



This is the accepted manuscript made available via CHORUS. The article has been published as:

Approximate Boltzmann distributions in quantum approximate optimization

Phillip C. Lotshaw, George Siopsis, James Ostrowski, Rebekah Herrman, Rizwanul Alam,
Sarah Powers, and Travis S. Humble

Phys. Rev. A **108**, 042411 — Published 12 October 2023

DOI: [10.1103/PhysRevA.108.042411](https://doi.org/10.1103/PhysRevA.108.042411)

Approximate Boltzmann Distributions in Quantum Approximate Optimization

Phillip C. Lotshaw,^{1,*} George Siopsis,² James Ostrowski,³ Rebekah Herrman,³ Rizwanul Alam,² Sarah Powers,⁴ and Travis S. Humble⁵

¹Quantum Computational Science Group, Oak Ridge National Laboratory, Oak Ridge, TN 37381, USA[†]

²Department of Physics and Astronomy, University of Tennessee, Knoxville, Tennessee 37996-1200 USA

³Department of Industrial and Systems Engineering,

University of Tennessee, Knoxville, Tennessee 37996-2315 USA

⁴Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37381, USA

⁵Quantum Science Center, Oak Ridge National Laboratory, Oak Ridge, TN 37381, USA

(Dated: August 31, 2023)

Approaches to compute or estimate the output probability distributions from the quantum approximate optimization algorithm (QAOA) are needed to assess the likelihood it will obtain a quantum computational advantage. We analyze output from QAOA circuits solving 7,200 random MaxCut instances, with $n = 14 - 23$ qubits and depth parameter $p \leq 12$, and find that the average basis state probabilities follow approximate Boltzmann distributions: The average probabilities scale exponentially with their energy (cut value), with a peak at the optimal solution. We describe the rate of exponential scaling or “effective temperature” in terms of a series with a leading order term $T \sim C_{\min}/n\sqrt{p}$, with C_{\min} the optimal solution energy. Using this scaling we generate approximate output distributions with up to 38 qubits and find these give accurate accounts of important performance metrics in cases we can simulate exactly.

I. INTRODUCTION

Recent studies have been motivated by the prospect of obtaining a quantum computational advantage in approximately solving NP-hard combinatorial optimization problems [1] using the quantum approximate optimization algorithm (QAOA) [2]. A quantum advantage with QAOA could have important scientific impacts in diverse fields including computer science, engineering, and operations research [3–5]. Theoretical [6–9] and numerical [10–14] approaches have provided performance benchmarks and bounds in certain cases, typically in terms of the approximation ratio metric, which quantifies average solution quality relative to the optimal. However, the average solution quality does not directly address the success probability of QAOA in generating a single near-optimal bitstring solution to a combinatorial problem. To address this, alternative approaches have begun considering full QAOA output distributions. A derivation showing that $p = 1$ QAOA pure states have Boltzmann-distributed solution probabilities has been presented in Ref. [15] based on correlations in the solution spectrums of random Ising models. Similar distributions have also been observed to describe ensemble averages of instances

at $p > 1$ in a “mean-field” approach for a state based on ensemble-averaged SAT instances [17, 18]; see also related work in Ref. [19], and on instantaneous quantum polynomial circuits in Ref. [16]. However, obtaining simplified descriptions of output from individual instances at $p > 1$, and benchmarking their accuracy, have remained as important steps towards understanding performance at large depths that are relevant for quantum advantage.

We analyze output probability distributions in optimized instances of QAOA solving the MaxCut problem, for ensembles of random Erdős-Rényi graphs with sizes $n \in \{14, 17, 20, 23\}$, and with QAOA depth parameters $p \leq 12$. We find the average probability per basis state scales exponentially with costs in the MaxCut objective, similar to Boltzmann distributions in statistical mechanics, as anticipated from previous works [15–18]. We observe systematic depth- and instance-dependent behaviors in the rate of exponential scaling or “effective temperature”, leading ultimately to a heuristic approach for generating approximate QAOA output. The estimated performance is found to accurately describe the true QAOA performance in median cases we can test exactly, while we also generate predictions for larger sizes $n \leq 38$ where we do not yet have exact results for comparison. In total, simple Boltzmann distributions are found to give satisfactory and unified accounts of observed QAOA output distributions across diverse cases we consider.

We further analyze the probability to obtain an optimal solution with QAOA. A previous study by Ashkay *et. al* has shown in simulations that a scaling $p \sim n$ suffices to obtain a fixed optimal solution probability in Max-2-SAT [12], while for MaxCut the simulated probability to measure an optimal solution has been shown to decrease exponentially with n , as expected from the growth of the 2^n -dimensional Hilbert space [10]. Here we analyze p and n dependence in the optimal solution

* lotshawpc@ornl.gov

[†] This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC0500OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan.

probability for MaxCut, which is found to differ from the Max-2-SAT case. Our analysis considers typical behavior within a specific ensemble of instances at limited sizes, and this dataset does not allow us to analyze complexity-theoretic worst-case scaling over all instances, which could vary from the typical behavior observed here.

A final component of the present work is an extension of a previous heuristic to rapidly identifying optimized variational parameters that are needed for high-performance QAOA circuits. This overcomes bottlenecks associated with variational parameter optimization, which have been suggested to limit quantum advantage in time-to-solution [20]. Building on Ref. [21], we show that a single set of parameters, taken from a previous study of the Sherrington-Kirkpatrick (SK) model at infinite size [6], can be rescaled and applied to the random graph ensembles we consider, with depth parameters $p \leq 12$. This extends the range of applicability of previous angle transfer heuristics, which have focused on transfer between generic or random instances at $p \leq 5$ [10, 21, 22] as well as large n limits or specific structured instances [23], such as 3-regular graphs [24–26] or the SK model [6], in some cases with larger p .

II. QUANTUM APPROXIMATE OPTIMIZATION ALGORITHM

An instance of a combinatorial optimization problem is defined by a cost function $C(\mathbf{z})$ with an argument $\mathbf{z} = (z_1, \dots, z_n)$ in terms of binary variables $z_i \in \{-1, 1\}$. Many scientifically relevant problems can be expressed in terms of quadratic unconstrained binary optimization for which [27]

$$C(\mathbf{z}) = \sum_i h_i z_i + \sum_{i,j} J_{i,j} z_i z_j \quad (1)$$

An optimal solution \mathbf{z}_{\min} globally minimizes the cost function

$$\mathbf{z}_{\min} \in \arg \min_{\mathbf{z}} C(\mathbf{z}) \quad (2)$$

However, finding optimal solutions is often intractable due to resource requirements that scale exponentially with problem size. To overcome this problem, approximation algorithms and heuristics seek approximate solutions with reduced compute time.

QAOA is a quantum heuristic designed to approximately solve combinatorial optimization problems. The cost function $C(\mathbf{z})$ is encoded into a quantum Ising Hamiltonian [27]

$$\hat{C} = \sum_{i,j} J_{i,j} \hat{Z}_i \hat{Z}_j + \sum_i h_i \hat{Z}_i, \quad (3)$$

where \hat{Z}_i are Pauli-Z operators, such that the eigenspectrum of \hat{C} matches the set of cost function values

$$\hat{C}|\mathbf{z}\rangle = C(\mathbf{z})|\mathbf{z}\rangle. \quad (4)$$

Throughout this work we use hat symbols to distinguish operators, such as \hat{C} , while reserving plain symbols for functions and numbers, such as the cost function $C(\mathbf{z})$ or a specific cost value C . A ground state of the operator \hat{C} is a computational basis state $|\mathbf{z}_{\min}\rangle = |(z_1, z_2, \dots, z_n)_{\min}\rangle$ that represents an optimal solution to the classical problem, while low-lying excited states represent approximate solutions.

QAOA uses a variational circuit ansatz to prepare a quantum state that is posited to return approximate ground states of \hat{C} upon measurement in the computational basis. The QAOA ansatz uses p layers that each alternate between Hamiltonian evolution under \hat{C} and under a “mixing” operator $\hat{B} = \sum_i \hat{X}_i$ with \hat{X}_i the Pauli-X operator on qubit i ,

$$|\gamma, \beta\rangle = \prod_{l=1}^p e^{-i\beta_l \hat{B}} e^{-i\gamma_l \hat{C}} |+\rangle, \quad (5)$$

where $|+\rangle = |+\rangle^{\otimes n}$ is the ground state of $-\hat{B}$. The $\beta = (\beta_1, \dots, \beta_p)$ and $\gamma = (\gamma_1, \dots, \gamma_p)$ are variational parameters chosen to minimize $\langle \hat{C} \rangle = \langle \gamma, \beta | \hat{C} | \gamma, \beta \rangle$; note the state notation $|\gamma, \beta\rangle$ depends implicitly on the number of algorithmic layers p , through γ and β . When suitable parameters have been identified, then repeated preparation and measurement of $|\gamma, \beta\rangle$ yields a set of candidate solutions $\{\mathbf{z}_{\text{cand}}\}$ and the bitstring $\mathbf{z}^* \in \{\mathbf{z}_{\text{cand}}\}$ producing the smallest cost $C(\mathbf{z})$ is taken as the final solution. The QAOA ansatz is used because 1) in the limit $p \rightarrow \infty$ it yields the ground state and 2) the solution quality can only improve as p increases [2].

Solution quality is often quantified through the approximation ratio

$$r = \frac{C_{\max} - \langle \hat{C} \rangle}{C_{\max} - C_{\min}} \quad (6)$$

where C_{\min} and C_{\max} are the extremal values of the cost function; we computed these by evaluating $C(\mathbf{z})$ for each \mathbf{z} . Here $0 \leq r \leq 1$ quantifies the expected cost; $r = 1$ signifies that QAOA prepares an optimal solution (when $\langle \hat{C} \rangle = C_{\min}$, in keeping with Eq. (2)), while $r = 0$ is the worst case. We will further consider the probabilities for individual solutions

$$\text{Pr}(\mathbf{z}) = |\langle \gamma, \beta | \mathbf{z} \rangle|^2. \quad (7)$$

We focus on the unweighted MaxCut problem, which is a standard benchmarking problem for QAOA. An instance of MaxCut is defined with respect to a graph $G = (V, E)$ with vertex labels V and edges $E = \{(i, j); i, j \in V\}$. The goal is to partition the vertices into two sets such that the number of edges with endpoints in different sets is maximized. For a set of $n = |V|$ binary variables $z_i \in \{-1, +1\}$, the MaxCut cost function is

$$C(\mathbf{z}) = \sum_{(i,j) \in E} z_i z_j \quad (8)$$

where the sum runs over the set of edges E in the problem graph. The cost Hamiltonian is

$$\hat{C} = \sum_{(i,j) \in E} \hat{Z}_i \hat{Z}_j. \quad (9)$$

We consider instance sets containing 300 random Erdős-Rényi graphs at each size $n \in \{14, 17, 20, 23\}$ [28], where in each instance each possible edge (i, j) is generated with probability 0.5. We chose Erdős-Rényi graphs rather than graphs with specified symmetries to avoid the possibility of non-generic symmetry-related behaviors.

III. RESULTS

A. Angle transfer

To apply QAOA to a given problem, it is necessary to choose angle parameters β, γ that define a QAOA circuit. While exhaustive brute force searches yield the best possible angles, these are impractical due to the search overhead [13, 20]. To remedy this problem, heuristics based on transferring angles between different problem instances have been developed and shown to achieve high performance with minimal compute time. These have included theoretical [23, 29] and empirical studies of transfer between graphs [10, 21, 22, 24, 26] as well as approaches for generating angles at layer $p + 1$ from angles at layer p [25, 30].

Our approach begins with angle parameters previously devised for the Sherrington-Kirkpatrick (SK) model at infinite size [6]. Similar angles were also identified in an analysis of large-girth graphs [31], with median deviations of 0.07% from the SK angles and a worst-case deviation of 9%. Thus, the SK angles approximate optimized angles for diverse instances, at least in the large size and degree limits of Refs. [6, 31]. Boulebnane and Montanaro [21] investigated an approach to rescaling these angles for applications to small Erdős-Rényi and Chung-Lu graphs at $p \leq 5$, and we extend their approach to $p \leq 12$. We use relationships discussed in Refs. [6, 21, 22] to rescale the SK angles using the average graph degree $\bar{d} = 2|E|/n$ of a particular instance with n qubits and $|E|$ edges in the problem graph. The specific relations we use are

$$\beta = \beta^{\text{SK}}, \quad \gamma = \frac{1}{\sqrt{\bar{d}}} \gamma^{\text{SK}} \quad (10)$$

where β^{SK} and γ^{SK} are taken from the analysis of the infinite-size SK model in Ref. [32]. We then perform a gradient-based optimization from these angles, using the NLOPT implementation of the BFGS optimization algorithm [33–35], to determine optimized angles β^* and γ^* for each individual instance.

Approximation ratios obtained from this procedure are shown in Fig. 1. Points show the median approximation ratio at each p for our sets of 300 graphs at each n , while error bars show the 0.1-0.9 quantiles over the




FIG. 1. Approximation ratios for Erdős-Rényi graphs at varying n and $p = 2, 4, \dots, 12$; offsets in p are included for clarity. Points show medians over the sets of 300 graphs at each n and error bars show 0.1 and 0.9 quantiles.

sets of graphs. The approximation ratios increase with p and approach unity for each graph, as expected for well-optimized angles in QAOA. This indicates qualitatively satisfactory behaviors are obtained from the rescaled and optimized SK angles.

We do not have optimal angles to compare against our transfer procedure, due to the computational expense of searching for optimal angles. Exhaustive searches in the $2p$ -dimensional parameter space have used hundreds of optimizations per instance at $p \leq 3$ [10, 22] and are expected to require significantly more optimizations for our cases at $p \leq 12$. As an alternative we perform a simpler comparison against a brute force search method with limited sampling of the parameter space, as described in Appendix A. We select 100 random initial angles per instance at $n = 14$ at each p , optimize these to find a local optimum, then take the best result for each instance to compare against the transfer procedure. We find the transferred angles obtain approximation ratios equivalent to or better than these optimized random samples. This indicates many brute force samples are needed to outperform the much simpler transfer approach, consonant with previous works [6, 10, 21, 22, 24–26, 30]. The transferred angles also give similar performance at varying sizes n , as seen in Fig. 1, further demonstrating their practical utility.

B. QAOA Probability Distributions

We consider probability distributions associated with the optimized QAOA states across our 300 graph ensembles at varying n and p , beginning here with an analysis of average distributions and optimal solution probabilities, then turning to specific instances in later sections.

To analyze QAOA output distributions, we compute probabilities $\Pr(\mathbf{z})$ for each solution \mathbf{z} and bin these to obtain the total probability to measure any solution with

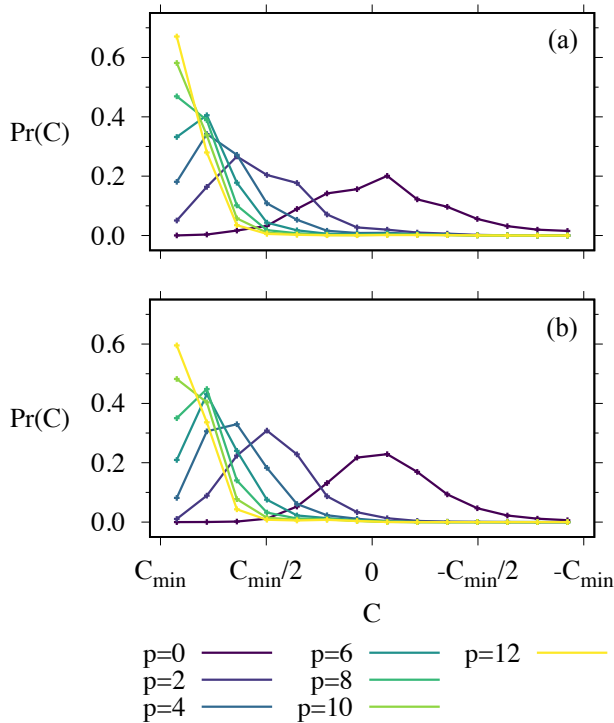


FIG. 2. Average probability to measure any state with a given cost C , binned in widths $|C_{\min}|/7$, averaged over the 300 graph instances at (a) $n = 14$ and (b) $n = 23$.

cost C ,

$$\Pr(C) = \sum_{z; C(z)=C} \Pr(z). \quad (11)$$

Here C refers to an individual cost value, not the operator \hat{C} . In combinatorial optimization the final result is often assessed only by its cost value and probabilities for different costs are quantified by $\Pr(C)$.

To obtain a coarse view of how the $\Pr(C)$ vary with C and approach the optimal C_{\min} as the number of layers p increases, we binned probabilities at varying costs in bins of width $|C_{\min}|/7$ for each instance. We plot the average binned probabilities in Fig. 2(a) for $n = 14$ and (b) for $n = 23$; we do not include standard deviation error bars as they obscure the appearance of the figure. For simplicity we plot intervals $[C_{\min}, -C_{\min}]$ though note the full distributions extend to $C_{\max} = |E|$ with very small probabilities not included in the figure.

Curves “ $p = 0$ ” show the distributions of the initial state, corresponding to a uniform random sampling of solutions. Here the probabilities $\Pr(C)$ are proportional to the densities of solutions

$$\varrho(C) = \sum_{z; C(z)=C} 1 \quad (12)$$

with $\Pr(C) = \varrho(C)/2^n$. The average $\varrho(C)$ and $p = 0$ output distributions are Gaussian-like, and the overlap with costs near the optimal C_{\min} decrease as n increases




FIG. 3. Probability to measure an optimal solution. Points show medians and error bars show 0.1-0.9 quantiles across 300 graphs at each n with $p \leq 12$. The medians are fit by $a \exp(-bn/p^{2/3})$, with $a = 2.75 \pm 0.09$ and $b = 0.502 \pm 0.009$, with \pm denoting the asymptotic standard error of the non-linear least squares fit.

from $n = 14$ in Fig. 2(a) to $n = 23$ in Fig. 2(b). Curves for $p \geq 2$ have improved solution quality, with quasi-Gaussian average distributions centered at costs that approach C_{\min} as p increases.

The probability to obtain an optimal solution

$$\Pr(C_{\min}) = \sum_{z_{\min} \in \arg \min_z C(z)} \Pr(z_{\min}) \quad (13)$$

is shown in Fig. 3, with points showing the median probability at each n and p and error bars showing the 0.1-0.9 quantiles. We make an approximate scaling of the median $\Pr(C_{\min}) = a e^{-bn/p^{2/3}}$, shown by the black line, with fit parameters a and b in the figure caption. This indicates that to obtain a fixed probability for an optimal solution in the median case, p should increase with size as $p \sim n^{3/2}$. This scaling differs from the scaling $p \sim n$ observed in optimized instances of Max2-SAT [12], which may be due to differences in the structure of the MaxCut problem, or possibly due to the choice of QAOA parameters. The optimal solution scaling $\Pr(C_{\min}) \sim e^{-k_p n}$ has been observed across MaxCut instances at sizes $n \leq 9$ in previous work [10]; here we find $k_p = b/p^{2/3}$.

To obtain a more complete understanding of the sources of the distributions in Figs. 2-3, we consider the average probability to measure a single basis state with a given cost C ,

$$\overline{\Pr}(z_C) = \frac{\Pr(C)}{\varrho(C)}. \quad (14)$$

We again consider the histogram binning procedure of previous paragraphs and plot the average $\overline{\Pr}(z_C)$ in Fig. 4(a),(b) for $n = 14, 23$ respectively, on a logarithmic scale. For the $p = 0$ initial state, the distribution is flat, since each individual basis state has probability $1/2^n$. At larger p , the distributions increase exponentially towards

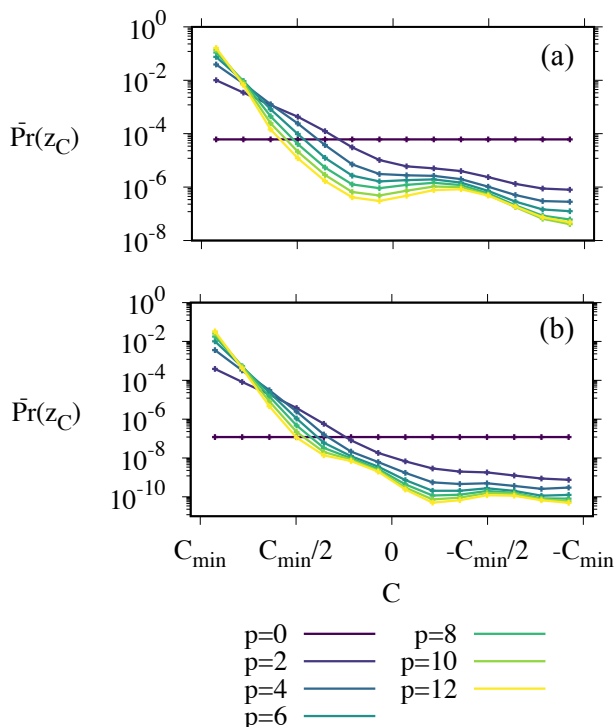


FIG. 4. Average probability to measure a single basis state with cost C in intervals of width $|C_{\min}|/7$, averaged over the 300 graph instances at (a) $n = 14$ and (b) $n = 23$.

C_{\min} , while they plateau to small values $\approx 10^{-6} - 10^{-10}$ at $C \gtrsim 0$. To a good approximation we can think of the average behavior as exponential, as solutions with $C \gtrsim 0$ contribute little overall probability.

The previous analysis gave a coarse picture of the average behavior of QAOA state distributions for random MaxCut instances, including exponential behavior of the average basis state probabilities $\overline{\text{Pr}}(z_C)$. However, Figs. 2 and 4 have considered averages over many different instances, and a more detailed account of individual instances is desirable. We address this in the next section.

C. Approximate Boltzmann distributions

The average exponential behavior of Fig. 4 raises a question of whether the exponential dependence arises from averaging many instances or whether individual QAOA instances show exponential distributions. Exponential distributions were shown to approximately characterize QAOA at $p = 1$ in Refs. [15, 16], though it is not clear if this should describe our cases with $p \geq 2$. We consider modeling individual QAOA instances with “Boltzmann distributions” describing the average basis state probabilities

$$\overline{\text{Pr}}_{\text{exp}}(z_C) = \frac{1}{Z} e^{-C/T}, \quad (15)$$

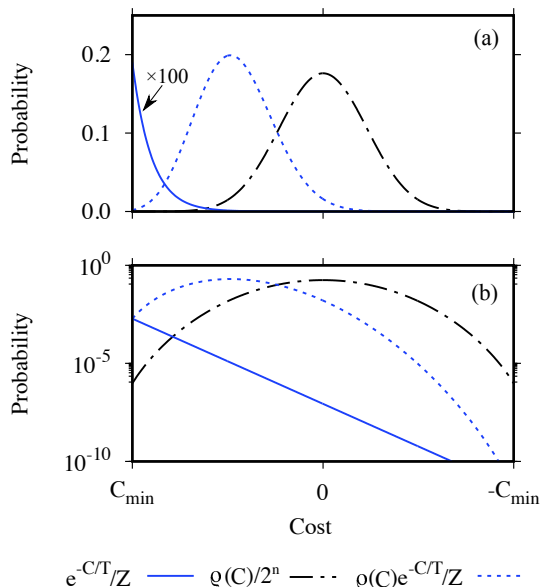


FIG. 5. Schematic exponential distributions, on (a) a linear scale and (b) a log scale, see text for details.

where the partition function $Z = \sum_C \varrho(C) e^{-C/T}$ is a normalizing constant. T controls the rate of change in basis state probabilities with varying costs. As $T \rightarrow \infty$ the distribution becomes uniform, giving the initial QAOA state distribution ($p = 0$), while the probabilities condense to the ground state as $T \rightarrow 0$ ($p \rightarrow \infty$). Eq. (15) is mathematically identical to the Boltzmann distribution in statistical mechanics, though with a different conceptual foundation, as discussed in Sec. IV. The total probability to obtain any solution at a given cost is then

$$\text{Pr}_{\text{exp}}(C) = \frac{\varrho(C)}{Z} e^{-C/T}. \quad (16)$$

Before comparing to QAOA, it may be useful to consider what to expect from the Boltzmann distribution as shown schematically in Fig. 5 on (a) a linear scale and (b) a log scale. The basis state probabilities $\overline{\text{Pr}}_{\text{exp}}(z_C)$ (solid blue) decay exponentially from C_{\min} (Eq. (15)). The number of states per cost $\varrho(C)$ (Eq. (12)) can be exponentially large depending on C ; for the schematic we take $\varrho(C)$ as a binomial distribution on 20 variables with costs ± 1 and show a scaled version $\varrho(C)/2^n$ (black). The total probability $\text{Pr}_{\text{exp}}(C)$ to measure any basis state with cost C (dotted blue) is the product of the exponentially small $e^{-C/T}/Z$ and exponentially large $\varrho(C)$ (Eq. (16)), with a peak at intermediate costs between zero and C_{\min} . This schematically captures the expected behavior in terms of the relationships between $\overline{\text{Pr}}_{\text{exp}}(z_C)$, $\text{Pr}_{\text{exp}}(C)$, and $\varrho(C)$. We now consider modeling QAOA probability distributions using these relations.

We used non-linear least squares to fit Eq. (16) to exact probabilities $\text{Pr}(C)$ for each graph instance at each p . We quantify agreement between the exponential fits $\text{Pr}_{\text{exp}}(C)$ and the true probability distributions from the

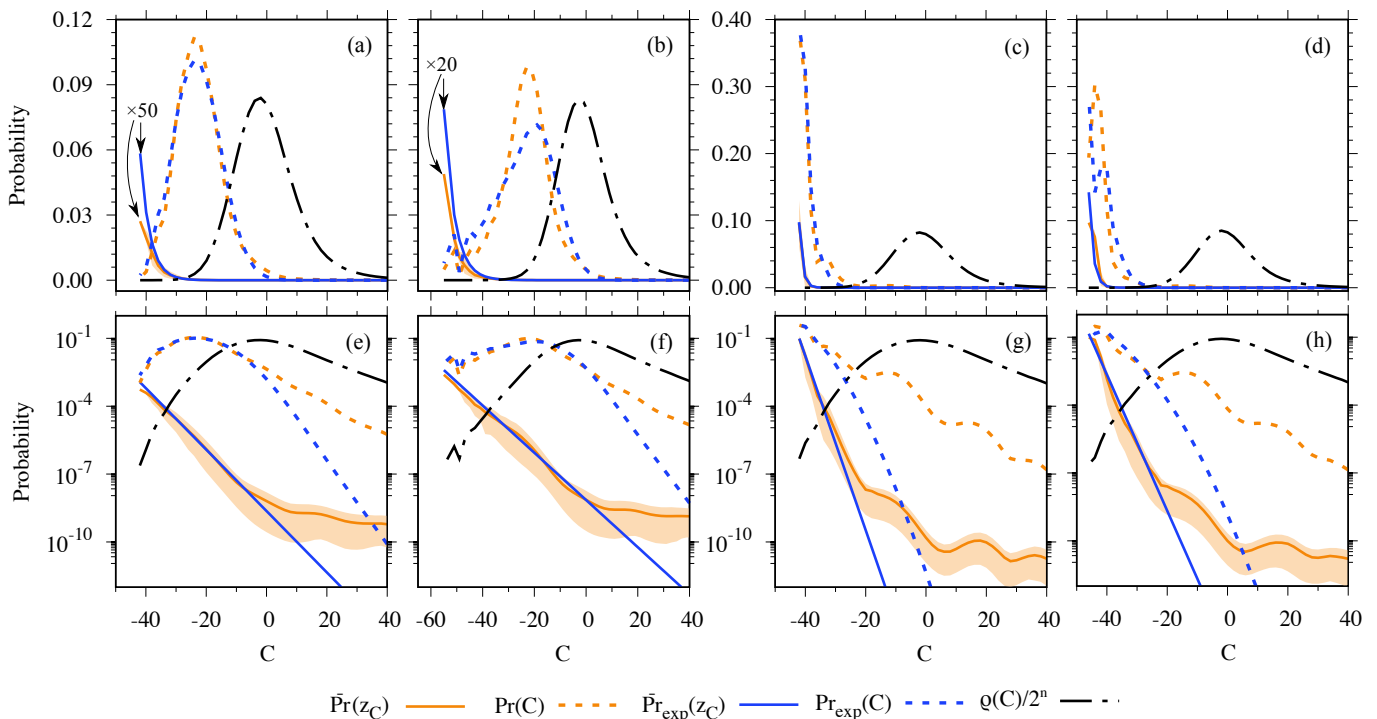


FIG. 6. Exact and best-fit exponential probability distributions for example instances at $n = 23$ on a linear scale (top row) and logarithmic scale (bottom row). (a) and (e) show the instance with median TVD between the true and exponential distributions at $p = 2$ with best-fit $T = 3.20$, (b) and (f) show the worst-case TVD at $p = 2$ with $T = 4.16$, (c) and (g) show the median TVD at $p = 12$ with $T = 1.13$, and (d) and (h) show the worst TVD instance at $p = 12$ with $T = 1.44$. Orange and black curves are related through Eq. (14), blue and black curves are related through Eqs. (15)-(16). Shaded orange bands show the 0.1-0.9 quantiles of the $\overline{\text{Pr}}(z_C)$ distributions at each C (too small to be visible in top panels).

wavefunction $\text{Pr}(C)$ using the total variation distance $\text{TVD} = \sum_C |\text{Pr}_{\text{exp}}(C) - \text{Pr}(C)|/2$. If the probability distributions overlap perfectly then $\text{TVD} = 0$ while $\text{TVD} = 1$ when the distributions are disjoint. We find significant overlaps with median TVD in the interval $0.034 - 0.095$ at each n and p . Thus Eq. (16) succeeds in accounting for much of the true behavior.

We consider specific examples of true and exponential distributions at $n = 23$ in Fig. 6. First we consider the $p = 2$ instance with the median $\text{TVD} = 0.056$ from among the 300 graphs at this n , shown in Fig. 6(a) on a linear scale and below in Fig. 6(e) on a log scale. The density $q(C)$ resembles the schematic density of Fig. 5, but with asymmetries related to correlations in the terms $Z_i Z_j$ in the cost Hamiltonian, which are not present in the simpler schematic distribution. The exponential fit distributions (blue) show behaviors analogous to Fig. 5 and can be understood following the discussion of that figure. The true QAOA distributions (orange) closely resemble the exponential distributions, as expected from the previous analysis of the TVD.

We now consider the worst case instance at $n = 23$ and $p = 2$ in Fig. 6(b)-(f), with $\text{TVD} = 0.149$. The deviations are larger, and this can be attributed to variations in the average basis state probabilities away from exponential behavior, as seen in Fig. 6(f); the $\overline{\text{Pr}}(z_C)$ are slightly

above $\overline{\text{Pr}}_{\text{exp}}(z_C)$ near the most probable cost $C = -23$ and slightly below the $\overline{\text{Pr}}_{\text{exp}}(z_C)$ near the optimal C_{min} . The exponential distribution describes the correct qualitative behavior with semi-quantitative agreement. Furthermore, we might expect that properties of the distribution as a whole will be reproduced with smaller error, for example, probabilities that are too large or too small may tend to cancel in the approximation ratio average. We will consider this further in the Sec. III F.

Fig. 6(c),(d),(g) and (h) are similar to the previous panels but show distributions for $n = 23$ with $p = 12$. Here, the instance with the median $\text{TVD} = 0.087$ is shown in (c) and (g) while the worst case instance with $\text{TVD} = 0.266$ is shown in (d) and (h). This latter instance has the largest TVD from among all cases at $n = 23$ (i.e., over all $n = 23$ graphs at all p).

For the median distribution in Fig. 6(c), the true probabilities are concentrated on and near the optimal solution, in agreement with the exponential approximation. For the worst case distribution in Fig. 6(d), the true and exponential distributions are again concentrated near the optimal solution. However, the individual probabilities at each cost are reproduced with lower accuracy by the exponential distribution, which is peaked at C_{min} and dips at the next value $C_{\text{min}} - 2$, while the true distribution shows an opposite trend; the differences in probabil-

ity near C_{\min} account for the large TVD for this instance. The exponential approximation fails to capture the individual probabilities per cost in the worst case, but it maintains a qualitative similarity to the true distribution.

D. Entropy

We have seen that QAOA states can be described approximately by best-fit Boltzmann distributions. An important feature of the Boltzmann distribution is that it obtains the maximum entropy consistent with a given average energy (or cost) [36]. Here we analyze whether the QAOA distributions match to this maximum-entropy Boltzmann expectation.

We define the Shannon entropy of the output distribution from QAOA as

$$S = - \sum_{\mathbf{z}} \Pr(\mathbf{z}) \log_2 \Pr(\mathbf{z}) \quad (17)$$

The entropy quantifies the uncertainty in the projective measurement distribution from QAOA [37]; similar entropies have also been considered in analyses of quantum pure-state thermodynamics [38–42]. We compare QAOA state entropies against Boltzmann distribution entropies with the same $\langle \hat{C} \rangle$; these are different from the “Boltzmann-fits” of the previous section, as a consistent $\langle \hat{C} \rangle$ is needed for a valid maximum-entropy comparison.

In Fig. 7 we compare the maximum Boltzmann entropy $S_{\text{Boltzmann}}$ against QAOA distribution entropies S_{QAOA} at $n = 23$ and $p = 2$, as well as entropies of randomized states S_{random} with the same $\langle \hat{C} \rangle$, see Appendix B for details. The QAOA states obtain sub-Boltzmann entropies (below the upper diagonal line), but these significantly exceed the entropies from the randomly generated states. Thus, the QAOA states are closer to maximum entropy than the random states, as expected from the approximate correspondence between average QAOA probabilities and the Boltzmann fits, seen previously in Fig. 6.

We attribute the entropy differences between the Boltzmann and QAOA distributions to two factors: First, the average QAOA state probabilities $\overline{\Pr}(\mathbf{z}_C)$ do not exactly follow Boltzmann distributions. This lowers the entropy from the maximum Boltzmann value. Second, there are fluctuations in the individual basis states probabilities $\Pr(\mathbf{z})$ about the averages $\overline{\Pr}(\mathbf{z}_C)$, and these fluctuations further lower the entropy.

To better assess how fluctuations influence the entropy, we perform a final comparison against a type of state with random complex Gaussian fluctuations about a Boltzmann distribution, as described in Appendix B. These states can be understood analytically [41] for a straightforward comparison. Their expected entropy is $S_{\text{Boltzmann, fluc}} = S_{\text{Boltzmann}} - (1 - \gamma_{\text{EM}})/\ln(2)$, with $\gamma_{\text{EM}} \approx 0.577$ the Euler-Mascheroni constant. Thus, for these states the fluctuations decrease the expected entropy by a constant factor $(1 - \gamma_{\text{EM}})/\ln(2)$ relative to the Boltzmann value, yielding the lower diagonal curve




FIG. 7. QAOA state entropies at $n = 23$ and $p = 2$, compared against entropies of random states with the same $\langle \hat{C} \rangle$. The QAOA state entropies are consistent with random fluctuations about a Boltzmann distribution, $S_{\text{QAOA}} \approx S_{\text{Boltzmann, fluc}}$, while the random state entropies are considerably lower.

in Fig. 7, which approximately matches the QAOA entropies. We conclude the QAOA entropy is close to what is expected for a class of states with fluctuations about average Boltzmann distributions.

E. Exponential Scaling

Here we consider the exponential scaling parameter T in the best-fit Boltzmann distributions. Our goal is to devise a heuristic scaling that accounts for QAOA probability distributions across our instances.

We motivate an empirical scaling formula for the temperature as follows. We begin with the factor $\beta_T \equiv 1/T$ that appears within the Boltzmann probability factors. To obtain a generic expression for β_T we assume it can be expanded in powers of the number of layers p , about the limit $\beta_T(p = 0) = 0$ which describes the initial ($p = 0$) uniform superposition state. We find an expansion in powers of $\sqrt{p/\tilde{p}}$ gives a good account of our results, where \tilde{p} is an assumed constant that is characteristic of depths well above the $p \leq 12$ we consider here. We therefore begin with the following expression

$$\beta_T(p/\tilde{p}) = \beta_T(0) + \beta'_T(0) \sqrt{\frac{p}{\tilde{p}}} + \beta''_T(0) \frac{p}{2\tilde{p}} + \dots \quad (18)$$

Using $\beta_T(0) = 0$ and inverting the previous expression




FIG. 8. Two-dimensional histogram with frequencies of exponential scaling parameters T compared against $C_{\min}/n\sqrt{p}$, for each of the 7,200 instances considered. The dotted line shows T_e from Eq. (20).

we find the temperature to order p/\tilde{p}

$$\begin{aligned} T &= \frac{1}{\beta'_T(0)} \sqrt{\frac{\tilde{p}}{p}} \left(1 - \frac{\beta'_T(0)}{2\beta''_T(0)} \sqrt{\frac{\tilde{p}}{p}} \right) \\ &= \frac{\sqrt{\tilde{p}}}{\beta'_T(0)} \frac{1}{\sqrt{p}} - \frac{1}{2\beta''_T(0)} \end{aligned} \quad (19)$$

We fit this function to determine values for the unknown quantities. We find good agreement across our 7,200 instances at varying n and p using the following empirical temperature relationship

$$T_e = c \frac{C_{\min}}{n\sqrt{p}} + d. \quad (20)$$

with $c = -2.738 \pm 0.005$ and $d = -0.255 \pm 0.003$, with \pm denoting the asymptotic standard error of the non-linear least-squares fit. The factor C_{\min}/n is empirical and gives an energy scale for the temperature in terms of the ground state energy per qubit; including a similar energy dependence with d does not significantly effect the quality of the fit.

We examine the best-fit T and the T_e of (20) across the total set of 7,200 instances at varying n and p . Given the large number of data points in our dataset it is difficult to visualize the results in a scatter plot. We therefore bin instances into a two-dimensional histogram and report frequencies of occurrences, resulting in Fig. 8. T_e is shown by the dotted line in Fig. 8; it captures the typical behavior and trend, though individual instances show some deviations scattered around T_e . Note the logarithmic color scale is chosen to make deviations evident, while the majority of cases are close to the heuristic T_e . We conclude T_e gives a satisfactory account of the typical behavior within our dataset.

Notably T_e does not depend on the specific angles γ and β for each instance. A general characterization would need to include γ and β dependence from the QAOA states $|\gamma, \beta\rangle$. For us the β and γ for each instance are generated systematically from local optimization of the SK angles as described in Sec. III A, which evidently leads to consistent scalings $T \approx T_e$ across our instances.

F. Approximating QAOA performance metrics

We have seen that exponential average basis state probabilities are successful in reproducing approximate QAOA distributions, and further that the exponential scaling factors show systematic behavior captured by the heuristic T_e . This suggests typical QAOA performance can be predicted with T_e and the density of solutions $\varrho(C)$, without quantum state simulations; we comment on the difficulty of computing $\varrho(C)$ in Sec. III G. In this section, we assess performance metric distributions generated from the T_e heuristic relative to true values from the wavefunction. We present a detailed error analysis in Appendix C and refer to its main conclusions.

First we consider approximation ratios r and r_{exp} from the exact and approximate distributions in Fig. 9(a), where r_{exp} is computed using T_e (Eq. (20)) in the exponential approximation (Eqs. (15)-(16)). To visualize the distributions, we use a two-dimensional histogram, showing frequencies of occurrences on a logarithmic scale. The r_{exp} slightly overestimates the true r , with median relative error of 0.6% at each n and p , as shown in Appendix C.

Next we consider the accuracy of the approximate exponential distributions for predicting the likelihood to obtain high quality results from QAOA. We consider the cumulative distribution function, which quantifies the total probability to measure any result with cost C less than some threshold C' ,

$$\Pr(C \leq C') = \sum_{C \leq C'} \Pr(C). \quad (21)$$

The cumulative distribution function is a more difficult test for our model, as it relies on predicting the total probability across a subset of costs, rather than an average over all costs as we had with r . These small probability subsets are less likely to contain mutually offsetting errors, so we expect the model to be less accurate than for r .

First we consider an example of the cumulative distribution functions at each p for the example instance seen previously in Fig. 2(a),(e). Recall this had the median TVD from among all instances at $p = 2$ and $n = 23$, when the probabilities were approximated by the best-fit T . Here we consider the distributions that are generated from the heuristic T_e . Fig. 10 shows the cumulative distribution from the quantum state in solid lines at each p as well as the distribution generated from T_e in dashed




FIG. 9. Two-dimensional histograms comparing exact QAOA performance metrics against results generated from the exponential distribution Eq. (16) with the heuristic T_e of Eq. (20); labels for the left-hand axes (computed with T_e) are shown above those axes. (a) the approximation ratio, (b) the cumulative distribution function with $C' = C_{0.08}$, (c) the cumulative distribution function with $C' = C_{0.04}$, and (d) the logarithm of the probability for an optimal solution. Dashed diagonal lines indicate where the exact and approximate results are equal.

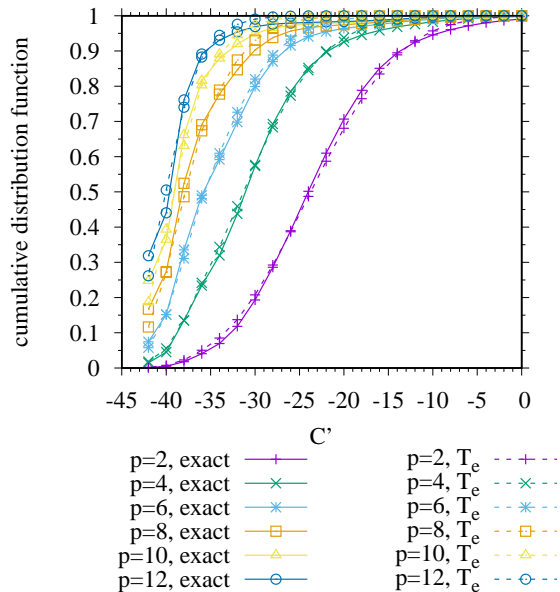


FIG. 10. Cumulative distribution functions for the instance of Fig. 2(a) and (e), at varying p . Solid lines show the exact result $\Pr(C \leq C')$ while dashed lines show heuristic results $\Pr_{\text{exp}}(C \leq C')$ computed with T_e .

lines. The T_e accurately captures the cumulative distribution for each p for this instance.

We assess typical performance in predicting the cumulative distribution for two values of C' that focus on costs near the optimal solution. To quantify closeness to the ground state, we define costs $C_\alpha = (1 - \alpha)C_{\min} + \alpha C_{\max}$. For a given α , any cost $C \leq C_\alpha$ is within α from the optimal relative to the worst case; for example, if $\alpha = 0.08$ then any cost $C \leq C_\alpha$ is within 8% of optimal.

Figure 10 (b) and (c) show the distributions of cumulative distribution function values for $C' = C_{0.08}$ and $C' = C_{0.04}$ respectively. Predictions from T_e correlate well with true values from the wavefunction, keeping

in mind that we are using a logarithmic color scale, so brighter colors are much more significant. Deviations are greater than we had for r , as anticipated. In Appendix C we show that the median relative error is $\leq 11\%$ at each n and p for $C' = C_{0.08}$, while for $C' = C_{0.04}$ the median error is $\leq 35\%$ at $p = 2$ for each n , while at $p > 2$ the median error is $\leq 12\%$.

Finally, we consider predictions of the likelihood to measure an optimal solution $\Pr(C_{\min})$. This is our most stringent test for T_e , as there is no chance for offsetting errors in computing a single probability. Distributions of $\log_{10}(\Pr(C_{\min}))$ are compared against results generated from T_e in Fig. 9(d); we use the logarithm to visualize the varying orders of magnitude that are present. The predicted results show clear correlation to the true results, but with greater errors than in the previous panels of Fig. 9. At $p = 2$ the median relative errors are as large as 65%, which we attribute in part to the small values of $\Pr(C_{\min})$ at this p , which lead to large relative errors under small additive errors. The relative error decreases as p increases, with median errors $\leq 29\%$ at $p > 2$ across all n , as shown in Appendix C. In the next section we further compare the scaling of $\Pr_{\text{exp}}(C_{\min})$ with the scaling observed in the exact results of Fig. 3.

The performance of T_e in capturing important QAOA metrics is notable given that the T_e distributions are far simpler than the exact QAOA states; they use only T_e and $\varrho(C)$ to describe QAOA distributions across the 7,200 instances, as opposed to deep quantum circuit simulations. These exponential distributions give accurate accounts of the approximation ratio, with an error of $\leq 1.5\%$ for 90% of cases at each n and p , while the errors are larger for probabilities in small collections of basis states. We conclude the exponential distributions with T_e capture important aspects of the true behavior.

G. Predicting performance at larger sizes

We have seen that QAOA probability distributions and performance metrics can be approximated by exponential probability distributions $\text{Pr}_{\text{exp}}(C)$, using only the density of solutions $\varrho(C)$ and the scaling parameter T_e . This opens the possibility of predicting performance at larger sizes n , using the approximate relations and without quantum state simulations.

We computed approximate QAOA performance metrics for collections of 300 random Erdős-Rényi graphs at each $n \in \{26, 29, 32, 35, 38\}$ with $p = 2, 4, \dots, 12$. These use $\text{Pr}_{\text{exp}}(C)$ of Eq. (16) with T_e from Eq. (20), with densities of solutions $\varrho(C)$ computed from exhaustive enumeration of the cost values. Exactly computing $\varrho(C)$ requires time $\sim 2^n$, but the time and memory requirements are small compared to what is needed for quantum state simulations. In Fig. 11 we plot the predicted approximation ratios and probabilities for the optimal solution at these n as well as results from the previous $n \in \{14, 17, 20, 23\}$, in all cases using the exponential distributions Eq. (16) with T_e from Eq. (20).

Approximation ratios r are shown in Fig. 11(a). The median r increase with n at small p , then transition to decreasing with n at large p , and the interquantile ranges decrease with n , all in accord with the exact results of Fig. 1(a). Hence, we verify that the model obtains similar results for the approximation ratio at $n \leq 23$ and also predict these same behaviors extend to $n \leq 38$.

Predicted probabilities for an optimal solution are plotted in Fig. 11(b), with the solid line showing the best-fit relation from the exact results in Fig. 3. The predicted results follow similar scaling to the fit from the exact results. It is not so clear how to derive this scaling from the exponential formalism, due to difficulties in computing the partition function Z at varying T , but nonetheless we see the exponential distributions reproduce the known behavior. There are deviations in the rightmost point at each n , corresponding to $p = 2$, which may be due to larger errors in the approximate $\text{Pr}_{\text{exp}}(C_{\min})$ at this p as noted in Sec. III F, though it is worth noting that the exact $p = 2$ medians in Fig. 3 also slightly exceed the best fit line. At $p \geq 4$ (all points except the rightmost points at each n) the median predicted probabilities are close to the trend line observed in the exact results in Fig. 3 (repeated as the black line in Fig. 11(b)). Hence, we confirm the known scaling behavior of $\text{Pr}(C_{\min})$ for $n \leq 23$ and predict similar scaling extends to $n \leq 38$. The predictions depicted in Fig. 11 can be tested in future work, through simulations or direct sampling from a high-fidelity quantum device.

IV. DISCUSSION

The systematic behavior in the approximate QAOA Boltzmann probabilities is surprising. One reason is that the Boltzmann factor $e^{-C/T}$ reflects global struc-

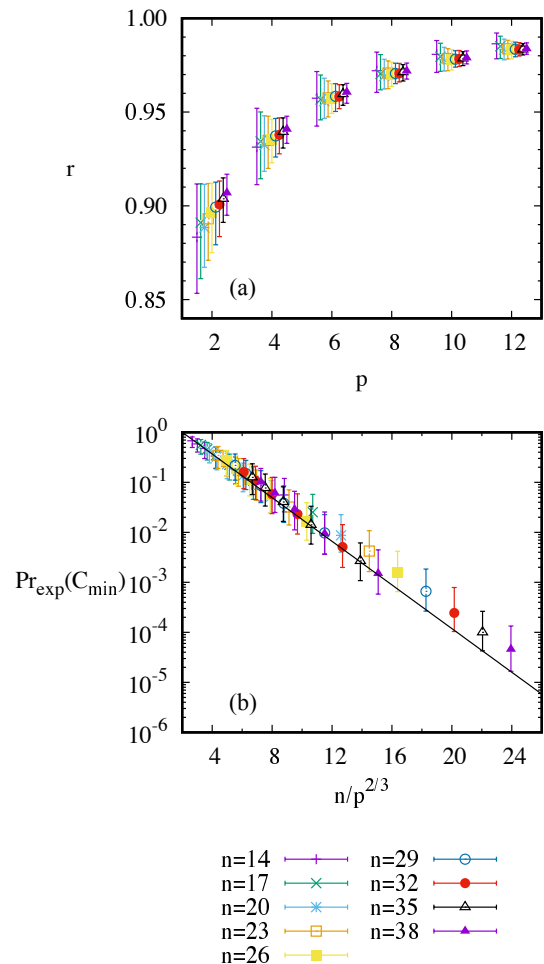


FIG. 11. Predictions for (a) approximation ratios and (b) probabilities for an optimal solution, computed with Eqs. (16) and (20). Points show medians across sets of 300 graphs at each n with $p = 2, 4, \dots, 12$ and error bar show the 0.1-0.9 quantiles. Offsets in p are included in (a) for clarity. In (b), the solid line shows the best-fit scaling from Fig. 3.

ture in the output distribution, in terms of the variable $C = C(\mathbf{z})$ that depends on all qubits, while QAOA is usually understood as a local algorithm. For example, it is well known that QAOA dynamics generate local correlations in terms of p -dependent “subgraphs”, and this has often been taken as a starting point for analyzing QAOA [2, 9, 23, 31, 45]. The current work and Refs. [15–19] have assessed the opposite extreme, in terms of global structure. Exploring the emergence of global structure from local dynamics appears likely to yield new insights into structure and behaviors of QAOA, potentially leading to new ways to predict performance similar to Sec. III G, or to new approaches for understanding prospects for quantum advantage.

There is an interesting connection between the present work and statistical mechanics, where the Boltzmann distribution plays a prominent role in describing microscopic state probabilities in the presence of a heat bath

at temperature T . An important difference here is that the QAOA distribution arises from pure state dynamics alone, without any exterior system or bath. The QAOA states resemble approximate “thermal pure quantum states” in the terminology of Ref. [46], which proposed such states for fundamental investigations of thermodynamic behavior.

The framework of statistical mechanics opens interesting possibilities for future analysis of QAOA. One interesting aspect of physical Boltzmann distributions is that they concentrate in energy, with fluctuations $\Delta E/E$ that tend to zero as the size of the system increases [47]. This is the statistical mechanical explanation for why macroscopic systems have an effectively constant energy at thermal equilibrium. It could be interesting to apply similar reasoning to QAOA, to argue for p -independent concentration at large n , which might provide an alternative route to establishing concentration results beyond small p analyses [2] or detailed considerations of specific instances [13, 23]. To establish such an analogy it appears necessary to derive exponential distributions from the QAOA ansatz, to better assess their prevalence and limitations. This is an interesting topic for future work and may build on initial efforts of Refs. [15, 17, 19].

It is important to clarify that we have analyzed a dataset at small n and p and that different behavior might be observed at large n and p . For worst-case instances it is expected that exponential time will be needed to reach exact solutions, which could arise for example if the number of algorithmic layers or the time for angle optimization scales exponentially. At the same time, we find systematic behavior across small typical instances, and it is certainly worth testing how well this extends to larger typical instances as well as comparing against worst-case behavior.

Although we cannot know the true large size scaling for QAOA given our dataset, it may nonetheless be useful to assess what would be expected if the scaling we observed were to extend to larger sizes. How many layers would be needed to identify an optimal solution at larger n with a modestly high probability? From our scaling we estimate that to obtain a fixed optimal solution probability, the estimated p at larger n entails resources that may be very difficult to achieve on near-term hardware, unless the true large-size scaling is actually much more forgiving than what we have observed at small sizes. Fig. 12 shows the size-dependent number of layers for varying optimal solution probabilities $\Pr(C_{\min})$ (or varying expected numbers of measurements $1/\Pr(C_{\min})$), based on the scaling in Figs. 3,11. We estimate hundreds of layers are expected at sizes of $n \sim 10^3 - 10^4$, in the absence of noise. Implementing these deep circuits with high fidelity would require very low noise levels. There may also be non-linear behaviors that were not detected in our dataset but that become dominant at larger n or p , which may further limit performance. Overall these extrapolations suggest demanding resource requirements for identifying optimal solutions with QAOA at scale,

unless the true large size scaling is more forgiving than what is suggested by our fitted small-size trend.




FIG. 12. Estimated number of layers p to obtain a fixed optimal solution probability $\Pr(C_{\min})$, based on the empirical scaling in Figs. 3,11.

Finally, it is worth noting that generating samples from exponential distributions is of independent interest for a variety of applications. These include understanding the thermodynamic properties of materials [46], optimization with simulated annealing [48], and machine learning with “Boltzmann machines” [49]. If QAOA can quickly generate exponentially distributed samples then it might find use in new quantum-classical algorithms.

DATA AVAILABILITY

Simulation results are available at: <https://code.ornl.gov/5ci/approximate-boltzmann-distributions-in-quantum-approximate-optimization>

ACKNOWLEDGMENTS

P.C.L. thanks Ryan Bennink for interesting discussions of combinatorial optimization, feedback, and encouragement. This work was supported by DARPA ONISQ program under award W911NF-20-2-0051. J. Ostrowski acknowledges the Air Force Office of Scientific Research award, AF-FA9550-19-1-0147. G. Siopsis acknowledges the Army Research Office award W911NF-19-1-0397. J. Ostrowski and G. Siopsis acknowledge the National Science Foundation award OMA-1937008. This research used resources of the Compute and Data Environment for Science (CADES) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

-
- [1] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [2] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028) (2014).
- [3] S. Harwood, C. Gambella, D. Tenev, A. Simonetto, D. Bernal, and D. Greenberg, Formulating and solving routing problems on quantum computers, *IEEE Transactions on Quantum Engineering* **2**, 1 (2021).
- [4] P. Vikstål, M. Grönkvist, M. Svensson, M. Andersson, G. Johansson, and G. Ferrini, Applying the quantum approximate optimization algorithm to the tail-assignment problem, *Physical Review Applied* **14**, 034009 (2020).
- [5] H. Jing, Y. Wang, and Y. Li, Data-driven quantum approximate optimization algorithm for cyber-physical power systems, [arXiv:2204.00738](https://arxiv.org/abs/2204.00738) (2022).
- [6] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, *Quantum* **6**, 759 (2022).
- [7] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: Worst case examples, [arXiv:2005.08747](https://arxiv.org/abs/2005.08747) (2020).
- [8] J. Basso, E. Farhi, K. Marwaha, B. Villalonga, and L. Zhou, The quantum approximate optimization algorithm at high depth for maxcut on large-girth regular graphs and the sherrington-kirkpatrick model, [arXiv:2110.14206](https://arxiv.org/abs/2110.14206) (2021).
- [9] K. Marwaha, Local classical MAX-CUT algorithm outperforms $p = 2$ QAOA on high-girth regular graphs, *Quantum* **5**, 437 (2021).
- [10] P. C. Lotshaw, T. S. Humble, R. Herrman, J. Ostrowski, and G. Siopsis, Empirical performance bounds for quantum approximate optimization, *Quantum Information Processing* **20**, 1 (2021).
- [11] V. Akshay, H. Philathong, M. E. Morales, and J. D. Biamente, Reachability deficits in quantum approximate optimization, *Physical review letters* **124**, 090504 (2020).
- [12] V. Akshay, H. Philathong, E. Campos, D. Rabinovich, I. Zacharov, X.-M. Zhang, and J. Biamente, Circuit depth scaling for quantum approximate optimization, *Phys. Rev. A* **106**, 042438 (2022).
- [13] D. Lykov, J. Wurtz, C. Poole, M. Saffman, T. Noel, and Y. Alexeev, Sampling frequency thresholds for quantum advantage of quantum approximate optimization algorithm, [arXiv:2206.03579](https://arxiv.org/abs/2206.03579) (2022).
- [14] H. Ushijima-Mwesigwa, R. Shaydulin, C. F. Negre, S. M. Mniszewski, Y. Alexeev, and I. Safro, Multilevel combinatorial optimization across quantum architectures, *ACM Transactions on Quantum Computing* **2**, 1 (2021).
- [15] P. Díez-Valle, D. Porras, and J. J. García-Ripoll, QAOA pseudo-Boltzmann states, [arXiv:2201.03358](https://arxiv.org/abs/2201.03358) (2022).
- [16] S. Leontica and D. Amaro, Quantum optimization with instantaneous quantum polynomial circuits, [arXiv:2210.05526](https://arxiv.org/abs/2210.05526) (2022).
- [17] T. Hogg, Quantum search heuristics, *Phys. Rev. A* **61**, 052311 (2000).
- [18] T. Hogg and D. Portnov, Quantum optimization, *Information Sciences* **128**, 181 (2000).
- [19] J. Sud, S. Hadfield, E. Rieffel, N. Tubman, and T. Hogg, A parameter setting heuristic for the quantum alternating operator ansatz, [arXiv:2211.09270](https://arxiv.org/abs/2211.09270) (2022).
- [20] G. G. Guerreschi and A. Y. Matsuura, QAOA for Max-Cut requires hundreds of qubits for quantum speed-up, *Scientific reports* **9**, 1 (2019).
- [21] S. Boulebnane and A. Montanaro, Predicting parameters for the quantum approximate optimization algorithm for MAX-CUT from the infinite-size limit, [arXiv:2110.10685](https://arxiv.org/abs/2110.10685) (2021).
- [22] R. Shaydulin, P. C. Lotshaw, J. Larson, J. Ostrowski, and T. S. Humble, Parameter transfer for quantum approximate optimization of weighted maxcut, [arXiv:2201.11785](https://arxiv.org/abs/2201.11785) (2022).
- [23] F. G. Brandao, M. Broughton, E. Farhi, S. Gutmann, and H. Neven, For fixed control parameters the quantum approximate optimization algorithm's objective function value concentrates for typical instances, [arXiv:1812.04170](https://arxiv.org/abs/1812.04170) (2018).
- [24] J. Wurtz and D. Lykov, Fixed-angle conjectures for the quantum approximate optimization algorithm on regular MaxCut graphs, *Physical Review A* **104**, 052419 (2021).
- [25] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices, *Physical Review X* **10**, 021067 (2020).
- [26] A. Galda, X. Liu, D. Lykov, Y. Alexeev, and I. Safro, Transferability of optimal QAOA parameters between random graphs, in *2021 IEEE International Conference on Quantum Computing and E* (IEEE, 2021) pp. 171–180.
- [27] A. Lucas, Ising formulations of many np problems, *Frontiers in physics* , 5 (2014).
- [28] A. Hagberg, P. Swart, and D. S Chult, *Exploring network structure, dynamics, and function using NetworkX*, Tech. Rep. (Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008).
- [29] J. Wurtz and P. J. Love, Counterdiabaticity and the quantum approximate optimization algorithm, *Quantum* **6**, 635 (2022).
- [30] X. Lee, N. Xie, D. Cai, Y. Saito, and N. Asai, A depth-progressive initialization strategy for quantum approximate optimization algorithm, [arXiv:2209.11348](https://arxiv.org/abs/2209.11348) (2022).
- [31] J. Basso, E. Farhi, K. Marwaha, B. Villalonga, and L. Zhou, The quantum approximate optimization algorithm at high depth for MaxCut on large-girth regular graphs and the Sherrington-Kirkpatrick model, [arXiv:2110.14206](https://arxiv.org/abs/2110.14206) (2022).
- [32] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, [arXiv:1910.08187v3](https://arxiv.org/abs/1910.08187v3) (2021).
- [33] S. G. Johnson, *The NLOpt nonlinear-optimization package*, accessed January 2022.
- [34] J. Nocedal, Updating quasi-Newton matrices with limited storage, *Math. Comput.* **35**, 773 (1980).
- [35] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, *Math. Programming* **45**, 503 (1989).
- [36] E. T. Jaynes, Information theory and statistical mechanics, *Physical review* **106**, 620 (1957).

- [37] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, 2010).
- [38] P. C. Lotshaw and M. E. Kellman, Quantum microcanonical entropy, Boltzmann’s equation, and the second law, *J. Phys. Chem. A* **123**, 831 (2018).
- [39] G. L. Barnes, P. C. Lotshaw, and M. E. Kellman, Quantum thermodynamics, entropy of the universe, free energy, and the second law, [arXiv:1511.06176](#) (2015).
- [40] P. C. Lotshaw and M. E. Kellman, Asymmetric temperature equilibration with heat flow from cold to hot in a quantum thermodynamic system, *Phys. Rev. E* **104**, 054101 (2021).
- [41] P. C. Lotshaw and M. E. Kellman, On quantum entropy and excess entropy production in a system-environment pure state, [arXiv:2211.14152](#) (2022).
- [42] A. Polkovnikov, Microscopic diagonal entropy and its connection to basic thermodynamic relations, *Ann. Phys.* **326**, 486 (2011).
- [43] M. Larocca, N. Ju, D. García-Martín, P. J. Coles, and M. Cerezo, Theory of overparametrization in quantum neural networks, [arXiv:2109.11676](#) (2021).
- [44] D. Gamarnik and Q. Li, On the Max-Cut of sparse random graphs, *Random Struct. Alg.* **52**, 219 (2018), [arXiv:1411.1698](#).
- [45] J. Wurtz and P. Love, MaxCut quantum approximate optimization algorithm performance guarantees for $p > 1$, *Physical Review A* **103**, 042612 (2021).
- [46] S. Sugiura and A. Shimizu, Thermal pure quantum states at finite temperature, *Phys. Rev. Lett.* **108**, 240401 (2012).
- [47] M. Plischke and B. Bergersen, *Equilibrium statistical physics*, 3rd ed. (World scientific, 2006).
- [48] S. Kirkpatrick, C. D. G. Jr., and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
- [49] G. Montufar, Restricted boltzmann machines: Introduction and review, [arXiv:1806.07066](#) (2018).
- [50] L. R. Mead and N. Papanicolaou, Maximum entropy in the problem of moments, *Journal of Mathematical Physics* **25**, 2404 (1984).

Appendix A: Angle transfer compared to brute force search

To assess performance of our angle transfer approach from Sec. III A, we compared against a more standard approach based on a brute force search with a modest amount of sampling. We generated 100 random initial angles for each instance at $n = 14$ and $p = 2, 4, \dots, 12$. We optimized each random set of angles with the BFGS algorithm to identify a nearby local optimum, then selected the angles β^* and γ^* that obtained the largest approximation ratio for each instance. We denote these approximation ratios r_{bf} and compare them against our transfer approach of Sec. III A, denoted here as r_{tr} . Similar brute force approaches have been considered in previous publications, which have concluded that angle transfer achieves similar performance to brute force optimization at $p \leq 5$ [10, 21, 22].

Figure 13 compares the approximation ratios obtained




FIG. 13. Approximation ratios from our angle transfer procedure r_{tr} are equivalent or larger than approximation ratios r_{bf} from a brute force approach that optimizes 100 random initial angles for each instance to identify a best set of parameters. The dashed diagonal line shows where $r_{\text{tr}} = r_{\text{bf}}$.

from the two approaches. At $p = 2$, the results are equivalent for all cases, while at $p \geq 4$ the transfer approach performs equivalently or better in all cases, $r_{\text{tr}} \geq r_{\text{bf}}$. This confirms that our angle selection procedure is competitive with a more standard brute force approach, consonant with previous analyses at $p \leq 5$.

A more exhaustive analysis of performance of our procedure would require greater numbers of random initial samples for the brute force case. When the number of samples is sufficiently large, this should identify the global optimal angles and yield results that cannot be worse than the transfer approach. However, this may require a very large number of samples, as the size of the γ, β parameter space increases exponentially with p . This makes the brute force searches impractical. By contrast, the angle selection procedure we have used here is much more scalable, and achieves impressive results in Fig. 13 and Fig. 1.

Appendix B: Entropy analysis

Here we present details of how each entropy is computed in Sec. III.D. The entropy of the QAOA states S_{QAOA} is evaluated directly from the wavefunction probabilities $\text{Pr}(\mathbf{z})$ following Eqs. (7) and (17). The Boltzmann entropy is defined using the Boltzmann distribution $\text{Pr}_{\text{Boltzmann}}(\mathbf{z}) = e^{-C(\mathbf{z})/T_B}/Z$ with the same expected cost $\langle \hat{C} \rangle$ as a given QAOA instance, and with $Z = \sum_{\mathbf{z}} e^{-C(\mathbf{z})/T_B}$. We compute the Boltzmann distribution and its entropy following the approach of minimizing the “ Γ potential” of Ref. [50] to identify a T_B

that satisfies the standard statistical mechanical relation

$$-\frac{\partial \ln Z}{\partial(1/T_B)} = \langle \hat{C} \rangle. \quad (\text{B1})$$

The Boltzmann entropy is then computed using the standard relation (expressed in units of \log_2)

$$S_{\text{Boltzmann}} = \frac{1}{\ln(2)} \left(\frac{\langle \hat{C} \rangle}{T_B} + \ln Z \right) \quad (\text{B2})$$

The entropy $S_{\text{Boltzmann, fluc}}$ is devised following the approach of Ref. [41]. We consider a quantum state with random independent complex Gaussian fluctuations about the Boltzmann distribution

$$|\psi\rangle = \sum_{\mathbf{z}} \tilde{g}_{\mathbf{z}} \sqrt{\frac{e^{-C(\mathbf{z})/T_B}}{Z}} |\mathbf{z}\rangle \quad (\text{B3})$$

Here $\tilde{g}_{\mathbf{z}} = (g_{\mathbf{z}} + ig'_{\mathbf{z}})/\sqrt{2}$, where each of $g_{\mathbf{z}}$ and $g'_{\mathbf{z}}$ are taken as random independent Gaussian variates, with probability density $\Pr(g_{\mathbf{z}}) = e^{-g_{\mathbf{z}}^2/2}/\sqrt{2\pi}$ and similarly for $g'_{\mathbf{z}}$. The expected entropy is then approximated by taking expectation values for the terms containing $\tilde{g}_{\mathbf{z}}$, which yields the correction factor $-(1 - \gamma_{EM})/\ln(2)$ in addition to the Boltzmann entropy [41].

We generate the random states with the same $\langle \hat{C} \rangle$ as the QAOA states and compute their entropies as follows. We consider a set F of all the probability assignments consistent with the constraints

$$\begin{aligned} \sum_C \Pr(C) \varrho(C) C &= \langle \hat{C} \rangle \\ \sum_C \Pr(C) \varrho(C) &= 1 \\ \Pr(C) &\geq 0 \quad \forall C \end{aligned} \quad (\text{B4})$$

To generate a randomized probability set $\{\Pr(C)\} \in F$, we first generate a uniformly random set of probabilities $\{\Pr'(C)\}$. We then find the closest point $\{\Pr(C)\} \in F$ with respect to the L_2 metric, by minimizing

$$\min_{\{\Pr(C)\} \in F} \sum_C (\Pr'(C) - \Pr(C))^2 \quad (\text{B5})$$

To perform this minimization we use the software package *Gurobi*. The result is a randomized probability assignment $\{\Pr(C)\}$ consistent with the QAOA expected cost $\langle \hat{C} \rangle$. Finally, for each cost C we generate $\varrho(C)$ different basis state probabilities $\Pr(\mathbf{z})$ uniformly at random and normalize these probabilities so they sum to $\Pr(C)$. We then compute the entropy Eq. (17) using these randomized $\Pr(\mathbf{z})$. The resulting entropy corresponds to a randomized state with the same $\langle \hat{C} \rangle$ as the QAOA state.

Appendix C: Error analysis of approximate QAOA performance

We define relative and absolute errors to compare the simple approximate QAOA distributions with T_e , from

Eqs. (16) and (20), to the true QAOA distributions in Eq. (11). We define the absolute relative error as

$$\epsilon_R(Q) = |1 - Q_{\text{exp}}(T_e)/Q|, \quad (\text{C1})$$

where Q is the quantity of interest, for example, the approximation ratio $Q = r$ or probability of the ground state $Q = \Pr(C_{\min})$. Similarly we define the absolute difference error as

$$\epsilon_D(Q) = |Q_{\text{exp}}(T_e) - Q|. \quad (\text{C2})$$

In each case $\epsilon = 0$ signifies zero error, while the amount of error increases with ϵ . In Fig. 14 we analyze errors for each of the quantities considered previously in Fig. 9 of Sec. III.F. We consider median errors, denoted by points in Fig. 14, with error bars showing the 0.1-0.9 quantiles. The top row shows the relative absolute error while the bottom row shows the absolute difference error.

In Fig. 14(a) and (e) we plot errors $\epsilon_R(r)$ and $\epsilon_D(r)$ respectively. The median error is $\leq 0.6\%$ at each n and p , while 90% of cases at each n and p have relative errors $\leq 1.5\%$, as shown by the error bars. The worst-case error is 2.9%. We conclude that T_e gives very accurate predictions for the approximation ratio across the 7,200 instances we consider.

In Fig. 14(b) and (f) we assess error in predicting the cumulative distribution function $\Pr(C \leq C_{0.08})$. The relative errors are largest at $p = 2$, and decrease as p and $\Pr(C \leq C_{0.08})$ increase. The absolute difference error in Fig. 14(f) shows only minor variations across n and p , consistent with larger relative errors at small p , when $\Pr(C \leq C_{0.08})$ is small. The median relative error in $\Pr(C \leq C_{0.08})$ from the approximation with T_e is $\leq 11\%$ for median cases, and $\leq 33\%$ in 90% of cases, as seen by the 0.9 quantile error bars. Thus, for the majority of instances at each n and p , we obtain an accurate prediction of $\Pr(C \leq C_{0.08})$ to within a factor of 1.33 or better. Larger errors are observed for the cumulative distribution function $\Pr(C \leq C_{0.04})$ in Fig. 14(c) and (g), as anticipated from the previous analysis of Fig. 9. At $p = 2$ the median relative error is $\leq 35\%$, when $\Pr(C \leq C_{0.04})$ is small. At $p > 2$ the median relative error decreases to $\leq 12\%$ across all n and p .

Finally, we consider predicted probabilities to observe an optimal solution $\Pr(C_{\min})$. Here the absolute error is close to zero at $p = 2$, when the probability in optimal solutions is small. These small absolute difference errors result in median relative errors $\leq 65\%$. The difference errors increase with p , as more population enters the optimal solution states, while the relative errors decrease for a similar reason. The median errors are $\leq 29\%$ at $p > 2$, indicating a better approximation for these p .

Overall, the median errors tend to increase with n , as do the spread of the errors, as reflected in the 0.1-0.9 quantile error bars. Nonetheless, the increase in median relative error with n is fairly modest for most cases, especially at $p > 2$, which indicates versatility of the approach in predicting median case behavior over our graph ensembles.

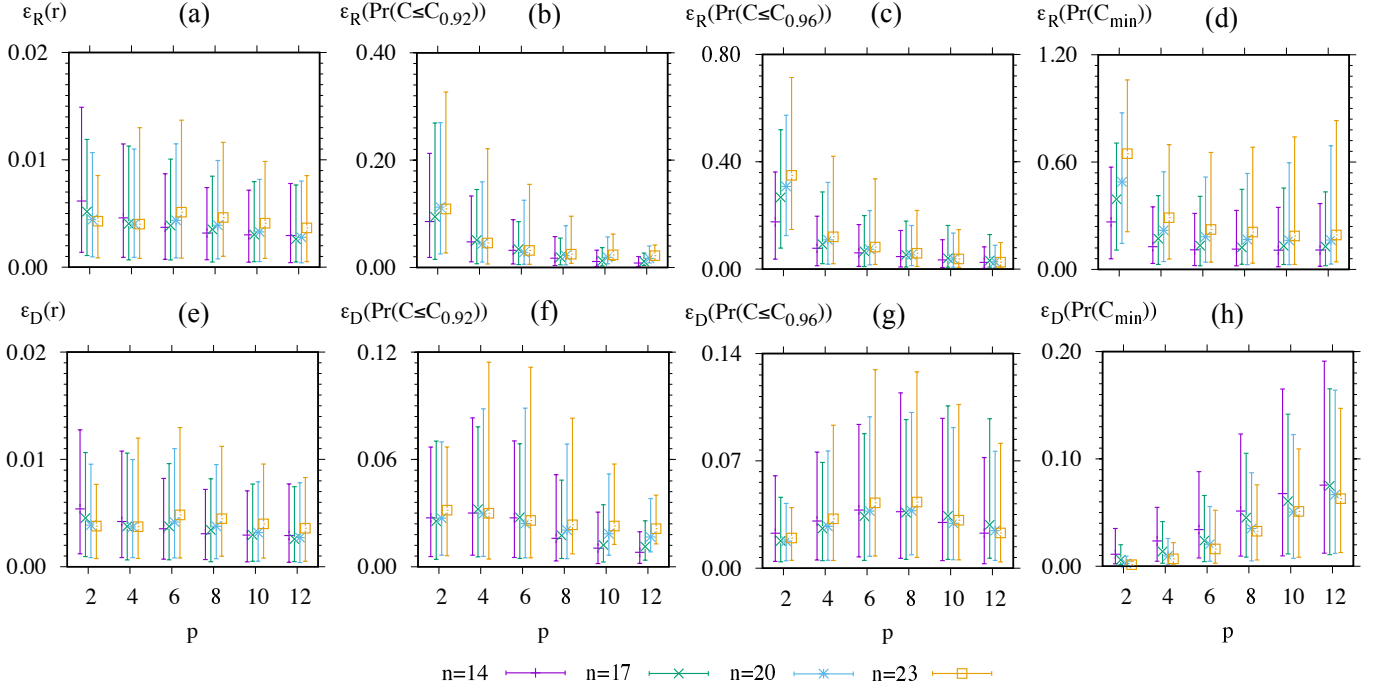


FIG. 14. Errors in approximating various QAOA performance metrics with exponential distributions using the empirical scaling factor T_e . The top row shows relative errors, ϵ_R of Eq. (C1), while the bottom row shows difference errors, ϵ_D of Eq. (C2). Each column shows different performance metrics: (a),(e) the approximation ratio r , (b),(f) the cumulative distribution function $\Pr(C \leq C_{0.92})$, (c),(g) the cumulative distribution function $\Pr(C \leq C_{0.96})$, and (d),(h) the probability of to measure an optimal solution. Points are medians over the sets of 300 graphs at each n and error bars show the 0.1-0.9 quantiles.

Appendix D: Supplemental Figures

In Fig. 2 we did not include error bars because they crowded the curves in the figure. Here in Fig. 15 we

present a version with error bars to gauge the size of deviations among varying instances.

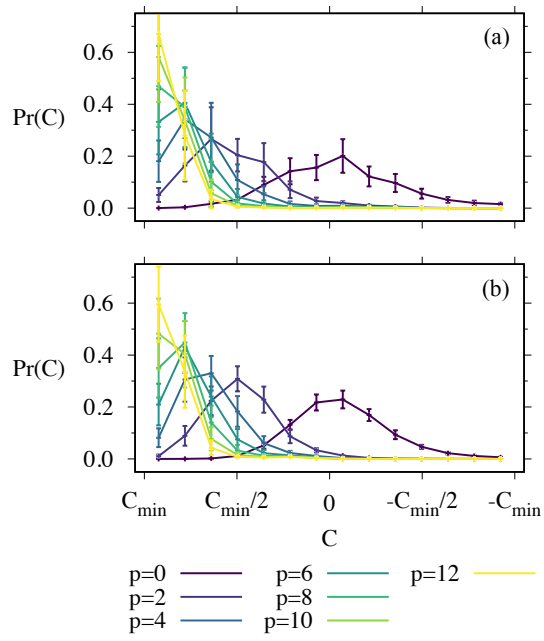


FIG. 15. Same as Fig. 2 but with standard deviation error bars.