

Energy Gap of Quantum Spin Glasses: A Projection Quantum Monte Carlo Study

L. Brodolini^{1,2}, G. E. Astrakharchik³, S. Giorgini⁴, and S. Pilati^{1,2}

¹*CQM group, School of Science and Technology, Physics Division, Università di Camerino, 62032 Camerino, Italy*

²*INFN, Sezione di Perugia, I-06123 Perugia, Italy*

³*Departament de Física, Universitat Politècnica de Catalunya, Campus Nord B4-B5, E-08034 Barcelona, Spain*

⁴*Pitaevskii BEC Center, CNR-INO and Dipartimento di Fisica, Università di Trento, I-38123 Trento, Italy*



(Received 27 February 2026; revised 18 June 2026; accepted 13 July 2026; published 3 August 2026)

The performance of quantum annealing for combinatorial optimization is fundamentally limited by the minimum energy gap Δ encountered at quantum phase transitions. We investigate the scaling of Δ with system size N for two paradigmatic quantum spin-glass models: the two-dimensional Edwards-Anderson (2D-EA) and the all-to-all Sherrington-Kirkpatrick (SK) models. Utilizing a newly proposed unbiased energy-gap estimator for continuous-time projection quantum Monte Carlo simulations, complemented by high-performance sparse eigenvalue solvers, we characterize the gap distributions across disorder realizations. It is found that, in the 2D-EA case, the inverse-gap distribution develops a fat tail with infinite variance as N increases. This indicates that the unfavorable superalgebraic scaling of Δ , recently reported for binary couplings [M. Bernaschi *et al.*, The quantum transition of the two-dimensional Ising spin glass, *Nature (London)* **631**, 749 (2024)], persists for the Gaussian disorder considered here, pointing to a universal feature of 2D spin glasses. Conversely, the SK model retains a finite-variance distribution, with the disorder-averaged gap following a rather slow power law, close to $\Delta \propto N^{-1/3}$. This finding provides a promising outlook for the potential efficiency of quantum annealers for optimization problems with dense connectivity.

DOI: [10.1103/fm8m-kz13](https://doi.org/10.1103/fm8m-kz13)

Introduction—The spectral gap Δ —the energy difference between the ground state and the first excited state—is a fundamental quantity across condensed-matter physics, quantum chemistry, and quantum information science. In the context of adiabatic quantum computing, the smallest value that Δ attains on the annealing path governs the computational complexity of solving combinatorial optimization problems. In fact, the time required to avoid nonadiabatic transitions and ensure an optimal solution scales as $1/\Delta^2$ [1]. As many quadratic optimization problems can be mapped onto Ising spin glasses, significant effort has focused on these models and, in particular, on the spin-glass quantum phase transition, where Δ is expected to vanish rapidly with the number of spins N [2–5]. Despite its importance, estimation of the energy gaps remains a formidable computational problem [6–8], leading to conflicting reports on the scaling of Δ in the literature [9–13]. This uncertainty has spurred the search for more efficient numerical algorithms capable of probing the large N regime [14–17].

Recent large-scale path-integral Monte Carlo (PIMC) simulations of the two-dimensional quantum Edwards-Anderson (2D-EA) model [18] have clarified the role of the parity symmetry. It was found that the smallest gap—separating the even ground state from the first odd excited state—exhibits a superalgebraic scaling with system size. Instead, for the symmetry restricted gap, a power-law scaling was found, indicating a polynomial computational complexity. However, these findings were restricted to square lattices with binary couplings, whereas many realistic optimization problems—e.g., mean-variance portfolio optimization or number partitioning problems—involve denser connectivity and continuous coupling distributions [19,20].

In this Letter, we investigate the scaling of Δ for both the 2D-EA and the all-to-all Sherrington-Kirkpatrick (SK) models featuring Gaussian couplings and uniform transverse field. Our analysis is powered by unbiased continuous-time projection quantum Monte Carlo (PQMC) simulations [21], which we extend here so that energy gaps also can be extracted. Notably, we show that the energy-gap estimator that we introduce is unbiased—meaning it is independent of the choice of the guiding wave function—ensuring high-fidelity results for systems up to $N \gtrsim 100$. Beyond PQMC simulations, we also resort to high-performance sparse eigenvalue solvers for smaller sizes ($N \lesssim 30$).

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

The above two computational approaches allow us to extend the previous study on the 2D-EA model [18] to more general, continuous couplings, as well as to establish via unbiased numerics the gap scaling in the SK case [22], addressing both symmetry restricted and unrestricted gaps. Specifically, our results for the 2D-EA model demonstrate that the distribution of the inverse symmetry-unrestricted gap $\eta = 1/\Delta$ acquires a fat tail with infinite variance for a finite N , indicating that the superalgebraic scaling, previously observed for binary couplings, is indeed a universal feature that persists for continuous couplings. Conversely, for the SK model, the variance of η remains finite, and the disorder-averaged gap follows a favorable power-law scaling with a rather small exponent, close to $\Delta \propto N^{-1/3}$. This suggests that the all-to-all connectivity of the SK model provides a significant advantage for quantum annealing over optimization problems that can be mapped to 2D sparse graphs.

Models and methods—The quantum Ising Hamiltonians that we address are defined using conventional Pauli matrices σ_i^x and σ_i^z as

$$\hat{H} = -\Gamma \sum_{i=1}^N \sigma_i^x - \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z. \quad (1)$$

The first term describes the effect of a uniform transverse field Γ . The second term accounts for interactions which, in the 2D-EA model, act only on the four nearest neighbors of a square lattice of side length $L = \sqrt{N}$, obeying periodic boundary conditions. The random couplings are sampled from a normal distribution with zero mean and variance J^2 , namely $J_{ij} \sim \mathcal{N}(0, J^2)$. In the SK model, the interactions act on every spin pair with couplings $J_{ij} \sim \mathcal{N}(0, J^2/N)$. In the following, the scale J is used as the energy unit.

Our main computational approach is the QMC algorithm. Hereafter, we show how an unbiased gap-estimator algorithm can be incorporated. We adopt the standard computational basis of eigenstates $|\mathbf{x}\rangle$ of σ_i^z , such that $\sigma_i^z|\mathbf{x}\rangle = x_i|\mathbf{x}\rangle$, and the 2^N spin configurations are denoted as $\mathbf{x} = (x_1, \dots, x_N)$, with $x_i = \pm 1$ for $i = 1, \dots, N$. Instead of directly solving the imaginary-time Schrödinger equation for the many-body wave function $\psi(\mathbf{x}, \tau) = \langle \mathbf{x} | \psi(\tau) \rangle$, we improve the convergence by performing the importance sampling with a guiding wave function $\psi_g(\mathbf{x}) = \langle \mathbf{x} | \psi_g \rangle$. The population of walkers then samples configurations proportionally to the distribution $f(\mathbf{x}, \tau) = \psi_g(\mathbf{x})\psi(\mathbf{x}, \tau)$, which evolves according to the imaginary-time ($\tau = it/\hbar$) Schrödinger equation, modified by the guiding function and a reference energy E_{ref} . The corresponding Green's function is

$$G(\mathbf{x}, \mathbf{x}_0, \tau) = \psi_g(\mathbf{x}) \langle \mathbf{x} | e^{-\tau \hat{H}_r} | \mathbf{x}_0 \rangle \psi_g^{-1}(\mathbf{x}_0), \quad (2)$$

where $\hat{H}_r = \hat{H} - E_{\text{ref}}$ is the energy-shifted Hamiltonian. The Hamiltonian (1) commutes with the parity operator

$\hat{P} = \prod_{i=1}^N \sigma_i^x$, whose eigenvalues are either -1 or 1 . Thus, the eigenstates separate into the odd and even sectors, with corresponding energy levels $E_{o,0}, E_{o,1}, \dots$, and $E_{e,0}, E_{e,1}, \dots$, respectively. The states $|\psi_0\rangle$ and $|\psi_g\rangle$ are even, and the ground-state energy is $E_0 = E_{e,0}$.

To derive the gap estimator, we evaluate the matrix element $\langle \psi_g | \hat{O}(\tau) \hat{O}(0) | \psi_0 \rangle = \langle \psi_g | \hat{O}(0) \hat{O}(-\tau) | \psi_0 \rangle$, where the equality follows from the time invariance. Here, $\hat{O}(\tau)$ denotes the Heisenberg representation of the observable $\hat{O}$ at imaginary time τ , defined as $\hat{O}(\tau) = e^{\tau \hat{H}_r} \hat{O} e^{-\tau \hat{H}_r}$. After equilibration from the initial state, E_{ref} provides an unbiased estimate of the ground-state energy E_0 . We can hence write $e^{\tau \hat{H}_r} |\psi_0\rangle = |\psi_0\rangle$ and, using two completeness relations, the matrix element can be written as

$$\langle \psi_g | \hat{O} \hat{O}(-\tau) | \psi_0 \rangle = \sum_{\mathbf{x}, \mathbf{x}_0} \psi_g(\mathbf{x}) \langle \mathbf{x} | \hat{O} e^{-\tau \hat{H}_r} \hat{O} | \mathbf{x}_0 \rangle \psi_0(\mathbf{x}_0).$$

For diagonal operators, i.e., such that $\hat{O}|\mathbf{x}\rangle = O(\mathbf{x})|\mathbf{x}\rangle$, the above equation can be recast as

$$\begin{aligned} \langle \psi_g | \hat{O} \hat{O}(-\tau) | \psi_0 \rangle &= \sum_{\mathbf{x}, \mathbf{x}_0} O(\mathbf{x}) O(\mathbf{x}_0) \psi_g(\mathbf{x}) \\ &\quad \times \langle \mathbf{x} | e^{-\tau \hat{H}_r} | \mathbf{x}_0 \rangle \psi_0(\mathbf{x}_0). \end{aligned}$$

Recalling that after equilibration for any τ one has $f(\mathbf{x}, \tau) = \psi_g(\mathbf{x})\psi_0(\mathbf{x})$ and using the definition of $G(\mathbf{x}, \mathbf{x}_0, \tau)$ in Eq. (2), one gets

$$\begin{aligned} \langle \psi_g | \hat{O} \hat{O}(-\tau) | \psi_0 \rangle &= \sum_{\mathbf{x}, \mathbf{x}_0} O(\mathbf{x}) O(\mathbf{x}_0) G(\mathbf{x}, \mathbf{x}_0, \tau) f(\mathbf{x}_0, 0) \\ &= \sum_{\mathbf{x}} O(\mathbf{x}) O(\mathbf{x}_0) f(\mathbf{x}, \tau). \end{aligned} \quad (3)$$

The right side in the last line is easily stochastically estimated in the QMC simulation by allowing walkers to inherit the observable values at previous times and averaging over all configurations $\mathbf{x}$ present at imaginary time τ that were in $\mathbf{x}_0$ at the earlier time $\tau = 0$.

Now, for an odd operator, such that $\hat{P} \hat{O} \hat{P}^\dagger = -\hat{O}$, one can write the expansions $\hat{O}|\psi_g\rangle = \sum_{n_o} d_{n_o} |\psi_{n_o}\rangle$ and $\hat{O}|\psi_0\rangle = \sum_{n_o} c_{n_o} |\psi_{n_o}\rangle$, where the index $n_o = 0, 1, \dots$ labels the odd eigenstates. Using these expansions, the matrix element can be rewritten as

$$\begin{aligned} \langle \psi_g | \hat{O} \hat{O}(-\tau) | \psi_0 \rangle &= \sum_{n_o} c_{n_o} d_{n_o}^* e^{-(E_{n_o} - E_0)\tau} \\ &\simeq c_{n_o} d_{n_o}^* e^{-(E_{o,0} - E_0)\tau}, \end{aligned}$$

where the last result holds for long enough τ so that higher (odd) excited states do not contribute. Hence, the long-time exponential decay of the matrix element allows one to determine the lowest odd energy gap $\Delta_{o,1} = E_{o,0} - E_0$.

Importantly, the choice of $\psi_g(\mathbf{x})$ does not affect the gap estimate, meaning that we have derived a pure estimator. Yet, a good approximation $\psi_g(\mathbf{x}) \simeq \psi_0(\mathbf{x})$ improves the statistics and eliminates the walker population control bias [23,24]. Here, we consider neural quantum states in the form of restricted Boltzmann machines [25] featuring (unless otherwise specified) $N_h = N$ hidden neurons. These states are optimized via variational energy minimization using the NetKet library [26,27]. Some metrics on their accuracy are reported in Supplemental Material (SM) [28].

An even operator, such that $\hat{P} \hat{O} \hat{P}^\dagger = \hat{O}$, allows for extracting the even gap $\Delta_{e,0} = E_{e,1} - E_0$ from the exponential decay to a finite asymptotic value. Demonstrative benchmarks for both odd and even operators are provided in SM [28]. Details on how we optimize the choice of the operators and on the fitting protocols are reported in the End Matter and in SM. Estimating $\Delta_{e,0}$ via PQMC simulations turns out to be more computationally demanding than $\Delta_{o,0}$, since the successive gap $\Delta_{e,1}$ is often of similar (large) magnitude. This implies that higher-level contributions to the imaginary-time correlation vanish in regimes where the signal-to-noise ratio is small. Hence, in this case we resort to a second computational approach, namely, to exact eigenvalue calculations based on the Lanczos algorithm. Specifically, we adopt the implementation provided by the sparse linear algebra library Cupy [35] for graphic processing units, exploiting the QuSpin library for efficient handling of large sparse Hamiltonian matrices [36]. Because of parity symmetry, block diagonalization of the odd and even sectors is possible. Conversely, for most disorder realizations one has $\Delta_{o,0} \ll \Delta_{o,1}$, allowing for a computationally efficient estimation of the odd gap via PQMC simulations for sufficiently many disorder realizations and sizes $N \lesssim 125$. Below we discuss results obtained via either PQMC or eigenvalue-solver calculations.

Results—A preliminary analysis based on the exact eigenvalue solver reveals that, for the quantum spin models addressed in this Letter, in the critical regime the lowest odd energy gap $\Delta_o \equiv \Delta_{o,0}$ is smaller than the lowest even energy gap $\Delta_e \equiv \Delta_{e,0}$, for all considered disorder realizations. Hence, hereafter, the smallest energy gap Δ , which is the main focus of this Letter, is identified with the former, $\Delta \equiv \Delta_o$. Hence, an odd operator $\hat{O}$ is adopted in the PQMC gap estimator described above.

We first focus on the Gaussian 2D-EA model, tuning the transverse field at $\Gamma = 1.98$, corresponding to the estimated critical point of the spin-glass quantum phase transition $\Gamma = 1.98(7)$ [37]. Since the annealing times scale with the square of the inverse gap η , it is interesting to analyze the distribution of this latter quantity across disorder realizations. The sample size ranges from $N_r \approx 5000$ disorder realizations in the case of eigenvalue-solver calculations for $N \approx 20$ to $N_r = 1000$ in PQMC simulations for $N \simeq 120$.




FIG. 1. Complementary empirical distribution function $1 - F(\eta)$ of the inverse odd gap $\eta = 1/\Delta$ of the Gaussian 2D-EA model at the spin-glass quantum phase transition. The various symbols correspond to the lattice sizes $L = \sqrt{N}$ reported in the legend. The three vertical lines indicate the smallest η accounted for in the Hill estimator for $L = 11$, for three sample fractions: $\kappa = 0.8$, $\kappa = 0.3$, and $\kappa = 0.15$ (from left to right).

Figure 1 shows the complementary empirical distribution function $1 - F(\eta)$, namely, the estimate of the probability of finding inverse gaps larger than a given threshold η . As expected, this probability decreases with η , but the slope rapidly diminishes with the lattice length L . The small fluctuations around the general trend are attributed to the large but finite number of realizations. To quantify this trend, we consider the Hill estimator for the tail index α , defined as

$$\frac{1}{\alpha} = \frac{1}{k} \sum_{i=1}^k \ln \left(\frac{\eta_{(N_r-i+1)}}{\eta_{(N_r-k)}} \right), \quad (4)$$

where the N_r realizations $\eta_1, \eta_2, \dots, \eta_{N_r}$ are arranged in ascending order, i.e., $\eta_{N_r} > \eta_{N_r-1} > \dots$, and the $k < N_r$ largest values are considered. Note that this estimator assumes a power-law tail of the probability density function of the form $f(\eta) = F'(\eta) \propto \eta^{-(\alpha+1)}$ for large η . We first consider a rather large sample fraction $\kappa = 0.80 \simeq k/N_r$, such that the whole bulk regime of the function $1 - F(\eta)$ —beyond the shoulder at small η —is accounted for. As shown in Fig. 2, the resulting tail index is relatively small, namely $\alpha \approx 2$, already for small sizes $L \simeq 4$, denoting a distribution with a rather fat tail. Notably, α further decreases with the system size L . In fact, if an empirical power-law fitting function is assumed, the threshold $\alpha = 1$, below which the disorder ensemble average $[\eta]$ is diverging, is crossed for finite sizes $L \simeq 12$. The extrapolation to $L \rightarrow \infty$ is sizably below the $\alpha = 1$ threshold. The same conclusion is reached also by directly performing power-law fits in the bulk region of $1 - F(\eta)$ (see End Matter). This finding is




FIG. 2. Hill estimator for the tail index α [see Eq. (4)] for the Gaussian 2D-EA Hamiltonian as a function of the lattice size L . The three symbols correspond to the sampled fractions κ reported in the legend. The continuous curves represent power-law fitting functions. The error bars are determined via a bootstrap method with $N_{\text{boot}} = 2000$ resamples. The arrows on the right point to the extrapolation $L \rightarrow \infty$. The two horizontal straight lines indicate the thresholds for finite mean and variance.

consistent with the finite-temperature PIMC simulations of the binary 2D-EA model of Ref. [18]. For the feasible sizes, the PQMC simulations also reach the smallest gaps, allowing us to focus in more detail on the large η regime of the empirical distribution function, as well as to assess the possible effect of the sample size. Considering smaller sample fractions $\kappa \lesssim 0.3$, such that only the large η regime is accounted for, the extrapolated tail index appears to saturate to somewhat larger values, in the range $1 \leq \alpha \leq 2$. This range corresponds to a distribution with infinite variance but finite average. Notice that an infinite variance precludes estimating the distribution average from a finite sample. Concerning the potential efficiency of quantum annealers, tail indices $\alpha \lesssim 2$ are problematic, denoting a divergent average required annealing time.

Since typical combinatorial optimization problems often lead to a much denser connectivity than the four nearest neighbors of the square lattice, it makes sense to address the opposite regime, considering the all-to-all connectivity of the quantum SK Hamiltonian. Its inverse-gap distribution at the spin-glass quantum critical point $\Gamma \approx 1.5$ [38–43] is analyzed in Fig. 3. In this case, the decay rate of $1 - F(\eta)$ at large η is less sensitive to the system size N . Indeed, for sizes $N \rightarrow \infty$, the estimated tail index saturates well above the threshold $\alpha = 2$ (see Fig. 4). Hence, the disorder averaged inverse gap $[\eta]$ is well defined for all finite sizes, and a stochastic estimate must converge for large sample size given that the variance also is finite. The scaling of $[\eta]$ with N is analyzed in SM [28]. In Fig. 5, the more common scaling of the disorder averaged gap $[\Delta]$ versus N is shown.




FIG. 3. Complementary empirical distribution function $1 - F(\eta)$ for the SK model. Symbols and lines are defined as in Fig. 1.

The error bars (often smaller than the symbol size) represent the estimated standard deviation of the average over disorder realizations. Importantly, it is found that at the critical point a power-law fit $[\Delta] \propto N^{-\theta}$ with $\theta = 0.32(1)$ accurately describes both the eigenvalue-solver and the PQMC data. The systematic uncertainty due to the finite lattice sizes are estimated considering three restricted fitting windows—removing three small-size data, three large-size data, or both—finding variations consistent with the reported statistical error bar. Furthermore, to further account for the finite number of realizations, the fitting is repeated on datasets of gaps averaged over fivefold splittings of the realization ensemble, obtaining variations compatible with the statistical uncertainty. It is worth mentioning that a $1/3$ exponent was previously conjectured, arguing that a result for the two-pattern Gaussian Hopfield model [22] should apply also to other mean-field models. Instead, another study based on the Hartree-Fock theory predicted $\theta \approx 0.66$ [44] for the quantum SK model. Our unbiased numerics put the prediction $\theta \approx 1/3$ on firm ground. It is also worth mentioning that some studies addressed a different but related problem, namely, the classical SK model and the scaling of the free-energy barrier height [45–47], finding a power-law growth with exponent $1/3$.

The inset of Fig. 5 compares the scaling exponents for odd and even gaps. For the latter, and for the odd gap at transverse fields other than $\Gamma = 1.5$, eigenvalue-solver data up to sizes $N = 26$ are considered. Assuming a power-law scaling even slightly away from the critical point, we find that the odd gap exponent is larger than the even gap exponent (slightly) within the spin-glass phase $\Gamma \lesssim 1.5$ and vice versa in the paramagnetic phase. Interestingly, the two exponents agree within statistical uncertainties at the critical point. This indicates that both symmetry-changing and symmetry-restricted excitations are suppressed by a power-law scaling gap.




FIG. 4. Hill estimator α as a function of the number of spins N for the SK Hamiltonian. Symbols and lines are defined as in Fig. 3.

Discussion—In summary, our investigation of the quantum 2D-EA model demonstrates that the superalgebraic scaling of the symmetry unrestricted (i.e., odd) energy gap—previously observed for binary disorder via finite-temperature PIMC simulations—is a universal feature that persists for Gaussian couplings and continuous imaginary time. By contrast, the quantum SK model exhibits a more benign power-law scaling with a favorable exponent $\theta \simeq 1/3$, which remains stable at least up to $N \simeq 125$. This finding offers a promising outlook for quantum annealing of combinatorial optimization problems featuring a dense connectivity. While our focus has been on the quantum critical point, we acknowledge that the critical region is not the sole bottleneck in quantum optimization. Exponentially small gaps can emerge deep within the spin-glass phase [22]. Future work should examine finite-dimensional systems with a notion of locality, bridging the gap between the 2D nearest-neighbor lattice and the all-to-all connectivity. Prominent examples are power-law random interactions relevant to Rydberg-atom arrays [48] or trapped ions [49].

Finally, this Letter relied on a continuous-time PQMC algorithm, which we augmented introducing a pure gap estimator. As we have shown, this estimator is independent of the guiding wave function. It provides a robust and versatile tool, which we expect will be instrumental in resolving spectral properties across diverse fields, from frustrated magnetism to quantum chemistry.

Acknowledgments—We thank M. Murdaca and E. Vitali for assistance on the LUMI supercomputer. Support from the following sources is acknowledged: PNRR MUR Project No. PE0000023-NQSTI; PRIN 2022 MUR Project “Hybrid algorithms for quantum simulators,” No. 2022H77XB7; PRIN-PNRR 2022 MUR Project “UEFA,” No. P2022NMB AJ; National Centre for HPC, Big Data and Quantum Computing (ICSC), CN00000013




FIG. 5. Odd energy gap Δ_o , averaged over disorder realizations, as a function of the number of spins N , for the SK model. The black line represents a power-law fit $\Delta = C/N^\theta$, providing $C = 3.06(3)$ and $\theta = 0.32(1)$. Inset: scaling exponent θ as a function of the transverse field Γ for the odd gap Δ_o and the even gap Δ_e . The dashed vertical line shows the position of the critical point.

Spoke 7—Materials & Molecular Sciences; CINECA Award No. INF25_lincoln; and EuroHPC Joint Undertaking for awarding access to the EuroHPC supercomputer LUMI, hosted by CSC (Finland), through EuroHPC Development and Regular Access calls.

Data availability—The data presented in this article are available online [50].

- [1] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [2] K. Takahashi and Y. Matsuda, Energy-gap analysis of quantum spin-glass transitions at zero temperature, *J. Phys. Conf. Ser.* **233**, 012008 (2010).
- [3] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and F. Zamponi, The quantum adiabatic algorithm applied to random optimization problems: The quantum spin glass perspective, *Phys. Rep.* **523**, 127 (2013).
- [4] M. Tikhanovskaya, S. Sachdev, and R. Samajdar, Equilibrium dynamics of infinite-range quantum spin glasses in a field, *PRX Quantum* **5**, 020313 (2024).
- [5] A. P. Young, S. Knysh, and V. N. Smelyanskiy, Size dependence of the minimum excitation gap in the quantum adiabatic algorithm, *Phys. Rev. Lett.* **101**, 170503 (2008).
- [6] A. Ambainis, On physical problems that are slightly more difficult than QMA, in *Proceedings of the 2014 IEEE 29th Conference on Computational Complexity (IEEE, Vancouver, BC, Canada, 2014)*, pp. 32–43, 10.1109/CCC.2014.12.
- [7] S. Gharibian, Y. Huang, Z. Landau, and S. W. Shin, Quantum Hamiltonian complexity, *Found. Trends Theor. Comput. Sci.* **10**, 159 (2014).
- [8] S. Gharibian and J. Yirka, The complexity of simulating local measurements on quantum systems, *Quantum* **3**, 189 (2019).

- [9] H. Rieger and A. P. Young, Zero-temperature quantum phase transition of a two-dimensional Ising spin glass, *Phys. Rev. Lett.* **72**, 4141 (1994).
- [10] M. Guo, R. N. Bhatt, and D. A. Huse, Quantum critical behavior of a three-dimensional Ising spin glass in a transverse magnetic field, *Phys. Rev. Lett.* **72**, 4137 (1994).
- [11] R. R. P. Singh and A. P. Young, Critical and Griffiths-McCoy singularities in quantum Ising spin glasses on d -dimensional hypercubic lattices: A series expansion study, *Phys. Rev. E* **96**, 022139 (2017).
- [12] R. Miyazaki and H. Nishimori, Real-space renormalization-group approach to the random transverse-field Ising model in finite dimensions, *Phys. Rev. E* **87**, 032154 (2013).
- [13] D. A. Matoz-Fernandez and F. Romá, Unconventional critical activated scaling of two-dimensional quantum spin glasses, *Phys. Rev. B* **94**, 024201 (2016).
- [14] I. Hen, Excitation gap from optimized correlation functions in quantum Monte Carlo simulations, *Phys. Rev. E* **85**, 036705 (2012).
- [15] K. Choo, G. Carleo, N. Regnault, and T. Neupert, Symmetries and many-body excitations with neural-network quantum states, *Phys. Rev. Lett.* **121**, 167204 (2018).
- [16] N. Reinić, D. Jaschke, D. Wanisch, P. Silvi, and S. Montangero, Finite-temperature Rydberg arrays: Quantum phases and entanglement characterization, *Phys. Rev. Res.* **6**, 033322 (2024).
- [17] M. Schmitt, M. M. Rams, J. Dziarmaga, M. Heyl, and W. H. Zurek, Quantum phase transition dynamics in the two-dimensional transverse-field Ising model, *Sci. Adv.* **8**, eabl6850 (2022).
- [18] M. Bernaschi, I. González-Adalid Pemartín, V. Martín-Mayor, and G. Parisi, The quantum transition of the two-dimensional Ising spin glass, *Nature (London)* **631**, 749 (2024).
- [19] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 74887 (2014).
- [20] D. Venturelli and A. Kondratyev, Reverse quantum annealing approach to portfolio optimization problems, *Quantum Mach. Intell.* **1**, 17 (2019).
- [21] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, England, 2017).
- [22] S. Knysh, Computational bottlenecks of quantum annealing, *arXiv:1506.08608*.
- [23] S. Pilati and P. Pieri, Simulating disordered quantum Ising chains via dense and sparse restricted Boltzmann machines, *Phys. Rev. E* **101**, 063308 (2020).
- [24] S. Pilati, E. M. Inack, and P. Pieri, Self-learning projective quantum Monte Carlo simulations guided by restricted Boltzmann machines, *Phys. Rev. E* **100**, 043301 (2019).
- [25] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [26] G. Carleo, K. Choo, D. Hofmann, J. E. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, M. Mauri, G. Mazzola, C. B. Mendl, E. van Nieuwenburg, O. O'Reilly, H. Thévéniaut, G. Torlai, F. Vicentini, and A. Wietek, NetKet: A machine learning toolkit for many-body quantum systems, *SoftwareX* **10**, 100311 (2019).
- [27] F. Vicentini, D. Hofmann, A. Szabó, D. Wu, C. Roth, C. Giuliani, G. Pescia, J. Nys, V. Vargas-Calderón, N. Astrakhantsev, and G. Carleo, NetKet 3: Machine learning toolbox for many-body quantum systems, *SciPost Phys. Codebases* **7** (2022).
- [28] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/fm8m-kz13> for additional computational details, benchmark data, and scaling analyses, which includes Refs. [29–34].
- [29] P. Pfeuty, The one-dimensional Ising model with a transverse field, *Ann. Phys. (N.Y.)* **57**, 79 (1970).
- [30] G. G. Cabrera and R. Jullien, Role of boundary conditions in the finite-size Ising model, *Phys. Rev. B* **35**, 7062 (1987).
- [31] A. P. Young and H. Rieger, Numerical study of the random transverse-field Ising spin chain, *Phys. Rev. B* **53**, 8486 (1996).
- [32] G. B. Mbeng, A. Russomanno, and G. E. Santoro, The quantum Ising chain for beginners, *SciPost Phys. Lect. Notes* **82** (2024).
- [33] G. X. Tang, J. Z. Zhuang, L. M. Duan, and Y. K. Wu, Stretched exponential scaling of parity-restricted energy gaps in a random transverse-field Ising model, *arXiv:2512.03526*.
- [34] S. Suzuki, J.-i. Inoue, and B. K. Chakrabarti, *Quantum Ising Phases and Transitions in Transverse Ising Models* (Springer, Berlin, Heidelberg, 2012), Vol. 862.
- [35] R. Nishino and S. H. C. Loomis, CuPy: A NumPy-compatible library for NVIDIA GPU calculations, in *31st Conference on Neural Information Processing Systems (NIPS, Long Beach, CA, USA, 2017)*, Vol. 151, http://learningsys.org/nips17/assets/papers/paper_16.pdf.
- [36] P. Weinberg and M. Bukov, QuSpin: A Python package for dynamics and exact diagonalisation of quantum many body systems. Part II: Bosons, fermions and higher spins, *SciPost Phys.* **7**, 020 (2019).
- [37] L. Brodoloni and S. Pilati, Zero-temperature Monte Carlo simulations of two-dimensional quantum spin glasses guided by neural network states, *Phys. Rev. E* **110**, 065305 (2024).
- [38] J. Miller and D. A. Huse, Zero-temperature critical behavior of the infinite-range quantum Ising spin glass, *Phys. Rev. Lett.* **70**, 3147 (1993).
- [39] D. Lancaster and F. Ritort, Solving the Schrödinger equation for the Sherrington—Kirkpatrick model in a transverse field, *J. Phys. A Math. Gen.* **30**, L41 (1997).
- [40] A. Das and B. K. Chakrabarti, Reaching the ground state of a quantum spin glass using a zero-temperature quantum Monte Carlo method, *Phys. Rev. E* **78**, 061121 (2008).
- [41] A. P. Young, Stability of the quantum Sherrington-Kirkpatrick spin glass model, *Phys. Rev. E* **96**, 032112 (2017).
- [42] A. Kiss, G. Zaránd, and I. Lovas, Complete replica solution for the transverse field Sherrington-Kirkpatrick spin glass model with continuous-time quantum Monte Carlo method, *Phys. Rev. B* **109**, 024431 (2024).
- [43] P. M. Schindler, T. Guaita, T. Shi, E. Demler, and J. I. Cirac, Variational ansatz for the ground state of the quantum Sherrington-Kirkpatrick model, *Phys. Rev. Lett.* **129**, 220401 (2022).
- [44] Y. W. Koh, Effects of low-lying excitations on ground-state energy and energy gap of the Sherrington-Kirkpatrick model in a transverse field, *Phys. Rev. B* **93**, 134202 (2016).
- [45] E. Bittner and W. Janke, Free-energy barriers in the Sherrington-Kirkpatrick model, *Europhys. Lett.* **74**, 195 (2006).

- [46] T. Aspelmeier and M. A. Moore, Free-energy barriers in the Sherrington-Kirkpatrick model, *Phys. Rev. E* **105**, 034138 (2022).
- [47] M. Bernaschi, A. Billoire, A. Maiorano, G. Parisi, and F. Ricci-Tersenghi, Strong ergodicity breaking in aging of mean-field spin glasses, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 17522 (2020).
- [48] L. Brodolini, J. Vovrosh, S. Julià-Farré, A. Dauphin, and S. Pilati, Spin-glass quantum phase transition in amorphous arrays of Rydberg atoms, *Phys. Rev. A* **112**, L051303 (2025).
- [49] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe, Many-body

- localization in a quantum simulator with programmable random disorder, *Nat. Phys.* **12**, 907 (2016).
- [50] L. Brodolini, G. Astrakharchik, S. Giorgini, and S. Pilati, Data used in “Energy gap of quantum spin glasses: a projection quantum Monte Carlo study”, [10.5281/zenodo.18705682](https://zenodo.org/record/18705682) (2026).
- [51] C. Monthus and T. Garel, Typical versus averaged overlap distribution in spin glasses: Evidence for droplet scaling theory, *Phys. Rev. B* **88**, 134204 (2013).
- [52] J. Casulleras and J. Boronat, Unbiased estimators in quantum Monte Carlo methods: Application to liquid ^4He , *Phys. Rev. B* **52**, 3654 (1995).

End Matter

Comparing the Gaussian and binary 2D Edwards-Anderson Hamiltonians—The energy gap is sensitive to all details of the corresponding Hamiltonian. On the other hand, properties such as the dynamical critical exponent z , defined from the putative power-law scaling of the disorder averaged odd gap $[\Delta_o] \propto L^{-z}$ at the critical point, is expected to depend only on the corresponding universality class and not on, e.g., short-range details of the Hamiltonian. This motivates us to compare the scaling obtained here for the quantum 2D-EA model with Gaussian couplings, with the corresponding results from Ref. [18] for the binary case where $J_{ij} = \pm J$ with 50% probability. As shown in Fig. 6, where an appropriate global multiplicative factor is applied to the binary-coupling results, the scalings of the datasets corresponding to Gaussian [$J_{ij} \propto \mathcal{N}(0, J^2)$]




FIG. 6. Disorder averaged energy odd gap $[\Delta_o]$ for the 2D-EA Hamiltonian, as a function of the lattice length L . The two datasets correspond to eigenvalue-solver and continuous-time PQMC calculations for Gaussian couplings $J_{ij} \propto \mathcal{N}(0, J^2)$ and to the discrete-time PIMC simulations for binary couplings $J_{ij} = \pm J$ from Ref. [18]. The deviations from the power-law fit $[\Delta_o] \propto L^{-z}$, performed on sizes up to $L = 7$, are consistent with the superalgebraic scaling discussed in the main text. The inset shows the scaling of the typical average Δ_o^{typ} .

and binary ($J_{ij} = \pm J$) are evidently consistent. At large L both datasets display sizable deviations from the power-law scaling. On the one hand, the binary-coupling results, obtained through finite-temperature PIMC simulations with discrete imaginary time, reach larger system sizes. On the other hand, the continuous-time PQMC simulations are not affected by the Trotter discretization, denoting its irrelevance for the scaling of $[\Delta_o]$. Notice that, for the binary couplings, few disorder instances were excluded—mostly for the largest sizes—due to an excessively small gap to be addressed via the finite temperature algorithm. The cross validation suggests that these few gaps do not significantly contribute to the scaling of the disorder average. Sizable deviations from the power-law scaling are also found considering the typical average, defined as $\Delta_o^{\text{typ}} = \exp[\ln(\Delta_o)]$ [51]; again, square parentheses denote the average over disorder realizations. The corresponding data for Gaussian and binary couplings are shown in the inset of Fig. 6.

To deepen the comparison between Gaussian and binary couplings, we inspect the width of the corresponding distributions of the inverse gaps $\eta = 1/\Delta_o$. Following Ref. [18], a power-law fit is performed on the complementary cumulative distribution $1 - F(\eta) = B/\eta^m$, excluding both the small η regime and the extreme-tail regime at large η . Specifically, the fitting range is chosen by minimizing the reduced chi-squared value. The fitted exponents m for different sizes L are shown in Fig. 7. The uncertainty due to the choice of fitting window is estimated by shifting its position by 10% in both directions, leading to a rather conservative estimate of the error bar. By extrapolating the Gaussian result to larger lattice sizes via an empirical power-law fitting function, the binary coupling result at $L = 24$ reported in Ref. [18], namely $m \simeq 0.8$, is recovered. It is worth stressing that decay rates $m < 1$ correspond to a distribution with infinite average. This analysis, which focuses on the bulk of the cumulative distribution function of η , confirms the conclusions drawn from the Hill estimator with sample fraction $\kappa = 0.8$. It is worth




FIG. 7. Power m of the decay of the complementary cumulative distribution function $1 - F(\eta) = B/\eta^m$, with B and m fitting parameters, as a function of the lattice length L of the 2D-EA model at the spin-glass quantum phase transition. The curve is an empirical power-law fit on the Gaussian $[\mathcal{N}(0, J^2)]$ results. The extrapolation agrees with the binary couplings ($50\% \pm J$) result at $L = 24$.

mentioning that, in the case of binary couplings, the small uncertainty in the critical transverse field was found not to affect the scaling analyses. The universality of the scaling in terms of couplings distribution indicates that the (comparably small) uncertainty is irrelevant also for the Gaussian distribution.

Extracting the gap from PQMC imaginary-time correlations—The PQMC simulations provide the imaginary-time correlation function

$$C(\tau) = \langle \hat{O}(0) \hat{O}(-\tau) \rangle, \quad (\text{B1})$$

where the brackets $\langle \rangle$ denote the Monte Carlo average defined in Eq. (3). When the observable $\hat{O}$ is odd under parity symmetry, in the long imaginary-time limit $\tau \rightarrow \infty$ the correlation decays exponentially as

$$C(\tau) \propto e^{-\Delta_o \tau}, \quad (\text{B2})$$

allowing us to extract the smallest odd gap $\Delta_o = E_{o,0} - E_0$ from an exponential fit at large τ . The selection of the fitting window has to be automatized, as detailed below. It is worth noting that, by adopting the standard forward-walking approach [52], the whole profile of the pure imaginary time correlation could be determined from

$$\langle \psi_0 | \hat{O}(0) \hat{O}(-\tau) | \psi_0 \rangle = \langle \hat{O}(-\Delta\tau) \hat{O}(-\Delta\tau - \tau) \rangle, \quad (\text{B3})$$

where the time interval $\Delta\tau$ is chosen large enough to eliminate the dependence on $\psi_g(\mathbf{x})$, which would

otherwise affect the behavior of the correlation function $C(\tau)$ for small values of τ .

With an even-parity operator $\hat{O}$, for which the ground-state expectation value is finite, the imaginary-time correlation function converges to a finite constant at large τ . To isolate the exponential decay associated with the lowest even-parity excitation, we define the offset correlation

$$\tilde{C}(\tau) = C(\tau) - \langle \hat{O}(-\tau_f) \rangle \langle \hat{O}(0) \rangle, \quad (\text{B4})$$

where the imaginary time τ_f is chosen to be sufficiently long to reach the uncorrelated plateau regime. The expectation values $\langle \hat{O}(-\tau_f) \rangle$ and $\langle \hat{O}(0) \rangle$ are independently evaluated in the PQMC simulation. The even gap $\Delta_e = E_{e,1} - E_0$ is hence extracted from an exponential fit $\tilde{C}(\tau) \propto e^{-\Delta_e \tau}$ in the large τ regime.

Choice of operators: For the odd gap we consider single-spin operators $\hat{O}_i = \sigma_i^z$. In an initial short stint of the simulation, all N sites are considered, and the site index i providing the largest signal at a suitably chosen imaginary time is selected on-the-fly for the remainder of the simulation. An illustrative example of the selection process is shown in SM. For even parity, we consider two-site operators of the form $\hat{O} = \sigma_i^z \sigma_j^z$. Again, the optimal pair i, j is selected on-the-fly from a batch of 10 randomly chosen pairs after an initial short simulation stint. Illustrative examples of the odd and even correlation functions for one random instance of the quantum SK model are shown in Fig. 8.

Selection of the fitting window: The imaginary-time correlations $[C(\tau)$ or $\tilde{C}(\tau)$ for odd or even operators, respectively] decay exponentially only after an initial transient at small τ in which higher excitations still contribute. On the other hand, at very long imaginary




FIG. 8. Odd and (offset) even imaginary-time correlation functions $C(\tau)$ and $\tilde{C}(\tau)$ for a random realization of the quantum SK model at transverse field $\Gamma = 1.5$ and system size $N = 100$. The bold segment indicates the asymptotic exponential scaling from which the lowest odd Δ_o is extracted.

times the signal-to-noise ratio deteriorates, eventually making the signal comparable with its statistical uncertainty $\sigma(\tau)$. To isolate the region dominated by a single exponential decay, we restrict the fitting window to $\tau \in [\tau_{\min}, \tau_{\max}]$. $\tau_{\max}$ is tuned so to exclude data where the signal-to-noise falls below a predefined threshold, typically $C(\tau)/\sigma(\tau) \approx 8$ [$\tilde{C}(\tau)$ is used for even gaps]. To

select the left extreme $\tau_{\min}$, we increase it until the fitted gap converges (from above) to a plateau value. The plateau region signals that the single-exponential assumption is valid. In practice, the procedure stops when the fitted gap remains constant within statistical uncertainty for consecutive windows or if it starts increasing.