

MIT Open Access Articles

Optimal stopping rules for sequential hypothesis testing

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Daskalakis, C and Kawase, Y. 2017. "Optimal stopping rules for sequential hypothesis testing." Leibniz International Proceedings in Informatics, LIPIcs, 87.

As Published: 10.4230/LIPIcs.ESA.2017.32

Persistent URL: <https://hdl.handle.net/1721.1/143459>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 3.0 unported license



Optimal Stopping Rules for Sequential Hypothesis Testing

Constantinos Daskalakis^{*1} and Yasushi Kawase^{†2}

- 1 Massachusetts Institute of Technology, Massachusetts, USA
costis@mit.edu
- 2 Tokyo Institute of Technology, Tokyo, Japan, and
RIKEN AIP Center, Tokyo, Japan
kawase.y.ab@m.titech.ac.jp

Abstract

Suppose that we are given sample access to an unknown distribution p over n elements and an explicit distribution q over the same n elements. We would like to reject the null hypothesis “ $p = q$ ” after seeing *as few samples as possible*, when $p \neq q$, while we never want to reject the null, when $p = q$. Well-known results show that $\Theta(\sqrt{n}/\epsilon^2)$ samples are necessary and sufficient for distinguishing whether p equals q versus p is ϵ -far from q in total variation distance. However, this requires the distinguishing radius ϵ to be fixed prior to deciding how many samples to request. Our goal is instead to design *sequential hypothesis testers*, i.e. online algorithms that request i.i.d. samples from p and stop as soon as they can confidently reject the hypothesis $p = q$, without being given a lower bound on the distance between p and q , when $p \neq q$. In particular, we want to minimize the number of samples requested by our tests as a function of the distance between p and q , and if $p = q$ we want the algorithm, with high probability, to never reject the null. Our work is motivated by and addresses the practical challenge of sequential A/B testing in Statistics.

We show that, when $n = 2$, any sequential hypothesis test must see $\Omega\left(\frac{1}{d_{\text{tv}}(p,q)^2} \log \log \frac{1}{d_{\text{tv}}(p,q)}\right)$ samples, with high (constant) probability, before it rejects $p = q$, where $d_{\text{tv}}(p, q)$ is the—unknown to the tester—total variation distance between p and q . We match the dependence of this lower bound on $d_{\text{tv}}(p, q)$ by proposing a sequential tester that rejects $p = q$ from at most $O\left(\frac{\sqrt{n}}{d_{\text{tv}}(p,q)^2} \log \log \frac{1}{d_{\text{tv}}(p,q)}\right)$ samples with high (constant) probability. The $\Omega(\sqrt{n})$ dependence on the support size n is also known to be necessary. We similarly provide two-sample sequential hypothesis testers, when sample access is given to both p and q , and discuss applications to sequential A/B testing.

1998 ACM Subject Classification F.2.2 Computations on discrete structures, G.3 Probability and Statistics

Keywords and phrases property testing, sequential hypothesis testing, A/B testing

Digital Object Identifier 10.4230/LIPIcs.ESA.2017.32

1 Introduction

A central problem in Statistics is testing how well observations of a stochastic phenomenon conform to a statistical hypothesis. A common scenario involves access to i.i.d. samples from an unknown distribution p over some set Σ and a hypothesis distribution q over the same

* Supported by NSF Awards CCF-1617730 and CCF-1650733 and ONR Grant N00014-12-1-0999.

† Supported by JSPS KAKENHI Grant Number JP16K16005.



set. The goal is to distinguish between $p = q$ and $p \neq q$. This problem, in myriads of forms, has been studied since the very beginnings of the field. Much of the focus has been on the asymptotic analysis of tests in terms the error exponents of their type I or type II errors.

More recently, the problem received attention from property testing, with emphasis on the finite sample regime. A formulation of the problem that is amenable to finite sample analysis is the following: given sample access to p and a hypothesis q as above, together with some $\epsilon > 0$, how many samples are needed to distinguish, correctly with probability at least $2/3$,¹ between $p = q$ and $d(p, q) > \epsilon$, for some distance of interest d ? For several distances d , we know tight answers on the number of samples required. For instance, when we take d to be the total variation distance, d_{tv} ,² we know that $\Theta(\sqrt{n}/\epsilon^2)$ samples are necessary and sufficient, where $n = |\Sigma|$ [7, 27, 33]. Tight answers are also known for other distances, variants of this problem, and generalizations [15, 8, 32, 9, 29, 13, 2, 1, 10, 12, 11], but our focus will be on distinguishing the identity of p and q under total variation distance.

While the existing literature gives tight upper and lower bounds for this problem, it still requires a lower bound ϵ on the distance between p and q when they differ, aiming for that level of distinguishing accuracy, when choosing the sample size. This has two implications:

1. Even when p and q are blatantly far from each other, the test will still request $\Theta(\sqrt{n}/\epsilon^2)$ samples, as the distance of p and q is unknown to the test when the sample size is determined.
2. When $p \neq q$, but $d_{\text{tv}}(p, q) \leq O(\epsilon)$, there are no guarantees about the output of the test, just because the sample is not big enough to confidently decide that $p \neq q$.

Both issues above are intimately related to the fact that these tests predetermine the number of samples to request, as a function of the support size n and the desired distinguishing radius ϵ .

In practice, however, samples are costly to acquire. Even, when they are in abundance, they may be difficult to process. As a result, it is a common practice in clinical trials or online experimentation to “peek” at the data before an experiment is completed, in the hopes that significant evidence is collected supporting or rejecting the hypothesis. Done incorrectly this may induce statistical biases invalidating the reported significance and power bounds of the experiment [25].

Starting with a demand for more efficient testing during World War II, there has been a stream of work in Statistics addressing the challenges of *sequential hypothesis testing*; see, e.g., [34, 35, 28, 30, 26, 24, 36, 23, 22, 6, 31, 18, 20, 3, 5] and their references. These methods include the classical *sequential probability ratio test* (SPRT) [34, 35] and its generalizations [22, 6], where the alternative hypothesis is either known exactly or is parametric (i.e. p either equals q , or p is different than q , but belongs in the same parametric class as q). An alternative to SPRT methods are methods performing *repeated significance tests* (RST) [28, 26, 24, 36, 18, 5]. These methods target scalar distributions and either make parametric assumptions about p and q (e.g. Bernoulli, Gaussian, or Exponential family assumptions), or compare moments of p and q (usually their means). In particular these methods are closely related to the task of choosing the best arm in bandit settings; see, e.g., [17] and its references.

In contrast to the existing literature, we want to study categorical random variables, and do not want to make any parametric assumptions about p and q . In particular, we do not

¹ As usual the probability “ $2/3$ ” in the definition of the problem can be boosted to any constant $1 - \delta$, at a cost of an extra $\log \frac{1}{\delta}$ factor in the sample complexity.

² See Section 2 for a definition.

want to make any assumptions about the alternatives. If the null hypothesis, $p = q$, fails, we do not know *how* it will fail. We are simply interested in determining whether $p = q$ or $p \neq q$, as soon as possible. Our goal is to devise an online policy such that, given any sequence of samples from p , the policy decides to

- (i) either continue, drawing another sample from p ; or
- (ii) stop and declare $p \neq q$.

We want that our policy:

1. has *small error rate*, i.e. for some user-specified constants $\alpha, \beta > 0$,
 - a. If $p = q$, the policy will stop, with probability at most α ; i.e. the type I error is α .
 - b. If $p \neq q$, the policy will stop (i.e. declare $p \neq q$), with probability at least $1 - \beta$; i.e. the type II error is β .
2. draws as few samples as possible, when $p \neq q$, in the event that it stops (which happens with probability at least $1 - \beta$).

In other words, we want to define a *stopping rule* such that, for as small a function $k = k(n, \cdot)$ as possible, the stopping time τ satisfies:

- (i) $\Pr_q[\tau = +\infty] \geq 1 - \alpha$, and
- (ii) $\Pr_p[\tau < k(n, d_{tv}(p, q))] \geq 1 - \beta$ for all $p \neq q$,

where n is the cardinality of the set Σ on which p and q are supported. Henceforth we will call a stopping rule *proper* if it satisfies property (i) above. We want to design proper stopping rules that satisfy (ii) for as small a function $k(\cdot, \cdot)$ as possible. That is, with probability at least $1 - \beta$, we want to reject the hypothesis “ $p = q$ ” as soon as possible. As we focus on the dependence of our stopping times on n and $d_{tv}(p, q)$, we only state and prove our results throughout this paper for $\alpha = \beta = 1/3$. Changing α and β to different constants will only change the constants in our bounds. Our results are the following:³

1. In Theorem 1, we show that, when $n = 2$, i.e. when p and q are Bernoulli, and even when q is uniform, there is no proper stopping rule such that $k(2, d_{tv}(p, q)) < \frac{1}{16d_{tv}(p, q)^2} \log \log \frac{1}{d_{tv}(p, q)}$.^{4,5} Our lower bound is reminiscent of the lower bound on the number of samples needed to identify the best of two arms in a bandit setting, proven in [17]. This was shown by an application of an information theoretic lower bound of Farrel for distinguishing whether an exponential family has positive or negative mean [14]. Farrel lower bounds the expected number of observations that are needed, while we show that not even a constant probability of stopping below our bound can be achieved. This is a weaker target, hence the lower bound is stronger. Finally, our goal is even weaker as we only want to determine whether $p \neq q$, but not to identify the Bernoulli with the highest mean. Our proof is combinatorial and concise.
2. In Theorem 2, we construct, for any q and n , a proper stopping rule satisfying $k(n, d_{tv}(p, q)) < \frac{c\sqrt{n}}{d_{tv}(p, q)^2} \log \log \frac{1}{d_{tv}(p, q)}$, for some constant c . By Theorem 1 the dependence of this bound on $d_{tv}(p, q)$ is optimal. Moreover, it follows from standard testing lower bounds, that the dependence on n is also optimal.⁶ In fact Theorem 2 achieves something stronger. It shows that, whenever $p \neq q$, with probability at least $2/3$, the stopping rule will actually stop later than $\Omega\left(\frac{\sqrt{n}}{\chi^2(p, q)} \log \log \frac{1}{\chi^2(p, q)}\right)$ and prior to $O\left(\frac{\sqrt{n}}{d_{tv}(p, q)^2} \log \log \frac{1}{d_{tv}(p, q)}\right)$.

³ The formal statements of our theorems are given in the notation introduced in Section 1.1.

⁴ Note that for Bernoulli's $d_{tv}(p, q)$ equals the difference of their means.

⁵ Throughout the paper, we assume that \log means logarithm to the base e .

⁶ In particular, as we have already noted, it is known that $\Omega(\sqrt{n}/\epsilon^2)$ samples are necessary to distinguish $p = q$ from $d_{tv}(p, q) > \epsilon$, for small enough constant ϵ . As our task is harder than distinguishing whether $p \neq q$ for a fixed radius of accuracy ϵ , we need to pay at least this many samples.

3. In Theorem 3, we study *two-sample sequential hypothesis testing*, where we are given sample access to both distributions p and q . Similarly to the one-sample case, our goal is to devise a stopping rule that is proper, i.e. when $p = q$, it does not stop with probability at least $2/3$, while also minimizing the samples it takes to determine that $p \neq q$. That is, when $p \neq q$, it stops, with probability at least $2/3$, after having seen as few samples as possible. We show that there is a proper stopping rule which, whenever $p \neq q$, stops after having seen $\Theta\left(\frac{n/\log n}{d_{\text{tv}}(p,q)^2} \log \log \frac{1}{d_{\text{tv}}(p,q)}\right)$ samples, with probability at least $2/3$. The dependence on $d_{\text{tv}}(p,q)$ is optimal from Theorem 1. As our tight upper and lower bounds on the number of samples allow us to estimate $d_{\text{tv}}(p,q)$ to within a constant factor, the lower bounds of [32] for estimating the distance between distributions imply that the dependence of our bounds on n is also optimal.
4. The dependence of our upper bounds on $d_{\text{tv}}(p,q)$ is reminiscent of recent work in the bandit literature [19, 17] and sequential non-parametric testing [5], where stopping times with *iterated log complexity* have appeared. These results are intimately related to the Law of Iterated Logarithm [21, 4]. Our results are instead obtained in a self-contained and purely combinatorial fashion. Moreover, as discussed earlier, our testing goals are different than those in these works. While both works study scalar distributions, distinguishing them in terms of their means, we study categorical random variables distinguishing them in terms of their total variation distance.

1.1 Model

Let p, q be discrete distributions over $\Sigma = [n]$, where $[n] = \{0, 1, \dots, n-1\}$. We assume that $n \geq 2$. In the *one-sample sequential hypothesis testing problem*, distributions q and sample access is provided to distribution p . Our goal is to distinguish between $p = q$ and $p \neq q$. Since p and q could be arbitrarily close even when they differ, our goal is to reject hypothesis $p = q$ as soon as possible when $p \neq q$, as explained below.

Let $[n]^*$ be the *Kleene star* of $[n]$, i.e., the set of all strings of finite length consisting of symbols in $[n]$. A function $T : [n]^* \rightarrow \{0, 1\}$ is called a *stopping rule* if $T(x_1 \dots x_k) = 1$ implies $T(x_1 \dots x_k x_{k+1} \dots x_{k+\ell}) = 1$ for any integers $k, \ell \geq 0$ and $x_i \in [n]$ ($i = 1, \dots, k + \ell$). For all sequences $x \in \{0, 1\}^*$, $T(x) = 1$ and $T(x) = 0$ mean respectively that the rule rejects hypothesis $p = q$ or it continues testing, after having seen x . For an infinite sequence $x = (x_1 x_2 \dots) \in [n]^{\mathbb{N}}$, we define the *stopping time* to be the $\min\{t \mid T(x_1 \dots x_t) = 1\}$. Let $N(a \mid x)$ be the number of times symbol $a \in [n]$ occurs in the sequence $x \in [n]^*$. Let $\tau(T, p)$ be a random variable that represents the stopping time when the sequence is generated by p , i.e. for all k :

$$\Pr[\tau(T, p) \leq k] = \sum_{x \in [n]^k} \left(T(x) \prod_{i=1}^k p_{x_i} \right) = \sum_{x \in [n]^k} \left(T(x) \prod_{i \in [n]} p_i^{N(i|x)} \right).$$

With the above notation, our goal in the one-sample sequential hypothesis testing problem is to find, for a given distribution q , a stopping rule T such that

- (a) $\Pr[\tau(T, q) \leq k] \leq 1/3$ for any k , and
- (b) $\Pr[\tau(T, p) \leq k] \geq 2/3$ for k as small as possible whenever $p \neq q$.

We call a stopping rule *proper* if it satisfies the condition (a).⁷

We also consider the *two-sample sequential hypothesis testing problem* where p and q are both unknown distributions over $[n]$, and sample access is given to both. For simplicity, this paper only studies stopping rules that use the same number of samples from each

⁷ As noted earlier there is nothing special with the constants “1/3” and “2/3” here. We could turn these to any constants α and $1 - \beta$ respectively at a cost of a constant factor in our sample complexity.

distribution. This assumption increases the sample complexity by a factor of at most 2. Then a stopping rule T is defined as a function from $\bigcup_{k \in \mathbb{N}} ([n]^k \times [n]^k)$ to $\{0, 1\}$ such that $T(x, y) = 1$ implies $T(xz, yw) = 1$ for any strings $x, y, z, w \in [n]^*$ with $|x| = |y|$ and $|z| = |w|$. Here, $|x|$ represents the length of $x \in [n]^*$. The stopping time for infinite sequences $x, y \in [n]^{\mathbb{N}}$ is given by $\min\{t \mid T(x_1 \cdots x_t, y_1 \cdots y_t) = 1\}$. Also, the stopping time $\tau(T, p, q)$ is a random variable such that

$$\Pr[\tau(T, p, q) \leq k] = \sum_{x \in [n]^k} \sum_{y \in [n]^k} \left(T(x, y) \prod_{i \in [n]} p_i^{N(i|x)} \prod_{j \in [n]} q_j^{N(j|y)} \right).$$

Now, our task is to find a stopping rule T such that

- (1) $\Pr[\tau(T, p, q) \leq k] \leq 1/3$ for any k whenever $p = q$, and
- (2) $\Pr[\tau(T, p, q) \leq k] \geq 2/3$ for k as small as possible whenever $p \neq q$.

Before describing our results, we briefly review notations and definitions used in the results. The *total variation distance* between p and q , denoted by $d_{\text{tv}}(p, q)$, is defined to be $d_{\text{tv}}(p, q) = \frac{1}{2} \sum_{i \in [n]} |p_i - q_i| = \frac{1}{2} \|p - q\|_1$. The χ^2 -distance between p and q (which is not a true distance) is given by $\chi^2(p, q) = \sum_{i \in [n]} (p_i - q_i)^2 / q_i = \left(\sum_{i \in [n]} p_i^2 / q_i \right) - 1$. Note that these two distances satisfy $d_{\text{tv}}(p, q)^2 \leq \frac{1}{4} \chi^2(p, q)$ for any distributions p, q by Cauchy–Schwartz inequality.

1.2 Our results

We first prove that any proper stopping rule for the one-sample sequential hypothesis testing problem, must see $\frac{1}{16 \cdot d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)}$ samples before it stops, even when $n = 2$, i.e. both distributions are Bernoulli, and the known distribution q is Bernoulli(0.5).

► **Theorem 1 (One-Sample Sequential Hypothesis Testing Lower Bound).** *Even when $n = 2$ and $q = (1/2, 1/2)$, there exist no proper stopping rule T and positive real ϵ_0 such that*

$$\Pr \left[\tau(T, p) \leq \frac{1}{16 \cdot d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)} \right] \geq 2/3 \quad (\text{whenever } 0 < d_{\text{tv}}(p, q) < \epsilon_0). \quad (1)$$

Here, we remark that $d_{\text{tv}}(p, q) = |1/2 - p_0| = |1/2 - p_1|$.

As we noted earlier, our lower bound involving the iterated logarithm appears similar to that of Farrel [14], but it is a slightly stronger statement. More precisely, he proved that $\limsup_{p \rightarrow q} \frac{d_{\text{tv}}(p, q)^2 \cdot \mathbb{E}[\tau(T, p)]}{\log \log \frac{1}{d_{\text{tv}}(p, q)}} \geq c$ for a certain positive constant c . Theorem 1 implies the result but not vice versa. Also, our proof is elementary and purely combinatorial. It is given in Section 3.

We next provide a black-box reduction, obtaining optimal sequential hypothesis testers from “robust” non-sequential hypothesis testers. In particular, we use algorithms for *robust identity testing* where the goal is, given some accuracy ϵ , to distinguish whether p and q are $O(\epsilon)$ -close in some distance versus $\Omega(\epsilon)$ -far in some (potentially) different distance [32, 1]. We propose a schedule for repeated significance tests, which perform robust identity testing with different levels of accuracy ϵ , ultimately compounding to optimal sequential testers. In the inductive step, given the current value of ϵ , we run the non-sequential test with accuracy ϵ for $\Theta(\log \log 1/\epsilon)$ times, and take the majority vote. If the majority votes ϵ -far, we stop the procedure. Otherwise, we decrease ϵ geometrically and continue. The accuracy improvement by the $\Theta(\log \log 1/\epsilon)$ -fold repetition allows the resulting stopping rule to be proper.

Our theorems for one-sample and two-sample sequential hypothesis testing are stated below and proven in Section 4. As noted earlier, stopping times involving the iterated logarithm have appeared in the multi-armed bandit and sequential hypothesis testing literature. As

32:6 Optimal Stopping Rules for Sequential Hypothesis Testing

explained, our testing goals are different than those in this prior work. While they study scalar distributions, distinguishing them in terms of their means, we study categorical random variables distinguishing them in terms of their total variation distance. Moreover, our results do not appeal to the law of the iterated logarithm and are obtained in a purely combinatorial fashion, using of course prior work on property testing.

► **Theorem 2 (One-Sample Sequential Hypothesis Testing Upper Bound).** *For any known distribution q over $[n]$, there exists a proper stopping rule T and positive reals ϵ_0 and c such that*

$$\Pr \left[\frac{\sqrt{n}}{c \cdot \chi^2(p, q)} \log \log \frac{1}{\chi^2(p, q)} \leq \tau(T, p) \leq \frac{c\sqrt{n}}{d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)} \right] \geq 2/3 \quad (2)$$

holds for any p satisfying $0 < \chi^2(p, q) < \epsilon_0$. Note that $0 < d_{\text{tv}}(p, q) < \sqrt{\epsilon_0}/2$ holds when $0 < \chi^2(p, q) < \epsilon_0$ since $d_{\text{tv}}(p, q)^2 \leq \frac{1}{4}\chi^2(p, q)$.

► **Theorem 3 (Two-Sample Sequential Hypothesis Testing Upper Bound).** *There exists a proper stopping rule T and positive reals ϵ_0 and c such that*

$$\Pr \left[\frac{n/\log n}{c \cdot d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)} \leq \tau(T, p, q) \leq \frac{c \cdot n/\log n}{d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)} \right] \geq 2/3 \quad (3)$$

holds for any unknown distributions p, q over $[n]$ satisfying $0 < d_{\text{tv}}(p, q) < \epsilon_0$.

Since the lower bounds on the stopping time in both (2) and (3) go to infinity as p goes to q , the stopping rules never stop with probability at least $2/3$ when $p = q$. Hence, the stopping rules are proper. As noted earlier, we can improve the confidence from $2/3$ to $1 - \delta$ at the cost of a multiplicative factor $\log(1/\delta)$ in the sample complexity. The dependence of both upper bounds on $d_{\text{tv}}(p, q)$ is tight as per Theorem 1. The \sqrt{n} dependence in Theorem 2 is tight because it is known that testing whether $d_{\text{tv}}(p, q) = 0$ or $d_{\text{tv}}(p, q) \geq 1/2$ requires $\Omega(\sqrt{n})$ samples [15, 8]. In addition, Theorem 3, allows us to estimate the total variation distance between p and q because the stopping time and the total variation distance satisfy the relation $\tau(T, p) = \Theta\left(\frac{n/\log n}{d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)}\right)$. This and the lower bounds for estimating the ℓ_1 distance of distributions provided in [32], imply that the dependence of Theorem 3 on n is also optimal.

As a simple corollary of the above results, we can also provide an efficient algorithm for sequential A/B testing, replicating the bounds obtainable from [19, 17, 5], without appealing to the Law of the Iterated Logarithm.

► **Theorem 4.** *There exists an algorithm that distinguishes between the cases (a) $p > q$ and (b) $q > p$, using $\Theta\left(\frac{1}{|p-q|^2} \log \log \frac{1}{|p-q|}\right)$ samples for any unknown Bernoulli distributions with success probabilities p and q .*

2 Known Results

In this section, we state known results for robust identity testing, which we use in our upper bounds.

► **Theorem 5 ([2]).** *For any known distribution q , there exists an algorithm with sample complexity $\Theta(\sqrt{n}/\epsilon^2)$ which distinguishes between the cases*

(a) $\sqrt{\chi^2(p, q)} \leq \epsilon/2$ and

(b) $d_{\text{tv}}(p, q) \geq \epsilon$,

with probability at least $2/3$.

► **Theorem 6** ([32]). *Given sample access to two unknown distributions p and q , there exists an algorithm with sample complexity $\Theta(\frac{n}{\epsilon^2 \log n})$ which distinguishes between the cases*

(a) $d_{\text{tv}}(p, q) \leq \epsilon/2$ and

(b) $d_{\text{tv}}(p, q) \geq \epsilon$,

with probability at least $2/3$.

We remark that, even though the proofs of Theorems 5 and 6 may use Poisson sampling, i.e., the sample complexities are Poisson distributed, we can assume that the numbers of samples are deterministically chosen. This is because the Poisson distribution is sharply concentrated around the expected value.

In our analysis of the upper and lower bounds of sample complexities, we use the following Hoeffding's inequality.

► **Theorem 7** (Hoeffding's inequality [16]). *Let X be a binomial distribution with n trials and probability of success p . Then, for any real ϵ , we have $\Pr[X \leq (p - \epsilon)n] = \sum_{i=0}^{\lfloor (p-\epsilon)n \rfloor} \binom{n}{i} p^i (1-p)^{n-i} \leq \exp(-2\epsilon^2 n)$.*

3 Lower bound

In this section, we prove Theorem 1, i.e., our lower bound on the sample complexity for the binary alphabet case $n = 2$. We abuse notation using p, q to denote the probabilities that our distributions output 1. In particular, $1 - p$ and $1 - q$ are the probabilities they output 0.

We first observe that, for any stopping rule T , the stopping times $\tau(T, p)$ and $\tau(T, q)$ take similar values when p, q are close.

► **Lemma 8.** *Let $p < 1/2$, $q = 1/2$, $1 > \alpha > 0$ and s, t be positive integers such that $s > t$. If $\Pr[t \leq \tau(T, p) \leq s] \geq \alpha$, then we have*

$$\Pr[t \leq \tau(T, q) \leq s] \geq (\alpha - \alpha^2) \cdot (1/e)^{4(\frac{1}{2}-p)^2 \cdot s + 4(\frac{1}{2}-p)\sqrt{s \log(1/\alpha)}}.$$

Proof. Let $A = \{x \in \{0, 1\}^s \mid T(x_1 \dots x_{t-1}) = 0 \text{ and } T(x_1 \dots x_s) = 1\}$. Then the stopping probability for p can be written as

$$\Pr[t \leq \tau(T, p) \leq s] = \sum_{x \in A} p^{N(1|x)} (1-p)^{N(0|x)}.$$

Recall that $N(a | x)$ is the number of times a symbol $a \in \{0, 1\}$ occurs in a string $x \in \{0, 1\}^s$. Note that $|x| = N(1 | x) + N(0 | x)$. Let $A_1 = \{x \in A \mid N(1 | x) < p \cdot s - \sqrt{s \log(1/\alpha)}\}$ and $A_2 = \{x \in A \mid N(1 | x) \geq p \cdot s - \sqrt{s \log(1/\alpha)}\}$. By using Hoeffding's inequality, we have

$$\begin{aligned} \sum_{x \in A_1} p^{N(1|x)} (1-p)^{N(0|x)} &\leq \sum_{x \in \{0, 1\}^s: N(1|x) < p \cdot s - \sqrt{s \log(1/\alpha)}} p^{N(1|x)} (1-p)^{N(0|x)} \\ &\leq \sum_{k=0}^{\lfloor p \cdot s - \sqrt{s \log(1/\alpha)} \rfloor} \binom{s}{k} p^k (1-p)^{s-k} \leq \exp\left(-2 \left(\frac{\sqrt{s \log(1/\alpha)}}{s}\right)^2 \cdot s\right) = \alpha^2. \end{aligned}$$

Hence, it holds that

$$\sum_{x \in A_2} p^{N(1|x)} (1-p)^{N(0|x)} = \sum_{x \in A \setminus A_1} p^{N(1|x)} (1-p)^{N(0|x)} \geq \alpha - \alpha^2. \quad (4)$$

32:8 Optimal Stopping Rules for Sequential Hypothesis Testing

In what follows, we bound the value $\Pr[t \leq \tau(T, q) \leq s]$. Since $A_2 \subseteq A$ and $s = N(1 | x) + N(0 | x)$, we have

$$\Pr[t \leq \tau(T, q) \leq s] = \sum_{x \in A} \frac{1}{2^s} \geq \sum_{x \in A_2} \frac{1}{2^s} = \sum_{x \in A_2} \frac{1}{4^{N(1|x)}} \cdot \frac{1}{2^{N(0|x) - N(1|x)}}.$$

Since $p(1-p) = -(p-1/2)^2 + 1/4 \leq 1/4$, it holds that

$$\begin{aligned} \sum_{x \in A_2} \frac{1}{4^{N(1|x)}} \cdot \frac{1}{2^{N(0|x) - N(1|x)}} &\geq \sum_{x \in A_2} p^{N(1|x)} (1-p)^{N(1|x)} \cdot \left(\frac{1}{2}\right)^{N(0|x) - N(1|x)} \\ &= \sum_{x \in A_2} p^{N(1|x)} (1-p)^{N(0|x)} \cdot \left(\frac{1/2}{1-p}\right)^{N(0|x) - N(1|x)}. \end{aligned} \quad (5)$$

Note that, for $x \in A_2$, we have $N(0 | x) - N(1 | x) = s - 2N(1 | x) \leq s - 2(p \cdot s - \sqrt{s \log(1/\alpha)}) = 2(1/2 - p)s + 2\sqrt{s \log(1/\alpha)}$ since $s = N(1 | x) + N(0 | x)$ and $N(1 | x) \geq p \cdot s - \sqrt{s \log(1/\alpha)}$. Also, we have $\frac{1/2}{1-p} = \frac{1}{1+(1-2p)} < 1$ since $p < 1/2$. Thus, we get

$$\left(\frac{1/2}{1-p}\right)^{N(0|x) - N(1|x)} \geq \left(\frac{1}{1+(1-2p)}\right)^{2(1/2-p)s + 2\sqrt{s \log(1/\alpha)}}. \quad (6)$$

Applying (6) and (4) to (5) yields

$$\begin{aligned} \sum_{x \in A_2} p^{N(1|x)} (1-p)^{N(0|x)} \cdot \left(\frac{1/2}{1-p}\right)^{N(0|x) - N(1|x)} \\ \geq \sum_{x \in A_2} p^{N(1|x)} (1-p)^{N(0|x)} \cdot \left(\frac{1}{1+(1-2p)}\right)^{2(1/2-p)s + 2\sqrt{s \log(1/\alpha)}} \\ \geq (\alpha - \alpha^2) \cdot \left(\frac{1}{1+(1-2p)}\right)^{2(1/2-p)s + 2\sqrt{s \log(1/\alpha)}}. \end{aligned}$$

Here, $1 + (1 - 2p) \leq e^{1-2p}$ holds since $1 + x \leq e^x$ for any x . Therefore, we conclude that

$$\Pr[t \leq \tau(T, q) \leq s] \geq (\alpha - \alpha^2) \cdot (1/e)^{4(1/2-p)^2 s + 4(1/2-p)\sqrt{s \log(1/\alpha)}},$$

which is our claim. \blacktriangleleft

Next, we see that the stopping time $\tau(T, p)$ is not so small when T is proper.

► **Lemma 9.** *Suppose that $1/4 < p < 1/2$, $q = 1/2$, and T is a proper stopping rule. Then we have $\Pr\left[\tau(T, p) \leq \frac{1}{10000 \cdot |p-1/2|^2}\right] \leq 1/2$.*

Proof. Let $s = \left\lfloor \frac{1}{10000 \cdot |p-1/2|^2} \right\rfloor$ and $B = \{x \in \{0, 1\}^s \mid T(x) = 1\}$. By the assumption that the rule is proper, we have $\Pr\left[\tau(T, q) \leq \frac{1}{10000 \cdot |p-1/2|^2}\right] = \Pr[\tau(T, q) \leq s] = |B|/2^s \leq 1/3$. Let $B_1 = \{x \in \{0, 1\}^s \mid T(x) = 1, |N(1 | x) - ps| > 2\sqrt{s}\}$ and $B_2 = \{x \in \{0, 1\}^s \mid T(x) = 1, |N(1 | x) - ps| \leq 2\sqrt{s}\}$. Then we have $\Pr[\tau(T, p) \leq s] = \sum_{x \in B} p^{N(1|x)} (1-p)^{N(0|x)} = \sum_{x \in B_1} p^{N(1|x)} (1-p)^{N(0|x)} + \sum_{x \in B_2} p^{N(1|x)} (1-p)^{N(0|x)}$. We bound the two terms separately. By using Hoeffding's inequality, we have

$$\sum_{x \in B_1} p^{N(1|x)} (1-p)^{N(0|x)} \leq 2 \exp\left(-2 \left(\frac{2\sqrt{s}}{s}\right)^2 \cdot s\right) = \frac{2}{e^8} < 0.1.$$

Also, we have

$$\begin{aligned}
\sum_{x \in B_2} p^{N(1|x)} (1-p)^{N(0|x)} &\leq \sum_{x \in B_2} p^{ps-2\sqrt{s}} (1-p)^{(1-p)s+2\sqrt{s}} \\
&\leq \sum_{x \in B} (p(1-p))^{ps} \cdot (1-p)^{(1-2p)s} \cdot \left(\frac{1-p}{p}\right)^{2\sqrt{s}} \\
&\leq \frac{2^s}{3} \cdot \left(\frac{1}{4}\right)^{ps} \cdot (1-p)^{(1-2p)s} \cdot \left(1 + \frac{1-2p}{p}\right)^{2\sqrt{s}} \\
&= \frac{1}{3} \cdot (1 + (1-2p))^{(1-2p)s} \cdot \left(1 + \frac{1-2p}{p}\right)^{2\sqrt{s}} \\
&\leq \frac{1}{3} \cdot \exp((1-2p)^2 s + 8 \cdot (1-2p)\sqrt{s}) \\
&\leq \frac{1}{3} \cdot \exp\left(\frac{4}{10000} + \frac{16}{100}\right) = \frac{e^{0.1604}}{3} < 0.4.
\end{aligned} \tag{7}$$

Here, (7) holds since $1+x \leq e^x$ for any $x \geq 0$ and $1/4 < p < 1/2$.

Therefore, we obtain

$$\Pr \left[\tau(T, p) \leq \frac{1}{10000 \cdot |p-1/2|^2} \right] < 0.1 + 0.4 = \frac{1}{2}. \quad \blacktriangleleft$$

Now we are ready to prove Theorem 1. Recall that $q = 1/2$.

Proof of Theorem 1. To obtain a contradiction, suppose that a proper stopping rule T satisfies Condition (1) for some ϵ_0 , i.e., $\Pr \left[\tau(T, p) \leq \frac{\log \log \frac{1}{|p-1/2|}}{16|p-1/2|^2} \right] \geq \frac{2}{3}$ holds for any p such that $0 < |p-1/2| < \epsilon_0$. By Lemma 9, we have $\Pr \left[\tau(T, p) > \frac{1}{10000 \cdot |p-1/2|^2} \right] \geq \frac{1}{2}$ holds for any p such that $1/4 < p < 1/2$. Hence, we have

$$\Pr \left[\frac{1}{10000 \cdot |p-1/2|^2} < \tau(T, p) \leq \frac{\log \log \frac{1}{|p-1/2|}}{16|p-1/2|^2} \right] \geq \frac{2}{3} + \frac{1}{2} - 1 = \frac{1}{6}$$

for any p such that $1/2 - \min\{\epsilon_0, 1/4\} < p < 1/2$.

Let $p(k) = 1/2 - 1/M^{k^2}$ where k is a natural number and M is a real number that satisfies $M > \max\{e^{e^{32}}, 1/\epsilon_0\}$. Since $0 < 1/2 - p(k) < 1/M < \min\{\epsilon_0, 1/e^{e^{32}}\} \leq \min\{\epsilon_0, 1/4\}$ for any $k \geq 1$, we have

$$\Pr \left[\frac{M^{2k^2}}{10000} < \tau(T, p(k)) \leq \frac{M^{2k^2}}{16} \log \log M^{k^2} \right] \geq \frac{1}{6}.$$

Let U_k be the interval $\left(\frac{M^{2k^2}}{10000}, \frac{M^{2k^2}}{16} \log \log M^{k^2}\right]$. Then $U_i \cap U_j = \emptyset$ holds, for any distinct natural numbers i, j , because we have

$$\frac{M^{2(k+1)^2}}{10000} = \frac{M^{2k^2+4k+2}}{10000} = \frac{M^2}{10000} \cdot M^{2k^2} \cdot M^{4k} > \frac{M^{2k^2}}{16} \log \log M^{k^2}.$$

Here, we use the facts that $M^2/10000 > 1/16$ and $M^{4k} > \log \log M^{k^2}$. The former fact holds by $M > e^{e^{32}} > 25 = \sqrt{10000/16}$. The later fact holds since $M^{4k} = M^{3k} \cdot M^k > 2M^k > M + M^k > \log \log M + \log k^2 = \log \log M^{k^2}$ by $M > e^{e^{32}} > 2$.

32:10 Optimal Stopping Rules for Sequential Hypothesis Testing

In what follows, we produce a contradiction by evaluating the probability

$$P(\ell) = \Pr \left[\tau(T, q) \leq \frac{M^{2\ell^2}}{16} \log \log M^{\ell^2} \right]$$

for a sufficiently large integer ℓ . As the intervals U_i are disjoint, we have

$$P(\ell) \geq \Pr \left[\tau(T, q) \in \bigcup_{k=1}^{\ell} U_k \right] = \sum_{k=1}^{\ell} \Pr [\tau(T, q) \in U_k].$$

Applying Lemma 8 with $p = 1/2 - 1/M^{k^2}$, $s = \frac{M^{2k^2}}{16} \log \log M^{k^2}$, $t = \frac{M^{2k^2}}{10000}$, and $\alpha = 1/6$, we have

$$\begin{aligned} \Pr [\tau(T, q) \in U_k] &\geq \frac{5}{36} \cdot \left(\frac{1}{e} \right)^{4 \cdot \frac{1}{M^{2k^2}} \cdot \frac{M^{2k^2}}{16} \log \log M^{k^2} + 4 \cdot \frac{1}{M^{k^2}} \sqrt{\frac{M^{2k^2}}{16} (\log \log M^{k^2}) \log 6}} \\ &\geq \frac{5}{36} \cdot \left(\frac{1}{e} \right)^{\frac{1}{4} \log \log M^{k^2} + \sqrt{2 \log \log M^{k^2}}}. \end{aligned}$$

Since $\frac{1}{4} \log \log x \geq \sqrt{2 \log \log x}$ holds for $\log \log x \geq 32$ (i.e., $x \geq e^{e^{32}}$), we have

$$\frac{1}{4} \log \log M^{k^2} + \sqrt{2 \log \log M^{k^2}} \leq \frac{1}{2} \log \log M^{k^2}.$$

Hence, we obtain

$$\begin{aligned} P(\ell) &\geq \sum_{k=1}^{\ell} \frac{5}{36} \cdot \left(\frac{1}{e} \right)^{\frac{1}{4} \log \log M^{k^2} + \sqrt{2 \log \log M^{k^2}}} \geq \sum_{k=1}^{\ell} \frac{5}{36} \cdot \left(\frac{1}{e} \right)^{\frac{1}{2} \log \log M^{k^2}} \\ &= \sum_{k=1}^{\ell} \frac{5}{36} \cdot \left(\frac{1}{\log M^{k^2}} \right)^{1/2} = \frac{5}{36\sqrt{\log M}} \sum_{k=1}^{\ell} \frac{1}{k} \geq \frac{5}{36\sqrt{\log M}} \int_1^{\ell+1} \frac{dx}{x} = \frac{5 \log(\ell+1)}{36\sqrt{\log M}}. \end{aligned}$$

By choosing $\ell = \lfloor M \rfloor$, we get $P(\lfloor M \rfloor) \geq \frac{5 \log M}{36\sqrt{\log M}} = \frac{5}{36} \sqrt{\log M} > \frac{5}{36} \sqrt{\log e^{e^{32}}} > 1$, which is a contradiction. \blacktriangleleft

4 Upper bounds

In this section, we give stopping rules for testing identity with small sample complexity.

4.1 The case when q is explicit but p is unknown

In this subsection, we first provide a framework to obtain stopping rules from algorithms for robust identity testing and then prove Theorem 2.

We state a lemma to improve the success probability of a test by repeatedly running the test and taking a majority vote.

► **Lemma 10.** *Suppose that we have an algorithm for a decision problem with success probability at least $2/3$. Then, by running the algorithm $\lceil 18 \log(3k) \rceil$ times and taking the majority, the success probability increases to at least $1 - \frac{1}{9k^2}$.*

Algorithm 1: Stopping rule T^q induced by $T^{q,\epsilon}$

input: $x_1 \cdots x_t \in [n]^*$, distributions q over $[n]$ **output:** 0 or 1

- 1 Let $s_0 = 0$;
- 2 **for** $k = 1, 2, \dots$ **do**
- 3 Let $\epsilon_k = 1/2^k$ and $s_k = s_{k-1} + f(q, \epsilon_k) \cdot \lceil 18 \log(3k) \rceil$;
- 4 **if** $s_k > t$ **then return** 0;
- 5 **else if** $T^{q,\epsilon_k}(x_{s_{k-1}+1} \cdots x_{s_k}) = 1$ **then return** 1;

Suppose that we have an algorithm, for a given q , with sample complexity $f(q, \epsilon)$ that distinguishes between the cases

(a) $d_1(p, q) \geq \epsilon$ and

(b) $d_2(p, q) \leq \epsilon/2$,

with probability at least $2/3$, where d_1 and d_2 are distance measures that depend on the application. Then, by Lemma 10, we can obtain a stopping rule $T^{q,\epsilon}$ such that

- $\Pr[\tau(T^{q,\epsilon}, p) \leq f(q, \epsilon) \cdot \lceil 18 \log(3k) \rceil] \geq 1 - \frac{1}{9k^2}$ if $d_1(p, q) \geq \epsilon$, and
- $\Pr[\tau(T^{q,\epsilon}, p) \leq f(q, \epsilon) \cdot \lceil 18 \log(3k) \rceil] \leq \frac{1}{9k^2}$ if $d_2(p, q) \leq \epsilon/2$.

We then formulate a stopping rule T^q for identity testing as follows. The tester guesses ϵ and then tests identity of p, q by using $T^{q,\epsilon}$. If $T^{q,\epsilon}$ does not stop with $f(q, \epsilon) \cdot \lceil 18 \log(3k) \rceil$ samples, it reduces ϵ to half and continue the procedure recursively. The stopping rule T^q is summarized as Algorithm 1.

We show that T^q is the desired stopping rule.

► **Lemma 11.** *If $p \neq q$, the stopping time $\tau(T^q, p)$ for T^q in Algorithm 1 satisfies $\Pr[s_a \leq \tau(T^q, p) \leq s_b] \geq 2/3$, where $a = \lfloor \log_2 \frac{1}{2d_2(p,q)} \rfloor$, $b = \lceil \log_2 \frac{1}{d_1(p,q)} \rceil$, and $s_\ell = \sum_{k=1}^{\ell} f(q, \epsilon_k) \cdot \lceil 18 \log(3k) \rceil$.*

Proof. Since $d_1(p, q) \geq 1/2^b = \epsilon_b$ by $b = \lceil \log_2 \frac{1}{d_1(p,q)} \rceil$, the stopping time is larger than s_b with probability at most

$$\begin{aligned} \Pr[\tau(T^q, p) > s_b] &= \Pr[\tau(T^q, p) \geq s_b] \cdot \Pr[\tau(T^q, p) \neq s_b \mid \tau(T^q, p) \geq s_b] \\ &= \Pr[\tau(T^q, p) \geq s_b] \cdot (1 - \Pr[\tau(T^q, p) = s_b \mid \tau(T^q, p) \geq s_b]) \\ &\leq 1 - \Pr[\tau(T^q, p) = s_b \mid \tau(T^q, p) \geq s_b] \\ &= 1 - \Pr[\tau(T^{q,\epsilon_b}, p) \leq f(q, \epsilon_b) \cdot \lceil 18 \log(3b) \rceil] \leq \frac{1}{9} \cdot \frac{1}{b^2} \leq \frac{1}{9}. \end{aligned}$$

On the other hand, since $d_2(p, q) \leq \frac{1}{2} \cdot \frac{1}{2^a} \leq \epsilon_k/2$ for any $1 \leq k \leq a$ by $a = \lfloor \log_2 \frac{1}{2d_2(p,q)} \rfloor$, the stopping time is smaller than s_a with probability at most

$$\begin{aligned} \Pr[\tau(T^q, p) < s_a] &= \sum_{k=1}^{a-1} \Pr[\tau(T^q, p) = s_k] \leq \sum_{k=1}^{a-1} \Pr[\tau(T^q, p) = s_k \mid \tau(T^q, p) \geq s_k] \\ &= \sum_{k=1}^{a-1} \Pr[\tau(T^{q,\epsilon_k}, p) \leq f(q, \epsilon_k) \cdot \lceil 18 \log(3k) \rceil] \\ &\leq \sum_{k=1}^{a-1} \frac{1}{9} \cdot \frac{1}{k^2} < \sum_{k=1}^{\infty} \frac{1}{9} \cdot \frac{1}{k^2} = \frac{1}{9} \cdot \frac{\pi^2}{6} < \frac{2}{9}. \end{aligned}$$

32:12 Optimal Stopping Rules for Sequential Hypothesis Testing

Hence, the stopping time of T^q satisfies

$$\Pr[s_a \leq \tau(T^q, p) \leq s_b] = 1 - \Pr[\tau(T^q, p) > s_b] - \Pr[\tau(T^q, p) < s_a] \geq 1 - \frac{1}{9} - \frac{2}{9} = \frac{2}{3},$$

which completes the proof. \blacktriangleleft

Next, we prove Theorem 2. To provide stopping rules, we use robust identity testing algorithm in Theorem 5. When T^q is the stopping rule induced by the algorithm in Theorem 5, we have $d_1(p, q) = d_{\text{tv}}(p, q)$, $d_2(p, q) = \sqrt{\chi^2(p, q)}$, and $f(q, \epsilon) = \lfloor \frac{c_q \sqrt{n}}{\epsilon^2} \rfloor$ for a constant c_q , which depends only on q . Note that $\max_q c_q = O(1)$. Then, we have

$$s_\ell = \sum_{k=1}^{\ell} \left\lfloor \frac{c_q \sqrt{n} \cdot \lceil 18 \log(3k) \rceil}{\epsilon_k^2} \right\rfloor = \sum_{k=1}^{\ell} \left\lfloor c_q \sqrt{n} \cdot 4^k \cdot \lceil 18 \log(3k) \rceil \right\rfloor = \Theta(\sqrt{n} \cdot 4^\ell \log \ell).$$

Here, the last equality holds since

$$4^\ell \log(3\ell) < \sum_{k=1}^{\ell} 4^k \log(3k) < \sum_{k=1}^{\ell} 4^k \log(3\ell) = \frac{4}{3}(4^\ell - 1) \log(3\ell) < \frac{4}{3} \cdot 4^\ell \log(3\ell).$$

By setting $a = \left\lceil \log_2 \frac{1}{2\sqrt{\chi^2(p, q)}} \right\rceil$ and $b = \left\lceil \log_2 \frac{1}{d_{\text{tv}}(p, q)} \right\rceil$, we have

$$s_a = \Theta\left(\frac{\sqrt{n}}{\chi^2(p, q)} \log \log \frac{1}{\chi^2(p, q)}\right) \quad \text{and} \quad s_b = \Theta\left(\frac{\sqrt{n}}{d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)}\right),$$

and hence, we obtain Theorem 2.

4.2 The case when p and q are both unknown

We next consider the case when p and q are both unknown. We build a similar framework for the case and then provide a stopping rule for Theorem 3.

Suppose that we have an algorithm with sample complexity $g(\epsilon)$ that distinguishes between the cases

- (a) $d_1(p, q) \geq \epsilon$ and
- (b) $d_2(p, q) \leq \epsilon/2$,

with probability at least $2/3$. Then, by Lemma 10, we can obtain a stopping rule T^ϵ such that

- $\Pr\left[\tau(T^{q, \epsilon}, p) \leq g(\epsilon) \cdot \lceil 18 \log(3k) \rceil\right] \geq 1 - \frac{1}{9k^2}$ if $d_1(p, q) \geq \epsilon$, and
- $\Pr\left[\tau(T^{q, \epsilon}, p) \leq g(\epsilon) \cdot \lceil 18 \log(3k) \rceil\right] \leq \frac{1}{9k^2}$ if $d_2(p, q) \leq \epsilon/2$.

Our framework is almost the same as Algorithm 1. The stopping rule T induced by T^ϵ is shown as Algorithm 2.

Then we can prove the following lemma in the same way as the proof of Lemma 11.

► **Lemma 12.** *If $p \neq q$, the stopping time $\tau(T, p, q)$ for T in Algorithm 2 satisfies*

$$\Pr[s_a \leq \tau(T, p, q) \leq s_b] \geq 2/3$$

where $a = \left\lceil \log_2 \frac{1}{2d_2(p, q)} \right\rceil$, $b = \left\lceil \log_2 \frac{1}{d_1(p, q)} \right\rceil$, and $s_\ell = \sum_{k=1}^{\ell} f(q, \epsilon_k) \cdot \lceil 18 \log(3k) \rceil$.

Algorithm 2: Stopping rule T induced by T^ϵ

input: $x_1 \cdots x_t \in [n]^*$ and $y_1 \cdots y_t \in [n]^*$ **output:** 0 or 1

- 1 Let $s_0 = 0$;
- 2 **for** $k = 1, 2, \dots$ **do**
- 3 Let $\epsilon_k = 1/2^k$ and $s_k = s_{k-1} + g(\epsilon_k) \cdot \lceil 18 \log(3k) \rceil$;
- 4 **if** $s_k > t$ **then return** 0;
- 5 **else if** $T^{\epsilon_k}(x_{s_{k-1}+1} \cdots x_{s_k}, y_{s_{k-1}+1} \cdots y_{s_k}) = 1$ **then return** 1;

When T^q is the stopping rule induced by the algorithm in Theorem 6, we have $d_1(p, q) = d_2(p, q) = d_{\text{tv}}(p, q)$ and $g(\epsilon) = \lfloor \frac{cn}{\epsilon^2 \log n} \rfloor$ for a constant c . Then, we have

$$s_\ell = \sum_{k=1}^{\ell} \left\lfloor \frac{cn \cdot \lceil 18 \log(3k) \rceil}{\epsilon_k^2 \log n} \right\rfloor = \sum_{k=1}^{\ell} \left\lfloor \frac{cn \cdot 4^k \cdot \lceil 18 \log(3k) \rceil}{\log n} \right\rfloor = \Theta \left(\frac{n \cdot 4^\ell \log \ell}{\log n} \right).$$

By setting $a = \left\lfloor \log_2 \frac{1}{2d_{\text{tv}}(p, q)} \right\rfloor$ and $b = \left\lceil \log_2 \frac{1}{d_{\text{tv}}(p, q)} \right\rceil$, we have

$$s_a = \Theta \left(\frac{n/\log n}{d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)} \right) \quad \text{and} \quad s_b = \Theta \left(\frac{n/\log n}{d_{\text{tv}}(p, q)^2} \log \log \frac{1}{d_{\text{tv}}(p, q)} \right).$$

Hence, we obtain Theorem 3.

References

- 1 Jayadev Acharya, Clément L. Canonne, and Gautam Kamath. A chasm between identity and equivalence testing with conditional queries. In *Proceedings of APPROX/RANDOM*, 2015.
- 2 Jayadev Acharya, Constantinos Daskalakis, and Gautam C. Kamath. Optimal Testing for Properties of Distributions. In *Proceedings of NIPS*, pages 3591–3599, 2015.
- 3 Akshay Balsubramani. Sharp finite-time iterated-logarithm martingale concentration. page 25, 2014. [arXiv:1405.2639](#).
- 4 Akshay Balsubramani. Sharp finite-time iterated-logarithm martingale concentration. *arXiv preprint arXiv:1405.2639*, 2014.
- 5 Akshay Balsubramani and Aaditya Ramdas. Sequential nonparametric testing with the law of the iterated logarithm. *arXiv preprint arXiv:1506.03486*, 2015.
- 6 Jay Bartroff, Tze Leung Lai, and Mei-Chiung Shih. *Sequential experimentation in clinical trials: design and analysis*, volume 298. Springer Science & Business Media, 2012.
- 7 Tuğkan Batu, Eldar Fischer, Lance Fortnow, Ravi Kumar, Ronitt Rubinfeld, and Patrick White. Testing random variables for independence and identity. In *Proceedings FOCS*, pages 442–451, 2001.
- 8 Tuğkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D Smith, and Patrick White. Testing that distributions are close. In *Proceedings of FOCS*, pages 259–269, 2000.
- 9 Arnab Bhattacharyya, Eldar Fischer, Ronitt Rubinfeld, and Paul Valiant. Testing monotonicity of distributions over general partial orders. In *Proceedings of ICS*, pages 239–252, 2011.
- 10 Clément L. Canonne. A survey on distribution testing: Your data is big. but is it blue? *Electronic Colloquium on Computational Complexity (ECCC)*, 22:63, 2015.
- 11 Clément L. Canonne, Ilias Diakonikolas, Themis Gouleakis, and Ronitt Rubinfeld. Testing shape restrictions of discrete distributions. In *Proceedings of STACS*, 2016.

- 12 Clément L. Canonne, Dana Ron, and Rocco A. Servedio. Testing probability distributions using conditional samples. *SIAM Journal on Computing*, 44(3):540–616, 2015.
- 13 Siu-On Chan, Ilias Diakonikolas, Gregory Valiant, and Paul Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of SODA*, pages 1193–1203, 2014.
- 14 Roger H. Farrell. Asymptotic behavior of expected sample size in certain one sided tests. *The Annals of Mathematical Statistics*, pages 36–72, 1964.
- 15 Oded Goldreich and Dana Ron. On Testing Expansion in Bounded-Degree Graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 7(20), 2000.
- 16 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- 17 Kevin Jamieson, Matt Malloy, Robert Nowak, and Sebastien Bubeck. li’ UCB : An Optimal Exploration Algorithm for Multi-Armed Bandits. In *Proceedings of COLT*, 2014.
- 18 Ramesh Johari, Leo Pekelis, and David J Walsh. Always valid inference: Bringing sequential analysis to A/B testing. *arXiv preprint arXiv:1512.04922*, 2015.
- 19 Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *Proceedings of ICML*, 2013.
- 20 Eugene Kharitonov, Aleksandr Vorobev, Craig Macdonald, Pavel Serdyukov, and Iadh Ounis. Sequential testing for early stopping of online experiments. In *Proceedings of SIGIR*, pages 473–482, 2015.
- 21 Aleksandr Khintchine. über einen satz der wahrscheinlichkeitsrechnung. *Fundamenta Mathematicae*, 6(1):9–20, 1924.
- 22 Martin Kulldorff, Robert L Davis, Margarette Kolczak, Edwin Lewis, Tracy Lieu, and Richard Platt. A maximized sequential probability ratio test for drug and vaccine safety surveillance. *Sequential analysis*, 30(1):58–78, 2011.
- 23 Tze Leung Lai. *Sequential analysis*. Wiley Online Library, 2001.
- 24 K.K. Gordon Lan and David L. DeMets. Discrete sequential boundaries for clinical trials. *Biometrika*, 70(3):659–663, 1983.
- 25 Evan Miller. How Not To Run An A/B Test. 2010. URL: <http://www.evanmiller.org/how-not-to-run-an-ab-test.html>.
- 26 Peter C. O’Brien and Thomas R Fleming. A multiple testing procedure for clinical trials. *Biometrics*, pages 549–556, 1979.
- 27 Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.
- 28 Stuart J. Pocock. Group sequential methods in the design and analysis of clinical trials. *Biometrika*, pages 191–199, 1977.
- 29 Ronitt Rubinfeld. Taming big probability distributions. *XRDS*, 19(1):24–28, 2012.
- 30 David Siegmund. Estimation following sequential tests. *Biometrika*, 65(2):341–349, 1978.
- 31 David Siegmund. *Sequential analysis: tests and confidence intervals*. Springer Science & Business Media, 2013.
- 32 Gregory Valiant and Paul Valiant. The power of linear estimators. In *Proceedings of FOCS*, pages 403–412, 2011.
- 33 Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. In *Proceedings of FOCS*, pages 51–60, 2014.
- 34 Abraham Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 1945.
- 35 Abraham Wald and Jacob Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, pages 326–339, 1948.
- 36 Samuel K. Wang and Anastasios A. Tsiatis. Approximately optimal one-parameter boundaries for group sequential trials. *Biometrics*, pages 193–199, 1987.