Supplementary Material for "The Composition Theorem for Differential Privacy"

5 Proof of Theorem 3.3

We first propose a simple mechanism and prove that the proposed mechanism dominates over all (ε, δ)-differentially private mechanisms. Analyzing the privacy region achieved by the k-fold composition of the proposed mechanism, we get a bound on the privacy region under the adaptive composition. This gives an exact characterization of privacy under composition, since we show both converse and achievability. We prove that no other family of mechanisms can achieve 'more degraded' privacy (converse), and that there is a mechanism that we propose which achieves the privacy region (achievability).

5.1 Achievability

We propose the following simple mechanism $\tilde{M}_i$ at the i-th step in the composition. Null hypothesis (b=0) outcomes $X^{i,0} = M_i(D^{i,0}, q_i)$'s which are independent and identically distributed as a discrete random variable $\tilde{X}_0 \sim \tilde{P}_0(\cdot)$, where

$$\mathbb{P}(\tilde{X}_0 = x) \;=\; \tilde{P}_0(x) \;\equiv\; \begin{cases} \delta & \text{for } x = 0\,, \\ \frac{(1-\delta)\,e^{\varepsilon}}{1+e^{\varepsilon}} & \text{for } x = 1\,, \\ \frac{1-\delta}{1+e^{\varepsilon}} & \text{for } x = 2\,, \\ 0 & \text{for } x = 3\,. \end{cases} \tag{12}$$

Alternative hypothesis (b=1) outcomes $X^{i,1} = M_i(D^{i,1}, q_i)$'s are independent and identically distributed as a discrete random variable $\tilde{X}_1 \sim \tilde{P}_1(\cdot)$, where

$$\mathbb{P}(\tilde{X}_1 = x) \;=\; \tilde{P}_1(x) \;\equiv\; \begin{cases} 0 & \text{for } x = 0\,, \\ \frac{1-\delta}{1+e^{\varepsilon}} & \text{for } x = 1\,, \\ \frac{(1-\delta)\,e^{\varepsilon}}{1+e^{\varepsilon}} & \text{for } x = 2\,, \\ \delta & \text{for } x = 3\,. \end{cases} \tag{13}$$

In particular, the output of this mechanism does not depend on the database $D^{i,b}$ or the query q_i, and only depends on the hypothesis b. The privacy region of a single access to this mechanism is $\mathcal{R}(\varepsilon, \delta)$ in Figure 1. Hence, by Theorem 2.5, all (ε, δ)-differentially private mechanisms are dominated by this mechanism.

In general, the privacy region $\mathcal{R}(M, D_0, D_1)$ of any mechanism can be represented as an intersection of multiple $\{(\tilde{\varepsilon}_j, \tilde{\delta}_j)\}$ privacy regions for $j \in \{1, 2, \ldots\}$. For a mechanism M, we can compute the $(\tilde{\varepsilon}_j, \tilde{\delta}_j)$ pairs representing the privacy region as follows. Given a null hypothesis database D_0, an alternative hypothesis database D_1, and a mechanism M whose output space is $\mathcal{X}$, let P_0 and P_1 denote the probability density function of the outputs $M(D_0)$ and $M(D_1)$ respectively. To simplify notations we assume that P_0 and P_1 are symmetric, i.e. there exists a permutation π over $\mathcal{X}$ such that $P_0(x) = P_1(\pi(x))$ and $P_1(x) = P_0(\pi(x))$. This ensures that we get a symmetric privacy region.

The privacy region $\mathcal{R}(M, D_0, D_1)$ can be described by its boundaries. Since it is a convex set, a tangent line on the boundary with slope $-e^{\tilde{\varepsilon}_j}$ can be represented by the smallest $\tilde{\delta}_j$ such that

$$P_{\mathrm{FA}} \;\;\geq\;\; -e^{\tilde{\varepsilon}_j} P_{\mathrm{MD}} + 1 - \tilde{\delta}_j\,, \tag{14}$$

for all rejection sets (cf. Figure 3). Letting S denote the complement of a rejection set, such that $P_{\rm FA} = 1 - P_0(S)$ and $P_{\rm MD} = P_1(S)$, the minimum shift $\tilde{\delta}_j$ that still ensures that the privacy region is above the line (14) is defined as $\tilde{\delta}_j = d_{\tilde{\varepsilon}_j}(P_0, P_1)$ where

$$d_{\tilde{\varepsilon}}(P_0, P_1) \;\equiv\; \max_{S \subseteq \mathcal{X}} \Big\{ P_0(S) - e^{\tilde{\varepsilon}}\, P_1(S) \Big\} .$$

The privacy region of a mechanism is completely described by the set of slopes and shifts, $\{(\tilde{\varepsilon}_j, \tilde{\delta}_j) : \tilde{\varepsilon}_j \in E \text{ and } \tilde{\delta}_j = d_{\tilde{\varepsilon}_j}(P_0, P_1)\}$, where

$$E \;\equiv\; \{\, 0 \leq \tilde{\varepsilon} < \infty \,:\, P_0(x) = e^{\tilde{\varepsilon}}\, P_1(x) \text{ for some } x \in \mathcal{X} \} .$$

Any $\tilde{\varepsilon} \notin E$ does not contribute to the boundary of the privacy region. For the above example distributions $\tilde{P}_0$ and $\tilde{P}_1$, $E = \{\varepsilon\}$ and $d_\varepsilon(\tilde{P}_0, \tilde{P}_1) = \delta$.

Remark 5.1. *For a database access mechanism M over a output space $\mathcal{X}$ and a pair of neighboring databases D_0 and D_1, let P_0 and P_1 denote the probability density function for random variables $M(D_0)$ and $M(D_1)$ respectively. Assume there exists a permutation π over $\mathcal{X}$ such that $P_0(x) = P_1(\pi(x))$. Then, the privacy region is*

$$\mathcal{R}(\, M, D_0, D_1\,) \;=\; \bigcap_{\tilde{\varepsilon} \in E} \mathcal{R}\big(\, \tilde{\varepsilon}, d_{\tilde{\varepsilon}}(P_0, P_1)\,\big) ,$$

where $\mathcal{R}(M, D, D')$ and $\mathcal{R}(\tilde{\varepsilon}, \tilde{\delta})$ are defined as in (3) and (2).

The symmetry assumption is to simplify notations, and the analysis can be easily generalized to deal with non-symmetric distributions.

Now consider a k-fold composition experiment, where at each sequential access $\tilde{M}_i$, we receive a random output $X^{i,b}$ independent and identically distributed as $\tilde{X}_b$. We can explicitly characterize the distribution of k-fold composition of the outcomes: $\mathbb{P}(X^{1,b} = x_1, \ldots, X^{k,b} = x_k) = \prod_{x=1}^{k} \tilde{P}_b(x_i)$. It follows form the structure of these two discrete distributions that, $E = \{e^{(k-2\lfloor k/2 \rfloor)\varepsilon}, e^{(k+2-2\lfloor k/2 \rfloor)\varepsilon}, \ldots, e^{(k-2)\varepsilon}, e^{k\varepsilon}\}$. After some algebra, it also follows that

$$d_{(k-2i)\varepsilon}\big(\, (\tilde{P}_0)^k, (\tilde{P}_1)^k\,\big) \;=\; 1 - (1-\delta)^k + (1-\delta)^k \frac{\sum_{\ell=0}^{i-1} \binom{k}{\ell} \big(e^{\varepsilon(k-\ell)} - e^{\varepsilon(k-2i+\ell)}\big)}{(1+e^{\varepsilon})^k} .$$

for $i \in \{0, \ldots, \lfloor k/2 \rfloor\}$. From Remark 5.1, it follows that the privacy region is $\mathcal{R}(\{\varepsilon_i, \delta_i\}) = \bigcap_{i=0}^{\lfloor k/2 \rfloor} \mathcal{R}\big(\varepsilon_i, \delta_i\big)$, where $\varepsilon_i = (k-2i)\varepsilon$ and δ_i's are defined as in (6). Figure 2 shows this privacy region for $k = 1, \ldots, 5$ and for $\varepsilon = 0.4$ and for two values of $\delta = 0$ and $\delta = 0.1$.

5.2 Converse

We will now prove that this region is the largest region achievable under k-fold adaptive composition of any (ε, δ)-differentially private mechanisms.

From Corollary 2.3, any mechanism whose privacy region is included in $\mathcal{R}(\{\varepsilon_i, \delta_i\})$ satisfies $(\tilde{\varepsilon}, \tilde{\delta})$-differential privacy. We are left to prove that for the family of all (ε, δ)-differentially private mechanisms, the privacy region of the k-fold composition experiment is included inside $\mathcal{R}(\{\varepsilon_i, \delta_i\})$.

To this end, consider the following composition experiment, which reproduces the *view of the adversary* from the original composition experiment.

At each time step i, we generate a random variable $X^{i,b}$ distributed as $\tilde{X}_b$ independent of any other random events, and call this the output of a database access mechanism $\tilde{M}_i$ such that $\tilde{M}_i(D^{i,b}, q_i) = X^{i,b}$. Since, $X^{i,b}$ only depends on b, and is independent of the actual database or the query, we use $\tilde{M}_i(b)$ to denote this outcome.

We know that $\tilde{M}_i(b)$ has privacy region $\mathcal{R}(\varepsilon, \delta)$ for any choices of $D^{i,0}$, $D^{i,1}$ and q_i. Now consider the mechanism M_i from the original experiment. Since it is (ε, δ)-differentially private, we know from Corollary 2.1 that $\mathcal{R}(M_i, D^{i,0}, D^{i,1}) \subseteq \mathcal{R}(\varepsilon, \delta)$ for any choice of neighboring databases $D^{i,0}$, $D^{i,1}$. Hence, from the converse of data processing inequality (Theorem 2.5), we know that there exists a mechanism T_i that takes as input $X^{i,b}$, q_i, $D^{i,0}$, $D^{i,1}$, and an (ε, δ)-differentially private mechanism M_i, and produces an output $Y^{i,b}$ which is distributed as $M_i(D^{i,b}, q_i)$ for all $b \in \{0, 1\}$. Hence, $Y^{i,b}$ is independent of the past conditioned on $X^{i,b}, D^{i,0}, D^{i,1}, q_i, M_i$. Precisely we have the following Markov chain:

$$(b, R, \{X^{\ell,b}, D^{\ell,0}, D^{\ell,1}, q_\ell, M_\ell\}_{\ell \in [i-1]}) \text{–} (X^{i,b}, D^{i,0}, D^{i,1}, q_i, M_i) \text{–} Y^{i,b} \ ,$$

where R is any internal randomness of the adversary $\mathcal{A}$. Since, (X, Y)–Z–W implies X–(Y, Z)-W, we have

$$b \text{–} (R, \{X^{\ell,b}, D^{\ell,0}, D^{\ell,1}, q_\ell, M_\ell\}_{\ell \in [i]}) \text{–} Y^{i,b} \ .$$

Notice that if we know R and the outcomes $\{Y^{\ell,b}\}_{\ell \in [i]}$, then we can reproduce the original experiment until time i. This is because the choices of $D^{i,0}, D^{i,1}, q_i, M_i$ are exactly specified by R and $\{Y^{\ell,b}\}_{\ell \in [i]}$. Hence, we can simplify the Markov chain as

$$b \text{–} (R, X^{i,b}, \{X^{\ell,b}, Y^{\ell,b}\}_{\ell \in [i-1]}) \text{–} Y^{i,b} \ . \tag{15}$$

Further, since $X^{i,b}$ is independent of the past conditioned on b, we have

$$X^{i,b} \text{–} b \text{–} (R, \{X^{\ell,b}, Y^{\ell,b}\}_{\ell \in [i-1]}) \ . \tag{16}$$

It follows that

$$\begin{aligned} \mathbb{P}(b, r, x_1 \ldots, x_k, y_1, \ldots, y_k) &= \mathbb{P}(b, r, x_1, \ldots, x_k, y_1, \ldots, y_{k-1}) \mathbb{P}(y_k | r, x_1, \ldots, x_k, y_1, \ldots, y_{k-1}) \\ &= \mathbb{P}(b, r, x_1, \ldots, x_{k-1}, y_1, \ldots, y_{k-1}) \mathbb{P}(x_k | b) \mathbb{P}(y_k | r, x_1, \ldots, x_k, y_1, \ldots, y_{k-1}) \ , \end{aligned}$$

where we used (15) in the first equality and (16) in the second. By induction, we get a decomposition of $\mathbb{P}(b, r, x_1, \ldots, x_k, y_1, \ldots, y_k) = \mathbb{P}(b | r, x_1, \ldots, x_k)\, \mathbb{P}(y_1, \ldots, y_k, r, x_1, \ldots, x_k)$. From the construction of the experiment, it also follows that the internal randomness R is independent of the hypothesis b and the outcomes $X^{i,b}$'s: $\mathbb{P}(b | r, x_1, \ldots, x_k) = \mathbb{P}(b | x_1, \ldots, x_k)$. Then, marginalizing over R, we get $\mathbb{P}(b, x_1, \ldots, x_k, y_1, \ldots, y_k) = \mathbb{P}(b | x_1, \ldots, x_k)\, \mathbb{P}(y_1, \ldots, y_k, x_1, \ldots, x_k)$. This implies the following Markov chain:

$$b \text{–} (\{X^{i,b}\}_{i \in [k]}) \text{–} (\{Y^{i,b}\}_{i \in [k]}) \ , \tag{17}$$

and it follows that a set of mechanisms $(M_1, \ldots, M_k)$ dominates $(\tilde{M}_1, \ldots, \tilde{M}_k)$ for two databases $\{D^{i,0}\}_{i \in [k]}$ and $\{D^{i,1}\}_{i \in [k]}$. By the data processing inequality for differential privacy (Theorem 2.4), this implies that

$$\mathcal{R}\big(\{M_i\}_{i \in [k]}, \{D^{i,0}\}_{i \in [k]}, \{D^{i,1}\}_{i \in [k]}\big) \ \subseteq \ \mathcal{R}\big(\{\tilde{M}_i\}_{i \in [k]}, \{D^{i,0}\}_{i \in [k]}, \{D^{i,1}\}_{i \in [k]}\big) \ = \ \mathcal{R}\big(\{\varepsilon_i, \delta_i\}\big) \ .$$

This finishes the proof of the desired claim.

Alternatively, one can prove (17), using a probabilistic graphical model. Precisely, the following Bayesian network describes the dependencies among various random quantities of the experiment described above. Since the set of nodes $(X^{1,b}, X^{2,b}, X^{3,b}, X^{4,b})$ d-separates node b from the rest of the bayesian network, it follows immediately from the Markov property of this Bayesian network that (17) is true (cf. [Lau96]).

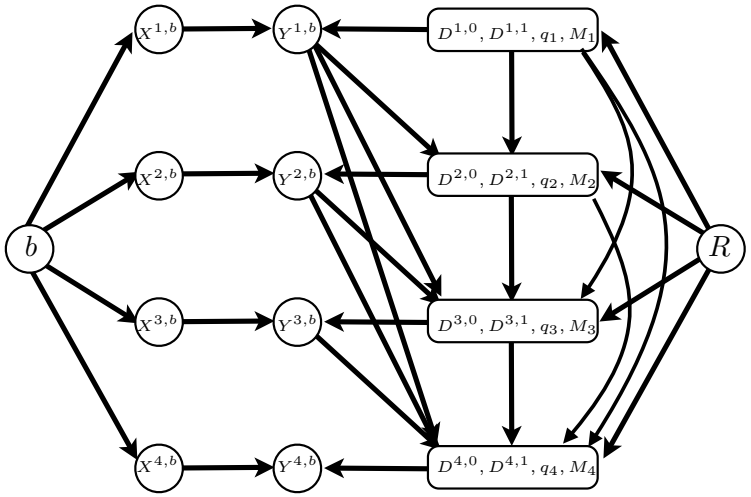

Figure 4: Bayesian network representation of the composition experiment. The subset of nodes $(X^{1,b}, X^{2,b}, X^{3,b}, X^{4,b})$ d-separates node b from the rest of the network.

6 Proof of Theorem 3.4

We need to provide an outer bound on the privacy region achieved by $\tilde{X}_0$ and $\tilde{X}_1$ defined in (12) and (13) under k-fold composition. Let P_0 denote the probability mass function of $\tilde{X}_0$ and P_1 denote the PMF of $\tilde{X}_1$. Also, let P_0^k and P_1^k denote the joint PMF of k i.i.d. copies of $\tilde{X}_0$ and $\tilde{X}_1$ respectively. Also, for a set $S \subseteq \mathcal{X}^k$, we let $P_0^k(S) = \sum_{x \in S} P_0^k(x)$. In our example, $\mathcal{X} = \{1, 2, 3, 4\}$, and

$$
\begin{aligned}
P_0 &= \begin{bmatrix} \delta & \frac{(1-\delta)e^\varepsilon}{1+e^\varepsilon} & \frac{1-\delta}{1+e^\varepsilon} & 0 \end{bmatrix}, \\
P_1 &= \begin{bmatrix} 0 & \frac{1-\delta}{1+e^\varepsilon} & \frac{(1-\delta)e^\varepsilon}{1+e^\varepsilon} & \delta \end{bmatrix}, \\
P_0^2 &= \begin{bmatrix} \delta^2 & \delta\frac{(1-\delta)e^\varepsilon}{1+e^\varepsilon} & \delta\frac{(1-\delta)}{1+e^\varepsilon} & 0 \\ \delta\frac{(1-\delta)e^\varepsilon}{1+e^\varepsilon} & \left(\frac{(1-\delta)e^\varepsilon}{1+e^\varepsilon}\right)^2 & \left(\frac{1-\delta}{1+e^\varepsilon}\right)^2 e^\varepsilon & 0 \\ \delta\frac{1-\delta}{1+e^\varepsilon} & \left(\frac{1-\delta}{1+e^\varepsilon}\right)^2 e^\varepsilon & \left(\frac{1-\delta}{1+e^\varepsilon}\right)^2 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \text{etc.}
\end{aligned}
$$

We can compute the privacy region from P_0^k and P_1^k directly, by computing the line tangent to the boundary. A tangent line with slope $-e^{\tilde{\varepsilon}}$ can be represented as

$$P_{\mathrm{FA}} = -e^{\tilde{\varepsilon}} P_{\mathrm{MD}} + 1 - d_{\tilde{\varepsilon}}(P_0^k, P_1^k) \; . \tag{18}$$

To find the tangent line, we need to maximize the shift, which is equivalent to moving the line downward until it is tangent to the boundary of the privacy region (cf. Figure 3).

$$d_{\tilde{\varepsilon}}(P_0^k, P_1^k) \;\; \equiv \;\; \max_{S \subseteq \mathcal{X}^k} \; P_0^k(S) - e^{\tilde{\varepsilon}} P_1^k(S) \; .$$

Notice that the maximum is achieved by a set $B \equiv \{x \in \mathcal{X}^k \,|\, P_0^k(x) \geq e^{\tilde{\varepsilon}} P_1^k(x)\}$. Then,

$$d_{\tilde{\varepsilon}}(P_0^k, P_1^k) \;\; = \;\; P_0^k(B) - e^{\tilde{\varepsilon}} P_1^k(B) \; .$$

For the purpose of proving the bound of the form (7), we separate the analysis of the above formula into two parts: one where either $P_0^k(x)$ or $P_1^k(x)$ is zero and the other when both are positive. Effectively, this separation allows us to treat the effects of $(\varepsilon, 0)$-differential privacy and $(0, \delta)$-differential privacy separately. In previous work [DRV10], they required heavy machinery from dense subsets of pseudorandom sets [RTTV08] to separate the analysis in a similar way. Here we provide a simple proof technique. Further, all the proof techniques we use naturally generalize to compositions of general (ε, δ)-differentially private mechanisms other than the specific example of $\tilde{X}_0$ and $\tilde{X}_1$ we consider in this section.

Let $\tilde{X}_0^k$ denote a k-dimensional random vector whose entries are independent copies of $\tilde{X}_0$. We partition B into two sets: $B = B_0 \bigcup B_1$ and $B_0 \bigcap B_1 = \emptyset$. Let $B_0 \equiv \{x \in \mathcal{X}^k \,:\, P_0^k(x) \geq e^{\tilde{\varepsilon}} P_1^k(x)$, and $P_1^k(x) = 0\}$ and $B_1 \equiv \{x \in \mathcal{X}^k \,:\, P_0^k(x) \geq e^{\tilde{\varepsilon}} P_1^k(x)$, and $P_1^k(x) > 0\}$. Then, it is not hard to see that $P_0^k(B_0) = 1 - \mathbb{P}(\tilde{X}_0^k \in \{1,2,3\}^k) = 1 - (1-\delta)^k$, $P_1^k(B_0) = 0$, $P_0^k(B_1) = P_0^k(B_1 | \tilde{X}_0^k \in \{1,2\}^k) \mathbb{P}(\tilde{X}_0^k \in \{1,2\}^k) = (1-\delta)^k \, P_0^k(B_1 | \tilde{X}_0^k \in \{1,2\}^k)$, and $P_1^k(B_1) = (1-\delta)^k \, P_1^k(B_1 | \tilde{X}_1^k \in \{1,2\}^k)$. It follows that

$$\begin{aligned} P_0^k(B_0) - e^{\tilde{\varepsilon}} P_1^k(B_0) &= 1 - (1-\delta)^k \; , \text{ and} \\ P_0^k(B_1) - e^{\tilde{\varepsilon}} P_1^k(B_1) &= (1-\delta)^k \big(P_0^k(B_1 | \tilde{X}_0^k \in \{1,2\}^k) - e^{\tilde{\varepsilon}} P_1^k(B_1 | \tilde{X}_1^k \in \{1,2\}^k) \big) \; . \end{aligned}$$

Let $\tilde{P}_0^k(x) \equiv P_0^k(x | x \in \{1,2\}^k)$ and $\tilde{P}_1^k(x) \equiv P_1^k(x | x \in \{1,2\}^k)$. Then, we have

$$\begin{aligned} d_{\tilde{\varepsilon}}(P_0^k, P_1^k) &= P_0^k(B_0) - e^{\tilde{\varepsilon}} P_1^k(B_0) + P_0^k(B_1) - e^{\tilde{\varepsilon}} P_1^k(B_1) \\ &= 1 - (1-\delta)^k + (1-\delta)^k \big(\tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}} \tilde{P}_1^k(B_1) \big) \; . \end{aligned} \tag{19}$$

Now, we focus on upper bounding $\tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}} \tilde{P}_1^k(B_1)$, using a variant of Chernoff's tail bound. Notice that

$$\begin{aligned} \tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}} \tilde{P}_1^k(B_1) &= \mathbb{E}_{\tilde{P}_0^k} \big[\mathbb{I}_{\left(\log(\tilde{P}_0^k(\tilde{X}^k)/\tilde{P}_1^k(\tilde{X}^k)) \geq \tilde{\varepsilon} \right)} \big] - e^{\tilde{\varepsilon}} \mathbb{E}_{\tilde{P}_0^k} \big[\mathbb{I}_{\left(\log(\tilde{P}_0^k(\tilde{X}^k)/\tilde{P}_1^k(\tilde{X}^k)) \geq \tilde{\varepsilon} \right)} \frac{\tilde{P}_1^k(\tilde{X}^k)}{\tilde{P}_0^k(\tilde{X}^k)} \big] \\ &= \mathbb{E}_{\tilde{P}_0^k} \Big[\mathbb{I}_{\left(\log(\tilde{P}_0^k(\tilde{X}^k)/\tilde{P}_1^k(\tilde{X}^k)) \geq \tilde{\varepsilon} \right)} \Big(1 - e^{\tilde{\varepsilon}} \frac{\tilde{P}_1^k(\tilde{X}^k)}{\tilde{P}_0^k(\tilde{X}^k)} \Big) \Big] \\ &\leq \mathbb{E}[e^{\lambda Z - \lambda\tilde{\varepsilon} + \lambda \log \lambda - (\lambda+1)\log(\lambda+1)}] \; , \end{aligned} \tag{20}$$

where we use a random variable $Z \equiv \log(\tilde{P}_0^k(\tilde{X}_0^k)/\tilde{P}_1^k(\tilde{X}_0^k))$ and the last line follows from $\mathbb{I}_{(x\geq\tilde{\varepsilon})}(1-e^{\tilde{\varepsilon}-x}) \leq e^{\lambda(x-\tilde{\varepsilon})+\lambda\log\lambda-(\lambda+1)\log(\lambda+1)}$ for any $\lambda \geq 0$. To show this inequality, notice that the right-hand side is always non-negative. So it is sufficient to show that the inequality holds, without the indicator on the left-hand side. Precisely, let $f(x) = e^{\lambda(x-\tilde{\varepsilon})+\lambda\log\lambda-(\lambda+1)\log(\lambda+1)} + e^{\tilde{\varepsilon}-x} - 1$. This is a convex function with $f(x^*) = 0$ and $f'(x^*) = 0$ at $x^* = \tilde{\varepsilon} + \log((\lambda+1)/\lambda)$. It follows that this is a non-negative function.

Next, we give an upper bound on the moment generating function of Z.

$$\begin{aligned}\mathbb{E}_{\tilde{P}_0}[e^{\lambda\log(P_0(X)/P_1(X))}] &= \frac{e^\varepsilon}{e^\varepsilon+1}e^{\lambda\varepsilon} + \frac{1}{e^\varepsilon+1}e^{-\lambda\varepsilon} \\ &\leq e^{\frac{e^\varepsilon-1}{e^\varepsilon+1}\lambda\varepsilon+\frac{1}{2}\lambda^2\varepsilon^2},\end{aligned}$$

for any λ, which follows from the fact that $pe^x + (1-p)e^{-x} \leq e^{(2p-1)x+(1/2)x^2}$ for any $x \in \mathbb{R}$ and $p \in [0,1]$ [AS04, Lemma A.1.5]. Substituting this into (20) with a choice of $\lambda = \frac{\tilde{\varepsilon}-k\varepsilon(e^\varepsilon-1)/(e^\varepsilon+1)}{k\varepsilon^2}$, we get

$$\begin{aligned}\tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}}\tilde{P}_1^k(B_1) &\leq \exp\Big\{\frac{e^\varepsilon-1}{e^\varepsilon+1}\lambda\varepsilon k + \frac{1}{2}\lambda^2\varepsilon^2 k - \lambda\tilde{\varepsilon} + \lambda\log\lambda - (\lambda+1)\log(\lambda+1)\Big\} \\ &\leq \exp\Big\{-\frac{1}{2k\varepsilon^2}\Big(\tilde{\varepsilon} - k\varepsilon\frac{e^\varepsilon-1}{e^\varepsilon+1}\Big)^2 - \log(\lambda+1)\Big\} \\ &\leq \frac{1}{1+\frac{\tilde{\varepsilon}-k\varepsilon(e^\varepsilon-1)/(e^\varepsilon+1)}{k\varepsilon^2}}\exp\Big\{-\frac{1}{2k\varepsilon^2}\Big(\tilde{\varepsilon} - k\varepsilon\frac{e^\varepsilon-1}{e^\varepsilon+1}\Big)^2\Big\} \\ &= \frac{1}{1+\frac{\sqrt{2k\varepsilon^2\log(e+(\sqrt{k\varepsilon^2}/\tilde{\delta}))}}{k\varepsilon^2}}\frac{1}{e+\frac{\sqrt{k\varepsilon^2}}{\tilde{\delta}}} \\ &\leq \frac{1}{\sqrt{k\varepsilon^2}+\sqrt{2\log(e+(\sqrt{k\varepsilon^2}/\tilde{\delta}))}}\frac{\tilde{\delta}}{\frac{e\tilde{\delta}}{\sqrt{k\varepsilon^2}}+1},\end{aligned}$$

for our choice of $\tilde{\varepsilon} = k\varepsilon(e^\varepsilon-1)/(e^\varepsilon+1) + \varepsilon\sqrt{2k\,\log(e+(\sqrt{k\varepsilon^2}/\tilde{\delta}))}$. The right-hand side is always less than $\tilde{\delta}$.

Similarly, one can show that the right-hand side is less than $\tilde{\delta}$ for the choice of $\tilde{\varepsilon} = k\varepsilon(e^\varepsilon-1)/(e^\varepsilon+1)+\varepsilon\sqrt{2k\,\log(1/\tilde{\delta})}$. We get that the k-fold composition is $(\tilde{\varepsilon}, 1-(1-\delta)^k(1-\tilde{\delta}))$-differentially private.

7 Proof of Theorem 3.5

In this section, we closely follow the proof of Theorem 3.4 in Section 6 carefully keeping the dependence on ℓ, the index of the composition step. For brevity, we omit the details which overlap with the proof of Theorem 3.4. By the same argument as in the proof of Theorem 3.3, we only need to provide an outer bound on the privacy region achieved by $\tilde{X}_0^{(\ell)}$ and $\tilde{X}_1^{(\ell)}$ under k-fold

composition, defined as

$$\mathbb{P}(\tilde{X}_0^{(\ell)} = x) \;=\; \tilde{P}_0^{(\ell)}(x) \;\equiv\; \begin{cases} \delta_\ell & \text{for } x = 0\,, \\ \frac{(1-\delta_\ell)\,e^{\varepsilon_\ell}}{1+e^{\varepsilon_\ell}} & \text{for } x = 1\,, \\ \frac{1-\delta_\ell}{1+e^{\varepsilon_\ell}} & \text{for } x = 2\,, \\ 0 & \text{for } x = 3\,. \end{cases} \text{, and}$$

$$\mathbb{P}(\tilde{X}_1^{(\ell)} = x) \;=\; \tilde{P}_1^{(\ell)}(x) \;\equiv\; \begin{cases} 0 & \text{for } x = 0\,, \\ \frac{1-\delta_\ell}{1+e^{\varepsilon_\ell}} & \text{for } x = 1\,, \\ \frac{(1-\delta_\ell)\,e^{\varepsilon_\ell}}{1+e^{\varepsilon_\ell}} & \text{for } x = 2\,, \\ \delta_\ell & \text{for } x = 3\,. \end{cases}$$

Using the similar notations as Section 6, it follows that under k-fold composition,

$$d_{\tilde{\varepsilon}}(P_0^k, P_1^k) \;=\; 1 - \prod_{\ell=1}^{k}(1-\delta_\ell) + \big(\tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}}\tilde{P}_1^k(B_1)\big)\prod_{\ell=1}^{k}(1-\delta_\ell)\,. \tag{21}$$

Now, we focus on upper bounding $\tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}}\tilde{P}_1^k(B_1)$, using a variant of Chernoff's tail bound. We know that

$$\begin{aligned} \tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}}\tilde{P}_1^k(B_1) &= \mathbb{E}_{\tilde{P}_0^k}\big[\mathbb{I}_{\left(\log(\tilde{P}_0^k(\tilde{X}^k)/\tilde{P}_1^k(\tilde{X}^k))\geq\tilde{\varepsilon}\right)}\big] - e^{\tilde{\varepsilon}}\mathbb{E}_{\tilde{P}_0^k}\big[\mathbb{I}_{\left(\log(\tilde{P}_0^k(\tilde{X}^k)/\tilde{P}_1^k(\tilde{X}^k))\geq\tilde{\varepsilon}\right)}\frac{\tilde{P}_1^k(\tilde{X}^k)}{\tilde{P}_0^k(\tilde{X}^k)}\big] \\ &= \mathbb{E}_{\tilde{P}_0^k}\Big[\mathbb{I}_{\left(\log(\tilde{P}_0^k(\tilde{X}^k)/\tilde{P}_1^k(\tilde{X}^k))\geq\tilde{\varepsilon}\right)}\Big(1 - e^{\tilde{\varepsilon}}\frac{\tilde{P}_1^k(\tilde{X}^k)}{\tilde{P}_0^k(\tilde{X}^k)}\Big)\Big] \\ &\leq \mathbb{E}[e^{\lambda Z - \lambda\tilde{\varepsilon} + \lambda\log\lambda - (\lambda+1)\log(\lambda+1)}]\,, \end{aligned} \tag{22}$$

where we use a random variable $Z \equiv \log(\tilde{P}_0^k(\tilde{X}_0^k)/\tilde{P}_1^k(\tilde{X}_0^k))$ and the last line follows from the fact that $\mathbb{I}_{(x\geq\tilde{\varepsilon})}(1 - e^{\tilde{\varepsilon}-x}) \leq e^{\lambda(x-\tilde{\varepsilon})+\lambda\log\lambda-(\lambda+1)\log(\lambda+1)}$ for any $\lambda \geq 0$.

Next, we give an upper bounds on the moment generating function of Z. From the definition of $\tilde{P}_0^{(\ell)}$ and $\tilde{P}_1^{(\ell)}$, $\mathbb{E}[e^{\lambda Z}] = \Big(\mathbb{E}_{\tilde{P}_0^{(\ell)}}[e^{\lambda\log(\tilde{P}_0^{(\ell)}(\tilde{X}_0^{(\ell)})/\tilde{P}_1^{(\ell)}(\tilde{X}_0^{(\ell)}))}]\Big)^k$. Let $\tilde{\varepsilon} = \sum_{\ell=1}^k(e^{\varepsilon_\ell}-1)\varepsilon_\ell/(e^{\varepsilon_\ell}+1) + \sqrt{2\sum_{\ell=1}^k\varepsilon_\ell^2\log\big(e + (\sqrt{\sum_{\ell=1}^k\varepsilon_\ell^2}/\tilde{\delta})\big)}$. Next we show that the k-fold composition is $(\tilde{\varepsilon}, 1-(1-\tilde{\delta})\prod_{\ell\in[k]}(1-\delta_\ell)\,)$-differentially private.

$$\mathbb{E}_{\tilde{P}_0^{(\ell)}}[e^{\lambda\log(P_0^{(\ell)}(X)/P_1^{(\ell)}(X))}] \;\leq\; e^{\frac{e^{\varepsilon_\ell}-1}{e^{\varepsilon_\ell}+1}\lambda\varepsilon_\ell + \frac{1}{2}\lambda^2\varepsilon_\ell{}^2}\,,$$

for any λ. Substituting this into (22) with a choice of $\lambda = \frac{\tilde{\varepsilon} - \sum_{\ell\in[k]}\varepsilon_\ell(e^{\varepsilon_\ell}-1)/(e^{\varepsilon_\ell}+1)}{\sum_{\ell\in[k]}\varepsilon_\ell^2}$, we get

$$\begin{aligned} \tilde{P}_0^k(B_1) - e^{\tilde{\varepsilon}}\tilde{P}_1^k(B_1) &\leq \frac{1}{1 + \frac{\tilde{\varepsilon} - \sum_{\ell\in[k]}\varepsilon_\ell(e^{\varepsilon_\ell}-1)/(e^{\varepsilon_\ell}+1)}{\sum_{\ell\in[k]}\varepsilon_\ell^2}}\exp\Big\{-\frac{1}{2\sum_{\ell\in[k]}\varepsilon_\ell^2}\Big(\tilde{\varepsilon} - \sum_{\ell\in[k]}\varepsilon_\ell\frac{e^{\varepsilon_\ell}-1}{e^{\varepsilon_\ell}+1}\Big)^2\Big\} \\ &\leq \;. \end{aligned}$$

Substituting $\tilde{\varepsilon}$, we get the desired bound.

Similarly, we can prove that with $\tilde{\varepsilon} = \sum_{\ell=1}^k(e^{\varepsilon_\ell}-1)\varepsilon_\ell/(e^{\varepsilon_\ell}+1) + \sqrt{2\sum_{\ell=1}^k\varepsilon_\ell^2\log\big(1/\tilde{\delta}\big)}$, the desired bound also holds.

8 Proofs

8.1 Proof of Theorem 2.4

Consider hypothesis testing between D_1 and D_2. If there is a point $(P_{\mathrm{MD}}, P_{\mathrm{FA}})$ achieved by M' but not by M, then we claim that this is a contradiction to the assumption that D–X–Y form a Markov chain. Consider a decision maker who have only access to the output of M. Under the Markov chain assumption, he can simulate the output of M' by generating a random variable Y conditioned on $M(D)$ and achieve every point in the privacy region of M' (cf. Theorem 2.2). Hence, the privacy region of M' must be included in the privacy region of M.

8.2 Proof of Theorem 2.1

First we prove that (ε, δ)-differential privacy implies (1). From the definition of differential privacy, we know that for all rejection set $S \subseteq \mathcal{X}$, $\mathbb{P}(M(D_0) \in \bar{S}) \leq e^{\varepsilon}\mathbb{P}(M(D_1) \in \bar{S}) + \delta$. This implies $1 - P_{\mathrm{FA}}(D_0, D_1, M, S) \leq e^{\varepsilon} P_{\mathrm{MD}}(D_0, D_1, M, S) + \delta$. This implies the first inequality of (1), and the second one follows similarly.

The converse follows analogously. For any set S, we assume $1 - P_{\mathrm{FA}}(D_0, D_1, M, S) \leq e^{\varepsilon} P_{\mathrm{MD}}(D_0, D_1, M, S) + \delta$. Then, it follows that $\mathbb{P}(M(D_0) \in \bar{S}) \leq e^{\varepsilon}\mathbb{P}(M(D_1) \in \bar{S}) + \delta$ for all choices of $S \subseteq \mathcal{X}$. Together with the symmetric condition $\mathbb{P}(M(D_1) \in \bar{S}) \leq e^{\varepsilon}\mathbb{P}(M(D_0) \in \bar{S}) + \delta$, this implies (ε, δ)-differential privacy.

8.3 Proof of Remark 2.2

We have a decision rule γ represented by a partition $\{S_i\}_{i\in\{1,\ldots,N\}}$ and corresponding accept probabilities $\{p_i\}_{i\in\{1,\ldots,N\}}$, such that if the output is in a set S_i, we accept with probability p_i. We assume the subsets are sorted such that $1 \geq p_1 \geq \ldots \geq p_N \geq 0$. Then, the probability of false alarm is

$$P_{\mathrm{FA}}(D_0, D_1, M, \gamma) \;=\; \sum_{i=1}^{N} p_i \, \mathbb{P}(M(D_0) \in S_i) \;=\; p_N + \sum_{i=2}^{N} (p_{i-1} - p_i) \, \mathbb{P}(M(D_0) \in \cup_{j<i} S_j) \; .$$

and similarly, $P_{\mathrm{MD}}(D_0, D_1, M, \gamma) = (1 - p_1) + \sum_{i=2}^{N} (p_{i-1} - p_i) \, \mathbb{P}(M(D_1) \notin \cup_{j<i} S_j)$. Recall that $P_{\mathrm{FA}}(D_0, D_1, M, S) = \mathbb{P}(M(D_0) \in S)$ and $P_{\mathrm{MD}}(D_0, D_1, M, S) = \mathbb{P}(M(D_1) \in \bar{S})$. So for any decision rule γ, we can represent the pair $(P_{\mathrm{MD}}, P_{\mathrm{FA}})$ as a convex combination:

$$\big(P_{\mathrm{MD}}(D_0, D_1, M, \gamma), P_{\mathrm{FA}}(D_0, D_1, M, \gamma) \big) \;=\; \sum_{i=1}^{N+1} (p_{i-1} - p_i) \big(P_{\mathrm{MD}}(D_0, D_1, M, \cup_{j<i} S_j), P_{\mathrm{FA}}(D_0, D_1, M, \cup_{j<i} S_j) \big)$$

where we used $p_0 = 1$ and $p_{N+1} = 0$, and hence it is included in the convex hull of the privacy region achieved by decision rules with hard thresholding.

9 Acknowledgement

The authors thank Maxim Raginsky for helpful discussions and for pointing out [Bla53], and Moritz Hardt for pointing out an error in an earlier version of this paper.

A Examples illustrating the strengths of graphical representation of differential privacy

Remark A.1. *The following statements are true.*

(a) If a mechanism is (ε, δ)-differentially private, then it is $(\tilde{\varepsilon}, \tilde{\delta})$-differentially private for all pairs of $\tilde{\varepsilon}$ and $\tilde{\delta} \geq \delta$ satisfying

$$\frac{1-\delta}{1+e^{\varepsilon}} \;\; \geq \;\; \frac{1-\tilde{\delta}}{1+e^{\tilde{\varepsilon}}} \, .$$

(b) For a pair of neighboring databases D and D', and all (ε, δ)-differentially private mechanisms, the total variation distance defined as $\|M(D) - M(D')\|_{\mathrm{TV}} = \max_{S \subseteq \mathcal{X}} \mathbb{P}(M(D') \in S) - \mathbb{P}(M(D) \in S)$ is bounded by

$$\sup_{(\varepsilon,\delta)\text{-}\textit{differentially private } M} \|M(D) - M(D')\|_{\mathrm{TV}} \;\; \leq \;\; 1 - \frac{2(1-\delta)}{1+e^{\varepsilon}} \, .$$

Proof. **Proof of (*a*).** From Figure 1, it is immediate that $\mathcal{R}(\varepsilon, \delta) \subseteq \mathcal{R}(\tilde{\varepsilon}, \tilde{\delta})$ when the conditions are satisfied. Then, for a (ε, δ)-private M, it follows from $\mathcal{R}(M) \subseteq \mathcal{R}(\varepsilon, \delta) \subseteq \mathcal{R}(\tilde{\varepsilon}, \tilde{\delta})$ that M is $(\tilde{\varepsilon}, \tilde{\delta})$-differentially private.

Proof of (*b*). By definition, $\|M(D) - M(D')\|_{\mathrm{TV}} = \max_{S \subseteq \mathcal{X}} \mathbb{P}(M(D') \in S) - \mathbb{P}(M(D) \in S)$. Letting S be the rejection region in our hypothesis testing setting, the total variation distance is defined by the following optimization problem:

$$\begin{aligned} \max_{S} \quad & 1 - P_{\mathrm{MD}}(S) - P_{\mathrm{FA}}(S) \\ \text{subject to} \quad & (P_{\mathrm{MD}}(S), P_{\mathrm{FA}}(S)) \in \mathcal{R}(\varepsilon, \delta), \text{ for all } S \subseteq \mathcal{X} \, . \end{aligned} \tag{23}$$

From Figure 1 it follows immediately that the total variation distance cannot be larger than $\delta + (1-\delta)(e^{\varepsilon} - 1)/(e^{\varepsilon} + 1)$. □

B Analysis of the Gaussian mechanism in Theorem 4.3

Following the analysis in Section 6, we know that the privacy region of a composition of mechanisms is described by a set of (ε, δ) pairs that satisfy the following:

$$\delta \;\; = \;\; \mu_0^k(B) - e^{\varepsilon} \mu_1^k(B) \, ,$$

where μ_0^k and μ_1^k are probability measures of the mechanism under k-fold composition when the data base is D_0 and D_1 respectively, and the subset $B = \arg\max_{S \subseteq \mathbb{R}^k} \mu_0^k(S) - e^{\varepsilon} \mu_1^k(S)$.

In the case of Gaussian mechanisms, we can assume without loss of generality that D_0 is such that $q_i(D_0) = 0$ and D_1 is such that $q_i(D_1) = \Delta$ for all $i \in \{1, \ldots, k\}$. When adding Gaussian noises with variances σ^2, we want to ask how small the variance can be and still ensure (ε, δ)-differential privacy under k-fold composition.

Let $f_0^k(x_1,\ldots,x_k) = \prod_{i=1}^k f_0(x_i) = (1/\sqrt{2\pi\sigma^2})^k e^{-\sum_{i=1}^k x_i^2/2\sigma^2}$ and $f_1^k(x_1\ldots,x_k) = \prod_{i=1}^k f_1(x_i) = (1/\sqrt{2\pi\sigma^2})^k e^{-\sum_{i=1}^k (x_i-\Delta)^2/2\sigma^2}$ be the probability density functions of Gaussians centered at zero and $\Delta\mathbb{1}_k$ respectively. Using a similar technique as in (20), we know that

$$\begin{aligned}
\mu_0^k(B) - e^\varepsilon \mu_1^k(B) &= \mathbb{E}_{\mu_0^k}\big[\mathbb{I}_{\left(\log(f_0^k(\tilde{X}^k)/f_1^k(\tilde{X}^k))\geq\varepsilon\right)}\big] - e^\varepsilon \mathbb{E}_{\mu_0^k}\big[\mathbb{I}_{\left(\log(f_0^k(\tilde{X}^k)/f_1^k(\tilde{X}^k))\geq\varepsilon\right)}\frac{f_1^k(\tilde{X}^k)}{f_0^k(\tilde{X}^k)}\big] \\
&= \mathbb{E}_{\mu_0^k}\Big[\mathbb{I}_{\left(\log(f_0^k(\tilde{X}^k)/f_1^k(\tilde{X}^k))\geq\varepsilon\right)}\Big(1 - e^\varepsilon \frac{f_1^k(\tilde{X}^k)}{f_0^k(\tilde{X}^k)}\Big)\Big] \\
&\leq \mathbb{E}[e^{\lambda Z - \lambda\varepsilon + \lambda\log\lambda - (\lambda+1)\log(\lambda+1)}]\,, \qquad (24)
\end{aligned}$$

where $\tilde{X}^k$ is a random vector distributed according to μ_0^k, $Z \equiv \log(f_0^k(\tilde{X}^k)/f_1^k(\tilde{X}^k))$, and the last line follows from $\mathbb{I}_{(x\geq\varepsilon)}(1-e^{\varepsilon-x}) \leq e^{\lambda(x-\varepsilon)+\lambda\log\lambda-(\lambda+1)\log(\lambda+1)}$ for any $\lambda \geq 0$.

Next, we give an upper bound on the moment generating function of Z.

$$\begin{aligned}
\mathbb{E}_{\mu_0}[e^{\lambda\log(f_0(X)/f_1(X))}] &= \mathbb{E}[e^{-\lambda\Delta X/\sigma^2}]e^{\lambda\Delta^2/2\sigma^2} \\
&\leq e^{(\Delta^2/2\sigma^2)\lambda^2 + (\Delta^2/2\sigma^2)\lambda}\,,
\end{aligned}$$

for any $\lambda \geq 0$. Substituting this into (24) with a choice of $\lambda = \frac{\sigma^2}{k\Delta^2}\big(\varepsilon - \frac{k\Delta^2}{2\sigma^2}\big)$, which is positive for $\varepsilon > k\Delta^2/2\sigma^2$, we get

$$\begin{aligned}
\mu_0^k(B) - e^\varepsilon \mu_1^k(B) &\leq \exp\Big\{(k\Delta^2/2\sigma^2)\lambda^2 + (k\Delta^2/2\sigma^2)\lambda - \varepsilon\lambda + \lambda\log\lambda - (\lambda+1)\log(\lambda+1)\Big\} \\
&\leq \frac{1}{1+\frac{\sigma^2}{k\Delta^2}\big(\varepsilon - \frac{k\Delta^2}{2\sigma^2}\big)} \exp\Big\{-\frac{\sigma^2}{2k\Delta^2}\Big(\varepsilon - \frac{k\Delta^2}{2\sigma^2}\Big)^2\Big\} \\
&\leq \frac{1}{1+\sqrt{\frac{2\sigma^2}{k\Delta^2}\log(e + \frac{1}{\delta}\sqrt{\frac{k\Delta^2}{\sigma^2}})}} \frac{1}{e + \frac{1}{\delta}\sqrt{\frac{k\Delta^2}{\sigma^2}}} \\
&\leq \frac{1}{\sqrt{\frac{k\Delta^2}{\sigma^2}} + \sqrt{2\log(e + (1/\delta)\sqrt{k\Delta^2/\sigma^2})}} \frac{\delta}{e\delta\sqrt{\frac{\sigma^2}{k\Delta^2}+1}}\,,
\end{aligned}$$

for our choice of σ^2 such that $\varepsilon \geq k\Delta^2/(2\sigma^2) + \sqrt{(2k\Delta^2/\sigma^2)\log(e + (1/\delta)\sqrt{k\Delta^2/\sigma^2})}$. The right-hand side is always less than δ.

With $\sigma^2 \geq (4k\Delta^2/\varepsilon^2)\log(e + (\varepsilon/\delta))$ and $\sigma^2 \geq k\Delta^2/(4\varepsilon)$, this ensures that the above condition is satisfied. This implies that we only need $\sigma^2 = O((k\Delta^2/\varepsilon^2)\log(e + (\varepsilon/\delta)))$.

C Analysis of the geometric mechanism in Theorem 4.4

Theorem 4.4 follows directly from the proof of Theorem 3.3, once the appropriate associations are made. Consider two databases D_0 and D_1, and a single query q such that $q(D_1) = q(D_0) + 1$. The geometric mechanism produces two random outputs $q(D_0) + Z$ and $q(D_1) + Z$ where Z is distributed accruing to the geometric distribution. Let P_0 and P_1 denote the distributions of the random output respectively. For $x \leq q(D_0)$, $P_0(x) = e^\varepsilon P_1(x)$, and for $x > q(D_0)$, $e^\varepsilon P_0(x) = P_1(x)$. Then, it is not difficult to see that the privacy region achieved by the geometric mechanism is equal

to the privacy region achieved by the canonical binary example of $\tilde{X}_0$ and $\tilde{X}_1$ in (12) and (13) with $\delta = 0$. This follows from the fact there is a stochastic transition from the pair $\tilde{X}_0$ and $\tilde{X}_1$ to $q(D_0) + Z$ and $q(D_1) + Z$; further, the converse is also true. Hence, from the perspective of hypothesis testing, those two (pairs of) outcomes are equivalent.

It now follows from the proof of Theorem 3.3 that the k-fold composition privacy region is exactly the optimal privacy region described in (5) with $\delta = 0$. We also know that this is the largest possible privacy region achieved by a class of $(\varepsilon, 0)$-differentially private mechanisms.

D Analysis of Johnson-Lindenstrauss mechanism

For cut queries, Johnson-Lindenstrauss mechanism proceeds as follows:

Johnson-Lindenstrauss mechanism for cut queries [BBDS12]

Input: A n-node graph G, parameters $\varepsilon, \delta, \eta, \nu > 0$

Output: An approximate Laplacian of G: $\tilde{L}$

1: Set $r = 8\log(2/\nu)/\nu^2$ and $w = \sqrt{32r\log(2/\delta)}\log(4r/\delta)/\varepsilon$

2: For every pair of nodes $I \neq j$, set new weights $w_{i,j} = w/n + (1 - w/n)w_{i,j}$

3: Randomly draw a matrix N of size $r \times \binom{n}{2}$, whose entries are i.i.d. samples of $\mathcal{N}(0, 1)$

4: Output $\tilde{L} = (1/r)E_G^T N^T N E_G$,
where E_G is an $\binom{n}{2} \times n$ matrix whose (i, j)-th row is $\sqrt{w_{i,j}}(e_i - e_j)$

Here e_i is the standard basis vector with one in the i-th entry. Given this synopsis of the sanitized graph Laplacian, a cut query $q(G, S)$ returns $1/(1 - w/n)(\mathbb{1}_S^T \tilde{L} \mathbb{1}_S - w|S|(n - |S|)/n)$, where $\mathbb{1}_S \in \{0, 1\}^n$ is the indicator vector for the set S. If the matrix N is an identity matrix, this returns the correct cut value of G.

We have the choice of $w \in \mathbb{R}$ and $r \in \mathbb{Z}$ to ensure that the resulting mechanism is (ε, δ)-differentially private, and satisfy (η, τ, ν)-approximation guarantees of (9). We utilize the following lemma from [BBDS12].

Lemma 1. *With the choice of*

$$w \;=\; \frac{4}{\varepsilon_0}\log(2/\delta_0) \quad \textit{and} \quad r \;=\; \frac{8\log(2/\nu)}{\eta^2}\,,$$

each row of NE_G satisfy $(\varepsilon_0, \delta_0)$-differential privacy, and the resulting Johnson-Lindenstrauss mechanism satisfy (η, τ, ν)-approximation guarantee with

$$\tau \;=\; 2|S|\,\eta\, w\,,$$

where $|S|$ is the size of the smaller partition S of the cut $(S, \bar{S})$.

The error bound in (10) follows from choosing

$$\varepsilon_0 \;=\; \frac{\varepsilon}{\sqrt{4r\log(2/\delta)}} \quad \text{and} \quad \delta_0 \;=\; \frac{\delta}{2r}\,,$$

and applying Theorem 3.2 to ensure that the resulting mechanism with r-composition of the r rows of ME_G is (ε, δ)-differentially private. Here it is assumed that $\varepsilon < 1$.

Now, with Theorem 3.3, we do not require ε_0 to be as small, which in turn allows us to add smaller noise w, giving us an improved error bound on τ. Precisely, using Theorem 3.4 it follows that a choice of

$$\varepsilon_0 \;=\; \frac{\varepsilon}{\sqrt{4r\log(e+2\varepsilon/\delta)}} \quad \text{and} \quad \delta_0 \;=\; \frac{\delta}{2r}\,,$$

suffices to ensure that after r-composition we get (ε,δ)-differential privacy. Resulting noise is bounded by $w \leq 4\sqrt{4r\log(e+2\varepsilon/\delta)}\log(4r/\delta)/\varepsilon$, which gives the error bound in (11). The proof follows analogously for the matrix variance queries.

References

[AS04] Noga Alon and Joel H Spencer, *The probabilistic method*, Wiley. com, 2004.

[BBDS12] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet, *The johnson-lindenstrauss transform itself preserves differential privacy*, Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on, IEEE, 2012, pp. 410–419.

[BDMN05] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim, *Practical privacy: the SuLQ framework*, Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, ACM, 2005, pp. 128–138.

[Bla53] David Blackwell, *Equivalent comparisons of experiments*, The Annals of Mathematical Statistics **24** (1953), no. 2, 265–272.

[Bla65] N Blachman, *The convolution inequality for entropy powers*, Information Theory, IEEE Transactions on **11** (1965), no. 2, 267–271.

[BLR13] Avrim Blum, Katrina Ligett, and Aaron Roth, *A learning theory approach to noninteractive database privacy*, Journal of the ACM (JACM) **60** (2013), no. 2, 12.

[CT88] Thomas M Cover and A Thomas, *Determinant inequalities via information theory*, SIAM journal on Matrix Analysis and Applications **9** (1988), no. 3, 384–392.

[CT12] Thomas M Cover and Joy A Thomas, *Elements of information theory*, John Wiley & Sons, 2012.

[DCT91] Amir Dembo, Thomas M Cover, and Joy A Thomas, *Information theoretic inequalities*, Information Theory, IEEE Transactions on **37** (1991), no. 6, 1501–1518.

[DKM+06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor, *Our data, ourselves: Privacy via distributed noise generation*, Advances in Cryptology-EUROCRYPT 2006, Springer, 2006, pp. 486–503.

[DL09] Cynthia Dwork and Jing Lei, *Differential privacy and robust statistics*, Proceedings of the 41st annual ACM symposium on Theory of computing, ACM, 2009, pp. 371–380.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith, *Calibrating noise to sensitivity in private data analysis*, Theory of Cryptography, Springer, 2006, pp. 265–284.

[DN03] Irit Dinur and Kobbi Nissim, *Revealing information while preserving privacy*, Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, ACM, 2003, pp. 202–210.

[DN04] Cynthia Dwork and Kobbi Nissim, *Privacy-preserving datamining on vertically partitioned databases*, Advances in Cryptology–CRYPTO 2004, Springer, 2004, pp. 528–544.

[DRV10] Cynthia Dwork, Guy N Rothblum, and Salil Vadhan, *Boosting and differential privacy*, Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on, IEEE, 2010, pp. 51–60.

[Dwo06] Cynthia Dwork, *Differential privacy*, Automata, languages and programming, Springer, 2006, pp. 1–12.

[GRS12] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan, *Universally utility-maximizing privacy mechanisms*, SIAM Journal on Computing **41** (2012), no. 6, 1673–1693.

[GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman, *Iterative constructions and private data release*, Theory of Cryptography, Springer, 2012, pp. 339–356.

[GT04] Ben Green and Terence Tao, *The primes contain arbitrarily long arithmetic progressions*, arXiv preprint math/0404188 (2004).

[GV12] Quan Geng and Pramod Viswanath, *Optimal noise-adding mechanism in differential privacy*, arXiv preprint arXiv:1212.1186 (2012).

[GV13] ———, *The optimal mechanism in (ϵ, δ)-differential privacy*, arXiv preprint arXiv:1305.1330 (2013).

[HLM10] Moritz Hardt, Katrina Ligett, and Frank McSherry, *A simple and practical algorithm for differentially private data release*, arXiv preprint arXiv:1012.4763 (2010).

[HR10] Moritz Hardt and Guy N Rothblum, *A multiplicative weights mechanism for privacy-preserving data analysis*, Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on, IEEE, 2010, pp. 61–70.

[HR13] Moritz Hardt and Aaron Roth, *Beyond worst-case analysis in private singular vector computation*, Proceedings of the 45th annual ACM symposium on Symposium on theory of computing, ACM, 2013, pp. 331–340.

[HT10] Moritz Hardt and Kunal Talwar, *On the geometry of differential privacy*, Proceedings of the 42nd ACM symposium on Theory of computing, ACM, 2010, pp. 705–714.

[Lau96] S. L. Lauritzen, *Graphical Models*, Oxford University Press, 1996.

[LV07] Tie Liu and Pramod Viswanath, *An extremal inequality motivated by multiterminal information-theoretic problems*, Information Theory, IEEE Transactions on **53** (2007), no. 5, 1839–1851.

[MN12] S Muthukrishnan and Aleksandar Nikolov, *Optimal private halfspace counting via discrepancy*, Proceedings of the 44th symposium on Theory of Computing, ACM, 2012, pp. 1285–1292.

[MT07] Frank McSherry and Kunal Talwar, *Mechanism design via differential privacy*, Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on, IEEE, 2007, pp. 94–103.

[RTTV08] Omer Reingold, Luca Trevisan, Madhur Tulsiani, and Salil Vadhan, *Dense subsets of pseudorandom sets*, Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on, IEEE, 2008, pp. 76–85.

[Sta59] AJ Stam, *Some inequalities satisfied by the quantities of information of fisher and shannon*, Information and Control **2** (1959), no. 2, 101–112.

[TZ08] Terence Tao and Tamar Ziegler, *The primes contain arbitrarily long polynomial progressions*, Acta Mathematica **201** (2008), no. 2, 213–305.

[VG06] Sergio Verdú and Dongning Guo, *A simple proof of the entropy-power inequality*, IEEE Transactions on Information Theory **52** (2006), no. 5, 2165–2166.

[WZ10] Larry Wasserman and Shuheng Zhou, *A statistical framework for differential privacy*, Journal of the American Statistical Association **105** (2010), no. 489, 375–389.

[Zam98] Ram Zamir, *A proof of the fisher information inequality via a data processing argument*, Information Theory, IEEE Transactions on **44** (1998), no. 3, 1246–1250.