

A Background on Martingale Concentration

In this section, we provide a background on the basics of martingale concentration needed throughout this paper. Central to all results in this section is Ville's inequality [Ville, 1939], which can be viewed as a time-uniform version of Markov's inequality for martingales.

Lemma A.1 (Ville's Inequality [Ville, 1939]). *Let $(X_t)_{t \geq 0}$ be a nonnegative supermartingale with respect to some filtration $(\mathcal{F}_t)_{t \geq 0}$. Then, for any confidence parameter $\delta \in (0, 1)$, we have*

$$\mathbb{P}\left(\exists t \geq 0 : X_t \geq \frac{\mathbb{E}X_0}{\delta}\right) \leq \delta.$$

While standard Brownian motion $(B_t)_{t \geq 0}$ is not a nonnegative supermartingale, geometric Brownian motion given by $Y_t^\lambda := \exp\left(\lambda B_t - \frac{\lambda^2}{2}t\right)$ is a nonnegative martingale for any $\lambda \in \mathbb{R}$, and hence Lemma A.1 can be applied. In fact, the probability in the lemma above becomes exactly δ when it is applied to a nonnegative martingale with continuous paths like Y_t^λ . From Ville's inequality, the following *line-crossing inequality* for Brownian motion can be obtained.

Lemma A.2 (Line-Crossing Inequality). *For $\delta \in (0, 1)$ and $a, b > 0$ satisfying $e^{-2ab} = \delta$, we have*

$$\mathbb{P}(\exists t \geq 0 : B_t \geq at + b) = \delta.$$

A proof of the above fact can be found in any standard book on continuous time martingale theory [Le Gall, 2016, Durrett, 2019]. The above also follows from a special case of the more general time-uniform Chernoff bound presented in Howard et al. [2020].

The above inequality can be seen as optimizing the tightness of the time-uniform boundary at one pre-selected point in time. However, due to the adaptive nature of the Brownian mechanism presented in Section 3, it is sometimes desirable to construct a time-uniform boundary which sacrifices tightness at a fixed point in time to obtain greater tightness over all of time.

The *method of mixtures* provides one such approach for constructing tighter time-uniform boundaries [Kaufmann and Koolen, 2021, Howard et al., 2021]. We discuss this concept briefly in the context of Brownian motion. Observe that, since $(Y_t^\lambda)_{t \geq 0}$ is a nonnegative martingale, for any probability measure π on \mathbb{R} , the process $(X_t^\pi)_{t \geq 0}$ given by

$$X_t^\pi := \int_{\mathbb{R}} Y_t^\lambda \pi(d\lambda)$$

is also nonnegative martingale. By appropriately choosing the probability measure π and applying Ville's inequality, one obtains the following concentration inequality [Howard et al., 2021].

Lemma A.3 (Mixture Inequality). *Let $\rho > 0$ and $\delta \in (0, 1)$ be arbitrary. Then,*

$$\mathbb{P}\left(\exists t \geq 0 : B_t \geq \sqrt{2(t + \rho) \log\left(\frac{1}{\delta} \sqrt{\frac{t + \rho}{\rho}}\right)}\right) = \delta.$$

We leverage Lemmas A.2 and A.3 to construct the privacy boundaries in Theorem 3.6 in Appendix B.

B Proofs From Section 3

Here, we prove the results from Section 3. We start by showing that BM is in fact a noise-reduction mechanism, per the condition in Definition 2.3.

Proof of Lemma 3.2. For any initial value $\mu \in \mathbb{R}^d$ and any $n \geq 1$, let $p_{1:n}^\mu$ denote the joint density of B_{T_n}, \dots, B_{T_1} where $(B_t)_{t \geq 0}$ is a Brownian motion started at μ and $T_n := T_n(x)$ is a time function. We have the decomposition

$$p_{1:n}^\mu(B_{T_n}, \dots, B_{T_1}) \propto \exp\left(-\frac{(B_{T_1} - \mu)^2}{2T_1}\right) \prod_{m=2}^n \exp\left(\frac{-(B_{T_m} - \mu - \frac{T_m}{T_{m-1}}(B_{T_{m-1}} - \mu))^2}{2(T_{m-1} - T_m)} \cdot \frac{T_{m-1}}{T_m}\right).$$




Figure 4: A comparison of the linear and mixture boundaries, both optimized for tightness at $\epsilon = 0.3$ with $\delta = 10^{-6}$. The first plot directly plots the corresponding bounds as in Theorem 3.6. The second plot inverts the boundaries, showing the variance necessary to meet a target privacy level.

This decomposition follows as $B_{T_1} \sim \mathcal{N}(\mu, T_1)$, and, given $B_{T_1}, \dots, B_{T_{m-1}}$, we have that $B_{T_m} \sim \mathcal{N}\left(\mu + \frac{T_m}{T_{m-1}}(B_{T_{m-1}} - \mu), \frac{(T_{m-1} - T_m)T_m}{T_{m-1}}\right)$, as B_{T_m} is conditionally distributed as a Brownian bridge. Now, a straightforward calculation yields the equivalence

$$\begin{aligned} p_{1:n}^\mu(B_{T_n}, \dots, B_{T_1}) &\propto \exp\left(-\frac{(B_{T_1} - \mu)^2}{2T_1}\right) \prod_{m=2}^n \exp\left(-\frac{(B_{T_m} - \mu - \frac{T_m}{T_{m-1}}(B_{T_{m-1}} - \mu))^2}{2(T_{m-1} - T_m)} \cdot \frac{T_{m-1}}{T_m}\right) \\ &= \exp\left(-\frac{(B_{T_n} - \mu)^2}{2T_n}\right) \prod_{m=2}^n \exp\left(-\frac{(B_{T_{m-1}} - B_{T_m})^2}{2(T_{m-1} - T_m)}\right). \end{aligned}$$

Hence, for any two mean vectors $\mu, \mu' \in \mathbb{R}^d$, we can decompose the ratio of densities as

$$\frac{p^\mu(B_{T_n}, \dots, B_{T_1})}{p^{\mu'}(B_{T_n}, \dots, B_{T_1})} = \frac{\exp\left(-\frac{(B_{T_n} - \mu)^2}{2T_n}\right)}{\exp\left(-\frac{(B_{T_n} - \mu')^2}{2T_n}\right)},$$

which is just the ratio between the density of a $\mathcal{N}(\mu, T_n)$ random variable evaluated at B_{T_n} and the density of a $\mathcal{N}(\mu', T_n)$ random variable evaluated at B_{T_n} , proving precisely the desired result. \square

We now prove Theorem 3.4, which gives a closed form characterization of the Brownian mechanism. In what follows, we use the same notation for the density of Brownian motion as in the above proof.

Proof of Theorem 3.4. The second statement of the theorem is trivial and follows from our assumption of bounded ℓ_2 sensitivity. Hence, we only prove the first statement below.

Without loss of generality, and for the sake of simplicity, we can consider the function $g(y) := f(y) - f(x)$, as then $g(x) = 0$. Observe that, for $y \in \mathcal{X}$, the vector $g(y) + B_t$ has Lebesgue density

$$p_t^{g(y)}(\beta) \propto \exp\left(-\frac{1}{2t}\|\beta - g(y)\|_2^2\right).$$

Consequently, the privacy loss can be written as

$$\begin{aligned}
\mathcal{L}_n^{\text{BM}}(x, x') &= \log \left(\frac{p_{T_n(x)}^0(B_{T_n(x)})}{p_{T_n(x)}^{g(x')}(B_{T_n(x)})} \right) = \frac{1}{2T_n(x)} (-\|B_{T_n(x)}\|_2^2 + \|B_{T_n(x)} - g(x')\|_2^2) \\
&= -\frac{1}{T_n(x)} \langle B_{T_n(x)}, g(x') \rangle + \frac{1}{2T_n(x)} \|g(x')\|_2^2 \\
&= -\frac{\|g(x')\|_2}{T_n(x)} \left\langle B_{T_n(x)}, \frac{g(x')}{\|g(x')\|_2} \right\rangle + \frac{1}{2T_n(x)} \|g(x')\|_2^2 \\
&= -\frac{\|g(x')\|_2}{T_n(x)} \left\langle B_{T_n(x)}, \frac{g(x')}{\|g(x')\|_2} \right\rangle + \frac{1}{2T_n(x)} \|g(x')\|_2^2 \\
&= -\frac{\|g(x')\|_2}{T_n(x)} W_{T_n(x)} + \frac{1}{2T_n(x)} \|g(x')\|_2^2.
\end{aligned}$$

Note that the last inequality follows from the fact that if $(B_t)_{t \geq 0}$ is a d -dimensional Brownian motion and $z \in \mathbb{R}^d$ is a unit vector under the ℓ_2 norm, then the process $W_t := \langle z, B_t \rangle$ is a standard Brownian motion. Noting that $g(x') = f(x') - f(x)$ and that $(-W_t)_{t \geq 0}$ is also a Brownian motion furnishes the result. \square

We now use the characterization of privacy loss in Theorem 3.4 alongside the time-uniform concentration results for continuous time martingales found in Appendix A to construct two general families of privacy boundaries. We now prove Theorem 3.6.

Proof of Theorem 3.6. Recall from Theorem 3.4 that we have the following bound

$$\mathcal{L}_n^{\text{BM}}(x, x') \leq \frac{\Delta^2}{2T_n(x)} + \frac{\Delta}{T_n(x)} W_{T_n(x)}^+,$$

where $A^+ := \max(A, 0)$. First, by leveraging Lemma A.3, we see that, with probability at least $1 - \delta$, for all $n \in \mathbb{N}$, we have

$$\mathcal{L}_n^{\text{BM}}(x, x') \leq \frac{\Delta^2}{2T_n(x)} + \frac{\Delta}{T_n(x)} \sqrt{2(T_n(x) + \rho) \log \left(\frac{1}{\delta} \sqrt{\frac{T_n(x) + \rho}{\rho}} \right)} = \psi_\rho^M(T_n(x)),$$

proving that ψ_ρ^M is a valid δ -privacy boundary. Likewise, by Lemma A.2, we have that

$$\mathcal{L}_n^{\text{BM}}(x, x') \leq \frac{\Delta^2}{2T_n(x)} + \frac{\Delta}{T_n(x)} (aT_n(x) + b) = \frac{\Delta}{T_n(x)} \left(\frac{\Delta}{2} + b \right) + \Delta a = \psi_{a,b}^L(T_n(x)),$$

showing $\psi_{a,b}^L$ is a valid δ -privacy boundary. \square

C Proofs From Section 5

In this appendix, we provide proofs of the results in Section 5. We start by proving the privacy guarantees for `ReducedAboveThreshold`.

Proof of Theorem 5.1. For `ReducedAboveThreshold` as described in Algorithm 1, on the event $\{N(x) = n\}$, all information leaked about the underlying private dataset is contained in $\text{Alg}_{1:n}(x)$ and $\alpha_{1:n}(x)$, where $\alpha_n(x)$ is defined to be the n th bit output by `ReducedAboveThreshold`. For any $y \in \mathcal{X}$, let $q_{1:n}^y$ denote the joint density of $(\text{Alg}_{1:n}(y), \alpha_{1:n}(y))$, $p_{1:n}^y$ the marginal density of $\text{Alg}_{1:n}(y)$, and $p_{1:n}^y(\cdot \mid \cdot)$ the conditional pmf of $\alpha_{1:n}(y)$ given the observed values of $\text{Alg}_{1:n}(y)$. As such, for any neighboring datasets $x \sim x'$, on the event $\{N(x) = n\}$, the privacy loss of

ReducedAboveThreshold, denoted by $\mathcal{L}^{\text{RAT}}(x, x')$, is given by

$$\begin{aligned}
\mathcal{L}^{\text{RAT}}(x, x') &= \log \left(\frac{q_{1:n}^x(\text{Alg}_{1:n}(x), \alpha_{1:n}(x))}{q_{1:n}^{x'}(\text{Alg}_{1:n}(x), \alpha_{1:n}(x))} \right) \\
&= \log \left(\frac{p_{1:n}^x(\text{Alg}_{1:n}(x))}{p_{1:n}^{x'}(\text{Alg}_{1:n}(x))} \right) + \log \left(\frac{p_{1:n}^x(\alpha_{1:n}(x) \mid \text{Alg}_{1:n}(x))}{p_{1:n}^{x'}(\alpha_{1:n}(x) \mid \text{Alg}_{1:n}(x))} \right) \\
&= \log \left(\frac{p_{1:n}^x(\text{Alg}_{1:n}(x))}{p_{1:n}^{x'}(\text{Alg}_{1:n}(x))} \right) + \log \left(\frac{p_{1:n}^x(0^{n-1}1 \mid \text{Alg}_{1:n}(x))}{p_{1:n}^{x'}(0^{n-1}1 \mid \text{Alg}_{1:n}(x))} \right) \\
&= \mathcal{L}_{1:n}^{\text{Alg}}(x, x') + L_n(x, x'),
\end{aligned}$$

where $0^{n-1}1$ denotes the string of $n-1$ 0's followed by a single 1. In the last line we leverage the definition of the privacy loss between $\text{Alg}_{1:n}(x)$ and $\text{Alg}_{1:n}(x')$ and define

$$L_n(x, x') := \log \left(\frac{p_{1:n}^x(0^{n-1}1 \mid \text{Alg}_{1:n}(x))}{p_{1:n}^{x'}(0^{n-1}1 \mid \text{Alg}_{1:n}(x))} \right).$$

Now, to finish the result, it suffices to prove that, for any n , $L_n(x, x') \leq \mathcal{E}_n(\text{Alg}_{1:n-1}(x))$. Without loss of generality, we can assume all thresholds take the same value τ across rounds, as we can always define the shifted function $u'_n(\text{Alg}_{1:n}(x), x) := u_n(\text{Alg}_{1:n}(x), x) - \tau_n + \tau$. To prove our desired inequality, we proceed largely in the same way as the proof of AboveThreshold found in Lyu et al. [2017], noting that conditioning on $\text{Alg}_{1:n}(x)$ serves to fix the utility functions $u_1(\text{Alg}_1(x), \cdot), \dots, u_n(\text{Alg}_{1:n}(x), \cdot)$ and the privacy levels $\mathcal{E}_1, \mathcal{E}_2(\text{Alg}_1(x)), \dots, \mathcal{E}_n(\text{Alg}_{1:n-1}(x))$. For simplicity, going forward, we refer to the former quantities as $u_1(\cdot), \dots, u_n(\cdot)$ and the latter quantities just as $\epsilon_1, \dots, \epsilon_n$. The only remaining caveat that we must take care in handling variable amount of noise on the threshold introduced by LNR. Going forward, let $\mathbb{P}_{1:n}$ denote the conditional probability $\mathbb{P}(\cdot \mid \text{Alg}_{1:n}(x))$. First, observe that we can write the numerator of $L_n(x, x')$ as

$$p^x(0^{n-1}1 \mid \text{Alg}_{1:n}(x)) = \int_{\mathbb{R}^n} g_{1:n}^\tau(s_1, \dots, s_n) \left(\prod_{i=1}^{n-1} \mathbb{P}_{1:n}(u_i(x) + \xi_i < s_i) \right) \mathbb{P}_{1:n}(u_n(x) + \xi_n \geq s_n) d\vec{s},$$

where $g_{1:n}^\tau$ represents the density for the joint distribution of $(\tau + Z(2\Delta/\epsilon_m))_{m=1}^n$, where $(Z(t))_{t \geq \eta}$ is as defined in Equation (5). We now need three inequalities. The first two are standard from the analysis of Lyu et al. [2017], so we do not provide a proof. The third inequality is a product of our novel ReducedAboveThreshold mechanism, and hence we provide a proof. The inequalities are:

1. For $i < n$ and fixed s_i , $\mathbb{P}_{1:n}(u_i(x) + \xi_i < s_i) \leq \mathbb{P}_{1:n}(u_i(x') + \xi_i < s_i + \Delta)$,
2. for $i = n$ and any s_n , $\mathbb{P}_{1:n}(u_n(x) + \xi_n \geq s_n) \leq e^{\epsilon_n/2} \mathbb{P}_{1:n}(u_n(x') + \xi_n \geq s_n + \Delta)$, and
3. for any $s_{1:n} \in \mathbb{R}^n$, $g_{1:n}^\tau(s_1, \dots, s_n) \leq e^{\epsilon_n/2} g_{1:n}^\tau(s_1 + \Delta, \dots, s_n + \Delta)$.

We now prove the third inequality. We have that

$$\begin{aligned}
\frac{g_{1:n}^\tau(s_1, \dots, s_n)}{g_{1:n}^{\tau-\Delta}(s_1, \dots, s_n)} &= \frac{g_n^\tau(s_n) g_{1:n-1}^\tau(s_1, \dots, s_{n-1} \mid s_n)}{g_n^{\tau-\Delta}(s_n) g_{1:n-1}^{\tau-\Delta}(s_1, \dots, s_{n-1} \mid s_n)} \\
&= \frac{g_n^\tau(s_n)}{g_n^{\tau-\Delta}(s_n)} \leq e^{\epsilon_n/2},
\end{aligned}$$

where the first equality follows from applying Bayes rule to the joint densities of the noisy thresholds, and the second equality follows from the fact that $(Z(t))$ forms a Markov process. This in particular implies that the density conditional density given the n th threshold satisfies $g_{1:n-1}^a(s_1, \dots, s_{n-1} \mid s_n) = g_{1:n-1}^b(s_1, \dots, s_{n-1} \mid s_n)$ for all $a, b \in \mathbb{R}$. The last inequality follows from examining the ratio of densities of $\text{Lap}(\tau, 2\Delta/\epsilon_n)$ and $\text{Lap}(\tau - \Delta, 2\Delta/\epsilon_n)$ random variables. Now, observe that by a simple shift of parameters we have

$$g_{1:n}^{\tau-\Delta}(s_1, \dots, s_n) = g_{1:n}^\tau(s_1 + \Delta, \dots, s_n + \Delta).$$

Plugging this in, we have

$$\begin{aligned}
& p^x(0^{n-1}1 \mid \text{Alg}_{1:n}(x)) \\
&= \int_{\mathbb{R}^n} g_{1:n}^\tau(s_1, \dots, s_n) \left(\prod_{i=1}^{n-1} \mathbb{P}_{1:n}(u_i(x) + \xi_i < s_i) \right) \mathbb{P}_{1:n}(u_n(x) + \xi_n \geq s_n) d\vec{s} \\
&\leq e^{\epsilon_n/2} \int_{\mathbb{R}^n} g_{1:n}^{T-\Delta}(s_1, \dots, s_n) \left(\prod_{i=1}^{n-1} \mathbb{P}_{1:n}(u_i(x) + \xi_i < s_i) \right) \mathbb{P}_{1:n}(u_n(x) + \xi_n \geq s_n) d\vec{s} \\
&\leq e^{\epsilon_n} \int_{\mathbb{R}^n} g_{1:n}^{T-\Delta}(s_1, \dots, s_n) \left(\prod_{i=1}^{n-1} \mathbb{P}_{1:n}(u_i(x') + \xi_i < s_i + \Delta) \right) \mathbb{P}_{1:n}(u_n(x') + \xi_n \geq s_n + \Delta) d\vec{s} \\
&= e^{\epsilon_n} \int_{\mathbb{R}^n} g_{1:n}^T(s_1, \dots, s_n) \left(\prod_{i=1}^{n-1} \mathbb{P}_{1:n}(u_i(x') + \xi_i < s_i) \right) \mathbb{P}_{1:n}(u_n(x') + \xi_n \geq s_n) d\vec{s} \\
&= e^{\epsilon_n} p^{x'}(0^{n-1}1 \mid \text{Alg}_{1:n}(x)).
\end{aligned}$$

Rearranging furnishes the desired result. \square

We can also prove a corresponding utility guarantee for `ReducedAboveThreshold`. As mentioned earlier, this utility guarantee is naive in the sense that it is derived from a union bound. Thus, instead of plotting the utility guarantee in our experiments in Section 6, we instead plot empirically observed loss/accuracy. Additionally, for the utility guarantee to hold, the sequence of privacy functions $(\mathcal{E}_n)_{n \geq 1}$ must be constant functions, i.e. $\mathcal{E}_n = \epsilon_n$ for each n . We now state the formal, high-probability utility guarantee in the following proposition.

Proposition C.1. *Let $(p_n)_{n \geq 1}$ be a sequence of non-negative numbers such that $\sum_{i=1}^{\infty} p_i = 1$, and let $\gamma \in (0, 1)$ be a confidence parameter. Define the sequence of parameters $(\eta_n)_{n \geq 1}$ by*

$$\eta_n := \frac{4\Delta}{\epsilon_n} \left(\log \left(\frac{2}{\gamma} \right) - \log(p_n) \right).$$

Then, if $N(x)$ is the time defined in Theorem 5.1, with probability at least $1 - \gamma$, we have

$$u_{N(x)}(x) \geq \tau_{N(x)} - \eta_{N(x)}.$$

Proof. The above utility guarantee follows from applying two simple union bounds. First, we have

$$\mathbb{P} \left(\bigcup_{n \geq 1} \{|\xi_n| \geq \eta_n/2\} \right) \leq \sum_{n \geq 1} \mathbb{P}(|\xi_n| \geq \eta_n/2) = \sum_{n \geq 1} \exp \left(\frac{-\epsilon_n \eta_n}{4\Delta} \right) = \frac{\gamma}{2} \sum_{n \geq 1} p_n = 1.$$

Second, we have that

$$\mathbb{P} \left(\bigcup_{n \geq 1} \{|\zeta_n| \geq \eta_n/2\} \right) \leq \sum_{n \geq 1} \mathbb{P}(|\zeta_n| \geq \eta_n/2) = \sum_{n \geq 1} \exp \left(\frac{-\epsilon_n \eta_n}{2\Delta} \right) \leq \frac{\gamma}{2} \sum_{n \geq 1} p_n = 1.$$

Thus, with probability at least $1 - \gamma$, we have simultaneously for all $n \geq 1$ that $|\xi_n| \leq \eta_n/2$ and $|\zeta_n| \leq \eta_n/2$. Thus, with the same probability, on round $N(x)$, we have

$$u_{N(x)}(x) \geq \tau_{N(x)} - \eta_{N(x)}.$$

\square

D Proofs From Section 4

We first prove that the process defined in Equation (5) has Laplace marginal distributions.

Theorem D.1. *Let $(Z_t)_{t \geq \eta}$ be the process defined in Equation (5). Then, for any $t \geq \eta$, we have*

$$Z_t \sim \text{Lap}(t).$$

In what follows, we sometimes use the notation $Z(t)$ interchangeably with Z_t for convenience.

Proof. Recall that if $X \sim \text{Lap}(s)$, then X has characteristic function φ_s given by

$$\varphi_s(\lambda) = \frac{1}{1 + \lambda^2 s^2}.$$

Let ϕ denote the characteristic function of $Z_t - Z_\eta$. Since Z_η and $Z_t - Z_\eta$ are independent, to show $Z_t \sim \text{Lap}(t)$, it suffices to show that

$$\phi(\lambda) = \frac{\varphi_t(\lambda)}{\varphi_\eta(\lambda)} = \frac{1 + \lambda^2 \eta^2}{1 + \lambda^2 t^2}.$$

Now, observe that the inhomogenous Poisson process $(P_t)_{t \geq \eta}$ can be written as $(\tilde{P}(e^{t/2}))_{t \geq \log(\eta^2)}$ where \tilde{P} is a homogeneous Poisson process with rate $\lambda = 1$ on $[\log(\eta^2), \infty)$. In terms of the process \tilde{P} , we can consider the process $(\tilde{Z}_t)_{t \geq \log(\eta^2)}$ given by

$$\tilde{Z}_t = \sum_{n \leq \tilde{P}_t} \text{Lap}\left(e^{\tilde{T}_n/2}\right),$$

where $\tilde{T}_n := \inf\{t \geq \log(\eta^2) : \tilde{P}_t \geq n\}$ and $\tilde{T}_0 = \log(\eta^2)$. It is easy to see that

$$\tilde{Z}(\log(t^2)) - \tilde{Z}(\log(\eta^2)) =_d Z_t - Z_\eta.$$

Leveraging this identity, it follows that we have

$$\begin{aligned} \phi(\lambda) &= \mathbb{E}\left[e^{i\lambda(Z_t - Z_\eta)}\right] = \mathbb{E}\left[e^{i\lambda(\tilde{Z}(\log(t^2)) - \tilde{Z}(\log(\eta^2)))}\right] \\ &= \sum_{n=0}^{\infty} \frac{\eta^2}{t^2} \frac{[\log(t^2/\eta^2)]^n}{n!} \int_{\log(\eta^2) \leq u_1 < u_2 < \dots < u_n \leq \log(t^2)} f^{(n)}(u_1, \dots, u_n) \prod_{i=1}^n \mathbb{E}\left[e^{i\lambda \text{Lap}(e^{u_i/2})}\right] d\mathbf{u} \\ &= \frac{\eta^2}{t^2} \sum_{n=0}^{\infty} \int_{\log(\eta^2) \leq u_1 < u_2 < \dots < u_n \leq \log(t^2)} \prod_{i=1}^n \frac{1}{1 + \lambda^2 e^{u_i}} d\mathbf{u}. \end{aligned} \quad (6)$$

In the above, $f^{(n)}(u_1, \dots, u_n) := \frac{n!}{[\log(t^2/\eta^2)]^n}$ is the distribution of the order statistics $(U_{(1)}, \dots, U_{(n)})$ of n i.i.d. random variables that are uniform on $[\log(\eta^2), \log(t^2)]$. Essentially, what we have done is *first* conditioned of the number of Poisson arrivals that occur in the interval $[\log(\eta^2), \log(t^2)]$. Then, on the event $\{N(t) = n\}$, we condition again on the location of the n arrivals, which we know to be uniformly distributed across the time interval. Once the arrival locations are known, we can compute the conditional characteristic function, which is the the product of characteristic functions as illustrated in the integral above.

Now, we show inductively that

$$\int_{\log(\eta^2) \leq u_1 < u_2 < \dots < u_n \leq \log(t^2)} \prod_{i=1}^n \frac{1}{1 + \lambda^2 e^{u_i}} d\mathbf{u} = \frac{1}{n!} \left[\log\left(\frac{t^2}{\eta^2} \frac{1 + \lambda^2 \eta^2}{1 + \lambda^2 t^2}\right) \right]^n.$$

The base case of $n = 1$ is trivially true. Now, we have that

$$\begin{aligned} &\int_{\log(\eta^2) \leq u_1 < u_2 < \dots < u_n \leq \log(t^2)} \prod_{i=1}^n \frac{1}{1 + \lambda^2 e^{u_i}} d\mathbf{u} \\ &= \int_{u_1 = \log(\eta^2)}^{\log(t^2)} \frac{1}{1 + \lambda^2 e^{u_1}} \int_{u_1 < u_2 < \dots < u_n \leq \log(t^2)} \prod_{i=2}^n \frac{1}{1 + \lambda^2 e^{u_i}} d\mathbf{u}_{-1} du_1 \\ &= \frac{1}{(n-1)!} \int_{u = \log(\eta^2)}^{\log(t^2)} \frac{1}{1 + \lambda^2 e^u} \left[\log\left(\frac{t^2}{e^u} \frac{1 + \lambda^2 e^u}{1 + \lambda^2 t^2}\right) \right]^{n-1} du \\ &= \frac{1}{n!} \int_{\log(\eta^2)}^{\log(t^2)} \frac{d}{du} \left[-\log\left(\frac{t^2}{e^u} \frac{1 + \lambda^2 e^u}{1 + \lambda^2 t^2}\right) \right]^n du = \frac{1}{n!} \left[\log\left(\frac{t^2}{\eta^2} \frac{1 + \lambda^2 \eta^2}{1 + \lambda^2 t^2}\right) \right]^n. \end{aligned}$$

Leveraging this identity and picking up from the expression for $\phi(\lambda)$ in Equation (6), we have that

$$\begin{aligned}\phi(\lambda) &= \frac{\eta^2}{t^2} \sum_{n=0}^{\infty} \frac{1}{n!} \left[\log \left(\frac{t^2}{\eta^2} \frac{1 + \lambda^2 \eta^2}{1 + \lambda^2 t^2} \right) \right]^n \\ &= \frac{\eta^2}{t^2} \exp \left(\log \left(\frac{t^2}{\eta^2} \frac{1 + \lambda^2 \eta^2}{1 + \lambda^2 t^2} \right) \right) = \frac{1 + \lambda^2 \eta^2}{1 + \lambda^2 t^2}.\end{aligned}$$

This proves the desired result. \square

The above proof can also be leveraged to show that, for any finite fixed sequence of times $(t_n)_{n \in [K]}$, $(Z(t_1), \dots, Z(t_K))$ has the same distribution as $(\zeta_1, \dots, \zeta_K)$, where $(\zeta_n)_{n \in [K]}$ is the Laplace process associated with times $(t_n)_{n \in [K]}$ as outlined in Equation (4). This justifies that the process $(Z(t))_{t \geq \eta}$ is in fact a continuous time generalization of the aforementioned discrete time process.

E Additional Experimental Details

Parameter settings: We set the regularization parameter to be $\lambda = 0.05$ and note that the ℓ_2 and ℓ_1 -sensitivity for the output perturbation of logistic regression are respectively $\frac{2}{n\lambda}$ and $\frac{2\sqrt{d}}{n\lambda}$. Likewise, for covariance perturbation in ridge regression, the ℓ_2 -sensitivities for privately releasing $X^T X$ and $X^T y$ are both 2.0, and the corresponding ℓ_1 -sensitivities for releasing these quantities are $2.0d$ and $2.0\sqrt{d}$ respectively [Ligett et al., 2017, Chaudhuri et al., 2011]. We set the failure probability for BM to be $\delta = 10^{-6}$, and in each task map privacy parameters (ϵ_n) to times (t_n) using the linear privacy boundary $\psi_{a,b}^L$ optimized for tightness at $\epsilon = 0.3$.

Optimizing privacy boundaries: We provide a high level description of how one may set the parameters associated with the privacy boundaries discussed in Theorem 3.6. Let us consider the case of the mixture boundary ψ_ρ^M for illustrative purposes.

Suppose a data analyst desires that the final level of privacy loss obtained by interacting with the Brownian mechanism should be approximately ϵ . Then, intuitively, the analyst should want to add the variance of the Gaussian noise added to be as small as possible when the privacy boundary takes value ϵ . In mathematical notation, the analyst wants to find a parameter ρ^* satisfying

$$\rho^* = \arg \min_{\rho} (\psi_\rho^M)^{-1}(\epsilon),$$

where we note that the inverse function $(\psi_\rho^M)^{-1}$ exists as ψ_ρ^M is strictly increasing. While this inverse has no closed form in general, the parameter ρ^* can be efficiently computed using a few lines of code. A similar, even more straightforward computation can be conducted for the linear privacy boundary.

Simulating Noise Reduction Mechanisms: We briefly describe how a data analyst can produce samples from the Brownian mechanism and the Laplace noise reduction mechanism. First, since $T_1(x)$ is a constant, we have $\text{BM}_1(x) \sim \mathcal{N}(f(x), T_1(x))$. Then, given $\text{BM}_{1:m-1}(x)$, we have $\text{BM}_m(x) \sim \mathcal{N}\left(f(x) + \frac{T_m(x)}{T_{m-1}(x)}(B_{T_{m-1}}(x) - f(x)), \frac{(T_{m-1}(x) - T_m(x))T_m(x)}{T_{m-1}(x)}\right)$. Since simulating the Brownian mechanism only requires normal samples, it can be efficiently computed.

Second, to sample from LNR, one can first generate the the points of arrival of the inhomogeneous Poisson process $(P_t)_{t \geq \eta}$ up to time $T_1(x)$. Let $\mathcal{T}_1, \dots, \mathcal{T}_N$ denote these arrival times, where we note that N , the number of arrivals up to time $T_1(x)$, is a random variable. Then, one can generate $Y_m \sim \text{Lap}(\mathcal{T}_m)$ for $m \leq N$. From this information, the process $(Z_t)_{\eta \leq t \leq T_1(x)}$ can be readily computed, as in Equation (5).