## A Other properties of differential privacy and RDP

RDP inherits and generalizes the information-theoretic properties of DP.
Lemma 15 (Selected Properties of RDP [Mironov, 2017]). If $\mathcal{M}$ obey $\epsilon_{\mathcal{M}}(\cdot)-R D P$, then

1. [Indistinguishability] For any measurable set $S \subset \operatorname{Range}(\mathcal{M})$, and any neighboring $D, D^{\prime}$

$$
e^{-\epsilon(\alpha)} \operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right) \in S\right]^{\frac{\alpha}{\alpha-1}} \leq \operatorname{Pr}[\mathcal{M}(D) \in S] \leq e^{\epsilon(\alpha)} \operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right) \in S\right]^{\frac{\alpha-1}{\alpha}}
$$

2. [Post-processing] For all function $f, \epsilon_{f \circ \mathcal{M}}(\cdot) \leq \epsilon_{\mathcal{M}}(\cdot)$.
3. [Composition] $\epsilon_{\left(\mathcal{M}_{1}, \mathcal{M}_{2}\right)}(\cdot)=\epsilon_{\mathcal{M}_{1}}(\cdot)+\epsilon_{\mathcal{M}_{2}}(\cdot)$.

This composition rule, together with Lemma 3, often allows for tighter calculations of $(\epsilon, \delta)$-DP for the composed mechanism than directly invoking the strong composition theorem below.
Lemma 16 (Strong composition [Kairouz et al., 2015]). For all $\epsilon, \delta, \delta^{\prime} \geq 0$, the (adaptive) composition of $k(\epsilon, \delta)$-DP mechanisms obey $\left(\epsilon^{\prime}, \delta^{\prime}+k \delta\right)$-DP for $\epsilon^{\prime}=\sqrt{2 k \log \left(1 / \delta^{\prime}\right)} \epsilon+k \frac{e^{\epsilon}-1}{e^{\epsilon}+1}$.

## B Proof of Theorem 8-RDP of the Generalized SVT for $c=1$

For unbounded sequences, the output space of the algorithm is $\left\{\perp^{k} \top \mid k=0,1, \ldots, \infty\right\}$. In the case when $k_{\max }<+\infty$, the output space is $\left\{\perp^{k} T \mid k=0,1, \ldots, k_{\max }-1\right\} \cup\left\{\perp^{k_{\max }}\right\}$. For notation convenience, we replace $\perp^{k_{\text {max }}}$ with $\perp^{k_{\text {max }}} T$, which can be thought of fixing a dummy query at time $k_{\max }+1$ which always outputs $+\infty$ regardless of inputs. In both cases, we can completely describe the output distribution the SVT with a positive random integer $K$. As a result, we will write $K \sim \mathcal{M}(D)$ and $K \sim \mathcal{M}\left(D^{\prime}\right)$ without loss of generality.
Also w.l.o.g., we assume thresholds $T_{i}$ are all zero. There are two types of random variables in the algorithm: the threshold noise $\rho$ and the query noise $\nu_{i}$ to each of the $i$ queries, $\left\{\nu_{i}\right\}_{i=1}^{k+1}$. We will use $p_{\rho}(z)$ to denote the probability density of $\rho$, evaluated at $z$, and we will use $p\left(\nu_{i}\right)$ as the pdf of $\nu_{i}$.
The probability of outputting $o$ (or $K=k+1$ ), can be written explicitly as follows:

$$
\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]=\int_{-\infty}^{+\infty} p_{\rho}(z)\left(\prod_{i \leq k} \int_{-\infty}^{z-q_{i}\left(D^{\prime}\right)} p\left(\nu_{i}\right) d \nu_{i}\right) \cdot \int_{z+q_{k+1}\left(D^{\prime}\right)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1} d z
$$

Our goal of is to bound $\mathbb{E}_{o \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left(\frac{\operatorname{Pr}[\mathcal{M}(D)=o]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]}\right)^{\alpha}\right]$ using the RDP functions of $\mathcal{M}_{\rho}$ and $\mathcal{M}_{\nu}$.
The key of the analysis relies on a sequence of fictitious queries $\tilde{q}_{1}, \tilde{q}_{2}, \ldots$ which mirrors the actual sequence of queries $q_{1}, q_{2}, \ldots$ that are adaptively selected. These fictitious queries satisfy for all $i=1,2,3, \ldots$

$$
\tilde{q}_{i}(x)=\left\{\begin{array}{lr}
q_{i}(D)+\triangle, & \text { when } x=D  \tag{4}\\
q_{i}\left(D^{\prime}\right) & \text { otherwise }
\end{array}\right.
$$

The following lemma establishes that we can decompose the problem into one that involves the Renyi-divergence between a distribution induced by these fictitious queries and another distribution induced of the actual queries.
Lemma 17. Consider Algorithm 2 with $c=1$, i.e., the output sequence $o \in\left\{\perp^{k} \top \mid k=0,1, \ldots, \infty\right\}$, then we have

$$
\begin{equation*}
\mathbb{E}_{o \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left(\frac{\operatorname{Pr}[\mathcal{M}(D)=o]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]}\right)^{\alpha}\right] \leq \mathbb{E}_{z \sim p_{\rho}}[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha} \underbrace{\mathbb{E}_{K \sim \mathcal{M}\left(D^{\prime}\right)}\left[\frac{(\operatorname{Pr}[\mathcal{M}(D)=K \mid z, \tilde{\mathcal{Q}}])^{\alpha}}{\left(\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=K \mid z\right]\right)^{\alpha} \mid}\right]}_{\text {denoted by (*) }}] \tag{5}
\end{equation*}
$$

where $K$ is a random variable, denotes the number of $\perp$ plus 1 when the algorithm stops, and the explicit conditioning on $\tilde{\mathcal{Q}}$ indicates that the probability is evaluated by hypothetically running the algorithm on the fictitious queries $\tilde{q}_{1}, \tilde{q}_{2}, \ldots \in \tilde{\mathcal{Q}}$.

Proof of Lemma 17. From the definition of Renyi DP, we have

$$
\begin{equation*}
\mathbb{E}_{o \sim D^{\prime}}\left[\frac{\operatorname{Pr}[\mathcal{M}(D)=o]^{\alpha}}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]^{\alpha}}\right]=\sum_{k=0}^{\infty} \frac{\operatorname{Pr}\left[\mathcal{M}(D)=\perp^{k} T\right]^{\alpha}}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=\perp^{k} T\right]^{\alpha-1}} \tag{6}
\end{equation*}
$$

Without loss of generality, we will replace $o$ with $k$ which measures the number of $\perp \mathrm{s}$ in $o$. By law of total expectation, we can condition on $\rho=z$

$$
\begin{aligned}
& \operatorname{Pr}[\mathcal{M}(D)=k]=\mathbb{E}_{z \sim p_{\rho}}[\operatorname{Pr}[\mathcal{M}(D)=k \mid z]] \\
= & \mathbb{E}_{z \sim p_{\rho}}\left[\prod_{i \leq k} \operatorname{Pr}\left[q_{i}(D)+\nu_{i}<z \mid z\right] \operatorname{Pr}\left[q_{k+1}(D)+\nu_{i} \geq z \mid z\right]\right] \\
= & \int_{-\infty}^{+\infty} p_{\rho}(z)\left(\prod_{i \leq k} \int_{-\infty}^{z-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right) \cdot \int_{z-q_{k+1}(D)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1} d z \\
u:=z+\triangle+\infty & \stackrel{\downarrow}{=} \int_{-\infty}^{+\infty} p_{\rho}(u-\triangle)\left(\prod_{i \leq k} \int_{-\infty}^{u-\triangle-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right) \cdot \int_{u-\Delta-q_{k+1}(D)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1} d u \\
= & \int_{-\infty}^{+\infty} p_{\rho}(u)\left(\frac{p_{\rho}(u-\triangle)}{p_{\rho}(u)}\right)\left(\prod_{i \leq k} \int_{-\infty}^{u-\triangle-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right) \cdot \int_{u-\triangle-q_{k+1}(D)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1} d u \\
= & \mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)\left(\prod_{i \leq k} \int_{-\infty}^{z-\triangle-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right) \cdot \int_{z-\triangle-q_{k+1}(D)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1}\right]
\end{aligned}
$$

where in the last line, we rename the variable $u$ back to $z$.
Substituting the above expression to the definition of RDP and apply Jensen's inequality

$$
\begin{align*}
&(6)=\sum_{k=0}^{\infty} \frac{\mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\Delta)}{p_{\rho}(z)}\right)\left(\prod_{i \leq k} \int_{-\infty}^{z-\triangle-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right) \cdot \int_{z-\triangle-q_{k+1}(D)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1}\right]^{\alpha}}{\mathbb{E}_{z \sim p_{\rho}}\left[\left(\prod_{i \leq k} \int_{-\infty}^{z-q_{i}\left(D^{\prime}\right)} p\left(\nu_{i}\right) d \nu_{i}\right) \cdot \int_{z--q_{k+1}\left(D^{\prime}\right)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1}\right]^{\alpha-1}} \\
& \leq \sum_{k=0}^{\infty} \mathbb{E}_{z \sim p_{\rho}} \frac{\left(\frac{p_{\rho}(z-\Delta)}{p_{\rho}(z)}\left(\prod_{i \leq k} \int_{-\infty}^{z-\triangle-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right) \int_{z-\triangle-q_{k+1}(D)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1}\right)^{\alpha}}{\left(\left(\prod_{i \leq k} \int_{-\infty}^{z-q_{i}\left(D^{\prime}\right)} p\left(\nu_{i}\right) d \nu_{i}\right) \int_{z-q_{k+1}\left(D^{\prime}\right)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1}\right)^{\alpha-1}} \tag{7}
\end{align*}
$$

The inequality applies Jensen's inequality to bivariate function $f(x, y)=x^{\alpha} y^{1-\alpha}$, which is jointly convex on $\mathcal{R}_{+}^{2}$ for $\alpha \in(1,+\infty)$.

Exchange the order of integral variable $z$ and $k$ in (7), we get

$$
\begin{align*}
\text { (7) } & =\mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha} \sum_{k=0}^{\infty} \frac{\left(\left(\prod_{i \leq k} \int_{-\infty}^{z-\triangle-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right) \int_{z-\Delta-q_{k+1}(D)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1}\right)^{\alpha}}{\left(\left(\prod_{i \leq k} \int_{-\infty}^{z-q_{i}\left(D^{\prime}\right)} p\left(\nu_{i}\right) d \nu_{i}\right) \int_{z-q_{k+1}\left(D^{\prime}\right)}^{\infty} p\left(\nu_{k+1}\right) d \nu_{k+1}\right)^{\alpha-1}}\right] \\
& =\mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha} \sum_{k=0}^{\infty} \frac{\left(\prod_{i=1}^{k} \operatorname{Pr}\left[q_{i}(D)+\Delta+\nu_{i}<z \mid z\right] \operatorname{Pr}\left[q_{k+1}(D)+\Delta+\nu_{k+1} \geq z \mid z\right]\right)^{\alpha}}{\left(D^{\prime}\right)}\right]  \tag{8}\\
& =\mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha} \mathbb{E}_{K \sim \mathcal{M}\left(D^{\prime}\right)}\left[\frac{\left(\prod_{i=1}^{K-1} \operatorname{Pr}\left[\tilde{q}_{i}(D)+\nu_{i}<z \mid z\right] \operatorname{Pr}\left[\tilde{q}_{K}(D)+\nu_{K} \geq z \mid z\right]\right)^{\alpha-1}}{\left(\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=K \mid z\right]\right)^{\alpha}}\right] z\right] \\
& =\mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha} \mathbb{E}_{K \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left.\frac{(\operatorname{Pr}[\mathcal{M}(D)=K \mid z, \tilde{\mathcal{Q}}])^{\alpha}}{\left(\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=K \mid z\right]\right)^{\alpha}} \right\rvert\, z\right]\right] \tag{9}
\end{align*}
$$

which completes the proof.
To understand the last step: recall our definition of fictitious query $\tilde{q}_{i}$, which obeys $\tilde{q}_{i}(D)=q_{i}(D)+\triangle$ and $\tilde{q}_{i}(\tilde{D})=q_{i}(\tilde{D})$ for all other dataset $\tilde{D} \neq D$. Observe that the expression in the numerator of (8) actually describes a valid probability distribution of $K$, which says the probability of $\mathcal{M}(D)$ stopping at time $K=k+1$ when the sequence of input is $\tilde{q}_{1}, \ldots, \tilde{q}_{k}+1, \ldots$, i.e.,

$$
\prod_{i=1}^{k} \operatorname{Pr}\left[q_{i}(D)+\Delta+\nu_{i}<z \mid z\right] \operatorname{Pr}\left[q_{k+1}(D)+\Delta+\nu_{i} \geq z \mid z\right]=\operatorname{Pr}_{\mathcal{M}(D)}\left[K=k+1 \mid z, \tilde{q}_{1}, \tilde{q}_{2}, \ldots\right]
$$

The conditioning on the sequence of queries might appear to be new, but recall that all our probabilities are conditioned on a sequence of queries that are chosen from $\mathcal{Q}(\triangle)$ or $\mathcal{Q}_{+}(\triangle)$ to begin with. They are just not written out explicitly. This is an instance, where we actually need to condition on a different set of queries to formally write down this valid probability distribution above.

Remark. The lemma de-convolves the moment of interests into the mixture of conditional moments of another two distributions that can be written down explicitly. The proof is delicate but informative, as it explicitly leveraging the fact that $\mathcal{M}_{\nu}$ is a noise-adding mechanism, so as to argue the implication of the randomization for a different query $\tilde{q}_{i}$ than the one that it seems to be intending for according to the algorithm $q_{i}$. There are several other novel components. We encourage readers to check it out in details.

A remarkable consequence of this lemma is that we can essentially cancel all factors concerning $\perp \mathrm{s}$.

$$
\begin{align*}
(*) & =\mathbb{E}_{K \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left.\frac{\left(\prod_{i=1}^{K-1} \operatorname{Pr}\left[q_{i}(D)+\triangle+\nu_{i}<z \mid z\right] \operatorname{Pr}\left[q_{K}(D)+\Delta+\nu_{i} \geq z \mid z\right]\right)^{\alpha}}{\left(\prod_{i=1}^{K-1} \operatorname{Pr}\left[q_{i}\left(D^{\prime}\right)+\nu_{i}<z \mid z\right] \operatorname{Pr}\left[q_{K}\left(D^{\prime}\right)+\nu_{i} \geq z \mid z\right]\right)^{\alpha}} \right\rvert\, z\right] \\
& \leq \mathbb{E}_{K \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left.\frac{\left(\operatorname{Pr}\left[q_{K}(D)+\triangle+\nu_{i} \geq z \mid z\right]\right)^{\alpha}}{\left(\operatorname{Pr}\left[q_{K}\left(D^{\prime}\right)+\nu_{i} \geq z \mid z\right]\right)^{\alpha}} \right\rvert\, z\right] \tag{10}
\end{align*}
$$

The inequality in the last line uses the fact that $q_{i}$ has a global sensitivity of $\triangle$, which implies that $\operatorname{Pr}\left[q_{i}(D)+\triangle+\nu_{i}<z \mid z\right] \leq \operatorname{Pr}\left[q_{i}\left(D^{\prime}\right)+\nu_{i}<z \mid z\right.$. for all $i$.
Further observe that $\tilde{q}_{K}$ has a sensitivity of $2 \triangle$ since $q_{K}$ has sensitivity $\triangle$. By the property of the noise-adding mechanism $\mathcal{M}_{\nu}$, it obeys $\epsilon_{\nu}(\alpha)$-RDP for all queries having having sensitivity $2 \triangle$. Therefore, if $\epsilon_{\nu}(\infty)<+\infty$, then we can bound (10) with $e^{\alpha \epsilon_{\nu}(\infty)}$. In fact, this bound can be improved slightly if we directly work with (5), which we state as a lemma.
Lemma 18. If $\epsilon_{\nu}(\infty)<+\infty$, then the expression $(*)$ in (5) obeys $(*) \leq e^{(\alpha-1) \epsilon_{\nu}(\infty)}$.
Proof. We use a trick due to [Bun and Steinke, 2016] with some modifications.

First, check that by our trivial bounds

$$
0 \leq \frac{\operatorname{Pr}[\mathcal{M}(D)=K \mid z, \tilde{\mathcal{Q}}]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=K \mid z\right]} \leq e^{\epsilon}
$$

Define random function $A(K)$ supported on $\left\{0, e^{\epsilon}\right\}$ such that $\mathbb{E}[A(K) \mid K]=\frac{\operatorname{Pr}[\mathcal{M}(D)=K \mid z, \tilde{\mathcal{Q}}]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=K \mid z\right]}$.
Note that when $\alpha=1$

$$
\operatorname{Pr}\left[A(K)=e^{\epsilon}\right] \cdot e^{\epsilon}=\mathbb{E}_{K}[\mathbb{E}[A(K) \mid K]]=\mathbb{E}_{K \sim \mathcal{M}\left(D^{\prime}\right)}\left[\frac{\operatorname{Pr}[\mathcal{M}(D)=K \mid z, \tilde{\mathcal{Q}}]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=K \mid z\right]}\right]=1
$$

The first moment is equal to 1 critically relies on our construction where the numerator in the expectation of $(*)$ is the $\alpha$ th power of a valid probability distributions.

This implies that $\operatorname{Pr}\left[A(K)=e^{\epsilon}\right]=e^{-\epsilon}$, therefore

$$
(*)=\mathbb{E}_{K}\left[\mathbb{E}[A(K) \mid K]^{\alpha}\right] \stackrel{\substack{\text { Jensen } \\ \perp}}{\leq}\left[\mathbb{E}\left[A(K)^{\alpha} \mid K\right]\right]=\mathbb{E}\left[A(K)^{\alpha}\right]=\operatorname{Pr}\left[A(K)=e^{\epsilon}\right] \cdot e^{\alpha \epsilon}=e^{(\alpha-1) \epsilon},
$$

which completes the proof.

Now we are ready to prove the three claims of Theorem 8.

The claim (3): Substitute the the above bound into Lemma 17, we get:

$$
\begin{aligned}
& \quad \mathbb{E}_{o \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left(\frac{\operatorname{Pr}[\mathcal{M}(D)=o]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]}\right)^{\alpha}\right] \\
& \text { Lemma } 17 \text { and } 18 \\
& \leq \mathbb{E}_{z \sim p_{\rho}}^{\perp}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha}\right] e^{(\alpha-1) \epsilon_{\nu}(\infty)} \leq e^{(\alpha-1) \epsilon_{\rho}(\alpha)} e^{(\alpha-1) \epsilon_{\nu}(\infty)} .
\end{aligned}
$$

where the second inequality in the last line uses the definition of RDP of $\mathcal{M}_{\rho}$, for a trivial query $q(D)=-\Delta, q\left(D^{\prime}\right)=0$. (3) follows by simply taking $\log (\cdot) /(\alpha-1)$ on both sides.

The claim (1) and (2): To get the other two bounds, we need an alternative analysis of (**). To avoid crowded notations, we drop the conditioning on $z$ from $\operatorname{Pr}[\cdot \mid \rho=z]$. By the definition of expectation,

$$
\begin{align*}
&(*) \leq(10)=\sum_{k=0}^{\infty} \prod_{i=0}^{k} \operatorname{Pr}\left[q_{i}\left(D^{\prime}\right)+\nu_{i}<z\right] \operatorname{Pr}\left[q_{k+1}\left(D^{\prime}\right)+\nu_{k+1} \geq z\right] \frac{\operatorname{Pr}\left[\tilde{q}_{k+1}(D)+\nu_{k+1} \geq z\right]^{\alpha}}{\operatorname{Pr}\left[\tilde{q}_{k+1}\left(D^{\prime}\right)+\nu_{k+1} \geq z\right]^{\alpha}} \\
& \stackrel{\tilde{q}=q \text { on } D^{\prime}}{\stackrel{\downarrow}{=}} \sum_{k=0}^{\infty} \prod_{i=0}^{k} \operatorname{Pr}\left[q_{i}\left(D^{\prime}\right)+\nu_{i}<z\right] \frac{\operatorname{Pr}\left[\tilde{q}_{k+1}(D)+\nu_{k+1} \geq z\right]^{\alpha}}{\operatorname{Pr}\left[\tilde{q}_{k+1}\left(D^{\prime}\right)+\nu_{k+1} \geq z\right]^{\alpha-1}} \\
& \stackrel{\text { Lemma }}{\perp}  \tag{11}\\
& \stackrel{\downarrow}{\leq}\left(\sum_{k=0}^{\infty} \prod_{i=0}^{k} \operatorname{Pr}\left[\tilde{q}_{i}\left(D^{\prime}\right)+\nu_{i}<z\right]\right) \cdot e^{\epsilon(\alpha)(\alpha-1)} .
\end{align*}
$$

In the last line, we applied the "indistinguishability" property of an RDP mechanism in Lemma 15 for the particular event $S=x \in \mathbb{R} \mid x \geq z$, for the random-variable $\mathcal{M}\left(D, \tilde{q}_{k+1}\right)$ and $\mathcal{M}\left(D^{\prime}, \tilde{q}_{k+1}\right)$ in the numerator and denominator respectively.

The issue is how to proceed. $\sum_{k=0}^{\infty} \prod_{i=0}^{k} \operatorname{Pr}\left[\tilde{q}_{i}\left(D^{\prime}\right)+\nu_{i}<z\right]$ does not sum to 1 because $\prod_{i=0}^{k} \operatorname{Pr}\left[\tilde{q}_{i}\left(D^{\prime}\right)+\nu_{i}<z\right]$ is not a probability distribution of $k$. The saving grace is the following alternative definition of expectation.
Lemma 19. For a non-negative random variable $X, \mathbb{E}[X]=\int_{0}^{\infty} \operatorname{Pr}[X>x] d x$.
Recall that $K$ is the first index of $\top$, we can rewrite $\prod_{i=0}^{k} \operatorname{Pr}\left[\tilde{q}_{i}\left(D^{\prime}\right)+\nu_{i}<z\right]$ as $\operatorname{Pr}_{D^{\prime}}[K>k \mid z]$. Thus

$$
\begin{equation*}
\sum_{k=0}^{\infty} \prod_{i=0}^{k} \operatorname{Pr}\left[\tilde{q}_{i}\left(D^{\prime}\right)+\nu_{i}<z\right]=\sum_{k=0}^{\infty} \operatorname{Pr}_{D^{\prime}}[K>k \mid z]=\mathbb{E}[K \mid z] \tag{12}
\end{equation*}
$$

It follows that

$$
\mathbb{E}_{o \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left(\frac{\operatorname{Pr}[\mathcal{M}(D)=o]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]}\right)^{\alpha}\right] \leq \mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha} \mathbb{E}[K \mid z]\right] e^{\epsilon_{\nu}(\alpha)(\alpha-1)}
$$

Claim (1) uses $\mathbb{E}[K \mid z] \leq k_{\max }+1$. By using a different Holder's inequality with conjugate pair $\gamma$ and $\gamma *:=\gamma /(\gamma-1)$, we obtain
$\mathbb{E}_{o \sim \mathcal{M}\left(D^{\prime}\right)}\left[\left(\frac{\operatorname{Pr}[\mathcal{M}(D)=o]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]}\right)^{\alpha}\right] \leq \mathbb{E}_{z \sim p_{\rho}}\left[\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\gamma^{*} \alpha}\right]^{1 / \gamma^{*}} \cdot\left(\mathbb{E}_{z \sim p_{\rho}}\left[\mathbb{E}[K \mid z]^{\gamma}\right]\right)^{1 / \gamma} \cdot e^{\epsilon_{\nu}(\alpha)(\alpha-1)}$
(2) follows by taking $\log (\cdot) /(\alpha-1)$ on both sides and applying the definition of RDP. This completes the proof of Theorem 8.

## C Other proofs of technical results

Proposition 20 (Restatement of Proposition 10 with mroe details). Let Algorithm 2 be instantiated with $\mathcal{Q}_{+}(\triangle), \mathcal{M}_{\rho}$ and $\mathcal{M}_{\nu}$ be Gaussian mechanism with parameter $\sigma_{1}$ and $\sigma_{2}$. Then for all $T<+\infty$ and $\gamma>1$ such that $\sigma_{2}>\sqrt{\gamma} \sigma_{1}$, Algorithm 2 with $c=1$ halts with $K$ rounds satisfying

$$
\begin{equation*}
\mathbb{E}_{\rho}\left[\mathbb{E}[K \mid \rho=z]^{\gamma}\right] \leq \int_{-\infty}^{\infty} \frac{1}{\sigma_{1}} \phi\left(z / \sigma_{1}\right)\left(\frac{\Phi\left((T+z) / \sigma_{2}\right)}{1-\Phi\left((T+z) / \sigma_{2}\right)}\right)^{\gamma} d z<+\infty \tag{13}
\end{equation*}
$$

where $\phi(x)=\frac{e^{-x^{2}}}{\sqrt{2 \pi}}$ and $\Phi(x)=\int_{-\infty}^{x} \phi(x) d x$ are the pdf and CDF of the standard normal distribution. If $\sigma_{2} \geq \sqrt{\gamma+1} \sigma_{1}$, then a more interpretable bound of the above is

$$
\mathbb{E}\left[\mathbb{E}[K \mid \rho=z]^{\gamma}\right] \leq 1+\left(c_{\gamma} \sqrt{2 \pi} \max \left\{\frac{T(1+\gamma)}{\sigma_{1}}, 1\right\}\right)^{\gamma}(1+\gamma)^{1 / 2} e^{\frac{\gamma T^{2}}{2 \sigma_{1}^{2}}}
$$

where $c_{\gamma}$ is a universal constant that comes from the moments bounds and depends only on $\gamma$. For the special case when $\gamma=2$, and $\sigma_{2}=\sqrt{3} \sigma_{1}$, we get $\mathbb{E}\left[\mathbb{E}[K \mid \rho=z]^{2}\right] \leq 1+2 \sqrt{3} \pi\left(1+\frac{9 T^{2}}{\sigma_{1}^{2}}\right) e^{\frac{T^{2}}{\sigma_{1}^{2}}}$.

Proof of Proposition 20. Consider the case when all queries are non-negative, and the threshold $T$ is given, then the datasets that maximizes all moments of $K \mid \rho=z$ for all $z$ are given by $f_{i}(D)=0$ for all $i$. Notice that $K \mid \rho=z$ follows a Negative Binomial Distribution, thus

$$
\mathbb{E}[K \mid z]=\frac{F_{v}[T+z]}{1-F_{v}[T+z]}
$$

where $F_{v}$ is the cumulative density function (CDF) of the noise $v$. The moments of $\mathbb{E}[K \mid z]$, when exists, can be computed by numerical integration. When $z \sim \mathcal{N}\left(0, \sigma_{1}^{2}\right)$ and $v \sim \mathcal{N}\left(0, \sigma_{2}^{2}\right)$ for $\sigma_{2}>\sigma_{1} \sqrt{\gamma}$, we can work out bounds of the $\gamma$ th moments of $\mathbb{E}[K \mid z]$.
Let $\phi$ be the standard normal density function and $\Phi$ be the CDF. There is a lower bound of the Gaussian tail for all $x>0$

$$
1-\Phi(x) \geq \frac{x}{x^{2}+1} \phi(x)
$$

Thus for $y \geq-T+\sigma_{2}$, we have

$$
\begin{aligned}
& \mathbb{E}\left[\mathbb{E}[K \mid z]^{\gamma}\right]=\int_{-\infty}^{\infty} \frac{1}{\sigma_{1}} \phi\left(z / \sigma_{1}\right)\left(\frac{\Phi\left((T+z) / \sigma_{2}\right)}{1-\Phi\left((T+z) / \sigma_{2}\right)}\right)^{\gamma} d z \\
&=\int_{-\infty}^{y} \frac{1}{\sigma_{1}} \phi\left(z / \sigma_{1}\right)\left(\frac{\Phi\left((T+z) / \sigma_{2}\right)}{1-\Phi\left((T+z) / \sigma_{2}\right)}\right)^{\gamma} d z+\int_{y}^{\infty} \frac{1}{\sigma_{1}} \phi\left(z / \sigma_{1}\right)\left(\frac{\Phi\left((T+z) / \sigma_{2}\right)}{1-\Phi\left((T+z) / \sigma_{2}\right)}\right)^{\gamma} d z \\
& \leq \int_{-\infty}^{y} \frac{1}{\sigma_{1}} \phi\left(z / \sigma_{1}\right) d z+\int_{y}^{\infty} \frac{1}{\sigma_{1}} \phi\left(z / \sigma_{1}\right)\left(\frac{\frac{(T+z)^{2}}{\sigma_{2}^{2}}+1}{\frac{(T+z)}{\sigma_{2}} \phi\left((T+z) / \sigma_{2}\right)}\right)^{\gamma} d z \\
& T+y \geq \sigma_{2} \\
& \leq \Phi\left(y / \sigma_{1}\right)+\int_{y}^{\infty} \frac{1}{\sqrt{2 \pi} \sigma_{1}} e^{-\frac{z^{2}}{2 \sigma_{1}^{2}}}(2 \pi)^{\gamma / 2} e^{\frac{\gamma(T+z)^{2}}{2 \sigma_{2}^{2}}} 2^{\gamma} \frac{(T+z)^{\gamma}}{\sigma_{2}^{\gamma}} d z \\
& u:=z+T, \tilde{\sigma}:=\left(\frac{1}{\sigma_{1}^{2}}-\frac{\gamma}{\sigma_{2}^{2}}\right)^{-1 / 2} \\
& \stackrel{\perp}{=} \Phi\left(y / \sigma_{1}\right)+(2 \pi)^{\frac{\gamma}{2}} \frac{\tilde{\sigma}}{\sigma_{1}} \int_{T+y}^{\infty} \frac{u^{\gamma}}{\sigma_{2}^{\gamma}} \frac{1}{\sqrt{2 \pi} \tilde{\sigma}} e^{-\frac{u^{2}}{2 \tilde{\sigma}^{2}}+\frac{2 u T}{2 \sigma_{1}^{2}}-\frac{T^{2}}{2 \sigma_{1}^{2}}} d u \\
&=\Phi\left(y / \sigma_{1}\right)+(2 \pi)^{\frac{\gamma}{2}} \frac{\tilde{\sigma}}{\sigma_{1}} e^{-\frac{T^{2}}{2 \sigma_{1}^{2}}+\frac{T^{2} \tilde{\sigma}^{2}}{2 \sigma_{1}^{4}}} \int_{T+y}^{\infty} \frac{u^{\gamma}}{\sigma_{2}^{\gamma}} \frac{1}{\sqrt{2 \pi} \tilde{\sigma}} e^{-\frac{\left(u-T \frac{\tilde{\sigma}^{2}}{\left.\sigma_{1}^{2}\right)^{2}}\right.}{2 \tilde{\sigma}^{2}}} d u
\end{aligned}
$$

take $|\cdot|$ and relax the range of integral

$$
\stackrel{\text { x the range of integral }}{\stackrel{\downarrow}{\leq} \Phi\left(y / \sigma_{1}\right)+(2 \pi)^{\frac{\gamma}{2}} \frac{\tilde{\sigma}}{\sigma_{1}} e^{\frac{\gamma T^{2}}{2\left(\sigma_{2}^{2}-\gamma \sigma_{1}^{2}\right)}} \mathbb{E}_{X \sim \mathcal{N}\left(\frac{T \tilde{\sigma}}{\sigma_{1}^{2}}, 1\right)}\left[|X|^{\gamma}\right]}
$$

where $\mathbb{E}_{X \sim \mathcal{N}\left(\frac{T \tilde{\tilde{2}}}{\sigma_{1}^{2}}, 1\right)}\left[|X|^{\gamma}\right]$ is the $m$ th non-central moments which is on the order of $\max \left\{\frac{T \tilde{\tilde{\sigma}}}{\sigma_{1}^{2}}, 1\right\}^{\gamma}-$ and can be evaluated in a closed-form. Finally, we can simply take $y-T+\sigma_{2}$.
Now, suppose we take $\sigma_{2}=\sqrt{1+\gamma} \sigma_{1}$, then we get $\tilde{\sigma}=\sigma_{2}$. The above bound simplifies to:

$$
\mathbb{E}\left[\mathbb{E}[K \mid z]^{\gamma}\right] \leq 1+\left(c_{\gamma} \sqrt{2 \pi} \max \left\{\frac{T(1+\gamma)}{\sigma_{1}}, 1\right\}\right)^{\gamma}(1+\gamma)^{1 / 2} e^{\frac{\gamma T^{2}}{2 \sigma_{1}^{2}}}
$$

where $c_{\gamma}$ is a universal constant that comes from the moments bounds and depends only on $\gamma$. If $\gamma=2$ and $\sigma_{2}=\sqrt{3} \sigma_{1}$, then $\tilde{\sigma}=\sigma_{2}$, and

$$
\mathbb{E}\left[\mathbb{E}[K \mid z]^{2}\right] \leq 1+2 \sqrt{3} \pi\left(1+\frac{9 T^{2}}{\sigma_{1}^{2}}\right) e^{\frac{T^{2}}{\sigma_{1}^{2}}}
$$

## C. 1 RDP analysis with $c \geq 1$, proof of Theorem 11

Theorem 21 (Restatement of Theorem 11, RDP for length-capped SVT with $c>1$ ). The generalized SVT with cut-off parameter $c>1$ and a maximum length is $k_{\max }$ obeys that

$$
\mathbb{D}_{\alpha}\left(\mathcal{M}(D) \| \mathcal{M}\left(D^{\prime}\right)\right) \leq \epsilon_{\rho}(\alpha)+c \epsilon_{\nu}(\alpha)+\frac{\log \sum_{k=0}^{c}\binom{k_{\max }}{k}}{\alpha-1}
$$

*After a careful revision, we found there is a minor typo in the statement of Theorem 11 (the $1 /(\alpha-1)$ term shouldn't be there) we provide the correct version above.
The proof follows a similar sequence of arguments to that we presented for $c=1$.
When $c>1$, the output space of the algorithm is $S=\{\top, \perp\}^{\ell}, \ell=0,1, \ldots, k_{\text {max }}$ with the additional restriction that the number of $\top_{\mathrm{s}}$ are smaller than $c$. Denote $I_{\perp}:=\left\{i: o_{i}=\perp\right\}$ and $I_{\top}:=\left\{j: o_{j}=\right.$ $\top\}$. Then we can write the probability of outputting $o$ as following:

$$
\operatorname{Pr}[\mathcal{M}(D)=o]=\int_{-\infty}^{+\infty} p_{\rho}(z)\left(\prod_{i \in I_{\perp}} \int_{-\infty}^{z-q_{i}(D)} p\left(\nu_{i}\right) d \nu_{i}\right)\left(\prod_{j \in I_{\top}} \int_{z+q_{j}(D)}^{\infty} p\left(\nu_{j}\right) d \nu_{j}\right) d z
$$

Similarly, we have

$$
\begin{equation*}
\mathbb{E}_{o}\left[\left(\frac{\operatorname{Pr}[\mathcal{M}(D)=o]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]}\right)^{\alpha}\right]=\sum_{o \in S} \operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]\left(\frac{\operatorname{Pr}[\mathcal{M}(D)=o]}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]}\right)^{\alpha} \tag{*}
\end{equation*}
$$

Apply the same logic from the proof for $c=1$, we can upper bound $\operatorname{Pr}[\mathcal{M}(D)=o]$ in the following

$$
\begin{equation*}
\operatorname{Pr}[\mathcal{M}(D)=o] \leq \mathbb{E}_{z \sim p_{\rho}}\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)} \prod_{i \in I_{\perp}} \int_{-\infty}^{z-q_{i}\left(D^{\prime}\right)} p\left(\nu_{i}\right) d \nu_{i}\right)\left(\prod_{j \in I_{\top}} \int_{z+\Delta+q_{j}(D)}^{\infty} p\left(\nu_{j}\right) d \nu_{j}\right) \tag{14}
\end{equation*}
$$

Then apply Jensen' inequality to $\frac{\operatorname{Pr}[\mathcal{M}(D)=o]^{\alpha}}{\operatorname{Pr}\left[\mathcal{M}\left(D^{\prime}\right)=o\right]^{\alpha-1}}$, we have

$$
\begin{equation*}
(14) \leq \sum_{o \in S} \mathbb{E}_{z \sim p_{\rho}} \frac{\left(\frac{p_{\rho}(z-\Delta)}{p_{\rho}(z)} \prod_{i \in I_{\perp}} \int_{-\infty}^{z-q_{i}\left(D^{\prime}\right)} p\left(\nu_{i}\right) d \nu_{i}\right)^{\alpha}\left(\prod_{j \in I_{\top}} \int_{z+\Delta+q_{j}(D)}^{\infty} p\left(\nu_{j}\right) d \nu_{j}\right)^{\alpha}}{\left(\prod_{i \in I_{\perp}} \int_{-\infty}^{z-q_{i}\left(D^{\prime}\right)} p\left(\nu_{i}\right) d \nu_{i}\right)^{\alpha-1}\left(\prod_{j \in I_{\top}} \int_{z+q_{j}\left(D^{\prime}\right)}^{\infty} p\left(\nu_{j}\right) d \nu_{j}\right)^{\alpha-1}} \tag{15}
\end{equation*}
$$

Exchange the order of integral in $z$ and $o$, we get

$$
(15)=\underbrace{\mathbb{E}_{z \sim p_{\rho}}\left(\frac{p_{\rho}(z-\triangle)}{p_{\rho}(z)}\right)^{\alpha} \sum_{o \in S} \operatorname{Pr}\left[\prod_{i \in I_{\perp}} q_{i}\left(D^{\prime}\right)+\nu_{i}<z\right]}_{\text {denote by }(* *)} \cdot \frac{\left(\prod_{j \in I_{\top}} \operatorname{Pr}\left[q_{j}(D)+\nu_{j}+\Delta \geq z\right]\right)^{\alpha}}{\left(\prod_{j \in I_{\top}} \operatorname{Pr}\left[q_{j}\left(D^{\prime}\right)+\nu_{j} \geq z\right]\right)^{\alpha-1}}
$$

In the case of length-capped SVT, the algorithm stops whenever $|o| \geq k_{\max }$ or $\left|I_{\top}\right| \geq c$. By the fact that probabilities $\leq 1$, we use the following crude bound

$$
\sum_{o \in S} \operatorname{Pr}\left[\prod_{i \in I_{\perp}} q_{i}\left(D^{\prime}\right)+\nu_{i}<z\right] \leq \sum_{o \in S} 1=|S|
$$

i.e., the cardinality of the output space, which is bounded from above by $\sum_{k=0}^{c}\binom{k_{\max }}{k}$.

Moreover, we can bound the $\mathbb{E}_{z \sim p_{\rho}}\left(\frac{p_{\rho}(z-\Delta)}{p_{\rho}(z)}\right)^{\alpha}$ term with $e^{(\alpha-1) \epsilon_{\rho}(\alpha)}$ using the definition of RDP (with a trivial query that outputs 0 and $\triangle$ for $D$ and $D^{\prime}$ as we constructed before). Therefore, ( $\left.* *\right)$ is bounded by $e^{(\alpha-1) \epsilon_{\rho}(\alpha)} \cdot \sum_{k=0}^{c}\binom{k_{\text {max }}}{k}$.

For the second part $\frac{\left(\prod_{j \in I_{\top}} \operatorname{Pr}\left[q_{j}(D)+\nu_{j}+\Delta \geq z\right]\right)^{\alpha}}{\left(\prod_{j \in I_{\top}} \operatorname{Pr}\left[q_{j}\left(D^{\prime}\right)+\nu_{j} \geq z\right]\right)^{\alpha-1}}$, we apply the same trick of defining a sequence of fictitious queries $\tilde{q}_{1}, \ldots, \tilde{q}_{k_{\text {max }}}$ as in 4 . For each $j \in I_{\top}, \frac{\operatorname{Pr}\left[\tilde{q}_{j}(D)+\nu_{j}+\Delta \geq z\right]}{\left.\operatorname{Pr}\left[\tilde{q}_{j}\left(D^{\prime}\right)+\nu_{j} \geq z\right]\right)^{\alpha-1}} \leq e^{\epsilon_{\nu}(\alpha)(\alpha-1)}$ using the "indistinguishability" property of an RDP mechanism defined in Lemma 15. Since $\left|I_{\top}\right| \leq c$, the second part is bounded by $e^{c \epsilon_{\nu}(\alpha)(\alpha-1)}$.

## D Different ways of doing composition

There are several different ways of achieving strong composition for SVT.

## D. 1 RDP Composition of generalized SVT with $c=1$

In the case when a finite $k_{\max }$ is enforced, SVT is able to process up to $c k_{\text {max }}$ queries, and has a total RDP bound of $c$ times

$$
\epsilon(\alpha) \leq c \epsilon_{\rho}(\alpha)+c \epsilon_{\nu}(\alpha)+\frac{c \log \left(1+k_{\max }\right)}{\alpha-1}
$$

Corollary $22((\epsilon, \delta)$-DP of Gaussian SVT with $c \geq 1$ with RDP composition). When both noise are Gaussian we get an overall $(\epsilon, \delta)$-DP with

$$
\epsilon(\delta) \leq \frac{c \Delta^{2}}{2 \sigma_{1}^{2}}+\frac{2 c \Delta^{2}}{\sigma_{2}^{2}}+2 \sqrt{c\left(\frac{\Delta^{2}}{2 \sigma_{1}^{2}}+\frac{2 \Delta^{2}}{\sigma_{2}^{2}}\right)\left(\log \left(\delta^{-1}\right)+c \log \left(\left(1+k_{\max }\right)\right)\right)}
$$

The results recover the strong composition that gives an $\left(O\left(\frac{\sqrt{c \log (1 / \delta)} \Delta}{\sigma_{1}}\right), \delta\right)$-DP guarantee up to $c=O\left(\sqrt{\sigma_{1}^{2} / \triangle^{2}}\right)$, provided that $\delta \leq\left(k_{\max }\right)^{-c}$. In the case of nonnegative query case, we may obtain a similar bound for each valid $\gamma$ which requires only $\delta \leq\left(O\left(\frac{T^{2}}{\sigma_{1}^{2}}\right)\right)^{-c}$. For larger $c$ or large $\delta$, the privacy losses increase linearly with $c$.

## D. 2 Directly use the RDP of generalized SVT with $c>1$

We discussed this in the remark underneath Theorem 11, which says that we obtain $(\epsilon, \delta)$-DP with

$$
\epsilon(\delta) \leq \frac{\Delta^{2}}{2 \sigma_{1}^{2}}+\frac{2 c \Delta^{2}}{\sigma_{2}^{2}}+\sqrt{2\left(\frac{\Delta^{2}}{2 \sigma_{1}^{2}}+\frac{2 c \Delta^{2}}{\sigma_{2}^{2}}\right)\left(\log \left(\delta^{-1}\right)+\log c\binom{k_{\max }}{c}\right)}
$$

The story is similar to the case for $c=1$, we get the $O(\sqrt{c})$-type strong composition when $\delta \propto k_{\max }^{-c}$, and $O(c)$-type composition when $\delta<k_{\max }^{-1}$.

## D. 3 Hybrid composition based on both RDP and KOV-composition

The restriction on small $\delta$ is quite limiting, which motivated us to consider the hybrid scheme for composition as in the Stagewise generalized SVT (Algorithm 3) which chooses $c^{\prime}$ according to the pre-specified $\delta$. This essentially allows us to replace $c$ with $\log (c / \delta)$ as in Theorem 12.

## E Connections of Noisy Screening and its Data-Dependent Privacy Bound

NoisyScreening and its data-dependent privacy loss computation is an alternative way of viewing the problem that SVT addresses. Similar to SVT, given a threshold, noisy screening output $\perp$ if $q_{t}(D)+\mathcal{N}\left(0, \sigma^{2}\right) \geq T$ and $\top$ otherwise. Notice that this is a fixed algorithm that obeys $\left(\alpha, \frac{\alpha k \triangle^{2}}{\sigma^{2}}\right)-$ RDP when running for $k$ iterations by the post-processing property and composition of Gaussian mechanisms. However, when $\left|q_{t}(D)-T\right|$ is large, then the post-processing of the output of Gaussian mechanism actually amplifies the privacy guarantee due to the overwhelming probability of outputting either $\perp$ or $\top$ on $D$ and all neighboring dataset $D^{\prime}$ to $D$, (see details in [Papernot et al., 2018, Theorem 6]).

The data-dependent analysis of Noisy screening behaves like SVT as it pays a privacy loss only when the query $q_{t}(D)$ is close to the threshold $T$. The notion of privacy loss that is achieved by noisyscreening has a similar notion of DP but it is data-dependent therefore the value itself is considered private information. Part of our motivation in this work is to investigate whether we can obtain formally differentially private methods that replicates the practical performance of data-dependent privacy analysis for noisy screening while satisfying a pre-specified privacy budget.

## F Applications

## F. 1 Adaptive data analysis

The Private-Guess-and-Check algorithm is an application of sparse vector technique that assumes that the analyst sends a "guess" with each query. The curator will then check if the guess is sufficiently accurate by SVT and return the "guess" without tapping into the data and move to the next round and only release something from the data using, e.g., Gaussian mechanism if the "guess" fails to be accurate. In this way, SVT will allow us to answer exponentially more queries while paying a privacy loss that essentially depends only on the number of "incorrect" guesses.
Example 23. "Reuseable-Holdout" is one way to generate these guesses in the context of adaptive data analysis. In the original version, the algorithm split the dataset into $c+1$ folds and use each fold to answer only one "bad" query. A better version splits the dataset into just two, and then give one of the fold to the analyst for exploratory data analysis and for coming up with queries and guesses, then answer those queries using private-guess-and-check.
Example 24. "Private-Multiplicative-weight" [Hardt and Rothblum, 2010] is a classical DP algorithm that can be thought of as another instantiation of the Private-Guess-and-Check. It aims at iteratively refine a vector of weights on all possible datasets (each one of them is an expert) using the Hedge algorithm from online learning to produce "guesses" in the form of an exponentially weighted averages. The algorithm using SVT such that it only releases the private answers to the queries if the answer is sufficiently different from the "guess". The regret bound of the Hedge algorithm yields a bound of $c$ that implies a bound on the wrong "guesses", which then allows answering exponentially many queries accurately.
Lemma 25 (Theorem 3.5 [Jung et al., 2020] ). Suppose $\mathcal{M}$ is $(\epsilon, \delta)$-differentially private and $(\alpha, \beta)$-accurate with respect to samples. Then for any analyst $\mathcal{A}$ that chooses a sequence of queries adaptively, we have that


Proof of Theorem 14. Let the threshold that we use $T=\frac{\sigma \sqrt{2 \log (k / \delta)}}{n}$, then it is clear that the answers are $(\sigma \sqrt{2 \log (k / \delta)}, \delta)$-sample-accurate using the concentration bound of Gaussian noise. Now by the composition of $c$-Gaussian mechanisms and the Gaussian-SVT with parameter $c=1, k, \sigma, \triangle=1 / n$, the whole procedure is $\left(O\left(\sqrt{\frac{c}{n^{2} \sigma^{2}} \log (1 / \delta)^{2}}\right), \delta\right)$-DP. By choosing $\sigma \asymp c^{1 / 4} n^{-1 / 2} \log (k / \delta)^{1 / 4}$, and then pick $v_{1}=v_{2}=O\left(c^{1 / 4} n^{-1 / 2} \log (k / \delta)^{3 / 4}\right)$, then we obtain the high-probability generalization bound for adaptive queries as claimed.

Remark. In comparison, if we use Laplace mechanism-based SVT for the same problem, the answers will only be $(\log (k / \delta) /(n \epsilon), \delta)$-sample-accurate, due to the perturbation of the threshold. It will be $(O(\epsilon \sqrt{c \log (1 / \delta)}), \delta)$-DP. By choosing $\epsilon$ appropriately to balance the two terms, we get that the standard SVT also achieves the same bound of $O\left(c^{1 / 4} \log (k / \delta)^{3 / 4} n^{-1 / 2}\right)$.

## F. 2 Model-agnostic private learning

Model-agnostic private learning is another cute application of the sparse-vector technique. In this problem, the learner has access to a private labeled dataset and a public unlabeled dataset. The algorithm leverages a blackbox learner, e.g., a deep learning algorithm, by training one classifier on each randomly split of the private dataset. Then it privately labels the public dataset by privately releasing the majority-votes of these classifiers' predictions.

This scheme has been shown to be practical [Papernot et al., 2017, 2018] by combining simple Gaussian mechanism for differential privacy with semi-supervised learning approaches. Bassily et al. [2018] substantially improves the algorithm by showing that, under a PAC-learning framework, that one can privately release the labels for all public data points, while spending the privacy budget only for those data points where the voters are labeling incorrectly (or labeling inconsistently, to be more general).
SVT is applied to test whether each query has received an overwhelming majority from the voters by testing if distance-to-stability is sufficiently large. If so, the exact answer $f(D)$ is released with $\perp$ and if not, only $T$ is released. Interestingly, this approach has a privacy loss that depends only on the number of Ts and it can be thought of as a composition of the SVT with the $(0, \delta)$-DP part of the event from the "Stability"-based argument.
While this approach provides a substantial benefit in theory, it has been observed in practice that it is often outperformed by simple Gaussian mechanism in practice, since the latter uses a moreconcentrated noise and also a much tighter composition.

In the experiment section, we demonstrate that the story is now different when Gaussian SVT is used as a drop-in replacement.

## G More details about the experiments

## G. 1 Calibrating noise to privacy

In this section, given a predetermined privacy budget $(\epsilon, \delta)$ and the cut-off $c$, we evaluate SVT variants by comparing how many queries each SVT algorithm. Suppose we have an infinite sequence of queries with ground truth at 0 (null-hypothesis) and a fixed margin/ threshold. For each of the SVT-like algorithm, we calibrate the noise (or the length for Gaussian-based SVT) according to the privacy budget.
We estimate the length of answered queries with Negative Binomial Distribution. For example, in the case of Gaussian-SVT, when $z \sim \mathcal{N}\left(0, \sigma_{1}^{2}\right)$ and $\nu \sim \mathcal{N}\left(0, \sigma_{2}^{2}\right)$, denote $K$ as the number of queries answered when hits $c$. Notice that $K \mid \rho=z$ follows a Negative Binomial Distribution, $\mathbb{E}[K \mid z]$ can be estimated with $\frac{c F_{\nu}[T+z]}{1-F_{\nu}[T+z]}$, where $F_{\nu}$ is the CDF of the noise $\nu$ and queries are all zeros. Then by law of expectation, we can estimate $\mathbb{E}[K]$ with

$$
\frac{1}{N} \sum_{i=1}^{N} \frac{c \cdot F_{\nu}\left[T+z_{i}\right]}{1-F_{\nu}\left[T+z_{i}\right]}
$$

$N$ is the number of trails and we sample $z_{i} \sim \mathcal{N}\left(0, \sigma_{1}^{2}\right)$ in each trail. We set $N=10^{5}$ in our experiment.
We present our results in Figure 1, where the predefined privacy budget $\epsilon$ is varied at the $x$-axis. Regarding the choice of $k_{\max }$, noting setting $k_{\max }$ too large would lead the algorithm to hit $\top$ far before it reaches $k_{\max }$ limit. Hence we set $k_{\max }$ individually for each algorithm with each $\epsilon$ budget. For example, we set $k_{\max } \approx 50 \cdot 2^{4 \epsilon}$ for Gaussian-SVT with $T=100$ and $k_{\max } \approx 20 \cdot 10^{5 \epsilon}$ with $T=700$ such that $k_{\max } \approx 1 /\left(1-F_{\nu}\left[T+z_{i}\right]\right)$. The curves are unsmooth for some algorithms (e.g., Gaussian-SVT), since they would start a new subroutine when $k_{\max }$ is achieved even before hitting $\top$. The left part is the low margin regime $(T=100)$ and $\delta$ is set to be $10^{-6}$. The purple line is the stage-wise generalized SVT, which can answer the largest number of queries across two regimes, especially in the high-margin regime. This is expected since when the margin $T$ is sufficiently high, the false positives will fall into the area of tails bound in either Laplace or Gaussian distribution. As discussed earlier, the Laplace distribution used in the SVT is heavy-tailed distribution, which would trigger the $c$ false positives sooner.

## G. 2 Privacy cost for answering a full sequence

In this section, we evaluate SVT variants by comparing the composed privacy loss for finishing a fixed length sequence of queries.
Given a fixed sequence of $|Q|=10^{5}$ queries with ground truth at 0 , we fix the margin with $T=1000$ for all algorithms, and align the standard deviation of Laplace noise and Gaussian noise, which is used

(a) Exp 2: Synthetic null sequence with aligned vari-(b) Exp 2: Synthetic null sequence with aligned tail ance of noise
bound
Figure 3: Total composed privacy loss a the algorithm progressed for $\delta=10^{-6}$. On the left, we fix the margin $T=1000$ and $\sigma_{1}=210$ and align the variance of noise to perturb the queries. On the right, we align the tail bound for all algorithms by varying the noise scale with a fixed threshold $T=1500$.
to perturb the query, i.e., $\sigma_{2}=\sqrt{2} \lambda_{2}=240$. In Figure 3(a), the black-dash line describes the privacy cost of the Gaussian mechanism, serving as our baseline. Notably, with the same level noise to perturb queries, the Gaussian Mechanism's global sensitivity is the half of that in Gaussian-SVT. The brown line reports the privacy cost of data-dependent noisy screening before applying smooth-sensitivity analysis, which itself is considered private information. When the \#iteration is small, all SVT-based algorithms have a flat region due to the number of $T$ is zero.
We now evaluate SVT variants with an aligned tail bound. As shown in Figure 3(b), we fix the margin $T=1500$ and vary $\gamma$ at the x-axis. For each $\gamma$, we choose the noise $\nu$ for each algorithm adaptively, such that the false positive rate is $\gamma$. For example, In the case of all Gaussian-based algorithms, the $\sigma_{2}$ is set such that $F_{\nu}[T]=1-\gamma$, where $\nu$ is drawn from $\mathcal{N}\left(0, \sigma_{2}^{2}\right)$ and $F_{\nu}$ is the CDF of the noise $\nu$. We see that Gaussian-SVT and Laplace-SVT perform similarly as their tail bounds are aligned and the stage-wise generalized SVT has the least privacy cost as it enjoys a sharper composition. The big gap between SVT-based algorithms and the Noisy Screening is due to the numerical issues in the calculation of the survival function: as $\gamma$ is small, e.g., $\gamma<10^{-6}$, the scipy-based calculation in python will output 0 for $\operatorname{Pr}\left[q_{t}(D)+\mathcal{N}\left(0, \sigma_{2}^{2}\right) \geq T\right]$, which implied no false positive is detected. However, for all SVT-based algorithms, we enforce the false-positive $c$ to be at least 1, which can potentially yield a significant gap when $\gamma$ is small.
Aligned variance or aligned FPR Note that there are two undetermined parameters ( $\sigma_{2}, \lambda_{2}$ ) for SVT variants if we compare the privacy cost of answering a full sequence with a fixed margin of $T$. To build a meaningful connection between Laplace variants and Gaussian variants, we consider two types of alignments - variance of noise and FPR. There are other choices of alignments (e.g., mean, the third moment). Our consideration of the choice is to investigate the regions where Gaussian variants are advantageous. More specifically, with an aligned variance, we can observe the advantage of Gaussian SVT due to a thinner tail bound while this advantage disappears when we align the tail bound (FPR).

## G. 3 Evaluation with real life data

In the application of model agnostic learning, Noisy screening is applied to test whether each query has received an overwhelming majority from the voters. We instantiate the task with the Private-kNN framework Zhu et al. [2020] using the CIFAR-10 dataset. The training set of CIFAR-10 is simulated as the private domain, and a sequence of queries are drawn from the public domain (testing set); for each query, we pick the top $K=300$ closest neighbors from the private domain and output $\top$ if the plurality of neighbors is above a predetermined threshold.

In the privacy-preserving screening, we set the threshold $T=210$ and fix the Gaussian noise $\sigma_{2}=80$ to perturb query for each algorithm. For Gaussian-SVT, we set $\sigma_{1}=40$ to perturb the threshold. The error bars are computed based on 10 independent run and has a correct $95 \%$ coverage.

## G. 4 Adaptive data analysis

"Reuseable-Holdout" is one application of SVT in the context of adaptive data analysis. We instantiate the task using the experimental setups from Dwork et al. [2015a]. In this experiment, the analyst is given a $d$-dimensional labeled data $S$ of size $2 n$, where each attribute is drawn independently from the normal distribution $\mathcal{N}(0,1)$. The analyst splits $S$ randomly into a training set $S_{t}$ and a holdout set $S_{h}$ of equal size. The labels $y \in\{-1,1\}$ are generated uniformly at random, so the data point and its label are not correlated. The goal of the analyst is to select variables to be included in the classifier. The analyst picks $k$ variables ( $k<d$ ) with the label's largest absolute correlations with the training set. Then she verifies the correlations on the holdout set and selects only those variables whose correlation agrees in sign with the correlation on the training set, and the correlations are greater than a predefined threshold $T$. Then the analyst generates a linear threshold classifier using the selected variables and tests it on the holdout set. Full details can be found in the supplementary materials of Dwork et al. [2015a]. We set $n=10000, d=10000$ and varied the number of selected variables $k$. After a careful inspectation of their code, we find that they use the Gaussian mechanism to select variables, rather than the Laplace-based SVT that they analyze. Note that this is also one motivation of this work - to find SVTs with thin tail bounds and makes SVT practical.
We provide the "Reuseable-Holdout" algorithm implemented with Gaussian Mechanism as follows.
Step 1 For each attribute $i \in[d]$ compute the correlation with the label on the training and holdout sets: $w_{i}^{t}=\sum_{(x, y) \in S_{t}} x_{i} y$ and $w_{i}^{h}=\sum_{(x, y) \in S_{h}} x_{i} y$.
Step 2 Sort $w_{i}^{t}$ with $k$ lagest values.
Step 3 For each of the $k$ features, if training correlation $\left|w_{i}^{t}\right| \geq 1 / \sqrt{n} ;\left|w_{i}^{h}\right| \geq 1 / \sqrt{n}$, test if $\left|w_{i}^{t}-w_{i}^{h}\right|+\mathcal{N}\left(0, \sigma_{2}^{2}\right)>T$. If so, return $\perp$, else return $\top$.
Step 4 Pick out the subset of features (denoted as $V_{k}$ ) with output $\perp$ and construct a linear classifier $f(x)=\operatorname{sign}\left(\sum_{i \in V_{k}} \operatorname{sign}\left(w_{i}^{t}\right) \cdot x_{i}\right)$ using the sign of the training correlation.

In the Gaussian-SVT based algorithm, we first perturb $T$ with $\tilde{T}=T+\mathcal{N}\left(0, \sigma_{1}^{2}\right)$. Then we modefiy the "step 3 " with testing if $\left|w_{i}^{t}-w_{i}^{h}\right|+\mathcal{N}\left(0, \sigma_{2}^{2}\right)>\tilde{T}$. If so, return $\perp$ and refresh $\tilde{T}$, else return $\top$.
We set $T=0.04$ and $2 \sigma_{1}=\sigma_{2}=0.01$. As shown in Figure 4, we provide the average and standard deviation of results with 100 independent trials. Noting there are no correlations between $x$ and $y$, no classifiers can achieve accuracy better than $50 \%$. However, a standard holdout (denoted as "training") results in accuracy $\gg 50 \%$, which is overfitting. "Fresh" refers to the classifier accuracy on another fresh data of size $n$, which results in an accuracy $50 \%$. Both the Gaussian mechanism and Gaussian SVT can prevent the algorithm from overfitting to the holdout set. Moreover, with the same noise scale to perturb queries, Gaussian-SVT pays privacy cost proportional to the size of features that outputs $\perp$. Hence it can answer exponentially more queries with the same privacy budget.

## H One example for SVT with unbounded length

We provide an example for Laplace SVT with unbounded length as follows.
Proposition 26. Suppose we have an infinite sequence of queries $q_{1}, \ldots, q_{\infty}$ with ground truth at 0 , threshold $T_{i}$ increases exponentially (i.e. $T_{i+1}=\theta T_{i}$ ) where $i$ is the index of the query. The Laplace noise $\nu_{i} \sim \operatorname{Lap}\left(\lambda_{2}\right)$ is used to perturb each query. Then the expected length of SVT with a cut-off $c=1$ is unbounded.

Proof. We first rewrite $T_{0}=\beta \cdot \lambda_{2}$. Recall that if $Y \sim \operatorname{Lap}(b)$, then $\operatorname{Pr}[Y \geq t \cdot b]=\frac{1}{2} \exp (-t)$. Therefore, we have $\operatorname{Pr}[K=1]=\frac{1}{2} \exp (-\beta)$ where $K$ is a random variable indicating the stopping time. Moreover, $\operatorname{Pr}[K=2]$ can be written as $\operatorname{Pr}\left[q_{2}+\nu_{2} \geq T_{2} \mid q_{1}+\nu_{1}<T_{1}\right] \cdot \operatorname{Pr}[K \neq 1]$, which is smaller than $\frac{1}{2} \exp (-\beta \cdot \theta)$. Noting the probability of stopping at all integer can be written as


Figure 4: Learning uncorrelated label. The $x$-axis is the number of variables selected for the classifier. The $y$-axis indicates average classification accuracy over 100 executions.
$\lim _{i \rightarrow \infty} \operatorname{Pr}[K<i+1]=\sum_{i=1}^{\infty} \operatorname{Pr}[K=i] \leq \sum_{i=1}^{\infty} \frac{1}{2} \exp \left(-\beta \theta^{i-1}\right)<1$. Hence the expected length of SVT with a cut-off $c=1$ is unbounded.

