Supplementary Material

A Proofs and Extensions for "Local Pseudo-Randomizers"

Lemma A.1 (Lemma 3.2 restated)**.** *For a t-samplable deletion ε-DP local randomizer $\mathcal{R}\colon X \to Y$ and $G\colon \{0,1\}^\ell \to \{0,1\}^t$, let $\mathcal{D}$ denote the following family of tests which take $r' \in \{0,1\}^t$ as an input:*

$$\mathcal{D} := \left\{ \texttt{ind}\left(\frac{\mathbf{Pr}[\mathcal{R}(x) = \mathcal{R}_\emptyset(r')]}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r) = \mathcal{R}_\emptyset(r')]} \geq \theta \right) \;\middle|\; x \in X, \theta \in [0, e^\varepsilon] \right\},$$

where ind *denotes the $\{0,1\}$ indicator function of a condition. If G β-fools $\mathcal{D}$ for $\beta < 1/(2e^\varepsilon)$ then $R[G]$ is a deletion $(\varepsilon + 2e^\varepsilon\beta)$-DP local randomizer. Furthermore, for every $\gamma > 0$, $\mathcal{R}[G,\gamma]$ is a deletion $(\varepsilon + 2e^\varepsilon\beta)$-DP local randomizer.*

Proof. We demonstrate that if $\mathcal{R}[G]$ is not a deletion $(\varepsilon+2e^\varepsilon\beta)$-DP randomizer then there exists a test in $\mathcal{D}$ that distinguishes the output of G from true randomness that succeeds with probability at least β. To analyze the privacy guarantees of $\mathcal{R}[G]$ we let the reference distribution ρ_G be the uniform distribution over $\{0,1\}^\ell$. For brevity, for $y \in Y$ we denote the density ratio of $\mathcal{R}(x)$ to ρ at y by

$$\pi_x(y) := \frac{\mathbf{Pr}[\mathcal{R}(x) = y]}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r) = y]}.$$

Then, $\mathcal{R}[G]$ outputs a seed s with probability:

$$\mu_x(s) := \frac{\pi_x(\mathcal{R}_\emptyset(G(s)))}{\sum_{s'\in\{0,1\}^\ell} \pi_x(\mathcal{R}_\emptyset(G(s')))}.$$

By definition of our reference distribution, $\rho_G(s) = 2^{-\ell}$ for all s. Therefore

$$\frac{\mu_x(s)}{\rho_G(s)} = \frac{\mu_x(s)}{2^{-\ell}} = \frac{\pi_x(\mathcal{R}_\emptyset(G(s)))}{\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))]}.$$

We observe that, by the fact that $\mathcal{R}(x)$ is ε-DP we have that $\pi_x(\mathcal{R}_\emptyset(G(s))) \in [e^{-\varepsilon}, e^\varepsilon]$. Therefore, to show that $\mathcal{R}[G]$ is $(\varepsilon + 2e^\varepsilon\beta)$-DP, it suffices to show that the denominator is in the range $[e^{-2e^\varepsilon\beta}, e^{2e^\varepsilon\beta}]$. To show this, we assume for the sake of contradiction that it is not true. Namely, that either

$$\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] > e^{2e^\varepsilon\beta} > 1 + e^\varepsilon\beta$$

or

$$\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] < e^{-2e^\varepsilon\beta} < 1 - e^\varepsilon\beta,$$

where we used the assumption that $\beta < 1/(2e^\varepsilon)$ in the second inequality.

We first deal with the case when $\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] > 1 + e^\varepsilon\beta$ (as the other case will be essentially identical). Observe that for true randomness we have:

$$\mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r'))] = \mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}\left[\frac{\mathbf{Pr}[\mathcal{R}(x) = \mathcal{R}_\emptyset(r')]}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r) = \mathcal{R}_\emptyset(r')]}\right] = 1.$$

Using the fact that $\pi_x(y) \in [0, e^\varepsilon]$ we have that

$$\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] = \int_0^{e^\varepsilon} \mathop{\mathbf{Pr}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s'))) \geq \theta] d\theta$$

and, similarly,

$$\mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r'))] = \int_0^{e^\varepsilon} \mathop{\mathbf{Pr}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta]\, d\theta.$$

Thus, by our assumption,

$$\int_0^{e^\varepsilon} \left(\mathop{\mathbf{Pr}}_{s' \sim \{0,1\}^\ell} [\pi_x(\mathcal{R}_\emptyset(G(s'))) \geq \theta] - \mathop{\mathbf{Pr}}_{r' \sim \{0,1\}^t} [\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta] \right) d\theta > e^\varepsilon \beta.$$

This implies that there exists $\theta \in [0, e^\varepsilon]$ such that

$$\mathop{\mathbf{Pr}}_{s' \sim \{0,1\}^\ell} [\pi_x(\mathcal{R}_\emptyset(G(s'))) \geq \theta] - \mathop{\mathbf{Pr}}_{r' \sim \{0,1\}^t} [\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta] > \beta.$$

Note that $\mathtt{ind}\left(\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta\right) \in \mathcal{D}$, for all $x \in X$ and $\theta \in [0, e^\varepsilon]$ contradicting our assumption on G. Thus we obtain that $\mathbf{E}_{s' \sim \{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] \leq 1 + e^\varepsilon \beta$. We can arrive at a contradiction in the case when $\mathbf{E}_{s' \sim \{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] < 1 - e^\varepsilon \beta$ in exactly the same way.

To show that $\mathcal{R}[G, \gamma]$ is a deletion $(\varepsilon + 2e^\varepsilon \beta)$-DP local randomizer we observe that for every x, conditioned on accepting one of the samples, $\mathcal{R}[G, \gamma](x)$ outputs a sample distributed exactly according to $\mathcal{R}[G](x)$. If $\mathcal{R}[G, \gamma](x)$ does not accept any samples than it samples from the reference distribution ρ_G. Thus given that $\mathcal{R}[G](x)$ is $(\varepsilon + 2e^\varepsilon \beta)$-close to ρ_G we have that the output distribution $\mathcal{R}[G, \gamma](x)$ is also $(\varepsilon + 2e^\varepsilon \beta)$-close to ρ_G. □

As the first step for proving Lemma 3.3 we show that when used with the identity G, the resulting randomizer is γ-close in total variation distance to $\mathcal{R}$.

Lemma A.2. *Let $\mathcal{R}$ be a deletion ε-DP t-samplable local randomizer. Then for the identity function $\mathrm{ID}_t \colon \{0,1\}^t \to \{0,1\}^t$ and any $\gamma > 0$ we have that $\mathcal{R}[\mathrm{ID}_t, \gamma]$ is a deletion ε-DP local randomizer and for every $x \in \mathcal{X}$, $\mathrm{TV}(\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t, \gamma](x)), \mathcal{R}(x)) \leq \gamma$.*

Proof. When applied with $G = \mathrm{ID}_t$, y is distributed according to the reference distribution of $\mathcal{R}$. Thus the algorithm performs standard rejection sampling until it accepts a sample or exceeds the bound J on the number of steps. Note that deletion DP implies that $\frac{\mathbf{Pr}[\mathcal{R}(x)=y]}{e^\varepsilon \mathbf{Pr}_{r \sim \{0,1\}^t}[\mathcal{R}_\emptyset(r)=y]} \leq 1$. At each step, conditioned on success the algorithm samples s such that $\mathcal{R}_\emptyset(s)$ is distributed identically to $\mathcal{R}(x)$. Further, the acceptance probability at each step is

$$\mathop{\mathbf{E}}_{y \sim \rho} \left[\frac{\mathbf{Pr}[\mathcal{R}(x) = y]}{e^\varepsilon \mathbf{Pr}_{r \sim \{0,1\}^t}[\mathcal{R}_\emptyset(r) = y]} \right] = \sum_{y \in Y} \frac{\mathbf{Pr}[\mathcal{R}(x) = y]}{e^\varepsilon} = \frac{1}{e^\varepsilon}.$$

Thus the probability that all the steps reject is $\leq (1 - e^{-\varepsilon})^J \leq \gamma$. This implies that $\mathrm{TV}(\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t, \gamma](x)), \mathcal{R}(x)) \leq \gamma$. □

We can now state the implications of using a sufficiently strong PRG on the output of the randomizer.

Lemma A.3 (Lemma 3.3 restated)**.** *Let $\mathcal{R}$ be a deletion ε-DP t-samplable local randomizer, let $G \colon \{0,1\}^\ell \to \{0,1\}^t$ be (T, β)-PRG. Let $T(\mathcal{R}, G, \gamma)$ denote the running time of $\mathcal{R}[G, \gamma]$ and assume that $T > T(\mathcal{R}, G, \gamma)$. Then for all $x \in X$, $\mathcal{R}_\emptyset(G(\mathcal{R}[G, \gamma](x)))$ is (T', β')-computationally indistinguishable from $\mathcal{R}(x)$, where $\beta' = \gamma + e^\varepsilon \ln(1/\gamma)\beta$ and $T' = T - T(\mathcal{R}, G, \gamma)$.*

Proof. By Lemma A.2, $\mathrm{TV}(\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t, \gamma](x)), \mathcal{R}(x)) \leq \gamma$ and thus it is sufficient to prove that $\mathcal{R}_\emptyset(G(\mathcal{R}[G, \gamma](x)))$ is $(T', e^\varepsilon \ln(1/\gamma)\beta)$-computationally indistinguishable from $\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t, \gamma](x))$. To prove this, assume for the sake of contradiction, that there exists a test D' running in time T' such that for some x,

$$\big| \mathbf{Pr}[D'(\mathcal{R}_\emptyset(G(\mathcal{R}[G, \gamma](x)))) = 1] - \mathbf{Pr}[D'(\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t, \gamma](x))) = 1] \big| \geq e^\varepsilon \ln(1/\gamma)\beta.$$

Then we claim that there exists a test for distinguishing $G(s)$ for $s \sim \{0,1\}^\ell$ from a truly random seed $r \sim \{0,1\}^t$. Note that $\mathcal{R}_\emptyset(G(\mathcal{R}[G, \gamma]))$ can be seen as $\mathcal{R}[G, \gamma]$ that outputs directly $y = \mathcal{R}_\emptyset(G(s))$ instead of s itself. Next we observe that $\mathcal{R}_\emptyset(G(\mathcal{R}[G, \gamma]))$ uses the output of G at most $J = e^\varepsilon \ln(1/\gamma)$ times in place of truly random t-bit string used by $\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t, \gamma])$. Thus, by the standard hybrid argument, one of those applications can be used to test G with success probability at least $e^\varepsilon \ln(1/\gamma)\beta/J = \beta$. This test requires running a hybrid between $\mathcal{R}_\emptyset(G(\mathcal{R}[G, \gamma]))$ and $\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t, \gamma])$ in addition to D' itself. We can assume that sampling a fresh t bits takes less time than sampling ℓ bits and applying G and therefore the running time of the resulting test is at most $T' + T(\mathcal{R}, G, \gamma) = T$. Thus we obtain a contradiction to G being (T, β)-PRG. □

Theorem A.4 (Theorem 3.4 restated)**.** *Let $\mathcal{R}$ be a deletion ε-DP t-samplable local randomizer, let $G\colon \{0,1\}^\ell \to \{0,1\}^t$ be (T,β)-PRG for $\beta < 1/(2e^\varepsilon)$. Let $T(\mathcal{R},G,\gamma)$ be the running time of $\mathcal{R}[G,\gamma]$ and assume that $T > T(\mathcal{R},G,\gamma)$. Then $\mathcal{R}[G,\gamma]$ is a deletion $(\varepsilon + 2e^\varepsilon\beta)$-DP local randomizer and for all $x \in X$, $\mathcal{R}_\emptyset(G(\mathcal{R}[G,\gamma](x)))$ is (T',β')-computationally indistinguishable from $\mathcal{R}(x)$, where $\beta' = \gamma + e^\varepsilon \ln(1/\gamma)\beta$ and $T' = T - T(\mathcal{R},G,\gamma)$.*

Proof. The second part of the claim is exactly Lemma 3.3. To see the first part of the claim note that by our assumption $T > T(\mathcal{R},G,\gamma)$ and computation of the ratio of densities $\frac{\mathbf{Pr}[\mathcal{R}(x)=\mathcal{R}_\emptyset(r')]}{e^\varepsilon \mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r)=\mathcal{R}_\emptyset(r')]}$ for any $r' \in \{0,1\}^t$ is part of $\mathcal{R}[G,\gamma]$. This implies that the test family $\mathcal{D}$ defined in Lemma 3.2 can be computed in time T. Now applying Lemma 3.2 gives us the privacy claim. □

Lemma A.5 (Lemma 3.5 restated)**.** *Let $\mathcal{R}$ be a deletion ε-DP t-samplable local randomizer, let $G\colon \{0,1\}^\ell \to \{0,1\}^t$ be an arbitrary function. Then $\mathcal{R}[G,\gamma]$ is a deletion 2ε-DP local randomizer.*

Proof. As in the proof of Lemma 3.2 we observe that if we take the reference distribution to be uniform over $\{0,1\}^\ell$ we will get that, conditioned on accepting a sample, the seed s is output with probability $\mu_x(s)$ such that

$$\frac{\mu_x(s)}{\rho_G(s)} = \frac{\mu_x(s)}{2^{-\ell}} = \frac{\pi_x(\mathcal{R}_\emptyset(G(s)))}{\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))]}.$$

By the fact that $\mathcal{R}(x)$ is ε-DP we have that for every $s' \in \{0,1\}^\ell$, $\pi_x(\mathcal{R}_\emptyset(G(s'))) \in [e^{-\varepsilon}, e^\varepsilon]$ and thus $\frac{\mu_x(s)}{\rho_G(s)} \in [e^{-2\varepsilon}, e^{2\varepsilon}]$. □

A.1 Extension to Replacement DP

We now show that the same approach can be used to compress a replacement ε_r-DP randomizer $\mathcal{R}$. To do this we first let ρ be some reference distribution relative to which $\mathcal{R}$ is deletion ε-DP for some $\varepsilon \le \varepsilon_r$. One possible way to define ρ is to pick some fixed $x_0 \in X$ and let ρ be the distribution of $\mathcal{R}(x_0)$. In this case $\varepsilon = \varepsilon_r$. But other choices of ρ are possible that give an easy to sample distribution and $\varepsilon < \varepsilon_r$. In fact, for some standard randomizers such as addition of Laplace noise we will get $\varepsilon = \varepsilon_r/2$.

Now assuming that ρ is t-samplable and given a PRG $G\colon \{0,1\}^\ell \to \{0,1\}^t$ we define $\mathcal{R}[G]$ as in Def. 3.1 and $\mathcal{R}[G,\gamma]$ us in Algorithm 1. The randomizer $\mathcal{R}$ is deletion ε-DP so all the results we proved apply to it as well (with the deletion ε and not the replacement ε_r). In addition we show that replacement privacy is preserved as well.

Lemma A.6. *For a t-samplable deletion ε-DP and replacement ε_r-DP local randomizer $\mathcal{R}\colon X \to Y$ and $G\colon \{0,1\}^\ell \to \{0,1\}^t$, let $\mathcal{D}$ denote the following family of tests which take $r' \in \{0,1\}^t$ as an input:*

$$\mathcal{D} := \left\{ \mathtt{ind}\left(\frac{\mathbf{Pr}[\mathcal{R}(x) = \mathcal{R}_\emptyset(r')]}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r) = \mathcal{R}_\emptyset(r')]} \ge \theta\right) \;\middle|\; x \in X, \theta \in [0, e^\varepsilon] \right\}.$$

If G β-fools $\mathcal{D}$ for $\beta < 1/(2e^\varepsilon)$ then $R[G]$ is a replacement $(\varepsilon_r + 4e^\varepsilon\beta)$-DP local randomizer. Furthermore, for every $\gamma > 0$, $\mathcal{R}[G,\gamma]$ is a replacement $(\varepsilon_r + 4e^\varepsilon\beta)$-DP local randomizer.

Proof. As in the proof of Lemma 3.2, for $y \in Y$, we denote the density ratio of $\mathcal{R}(x)$ to ρ at y by

$$\pi_x(y) := \frac{\mathbf{Pr}[\mathcal{R}(x) = y]}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r) = y]}$$

and note that $\mathcal{R}[G]$ outputs a seed s with probability:

$$\mu_x(s) := \frac{\pi(\mathcal{R}_\emptyset(G(s)))}{\sum_{s'\in\{0,1\}^\ell} \pi(\mathcal{R}_\emptyset(G(s')))}.$$

Thus for two inputs $x, x' \in X$ and any $s \in \{0,1\}^\ell$ we have that

$$\begin{aligned}\frac{\mu_x(s)}{\mu_{x'}(s)} &= \frac{\pi_x(\mathcal{R}_\emptyset(G(s)))}{\pi_{x'}(\mathcal{R}_\emptyset(G(s)))} \cdot \frac{\sum_{s'\in\{0,1\}^\ell} \pi_{x'}(\mathcal{R}_\emptyset(G(s')))}{\sum_{s'\in\{0,1\}^\ell} \pi_x(\mathcal{R}_\emptyset(G(s')))} \\ &= \frac{\mathbf{Pr}[\mathcal{R}(x) = \mathcal{R}_\emptyset(G(s'))]}{\mathbf{Pr}[\mathcal{R}(x') = \mathcal{R}_\emptyset(G(s'))]} \cdot \frac{\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_{x'}(\mathcal{R}_\emptyset(G(s')))]}{\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))]}\end{aligned}$$

Now $\mathcal{R}$ is ε_r-replacement-DP and therefore the first term satisfies:

$$\frac{\mathbf{Pr}[\mathcal{R}(x) = \mathcal{R}_\emptyset(G(s'))]}{\mathbf{Pr}[\mathcal{R}(x') = \mathcal{R}_\emptyset(G(s'))]} \in \left[e^{-\varepsilon_r}, e^{\varepsilon_r}\right].$$

At the same time, we showed in Lemma 3.2 that $\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] \in [e^{-2e^\varepsilon\beta}, e^{2e^\varepsilon\beta}]$ and also $\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_{x'}(\mathcal{R}_\emptyset(G(s')))] \in [e^{-2e^\varepsilon\beta}, e^{2e^\varepsilon\beta}]$. Therefore $\frac{\mu_x(s)}{\mu_{x'}(s)} \in [e^{-\varepsilon_r - 4e^\varepsilon\beta}, e^{\varepsilon_r + 4e^\varepsilon\beta}]$.

To show that $\mathcal{R}[G,\gamma]$ is a replacement $(\varepsilon_r + 4e^\varepsilon\beta)$-DP local randomizer we observe that for every x, $\mathcal{R}[G,\gamma](x)$ is a mixture of $\mathcal{R}[G](x)$ and ρ_G. As we showed, $\mathcal{R}[G](x)$ is $(\varepsilon_r + 4e^\varepsilon\beta)$-close to $\mathcal{R}[G](x')$ and we also know from Lemma 3.2 that ρ_G is $(\varepsilon + 2e^\varepsilon\beta)$-close to $\mathcal{R}[G](x')$. By quasi-convexity we obtain that $\mathcal{R}[G,\gamma](x)$ is $(\varepsilon_r + 4e^\varepsilon\beta)$-close to $\mathcal{R}[G](x')$. We also know that $\mathcal{R}[G,\gamma](x)$ is $(\varepsilon + 2e^\varepsilon\beta)$-close to ρ_G. Appealing to quasi-convexity again, we obtain that $\mathcal{R}[G,\gamma](x)$ is $(\varepsilon_r + 4e^\varepsilon\beta)$-close to $\mathcal{R}[G,\gamma](x')$. □

A.2 Extension to (ε,δ)-DP

We next extend our approach to (ε,δ)-DP randomizers. The approach here is similar, except that we for some outputs $y = \mathcal{R}_\emptyset(G(s))$, the prescribed "rejection probability" in the original approach would be larger than one. To handle this, we simply truncate this ratio at 1 to get a probability. Algorithm 4 is identical to Algorithm 1 except for this truncation in the step where we sample b.

Algorithm 4 $\mathcal{R}[G,\gamma]$: PRG compression of deletion (ε,δ)-DP $\mathcal{R}$

Input: $x \in X$, $\varepsilon, \gamma > 0$; seeded PRG $G\colon \{0,1\}^\ell \to \{0,1\}^t$; t-samplable ε-DP randomizer $\mathcal{R}$.

$j = 0$; $J = e^\varepsilon \ln(1/\gamma)/(1-\delta)$
Sample a random seed $s \in \{0,1\}^\ell$.
while $j < J$ **do**
  $y = \mathcal{R}_\emptyset(G(s))$
  Sample b from Bernoulli $\left(\min\left(1, \frac{\mathbf{Pr}[\mathcal{R}(x)=y]}{e^\varepsilon \mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r)=y]}\right)\right)$
  if $b == 1$ **then**
    BREAK
  end if
  $j = j + 1$
  Sample a random seed $s \in \{0,1\}^\ell$.
end while
Send s

The proof is fairly similar to that for the pure DP randomizer. We start with a lemma that relates the properties of the PRG to the properties of the randomizer that need to be preserved in order to ensure that it satisfies deletion (ε',δ')-LDP.

Lemma A.7. *For a t-samplable deletion (ε,δ)-DP local randomizer $\mathcal{R}\colon X \to Y$ and $G\colon \{0,1\}^\ell \to \{0,1\}^t$, let $\mathcal{D}$ denote the following family of tests which take $r' \in \{0,1\}^t$ as an input:*

$$\mathcal{D} := \left\{ \mathtt{ind}\left(\frac{\mathbf{Pr}[\mathcal{R}(x) = \mathcal{R}_\emptyset(r')]}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r) = \mathcal{R}_\emptyset(r')]} \geq \theta\right) \;\middle|\; x \in X, \theta \in [0, e^\varepsilon] \right\}.$$

Suppose that G β-fools $\mathcal{D}$ and let $\pi_x(y) := \frac{\min(\mathbf{Pr}[\mathcal{R}(x)=y], e^\varepsilon \mathbf{Pr}[\mathcal{R}(\emptyset)=y])}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r)=y]}$. Then

$$\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] \in [1 - \delta - e^\varepsilon\beta, 1 + e^\varepsilon\beta]$$

and

$$\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[|1 - e^\varepsilon \pi_x(\mathcal{R}_\emptyset(G(s')))|_+\right] \leq \delta + \beta.$$

Proof. Let ρ_G be the uniform distribution over $\{0,1\}^\ell$. Let $\nu_x(y) := \mathbf{Pr}[\mathcal{R}(x) = y]$ and let $\tilde{\nu}_x(y) := \min(\nu_x(y), e^\varepsilon \mathbf{Pr}[\mathcal{R}(\emptyset) = y])$. Note that $\tilde{\nu}_x(\cdot)$ does not necessarily define a probability distribution. For $S = \{y :$

$\tilde{\nu}_x(y) < \nu_x(y)\}$, we have

$$\begin{aligned}\nu_x(S) &= \sum_{y\in S} \nu_x(y) \\ &= \sum_{y\in S} \tilde{\nu}_x(y) + \sum_{y\in S} (\nu_x(y) - \tilde{\nu}_x(y)) \\ &= \sum_{y\in S} e^{\varepsilon}\rho(y) + \sum_{y} (\nu_x(y) - \tilde{\nu}_x(y)) \\ &= e^{\varepsilon}\rho(S) + (1 - \sum_{y} \tilde{\nu}_x(y)).\end{aligned}$$

Then deletion (ε, δ)-DP of $\mathcal{R}$ implies that $\sum_y \tilde{\nu}_x(y) \geq 1 - \delta$. Observe that this implies that for true randomness we have:

$$\begin{aligned}\mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r'))] &= \mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}\left[\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(r'))}{\mathbf{Pr}_{r\sim\{0,1\}^t}[\mathcal{R}_\emptyset(r) = \mathcal{R}_\emptyset(r')]}\right] \\ &= \mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}\left[\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))}\right] \\ &= \mathop{\mathbf{E}}_{y\sim\rho}\left[\frac{\tilde{\nu}_x(y)}{\rho(y)}\right] \\ &= \sum_{y\in Y} \rho(y)\cdot\frac{\tilde{\nu}_x(y)}{\rho(y)} \\ &= \sum_{y\in Y} \tilde{\nu}_x(y) \quad \in \quad [1-\delta, 1].\end{aligned}$$

Using the fact that $\pi_x(y) \in [0, e^{\varepsilon}]$ we have that

$$\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] = \int_0^{e^{\varepsilon}} \mathop{\mathbf{Pr}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s'))) \geq \theta] d\theta$$

and, similarly,

$$\mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r'))] = \int_0^{e^{\varepsilon}} \mathop{\mathbf{Pr}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta]\, d\theta.$$

Thus, it follows that

$$\begin{aligned}\left|\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] - \mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r'))]\right| &= \left|\int_0^{e^{\varepsilon}} \left(\mathop{\mathbf{Pr}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s'))) \geq \theta] - \mathop{\mathbf{Pr}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta]\right) d\theta\right| \\ &\leq \int_0^{e^{\varepsilon}} \left|\mathop{\mathbf{Pr}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s'))) \geq \theta] - \mathop{\mathbf{Pr}}_{r'\sim\{0,1\}^t}[\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta]\right| d\theta \\ &\leq e^{\varepsilon}\beta,\end{aligned}$$

where in the last step, we have used the property of the pseudorandom generator that it fools $\mathcal{D}$, and the fact that for $\theta \in [0, e^{\varepsilon})$, $\frac{\tilde{\nu}_x(y)}{\mathbf{Pr}[\mathcal{R}(\emptyset)=y]} < \theta$ if and only if $\frac{\nu_x(y)}{\mathbf{Pr}[\mathcal{R}(\emptyset)=y]} < \theta$. The first part of the claim follows.

For the second part of the claim we first note that deletion (ε, δ)-DP of $\mathcal{R}$ implies that

$$\begin{aligned}\mathop{\mathbf{E}}_{r' \sim \{0,1\}^t}[|1 - e^{\varepsilon}\pi_x(\mathcal{R}_\emptyset(r'))|_+] &= \mathop{\mathbf{E}}_{y \sim \rho}[|1 - e^{\varepsilon}\pi_x(y)|_+] \\ &= \mathop{\mathbf{E}}_{y \sim \rho}[|1 - e^{\varepsilon}\pi_x(y)|_+] \\ &= \sum_{y \in Y} \rho(y)|1 - e^{\varepsilon}\pi_x(y)|_+ \\ &= \sum_{y \in Y} |\rho(y) - e^{\varepsilon}\tilde{\nu}_x(y)|_+ \\ &= \sum_{y \in Y} |\rho(y) - e^{\varepsilon}\nu_x(y)|_+ \leq \delta.\end{aligned}$$

Also note that

$$\mathop{\mathbf{E}}_{r' \sim \{0,1\}^t}[|1 - e^{\varepsilon}\pi_x(\mathcal{R}_\emptyset(r'))|_+] = \int_0^1 \mathop{\mathbf{Pr}}_{r' \sim \{0,1\}^t}[1 - e^{\varepsilon}\pi_x(\mathcal{R}_\emptyset(r')) \geq \theta]\, d\theta = 1 - \int_0^1 \mathop{\mathbf{Pr}}_{r' \sim \{0,1\}^t}\left[\pi_x(\mathcal{R}_\emptyset(r')) \geq \frac{\theta}{e^{\varepsilon}}\right] d\theta.$$

Similarly,

$$\mathop{\mathbf{E}}_{s' \sim \{0,1\}^\ell}[|1 - e^{\varepsilon}\pi_x(\mathcal{R}_\emptyset(G(s')))|_+] = 1 - \int_0^1 \mathop{\mathbf{Pr}}_{s' \sim \{0,1\}^\ell}\left[\pi_x(\mathcal{R}_\emptyset(G(s'))) \geq \frac{\theta}{e^{\varepsilon}}\right] d\theta.$$

Thus by the same argument as before, the fact that G, β-fools $\mathcal{D}$ implies that

$$\mathop{\mathbf{E}}_{s' \sim \{0,1\}^\ell}\left[|1 - e^{\varepsilon}\pi_x(\mathcal{R}_\emptyset(G(s')))|_+\right] \leq \mathop{\mathbf{E}}_{r' \sim \{0,1\}^t}[|1 - e^{\varepsilon}\pi_x(\mathcal{R}_\emptyset(r'))|_+] + \beta \leq \delta + \beta.$$

□

We can now give an analogue of Lemma 3.2 for deletion (ε, δ)-DP randomizers.

Lemma A.8. *For a t-samplable deletion (ε, δ)-DP local randomizer $\mathcal{R}\colon X \to Y$ and $G\colon \{0,1\}^\ell \to \{0,1\}^t$, let $\mathcal{D}$ denote the following family of tests which take $r' \in \{0,1\}^t$ as an input:*

$$\mathcal{D} := \left\{ \mathtt{ind}\left(\frac{\mathbf{Pr}[\mathcal{R}(x) = \mathcal{R}_\emptyset(r')]}{\mathbf{Pr}_{r \sim \{0,1\}^t}[\mathcal{R}_\emptyset(r) = \mathcal{R}_\emptyset(r')]} \geq \theta\right) \;\middle|\; x \in X, \theta \in [0, e^{\varepsilon}] \right\}.$$

If G β-fools $\mathcal{D}$ where $\delta + e^{\varepsilon}\beta < 1/2$ then $R[G]$ is a deletion $(\varepsilon + 2\delta + 2e^{\varepsilon}\beta, \delta + \beta)$-DP local randomizer. Furthermore, for every $\gamma > 0$, $\mathcal{R}[G, \gamma]$ is a deletion $(\varepsilon + 2\delta + 2e^{\varepsilon}\beta, \delta + \beta)$-DP local randomizer.

Proof. As before, we let the reference distribution ρ_G be the uniform distribution over $\{0,1\}^\ell$. Using the definitions in the proof of Lemma A.7 we observe that $\mathcal{R}[G](x)$ outputs s with probability:

$$\mu_x(s) := \frac{\pi(\mathcal{R}_\emptyset(G(s)))}{\sum_{s' \in \{0,1\}^\ell} \pi(\mathcal{R}_\emptyset(G(s')))} = \frac{\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s)))}{\rho(\mathcal{R}_\emptyset(G(s)))}}{\frac{\mathbf{E}_{s' \sim \{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))]}{2^\ell}} = \rho_G(s) \cdot \frac{\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s)))}{\rho(\mathcal{R}_\emptyset(G(s)))}}{\mathbf{E}_{s' \sim \{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))]}.$$

Now, by definition of $\tilde{\nu}_x$ we have that the numerator satisfies $\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s)))}{\rho(\mathcal{R}_\emptyset(G(s)))} \leq e^{\varepsilon}$. In addition, by Lemma A.7 the denominator $\mathbf{E}_{s' \sim \{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] \geq 1 - \delta - e^{\varepsilon}\beta$. Therefore

$$\mu_x(s) \leq \rho_G(s) \cdot \frac{e^{\varepsilon}}{1 - \delta - e^{\varepsilon}\beta} \leq e^{\varepsilon + 2\delta + e^{\varepsilon}\beta}\rho_G(s).$$

Now for the other side of (ε, δ)-closeness we simply observe that by the Lemma A.7,

$$\begin{aligned}\sum_{s\in\{0,1\}^\ell} \left|\rho_G(s) - e^{\varepsilon+e^\varepsilon\beta}\mu_x(s)\right|_+ &= \sum_{s\in\{0,1\}^\ell} \left|\rho_G(s) - e^{\varepsilon+e^\varepsilon\beta}\rho_G(s)\cdot\frac{\pi_x(\mathcal{R}_\emptyset(G(s)))}{\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))]}\right|_+ \\ &\le \sum_{s\in\{0,1\}^\ell} |\rho_G(s) - e^\varepsilon\rho_G(s)\cdot\pi_x(\mathcal{R}_\emptyset(G(s)))|_+ \\ &= \mathop{\mathbf{E}}_{s\sim\{0,1\}^\ell}\left[|1 - e^\varepsilon\pi_x(\mathcal{R}_\emptyset(G(s)))|_+\right] \le \delta+\beta.\end{aligned}$$

Now, to establish that $R[G,\gamma]$ is $(\varepsilon+2\delta+2e^\varepsilon\beta, \delta+\beta)$ we, as before, appeal to quasi-convexity. □

To establish the utility guarantees for $\mathcal{R}[G,\gamma]$ we follow the same approach by establishing the utility guarantees for $\mathcal{R}[\mathrm{ID}_t,\gamma]$ and then using the properties of G.

Lemma A.9. *Let $\mathcal{R}$ be a deletion ε-DP t-samplable local randomizer. Then for the identity function $\mathrm{ID}_t\colon \{0,1\}^t \to \{0,1\}^t$ and any $\gamma > 0$ we have that $\mathcal{R}[\mathrm{ID}_t,\gamma]$ is a deletion ε-DP local randomizer and for every $x \in \mathcal{X}$, $\mathrm{TV}(\mathcal{R}_\emptyset(\mathcal{R}[\mathrm{ID}_t,\gamma](x)), \mathcal{R}(x)) \le \delta+\gamma$.*

Proof. Conditioned on accepting a sample, $\mathcal{R}[\mathrm{ID}_t,\gamma]$ outputs a sample from the truncated version of the distribution of $\mathcal{R}(x)$. Specifically, y is output with probability $\bar\nu_x(y) := \frac{\tilde\nu(y)}{\sum_{y\in Y}\tilde\nu(y)}$, where $\nu_x(y) := \mathbf{Pr}[\mathcal{R}(x)=y]$ and $\tilde\nu_x(y) := \min(\nu_x(y), e^\varepsilon\,\mathbf{Pr}[\mathcal{R}(\emptyset)=y])$. From the proof of Lemma A.7, we know that $\sum_{y\in Y}\tilde\nu_x(y) \ge 1-\delta$. Thus

$$\begin{aligned}\mathrm{TV}(\nu_x,\bar\nu_x) &= \frac{1}{2}\sum_{y\in Y}|\nu_x(y)-\bar\nu_x(y)| \\ &\le \frac{1}{2}\sum_{y\in Y}(|\nu_x(y)-\tilde\nu_x(y)| + |\tilde\nu_x(y)-\bar\nu_x(y)|) \\ &= \frac{1}{2}\sum_{y\in Y}(\nu_x(y)-\tilde\nu_x(y)+\bar\nu_x(y)-\tilde\nu_x(y)) \le \delta.\end{aligned}$$

Truncation of the distribution also reduces the probability that a sample is accepted. Specifically,

$$\mathop{\mathbf{E}}_{y\sim\rho}\left[\frac{\tilde\nu_x(y)}{e^\varepsilon\rho(y)}\right] = \sum_{y\in Y}\frac{\tilde\nu_x(y)}{e^\varepsilon} \ge \frac{1-\delta}{e^\varepsilon}.$$

$\mathcal{R}[G,\gamma]$ tries at least $e^\varepsilon\ln(1/\gamma)/(1-\delta)$ samples and therefore, as in the proof of Lemma A.2, failure to accept any samples adds at most γ to the total variation distance. □

From here we can directly obtain the analogues of Lemma 3.3 and Theorem3.4.

Finally, to deal with the replacement version of (ε,δ)-DP we combine the ideas we used in Lemmas A.6 and A.8. The main distinction is a somewhat stronger test that we need to fool in this case.

Lemma A.10. *For a t-samplable replacement (ε_r,δ_r)-DP and deletion (ε,δ)-DP local randomizer $\mathcal{R}\colon X\to Y$ and $G\colon\{0,1\}^\ell\to\{0,1\}^t$, let $\mathcal{D}$ and $\mathcal{D}_r$ denote the following families of tests which take $r'\in\{0,1\}^t$ as an input:*

$$\mathcal{D} := \left\{ \mathtt{ind}\left(\frac{\mathbf{Pr}[\mathcal{R}(x)=\mathcal{R}_\emptyset(r')]}{\rho(\mathcal{R}_\emptyset(r'))}\ge\theta\right) \;\middle|\; x\in X,\theta\in[0,e^\varepsilon]\right\};$$

$$\mathcal{D}_r := \left\{ \mathtt{ind}\left(\frac{\tilde\nu_x(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))} - e^\varepsilon\frac{\tilde\nu_{x'}(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))}\ge\theta\right) \;\middle|\; x,x'\in X,\theta\in[0,e^\varepsilon]\right\},$$

where ρ is the reference distribution of $\mathcal{R}$ and $\tilde\nu_x(y) := \min(\mathbf{Pr}[\mathcal{R}(x)=y], e^\varepsilon\rho(y))$. If G β-fools $\mathcal{D}\cup\mathcal{D}_r$ where $\delta+e^\varepsilon\beta<1/2$ then $R[G]$ is a replacement $(\varepsilon_r+2\delta+3e^\varepsilon\beta, 2\delta_r+2e^\varepsilon\beta)$-DP local randomizer. Furthermore, for every $\gamma>0$, $\mathcal{R}[G,\gamma]$ is a $(\varepsilon_r+2\delta+3e^\varepsilon\beta, 2\delta_r+2e^\varepsilon\beta)$-DP local randomizer.

Proof. First we observe that $\mathcal{R}$ being $(\varepsilon_r, \delta_r)$ replacement DP implies that $\tilde{\nu}_x$ and $\tilde{\nu}_{x'}$ are $(\varepsilon_r, \delta_r)$ close in the following sense:

$$\begin{aligned}\mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}\left[\left|\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))} - e^{\varepsilon_r}\frac{\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))}\right|_+\right] &= \mathop{\mathbf{E}}_{y\sim\rho}\left[\left|\frac{\tilde{\nu}_x(y)}{\rho(y)} - e^{\varepsilon_r}\frac{\tilde{\nu}_{x'}(y)}{\rho(y)}\right|_+\right] \\ &= \sum_{y\in Y}\left|\tilde{\nu}_x(y) - e^{\varepsilon_r}\tilde{\nu}_{x'}(y)\right|_+ \\ &\le \sum_{y\in Y}\left|\nu_x(y) - e^{\varepsilon_r}\nu_{x'}(y)\right|_+ \le \delta_r,\end{aligned}$$

where we used the fact that if $\nu_{x'}(y) > \tilde{\nu}_{x'}(y)$ then $\tilde{\nu}_{x'}(y) = e^{\varepsilon}\rho(y) \ge \tilde{\nu}_x(y)$ and so

$$\left|\tilde{\nu}_x(y) - e^{\varepsilon_r}\tilde{\nu}_{x'}(y)\right|_+ = \left|\tilde{\nu}_x(y) - e^{\varepsilon_r}\nu_{x'}(y)\right|_+.$$

Using the decomposition

$$\mathop{\mathbf{E}}_{r'\sim\{0,1\}^t}\left[\left|\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))} - e^{\varepsilon_r}\frac{\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))}\right|_+\right] = \int_0^{e^\varepsilon}\mathop{\mathbf{Pr}}_{r'\sim\{0,1\}^t}\left[\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))} - e^{\varepsilon_r}\frac{\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(r'))}{\rho(\mathcal{R}_\emptyset(r'))} \ge \theta\right]d\theta$$

and the fact that G β fools $\mathcal{D}_r$ we obtain that

$$\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[\left|\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s')))}{\rho(\mathcal{R}_\emptyset(G(s')))} - e^{\varepsilon_r}\frac{\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(G(s')))}{\rho(\mathcal{R}_\emptyset(G(s')))}\right|_+\right] \le \delta_r + e^\varepsilon\beta. \tag{2}$$

By Lemma A.7 we have that for $\pi_x(y) := \frac{\tilde{\nu}_x(y)}{\rho(y)}$ it holds that

$$\zeta_x := \mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))] \in [1-\delta-e^\varepsilon\beta, 1+e^\varepsilon\beta].$$

Now following the notation in Lemma A.8 we know that the distribution of $\mathcal{R}[G](x)$ is

$$\mu_x(s) = \rho_G(s)\cdot\frac{\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s)))}{\rho(\mathcal{R}_\emptyset(G(s)))}}{\mathbf{E}_{s'\sim\{0,1\}^\ell}[\pi_x(\mathcal{R}_\emptyset(G(s')))]} = \frac{\rho_G(s)\cdot\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s)))}{\zeta_x\cdot\rho(\mathcal{R}_\emptyset(G(s)))}.$$

Thus setting $\varepsilon' = \varepsilon_r + 2\delta + 3e^\varepsilon\beta$ we obtain:

$$\begin{aligned}\sum_{s'\in\{0,1\}^\ell}\left|\mu_x s - e^{\varepsilon'}\mu_{x'}(s)\right|_+ &= \mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[\left|\frac{\mu_x s}{\rho_G(s')} - e^{\varepsilon'}\frac{\mu_{x'}}{\rho_G(s')}\right|_+\right] \\ &= \mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[\left|\frac{\mu_x s}{\rho_G(s')} - e^{\varepsilon'}\frac{\mu_{x'}}{\rho_G(s')}\right|_+\right] \\ &= \mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[\left|\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s')))}{\zeta_x\cdot\rho(\mathcal{R}_\emptyset(G(s')))} - e^{\varepsilon'}\frac{\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(G(s')))}{\zeta_{x'}\cdot\rho(\mathcal{R}_\emptyset(G(s')))}\right|_+\right] \\ &= \frac{1}{\zeta_x}\cdot\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[\left|\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s')))}{\rho(\mathcal{R}_\emptyset(G(s')))} - e^{\varepsilon'}\frac{\zeta_x\cdot\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(G(s')))}{\zeta_{x'}\cdot\rho(\mathcal{R}_\emptyset(G(s')))}\right|_+\right] \\ &\le \frac{1}{1-\delta-e^\varepsilon\beta}\cdot\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[\left|\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s')))}{\rho(\mathcal{R}_\emptyset(G(s')))} - e^{\varepsilon'}\frac{(1-\delta-e^\varepsilon\beta)\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(G(s')))}{(1+e^\varepsilon\beta)\rho(\mathcal{R}_\emptyset(G(s')))}\right|_+\right] \\ &\le \frac{1}{1-\delta-e^\varepsilon\beta}\cdot\mathop{\mathbf{E}}_{s'\sim\{0,1\}^\ell}\left[\left|\frac{\tilde{\nu}_x(\mathcal{R}_\emptyset(G(s')))}{\rho(\mathcal{R}_\emptyset(G(s')))} - e^{\varepsilon_r}\frac{\tilde{\nu}_{x'}(\mathcal{R}_\emptyset(G(s')))}{\rho(\mathcal{R}_\emptyset(G(s')))}\right|_+\right] \\ &\le 2(\delta_r + e^\varepsilon\beta),\end{aligned}$$

where we used that $\frac{1+e^\varepsilon\beta}{1-\delta-e^\varepsilon\beta} \le e^{2\delta+3e^\varepsilon\beta}$ and $\frac{1}{1-\delta-e^\varepsilon\beta} \le 2$ whenever $\delta + e^\varepsilon\beta < 1/2$. □

B Proofs and Details for "Frequency Estimation"

Lemma B.1 (Lemma 4.1 restated)**.** *PI-RAPPOR randomizer (Alg. 2) is deletion* $\max\left\{\frac{\alpha_1}{\alpha_0}, \frac{1-\alpha_0}{1-\alpha_1}\right\}$*-DP and replacement* $\frac{\alpha_1(1-\alpha_0)}{\alpha_0(1-\alpha_1)}$*-DP.*

Proof. While it is easy to analyze the privacy guarantees of PI-RAPPOR directly it is instructive to show that these guarantees follow from our general compression technique. Specifically, there is a natural way to sample from the reference distribution of RAPPOR relative to which our pairwise PRG fools the density tests given in Lemma 3.2.

To sample from the reference distribution of RAPPOR we pick k values $z_1, \ldots, z_k$ randomly independently and uniformly from $\mathbb{F}_p$ and then output $\texttt{bool}(z_1), \texttt{bool}(z_2), \ldots, \texttt{bool}(z_k)$ (we note that samplability is defined using uniform distribution over binary strings length t as an input but any other distribution can be used instead). By our choice of parameter p and definition of $\texttt{bool}$, this gives k i.i.d. samples from $\mathrm{Bernoulli}(\alpha_0)$, which is the reference distribution for RAPPOR. Let $\mathcal{R}$ denote the RAPPOR randomizer. For any $j \in [k]$ and $z' \in \mathbb{F}_p^k$ the ratio of densities at z' satisfies:

$$\frac{\mathbf{Pr}[\mathcal{R}(j) = \mathcal{R}_\emptyset(z')]}{\mathbf{Pr}_{z\sim\mathbb{F}_p^k}[\mathcal{R}_\emptyset(z) = \mathcal{R}_\emptyset(z')]} = \begin{cases} \frac{\alpha_1}{\alpha_0}, & \text{if } \texttt{bool}(z'_j) = 1 \\ \frac{1-\alpha_1}{1-\alpha_0}, & \text{otherwise.} \end{cases}.$$

With probability α_1, PI-RAPPOR algorithm samples ϕ uniformly from $\Phi_{j,1}$ and with probability $1-\alpha_1$ PI-RAPPOR algorithm samples ϕ uniformly from $\Phi_{j,0}$. This means that PI-RAPPOR is exactly equal to $\mathcal{R}[G]$, where $G\colon \mathbb{F}_p^2 \to \mathbb{F}_p^k$ is defined as $G(\phi) = \phi(1), \phi(2), \ldots, \phi(k)$.

Now to prove that PI-RAPPOR has the same deletion privacy guarantees as RAPPOR it suffices to prove that G 0-fools the tests based on the ratio of densities above. This follows immediately from the fact that $\texttt{bool}(\phi(j))$ for $\phi \sim \Phi$ is distributed in the same way as $\texttt{bool}(z_j)$ for $z \sim \mathbb{F}_p^k$.

To prove that PI-RAPPOR has the same replacement privacy guarantees as RAPPOR we simply use the same reference distribution and apply Lemma A.6. □

Lemma B.2 (Lemma 4.2 restated)**.** *For any dataset* $S \in [k]^n$*, the estimate* $\tilde{c}$ *computed by PI-RAPPOR algorithm (Algs. 2,3) satisfies:*

- $\mathbf{E}[\tilde{c}] = c(S)$
- *For all* $j \in [k]$*,*

$$\mathbf{Var}[\tilde{c}_j] = c(S)_j \frac{1-\alpha_0-\alpha_1}{\alpha_1-\alpha_0} + n\frac{\alpha_0(1-\alpha_0)}{(\alpha_1-\alpha_0)^2}$$

  For the symmetric case $\alpha_0 = 1-\alpha_1$ *this simplifies to* $\mathbf{Var}[\tilde{c}_j] = n\frac{\alpha_0(1-\alpha_0)}{(1-2\alpha_0)^2}$.

In particular, the expected ℓ_2 *squared error is*

$$\mathbf{E}\left[\|\tilde{c} - c(S)\|_2^2\right] = n\frac{1-\alpha_0-\alpha_1}{\alpha_1-\alpha_0} + nk\frac{\alpha_0(1-\alpha_0)}{(\alpha_1-\alpha_0)^2}.$$

Proof. We first note that

$$\tilde{c}_j = \sum_{i\in[n]} \frac{\texttt{bool}(\phi^i(j)) - \alpha_0}{\alpha_1-\alpha_0},$$

where ϕ^i is the output of the PI-RAPPOR randomizer on input x_i. Thus to prove the claim about the expectation it is sufficient to prove that for every i,

$$\mathbf{E}\left[\frac{\texttt{bool}(\phi^i(j)) - \alpha_0}{\alpha_1-\alpha_0}\right] = \texttt{ind}\,(x_i = j)$$

and to prove the claim for variance it is sufficient to prove that

$$\mathbf{Var}\left[\frac{\texttt{bool}(\phi^i(j)) - \alpha_0}{\alpha_1-\alpha_0}\right] = \texttt{ind}\,(x_i = j)\,\frac{1-\alpha_0-\alpha_1}{\alpha_1-\alpha_0} + \frac{\alpha_0(1-\alpha_0)}{(\alpha_1-\alpha_0)^2}.$$

If $x_i = j$ then both of these claims follow directly from the fact that, by definition of PI-RAPPOR randomizer, in this case the distribution of $\texttt{bool}(\phi^i(j))$ is $\text{Bernoulli}(\alpha_1)$.

If, on the other hand $x_i \neq j$, we use pairwise independence of $\texttt{bool}(\phi(x_i))$ and $\texttt{bool}(\phi(j))$ for $\phi \sim \Phi$ to infer that conditioning the distribution $\texttt{bool}(\phi(x_i)) = b$ (for any b) does not affect the distribution of $\texttt{bool}(\phi(x_i))$. Thus, if $x_i \neq j$ then $\texttt{bool}(\phi^i(j))$ is distributed as $\text{Bernoulli}(\alpha_0)$ and we can verify the desired property directly.

Finally,

$$\mathbf{E}\left[\|\tilde{c} - c(S)\|_2^2\right] = \mathbf{E}\left[\sum_{j\in[k]} (\tilde{c}_j - c(S)_j)^2\right] = \sum_{j\in[k]} \mathbf{Var}[\tilde{c}_j] = n\frac{1-\alpha_0-\alpha_1}{\alpha_1-\alpha_0} + nk\frac{\alpha_0(1-\alpha_0)}{(\alpha_1-\alpha_0)^2}$$

□

Below, we analyze the computational and communication cost of PI-RAPPOR and discuss the choice of p.

Lemma B.3. *PI-RAPPOR randomizer (Alg. 2) can be implemented in $\tilde{O}(\log p)$ time and uses $2\lceil \log_2 p \rceil$ bits of communication.*

Proof. The running time of PI-RAPPOR is dominated by the time to pick a random and uniform element in $\Phi_{j,b}$. This can be done by picking $\phi_1 \in \mathbb{F}_p$ randomly and uniformly. We then need to pick ϕ_0 randomly and uniformly from the set $\{\phi_0 \mid \texttt{bool}(\phi(j)) = b\}$. Given the result of multiplication $j\phi_1$ this can be done in $O(\log p)$ time. For example for $b = 1$ this set is equal to $\{-j\phi_1, -j\phi_1 + 1, \ldots, -j\phi_1 + \alpha_0 p - 1\}$ where all arithmetic operations are in $\mathbb{F}_p$. The set consists of at most two contiguous ranges of integers and thus a random and uniform element can be chosen in $O(\log p)$ time. Multiplication in $\mathbb{F}_p$ can be done in $O(\log(p) \cdot (\log\log p)^2)$ (*e.g.* (Menezes et al., 2018)) but in most practical settings standard Montgomery modular multiplication that takes $O(\log^2(p))$ time would be sufficiently fast. □

The analysis of the running time of decoding and aggregation is similarly straightforward since decoding every bit of message takes time that is dominated by the time of a single multiplication in $\mathbb{F}_p$.

Lemma B.4. *For every $j \in k$, the server-side of PI-RAPPOR (Alg. 3) computes $\tilde{c}_j$ in time $\tilde{O}(n \log p)$. In particular, the entire histogram is computed in time $\tilde{O}(kn \log p)$.*

Note that the construction of the entire histogram on the server is relatively expensive. For comparison we note that aggregation in the compression schemes in (Acharya et al., 2019) and (Chen et al., 2020) can be done in $\tilde{O}(n + k)$. However these schemes require $\Omega(k)$ computation on each client and thus the entire system also performs $\Omega(nk)$ computation. They also do not give a frequency oracle since the decoding time of even a single message is linear in k.

Finally, we need to discuss how to pick p. In addition to the condition that is p a prime larger than k, our algorithm requires that $\alpha_0 p$ be an integer. We observe that while, in general, we cannot always guarantee that $\alpha_0 = p/(e^\varepsilon + 1)$, by picking p that is a sufficiently large multiple of $\max\{e^\varepsilon, 1/\varepsilon\}$ we get an ε'-DP PI-RAPPOR algorithm for ε' that is slightly smaller than ε (which also implies that its utility is slightly worse). We make this formal below.

Lemma B.5. *There exists a constant c_0 such that for any $\varepsilon > 0$, $k \in \mathbb{N}$, $\Delta > 0$ and any prime $p \geq c_0 \max\{e^\varepsilon, 1/\varepsilon\}/\Delta$ we have that symmetric PI-RAPPOR with parameter $\alpha_0 = \lceil p/(e^\varepsilon + 1)\rceil/p$ satisfies deletion ε-DP and outputs an estimate that satisfies: for all $j \in [k]$, $\mathbf{Var}[\tilde{c}_j] \leq n\frac{(1+\Delta)e^\varepsilon}{(e^\varepsilon-1)^2}$. Further, PI-RAPPOR with $\alpha_0 = \lceil p/(e^\varepsilon + 1)\rceil/p$ and $\alpha_1 = 1/2$ satisfies replacement ε-DP and outputs an estimate that satisfies: for all $j \in [k]$, $\mathbf{Var}[\tilde{c}_j] = c(S)_j + n\frac{4(1+\Delta)e^\varepsilon}{(e^\varepsilon-1)^2}$.*

Proof. We first note that by our definition, $\alpha_0 p = \lceil p/(e^\varepsilon + 1)\rceil$ and therefore is an integer (as required by PI-RAPPOR). We denote by $\varepsilon' = \ln(1 - p/\alpha_0)$ (so that $\alpha_0 = 1/(e^{\varepsilon'} + 1)$ and note that $\varepsilon' \leq \varepsilon$. Thus the symmetric PI-RAPPOR satisfies ε-DP. We now note that $|1/(e^{\varepsilon'} + 1) - 1/(e^\varepsilon + 1)| \leq 1/p$. This implies that the bound on variance of PI-RAPPOR satisfies:

$$\begin{aligned}\mathbf{Var}[\tilde{c}_j] &= n\frac{\alpha_0(1-\alpha_0)}{(1-2\alpha_0)^2} = n\frac{\frac{1}{e^{\varepsilon'}+1}(1-\frac{1}{e^{\varepsilon'}+1})}{(1-2\frac{1}{e^{\varepsilon'}+1})^2} \\ &\leq n\frac{(\frac{1}{e^\varepsilon+1}+\frac{1}{p})(1-\frac{1}{e^\varepsilon+1})}{(1-2\frac{1}{e^{\varepsilon'}+1}-\frac{2}{p})^2}.\end{aligned}$$

$$\mathbf{Var}[\tilde{c}_j] = n\frac{\alpha_0(1-\alpha_0)}{(1-2\alpha_0)^2} = n\frac{\frac{1}{e^{\varepsilon'}+1}(1-\frac{1}{e^{\varepsilon'}+1})}{(1-2\frac{1}{e^{\varepsilon'}+1})^2} \leq n\frac{(\frac{1}{e^{\varepsilon}+1}+\frac{1}{p})(1-\frac{1}{e^{\varepsilon}+1})}{(1-2\frac{1}{e^{\varepsilon'}+1}-\frac{2}{p})^2}.$$

If $\varepsilon \leq 1$ then $\frac{1}{e^{\varepsilon'}+1} \geq \frac{1}{e+1}$ and $1-2\frac{1}{e^{\varepsilon'}+1} \geq \frac{\varepsilon}{e+1}$. Thus the addition/subtraction of $1/p$ to these quantities for $p \geq c_0/(\varepsilon\Delta)$ increases the bound by at most a multiplicative factor $(1+\Delta)$ (for a sufficiently large constant c_0).

Otherwise (if $\varepsilon > 1$), then $\frac{1}{e^{\varepsilon'}+1} \geq \frac{1}{e^{\varepsilon}}$ and $1-2\frac{1}{e^{\varepsilon'}+1} \geq \frac{e-1}{e+1}$. Thus the addition/subtraction of $1/p$ to these quantities for $p \geq c_0 e^{\varepsilon}/\Delta$ increases the bound by at most a multiplicative factor $(1+\Delta)$ (for a sufficiently large constant c_0).

The analysis for replacement DP is analogous. □

In practice, setting $\Delta = 1/100$ will make the loss of accuracy insignificant. Thus we can conclude that PI-RAPPOR with $p \geq c_1 \max\{k, e^{\varepsilon}, 1/\varepsilon\}$ for a sufficiently large constant c_1 achieves essentially the same guarantees as RAPPOR. This means that the communication cost of PI-RAPPOR is $2\log_2(\max\{k, e^{\varepsilon}, 1/\varepsilon\}) + O(1)$. Also we are typically interested in compression when $k \gg \max\{e^{\varepsilon}, 1/\varepsilon\}$ and in such case the communication cost is $2\log_2(k) + O(1)$.

C Details and Empirical Results for "Mean Estimation"

Below we describe the simple reduction by repetition and state the resulting guarantees.

Lemma C.1. *Assume that for some $\varepsilon > 0$ there exists a local ε-DP randomizer $\mathcal{R}_\varepsilon \colon \mathbb{B}^d \to Y$ and a decoding procedure $\mathtt{decode} \colon Y \to \mathcal{R}^d$ that for all $\mathbf{x} \in \mathbb{B}^d$, satisfies: $\mathbf{E}[\mathtt{decode}(\mathcal{R}_\varepsilon(\mathbf{x}))] = \mathbf{x}$ and $\mathbf{E}[\|\mathtt{decode}(\mathcal{R}_\varepsilon(\mathbf{x})) - \mathbf{x}\|_2^2] \leq \alpha_\varepsilon$. Further assume that $\mathcal{R}_\varepsilon$ uses ℓ bits of communication and runs in time T. Then for every integer $m \geq 2$ there is a local $(m\varepsilon)$-DP randomizer $\mathcal{R}_\varepsilon^m \colon \mathbb{B}^d \to Y^m$ and decoding procedure $\mathtt{decode}^m \colon Y^m \to \mathcal{R}^d$ that uses $m\ell$ bits of communication, runs in time mT and for every $\mathbf{x} \in \mathbb{B}^d$ satisfies: $\mathbf{E}[\mathtt{decode}^m(\mathcal{R}'_\varepsilon(\mathbf{x}))] = \mathbf{x}$ and $\mathbf{E}[\|\mathtt{decode}^m(\mathcal{R}_\varepsilon^m(\mathbf{x})) - \mathbf{x}\|_2^2] \leq \frac{\alpha_\varepsilon}{m}$.*

In particular, if for every $\varepsilon \in (1/2, 1]$, $\alpha_\varepsilon \leq \frac{cd}{\varepsilon^2}$ for some constant c, then for every $\varepsilon > 0$ there is a local ε-DP randomizer $\mathcal{R}'_\varepsilon$ and decoding procedure $\mathtt{decode}'$ that uses $\lceil\varepsilon\rceil\ell$ bits of communication, runs in time $\lceil\varepsilon\rceil T$ and for every $\mathbf{x} \in \mathbb{B}^d$ satisfies: $\mathbf{E}[\mathtt{decode}'(\mathcal{R}'_\varepsilon(\mathbf{x}))] = \mathbf{x}$ and $\mathbf{E}[\|\mathtt{decode}'(\mathcal{R}'_\varepsilon(\mathbf{x})) - \mathbf{x}\|_2^2] \leq \frac{2cd}{\min\{\varepsilon, \varepsilon^2\}}$.

Proof. The randomizer $\mathcal{R}_\varepsilon^m(\mathbf{x})$ runs $\mathcal{R}_\varepsilon(\mathbf{x})$ m times independently to obtain $y_1, \ldots, y_m$ and outputs these values. To decode we define $\mathtt{decode}^m(y_1, \ldots, y_m) := \frac{1}{m}(\mathtt{decode}(y_1) + \cdots + \mathtt{decode}(y_m))$. By (simple) composition of differential privacy, $\mathcal{R}_\varepsilon^m$ is (εm)-DP. The utility claim follows directly from linearity of expectation and independence of the estimates:

$$\mathbf{E}[\|\mathtt{decode}^m(\mathcal{R}_\varepsilon^m(\mathbf{x})) - \mathbf{x}\|_2^2] = \frac{1}{m} \cdot \mathbf{E}[\|\mathtt{decode}(\mathcal{R}_\varepsilon(\mathbf{x})) - \mathbf{x}\|_2^2] \leq \frac{\alpha_\varepsilon}{m}.$$

For the second part of the claim we define $\mathcal{R}'_\varepsilon$ as follows. For $\varepsilon \leq 1$, $\mathcal{R}'_\varepsilon(\mathbf{x})$ just outputs $\mathcal{R}_\varepsilon(\mathbf{x})$ and in this case $\mathtt{decode}'$ is the same as $\mathtt{decode}$. For $\varepsilon > 1$, we let $m = \lceil\varepsilon\rceil$ and apply the lemma to $\mathcal{R}_{\varepsilon'}$ for $\varepsilon' = \varepsilon/\lceil\varepsilon\rceil$. Note that $\varepsilon' \in (1/2, 1)$ and therefore the resulting bound on variance is

$$\mathbf{E}[\|\mathtt{decode}'(\mathcal{R}'_\varepsilon(\mathbf{x})) - \mathbf{x}\|_2^2] \leq \frac{1}{\lceil\varepsilon\rceil}\frac{cd}{\varepsilon'^2} = \frac{cd}{\varepsilon\varepsilon'} \leq \frac{2cd}{\varepsilon}.$$

□

For example, by using the reduction in Lemma C.1, we can reduce the computational cost to $\tilde{O}(\lceil\varepsilon\rceil d)$ while increasing the communication to $O(\lceil\varepsilon\rceil \log d)$. The server side reconstruction now requires sampling and averaging $n\lceil\varepsilon\rceil$ d-dimensional vectors. Thus the running time is $\tilde{O}(nd\varepsilon)$.

C.1 Empirical Comparison of Budget Splitting for Mean Estimation Algorithms

Figure 2 shows the results of applying Lemma C.1 to `PrivHS`, `PrivUnit` and `PrivUnitOptimized`. We plot the single repetition version of `SQKR` for comparison. The `SQKR` algorithm does not get more efficient for smaller ε and thus splitting it makes it worse in every aspect. As its error grows quickly with splitting, we do not plot the split version of `SQKR` in these plots. The results demonstrate that splitting does have some cost in terms of expected squared error, and going from $\varepsilon = 8$ to two runs of $\varepsilon = 4$ costs us about $2\times$ in expected squared error, and that the error continues to increase as we split more. These results can inform picking an appropriate point on the computation cost-error tradeoff and suggest that for ε around 8,

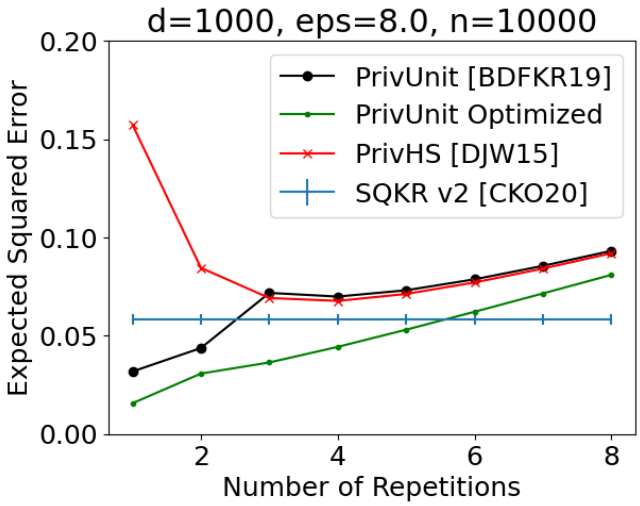

Figure 2. Expected ℓ_2^2 error of mechanisms `PrivHS`, `PrivUnit` and `PrivUnitOptimized` for a total $\varepsilon = 8$, as a function of the number of repetitions of the mechanism with a proportionately smaller ε. The `SQKR` v2 line is for a single run with $\varepsilon = 8$ without splitting.

the choice in most cases will be between not splitting and splitting into two mechanisms. Note that even with two or three repetitions, `PrivUnitOptimized` has $2 - 3\times$ smaller error compared to `PrivHS` and `SQKR`. For `PrivHS`, the sweet spot seems to be splitting into multiple mechanisms each with $\varepsilon \approx 2$.