A Additional proof details

Below we introduce Chernoff's multiplicative bound, that we use in the proof of theorem theorem 4.3

Theorem A.1 (Chernoff multiplicative bound, Theorem D.4. in [26]). *Let $X_1, ..., X_n$ be independent random variables drawn according to some distribution $\mathcal{D}$ with mean μ and support in $[0, M]$. Then for any $\gamma \in [0, \frac{M}{\mu} - 1]$ the following inequalities hold:*

$$P\left(\frac{1}{n}\sum_{i=1}^{n} X_i \geq (1+\gamma)\mu\right) \leq e^{\frac{-n\mu\gamma^2}{3M}} \tag{8}$$

$$P\left(\frac{1}{n}\sum_{i=1}^{n} X_i \leq (1-\gamma)\mu\right) \leq e^{\frac{-n\mu\gamma^2}{2M}} \tag{9}$$

Theorem (theorem 4.3). *Let P and Q be two distributions. Let $h \colon \Omega \subseteq \mathcal{Y} \to \mathbb{R}$ be a function such that $\sup_{x\in\Omega} h(x) < C$, $\mathbf{x} = (x_1, ..., x_n)$ and $\mathbf{y} = (y_1, ..., y_n)$ be n realizations of P and Q, respectively, $\mu_1 = \mathbb{E}_P\left[e^{(\alpha-1)h(x)}\right]$, and $\mu_2 = \mathbb{E}_Q\left[e^{\alpha h(x)}\right]$. Define also $M_1 = e^{(\alpha-1)C}$ and $M_2 = e^{\alpha C}$. Then, if $\gamma \in [0, \min(\frac{M_1}{\mu_1}, \frac{M_2}{\mu_2})]$, and $n \geq \max\left(\frac{3M_1\log(2/\beta)}{\mu_1\gamma^2}, \frac{2M_2\log(2/\beta)}{\mu_2\gamma^2}\right)$, with probability at least $1-\beta$, we have*

$$R_\alpha^h(P||Q) \geq R_\alpha^{h,n}(\mathbf{x}||\mathbf{y}) - \log\left(\frac{1+\gamma}{1-\gamma}\right) \tag{10}$$

Proof. From Chernoff's multiplicative bound (theorem A.1, eq. (8)) we know that for $\gamma_1 \in [0, \frac{M_1}{\mu_1} - 1]$, with probability less than $e^{\frac{-n\mu_1\gamma_1^2}{3M_1}}$, we have

$$\frac{\frac{1}{n}\sum_{i=1}^{n} e^{(\alpha-1)h(x_i)}}{\mathbb{E}[e^{(\alpha-1)h(X)}]} \geq (1+\gamma_1)$$

This implies that

$$\begin{aligned}&\log\left[\frac{1}{n}\sum_{i=1}^{n} e^{(\alpha-1)h(x_i)}\right] - \log\mathbb{E}_P\left[e^{(\alpha-1)h(X)}\right]\\ &\geq \log(1+\gamma_1) \geq \tfrac{\alpha-1}{\alpha}\log(1+\gamma),\end{aligned}$$

or equivalently,

$$\begin{aligned}&\tfrac{\alpha}{\alpha-1}\log\frac{1}{n}\sum_{i=1}^{n} e^{(\alpha-1)h(x_i)} - \tfrac{\alpha}{\alpha-1}\log\mathbb{E}_P\left[e^{(\alpha-1)h(X)}\right]\\ &\geq \log(1+\gamma_1)\end{aligned} \tag{11}$$

Note that by setting $n \geq \frac{3e^{(\alpha-1)C}\log(2/\beta)}{\mu_1\gamma_1^2}$ the above bound holds with probability at most $\beta/2$. A similar analysis (using eq. (9)) shows that for $n \geq \frac{2e^{\alpha C}\log(2/\beta)}{\mu_2\gamma_2^2}$ with probability at most $\beta/2$ the following bound holds:

$$\begin{aligned}&\log\mathbb{E}_Q\left[e^{\alpha h(y)}\right] - \log\left[\frac{1}{n}\sum_{i=1}^{n} e^{\alpha h(y_i)}\right]\\ &\geq \log\left(\frac{1}{1-\gamma_2}\right).\end{aligned} \tag{12}$$

Finally, summing eq. (11) and eq. (12), and using the union bound, with probability $1-\beta$ we have that

$$R^{h,n}(\mathbf{x}||\mathbf{y}) - \log\left(\frac{1}{1-\gamma_2}\right) - \log(1+\gamma_1) \leq R^h(P||Q).$$

The proof follows by letting $\gamma = \min(\gamma_1, \gamma_2)$. □

B Experiment details

B.1 Approximate DP tester [16]

We include the approximate DP tester in

Algorithm 2 ADPTESTER [16]

Input: Universe size m, $\epsilon, \delta, \eta > 0$ approximate DP privacy parameters and approximation error respectively.
$\lambda = \max\{\frac{4m(1+e^{2\epsilon})}{\eta^2}, \frac{12(1+e^{2\epsilon})}{\eta^2}\}$
Sample $r \sim \text{Poisson}(\lambda)$
$X_i \sim \mathcal{M}(D), Y_i \sim \mathcal{M}(D')$ for $i \in [r]$
for $j = 1$ **to** m **do**
    $x_j =$ Number of j's in $\{X_i\}_{i=1}^r$
    $y_j =$ Number of j's in $\{Y_i\}_{i=1}^r$
    $z_i = \frac{1}{r}(x_i - e^{\epsilon} y_i)$
end for
$z = \sum_{i=1}^r \max\{0, z_i\}$
if $z < \delta + \eta$ **then**
    Return `Passed`
else
    Return `False`
end if

B.2 Further comparison with DP-Sniper

Below we include a more complete comparison against DPSNIPER on SVT mechanisms with add/remove and ℓ_∞ neighboring relations and with different sample complexities. DPSNIPER only has an adavntage over RENYITESTER when using the ℓ_∞ relation and using at least 10 million samples. For the more common add/remove definition RENYITESTER has the same performance as DPSNIPER .

Table 3: Comparison between DP-Sniper [5] and RENYITESTER with different sample complexities using add/remove neighboring relation.

mechanism	Renyi Tester	DP-Sniper	Renyi Tester	DP-Sniper
	400K samples		10M samples	
SVT1	passed	passed	passed	passed
SVT2	passed	passed	passed	passed
SVT3	passed	passed	passed	failed
SVT4	passed	passed	passed	passed
SVT5	failed	failed	failed	failed
SVT6	passed	failed	passed	failed

B.3 RENYITESTER implementation details

Function class. For all auditing mechanisms we used as the underlying function class Φ the family of fully connected neural network with two dense layers, each with 100 units. To ensure $h \in \Phi$ are bounded but contain the real value of the divergence (ϵ for which we test) we add scaled hyperbolic tangent activations scaled to 16ϵ.

Privacy parameters. Below we show results with different selections of hyperparameters ϵ and α used for auditing mechanisms. The range of ϵ and α was selected based on sample complexity sizes that allowed us to run the tests in an efficient manner. Besides NONDPMEAN2, we notice that results tend to be consistent across the selections of these parameters.

B.4 Exploring the space of datasets

Average number of trials to detect privacy violations. In **??** we provide details on the number of datasets our algorithm needs to test before finding a dataset where the privacy guarantee is broken.