

Quantum Sub-Gaussian Mean Estimator

Yassine Hamoudi

Université de Paris, IRIF, CNRS, F-75006 Paris, France.
hamoudi@irif.fr

September 17, 2021

Abstract

We present a new quantum algorithm for estimating the mean of a real-valued random variable obtained as the output of a quantum computation. Our estimator achieves a nearly-optimal quadratic speedup over the number of classical i.i.d. samples needed to estimate the mean of a heavy-tailed distribution with a sub-Gaussian error rate. This result subsumes (up to logarithmic factors) earlier works on the mean estimation problem that were not optimal for heavy-tailed distributions [BHMT02; BDGT11], or that require prior information on the variance [Hei02; Mon15; HM19]. As an application, we obtain new quantum algorithms for the (ϵ, δ) -approximation problem with an optimal dependence on the coefficient of variation of the input random variable.

1 Introduction

The problem of estimating the mean μ of a real-valued random variable X given *i.i.d.* samples from it is one of the most basic tasks in statistics and in the Monte Carlo method. The properties of the various classical *mean estimators* are well understood. The standard non-asymptotic criterion used to assess the quality of an estimator is formulated as the following *high probability deviation bound*: upon performing n random experiments that return n samples from X , and given a failure probability $\delta \in (0, 1)$, what is the smallest error $\epsilon(n, \delta, X)$ such that the output $\tilde{\mu}$ of the estimator satisfies $|\tilde{\mu} - \mu| > \epsilon(n, \delta, X)$ with probability at most δ ? Under the standard assumption that the unknown random variable X has a finite variance σ^2 , the best possible performances are obtained by the so-called *sub-Gaussian estimators* [LM19] that achieve the following deviation bound

$$\Pr \left[|\tilde{\mu} - \mu| > L \sqrt{\frac{\sigma^2 \log(1/\delta)}{n}} \right] \leq \delta \quad (1)$$

for some constant L . The term “sub-Gaussian” reflects that these estimators have a Gaussian tail even for non-Gaussian distributions. The most well-known sub-Gaussian estimator is arguably the *median-of-means* [NY83; JVV86; AMS99], which consists of partitioning the n samples into roughly $\log(1/\delta)$ groups of equal size, computing the empirical mean over each group, and returning the median of the obtained means.

The process of generating a random sample from X is generalized in the quantum model by assuming the existence of a unitary operator U where $U|0\rangle$ coherently encodes the distribution of X . A *quantum experiment* is then defined as one application of this operator or its inverse. The celebrated quantum amplitude estimation algorithm [BHMT02] provides a way to estimate the mean of any *Bernoulli* random variable by performing fewer experiments than with any classical estimator. Yet, for general distributions, the existing quantum mean estimators either

require additional information on the variance [Hei02; Mon15; HM19] or are less performant than the classical sub-Gaussian estimators when the distribution is heavy tailed [BHMT02; Ter99; BDGT11; Mon15]. These results leave open the existence of a general quantum speedup for the mean estimation problem. We address this question by introducing the concept of *quantum sub-Gaussian estimators*, defined through the following deviation bound

$$\Pr\left[|\tilde{\mu} - \mu| > L \frac{\sigma \log(1/\delta)}{n}\right] \leq \delta \quad (2)$$

for some constant L . We give the first construction of a quantum estimator that achieves this bound up to a logarithmic factor in n . Additionally, we prove that it is impossible to go below that deviation level. This result provides a clear equivalent of the concept of sub-Gaussian estimator in the quantum setting.

A second important family of mean estimators addresses the (ϵ, δ) -approximation problem, where given a fixed relative error $\epsilon \in (0, 1)$ and a failure probability $\delta \in (0, 1)$ the goal is to output a mean estimate $\tilde{\mu}$ such that

$$\Pr[|\tilde{\mu} - \mu| > \epsilon|\mu|] \leq \delta. \quad (3)$$

The aforementioned sub-Gaussian estimators do not quite answer this question since the number of experiments they require (respectively $n = \Omega\left(\left(\frac{\sigma}{\epsilon\mu}\right)^2 \log(1/\delta)\right)$ and $n = \tilde{\Omega}\left(\frac{\sigma}{\epsilon|\mu|} \log(1/\delta)\right)$) depends on the *unknown* quantities σ and μ . Sometimes a good upper bound is known on the *coefficient of variation* $|\sigma/\mu|$ and can be used to parametrize a sub-Gaussian estimator. Otherwise, the standard approach is based on *sequential analysis* techniques, where the number of experiments is chosen adaptively depending on the results of previous computations. Given a random variable distributed in $[0, 1]$, the optimal classical estimators perform $\Theta\left(\left(\left(\frac{\sigma}{\epsilon\mu}\right)^2 + \frac{1}{\epsilon\mu}\right) \log(1/\delta)\right)$ random experiments *in expectation* [DKLR00] for computing an (ϵ, δ) -approximation of μ . We construct a quantum estimator that reduces this number to $\Theta\left(\left(\frac{\sigma}{\epsilon\mu} + \frac{1}{\sqrt{\epsilon\mu}}\right) \log(1/\delta)\right)$ and we prove that it is optimal.

1.1 Related work

There is an extensive literature on classical sub-Gaussian estimators and we refer the reader to [LM19; Cat12; BCL13; DLLO16; LV20] for an overview of the main results and recent improvements. We point out that the *empirical mean* estimator is not sub-Gaussian, although it is optimal for Gaussian random variables [SV05; Cat12]. The non-asymptotic performances of the empirical mean estimator are captured by several standard concentration bounds such as the Chebyshev, Chernoff and Bernstein inequalities.

There is a series of quantum mean estimators [Gro98; AW99; BDGT11] that get close to the bound $\Pr\left[|\tilde{\mu} - \mu| > L \frac{\log(1/\delta)}{n}\right] \leq \delta$ for any random variable distributed in $[0, 1]$ and some constant L . Similar results hold for numerical integration problems [AW99; Nov01; Hei02; TW02; Hei03]. The amplitude estimation algorithm [BHMT02; Ter99] leads to a sharper bound of $\Pr\left[|\tilde{\mu} - \mu| > L \left(\frac{\sqrt{\mu(1-\mu)} \log(1/\delta)}{n} + \frac{\log(1/\delta)^2}{n^2}\right)\right] \leq \delta$ (see Proposition 4.1) when X is distributed in $[0, 1]$. Nevertheless, the quantity $\mu(1-\mu)$ is always larger than or equal to the variance σ^2 . The question of improving the dependence on σ^2 was considered in [Hei02; Mon15; HM19]. The estimators of [Hei02; Mon15] require to know an upper bound Σ on the standard deviation σ , whereas [HM19] needs an upper bound Δ on the coefficient of variation σ/μ (for non-negative random variables). The performances of these estimators are captured (up to logarithmic factors) by the deviation bound given in Equation (2) with σ replaced by Σ and $\mu\Delta$ respectively.

The (ϵ, δ) -approximation problem has been addressed by several classical works such as [DKLR00; MSA08; GNP13; Hub19]. In the quantum setting, there is a variant [BHMT02, Theorem 15] of the amplitude estimation algorithm that performs $O(\log(1/\delta)/(\epsilon\sqrt{\mu}))$ experiments

in expectation to compute an (ϵ, δ) -approximate of the mean of a random variable distributed in $[0, 1]$ (see Theorem A.4 and Proposition 5.2). However, the complexity of this estimator does not scale with σ . Given an upper bound Δ on σ/μ , the estimator of [HM19] can be used to compute an (ϵ, δ) -approximate with roughly $\tilde{O}(\Delta \log(1/\delta)/\epsilon)$ quantum experiments if the random variable is non-negative.

We note that the related problem of estimating the mean with *additive* error ϵ , that is $\Pr[|\tilde{\mu} - \mu| > \epsilon] \leq \delta$, has also been considered by several authors. The optimal number of experiments is $\Theta(\log(1/\delta)/\epsilon^2)$ classically [CEG95] and $\Theta(1/\epsilon)$ quantumly [NW99] (with failure probability $\delta = 1/3$). These bounds do not depend on unknown parameters (as opposed to the relative error case), thus sequential analysis techniques are unnecessary here. Montanaro [Mon15] also described an estimator that performs $\tilde{O}(\Sigma \log(1/\delta)/\epsilon)$ quantum experiments given an upper bound Σ on the standard deviation σ .

1.2 Contributions and organization

We first formally define the input model in Section 2. We introduce the concept of “q-random variable” (Definition 2.3) to describe a random variable that corresponds to the output of a quantum computation. We measure the complexity of an algorithm by counting the number of *quantum experiments* (Definition 2.4) it performs with respect to a q-random variable.

We construct a quantum algorithm for estimating the quantiles of a q-random variable in Section 3, and we use it in Section 4 to design the following quantum sub-Gaussian estimator.

Theorem 4.2 (Restated). *There exists a quantum algorithm with the following properties. Let X be a q-random variable with mean μ and variance σ^2 , and set as input a time parameter n and a real $\delta \in (0, 1)$ such that $n \geq \log(1/\delta)$. Then, the algorithm outputs a mean estimate $\tilde{\mu}$ such that, $\Pr\left[|\tilde{\mu} - \mu| > \frac{\sigma \log(1/\delta)}{n}\right] \leq \delta$, and it performs $O(n \log^{3/2}(n) \log \log(n))$ quantum experiments.*

Then we turn our attention to the (ϵ, δ) -approximation problem in Section 5. In case we have an upper bound Δ on the coefficient of variation $|\sigma/\mu|$, we directly use our sub-Gaussian estimator to obtain an algorithm that performs $\tilde{O}(\frac{\Delta}{\epsilon} \log(1/\delta))$ quantum experiments (Corollary 5.1). Next, we consider the more subtle parameter-free setting where there is no prior information about the input random variable, except that it is distributed in $[0, 1]$. In this case, the number of experiments is chosen *adaptively*, and the bound we get is stated in expectation.

Theorem 5.3 (Restated). *There exists a quantum algorithm with the following properties. Let X be a q-random variable distributed in $[0, 1]$ with mean μ and variance σ^2 , and set as input two reals $\epsilon, \delta \in (0, 1)$. Then, the algorithm outputs a mean estimate $\tilde{\mu}$ such that $\Pr[|\tilde{\mu} - \mu| > \epsilon\mu] \leq \delta$, and it performs $\tilde{O}\left(\left(\frac{\sigma}{\epsilon\mu} + \frac{1}{\sqrt{\epsilon\mu}}\right) \log(1/\delta)\right)$ quantum experiments in expectation.*

Finally, we prove several lower bounds in Section 6 that match the complexity of the above estimators. We also consider the weaker input model where one is given copies of a quantum state encoding the distribution of X . We prove that no quantum speedup is achievable in this setting (Theorem 6.6).

1.3 Proof overview

Sub-Gaussian estimator. Our approach (Theorem 4.2) combines several ideas used in previous classical and quantum mean estimators. In this section, we simplify the exposition by assuming that the random variable X is non-negative and by replacing the variance σ^2 with the second moment $\mathbb{E}[X^2]$. We also take the failure probability δ to be a small constant. Our starting point is a variant of the *truncated mean estimators* [Bic65; BCL13; LM19]. Truncation is a process that consists of replacing the samples larger than some threshold value with a smaller number. This has the effect of reducing the tail of the distribution, but also of changing its expectation. Here

we study the effect of replacing the values larger than some threshold b with 0, which corresponds to the new random variable $Y = X\mathbb{1}_{X \leq b}$. We consider the following classical sub-Gaussian estimator that we were not able to find in the literature: set $b = \sqrt{n\mathbb{E}[X^2]}$ and compute the empirical mean of n samples from Y . By a simple calculation, one can prove that the expectation of the removed part is at most $\mathbb{E}[X - Y] \leq \mathbb{E}[X^2]/b = \sqrt{\mathbb{E}[X^2]}/n$. Moreover, using Bernstein's inequality and the boundedness of Y , the error between the output estimate and $\mathbb{E}[Y]$ is on the order of $\sqrt{\mathbb{E}[X^2]}/n$. These two facts together imply that the overall error for estimating $\mathbb{E}[X]$ is indeed of a sub-Gaussian type. This approach can be carried out in the quantum model by performing the truncation in superposition. This is similar to what is done in previous quantum mean estimators [Hei02; Mon15; HM19]. In order to obtain a quantum speedup, one must balance the truncation level differently by taking $b = n\sqrt{\mathbb{E}[X^2]}$. Then, by a clever use of amplitude estimation discovered by Heinrich [Hei02] (see also [HM18, Proposition A.1]), the expectation of Y can be estimated with an error on the order of $\sqrt{\mathbb{E}[Y^2]}/n + \max(Y)/n^2 \leq 2\sqrt{\mathbb{E}[X^2]}/n$. The main drawback of this estimator is that it requires the knowledge of $\mathbb{E}[X^2]$ to perform the truncation. In previous work [Hei02; Mon15; HM19], the authors made further assumptions on the variance to be able to approximate b . Here, we overcome this issue by choosing the truncation level b differently. Borrowing ideas from classical estimators [LM19], we define b as the quantile value that satisfies $\Pr[X \geq b] = 1/n^2$. This quantile is always smaller than the previous threshold value $n\sqrt{\mathbb{E}[X^2]}$. Moreover, it can be shown that the removed part $\mathbb{E}[X - Y]$ is still on the order of $\sqrt{\mathbb{E}[X^2]}/n$. We give a new quantum algorithm for approximating this quantile with roughly n quantum experiments (Theorem 3.4), whereas it would require n^2 random experiments classically. Our quantile estimation algorithm builds upon the quantum minimum finding algorithm of Dürr and Høyer [DH96; AGGW20] and the k th-smallest element finding algorithm of Nayak and Wu [NW99]. Importantly, it does not require any knowledge about $\mathbb{E}[X^2]$.

(ϵ, δ) -Approximation without side information. We follow an approach similar to that of a classical estimator described in [DKLR00]. Our algorithm (Theorem 5.3) uses the quantum sub-Gaussian estimator and the quantum *sequential Bernoulli estimator* described in Proposition 5.2. The latter estimator can estimate the mean μ of a random variable X distributed in $[0, 1]$ with constant relative error by performing $O(1/\sqrt{\mu})$ quantum experiments in expectation. The first step of the (ϵ, δ) -approximation algorithm is to compute a rough estimate $\hat{\mu}$ of μ with the sequential Bernoulli estimator. Then, the variance σ^2 of X is estimated by using again the sequential Bernoulli estimator on the random variable $(X - X')/2$ (where X' is an independent copy of X). The latter estimation is stopped if it uses more than $O(1/\sqrt{\epsilon\hat{\mu}})$ quantum experiments. We show that if $\sigma^2 \geq \Omega(\epsilon\mu)$ then the computation is not stopped and the resulting estimate $\tilde{\sigma}^2$ is close to σ^2 with high probability. Otherwise, it is stopped with high probability and we set $\tilde{\sigma} = 0$. Finally, the quantum sub-Gaussian estimator is used with the parameter $n \approx \max(\frac{\tilde{\sigma}}{\epsilon\hat{\mu}}, \frac{1}{\sqrt{\epsilon\hat{\mu}}})$ to obtain a refined estimate $\tilde{\mu}$ of μ . The choice of the first (resp. second) term in the maximum value implies that $|\tilde{\mu} - \mu| \leq \epsilon\mu$ with high probability when the variance σ^2 is larger (resp. smaller) than $\epsilon\mu$. In order to upper bound the expected number of experiments performed by this estimator, we show in Proposition 5.2 that the estimates $\hat{\mu}$ and $\tilde{\sigma}$ obtained with the sequential Bernoulli estimator satisfy the expectation bounds $\mathbb{E}[1/\hat{\mu}] \leq 1/\mu$, $\mathbb{E}[\tilde{\sigma}] \leq \sigma$ and $\mathbb{E}[1/\sqrt{\hat{\mu}}] \leq 1/\sqrt{\mu}$.

Lower bounds. We sketch the proof of optimality of the quantum sub-Gaussian estimator (Theorem 6.2). The lower bound is proved in the stronger quantum query model, which allows us to extend it to all the other models mentioned in Section 2. Our approach is inspired by the truncation level chosen in the algorithm. Given σ and n , we consider the two distributions p_0 and p_1 that output respectively $\frac{n\sigma}{\sqrt{1-1/n^2}}$ and $\frac{-n\sigma}{\sqrt{1-1/n^2}}$ with probability $1/n^2$, and 0 otherwise. The two distributions have variance σ^2 and the distance between their means is larger than $\frac{2\sigma}{n}$.

Thus, any estimator that satisfies the bound $\Pr[|\tilde{\mu} - \mu| > \frac{\sigma}{n}] \leq \frac{1}{3}$ can distinguish between p_0 and p_1 with constant success probability. However, we show by a reduction to Quantum Search that it requires at least $\Omega(n)$ quantum experiments to distinguish between two distributions that differ with probability at most $1/n^2$.

2 Model of input

The input to the mean estimation problem is represented by a real-valued random variable X defined on some probability space. A classical estimator accesses this input by obtaining n *i.i.d* samples of X . In this section, we describe the access model for quantum estimators and we compare it to previous models suggested in the literature. We only consider finite probability spaces for finite encoding reasons. First, we recall the definition of a random variable, and we define a classical model of access called a *random experiment*.

Definition 2.1 (RANDOM VARIABLE). A finite *random variable* is a function $X : \Omega \rightarrow E$ for some probability space (Ω, p) , where Ω is a finite sample set, $p : \Omega \rightarrow [0, 1]$ is a probability mass function and $E \subset \mathbb{R}$ is the support of X . As is customary, we will often omit to mention (Ω, p) when referring to the random variable X .

Definition 2.2 (RANDOM EXPERIMENT). Given a random variable X on a probability space (Ω, p) , we define a *random experiment* as the process of drawing a sample $\omega \in \Omega$ according to p and observing the value of $X(\omega)$.

We now introduce the concept of “q-random variable” to represent a quantum process that outputs a real number.

Definition 2.3 (Q-RANDOM VARIABLE). A *q-variable* is a triple (\mathcal{H}, U, M) where \mathcal{H} is a finite-dimensional Hilbert space, U is a unitary transformation on \mathcal{H} , and $M = \{M_x\}_{x \in E}$ is a projective measurement on \mathcal{H} indexed by a finite set $E \subset \mathbb{R}$. Given a random variable X on a probability space (Ω, p) , we say that a q-variable (\mathcal{H}, U, M) *generates* X when,

- (1) \mathcal{H} is a finite-dimensional Hilbert space with some basis $\{|\omega\rangle\}_{\omega \in \Omega}$ indexed by Ω .
- (2) U is a unitary transformation on \mathcal{H} such that $U|\mathbf{0}\rangle = \sum_{\omega \in \Omega} \sqrt{p(\omega)}|\omega\rangle$.
- (3) $M = \{M_x\}_x$ is the projective measurement on \mathcal{H} defined by $M_x = \sum_{\omega: X(\omega)=x} |\omega\rangle\langle\omega|$.

A random variable X is a *q-random variable* if it is generated by some q-variable (\mathcal{H}, U, M) .

We stress that the sample space Ω may not be known explicitly, and we do not assume that it is easy to perform a measurement in the $\{|\omega\rangle\}_{\omega \in \Omega}$ basis for instance. Often, we are given a unitary U such that $U|\mathbf{0}\rangle = \sum_{x \in E} \sqrt{p(x)}|\psi_x\rangle|x\rangle$ for some unknown garbage unit state $|\psi_x\rangle$, together with the measurement $M = \{I \otimes |x\rangle\langle x|\}_{x \in E}$. In this case, we can consider the q-random variable X defined on the probability space (Ω, p) where $\Omega = \{|\psi_x\rangle|x\rangle\}_{x \in E}$ and $X(|\psi_x\rangle|x\rangle) = x$.

We further assume that there exist two quantum oracles, defined below, for obtaining information on the function $X : \Omega \rightarrow E$. These two oracles can be efficiently implemented if we have access to a quantum evaluation oracle $|\omega\rangle|\mathbf{0}\rangle \mapsto |\omega\rangle|X(\omega)\rangle$ for instance. The rotation oracle (Assumption B) has been extensively used in previous quantum mean estimators [Ter99; BDGT11; Mon15; HM19]. The comparison oracle (Assumption A) is needed in our work to implement the quantile estimation algorithm.

Assumption A (COMPARISON ORACLE). Given a q-random variable X on a probability space (Ω, p) , and any two values $a, b \in \mathbb{R} \cup \{-\infty, +\infty\}$ such that $a < b$, there is a unitary operator $C_{a,b}$ acting on $\mathcal{H} \otimes \mathbb{C}^2$ such that for all $\omega \in \Omega$,

$$C_{a,b}(|\omega\rangle|0\rangle) = \begin{cases} |\omega\rangle|1\rangle & \text{when } a < X(\omega) \leq b, \\ |\omega\rangle|0\rangle & \text{otherwise.} \end{cases}$$

Assumption B (ROTATION ORACLE). Given a q -random variable X on a probability space (Ω, p) , and any two reals $0 \leq a < b$, there is a unitary operator $R_{a,b}$ acting on $\mathcal{H} \otimes \mathbb{C}^2$ such that for all $\omega \in \Omega$,

$$R_{a,b}(|\omega\rangle|0\rangle) = \begin{cases} |\omega\rangle \left(\sqrt{1 - \left| \frac{X(\omega)}{b} \right|} |0\rangle + \sqrt{\left| \frac{X(\omega)}{b} \right|} |1\rangle \right) & \text{when } a < X(\omega) \leq b, \\ |\omega\rangle|0\rangle & \text{otherwise.} \end{cases}$$

We now define the measure of complexity used to count the number of accesses to a q -random variable, which are referred to as *quantum experiments*.

Definition 2.4 (QUANTUM EXPERIMENT). Let X be a q -random variable that satisfies Assumptions A and B. Let (\mathcal{H}, U, M) be a q -variable that generates X . We define a *quantum experiment* as the process of applying any of the unitaries $U, C_{a,b}, R_{a,b}$ (for any values of $a < b$), their inverses or their controlled versions, or performing a measurement according to M .

Note that a random experiment (Definition 2.2) can be simulated with two quantum experiments by computing the state $U|0\rangle$ and measuring it according to M . We briefly mention two other possible input models. First, some authors [Gro98; NW99; Hei02; BHH11; CFMW10; BDGT11; LW19] consider the stronger query model where p is the uniform distribution and a quantum evaluation oracle is provided for the function $\omega \mapsto X(\omega)$. A second model tackles the problem of *learning from quantum states* [BJ99; AW18; ABC+20], where the input consists of several copies of $\sum_{x \in E} \sqrt{\Pr[X = x]} |x\rangle$ (we do not have access to a unitary preparing that state). We show in Theorem 6.6 that no quantum speedup is achievable for our problem in the latter setting.

3 Quantile estimation

In this section, we present a quantum algorithm for estimating the quantiles of a finite random variable X . This is a key ingredient for the sub-Gaussian estimator of Section 4. For the convenience of reading, we define a quantile in the following non-standard way (the cumulative distribution function is replaced with its complement).

Definition 3.1 (QUANTILE). Given a discrete random variable X and a real $p \in [0, 1]$, the *quantile* of order p is the number $Q(p) = \sup\{x \in \mathbb{R} : \Pr[X \geq x] \geq p\}$.

Our result is inspired by the minimum finding algorithm of Dürr and Høyer [DH96] and its generalization in [AGGW20]. The problem of estimating the quantiles of a set of numbers under the *uniform* distribution was studied before by Nayak and Wu [NW99; Nay99]. We differ from that work by allowing arbitrary distributions, and by not using the amplitude estimation algorithm. On the other hand, we restrict ourselves to finding a constant factor estimate, whereas [NW99; Nay99] can achieve any wanted accuracy.

The idea behind our algorithm is rather simple: if we compute a sequence of values $-\infty = y_0 \leq y_1 \leq y_2 \leq y_3 \leq \dots$ where each y_{j+1} is sampled from the distribution of X conditioned on $y_{j+1} \geq y_j$, then when $j \simeq \log(1/p)$ the value of y_j should be close to the quantile $Q(p)$. The complexity of sampling each y_j is on the order of $1/\Pr[X \geq y_j]$ classically, but it can be done quadratically faster in the quantum setting. We analyze a slightly different algorithm, where the sequence of samples is strictly increasing and instead of stopping after roughly $\log(1/p)$ iterations we count the number of experiments performed by the algorithm and stop when it reaches a value close to $1/\sqrt{p}$. This requires showing that the times T_j spent on sampling y_j is neither too large nor too small with high probability, which is proved in the next lemma.

Lemma 3.2. *There is a quantum algorithm such that, given a q -random variable X and a value $x \in \mathbb{R} \cup \{-\infty, +\infty\}$, it outputs a sample y from the probability distribution of X*

conditioned on $y > x$. If we let T denote the number of quantum experiments performed by this algorithm, then there exist two universal constants $c_0 < c_1$ such that $\mathbb{E}[T] \leq c_1/\sqrt{\Pr[X > x]}$ and $\Pr[T < c_0/\sqrt{\Pr[X > x]}] \leq 1/10$.

Proof. Let (\mathcal{H}, U, M) be a q -variable generating X . We use the comparison oracle $C_{x, +\infty}$ from Assumption A to construct the unitary $V = C_{x, +\infty}(U \otimes I)$ acting on $\mathcal{H} \otimes \mathbb{C}^2$. By definition of $C_{x, +\infty}$ and U (Section 2), we have that $V|\mathbf{0}\rangle = \sum_{\omega \in \Omega: X(\omega) \leq x} \sqrt{p(\omega)}|\omega\rangle|0\rangle + \sum_{\omega \in \Omega: X(\omega) > x} \sqrt{p(\omega)}|\omega\rangle|1\rangle = \sqrt{1 - \Pr[X > x]}|\phi_0\rangle|0\rangle + \sqrt{\Pr[X > x]}|\phi_1\rangle|1\rangle$ for some unit states $|\phi_0\rangle, |\phi_1\rangle$ where $|\phi_1\rangle = \frac{1}{\sqrt{\Pr[X > x]}} \sum_{\omega: X(\omega) > x} \sqrt{p(\omega)}|\omega\rangle$. The algorithm for sampling y conditioned on $y > x$ consists of two steps. First, we use the sequential amplitude amplification algorithm Seq-AAmp($V, I \otimes |1\rangle\langle 1|$) from Theorem A.2 on V to obtain the state $|\phi_1\rangle$. Next, we measure $|\phi_1\rangle$ according to M . The claimed properties follow directly from Theorem A.2. \square

We use the next formula for the probability that a value x occurs in the sequence $(y_j)_j$ defined before. This lemma is adapted from [DH96, Lemma 1].

Lemma 3.3 (Lemma 47 in [AGGW20]). *Let X be a discrete random variable. Consider the increasing sequence of random variables Y_0, Y_1, Y_2, \dots where Y_0 is a fixed value and Y_{j+1} for $j \geq 0$ is a sample drawn from X conditioned on $Y_{j+1} > Y_j$. Then, for any $x, y \in \mathbb{R}$,*

$$\Pr[x \in \{Y_1, Y_2, \dots\} \mid Y_0 = y] = \begin{cases} \Pr[X = x \mid X \geq x] & \text{when } x > y, \\ 0 & \text{otherwise.} \end{cases}$$

The quantile estimation algorithm is described in Algorithm 1 and its analysis is provided in the next theorem.

1. Repeat the following steps for $i = 1, 2, \dots, \lceil 6 \log(1/\delta) \rceil$.
 - (a) Set $y_0 = -\infty$ and initialize a counter $C = 0$ that is incremented each time a quantum experiment is performed.
 - (b) Set $j = 1$. Repeat the following process and interrupt it when $C = c'/\sqrt{p}$ (where c' is a constant chosen in the proof of Theorem 3.4): sample an element y_{j+1} from X conditioned on $y_{j+1} > y_j$ by using the algorithm of Lemma 3.2, set $j \leftarrow j + 1$.
 - (c) Set $\tilde{Q}^{(i)} = y_j$.
2. Output $\tilde{Q} = \text{median}(\tilde{Q}^{(1)}, \dots, \tilde{Q}^{(\lceil 6 \log(1/\delta) \rceil)})$.

Algorithm 1: Quantile estimation algorithm, $\text{Quantile}(X, p, \delta)$.

Theorem 3.4 (QUANTILE ESTIMATION). *Let X be a q -random variable. Given two reals $p, \delta \in (0, 1)$, the approximate quantile \tilde{Q} produced by the quantile estimation algorithm $\text{Quantile}(X, p, \delta)$ (Algorithm 1) satisfies*

$$Q(p) \leq \tilde{Q} \leq Q(cp)$$

with probability at least $1 - \delta$, where $c < 1$ is a universal constant. The algorithm performs $O\left(\frac{\log(1/\delta)}{\sqrt{p}}\right)$ quantum experiments.

Proof. Let c_0, c_1 be the universal constants mentioned in Lemma 3.2, and set $c = c_0^2/(c_1^2\sqrt{191})$ and $c' = 190c_1$. Fix i and consider the sequence $(y_j)_{j \geq 0}$ that would be computed during the i -th execution of steps 1.a-1.c if the stopping condition on C was removed. We prove that immediately

after the c'/\sqrt{p} -th quantum experiment is performed (which may occur during the computation of y_{j+1}), the current value of y_j satisfies $Q(p) \leq y_j \leq Q(cp)$ with probability at least $(9/10)^2$. The analysis is done in two parts.

First, let $x^- = Q(p)$ and denote by T^- the number of experiments performed until y_j becomes larger than or equal to x^- . According to Lemma 3.3, the probability that a given x occurs in the sequence $(y_j)_{j \geq 0}$ is equal to $\Pr[X = x | X \geq x]$. Moreover, using Lemma 3.2, the expected number of experiments performed at step 1.b when $y_j = x$ is at most $c_1/\sqrt{\Pr[X > x]}$. Consequently, we have

$$\mathbb{E}[T^-] \leq c_1 \sum_{x < x^-} \frac{\Pr[X = x | X \geq x]}{\sqrt{\Pr[X > x]}}.$$

Suppose that $Q(1) \neq x^-$ (otherwise $T^- = 0$). We upper bound the above sum by splitting it into several parts as follows. Define $Q_k = Q(2^{-k})$ for $k \geq 0$ and let ℓ be the largest integer such that $Q(2^{-\ell}) < x^-$. For each $1 \leq k \leq \ell$ such that $Q_{k-1} \neq Q_k$, we have

$$\begin{aligned} \sum_{Q_{k-1} \leq x < Q_k} \frac{\Pr[X = x | X \geq x]}{\sqrt{\Pr[X > x]}} &\leq \frac{1}{\sqrt{\Pr[X > Q_{k-1}]}} + \sum_{Q_{k-1} < x < Q_k} \frac{\Pr[X = x]}{\Pr[X > x]^{3/2}} \\ &\leq \frac{1}{\sqrt{\Pr[X \geq Q_k]}} + \frac{\Pr[X > Q_{k-1}]}{\Pr[X \geq Q_k]^{3/2}} \\ &\leq \frac{1}{\sqrt{2^{-k}}} + \frac{2^{-(k-1)}}{2^{-3k/2}} \\ &\leq 2^{k/2+2}. \end{aligned}$$

Similarly, $\sum_{Q_\ell \leq x < x^-} \frac{\Pr[X = x | X \geq x]}{\sqrt{\Pr[X > x]}} \leq 2^{\ell/2} + 2^{-\ell+1}/p^{3/2}$. Thus, $\mathbb{E}[T^-] \leq c_1(\sum_{k=1}^{\ell} 2^{k/2+2} + 2^{\ell/2} + 2^{-\ell+1}/p^{3/2}) \leq 19c_1/\sqrt{p}$ where we used that $\log(1/p) - 1 \leq \ell < \log(1/p)$ since $Q_\ell < Q(p) \leq Q_{\ell+1}$. By Markov's inequality, $\Pr[T^- \leq 190c_1/\sqrt{p}] \geq 9/10$.

Secondly, let $x^+ = Q(cp)$ and denote by T^+ the number of experiments performed at step 1.b to sample y_{j+1} when $y_j \geq x^+$. According to Lemma 3.2, we have $\Pr[T^+ \geq c_0/\sqrt{\Pr[X > y_j]}] \geq 9/10$. Moreover, $\Pr[X > y_j] \leq cp = c_0^2/(c_1^2\sqrt{191})p$ by definition of x^+ . Thus, $\Pr[T^+ \geq 191c_1/\sqrt{p}] \geq 9/10$.

We conclude that step 1.b is interrupted when the value $\tilde{Q}^{(i)}$ satisfies $Q(p) \leq \tilde{Q}^{(i)} \leq Q(cp)$ with probability at least $(9/10)^2$. Thus, by the Chernoff bound, the output \tilde{Q} satisfies $Q(p) \leq \tilde{Q} \leq Q(cp)$ with probability at least $1 - \delta$. The total number of experiments is guaranteed to be $O(\log(1/\delta)/\sqrt{p})$ by our use of the counter C . \square

4 Sub-Gaussian estimator

In this section, we present the main quantum algorithm for estimating the mean of a random variable with a near-quadratic speedup over the classical sub-Gaussian estimators. Our result uses the following *Bernoulli estimator*, which is a well-known adaptation of the amplitude estimation algorithm to the mean estimation problem [BHMT02; Ter99; Mon15]. The Bernoulli estimator allows us to estimate the mean of the truncated random variable $X\mathbb{1}_{a < X \leq b}$ for any a, b .

Proposition 4.1 (BERNOULLI ESTIMATOR). *There exists a quantum algorithm, called the Bernoulli estimator, with the following properties. Let X be a q -random variable and set as input a time parameter $n \geq 0$, two range values $0 \leq a < b$, and a real $\delta \in (0, 1)$ such that $n \geq \log(1/\delta)$. Then, the Bernoulli estimator $\text{BernEst}(X, n, a, b, \delta)$ outputs a mean estimate $\tilde{\mu}_{a,b}$ of $\mu_{a,b} = \mathbb{E}[X\mathbb{1}_{a < X \leq b}]$ such that $|\tilde{\mu}_{a,b} - \mu_{a,b}| \leq \frac{\sqrt{b\mu_{a,b}} \log(1/\delta)}{n} + \frac{b \log(1/\delta)^2}{n^2}$. It performs $O(n)$ quantum experiments.*

Proof. Let (\mathcal{H}, U, M) be a q-variable generating X . Using the rotation oracle $R_{a,b}$ from Assumption B, we define the unitary algorithm $V = R_{a,b}(U \otimes I)$ acting on $\mathcal{H} \otimes \mathbb{C}^2$. In order to simplify notations, let us first assume that the random variable X is distributed in the interval (a, b) . Then, $\mu = \mu_{a,b}$ and by definition of $R_{a,b}$ and U (Section 2) we have,

$$\begin{aligned} V|\mathbf{0}\rangle &= \sum_{\omega \in \Omega} \sqrt{p(\omega)} |\omega\rangle \left(\sqrt{1 - \frac{X(\omega)}{b}} |0\rangle + \sqrt{\frac{X(\omega)}{b}} |1\rangle \right) \\ &= \sqrt{1 - \frac{\mu}{b}} \left(\sum_{\omega \in \Omega} \sqrt{\frac{p(\omega)(b - X(\omega))}{b - \mu}} |\omega\rangle \right) |0\rangle + \sqrt{\frac{\mu}{b}} \left(\sum_{\omega \in \Omega} \sqrt{\frac{p(\omega)X(\omega)}{\mu}} |\omega\rangle \right) |1\rangle. \end{aligned}$$

Thus, there exist some unit states $|\psi_0\rangle, |\psi_1\rangle$ such that $V|\mathbf{0}\rangle = \sqrt{1 - \frac{\mu}{b}} |\psi_0\rangle + \sqrt{\frac{\mu}{b}} |\psi_1\rangle$ and $(I \otimes |1\rangle\langle 1|)V|\mathbf{0}\rangle = \sqrt{\frac{\mu}{b}} |\psi_1\rangle$. If X takes values outside the interval (a, b) then the same result holds with $\mu_{a,b}$ in place of μ and a different definition of $|\psi_0\rangle, |\psi_1\rangle$.

Consider the output \tilde{v} of the amplitude estimation algorithm $\text{AEst}(V, \Pi, \lceil \frac{2\pi n}{\log(1/\delta)} \rceil)$ (Theorem A.3) where $\Pi = I \otimes |1\rangle\langle 1|$. Then, the estimate $b\tilde{v}$ satisfies the statement of the proposition with probability $8/\pi^2$ by Theorem A.3. The Bernoulli estimator consists of running $\lceil 6 \log(1/\delta) \rceil$ copies of $\text{AEst}(V, \Pi, \lceil \frac{2\pi n}{\log(1/\delta)} \rceil)$ and outputting the median of the results. The success probability is at least $1 - \delta$ by the Chernoff bound. \square

The Bernoulli estimator can estimate the mean of a non-negative q-random variable X by setting $a = 0$ and $b = \max X$. However, its performance is worse than that of the classical sub-Gaussian estimators when the maximum of X is large compared to its variance. Our quantum sub-Gaussian estimator (Algorithm 2) uses the Bernoulli estimator in a more subtle way, and in combination with the quantile estimation algorithm.

1. Set $k = \log n$ and $m = dn\sqrt{\log n} \frac{\log(9k/\delta)}{\log(1/\delta)}$, where $d > 1$ is a constant chosen in the proof of Theorem 4.2 (if k is not an integer, round n to the next power of two).
2. Compute the median η of $\lceil 30 \log(2/\delta) \rceil$ classical samples from X and define the non-negative random variables

$$Y^+ = (X - \eta) \mathbb{1}_{X \geq \eta} \quad \text{and} \quad Y^- = -(X - \eta) \mathbb{1}_{X \leq \eta}.$$

3. Compute an estimate $\tilde{\mu}_{Y^+}$ of $\mathbb{E}[Y^+]$ and an estimate $\tilde{\mu}_{Y^-}$ of $\mathbb{E}[Y^-]$ by executing the following steps with $Y := Y^+$ and $Y := Y^-$ respectively:
 - (a) Compute an estimate \tilde{Q} of the quantile of order $p = \left(\frac{\log(1/\delta)}{6n}\right)^2$ of Y with failure probability $\delta/8$ by using the **quantile estimation** algorithm $\text{Quantile}(Y, p, \delta/8)$.
 - (b) Define $a_{-1} = 0$ and $a_\ell = \frac{2^\ell}{n} \tilde{Q}$ for $\ell \geq 0$. Compute an estimate $\tilde{\mu}_\ell$ of $\mathbb{E}[Y \mathbb{1}_{a_{\ell-1} < Y \leq a_\ell}]$ with failure probability $\delta/(9k)$ for each $0 \leq \ell \leq k$, by using the **Bernoulli estimator** $\text{BernEst}(Y, m, a_{\ell-1}, a_\ell, \delta/(9k))$ with m quantum experiments.
 - (c) Set $\tilde{\mu}_Y = \sum_{\ell=0}^k \tilde{\mu}_\ell$.
4. Output $\tilde{\mu} = \eta + \tilde{\mu}_{Y^+} - \tilde{\mu}_{Y^-}$.

Algorithm 2: Sub-Gaussian estimator, $\text{SubGaussEst}(X, n, \delta)$.

Theorem 4.2 (SUB-GAUSSIAN ESTIMATOR). *Let X be a q -random variable with mean μ and variance σ^2 . Given a time parameter n and a real $\delta \in (0, 1)$ such that $n \geq \log(1/\delta)$, the sub-Gaussian estimator $\text{SubGaussEst}(X, n, \delta)$ (Algorithm 2) outputs a mean estimate $\tilde{\mu}$ such that,*

$$\Pr \left[|\tilde{\mu} - \mu| \leq \frac{\sigma \log(1/\delta)}{n} \right] \geq 1 - \delta.$$

The algorithm performs $O(n \log^{3/2}(n) \log \log(n))$ quantum experiments.

Proof. First, by standard concentration inequalities, the median η computed at step 2 satisfies $|\eta - \mu| \leq 2\sigma$ with probability at least $1 - \delta/2$. Moreover, if $|\eta - \mu| \leq 2\sigma$ then $\sqrt{\mathbb{E}[(X - \eta)^2]} = \sqrt{\mathbb{E}[(X - \mu + \mu - \eta)^2]} \leq \sqrt{\mathbb{E}[(X - \mu)^2]} + |\mu - \eta| \leq 3\sigma$, by using the triangle inequality. Below we prove that for any non-negative random variable Y the estimate $\tilde{\mu}_Y$ of $\mu_Y = \mathbb{E}[Y]$ computed at step 3 satisfies

$$|\tilde{\mu}_Y - \mu_Y| \leq \frac{\sqrt{\mathbb{E}[Y^2]} \log(1/\delta)}{5n} \quad (4)$$

with probability at least $1 - \delta/4$. Using the fact that $X = \eta + Y_+ - Y_-$ and $(X - \eta)^2 = Y_+^2 + Y_-^2$, we can conclude that

$$|\tilde{\mu} - \mu| \leq \frac{\left(\sqrt{\mathbb{E}[Y_+^2]} + \sqrt{\mathbb{E}[Y_-^2]} \right) \log(1/\delta)}{5n} \leq \frac{\sqrt{2\mathbb{E}[(X - \eta)^2]} \log(1/\delta)}{5n} \leq \frac{\sigma \log(1/\delta)}{n}$$

with probability at least $1 - \delta$. The algorithm performs $O(\log(1/\delta)) \leq O(n)$ classical experiments during step 2, $O(\log(1/\delta)/\sqrt{p}) \leq O(n)$ quantum experiments during step 3.a, and $O(km) \leq O(n \log^{3/2}(n) \log \log(n))$ quantum experiments during step 3.b.

We now turn to the proof of Equation (4). We make the assumption that all the subroutines used in step 3 are successful, which is the case with probability at least $(1 - \delta/8)(1 - \delta/(9k))^{k+1} \geq 1 - \delta/4$. First, according to Theorem 3.4, we have $Q(p) \leq \tilde{Q} \leq Q(cp)$ for some universal constant c . It implies that $cp \leq \Pr[Y \geq Q(cp)] \leq \Pr[Y \geq \tilde{Q}] \leq \mathbb{E}[Y^2]/\tilde{Q}^2$, where the first two inequalities are by definition of the quantile function Q , and the last inequality is a standard fact. Consequently, by our choice of p ,

$$\tilde{Q} \leq \frac{6n\sqrt{\mathbb{E}[Y^2]}}{\sqrt{c} \log(1/\delta)}. \quad (5)$$

Next, we upper bound the expectation of the part of Y that is above the largest threshold $a_k = \tilde{Q}$ considered in step 3.b. By Cauchy-Schwarz' inequality, we have $\mathbb{E}[Y \mathbb{1}_{Y > \tilde{Q}}] \leq \sqrt{\mathbb{E}[Y^2] \Pr[Y > \tilde{Q}]}$. Moreover, by definition of Q , $\Pr[Y > \tilde{Q}] \leq \Pr[Y > Q(p)] \leq p$. Thus,

$$\mathbb{E}[Y \mathbb{1}_{Y > \tilde{Q}}] \leq \frac{\sqrt{\mathbb{E}[Y^2]} \log(1/\delta)}{6n}. \quad (6)$$

The expectation of Y is decomposed into the sum $\mu_Y = \sum_{\ell=0}^k \mu_\ell + \mathbb{E}[Y \mathbb{1}_{Y > a_k}]$, where $\mu_\ell = \mathbb{E}[Y \mathbb{1}_{a_{\ell-1} < Y \leq a_\ell}]$ is estimated at step 3.b. We have $|\tilde{\mu}_\ell - \mu_\ell| \leq \frac{\sqrt{a_\ell \mu_\ell} \log(1/\delta)}{dn\sqrt{\log n}} + \frac{a_\ell \log(1/\delta)^2}{d^2 n^2 \log n}$ for all

$0 \leq \ell \leq k$ according to Proposition 4.1. Thus, by the triangle inequality,

$$\begin{aligned}
|\tilde{\mu}_Y - \mu_Y| &\leq \sum_{\ell=0}^k |\tilde{\mu}_\ell - \mu_\ell| + \mathbb{E}[Y \mathbb{1}_{Y > a_k}] \\
&\leq \sum_{\ell=0}^k \frac{\sqrt{a_\ell \mu_\ell} \log(\frac{1}{\delta})}{dn \sqrt{\log n}} + \sum_{\ell=0}^k \frac{a_\ell \log(\frac{1}{\delta})^2}{d^2 n^2 \log n} + \mathbb{E}[Y \mathbb{1}_{Y > a_k}] \\
&\leq \frac{\tilde{Q} \log(\frac{1}{\delta})}{dn^2 \sqrt{\log n}} + \sum_{\ell=1}^k \frac{\sqrt{2\mathbb{E}[Y^2 \mathbb{1}_{a_{\ell-1} < Y \leq a_\ell}] \log(\frac{1}{\delta})}}{dn \sqrt{\log n}} + \frac{2\tilde{Q} \log(\frac{1}{\delta})^2}{d^2 n^2 \log n} + \mathbb{E}[Y \mathbb{1}_{Y > a_k}] \\
&\leq \frac{\sqrt{2k} \sqrt{\sum_{\ell=1}^k \mathbb{E}[Y^2 \mathbb{1}_{a_{\ell-1} < Y \leq a_\ell}] \log(\frac{1}{\delta})}}{dn \sqrt{\log n}} + \frac{3\tilde{Q} \log(\frac{1}{\delta})^2}{dn^2 \sqrt{\log n}} + \mathbb{E}[Y \mathbb{1}_{Y > a_k}] \\
&\leq \frac{\sqrt{2k} \sqrt{\mathbb{E}[Y^2]} \log(\frac{1}{\delta})}{dn \sqrt{\log n}} + \frac{3\tilde{Q} \log(\frac{1}{\delta})^2}{dn^2 \sqrt{\log n}} + \mathbb{E}[Y \mathbb{1}_{Y > a_k}] \\
&\leq \frac{\sqrt{2} \sqrt{\mathbb{E}[Y^2]} \log(\frac{1}{\delta})}{dn} + \frac{18 \sqrt{\mathbb{E}[Y^2]} \log(\frac{1}{\delta})}{\sqrt{c} dn \sqrt{\log n}} + \frac{\sqrt{\mathbb{E}[Y^2]} \log(\frac{1}{\delta})}{6n} \\
&\leq \frac{\sqrt{\mathbb{E}[Y^2]} \log(\frac{1}{\delta})}{5n}
\end{aligned}$$

where the third step uses $a_0 \mu_0 \leq a_0^2 = (\tilde{Q}/n)^2$ and $a_\ell \mu_\ell \leq (a_\ell/a_{\ell-1}) \mathbb{E}[Y^2 \mathbb{1}_{a_{\ell-1} < Y \leq a_\ell}] \leq 2\mathbb{E}[Y^2 \mathbb{1}_{a_{\ell-1} < Y \leq a_\ell}]$ when $\ell \geq 1$, the fourth step uses the Cauchy–Schwarz inequality, the sixth step uses Equations (5) and (6), and in the last step we choose $d = 600/\sqrt{c}$. \square

5 (ϵ, δ) -Estimators

We study the (ϵ, δ) -approximation problem under two different scenarios. First, we consider the case where we know an upper bound Δ on the coefficient of variation $|\sigma/\mu|$. As a direct consequence of Theorem 4.2 we obtain the following estimator that subsumes a similar result shown in [HM19] for non-negative random variables.

Corollary 5.1 (Relative estimator). *There exists a quantum algorithm with the following properties. Let X be a q -random variable with mean μ and variance σ^2 , and set as input a value $\Delta \geq |\sigma/\mu|$ and two reals $\epsilon, \delta \in (0, 1)$. Then, the algorithm outputs a mean estimate $\tilde{\mu}$ such that $\Pr[|\tilde{\mu} - \mu| > \epsilon|\mu|] \leq \delta$ and it performs $\tilde{O}(\frac{\Delta}{\epsilon} \log(1/\delta))$ quantum experiments.*

Proof. The algorithm runs the **sub-Gaussian estimator** $\text{SubGaussEst}(X, \frac{\Delta}{\epsilon} \log(1/\delta), \delta)$. \square

Next, we construct a parameter-free estimator that performs $\tilde{O}((\frac{\sigma}{\epsilon\mu} + \frac{1}{\sqrt{\epsilon\mu}}) \log(1/\delta))$ quantum experiments in expectation for any random variable distributed in $[0, 1]$. We follow an approach similar to the classical \mathcal{AA} algorithm described in [DKLR00]. We first give a sequential estimator that approximates the mean with constant relative error and that performs $O(1/\sqrt{\mu})$ quantum experiments in expectation. We use the term “sequential” in reference to sequential analysis techniques. The classical counterpart of this estimator is the Stopping Rule Algorithm in [DKLR00].

Proposition 5.2 (SEQUENTIAL BERNOULLI ESTIMATOR). *There is an algorithm, called the sequential Bernoulli estimator, with the following properties. Let X be a q -random variable distributed in $[0, 1]$ with mean μ . Then, the sequential Bernoulli estimator $\text{Seq-BernEst}(X)$ outputs an estimate $\tilde{\mu}$ and performs a number T of quantum experiments such that,*

1. *There is a universal constant $c \in (0, 1)$ such that $\Pr[|\tilde{\mu} - \mu| \leq c\mu] \geq 7/8$.*

2. There is a universal constant c' such that $\mathbb{E}[T^2] = \mathbb{E}[1/\tilde{\mu}] \leq c'/\mu$.
3. There is a universal constant c'' such that $\mathbb{E}[\sqrt{\tilde{\mu}}] \leq c''\sqrt{\mu}$.

Proof. The algorithm is identical to the one of Proposition 4.1 with $a = 0$ and $b = 1$, except that the amplitude estimation algorithm is replaced with the sequential amplitude estimation algorithm (Theorem A.4). The algorithm inherits the properties proved in Theorem A.4. \square

The expected number of experiments performed by the sequential Bernoulli estimator is $\mathbb{E}[T] \leq \sqrt{\mathbb{E}[T^2]} \leq 1/\sqrt{\mu}$. The output $\tilde{\mu}$ of the sequential Bernoulli estimator can be used in the Bernoulli estimator (Proposition 4.1) with parameter $n = 8 \log(1/\delta)/(\epsilon\sqrt{\tilde{\mu}})$ to solve the (ϵ, δ) -approximation problem. However, the expected number of experiments performed with this approach is $O(\log(1/\delta)/(\epsilon\sqrt{\mu}))$. We propose a better algorithm with an improved dependence on ϵ . The algorithm uses the sequential Bernoulli estimator and the sub-Gaussian estimator.

1. For $i = 1, \dots, 32 \log(1/\delta)$:
 - (a) Compute an estimate $\tilde{\mu}_X$ of $\mu = \mathbb{E}[X]$ by using the sequential Bernoulli estimator $\text{Seq-BernEst}(X)$ (Proposition 5.2).
 - (b) Let Y denote the random variable $(X - X')^2/2$ where X' is independent from X and identically distributed. Compute an estimate $\tilde{\mu}_Y$ of $\mu_Y = \mathbb{E}[Y]$ by using the sequential Bernoulli estimator $\text{Seq-BernEst}(Y)$ (Proposition 5.2). Stop the computation if it performs more than $\frac{c_1}{\sqrt{\epsilon\tilde{\mu}_X}}$ quantum experiments (where c_1 is a constant chosen in the proof of Theorem 5.3) and set $\tilde{\mu}_Y = 0$.
 - (c) Compute a second estimate $\tilde{\mu}_X^{(i)}$ of μ by using the sub-Gaussian estimator $\text{SubGaussEst}(X, n, 15/16)$ (Theorem 4.2) with $n = c_2 \max\left(\frac{\sqrt{\tilde{\mu}_Y}}{\epsilon\tilde{\mu}_X}, \frac{1}{\sqrt{\epsilon\tilde{\mu}_X}}\right)$ (where c_2 is a constant chosen in the proof of Theorem 5.3).
2. Output $\tilde{\mu} = \text{median}\left(\tilde{\mu}_X^{(1)}, \dots, \tilde{\mu}_X^{(32 \log(1/\delta))}\right)$.

Algorithm 3: Sequential (ϵ, δ) -estimator.

Theorem 5.3 (SEQUENTIAL RELATIVE ESTIMATOR). *Let X be a q -random variable distributed in $[0, 1]$ with mean μ and variance σ^2 . Given two reals $\epsilon, \delta \in (0, 1)$ the estimate $\tilde{\mu}$ output by the sequential relative estimator (Algorithm 3) satisfies $\Pr[|\tilde{\mu} - \mu| > \epsilon\mu] \leq \delta$. The algorithm performs $\tilde{O}\left(\left(\frac{\sigma}{\epsilon\mu} + \frac{1}{\sqrt{\epsilon\mu}}\right) \log(1/\delta)\right)$ quantum experiments in expectation.*

Proof. We prove that, for a fixed value of i , the estimate $\tilde{\mu}_X^{(i)}$ computed at step 1.c satisfies $\Pr[|\tilde{\mu}_X^{(i)} - \mu| \leq \epsilon\mu] \geq 5/8$ and the number of experiments performed during its computation is $\tilde{O}\left(\left(\frac{\sigma}{\epsilon\mu} + \frac{1}{\sqrt{\epsilon\mu}}\right)\right)$ in expectation. The theorem follows by the Chernoff bound and the linearity of expectation.

Let c, c', c'' denote the constants mentioned in Proposition 5.2, and set $c_1 = 16c'\sqrt{(1+c)}$ and $c_2 = 4(1+c)/\sqrt{1-c}$. We assume that $|\tilde{\mu}_X - \mu| \leq c\mu$ at step 1.a, which is the case with probability at least $7/8$ by Proposition 5.2. The analysis of steps 1.b and 1.c is split into two cases to show that $\Pr[|\tilde{\mu}_X^{(i)} - \mu| \leq \epsilon\mu] \geq 5/8$. First, if $\sigma \leq \sqrt{\epsilon\mu}$, then we can ignore step 1.b and consider the second term in the max at step 1.c. By Theorem 4.2, the estimate $\tilde{\mu}_X^{(i)}$ satisfies $|\tilde{\mu}_X^{(i)} - \mu| \leq \frac{4\sigma}{c_2/\sqrt{\epsilon\tilde{\mu}_X}} \leq \frac{4\sqrt{1+c}}{c_2}\epsilon\mu \leq \epsilon\mu$ with probability $15/16$. Secondly, if $\sigma \geq \sqrt{\epsilon\mu}$, then by

Proposition 5.2 and the fact that $\mu_Y = \sigma^2$, the estimate $\tilde{\mu}_Y$ computed at step 1.b satisfies $|\tilde{\mu}_Y - \sigma^2| \leq c\sigma^2$ with probability 7/8 if we remove the stopping condition. Since we assumed that $\tilde{\mu}_X \leq (1+c)\mu$, the computation is interrupted if it performs more than $\frac{c_1}{\sqrt{\epsilon\tilde{\mu}_X}} \geq \frac{c_1}{\sqrt{(1+c)\mu_Y}} = \frac{16c'}{\sqrt{\mu_Y}}$ experiments. However, by Proposition 5.2 and Markov's inequality, the number of experiments performed by the sequential Bernoulli estimator at step 1.b is at most $16c'/\sqrt{\mu_Y}$ with probability at least 15/16. Consequently, we can assume that $\tilde{\mu}_Y \geq (1-c)\sigma^2$ with success probability at least $7/8 \cdot 15/16$. In this case, by considering the first term in the max at step 1.c, the estimate $\tilde{\mu}_X^{(i)}$ satisfies $|\tilde{\mu}_X^{(i)} - \mu| \leq \frac{4\sigma}{c_2\sqrt{\tilde{\mu}_Y/(\epsilon\tilde{\mu}_X)}} \leq \frac{4(1+c)}{c_2\sqrt{1-c}}\epsilon\mu \leq \epsilon\mu$ with probability 15/16. The overall success probability is at least $(7/8)^2(15/16)^2 \geq 5/8$.

We now analyse the expected number of quantum experiments performed during the computation of $\tilde{\mu}_X^{(i)}$. Step 1.a performs $O(1/\sqrt{\mu})$ experiments in expectation by Proposition 5.2. Step 1.b is stopped after $O(1/(\sqrt{\epsilon\tilde{\mu}}))$ experiments in expectation since $\mathbb{E}[1/\sqrt{\tilde{\mu}_X}] \leq O(1/\sqrt{\mu})$ by Proposition 5.2. Step 1.c performs $\tilde{O}\left(\max\left(\frac{\sqrt{\tilde{\mu}_Y}}{\epsilon\tilde{\mu}_X}, \frac{1}{\sqrt{\epsilon\tilde{\mu}_X}}\right)\right)$ experiments by Theorem 4.2. The estimates $\tilde{\mu}_Y$ and $\tilde{\mu}_X$ are independent if we ignore the stopping condition at step 1.b, in which case $\mathbb{E}\left[\frac{\sqrt{\tilde{\mu}_Y}}{\tilde{\mu}_X}\right] = \mathbb{E}\left[\frac{1}{\tilde{\mu}_X}\right]\mathbb{E}[\sqrt{\tilde{\mu}_Y}] \leq O\left(\frac{\sigma}{\mu}\right)$ by Proposition 5.2. The stopping condition can only decrease this quantity. Thus, step 1.c performs $\tilde{O}\left(\max\left(\frac{\sigma}{\epsilon\mu}, \frac{1}{\sqrt{\epsilon\mu}}\right)\right)$ experiments in expectation \square

6 Lower bounds

We prove several lower bounds for the mean estimation problem under different scenarios. In Section 6.1, we study the number of experiments that must be performed to estimate the mean with a sub-Gaussian error rate. In Section 6.2, we study the number of experiments needed to solve the (ϵ, δ) -approximation problem. Finally, in Section 6.3, we consider the mean estimation problem in the state-based model, where the input consists of several copies of a quantum state encoding a distribution.

6.1 Sub-Gaussian estimation

We show that the quantum sub-Gaussian estimator described in Theorem 4.2 is optimal up to a polylogarithmic factor. We make use of the following lower bound for Quantum Search in the small-error regime.

Proposition 6.1 (Theorem 4 in [BCWZ99]). *Let $N > 0$, $1 \leq K \leq 0.9N$ and $\delta \geq 2^{-N}$. Let $T(N, K, \delta)$ be the minimum number of quantum queries any algorithm must use to decide with failure probability at most δ whether a function $f : [N] \rightarrow \{0, 1\}$ has 0 or K preimages of 1. Then, $T(N, K, \delta) \geq \Omega(\sqrt{N/K} \log(1/\delta))$.*

We construct two particular probability distributions that allow us to reduce the Quantum Search problem to the sub-Gaussian mean estimation problem.

Theorem 6.2. *Let $n > 1$ and $\delta \in (0, 1)$ such that $n \geq 2 \log(1/\delta)$. Fix $\sigma > 0$ and consider the family \mathcal{P}_σ of all q -random variables with variance σ^2 . Let $T(n, \sigma, \delta)$ be the minimum number of quantum experiments any algorithm must perform to compute with failure probability at most δ a mean estimate $\tilde{\mu}$ such that $|\tilde{\mu} - \mu| \leq \frac{\sigma \log(1/\delta)}{n}$ for any $X \in \mathcal{P}_\sigma$ with mean μ . Then, $T(n, \sigma, \delta) \geq \Omega(n)$.*

Proof. Let $m = \frac{n}{\log(1/\delta)}$ and $b = \frac{m}{\sqrt{1-1/m^2}}\sigma$. We define the probability distribution p_0 with support $\{0, b\}$ that takes value b with probability $\frac{1}{m^2}$. Similarly, we define the probability distribution p_1 with support $\{0, -b\}$ that takes value $-b$ with probability $\frac{1}{m^2}$. The variance of

each distribution is equal to σ^2 . Moreover, the means μ_0 and μ_1 of the two distributions satisfy that,

$$\mu_0 - \mu_1 > 2 \frac{\sigma \log(1/\delta)}{n}. \quad (7)$$

Let N, K be two integers such that $N \geq \log(1/\delta)$ and $K/N = 1/m^2$ (assuming m is rational). Let F_0 be the family of all functions $f : [N] \rightarrow \{0, 1\}$ with exactly K preimages of 1. Similarly, let F_1 be the family of all functions $f : [N] \rightarrow \{-1, 0\}$ with exactly K preimages of -1 . By using Proposition 6.1, it is easy to see that any algorithm that can distinguish between $f \in F_0$ and $f \in F_1$ with success probability $1 - \delta$ must use at least $\Omega(\sqrt{N/K} \log(1/\delta)) = \Omega(m \log(1/\delta)) = \Omega(n)$ quantum queries to f . We associate with each function $f \in F_0 \cup F_1$ the q-variable $(\mathcal{H}, U, M)_f$ where $\mathcal{H} = \mathbb{C}^{N+2}$, $U|0\rangle = \frac{1}{\sqrt{N}} \sum_{x \in [N]} |x\rangle |f(x)\rangle$, and $M = \{I \otimes |0\rangle\langle 0|, I \otimes |-1\rangle\langle -1|, I \otimes |1\rangle\langle 1|\}$. The random variable X generated by $(\mathcal{H}, U, M)_f$ is distributed according to p_0 if $f \in F_0$, and according to p_1 if $f \in F_1$. Moreover, one quantum experiment with respect to X can be simulated with one quantum query to f . Consequently, any algorithm that can distinguish between a random variable distributed according to p_0 or p_1 with success probability $1 - \delta$ must perform at least $\Omega(n)$ quantum experiments. On the other hand, by Equation (7), if an algorithm can estimate the mean with an error rate smaller than $\frac{\sigma \log(1/\delta)}{n}$ then it can distinguish between $f \in F_0$ and $f \in F_1$. Thus, $T(n, \sigma, \delta) \geq \Omega(n)$. \square

6.2 (ϵ, δ) -Estimation

We consider the (ϵ, δ) -estimation problem in the parameter-free setting, when the coefficient of variation is unknown. We make use of the next lower bound for Quantum Counting.

Proposition 6.3 (Theorem 4.2.6 in [Nay99]). *Let $N > 0$, $1 < K \leq N$ and $\epsilon \in (\frac{1}{4K}, 1)$. Consider the set of all quantum algorithms such that, given a query oracle to any function $f : [N] \rightarrow \{0, 1\}$, they return an estimate \tilde{C} of the number C of preimages of 1 in f such that $|\tilde{C} - C| \leq \epsilon C$ with probability at least $2/3$. Let $T_K(N, \epsilon)$ be the minimum number of quantum queries any such algorithm must use when the oracle has exactly K preimages of 1. Then,*

$$T_K(N, \epsilon) \geq \Omega\left(\frac{\sqrt{K(N-K)}}{\epsilon K+1} + \sqrt{\frac{N}{\epsilon K+1}}\right).$$

We obtain by a simple reduction to the above problem that the result described in Theorem 5.3 is nearly optimal.

Proposition 6.4. *Let $\epsilon \in (0, 1)$. Let $\mathcal{P}_{\mathcal{B}}$ denote the family of all q-random variables that follow a Bernoulli distribution. Consider any algorithm that takes as input $X \in \mathcal{P}_{\mathcal{B}}$ and that outputs a mean estimate $\tilde{\mu}$ such that $|\tilde{\mu} - \mathbb{E}[X]| \leq \epsilon \mathbb{E}[X]$ with probability at least $2/3$. Then, for any $\mu \in (0, 1)$, there exists $X \in \mathcal{P}_{\mathcal{B}}$ with mean μ such that the algorithm performs at least $\Omega\left(\frac{\sigma}{\epsilon \mu} + \frac{1}{\sqrt{\epsilon \mu}}\right)$ quantum experiments on input X , where $\sigma^2 = \text{Var}[X]$.*

Proof. Given $\epsilon \in (0, 1)$ and $\mu \in (0, 1)$, we choose two integers K and N such that $K > 1/(4\epsilon)$ and $K/N = \mu$ (assuming μ is rational). Similarly to the proof of Theorem 6.2, we associate with each function $f : [N] \rightarrow \{0, 1\}$ the q-variable $(\mathcal{H}, U, M)_f$ where $\mathcal{H} = \mathbb{C}^{N+2}$, $U|0\rangle = \frac{1}{\sqrt{N}} \sum_{x \in [N]} |x\rangle |f(x)\rangle$, and $M = \{I \otimes |0\rangle\langle 0|, I \otimes |1\rangle\langle 1|\}$. If a quantum algorithm can estimate the mean of any Bernoulli random variable with error ϵ and success probability $2/3$, then it can be used to count the number of preimages of 1 in f with the same accuracy. Thus, by Proposition 6.3, any such algorithm must perform at least $\Omega\left(\frac{\sqrt{K(N-K)}}{\epsilon K+1} + \sqrt{\frac{N}{\epsilon K+1}}\right) = \Omega\left(\frac{\sqrt{\mu(1-\mu)}}{\epsilon \mu+1/N} + \frac{1}{\sqrt{\epsilon \mu+1/N}}\right) = \Omega\left(\frac{\sigma}{\epsilon \mu} + \frac{1}{\sqrt{\epsilon \mu}}\right)$ quantum experiments on a q-random variable with mean μ and variance $\sigma^2 = \mu(1-\mu)$. \square

6.3 State-based estimation

We consider the *state-based* model where the input consists of several copies of a quantum state $|p\rangle = \sum_{x \in E} \sqrt{p(x)}|x\rangle$ encoding a distribution p over E . This model is weaker than the one described before, since it does not provide access to a unitary algorithm preparing $|p\rangle$. We prove that no quantum speedup is achievable in this setting. Our result uses the next lower bound on the number of copies needed to distinguish two states.

Lemma 6.5. *Let $\delta \in (0, 1)$ and consider two probability distributions p_0 and p_1 with the same finite support $E \subset \mathbb{R}$. Define the states $|\phi_0\rangle = \sum_{x \in E} \sqrt{p_0(x)}|x\rangle$ and $|\phi_1\rangle = \sum_{x \in E} \sqrt{p_1(x)}|x\rangle$. Then, the smallest integer T such that there is an algorithm that can distinguish $|\phi_0\rangle^{\otimes T}$ from $|\phi_1\rangle^{\otimes T}$ with success probability at least $1 - \delta$ satisfies $T \geq \frac{\ln(1/(4\delta))}{D(p_0\|p_1)}$, where $D(p_0\|p_1) = \sum_{x \in E} p_0(x) \ln\left(\frac{p_0(x)}{p_1(x)}\right)$ is the KL-divergence from p_0 to p_1 .*

Proof. According to Helstrom's bound [Hel69] the best success probability to distinguish between two states $|\phi\rangle$ and $|\phi'\rangle$ is $\frac{1}{2}(1 + \sqrt{1 - |\langle\phi|\phi'\rangle|^2})$. Thus, the smallest number T needed to distinguish $|\phi_0\rangle^{\otimes T}$ from $|\phi_1\rangle^{\otimes T}$ must satisfy $\frac{1}{2}(1 + \sqrt{1 - \langle\phi_0|\phi_1\rangle^{2T}}) \geq 1 - \delta$. It implies that $T \geq \frac{-\ln(1 - (1 - 2\delta)^2)}{-2\ln(\langle\phi_0|\phi_1\rangle)} \geq \frac{\ln(1/(4\delta))}{-2\ln\left(\sum_{x \in E} p_0(x) \sqrt{\frac{p_1(x)}{p_0(x)}}\right)} \geq \frac{\ln(1/(4\delta))}{\sum_{x \in E} p_0(x) \ln\left(\frac{p_0(x)}{p_1(x)}\right)} = \frac{\ln(1/(4\delta))}{D(p_0\|p_1)}$ where the second inequality uses the concavity of the logarithm function. \square

We use the above lemma to show that no quantum mean estimator can perform better than the classical sub-Gaussian estimators in the state-based input model.

Theorem 6.6. *Let $n > 1$ and $\delta \in (0, 1)$ such that $n \geq 2 \log(1/\delta)$. Fix $\sigma > 0$ and consider the family \mathcal{P}_σ of all distributions with finite support whose variance lies in the interval $[\sigma^2, 4\sigma^2]$. For any $p \in \mathcal{P}_\sigma$ with support $E \subset \mathbb{R}$, define the state $|p\rangle = \sum_{x \in E} \sqrt{p(x)}|x\rangle$. Let $T(n, \sigma, \delta)$ be the smallest integer such that there exists an algorithm that receives the state $|p\rangle^{\otimes T(n, \sigma, \delta)}$ for any $p \in \mathcal{P}_\sigma$, and that outputs an estimate $\tilde{\mu}$ of the mean μ of p such that $\Pr\left[|\tilde{\mu} - \mu| > \sqrt{\frac{\sigma^2 \log(1/\delta)}{n}}\right] \leq \delta$. Then, $T(n, \sigma, \delta) \geq \Omega(n)$.*

Proof. Let $m = \frac{n}{\log(1/\delta)}$, $b = \frac{m}{\sqrt{m-1}}\sigma$ and $\alpha = 2 \ln\left(1 + \sqrt{1 - \frac{1}{m}}\right)$. We define the two distributions p_0 and p_1 with support $E = \{0, b\}$ such that $p_0(b) = \frac{e^\alpha}{m}$ and $p_1(b) = \frac{1}{m}$. Let μ_0 and σ_0^2 (resp. μ_1 and σ_1^2) denote the expectation and the variance of p_0 (resp. p_1). Observe that $p_0, p_1 \in \mathcal{P}_\sigma$ since $\sigma_0 \in [\sigma, 2\sigma]$ and $\sigma_1 = \sigma$. Moreover, $\mu_0 - \mu_1 = \sigma \frac{e^\alpha - 1}{\sqrt{m-1}} = \sigma(e^{\alpha/2} + 1) \frac{e^{\alpha/2} - 1}{\sqrt{m-1}} > 2\sqrt{\frac{\sigma^2 \log(1/\delta)}{n}}$. Thus, we can distinguish $|p_0\rangle^{\otimes T(n, \sigma, \delta)}$ from $|p_1\rangle^{\otimes T(n, \sigma, \delta)}$ with failure probability δ by using any optimal algorithm that satisfies the error bound stated in the theorem. Since the KL-divergence from p_0 to p_1 is $D(p_0\|p_1) \leq p_0(b) \ln\left(\frac{p_0(b)}{p_1(b)}\right) = \frac{\alpha e^\alpha}{m^2} \leq \frac{6}{m}$, we must have $T(n, \sigma, \delta) \geq \Omega\left(\frac{\log(1/\delta)}{D(p_1\|p_0)}\right) = \Omega(n)$ by Lemma 6.5. \square

7 Discussion

One interesting open question is to find a quantum mean estimator that achieves the deviation bound $\Pr\left[|\tilde{\mu} - \mu| > \frac{\sigma \log(1/\delta)}{n}\right] \leq \delta$ by performing a number of experiments that is *linear* in n . The current best upper bound (Theorem 4.2) is $O(n \log^{3/2}(n) \log \log(n))$, and the lower bound is $\Omega(n)$ (Theorem 6.2). A first step toward this goal could be to obtain a better algorithm for the restricted case of Gaussian distributions. An equivalent goal is to find the smallest value L such that the deviation bound $\Pr\left[|\tilde{\mu} - \mu| > L \frac{\sigma \log(1/\delta)}{n}\right] \leq \delta$ can be achieved by a quantum mean estimator that performs *at most* n quantum experiments. Classically, for the sub-Gaussian deviation bound of Equation (1), the optimal value is $L = \sqrt{2}(1 + o(1))$ [Cat12; LV20].

There exist many variants of the quantum mean estimation problem that have not been completely explored in the quantum model yet. Let us mention for instance the multivariate setting [LM19], where the objective is to estimate the mean of a random variable taking values in \mathbb{R}^d . Heinrich [Hei04] proved that no quantum speed-up is achievable under some condition on the largeness of d . On the other hand, Cornelissen and Jerbi [CJ21] obtained partial quantum speed-ups for some parameter settings. We also note that the first polynomial-time classical algorithm with a sub-Gaussian error rate for this problem was only found recently by Hopkins [Hop20].

References

- [ABC+20] S. Arunachalam, A. Belovs, A. M. Childs, R. Kothari, A. Rosmanis, and R. de Wolf. “Quantum Coupon Collector”. In: *Proceedings of the 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)*. 2020, 10:1–10:17 (cit. on p. 6).
- [AGGW20] J. van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf. “Quantum SDP-Solvers: Better Upper and Lower Bounds”. In: *Quantum* 4 (2020), p. 230 (cit. on pp. 4, 6, 7).
- [AMS99] N. Alon, Y. Matias, and M. Szegedy. “The Space Complexity of Approximating the Frequency Moments”. In: *Journal of Computer and System Sciences* 58.1 (1999), pp. 137–147 (cit. on p. 1).
- [AW18] S. Arunachalam and R. de Wolf. “Optimal Quantum Sample Complexity of Learning Algorithms”. In: *Journal of Machine Learning Research* 19.1 (2018), pp. 2879–2878 (cit. on p. 6).
- [AW99] D. S. Abrams and C. P. Williams. *Fast Quantum Algorithms for Numerical Integrals and Stochastic Processes*. [arXiv:quant-ph/9908083](https://arxiv.org/abs/quant-ph/9908083). 1999 (cit. on p. 2).
- [BBHT98] M. Boyer, G. Brassard, P. Høyer, and A. Tapp. “Tight Bounds on Quantum Searching”. In: *Fortschritte der Physik* 46.4-5 (1998), pp. 493–505 (cit. on p. 19).
- [BCL13] S. Bubeck, N. Cesa-Bianchi, and G. Lugosi. “Bandits with Heavy Tail”. In: *IEEE Transactions on Information Theory* 59.11 (2013), pp. 7711–7717 (cit. on pp. 2, 3).
- [BCWZ99] H. Buhrman, R. Cleve, R. de Wolf, and C. Zalka. “Bounds for Small-Error and Zero-Error Quantum Algorithms”. In: *Proceedings of the 40th Symposium on Foundations of Computer Science (FOCS)*. 1999, pp. 358–368 (cit. on p. 13).
- [BDGT11] G. Brassard, F. Dupuis, S. Gambs, and A. Tapp. *An Optimal Quantum Algorithm to Approximate the Mean and its Application for Approximating the Median of a Set of Points over an Arbitrary Distance*. [arXiv:1106.4267](https://arxiv.org/abs/1106.4267) [quant-ph]. 2011 (cit. on pp. 1, 2, 5, 6).
- [BHH11] S. Bravyi, A. W. Harrow, and A. Hassidim. “Quantum Algorithms for Testing Properties of Distributions”. In: *IEEE Transactions on Information Theory* 57.6 (2011), pp. 3971–3981 (cit. on p. 6).
- [BHMT02] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. “Quantum Amplitude Amplification and Estimation”. In: *Contemporary Mathematics* 305 (2002), pp. 53–74 (cit. on pp. 1, 2, 8, 18–20).
- [Bic65] P. J. Bickel. “On Some Robust Estimates of Location”. In: *The Annals of Mathematical Statistics* 36.3 (1965), pp. 847–858 (cit. on p. 3).
- [BJ99] N. H. Bshouty and J. C. Jackson. “Learning DNF over the Uniform Distribution Using a Quantum Example Oracle”. In: *SIAM Journal on Computing* 28.3 (1999), pp. 1136–1153 (cit. on p. 6).

- [Cat12] O. Catoni. “Challenging the Empirical Mean and Empirical Variance: A Deviation Study”. In: *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* 48.4 (2012), pp. 1148–1185 (cit. on pp. 2, 15).
- [CEG95] R. Canetti, G. Even, and O. Goldreich. “Lower Bounds for Sampling Algorithms for Estimating the Average”. In: *Information Processing Letters* 53.1 (1995), pp. 17–25 (cit. on p. 3).
- [CFMW10] S. Chakraborty, E. Fischer, A. Matsliah, and R. de Wolf. “New Results on Quantum Property Testing”. In: *Proceedings of the 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. 2010, pp. 145–156 (cit. on p. 6).
- [CJ21] A. Cornelissen and S. Jerbi. *Quantum Algorithms for Multivariate Monte Carlo Estimation*. [arXiv:2107.03410](https://arxiv.org/abs/2107.03410) [quant-ph]. 2021 (cit. on p. 16).
- [DH96] C. Dürr and P. Høyer. *A Quantum Algorithm for Finding the Minimum*. [arXiv:quant-ph/9607014](https://arxiv.org/abs/quant-ph/9607014). 1996 (cit. on pp. 4, 6, 7).
- [DKLR00] P. Dagum, R. Karp, M. Luby, and S. Ross. “An Optimal Algorithm for Monte Carlo Estimation”. In: *SIAM Journal on Computing* 29.5 (2000), pp. 1484–1496 (cit. on pp. 2, 4, 11).
- [DLLO16] L. Devroye, M. Lerasle, G. Lugosi, and R. I. Oliveira. “Sub-Gaussian Mean Estimators”. In: *The Annals of Statistics* 44.6 (2016), pp. 2695–2725 (cit. on p. 2).
- [GNP13] L. Gajek, W. Niemirow, and P. Pokarowski. “Optimal Monte Carlo Integration with Fixed Relative Precision”. In: *Journal of Complexity* 29.1 (2013), pp. 4–26 (cit. on p. 2).
- [Gro98] L. K. Grover. “A Framework for Fast Quantum Mechanical Algorithms”. In: *Proceedings of the 30th Symposium on Theory of Computing (STOC)*. 1998, pp. 53–62 (cit. on pp. 2, 6).
- [Hei02] S. Heinrich. “Quantum Summation with an Application to Integration”. In: *Journal of Complexity* 18.1 (2002), pp. 1–50 (cit. on pp. 1, 2, 4, 6).
- [Hei03] S. Heinrich. “From Monte Carlo to Quantum Computation”. In: *Mathematics and Computers in Simulation* 62.3–6 (2003), pp. 219–230 (cit. on p. 2).
- [Hei04] S. Heinrich. “On the Power of Quantum Algorithms for Vector Valued Mean Computation”. In: *Monte Carlo Methods and Applications* 10.3–4 (2004), pp. 297–310 (cit. on p. 16).
- [Hel69] C. W. Helstrom. “Quantum Detection and Estimation Theory”. In: *Journal of Statistical Physics* 1.2 (1969), pp. 231–252 (cit. on p. 15).
- [HM18] Y. Hamoudi and F. Magniez. *Quantum Chebyshev’s Inequality and Applications*. [arXiv:1807.06456](https://arxiv.org/abs/1807.06456) [quant-ph]. 2018 (cit. on p. 4).
- [HM19] Y. Hamoudi and F. Magniez. “Quantum Chebyshev’s Inequality and Applications”. In: *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*. 2019, 69:1–69:16 (cit. on pp. 1–5, 11).
- [Hop20] S. B. Hopkins. “Mean Estimation with Sub-Gaussian Rates in Polynomial Time”. In: *The Annals of Statistics* 48.2 (2020), pp. 1193–1213 (cit. on p. 16).
- [Hub19] M. Huber. “An Optimal (ϵ, δ) -Randomized Approximation Scheme for the Mean of Random Variables with Bounded Relative Variance”. In: *Random Structures & Algorithms* 55.2 (2019), pp. 356–370 (cit. on p. 2).
- [JV86] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. “Random Generation of Combinatorial Structures from a Uniform Distribution”. In: *Theoretical Computer Science* 43 (1986), pp. 169–188 (cit. on p. 1).

- [LM19] G. Lugosi and S. Mendelson. “Mean Estimation and Regression Under Heavy-Tailed Distributions: A Survey”. In: *Foundations of Computational Mathematics* 19.5 (2019), pp. 1145–1190 (cit. on pp. 1–4, 16).
- [LV20] J. C. H. Lee and P. Valiant. *Optimal Sub-Gaussian Mean Estimation in \mathbb{R}* . [arXiv:2011.08384 \[math.ST\]](#). 2020 (cit. on pp. 2, 15).
- [LW19] T. Li and X. Wu. “Quantum Query Complexity of Entropy Estimation”. In: *IEEE Transactions on Information Theory* 65.5 (2019), pp. 2899–2921 (cit. on p. 6).
- [Mon15] A. Montanaro. “Quantum Speedup of Monte Carlo Methods”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 471.2181 (2015), p. 20150301 (cit. on pp. 1–5, 8).
- [MSA08] V. Mnih, C. Szepesvári, and J.-Y. Audibert. “Empirical Bernstein Stopping”. In: *Proceedings of the 25th International Conference on Machine Learning (ICML)*. 2008, pp. 672–679 (cit. on p. 2).
- [Nay99] A. Nayak. “Lower Bounds for Quantum Computation and Communication”. PhD thesis. University of California, Berkeley, 1999 (cit. on pp. 6, 14).
- [Nov01] E. Novak. “Quantum Complexity of Integration”. In: *Journal of Complexity* 17.1 (2001), pp. 2–16 (cit. on p. 2).
- [NW99] A. Nayak and F. Wu. “The Quantum Query Complexity of Approximating the Median and Related Statistics”. In: *Proceedings of the 31st Symposium on Theory of Computing (STOC)*. 1999, pp. 384–393 (cit. on pp. 3, 4, 6).
- [NY83] A. S. Nemirovsky and D. B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. John Wiley & Sons, 1983 (cit. on p. 1).
- [SV05] L. J. Schulman and V. V. Vazirani. “A Computationally Motivated Definition of Parametric Estimation and its Applications to the Gaussian Distribution”. In: *Combinatorica* 25.4 (2005), pp. 465–486 (cit. on p. 2).
- [Ter99] B. M. Terhal. “Quantum Algorithms and Quantum Entanglement”. PhD thesis. University of Amsterdam, 1999 (cit. on pp. 2, 5, 8).
- [TW02] J. F. Traub and H. Wozniakowski. “Path Integration on a Quantum Computer”. In: *Quantum Information Processing* 1.5 (2002), pp. 365–388 (cit. on p. 2).

A Auxiliary algorithms

A.1 Amplitude amplification

The amplitude amplification algorithm [BHMT02] is a generalization of Quantum Searching to the problem of boosting the success probability of a quantum algorithm that performs no intermediate measurement. The next result corresponds to Equation (8) in [BHMT02].

Theorem A.1 (AMPLITUDE AMPLIFICATION, [BHMT02]). *Let U be a unitary quantum algorithm and let Π be a projection operator. Consider the angle $\theta \in [0, \frac{\pi}{2}]$ and two unit states $|\psi_0\rangle, |\psi_1\rangle$ such that $\sin(\theta)|\psi_1\rangle = \Pi U|\mathbf{0}\rangle$ and $U|\mathbf{0}\rangle = \cos(\theta)|\psi_0\rangle + \sin(\theta)|\psi_1\rangle$. Then, for any integer $n \geq 0$, the amplitude amplification algorithm $\text{AAmp}(U, \Pi, n)$ satisfies*

$$\text{AAmp}(U, \Pi, n)|\mathbf{0}\rangle = \cos((2n + 1)\theta)|\psi_0\rangle + \sin((2n + 1)\theta)|\psi_1\rangle.$$

The algorithm uses $n + 1$ applications of U , n applications of U^\dagger , and n applications of the reflection operator $I - 2\Pi$.

Next, we present a variant of the amplitude amplification algorithm that does not use a pre-defined number of applications of U and U^\dagger . We call it the “sequential amplitude amplification” algorithm in reference to *sequential analysis*. The original version of this algorithm was analysed in Theorem 3 of [BBHT98; BHMT02], with a bound on the expected time complexity $\mathbb{E}[T]$. We propose a slightly different version that allows us to bound $\mathbb{E}[T^2]$ and $\mathbb{E}[1/T]$ (note that $\mathbb{E}[T] \leq \sqrt{\mathbb{E}[T^2]}$). These bounds will be used in Theorem A.4 and Theorem 5.3.

1. Set $\ell = 0$ and $\lambda = 1.1$.
2. Increase ℓ by 1 and pick an integer n between $\lambda^{\ell-1}$ and $\lambda^\ell - 1$ uniformly at random.
3. Apply the amplitude amplification algorithm $\text{AAmp}(U, \Pi, n)$ (Theorem A.1) on $|0\rangle$ and measure the state by using the projective measurement $\{I - \Pi, \Pi\}$. If the outcome is “ Π ” then stop and output the obtained state. Otherwise, go to step 2.

Algorithm 4: Sequential amplitude amplification.

Theorem A.2 (SEQUENTIAL AMPLITUDE AMPLIFICATION). *Let U be a unitary quantum algorithm and let Π be a projection operator. Define the number $p \in [0, 1]$ and the two unit states $|\psi_0\rangle, |\psi_1\rangle$ such that $U|0\rangle = \sqrt{1-p}|\psi_0\rangle + \sqrt{p}|\psi_1\rangle$ and $\Pi U|0\rangle = \sqrt{p}|\psi_1\rangle$. If $p > 0$ then the sequential amplitude amplification algorithm $\text{Seq-AAmp}(U, \Pi)$ outputs the state $|\psi_1\rangle$ with probability 1. Moreover, if we let T denote the number of applications of U , U^\dagger and $I - 2\Pi$ used by the algorithm, then $\mathbb{E}[T^2] \leq O(1/p)$ and $\mathbb{E}[1/T] \leq O(\sqrt{p})$.*

Proof. Let $0 < \theta \leq \pi/2$ be the angle such that $\sqrt{p} = \sin \theta$. We show the theorem in the case where $\theta < \pi/4$ (the case $\theta \geq \pi/4$ is easy to handle separately). We first prove the inequality $\mathbb{E}[T^2] \leq O(1/p)$. Let P_ℓ denote the probability of obtaining $|\psi_1\rangle$ when n is picked uniformly at random between $\lambda^{\ell-1}$ and $\lambda^\ell - 1$ and the state $\mathcal{A}(U, \Pi, n)|0\rangle$ is measured with respect to $\{I - \Pi, \Pi\}$. Let $\ell_+ = \left\lceil \log_\lambda \left(\frac{\lambda}{(\lambda-1)\sin(2\theta)} \right) \right\rceil$. If $\ell \geq \ell_+$ then,

$$\begin{aligned}
P_\ell &= \frac{1}{(\lambda-1)\lambda^{\ell-1}} \sum_{n=\lambda^{\ell-1}}^{\lambda^\ell-1} \sin^2((2n+1)\theta) && \text{by Theorem A.1} \\
&= \frac{1}{(\lambda-1)\lambda^{\ell-1}} \sum_{n=\lambda^{\ell-1}}^{\lambda^\ell-1} \frac{1 - \cos^2((2n+1)2\theta)}{2} && \text{by a trigonometric identity} \\
&\geq \frac{1}{2} - \frac{1}{2(\lambda-1)\lambda^{\ell-1}} \sum_{n=0}^{\lambda^\ell-1} \cos^2((2n+1)2\theta) \\
&= \frac{1}{2} - \frac{\sin(4\lambda^\ell\theta)}{4(\lambda-1)\lambda^{\ell-1}\sin(2\theta)} && \text{by a trigonometric identity} \\
&\geq \frac{1}{4} && \text{by } \ell \geq \left\lceil \log_\lambda \left(\frac{\lambda}{(\lambda-1)\sin(2\theta)} \right) \right\rceil
\end{aligned}$$

Moreover, the algorithm has used at most $\sum_{n=1}^{\ell} \lambda^n \leq 10\lambda^{\ell+1}$ applications of U , U^\dagger and $I - 2\Pi$ after ℓ iterations of step 3. Consequently, $\mathbb{E}[T^2] \leq \sum_{\ell \geq \ell_+} (10\lambda^{\ell+1})^2 (3/4)^{\ell-\ell_+} \leq O(\lambda^{2\ell_+}) \leq O(1/p)$.

Next, we prove the inequality $\mathbb{E}[1/T] \leq O(\sqrt{p})$. Let $\ell_- = \left\lceil \log_\lambda \left(\frac{1}{5\theta} \right) \right\rceil$. We will use that $(2/\pi)x \leq \sin(x) \leq x$ for all $x \in [0, \pi/2]$. According to Theorem A.1, if n is an integer between $\lambda^{\ell-1}$ and $\lambda^\ell - 1$ for some $1 \leq \ell \leq \ell_-$, then the probability of obtaining $|\psi_1\rangle$ when $\mathcal{A}(U, \Pi, n)|0\rangle$ is measured is $\sin^2((2n+1)\theta) \leq (2\lambda^\ell - 1)^2 \theta^2$. Moreover, after ℓ iterations of

step 3, the algorithm used at least $\lambda^{\ell-1}$ applications of $U, U^\dagger, I - 2\Pi$ in total. Consequently, $\mathbb{E}[1/T] \leq \sum_{\ell=1}^{\ell_-} \frac{(2\lambda^{\ell-1})^2 \theta^2}{\lambda^{\ell-1}} + \frac{1}{\lambda^{\ell_-}} \leq O(\lambda^{\ell_-} \theta^2 + \theta) \leq O(\sqrt{p})$. \square

A.2 Amplitude estimation

The amplitude estimation algorithm [BHMT02] is a generalization of Quantum Counting to the problem of estimating the success probability of an algorithm. The next result corresponds to Theorems 11 and 12 in [BHMT02].

Theorem A.3 (AMPLITUDE ESTIMATION, [BHMT02]). *Let U be a unitary quantum algorithm and let Π be a projection operator. Define the number $p \in [0, 1]$ such that $p = \|\Pi U |\mathbf{0}\rangle\|^2$. Then, for any integer $n \geq 0$, the amplitude estimation algorithm $\text{AEst}(U, \Pi, n)$ outputs an amplitude estimate \tilde{p} such that,*

$$\Pr \left[|\tilde{p} - p| \leq \frac{2\pi\sqrt{p(1-p)}}{n} + \frac{\pi^2}{n^2} \right] \geq 8/\pi^2.$$

The algorithm uses n applications of $U, U^\dagger, I - 2\Pi$ and $O(\log^2(n))$ 2-qubit quantum gates.

We present a sequential version of the amplitude estimation algorithm that does not need a time parameter n as input. This result was first obtained by [BHMT02, Theorem 15]. We describe a variant with additional properties that is based on the sequential amplitude amplification algorithm. It is used in Proposition 5.2 and Theorem 5.3.

Theorem A.4 (SEQUENTIAL AMPLITUDE ESTIMATION). *There exists an algorithm, called the sequential amplitude estimation algorithm Seq-AEst , with the following properties. Let U be a unitary quantum algorithm and let Π be a projection operator. Define the number $p \in [0, 1]$ such that $p = \|\Pi U |\mathbf{0}\rangle\|^2$. Then, the algorithm $\text{Seq-AEst}(U, \Pi)$ outputs an amplitude estimate \tilde{p} and uses a number T of applications of $U, U^\dagger, I - 2\Pi$ such that,*

1. *There is a universal constant $c \in (0, 1)$ such that $\Pr[|\tilde{p} - p| \leq cp] \geq 7/8$.*
2. *There is a universal constant c' such that $\mathbb{E}[T^2] = \mathbb{E}[1/\tilde{p}] \leq c'/p$.*
3. *There is a universal constant c'' such that $\mathbb{E}[1/T] = \mathbb{E}[\sqrt{\tilde{p}}] \leq c''\sqrt{p}$.*

Proof. The algorithm $\text{Seq-AEst}(U, \Pi)$ consists of recording the number T of applications of $U, U^\dagger, I - 2\Pi$ used by the sequential amplitude amplification algorithm $\text{Seq-AAmp}(U, \Pi)$ (Theorem A.2), and choosing the estimate $\tilde{p} = 1/T^2$. The results follow immediately from Theorem A.2 and Markov's inequality. \square