Quantum Speedups for Computing Expectation Values and Partition Functions

Yassine Hamoudi

UC Berkeley

APS March Meeting 2023

Invited talk (Breiman Lecture) at NeurIPS 2021

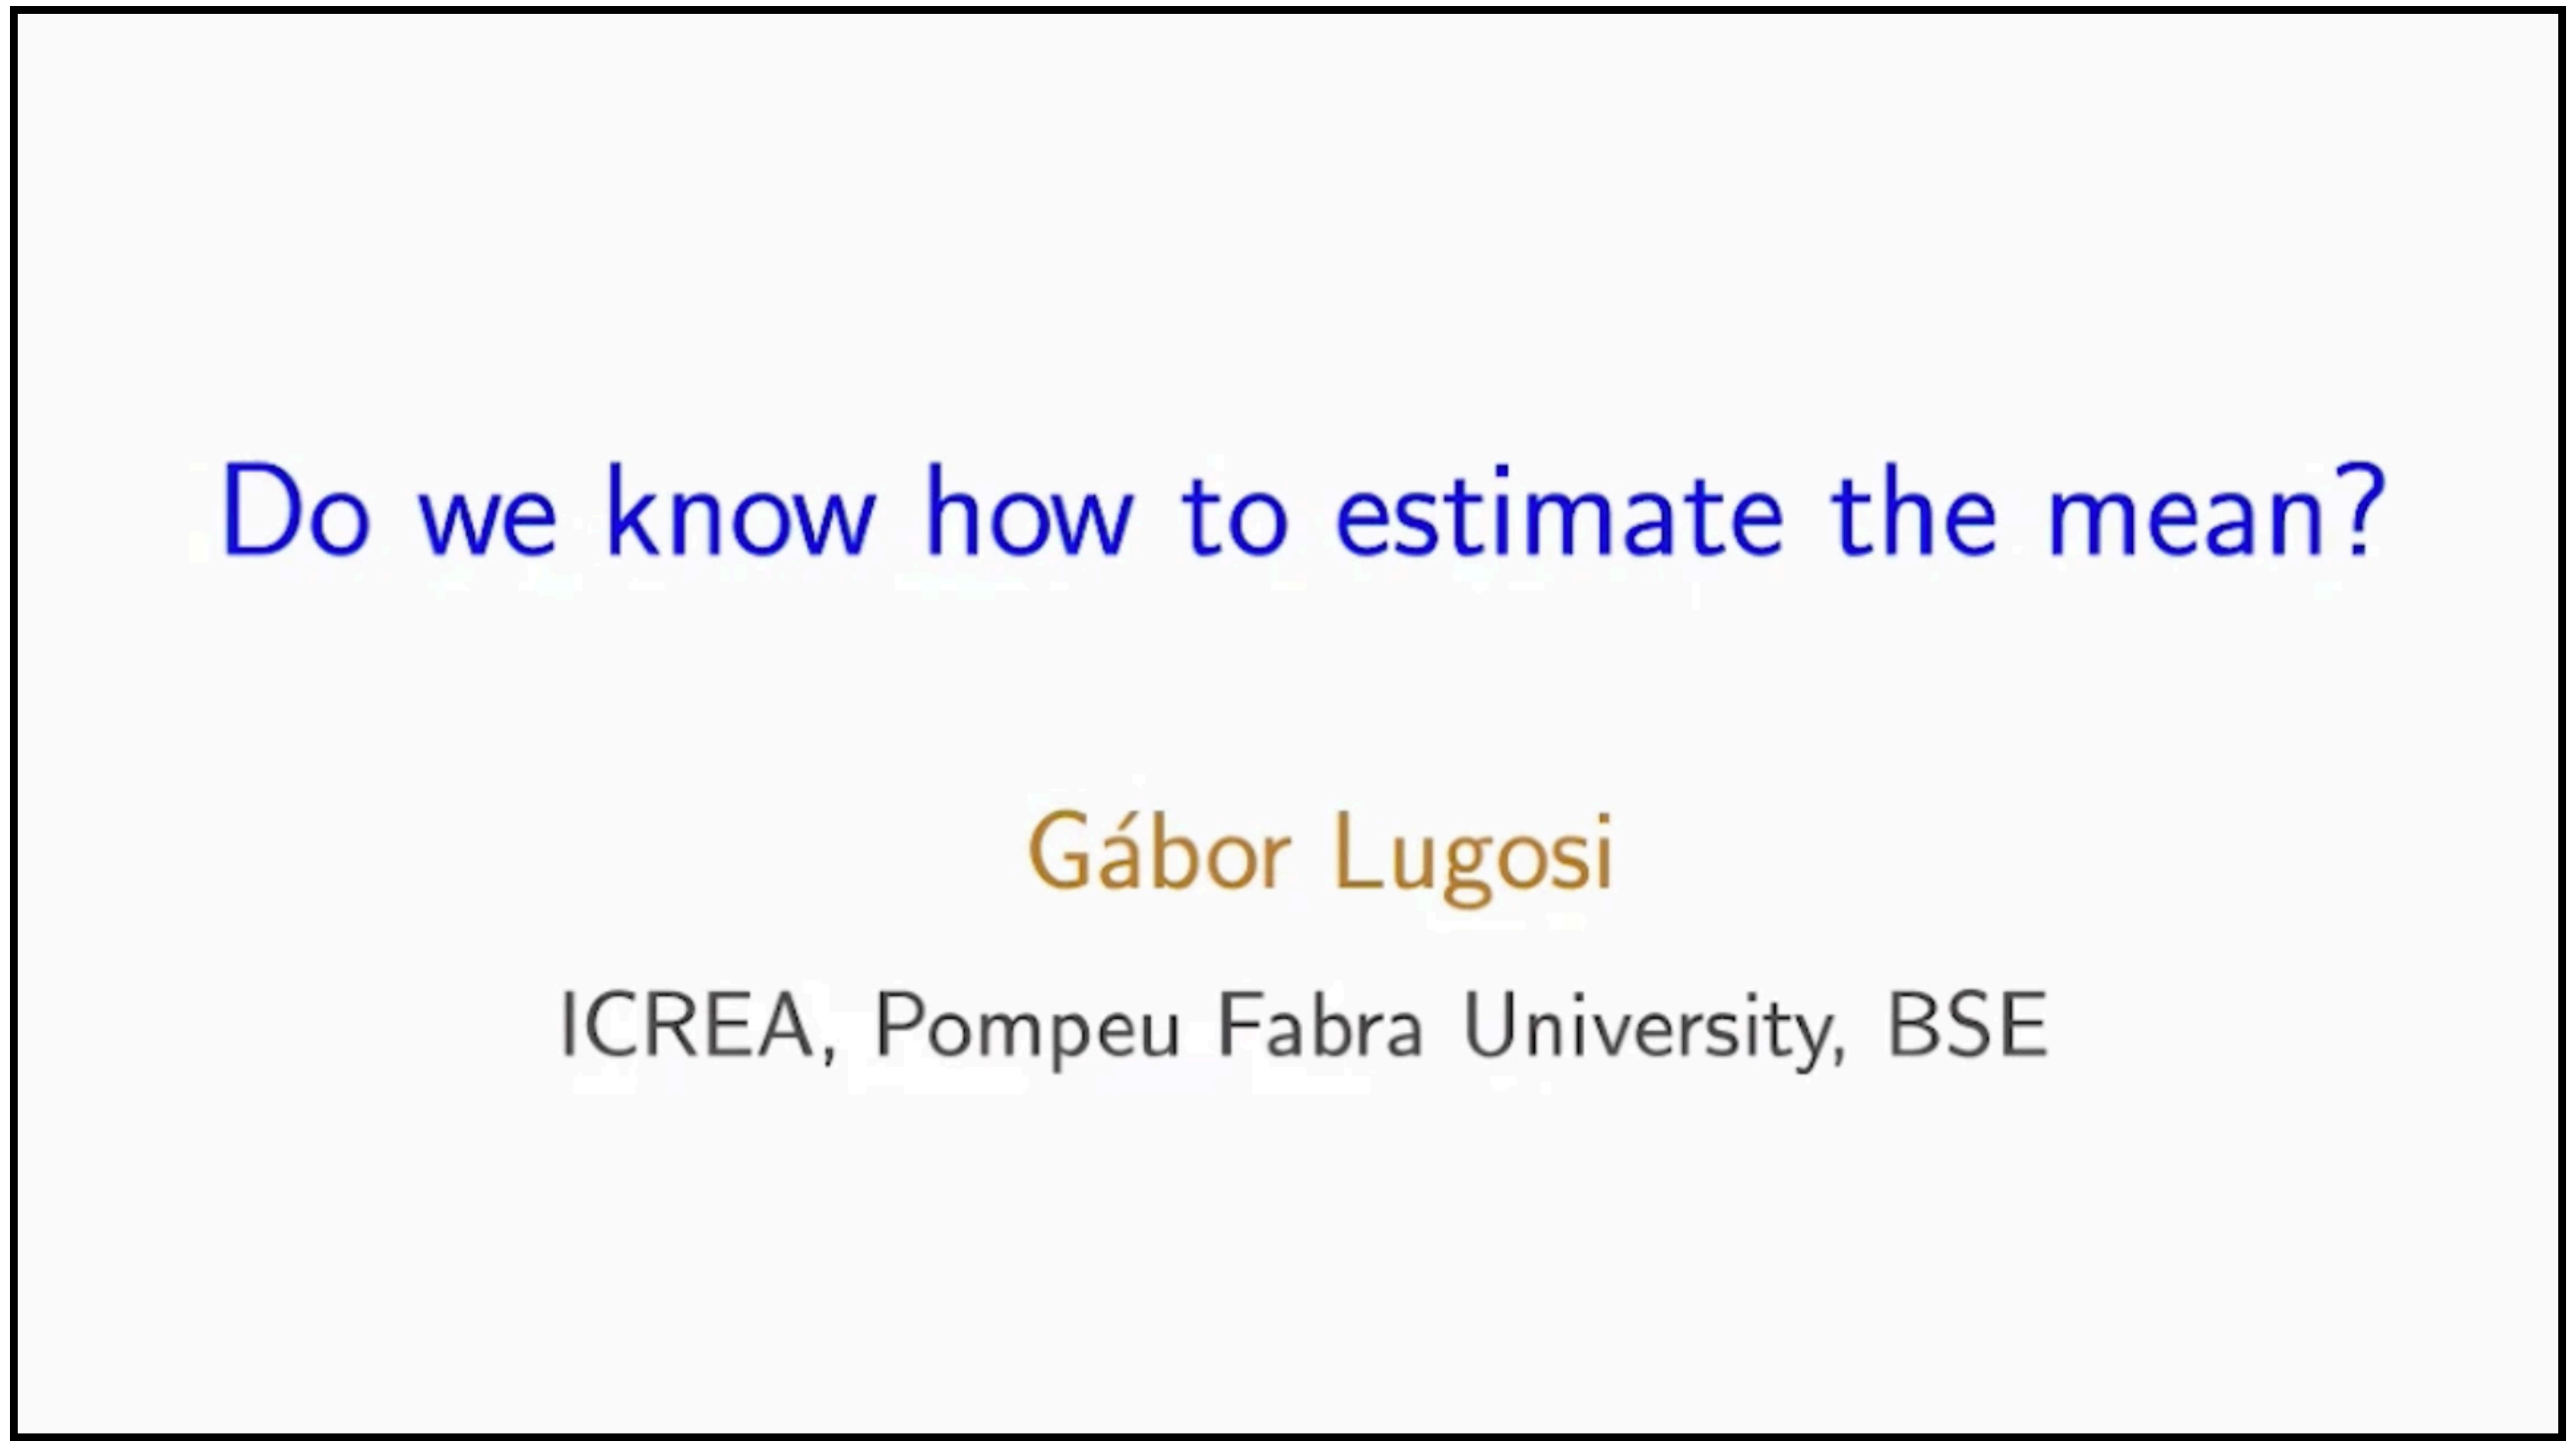

"Despite its long history, the subject has attracted a flurry of renewed activity. Motivated by applications in machine learning and data science, the problem has been viewed from new angles both from statistical and computational points of view."

Scenarios in quantum computing

Expectation value of an observable

$$\langle \psi | O | \psi \rangle$$

Mean of a quantum probability oracle

$$|0\rangle \mapsto \sum_{\omega} \sqrt{p_\omega}\, |\omega\rangle\, |X(\omega)\rangle$$

Shadow tomography

$$|\psi\rangle^{\otimes k} \to 01101\ldots \to \left\{ \langle \psi | O_i | \psi \rangle \right\}_i$$

Partition function of a Hamiltonian

$$Z(\beta) = \mathrm{Tr}\left(e^{-\beta H}\right)$$

Minimize the use of sample data

Can we do better than measuring repeated copies of $|\psi\rangle$?

The case of projectors

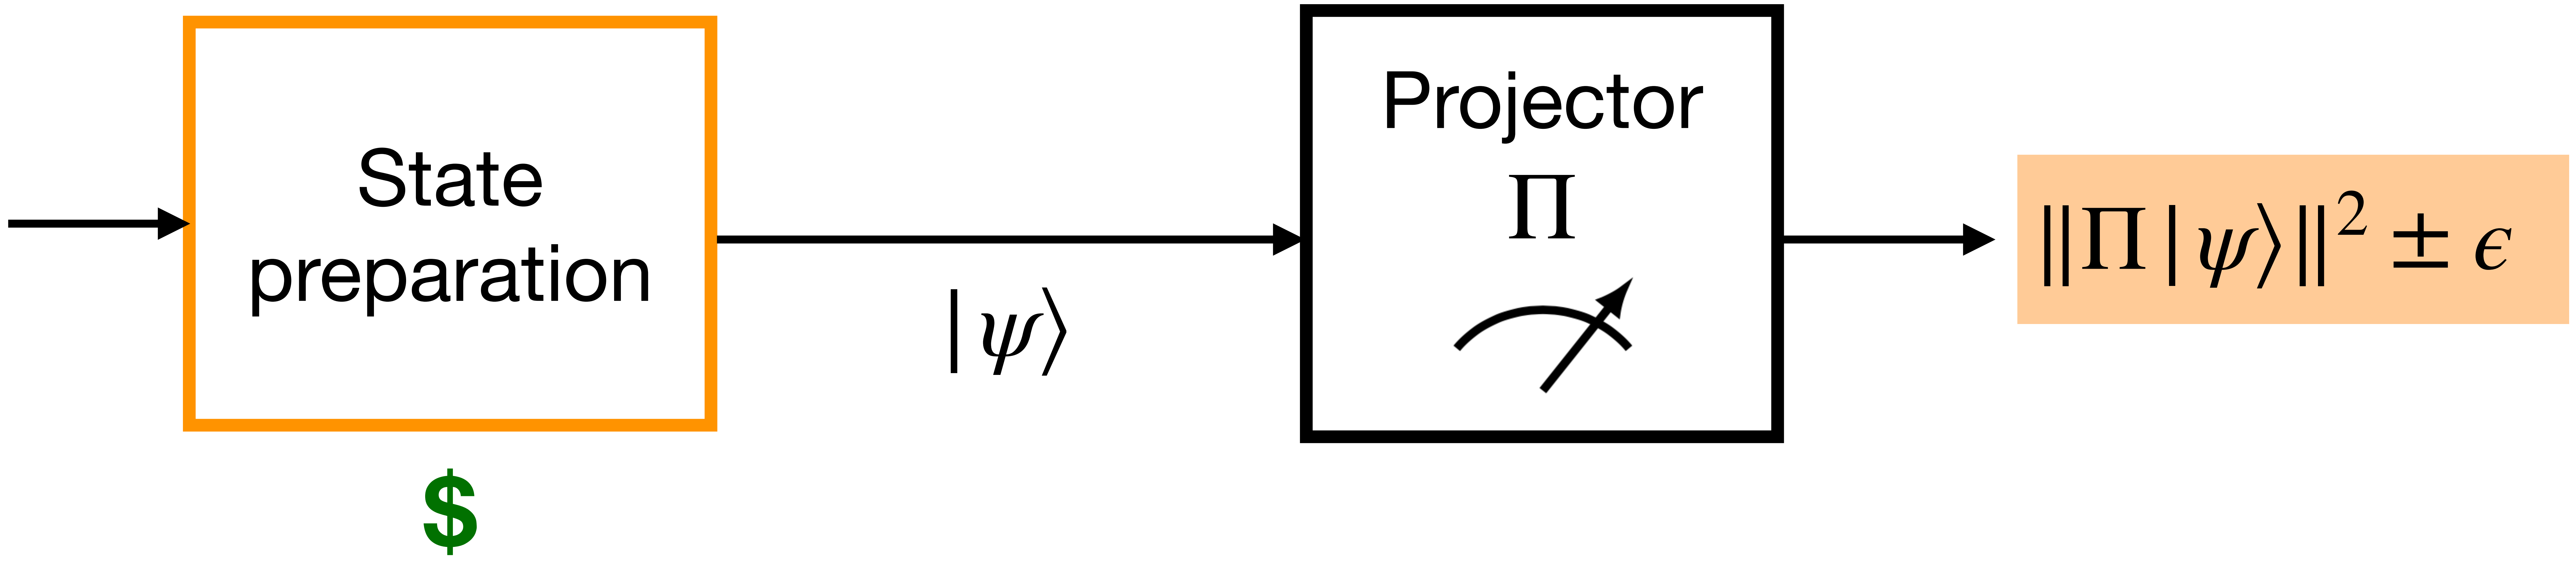

Repeated measurements

$\sim 1/\epsilon^2 \times \$$

(standard quantum limit)

Quantum Phase estimation

$\sim 1/\epsilon \times \$$

(Heisenberg limit)

The case of projectors

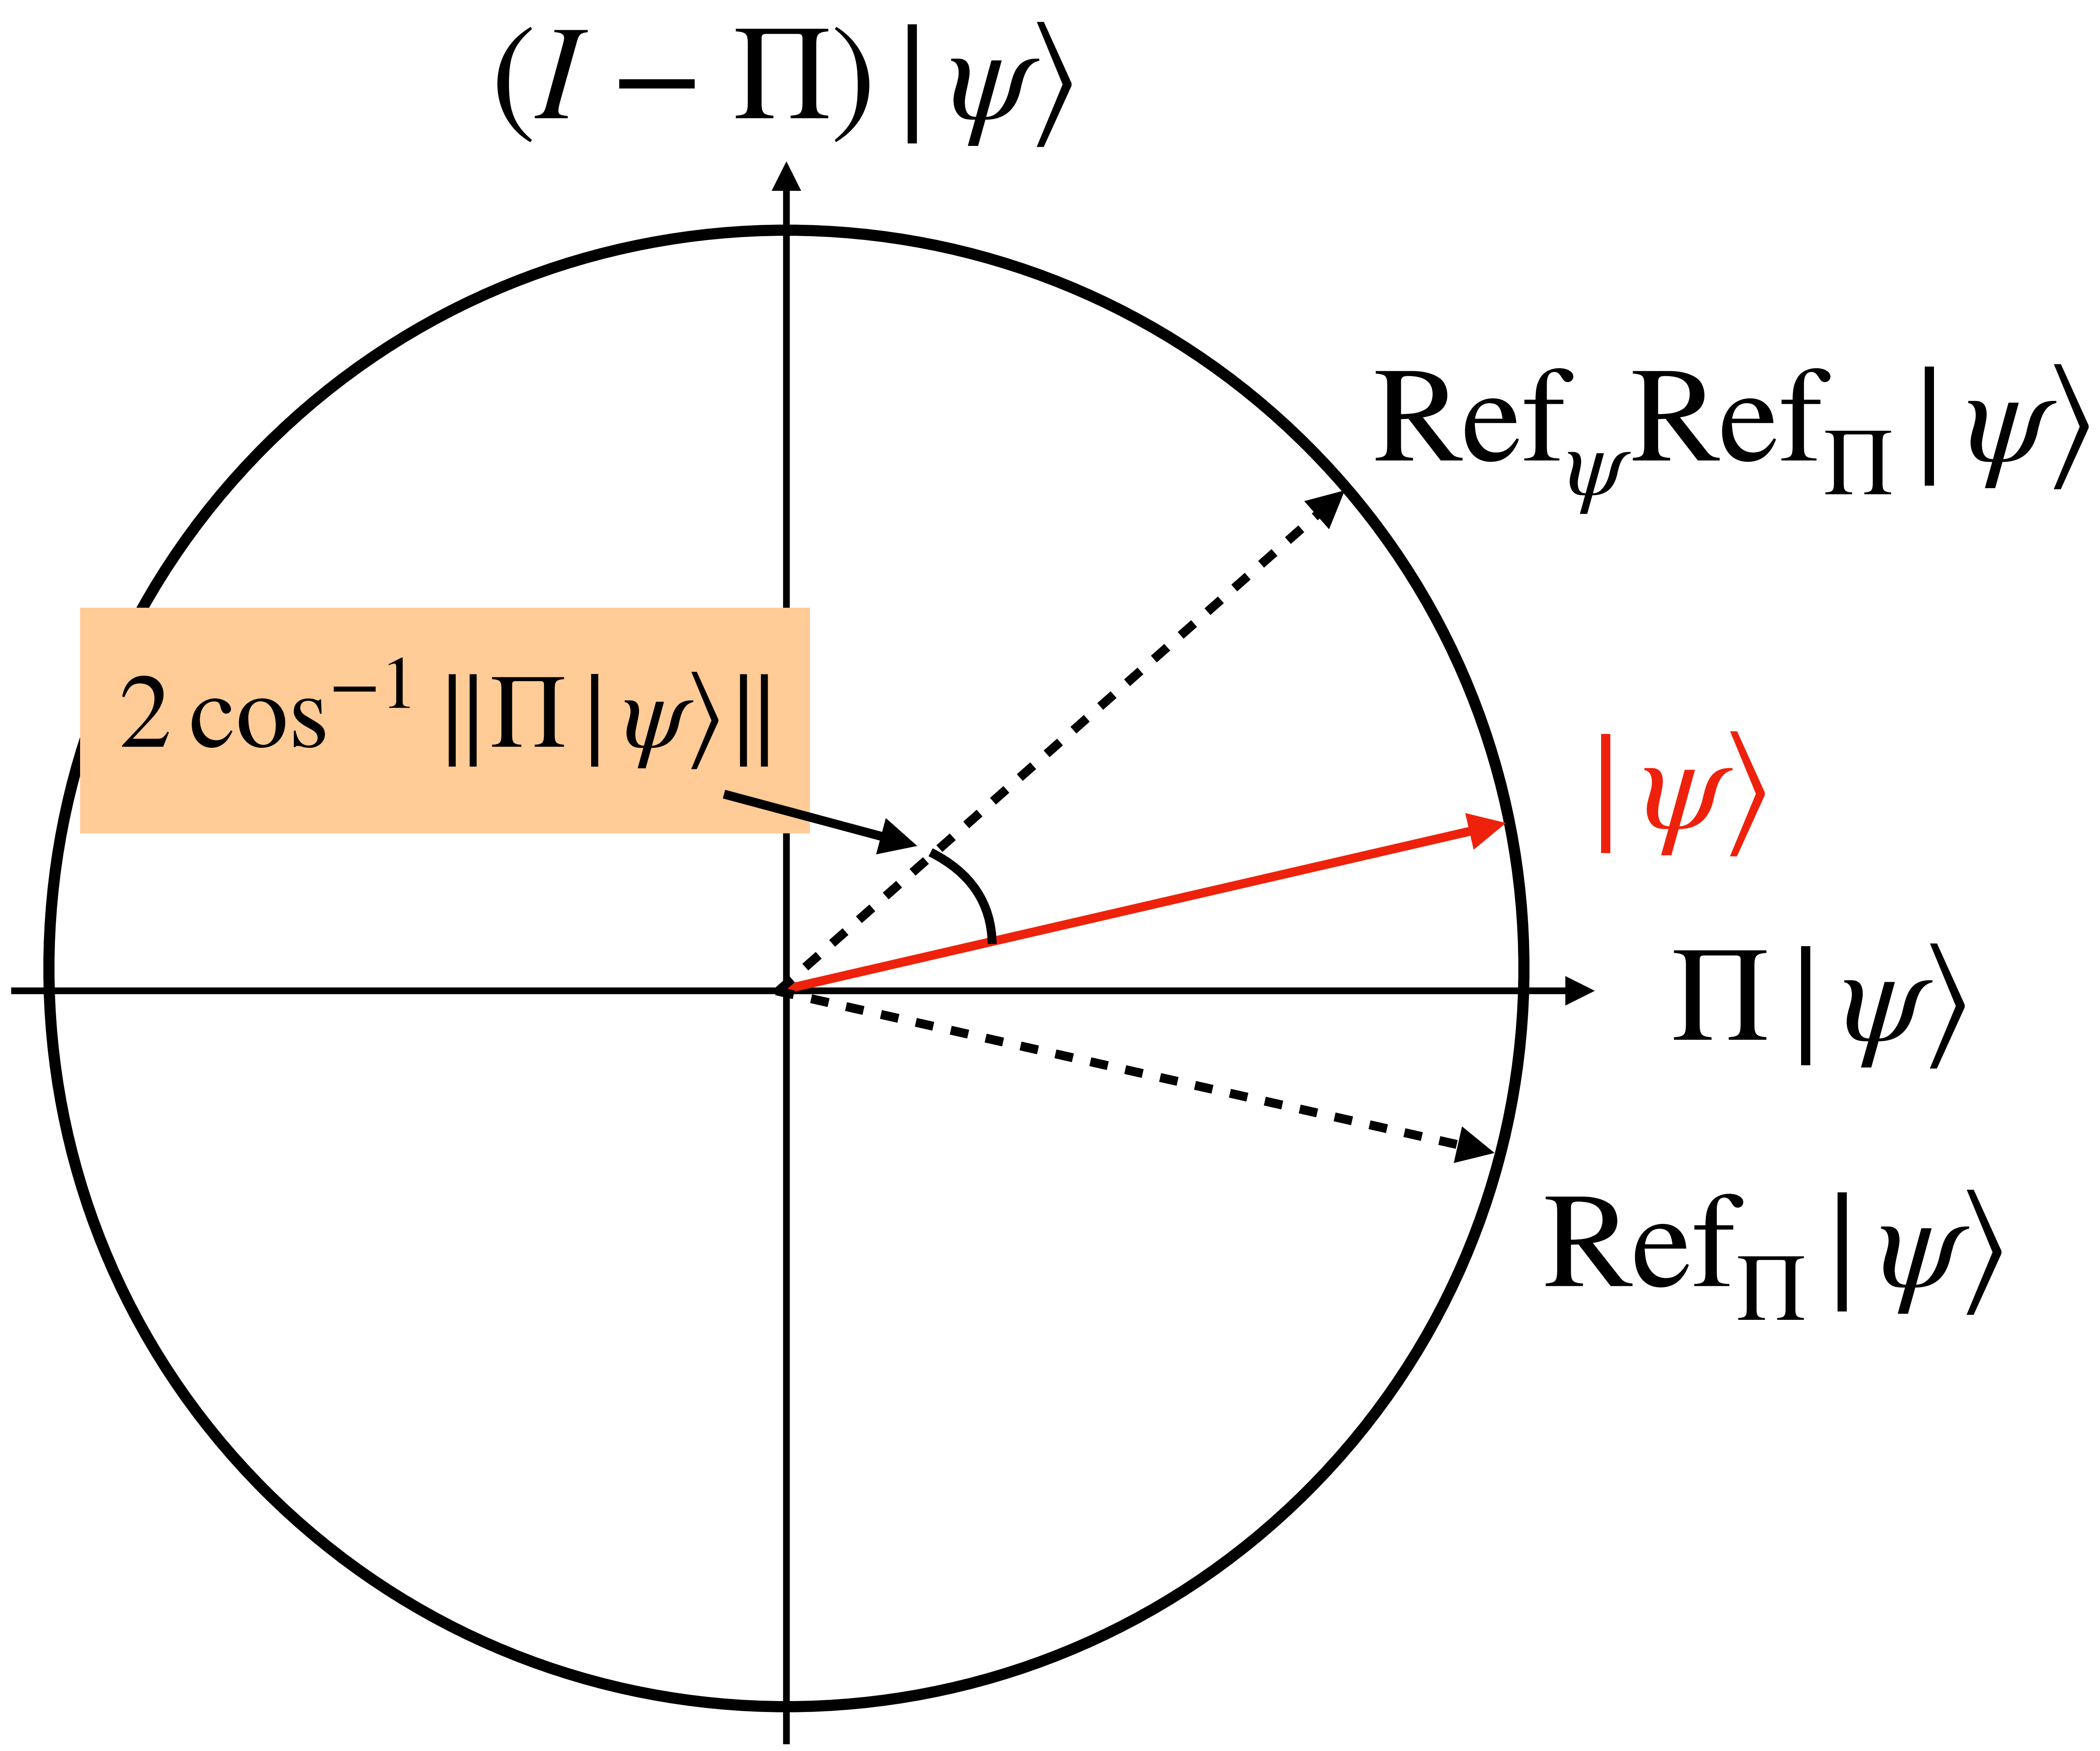

Encode $\|\Pi|\psi\rangle\|$ into the rotation angle of $R = \mathrm{Ref}_\psi \mathrm{Ref}_\Pi$

→ Phase estimation on R

The case of general observables

$$O = \sum_i \lambda_i \Pi_i$$

Can estimate $\langle \psi | O | \psi \rangle$ in time

$$\sim \|O\|/\epsilon \times \$$$

Can we do better?

Classical Mean estimation

If second moment exists, the error scales with the variance:

$$\sigma^2 = \langle \psi | O^2 | \psi \rangle - \langle \psi | O | \psi \rangle^2$$

Repeated measurements

$\sim (\sigma/\epsilon)^2 \times$ $

(Chebyshev inequality)

⟷ Incomparable

Quantum Phase estimation

$\sim \|O\|/\epsilon \times$ $

Penalized by outliers
in the spectrum of O

Outliers

Spectrum(O) :

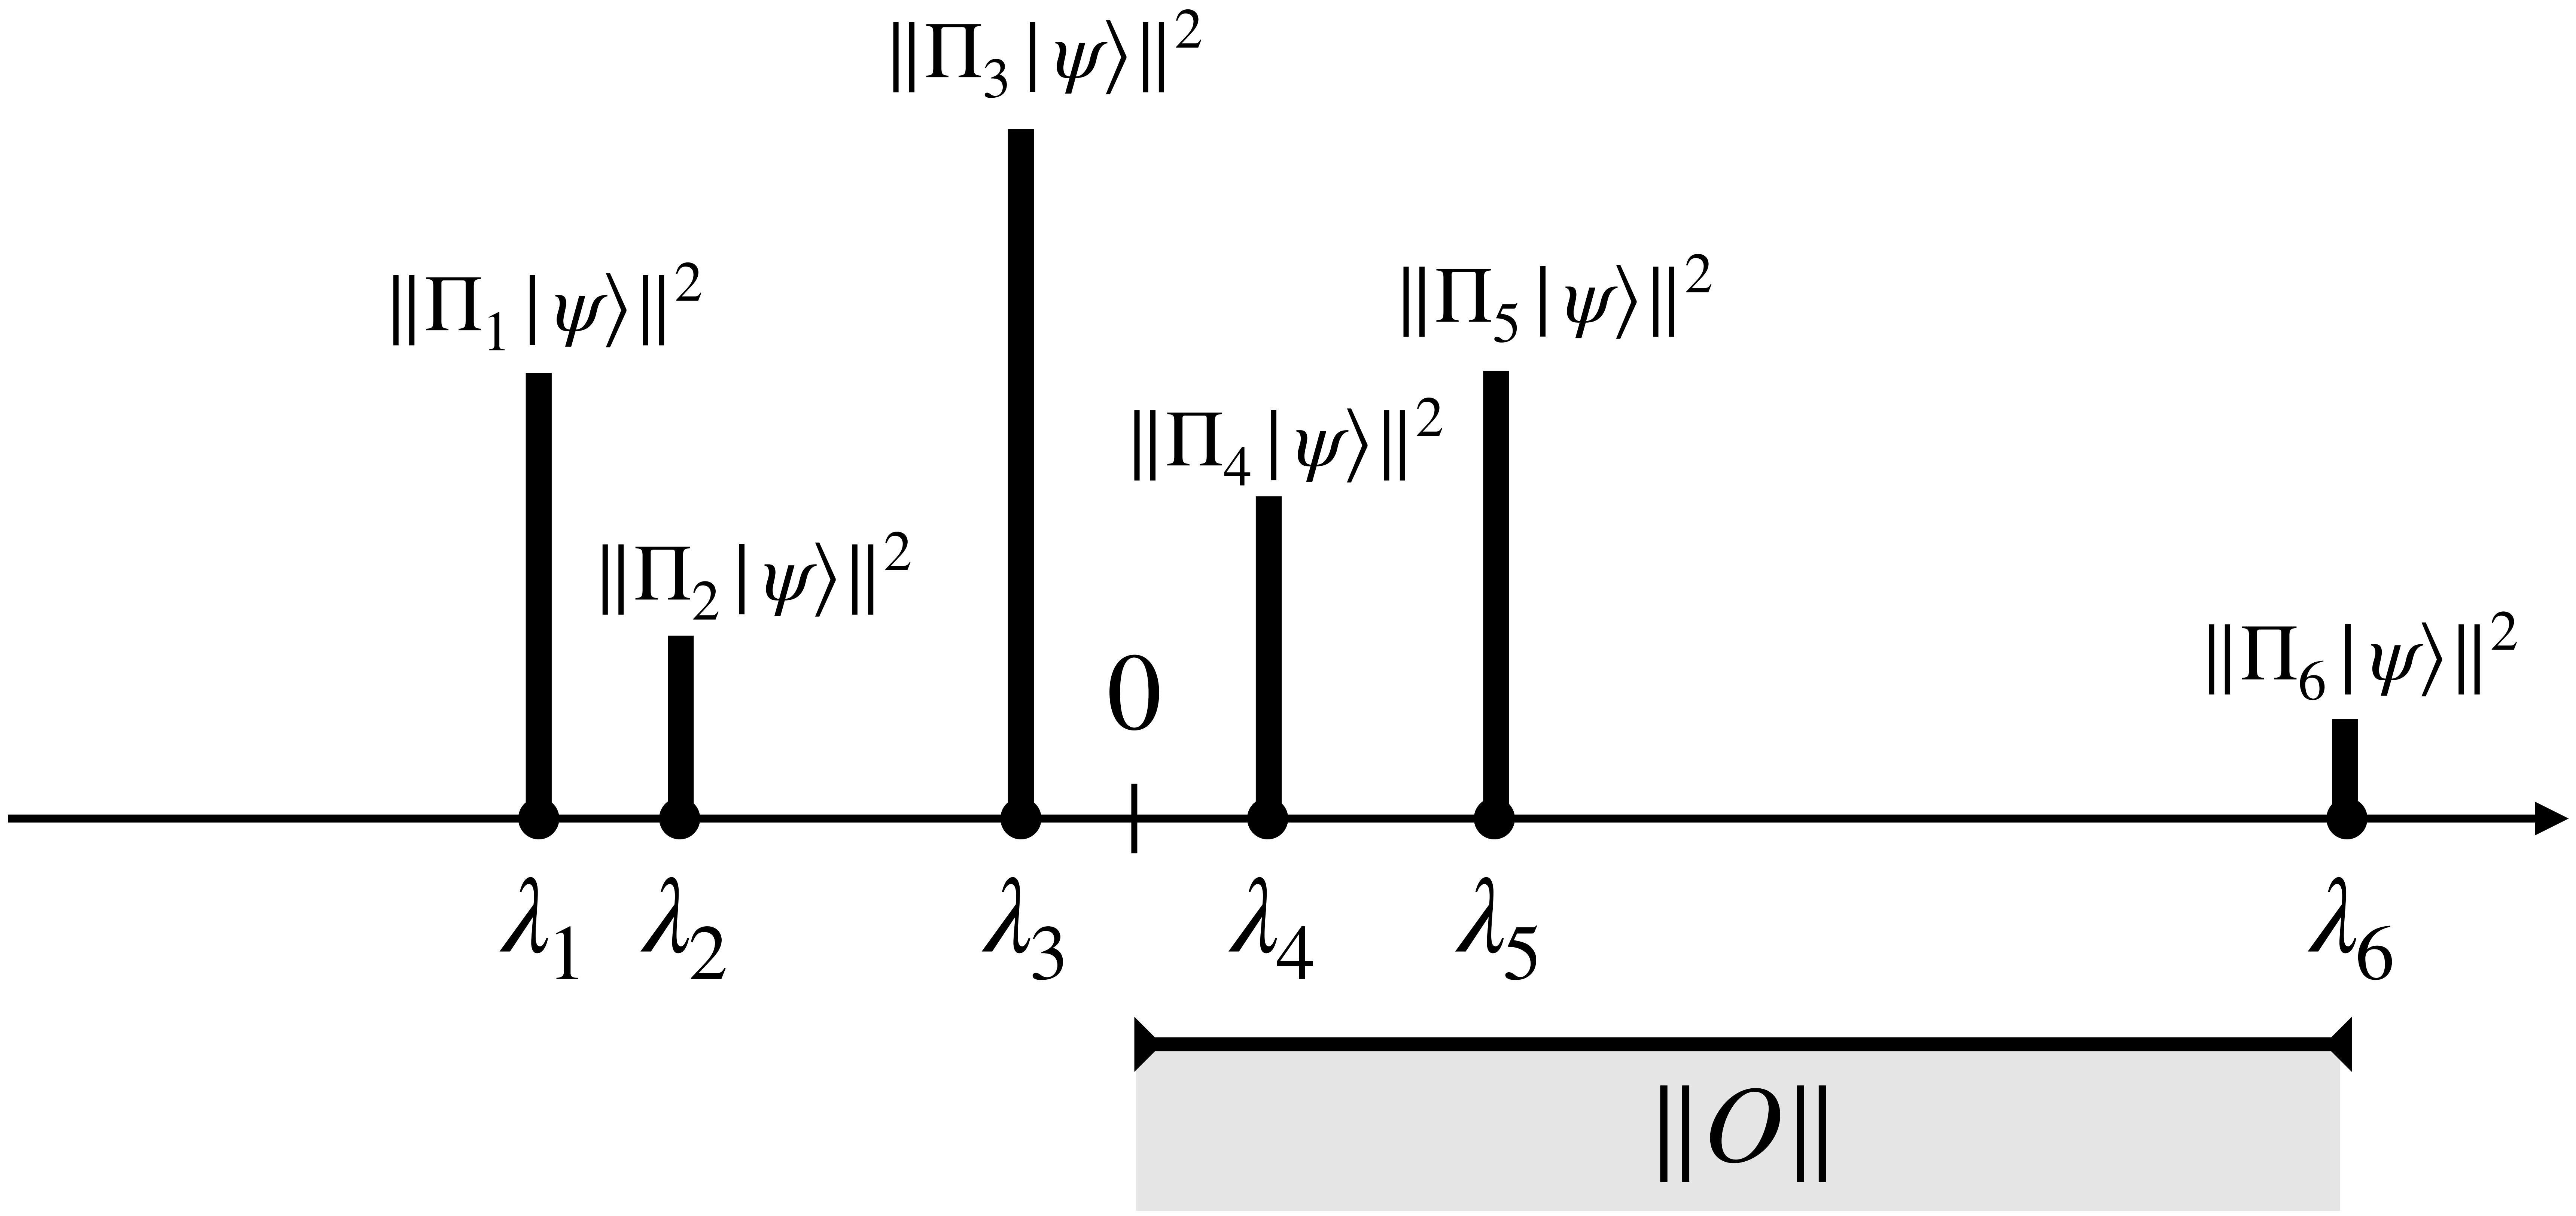

$$\langle\psi|O|\psi\rangle = \sum_i \lambda_i \cdot \|\Pi_i|\psi\rangle\|^2$$

Truncated expectation

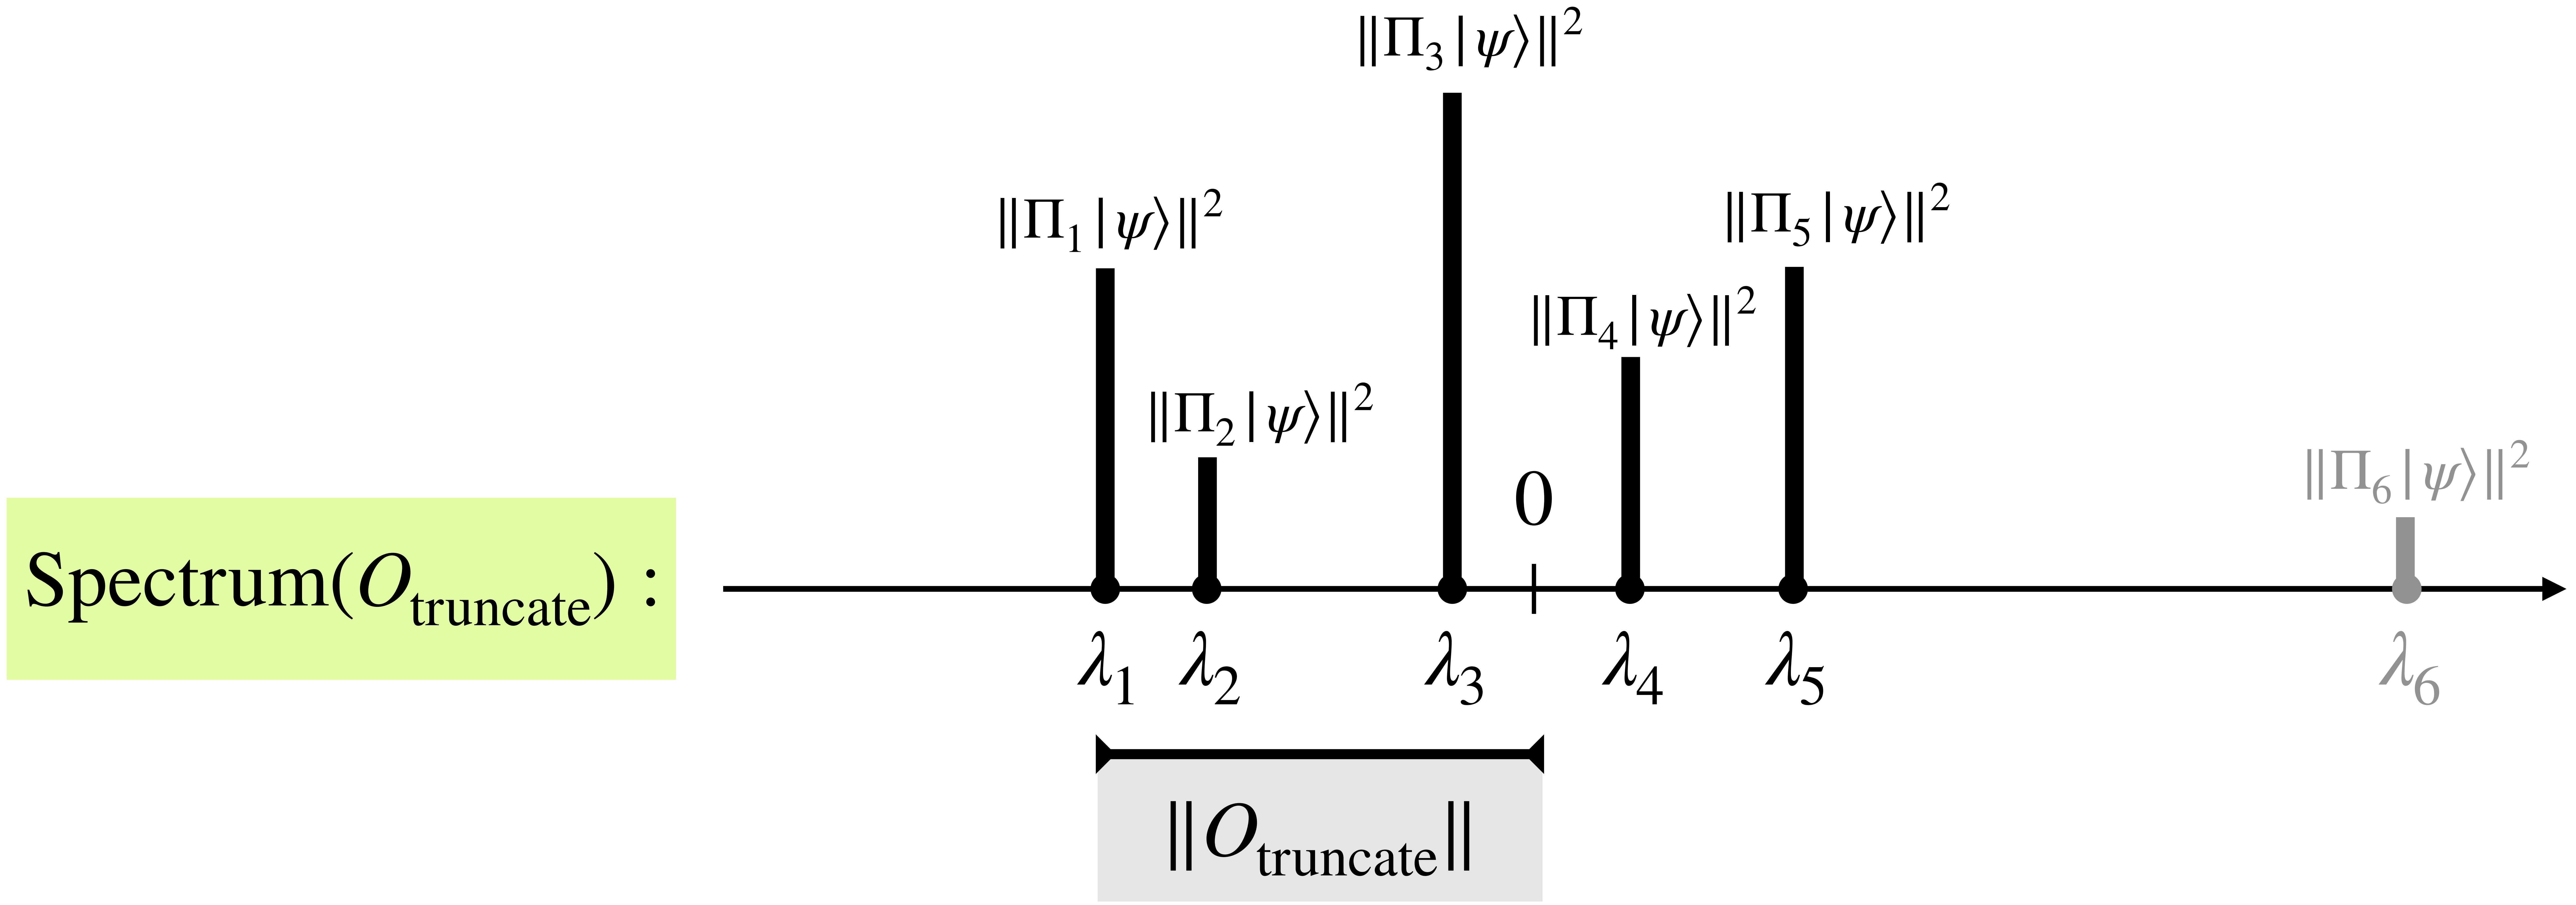

$$\langle\psi|O|\psi\rangle = \langle\psi|O_{\text{truncate}}|\psi\rangle \pm \epsilon/2$$

Optimal quantum estimator

Step 1: identify outliers by **quantile** estimation (amplitude amplification)

Step 2: estimate $\langle \psi | O_{\text{truncate}} | \psi \rangle$ (phase estimation)

Balance the cost of each step to

$$\sim \sigma/\epsilon \times \$$$

Full quadratic speedup over classical concentration inequalities

Multivariate estimator

Estimator for d observables $\langle \psi | O_1 | \psi \rangle, \ldots, \langle \psi | O_d | \psi \rangle$

Classical estimators

At most $\log(d)$ overhead

(reuse same samples for all estimates)

Quantum estimators

d overhead?

No “parallel” phase estimation?

Multivariate estimator

Estimator for d observables $\langle \psi | O_1 | \psi \rangle, \ldots, \langle \psi | O_d | \psi \rangle$

Average along a direction $u \in \mathbb{R}^d$

Commuting observables

$$u \mapsto u_1 \cdot \langle \psi | O_1 | \psi \rangle + \ldots + u_d \cdot \langle \psi | O_d | \psi \rangle$$

Non-commuting observables

$$u \mapsto \langle \psi | e^{-iu_1 O_1} \ldots e^{-iu_d O_d} | \psi \rangle$$

Doesn't scale with variance

Quantum gradient estimation

(variant of Phase estimation)

Multivariate estimator

Estimator for d observables $\langle \psi | O_1 | \psi \rangle, \ldots, \langle \psi | O_d | \psi \rangle$

Classical estimators

At most $\log(1/d)$ overhead

(reuse same samples for all estimates)

Quantum estimators

$\sqrt{d}$ overhead

Limited speedup in high-dimension

Partition functions

$$H : \Omega \to \{0,1,\ldots,n\}$$

Partition function: $Z(\beta) = \mathrm{Tr}\left(e^{-\beta H}\right)$

Gibbs sample: $|\pi_\beta\rangle \propto \sum_\sigma \sqrt{e^{-\beta H(\sigma)}}\,|\sigma\rangle$

Naive estimator

$$Z(\beta) = |\Omega| \cdot \underline{\langle \pi_0 | e^{-\beta H} | \pi_0 \rangle}$$

Slowly-evolving estimators

$$Z(\beta_{k+1}) = Z(\beta_k) \cdot \underline{\langle \pi_{\beta_k} | e^{-(\beta_{k+1}-\beta_k)H} | \pi_{\beta_k} \rangle}$$

Exponentially smaller variance

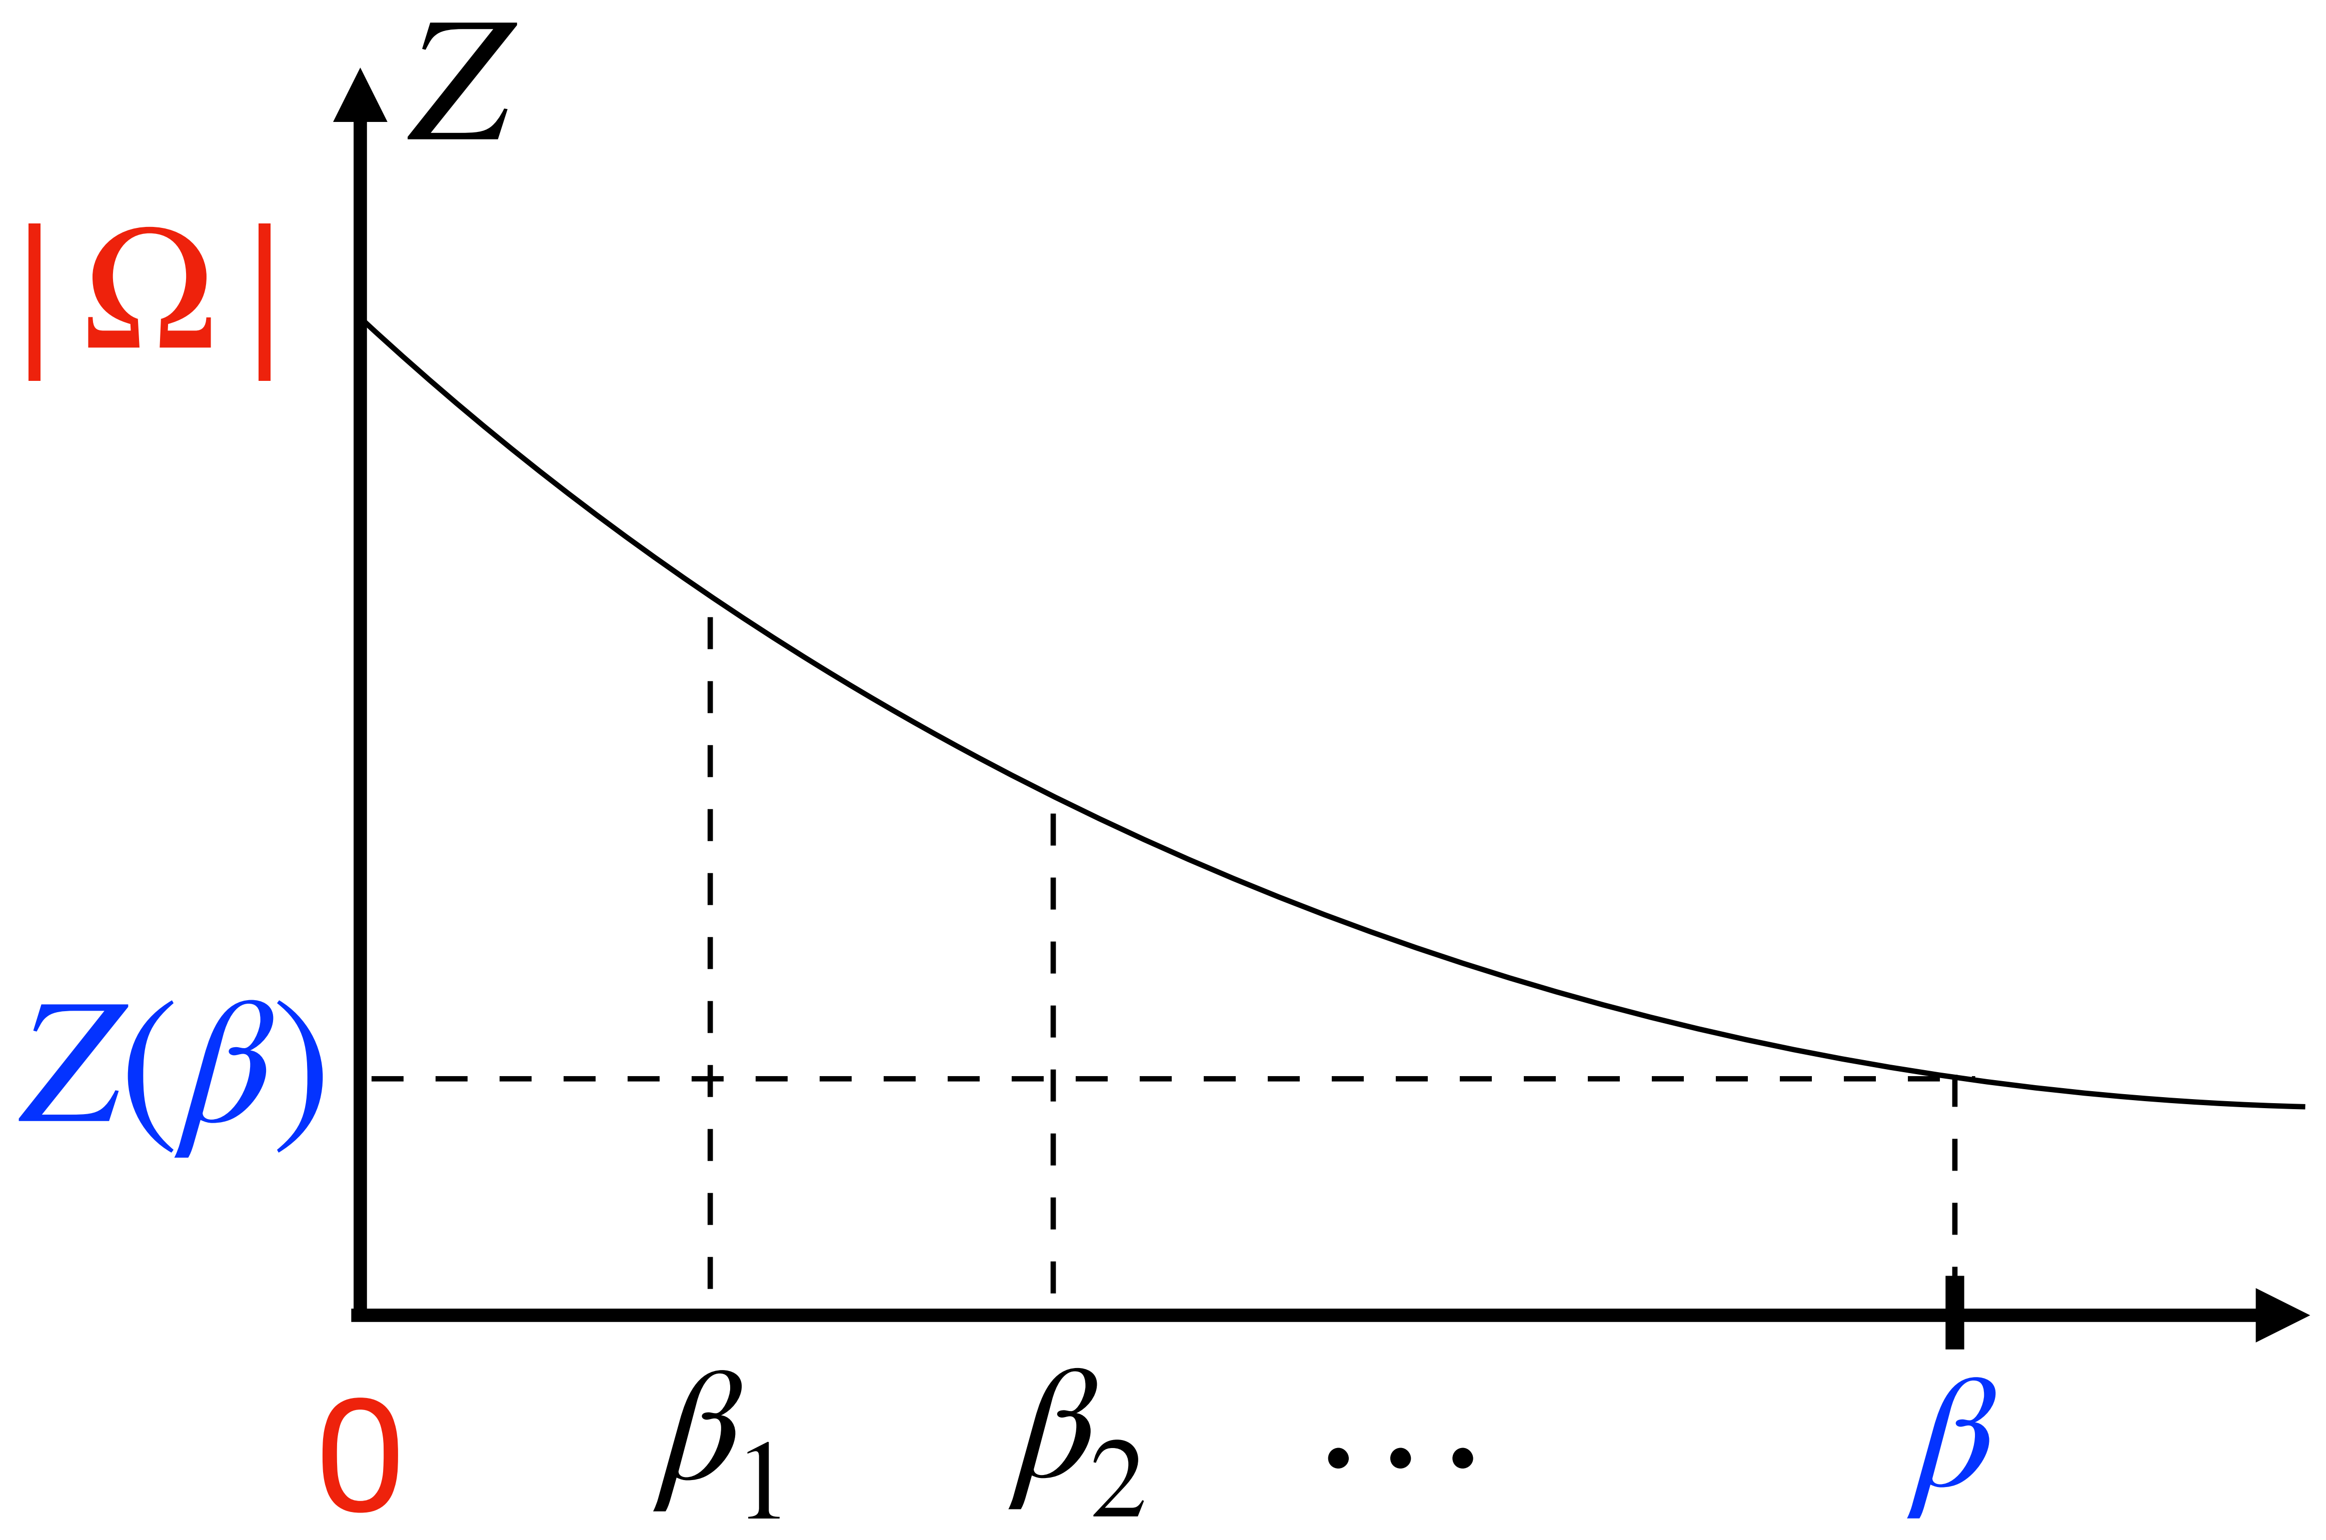

Example: Potts model

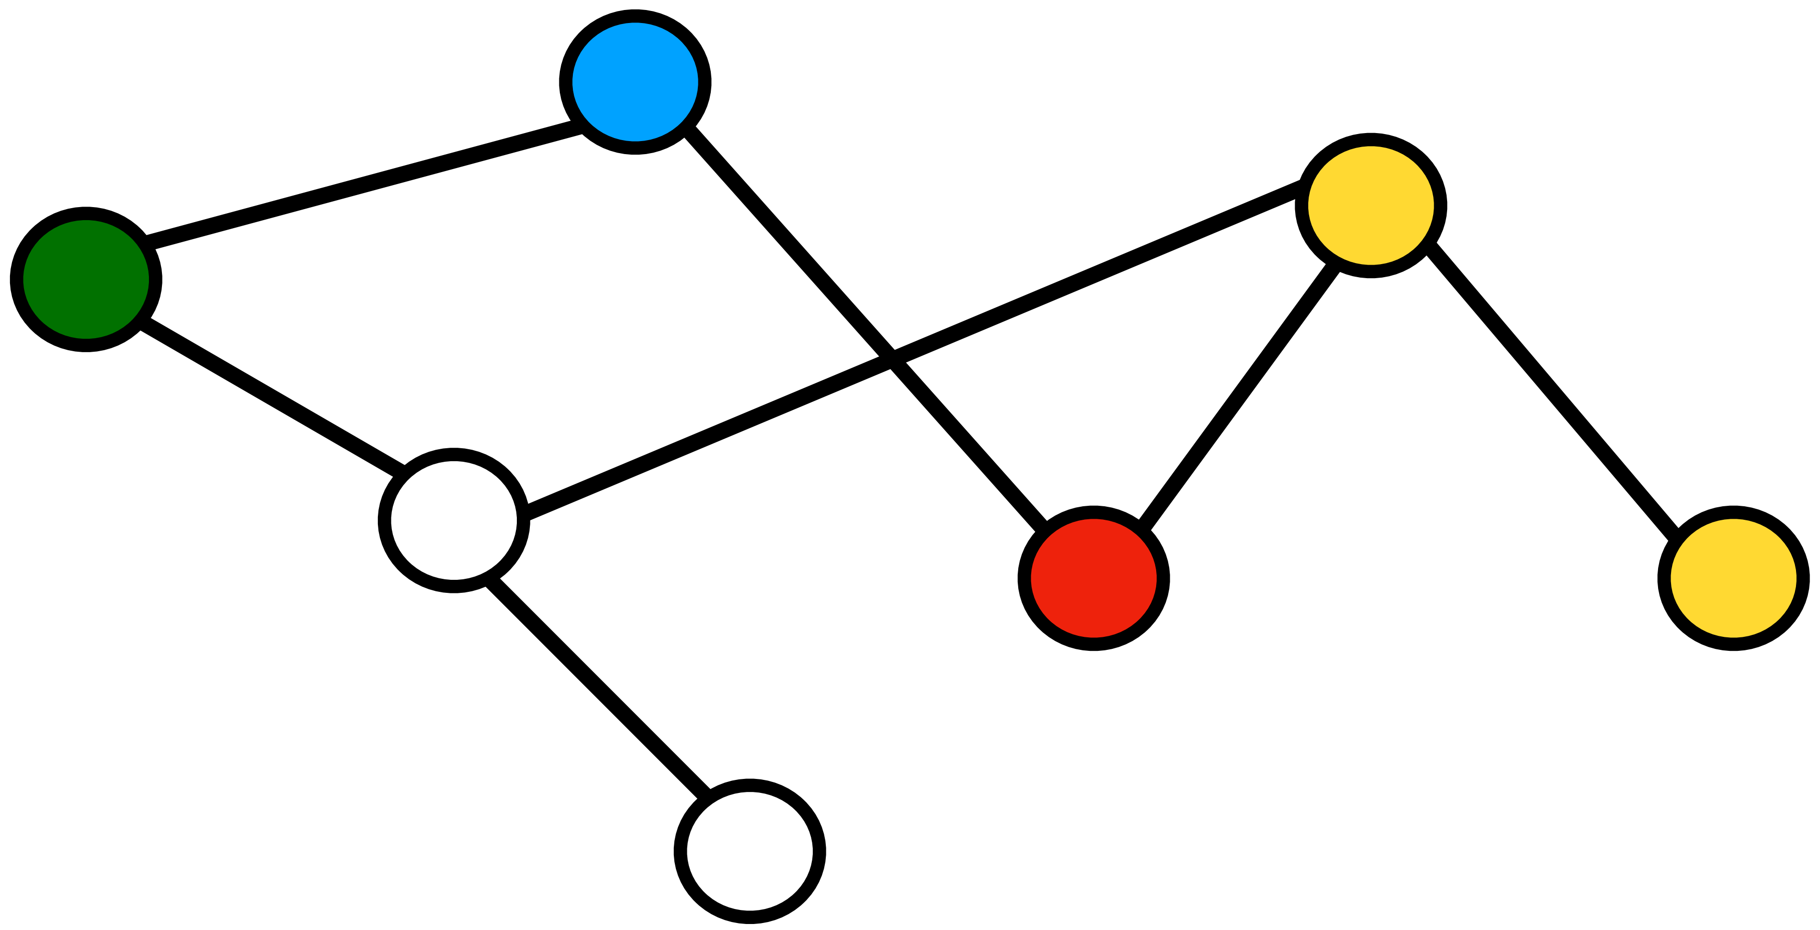

$H(\text{coloring}) = \#\text{monochromatic edges}$

Classical estimators

$$\sim \frac{(\#\text{vertices})^{2}}{\epsilon^{2}}$$

Quantum estimators

$$\sim \frac{(\#\text{vertices})^{5/4}}{\epsilon}$$

Szegedy quantum walk
\+ Quantum simulated annealing
\+ Unbiased quantum estimators

Future directions

- Quantum estimators with new features (robustness, differential privacy, …)
- Optimal variance-scaling (non-commuting observables, shadow tomography, …)
- Full quadratic speedup for estimating (classical) partition functions

Further readings

Classical estimators

- Lugosi. https://slideslive.com/38969196/do-we-know-how-to-estimate-the-mean. *NeurIPS*, 2021.
- Lugosi, Mendelson. "Mean Estimation and Regression Under Heavy-Tailed Distributions: A Survey". *FoCM*, 2019.

Univariate quantum estimators

- H. "Quantum Sub-Gaussian Mean Estimator". *ESA*, 2021.
- Knill, Ortiz, Somma. "Optimal quantum measurements of expectation values of observables". *PRA*, 2007.
- Kothari, O'Donnell. "Mean Estimation when You Have the Source Code; Or, Quantum Monte Carlo Methods". *SODA*, 2023.
- Rall. "Quantum algorithms for estimating physical quantities using block encodings". *PRA*, 2021.

Multivariate quantum estimators

- Cornelissen, H., Jerbi. "Near-Optimal Quantum Algorithms for Multivariate Mean Estimation". *STOC*, 2022.
- Huggins et al. "Nearly Optimal Quantum Algorithm for Estimating Multiple Expectation Values". *PRL,* 2022.

Partition function estimators

- Cornelissen, H. "A Sublinear-Time Quantum Algorithm for Approximating Partition Functions". *SODA,* 2023.
- Harrow, Wei. "Adaptive Quantum Simulated Annealing for Bayesian Inference and Estimating Partition Functions". *SODA*, 2020.
- Montanaro. "Quantum Speedup of Monte Carlo Methods". *Proc. R. Soc. A,* 2015.