
DP-SPRT: Differentially Private Sequential Probability Ratio Tests

Thomas Michel

Debabrota Basu

Emilie Kaufmann

Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9198-CRISTAL, F-59000 Lille, France

Abstract

We revisit Wald’s celebrated Sequential Probability Ratio Test for sequential tests of two simple hypotheses, under privacy constraints. We propose **DP-SPRT**, a wrapper that can be calibrated to achieve desired error probabilities and privacy constraints, addressing a significant gap in previous work. **DP-SPRT** relies on a private mechanism that processes a sequence of queries and stops after privately determining when the query results fall outside a predefined interval. This **OutsideInterval** mechanism improves upon naive composition of existing techniques like `AboveThreshold`, achieving a factor-of-2 privacy improvement and thus potentially benefiting other continual monitoring procedures. We prove generic upper bounds on the error and sample complexity of **DP-SPRT** that can accommodate various noise distributions based on the practitioner’s privacy needs. We exemplify them in two settings: Laplace noise (pure Differential Privacy) and Gaussian noise (Rényi differential privacy). In the former setting, by providing a lower bound on the sample complexity of any ϵ -DP test with prescribed type I and type II errors, we show that **DP-SPRT** is near optimal when both errors are small and the two hypotheses are close. Moreover, we conduct an experimental study revealing its good practical performance.

1 INTRODUCTION

Sequential hypothesis testing offers an elegant solution when data collection is costly or time-consuming, allowing decisions to be reached with fewer observations on

average compared to fixed-sample approaches. The Sequential Probability Ratio Test (SPRT), developed by [Wald \(1945\)](#), is the foundational algorithm in this field, achieving optimal performance for simple hypothesis testing with independent and identically distributed observations.

Sequential testing procedures are evaluated based on their error rates (Type I and Type II errors) and expected sample size (or *sample complexity*) under each hypothesis. [Wald \(1948\)](#) proved that SPRT minimizes sample complexities subject to error probability constraints—a remarkably strong optimality result in statistical decision theory. Despite its theoretical optimality and applications in clinical trials ([Armitage, 1950, 1954](#)), quality control, and manufacturing ([Wald, 2004](#); [Siegmund, 2013](#)), SPRT faces a significant modern challenge: *privacy*. When applied to sensitive data, such as patient records in medical trials, financial information, or personal behavior patterns, sequential testing procedures risk exposing confidential information. Privacy concerns are particularly acute for the most recent observations, as the decision to terminate data collection can reveal information about them.

Consider a clinical trial where patients sequentially receive treatments and their responses are recorded. With the classical SPRT, if the test concludes that a drug increases the chance of observing an undesirable effect, then we can deduce the effect actually happened to the last participant in the study. More sophisticated attacks could deduce many other pieces of private information about participants throughout the trial. Even if individual patient responses are kept confidential, the pattern of treatment assignments and the stopping decision itself can leak sensitive information about participants. This privacy risk requires a formal framework for privacy-preserving sequential testing.

Related Work. Differential privacy (DP) ([Dwork et al., 2006](#)) has emerged as the standard framework for addressing privacy concerns in statistical analyses. By ensuring that the presence or absence of any single data point cannot significantly affect the outcome, DP offers a principled method to protect individual data while enabling meaningful statistical infer-

ence. Researchers have explored DP hypothesis testing—primarily in the fixed-sample setting. For example, Gaboardi et al. (2016) studies private chi-squared testing, Canonne et al. (2019) investigates optimal private tests for simple hypotheses, Couch et al. (2019) studies nonparametric hypothesis testing with DP, and Kazan et al. (2023) proposes a generic framework for any fixed-sample hypothesis test. These approaches, however, are not directly applicable to sequential testing due to the adaptive nature of sample collection.

For sequential settings, existing differentially private mechanisms like AboveThreshold (Dwork et al., 2006) and the sparse vector technique enable monitoring whether query results exceed specified thresholds while preserving privacy. However, many applications—including sequential hypothesis testing—require monitoring both upper and lower bounds simultaneously, leading to suboptimal privacy guarantees when using naive composition.

Wang et al. (2022) explored the privatization of SPRT by adding Laplace noise to thresholds and using the exponential mechanism for the final decision. However, due to their specific use case, their approach provides a weaker notion of privacy that only requires that critical privacy failures do not happen only for a subset of plausible datasets. In contrast, our work maintains DP throughout the testing process while working to minimize stopping time. Zhang et al. (2022) developed PrivSPRT, a private version of the SPRT using Gaussian noise with Renyi DP guarantees (Mironov, 2017). The authors provide upper bounds on the error probabilities given fixed thresholds and noise level, but do not specify how to calibrate these thresholds to achieve desired Type I and Type II error rates while maintaining a desired privacy guarantee. In practice, this forces to rely on Monte Carlo simulation to determine appropriate thresholds, a process that may itself leak privacy in certain scenarios (Papernot and Steinke, 2021).

Another related line of work is Differentially Private Best Arm Identification (DP-BAI) in bandits (Sajed and Sheffet, 2019; Kalogerias et al., 2021; Azize et al., 2023, 2024). BAI is a generalization of sequential hypothesis testing, where we simultaneously test means of multiple distributions to identify the largest one. (Azize et al., 2023, 2024) design algorithms for ε -DP-BAI whose sample complexity has irreducible problem-dependent constants compared to the lower bound, keeping the optimal algorithm design open. Thus, we focus on the archetypal setup of testing two hypotheses, and aim to explicate the privacy-utility trade-offs.

Contributions. (1) We introduce a general framework for differentially private SPRTs, called **DP-SPRT**,

that supports *various noise mechanisms* while providing privacy guarantees and error control. We propose a generic sample complexity analysis of **DP-SPRT** for Bernoulli observations and particularize it to Gaussian noise (for Renyi DP) and Laplace noise (for ε -DP). (2) This framework is based on the novel **OutsideInterval** mechanism for differentially privately comparing a sequence of queries to two thresholds. This might be of independent interest for applications to DP problems with or without continual observations. (3) We further prove a lower bound on the sample complexity of any ε -DP algorithm that achieves prescribed type I and type II errors. This lower bound reveals two distinct regimes: one where privacy has minimal impact on sample complexity, and another where higher privacy requirements substantially increase the expected sample size. It also shows that *DP-SPRT with Laplace noise is near optimal in a regime of small errors and when the two hypotheses are close*. (4) On the practical side, we use subsampling to propose a practical improvement of **DP-SPRT** with better privacy-utility trade-off, especially in high privacy regimes.

Our analysis of **DP-SPRT** is conducted for Bernoulli observations, which model binary outcomes present in applications such as clinical trials, quality control, and A/B testing. Bernoulli distributions form a one-dimensional exponential family with bounded support in $[0, 1]$, making them suitable for DP mechanisms requiring bounded sensitivity. This setting provides an illustration of privacy-utility trade-offs while maintaining mathematical tractability.

2 PRELIMINARIES: SEQUENTIAL TESTS, DIFFERENTIAL PRIVACY, AND SPRT

Sequential hypothesis testing allows collecting data incrementally and making decisions as soon as sufficient evidence accumulates. Given a stream of i.i.d. observations X_i from some parametric distribution ν_θ with $\theta \in \Theta$, we consider two simple hypotheses $\mathcal{H}_0 : (\theta = \theta_0)$ and $\mathcal{H}_1 : (\theta = \theta_1)$, where θ_0 and θ_1 are two distinct parameters in Θ . We let \mathbb{P}_θ be the probability space under which the sequence $(X_i)_{i \in \mathbb{N}}$ is i.i.d. under ν_θ .

A sequential test consists of a *stopping rule* that determines when to stop collecting data, and a *decision rule* that selects one of the hypotheses. Specifically, a sequential test yields a pair (τ, \hat{d}) , where τ is a stopping time with respect to the filtration $\mathcal{F}_t = \sigma(X_1, \dots, X_t)$, and the final decision $\hat{d} \in \{0, 1\}$ is \mathcal{F}_τ -measurable. Given prescribed error levels $\alpha, \beta \in (0, 1)$, we seek to design tests that have a type I error smaller than α and a type II error smaller than β . These tests should have a small *sample complexity*, $\mathbb{E}_{\theta_i}[\tau]$, under either

hypotheses.

Definition 1 ((α, β) -Correct Sequential Test). *A sequential test (τ, \hat{d}) is (α, β) -correct if it satisfies $\mathbb{P}_{\theta_0}(\hat{d} = 1) \leq \alpha$ and $\mathbb{P}_{\theta_1}(\hat{d} = 0) \leq \beta$.*

Differentially Private Sequential Tests. Differential Privacy (DP) (Dwork et al., 2006) is usually defined for a randomized mechanism based on a dataset of fixed size. It is expressed as a constraint on the randomness used by the mechanism: the distribution of its output cannot change too much if a single entry in its input (the dataset) is modified. It is possible to extend DP to sequential tests by considering as input an *infinite observation sequence* $\underline{X} := (X_1, X_2, \dots)$. The mechanism $\mathcal{M}(\underline{X}) := (\tau(\underline{X}), \hat{d}(\underline{X}))$ outputs the stopping time and decision of the sequential test run on the data stream \underline{X} (regardless of the distribution of that stream). We formalize below for sequential tests different notions of privacy commonly used in the literature.

Definition 2 (Neighboring Observation Sequences). *Two observation sequences $\underline{X} := (X_1, X_2, \dots)$ and $\underline{X}' := (X'_1, X'_2, \dots)$ are neighboring if they differ in at most one observation, i.e., if there exists at most one index i such that $X_i \neq X'_i$.*

Definition 3 (Differentially Private Sequential Test). *A sequential test with associated randomized mechanism $\mathcal{M}(\underline{X}) := (\tau(\underline{X}), \hat{d}(\underline{X}))$ satisfies:*

(i) (ε, δ) -**DP** if for all neighboring observation sequences \underline{X} and \underline{X}' , and for all possible outputs $(t, d) \in \mathbb{N} \times \{0, 1\}$,

$$\Pr[\mathcal{M}(\underline{X}) = (t, d)] \leq e^\varepsilon \Pr[\mathcal{M}(\underline{X}') = (t, d)] + \delta.$$

We say that the test satisfies ε -pure DP if $\delta = 0$.

(ii) (α, ε) -**Rényi DP** if for all neighboring observation sequences \underline{X} and \underline{X}' :

$$D_\alpha(\mathcal{M}(\underline{X}) \parallel \mathcal{M}(\underline{X}')) \leq \varepsilon$$

where $D_\alpha(P \parallel Q) = \frac{1}{\alpha-1} \log \mathbb{E}_{x \sim Q} \left[\left(\frac{P(x)}{Q(x)} \right)^\alpha \right]$ is the Rényi divergence of order α . A test has an RDP profile $\tilde{\varepsilon}(\alpha)$ if it satisfies $(\alpha, \tilde{\varepsilon}(\alpha))$ -RDP for all $\alpha > 1$.

Rényi Differential Privacy (RDP) (Mironov, 2017) provides a useful alternative framework for privacy analysis that often leads to tighter composition bounds than standard differential privacy. (α, ε') -RDP can be converted to approximate (ε, δ) -DP (Dwork et al., 2014) using standard conversion formula, i.e., for $\varepsilon > \varepsilon'$, $\delta = \exp(-(\alpha-1)(\varepsilon - \varepsilon'))$ (Abadi et al., 2016).

Definition 3 captures the fact that the output of a sequential test consists of both when to stop and what decision to make. Thus, we must ensure that a single observation cannot substantially affect either the stopping time or the final decision. Our goal is to develop

(α, β) -correct tests with privacy guarantees, expressed either in DP or in RDP. To do so, it is natural to rely on an optimal (non-private) (α, β) -correct test—the SPRT (Wald, 1948).

SPRT: The Sequential Probability Ratio Test. Let f_θ denote the density of ν_θ with respect to some reference measure. Then, SPRT with thresholds γ_0 and γ_1 is defined as:

$$\tau := \inf \left\{ t \in \mathbb{N} : \Gamma_t := \prod_{i=1}^t \frac{f_{\theta_1}(X_i)}{f_{\theta_0}(X_i)} \notin (\gamma_0, \gamma_1) \right\} \text{ and}$$

$$\hat{d} = \begin{cases} 0 & \text{if } \Gamma_\tau \leq \gamma_0 \\ 1 & \text{if } \Gamma_\tau \geq \gamma_1 \end{cases}.$$

What makes SPRT remarkable is its optimality. As shown by Wald (1948), for any test with the same error probabilities as SPRT, the expected sample size cannot be smaller under either hypothesis. Still, using SPRT in practice requires calibrating the thresholds γ_0 and γ_1 to make it (α, β) -correct. In this work, we rely on an exact calibration¹, i.e., $\gamma_0 = \beta$ and $\gamma_1 = 1/\alpha$, for which SPRT can be proved to be (α, β) -correct using an elegant change-of-measure argument (see Appendix D).

For one-parameter canonical exponential families (Brown, 1986), in which $\Theta \subseteq \mathbb{R}$ and the density can be written $f_\theta(x) = \exp(\theta x - b(\theta))$ where $b : \Theta \rightarrow \mathbb{R}$ is called the log-partition function, SPRT takes a more explicit form (see Appendix A). To fix the ideas, we assume in sequel that $\theta_0 < \theta_1$. Letting $\bar{X}_n = \sum_{i=1}^n X_i/n$ denote the empirical mean, μ_i be the mean of the distribution ν_{θ_i} , and KL the Kullback-Leibler divergence, we have $\tau := \min(\tau_0, \tau_1)$, where

$$\tau_0 := \inf \left\{ n \in \mathbb{N}_{>0} : \bar{X}_n \leq \ell_0^n \right\},$$

$$\tau_1 := \inf \left\{ n \in \mathbb{N}_{>0} : \bar{X}_n \geq \ell_1^n \right\} \quad (1)$$

with thresholds:

$$\ell_0^n = \mu_0 + \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \log(1/\beta)/n}{\theta_1 - \theta_0},$$

$$\ell_1^n = \mu_1 - \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) - \log(1/\alpha)/n}{\theta_1 - \theta_0}.$$

and $\hat{d} = i$ if $\tau = \tau_i$. This class of distributions encompasses several well-known distributions such as Bernoulli, Poisson or Gaussian distributions with known variance. We propose in Section 4 a private version of this test, and analyze it for Bernoulli distributions. SPRT relies on the empirical mean, and

¹Wald (1945) rather advocates the use of the heuristic thresholds $\gamma_0 = \frac{\beta}{1-\alpha}$ and $\gamma_1 = \frac{1-\beta}{\alpha}$ but they cannot guarantee both type I and type II errors

private versions of the empirical mean are *mostly* available under a bounded sensitivity assumption², which holds for bounded distributions only. Note that, we have approaches that bound the sensitivity of subGaussian variables with high probability. But it introduces a sample-dependent $\sqrt{\log n}$ term in the sensitivity and loss in the probability of correctness (Kamath et al., 2022). To avoid this nuance and focus on explicating tight privacy-utility trade-off, we focus on the Bernoulli distributions.

Equation (1) reveals a clear structure of SPRT. We have two thresholds that determine when to stop collecting data— one for accepting \mathcal{H}_0 and one for accepting \mathcal{H}_1 . When adapting this test to ensure DP, a naïve approach might compose two instances of the AboveThreshold mechanism (Kaplan et al., 2021), one for each decision boundary. Instead, we present a tighter mechanism that jointly handles both thresholds with a smaller privacy cost.

3 MONITORING MECHANISM

Now, we introduce a novel private mechanism that forms the foundation of our approach to differentially private sequential testing. Let $f_i : D \rightarrow \mathbb{R}$ be a sequence of queries with sensitivity Δ (D is either a finite dataset or an observation sequence), and $(T_0^i, T_1^i) \in \mathbb{R}^2$ with $T_0^i \leq T_1^i$. The **OutsideInterval** mechanism takes D as input and outputs a sequence \underline{a} of the form $(\perp, \dots, \perp, \top_0)$ or $(\perp, \dots, \perp, \top_1)$ where $a_i = \perp$ indicates that $f_i(D)$ should belong to $[T_0^i, T_1^i]$ and $a_i = \top_0$ (resp. $a_i = \top_1$) indicates that $f_i(D)$ should be smaller than T_0^i (resp. larger than T_1^i). The length of \underline{a} is the first query falling outside of the interval. This design is directly motivated by SPRT, which monitors whether a test statistic falls outside an interval defined by two thresholds, rather than crossing a single threshold.

OutsideInterval (Algorithm 1) proceeds by adding noise to both the query results and the thresholds, and returns either \top_0 , \top_1 , or \perp depending on whether the noisy query result is below the lower threshold, above the upper threshold, or within the interval. We emphasize that the algorithm adds the same noise realizations to both threshold comparisons, which allows for tighter privacy guarantees.

This mechanism extends AboveThreshold (Dwork et al., 2006) algorithm to simultaneously handle two thresholds. Unlike the approach of Zhang et al. (2022), which composes two separate instances of AboveThreshold for the upper and lower thresholds, **OutsideInterval** han-

²The sensitivity of a function $f : \mathcal{X}^n \rightarrow \mathbb{R}$ is the maximum change in the function’s value when a single observation is modified: $\Delta_f = \max_{\underline{X}, \underline{X}' \text{ neighboring}} |f(\underline{X}) - f(\underline{X}')|$.

Algorithm 1 **OutsideInterval**

Require: Private database D , adaptively chosen stream of queries f_1, \dots where each f_i has sensitivity Δ , sequences of thresholds $(T_0^i)_{i=1}^\infty, (T_1^i)_{i=1}^\infty$, noise-distributions $\mathcal{D}_Z, \mathcal{D}_Y$

- 1: $Z \leftarrow$ Sample from \mathcal{D}_Z
- 2: **for** Each query i **do**
- 3: $Y_i \leftarrow$ Sample from \mathcal{D}_Y
- 4: **if** $f_i(D) + Y_i \leq T_0^i - Z$ **then**
- 5: Halt and output $a_i = \top_0$
- 6: **else if** $f_i(D) + Y_i \geq T_1^i + Z$ **then**
- 7: Halt and output $a_i = \top_1$
- 8: **else**
- 9: Output $a_i = \perp$
- 10: **end if**
- 11: **end for**

dles both thresholds simultaneously, providing better privacy guarantees.

We express the privacy of **OutsideInterval** in terms of the DP and RDP guarantees of the noise-adding mechanisms associated with Y_i and Z . This allows us to provide a flexible analysis that can accommodate different noise distributions.

Definition 4 (Noise-adding Mechanism). *Given a continuous noise distribution \mathcal{D} , a noise-adding mechanism \mathcal{M} is a randomized algorithm that takes as input a dataset D and a query $f : D \rightarrow \mathbb{R}$, and outputs a noisy answer $\mathcal{M}(D, f) = f(D) + X$, where X is drawn from \mathcal{D} .*

Theorem 1. *Let τ be the (random) length of the sequence output by **OutsideInterval**.*

(i) *If the noise-adding mechanisms corresponding to \mathcal{D}_Z and \mathcal{D}_Y satisfy ϵ_Z -DP for queries with sensitivity Δ , and ϵ_Y -DP for queries with sensitivity 2Δ , then **OutsideInterval** is $(\epsilon_Z + \epsilon_Y)$ -DP.*

(ii) *If the noise-adding mechanism corresponding to \mathcal{D}_Z has an RDP profile $\epsilon_Z(\alpha)$ for queries with sensitivity Δ and the one for \mathcal{D}_Y has an RDP profile $\epsilon_Y(\alpha)$ for queries with sensitivity 2Δ , then **OutsideInterval** (denoted by \mathcal{A}) satisfies: for all $\alpha > 1$,*

$$D_\alpha(\mathcal{A}(D) \parallel \mathcal{A}(D')) \leq \frac{\alpha - 1/2}{\alpha - 1} \epsilon_Z(2\alpha) + \epsilon_Y(\alpha) + \frac{\log \left(2 \mathbb{E}_{z \sim \mathcal{D}_Z} [\mathbb{E}_{(\tau, \delta) \sim \mathcal{A}(D')} [\tau | Z = z]^{2\alpha}] \right)}{2(\alpha - 1)}.$$

This result shows that the privacy of our mechanism depends on both the threshold noise Z and the query noise Y . For ϵ -DP, we achieve $(\epsilon_Z + \epsilon_Y)$ -DP, which improves by a factor of 2 over using two separate instances of AboveThreshold. For RDP, the result has a similar flavor as in (Zhang et al., 2022), but our specific analysis also reduces each of the terms by a factor

of 2 compared to their. The full proof is provided in Appendix B.

This privacy analysis forms the basis for our differentially private sequential tests in the next section. Unlike prior work by Zhang et al. (2022), which only addressed Rényi DP for Gaussian noise, our approach supports both pure ϵ -DP and RDP guarantees. Beyond hypothesis testing, `OutsideInterval` can serve as a drop-in replacement for applications that use `AboveThreshold` where we want to ensure privacy under continual observations (Dwork et al., 2010), allowing monitoring of two thresholds simultaneously and reporting which one was crossed without additional privacy cost, contrary to the approach that would compose two separate instances of `AboveThreshold`.

4 DIFFERENTIALLY PRIVATE SPRT

Building on the threshold mechanism presented in Section 3, we now introduce a general framework for differentially private sequential probability ratio tests (`DP-SPRT`) that supports various noise distributions while maintaining error control and privacy guarantees.

4.1 General Definition

We consider the sequential hypothesis test between $\mathcal{H}_0 : (\theta = \theta_0)$ and $\mathcal{H}_1 : (\theta = \theta_1)$ (with $\theta_0 < \theta_1$) based on i.i.d. samples X_1, X_2, \dots from ν_θ , that belongs to a one-dimensional exponential family with natural parameter θ . We seek to design a sequential test (τ, \hat{d}) that (i) satisfies differential privacy, (ii) controls type I and type II error probabilities at specified levels α and β , and (iii) minimizes the expected sample size.

The idea of our method is to privatize the classical SPRT by adding carefully calibrated noise to both the thresholds and test statistic. `DP-SPRT` requires two noise distributions \mathcal{D}_Y and \mathcal{D}_Z , a correction function $C(n, \delta)$ and a parameter $\gamma \in (0, 1)$ allocating the error probability between the SPRT error and the privacy noise. It is defined as $\tau := \min(\tau_0, \tau_1)$ with

$$\tau_0 := \inf \left\{ n \in \mathbb{N}_{>0} : \bar{X}_n + \frac{Y_n}{n} \leq \tilde{\ell}_0^n - \frac{Z}{n} \right\},$$

$$\tau_1 := \inf \left\{ n \in \mathbb{N}_{>0} : \bar{X}_n + \frac{Y_n}{n} \geq \tilde{\ell}_1^n + \frac{Z}{n} \right\}$$

where the private thresholds are:

$$\tilde{\ell}_0^n = \mu_0 + \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \frac{\log(1/(\gamma\beta))}{n}}{\theta_1 - \theta_0} - C(n, (1 - \gamma)\beta),$$

$$\tilde{\ell}_1^n = \mu_1 - \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) - \frac{\log(1/(\gamma\alpha))}{n}}{\theta_1 - \theta_0} + C(n, (1 - \gamma)\alpha).$$

Algorithm 2 SPRT and `DP-SPRT`

Require: Parameters θ_0, θ_1 , error probabilities α, β
Require: Noise distributions $\mathcal{D}_Z, \mathcal{D}_Y$, error allocation $\gamma \in (0, 1)$, correction function C

- 1: Initialize $n = 0$
- 2: Sample threshold noise $Z \sim \mathcal{D}_Z$
- 3: **while** true **do**
- 4: $n \leftarrow n + 1$
- 5: Observe sample X_n
- 6: Compute $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$
- 7: Sample query noise $Y_n \sim \mathcal{D}_Y$
- 8: $\hat{T}_0^n \leftarrow \mu_0 + \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}) + \log(\gamma\beta)/n}{\theta_1 - \theta_0} - C(n, (1 - \gamma)\beta)$
- 9: $\hat{T}_1^n \leftarrow \mu_1 - \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) + \log(\gamma\alpha)/n}{\theta_1 - \theta_0} + C(n, (1 - \gamma)\alpha)$
- 10: **if** $\bar{X}_n + \frac{Y_n}{n} \leq \hat{T}_0^n - \frac{Z}{n}$ **then**
- 11: Halt and accept \mathcal{H}_0 ($\hat{d} = 0$)
- 12: **else if** $\bar{X}_n + \frac{Y_n}{n} \geq \hat{T}_1^n + \frac{Z}{n}$ **then**
- 13: Halt and accept \mathcal{H}_1 ($\hat{d} = 1$)
- 14: **end if**
- 15: **end while**

Here, Y_n *i.i.d.* \mathcal{D}_Y are the successive query noise and $Z \sim \mathcal{D}_Z$ is the threshold noise. The decision rule is defined as $\hat{d} = i$ if $\tau = \tau_i$. `DP-SPRT` is described formally in Algorithm 2, where the blue terms highlight the differences with the vanilla SPRT.

4.2 Generic Analysis

We now propose a generic analysis of `DP-SPRT`, which depends on properties of the noise distributions \mathcal{D}_Y and \mathcal{D}_Z as well as the correction function.

(i) **Privacy.** The privacy guarantees follow from Theorem 1 by noting that `DP-SPRT` can be seen as an instance of `OutsideInterval`. Specifically, the sensitivity of empirical means of Bernoulli observations is n^{-1} . By multiplying both sides of the stopping conditions by n , we can express `DP-SPRT` in terms of the `OutsideInterval` mechanism with queries $f_n(D) = \sum_{i=1}^n X_i$, whose sensitivity is $\Delta = 1$, noise Y_n , and appropriately defined thresholds.

Theorem 2 (Privacy). *Let \mathcal{A} be the `DP-SPRT` algorithm and τ the (random) stopping time.*

(i) **ϵ -DP:** *If the noise-adding mechanisms corresponding to distributions \mathcal{D}_Z and \mathcal{D}_Y satisfy ϵ_Z -DP for queries with sensitivity 1, and ϵ_Y -DP for queries with sensitivity 2, respectively, then `DP-SPRT` satisfies $(\epsilon_Z + \epsilon_Y)$ -Differential Privacy.*

(ii) **(α, ϵ) -RDP:** *If the noise-adding mechanism corresponding to distribution \mathcal{D}_Z and \mathcal{D}_Y have an RDP profile $\epsilon_Z(\alpha)$ for queries with sensitivity 1 and $\epsilon_Y(\alpha)$ for queries with sensitivity 2, then `DP-SPRT` satisfies:*

for all $\alpha > 1$,

$$\mathbb{D}_\alpha(\mathcal{A}(D)\|\mathcal{A}(D')) \leq \frac{\alpha - 1/2}{\alpha - 1} \varepsilon_Z(2\alpha) + \varepsilon_Y(\alpha) + \frac{\log\left(2\mathbb{E}_{z \sim \mathcal{D}_Z}[\mathbb{E}_{(\tau, \hat{d}) \sim \mathcal{A}(D')}[\tau|z]^2]\right)}{2(\alpha - 1)}. \quad (2)$$

(ii) Correctness. We establish correctness of **DP-SPRT** with the following theorem.

Theorem 3 (Correctness of **DP-SPRT**). *For any error allocation $\gamma \in (0, 1)$, **DP-SPRT** is (α, β) -correct if the noises and correction function satisfy:*

$$\forall \delta \in (0, 1), \sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) \leq \delta, \quad (3)$$

$$\forall \delta \in (0, 1), \sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} + \frac{Z}{n} < -C(n, \delta)\right) \leq \delta. \quad (4)$$

In particular, if the distribution \mathcal{D}_Y is symmetric around zero, we remark that condition (3) alone is sufficient, as it implies condition (4).

The main idea is to decompose the error probabilities using a union bound into the error from the standard SPRT (bounded via a change-of-measure argument) and the error from privacy noise (controlled through the correction function). The parameter γ allocates the total error budget between these two sources. The symmetry condition ensures that both tail bounds needed for Type I and Type II error control are satisfied simultaneously. See Appendix D for the full proof.

(iii) Sample Complexity. The following theorem provides a generic upper bound on the expected sample complexity for any noise distribution satisfying appropriate conditions.

Theorem 4 (Sample Complexity of **DP-SPRT**). *Let us assume conditions (3) and (4) hold. For any error allocation $\gamma \in (0, 1)$, the expected stopping time of **DP-SPRT** satisfies:*

$$\mathbb{E}_{\theta_0}[\tau] \leq 1 + (1 - \gamma)\beta + \frac{1}{1 - e^{-\frac{(\text{TV}(\nu_{\theta_0}, \nu_{\theta_1}))^4}{2(\theta_1 - \theta_0)^2}}} + N_0,$$

$$\mathbb{E}_{\theta_1}[\tau] \leq 1 + (1 - \gamma)\alpha + \frac{1}{1 - e^{-\frac{(\text{TV}(\nu_{\theta_0}, \nu_{\theta_1}))^4}{2(\theta_1 - \theta_0)^2}}} + N_1$$

where $N_0 = N(\theta_0, \theta_1, \beta, \gamma)$, $N_1 = N(\theta_1, \theta_0, \alpha, \gamma)$, and

$$N(\theta, \theta', \delta, \gamma) = \inf \left\{ n : \frac{\log(1/(\delta\gamma))}{n|\theta - \theta'|} + 2C(n, (1 - \gamma)\delta) \leq \frac{\text{KL}(\nu_\theta, \nu_{\theta'})}{2|\theta - \theta'|} \right\}.$$

This result shows how the required number of samples depends on the divergence between the distributions,

the desired error rates, and the noise added for privacy. Specific upper bounds on N can be derived for different noise mechanisms. Note that this bound leverages Hoeffding's inequality (Boucheron et al., 2013), which applies to bounded random variables like Bernoulli distributions. For other exponential families, one could replace it by a Chernoff-type inequality, resulting in more complex expressions than the form presented here, which is specific to our Bernoulli setting. See Appendix E for the proof and further analysis.

4.3 Variants of DP-SPRT

The generic **DP-SPRT** framework provides flexibility in selecting noise distributions based on privacy requirements. Table 1 summarizes two key implementations: Laplace noise for pure ε -DP and Gaussian noise for Rényi DP.

Table 1: Comparison of **DP-SPRT** instantiations

	Laplace	Gaussian
Query noise	$Y_n \sim \text{Lap}(4/\varepsilon)$	$Y_n \sim \mathcal{N}(0, \sigma_Y^2)$
Threshold noise	$Z \sim \text{Lap}(2/\varepsilon)$	$Z \sim \mathcal{N}(0, \sigma_Z^2)$
Correction	$\frac{6 \log(n^s \zeta(s)/\delta)}{n\varepsilon}$	$\frac{\sqrt{2\sigma^2 \log(n^s \zeta(s)/2\delta)}}{n}$
Privacy	ε -DP	$(\alpha, \varepsilon(\alpha))$ -RDP

Here, $\sigma^2 = \sigma_Y^2 + \sigma_Z^2$ and $\varepsilon(\alpha) = \frac{\alpha - 1/2}{\alpha - 1} \cdot \frac{\alpha}{\sigma_Z^2} + \frac{\alpha}{2\sigma_Y^2} + \frac{\log(2\mathbb{E}[\mathbb{E}[\tau|Z]^2])}{2(\alpha - 1)}$.

Both variants satisfy conditions (3) and (4), ensuring correctness with the specified error probabilities. The Laplace noise variant provides the strongest privacy guarantee (pure ε -DP), while the Gaussian implementation offers Rényi differential privacy that can be converted to approximate (ε, δ) -DP using standard techniques (Mironov, 2017). Detailed sample complexity analyses for these implementations are presented in Appendix E.

5 ON DP-SPRT WITH LAPLACE NOISE

While using **DP-SPRT** with Gaussian noise provides RDP guarantees, some applications require the stronger guarantee of pure ε -DP. By examining a lower bound on the sample complexity, we argue that our instance of **DP-SPRT** with Laplace noise is actually near-optimal in some regimes. We then propose further enhancements using sub-sampling.

Corollary 1 (Sample Complexity of **DP-SPRT** with Laplace Noise). *For **DP-SPRT** with Laplace noise, using the parameters from Table 1, the function N in Theorem*

4 satisfies:

$$N(\theta_0, \theta_1, \beta, \gamma) \leq \frac{24(\theta_1 - \theta_0) \log(\zeta(s)/(1 - \gamma)\beta)}{\varepsilon \text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} + \frac{2 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} + o_{\beta \rightarrow 0}(\log(1/\beta)).$$

The bound has two components: one reflecting the statistical difficulty of hypothesis testing (as in non-private settings) and another quantifying the cost of ensuring differential privacy, with γ controlling their trade-off. The full derivation appears in Appendix E.

5.1 Comparison with a Lower Bound

In Appendix F, we prove the following lower bound on the sample complexity of any (α, β) -correct test that is ε -DP. This lower bound is valid for any distributions ν_θ , and follows from techniques used to prove lower bounds under DP constraints in the multi-armed bandit literature (e.g., Azize et al. (2023); Azize (2024)).

Theorem 5 (Lower Bound for Private Sequential Tests). *Let (τ, \hat{d}) be an ε -differentially private test with $\mathbb{P}_{\theta_0}(\hat{d} = 1) \leq \alpha$ and $\mathbb{P}_{\theta_1}(\hat{d} = 0) \leq \beta$. Then:*

$$\mathbb{E}_{\theta_0}[\tau] \geq \frac{\text{kl}(\alpha, 1 - \beta)}{\min(\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}), \varepsilon \cdot \text{TV}(\nu_{\theta_0}, \nu_{\theta_1}))} \text{ and}$$

$$\mathbb{E}_{\theta_1}[\tau] \geq \frac{\text{kl}(\beta, 1 - \alpha)}{\min(\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}), \varepsilon \cdot \text{TV}(\nu_{\theta_0}, \nu_{\theta_1}))}$$

where TV is the total variation distance between distributions and $\text{kl}(x, y) := \text{KL}(\mathcal{B}(x), \mathcal{B}(y)) = x \log(x/y) + (1 - x) \log((1 - x)/(1 - y))$ is the binary relative entropy.

We now compare this lower bound to our upper bound for Laplace DP-SPRT, in the asymptotic regime in which α and β go to zero, for which we have $\text{kl}(\alpha, 1 - \beta) \simeq \log(1/\beta)$ and $\text{kl}(\beta, 1 - \alpha) \simeq \log(1/\alpha)$. Combining Theorem 4 with Corollary 1, we have

$$\limsup_{\beta \rightarrow 0} \frac{\mathbb{E}_{\theta_0}[\tau]}{\log(1/\beta)} \leq \frac{48}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} \max\left(1, \frac{\theta_1 - \theta_0}{\varepsilon}\right),$$

where ν_{θ_i} is the Bernoulli distribution with natural parameter θ_i and mean $p_i = \frac{e^{\theta_i}}{1 + e^{\theta_i}}$. Using the simple expression of the TV for Bernoulli distributions in Theorem 5 and $p_1 > p_0$,

$$\liminf_{\beta \rightarrow 0} \frac{\mathbb{E}_{\theta_0}[\tau]}{\log(1/\beta)} \geq \max\left(\frac{1}{\text{kl}(p_0, p_1)}, \frac{1}{(p_1 - p_0)\varepsilon}\right).$$

By definition, we have $\text{kl}(p_0, p_1) = \text{KL}(\nu_{\theta_0}, \nu_{\theta_1})$. As for the second term, using the expression of KL divergence for exponential families (see Eq. (5) in Appendix A) we have $\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})/(\theta_1 - \theta_0) = (b(\theta_1) - b(\theta_0))/(\theta_1 - \theta_0) - p_0$ where b is the log-partition function, which goes to $b'(\theta_1) - p_0 = p_1 - p_0$ when p_1 goes to p_0 . Hence,

when p_1 and p_0 are close, the upper and lower bounds are matching, up to a multiplicative constant. By symmetry, we can get similar nearly matching upper and lower bounds on $\mathbb{E}_{\theta_1}[\tau]/\log(1/\alpha)$. This makes DP-SPRT with Laplace noise near-optimal in the regime of small errors and close hypotheses.

5.2 Privacy Amplification by Subsampling

DP-SPRT with Laplace noise can be further enhanced using subsampling (Steinke, 2022). That is, in each round the empirical mean is computed based on a random subset of the available observations instead of all observations. This approach allows adding less noise for the same privacy guarantee, and in turn, a smaller correction factor. This is particularly effective in high privacy regimes as demonstrated in the experiments. Details of the implementation of this approach are provided in Appendix G.

6 EXPERIMENTAL ANALYSIS

6.1 Experimental Setup

We conduct experiments on simulated Bernoulli data to evaluate the practical performance of our DP-SPRT methods. We implement the classical SPRT, PrivSPRT (Zhang et al., 2022), and three variants of our approach: DP-SPRT with Laplace noise, with Gaussian noise, and with Laplace noise and subsampling (Laplace-Sub). For all methods, we set target error rates $\alpha = \beta = 0.05$ and perform 1000 independent trials to estimate the error probabilities and the sample complexity. We present results for Bernoulli distributions with parameter $p_0 = 0.3$ and $p_1 = 0.7$.

We use parameters of Table 1 for DP-SPRT with Laplace noise to ensure ε -DP. For DP-SPRT with Gaussian noise, following Zhang et al. (2022), we set the variance parameters as $\sigma_Y^2 = 32 \log(1.25/\delta)/\varepsilon^2$ and $\sigma_Z^2 = 8 \log(1.25/\delta)/\varepsilon^2$, ensuring $(\varepsilon/2, \delta)$ -DP for queries of sensitivity 2 and 1 respectively. To compare with an instance of PrivSPRT having a similar RDP guarantee, we set $\sigma_1 = 2\sqrt{2}A\sigma_Z$ and $\sigma_2 = 2\sqrt{2}A\sigma_Y$ where $2A$ is the sensitivity due to their truncation parameter. A is set to 1. The rescaling accounts for the fact that their RDP bound has an additional factor 2 in the first two terms of Eq. (2) compared to ours.

The error allocation parameter of DP-SPRT is set to $\gamma(\varepsilon) = \min(1/2, 1 - 1/\varepsilon)$, allowing to recover the classical SPRT as ε increases. For the subsampling version, we adaptively set the subsampling rate to $\min(1, \sqrt{\varepsilon/10})$, depending on the privacy parameter. More discussion on alternative choices of subsampling rate is in Appendix G. We calibrate the thresholds of

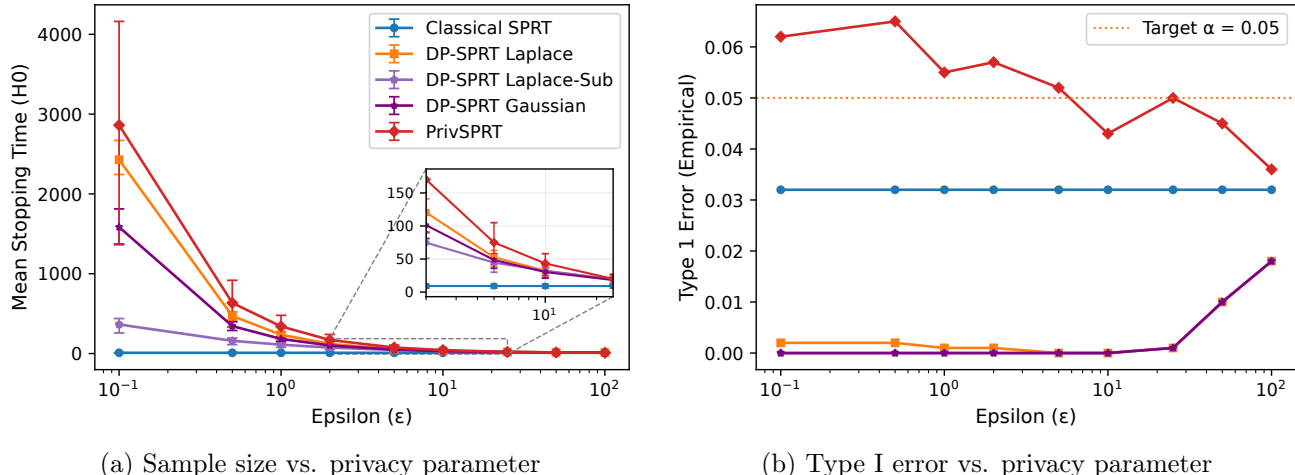


Figure 1: Experimental results for different DP-SPRT variants: (a) Average sample size as privacy parameter varies for $\alpha = \beta = 0.05$; (b) Empirical Type I error probability vs. privacy parameter. Error bars represent 95% percentile intervals over 1000 trials.

PrivSPRT through grid search with Monte Carlo simulation (100 runs), while our methods use the theoretical calibration derived in Section 4.

6.2 Results

(1) **Smaller Stopping Time across ϵ .** Figure 1(a) shows that as privacy requirements become stringent (smaller ϵ), all private methods require more samples. Our DP-SPRT variants consistently outperform or match PrivSPRT at equivalent privacy levels, with subsampling providing substantial benefits in high-privacy regimes. This contradicts prior beliefs that Laplace noise is unsuitable for sequential testing (Zhang et al., 2022).

(2) **Smaller Probability of Error across ϵ .** Figure 1(b) demonstrates that our DP-SPRT variants maintain error probabilities (significantly) below target across all privacy levels. In contrast, PrivSPRT occasionally exceeds the target rate, likely due to the noise introduced during its Monte Carlo calibration process. The fact that the empirical type I error of DP-SPRT is much below the target α suggests that our proposed correction function $C(n, \delta)$ is quite conservative. In Appendix H, we provide a more detailed experimental analysis with different p_0 and p_1 , the variance of the stopping time, its dependence on the privacy parameters and the target error probabilities, and a tuning of the correction function.

In brief, *our DP-SPRT methods achieve better efficiency and accuracy than PrivSPRT while offering theoretically-calibrated thresholds guaranteeing error control without empirical tuning.*

7 DISCUSSION AND FUTURE WORK

We introduced DP-SPRT, an adaptation of the classical SPRT to incorporate various types of privacy protections in sequential tests of simple hypotheses. This general framework supports different noise mechanisms, such as Laplace noise for ϵ -differential privacy and Gaussian noise for Rényi differential privacy. Its design is based on a specialized threshold mechanism, `OutsideInterval`, that concurrently handles both decision boundaries instead of treating them separately. Notably, DP-SPRT with Laplace noise is the first ϵ -DP sequential test that can be calibrated to attain prescribed type I and type II errors without any hyperparameter tuning. Moreover, for Bernoullis, its sample complexity is close to the lower bound that we have derived for the minimal sample complexity of any ϵ -DP test. In future work, we will seek to strengthen these optimality guarantees beyond Bernoulli, and extend our findings to composite hypothesis testing, for which `OutsideInterval` could also be useful.

Acknowledgements

This work has been partially supported by the French National Research Agency (ANR) in the framework of the PEPR IA FOUNDRY project (ANR-23-PEIA-0003) and the ANR BIP-UP project (ANR-22-CE23-003). We also acknowledge the Inria-ISI, Kolkata Associate Team “SeRAI”, and the ANR JCJC for the REPUBLIC project (ANR-22-CE23-0003-01) for partially supporting the project.

References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318.
- Armitage, P. (1950). Sequential analysis with more than two alternative hypotheses, and its relation to discriminant function analysis. *Journal of the Royal Statistical Society. Series B (Methodological)*, 12(1):137–144.
- Armitage, P. (1954). Sequential tests in prophylactic and therapeutic trials. *The Quarterly journal of medicine*, 23(91):255–274.
- Azize, A. (2024). *Privacy-Utility Trade-offs in Sequential Decision-Making under Uncertainty*. PhD thesis, Université de Lille.
- Azize, A., Jourdan, M., Al Marjani, A., and Basu, D. (2023). On the complexity of differentially private best-arm identification with fixed confidence. *Advances in Neural Information Processing Systems*, 36:71150–71194.
- Azize, A., Jourdan, M., Marjani, A. A., and Basu, D. (2024). Differentially private best-arm identification. *arXiv preprint arXiv:2406.06408*.
- Boucheron, S., Lugosi, G., and Massart, P. (2013). *Concentration inequalities. A non asymptotic theory of independence*. Oxford University Press.
- Brown, L. D. (1986). Fundamentals of statistical exponential families with applications in statistical decision theory. *Institute of Mathematical Statistics, Lecture Notes-Monograph Series*.
- Canonne, C. L., Kamath, G., McMillan, A., Smith, A., and Ullman, J. (2019). The structure of optimal private tests for simple hypotheses. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 310–321.
- Cappé, O., Garivier, A., Maillard, O.-A., Munos, R., and Stoltz, G. (2013). Kullback-Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541.
- Couch, S., Kazan, Z., Shi, K., Bray, A., and Groce, A. (2019). Differentially private nonparametric hypothesis testing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 737–751.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer.
- Dwork, C., Naor, M., Pitassi, T., and Rothblum, G. N. (2010). Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724.
- Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.
- Gaboardi, M., Lim, H., Rogers, R., and Vadhan, S. (2016). Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. In *International conference on machine learning*, pages 2111–2120. PMLR.
- Garivier, A. and Kaufmann, E. (2016). Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory*.
- Garivier, A., Ménard, P., and Stoltz, G. (2019). Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operation Research*, 44(2):377–399.
- Kalogerias, D. S., Nikolakakis, K. E., Sarwate, A. D., and Sheffet, O. (2021). Quantile multi-armed bandits: Optimal best-arm identification and a differentially private scheme. *IEEE Journal on Selected Areas in Information Theory*, 2(2):534–548.
- Kamath, G., Mouzakis, A., Singhal, V., Steinke, T., and Ullman, J. (2022). A private and computationally-efficient estimator for unbounded gaussians. In *Conference on Learning Theory*, pages 544–572. PMLR.
- Kaplan, H., Mansour, Y., and Stemmer, U. (2021). The sparse vector technique, revisited. In *Conference on Learning Theory*, pages 2747–2776. PMLR.
- Kazan, Z., Shi, K., Groce, A., and Bray, A. P. (2023). The test of tests: A framework for differentially private hypothesis testing. In *International Conference on Machine Learning*, pages 16131–16151. PMLR.
- Mironov, I. (2017). Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE.
- Papernot, N. and Steinke, T. (2021). Hyperparameter tuning with renyi differential privacy. *arXiv preprint arXiv:2110.03620*.
- Sajed, T. and Sheffet, O. (2019). An optimal private stochastic-mab algorithm based on optimal private stopping rule. In *International Conference on Machine Learning*, pages 5579–5588. PMLR.
- Siegmund, D. (2013). *Sequential analysis: tests and confidence intervals*. Springer Science & Business Media.

Steinke, T. (2022). Composition of differential privacy & privacy amplification by subsampling. *arXiv preprint arXiv:2210.00597*.

Tirinzi, A., Marjani, A. A., and Kaufmann, E. (2022). Near instance-optimal PAC reinforcement learning for deterministic mdps. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Wald, A. (1945). Sequential Tests of Statistical Hypotheses. *Annals of Mathematical Statistics*, 16(2):117–186.

Wald, A. (1948). Optimum character of the Sequential Probability Ratio Test. *Annals of Mathematical Statistics*, 19(3):326–339.

Wald, A. (2004). *Sequential analysis*. Courier Corporation.

Wang, Y., Sibai, H., Yen, M., Mitra, S., and Dullerud, G. E. (2022). Differentially private algorithms for statistical verification of cyber-physical systems. *IEEE Open Journal of Control Systems*, 1:294–305.

Zhang, W., Mei, Y., and Cummings, R. (2022). Private sequential hypothesis testing for statisticians: Privacy, error rates, and sample size. In *AISTATS*.

Zhu, Y. and Wang, Y.-X. (2020). Improving sparse vector technique with renyi differential privacy. *Advances in neural information processing systems*, 33:20249–20258.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
Mathematical setting and assumptions are detailed in Section 2 and 3. The DP-SPRT algorithm is formally presented in Algorithm 2, with the **OutsideInterval** mechanism in Algorithm 1.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
Privacy properties are analyzed in Theorem 2 and Theorem 1. Sample complexity bounds are provided in Theorem 4 and Corollary 1.
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
Anonymized source code is provided in the supplementary materials.
2. For any theoretical claim, check if you include:

- (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
Complete proofs are provided in the appendix.
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
Code is provided in the supplementary material.
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
No training required.
 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
 5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Appendix

Table of Contents

A	A Primer on SPRT for Exponential Families	12
B	Privacy of <code>OutsideInterval</code>: Proof of Theorem 1	13
	B.1 ϵ -DP Guarantee	15
	B.2 (α, ϵ) -Rényi DP Guarantee	17
C	The Variants of DP-SPRT: Laplace and Gaussian Noises	20
	C.1 DP-SPRT with Laplace Noise	20
	C.2 DP-SPRT with Gaussian Noise	21
D	Correctness of DP-SPRT for Exponential Families: Proof of Theorem 3	22
E	Sample Complexity of DP-SPRT for Bernoulli: Proof of Theorem 4	24
	E.1 Analysis for Laplace Noise	25
	E.2 Analysis for Gaussian Noise	26
F	Lower Bound: Proof of Theorem 5	27
G	Enhancing DP-SPRT with Subsampling	29
H	Complementary Experimental Analysis	30
	H.1 Comparative Performance Across Problem Difficulty	31
	H.2 Effect of Target Error Probability	31
	H.3 Privacy-Sample Complexity Trade-off	32
	H.4 Tuned Correction Function and Distribution of Stopping Times	32
	H.5 Summary of Findings	32

A A Primer on SPRT for Exponential Families

A one-dimensional canonical exponential family is a class of distributions

$$\mathcal{F} = \{(\nu_\theta)_{\theta \in \Theta} : \nu_\theta \text{ has a density } f_\theta(x) = \exp(\theta x - b(\theta)) \text{ w.r.t. } \nu\},$$

where ν is the reference measure, Θ is the natural parameter space and $b : \Theta \rightarrow \mathbb{R}$ is a convex, twice differentiable function called the log-partition function.

For exponential families, it can be shown that $\mu(\theta) = \mathbb{E}_{X \sim \nu_\theta}[X]$ satisfies $\mu(\theta) = b'(\theta)$, and $\sigma^2(\theta) = \text{Var}_{X \sim \nu_\theta}[X]$ satisfies $\sigma^2(\theta) = b''(\theta)$. As a consequence, the mean is an increasing function of the natural parameter, and there is a one-to-one mapping between mean and natural parameter.

We recall the expression of the KL divergence in an exponential family (see, e.g., [Cappé et al. \(2013\)](#)). For all $\theta, \theta' \in \Theta$, we have

$$\text{KL}(\nu_\theta, \nu_{\theta'}) = \mathbb{E}_{X \sim \nu_\theta} \left[\log \frac{f_\theta(X)}{f_{\theta'}(X)} \right] = \mu(\theta)(\theta - \theta') - b(\theta) + b(\theta'). \quad (5)$$

Specific form of the SPRT. The general SPRT

$$\tau = \inf \left\{ n \in \mathbb{N} : \prod_{i=1}^n \frac{f_{\theta_1}(X_i)}{f_{\theta_0}(X_i)} \notin (\beta, 1/\alpha) \right\}$$

particularizes to

$$\tau = \inf \left\{ n \in \mathbb{N} : \exp \left((\theta_1 - \theta_0) \sum_{i=1}^n X_i - n(b(\theta_1) - b(\theta_0)) \right) \notin (\beta, 1/\alpha) \right\},$$

which can be re-written $\tau = \min(\tau_0, \tau_1)$ with

$$\begin{aligned} \tau_0 &= \inf \left\{ n \in \mathbb{N} : (\theta_1 - \theta_0) \sum_{i=1}^n X_i - n(b(\theta_1) - b(\theta_0)) < \log(\beta) \right\}, \\ \tau_1 &= \inf \left\{ n \in \mathbb{N} : (\theta_1 - \theta_0) \sum_{i=1}^n X_i - n(b(\theta_1) - b(\theta_0)) > \log(1/\alpha) \right\}. \end{aligned}$$

Dividing by $\theta_1 - \theta_0 > 0$ and re-arranging the terms, we get

$$\begin{aligned} \tau_0 &= \inf \left\{ n \in \mathbb{N} : \frac{\sum_{i=1}^n X_i}{n} < \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} - \frac{\frac{1}{n} \log(1/\beta)}{\theta_1 - \theta_0} \right\}, \\ \tau_1 &= \inf \left\{ n \in \mathbb{N} : \frac{\sum_{i=1}^n X_i}{n} > \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} + \frac{\frac{1}{n} \log(1/\alpha)}{\theta_1 - \theta_0} \right\}. \end{aligned}$$

The expression given in Section 2 can be obtained by noting that $b(\theta_1) - b(\theta_0) = \mu_0(\theta_1 - \theta_0) + \text{KL}(\nu_{\theta_0}, \nu_{\theta_1})$ and $b(\theta_1) - b(\theta_0) = \mu_1(\theta_1 - \theta_0) - \text{KL}(\nu_{\theta_1}, \nu_{\theta_0})$.

However, the above expression is also interesting, as it shows that the two thresholds

$$T_0^n = \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} - \frac{\frac{1}{n} \log(1/\beta)}{\theta_1 - \theta_0} \text{ and } T_1^n = \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} + \frac{\frac{1}{n} \log(1/\alpha)}{\theta_1 - \theta_0}$$

to which the empirical mean is compared indeed satisfies $T_0^n \leq T_1^n$.

DP-SPRT can thus be alternatively written

$$\begin{aligned} \tau_0 &= \inf \left\{ n \in \mathbb{N} : \frac{\sum_{i=1}^n X_i}{n} + \frac{Y_n}{n} < \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} - \frac{\frac{1}{n} \log(1/\gamma\beta)}{\theta_1 - \theta_0} - C(n, (1-\gamma)\beta) - \frac{Z}{n} \right\}, \\ \tau_1 &= \inf \left\{ n \in \mathbb{N} : \frac{\sum_{i=1}^n X_i}{n} + \frac{Y_n}{n} > \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} + \frac{\frac{1}{n} \log(1/\gamma\alpha)}{\theta_1 - \theta_0} + C(n, (1-\gamma)\alpha) + \frac{Z}{n} \right\}. \end{aligned}$$

Bernoulli distributions. Bernoulli distribution form a one-dimensional exponential family if we set the natural parameter of the Bernoulli distribution with probability of success p_i to be $\theta_i = \log \frac{p_i}{1-p_i}$. The log-partition function is $b(\theta) = \log(1 + e^\theta)$.

On DP-SPRT versus PrivSPRT. PrivSPRT (Zhang et al., 2022) does not rely on the specific form of SPRT for exponential family based on its sufficient statistic (the empirical mean). It directly privatizes the sum of (truncated) log-likelihood ratios. For Bernoulli distribution with means p_0 and p_1 , the stopping rule of PrivSPRT is $\tilde{\tau} = \min(\tilde{\tau}_0, \tilde{\tau}_1)$ with

$$\begin{aligned}\tilde{\tau}_0 &= \inf \left\{ n \in \mathbb{N} : \sum_{i=1}^n \left[X_i \log \frac{p_1}{p_0} + (1 - X_i) \log \frac{1-p_1}{1-p_0} \right]_{-A}^A + Y_n^1 > b + Z^1 \right\}, \\ \tilde{\tau}_1 &= \inf \left\{ n \in \mathbb{N} : \sum_{i=1}^n \left[X_i \log \frac{p_1}{p_0} + (1 - X_i) \log \frac{1-p_1}{1-p_0} \right]_{-A}^A + Y_n^2 < -a + Z^2 \right\},\end{aligned}$$

where Y_n^1, Y_n^2 are iid from $\mathcal{N}(0, \sigma_2^2)$, and Z^1, Z^2 are both drawn from $\mathcal{N}(0, \sigma_1^2)$. $[x]_{-A}^A$ stands for the truncation to $[-A, A]$. On some instances on which the truncation is immaterial, that is when

$$\max \left(\log \frac{p_1}{p_0}, \log \frac{1-p_0}{1-p_1} \right) \leq A$$

(which is the case for $p_0 = 0.3, p_1 = 0.7$ and $A = 1$), PrivSPRT can also be re-expressed with the empirical mean. We obtain

$$\begin{aligned}\tilde{\tau}_0 &= \inf \left\{ n \in \mathbb{N} : \frac{\sum_{i=1}^n X_i}{n} + \frac{Y_n^1}{n(\theta_1 - \theta_0)} > \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} + \frac{b}{n(\theta_1 - \theta_0)} + \frac{Z^1}{n(\theta_1 - \theta_0)} \right\}, \\ \tilde{\tau}_1 &= \inf \left\{ n \in \mathbb{N} : \frac{\sum_{i=1}^n X_i}{n} + \frac{Y_n^2}{n(\theta_1 - \theta_0)} < \frac{b(\theta_1) - b(\theta_0)}{\theta_1 - \theta_0} - \frac{a}{n(\theta_1 - \theta_0)} + \frac{Z^2}{n(\theta_1 - \theta_0)} \right\}.\end{aligned}$$

The differences with DP-SPRT are the following: (1) PrivSPRT uses a and b as parameters to be tuned (empirically) to attain desired error levels while DP-SPRT uses an exact, theoretically valid calibration. (2) PrivSPRT uses two streams of noise, whereas the noise are common for τ_0 and τ_1 in DP-SPRT. (3) The noise is not scaled in the same way in both approaches.

B Privacy of `OutsideInterval`: Proof of Theorem 1

Two useful lemmas. We will use the following two lemmas in the proof of Theorem 1. The first lemma is a direct consequence of the definition of Rényi differential privacy.

Lemma 1 (Indistinguishability Property of RDP (from Mironov (2017))). *Let \mathcal{M} be a mechanism that satisfies $(\alpha, \tilde{\varepsilon}(\alpha))$ -RDP. Then for any measurable set $S \subseteq \text{Range}(\mathcal{M})$ and any neighboring datasets D, D' :*

$$e^{-\tilde{\varepsilon}(\alpha) \frac{\alpha}{\alpha-1}} \cdot \Pr[\mathcal{M}(D') \in S]^{\frac{\alpha}{\alpha-1}} \leq \Pr[\mathcal{M}(D) \in S] \leq e^{\tilde{\varepsilon}(\alpha) \frac{\alpha}{\alpha-1}} \cdot \Pr[\mathcal{M}(D') \in S]^{\frac{\alpha-1}{\alpha}}. \quad (6)$$

This second lemma is a direct consequence of the definition of differential privacy for noise-adding mechanisms with continuous noise.

Lemma 2. *Let $M(D) = f(D) + Z$ be a mechanism that adds noise $Z \in \mathbb{R}$ with probability density function p_Z to the output of a query f . Let Δ be the sensitivity of f , i.e., $\Delta = \max_{D \sim D'} |f(D) - f(D')|$, where $D \sim D'$ denotes that D and D' are adjacent datasets. If M is ε -differentially private, then*

$$\forall z \in \mathbb{R}, p_Z(z - \Delta) \leq e^\varepsilon p_Z(z).$$

Proof. By definition, M is ε -differentially private if for all adjacent datasets $D \sim D'$ and all measurable sets $S \subseteq \mathbb{R}$:

$$\Pr[M(D) \in S] \leq e^\varepsilon \Pr[M(D') \in S].$$

This must hold for all measurable sets S , which implies that it must hold pointwise for the probability density functions:

$$\forall x \in \mathbb{R}, p_{M(D)}(x) \leq e^\varepsilon p_{M(D')}(x).$$

Since $p_{M(D)}(x) = p_Z(x - f(D))$ and $p_{M(D')}(x) = p_Z(x - f(D'))$, we have, for all query f of sensitivity Δ and for all neighboring D and D' ,

$$\forall x \in \mathbb{R}, p_Z(x - f(D)) \leq e^\varepsilon p_Z(x - f(D')).$$

Consider the worst case where $f(D) = f(D') + \Delta$. Substituting we get

$$\forall x \in \mathbb{R}, p_Z(x - f(D') - \Delta) \leq e^\varepsilon p_Z(x - f(D')).$$

Letting $z = x - f(D')$ yields

$$\forall z \in \mathbb{R}, p_Z(z - \Delta) \leq e^\varepsilon p_Z(z).$$

□

We prove below a more general version of Theorem 1 for the **OutsideInterval** algorithm. Theorem 1 can be obtained by setting $\gamma = 2$ in the following theorem.

Theorem 6 (General Form of Theorem 1). *Let τ be a random variable indicating the stopping time of the **OutsideInterval** mechanism.*

(i) *If the noise-adding mechanism associated with \mathcal{D}_Z satisfies ε_Z -DP, and the mechanism associated with \mathcal{D}_Y satisfies ε_Y -DP for queries with sensitivity 2Δ , then **OutsideInterval** is $(\varepsilon_Z + \varepsilon_Y)$ -DP.*

(ii) *If the noise-adding mechanism associated with \mathcal{D}_Z has an RDP profile $\varepsilon_Z(\alpha)$ for queries with sensitivity Δ and the one associated with \mathcal{D}_Y has an RDP profile $\varepsilon_Y(\alpha)$ for queries with sensitivity 2Δ , then **OutsideInterval** (denoted by \mathcal{A}) satisfies:*

$$\mathbb{D}_\alpha(\mathcal{A}(D) \parallel \mathcal{A}(D')) \leq \frac{\alpha - (\gamma - 1)/\gamma}{\alpha - 1} \varepsilon_Z \left(\frac{\gamma}{\gamma - 1} \alpha \right) + \varepsilon_Y(\alpha) + \frac{\log(\mathbb{E}_{z \sim \mathcal{D}_Z}[\mathbb{E}[\tau|z]^\gamma])}{\gamma(\alpha - 1)},$$

for any $\gamma > 1$.

The proof is adapted from the one by [Zhu and Wang \(2020\)](#) for GeneralAboveThreshold.

Consider an arbitrary output of the algorithm $\mathbf{a} = (\perp, \dots, \perp, \top_d)$ with \perp appearing $k - 1$ times and $d \in \{0, 1\}$ indicating whether the k -th query falls below the lower threshold ($d = 0$) or exceeds the upper threshold ($d = 1$). Let

$$g_i(D, z) = \Pr[f_i(D) + Y_i > T_0^i - z \wedge f_i(D) + Y_i < T_1^i + z],$$

$$h_k^{(d)}(D, z) = \begin{cases} \Pr[f_k(D) + Y_k \geq T_1^k + z] & \text{if } d = 1 \\ \Pr[f_k(D) + Y_k \leq T_0^k - z] & \text{if } d = 0 \end{cases},$$

where the randomness in \Pr is on the distribution of the Y s only. Letting p_Z be the pdf of Z , we have

$$\Pr[\mathcal{A}(D) = \mathbf{a}] = \int_{-\infty}^{\infty} p_Z(z) \prod_{i=1}^{k-1} g_i(D, z) \cdot h_k^{(d)}(D, z) dz.$$

Step 1: Bounding output probabilities.

Using a change of variable $z \rightarrow z - \Delta$:

$$\begin{aligned} \Pr[\mathcal{A}(D) = \mathbf{a}] &= \int_{-\infty}^{\infty} p_Z(z - \Delta) \prod_{i=1}^{k-1} g_i(D, z - \Delta) \cdot h_k^{(d)}(D, z - \Delta) dz \\ &\leq \int_{-\infty}^{\infty} \frac{p_Z(z - \Delta)}{p_Z(z)} p_Z(z) \prod_{i=1}^{k-1} g_i(D, z - \Delta) \cdot h_k^{(d)}(D, z - \Delta) dz. \end{aligned}$$

Note that $p_Z(z - \Delta)/p_Z(z)$ is well-defined for DP or RDP mechanism, as from Lemma 1 or Lemma 2 it can be shown that the density has to be supported on \mathbb{R} .

For queries before k , we establish an upper bound on $g_i(D, z - \Delta)$ through inequality analysis. For any $i < k$, we have:

$$g_i(D, z - \Delta) = \Pr[f_i(D) + Y_i > T_0^i - (z - \Delta) \wedge f_i(D) + Y_i < T_1^i + (z - \Delta)].$$

Since f_i has sensitivity Δ , we know that $|f_i(D) - f_i(D')| \leq \Delta$. This gives us two inequalities:

$$\begin{aligned} f_i(D) &\leq f_i(D') + \Delta, \\ f_i(D) &\geq f_i(D') - \Delta. \end{aligned}$$

For the first inequality condition, we have:

$$\begin{aligned} \Pr[f_i(D) + Y_i > T_0^i - (z - \Delta)] &\leq \Pr[f_i(D') + \Delta + Y_i > T_0^i - (z - \Delta)] \\ &= \Pr[f_i(D') + Y_i > T_0^i - z]. \end{aligned}$$

For the second inequality condition:

$$\begin{aligned} \Pr[f_i(D) + Y_i < T_1^i + (z - \Delta)] &\leq \Pr[f_i(D') - \Delta + Y_i < T_1^i + (z - \Delta)] \\ &= \Pr[f_i(D') + Y_i < T_1^i + z]. \end{aligned}$$

Since both conditions in the conjunction are upper-bounded, we obtain:

$$\begin{aligned} g_i(D, z - \Delta) &\leq \Pr[f_i(D') + Y_i > T_0^i - z \wedge f_i(D') + Y_i < T_1^i + z] \\ &= g_i(D', z). \end{aligned}$$

Therefore:

$$\begin{aligned} \Pr[\mathcal{A}(D) = \mathbf{a}] &\leq \int_{-\infty}^{\infty} \frac{p_Z(z - \Delta)}{p_Z(z)} p_Z(z) \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D, z - \Delta) dz \\ &= \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\frac{p_Z(z - \Delta)}{p_Z(z)} \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D, z - \Delta) \right]. \end{aligned} \quad (7)$$

With our notation, we can also write

$$\Pr[\mathcal{A}(D') = \mathbf{a}] = \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D', z) \right]. \quad (8)$$

We will now use the properties of the private mechanism to relate the two probabilities. First we will prove the results when the mechanism is ε -DP.

B.1 ε -DP Guarantee

Proof of Theorem 1 (i). We have established in Equation (7) that for any output $\mathbf{a} = (\perp, \dots, \perp, \top_d)$ with \perp appearing $k - 1$ times:

$$\Pr[\mathcal{A}(D) = \mathbf{a}] \leq \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\frac{p_Z(z - \Delta)}{p_Z(z)} \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D, z - \Delta) \right].$$

Step 2: Bounding the density ratio and stopping time query.

Since the noise-adding mechanism corresponding to \mathcal{D}_Z satisfies ε_Z -DP for queries with sensitivity Δ , by Lemma 2, we have:

$$\frac{p_Z(z - \Delta)}{p_Z(z)} \leq e^{\varepsilon_Z}.$$

For $d \in \{0, 1\}$, we need to establish a bound on $h_k^{(d)}(D, z - \Delta)$ in terms of $h_k^{(d)}(D', z)$. We first consider the case $d = 1$ (upper threshold crossed).

Let us define a new query:

$$\tilde{f}_k(D'') = \begin{cases} f_k(D) + \Delta & \text{if } D'' = D \\ f_k(D') & \text{if } D'' = D' \end{cases}.$$

To establish the sensitivity of \tilde{f}_k , we note that since f_k has sensitivity Δ , we have $|f_k(D) - f_k(D')| \leq \Delta$. Thus:

$$\begin{aligned} |\tilde{f}_k(D) - \tilde{f}_k(D')| &= |f_k(D) + \Delta - f_k(D')| \\ &\leq |f_k(D) - f_k(D')| + \Delta \\ &\leq \Delta + \Delta = 2\Delta. \end{aligned}$$

Therefore, \tilde{f}_k has sensitivity 2Δ . Now, we can rewrite:

$$\begin{aligned} h_k^{(1)}(D, z - \Delta) &= \Pr[f_k(D) + Y_k \geq T_1^k + (z - \Delta)] \\ &= \Pr[\tilde{f}_k(D) - \Delta + Y_k \geq T_1^k + (z - \Delta)] \\ &= \Pr[\tilde{f}_k(D) + Y_k \geq T_1^k + z]. \end{aligned}$$

Similarly, $h_k^{(1)}(D', z) = \Pr[\tilde{f}_k(D') + Y_k \geq T_1^k + z]$.

Since the noise-adding mechanism corresponding to \mathcal{D}_Y satisfies ε_Y -DP for queries with sensitivity 2Δ , we can apply the privacy guarantee:

$$\Pr[\tilde{f}_k(D) + Y_k \geq T_1^k + z] \leq e^{\varepsilon_Y} \Pr[\tilde{f}_k(D') + Y_k \geq T_1^k + z].$$

Therefore: $h_k^{(1)}(D, z - \Delta) \leq e^{\varepsilon_Y} h_k^{(1)}(D', z)$.

A similar analysis for $d = 0$ (lower threshold crossed) would define $\tilde{f}_k(D) = f_k(D) - \Delta$ and $\tilde{f}_k(D') = f_k(D')$, yielding the same bound: $h_k^{(0)}(D, z - \Delta) \leq e^{\varepsilon_Y} h_k^{(0)}(D', z)$.

Step 3: Deriving the final DP bound.

Using these inequalities in our original expression:

$$\begin{aligned} \Pr[\mathcal{A}(D) = \mathbf{a}] &\leq \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\frac{p_Z(z - \Delta)}{p_Z(z)} \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D, z - \Delta) \right] \\ &\leq \mathbb{E}_{z \sim \mathcal{D}_Z} \left[e^{\varepsilon_Z} \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D, z - \Delta) \right] \\ &\leq \mathbb{E}_{z \sim \mathcal{D}_Z} \left[e^{\varepsilon_Z} \prod_{i=1}^{k-1} g_i(D', z) \cdot e^{\varepsilon_Y} h_k^{(d)}(D', z) \right] \\ &= e^{\varepsilon_Z + \varepsilon_Y} \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D', z) \right]. \end{aligned}$$

From Equation (8), we recognize that the expectation in the last line equals $\Pr[\mathcal{A}(D') = \mathbf{a}]$, giving us:

$$\Pr[\mathcal{A}(D) = \mathbf{a}] \leq e^{\varepsilon_Z + \varepsilon_Y} \Pr[\mathcal{A}(D') = \mathbf{a}].$$

By symmetry of the definitions of neighboring datasets, we can also show that $\Pr[\mathcal{A}(D') = \mathbf{a}] \leq e^{\varepsilon_Z + \varepsilon_Y} \Pr[\mathcal{A}(D) = \mathbf{a}]$.

Since this bound holds for any output \mathbf{a} , we have established that **OutsideInterval** satisfies $(\varepsilon_Z + \varepsilon_Y)$ -differential privacy. \square

B.2 (α, ε) -Rényi DP Guarantee

Proof of Theorem 1 (ii). Starting again from Equation (7), we proceed with proving the RDP guarantee.

Step 2: Computing the Rényi divergence and applying Jensen's inequality.

Let us consider the expectation in the definition of Rényi divergence:

$$\begin{aligned} \mathbb{E}_{\mathbf{a} \sim \mathcal{A}(D')} \left[\frac{\Pr[\mathcal{A}(D) = \mathbf{a}]^\alpha}{\Pr[\mathcal{A}(D') = \mathbf{a}]^\alpha} \right] &= \sum_{\mathbf{a} \in \mathcal{O}} \frac{\Pr[\mathcal{A}(D) = \mathbf{a}]^\alpha}{\Pr[\mathcal{A}(D') = \mathbf{a}]^{\alpha-1}} \\ &\leq \sum_{k=1}^{\infty} \sum_{d \in \{0,1\}} \frac{\mathbb{E}_{z \sim \mathcal{D}_Z} \left[\frac{p_Z(z-\Delta)}{p_Z(z)} \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D, z-\Delta) \right]^\alpha}{\mathbb{E}_{z \sim \mathcal{D}_Z} \left[\prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D', z) \right]^{\alpha-1}}, \end{aligned}$$

where we obtain the inequality using Equations (7) and (8).

Let us consider the bivariate function $\varphi_\alpha : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ defined as $\varphi_\alpha(x, y) = \frac{x^\alpha}{y^{\alpha-1}}$. Notice that for any $\alpha > 1$, φ_α is jointly convex over \mathbb{R}_+^2 . By Jensen's inequality, we have $\varphi_\alpha(\mathbb{E}[X], \mathbb{E}[Y]) \leq \mathbb{E}[\varphi_\alpha(X, Y)]$ for any positive random variables X and Y . Applying this inequality to our case:

$$\begin{aligned} &\mathbb{E}_{\mathbf{a} \sim \mathcal{A}(D')} \left[\frac{\Pr[\mathcal{A}(D) = \mathbf{a}]^\alpha}{\Pr[\mathcal{A}(D') = \mathbf{a}]^\alpha} \right] \\ &\leq \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\sum_{k=1}^{\infty} \sum_{d \in \{0,1\}} \frac{\left(\frac{p_Z(z-\Delta)}{p_Z(z)} \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D, z-\Delta) \right)^\alpha}{\left(\prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D', z) \right)^{\alpha-1}} \right] \\ &= \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\left(\frac{p_Z(z-\Delta)}{p_Z(z)} \right)^\alpha \mathbb{E}_{(\tau, \hat{d}) \sim \mathcal{A}(D')} \left[\frac{\left(\prod_{i=1}^{\tau-1} g_i(D', z) \cdot h_\tau^{(\hat{d})}(D, z-\Delta) \right)^\alpha}{\left(\prod_{i=1}^{\tau-1} g_i(D', z) \cdot h_\tau^{(\hat{d})}(D', z) \right)^\alpha} \middle| Z = z \right] \right] \\ &= \mathbb{E}_{z \sim \mathcal{D}_Z} \left[\left(\frac{p_Z(z-\Delta)}{p_Z(z)} \right)^\alpha \mathbb{E}_{(\tau, \hat{d}) \sim \mathcal{A}(D')} \left[\frac{\left(h_\tau^{(\hat{d})}(D, z-\Delta) \right)^\alpha}{\left(h_\tau^{(\hat{d})}(D', z) \right)^\alpha} \middle| Z = z \right] \right]. \end{aligned}$$

Step 3: Applying RDP guarantees and relating to stopping time.

We now establish a relationship between $h_k^{(d)}(D, z-\Delta)$ and $h_k^{(d)}(D', z)$ under Rényi DP for any $d \in \{0,1\}$ and $k > 0$. We prove that

$$\frac{\left(h_k^{(d)}(D, z-\Delta) \right)^\alpha}{\left(h_k^{(d)}(D', z) \right)^{\alpha-1}} \leq e^{\varepsilon_Y(\alpha)(\alpha-1)}. \quad (9)$$

To prove it for $d = 1$ (the upper threshold case) we introduce the query

$$\tilde{f}_k(D'') = \begin{cases} f_k(D) + \Delta & \text{if } D'' = D \\ f_k(D') & \text{if } D'' = D' \end{cases}.$$

To establish the sensitivity of \tilde{f}_k , we note that since f_k has sensitivity Δ , we have $|f_k(D) - f_k(D')| \leq \Delta$. Thus:

$$\begin{aligned} |\tilde{f}_k(D) - \tilde{f}_k(D')| &= |f_k(D) + \Delta - f_k(D')| \\ &\leq |f_k(D) - f_k(D')| + \Delta \\ &\leq \Delta + \Delta = 2\Delta. \end{aligned}$$

Therefore, \tilde{f}_k has sensitivity 2Δ . Applying the indistinguishability property of Rényi DP (Lemma 1) to the noise-adding mechanism with noise distribution Y_k , we obtain

$$\frac{\left(\Pr[\tilde{f}_k(D) + Y_k \geq T_1^k + z]\right)^\alpha}{\left(\Pr[\tilde{f}_k(D') + Y_k \geq T_1^k + z]\right)^{\alpha-1}} \leq e^{\varepsilon_Y(\alpha)(\alpha-1)}.$$

Since $\tilde{f}_k(D) + Y_k = f_k(D) + \Delta + Y_k$ and $\tilde{f}_k(D') + Y_k = f_k(D') + Y_k$:

$$\frac{\left(\Pr[f_k(D) + \Delta + Y_k \geq T_1^k + z]\right)^\alpha}{\left(\Pr[f_k(D') + Y_k \geq T_1^k + z]\right)^{\alpha-1}} \leq e^{\varepsilon_Y(\alpha)(\alpha-1)}.$$

Since $f_k(D) + \Delta + Y_k \geq T_1^k + z$ is equivalent to $f_k(D) + Y_k \geq T_1^k + (z - \Delta)$:

$$\frac{\left(\Pr[f_k(D) + Y_k \geq T_1^k + (z - \Delta)]\right)^\alpha}{\left(\Pr[f_k(D') + Y_k \geq T_1^k + z]\right)^{\alpha-1}} \leq e^{\varepsilon_Y(\alpha)(\alpha-1)},$$

which yields (9) for $d = 1$. The proof for $d = 0$ follows the same lines.

Using Equation (9),

$$\begin{aligned} &\mathbb{E}_{(\tau, \hat{d}) \sim \mathcal{A}(D')} \left[\frac{\left(h_\tau^{(\hat{d})}(D, z - \Delta)\right)^\alpha}{\left(h_\tau^{(\hat{d})}(D', z)\right)^\alpha} \middle| Z = z \right] \leq \mathbb{E}_{(\tau, \hat{d}) \sim \mathcal{A}(D')} \left[\frac{e^{(\alpha-1)\varepsilon_Y(\alpha)}}{\left(h_\tau^{(\hat{d})}(D', z)\right)} \middle| Z = z \right] \\ &= e^{(\alpha-1)\varepsilon_Y(\alpha)} \sum_{k=1}^{\infty} \sum_{d \in \{0,1\}} \mathbb{P}_{(\tau, \hat{d}) \sim \mathcal{A}(D')}(\tau = k, \hat{d} = d | Z = z) \frac{1}{h_k^{(d)}(D', z)} \\ &= e^{(\alpha-1)\varepsilon_Y(\alpha)} \sum_{k=1}^{\infty} \sum_{d \in \{0,1\}} \prod_{i=1}^{k-1} g_i(D', z) \cdot h_k^{(d)}(D', z) \cdot \frac{1}{h_k^{(d)}(D', z)} \\ &= 2e^{(\alpha-1)\varepsilon_Y(\alpha)} \sum_{k=1}^{\infty} \prod_{i=1}^{k-1} g_i(D', z). \end{aligned}$$

To relate this sum to an expectation, we note that $\prod_{i=1}^{k-1} g_i(D', z)$ represents the probability that the mechanism $\mathcal{A}(D')$ outputs \perp for each of the first $k-1$ queries, which means the stopping time τ is at least k . Therefore:

$$\prod_{i=1}^{k-1} g_i(D', z) = \mathbb{P}_{\mathcal{A}(D')}(\tau \geq k | Z = z).$$

Using that for a positive integer-valued random variable X , we can express its expectation using $\mathbb{E}[X] = \sum_{k=1}^{\infty} \Pr(X \geq k)$, we have

$$\sum_{k=1}^{\infty} \prod_{i=1}^{k-1} g_i(D', z) = \sum_{k=1}^{\infty} \mathbb{P}_{\mathcal{A}(D')}(\tau \geq k | Z = z) = \mathbb{E}_{\mathcal{A}(D')}[\tau | Z = z].$$

Thus, we have:

$$\mathbb{E}_{\mathcal{A}(D')} \left[\left. \frac{\left(h_{\tau}^{(\hat{d})}(D, z - \Delta) \right)^{\alpha}}{\left(h_{\tau}^{(\hat{d})}(D', z) \right)^{\alpha}} \right| Z = z \right] \leq 2e^{(\alpha-1)\varepsilon_Y(\alpha)} \cdot \mathbb{E}_{\mathcal{A}(D')}[\tau|Z = z].$$

Step 4: Combining results and applying Hölder's inequality.

Now we have:

$$\mathbb{E}_{\mathbf{a} \sim \mathcal{A}(D')} \left[\frac{\Pr[\mathcal{A}(D) = \mathbf{a}]^{\alpha}}{\Pr[\mathcal{A}(D') = \mathbf{a}]^{\alpha}} \right] \leq \mathbb{E}_{z \sim \mathcal{D}_Z} \left[2 \left(\frac{p_Z(z - \Delta)}{p_Z(z)} \right)^{\alpha} \cdot \mathbb{E}_{\mathcal{A}(D')}[\tau|Z = z] \cdot e^{(\alpha-1)\varepsilon_Y(\alpha)} \right].$$

We apply Hölder's inequality to separate the terms involving Z and τ . For any $\gamma > 1$, let γ^* be its Hölder conjugate such that $\frac{1}{\gamma^*} + \frac{1}{\gamma} = 1$. Let $X = \left(\frac{p_Z(z - \Delta)}{p_Z(z)} \right)^{\alpha}$ and $Y = \mathbb{E}_{\mathcal{A}(D')}[\tau|Z = z]$. By Hölder's inequality:

$$\begin{aligned} \mathbb{E}_{z \sim \mathcal{D}_Z}[X \cdot Y] &\leq \left(\mathbb{E}_{z \sim \mathcal{D}_Z}[X^{\gamma^*}] \right)^{1/\gamma^*} \cdot \left(\mathbb{E}_{z \sim \mathcal{D}_Z}[Y^{\gamma}] \right)^{1/\gamma} \\ &= \left(\mathbb{E}_{z \sim \mathcal{D}_Z} \left[\left(\frac{p_Z(z - \Delta)}{p_Z(z)} \right)^{\alpha\gamma^*} \right] \right)^{1/\gamma^*} \cdot \left(\mathbb{E}_{z \sim \mathcal{D}_Z}[\mathbb{E}[\tau|Z = z]^{\gamma}] \right)^{1/\gamma}. \end{aligned}$$

Therefore:

$$\begin{aligned} \mathbb{E}_{\mathbf{a} \sim \mathcal{A}(D')} \left[\frac{\Pr[\mathcal{A}(D) = \mathbf{a}]^{\alpha}}{\Pr[\mathcal{A}(D') = \mathbf{a}]^{\alpha}} \right] &\leq 2e^{(\alpha-1)\varepsilon_Y(\alpha)} \cdot \left(\mathbb{E}_{z \sim \mathcal{D}_Z} \left[\left(\frac{p_Z(z - \Delta)}{p_Z(z)} \right)^{\alpha\gamma^*} \right] \right)^{1/\gamma^*} \\ &\quad \cdot \left(\mathbb{E}_{z \sim \mathcal{D}_Z}[\mathbb{E}_{\mathcal{A}(D')}[\tau|Z = z]^{\gamma}] \right)^{1/\gamma}. \end{aligned}$$

The first expectation term is related to the RDP guarantee of the noise-adding mechanism with distribution \mathcal{D}_Z . Specifically, since this mechanism satisfies $\varepsilon_Z(\alpha\gamma^*)$ -RDP for queries with sensitivity Δ , we have:

$$\begin{aligned} \left(\mathbb{E}_{z \sim \mathcal{D}_Z} \left[\left(\frac{p_Z(z - \Delta)}{p_Z(z)} \right)^{\alpha\gamma^*} \right] \right)^{1/\gamma^*} &\leq \left(e^{(\alpha\gamma^*-1)\varepsilon_Z(\alpha\gamma^*)} \right)^{1/\gamma^*} \\ &= e^{\frac{\alpha\gamma^*-1}{\gamma^*}\varepsilon_Z(\alpha\gamma^*)}. \end{aligned}$$

Substituting $\gamma^* = \frac{\gamma}{\gamma-1}$ (from the Hölder conjugate relation), and simplifying the exponent through algebraic manipulations:

$$\begin{aligned} e^{\frac{\alpha\gamma^*-1}{\gamma^*}\varepsilon_Z(\alpha\gamma^*)} &= e^{\frac{\alpha \frac{\gamma}{\gamma-1} - 1}{\frac{\gamma}{\gamma-1}} \varepsilon_Z(\alpha \frac{\gamma}{\gamma-1})} \\ &= e^{\frac{(\alpha\gamma - \alpha + 1)(\gamma-1)}{\gamma} \varepsilon_Z(\alpha \frac{\gamma}{\gamma-1})} \\ &= e^{\frac{(\alpha-1)\gamma - \alpha + 1}{\gamma} \varepsilon_Z(\alpha \frac{\gamma}{\gamma-1})} \\ &= e^{\frac{(\alpha-1) - \frac{\alpha-1}{\gamma}}{\alpha-1} (\alpha-1) \varepsilon_Z(\alpha \frac{\gamma}{\gamma-1})} \\ &= e^{\frac{\alpha - (\gamma-1)/\gamma}{\alpha-1} (\alpha-1) \varepsilon_Z(\alpha \frac{\gamma}{\gamma-1})}. \end{aligned}$$

Taking logarithms and dividing by $(\alpha - 1)$, we obtain the bound:

$$D_{\alpha}(\mathcal{A}(D) \parallel \mathcal{A}(D')) \leq \frac{\alpha - (\gamma-1)/\gamma}{\alpha-1} \varepsilon_Z \left(\frac{\gamma}{\gamma-1} \alpha \right) + \varepsilon_Y(\alpha) + \frac{\log \left(2 \mathbb{E}_{z \sim \mathcal{D}_Z}[\mathbb{E}_{\mathcal{A}(D')}[\tau|Z = z]^{\gamma}] \right)}{\gamma(\alpha-1)}.$$

In particular, if we choose $\gamma = 2$ then we have:

$$D_\alpha(\mathcal{A}(D)\|\mathcal{A}(D')) \leq \frac{\alpha - 1/2}{\alpha - 1} \varepsilon_Z(2\alpha) + \varepsilon_Y(\alpha) + \frac{\log(2\mathbb{E}_{z \sim \mathcal{D}_Z}[\mathbb{E}_{\mathcal{A}(D')}[\tau|Z=z]^2])}{2(\alpha - 1)}.$$

This completes the proof of Theorem 1 (ii). □

C The Variants of DP-SPRT: Laplace and Gaussian Noises

C.1 DP-SPRT with Laplace Noise

For this variant, we use the following noise distributions and correction function:

$$\begin{aligned} Y_n &\sim \text{Lap}\left(\frac{4}{\varepsilon}\right), \\ Z &\sim \text{Lap}\left(\frac{2}{\varepsilon}\right), \\ C(n, \delta) &= \frac{6 \log(n^s \zeta(s)/\delta)}{n\varepsilon}. \end{aligned}$$

We now prove that this choice satisfies the condition in Equation (3):

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) \leq \delta.$$

First, we rewrite the inequality inside the probability:

$$\begin{aligned} \frac{Y_n}{n} - \frac{Z}{n} &> \frac{6 \log(n^s \zeta(s)/\delta)}{n\varepsilon} \\ \Rightarrow Y_n - Z &> \frac{6 \log(n^s \zeta(s)/\delta)}{\varepsilon}. \end{aligned}$$

Using the union bound, we have:

$$\mathbb{P}\left(Y_n - Z > \frac{6 \log(n^s \zeta(s)/\delta)}{\varepsilon}\right) \leq \mathbb{P}\left(Y_n > \frac{4 \log(n^s \zeta(s)/\delta)}{\varepsilon}\right) + \mathbb{P}\left(Z < -\frac{2 \log(n^s \zeta(s)/\delta)}{\varepsilon}\right).$$

For Laplace random variables, we have the tail bounds:

$$\text{If } X \sim \text{Lap}(b), \text{ then } \mathbb{P}(X \geq t) = \mathbb{P}(X \leq -t) = \frac{1}{2}e^{-t/b}.$$

Therefore:

$$\begin{aligned} \mathbb{P}\left(Y_n > \frac{4 \log(n^s \zeta(s)/\delta)}{\varepsilon}\right) &= \frac{1}{2}e^{-\frac{4 \log(n^s \zeta(s)/\delta)}{\varepsilon} \cdot \frac{\varepsilon}{4}} \\ &= \frac{1}{2}e^{-\log(n^s \zeta(s)/\delta)} \\ &= \frac{1}{2} \cdot \frac{\delta}{n^s \zeta(s)} \\ &= \frac{\delta}{2n^s \zeta(s)}. \end{aligned}$$

Similarly:

$$\begin{aligned} \mathbb{P}\left(Z < -\frac{2\log(n^s\zeta(s)/\delta)}{\varepsilon}\right) &= \frac{1}{2}e^{-\frac{2\log(n^s\zeta(s)/\delta)}{\varepsilon} \cdot \frac{\varepsilon}{2}} \\ &= \frac{1}{2}e^{-\log(n^s\zeta(s)/\delta)} \\ &= \frac{\delta}{2n^s\zeta(s)}. \end{aligned}$$

Combining these results:

$$\begin{aligned} \mathbb{P}\left(Y_n - Z > \frac{6\log(n^s\zeta(s)/\delta)}{n\varepsilon}\right) &\leq \frac{\delta}{2n^s\zeta(s)} + \frac{\delta}{2n^s\zeta(s)} \\ &= \frac{\delta}{n^s\zeta(s)}. \end{aligned}$$

Summing over all n :

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) &\leq \sum_{n=1}^{\infty} \frac{\delta}{n^s\zeta(s)} \\ &= \delta \sum_{n=1}^{\infty} \frac{1}{n^s\zeta(s)} \\ &= \delta \cdot \frac{1}{\zeta(s)} \sum_{n=1}^{\infty} \frac{1}{n^s} \\ &= \delta \cdot \frac{\zeta(s)}{\zeta(s)} \\ &= \delta, \end{aligned}$$

where we used the definition of the Riemann zeta function: $\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$.

C.2 DP-SPRT with Gaussian Noise

For this variant, we use the following noise distributions and correction function:

$$\begin{aligned} Y_n &\sim \mathcal{N}(0, \sigma_Y^2), \\ Z &\sim \mathcal{N}(0, \sigma_Z^2), \\ C(n, \delta) &= \frac{\sqrt{2(\sigma_Y^2 + \sigma_Z^2) \log(n^s\zeta(s)/\delta)}}{n}. \end{aligned}$$

We now prove that this choice satisfies the condition in Equation (3):

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) \leq \delta.$$

Since Y_n and Z are independent Gaussian random variables, their difference is also Gaussian. Specifically:

$$Y_n - Z \sim \mathcal{N}(0, \sigma_Y^2 + \sigma_Z^2).$$

The probability we need to bound is:

$$\mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > \frac{\sqrt{2(\sigma_Y^2 + \sigma_Z^2) \log(n^s\zeta(s)/\delta)}}{n}\right).$$

Multiplying both sides by n :

$$\mathbb{P}\left(Y_n - Z > \sqrt{2(\sigma_Y^2 + \sigma_Z^2) \log(n^s \zeta(s)/\delta)}\right).$$

For a Gaussian random variable $X \sim \mathcal{N}(0, \sigma^2)$, we have the standard tail bound:

$$\mathbb{P}(X > t) \leq e^{-t^2/2\sigma^2}.$$

Applying this bound with $X = Y_n - Z$, $\sigma^2 = \sigma_Y^2 + \sigma_Z^2$, and $t = \sqrt{2(\sigma_Y^2 + \sigma_Z^2) \log(n^s \zeta(s)/\delta)}$:

$$\begin{aligned} \mathbb{P}(Y_n - Z > t) &\leq e^{-t^2/2(\sigma_Y^2 + \sigma_Z^2)} \\ &= e^{-2(\sigma_Y^2 + \sigma_Z^2) \log(n^s \zeta(s)/\delta)/2(\sigma_Y^2 + \sigma_Z^2)} \\ &= e^{-\log(n^s \zeta(s)/\delta)} \\ &= \frac{\delta}{n^s \zeta(s)}. \end{aligned}$$

Therefore:

$$\mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) \leq \frac{\delta}{n^s \zeta(s)}.$$

Summing over all n :

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) &\leq \sum_{n=1}^{\infty} \frac{\delta}{n^s \zeta(s)} \\ &= \delta \sum_{n=1}^{\infty} \frac{1}{n^s \zeta(s)} \\ &= \delta \cdot \frac{1}{\zeta(s)} \sum_{n=1}^{\infty} \frac{1}{n^s} \\ &= \delta \cdot \frac{\zeta(s)}{\zeta(s)} \\ &= \delta. \end{aligned}$$

This demonstrates that the Gaussian noise implementation with the specified correction function satisfies condition (3).

D Correctness of DP-SPRT for Exponential Families: Proof of Theorem 3

In this appendix, we provide the detailed proof of Theorem 3, which establishes the error guarantees of our DP-SPRT. This result is central to our paper as it formally demonstrates that the algorithm achieves the desired privacy-utility trade-off while maintaining statistical validity.

Proof. Type I Error. We prove the bound on type I error first. Let $\gamma \in (0, 1)$ be an arbitrary parameter that allocates the error probability.

$$\begin{aligned} \mathbb{P}_{\theta_0}(\hat{d} = 1) &\leq \mathbb{P}_{\theta_0}(\tau_1 < \infty) \\ &= \mathbb{P}_{\theta_0}\left(\exists n \in \mathbb{N}_{>0} : \frac{\sum_{i=1}^n X_i}{n} - \mu_1 + \frac{Y_n}{n} \geq \frac{-\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) + \log(1/(\gamma\alpha))/n}{\theta_1 - \theta_0} + \frac{Z}{n} + C(n, (1 - \gamma)\alpha)\right). \end{aligned}$$

By the union bound, we can separate this into two events:

$$\begin{aligned} \mathbb{P}_{\theta_0}(\hat{d} = 1) &\leq \mathbb{P}_{\theta_0}\left(\exists n \in \mathbb{N}_{>0} : \frac{\sum_{i=1}^n X_i}{n} - \mu_1 \geq \frac{-\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) + \log(1/(\gamma\alpha))/n}{\theta_1 - \theta_0}\right) \\ &\quad + \mathbb{P}_{\theta_0}\left(\exists n \in \mathbb{N}_{>0} : \frac{Y_n}{n} - \frac{Z_1}{n} \geq C(n, (1-\gamma)\alpha)\right). \end{aligned}$$

For the first term, we can recognize the probability of eventually rejecting H_0 in the non-private SPRT. Let us define τ_{SPRT} as the first time the likelihood ratio exceeds the threshold for rejecting H_0 :

$$\tau_{\text{SPRT}} = \inf \left\{ n \in \mathbb{N}_{>0} : \log \left(\prod_{i=1}^n \frac{f_{\theta_1}(X_i)}{f_{\theta_0}(X_i)} \right) \geq \log \left(\frac{1}{\gamma\alpha} \right) \right\}.$$

For exponential families, this is equivalent to:

$$\tau_{\text{SPRT}} = \inf \left\{ n \in \mathbb{N}_{>0} : \frac{\sum_{i=1}^n X_i}{n} - \mu_1 \geq \frac{-\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) + \log(1/(\gamma\alpha))/n}{\theta_1 - \theta_0} \right\}.$$

Thus, the first term equals $\mathbb{P}_{\theta_0}(\tau_{\text{SPRT}} < \infty)$, which is the probability under θ_0 that we eventually reject H_0 in the non-private SPRT. Let $L_n = \prod_{i=1}^n \frac{f_{\theta_1}(X_i)}{f_{\theta_0}(X_i)}$ be the likelihood ratio.

Applying Wald's likelihood ratio identity (i.e., a change-of-measure argument) to this specific stopping event, we get

$$\mathbb{P}_{\theta_0}(\tau_{\text{SPRT}} < \infty) = \mathbb{E}_{\theta_1} \left[\frac{1}{L_{\tau_{\text{SPRT}}}} \cdot \mathbf{1}_{\{\tau_{\text{SPRT}} < \infty\}} \right].$$

By definition of τ_{SPRT} , on the event $\{\tau_{\text{SPRT}} < \infty\}$, we have $L_{\tau_{\text{SPRT}}} \geq \frac{1}{\gamma\alpha}$, which implies $\frac{1}{L_{\tau_{\text{SPRT}}}} \leq \gamma\alpha$. Therefore,

$$\begin{aligned} \mathbb{P}_{\theta_0}(\tau_{\text{SPRT}} < \infty) &= \mathbb{E}_{\theta_1} \left[\frac{1}{L_{\tau_{\text{SPRT}}}} \cdot \mathbf{1}_{\{\tau_{\text{SPRT}} < \infty\}} \right] \\ &\leq \mathbb{E}_{\theta_1} [\gamma\alpha \cdot \mathbf{1}_{\{\tau_{\text{SPRT}} < \infty\}}] \\ &= \gamma\alpha \cdot \mathbb{P}_{\theta_1}(\tau_{\text{SPRT}} < \infty) \\ &\leq \gamma\alpha. \end{aligned}$$

For the second term, using condition (3) and a union bound,

$$\begin{aligned} \mathbb{P}_{\theta_0}\left(\exists n \in \mathbb{N}_{>0} : \frac{Y_n}{n} - \frac{Z_1}{n} \geq C(n, (1-\gamma)\alpha)\right) &\leq \sum_{n=1}^{\infty} \mathbb{P}_{\theta_0}\left(\frac{Y_n}{n} - \frac{Z_1}{n} \geq C(n, (1-\gamma)\alpha)\right) \\ &\leq (1-\gamma)\alpha. \end{aligned}$$

Therefore

$$\begin{aligned} \mathbb{P}_{\theta_0}(\hat{d} = 1) &\leq \gamma\alpha + (1-\gamma)\alpha \\ &= \alpha. \end{aligned}$$

Type II Error. We now prove the bound on type II error. We have

$$\begin{aligned} \mathbb{P}_{\theta_1}(\hat{d} = 0) &\leq \mathbb{P}_{\theta_1}(\tau_0 < \infty) \\ &= \mathbb{P}_{\theta_1}\left(\exists n \in \mathbb{N}_{>0} : \frac{\sum_{i=1}^n X_i}{n} - \mu_0 + \frac{Y_n}{n} \leq \right. \\ &\quad \left. \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} - \frac{Z}{n} - C(n, (1-\gamma)\beta)\right) \end{aligned}$$

By the union bound, we can separate this into two events:

$$\begin{aligned} \mathbb{P}_{\theta_1}(\hat{d} = 0) &\leq \mathbb{P}_{\theta_1}\left(\exists n \in \mathbb{N}_{>0} : \frac{\sum_{i=1}^n X_i}{n} - \mu_0 \leq \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \log(1/(\gamma\beta))/n}{\theta_1 - \theta_0}\right) \\ &\quad + \mathbb{P}_{\theta_1}\left(\exists n \in \mathbb{N}_{>0} : \frac{Y_n}{n} + \frac{Z}{n} \leq -C(n, (1-\gamma)\beta)\right) \end{aligned}$$

For the first term, we recognize the probability of eventually accepting H_0 in the non-private SPRT when the truth is H_1 . By a similar change-of-measure argument as for type I error, this probability is bounded by $\gamma\beta$.

For the second term, using condition (4), we have:

$$\begin{aligned} \mathbb{P}_{\theta_1}\left(\exists n \in \mathbb{N}_{>0} : \frac{Y_n}{n} + \frac{Z}{n} \leq -C(n, (1-\gamma)\beta)\right) &\leq \sum_{n=1}^{\infty} \mathbb{P}_{\theta_1}\left(\frac{Y_n}{n} + \frac{Z}{n} < -C(n, (1-\gamma)\beta)\right) \\ &\leq (1-\gamma)\beta \end{aligned}$$

Therefore $\mathbb{P}_{\theta_1}(\hat{d} = 0) \leq \gamma\beta + (1-\gamma)\beta = \beta$.

Finally, we note that if the distribution \mathcal{D}_Y is symmetric around zero and condition (3) holds, then condition (4) is automatically satisfied. This is because:

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} + \frac{Z}{n} < -C(n, \delta)\right) &= \sum_{n=1}^{\infty} \mathbb{P}\left(-\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) \\ &= \sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) \leq \delta \end{aligned}$$

where the second equality uses the symmetry of Y_n around zero. \square

E Sample Complexity of DP-SPRT for Bernoulli: Proof of Theorem 4

We begin with the proof of the generic sample complexity result (Theorem 4), which provides an upper bound on the expected stopping time for any noise distribution satisfying appropriate conditions.

Proof of Theorem 4. We first derive the upper bound on $\mathbb{E}_{\theta_0}[\tau]$. Using the formula for the expectation of a non-negative integer-valued random variable,

$$\begin{aligned} \mathbb{E}_{\theta_0}[\tau] &\leq \mathbb{E}_{\theta_0}[\tau_0] = 1 + \sum_{n=1}^{\infty} \mathbb{P}_{\theta_0}(\tau_0 > n) \\ &\leq 1 + \sum_{n=1}^{\infty} \mathbb{P}_{\theta_0}\left(\frac{\sum_{i=1}^n X_i}{n} + \frac{Y_n}{n} > \mu_0 + \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0} - \frac{\log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} + \frac{Z}{n} - C(n, (1-\gamma)\beta)\right) \\ &\leq 1 + (1-\gamma)\beta + \sum_{n=1}^{\infty} \mathbb{P}_{\theta_0}\left(\frac{\sum_{i=1}^n X_i}{n} > \mu_0 + \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0} - \frac{\log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} - 2C(n, (1-\gamma)\beta)\right), \end{aligned}$$

where in the last inequality we have used a union bound together with the assumption that $\sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, \delta)\right) \leq \delta$.

Using the definition of $N(\theta_0, \theta_1, \beta, \gamma)$, which is the minimum number of samples needed to ensure that

$$\frac{\log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} + 2C(n, (1-\gamma)\beta) \leq \frac{1}{2} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}$$

we can split the sum and write

$$\begin{aligned} \mathbb{E}_{\theta_0}[\tau] &\leq 1 + (1-\gamma)\beta + N(\theta_0, \theta_1, \beta, \gamma) + \sum_{n=N_0}^{\infty} \mathbb{P}_{\theta_0}\left(\frac{\sum_{i=1}^n X_i}{n} > \mu_0 + \frac{1}{2} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}\right) \\ &\leq 1 + (1-\gamma)\beta + N(\theta_0, \theta_1, \beta, \gamma) + \sum_{n=1}^{\infty} \exp\left(-2 \left(\frac{1}{2} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}\right)^2 n\right), \end{aligned}$$

where the last step uses Hoeffding's inequality, as Bernoulli distributions are bounded in $[0, 1]$. Computing the series on the right-hand side and using Pinsker's inequality yields the bounded stated in Theorem 4.

We now upper bound $\mathbb{E}_{\theta_1}[\tau]$ following similar lines:

$$\begin{aligned} \mathbb{E}_{\theta_1}[\tau] &\leq \mathbb{E}_{\theta_1}[\tau_1] = 1 + \sum_{n=1}^{\infty} \mathbb{P}_{\theta_1}(\tau_1 > n) \\ &\leq 1 + \sum_{n=1}^{\infty} \mathbb{P}_{\theta_1} \left(\frac{\sum_{i=1}^n X_i}{n} + \frac{Y_n}{n} < \mu_1 - \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0})}{\theta_1 - \theta_0} + \frac{\log(1/(\gamma\alpha))/n}{\theta_1 - \theta_0} + \frac{Z}{n} + C(n, (1-\gamma)\alpha) \right) \\ &\leq 1 + (1-\gamma)\alpha + \sum_{n=1}^{\infty} \mathbb{P}_{\theta_1} \left(\frac{\sum_{i=1}^n X_i}{n} < \mu_1 - \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0})}{\theta_1 - \theta_0} + \frac{\log(1/(\gamma\alpha))/n}{\theta_1 - \theta_0} + 2C(n, (1-\gamma)\alpha) \right) \end{aligned}$$

where in the last inequality we have used that $\sum_{n=1}^{\infty} \mathbb{P} \left(\frac{Z}{n} - \frac{Y_n}{n} > C(n, \delta) \right) \leq \delta$.

Using this time the definition of $N(\theta_1, \theta_0, \alpha, \gamma)$ yields

$$\begin{aligned} \mathbb{E}_{\theta_1}[\tau] &\leq 1 + (1-\gamma)\alpha + N(\theta_1, \theta_0, \alpha, \gamma) + \sum_{n=1}^{\infty} \mathbb{P}_{\theta_1} \left(\frac{\sum_{i=1}^n X_i}{n} < \mu_1 - \frac{1}{2} \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0})}{\theta_1 - \theta_0} \right) \\ &\leq 1 + (1-\gamma)\alpha + N(\theta_1, \theta_0, \alpha, \gamma) + \sum_{n=1}^{\infty} \exp \left(-2 \left(\frac{1}{2} \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0})}{\theta_1 - \theta_0} \right)^2 n \right) \end{aligned}$$

The final bound on $\mathbb{E}_{\theta_1}[\tau]$ follows by computing the last series and applying Pinsker's inequality. □

E.1 Analysis for Laplace Noise

For [DP-SPRT](#) with Laplace noise, the correction function is:

$$C(n, \delta) = \frac{6 \log(n^s \zeta(s)/2\delta)}{n\varepsilon}.$$

To upper bound $N(\theta_0, \theta_1, \beta, \gamma)$, we consider n such that:

$$\frac{\log(1/(\gamma\beta))}{\theta_1 - \theta_0} \frac{1}{n} + \frac{12 \log(n^s \zeta(s)/2(1-\gamma)\beta)}{n\varepsilon} \leq \frac{1}{2} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0},$$

or equivalently:

$$\frac{2 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} + \frac{24(\theta_1 - \theta_0) \log(\zeta(s)/2(1-\gamma)\beta)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} + \frac{24s(\theta_1 - \theta_0)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \log(n) \leq n.$$

Using Lemma 3 stated below, such an n , and therefore $N(\theta_0, \theta_1, \beta, \gamma)$, is upper bounded by:

$$\begin{aligned} &\frac{2 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} + \frac{24(\theta_1 - \theta_0) \log(\zeta(s)/2(1-\gamma)\beta)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \\ &+ \frac{24s(\theta_1 - \theta_0)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \log \left(\left[\frac{24s(\theta_1 - \theta_0)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \right]^2 + \frac{4 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} \right) \\ &+ \frac{48(\theta_1 - \theta_0) \log(\zeta(s)/2(1-\gamma)\beta)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon}. \end{aligned}$$

Hence it follows from Theorem 4 that

$$\begin{aligned}
 \mathbb{E}_{\theta_0}[\tau] &\leq \frac{2 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} + \frac{24(\theta_1 - \theta_0) \log(\zeta(s)/2(1-\gamma)\beta)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \\
 &+ \frac{24s(\theta_1 - \theta_0)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \log \left(\left[\frac{24s(\theta_1 - \theta_0)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \right]^2 + \frac{4 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})} + \frac{48(\theta_1 - \theta_0) \log(\frac{\zeta(s)}{2(1-\gamma)\beta})}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \right) \\
 &+ 1 + (1-\gamma)\beta + \frac{1}{1 - \exp\left(-\frac{(\mu_1 - \mu_0)^4}{2(\theta_1 - \theta_0)^2}\right)}.
 \end{aligned}$$

A similar bound can be obtained for $\mathbb{E}_{\theta_1}[\tau]$ by upper bounding $N(\theta_1, \theta_0, \alpha, \gamma)$.

Looking at the leading terms in these bounds when α and β go to zero yields

$$\limsup_{\beta \rightarrow 0} \frac{\mathbb{E}_{\theta_0}[\tau]}{\log(1/\beta)} \leq 48 \max \left(\frac{1}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}, \frac{\theta_1 - \theta_0}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \right)$$

and

$$\limsup_{\alpha \rightarrow 0} \frac{\mathbb{E}_{\theta_1}[\tau]}{\log(1/\alpha)} \leq 48 \max \left(\frac{1}{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0})}, \frac{\theta_1 - \theta_0}{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0})\varepsilon} \right).$$

Lemma 3 (Lemma 19 from [Tirinzoni et al. \(2022\)](#)). *Let $B, C \geq 1$. If $k \leq B \log(k) + C$, then*

$$k \leq B \log(B^2 + 2C) + C.$$

E.2 Analysis for Gaussian Noise

For DP-SPRT with Gaussian noise, the correction function is:

$$C(n, \delta) = \frac{\sqrt{2(\sigma_Y^2 + \sigma_Z^2) \log(n^s \zeta(s)/2\delta)}}{n}.$$

To upper bound $N(\theta_0, \theta_1, \beta, \gamma)$, we need to find the minimum n such that:

$$\frac{\log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} + 2C(n, (1-\gamma)\beta) \leq \frac{1}{2} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}.$$

It is sufficient to enforce that each term on the left-hand side is less than half of the right-hand side:

$$\begin{aligned}
 \frac{\log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} &\leq \frac{1}{4} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}, \\
 2C(n, (1-\gamma)\beta) &\leq \frac{1}{4} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}.
 \end{aligned}$$

From the first condition, we directly get:

$$n \geq \frac{4 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}.$$

For the second condition, substituting the correction function:

$$2 \cdot \frac{\sqrt{2(\sigma_Y^2 + \sigma_Z^2) \log(n^s \zeta(s)/2(1-\gamma)\beta)}}{n} \leq \frac{1}{4} \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}.$$

Let $\mathfrak{C} = \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0}$. Squaring both sides and making the change of variable $m = n^2$:

$$8(\sigma_Y^2 + \sigma_Z^2) \log(m^{s/2} \zeta(s)/2(1-\gamma)\beta) \leq \frac{m \mathfrak{C}^2}{16}.$$

Rearranging:

$$m \geq \frac{128(\sigma_Y^2 + \sigma_Z^2)}{\mathfrak{e}^2} \log(m^{s/2} \zeta(s)/2(1-\gamma)\beta).$$

This is of the form $k \geq B \log(k) + C$ where:

$$\begin{aligned} k &= m, \\ B &= \frac{64s(\sigma_Y^2 + \sigma_Z^2)}{\mathfrak{e}^2}, \\ C &= \frac{128(\sigma_Y^2 + \sigma_Z^2) \log(\zeta(s)/2(1-\gamma)\beta)}{\mathfrak{e}^2}. \end{aligned}$$

Using Lemma 3:

$$m \geq B \log(B^2 + 2C) + C.$$

Converting back to n with $m = n^2$:

$$\begin{aligned} n &\geq \frac{8\sqrt{(\sigma_Y^2 + \sigma_Z^2)}}{|\mathfrak{C}|} \\ &\cdot \sqrt{s \log \left(\frac{(64s)^2(\sigma_Y^2 + \sigma_Z^2)^2}{\mathfrak{e}^4} + \frac{256(\sigma_Y^2 + \sigma_Z^2) \log(\zeta(s)/2(1-\gamma)\beta)}{\mathfrak{e}^2} \right) + 2 \log(\zeta(s)/2(1-\gamma)\beta)}. \end{aligned}$$

Therefore, $N(\theta_0, \theta_1, \beta, \gamma)$ is upper bounded by the maximum of:

$$\frac{4 \log(1/(\gamma\beta))}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}$$

and

$$\begin{aligned} &\frac{8\sqrt{(\sigma_Y^2 + \sigma_Z^2)}}{|\mathfrak{C}|} \\ &\cdot \sqrt{s \log \left(\frac{(64s)^2(\sigma_Y^2 + \sigma_Z^2)^2}{\mathfrak{e}^4} + \frac{256(\sigma_Y^2 + \sigma_Z^2) \log(\zeta(s)/2(1-\gamma)\beta)}{\mathfrak{e}^2} \right) + 2 \log(\zeta(s)/2(1-\gamma)\beta)}. \end{aligned}$$

F Lower Bound: Proof of Theorem 5

As the proof could generalize to other than Bernoulli distributions, we denote by ν_{θ_0} (resp. ν_{θ_1}) the distribution of each X_i under \mathcal{H}_0 (resp. \mathcal{H}_1) and by \mathbb{P}_{θ_0} the probability space under which $(X_i)_{i \in \mathbb{N}}$ is an iid stream distribution under ν_{θ_0} (resp. ν_{θ_1}).

We let $\mathbb{P}_{\theta_i}^Y$ denote the distribution of a random vector Y under the probability model \mathbb{P}_{θ_i} . To ease the notation, we will sometimes denote by $X_{1:t}$ the random vector (X_1, \dots, X_t) .

We establish the lower bound by combining information-theoretic inequalities with privacy constraints. Let (τ, d) be a (α, β) -correct test. We assume that τ is almost surely finite under both \mathbb{P}_{θ_0} and \mathbb{P}_{θ_1} , otherwise any lower bound on its expectation is trivially true.

Step 1: Lower bounding a KL divergence for any (α, β) -correct test. Using the data processing inequality, one can write

$$\begin{aligned} \text{KL} \left(\mathbb{P}_{\theta_0}^{(X_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\theta_1}^{(X_{1:\tau}, \tau, \hat{d})} \right) &\geq \text{KL} \left(\mathbb{P}_{\theta_0}^{\mathbf{1}(\hat{d}=1)}, \mathbb{P}_{\theta_1}^{\mathbf{1}(\hat{d}=1)} \right) \\ &= \text{kl} \left(\mathbb{P}_{\theta_0}(\hat{d} = 1), \mathbb{P}_{\theta_1}(\hat{d} = 1) \right) \\ &\geq \text{kl}(\alpha, 1 - \beta), \end{aligned}$$

where the last step uses the correctness of the test and monotonicity properties of the Bernoulli KL divergence, $\text{kl}(x, y) = x \log(x/y) + (1-x) \log((1-x)/(1-y))$.

This argument is quite common to prove lower bounds on regret or sample complexity in the multi-armed bandit literature, and was popularized by, e.g. [Garivier et al. \(2019\)](#); [Garivier and Kaufmann \(2016\)](#).

Step 2: Computing the KL for any test. Still using standard arguments, we can establish the following equality:

$$\text{KL} \left(\mathbb{P}_{\theta_0}^{(X_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\theta_1}^{(X_{1:\tau}, \tau, \hat{d})} \right) = \mathbb{E}_{\theta_0}[\tau] \cdot \text{KL}(\nu_{\theta_0}, \nu_{\theta_1}). \quad (10)$$

For completeness, we provide the proof of this result. Assuming ν_{θ_i} has a density f_{θ_i} with respect to a common reference measure, we can express:

$$\text{KL} \left(\mathbb{P}_{\theta_0}^{(X_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\theta_1}^{(X_{1:\tau}, \tau, \hat{d})} \right) = \mathbb{E}_{\theta_0} \left[\log \prod_{i=1}^{\tau} \frac{f_{\theta_0}(X_i)}{f_{\theta_1}(X_i)} \right] = \mathbb{E}_{\theta_0} \left[\sum_{i=1}^{\tau} \log \frac{f_{\theta_0}(X_i)}{f_{\theta_1}(X_i)} \right].$$

Let $Z_i = \log \frac{f_{\theta_0}(X_i)}{f_{\theta_1}(X_i)}$. Since the observations are i.i.d., each Z_i has mean: $\mathbb{E}_{\theta_0}[Z_i] = \text{KL}(\nu_{\theta_0}, \nu_{\theta_1})$. By Wald's identity [Wald \(1945\)](#) (as τ has finite expectation), we have

$$\mathbb{E}_{\theta_0} \left[\sum_{i=1}^{\tau} Z_i \right] = \mathbb{E}_{\theta_0}[\tau] \cdot \mathbb{E}_{\theta_0}[Z_1] = \mathbb{E}_{\theta_0}[\tau] \cdot \text{KL}(\mathbb{P}_{\theta_0}, \mathbb{P}_{\theta_1}),$$

and we obtain (10).

Step 3: Upper Bounding the KL for Private Tests. Using the privacy constraints, we now prove that for an ε -DP test we further have

$$\text{KL} \left(\mathbb{P}_{\theta_0}^{(X_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\theta_1}^{(X_{1:\tau}, \tau, \hat{d})} \right) \leq \varepsilon \cdot \text{TV}(\nu_{\theta_0}, \nu_{\theta_1}) \cdot \mathbb{E}_{\theta_0}[\tau]. \quad (11)$$

The argument is similar to the proof of Theorem 4.9 from [Azize \(2024\)](#), which applies to a more complex bandit model but with a finite horizon T . The idea is to build a coupling between the data collected under \mathbb{P}_{θ_0} and \mathbb{P}_{θ_1} .

Specifically, we let $\underline{\nu}$ be the maximal coupling between ν_{θ_0} and ν_{θ_1} : $(X, X') \sim \underline{\nu}$ is such that the marginal distribution of X is ν_{θ_0} , the marginal distribution of X' is ν_{θ_1} and $\mathbb{E}_{(X, X') \sim \underline{\nu}}[\mathbb{1}(X \neq X')] = \text{TV}(\nu_{\theta_0}, \nu_{\theta_1})$. We now define the probability space $\mathbb{P}_{\underline{\nu}}$ in which we collect iid pairs $(X_i, X'_i) \sim \underline{\nu}$ and let (τ, \hat{d}) be the stopping and decision rules of the test based on the iid stream X_i (distributed under ν_{θ_0}) and (τ', \hat{d}') be the stopping and decision rules of the test based on the iid stream X'_i (distributed under ν_{θ_1}).

We observe that the marginal distribution of $(X_{1:\tau}, \tau, \hat{d})$ in the probability space $\mathbb{P}_{\underline{\nu}}$ is $\mathbb{P}_{\theta_0}^{(X_{1:\tau}, \tau, \hat{d})}$ while that of $(X'_{1:\tau'}, \tau', \hat{d}')$ is $\mathbb{P}_{\theta_1}^{(X'_{1:\tau'}, \tau', \hat{d}')$. Using again the data-processing inequality,

$$\text{KL} \left(\mathbb{P}_{\theta_0}^{(X_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\theta_1}^{(X_{1:\tau}, \tau, \hat{d})} \right) \leq \text{KL} \left(\mathbb{P}_{\underline{\nu}}^{(X_{1:\tau}, X'_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\underline{\nu}}^{(X_{1:\tau'}, X'_{1:\tau'}, \tau', \hat{d}')} \right).$$

To compute this last KL-divergence, for clarity we will consider the special case of Bernoulli distributions that is the main focus of this paper. In that case the distribution of $(X_{1:\tau}, X'_{1:\tau}, \tau, \hat{d})$ and $(X_{1:\tau'}, X'_{1:\tau'}, \tau', \hat{d}')$ are discrete and we have

$$\begin{aligned} & \text{KL} \left(\mathbb{P}_{\underline{\nu}}^{(X_{1:\tau}, X'_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\underline{\nu}}^{(X_{1:\tau'}, X'_{1:\tau'}, \tau', \hat{d}')} \right) \\ &= \sum_{\substack{T > 0, d \in \{0,1\} \\ x_{1:T} \in \{0,1\}^T \\ x'_{1:T} \in \{0,1\}^T}} \mathbb{P}_{\underline{\nu}} \left(X_{1:T} = x_{1:T}, X'_{1:T} = x'_{1:T}, \tau = T, \hat{d} = d \right) \log \frac{\mathbb{P}_{\underline{\nu}} \left(X_{1:T} = x_{1:T}, X'_{1:T} = x'_{1:T}, \tau = T, \hat{d} = d \right)}{\mathbb{P}_{\underline{\nu}} \left(X_{1:T} = x_{1:T}, X'_{1:T} = x'_{1:T}, \tau' = T, \hat{d}' = d \right)} \\ &\leq \sum_{\substack{T > 0, d \in \{0,1\} \\ x_{1:T} \in \{0,1\}^T \\ x'_{1:T} \in \{0,1\}^T}} \mathbb{P}_{\underline{\nu}} \left(X_{1:T} = x_{1:T}, X'_{1:T} = x'_{1:T}, \tau = T, \hat{d} = d \right) \log \frac{\mathbb{P}_{\underline{\nu}} \left(\tau = T, \hat{d} = d | X_{1:T} = x_{1:T} \right)}{\mathbb{P}_{\underline{\nu}} \left(\tau' = T, \hat{d}' = d | X'_{1:T} = x'_{1:T} \right)}. \end{aligned}$$

Using the fact that the test is ε -DP we obtain by the group-privacy property that

$$\frac{\mathbb{P}_{\underline{\nu}}\left(\tau = T, \hat{d} = d | X_{1:T} = x_{1:T}\right)}{\mathbb{P}_{\underline{\nu}}\left(\tau' = T, \hat{d}' = d | X'_{1:T} = x'_{1:T}\right)} \leq e^{\varepsilon \sum_{i=1}^T \mathbb{1}(x_i \neq x'_i)},$$

hence

$$\begin{aligned} & \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{(X_{1:\tau}, X'_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\underline{\nu}}^{(X_{1:\tau'}, X'_{1:\tau'}, \tau', \hat{d}')}\right) \\ & \leq \sum_{\substack{T > 0, d \in \{0,1\} \\ x_{1:T} \in \{0,1\}^T \\ x'_{1:T} \in \{0,1\}^T}} \mathbb{P}_{\underline{\nu}}\left(X_{1:T} = x_{1:T}, X'_{1:T} = x'_{1:T}, \tau = T, \hat{d} = d\right) \varepsilon \sum_{i=1}^T \mathbb{1}(x_i \neq x'_i) \\ & = \varepsilon \mathbb{E}_{\underline{\nu}}\left[\sum_{t=1}^{\tau} \mathbb{1}(X_t \neq X'_t)\right] \\ & = \varepsilon \mathbb{E}_{\underline{\nu}}[\tau] \text{TV}(\nu_{\theta_0}, \nu_{\theta_1}) \\ & = \varepsilon \mathbb{E}_{\theta_0}[\tau] \text{TV}(\nu_{\theta_0}, \nu_{\theta_1}), \end{aligned}$$

where the last but one step uses Wald's equality and the fact that $\mathbb{E}_{\underline{\nu}}[\mathbb{1}(X_i \neq X'_i)] = \text{TV}(\nu_{\theta_0}, \nu_{\theta_1})$ by definition of the coupling. This proves (11).

Step 4: Putting things together. Combining the lower bound with (10) and the private upper bound (11), we have

$$\text{kl}(\alpha, 1 - \beta) \leq \text{KL}\left(\mathbb{P}_{\theta_0}^{(X_{1:\tau}, \tau, \hat{d})}, \mathbb{P}_{\theta_1}^{(X_{1:\tau}, \tau, \hat{d})}\right) \leq \min(\mathbb{E}_{\theta_0}[\tau] \cdot \text{KL}(\nu_{\theta_0}, \nu_{\theta_1}), \varepsilon \cdot \text{TV}(\nu_{\theta_0}, \nu_{\theta_1}) \cdot \mathbb{E}_{\theta_0}[\tau]).$$

Rearranging, we obtain:

$$\mathbb{E}_{\theta_0}[\tau] \geq \frac{\text{kl}(\alpha, 1 - \beta)}{\min(\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}), \varepsilon \cdot \text{TV}(\nu_{\theta_0}, \nu_{\theta_1}))}.$$

The lower bound on $\mathbb{E}_{\theta_1}[\tau]$ follows in the same way by swapping the roles of θ_0 and θ_1 .

G Enhancing DP-SPRT with Subsampling

Privacy amplification by subsampling (Steinke, 2022) provides an approach to enhance the privacy-utility trade-off in differentially private algorithms. When a dataset is randomly subsampled before applying a private mechanism, the effective privacy cost can be reduced roughly proportionally to the sampling rate for pure differentially private mechanisms.

In our DP-SPRT context, subsampling allows us to improve privacy for a given noise level, or equivalently, to maintain the same level of privacy while requiring less noise. At each step when a new observation X_i arrives, we implement Bernoulli subsampling by including it in our computation with probability $r \in (0, 1)$, independent of all previous decisions. We maintain a running sum S_n and count M_n of the included observations:

1. For each arriving observation X_i , include it with probability r (draw $B_i \sim \text{Bernoulli}(r)$ independently);
2. Update $S_n = \sum_{i=1}^n B_i X_i$ and $M_n = \sum_{i=1}^n B_i$;
3. Compute the subsampled empirical mean $\bar{X}_n^r = S_n / M_n$ when $M_n > 0$.

We chose this Bernoulli subsampling approach over alternatives like sampling without replacement because it requires only constant space and time complexity at each time step and thus can directly handle streaming observations with limited memory, providing significant computational advantages for long trials or challenging problem instances.

We incorporate subsampling into our DP-SPRT with Laplace noise by using noise scaled for ε/r -DP, since subsampling with rate r amplifies privacy by this factor, resulting in ε -DP overall. The subsampling mechanism is defined as follows:

$$\begin{aligned}\tau^{\varepsilon,r} &= \min(\tau_0^{\varepsilon,r}, \tau_1^{\varepsilon,r}) \\ \tau_0^{\varepsilon,r} &= \inf \left\{ n \in \mathbb{N}_{>0} : \bar{X}_n^r + r \frac{Y_n}{n} \leq \mu_0 + \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \log(1/(\gamma\beta))/M_n}{\theta_1 - \theta_0} \right. \\ &\quad \left. - r \frac{Z}{n} - r \frac{6 \log(n^s \zeta(s)/(1-\gamma)\beta)}{n\varepsilon} \right\} \\ \tau_1^{\varepsilon,r} &= \inf \left\{ n \in \mathbb{N}_{>0} : \bar{X}_n^r + r \frac{Y_n}{n} \geq \mu_1 + \frac{-\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) + \log(1/(\gamma\alpha))/M_n}{\theta_1 - \theta_0} \right. \\ &\quad \left. + r \frac{Z}{n} + r \frac{6 \log(n^s \zeta(s)/(1-\gamma)\alpha)}{n\varepsilon} \right\},\end{aligned}$$

where $M_n = \sum_{i=1}^n B_i$ is the actual count of included observations, $\bar{X}_n^r = S_n/M_n$ is the subsampled empirical mean, and $Y_n \sim \text{Lap}(4/\varepsilon)$, $Z \sim \text{Lap}(2/\varepsilon)$ as in standard DP-SPRT. The decision rule remains $\hat{d}^{\varepsilon,r} = i$ if $\tau^{\varepsilon,r} = \tau_i^{\varepsilon,r}$. The terms in red are the modifications compared to the standard DP-SPRT algorithm.

Privacy and Correctness: The subsampled DP-SPRT maintains ε -differential privacy and satisfies the error bounds $\mathbb{P}_{\theta_0}(\hat{d}^{\varepsilon,r} = 1) \leq \alpha$ and $\mathbb{P}_{\theta_1}(\hat{d}^{\varepsilon,r} = 0) \leq \beta$. The privacy follows from the subsampling amplification applied to the standard DP-SPRT analysis, while correctness uses appropriately scaled thresholds.

Sample Complexity: For the asymptotic regime where $\alpha, \beta \rightarrow 0$, the expected stopping time will be of the form:

$$\limsup_{\beta \rightarrow 0} \frac{\mathbb{E}_{\theta_0}[\tau_0^{\varepsilon,r}]}{\log(1/\beta)} \leq 48 \max \left(\frac{1}{r \cdot \text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}, \frac{r \cdot (\theta_1 - \theta_0)}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon} \right).$$

This bound reveals a fundamental trade-off: smaller subsampling rates r lead to worse statistical efficiency—the first term $1/(r \cdot \text{KL})$ increases due to higher variance of the subsampled empirical mean—but provide better privacy amplification as the second term $r \cdot (\theta_1 - \theta_0)/(\text{KL} \cdot \varepsilon)$ decreases, reducing the impact of privacy noise. The optimal sampling rate depends on which effect dominates, determined by the privacy parameter ε and how hard the hypotheses are to distinguish. When ε is small and high privacy is required, aggressive subsampling with small r can be beneficial as privacy amplification becomes more valuable. When ε is large and privacy is relaxed, less aggressive subsampling is preferable to maintain statistical efficiency. As demonstrated in our experiments (Section 6), this trade-off is particularly pronounced in high-privacy regimes.

For practical applications, the bound suggests setting the sampling rate as $r \approx \min(1, \sqrt{\varepsilon/(\theta_1 - \theta_0)})$ to balance the statistical and privacy terms. However, in our experiments we used the arbitrary choice $r = \min(1, \sqrt{\varepsilon/10})$, which already worked very well in practice for the range of parameter values we considered.

H Complementary Experimental Analysis

This section extends our experimental analysis beyond Section 6, examining: (1) performance across problem instances of varying difficulty, (2) stopping time distributions, (3) effects of target error probabilities and privacy parameters, and (4) a tuned version of our DP-SPRT with improved sample complexity.

We evaluate our methods across three problem instances of varying difficulty:

- **Instance 1 (Standard):** $p_0 = 0.3$, $p_1 = 0.7$ - Well-separated parameters (as in the main paper)
- **Instance 2 (Close):** $p_0 = 0.45$, $p_1 = 0.55$ - Parameters close to each other, requiring more samples
- **Instance 3 (Boundary):** $p_0 = 0.05$, $p_1 = 0.25$ - One parameter near boundary, affecting observation variance

H.1 Comparative Performance Across Problem Difficulty

Figure 2 presents the sample complexity across the three problem instances.

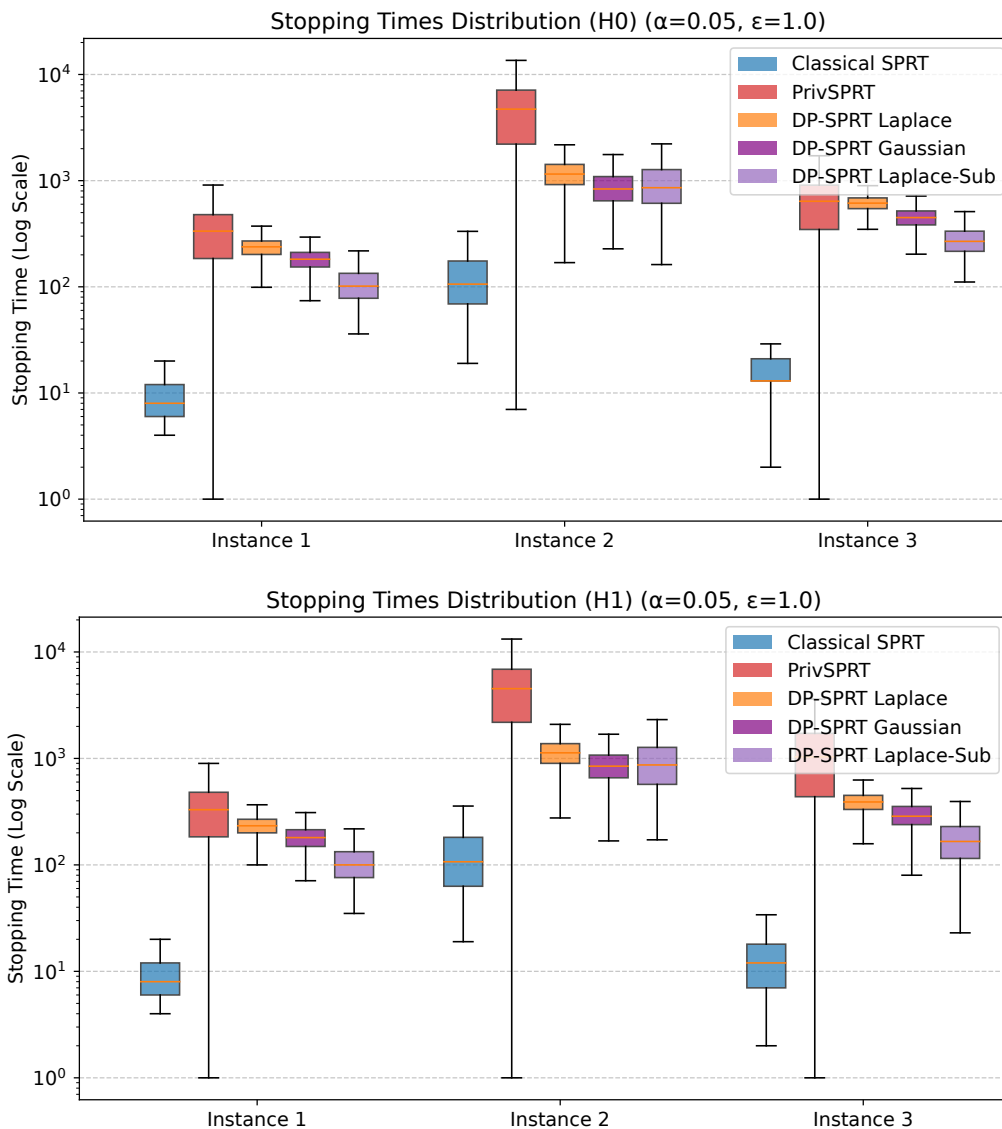


Figure 2: Sample complexity for different methods across problem instances with $\alpha = \beta = 0.05$ and $\epsilon = 1$. Top: Under \mathcal{H}_0 . Bottom: Under \mathcal{H}_1 .

As difficulty increases from Instance 1 to Instance 2, all methods require more samples. The advantage of our DP-SPRT variants over PrivSPRT becomes more pronounced in more challenging instances, with subsampling showing particular benefits when parameters are close (Instance 2).

Instance 3 exhibits asymmetry due to p_0 being close to zero, with decisions under \mathcal{H}_0 requiring fewer samples than under \mathcal{H}_1 for PrivSPRT. For DP-SPRT, we observe the opposite effect, although it is to a lesser extent. Our methods maintain their advantage in both scenarios, more prominently under \mathcal{H}_1 .

H.2 Effect of Target Error Probability

Figure 3 shows how sample complexity varies with different target error probabilities.

Stricter error requirements (smaller α and β) predictably increase sample complexity for all methods. At

$\alpha = \beta = 0.01$, we observe a wider gap between private and non-private methods.

H.3 Privacy-Sample Complexity Trade-off

Figure 4 demonstrates the relationship between privacy parameters and sample complexity.

The privacy-utility trade-off is evident as ϵ varies from 0.1 (high privacy) to 5 (relaxed privacy). At $\epsilon = 0.1$, all private methods require substantially more samples than non-private SPRT, with DP-SPRT with subsampling maintaining a clear advantage. As privacy constraints relax with increasing ϵ , all DP-SPRT variants converge toward classical SPRT performance.

PrivSPRT exhibits higher variance in stopping time, particularly at lower ϵ values. Our DP-SPRT variants show more consistent performance across trials, beneficial for applications requiring predictable resource allocation.

H.4 Tuned Correction Function and Distribution of Stopping Times

Our theoretical correction function $C(n, \delta)$ guarantees desired error probabilities but may be conservative. We explored a tuned version by scaling the function by the factor that maintains error probability targets, estimated through simulation.

The tuning process involves finding a scaling factor $\kappa \in (0, 1)$ to replace $C(n, \delta)$ with $\kappa \cdot C(n, \delta)$. For Instance 1 with $\epsilon = 1$ and $\alpha = \beta = 0.1$, we found $\kappa \approx 0.5$ maintains error probabilities below target. However, like PrivSPRT’s threshold, this value must be recomputed for new parameters.

Figure 5 compares stopping time distributions. The tuned version of DP-SPRT significantly reduces the number of samples compared to our standard version, however, it does not have the same theoretical backing. PrivSPRT exhibits both a larger mean and more spread out distribution, indicating more frequent occurrences of very long stopping times.

This visualization highlights not just average performance differences but also consistency and predictability. The lower variance of our methods represents a practical advantage for resource planning. The tuning approach offers a middle ground, leveraging theoretical guarantees while improving empirical performance, though like PrivSPRT, it requires empirical tuning for each specific setting.

H.5 Summary of Findings

Our complementary experiments reveal several important insights:

1. The advantage of our DP-SPRT methods becomes more pronounced as problem difficulty increases, making them particularly valuable for challenging testing scenarios with parameters that are close to each other.
2. All DP-SPRT variants exhibit lower variance in stopping times compared to PrivSPRT, leading to more predictable performance.
3. In asymmetric scenarios (e.g., Instance 3 with parameters near boundaries), decisions under different hypotheses require substantially different sample sizes for PrivSPRT. Our method is less sensitive to this asymmetry.
4. The subsampling approach provides substantial benefits in high-privacy regimes ($\epsilon = 0.1$).
5. Our theoretical correction function is conservative, and a tuned version scaling the function by a factor $\kappa \approx 0.5$ can significantly improve sample complexity while maintaining error guarantees.

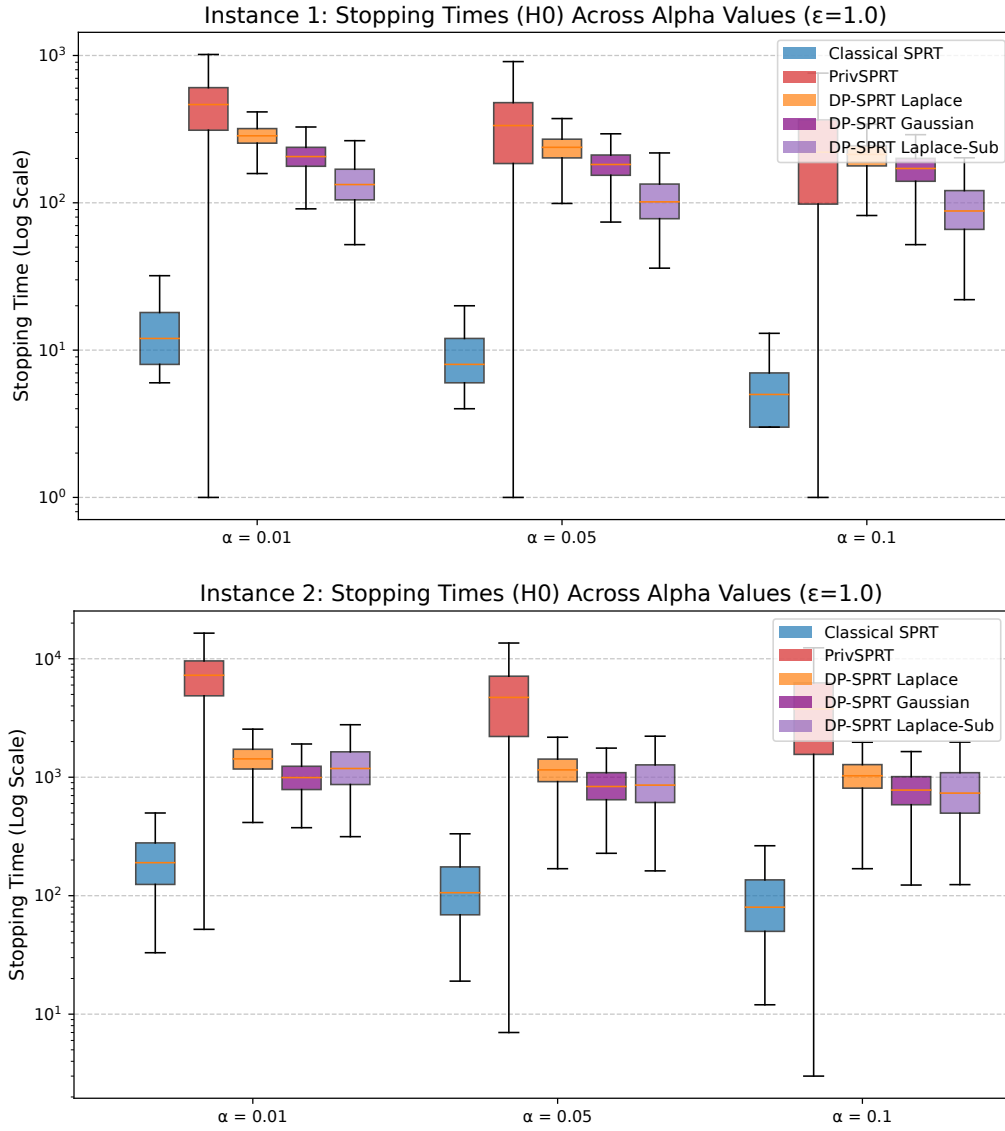


Figure 3: Sample complexity for different methods with varying target error probabilities at $\epsilon = 1$ and $\alpha = \beta \in \{0.01, 0.05, 0.1\}$. Top: Instance 1 under \mathcal{H}_0 . Bottom: Instance 2 under \mathcal{H}_0 .

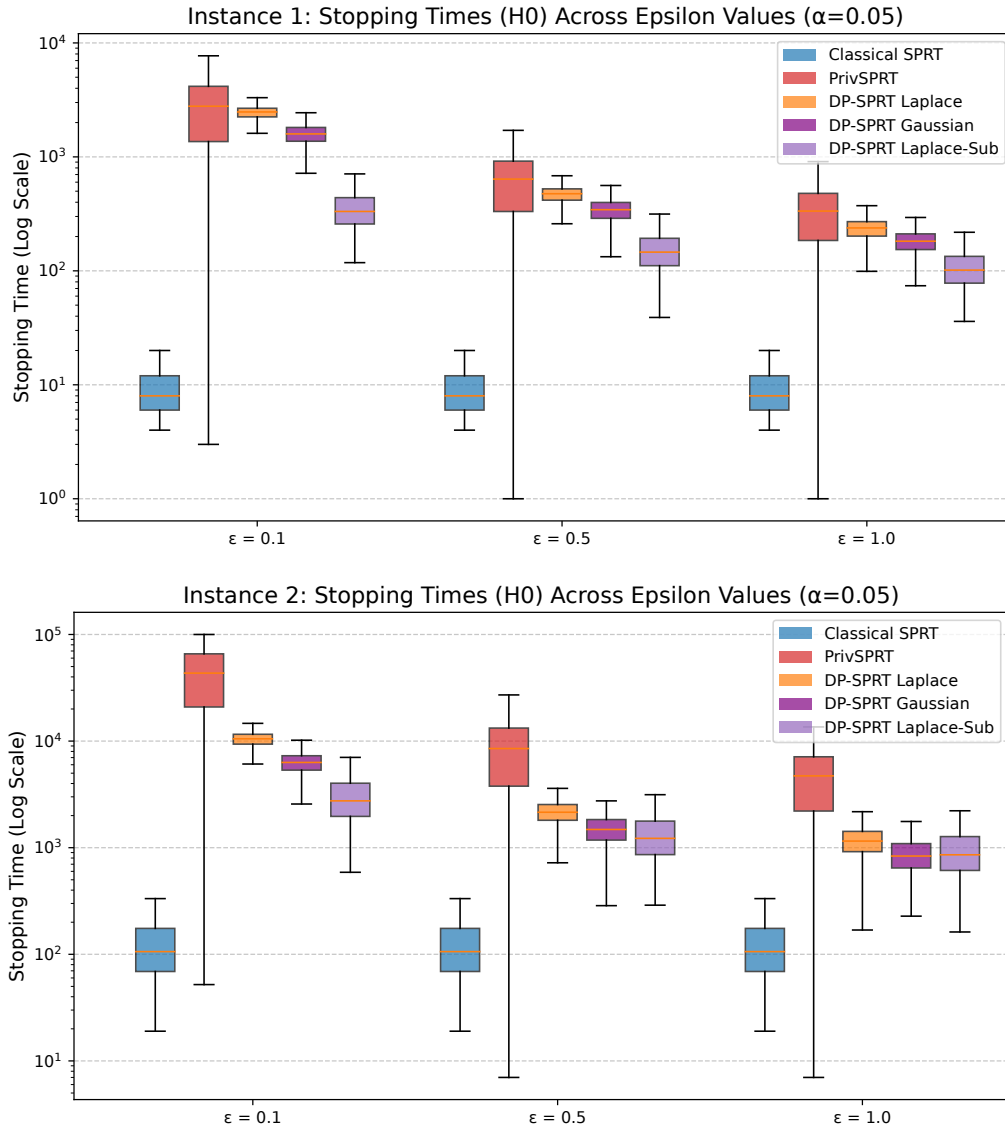


Figure 4: Sample complexity for different privacy parameters with $\alpha = \beta = 0.05$ and $\epsilon \in \{0.1, 1, 5\}$. Top: Instance 1 under \mathcal{H}_0 . Bottom: Instance 2 under \mathcal{H}_0 .

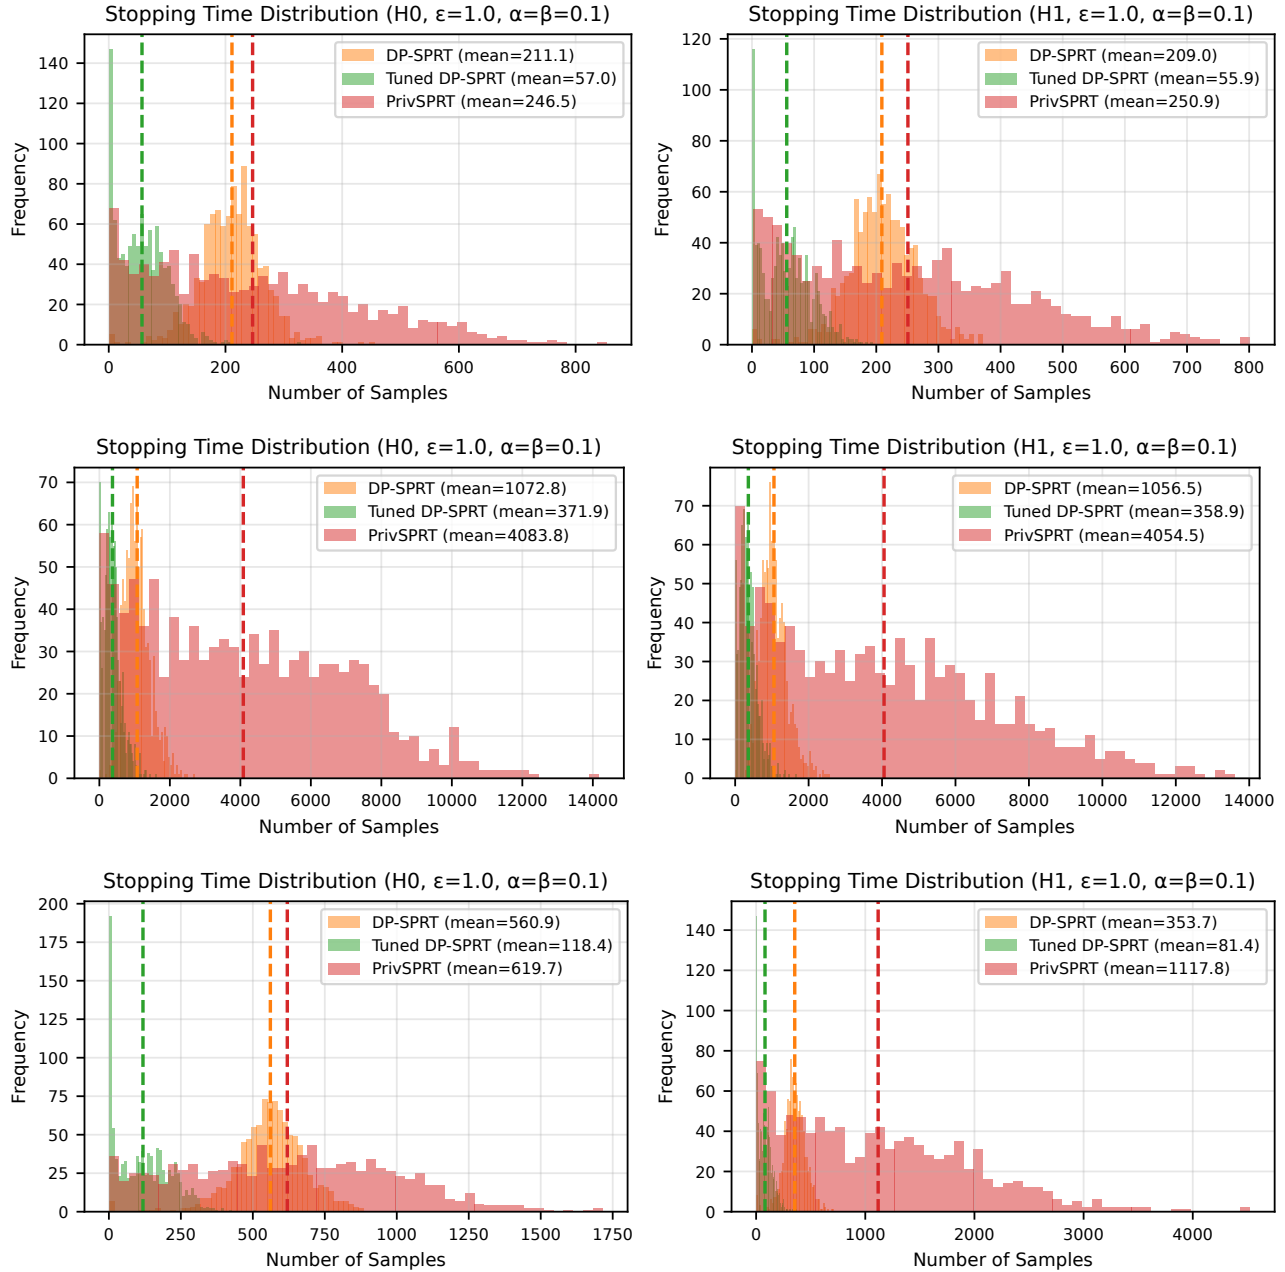


Figure 5: Histogram of stopping times for DP-SPRT with Laplace noise, tuned DP-SPRT, and PrivSPRT with $\alpha = \beta = 0.1$ and $\epsilon = 1$. Each row shows results under \mathcal{H}_0 (left) and \mathcal{H}_1 (right) for different instances. Top row: Instance 1. Middle row: Instance 2. Bottom row: Instance 3. Dashed lines indicate mean stopping times.