Quantum Algorithms for the Mean Estimation Problem

Yassine Hamoudi

UC Berkeley

Mean Estimation problem (over $\mathbb{R}$)

Experiment with unknown outcome distribution $D = (p_x)_x$

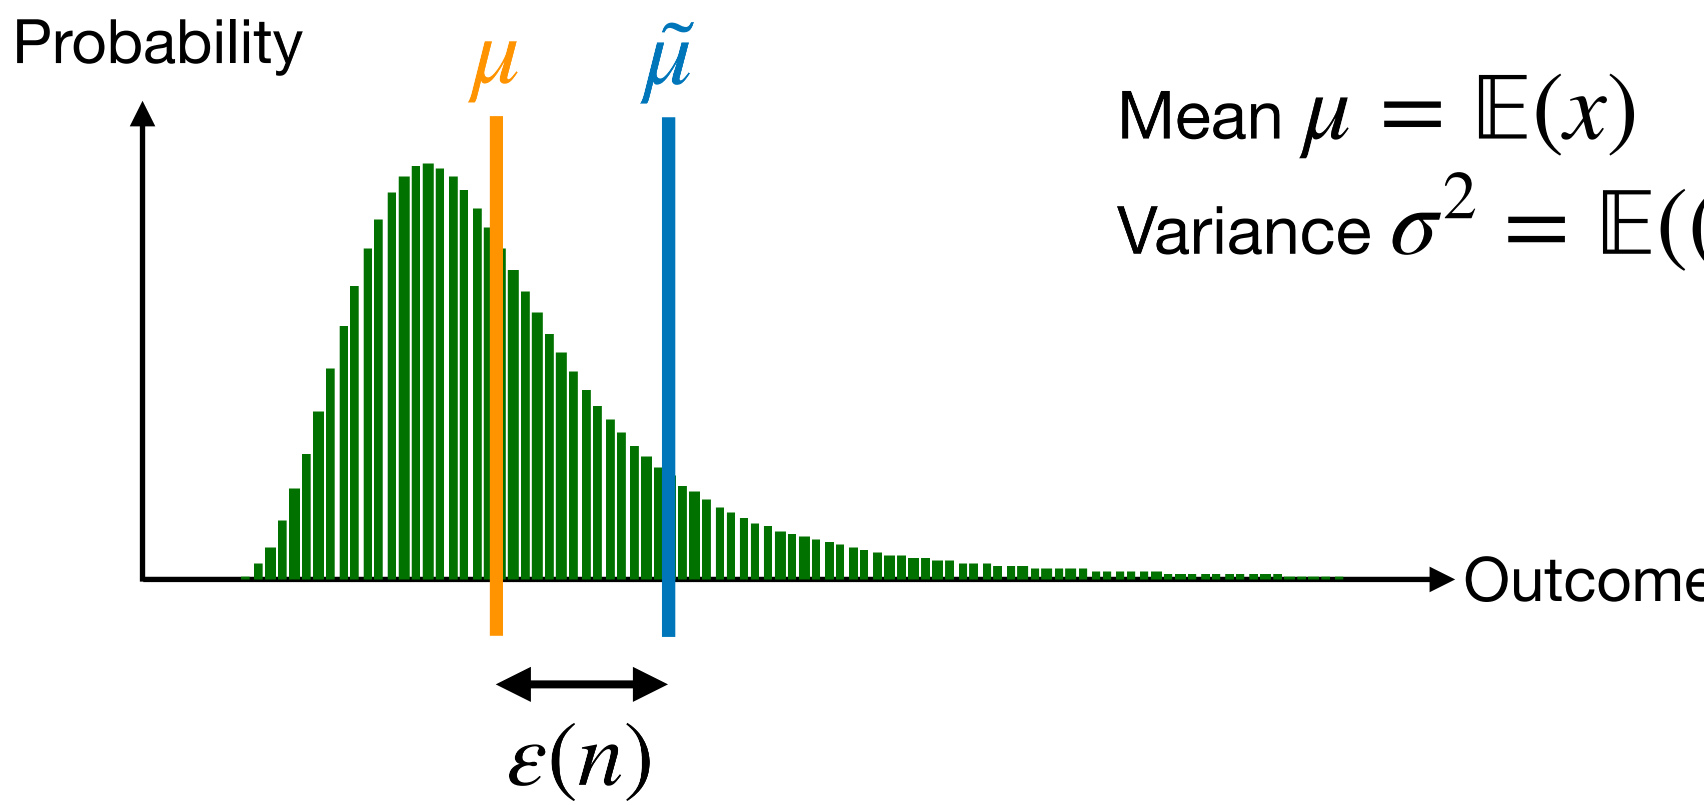

Mean $\mu = \mathbb{E}(x)$

Variance $\sigma^2 = \mathbb{E}((x - \mu)^2)$

Complexity parameter: number of times n the experiment is run

Goal: compute $\tilde{\mu}$ that minimizes the error $\varepsilon(n)$ such that

$$\Pr\left[\,|\mu - \tilde{\mu}| > \varepsilon(n)\right] < \delta \qquad \text{given } \delta \in (0,1)$$

1. The classics
2. Quantum input model(s)
3. Quantum “Bernoulli” estimator
4. Quantum “truncated” estimator
5. Quantum “sub-Gaussian” estimator
6. Future work

1

The *classics*

Sample-Mean estimator

3 7 4 8 4 5 4 1 0 2 2 6

$$\tilde{\mu} = \frac{3+7+4+8+4+5+4+1+0+2+2+6}{12} = 3.83$$

Optimal for **Gaussian** distributions: $\varepsilon_G(n) = \Theta\left(\sqrt{\frac{\sigma^2 \log(1/\delta)}{n}}\right)$

Central Limit Theorem: $\lim_{n\to\infty} \Pr\left[\,|\,\tilde{\mu} - \mu\,| > \varepsilon_G(n)\right] = \delta$ for any distribution

→ no guarantee for fixed n

→ non-asympotic error captured by Chebyshev inequality: $\varepsilon(n) = O\left(\sqrt{\frac{\sigma^2}{\delta n}}\right)$

Is there a **better** estimator?

Median-of-Means estimator

Partition the samples in $\sim \log(1/\delta)$ blocks:

3 7 4 8 4 5 4 1 0 2 2 6

$$\tilde{\mu}_1 = \frac{3+7+4+8}{4} \qquad \tilde{\mu}_2 = \frac{4+5+4+1}{4} \qquad \tilde{\mu}_3 = \frac{0+2+2+6}{4}$$

$$\tilde{\mu} = \text{median}(\tilde{\mu}_1, \tilde{\mu}_2, \tilde{\mu}_3) = 3.5$$

Error for **any** distribution: $\varepsilon(n) = O\left(\sqrt{\frac{\sigma^2 \log(1/\delta)}{n}}\right)$

Same as for Gaussian distribution!

An estimator is called **sub-Gaussian** if it satisfies

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

for any distribution with finite variance.

Examples: Median-of-Means, [Catoni'12], [Lee,Valiant'21], …

optimal $\sqrt{2} + o(1)$ prefactor

→ We need δ to be part of the input.

→ Fixed guarantees (ex: $|\mu - \tilde{\mu}| \leq \varepsilon\mu$ given ε) are typically achieved by finding upper-bounds on relevant quantities (ex: $N \geq (\sigma/(\varepsilon\mu))^2 \log(1/\delta)$).

2

Quantum input model(s)

Qsamples model

Unknown distribution $D = (p_x)_x$

n experiments = receive n copies of the **qsample** $\sum_x \sqrt{p_x}\,|x\rangle$

No advantage over the classical setting!

→ Measure the qsamples and run any **classical** sub-Gaussian estimators on the results

→ This is called the **Standard Quantum Limit** in quantum metrology

Oracle model

More powerful model: Black-box access to a quantum process generating qsamples

Formally: fix any **unitary** U_D such that $U_D|0\rangle = \sum_x \sqrt{p_x}\,|x\rangle\,|\text{garbage}_x\rangle$

weaker assumption

n experiments = n **applications** of U_D or U_D^{-1}

"reverse" the circuit computing U_D

The **Heisenberg Limit** predicts a $1/n$ error rate in this model (vs $1/\sqrt{n}$ before)

→ **Goal:** understand the dependence on other parameters (σ, δ, …)

Optimal error rates:

$$\Theta\left(\sqrt{\frac{\sigma^2 \log(1/\delta)}{n}}\right) \quad \text{vs} \quad \tilde{\Theta}\left(\frac{\sigma \log(1/\delta)}{n}\right)$$

Classical | Quantum

For most of this talk: $\tilde{O}\left(\frac{\sqrt{\mathbb{E}_D(x^2)}}{n}\right)$ with $\delta = 1/3$

+ distribution supported on non-negative values

3

Quantum “Bernoulli” estimator

Bernoulli distribution

Bernoulli(p): 1 with probability p 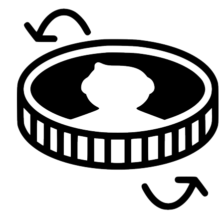

Given a distribution D over $[0,B]$ with mean μ, simulate Bernoulli(μ/B):

→ Sample $x \sim D$, sample $y \sim [0,B]$, return 1 if $y < x$ and 0 otherwise.

$$p = \sum_x p_x \cdot \frac{x}{B} = \frac{\mu}{B}$$

→ If we can estimate p with error $\varepsilon(n,p)$ then we can estimate the mean μ of D with error $B \cdot \varepsilon(n, \mu/B)$

Bernoulli distribution

Bernoulli(p): 1 with probability p 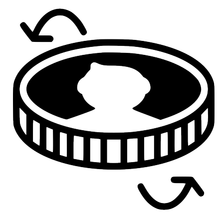

Given a distribution D over $[0,B]$ with mean μ, simulate Bernoulli(μ/B):

→ A similar reduction holds in the quantum model:

$$|0,0\rangle \xrightarrow{U_D} \sum_x \sqrt{p_x}\,|x\rangle\,|0\rangle$$

$$\xrightarrow[\text{rotation}]{\text{Controlled}} \sum_x \sqrt{p_x}\,|x\rangle\left(\sqrt{1-\frac{x}{B}}\,|0\rangle + \sqrt{\frac{x}{B}}\,|1\rangle\right)$$

$$= \sqrt{1-\frac{\mu}{B}}(\ldots)\,|0\rangle + \sqrt{\frac{\mu}{B}}(\ldots)\,|1\rangle$$

Amplitude estimation

Goal: estimate p given access to $U|0\rangle = \sqrt{1-p}\,|0\rangle + \sqrt{p}\,|1\rangle$

Grover operator: $G = U(2|0\rangle\langle 0| - I)U^{-1}(2|0\rangle\langle 0| - I)$

Two eigenvalues: $e^{-2i\theta}$ and $e^{+2i\theta}$ where $\sin^2(\theta) = p$

Apply n steps of
Phase Estimation
on G and $U|0\rangle$

↓

Take $\tilde{p} = \sin^2(\tilde{\theta})$

$$|\tilde{\theta} - \theta| \leq \frac{1}{n} \quad \text{or} \quad |\tilde{\theta} - (\pi - \theta)| \leq \frac{1}{n}$$

↓ *(trigonometric identities)*

$$|\tilde{p} - p| \lesssim \frac{\sqrt{p}}{n} + \frac{1}{n^2}$$

[Brassard,Höyer,Mosca,Tapp'02]

Quantum Bernoulli estimator

How good is the Bernoulli estimator for bounded distributions?

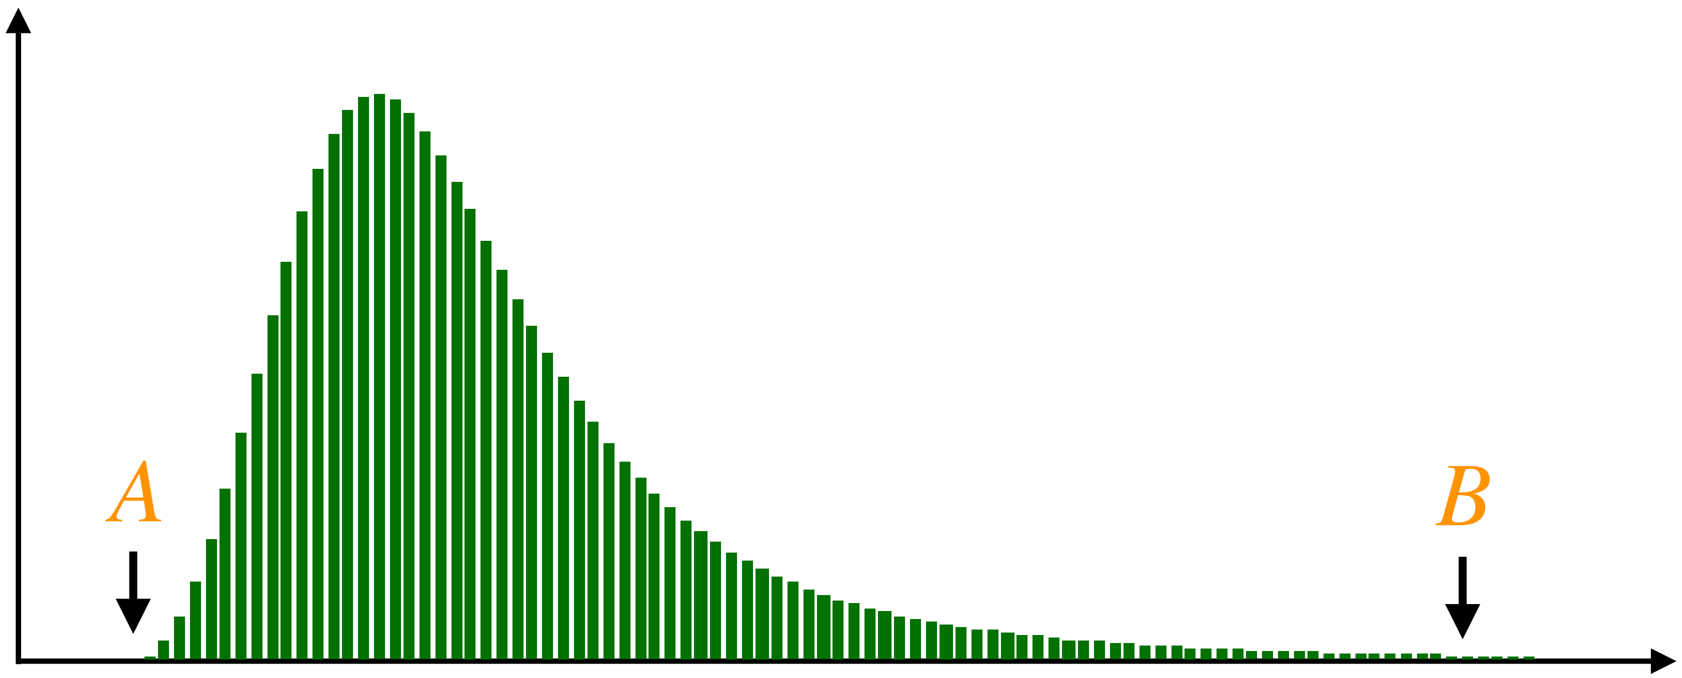

Error: $|\tilde{\mu} - \mu| \leq \frac{\sqrt{B\mu}}{n} + \frac{B}{n^2}$

$$\leq \frac{1}{n}\sqrt{\frac{B}{A}\mathbb{E}(x^2)} + \frac{B}{n^2}$$

Very sensitive to **outliers**!

Quantum Bernoulli estimator

How good is the Bernoulli estimator for bounded distributions?

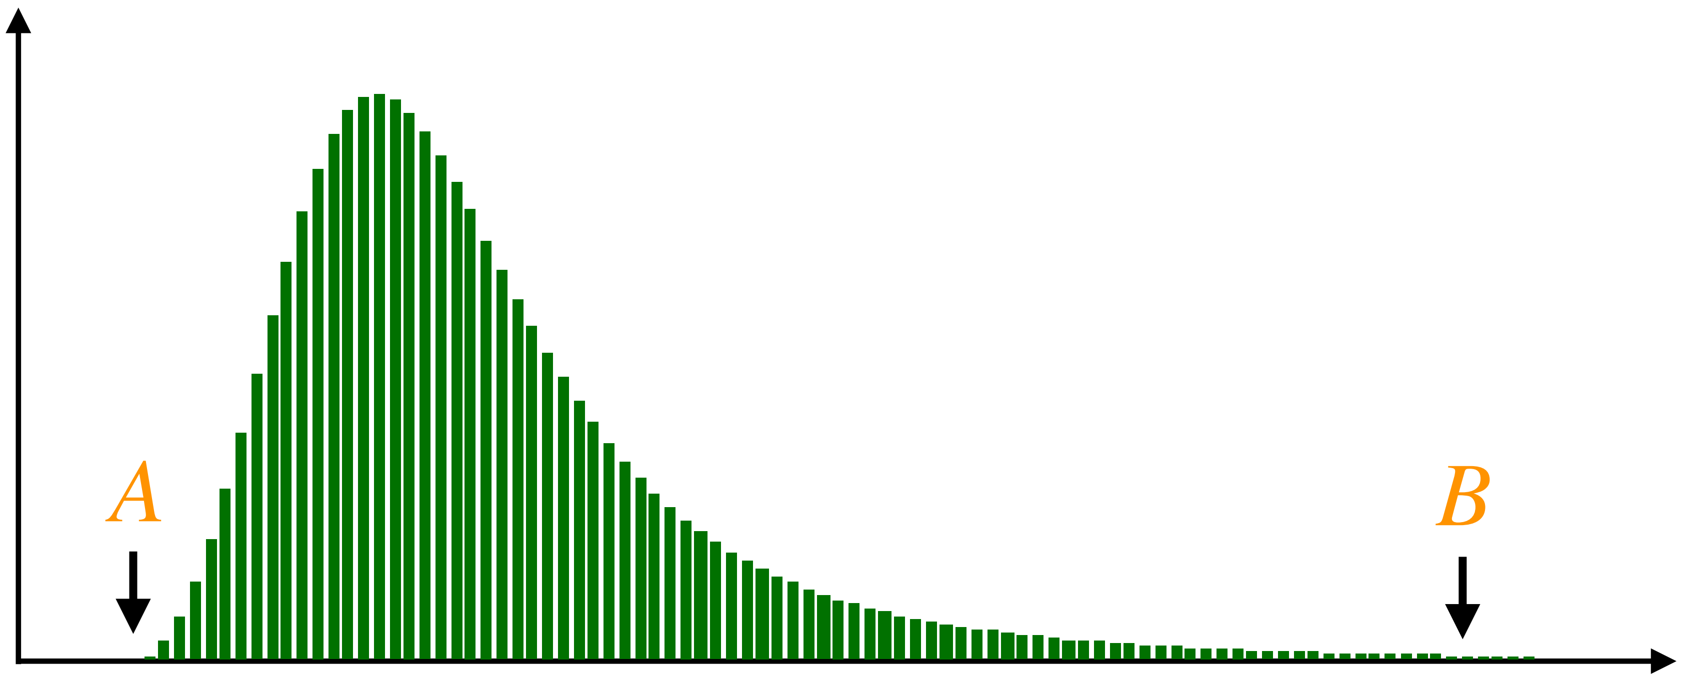

Error: $$|\tilde{\mu} - \mu| \leq \frac{\sqrt{B\mu}}{n} + \frac{B}{n^2}$$

$$\leq \frac{1}{n}\sqrt{\frac{B}{A}\mathbb{E}(x^2)} + \frac{B}{n^2}$$

Very sensitive to **outliers**!

④

Quantum “truncated” estimator

Truncated mean

Truncated distribution: $x \longmapsto x \cdot 1_{a<x\leq b}$

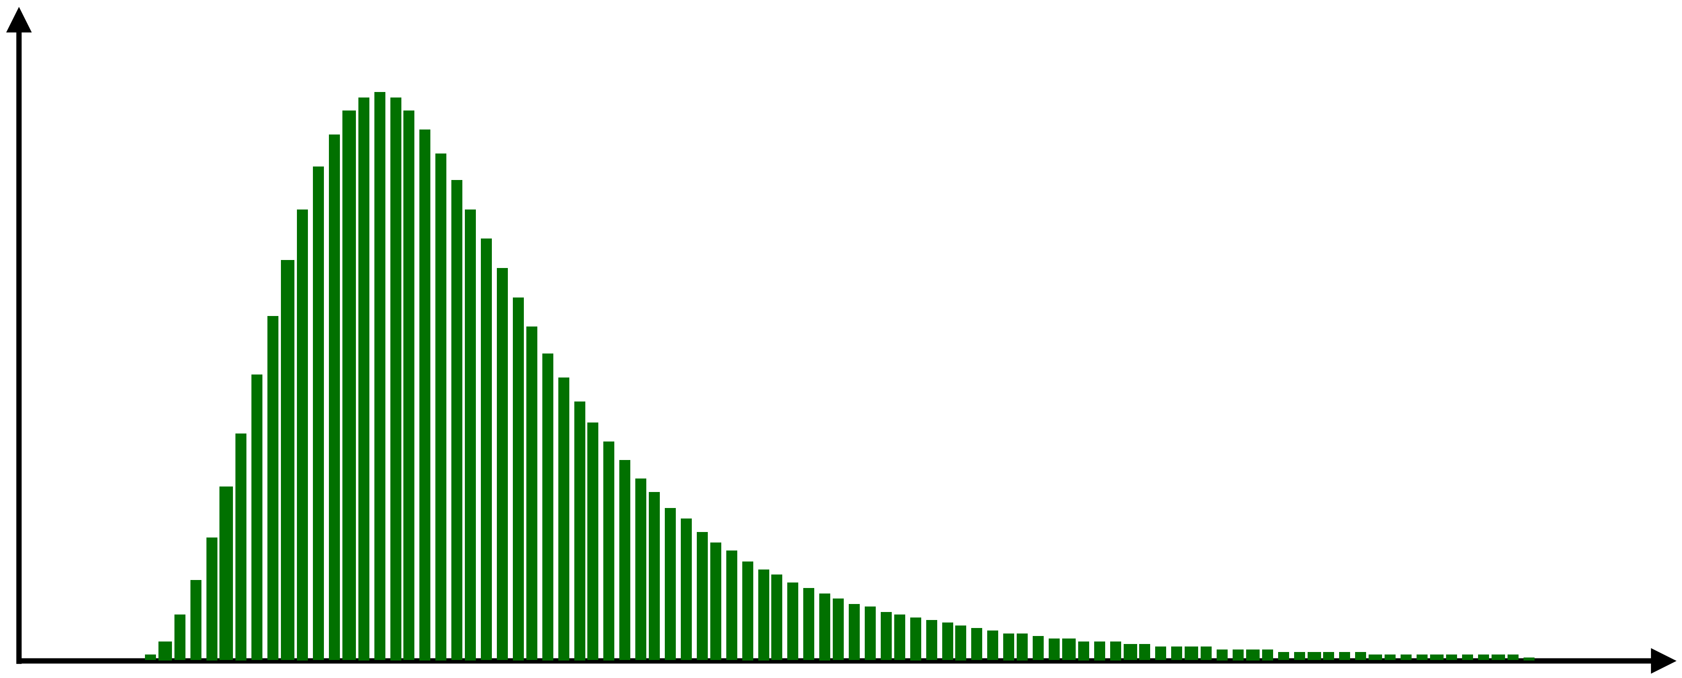

Truncated mean

Truncated distribution: $x \longmapsto x \cdot 1_{a<x\leq b}$

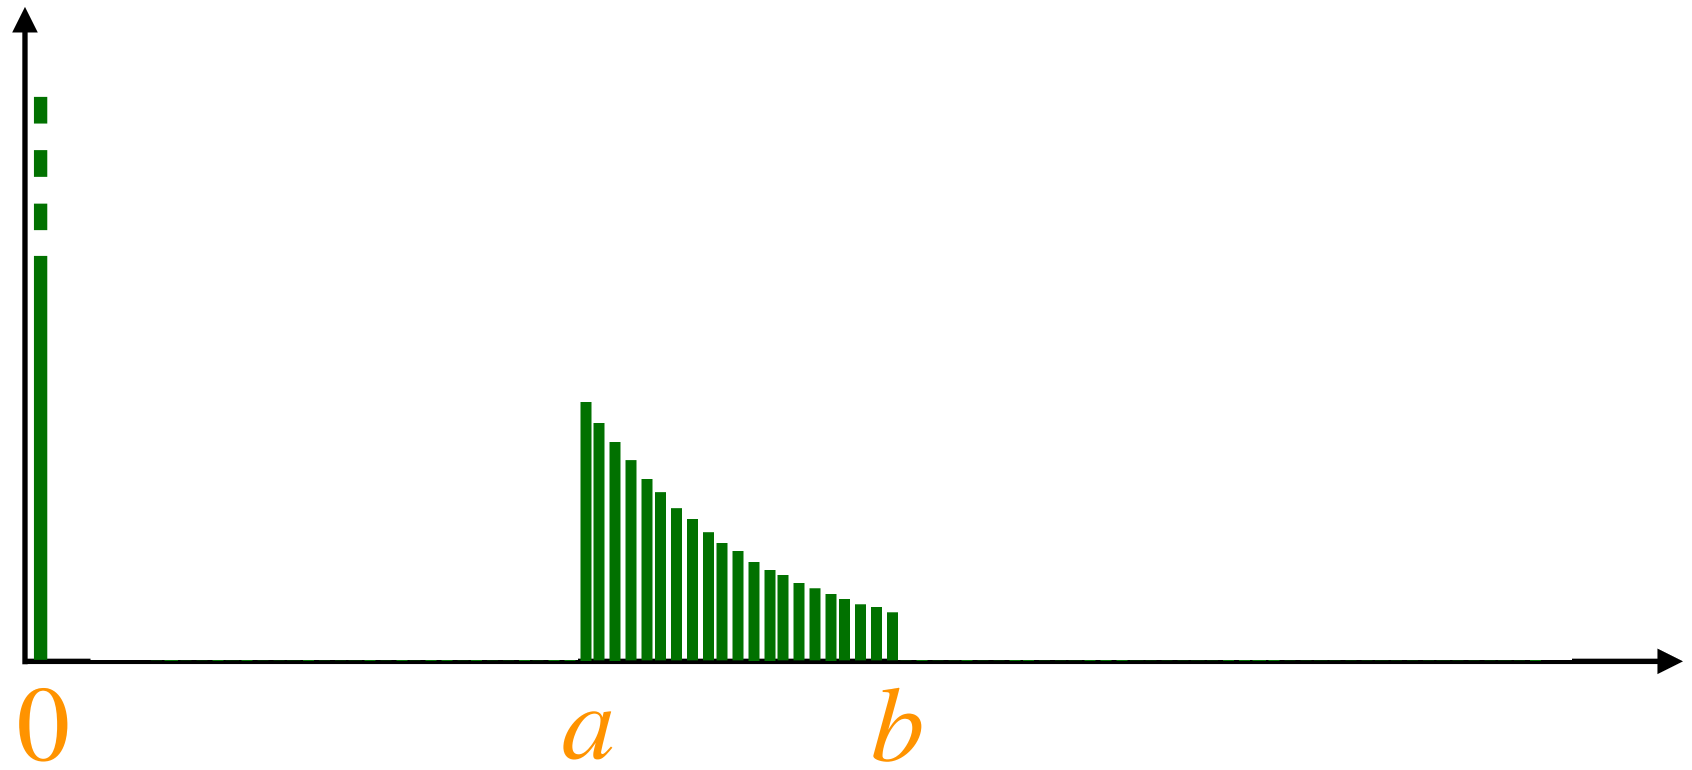

Truncated mean

Truncated distribution: $x \longmapsto x \cdot 1_{a<x\leq b}$

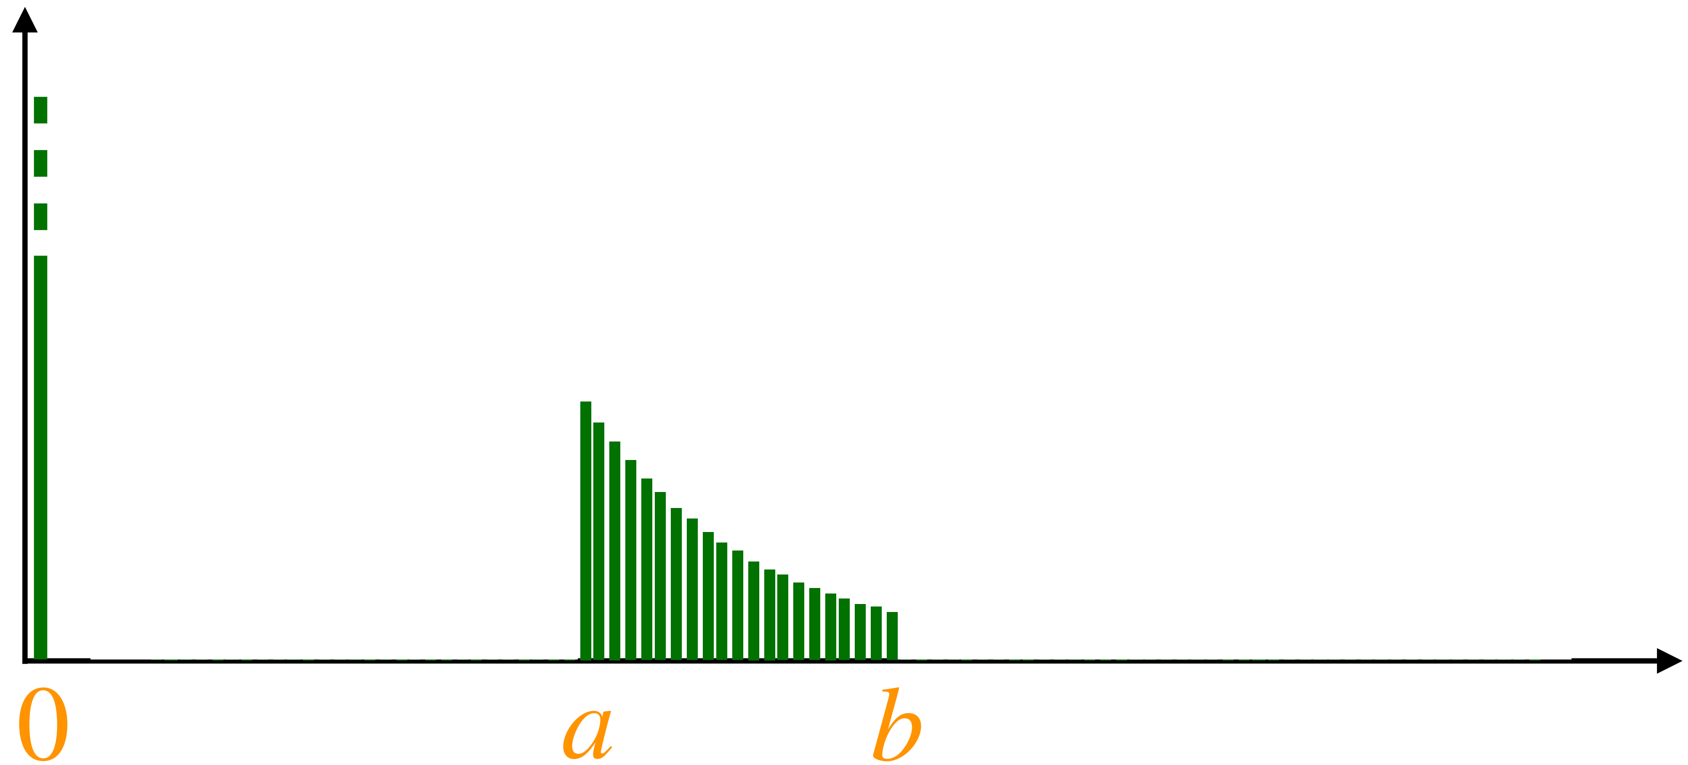

For any sequence $0 = a_0 < a_1 < a_2 < \ldots < a_k$:

$$\mathbb{E}(x) = \mathbb{E}\left(x \cdot 1_{a_0<x\leq a_1}\right) + \ldots + \mathbb{E}\left(x \cdot 1_{a_{k-1}<x\leq a_k}\right) + \mathbb{E}\left(x \cdot 1_{x>a_k}\right)$$

Quantum truncated estimator

Fix a threshold T :

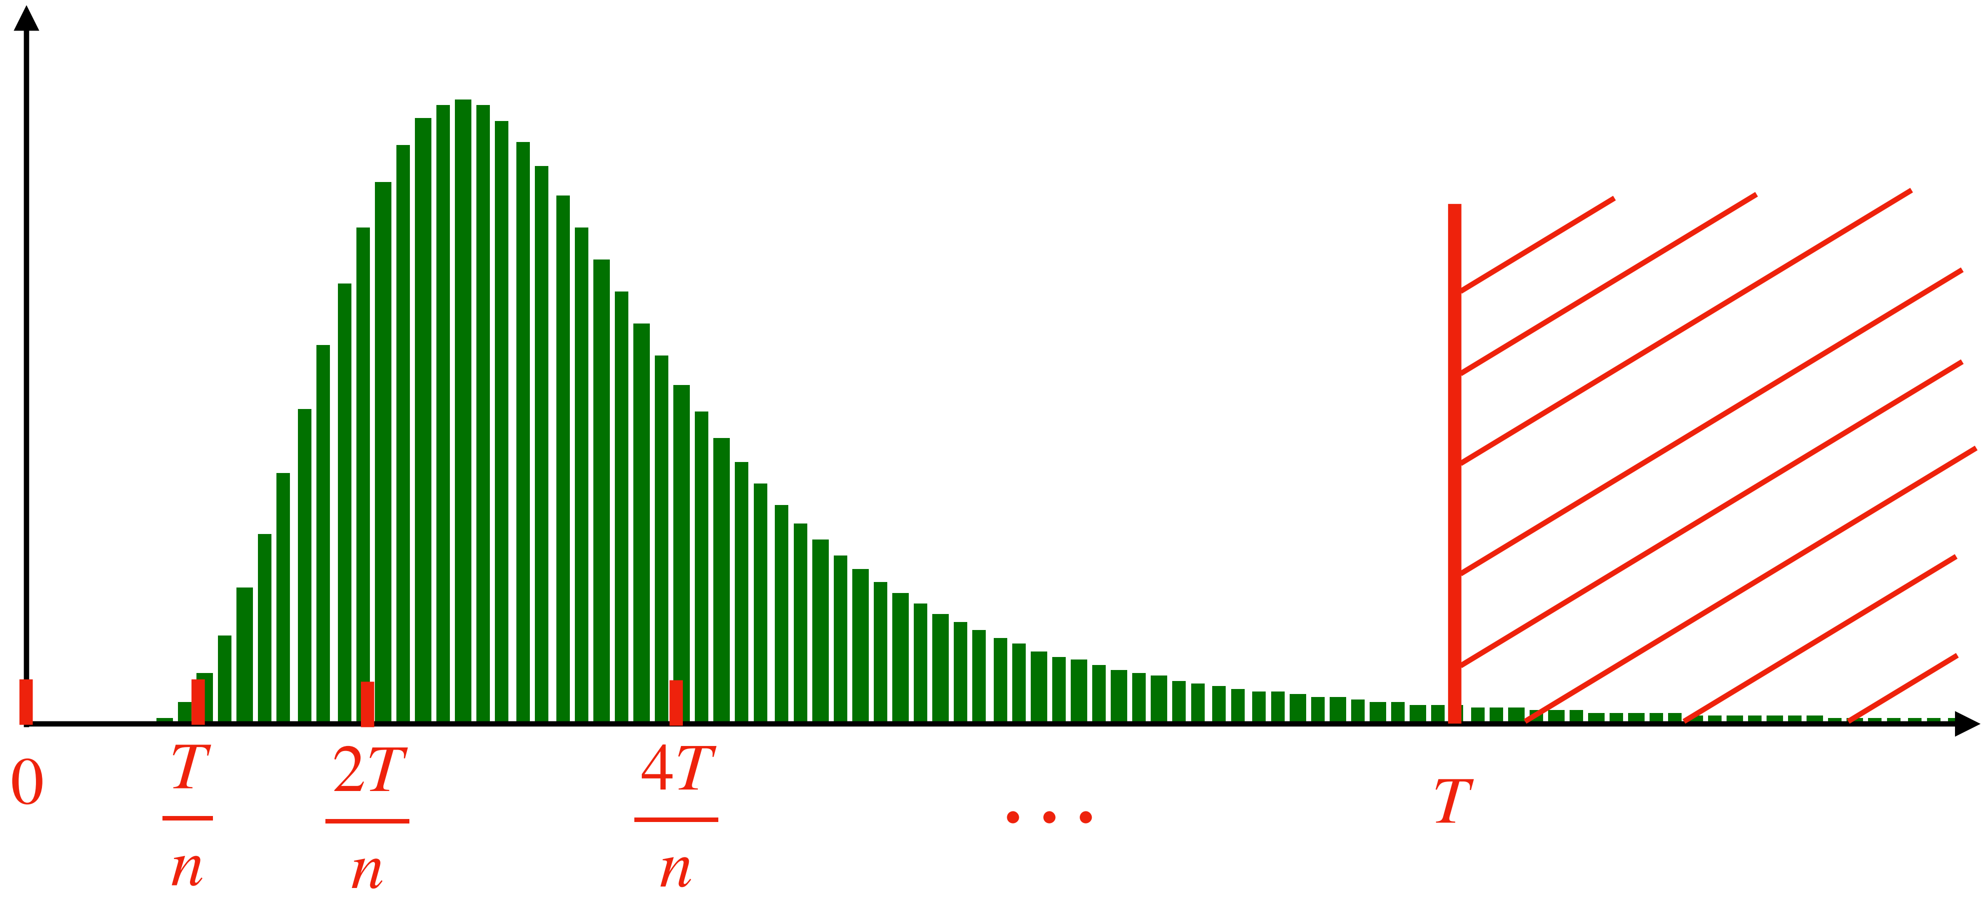

[Heinrich'02]
[Montanaro'15]
[H.,Magniez'19]

Quantum truncated estimator

Fix a threshold T :

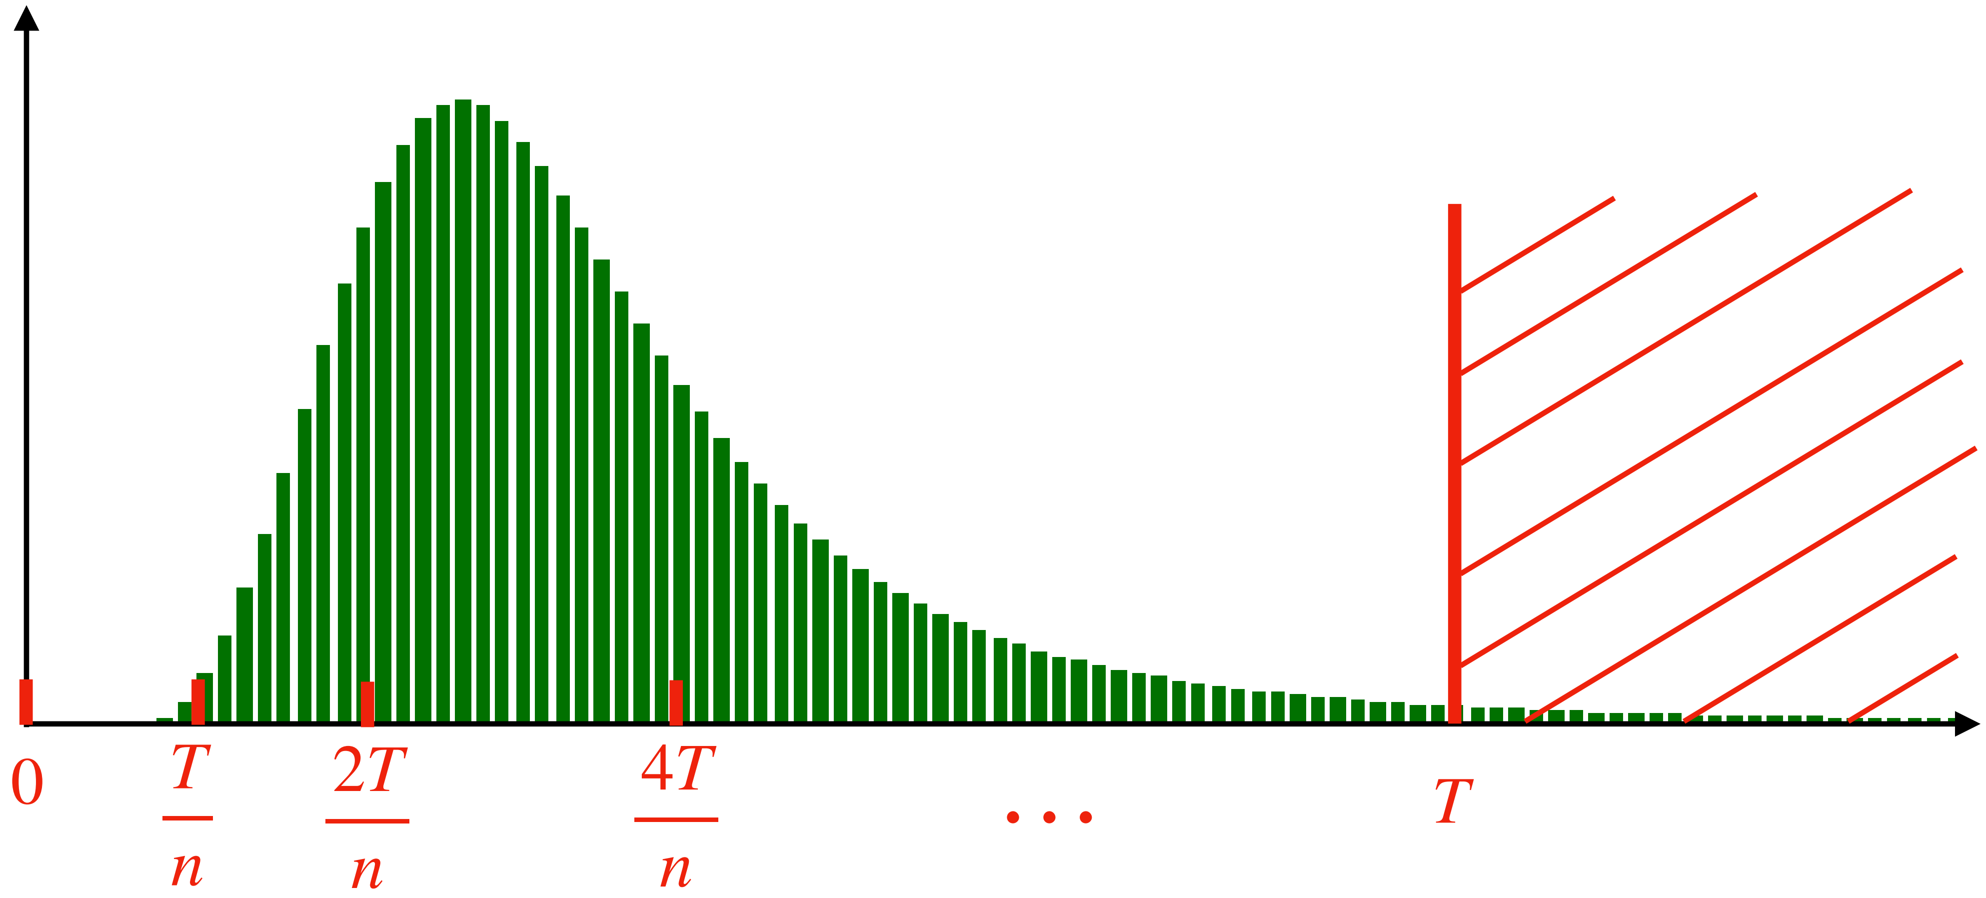

$\tilde{\mu} \;=$

[Heinrich'02]
[Montanaro'15]
[H.,Magniez'19]

Quantum truncated estimator

Fix a threshold T :

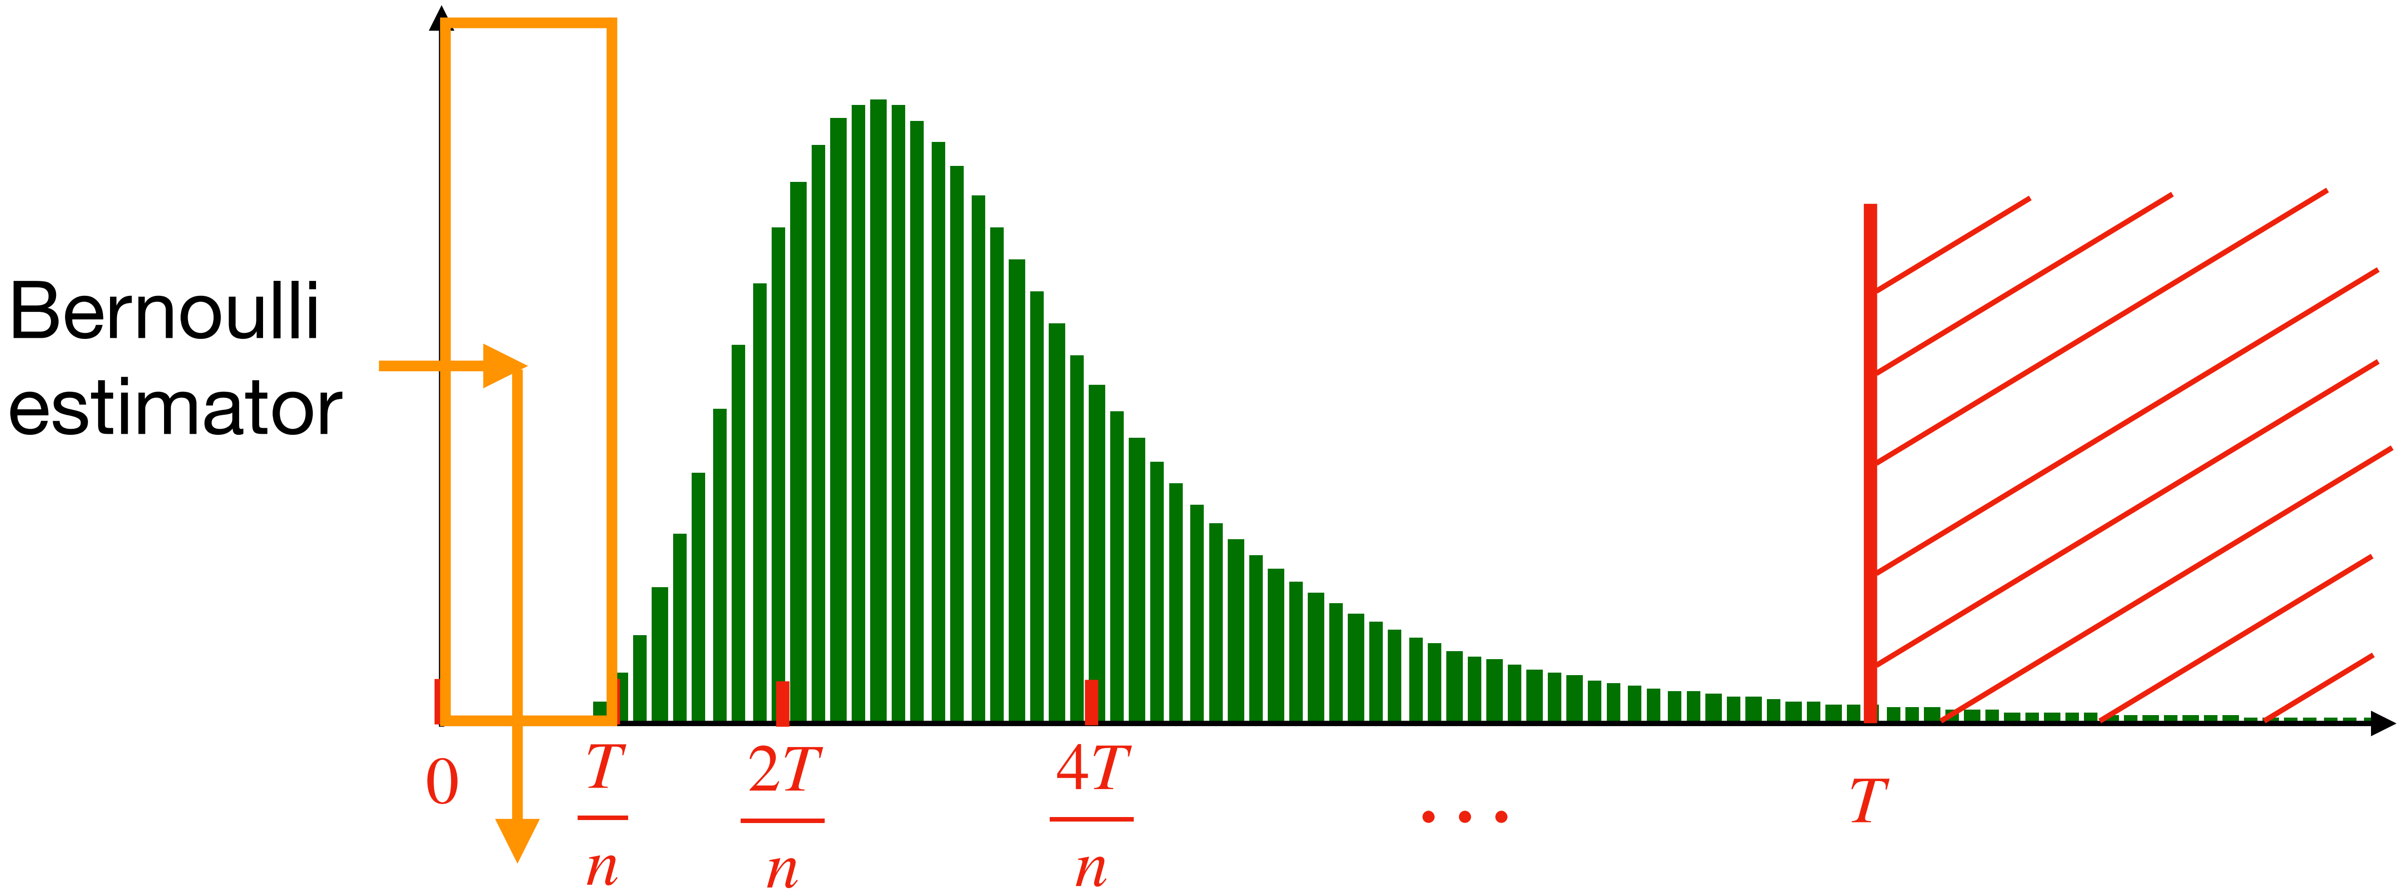

$$\tilde{\mu} \; = \; \tilde{\mu}_1 +$$

[Heinrich'02]
[Montanaro'15]
[H.,Magniez'19]

Quantum truncated estimator

Fix a threshold T :

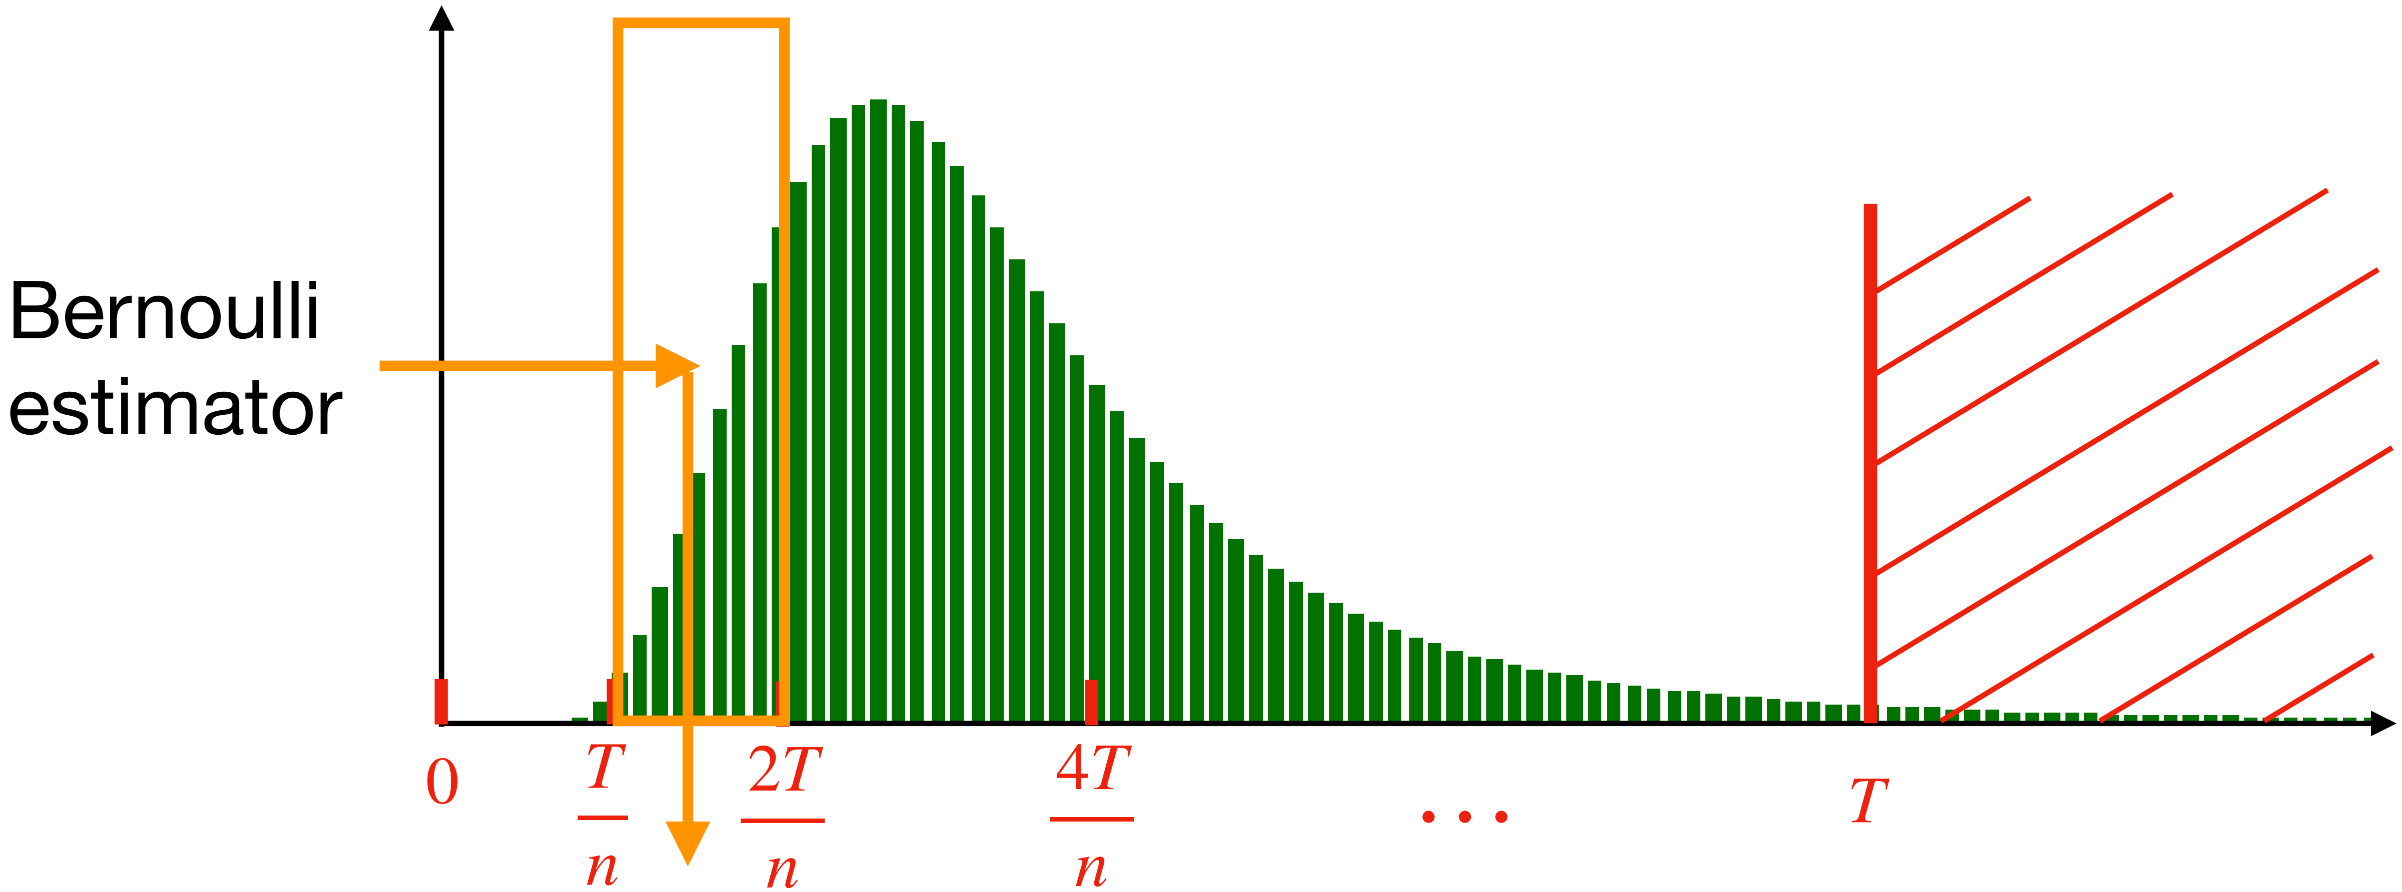

$$\tilde{\mu} = \tilde{\mu}_1 + \tilde{\mu}_2 +$$

[Heinrich'02]
[Montanaro'15]
[H.,Magniez'19]

Quantum truncated estimator

Fix a threshold T :

Bernoulli estimator

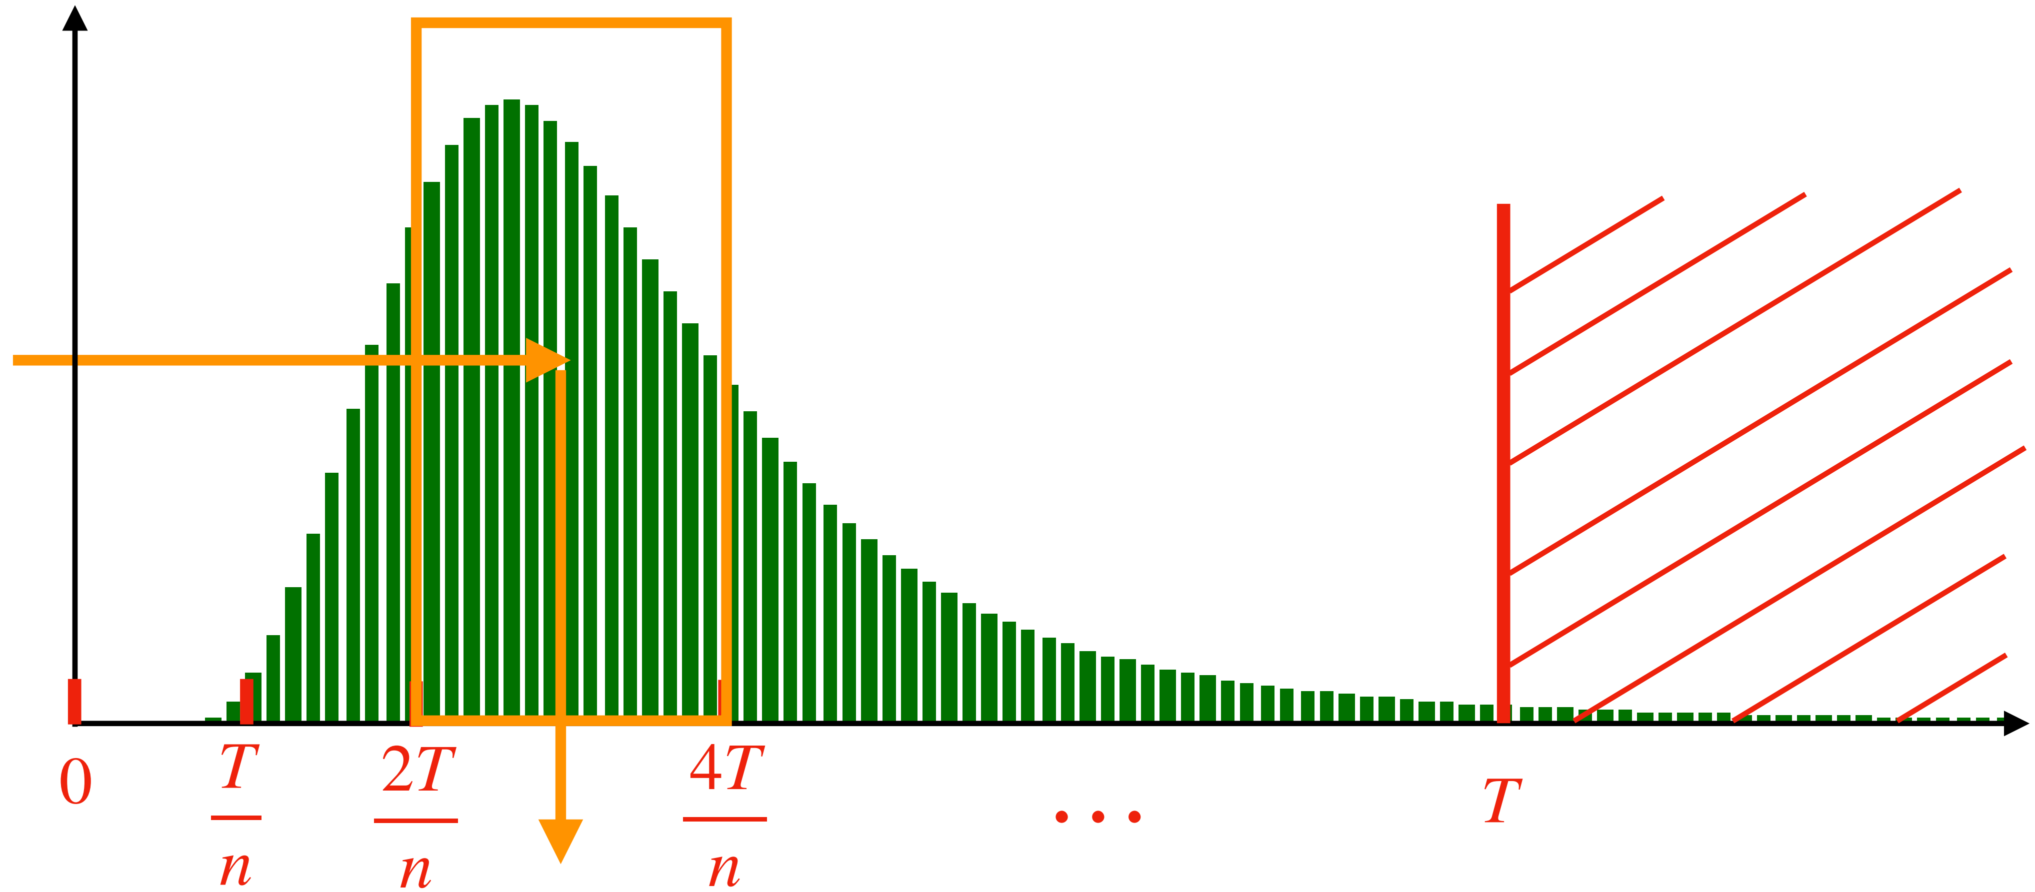

$$\tilde{\mu} = \tilde{\mu}_1 + \tilde{\mu}_2 + \tilde{\mu}_3 +$$

[Heinrich'02]
[Montanaro'15]
[H.,Magniez'19]

Quantum truncated estimator

Fix a threshold T :

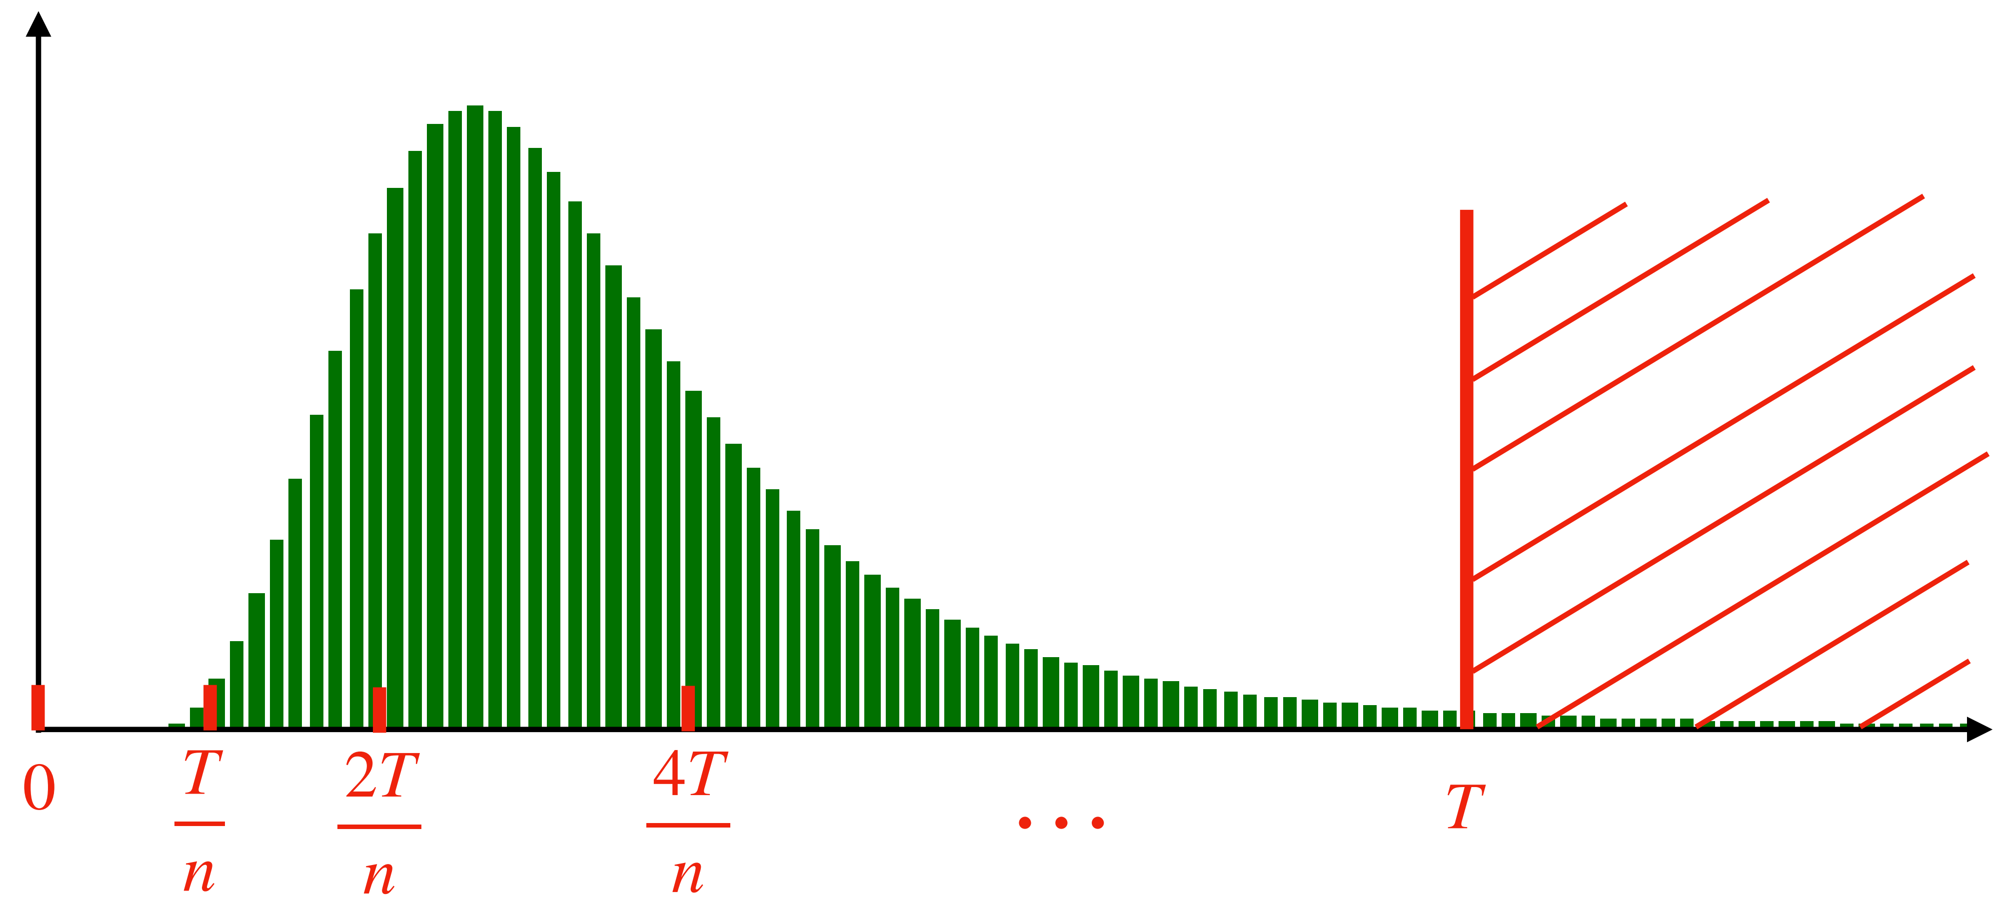

$$\tilde{\mu} = \tilde{\mu}_1 + \tilde{\mu}_2 + \tilde{\mu}_3 + \cdots + 0$$

[Heinrich'02]
[Montanaro'15]
[H.,Magniez'19]

Quantum truncated estimator

Fix a threshold T :

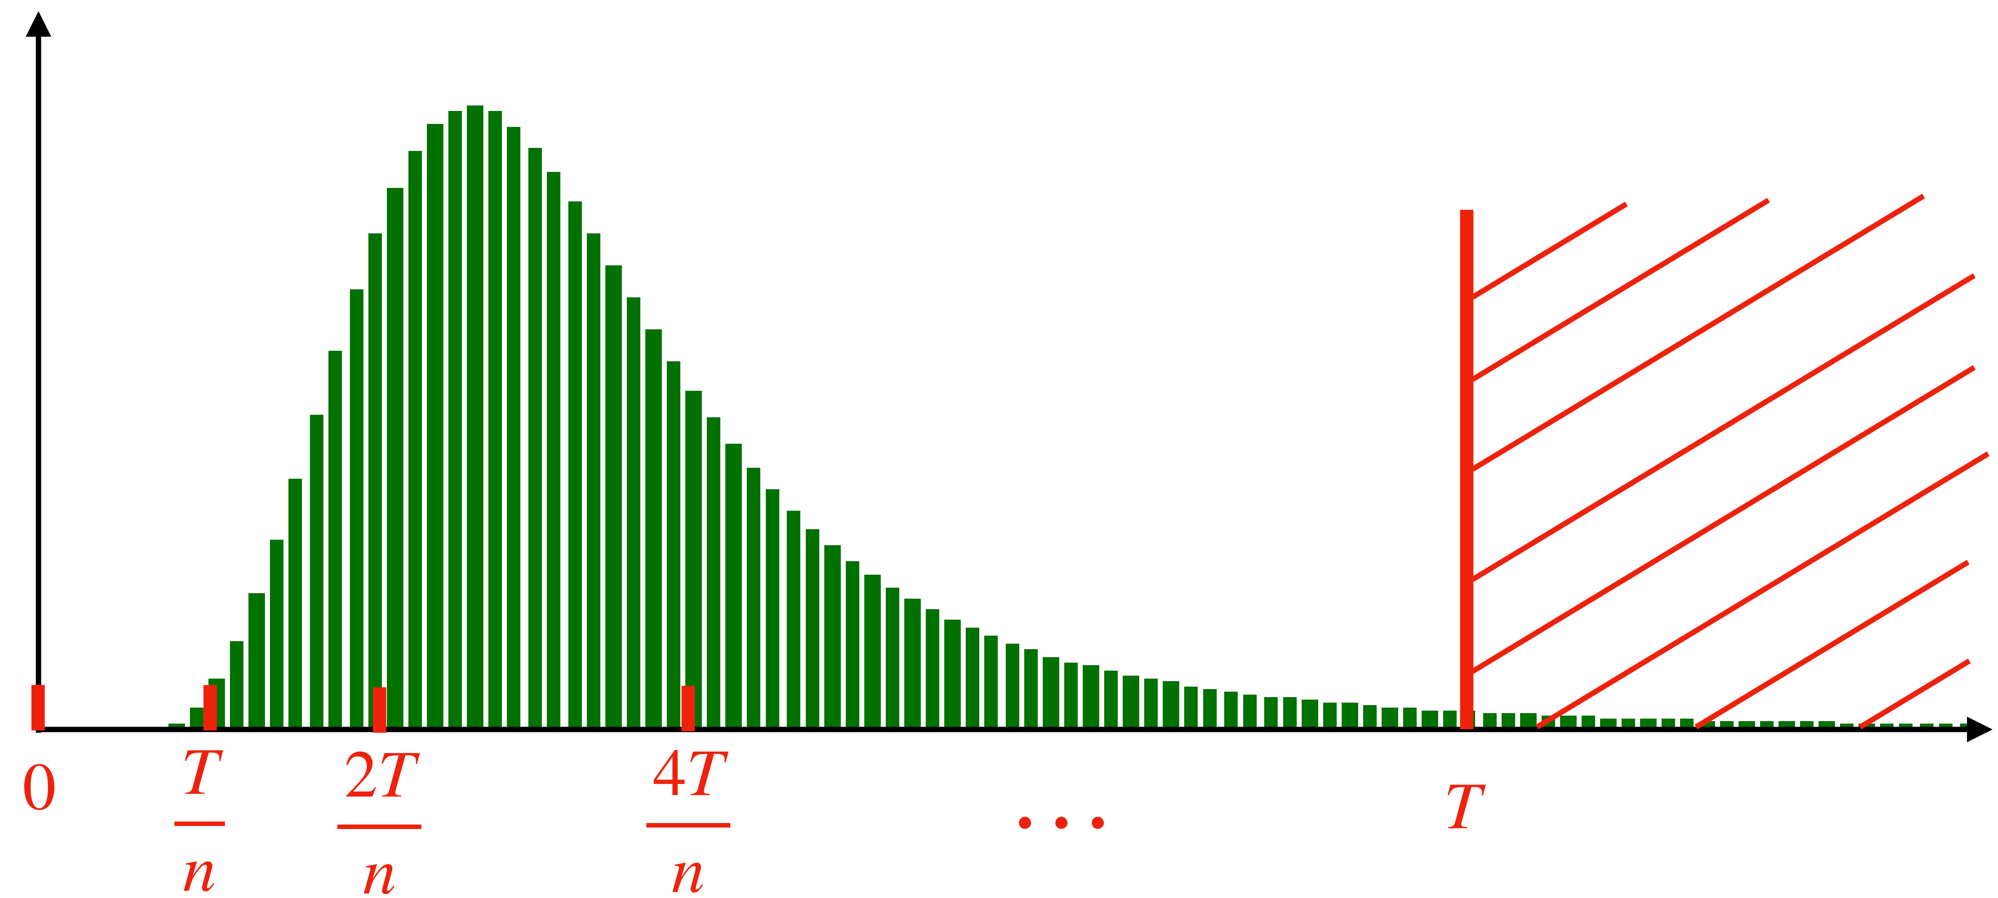

$$\tilde{\mu} \;=\; \tilde{\mu}_1 + \tilde{\mu}_2 + \tilde{\mu}_3 + \cdots + 0$$

Error: $$|\tilde{\mu} - \mu| \lesssim \sum_i \frac{\sqrt{2\mathbb{E}(x^2 \cdot 1_{a_i < x \leq a_{i+1}})}}{n} + \frac{T}{n^2} + \mathbb{E}(x \cdot 1_{x>T})$$

$$\lesssim \frac{\sqrt{\mathbb{E}_D(x^2)}}{n} + \frac{T}{n^2} + \mathbb{E}(x \cdot 1_{x>T})$$

[Heinrich'02]
[Montanaro'15]
[H.,Magniez'19]

5

Quantum “sub-Gaussian” estimator

Choosing the truncation level

For any T: $|\tilde{\mu} - \mu| \lesssim \frac{\sqrt{\mathbb{E}(x^2)}}{n} + \frac{T}{n^2} + \mathbb{E}(x \cdot 1_{x>T})$

First choice: $T \approx n\sqrt{\mathbb{E}(x^2)}$ [Heinrich'02] [Montanaro'15] [H.,Magniez'19]

✓ The "omitted" part is small: $\mathbb{E}(x \cdot 1_{x>T}) \leq \frac{\mathbb{E}(x^2)}{T} \leq \frac{\sqrt{\mathbb{E}(x^2)}}{n}$

✗ T depends on an unknown quantity

Second choice: $T \approx$ quantile such that $\Pr[x > T] = 1/n^2$ [H.'21]

✓ T is small (*Markov*): $\Pr[x > T] \leq \mathbb{E}(x^2)/T^2 \Rightarrow T \leq n\sqrt{\mathbb{E}(x^2)}$

✓ The "omitted" part is small (*Cauchy-Schwarz*): $\ldots \leq \sqrt{\mathbb{E}(x^2) \cdot \Pr(x > T)}$

✓ T can be computed without any prior knowledge about D

Quantile finding algorithm

Goal: find T such that $\Pr[x > T] = 1/n^2$

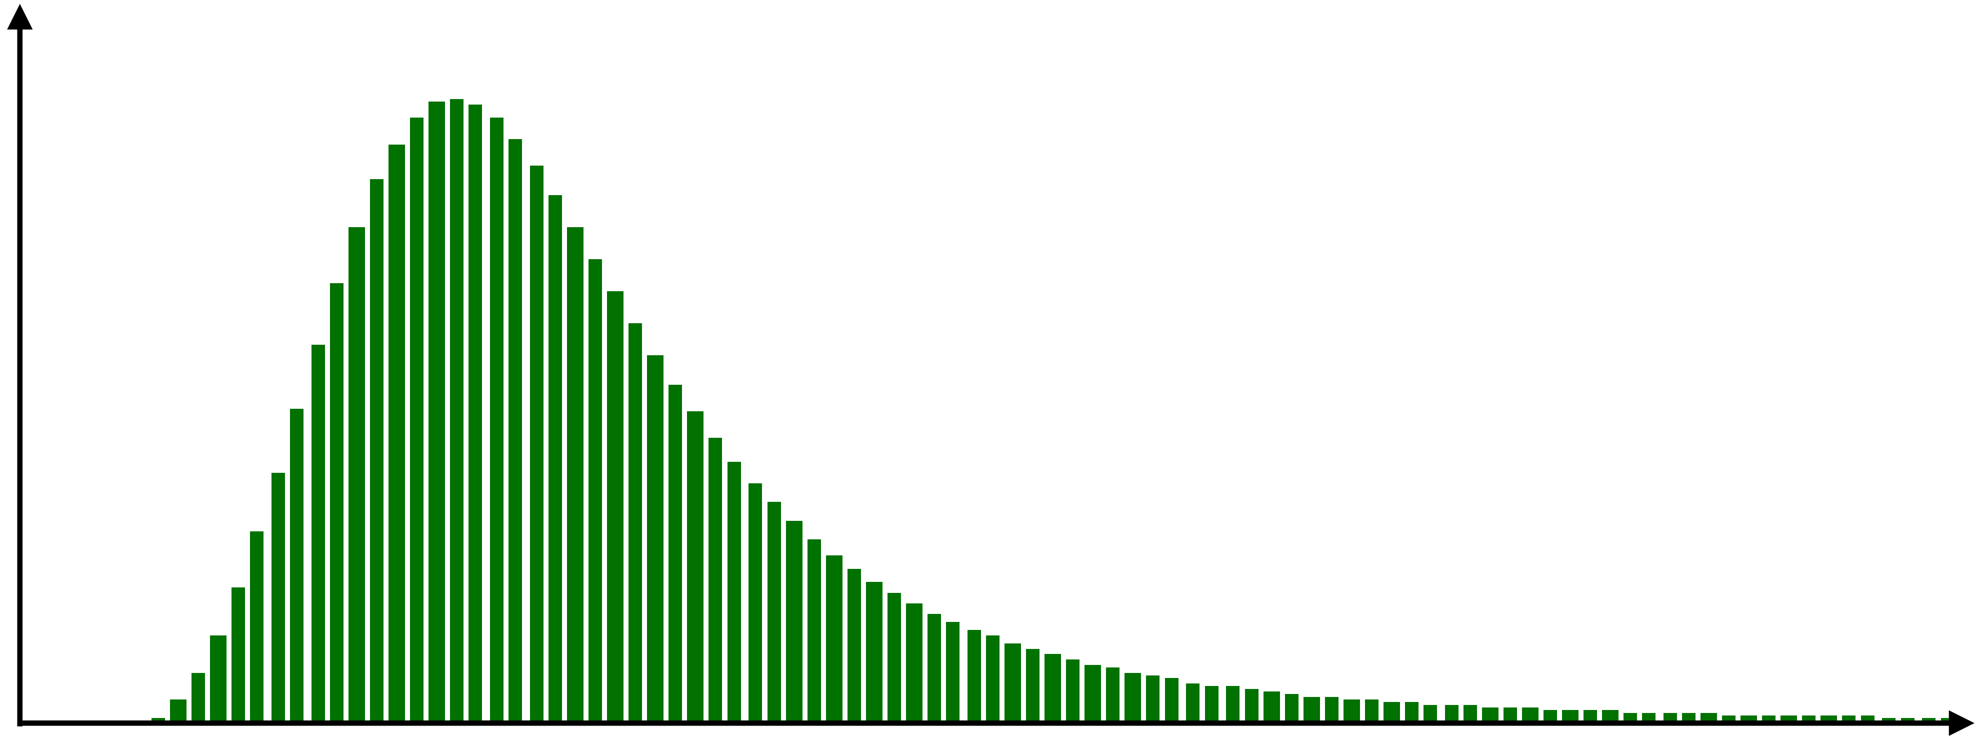

Quantile finding algorithm

Goal: find T such that $\Pr[x > T] = 1/n^2$

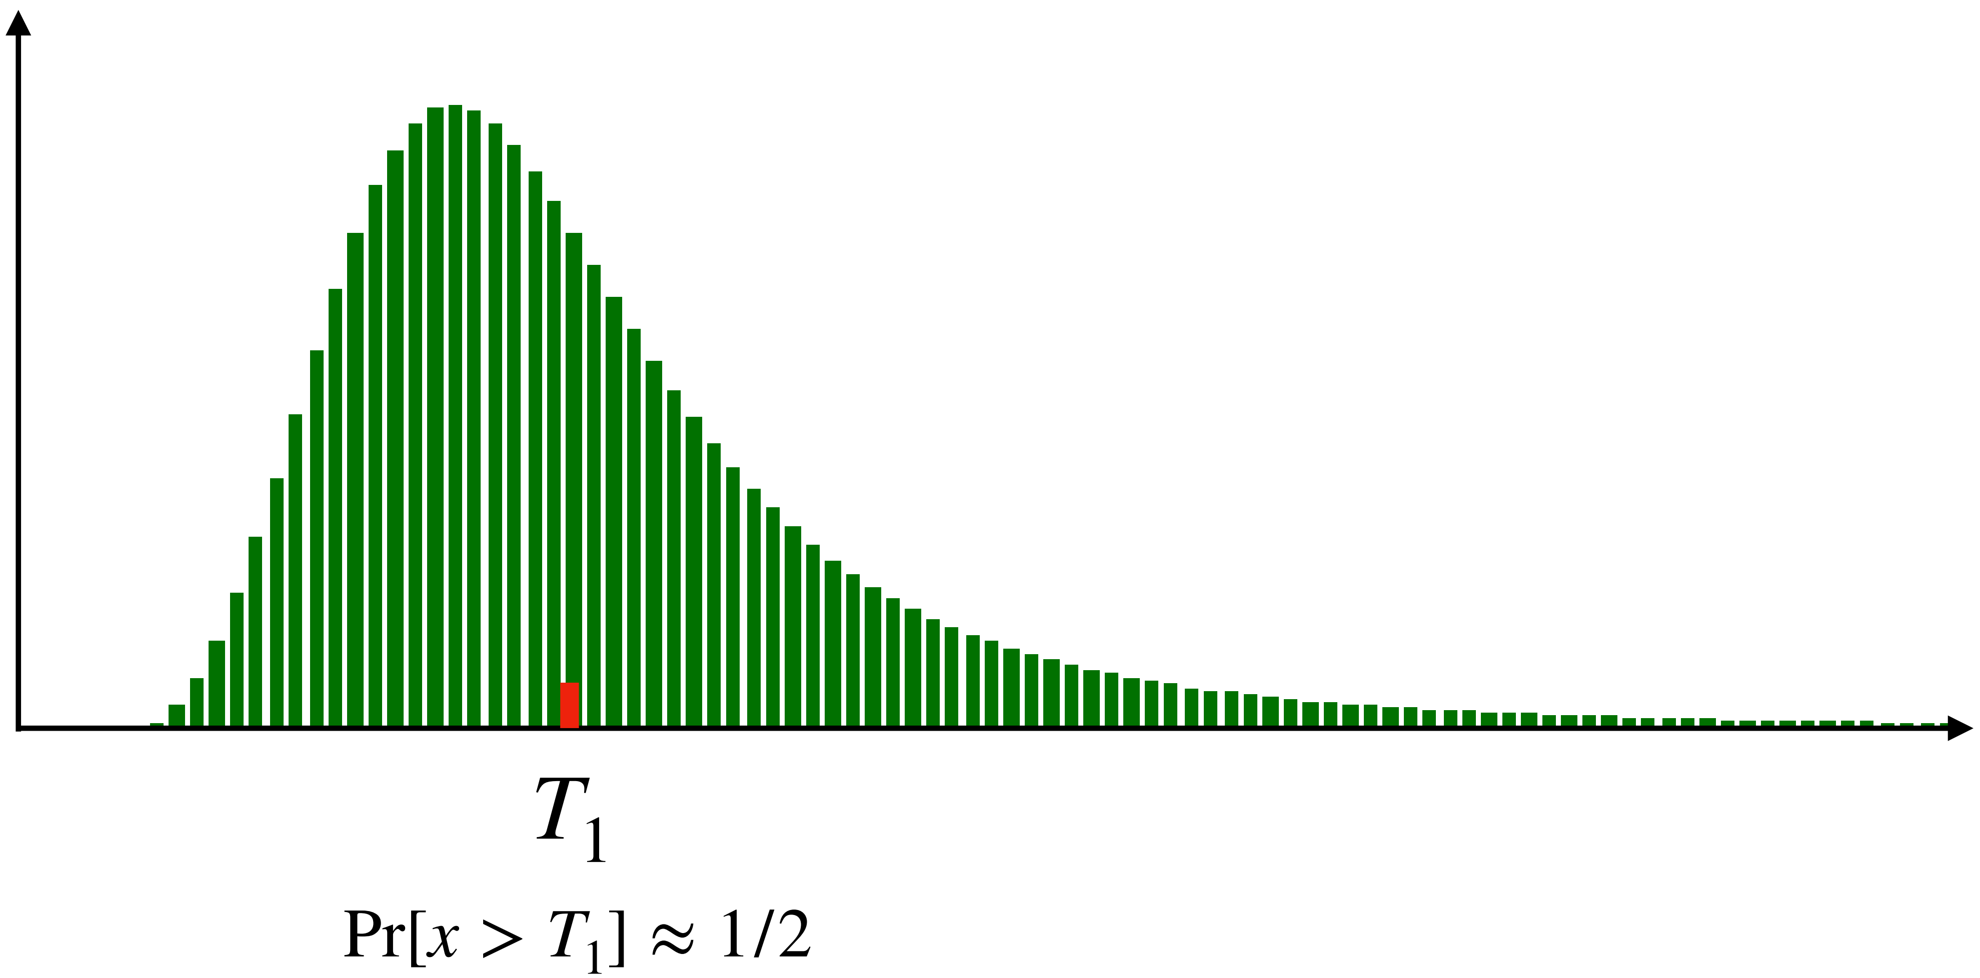

Quantile finding algorithm

Goal: find T such that $\Pr[x > T] = 1/n^2$

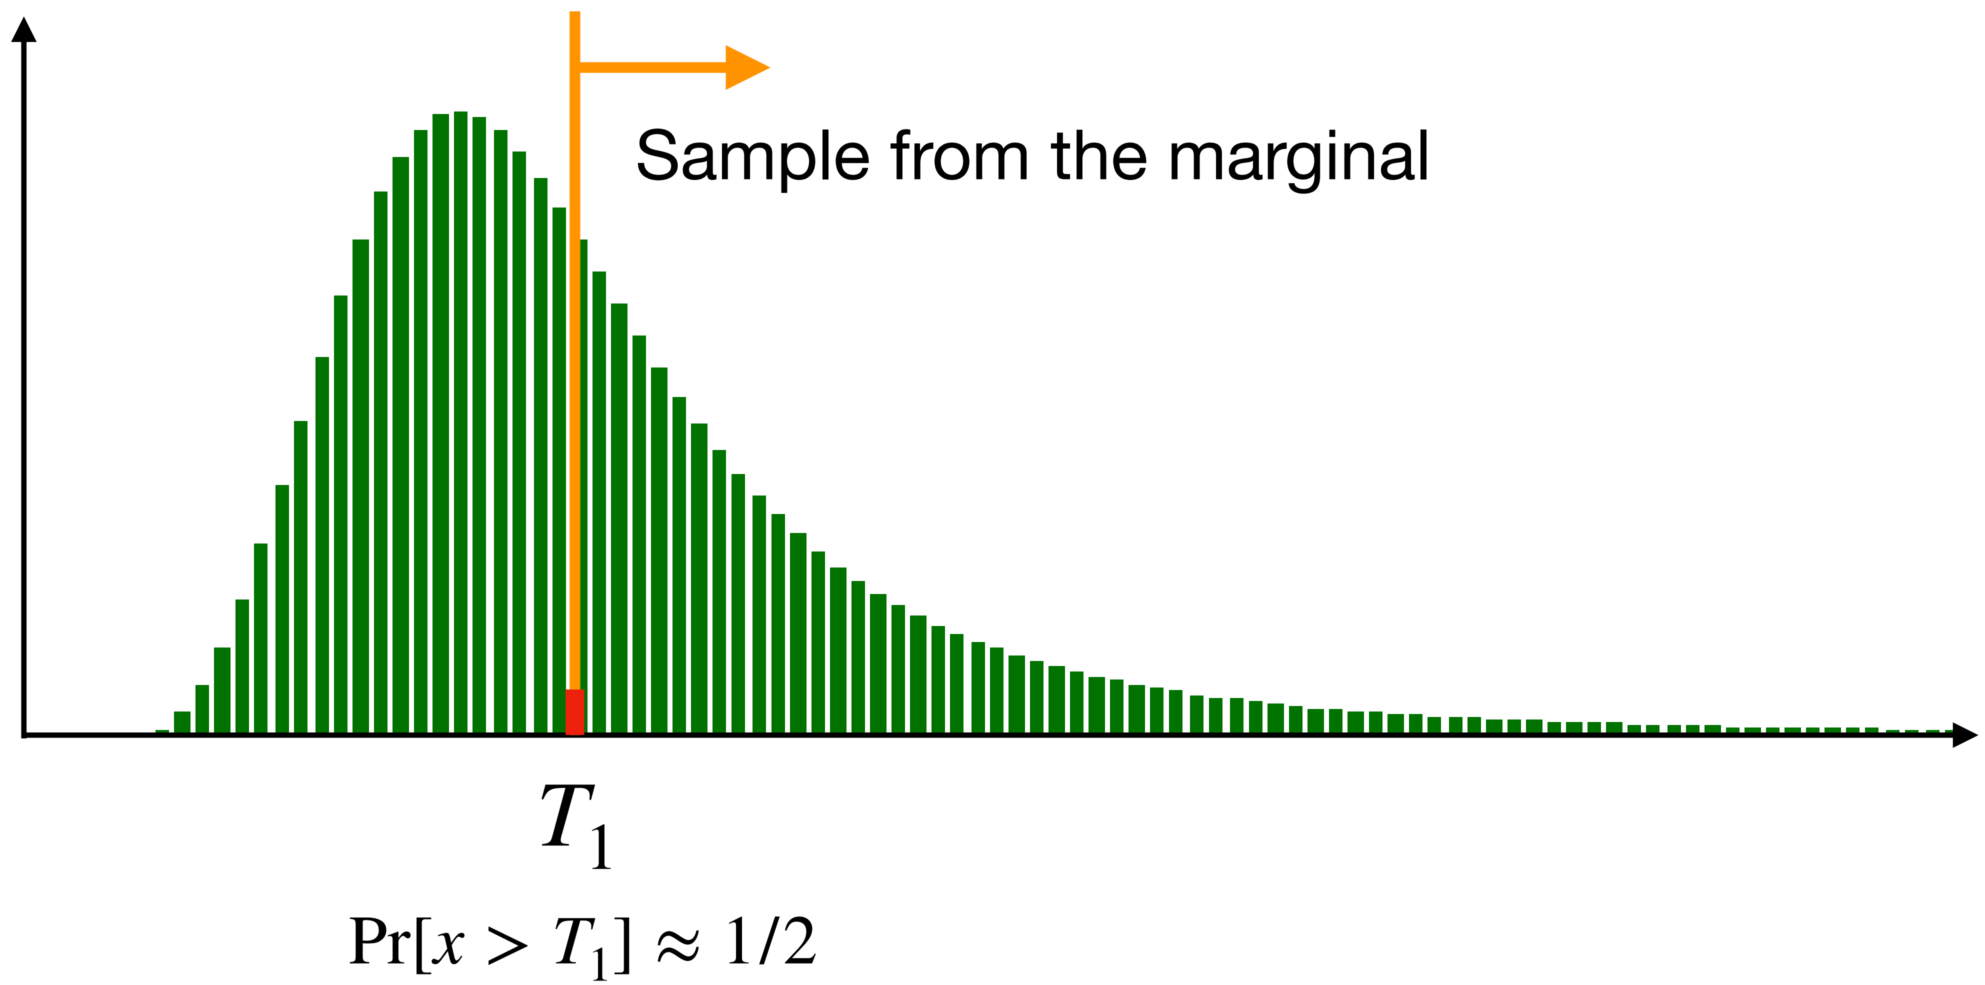

Quantile finding algorithm

Goal: find T such that $\Pr[x > T] = 1/n^2$

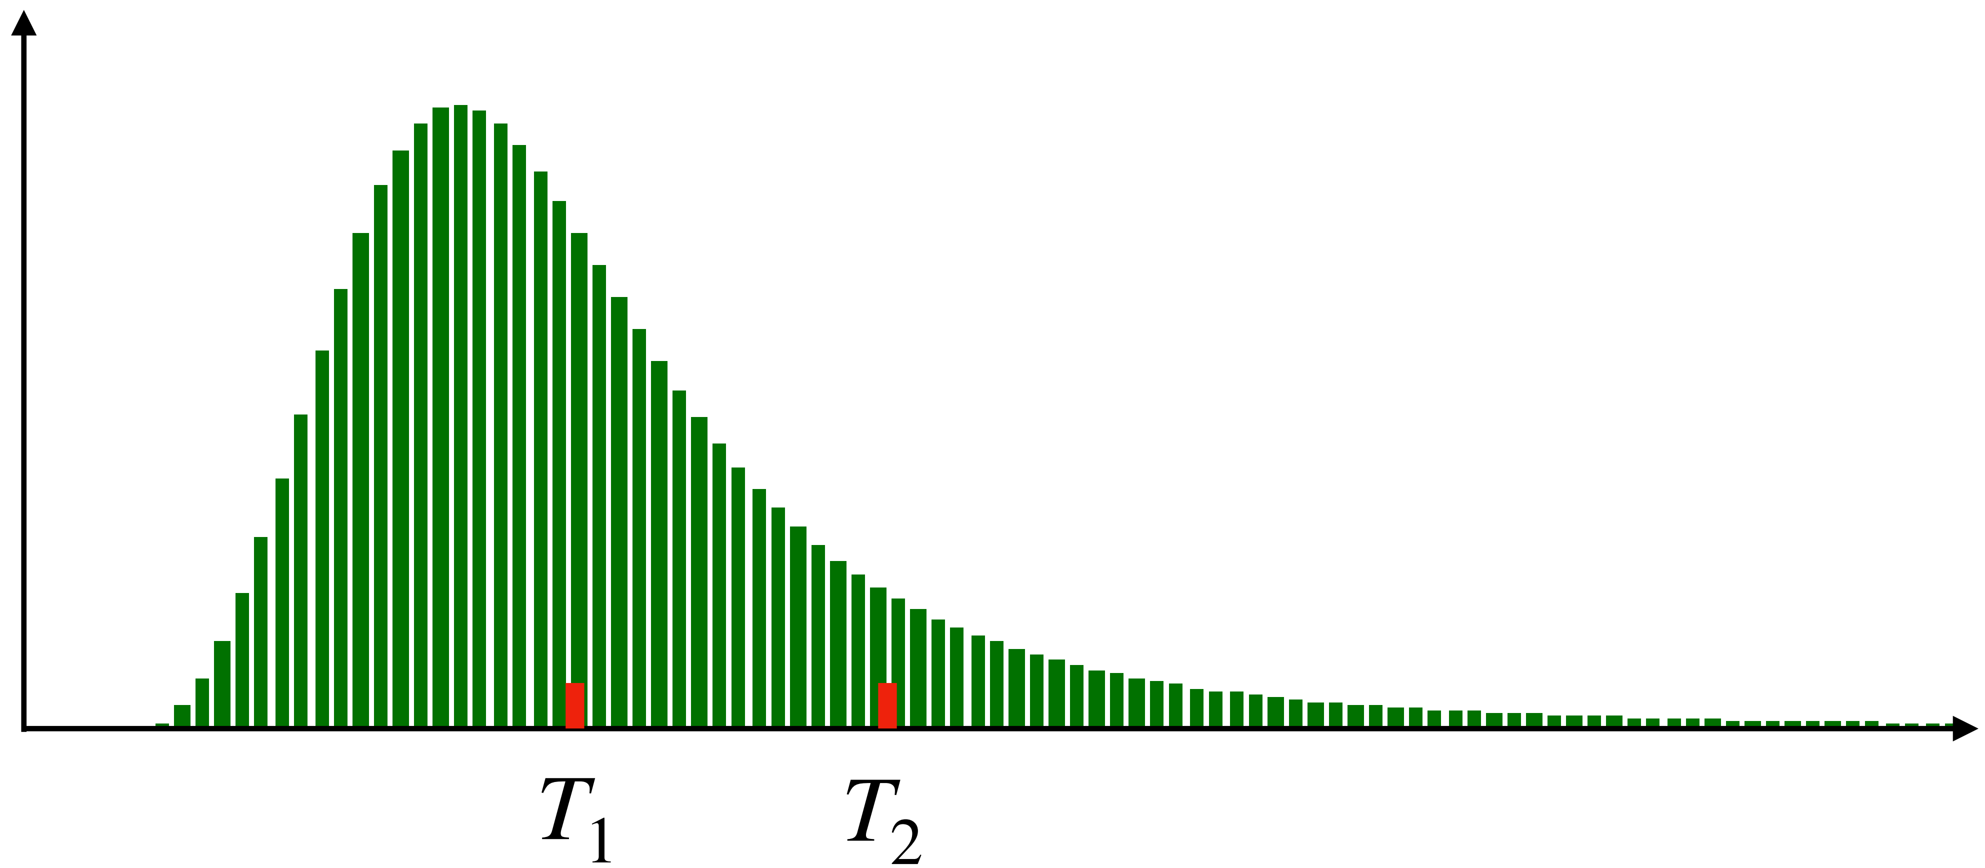

$\Pr[x > T_2] \approx 1/4$

Quantile finding algorithm

Goal: find T such that $\Pr[x > T] = 1/n^2$

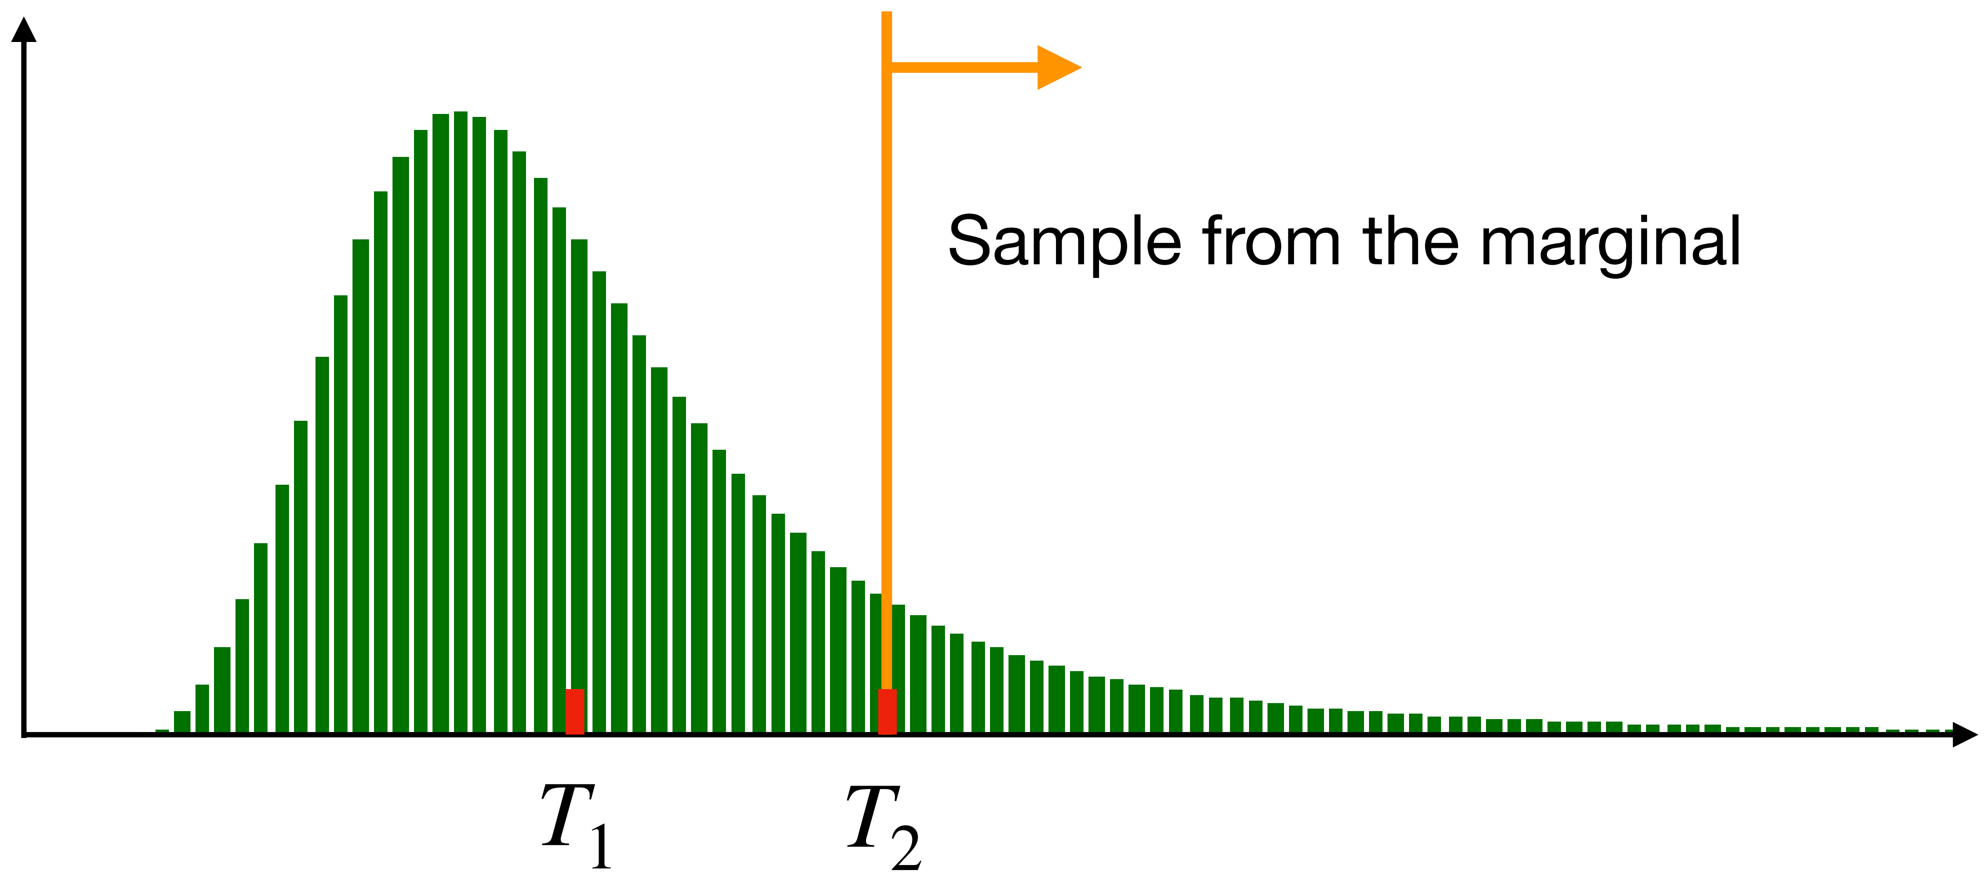

$\Pr[x > T_2] \approx 1/4$

Quantile finding algorithm

Goal: find T such that $\Pr[x > T] = 1/n^2$

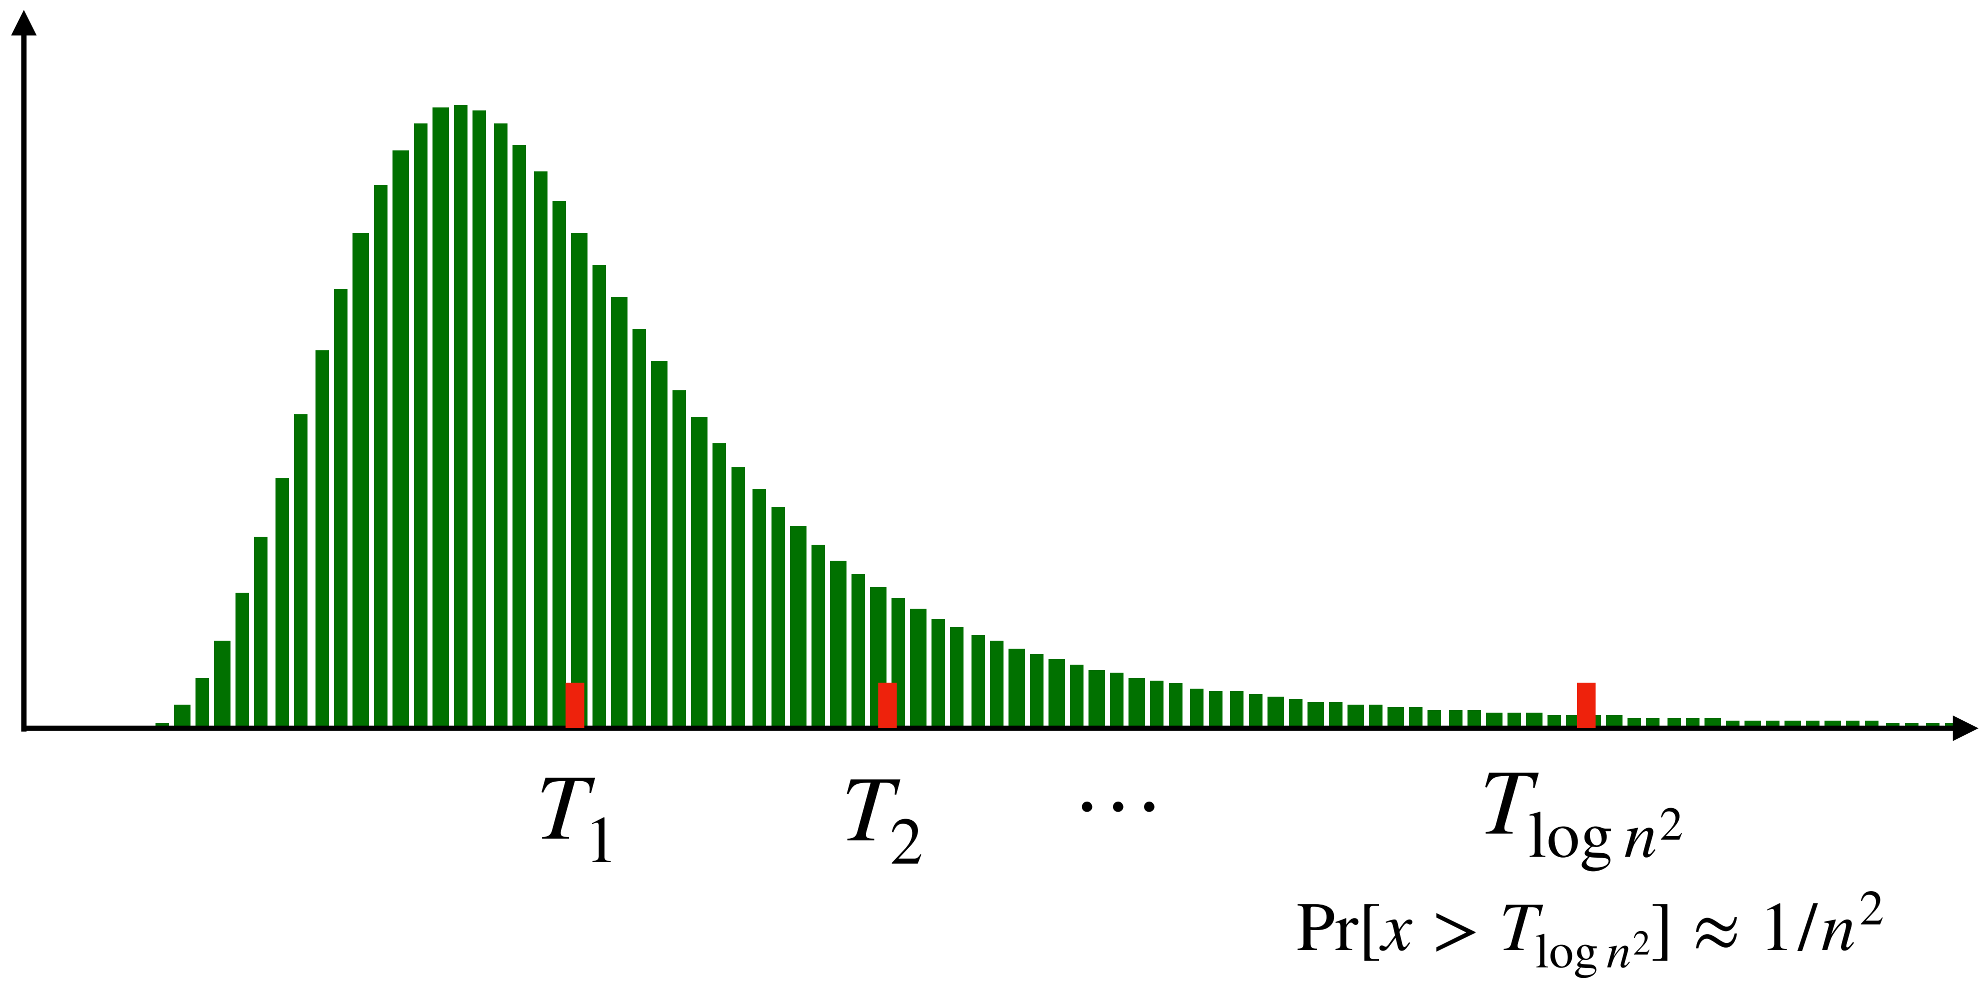

Quantile finding algorithm

Goal: find T such that $\Pr[x > T] = 1/n^2$

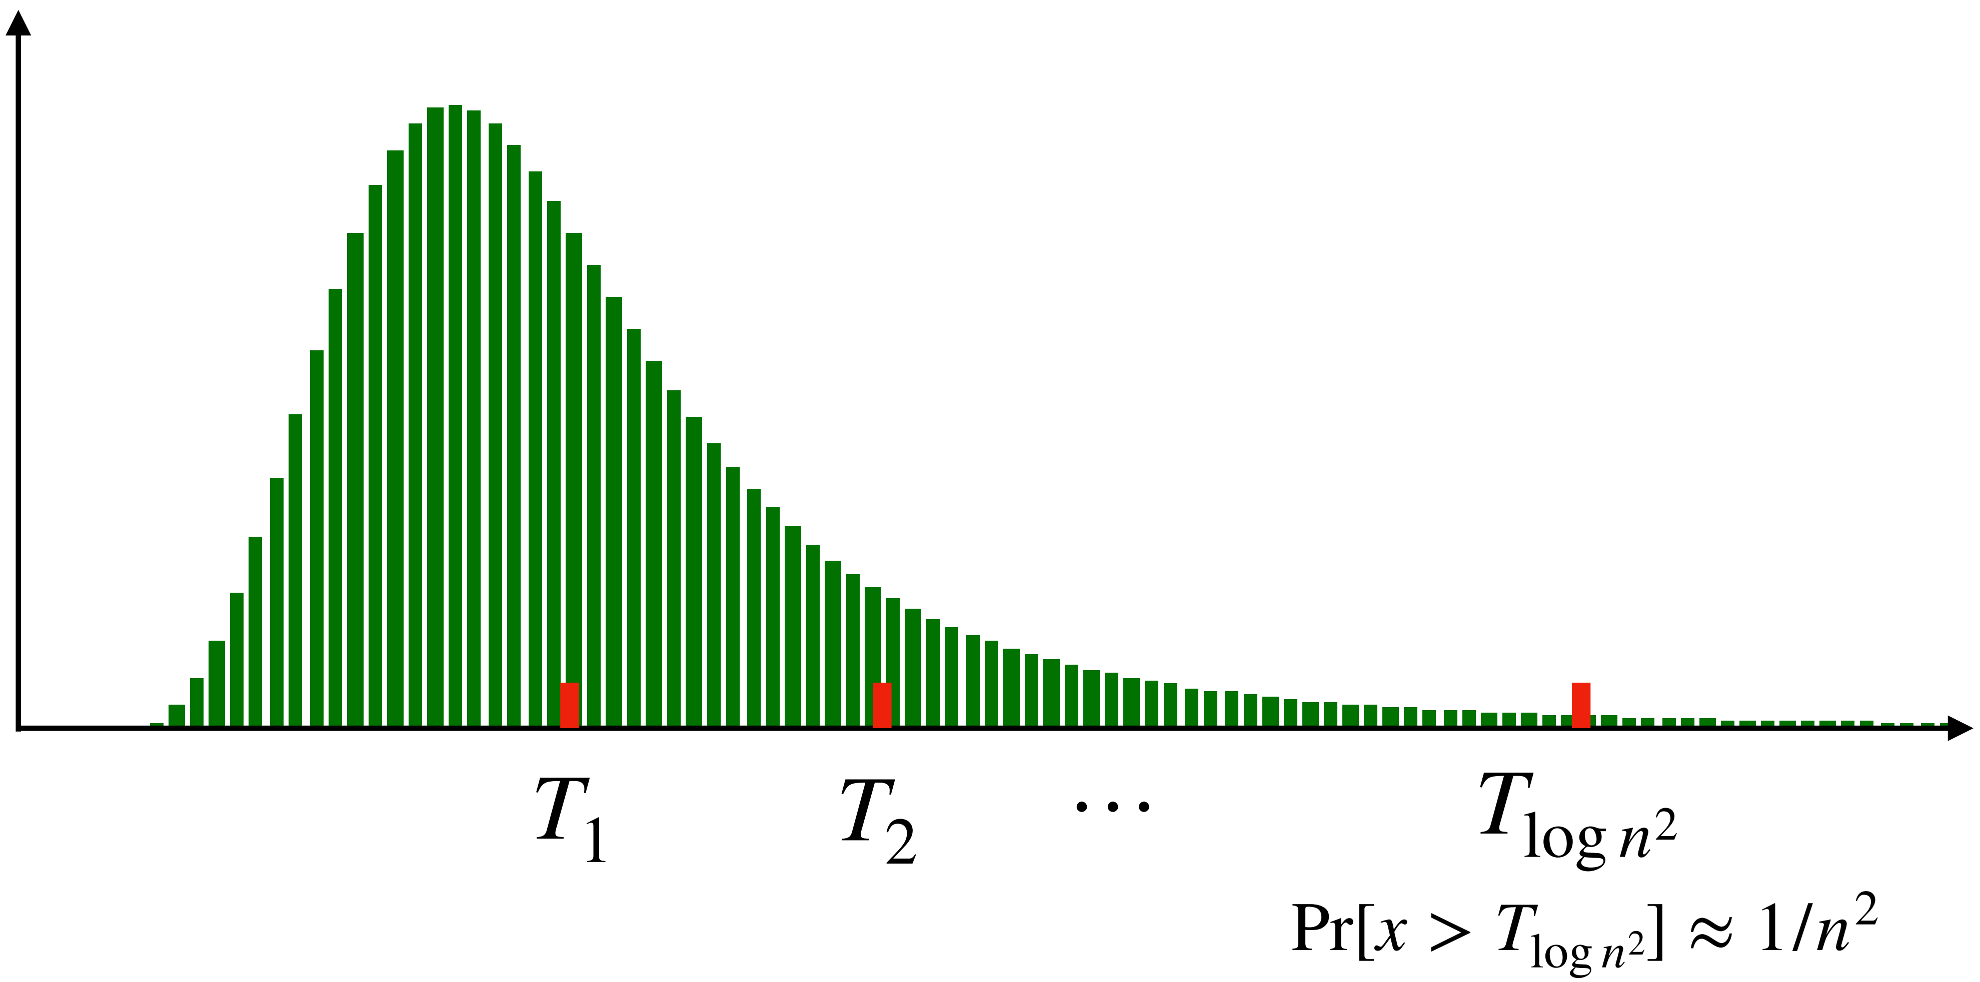

Only $\approx \log n$ steps but sampling from the marginal is harder for larger T_i's

Cost of last step: sampling $x \sim D$ conditioned on an event of proba. $\approx 1/n^2$

→ $\tilde{O}(n^2)$ with classical samples, $\tilde{O}(n)$ with **Amplitude Amplification**

6

Future work

Variance reduction

Distribution D with mean μ and variance σ^2

Average of n samples →

Distribution D' with mean μ and variance σ^2/n

↓ **n quantum experiments**

Distribution D'' with mean μ and variance σ^2/n^2 ?

- Alternative route to a quantum sub-Gaussian estimator
- Application in statistical physics for estimating **partition functions**

Multivariate mean estimation

Distribution supported over $\mathbb{R}^d$ where $d > 1$.

covariance matrix

Best classical error rate (in L_2-norm): $O\left(\sqrt{\frac{\mathrm{Tr}(\Sigma)}{n}} + \sqrt{\frac{\lambda_{\max}(\Sigma)\log(1/\delta)}{n}}\right)$

(+ computationally efficient [Hopkins'18])

Best quantum error rate: ?

[Heinrich'04] No possible speedup when $n \leq d$ for some distributions

[Cornelissen,Jerbi'21] Speedups when $n \geq d$ for some distributions

$$\min\left(\text{classical}, \frac{\sqrt{d\mathrm{Tr}(\Sigma)}\log(1/\delta)}{n}\right) ?$$

Multivariate mean estimation

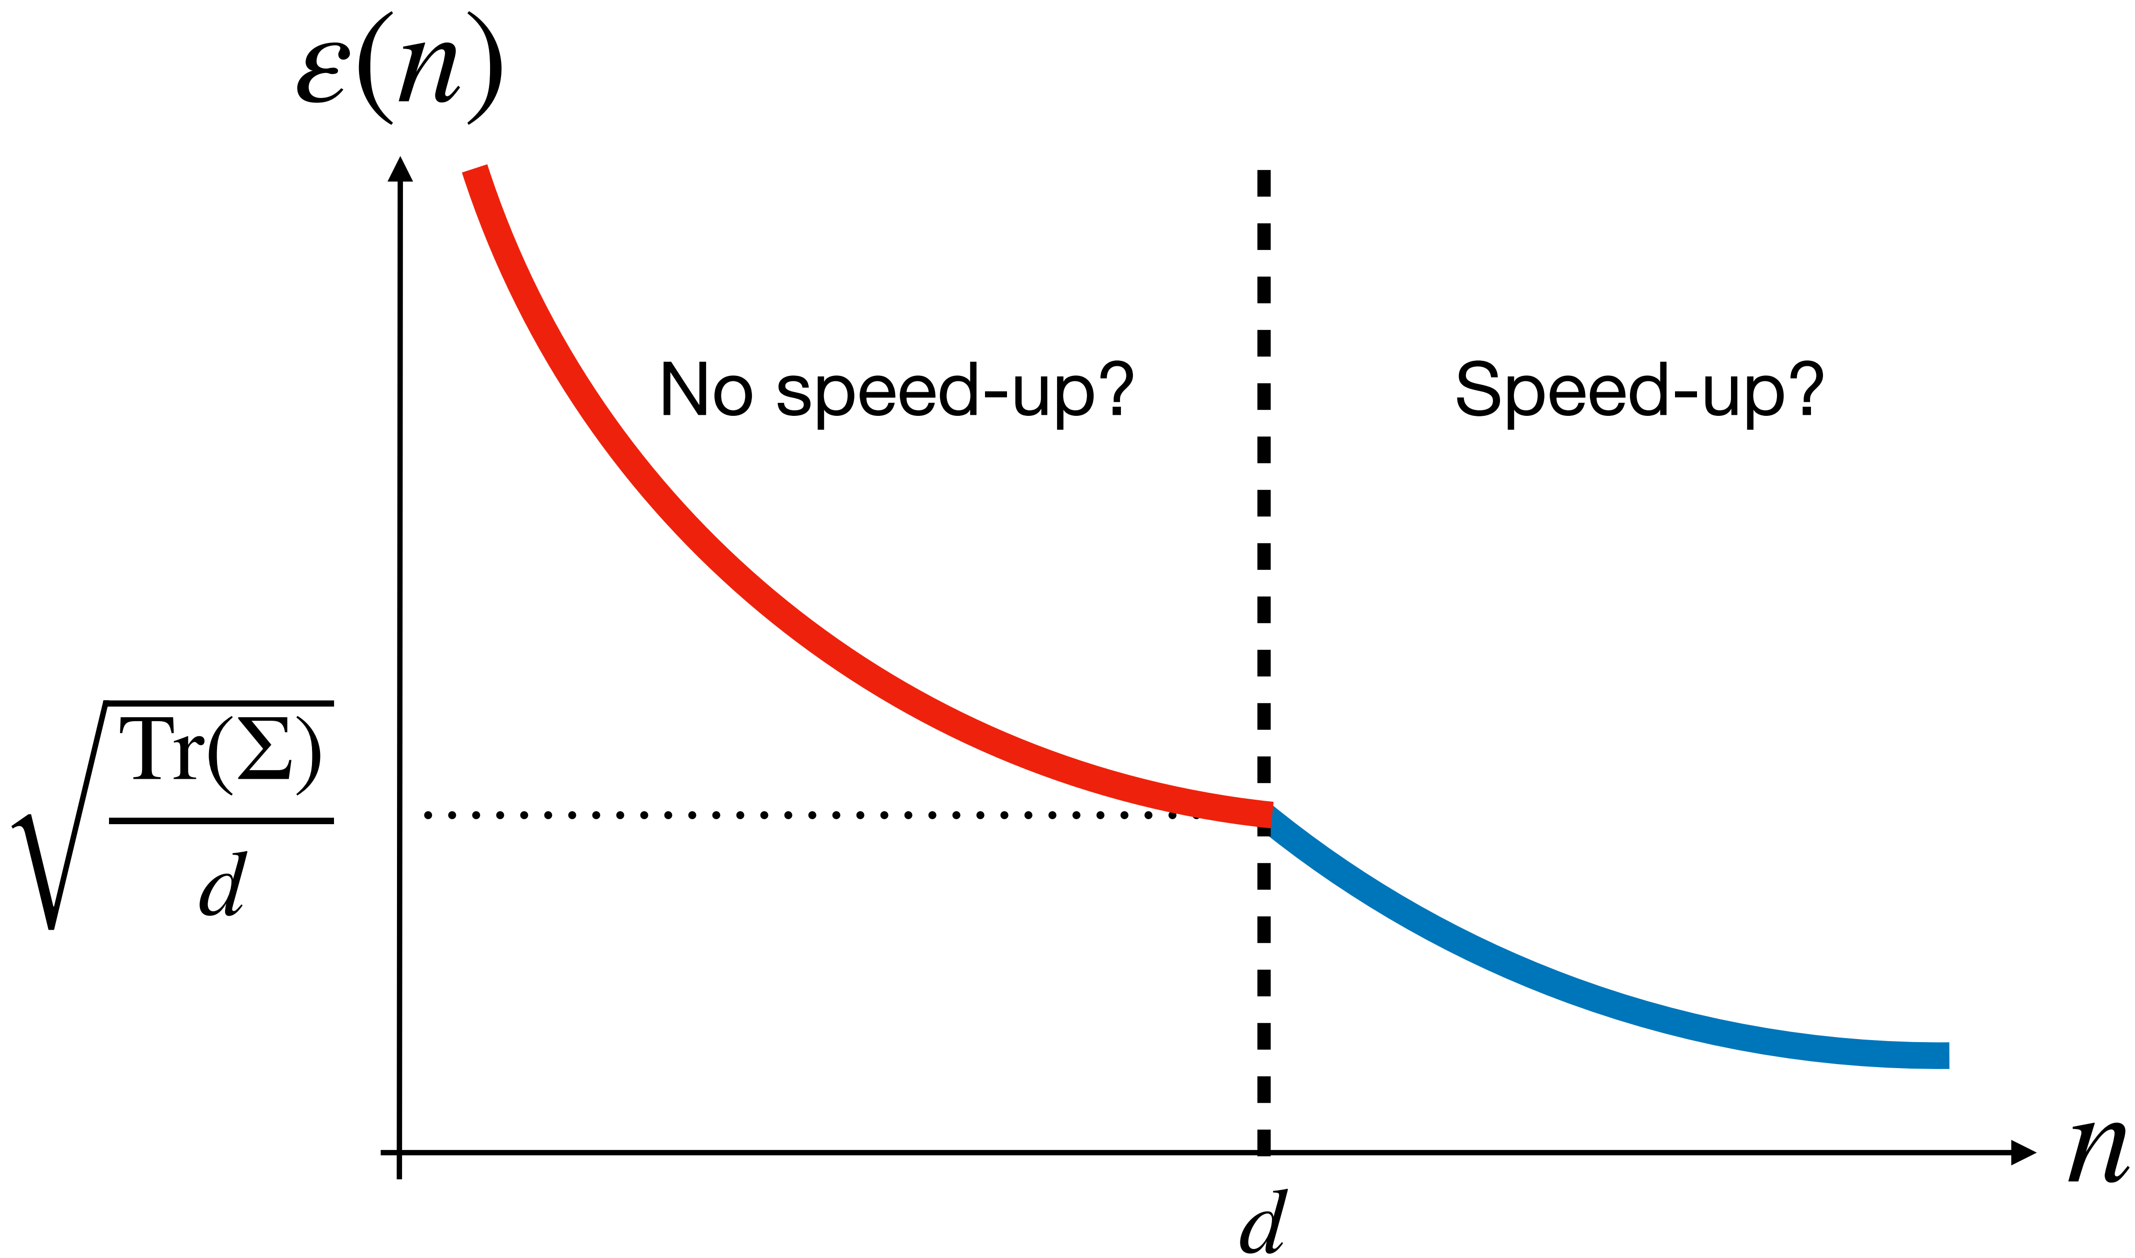

$$\min\left(\text{classical}, \frac{\sqrt{d\,\mathrm{Tr}(\Sigma)}\log(1/\delta)}{n}\right) ?$$

Further readings

[Brassard,Höyer, Mosca,Tapp'02] *Quantum Amplitude Amplification and Estimation*

[Heinrich'02] *Quantum Summation with an Application to Integration*

[Montanaro'15] *Quantum Speedup of Monte Carlo Methods*

[Lugosi,Mendelson'19] *Mean Estimation and Regression Under Heavy-Tailed Distributions: A Survey*

[H.,Magniez'19] *Quantum Chebyshev's Inequality and Applications*

[Harrow,Wei'20] *Adaptive Quantum Simulated Annealing for Bayesian Inference and Estimating Partition Functions*

[H.'21] *Quantum Sub-Gaussian Mean Estimator*

[Cornelissen,Jerbi'21] *Quantum algorithms for multivariate Monte Carlo estimation*