

and noise variance

$$V = \frac{\sum_{z=1}^{D} 2z^2 e^{-\gamma z^2}}{2\sum_{z=1}^{D} e^{-\gamma z^2}+1}, \tag{2.26}$$

will achieve (ε, δ)-DP such that

$$\delta = \frac{e^{-\gamma D^2}}{2\sum_{z=1}^{D} e^{-\gamma z^2}+1}. \tag{2.27}$$

Figure 2a plots the analytical asymptotic expression for δ (as γ tends to $\varepsilon/(2D-1)$ from below) versus ε for two noise support parameters $D = 11$ and $D = 15$. The plateaus in Figure 1 have disappeared and as D increases, δ decreases. For comparison, we also plot the best possible δ, which is found numerically by varying γ from 0.0001 to 0.3 in linear steps of 0.0001, evaluating $\delta_{\gamma,D}(\varepsilon)$ using (2.19), and choosing the minimum δ possible. The gaps vary from being small to zero and corroborate our intuition that focusing on the first case of (2.19) and optimising it as described above is a good design strategy. Figure 2b plots the expression for variance V versus ε (again as γ tends to $\varepsilon/(2D-1)$ from below) for the corresponding two noise support parameters $D = 11$ and $D = 15$.

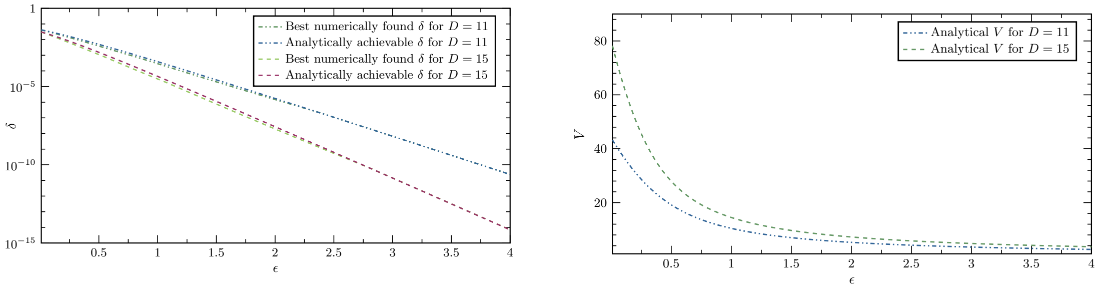

(A) δ versus ε: analytical δ versus numerically optimised values using (2.19).

(B) Analytical TableBuilder variance to achieve the corresponding analytical δ in Figure 2a.

FIGURE 2. Plots of δ and V versus ε for two TableBuilder noise support parameters $D = 11$ and $D = 15$.

2.4. **A TableBuilder Noise Design Guide.** In some applications, it may be desirable to achieve a specific (ε, δ)-DP measure for the ABS perturbation methodology. In this subsection, we use the results in Subsection 2.3 to prescribe a simple method for analytically choosing the parameters of the perturbation, that is, the support D and the variance V to achieve a desired (ε, δ)-DP.

(1) Start with the desired $\varepsilon > 0$ and $0 \le \delta \le 1$ as inputs.
(2) For the desired ε, linearly increase the support $D = 1, 2, \cdots$ and evaluate the expression in (2.27) with a value for γ in the range $[\varepsilon/(2D+1), \varepsilon/(2D-1))$, until the desired δ—or the first value smaller than δ—is reached. For simplicity, we can assume that $\gamma = \varepsilon/(2D-1) - 2\varepsilon/(10(4D^2-1))$ is chosen to ensure that $\varepsilon/(2D+1) < \gamma < \varepsilon/(2D-1)$.
(3) Select the last evaluated D, denoted by D^*, as the perturbation noise support parameter. Hence, $\mathcal{Z} = [-D^*, D^*]$. Denote the chosen value of γ as γ^*.
(4) The TableBuilder noise variance V to support the desired (ε, δ) is given by (2.26) using the γ^* found in the previous step.
(5) The TableBuilder noise pmf is given by

$$p_Z(Z = z) = \frac{e^{-\gamma^* z^2}}{2\sum_{z=1}^{D^*} e^{-\gamma^* z^2} + 1}, \quad z \in [-D^*, D^*]. \tag{2.28}$$

We now demonstrate how this routine works via an example.

Example 1. Let us assume the desired privacy target is $\varepsilon = 0.5$ and $\delta = 10^{-4}$. We find that the smallest D that satisfies (2.27) with $\gamma = \varepsilon/(2D-1) - 2\varepsilon/[10(4D^2-1)]$ is $D^* = 25$, resulting in $\delta \approx 9.91 \times 10^{-5}$. The corresponding perturbation variance is $V \approx 49.00$ and $\gamma = 0.5/(2 \times 25 - 1) - 1/(10(2500 - 1))$. So the overall perturbation pmf using (2.28) is

$$p_Z(Z = z) = \frac{e^{-(\frac{0.5}{2\times 25-1} - \frac{1}{10(2500-1)})z^2}}{2\sum_{z=1}^{25} e^{-(\frac{0.5}{2\times 25-1} - \frac{1}{10(2500-1)})z^2} + 1, \quad z \in [-25, 25]. \tag{2.29}$$

For example, if we evaluate the above pmf at $Z = 0$, $Z = \pm 1$, $Z = \pm 2$, $\cdots$, $Z = \pm 12$, $\cdots$, $Z = \pm 24$, and $Z = \pm 25$, we get

$$\begin{aligned}
p_Z(Z = 0) &\approx 0.056895481243871, \\
p_Z(Z = -1) = p_Z(Z = 1) &\approx 0.056320120792644, \\
p_Z(Z = -2) = p_Z(Z = 2) &\approx 0.054628714970934, \\
&\vdots \\
p_Z(Z = -12) = p_Z(Z = 12) &\approx 0.016632589297126, \\
&\vdots \\
p_Z(Z = -24) = p_Z(Z = 24) &\approx 0.000163117271714, \\
p_Z(Z = -25) = p_Z(Z = 25) &\approx 0.000099129808160.
\end{aligned} \tag{2.30}$$

Note that $P(Z = -z) = P(Z = z)$. The plot of the pmf is shown in Figure 3.

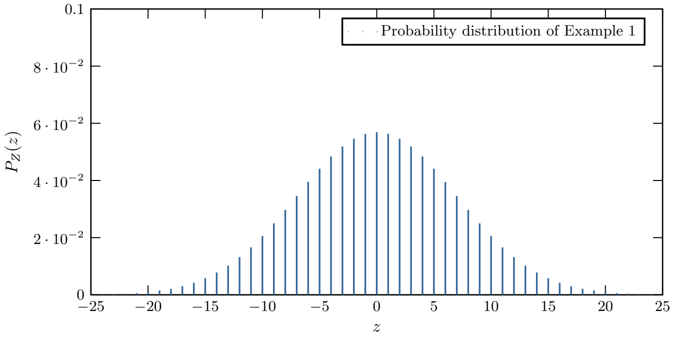

Figure 3. Perturbation distribution in Example 1.

3. CELL KEY METHODOLOGY

The ABS developed the cell key method to ensure that users cannot circumvent perturbation by making repeated requests for the same table. If the disclosure protection mechanism failed to deliver a consistent random perturbation, then a user could obtain different versions of the same table. Comparing the cell values across these different versions might reveal some information about the original table. This risk is particularly important to address in the context of the ABS TableBuilder, where there is no restriction to prevent a user requesting the same table many times [Lea09].

The cell key method assigns a pseudo-random number (also known as record key) to each record of the micro dataset. Record keys, $\texttt{Rkey}_i$, are positive integers less than 2^{32} generated from uniform distribution. In [TBE13], record keys of size 2^{32} were further processed (were combined byte-by-byte) to give cell keys of size 2^8. But this low cell key size was mainly implemented to reduce the complexity of lookup tables for sampling from a quantised perturbation noise distribution. However, this small cell key size is not strictly necessary. As we will see in the next section, larger cell key sizes are needed to maintain desired DP measures.

Therefore, in this section, we extend the cell key described in [TBE13] to allow cell keys to be a power of 2, which can be as high as 2^{32}. The cell key size is denoted by `KEYSIZE`. When a table is constructed, the record keys are summed over each cell, to give

$$\texttt{CellKey}^j = \sum_{i=1}^{N} \texttt{Rkey}_i^j (\text{ modulo } \texttt{bigN}), \tag{3.1}$$

where the cell key has four components $j = 1, \cdots, 4$, `bigN` is a large prime number that is very close to `KEYSIZE`, and we take the modulo to prevent integer overflows when we sum the pseudo-random numbers.[3] The final `CellKey` is determined as follows

$$\texttt{CellKey} = \texttt{CellKey}^1 \oplus \texttt{CellKey}^2 \oplus \texttt{CellKey}^3 \oplus \texttt{CellKey}^4, \tag{3.2}$$

where $\oplus$ is the bitwise XOR operator. The values $\texttt{CellKey}^1$, $\texttt{CellKey}^2$, $\texttt{CellKey}^3$, $\texttt{CellKey}^4$ are the four binary components derived from representing cell key as a binary number up to 32 bits.

The XOR operation is an important feature to remove the additive relationships between `CellKey` for the interior and marginal cells to ensure strong protection [Jon21]. We will use this `CellKey` and its size `KEYSIZE` in the next section for direct sampling from the perturbation noise.

To summarise, we assume that `CellKey` values are uniformly generated in the range $[0, \texttt{KEYSIZE} - 1]$, where `KEYSIZE` is a power of 2. Typical values 2^8, 2^{16}, or 2^{32} will be studied here, but other values are also possible.

4. SAMPLING AND ITS IMPACT ON (ε, δ)-DP

4.1. **Sampling.** For sampling, we first scale and quantise the cumulative mass function (cmf) of the proposed perturbation method in Proposition 3 according to the following procedure:

(1) For a given ϵ, D, $\gamma \in [\varepsilon/(2D+1), \varepsilon/(2D-1))$, the pmf of ABS perturbation method, p_Z, is given by (2.25). We first compute its cmf as

$$c_Z(Z = z) = \mathbb{P}[Z \leq z] = \sum_{z'=-\infty}^{z} p_Z(Z = z'), \quad z \in [-D, D], \tag{4.1}$$

where clearly $c_Z(Z = z) = 0$ for any $z < -D$ and $c_Z(Z = z) = 1$ for any $z \geq D$.

(2) Then, given the maximum cell key size `KEYSIZE`, we scale and quantise c_Z into c_Z^Q as follows:

$$c_Z^Q(Z = z) = \lceil c_Z(Z = z) \times \texttt{KEYSIZE} \rceil, \tag{4.2}$$

where $\lceil \cdot \rceil$ is the integer ceiling function. This will ensure that the minimum and maximum bounds 0, and 1 in c_Z will correspond to 0 and `KEYSIZE` in c_Z^Q, respectively.

(3) The values of c_Z^Q are stored in a lookup table of size $2D + 1$. Since D is usually small, this lookup table can be saved in a memory-efficient manner.

[3] The prime number `bigN` should be chosen as close to `KEYSIZE` as possible to ensure that the `CellKey` is also essentially uniformly distributed.

(4) When a `CellKey` is generated according to (3.2), we use the lookup table to get a sample from the distribution as follows. For a given value `CellKey` in the range $[0, \texttt{KEYSIZE} - 1]$, we output the sample S as follows: since

$$c_Z^Q(Z = z) \leq \texttt{CellKey} < c_Z^Q(Z = z + 1),$$

then

$$S = z + 1. \tag{4.3}$$

(5) If the cell key size, `KEYSIZE` is small, it may happen that two or more consecutive c_Z^Q may become identical. This means that some perturbation noise values z can never be achieved. If this happens, `KEYSIZE` must be increased or the parameters of the distribution must be adjusted to ensure the full support of the distribution can be achieved.

Example 2. We use the results from (2.30) in Example 1. Assume the cell key size is $\texttt{KEYSIZE} = 2^{32}$. We compute the scaled and quantised cmf C_Z^Q according to (4.2). For example, C_Z^Q values at $Z = -26, -25, -24, -23$, and $Z = 25$ are given as follows:

$$\begin{gathered} c_Z^Q(Z = -26) = 0, \\ c_Z^Q(Z = -25) = \lceil 0.000099129808160 \times 2^{32} \rceil = 425760, \\ c_Z^Q(Z = -24) = \lceil 0.0002622470798742925 \times 2^{32} \rceil = 1126343, \\ c_Z^Q(Z = -23) = \lceil 0.0005252540370388639 \times 2^{32} \rceil = 2255949, \\ \vdots \\ c_Z^Q(Z = 25) = \lceil 1 \times 2^{32} \rceil = 2^{32}. \end{gathered} \tag{4.4}$$

The values of c_Z^Q will be stored in a lookup table of size $2D + 1 = 51$. Now imagine that $\texttt{CellKey} = 2552$ is given according to (3.2). Since $c_Z^Q(Z = -26) \leq 2552 < c_Z^Q(Z = -25) = 425760$, we output $S = -25$ as the ABS perturbation noise. As another example, assume $\texttt{CellKey} = 1200124$ is given. Since $c_Z^Q(Z = -24) \leq 1200124 < c_Z^Q(Z = -23) = 2255949$, we should output $S = -23$ as the ABS perturbation noise, and so on.

Now assume that $\texttt{KEYSIZE} = 2^8$ is given instead. We can see that $c_Z^Q(Z = -25) = c_Z^Q(Z = -24) = c_Z^Q(Z = -23) = 1$. This means not all values in the support $[-D, D]$ can be realised in practice. Hence, we conclude that $\texttt{KEYSIZE} = 2^8$ is not a sufficient cell key size for this perturbation distribution.

4.2. **Evaluating Post-Sampling Utility and Privacy Measures.** It now remains to verify the properties of the scaled and quantised distribution in terms of bias, variance and $(\varepsilon^Q, \delta^Q)$-DP, where the superscript Q signifies values post sampling. To this end, we follow the procedures below.

(1) We convert the scaled and quantised cmf c_Z^Q in (4.2) into the scaled and quantised pmf p_Z^Q as follows:

$$p_Z^Q(Z = z) = \frac{c_Z^Q(Z = z) - c_Z^Q(Z = z - 1)}{\texttt{KEYSIZE}}, \quad z \in [-D : D]. \tag{4.5}$$

Note that we assume `KEYSIZE` is chosen sufficiently large to ensure that p_Z^Q has full support over $[-D, D]$. In steps below, we use the shorthand $p^Q(z) := p_Z^Q(Z = z)$.

(2) The resulting bias and variance of p_Z^Q are computed as

$$B^Q := \sum_{z=-D}^{D} z p^Q(z), \tag{4.6}$$

$$V^Q := \sum_{z=-D}^{D} z^2 p^Q(z) - (B^Q)^2. \tag{4.7}$$

These metrics clearly depend on the cell key size, `KEYSIZE`. Intuitively, the larger the `KEYSIZE`, the finer the quantisation will be and the closer the bias and variance of p_Z^Q should be to its original version obtained from p_Z.

To understand the effective $(\varepsilon^Q, \delta^Q)$-DP metric as a result of scaling and quantisation in p_Z^Q, we propose the following method.

(1) We first define and compute ε_{-1}^Q as follows:

$$\varepsilon_{-1}^Q := \arg\min \left\{ \{\varepsilon : \frac{p^Q(z)}{p^Q(z-1)} < e^{\varepsilon}, z \in [-D+1, D] \right\}. \tag{4.8}$$

The quantity ε_{-1}^Q is the minimum required value to ensure the ratio $p^Q(z)/p^Q(z-1)$ is bounded for support values $z \in [-D+1 : D]$. Similarly, we define and compute ε_{+1}^Q as follows:

$$\varepsilon_{+1}^Q := \arg\min \left\{ \varepsilon : \frac{p^Q(z)}{p^Q(z+1)} < e^{\varepsilon}, z \in [-D, D-1] \right\}. \tag{4.9}$$

The quantity ε_{+1}^Q is the minimum required value to ensure the ratio $p^Q(z)/p^Q(z+1)$ is bounded for support values $z \in [-D : D-1]$. The effective ε^Q is the maximum of the two quantities:

$$\varepsilon^Q = \max \left\{ \varepsilon_{-1}^Q, \varepsilon_{+1}^Q \right\}. \tag{4.10}$$

(2) Once the effective ε^Q is obtained as above, the effective δ^Q will be the maximum of the pmf p_Z^Q at the two extreme support values and is given by

$$\delta^Q = \max \left\{ p^Q(-D), p^Q(D) \right\}. \tag{4.11}$$

Again, the cell key size, `KEYSIZE` will play a main role on the resulting $(\varepsilon^Q, \delta^Q)$ metric. The larger the `KEYSIZE`, the closer $(\varepsilon^Q, \delta^Q)$ can get to the original (ε, δ) metrics for the continuous case. Also, 1/`KEYSIZE` will pose a lower bound on how small δ^Q can get, since this is the smallest value that $p_Z^Q(Z = -D)$ or $p_Z^Q(Z = D)$ can have.

Example 3. Continuing on Example 1 and Example 2, we can convert the scaled and quantised cmf C_Z^Q, according to (4.5), back to quantised pmf p_Z^Q. For example, at $Z = -25, -24, -23$, and $Z = 25$,

$$\begin{aligned} p_Z^Q(Z=-25) &= \tfrac{425760}{2^{32}} \approx 0.000099129974842, \\ p_Z^Q(Z=-24) &= \tfrac{1126343-425760}{2^{32}} \approx 0.000163117190823, \\ p_Z^Q(Z=-23) &= \tfrac{2255949-1126343}{2^{32}} \approx 0.0002630068920552731, \\ &\vdots \\ p_Z^Q(Z=25) &= \tfrac{2^{32}-4294541537}{2^{32}} = 0.00009912974201142788. \end{aligned} \tag{4.12}$$

Note that the quantised pmf has lost its complete symmetry, compared to the original pmf in Example 2. Its bias can be calculated from (4.6) to be $B^Q = -5.820766091346741 \times 10^{-9}$. Its variance can be calculated from (4.7) to be $V^Q = 49.002167175291106$, which are very close to the original zero-bias and design variance, respectively.

Now, we compute ε^Q according to (4.8)-(4.10), which gives $\varepsilon^Q \approx 0.498039387067656$. Interestingly, this is slightly smaller than the design target $\varepsilon = 0.5$. This is not unusual, since the quantisation is a nonlinear operation and ε^Q can be lower or higher than ε. We will investigate this further in the upcoming experiments. Finally, $\delta^Q = p_Z^Q(-D) \approx 9.9129974842 \times 10^{-5}$ is computed according to (4.11), and is only slightly larger than the original δ in Example 1.

4.3. **Experiments.** To study the effect of KEYSIZE on the perturbation bias, variance and DP measures more systematically, we consider the following scenario. We set $D = 10$, vary $\varepsilon \in [0.1, 2.5]$ in 0.1 steps, let $\gamma = \varepsilon/(2D-1) - 2\varepsilon/[10(4D^2-1)]$, and follow the proposed quantisation procedure we described in the previous subsections.

First, we find that bias $B^Q \approx -2.3 \times 10^{-9}$ is lowest when $\texttt{KEYSIZE} = 2^{32}$. This deteriorates to $B^Q \approx -1.5 \times 10^{-4}$ when $\texttt{KEYSIZE} = 2^{16}$ and to $B^Q \approx -0.04$ when $\texttt{KEYSIZE} = 2^8$. This confirms that KEYSIZE has a clear effect on the post-sampled perturbation measures. Furthermore, for $\texttt{KEYSIZE} = 2^8$ and $\texttt{KEYSIZE} = 2^{16}$ not all values of ε result in distributions with full support.

Next, we define the normalised error in variance after quantisation as

$$\frac{V^Q - V}{V}.$$

This normalised variance error is in the order of 10^{-9} and 10^{-4} for $\texttt{KEYSIZE} = 2^{32}$ and $\texttt{KEYSIZE} = 2^{16}$, respectively. However, when $\texttt{KEYSIZE} = 2^8$ the normalised variance error can be as high as ≈ 0.007.

Figure 4 shows the relation between ε at the time of design and the resulting ε^Q post-sampling for three different values of KEYSIZE. Having $\varepsilon \geq \varepsilon^Q$ is desirable and $\varepsilon^Q > \varepsilon$ is not desirable. We see that when $\texttt{KEYSIZE} = 2^8$ or $\texttt{KEYSIZE} = 2^{16}$, $\varepsilon^Q > \varepsilon$. When $\texttt{KEYSIZE} = 2^{32}$, ε^Q is either very close to ε or slightly lower. The nonlinear/jittery behaviour is not unusual and is due to the nonlinear sampling scheme, which involves the integer ceiling function. The other main problem with both $\texttt{KEYSIZE} = 2^8$ and $\texttt{KEYSIZE} = 2^{16}$ values is that the quantised perturbation p_Z^Q cannot provide full support due to the nonlinear quantisation and insufficiently large KEYSIZE. This is not acceptable since the designed support of $[-10, 10]$ cannot be maintained, which is the original design criterion. For $\texttt{KEYSIZE} = 2^8$, this happens after $\varepsilon = 0.6$ and for $\texttt{KEYSIZE} = 2^{16}$, this happens after $\varepsilon = 1.7$. Whereas, when $\texttt{KEYSIZE} = 2^{32}$, we see that $\varepsilon^Q \approx \varepsilon$ as desired, and the full support is maintained for all ε values under consideration.

Figure 5 shows the relation between δ at the time of design and the resulting δ^Q post sampling for three different values of KEYSIZE. $\delta \geq \delta^Q$ is desirable and $\delta^Q > \delta$ is not desirable. We see that when $\texttt{KEYSIZE} = 2^8$, $\delta^Q > \delta$. When $\texttt{KEYSIZE} = 2^{16}$ or $\texttt{KEYSIZE} = 2^{32}$, δ^Q is almost identical to the original δ. However, as mentioned before, the main problem in using $\texttt{KEYSIZE} = 2^8$ or $\texttt{KEYSIZE} = 2^{16}$ is that the full support of perturbation noise maintained cannot be maintained for all ε values under consideration.

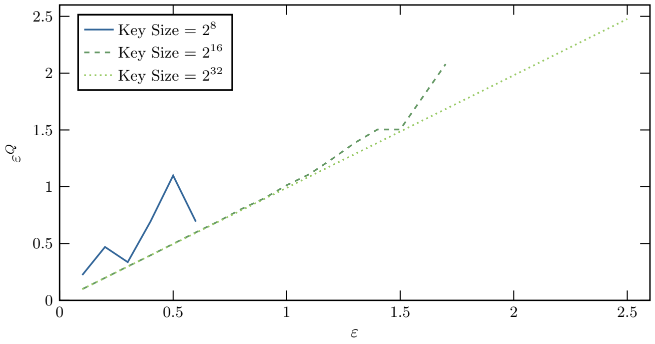

FIGURE 4. ε^Q versus ε for three values of cell key size. Overall, only `KEYSIZE` $= 2^{32}$ can closely follow the original ε across its entire range.

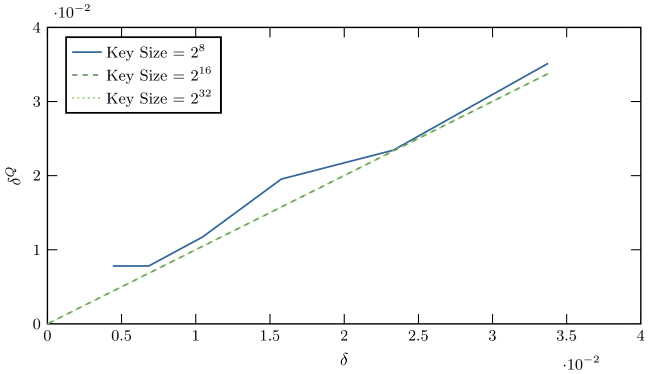

FIGURE 5. δ^Q versus δ for three values of cell key size. Overall, only `KEYSIZE` $= 2^{32}$ can closely follow the original δ across the entire range of ε.

5. CONCLUSION

The DP framework provides an opportunity to better quantify the confidentiality protection and data utility of the ABS perturbation methodology. We have proposed an alternative entropy maximisation approach which incorporates (ε, δ)-DP parameters for symmetric support.

We have proposed an approach to expand the cell key row index size. We have shown the importance of having a larger cell key size to achieve the desired (ε, δ)-DP parameters in our quantised sampling approach.

There are several potential directions for future research. This includes (1) extending the method to consider asymmetrical perturbation distributions; (2) developing a framework to consider ε and δ parameters for dynamic table environments; and (3) evaluating the performance against different types of perturbation distributions. Furthermore, the perturbation methodology that we studied in this paper is data-dependent in the sense that the noise distribution depends on the range of counts $[-D, D]$. It will be interesting to consider a combination of scope perturbation, such as those in [TBE13], and data perturbation methodologies from the lens of DP. Finally, it will be useful to study and design the

ABS perturbation methodology using generalised notions of DP such as Rényi DP [Mir17], concentrated DP [BS16], and probabilistically bounded DP [KAA$^+$22].

References

[ABS22] ABS. 1005.0—ABS Corporate Plan, 2021-22. https://www.abs.gov.au/ausstats/abs@.nsf/Lookup/by%20Subject/1005.0~2021-22~Main%20Features~Objectives~6#Objective3, 2022. Accessed: 2022-05-01. URL: https://www.abs.gov.au/ausstats/abs@.nsf/Lookup/by%20Subject/1005.0%7E2021-22%7EMain%20Features%7EObjectives%7E6#Objective3.

[BC19] James Bailie and Chien-Hung Chien. ABS perturbation methodology through the lens of differential privacy. https://unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2019/mtg1/SDC2019_S2_ABS_Bailie_D.pdf, 2019. URL: https://unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2019/mtg1/SDC2019_S2_ABS_Bailie_D.pdf.

[BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. https://arxiv.org/pdf/1605.02065, 2016. URL: https://arxiv.org/abs/1605.02065.

[BW18] Borja Balle and Yu-Xiang Wang. Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. https://proceedings.mlr.press/v80/balle18a/balle18a.pdf, 10-15 Jul 2018. URL: https://proceedings.mlr.press/v80/balle18a/balle18a.pdf.

[CKS22] Clement Canonne, Gautam Kamath, and Thomas Steinke. The discrete Gaussian for differential privacy. *Journal of Privacy and Confidentiality*, 12(1), Jul. 2022. https://doi.org/10.29012/jpc.784.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. https://journalprivacyconfidentiality.org/index.php/jpc/article/view/405, 2006. https://doi.org/https://doi.org/10.1007/11681878_14.

[DR$^+$14] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. https://www.nowpublishers.com/article/DownloadSummary/TCS-042, 2014. https://doi.org/http://dx.doi.org/10.1561/0400000042.

[FW05] Bruce Fraser and Janice Wooton. A proposed method for confidentialising tabular output to protect against differencing. https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=e60650e366ec84eade1c1dc919b1a9ac89e16850#page=285, 2005. URL: https://unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2005/wp.35.e.pdf.

[Jon21] Francis Jones. Statistical disclosure control for Caribbean census tables: A proposal to expand the availability of disaggregated census data. https://repositorio.cepal.org/server/api/core/bitstreams/ccea09b6-c57a-4124-8b01-0af6185b3ecf/content, 2021. URL: https://www.cepal.org/en/publications/46628-statistical-disclosure-control-caribbean-census-tables-proposal-expand.

[KAA$^+$22] Daniel Kifer, John M. Abowd, Robert Ashmead, Ryan Cumings-Menon, Philip Leclerc, Ashwin Machanavajjhala, William Sexton, and Pavel Zhuravlev. Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 Census. https://arxiv.org/pdf/2209.03310, 2022. URL: https://arxiv.org/abs/2209.03310.

[Lea09] Victoria Leaver. Implementing a method for automatically protecting user-defined Census tables. https://unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2009/wp.22.e.pdf, 2009.

[Mir17] Ilya Mironov. Rényi differential privacy. https://arxiv.org/pdf/1702.07476, 2017. https://doi.org/10.1109/CSF.2017.11.

[ML11] Jennifer K. Marley and Victoria L. Leaver. A method for confidentialising user-defined tables: statistical properties and a risk-utility analysis. https://2011.isiproceedings.org/papers/450007.pdf, 2011. URL: https://2011.isiproceedings.org/papers/450007.pdf.

[TBE13] Gwenda Thompson, Stephen Broadfoot, and Daniel Elazar. Methodology for the automatic confidentialisation of statistical outputs from remote servers at the australian bureau of statistics. https://unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2013/Topic_1_ABS.pdf, 2013. URL: https://unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2013/Topic_1_ABS.pdf.

Appendix A. Lagrange Optimsation of (2.1)

We first write the Lagrange function for the optimisation problem (2.1) subject to the conditions in (2.2) as

$$L(p_Z) = H(Z) - \underbrace{\alpha\,(\sum_z p_Z(z) - 1)}_{\text{valid pmf}} - \underbrace{\beta \sum_z z p_Z(z)}_{\text{zero bias}} - \underbrace{\gamma \left(\sum_z z^2 p_Z(z) - V\right)}_{\text{variance}}. \tag{A.1}$$

We can drop the condition of valid pmf by normalising the obtained result at the last step. Therefore, the simplified function becomes

$$L(p_Z) = -\sum_{z=-D}^{D} p(z)\log p(z) - \beta \sum_{z=-D}^{D} z p_Z(z) - \gamma \left(\sum_{z=-D}^{D} z^2 p_Z(z) - V\right). \tag{A.2}$$

We will also check that the solution is positive after the solution is derived. Taking the derivative of L with respect to $p_Z(z)$ and setting it to zero gives

$$\frac{\partial L}{\partial p_Z(z)} = -\log(p_Z(z)) - p_Z(z)\frac{1}{p_Z(z)} - \beta z - \gamma z^2 = 0,$$

so that

$$p_Z(z) = e^{(-1-\beta z-\gamma z^2)},$$

which is positive as expected from a probability mass. Next, we consider the bias. It is clear that $\beta = 0$ will ensure bias is zero. Hence, p_Z simplifies to

$$p_Z(z) = e^{-1-\gamma z^2}, \qquad z \in [-D, D]. \tag{A.3}$$

Normalising $p_Z(\cdot)$ so that it sums to one gives the desired result in (2.4).

Appendix B. More Details on Derivation of δ

For the mechanism $M(x)$, let $p_{M(x)}(y)$ denote the pmf of the random variable $Y = M(x)$ and $p_{M(x')}(y)$ denote the pmf of the random variable $Y' = M(x')$, where x, x' are neighbouring datasets. The privacy loss random variable is then defined as

$$\ell_{M,x,x'}(y) := \log\left(\frac{p_{M(x)}(y)}{p_{M(x')}(y)}\right).$$

[CKS22] show the worst-case event E^* that achieves the supreme in (2.9) is the subset of $\mathcal{Y}$ for which the privacy loss random variable is greater than ε. That is,

$$E^* = \left\{y \in \mathcal{Y} : \log\left(\frac{p_{M(x)}(y)}{p_{M(x')}(y)}\right) > \varepsilon\right\} \tag{B.1}$$

Referring to (1.1) and (2.3), we have $p_{M(x)}(y) = p_Z(y - q(x)) = Ce^{-\gamma(y-q(x))^2}$ and $p_{M(x')}(y) = p_Z(y - q(x')) = Ce^{-\gamma(y-q(x'))^2}$. Note that in counting queries, the largest absolute query sensitivity to neighbouring datasets is 1. That is, $q(x) - q(x') \in \{-1, 0, 1\}$ for two neighbouring datasets x, x'. Given the symmetry of $p(z)$ for the considered symmetric support $[-D, D]$, we need to only consider either $q(x) - q(x') = 1$ or $q(x) - q(x') = -1$

to characterise $\delta_{\gamma,D}(\varepsilon)$. Without loss of generality, let us set $q(x) = 0$ and $q(x') = -1$. Therefore, E^* is equivalently defined as

$$E^* = \left\{ z : z \in [-D, D+1], \frac{p(z)}{p(z-1)} > e^{\varepsilon} \right\}. \tag{B.2}$$

Appendix C. Derivation of δ for $D = \infty$

For ease of reference, we repeat equation (2.19) here.

$$\delta = \delta_{\gamma,D}(\varepsilon) = \begin{cases} Ce^{-\gamma D^2}, & \lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor \leq -D, \\ Ce^{-\gamma D^2} + C\sum_{z=-D+1}^{\lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor} (e^{-\gamma z^2} - e^{\varepsilon} e^{-\gamma (z-1)^2}), & -D < \lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor \leq 0, \end{cases} \tag{C.1}$$

For a finite ε and $\gamma := 1/2\sigma^2$, when $D = \infty$, the first case above becomes inactive and $Ce^{-\gamma D^2} = 0$. The second case becomes active and the summation becomes

$$\delta = C \sum_{z=-\infty}^{\lfloor 0.5 - \frac{\varepsilon}{2\gamma} \rfloor} (e^{-\gamma z^2} - e^{\varepsilon} e^{-\gamma (z-1)^2}) \tag{C.2}$$

$$= C \sum_{z=-\infty}^{\lfloor 0.5 - \varepsilon\sigma^2 \rfloor} e^{-\frac{z^2}{2\sigma^2}} - e^{\varepsilon} C \sum_{z=-\infty}^{\lfloor 0.5 - \varepsilon\sigma^2 \rfloor} e^{-\frac{(z-1)^2}{2\sigma^2}} \tag{C.3}$$

$$= C \sum_{z=-\infty}^{\lfloor 0.5 - \varepsilon\sigma^2 \rfloor} e^{-\frac{z^2}{2\sigma^2}} - e^{\varepsilon} C \sum_{z'=-\infty}^{\lfloor -0.5 - \varepsilon\sigma^2 \rfloor} e^{-\frac{z'^2}{2\sigma^2}}, \tag{C.4}$$

where we change the variable $z = z - 1$ in the second sum. If we carefully compare the last expression with the expression given for δ in equation (2.3) in Theorem 2.6 in [CKS22], we find that they are indeed identical for query sensitivity $\Delta = 1$.[4] We remark that this is not surprising. The reason is that we follow the work of [BW18] and [CKS22] by applying their methodology and using first-principle to derive δ. Our method does not provide an upper bound for δ, but computes it exactly. Indeed, a similar technique was used in Theorem 2.6 of [CKS22]. Hence, we get identical results when $D = \infty$. For finite D, our derivations are still exact for the pmf.

[4] Note that due to the symmetry of the discrete Gaussian pmf, one can replace $\Pr(Z > z)$ with $\Pr(Z < -z)$ and also since $Z \in \mathbb{Z}$, we have $\Pr(Z < -z) = \Pr(Z \leq \lfloor -z \rfloor)$.