# On the Privacy of Selection Mechanisms with Gaussian Noise

**Jonathan Lebensold**
McGill University, Mila

**Doina Precup**
McGill University, Mila, Google DeepMind

**Borja Balle**
Google DeepMind

## Abstract

Report Noisy Max and Above Threshold are two classical differentially private (DP) selection mechanisms. Their output is obtained by adding noise to a sequence of low-sensitivity queries and reporting the identity of the query whose (noisy) answer satisfies a certain condition. Pure DP guarantees for these mechanisms are easy to obtain when Laplace noise is added to the queries. On the other hand, when instantiated using Gaussian noise, standard analyses only yield approximate DP guarantees despite the fact that the outputs of these mechanisms lie in a discrete space. In this work, we revisit the analysis of Report Noisy Max and Above Threshold with Gaussian noise and show that, under the additional assumption that the underlying queries are bounded, it is possible to provide pure ex-ante DP bounds for Report Noisy Max and pure ex-post DP bounds for Above Threshold. The resulting bounds are tight and depend on closed-form expressions that can be numerically evaluated using standard methods. Empirically we find these lead to tighter privacy accounting in the high privacy, low data regime. Further, we propose a simple privacy filter for composing pure ex-post DP guarantees, and use it to derive a fully adaptive Gaussian Sparse Vector Technique mechanism. Finally, we provide experiments on mobility and energy consumption datasets demonstrating that our Sparse Vector Technique is practically competitive with previous approaches and requires less hyper-parameter tuning.

Proceedings of the 27[th] International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. 

## 1 INTRODUCTION

Differential Privacy (DP) (Dwork, 2006) has become the standard framework used for the private release of sensitive statistics. In particular, DP has been embraced by industry and governments to guarantee that potentially sensitive statistics cannot be linked back to individual users. For example, during the COVID-19 pandemic, Google Maps mobility data was published with DP (Aktay et al., 2020) in order for public health authorities to better understand various curve-flattening measures (such as work-from-home, or shelter-in-place). Recently, the Wikimedia Foundation deployed DP for their page-level visit statistics (Desfontaines, 2023).

Underpinning the design of differentially private mechanisms is a fundamental trade-off between privacy and utility characterized by the Fundamental Law of Information Recovery stating that "overly accurate answers to too many questions will destroy privacy in a spectacular way" (Dwork and Roth, 2014). In practice this imposes a limit on how many statistics can be privately released to a desired accuracy within a prespecified privacy budget. In applications where the space of possible statistics is too large to allow for a full private and accurate release, analysts can overcome the privacy-utility trade-off by identifying and releasing only those statistics that contain relevant information. This is necessary for example in periodic data collection where users contribute many times to a dataset and relevant statistics must be released repeatedly (e.g. to report temporal trends, change points, extreme events, etc.) (Hu et al., 2021; Wang et al., 2016; Xu et al., 2017). A notable example is the use of smart meters in energy grids, where statistics can help manage electrical demand and encourage smoother consumption but can also be used to infer information like income, occupancy, etc. (Ács and Castelluccia, 2011; Bohli et al., 2010; Haji Mirzaee et al., 2022).

Private selection mechanisms aim to identify relevant statistics. This problem is usually framed as query selection: an analyst defines a collection of queries against the target dataset, and a private mechanism is

used to identify queries returning "abnormally" large values. Two notable settings arise: the offline case, where all the queries can be specified in advance, and the online case, where the analyst can adaptively select queries based on the previous ones. In the offline setting, one of the best known private selection mechanisms is *Report Noisy Max* whereby the analyst submits a collection of low-sensitivity scalar queries. The mechanism then adds noise to the values of all the queries and returns the index of the query attaining the largest (noisy) value. The Exponential Mechanism (Dwork and Roth, 2014) and Permute-and-flip are commonly used for offline selection and are generally considered best-in-class (McKenna and Sheldon, 2020).

In the online setting, the cornerstone selection mechanism is *Above Threshold*, where the analyst submits a threshold and a sequence of (potentially adaptive) low-sensitivity scalar queries to which the mechanism iteratively computes answers, until a noisy value exceeds the target threshold. Running Above Threshold repeatedly is called the Sparse Vector Technique (Dwork et al., 2009), and is a common privacy primitive in change-point detection, empirical risk minimization, density estimation and other online learning algorithms (Zhang et al., 2021; Ligett et al., 2017; Dwork and Roth, 2014; Cummings et al., 2018).

When Laplace noise is used, Above Threshold and Report Noisy Max offer *pure* privacy guarantees – the strongest type of DP guarantee which does not have to account for some small failure probability. The use of the Laplace distribution also has the advantage of making the mathematical analysis of the privacy guarantees relatively simple; however, since the noise distribution is less concentrated than the Gaussian distribution, it can lead to a less accurate mechanism. Replacing the Laplace distribution with a Gaussian distribution is advantageous in many DP mechanisms, including online query selection (Abadi et al., 2016; Zhu and Wang, 2020; Papernot et al., 2018).

**Contributions.** In this paper we revisit the privacy analysis of the Above Threshold mechanism instantiated with Gaussian noise. Our main observation is that under the mild assumption that the queries submitted are uniformly bounded (in addition to low-sensitivity), we can provide pure DP guarantees that can be computed using standard numerical tools. In particular, we provide pure *ex-post*[1] DP guarantees for Gaussian Above Threshold. In the process, we also develop an *ex-ante* DP guarantee for Gaussian Report Noisy Max. Our analysis relies on identifying the worst case values for the query answers on a pair of neighboring datasets, a technique which might be of independent interest. Empirically, we find that these privacy bounds lead to tighter privacy accounting in the high privacy, low data regime, when compared to other Gaussian-based mechanisms.

Further, we define a meta-algorithm that composes Gaussian Above Threshold with ex-post DP guarantees. Thus, we derive a fully-adaptive Sparse Vector Technique (SVT), which we call *Filtered Self-Reporting Composition* (FSRC). Our method is particularly appealing since an analyst need not choose the number of releases or the maximum number of queries up front.

Finally, we provide experiments on mobility and energy consumption datasets demonstrating that our analyses yield mechanisms that in practice match or outperform previous approaches.

[1]This is a guarantee that depends on the output produced by the mechanism; see Section 3 for details.

## 2 PRELIMINARIES

**Differential Privacy.** Throughout we will be considering randomized mechanisms operating on some dataset $D \in \mathcal{X}$. Two datasets $D$ and $D'$ are said to be *neighboring*, denoted $D \simeq D'$, if they differ in the data of a single individual (e.g. one user is added, removed, or replaced by another).

**Definition 1** (DP (Dwork et al., 2006))**.** Let $\epsilon \geq 0$ and $\delta \geq 0$. A randomized mechanism $M : \mathcal{X} \to \mathcal{O}$, is $(\epsilon, \delta)$-DP if for every pair of neighboring datasets $D \simeq D'$ and every subset $S \subseteq \mathcal{O}$, we have:

$$\Pr[M(D) \in S] \leq e^{\epsilon} \Pr[M(D') \in S] + \delta \ .$$

When $\delta = 0$, we write $\epsilon$-DP and say the mechanism satisfies *pure DP*; otherwise we say that it satisfies *approximate DP*. If the output space $\mathcal{O}$ is discrete, then a mechanism is $\epsilon$-DP if and only if $\Pr[M(D) = o] \leq e^{\epsilon} \Pr[M(D') = o]$ for all $o \in \mathcal{O}$ and $D \simeq D'$.

A key feature of DP is its resilience to *post-processing*: any function of the output of a DP mechanism still satisfies DP with the same (or better) parameters. A common way to establish that a mechanism satisfies differential privacy is through a high-probability bound on the privacy loss random variable.

**Definition 2** (Privacy Loss (Dinur and Nissim, 2003))**.** Let $M : \mathcal{X} \to \mathcal{O}$ be a randomized mechanism and consider neighboring datasets $D \simeq D'$. The *privacy loss* of the pair $D$ and $D'$ under $M$ for a given output $o \in \mathcal{O}$ is defined as

$$\mathcal{L}_{M,D,D'}(o) = \log \frac{\Pr[M(D) = o]}{\Pr[M(D') = o]} \ .$$

**Definition 3** (pDP(Kasiviswanathan and Smith, 2014))**.** A mechanism $M : \mathcal{X} \to \mathcal{O}$ satisfies $(\epsilon, \delta)$-pDP if for any $D \simeq D' \in \mathcal{X}$,

$$\Pr_{o \sim M(D)} [\mathcal{L}_{M,D,D'}(o) > \epsilon] \leq \delta \ .$$

If a mechanism is $(\epsilon, \delta)$-pDP, then it is also $(\epsilon, \delta)$-DP (Kasiviswanathan and Smith, 2014). A simple way to obtain pDP guarantees is through bounds on the moment-generating function of the privacy loss random variable; this is one of the motivations for Rényi Differential Privacy (RDP), which relies on a bound of the Rényi divergence between two distributions.

**Definition 4** (Rényi divergence (Rényi, 1961))**.** Let $\alpha > 1$. The Rényi divergence of order $\alpha$ between two probability distributions $P$ and $Q$ on $\mathcal{X}$ is defined by:

$$\mathbb{D}_\alpha(P||Q) \triangleq \frac{1}{\alpha - 1} \log \mathbb{E}_{o \sim Q} \left[ \frac{P(o)}{Q(o)} \right]^\alpha .$$

**Definition 5** (Rényi DP (Mironov, 2017))**.** Let $\alpha > 1$ and $\epsilon \geq 0$. A randomized mechanism $M$ is $(\alpha, \epsilon)$-RDP for all $D \simeq D'$ if, $\mathbb{D}_\alpha(M(D) \| M(D')) \leq \epsilon$.

It is possible to convert RDP guarantees into probabilistic and approximate DP guarantees (Mironov, 2017; Balle et al., 2020). In particular, any $(\alpha, \epsilon)$-RDP mechanism satisfies $(\epsilon_p, \delta)$-pDP guarantees with $\epsilon_p = \epsilon + \log(1/\delta)/(\alpha - 1)$ by Markov's inequality. RDP greatly simplifies the analysis of mechanisms based on Gaussian noise because the Rényi divergence between Gaussian distributions has a simple expression, as well as a simple analysis under composition.

**Gaussian Mechanism.** Adding Gaussian noise to the result of a low-sensitivity query is a staple of differentially private mechanism design. Let $q : \mathcal{X} \to \mathbb{R}^d$ be a query with global sensitivity $\Delta_q = \sup_{D \simeq D'} \|q(D) - q(D')\|_2$ and $Z \sim \mathcal{N}(0, \sigma^2 I)$. The Gaussian Mechanism defined as $M(D) = q(D) + Z$ satisfies $(\alpha, \epsilon)$-RDP with $\epsilon = \frac{\alpha \Delta_q^2}{2\sigma^2}$ for every $\alpha > 1$ (Mironov, 2017). It is possible to convert this RDP guarantee into an approximate DP guarantee, although tighter approximate DP bounds can be obtained directly (Balle and Wang, 2018, Theorem 8).

**Gaussian Report Noisy Max.** In some applications it is useful to privately select which among a collection of queries $q_1, \ldots q_d : \mathcal{X} \to \mathbb{R}$ (approximately) attains the largest value on a given dataset. This leads to the *report noisy max* mechanism – when instantiated using Gaussian noise, the mechanism is given by, $M(D) = \arg\max_{i \in [d]} q_i(D) + Z_i$, where $Z_1, \ldots, Z_d \sim \mathcal{N}(0, \sigma^2)$. Since the arg max operation is merely a post-processing of the Gaussian mechanism applied to the $d$-dimensional query $q = (q_1, \ldots, q_d) : \mathcal{X} \to \mathbb{R}^d$ with sensitivity $\Delta_q$, the Gaussian Report Noisy Max mechanism inherits the same privacy guarantees as the Gaussian mechanism above (e.g. the bounds provided by Balle and Wang (2018, Theorem 8) or the RDP to DP conversion). Note that if each of the individual queries has sensitivity bounded by $\Delta$, then we have $\Delta_q \leq \sqrt{d}\Delta$.

**Gaussian Above Threshold.** The Above Threshold mechanism receives a sequence of low-sensitivity queries and privately identifies the first query (approximately) exceeding a given threshold. The mechanism was introduced by Dwork et al. (2009) and forms the basis of the Sparse Vector Technique, a composition of Above Threshold algorithms to find a sparse set of relevant queries among a large set. The standard version of the Above Threshold algorithm uses Laplace noise, and its privacy analysis is notoriously subtle (Lyu et al., 2016). Zhu and Wang (2020) recently proposed an RDP analysis which can be applied to the Gaussian version of Above Threshold (see Algorithm 3).

**Algorithm 1** Gaussian Above Threshold (Zhu and Wang, 2020)

**input:** dataset $D$; noise parameters $\sigma_X, \sigma_Z$; a stream of queries $q_1, q_2, \ldots$; threshold $\rho$.
$\hat{\rho} = \rho + \mathcal{N}(0, \sigma_X^2)$
**for** $t = 1, 2, \ldots$ **do**
  $\hat{q}_t = q_t(D) + \mathcal{N}(0, \sigma_Z^2)$
  **if** $\hat{q}_t \geq \hat{\rho}$ **then**
    Output $x_t = \top$ and HALT
  **else**
    Output $x_t = \bot$
  **end**
**end**

**Theorem 6** (General RDP Bound on Gaussian Above Threshold (Zhu and Wang, 2020))**.** *Suppose all queries given to the mechanism $M$ in Algorithm 3 have sensitivity bounded by $\Delta$. Then for $\gamma > 1$, $\infty > \alpha > 1$,*

$$\mathbb{D}_\alpha \left( M(D) \| M(D') \right) \leq \left( \frac{\gamma}{\gamma - 1} \right) \left( \frac{\alpha \Delta^2}{2\sigma_X^2} \right) + \frac{2\alpha\Delta^2}{\sigma_Z^2} + \frac{\log \mathbb{E}_{x \sim \mathcal{N}(0, \sigma_X^2)} \left[ \mathbb{E}[T \mid x]^\gamma \right]}{\gamma(\alpha - 1)},$$

*where $T$ is a random variable indicating the stopping time of $M(D)$.*

Unlike in the Laplace-based Above Threshold, the privacy bound for the Gaussian case depends on how long it takes the mechanism to stop (the third term in the expression above). When the queries are known a priori to be non-negative, it is possible to obtain bounds

on the running time to provide the RDP guarantee below.

**Theorem 7** (Gaussian Above Threshold RDP (Zhu and Wang, 2020))**.** *Gaussian Above Threshold, with $\sigma_Z \geq \sqrt{3}\sigma_X$, threshold $\rho \geq 0$, and $\Delta$-sensitive, non-negative queries, satisfies $(\alpha, \epsilon)$-RDP with*

$$\epsilon = \frac{\alpha \triangle^2}{\sigma_X^2} + \frac{2\alpha \triangle^2}{\sigma_Z^2} + \frac{\log\left(1 + 2\sqrt{3}\pi\left(1 + \frac{9\rho^2}{\sigma_X^2}\right)e^{\frac{\rho^2}{\sigma_X^2}}\right)}{2(\alpha - 1)} \quad .$$

Alternative methods to bound the running time often include the analyst supplying a bound $k$ on the running time to the algorithm, forcing it to stop after a certain number of steps even if the threshold has not been exceeded. Zhu and Wang (2020) also propose a number of composition results for the Sparse Vector Technique. One version allows the analyst to continue releasing queries without having to re-sample the threshold noise $\sigma_X$, but in each case the analyst must know ahead of time the number of times they wish to release a "$\top$", and have an upper bound on the maximum number of queries they intend to run.

# 3 PURE PRIVATE GAUSSIAN MECHANISMS

To introduce our main contribution, we begin with a warm-up. We propose a pure DP bound for Gaussian Report Noisy Max. The proof techniques are the same for Above Threshold, but the analysis is simpler since all the queries are symmetric.

## 3.1 Warm-Up: Gaussian Report Noisy Max

The following is a pure DP bound for Gaussian Report Noisy Max for queries with bounded range.

**Theorem 8** (Pure DP for Gaussian Report Noisy Max)**.** *Let $M$ be the Gaussian Report Noisy Max mechanism with standard deviation $\sigma$ applied to $d > 1$ bounded queries $q_1, \ldots, q_d : \mathcal{X} \to [a, b]$, each with sensitivity bounded by $\Delta$. Let $c = b - a$. Let $\Phi(\cdot)$ be the standard Gaussian CDF. Then $M$ satisfies $\epsilon$-DP with*

$$\epsilon = \frac{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\Phi\left(z - \frac{c-2\Delta}{\sigma}\right)^{d-1}\right]}{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\Phi\left(z - \frac{c}{\sigma}\right)^{d-1}\right]} \quad .$$

Notably, with the standard analysis of Gaussian Report Noisy Max, the privacy guarantee only depends on the ratio $\Delta/\sigma$. However in our case, the bound on privacy additionally depends on the range of the queries through $c/\sigma$. In fact, it is easy to see that if $c \to \infty$ then Gaussian Report Noisy Max cannot admit a pure DP bound.

The proof, deferred to the supplemental, relies on identifying the worst-case values that uniformly bounded queries can attain on a pair of neighboring datasets. In particular, we show that the worst case is obtained on a pair of neighboring datasets $D$ and $D'$ such that $q_1(D) = \cdots = q_{d-1}(D) = b - \Delta$, $q_d(D) = a + \Delta$, $q_1(D') = \cdots = q_{d-1}(D') = b$, and $q_d(D') = a$.

Note that the classical approach sketched in the previous section applied to our setting would consider the privacy of a Gaussian mechanism with a $d$-dimensional query of sensitivity $\Delta_q = \Delta\sqrt{d}$, and treat the arg max operation as a post-processing step. While our approach gives a pure DP guarantee, the classical approach does not because Gaussian noise by itself cannot provide that guarantee. However, the classical approach gives a bound that is easy to compute.

The RDP approach gives a simple closed form expression, while the direct approximate DP approach gives a bound that has a simple dependence on the CDF of a standard Gaussian (Balle and Wang, 2018). In contrast, the bound provided by our result requires estimating Gaussian expectations of a complex function; unfortunately, this does not admit a simple closed form expression. In Section 3.4 we evaluate numerical methods for computing this quantity and compare the resulting values of $\epsilon$ with those provided by the classical approach.

## 3.2 Ex-Post Gaussian Above Threshold

In the preceding section, we introduced a method to judge the privacy of Above Threshold before execution (Theorem 7) – these are often referred to as *ex-ante* privacy guarantees. Alternatively, we can measure the privacy loss after execution, leading to so-called *ex-post* privacy guarantees.

**Definition 9** (Ex-Post DP (Ligett et al., 2017))**.** Consider a function $\epsilon_{\text{p}} : \mathcal{O} \to \mathbb{R}_+ \cup \{\infty\}$. A randomized mechanism $M : \mathcal{X} \to \mathcal{O}$ satisfies $\epsilon_{\text{p}}$-ex-post-DP if for any possible output $o \in \mathcal{O}$ and any pair of neighboring datasets we have $\mathcal{L}_{M,D,D'}(o) \leq \epsilon_{\text{p}}(o)$.

Since Above Threshold outputs the (approximate) halting time, we can exploit the difference between when the ex-post and ex-ante privacy guarantee. In Fig. 1, we observe when the ex-post analysis is most beneficial, and when it is most likely to occur. As the public threshold increases, so does the separation between the ex-ante analysis and the ex-post analysis. With Above Threshold, we also need to consider how the mechanism might perform against a worst-case dataset maximizing the stopping time. Since queries are non-negative and bounded, this occurs when $q_1(D) = \cdots = q_t(D) = 0$. We then compute the median stopping time (dashed blue line) and the 80th

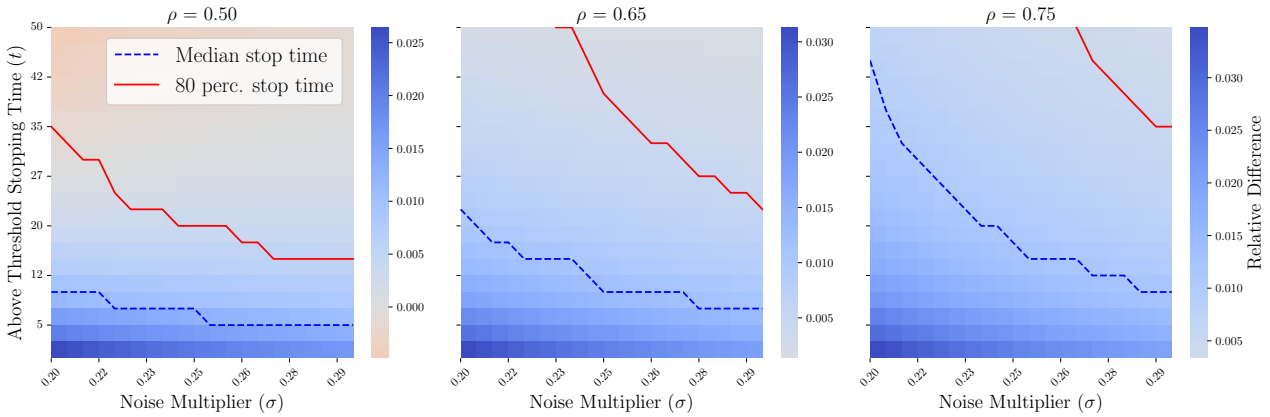


Figure 1: Privacy accounting for $\Delta = 1e{-}3, \delta = 1e{-}5$. The heatmap shows where the ex-post Above Threshold Analysis offers an improvement over the Gaussian Above Threshold. As a comparison, we simulate expected stopping times for a range of multipliers, for $[a, b] = [0, 1]$. The blue dotted line corresponds to the median stopping time when simulating 10k trials with a worst-case dataset. The red line corresponds to the 80th percentile. The plot shows a range of hyper-parameters as well as where the worst-case dataset is likely to halt. Our bounds provide improvements over the baseline below the blue line when squares are blue.

percentile (solid red line) for the worst case dataset. For bounded range between zero and one, we see the greatest effect when the mechanism halts early and $\rho$ is calibrated to be closer to one.

When the queries have bounded range, Gaussian Above Threshold (Algorithm 3) admits the following ex-post privacy bound.

**Theorem 10** (Pure Ex-post Gaussian Above Threshold)**.** *Let*

$$\psi_\xi(x) = \Phi\left(\left(x + \rho - (b + a) + \xi\right)/\sigma_Z\right) \quad \text{, and,}$$
$$\beta_\xi(x) = \Phi\left(\left(x - \rho + a + \xi\right)/\sigma_Z\right) \quad .$$

*Given a stream $q_1, q_2, \ldots : \mathcal{X} \to [a, b]$ with global sensitivity $\Delta$, the Gaussian Above Threshold mechanism (Algorithm 3), halting at time step $t$ with $o = \{\perp^{t-1}\top\}$, satisfies $\epsilon_p$-ex-post-DP with*

$$\epsilon_p(o) = \frac{\mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\psi_\Delta(\sigma_X x)^{(t-1)} \cdot \beta_\Delta(\sigma_X x)\right]}{\mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\psi_0(\sigma_X x)^{(t-1)} \cdot \beta_0(\sigma_X x)\right]} \quad .$$

Note that this formulation is very similar to the ratio of expectations in our pure DP analysis of Report Noisy Max (Theorem 12). The proof, deferred to the supplemental, follows from Pure DP Gaussian Report Noisy Max, where we find that the point where the ratio is maximized is the same in all but the final time step. In the last step, the query bound is negated. Hence, the ratio is largest when $q_t = a$. A final remark is that the mechanism's greatest privacy loss occurs when each prior step would have been more optimal than the step at which it halts.

### 3.3 Fully-Adaptive Composition with Ex-Post Privacy Guarantees

Above Threshold stops the first time the threshold is exceeded. The Sparse Vector Technique (SVT) applies a sequence of Above Threshold mechanisms to find a set of queries that (approximately) exceed the pre-defined threshold. In the case of Laplace noise, the privacy analysis of SVT can be performed either directly (if the noise applied to the threshold is not refreshed after each Above Threshold terminates) or via composition (Dwork and Roth, 2014; Lyu et al., 2016). Something similar holds for SVT with Gaussian noise, although in this case the analysis without noise resampling only applies to a range of parameters (Zhu and Wang, 2020). Here we present a simple technique for fully adaptive composition of mechanisms that have both pDP and ex-post-DP guarantees, which can be combined with our analysis of Gaussian Above Threshold to yield a fully adaptive SVT with Gaussian noise algorithm without additional hyper-parameters like the maximum number of queries per Above Threshold or the maximum number of invocations of the Above Threshold mechanism.

Our method (Algorithm 4) sequentially composes a stream of adaptive mechanisms and applies a stopping time rule that limits over-spending a predefined privacy budget. The privacy expenditure of each mechanism is tracked based on their output, using the ex-post privacy guarantee. The stopping rule uses the probabilistic DP guarantee to halt when there is a high enough probability of exceeding the privacy budget.

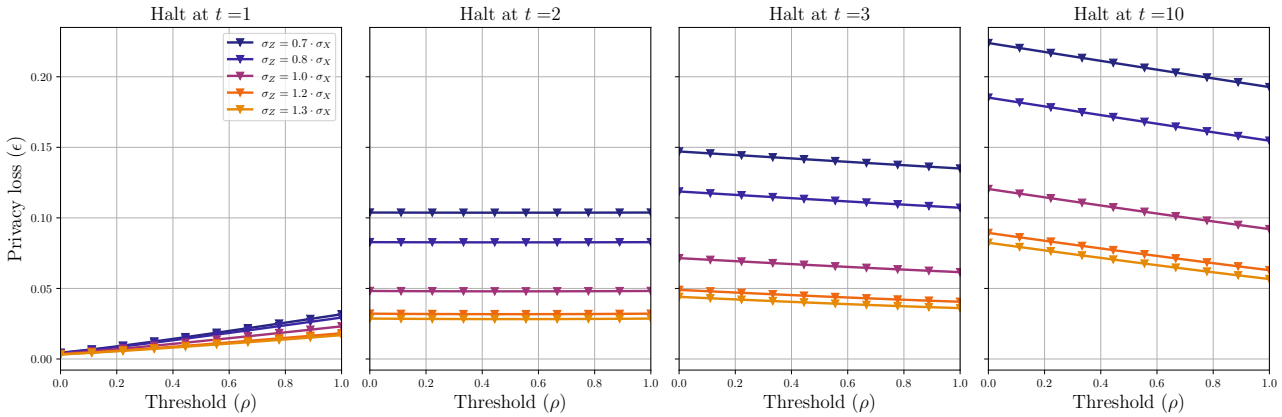


Figure 2: Ex-Post Above Threshold Privacy Loss ($\Delta = 0.001, \sigma_X = 0.15$). The ex-post privacy loss also changes as a function of the public threshold $\rho$. Note that if the mechanism halts after two timesteps, the minimum is observed when $\rho = 0.5$. As $t$ increases, the privacy loss decreases as $\rho \to 1$.

**Algorithm 2** Filtered Composition, Ex-Post Privacy

**input:** dataset $D$; privacy budget $\epsilon$; a stream of adaptive mechanisms $M_1, M_2, \ldots$; a stream of values $\epsilon_{\max,1}, \epsilon_{\max,2}, \ldots$; and stream of functions $\epsilon_{\mathrm{p},1}, \epsilon_{\mathrm{p},2}, \ldots$

**for** $t = 1, 2, \ldots$ **do**
  **if** $\sum_{i=1}^{t-1} \epsilon_i + \epsilon_{\max,t} \geq \epsilon$ **then**
    HALT
  **else**
    $o_t = M_t(D; o_{1:t-1})$
    $\epsilon_t = \epsilon_{\mathrm{p},t}(o_t)$
    Release $o_t$
  **end**
**end**

**Theorem 11.** *Suppose the mechanisms $M_1, M_2, \ldots$ provided to Algorithm 4 are such that $M_t$ is $(\epsilon_{\max,t}, \delta)$-pDP and $\epsilon_{p,t}$-ex-post-DP for all $t$. Then Algorithm 4 satisfies $(\epsilon, \delta)$-DP.*

In Fig. 2 we plot the ex-post privacy loss bound as a function of the public threshold parameter, $\rho$. Given a stream of queries $i = 1, 2, \ldots$, bounded by $a \leq q_i \leq b$, the mechanism will report less privacy loss as $\rho \to b$.

### 3.4 Numerical Computation of Bounds

We evaluate two methods for producing numerical estimates of the bounds: Monte Carlo and numerical integration. Monte Carlo density estimation is a natural starting point for computing bounds for Theorem 12 and Theorem 17. However, Monte Carlo methods require a tremendous amount of samples to yield accurate results. Note that in both cases, the CDF is taken to an exponential power in the numerator and the denominator, causing numerical instabilities. To improve runtime performance and stability, we use a scientific library that can compute integrals with high precision. The Python mpmath (mpmath, 2023) library is able to compute each density and the outer integral with arbitrary degree of precision.

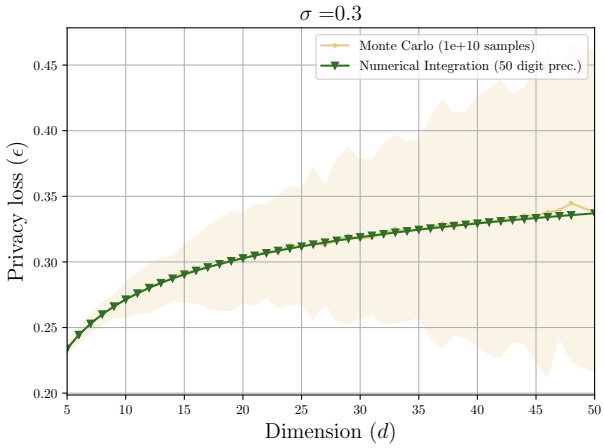


Figure 3: Gaussian Report Noisy Max for $\Delta = 0.01$. Numerical integration (green) compared to Monte Carlo estimate (beige) with 10B samples. Shaded region is the standard deviation over 100 trials. Numerical integration methods are deterministic; error bars only apply to Monte Carlo estimates, which are known to converge to the true estimate with infinite samples.

As shown in Fig. 3, for low sensitivity queries $\Delta$ and $\sigma = 0.3$, the variance between trials becomes unwieldy, even when we rely on common open source libraries that can precisely estimate the CDF of a univariate Gaussian.

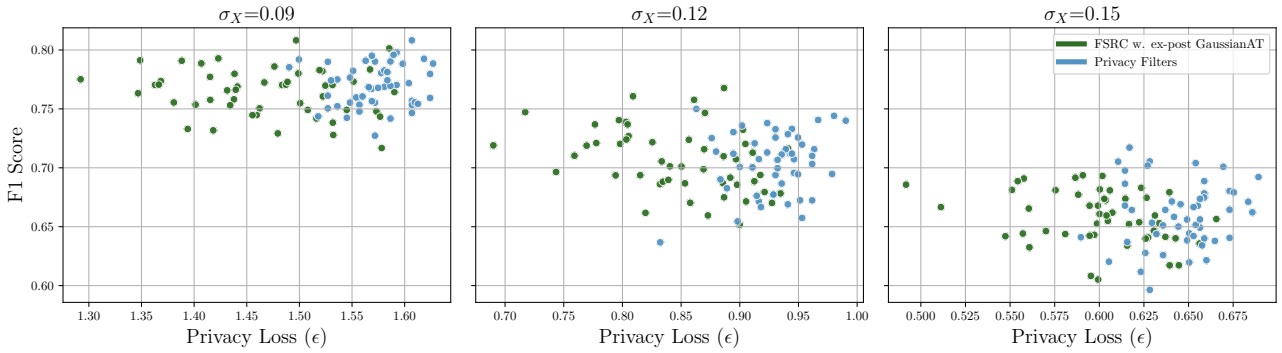


Figure 4: Scatter plot indicating the accuracy for UCI Bikes with $\rho = 0.575$. and final privacy spend over a range of noise multipliers. Final privacy loss ($\epsilon$) is reported for FSRC (green). Threshold noise, $\sigma_X$ in the range of $[0.09, 0.15]$. A clear separation in privacy accounting occurs over a range of noise multipliers.

## 4 EMPIRICAL RESULTS

We benchmark our SVT-like method (FSRC and Gaussian Above Threshold) on mobility and energy datasets. Additional experiments with the Pure DP Gaussian Report Noisy Max bound are included in the appendix. Experiments were done on a Apple M1 processor (32 GB), except for the Monte Carlo numerical estimation, which was done on a NVIDIA V100 GPU with 32 GB VRAM.

**The UCI Bikes Dataset** captures the utilization of shared bikes in a geographic area over the course of a year (Fanaee-T and Gama, 2013). Since we do not know the upper bound on registered customers, we take the maximum (6,946 users) and assume that this is a public value. Note that our analysis is still worst-case, in that a user is assumed to be contributing to each daily (normalized) count query. We set Above Threshold to HALT when bike sharing exceeds a threshold $\rho \in [0, 1]$. In Fig. 17 we plot a range of calibrations, and in the supplemental we show the privacy spend over time. **The LCL London Energy Dataset** (Greater London Authority, 2012), consists of energy usage for $N = 5,564$ customers over $d = 829$ days. The larger number of queries *increases* the privacy cost; however this is balanced by a *decrease* in individual contribution to each query due to each query having $\Delta = 1/N$. In Fig. 5 we plot a range of calibrations. As the threshold decreases, we witness more queries released and therefore a greater privacy spend.

### 4.1 Online Selection with Filtered Self-Reporting Composition

Answering sequential questions in an interactive setting typically requires the up-front selection of privacy parameters. The most common method to answer a large number of queries is the Sparse Vector Technique. In this setting, Above Threshold serves as a privacy primitive in more complex algorithms (Hardt and Rothblum, 2010). First, the curator decides how many times they expect to release a query in order to allocate the privacy budget. The budget is further split across two noise adding mechanisms. Noise is added to each query, as well as a public threshold parameter, which centers whether a point is flagged as a "$\top$" (interesting) or a "$\perp$" (not interesting). One solution—and our baseline for consideration—restricts a curator to fully-adaptive composition using a privacy filter (Rogers et al., 2023) and a novel result by Zhu and Wang (2020) which places no restriction on the number of queries. We are primarily concerned with satisfying the following two requirements: (1) the analyst should be able to modify queries after Above Threshold halts, and (2) they should be able to guarantee that an up-front privacy budget is respected. Therefore, we combine FSRC with the privacy analysis in Theorem 17. To calibrate $\epsilon_{\max}$, we take an RDP bound of Gaussian Above Threshold (Theorem 7) with $\delta = 1/N$.

**Experiment Parameters.** In each case, the datasets have a temporal axis, meaning that when the mechanism halts, we restart at the $t+1$' timestep and continue accumulating privacy spend. To evaluate FSRC, we compare Algorithm 4 using Gaussian Above Threshold (Algorithm 3), to a Privacy Filter (Whitehouse et al., 2023) with sequential composition of the same mechanism. In FSRC, we apply Theorem 17 in our privacy accounting. For the baseline, we compute an RDP bound using Zhu and Wang (2020). For each pair of noise multipliers and thresholds, we plot each result in a scatter plot, and the privacy spend over 50 runs. We use AutoDP (Wang, 2023) with support for global sensitivity calibration to apply their bound with $\Delta < 1$. As in Zhu and Wang (2020), we fix $\sigma_Z = \sqrt{3}\sigma_X$. We ran all our experiments over 50 runs

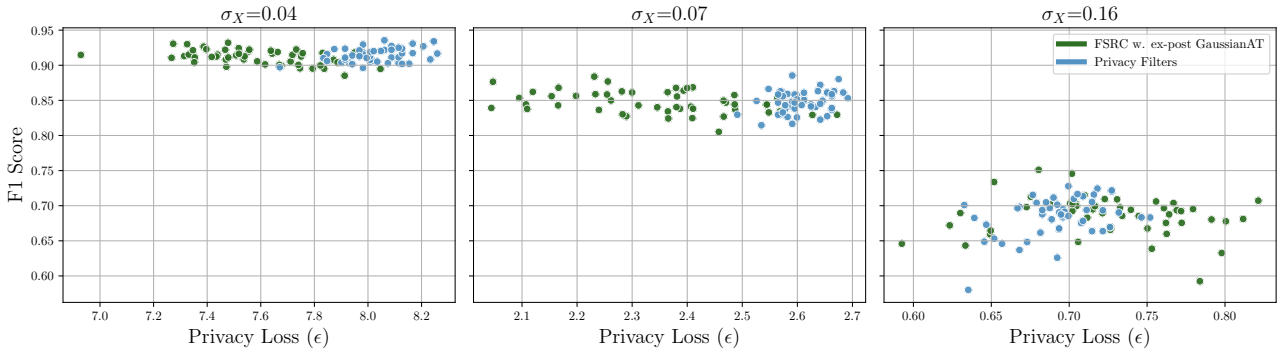


Figure 5: Scatter plot indicating the accuracy and final privacy spend over a range of noise multipliers for $\rho = 0.33$ with the LCL London Energy dataset. Privacy loss ($\epsilon$) is reported for FSRC (green). Threshold noise, $\sigma_X$, was evaluated in the range of $[0.04, 0.16]$. Our accounting method provides benefits when $\sigma_X = 0.04$.

with a range of values for $\sigma_X$. Importantly, the privacy spend using a DP filter cannot be disclosed without leaking privacy, and must therefore be replaced with a privacy odometer if the data curator wishes to share the spent budget (Rogers et al., 2016). **Utility.** We compute the F1 Score to measure utility for both algorithms. Since each algorithm outputs a binary vector, we can measure how many times the mechanism matches with a ground truth algorithm where no noise is added to the queries or the threshold.

## 5 RELATED WORK

Our work spans four areas of privacy research, (1) accounting methods for Gaussian mechanisms, (2) maximizing the number of interactive, user-level queries, (3) privacy filters and fully adaptive composition, and (4) ex-post privacy analysis.

**The Gaussian Mechanism** adds calibrated noise to the output of a query. In the learning setting, a number of privacy accounting techniques exist (Abadi et al., 2016). However such approaches do not map directly to query selection. In the online setting, ex-post DP accounting already exists for the Laplace Mechanism (Ligett et al., 2017); however, many successful deployments of DP rely on Gaussian mechanisms. This is likely due to the greater noise concentration around the mean and the thinner tails exhibited by Gaussian mechanisms (Dwork and Roth, 2014).

**The Sparse Vector Technique (SVT)** (Dwork et al., 2009) is a foundational differentially private algorithm. By splitting the privacy budget between the cost of returning a binary vector and the query of interest, utility is increased.

A Gaussian version offers better utility over the Laplace mechanism and can, surprisingly, support a potentially infinite number of queries (Zhu and Wang, 2020).

Hartmann et al. (2022) introduce *Output DP* as a means of analyzing SVT and the Propose-Test-Release with Laplace noise. Their work generalizes results from Ligett et al. (2017) and they show how basic composition bounds can (for few queries) offer better utility over advanced composition results.

**Ex-Post Privacy.** Ligett et al. (2017) defined ex-post privacy in terms of $(\epsilon, 0)$-DP. In common with this work, they studied SVT, but with a focus on the Laplace Mechanism. Redberg and Wang (2021) consider ex-post privacy with Gaussian mechanisms for data-dependent privacy parameters with the Gaussian mechanism over real-valued queries.

**Privacy Filters**, first proposed by Rogers et al. (2016), are a key ingredient in allowing adaptive composition of privacy-preserving mechanisms as well as the privacy spend. Feldman and Zrnic (2021), and Lécuyer (2021) extended these results to Rényi DP, with the caveat that the higher order parameters needed to be pre-defined. Recently, Whitehouse et al. (2023) tightened these results and Rogers et al. (2023) then applied ex-post privacy to probabilistic DP mechanisms. We consider their efforts to be closest to ours; however, they do not consider query online selection.

## 6 CONCLUSION

We provided pure-DP bounds on Gaussian selection mechanisms with bounded queries. Additionally, we developed new composition tools for ex-ante and ex-post privacy analysis. We demonstrated increased query accuracy for an equivalent privacy budget in energy and mobility datasets in the online setting.

We consider ex-post privacy analysis a promising

method tightening privacy accounting in online algorithms, and our composition result, FSRC, could be applied to other sequential algorithms.

Accounting for user-level privacy over long time horizons is necessary in a number of sequential and interactive decision-making tasks. In particular, we foresee direct benefit in applying these methods to model selection.

### Acknowledgements

J.L. is supported by the Google DeepMind Graduate Fund and the Fonds de Recherche du Québec. We wish to thank Thomas Steinke for his comments on an early draft. We also thank Guillaume Rabusseau, Jose Gallego-Posada, Maxime Wabartha and Vincent Luczkow for many fruitful discussions. Finally, we thank Iosif Pinelis for bringing l'Hôpital's Monotone Rule to our attention.

### References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318.

Ács, G. and Castelluccia, C. (2011). I have a dream!(differentially private smart metering). In *International Workshop on Information Hiding*, pages 118–132. Springer.

Aktay, A., Bavadekar, S., Cossoul, G., Davis, J., Desfontaines, D., Fabrikant, A., Gabrilovich, E., Gadepalli, K., Gipson, B., Guevara, M., et al. (2020). Google covid-19 community mobility reports: anonymization process description (version 1.1). *arXiv preprint arXiv:2004.04145*.

Anderson, G., Vamanamurthy, M., and Vuorinen, M. (1993). Inequalities for quasiconformal mappings in space. *Pacific J. Math.*, 160(1):1–18.

Balle, B., Barthe, G., Gaboardi, M., Hsu, J., and Sato, T. (2020). Hypothesis testing interpretations and renyi differential privacy. In Chiappa, S. and Calandra, R., editors, *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pages 2496–2506. PMLR.

Balle, B. and Wang, Y.-X. (2018). Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning*, pages 394–403. PMLR.

Bohli, J.-M., Sorge, C., and Ugus, O. (2010). A privacy model for smart metering. In *2010 IEEE International Conference on Communications Workshops*, pages 1–5.

Cummings, R., Krehbiel, S., Lai, K. A., and Tantipongpipat, U. (2018). Differential privacy for growing databases. *Advances in Neural Information Processing Systems*, 31.

Desfontaines, D. (2023). Publishing wikipedia usage data with strong privacy guarantees.

Dinur, I. and Nissim, K. (2003). Revealing information while preserving privacy. In *Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '03, page 202–210. Association for Computing Machinery.

Dwork, C. (2006). Differential privacy. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 4052 LNCS, pages 1–12. Springer Verlag.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 3876 LNCS, pages 265–284.

Dwork, C., Naor, M., Reingold, O., Rothblum, G. N., and Vadhan, S. (2009). On the complexity of differentially private data release. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, New York, NY, USA. ACM.

Dwork, C. and Roth, A. (2014). *The Algorithmic Foundations of Differential Privacy*. Foundations and trends in theoretical computer science. Now.

Fanaee-T, H. and Gama, J. (2013). Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence*, pages 1–15.

Feldman, V. and Zrnic, T. (2021). Individual privacy accounting via a renyi filter. *Advances in Neural Information Processing Systems*, 34:28080–28091.

Greater London Authority (2012). Smartmeter Energy Use Data in London Households.

Haji Mirzaee, P., Shojafar, M., Cruickshank, H., and Tafazolli, R. (2022). Smart grid security and privacy: From conventional to machine learning issues (threats and countermeasures). *IEEE Access*, 10:52922–52954.

Hardt, M. and Rothblum, G. N. (2010). A multiplicative weights mechanism for privacy-preserving data

analysis. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE.

Hartmann, V., Bindschaedler, V., Bentkamp, A., and West, R. (2022). Privacy accounting $\varepsilon$ conomics: Improving differential privacy composition via a posteriori bounds. *arXiv preprint arXiv:2205.03470*.

Hu, T., Wang, S., She, B., Zhang, M., Huang, X., Cui, Y., Khuri, J., Hu, Y., Fu, X., Wang, X., et al. (2021). Human mobility data in the covid-19 pandemic: characteristics, applications, and challenges. *International Journal of Digital Earth*, 14(9):1126–1147.

Kasiviswanathan, S. P. and Smith, A. (2014). On the 'semantics' of differential privacy: A bayesian formulation. *Journal of Privacy and Confidentiality*, 6(1).

Lécuyer, M. (2021). Practical privacy filters and odometers with r\'enyi differential privacy and applications to differentially private deep learning. *arXiv preprint arXiv:2103.01379*.

Ligett, K., Neel, S., Roth, A., Waggoner, B., and Wu, S. Z. (2017). Accuracy first: Selecting a differential privacy level for accuracy constrained erm. *Advances in Neural Information Processing Systems*, 30.

Lyu, M., Su, D., and Li, N. (2016). Understanding the sparse vector technique for differential privacy. *arXiv preprint arXiv:1603.01699*.

McKenna, R. and Sheldon, D. R. (2020). Permute-and-flip: A new mechanism for differentially private selection. *Advances in Neural Information Processing Systems*, 33:193–203.

Mironov, I. (2017). Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE.

mpmath (2023). *mpmath: a Python library for arbitrary-precision floating-point arithmetic (version 1.3.0)*. `http://mpmath.org/`.

Narayanan, A. and Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125.

Papernot, N., Song, S., Mironov, I., Raghunathan, A., Talwar, K., and Erlingsson, Ú. (2018). Scalable private learning with pate. *arXiv preprint arXiv:1802.08908*.

Pinelis, I. (2002). L'hospital type rules for monotonicity, with applications. *J. Inequal. Pure Appl. Math*, 3(1).

Ratnam, E. L., Weller, S. R., Kellett, C. M., and Murray, A. T. (2017). Residential load and rooftop pv generation: an australian distribution network dataset. *International Journal of Sustainable Energy*, 36(8):787–806.

Redberg, R. and Wang, Y.-X. (2021). Privately publishable per-instance privacy. *Advances in Neural Information Processing Systems*, 34:17335–17346.

Rényi, A. (1961). On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, volume 4, pages 547–562.

Rogers, R., Samorodnitsky, G., Wu, Z. S., and Ramdas, A. (2023). Adaptive privacy composition for accuracy-first mechanisms. *arXiv preprint arXiv:2306.13824*.

Rogers, R. M., Roth, A., Ullman, J., and Vadhan, S. (2016). Privacy odometers and filters: Pay-as-you-go composition. *Advances in Neural Information Processing Systems*, 29.

Wang, Q., Zhang, Y., Lu, X., Wang, Z., Qin, Z., and Ren, K. (2016). Real-time and spatio-temporal crowd-sourced social network data publishing with differential privacy. *IEEE Trans. Dependable Secure Comput.*, pages 1–1.

Wang, Y.-X. (2023). autodp: autodp: A flexible and easy-to-use package for differential privacy.

Whitehouse, J., Ramdas, A., Rogers, R., and Wu, S. (2023). Fully-adaptive composition in differential privacy. In *International Conference on Machine Learning*, pages 36990–37007. PMLR.

Xu, F., Tu, Z., Li, Y., Zhang, P., Fu, X., and Jin, D. (2017). Trajectory recovery from ash: User privacy is not preserved in aggregated mobility data. In *Proceedings of the 26th international conference on world wide web*, pages 1241–1250.

Zhang, W., Krehbiel, S., Tuo, R., Mei, Y., and Cummings, R. (2021). Single and multiple Change-Point detection with differential privacy. *J. Mach. Learn. Res.*, 22(29):1–36.

Zhu, Y. and Wang, Y.-X. (2020). Improving sparse vector technique with renyi differential privacy. *Advances in Neural Information Processing Systems*, 33:20249–20258.

# A PROOFS FROM SECTION 2

Our main contribution relies on a number of proof techniques which are extended from the simpler Gaussian Report Noisy Max setting.

## A.1 Pure DP Gaussian Report Noisy Max

**Theorem 12** (Pure DP for Gaussian Report Noisy Max)**.** *Let $M$ be the Gaussian Report Noisy Max mechanism with standard deviation $\sigma$ applied to $d > 1$ bounded queries $q_1, \ldots, q_d : \mathcal{X} \to [a, b]$, each with sensitivity bounded by $\Delta$. Let $c = b - a$. Then $M$ satisfies $\epsilon$-DP with*

$$\epsilon = \frac{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\Phi\left(z - \frac{c-2\Delta}{\sigma}\right)^{d-1}\right]}{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\Phi\left(z - \frac{c}{\sigma}\right)^{d-1}\right]} \quad . \tag{1}$$

To prove our Pure DP Gaussian Report Noisy Max bound, we first analyze the effect of the sensitivity in the privacy loss of the Gaussian Report Noisy Max mechanism instantiated with uniformly bounded queries.

**Lemma 13.** *Let $M$ be the Gaussian Report Noisy Max mechanism with standard deviation $\sigma$ applied to $d > 1$ bounded queries $q_1, \ldots, q_d : \mathcal{X} \to [a, b]$, each with sensitivity bounded by $\Delta$. Then we have*

$$\sup_{D\simeq D'}\sup_{o\in[d]}\frac{\Pr[M(D)=o]}{\Pr[M(D')=o]} \leq \sup_{y_1,\ldots,y_{d-1}\in[a-b+2\Delta,b-a]}\frac{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{d-1}\Phi\left(z-\frac{y_i-2\Delta}{\sigma}\right)\right]}{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{d-1}\Phi\left(z-\frac{y_i}{\sigma}\right)\right]} \quad . \tag{2}$$

*Proof.* Fix a pair of neighbouring datasets $D$ and $D'$. With a slight abuse of notation we let $q_i = q_i(D)$ and $q'_i = q_i(D')$ for all $i \in [d]$. Recall that $M(D) = \arg\max_{i\in[d]} q_i + Z_i$ with $Z_i \sim \mathcal{N}(0, \sigma^2)$, and $|q_i - q'_i| \leq \Delta$ for all $i$. Without loss of generality fix the output $o = d$. Then we have:

$$\begin{aligned}
\Pr[M(D)=d] &= \Pr[\wedge_{i=1}^{d-1} q_d + Z_d > q_i + Z_i] && (3)\\
&= \Pr[\wedge_{i=1}^{d-1} Z_i < Z_d + q_d - q_i] && (4)\\
&= \Pr[\wedge_{i=1}^{d-1} Z_i < Z_d + q'_d - q'_i + (q_d - q'_d) - (q_i - q'_i)] && (5)\\
&\leq \Pr[\wedge_{i=1}^{d-1} Z_i < Z_d + q'_d - q'_i + 2\Delta] && (6)\\
&= \mathbb{E}_{z\sim\mathcal{N}(0,\sigma^2)}\left[\Pr[\wedge_{i=1}^{d-1} Z_i < z + q'_d - q'_i + 2\Delta]\right] && (7)\\
&= \mathbb{E}_{z\sim\mathcal{N}(0,\sigma^2)}\left[\prod_{i=1}^{d-1}\Pr[Z_i < z + q'_d - q'_i + 2\Delta]\right] && (8)\\
&= \mathbb{E}_{z\sim\mathcal{N}(0,\sigma^2)}\left[\prod_{i=1}^{d-1}\Phi\left(\frac{z + q'_d - q'_i + 2\Delta}{\sigma}\right)\right] && (9)\\
&= \mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{d-1}\Phi\left(z + \frac{q'_d - q'_i + 2\Delta}{\sigma}\right)\right] \quad . && (10)
\end{aligned}$$

Therefore, writing $y_i = q'_i - q'_d$, we get

$$\frac{\Pr[M(D)=o]}{\Pr[M(D')=o]} \leq \frac{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{d-1}\Phi\left(z-\frac{y_i-2\Delta}{\sigma}\right)\right]}{\mathbb{E}_{z\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{d-1}\Phi\left(z-\frac{y_i}{\sigma}\right)\right]} \quad . \tag{11}$$

Note that a priori we have $q_i, q'_i \in [a, b]$ for all $i \in [d]$. However, for the above inequality to hold we set $q_i - q'_i = -\Delta$ for $i < d$ and $q_d - q'_d = \Delta$. These imply $q'_i = q_i + \Delta \in [a + \Delta, b]$ and $q'_d = q'_d - \Delta \in [a, b - \Delta]$, so $y_i = q'_i - q'_d \in [a - b + 2\Delta, b - a]$ for $i \in [d - 1]$. □

Given the result above, all that remains is to identify where the supremum over the $y_i$ is attained in Equation (2). To that end we will show that the ratio is non-decreasing in each of the $y_i$. The following Hôpital-like monotonicity rule and property of ratios of moment generation functions will be useful.

**Lemma 14** ((Pinelis, 2002; Anderson et al., 1993)). *Let $-\infty \leq a < b \leq \infty$ and let $f, g : [a, b] \to \mathbb{R}$ be continuous differentiable functions on $(a, b)$, with $f(a) = g(a) = 0$ or $f(b) = g(b) = 0$, and $g'(x) \neq 0$ for $x \in (a, b)$. If $f'/g'$ is non-decreasing in $(a, b)$, then so is $f/g$.*

We say that a random variable $Z$ has *tails majorized by an exponential decay* if the cumulative distribution function $F_Z$ of $Z$ is such that there exist positive constants $c_1, c_2$ satisfying $F_Z(z) = O(e^{c_1 z})$ for $z \to -\infty$ and $1 - F_Z(z) = O(e^{-c_2 z})$ for $z \to \infty$. This is a well-known sufficient condition for the cumulant generating function $K_Z(s) = \log \mathbb{E}\left[e^{sZ}\right]$ to exist.

**Lemma 15.** *Suppose $Z$ is a random variable with tails majorized by an exponential decay. Then for any $t > 0$ the function*

$$R(s) = \frac{\mathbb{E}\left[e^{(t+s)Z}\right]}{\mathbb{E}\left[e^{sZ}\right]} \ , \tag{12}$$

*is non-decreasing for all $s \in \mathbb{R}$.*

*Proof.* Let $\xi(s) = \mathbb{E}\left[e^{sZ}\right]$ be the moment generating function of $Z$. Taking the derivative of $R$ we get:

$$R'(s) = \frac{\xi'(t+s)\xi(s) - \xi(t+s)\xi'(s)}{\xi(s)^2} \ . \tag{13}$$

Thus, $R'(s) \geq 0$ if and only if $\xi'(t+s)\xi(s) - \xi(t+s)\xi'(s) \geq 0$, which is equivalent to

$$\frac{\partial}{\partial s} \log \xi(t+s) = \frac{\xi'(t+s)}{\xi(t+s)} \geq \frac{\xi'(s)}{\xi(s)} = \frac{\partial}{\partial s} \log \xi(s) \ . \tag{14}$$

Therefore $R$ is non-decreasing if the derivative of the cumulant generating function $K'_Z(s)$ (which exists by assumption) is non-decreasing. But note that when $K_Z(s)$ exists it is known to be convex and infinitely differentiable, and therefore its derivative is non-decreasing. □

Note that by symmetry of the expression of interest in the $y_i$ it suffices to establish monotonicity with respect to a single variable. Thus we define the following functions:

$$f(x) = \mathbb{E}_{z \sim \mathcal{N}(t,1)} \left[ \Phi(z - x) \prod_{i=1}^{k} \Phi(z - c_i) \right] \ , \tag{15}$$

$$g(x) = \mathbb{E}_{z \sim \mathcal{N}(0,1)} \left[ \Phi(z - x) \prod_{i=1}^{k} \Phi(z - c_i) \right] \ , \tag{16}$$

where $t > 0$ and $c_1, \ldots, c_k \in \mathbb{R}$ are constants.

**Lemma 16.** *The function $F(x) = f(x)/g(x)$ is non-decreasing for all $x \in \mathbb{R}$.*

*Proof.* We are going to apply Lemma 14. First note that since $\lim_{x\to\infty} \Phi(z - x) = 0$ we have $\lim_{x\to\infty} f(x) = \lim_{x\to\infty} g(x) = 0$. Next we observe that, writing $\phi(u) = e^{-\frac{u^2}{2}}/\sqrt{2\pi}$ for the density of a standard Gaussian random variable, we have

$$\frac{\partial}{\partial x}\Phi(z - x) = \frac{\partial}{\partial x} \int_{-\infty}^{z-x} \phi(u)du = -\phi(z - x) \ . \tag{17}$$

Thus, when computing $f'(x)$ we will obtain a term of the form $\phi(z - t)\phi(z - x)$, which can be simplified to

$$\phi(z - t)\phi(z - x) = \frac{1}{2\pi} e^{-\frac{(z-t)^2}{2}} e^{-\frac{(z-x)^2}{2}} = \frac{1}{2\pi} e^{-z^2} e^{z(x+t)} e^{-\frac{t^2+x^2}{2}} \ . \tag{18}$$

From this we can now show that $f'(x)$ is proportional to the moment generating function of a random variable $Z$ with density $p(z) = e^{-z^2} \prod_{i=1}^{k} \Phi(z - c_i) / (2\pi N)$ where $N$ is a normalizing constant (independent of $x$ and $t$):

$$f'(x) = -\mathbb{E}_{z\sim\mathcal{N}(t,1)} \left[ \phi(z-x) \prod_{i=1}^{k} \Phi(z - c_i) \right] \tag{19}$$

$$= -\int_{-\infty}^{\infty} \phi(z-t)\phi(z-x) \prod_{i=1}^{k} \Phi(z - c_i)\, dz \tag{20}$$

$$= -e^{-\frac{t^2+x^2}{2}} \cdot N \cdot \int_{-\infty}^{\infty} e^{z(x+t)} p(z) dz \tag{21}$$

$$= -e^{-\frac{t^2+x^2}{2}} \cdot N \cdot \mathbb{E}_{Z\sim p} \left[ e^{(x+t)Z} \right] \ . \tag{22}$$

A similar derivation also yields the following expression for $g'(x)$:

$$g'(x) = -e^{-\frac{x^2}{2}} \cdot N \cdot \mathbb{E}_{Z\sim p} \left[ e^{xZ} \right] \ . \tag{23}$$

Putting these two derivations together we obtain the following expression for the test function in Lemma 14:

$$H(x) := \frac{f'(x)}{g'(x)} = \frac{e^{-\frac{t^2}{2}} \mathbb{E}_{Z\sim p} \left[ e^{(x+t)Z} \right]}{\mathbb{E}_{Z\sim p} \left[ e^{xZ} \right]} \ . \tag{24}$$

Noting that the distribution with density $p(z)$ satisfies the condition of Lemma 15 we see that $H$ is non-decreasing and therefore $F$ is non-decreasing. □

With all these ingredients in place we see that Theorem 12 follows by using Lemma 16 to show that the supremum for each $y_i$ in Lemma 13 is attained at $y_i = b - a$.

### A.2 Pure DP Above Threshold

**Algorithm 3** Gaussian Above Threshold (Zhu and Wang, 2020)

**input:** dataset $D$; noise parameters $\sigma_X, \sigma_Z$; a stream of queries $q_1, q_2, \ldots$; threshold $\rho$.
$\hat{\rho} = \rho + \mathcal{N}(0, \sigma_X^2)$
**for** $t = 1, 2, \ldots$ **do**
  $\hat{q}_t = q_t(D) + \mathcal{N}(0, \sigma_Z^2)$
  **if** $\hat{q}_t \geq \hat{\rho}$ **then**
    Output $x_t = \top$ and HALT
  **else**
    Output $x_t = \bot$
  **end**
**end**

**Theorem 17** (Pure Ex-post Gaussian Above Threshold)**.** *Given a stream $q_1, q_2, \ldots : \mathcal{X} \to [a, b]$ with global sensitivity $\Delta$, the Gaussian Above Threshold mechanism (Algorithm 3) satisfies $\epsilon_{post}$-ex-post-DP with*

$$\epsilon_{post}(\{\bot^{t-1}\top\}) = \frac{\mathbb{E}_{x\sim\mathcal{N}(0,1)} \left[ \Phi\left( \frac{\sigma_X x + \rho - (b-a) + \Delta}{\sigma_Z} \right)^{(t-1)} \Phi\left( \frac{-\sigma_X x - \rho + a + \Delta}{\sigma_Z} \right) \right]}{\mathbb{E}_{x\sim\mathcal{N}(0,1)} \left[ \Phi\left( \frac{\sigma_X x + \rho - (b-a)}{\sigma_Z} \right)^{(t-1)} \Phi\left( \frac{-\sigma_X x - \rho + a}{\sigma_Z} \right) \right]} \ . \tag{25}$$

Overall, our proof follows a similar strategy to the proof in the previous section.

**Lemma 18.** *Let $M$ be the Gaussian Above Threshold mechanism, with public threshold $\rho \geq 0$, threshold noise of standard deviation $\sigma_X$, and query noise of standard deviation $\sigma_Z$ applied to a stream of bounded queries $q_1, q_2, \ldots : \mathcal{X} \to [a, b]$. Then for any stopping time $t \geq 1$ we have*

$$\sup_{D \simeq D'} \frac{\Pr[M(D) = t]}{\Pr[M(D') = t]} \leq \sup_{y_1,\ldots,y_{t-1} \in [a+\Delta, b], y_t \in [a, b-\Delta]} \frac{\mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{t-1} \Phi\left(\frac{\sigma_X x + \rho - y_i + \Delta}{\sigma_Z}\right) \cdot \Phi\left(\frac{-\sigma_X x - \rho + y_t + \Delta}{\sigma_Z}\right)\right]}{\mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{t-1} \Phi\left(\frac{\sigma_X x + \rho - y_i}{\sigma_Z}\right) \cdot \Phi\left(\frac{-\sigma_X x - \rho + y_t}{\sigma_Z}\right)\right]} \,. \tag{26}$$

*Proof.* Fix a pair of datasets $D \simeq D'$. Like before, we let $q_t = q_t(D)$ and $q'_t = q_t(D')$; they satisfy $|q_t - q'_t| \leq \Delta$. First of all we bound the probability that the mechanism on input $D$ stops at time $t$ as follows:

$$\begin{aligned}
\Pr[M(D) = t] &= \Pr[\wedge_{i=1}^{t-1}\{q_i + Z_i < X + \rho\} \wedge \{X + \rho < q_t + Z_t\}] && (27)\\
&= \Pr[\wedge_{i=1}^{t-1}\{Z_i < X + \rho - q_i\} \wedge \{X + \rho - q_t < Z_t\}] && (28)\\
&= \Pr[\wedge_{i=1}^{t-1}\{Z_i < X + \rho - q'_i - (q_i - q'_i)\} \wedge \{X + \rho - q'_t - (q_t - q'_t) < Z_t\}] && (29)\\
&\leq \Pr[\wedge_{i=1}^{t-1}\{Z_i < X + \rho - q'_i + \Delta\} \wedge \{X + \rho - q'_t - \Delta < Z_t\}] && (30)\\
&= \mathbb{E}_{x\sim\mathcal{N}(0,\sigma_X^2)}\left[\Pr[\wedge_{i=1}^{t-1}\{Z_i < x + \rho - q'_i + \Delta\} \wedge \{x + \rho - q'_t - \Delta < Z_t\}]\right] && (31)\\
&= \mathbb{E}_{x\sim\mathcal{N}(0,\sigma_X^2)}\left[\prod_{i=1}^{t-1} \Pr[Z_i < x + \rho - q'_i + \Delta] \cdot \Pr[x + \rho - q'_t - \Delta < Z_t]\right] && (32)\\
&= \mathbb{E}_{x\sim\mathcal{N}(0,\sigma_X^2)}\left[\prod_{i=1}^{t-1} \Pr[Z_i < x + \rho - q'_i + \Delta] \cdot \Pr[-Z_t < -x - \rho + q'_t + \Delta]\right] && (33)\\
&= \mathbb{E}_{x\sim\mathcal{N}(0,\sigma_X^2)}\left[\prod_{i=1}^{t-1} \Pr[Z_i < x + \rho - q'_i + \Delta] \cdot \Pr[Z_t < -x - \rho + q'_t + \Delta]\right] && (34)\\
&= \mathbb{E}_{x\sim\mathcal{N}(0,\sigma_X^2)}\left[\prod_{i=1}^{t-1} \Phi\left(\frac{x + \rho - q'_i + \Delta}{\sigma_Z}\right) \cdot \Phi\left(\frac{-x - \rho + q'_t + \Delta}{\sigma_Z}\right)\right] && (35)\\
&= \mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{t-1} \Phi\left(\frac{\sigma_X x + \rho - q'_i + \Delta}{\sigma_Z}\right) \cdot \Phi\left(\frac{-\sigma_X x - \rho + q'_t + \Delta}{\sigma_Z}\right)\right] \,. && (36)
\end{aligned}$$

A similar derivation also yields:

$$\Pr[M(D') = t] = \mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{t-1} \Phi\left(\frac{\sigma_X x + \rho - q'_i}{\sigma_Z}\right) \cdot \Phi\left(\frac{-\sigma_X x - \rho + q'_t}{\sigma_Z}\right)\right] \,. \tag{37}$$

Thus, the ratio of probabilities of $M$ stopping at time $t$ on $D$ and $D'$ can be bounded as:

$$\frac{\Pr[M(D) = t]}{\Pr[M(D') = t]} \leq \frac{\mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{t-1} \Phi\left(\frac{\sigma_X x + \rho - q'_i + \Delta}{\sigma_Z}\right) \cdot \Phi\left(\frac{-\sigma_X x - \rho + q'_t + \Delta}{\sigma_Z}\right)\right]}{\mathbb{E}_{x\sim\mathcal{N}(0,1)}\left[\prod_{i=1}^{t-1} \Phi\left(\frac{\sigma_X x + \rho - q'_i}{\sigma_Z}\right) \cdot \Phi\left(\frac{-\sigma_X x - \rho + q'_t}{\sigma_Z}\right)\right]} \,. \tag{38}$$

Now note that in the bound we set $q_i - q'_i = -\Delta$ for $i < t$ and $q_t - q'_t = \Delta$. Since all the queries are bounded in $[a, b]$ we have that $y_i = q'_i = q_i + \Delta \in [a + \Delta, b]$ for $i < t$ and $y_t = q'_t = q_t - \Delta \in [a, b - \Delta]$. Taking the supremum over these ranges completes the proof. □

To conclude, we show that the supremum in Lemma 18 is attained at $y_1 = \cdots = y_{t-1} = b$ and $y_t = a$. This follows from observing that the ratio is increasing in $y_1, \ldots, y_{t-1}$ and decreasing in $y_t$, which follows from the same argument used in Lemma 16.

### A.3 Filtered Self-Reporting Composition

**Theorem 19.** *Suppose the mechanisms $M_1, M_2, \ldots$ provided to Algorithm 4 are such $M_t$ is $(\epsilon_{\max,t}, \delta)$-pDP and $\epsilon_{post,t}$-ex-post-DP for all $t$. Then Algorithm 4 satisfies $(\epsilon, \delta)$-DP.*

**Algorithm 4** Filtered Composition, Ex-Post Privacy

**input:** dataset $D$; privacy budget $\epsilon$; a stream of adaptive mechanisms $M_1, M_2, \ldots$; a stream of values $\epsilon_{\max,1}, \epsilon_{\max,2}, \ldots$; a stream of functions $\epsilon_{\text{post},1}, \epsilon_{\text{post},2}, \ldots$

**for** $t = 1, 2, \ldots$ **do**
  **if** $\sum_{i=1}^{t-1} \epsilon_i + \epsilon_{\max,t} \geq \epsilon$ **then**
    HALT
  **else**
    $o_t = M_t(D; o_{1:t-1})$
    $\epsilon_t = \epsilon_{\text{post},t}(o_t)$
    Release $o_t$
  **end**
**end**

*Proof.* We will prove that the mechanism satisfies $(\epsilon, \delta)$-pDP and therefore also $(\epsilon, \delta)$-DP. Let $M$ denote the mechanism and fix a pair of neighboring datasets $D$ and $D'$. Let $o = o_1, \ldots, o_T \sim M(D)$ be sampled from the mechanism (here $T$ is a random stopping time) and consider the privacy loss random variable

$$\begin{aligned}\mathcal{L}(o) &= \log \frac{\Pr[M(D) = o]}{\Pr[M(D') = o]} \\ &= \sum_{i=1}^{T} \log \frac{\Pr[M_i(D; o_{1:i-1}) = o_i]}{\Pr[M_i(D'; o_{1:i-1}) = o_i]} + \log \frac{\Pr[M(D) \text{ halts after outputting } o_{1:T}]}{\Pr[M(D') \text{ halts after outputting } o_{1:T}]} \ .\end{aligned}$$

First of all we observe that given the outputs up to a certain time step $t$, the stopping condition is deterministic and independent of the dataset. Thus, the last term above vanishes. Furthermore, since each of the mechanisms $M_i$ satisfies $\epsilon_{\text{post},i}$-ex-post-DP, for all $i \in [T-1]$ we have

$$\log \frac{\Pr[M_i(D; o_{1:i-1}) = o_i]}{\Pr[M_i(D'; o_{1:i-1}) = o_i]} \leq \epsilon_{\text{post},i}(o_i) \ .$$

In addition, since $M_T$ is $(\epsilon_{\max,T}, \delta)$-pDP, we also have

$$\Pr\left[\log \frac{\Pr[M_T(D; o_{1:T-1}) = o_T]}{\Pr[M_T(D'; o_{1:T-1}) = o_T]} > \epsilon_{\max,T}\right] \leq \delta \ .$$

Now, since the stopping rule guarantees that $\sum_{i=1}^{T-1} \epsilon_{\text{post},i}(o_i) + \epsilon_{\max,T} \leq \epsilon$, we get

$$\Pr[\mathcal{L}(o) > \epsilon] \leq \Pr\left[\sum_{i=1}^{t-1} \epsilon_{\text{post},i}(o_i) + \log \frac{\Pr[M_T(D; o_{1:T-1}) = o_T]}{\Pr[M_T(D'; o_{1:T-1}) = o_T]} > \epsilon\right] < \delta \ .$$

□

# B PURE DP OFFLINE SELECTION MECHANISMS

In Fig. 6 we illustrate the variance observed when taking a Monte-Carlo estimate over a range of values $\sigma$ when computing our Pure DP Gaussian Report Noisy Max bound.

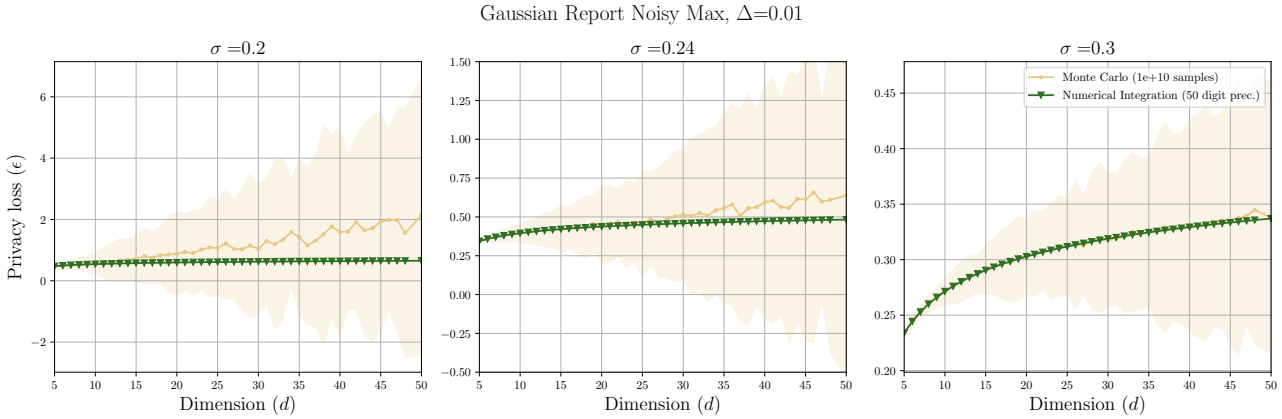


Figure 6: Comparison of numerical integration (in green) and Monte Carlo estimation (in beige, with 100 trials, each with 10B samples) are reported. Note Monte Carlo sampling methods are much slower than numerical methods and require greater numbers of samples as the dimension increases. We observe higher variance in the privacy loss estimate as dimension increases. The shaded region represents the standard deviation.

## B.1 Numerical evaluation compared to post-processing bounds

We study how privacy loss changes as noise (and privacy) are increased in Fig. 7 for Gaussian Report Noisy Max. The standard deviation of the query noise, $\sigma$, and the number of queries (dimension $d$) on the privacy loss estimate. We note a slow increase in the privacy loss as the number of queries increases, in particular when compared to the ex-ante analysis with $\Delta_q = \Delta\sqrt{d}$.

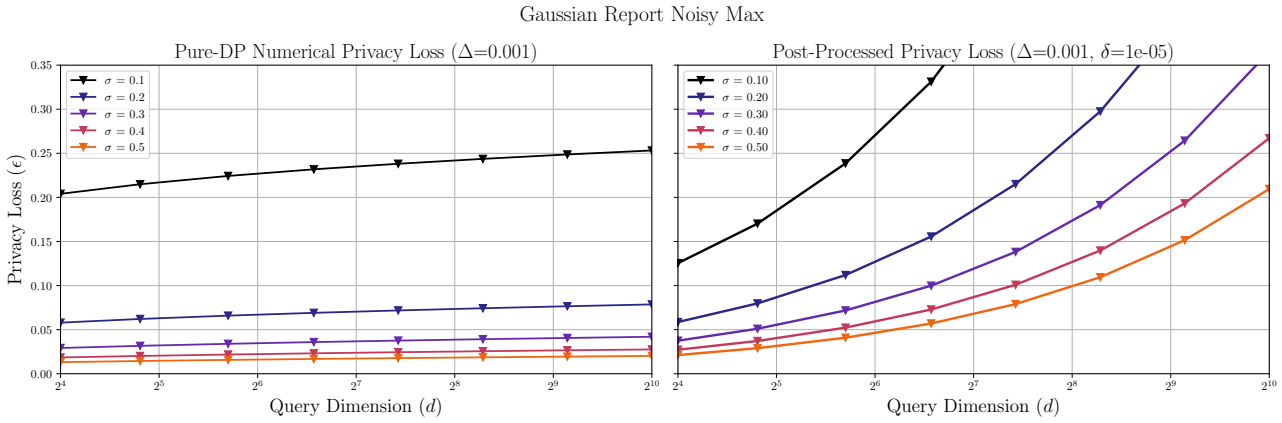


Figure 7: Left: numerical evaluation of the privacy loss for Gaussian Report Noisy Max. Privacy loss increase is very slowly as the number of queries increases. Right: we calculate privacy loss with a classic $(\epsilon, \delta)$ post-processing bound.

### B.2 Where pure DP offers tighter accounting for Report Noisy Max

To assess whether there are any utility gains from performing a numerical evaluation of the privacy loss, we produce a heatmap of the difference reported between the standard post-processing bound under RDP and our method in Fig. 8. There exists a smooth region where the privacy accounting difference is significant, particularly in the the high privacy, high query setting. As the queries become less sensitive ($\Delta \to 0$), this region becomes more pronounced.

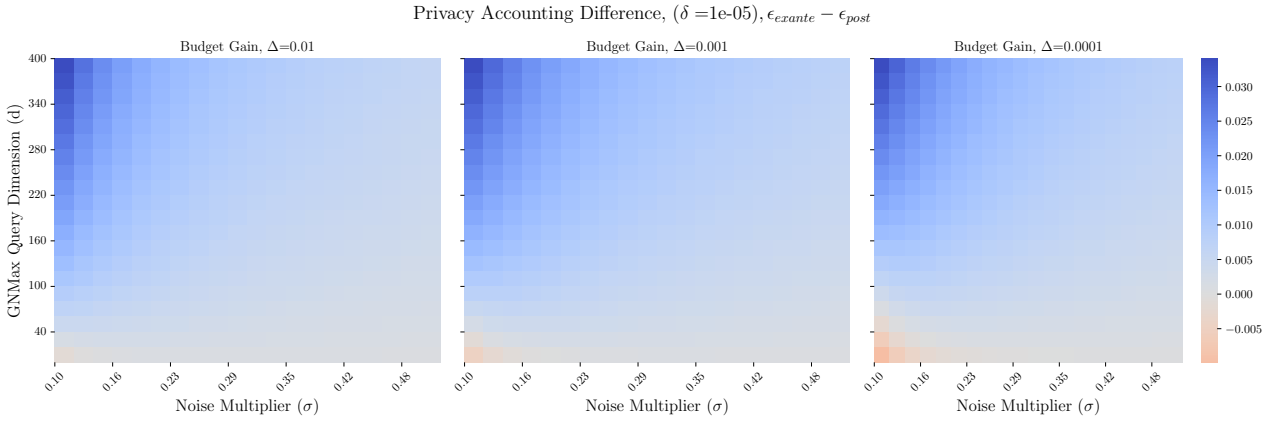


Figure 8: Comparison of the reported privacy loss difference between the standard Gaussian Report Noisy Max analysis and an ex-post evaluation for different levels of privacy ($\sigma$) and number of queries ($d$). For large $d$ and small $\sigma$, we observe the greatest difference in reported privacy loss.

### B.3 Offline Selection with Gaussian Report Noisy Max

Each accuracy/privacy loss plot reports the mean value over 1,000 trials, with the shaded region covering the standard error.

In our experiments, we normalize the dataset to values between zero and one. The shaded region represents the standard error across several runs. The mean value is represented in as a line. A classic example of this sort of problem is query selection over a long time horizon. Our experiments center on energy and mobility datasets, which have been known to leak privacy (Narayanan and Shmatikov, 2008). In both instances, user behavior, such as whether power consumption deviates from historical levels, can be joined with other datasets to infer changes to a person's socio-economic situation.

**Dataset Pre-Processing** The datasets we consider are temporal event logs over a fixed period of time with a regular reporting interval (e.g. hour, day, month). In each instance, we re-scale the reported values to be between zero and one. These transformations do not affects the task result since we are interesting in reporting a discrete value. In the offline selection task, we wish to find the time step whose reported value is largest, and in the online selection task, the goal is to produce a binary vector indicating which step exceeds a fixed threshold.

**Utility** We measure the accuracy of Gaussian Report Noisy Max by comparing the distance between the true answer and the noisy (private) answer step. Let $q_i^*$ be the true query answer and $\hat{q}_i$ be the answer selected by Report Noisy Max. Then accuracy $\mathrm{Acc}_\sigma$ is defined as

$$\mathrm{Acc}_\sigma(D) = 1 - \mathbb{E}\left[\frac{|q_{i^*}(D) - q_{\hat{i}}(D)|}{q_{\max}}\right] \quad . \tag{39}$$

where $q_{\max}$ is the maximum value of any query. In our experiments, $q_{\max} = 1$. In this construction, perfect utility occurs when $\mathrm{Acc}_\sigma(D) = 1$. This is a reasonable measure of utility since we wish to measure how close on average we are to the max query value.

**NEAR Energy Dataset.** The NEAR Energy Dataset (Ratnam et al., 2017) includes the annual energy consumption of 300 customers, we formulate the following question: Which day was consumption highest?

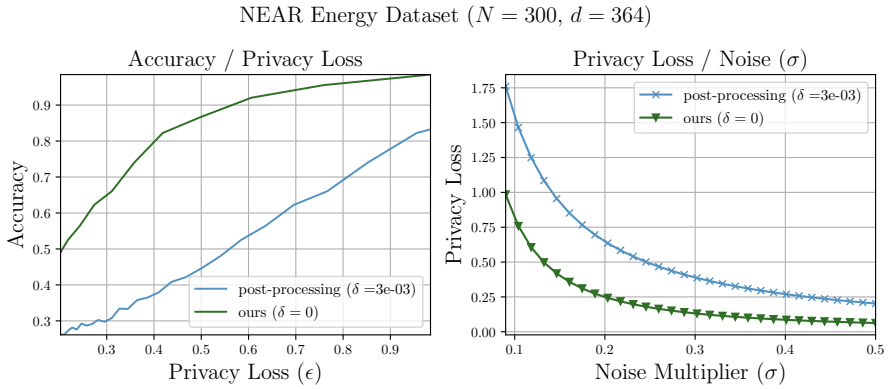


Figure 9: Privacy / utility when performing Gaussian Report Noisy Max on a query over $d = 364$. 90% Accuracy is attainable with our method whereas the same would only be possible with over double the privacy budget.

**London Energy Dataset** Here we follow the same analysis as in the case of the NEAR dataset (Greater London Authority, 2012), however the dataset comprises of $N = 5,564$ customers over $d = 829$ days. The larger number of queries *increases* the privacy cost, however this is balanced by a *decrease* in individual contribution to each query due to each query having $\Delta = 1/N$.

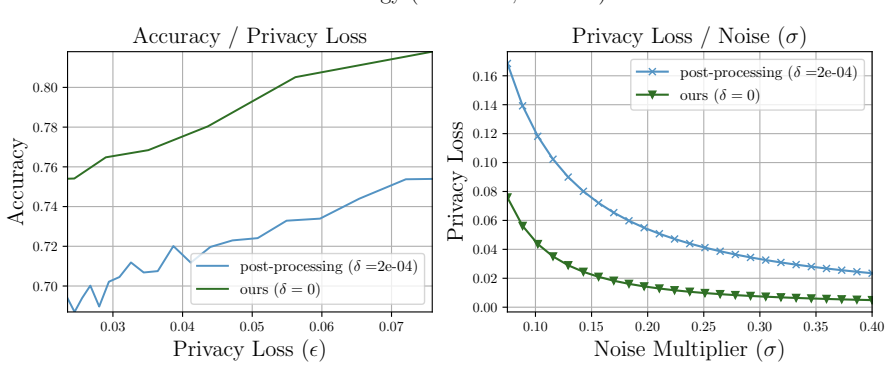


Figure 10: Privacy / utility when performing Gaussian Report Noisy Max on a query over $d = 829$ with the London Energy Dataset. As $N$ and $d$ increase, the difference in reported privacy loss widens.

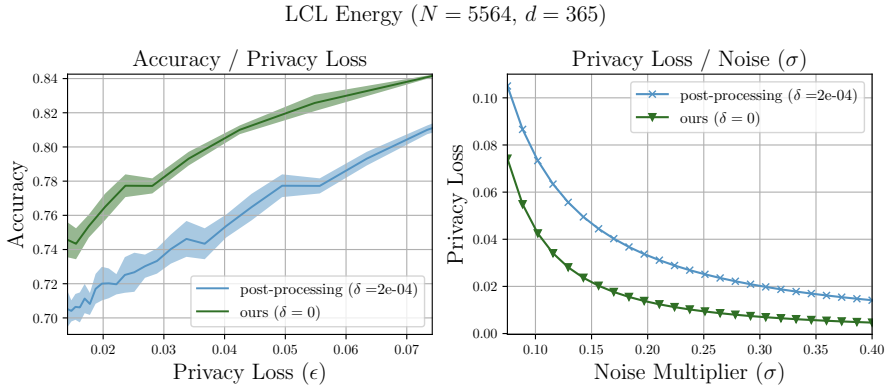


Figure 11: Privacy / utility when performing Gaussian Report Noisy Max on a query over $d = 365$ with the London Energy Dataset.

We remark that there are other offline selection mechanisms that guarantee pure DP without Gaussian noise. The exponential mechanism, which calibrates an exponential distribution with respect to a utility function, is commonly used solution to the offline query selection problem (Dwork and Roth, 2014). Despite a relatively simple implementation, the exponential mechanism requires that a sum over all query values be computed to fix the probability distribution. Permute-and-flip is a drop-in replacement for the exponential mechanism which has been shown to provide some utility benefits, but has the same calibration requirement (McKenna and Sheldon, 2020). In both instances the global sensitivity is no longer dependent on the number of queries since the maximum contribution for each user is taken with respect to each query individually.

In Fig. 12 we measure compare our method to pure DP baselines with other noise distributions. The global sensitivity for these exponential-based mechanisms is $\Delta := 1/N$, where $N$ is the number of users, compared to our bound, which depends on $d$ queries. One could make apply the same post-processing argument for Gaussian Report Noisy Max with noise drawn from the Laplace distribution however, this approach was omitted from our baselines since it did not appear competitive in our setting. Finally, we remark that our method may be easier to calibrate since each noisy query can be computed separately.

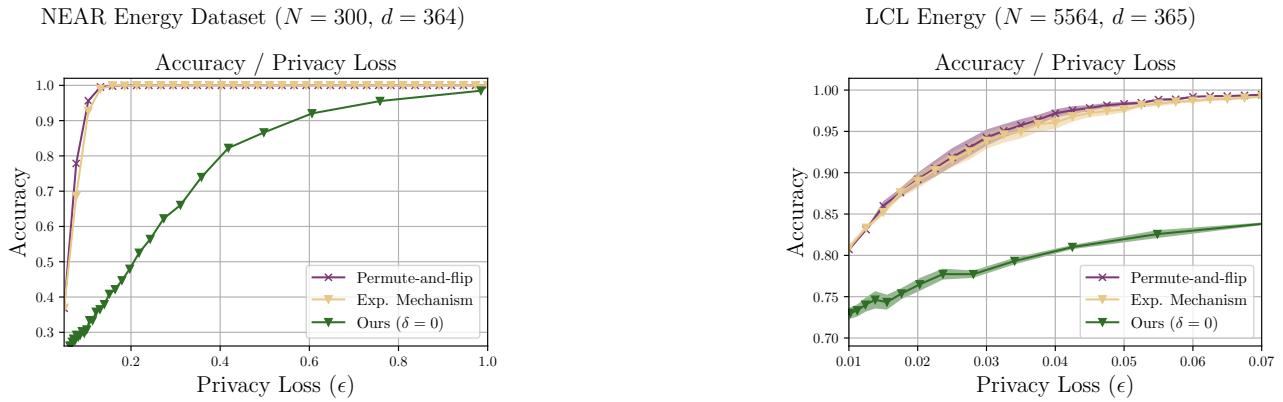


(a) Comparison for $d = 364$ on the NEAR Dataset.

(b) Comparison for $d = 365$ on the LCL Energy Dataset.

Figure 12: Comparison between our method and two classical offline selection mechanisms.

**UCI Bikes Dataset** The UCI Bike Sharing Dataset captures the utilization of shared bikes in a geographic area over the course of a year (Fanaee-T and Gama, 2013). We apply the same definition of accuracy and report on the day with peak consumption. Since we do not know the upper bound on registered customers, we take the maximum (6,946 users) and assume that this is a public value. We note that our analysis is still worst-case, in that a user is assumed to be contributing to each daily (normalized) count query. We limit our Report Noisy Max query to 365 days.

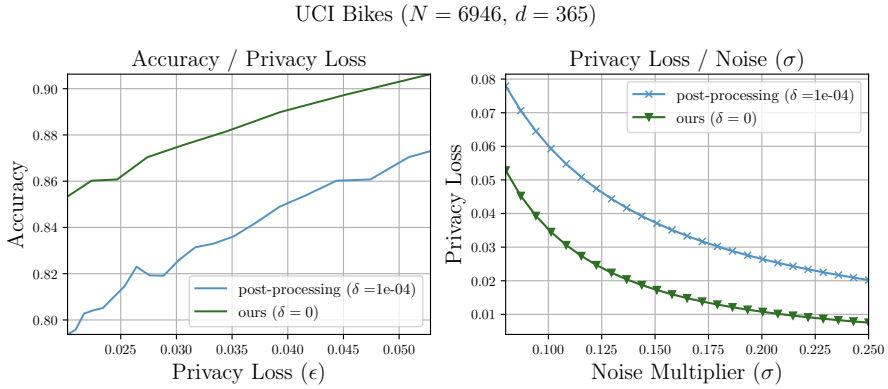


Figure 13: Privacy / utility when performing Gaussian Report Noisy Max on a query over $d = 365$ with the UCI Bike Sharing Dataset.

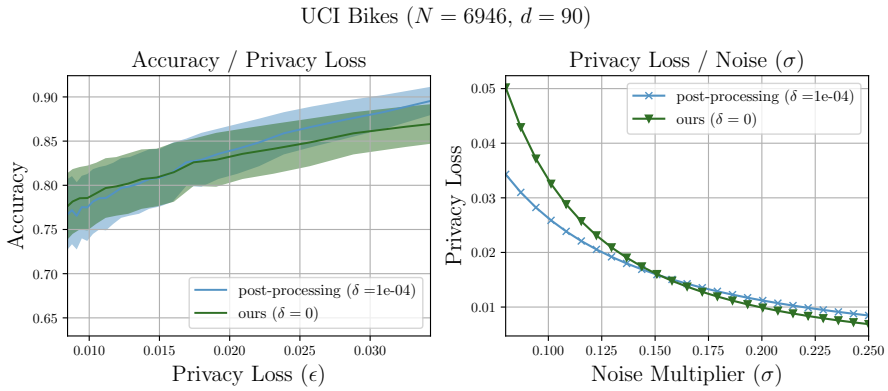


Figure 14: Privacy / utility when performing Gaussian Report Noisy Max on a query over $d = 90$ with the UCI Bike Sharing Dataset. Note that as the number of queries, $d$, decreases, baseline methods become more competitive.

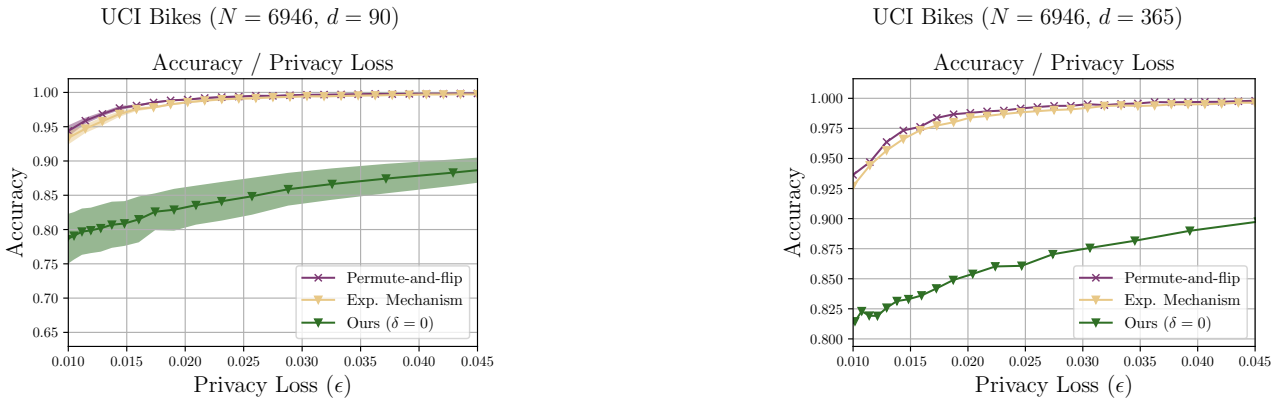


(a) Comparison for $d = 90$.

(b) Comparison for $d = 365$.

Figure 15: Comparison between our method and two classical offline selection mechanisms.

As shown in Fig. 15, the Exponential Mechanism and Permute-and-flip offer a competitive baseline. Fig. 16 provides a comparison of our method with Laplace Report Noisy Max. We observe a clear improvement in the pure DP setting using Gaussian noise and our bounds.

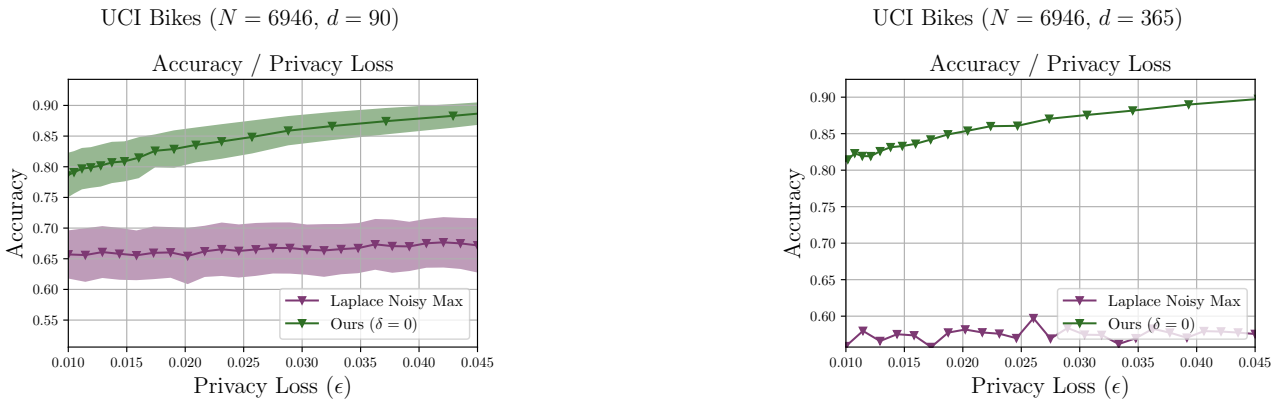


(a) Comparison for $d = 90$.

(b) Comparison for $d = 365$.

Figure 16: Comparison between our method and the Laplace Report Noisy Max on the UCI Bike Sharing Dataset.

# C COMPARISON WITH PRIVACY FILTERS

Our baseline for comparison are recent results by (Rogers et al., 2023), who propose a privacy filter which introduces an additive bound in both $\epsilon$ and $\delta$, thereby creating a result similar to advanced composition. Their construction does not rely on computing some $\epsilon_{\max}$.

**Theorem 20** ($(\epsilon, \delta)$-DP Filters (Whitehouse et al., 2023))**.** *Suppose $(M_t) \geq 1$ is a sequence of mechanisms such that, for any $t \geq 1$, $M_t$ is $(\epsilon_t, \delta_t)$-differentially private conditioned on $M_{1:t-1}$. Let $\epsilon > 0$, and $\delta = \delta' + \delta''$ be max privacy parameters s.t. $\delta' > 0, \delta'' \geq 0$. Let the stopping time function $N : \mathbb{R}_{\geq 0}^{\infty} \times \mathbb{R}_{\geq 0}^{\infty} \to \mathbb{N}$ be given by,*

$$N((\epsilon_t)_{t\geq 1}, (\delta_t)_{t\geq 1}) = \inf \left\{ n : \epsilon < \sqrt{2 \log\left(\frac{1}{\delta'}\right) \sum_{m \leq t+1} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq t+1} \epsilon_m^2 \quad or \quad \delta'' < \sum_{m \leq t+1} \delta_m \right\} . \tag{40}$$

*Then $M_{1:N(\cdot)}(\cdot) : \mathcal{X} \to \mathcal{O}^{\infty}$ is $(\epsilon, \delta)$-DP.*

To compare Filtered Self-Reporting Composition, we apply the following privacy filter, (Definition 21) with Algorithm 5. The bound we compute for Gaussian Above Threshold comes from (Zhu and Wang, 2020).

**Definition 21** (DP Privacy Filter (Rogers et al., 2016))**.** Let $N(\cdot, \cdot)$ be a $(\epsilon, \delta)$-DP Filter. Then the DP Composition Privacy Filter $\mathcal{F}(\cdot)$ is given by

$$\mathcal{F}_{\epsilon,\delta}((\epsilon_t)_{t\geq 1}, (\delta_t)_{t\geq 1}) = \begin{cases} \text{HALT} & \text{if } N((\epsilon_t)_{t\geq 1}, (\delta_t)_{t\geq 1}) > t \\ \text{CONT.} & \text{otherwise} \end{cases} . \tag{41}$$

**Algorithm 5** Composition with Privacy Filter

**input:** Dataset $D$, Privacy budget $\epsilon_{\max}$, $\delta_{\max}$, $(\epsilon_t, \delta_t)$-differentially private mechanisms $M_t$, for $t = 1, \ldots, T$, Privacy Filter $\mathcal{F}_{\epsilon,\delta}$.

**for** $t = 1, \ldots, T$ **do**
  $o_t = M_t(D)$
  **if** $\mathcal{F}_{\epsilon_{max},\delta_{max}}((\epsilon_t)_{t\geq 1}, (\delta)_{t\geq 1}) = HALT$ **then**
    $v_t = \perp$
  **else**
    $v_t = o_t$
  **end**
**end**
**return** $(v_1, \ldots, v_T)$

# D PURE DP ONLINE SELECTION: ADDITIONAL EXPERIMENTS

**UCI Bikes Dataset** We use the same UCI dataset as before, however this time configure Above Threshold to HALT when bike sharing exceeds a certain threshold between zero and one. In Fig. 17 we plot a range of calibrations, and in Fig. 18 we show the privacy spend over time.

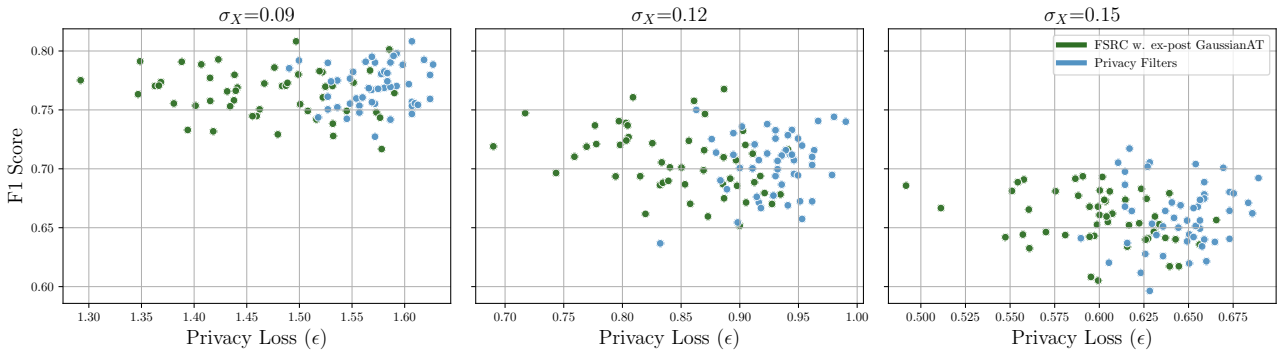


Figure 17: Scatter plot indicating the accuracy and final privacy spend over a range of noise multipliers for thresholds with the UCI Bikes dataset. Final privacy loss ($\epsilon$) is reported for Filtered Self-Reporting Composition (green). Threshold noise, $\sigma_X$, evaluated in the range of $[0.09, 0.15]$. We note a clear separation in privacy accounting over a range of noise multipliers.

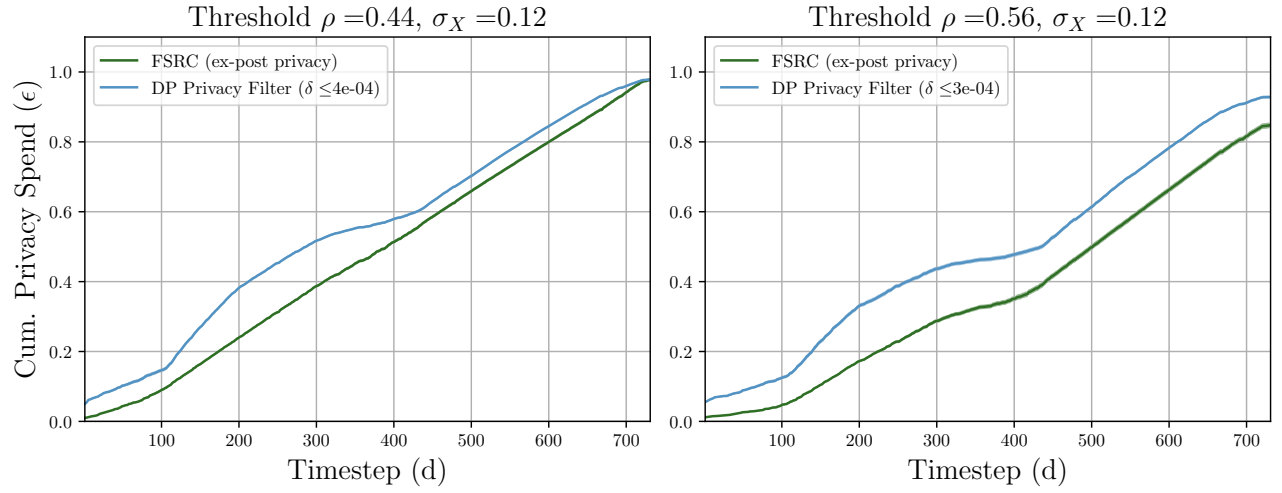


Figure 18: Privacy spend comparison between Gaussian Above Threshold and Filtered Self-Reporting Composition and Privacy Filters with Gaussian Above Threshold on the UCI Bikes dataset.

**London Energy Dataset.** We plot Gaussian Above Threshold under the same utility metric on the LCL London energy dataset, with 5,564 customers. In Fig. 19 we plot a range of calibrations. As the threshold decreases, we witness more queries released and therefore a greater privacy spend. we show the effect over time in Fig. 20.

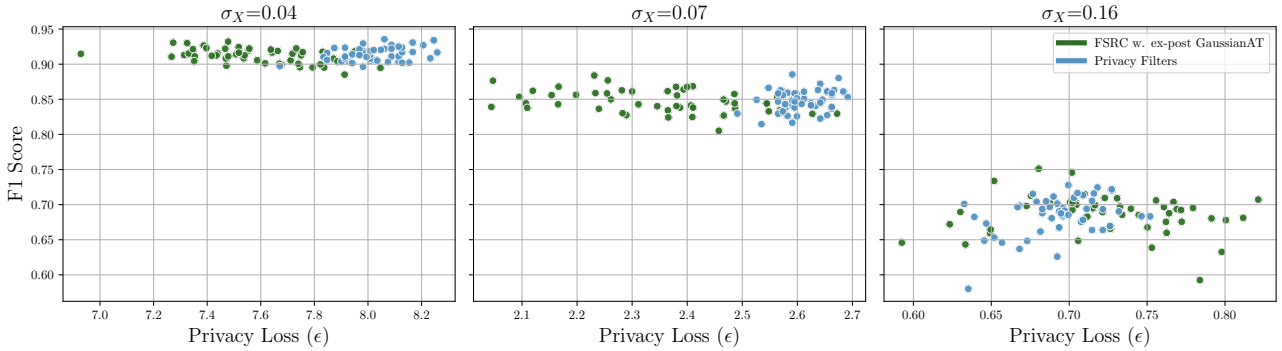


Figure 19: Scatter plot indicating the accuracy and final privacy spend over a range of noise multipliers for thresholds with the LCL London dataset. Final privacy loss ($\epsilon$) is reported for FSRC (green). Threshold noise, $\sigma_X$, was evaluated in the range of $[0.04, 0.16]$. With this dataset, our accounting method provides benefits when $\sigma_X > 0.06$.

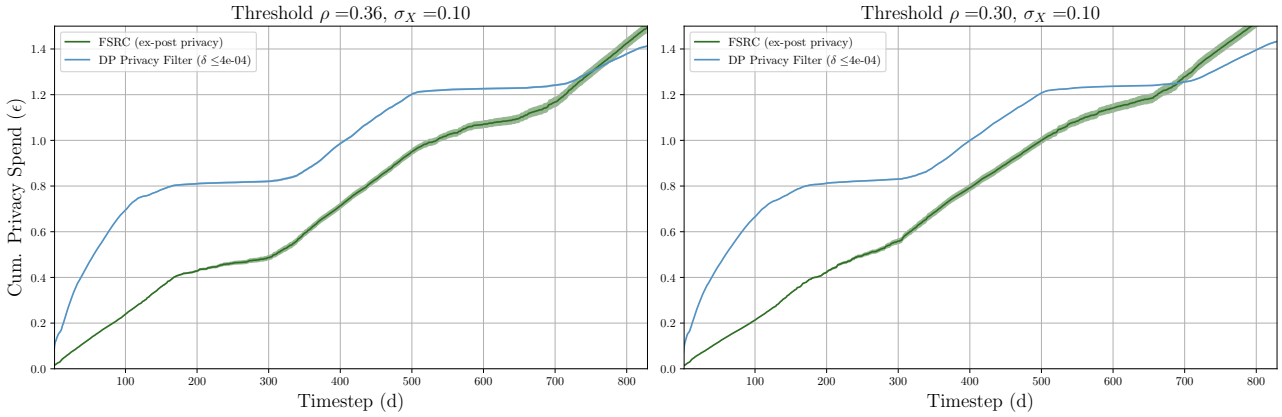


Figure 20: Privacy spend comparison between Gaussian Above Threshold and FSRC and Privacy Filters with Gaussian Above Threshold on the LCL London Energy dataset.