

Real-time simulations of quantum spin chains: Density of states and reweighting approaches

Pavel Buividovich and Johann Ostmeyer ^{*}*Department of Mathematical Sciences, University of Liverpool, Liverpool L69 3BX, United Kingdom*

(Received 5 October 2022; revised 12 December 2022; accepted 21 December 2022; published 10 January 2023)

We put the density of states (DoS) approach to Monte Carlo (MC) simulations under a stress test by applying it to a physical problem with the worst possible sign problem: the real-time evolution of a nonintegrable quantum spin chain. Benchmarks against numerical exact diagonalization and stochastic reweighting are presented. Both MC methods, the DoS approach and reweighting, allow for simulations of spin chains as long as $L = 40$, far beyond exact diagonalizability, though only for short evolution times $t \lesssim 1$. We identify discontinuities of the DoS as one of the key problems in the MC simulations and propose calculating some of the dominant contributions analytically, increasing the precision of our simulations by several orders of magnitude. Even after these improvements, the DoS is found highly nonsmooth, and therefore, the DoS approach cannot outperform reweighting. We prove this implication theoretically and provide numerical evidence, concluding that the DoS approach is not well suited for quantum real-time simulations with discrete degrees of freedom.

DOI: [10.1103/PhysRevB.107.024302](https://doi.org/10.1103/PhysRevB.107.024302)

I. INTRODUCTION

The simulation of the real-time evolution of a many-body quantum system up to some finite physical time t is a problem of utmost importance across many areas of physics. However, it is an NP-hard problem which, in general, can only be solved on a classical computer in a time

$$T_{\text{sim}} \sim e^{\alpha(t)n_{\text{dof}}} \quad (1)$$

that grows exponentially with the number of degrees of freedom (dof) n_{dof} in a system. Polynomial-time solutions are only possible on a quantum computer [1,2]. With the current state of technology, only relatively small quantum systems can be reliably simulated on quantum computers. Therefore, at least for some time, we must rely on classical computers to simulate the real-time evolution of quantum systems. While we certainly cannot simulate the real-time evolution in polynomial time on a classical computer, we can try to reduce the coefficient $\alpha(t)$ in the exponent in Eq. (1).

Conceptually, the simplest approach to simulate real-time evolution is to find all eigenstates of the Hamiltonian using numerical exact diagonalization (ED). This approach works well for systems with finite-dimensional Hilbert spaces, such as quantum spin chains, and allows us to calculate any real-time evolution [3] in a straightforward way. For a chain of L spin- $\frac{1}{2}$ dof, such as Ising spins, the Hilbert space dimension is $N = 2^L$. Correspondingly, the computational cost of the plain vanilla ED scales as $\mathcal{O}(N^3) = \mathcal{O}(2^{3L})$. Physical results found in most of the literature scale up to $L \leq 18$, see, e.g., Refs. [4–7]. More advanced methods, such as the shift-invert method of ED (SIMED) [8] and polynomially filtered ED (POLFED) [9] allow for extremely efficient ED. POLFED has only $\mathcal{O}(N^2) = \mathcal{O}(2^{2L})$ runtime and $\mathcal{O}(N) = \mathcal{O}(2^L)$ memory

requirement, both up to polynomial corrections in L and for a fixed number of eigenvalues. The largest systems simulated to date have $L = 26$, and the largest ones used for physics beyond mere algorithmic proof of principle have $L = 24$. Unfortunately, the polynomial corrections are not specified in the literature, but they appear to be quite significant if these are the largest sizes possible. In addition, it is only claimed that the limited number of eigenvalues from the bulk yield the correct results. The truncation in the eigenspectrum might well cause relevant artifacts.

Another possible approach is to approximate the quantum evolution $\exp(i\hat{H}t)$ as a product of many exactly treatable factors using a Suzuki-Trotter decomposition. Instead of ED, in this case, we must perform multiple matrix-matrix multiplications, with a computational cost that scales as $\mathcal{O}(N^2) = \mathcal{O}(2^{2L})$ for local Hamiltonians. This allows us to simulate spin chains with $L \leq 24$ [10,11].

An alternative to methods based on calculations in the entire Hilbert space is provided by Monte Carlo (MC) methods, which replace exact summation over all states with importance sampling. For finite-temperature equilibrium partition functions, MC methods typically reduce the exponential scaling in the system size L down to a polynomial one. MC methods can also be used for bosonic systems in a straightforward way, in contrast to ED which necessarily relies on truncations of an infinite-dimensional Hilbert space.

In the absolute majority of cases, real-time evolution problems have path integral representations with oscillatory integrands, either real or complex valued. The most straightforward approach to deal with nonpositive-definite integrands in MC simulations is reweighting, whereby one performs importance sampling with the weight proportional to the absolute value of the integrand, and reweights the contribution of any field configuration with the corresponding complex phase of the integrand. However, in this case, the NP-hardness of the real-time evolution problem reveals itself in the

*j.ostmeyer@liverpool.ac.uk

exponentially quick deterioration of statistical power with n_{dof} and the evolution time t . Many contributions with different complex phases cancel each other, and exponentially many MC samples are needed to calculate the path integral with desired precision. This exponential growth of computational complexity is referred to as the *sign problem*.

In the past decade, the density of states (DoS) approach has been advocated as an efficient way to make the sign problem milder [12–15]. Starting from a generic oscillatory path integral representation of a partition function Z of the form:

$$Z = \int \mathcal{D}\phi \exp(-S_R[\phi] + iS_I[\phi]), \quad (2)$$

the basic idea of the DoS approach is to construct a numerical approximation to the DoS, or probability distribution, of the complex phase $S_I[\phi]$:

$$\rho(E) = \int \mathcal{D}\phi \exp(-S_R[\phi])\delta(E - S_I[\phi]). \quad (3)$$

Here, $\int \mathcal{D}\phi$ denotes integration or summation over any continuous or discrete variables ϕ . Knowing $\rho(E)$, which can be determined without encountering a sign problem, we can calculate the full partition function Z as a simple one-dimensional integral over E :

$$Z = \int dE \rho(E) e^{iE}. \quad (4)$$

An important development was the logarithmic linear relaxation (LLR) method [16–18] based on Robbins-Monro (RM) [19] iterations, which allows obtaining high-precision results for the DoS $\rho(E)$ for lattice models with continuous [20,21] as well as discrete [22] dof. Applications include the compact $U(1)$ lattice gauge theory [20], heavy-dense QCD at finite quark density [12,13], finite-density Bose gas [14], finite-density Hubbard model [21], and the Potts model [22]. It is fair to say that, in all cases, the computational complexity of evaluating the original multidimensional oscillatory path integral in Eq. (2) translates into the complexity of evaluating the one-dimensional oscillatory integral in Eq. (4), which is very sensitive to statistical errors in the DoS $\rho(E)$. This problem was tackled in Refs. [12–14,21,22] by constructing various analytic approximations to $\rho(E)$ and using them to calculate the integral. The approximations included polynomial fits and splines. Resummations based on Fourier transforms of $\rho(E)$ were considered in Ref. [21]. This allowed us to obtain results of practical significance for lattice sizes that are intractable with other methods, e.g., the straightforward reweighting.

In this paper, we put the DoS/LLR approach under a stress test by using it to simulate the real-time evolution, for which the partition function in Eq. (2) is strongly dominated by the imaginary part of the action S_I . This is hence a physical problem for which the sign problem is expected to be maximally strong. As a model of choice, we use a nonintegrable quantum Ising chain with quenched disorder [4]. We explicitly quantify the computational cost required to obtain a fixed error of our measurements and compare it with that of the reweighting and ED methods. We find that, without further analytic approximations, the computational cost of evaluating the one-dimensional integral in Eq. (4) of the DoS is not better

than that of simple reweighting. We also critically examine the smoothness of the DoS and find an efficient way to remove some of the dominant discontinuities, which improves the performance of both the DoS and the reweighting approaches by orders of magnitude. This also allows us to construct reasonably good polynomial approximations for the DoS $\rho(E)$; however, they turn out to hardly improve the quality of the results because of the high number of parameters required. Therefore, for our particular system, the DoS approach only outperforms reweighting in very rare cases. Our overall conclusion is that both MC approaches, LLR and reweighting, are advantageous compared with ED when simulating short-time evolutions for large spatial system sizes, even allowing for sizes completely inaccessible to ED. On the other hand, ED is clearly better for long times.

We chose the nonintegrable quantum Ising chain with quenched disorder [4] as our test model because it is one of the simplest nonintegrable models which exhibit ergodic and many-body localization regimes. The Hamiltonian of the model reads, in its most general form,

$$H = -\sigma^z \cdot \mathbb{J} \cdot \sigma^z - \mathfrak{h} \cdot \sigma^x, \quad (5)$$

where $\sigma^{x,z} = (\sigma_1^{x,z}, \dots, \sigma_L^{x,z})$ is the spatial vector of Pauli x - and z -spin matrices, respectively, $\mathbb{J} \in \mathbb{R}^{L \times L}$ is some coupling matrix, and $\mathfrak{h} \in \mathbb{R}^L$ is a vector of local magnetic fields \mathfrak{h}_i .

We work in the ergodic regime, where no symmetry can accidentally make the sign problem milder. Furthermore, the model has a straightforward path integral representation. Its energy spectrum can be easily obtained using exact numerical diagonalization for chain lengths of order $L \lesssim 20$, which allows us to obtain benchmark results for any real-time observables without any statistical errors. In this paper, we focus on one of the simplest real-time quantities, the infinite-temperature spectral form factor (SFF):

$$\mathcal{K}(t) = \langle K(t) \rangle_{\mathbb{J}} = \langle |\text{Tr} e^{iHt}|^2 \rangle_{\mathbb{J}}, \quad (6)$$

where $\langle \cdot \rangle_{\mathbb{J}}$ denotes the average over different disorder realizations, and $U(t)$ is the real-time evolution operator for given disorder \mathbb{J} . Universal late-time features of SFFs, such as the ramp [23–25], have proven useful as probes of quantum chaos in strongly interacting many-body systems. SFFs are useful for distinguishing ergodic and many-body localized regimes [26].

Originally, our hope was to apply the DoS/LLR approach to resolve the ongoing debate about the existence and the properties of the many-body localization phase in quantum spin chains [4–7,10,11,27–33]. The debate is strongly focused on extrapolations toward the thermodynamic limit; therefore, it is useful to devise methods that could work for lattice sizes that are currently inaccessible for numerical ED. While our improvements to the DoS/LLR approach allow us to simulate spin chains of lengths as large as $L = 40$, well beyond the reach of any ED methods, such simulations are bound to early times $t \lesssim 1$. This is by far insufficient to resolve the late-time behavior of SFFs, which starts distinguishing the ergodic and many-body localized regimes at the Heisenberg timescale $t \gtrsim 2^L$. Let us note in passing that the necessity to simulate up to times of order of 2^L is very likely to make the calculation of the SFF an NP-hard problem even for a quantum computer. Indeed, the no fast-forwarding theorem

[34] is likely to hold for our Hamiltonian. Therefore, quantum computers will need an exponentially large computational time $T_{\text{sim}} \gtrsim 2^L$ to simulate physical evolution up to physical times $t \gtrsim 2^L$.

The rest of this paper is structured as follows. We provide the details of the model in Eq. (5) and its classical 1 + 1-dimensional counterpart in Sec. II. Next, we introduce the DoS/LLR algorithm in Sec. III, and we derive how the algorithm can be optimized in Sec. IV. Analytic estimates for the runtimes of all the different approaches are discussed in Sec. V, and finally, the results of numerical experiments are presented in Sec. VI.

II. THE MODEL

In this paper, we follow Ref. [4] and consider the Hamiltonian in Eq. (5) with both the nearest neighbor and next-to-nearest neighbor couplings:

$$\mathbb{J}_{ij} = (\mathbb{J}_0 + \Delta\mathbb{J}_i)\delta_{i,j+1} + \mathbb{J}_2\delta_{i,j+2}, \quad (7)$$

where the nearest neighbor coupling contains quenched disorder $\Delta\mathbb{J}_i$ that is sampled uniformly from $\Delta\mathbb{J}_i \in [-\Delta\mathbb{J}, \Delta\mathbb{J}]$. Periodic boundary conditions (pbc) are applied. The magnetic field term $\mathbb{h}_i = \mathbb{h}_0$ is constant. Throughout this paper, we choose $\mathbb{h}_0 = 0.6$, $\mathbb{J}_2 = 0.3$, $\mathbb{J}_0 = 1$, and $\Delta\mathbb{J} = 1$ deeply in the chaotic regime.

The real-time dynamics of the system as well as the severity of its sign problem are captured in the SFF:

$$\mathcal{K}(t) \equiv \langle K(t) \rangle_{\mathbb{J}} \equiv \langle |\text{Tr}U(t)|^2 \rangle_{\mathbb{J}}, \quad (8)$$

where $\langle \cdot \rangle_{\mathbb{J}}$ denotes the average over different disorder realizations, and $U(t)$ is the real-time evolution operator for given disorder \mathbb{J} :

$$\begin{aligned} U(t) &= e^{-iHt} = \prod_{j=1}^{N_t} e^{-i\delta H} \\ &= \prod_{j=1}^{N_t} e^{i\delta\sigma^z \cdot \mathbb{J} \cdot \sigma^z} e^{i\delta\mathbb{h} \cdot \sigma^x} + \mathcal{O}(\delta), \end{aligned} \quad (9)$$

with the physical evolution time t and the Suzuki-Trotter discretization step size $\delta \equiv t/N_t$. This Suzuki-Trotter decomposition becomes exact in the limit where the number of time steps $N_t \rightarrow \infty$.

The transfer-matrix method allows us to relate the Trotterized version to a two-dimensional system of classical spins [35]:

$$K(t) \propto Z \equiv \sum_{\{s\}} e^{-S(s)}, \quad (10)$$

where we sum over all possible spin $s_{i,k} = \pm 1$ combinations, arriving at the action:

$$S = -i \sum_k \sum_{i,j} s_{i,k} J_{ij} s_{j,k} - \sum_i \sum_k \left(h_i - \frac{\pi}{4} i \right) s_{i,k} s_{i,k+1}, \quad (11)$$

$$J_{ij} := \delta\mathbb{J}_{ij}, \quad h_i := -\frac{1}{2} \log \tan(\delta\mathbb{h}_i), \quad (12)$$

where J and h_i are real constants. See Appendix A for details on the derivation.

Any spin flip results in a change of S by $\lambda(h_i - \frac{\pi}{4}i)$ with $\lambda \in \{\pm 4, 0\}$ and some contributions of order δ . Therefore, both real and imaginary parts are significantly changed even for small δ . Here, $\lambda = \pm 4$ flips the sign of the action.

We split the action into three separate parts:

$$S = iS_I + S_R + i\pi S_S, \quad (13)$$

$$S_I := - \sum_k \sum_{i,j} s_{i,k} J_{ij} s_{j,k}, \quad (14)$$

$$S_R := - \sum_i \sum_k h_i s_{i,k} s_{i,k+1}, \quad (15)$$

$$S_S := \frac{1}{4} \sum_i \sum_k s_{i,k} s_{i,k+1}. \quad (16)$$

Now each of S_I , S_R , and S_S is real. Furthermore, S_S is an integer, and therefore,

$$e^{i\pi S_S} = \pm 1 \quad (17)$$

contributes only a sign to the total measure.

A. Probability weighted DoS

MC-based algorithms can easily implement the probability density induced by S_R , whereas the complex phase and real sign contributions by S_I and S_S , respectively, pose severe issues. As outlined in Eqs. (2)–(4), in this paper, we employ the DoS approach, rewriting the partition sum as

$$\begin{aligned} Z &= \sum_{\{s\}} e^{-S_R} e^{-iS_I} e^{-i\pi S_S} \\ &= \sum_{\{s\}} e^{-S_R} \int_{\mathbb{R}} dE e^{-iE} \delta(E - S_I) \sum_{z=\pm 1} z \delta_{z,e^{-i\pi S_S}} \\ &= \sum_{z=\pm 1} z \int_{\mathbb{R}} dE e^{-iE} \sum_{\{s\}} e^{-S_R} \delta(E - S_I) \delta_{z,e^{-i\pi S_S}} \\ &= \sum_{z=\pm 1} z \int_{\mathbb{R}} dE e^{-iE} \rho_z(E). \end{aligned} \quad (18)$$

In the last step, we defined the DoS $\rho_z(E)$ as

$$\rho_z(E) := \sum_{\{s\}} e^{-S_R} \delta(E - S_I) \delta_{z,e^{-i\pi S_S}}. \quad (19)$$

B. Normalization

Let us define the auxiliary partition sum:

$$Z_R = \sum_{\{s\}} e^{-S_R} = \sum_{z=\pm 1} \int_{\mathbb{R}} dE \rho_z(E), \quad (20)$$

which is a sum of manifestly positive contributions and therefore does not suffer from a sign problem. Note that all terms containing the coupling \mathbb{J} have been dropped, and Z_R therefore completely decouples into its one-dimensional (timelike)

parts, i.e.,

$$Z_R = \prod_i Z_{R,i}, \quad (21)$$

$$\begin{aligned} Z_{R,i} &= \text{Tr}[(T_2)_i^{N_i}] \\ &= 2^{N_i} (\cosh^{N_i} h_i + \sinh^{N_i} h_i), \end{aligned} \quad (22)$$

where the last line is the well-known solution of the one-dimensional classical Ising model and can be obtained by diagonalizing the 2×2 transfer matrices $(T_2)_i$ from Eq. (A7).

In the full simulation, one would usually normalize the sum over the DoS to one so that the MC estimator (denoted by $\langle \cdot \rangle$) of the partition sum becomes

$$\langle Z \rangle = Z_R \left\langle \frac{Z}{Z_R} \right\rangle, \quad (23)$$

where Z_R is known analytically, and $\langle \frac{Z}{Z_R} \rangle$ can be found by evaluating Eq. (18) with the normalized estimator of $\rho_z(E)$. Combining this result with Eq. (A12), we finally obtain a formula for the trace of the quantum mechanical time evolution operator:

$$\text{Tr}U(t) = AZ_R \left\langle \frac{Z}{Z_R} \right\rangle. \quad (24)$$

III. THE ALGORITHMS

The canonical choice and baseline algorithm for MC simulations with a sign problem is reweighting. In this approach, all contributions stemming from the imaginary part of the action are simply considered part of the observable. In this case, this means that configurations are sampled purely from the probability distribution $\mathbb{P}(S_R)$ induced by S_R (using, for instance, local Metropolis-Hastings updates), and the partition sum is estimated by

$$\begin{aligned} \left\langle \frac{Z}{Z_R} \right\rangle &= \frac{\sum_{\{s\}} e^{-S_R} e^{-iS_I} e^{-i\pi S_S}}{\sum_{\{s\}} e^{-S_R}} \\ &= \sum_{\{s\}} \mathbb{P}(S_R) e^{-iS_I} e^{-i\pi S_S} \\ &= \left\langle e^{-iS_I - i\pi S_S} \right\rangle_{S_R}. \end{aligned} \quad (25)$$

Another natural approach is to estimate the DoS ρ prior to determining any observables. This can be done efficiently using the LLR algorithm. The principle idea of LLR was proposed by Wang and Landau [16] and applied to the classical Ising model with uniform coupling in two dimensions. Later on, it was optimized through the introduction of the $1/t$ algorithm [17, 18] in accordance with the theoretically optimal scheme derived by Robbins and Monro [19].

Let us consider a classical Ising system of $N \equiv LN_i$ spins. The auxiliary variable E [typically identified as an energy, see Eq. (19)] is divided into M bins E_i , $i = 1, \dots, M$, so that we discretize the DoS:

$$\rho_i := \int_{\frac{E_i + E_{i-1}}{2}}^{\frac{E_i + E_{i+1}}{2}} dE \rho(E), \quad i = 2, \dots, M-1, \quad (26)$$

and for $i = 1$ ($i = M$), the lower (upper) bound of the integral takes the minimal (maximal) value of E . Of course, the arithmetic mean can be exchanged for any other value between

E_i and its neighbors without changing the general argument. Discretization errors are of order $\mathcal{O}(h^2)$, and further details are discussed in Appendix B.

We write

$$\rho_i \equiv e^{\alpha_i}, \quad (27)$$

and from now, on we aim to estimate α_i as accurately as possible (hence the logarithmic in the name of the method). Starting from uniform initial conditions ($\alpha_i = 0$ for all i), we employ an update scheme of Λ steps in total where, at the k th step, some configuration update $s \mapsto s'$ is proposed and accepted with the probability:

$$\begin{aligned} p_{\text{acc}} &= \frac{e^{-S_R[s']}}{e^{-S_R[s]}} \cdot \frac{\rho_{i(s)}}{\rho_{i(s')}} \\ &= e^{S_R[s] - S_R[s']} e^{\alpha_{i(s)} - \alpha_{i(s')}}, \end{aligned} \quad (28)$$

where $i(s)$ and $i(s')$ are the bins that correspond to the imaginary part of the action S_I on configurations s and s' , respectively. The DoS of the accepted bin, say i without loss of generality, must be updated:

$$\alpha_i \mapsto \alpha_i + \beta_{N,M,\Lambda}(k, i), \quad (29)$$

where $\beta \geq 0$. Without prior knowledge of the DoS, we drop the explicit dependence on i . Furthermore, we know from Ref. [19] that the conditions:

$$\sum_{k=0}^{\infty} \beta_{N,M,\Lambda}(k) = \infty, \quad (30)$$

$$\sum_{k=0}^{\infty} \beta_{N,M,\Lambda}(k)^2 < \infty, \quad (31)$$

must hold [36]. They propose a function asymptotically scaling as $1/k$ since this is the fastest decaying function (up to logarithmic factors) of the required class and therefore promises best convergence for large k .

The coefficients in the explicit form:

$$\beta_{N,M,\Lambda}(k) = \frac{a_{N,M,\Lambda}}{b_{N,M,\Lambda} + k}, \quad (32)$$

are of crucial importance, and we derive in Appendix C why

$$a_{N,M,\Lambda} = \ln 2 \frac{MN}{\ln \frac{\Lambda}{b_{N,M,\Lambda}}}, \quad (33)$$

$$b_{N,M,\Lambda} = 3M \quad (34)$$

are a particularly good choice. We will use them throughout the rest of this paper.

LLR iterations do not form a Markov chain because the accept/reject probability depends on the history of the previous updates. Therefore, statistical errors cannot be calculated via straightforward bootstrap procedures or similar. Instead, every simulation presented hereafter has been repeated $N_{\text{runs}} = 40$ times, and the error has been estimated from the resulting distribution. In practice, LLR simulations are not always stable but are prone to outliers. To counteract this effect, we used the median as the best estimator for observables instead of the more commonly used mean. Similarly, the errors have been approximated using the 16 and 84% quantiles instead of the standard deviation. For a normal distribution, this choice of median and quantiles is fully compatible with

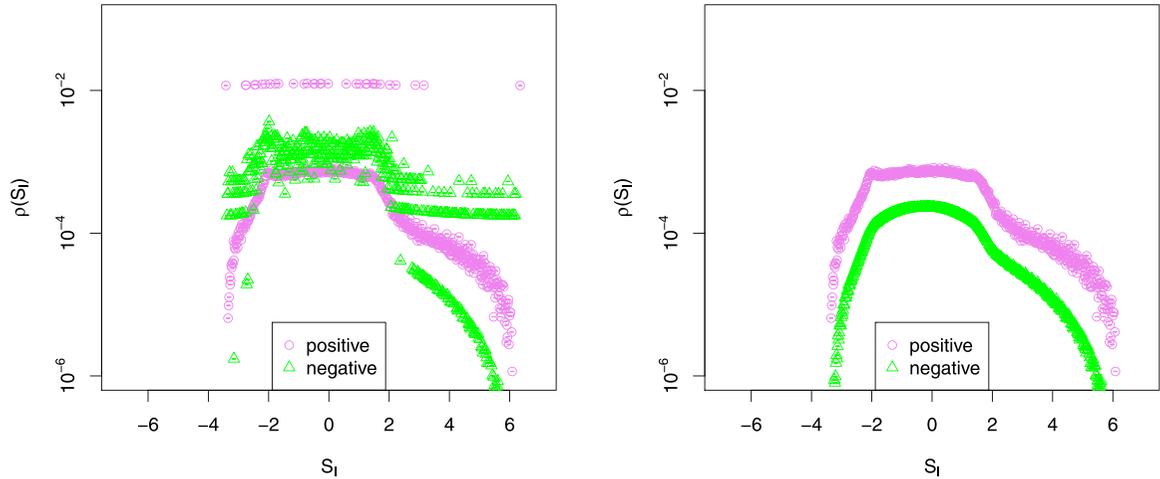


FIG. 1. Densities of states in the even (positive) and odd (negative) parity sectors, respectively, calculated with LLR_0 (left) and LLR_2 (right). The spin chain length is $L = 6$, the evolution time is $t = 1$, the number of Monte Carlo sweeps is $N_s = 10^8$.

the default choice of mean and standard deviation, but for other distributions, it is more stable.

IV. SMALL TRANSVERSE COMPONENT EXPANSION

The representation in Eq. (10) of the real-time partition function in terms of a finite sum over classical Ising spins implies a finite number of states, so that technically speaking the DoS in Eq. (19) is a collection of delta peaks. However, in the presence of quenched disorder, there are in general $\mathcal{O}(2^L)$ delta peaks spread randomly over the range of width $\mathcal{O}(L)$ of possible values of S_I . As a result, the bulk of the DoS approaches the continuous function in the thermodynamic limit $L \rightarrow +\infty$. For finite length L , the apparent continuous nature is a mere consequence of the finite resolution following from the fixed number M of bins defining the DoS, i.e., the true distribution is convoluted with a function that is 1 within a bin and 0 everywhere else.

Looking at a typical distribution of the DoS (see Fig. 1), it appears to be fractured into bands. The top one corresponds to the even sector and consists of a few very narrow peaks. The second (odd) band is closer to a smooth distribution, though it still features some irregularities. An alternating descent of smooth bands with roughly equal step size on a logarithmic scale follows. With the technically available precision, we usually cannot resolve any but the four highest bands.

It is straightforward to interpret these bands, as they directly correspond to a particular parity sector each. Every spin-flip pair in the temporal direction (see Fig. 2) changes this sector. It comes with a sign flip and a Boltzmann suppression, i.e.,

$$S_S \mapsto S_S - 1, \quad (35)$$

$$S_R \mapsto S_R + 4h, \quad (36)$$

$$\Rightarrow p \mapsto -p \left(\tan \frac{t\hbar}{N_t} \right)^2, \quad (37)$$

where $p = e^{-S_R - i\pi S_S}$ is the sign-full probability weight, and $S_{S,R}$ as in Eq. (13). This provides a canonical way of expanding the DoS $\rho_z(E)$ and the partition sum Z in powers of $\tan \frac{t\hbar}{N_t}$, that is, around small times t or transverse fields \hbar .

From now on, we will denote simulations restricted to parity sectors with at least n spin-flip pairs by LLR_n (or REW_n for reweighting). In practice, this means that the acceptance probability in Eq. (28) is set to zero for any proposed configuration with less than n spin-flip pairs. Otherwise, the sampling algorithm proceeds unchanged.

Thus, the plain vanilla algorithm allowing us to visit the full phase space will be denoted LLR_0 , and a lot of our studies will be focused on LLR_2 as a good tradeoff between precision and performance. The remaining contributions from all the configurations with less than n spin-flip pairs are calculated analytically as follows.

A. Leading order: $n = 0$ spin-flip pairs

Without any spin flips in the temporal direction, the phase space is confined to one of the 2^L spatial configurations repeated N_t times. Therefore, the DoS and the classical partition sum reduce to

$$\rho_z(E)|_{n=0} \propto \delta_{z,1} \sum_{\{s\}} \delta \left(E - t \sum_{i,j} s_i \mathbb{J}_{ij} s_j \right), \quad (38)$$

$$\begin{aligned} Z|_{n=0} &= \sum_{\{s\}} e^{iN_t \sum_{i,j} s_i \mathbb{J}_{ij} s_j + LN_t h} \\ &= \left(\tan \frac{t\hbar}{N_t} \right)^{-\frac{1}{2}LN_t} \sum_{\{s\}} e^{it \sum_{i,j} s_i \mathbb{J}_{ij} s_j}, \end{aligned} \quad (39)$$

and thus,

$$\begin{aligned} \text{Tr}U(t)|_{n=0} &= \left(\frac{\sin \delta\hbar \cos \delta\hbar}{\tan \delta\hbar} \right)^{\frac{1}{2}LN_t} \sum_{\{s\}} e^{it \sum_{i,j} s_i \mathbb{J}_{ij} s_j} \\ &= (\cos \delta\hbar)^{LN_t} \sum_{\{s\}} e^{it \sum_{i,j} s_i \mathbb{J}_{ij} s_j}. \end{aligned} \quad (40)$$

The prefactor $Z_0 \equiv (\cos \delta\hbar)^{LN_t} \xrightarrow{\delta \rightarrow 0} 1$ is clearly an artifact from Trotterization, but it cannot be dropped in realistic simulation since it crucially counteracts finite δ effects in the LLR_n simulation itself.

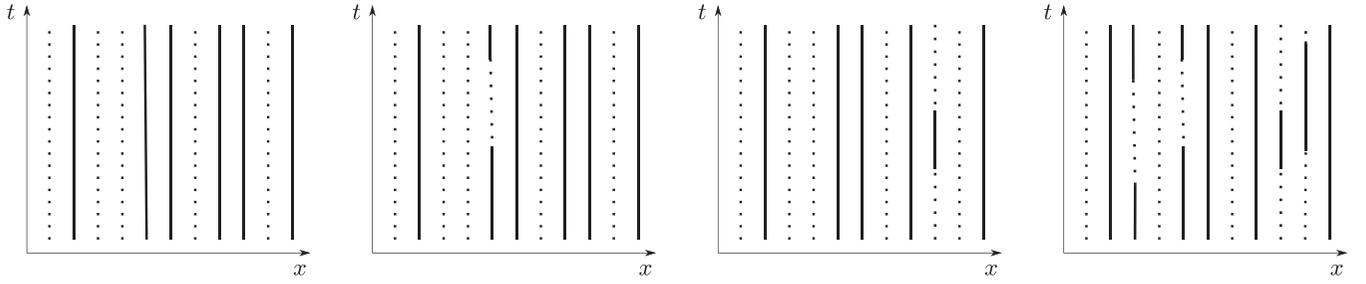


FIG. 2. Visualization of configurations with different numbers n of spin-flip pairs in the time continuum. Periodic boundary conditions are applied, so flips come in pairs only. Without loss of generality the solid lines represent spin-up and the dotted lines spin-down. From left to right: $n = 0$, $n = 1$, $n = 1$, and $n = 4$.

Here, we assume that the couplings \mathbb{J}_{ij} are nonuniform, and hence, the full sum over all configurations $\{s\}$ must be calculated. In simpler cases, e.g., $\mathbb{J}_{ij} = \mathbb{J} \delta_{i+1,j}$, the sum can be simplified and therefore the computational complexity reduced dramatically.

Note that the continuum limit of this result corresponds to the $\hbar = 0$ approximation of the quantum mechanical time evolution operator:

$$\begin{aligned} \text{Tr}U(t)|_{n=0} &= \text{Tr}\{e^{it\sigma^z \cdot \mathbb{J} \cdot \sigma^z} e^{\sigma^x \mathcal{O}(ht)}\} \\ &= \sum_{\{s\}} e^{it \sum_{i,j} s_i \mathbb{J}_{ij} s_j} + \mathcal{O}[(ht)^2], \end{aligned} \quad (41)$$

where the linear order in ht can be dropped because $\text{Tr}\sigma^x = 0$.

B. Next-to-leading odd order: $n = 1$ spin-flip pair

Let us now consider the case of a single spin-flip pair, that is, exactly one of the quantum spins has exactly two different values over time in the z basis. Some visualizations can be found in Fig. 2. We absorb the probability weight $-(\tan \delta \hbar)^2$ and the combinatorial factor $N_t/2$ (temporal translational invariance yields a factor N_t , exchangeability of start and end point of the flipped region a factor $\frac{1}{2}$) into a new prefactor:

$$Z_1 \equiv -\frac{N_t}{2} (\tan \delta \hbar)^2 Z_0. \quad (42)$$

Now for every configuration s without flips, there are L different candidates l for a flip with $N_t - 1$ different possible lengths τ each. We can write the resulting phase as

$$\begin{aligned} S_l|_{n=1} &= -(N_t - \tau) \sum_{i,j} s_i J_{ij} s_j - \tau \sum_{i,j} s_i (1 - 2\delta_{il}) J_{ij} (1 - 2\delta_{lj}) s_j \\ &= -\left(t \sum_{i,j} s_i \mathbb{J}_{ij} s_j - 4\tau \sum_j s_l J_{lj} s_j \right). \end{aligned} \quad (43)$$

Combining the prefactor and the phase leads to

$$\rho_z(E)|_{n=1} = Z_1 \delta_{z,-1} \sum_{\{s\}} \sum_l \sum_{\tau=1}^{N_t-1} \delta\left(E - t \sum_{i,j} s_i \mathbb{J}_{ij} s_j + 4\tau \sum_j s_l J_{lj} s_j\right), \quad (44)$$

$$\begin{aligned} \text{Tr}U(t)|_{n=1} &= Z_1 \sum_{\{s\}} e^{it \sum_{i,j} s_i \mathbb{J}_{ij} s_j} \sum_l \sum_{\tau=1}^{N_t-1} e^{-4i\tau \sum_j s_l J_{lj} s_j} \\ &= Z_1 \sum_{\{s\}} e^{it \sum_{i,j} s_i \mathbb{J}_{ij} s_j} \sum_l e^{-2it \sum_j s_l \mathbb{J}_{lj} s_j} \frac{\sin[2(N_t - 1) \sum_j s_l J_{lj} s_j]}{\sin(2 \sum_j s_l J_{lj} s_j)}, \end{aligned} \quad (45)$$

with the well-defined continuum limit:

$$\rho_z(E)|_{n=1} = \frac{(ht)^2}{2} \delta_{z,-1} \sum_{\{s\}} \sum_l \theta\left(-\left|E - t \sum_{i,j} s_i \mathbb{J}_{ij} s_j\right| + \left|4\tau \sum_j s_l J_{lj} s_j\right|\right), \quad (46)$$

$$\lim_{\delta \rightarrow 0} \text{Tr}U(t)|_{n=1} = -\frac{(ht)^2}{2} \sum_{\{s\}} e^{it \sum_{i,j} s_i \mathbb{J}_{ij} s_j} \sum_l e^{-2it \sum_j s_l \mathbb{J}_{lj} s_j} \text{sinc}\left(2t \sum_j s_l \mathbb{J}_{lj} s_j\right), \quad (47)$$

where θ denotes the Heaviside step function.

Note that the temporal sum in Eq. (45) could be calculated exactly since it reduced to a geometric series. This allows an evaluation of the remaining sum in $\mathcal{O}(L2^L)$, i.e., the same runtime as the contribution without flips.

Again, we could have obtained the continuum result up to leading order directly from the time evolution operator using the Zassenhaus formula:

$$\begin{aligned} \text{Tr}U(t)|_{n=1} &= \text{Tr}\left[e^{it\sigma^z \cdot \mathbb{J} \cdot \sigma^z} e^{it\mathbb{h} \sum_i \sigma_i^x}\right] + \mathcal{O}[(ht)^4] \\ &= \text{Tr}\left\{e^{it\sigma^z \cdot \mathbb{J} \cdot \sigma^z} \prod_i \left[\mathbb{1} + iht\sigma_i^x - \frac{1}{2}(ht)^2 \mathbb{1}\right]\right\} + \mathcal{O}[(ht)^4] \\ &= \left[1 - \frac{1}{2}L(ht)^2\right] \text{Tr}\left[e^{it\sigma^z \cdot \mathbb{J} \cdot \sigma^z}\right] + \mathcal{O}[(ht)^4]. \end{aligned} \quad (48)$$

Here, the cyclic property of the trace ensures that the simple first-order decomposition into σ^z and σ^x parts is equivalent to a symmetric decomposition with an error of order $\mathcal{O}[(ht)^3]$. Additionally, $\text{Tr}(\sigma^x)^3 = 0$ reduces the residual error to $\mathcal{O}[(ht)^4]$.

C. Higher orders: $n \geq 2$ spin-flip pairs

In principle, we can follow the steps of the previous section to write down the contribution of the n th order (n spin flips) for arbitrary n . The prefactor (without combinatorics) reads

$$Z_n \equiv (-\tan^2 \delta\mathbb{h})^n Z_0, \quad (49)$$

and the phase (including combinatorics) generalizes to

$$\begin{aligned} S_l &= \sum_k \sum_{i,j} s_i \left\{ \prod_l [1 - 2\delta_{il}\theta(k - \tau_1^l)\theta(\tau_2^l - k)] \right\} \\ &\times J_{ij} \left\{ \prod_l [1 - 2\delta_{lj}\theta(k - \tau_1^l)\theta(\tau_2^l - k)] \right\} s_j, \end{aligned} \quad (50)$$

where θ denotes the Heaviside step function, and $\tau_{1,2}^l$ are successive flipping times of spins at spatial position l . The partition sum must include all possible combinations of $\{l\}$ and corresponding $\tau_{1,2}^l$.

The evaluation of the partition sum for any $n \geq 2$, however, is prohibitively expensive because the temporal sum can no longer be evaluated analytically. Two different regions flipped at overlapping times introduce nontrivial entanglement (see right panel of Fig. 2). Another way to see this is the fact that the first-order Trotter decomposition of the time evolution operator results in an error of order $\mathcal{O}[(ht)^4]$. A higher order treatment therefore necessitates more complicated decompositions with nonzero commutators considered.

D. Spline fitting

In LLR_2 simulations, all the leading peaks have been removed from the DoS so that the remaining function is relatively smooth and can be approximated by a fit. Since we do not know anything about the global structure of the DoS in general, a smooth local approximation like a cubic spline [37]

is the canonical choice. Such a spline with a fixed number of free parameters or dof, where dof are smaller or equal to the number of data points, is fitted to the log-DoS, as depicted in Fig. 3. The fit can then be numerically integrated to high precision.

A smooth fit comes with two advantages. First, discretization errors in S_l due to a finite bin size in Eq. (26) are reduced. In practice, this effect is usually negligible, and if it is not, a finer discretization should be chosen in the first place. More relevantly, fluctuations of neighboring bins with large uncertainties are averaged out. This allows decent estimations of the SFF at relatively high noise levels.

The single but crucial disadvantage of fitting the data is that *a priori* the optimal number k_{dof} of dof is completely unclear. In Fig. 3, we show the same data fitted with different order splines to the effect that the leftmost is underfitted and will yield completely wrong results, whereas the rightmost is overfitted and the desired smoothing effect is absent. We use the Bayesian information criterion (BIC) [39]:

$$\text{BIC} = M \ln \left(\frac{\chi^2}{M} \right) + k_{\text{dof}} \ln(M), \quad (51)$$

to find the optimal number of dof. Here, M is again the number of bins of the DoS, and χ^2 is the usual chi-squared value. The BIC is supposed to be minimal for the best choice of k_{dof} .

Overall, the fitting allows us to extract results from otherwise too noisy data, but they come at the cost of additional uncertainty and possible bias so that high-precision data is better evaluated without any fitting. This becomes clear from the comparison of the two top plots on Fig. 4, where the errors obtained from spline-fitted results are only considerably smaller for small numbers of MC updates.

E. Origin of the nonsmooth DoS

The highly irregular shape of the DoS is the crucial difference to the DoS of models with continuous (and compact) variables where the LLR approach has been shown to outperform reweighting significantly [12]. It is therefore important to understand the origin of this shape. Naïvely, one might think that it has to do with the random coupling and the chaotic nature of the system, but it turns out that this is not the case. Even the simplest transverse Ising model with constant coupling and no next-to-nearest neighbor interaction features a nonsmooth DoS.

Qualitatively, the origin of the strange form of the DoS can be understood starting with the spikes in LLR_0 forming the highest band. In the case when the nearest neighbor couplings contain quenched disorder $\Delta \mathbb{J}_i$, there are 2^{L-1} spikes of equal height. In the case of constant couplings, for example, for the conventional transverse field quantum Ising model without quenched disorder, there are $\lfloor L/2 \rfloor + 1$ spikes with binomially distributed heights. Now the second band, i.e., the single flip-pair expansion, can be seen (up to a proportionality factor) as a convolution of the first band with a rectangular function, taking the value 1 for every S_l that can be reached by a single flip pair and the value 0 everywhere else. Thus, every delta peak smears out to a superposition of rectangle-shaped functions at the same position, as can be seen in Eq. (46). Its width relative to the total range of S_l scales as $\mathcal{O}(1/L)$. The

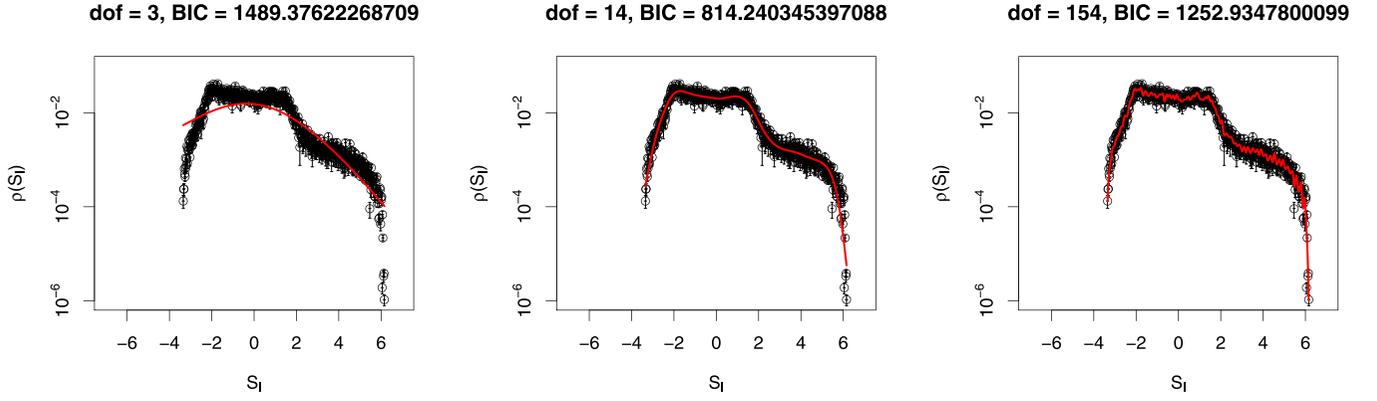


FIG. 3. Spline fits to $\log[\rho_+(S_I) - \rho_-(S_I)]$ with different degrees of freedom (dof) for 538 data points and respective Bayesian information criterion (BIC), Eq. (51). The spin chain length is $L = 6$, the evolution time is $t = 1$, the number of Monte Carlo sweeps is $N_s = 10^4$.

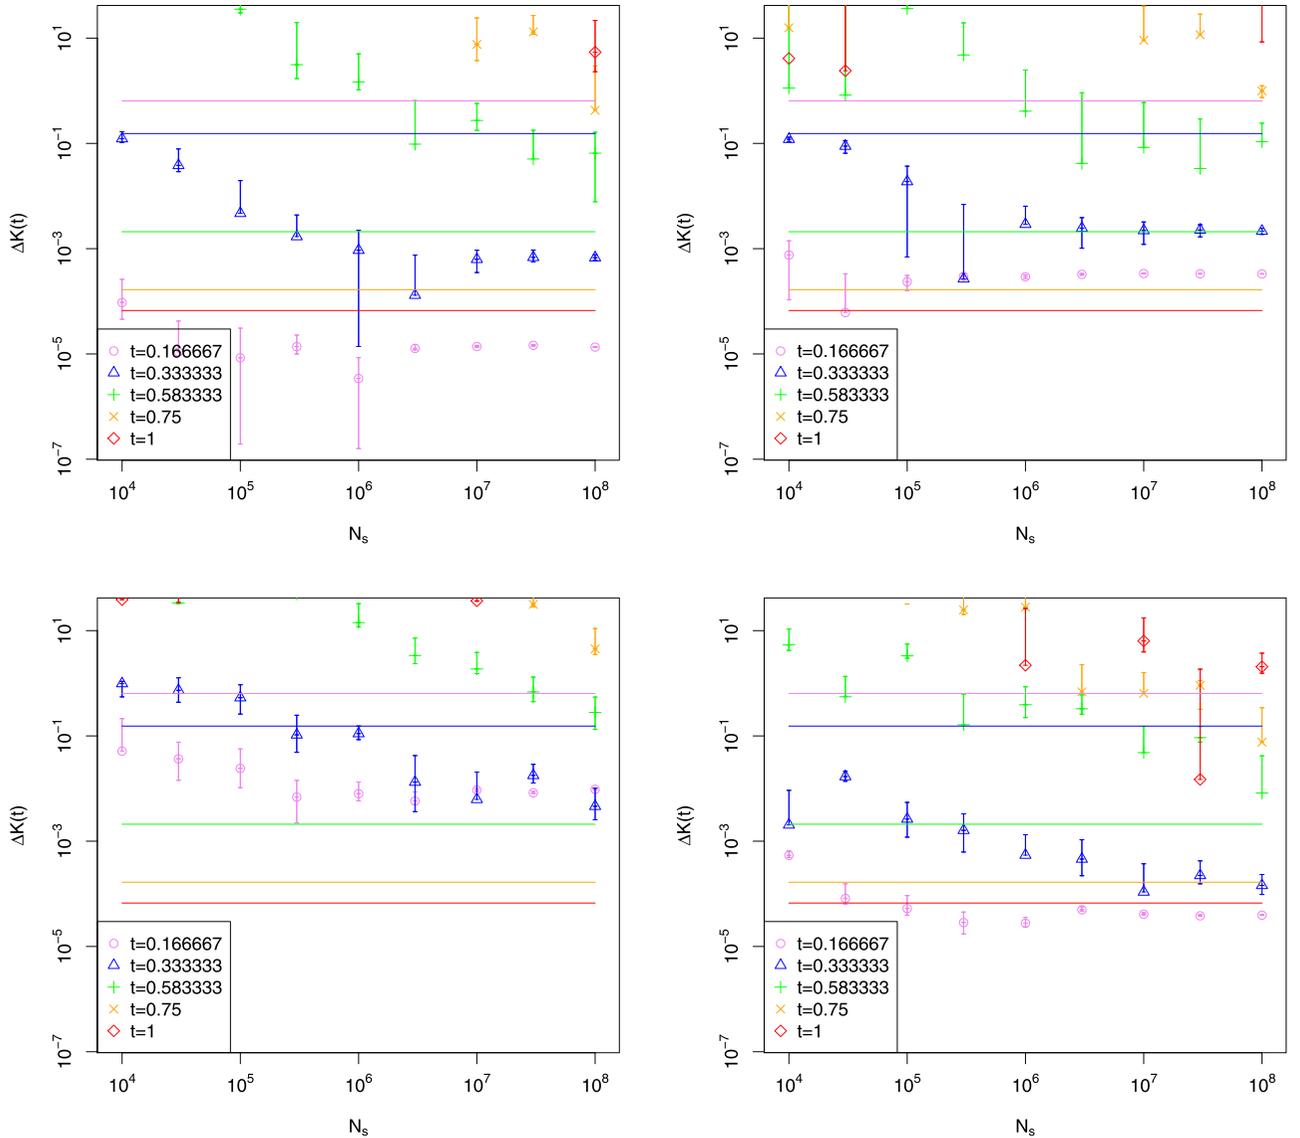


FIG. 4. Relative errors of the spectral form factor $\Delta K(t) = |K(t)/K_{ED}(t) - 1|$, where $K_{ED}(t)$ has been calculated with exact diagonalization, as a function of the number of sweeps N_s . For small evolution times $t \lesssim 0.3$, $\Delta K(t)$ reaches a plateau value that is entirely saturated by nonstatistical Trotterization errors. Error bars in negative direction that would extend below 0 are omitted. Lines show the respective statistical powers. Top left to bottom right: LLR_2 , LLR_2 with spline fitting, LLR_0 , and REW_2 . The chain length is $L = 16$ for all plots.

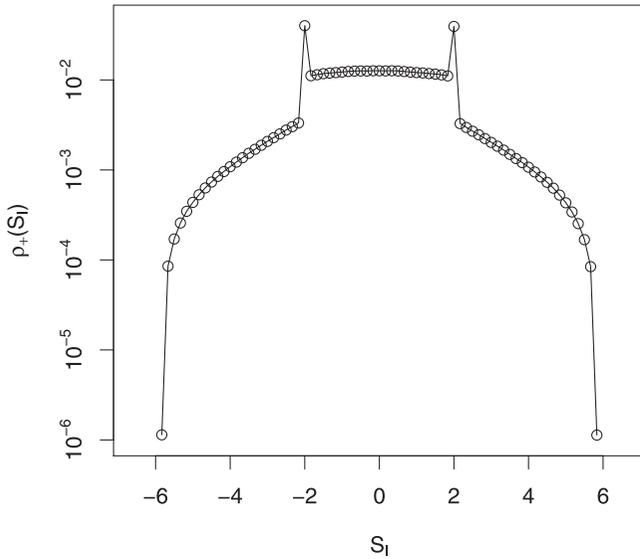


FIG. 5. Density of states in the even (positive) parity sector, calculated with LLR₂ for the pure transverse Ising model with constant coupling, i.e., $\mathbb{J}_2 = \Delta\mathbb{J} = 0$. Chain length $L = 6$; time, field strength and coupling $t = \hbar = \mathbb{J}_0 = 1$; number of Monte Carlo sweeps $N_s = 10^7$.

generalization is straightforward. With every additional flip pair, the previous distribution is convoluted with a rectangular function. For instance, the DoS of LLR₂ looks like the original spikes of LLR₀ convoluted with a triangle function, yielding the batman profile in Fig. 5 for the simple transverse field Ising model with constant coupling $\mathbb{J}_{ij} = \delta_{i,j+1}$ and $L = 6$.

Note that the notion of successive convolutions is only exact up to a single flip pair. Beyond, it still allows us to understand the shape of the DoS qualitatively, but crucial contributions from interferences between two or more flip pairs are neglected (e.g., two flip pairs might be on the same spatial site) so that the DoS is not reproduced correctly.

With larger L , the DoS smooths out in the disordered case. The overall smoothing stems from an exponentially large number of peaks constituting LLR₀, not from a broader smearing of the lower bands, since the convolution always acts on scales of $\mathcal{O}(1/L)$.

V. RUNTIME COMPARISON

A. ED

The straightforward solution for arbitrary times is the ED of the full 2^L -dimensional system. Since the diagonalization

of a matrix has cubic scaling, we end up with a runtime in $\mathcal{O}(2^{3L})$.

In our benchmarks, we used the LAPACKE_DSYEV routine provided by the LAPACKE package in C to calculate the complete spectrum of the respective Hamiltonian to machine precision.

For completeness, we also mention that the brute force approach of simply trying all the combinations in the classical system to calculate the exact partition sum has a runtime of $\mathcal{O}(2^{LN_i})$.

B. Reweighting

Stochastic methods like reweighting cannot be compared with exact methods like ED directly since the computational effort crucially depends on the desired precision ε . It is, however, reasonable to expect square-root convergence in the amount of statistics due to the central limit theorem and to set a sweep over the full lattice of $N = LN_i$ sites as the smallest unit of data. With these assumptions, we get a runtime scaling of $\mathcal{O}(LN_i\varepsilon^{-2})$ for a system without a sign problem. It is important to keep in mind that this is the best-case scenario, and phenomena like critical slowing down are completely neglected.

For the sake of simplicity, we will assume a small enough and constant Trotter step size δ , so that $N_i \propto t$. Note that this implies exponential runtime even without a sign problem for physically interesting long times $t \sim 2^L$.

The crucial change we must introduce for systems with a sign problem is a factor of the statistical power:

$$\Sigma = \left| \frac{Z}{Z_R} \right|, \quad (52)$$

that is, the modulus of the expectation value of the complex phase. This modification leads to a runtime scaling of

$$T_{\text{sim}} \sim \mathcal{O}(Lt\varepsilon^{-2}\Sigma^{-2}), \quad (53)$$

so in the following, we must estimate the magnitude of Σ as best we can. It is clear that we cannot calculate it exactly since that would allow us to solve the complete initial quantum mechanical problem.

We start with the observation that

$$|\text{Tr}U(t)| = |AZ_R| \Sigma, \quad (54)$$

where the left-hand side can be estimated from universal properties of the SFF, and the product AZ_R is known analytically. Let us expand the latter product near the continuum limit:

$$\begin{aligned} |AZ_R| &= \prod_i (\sin \delta\hbar_i \cos \delta\hbar_i)^{N_i/2} \left[\cosh \left(-\frac{1}{2} \ln \tan \delta\hbar_i \right)^{N_i} + \sinh \left(-\frac{1}{2} \ln \tan \delta\hbar_i \right)^{N_i} \right] \\ &= \prod_i (\delta\hbar_i)^{N_i/2} \left[\left(\frac{1}{\sqrt{\delta\hbar_i}} + \sqrt{\delta\hbar_i} \right)^{N_i} + \left(\frac{1}{\sqrt{\delta\hbar_i}} - \sqrt{\delta\hbar_i} \right)^{N_i} \right] [1 + \mathcal{O}(\delta\hbar_i)] \end{aligned}$$

$$\begin{aligned}
&= \prod_i [(1 + \delta h_i)^{N_i} + (1 - \delta h_i)^{N_i}] [1 + \mathcal{O}(\delta h_i)] \\
&= \prod_i \left[\left(1 + \frac{t h_i}{N_i}\right)^{N_i} + \left(1 - \frac{t h_i}{N_i}\right)^{N_i} \right] [1 + \mathcal{O}(\delta h_i)] \\
&= \prod_i 2 \cosh(t h_i) [1 + \mathcal{O}(\delta h_i)]. \tag{55}
\end{aligned}$$

In the common case of all fields equal $h_i \equiv h$, we thus obtain the formula:

$$\Sigma = \frac{|\text{Tr}U(t)|}{2^L \cosh(t h)^L}, \tag{56}$$

which is exact as $\delta \rightarrow 0$.

Since the SFF is bound from above $|\text{Tr}U(t)| \leq 2^L$, the best possible runtime scaling amounts to

$$T_{\text{sim}} \sim \mathcal{O}[L t \varepsilon^{-2} \cosh(t h)^{2L}], \tag{57}$$

so that reweighting can never outperform ED when $t h \gtrsim \text{acosh}2 \approx 1.3$.

Empirically, we find that, for small times, $|\text{Tr}U(t)| \sim 2^L e^{-L t^2}$ until it reaches the plateau at 2^L . Moreover, to maintain high precision, usually we have to scale the Trotter step with the system size $\delta \sim 1/L$, so that a more realistic runtime estimate is given by

$$T_{\text{sim}} \sim \mathcal{O}\left[\frac{1}{2^{-L} + e^{-L t^2}} L^2 t \varepsilon^{-2} \cosh(t h)^{2L}\right]. \tag{58}$$

C. LLR

In what follows, we show that, for both reweighting and LLR, the error is strongly dominated by the bulk, that is, the region with highest DoS. Since LLR is only advantageous in sampling the tails of the DoS distribution more accurately, it cannot significantly outperform reweighting. This conclusion is also confirmed by our MC data.

For a more formal comparison, we must introduce several assumptions. First, we demand that all the LLR parameters have been tuned properly. The bin size used for LLR is small enough to have negligible discretization errors. Such a choice is always possible, though the need to tune the parameters is a serious disadvantage of LLR. Secondly, we bin the data accumulated for reweighting in the same fashion, again without significant loss of precision. Now the normalized DoS obtained by either method can be interpreted as a probability distribution ρ_i on the same discrete and finite set $i \in \{1, \dots, M\}$. The same as in Eq. (26), M denotes the number of bins the DoS has been split into.

It is easy to see that, with reweighting every bin, i is visited $N_i = \rho_i \Lambda$ times on average, where Λ is the total number of samples. Therefore, in the large Λ limit, according to the central limit theorem (or naturally assuming a Poisson distribution of bin counts), the absolute uncertainty of the number of counts in said bin is $\sqrt{N_i}$ yielding an error of

$$\Delta \rho_i \equiv \frac{\rho_i}{\sqrt{N_i}} = \sqrt{\frac{\rho_i}{\Lambda}}. \tag{59}$$

Since LLR is a special case of the RM algorithm with a nondifferentiable target function, the coefficients in the exponent converge strictly as $\mathcal{O}(1/\sqrt{\Lambda})$ [40]. Thus, the relative error of ρ_i lies in $\mathcal{O}[e^{1/\sqrt{\Lambda}} - 1] = \mathcal{O}(1/\sqrt{\Lambda})$. We can obtain some bounds on the proportionality factor as well, though these results are by no means as precise as for reweighting.

In the early stages of the RM iteration, the probability to update the i th bin is simply ρ_i (the same as for reweighting). As RM iterations proceed and the coefficients α_i are updated, at asymptotically large Λ , all bins are visited with equal probabilities $\sim 1/M$. Therefore, in the bulk of the DoS defined as all the regions with $\rho_i \gg 1/M$, the probability to visit bin i decreases with the number of iterations Λ and eventually approaches $1/M$ from above. We can thus provide bounds on N_i in the bulk region:

$$\rho_i \Lambda \geq N_i \geq \frac{\Lambda}{M}. \tag{60}$$

On the other hand, in the tail of the distribution where $\rho_i \ll 1/M$, the probability to visit the bin i increases with the number of RM iterations and eventually approaches $1/M$ from below. The expectation value of the iteration count N_i in this region is therefore bounded by

$$\rho_i \Lambda \leq N_i \leq \frac{\Lambda}{M}. \tag{61}$$

In the following, we assume the best possible scenario, using the upper bounds on N_i , in which case the statistical error is given by

$$\Delta \rho_i = \begin{cases} \sqrt{\frac{\rho_i}{\Lambda}} & \text{if } i \text{ in bulk,} \\ \rho_i \sqrt{\frac{M}{\Lambda}} & \text{if } i \text{ in tail,} \end{cases} \tag{62}$$

and shows that LLR cannot outperform reweighting significantly even in that case.

Let us consider some probability ρ_0 from the bulk and another $\rho_\infty = \zeta \rho_0$ from the tail, where ζ is some small parameter. For reweighting, we need $\Lambda \sim \zeta^{-2}$ samples to resolve ρ_∞ at all. For LLR, on the other hand, the error in the tail is much smaller. However, even if we knew ρ_∞ exactly, this knowledge is useless if $\Delta \rho_0 > \rho_\infty$, i.e., the error is bulk dominated. From this follows the condition:

$$\Delta \rho_0 < \rho_\infty \Rightarrow \sqrt{\frac{\rho_0}{\Lambda}} < \zeta \rho_0 \Leftrightarrow \Lambda > \frac{\zeta^{-2}}{\rho_0}. \tag{63}$$

Thus, not only does LLR follow the same scaling as for reweighting in Eq. (53), but it even has very similar errors quantitatively since the statistical uncertainties in the bulk are almost identical for reweighting and LLR, and these uncertainties dominate in both cases.

We remark that this argument is applicable independently of the specific form of the DoS and the conclusion therefore generalizable beyond its application to spin chains. Generally, complexity of a given problem is encoded in the magnitude of the parameter ζ . For instance, in our case of real-time spin chain simulations, $\zeta \sim e^{-Lt}$ is exponentially suppressed with the system size providing an alternative perspective on the origin of the NP-hardness of the sign problem.

D. Augmented LLR₂ or REW₂

Removing the two leading-order contributions from the stochastic calculations and evaluating them exactly comes with an additional cost of

$$\mathcal{O}(L2^L) \quad (64)$$

as compared with plain vanilla LLR₀ or REW₀. However, it allows us to achieve the same precision with a significantly lower accuracy goal since the stochastic part contributes with a reduced weight. More specifically,

$$\varepsilon^{-1} \mapsto \varepsilon^{-1} \left\{ 1 - \cosh(t\hbar)^{-L} \left[1 + \frac{1}{2}L(t\hbar)^2 \right] \right\} \quad (65)$$

results in an overall runtime of the approximate order

$$T_{\text{sim}} \sim \mathcal{O} \left\{ L2^L + Lt\varepsilon^{-2} \left[\cosh(t\hbar)^L - 1 - \frac{L(t\hbar)^2}{2} \right]^2 \right\}. \quad (66)$$

This estimate should be compared with Eq. (57). We see that summation over the two leading-order contributions allows us to obtain equally precise results at short times with significantly smaller statistics (i.e., computational effort) and to extend the stochastically calculable time region toward somewhat larger values.

VI. RESULTS

All the simulations presented in this paper have been performed using R as a frontend and C as a backend. The complete code and most of the data are publicly available at Ref. [41]. Some of the data are rather large and have therefore not been published in the same way, but we are happy to provide them in case of interest.

We tested our algorithm on chains of up to $L = 50$ with LLR₀/REW₀ and up to $L = 40$ with LLR₂/REW₂. Small lengths $L \leq 16$ have been benchmarked against results from ED. We emphasize that the larger lattices with $L \gtrsim 24$ are far out of the reach of any exact calculations. For instance, $L = 40$ would require 16 TB of memory to even store a single complex state vector in double precision. We also stress that, while LLR₂/REW₂ has a runtime scaling as $\mathcal{O}(2^L)$ and therefore $L > 40$ quickly becomes unfeasible, there are no such limitations for LLR₀/REW₀, and L can be chosen practically arbitrarily large at the cost of somewhat reducing the evolution time t .

In each simulation, we chose a single random realization of the coupling J . The disorder average over J is technically straightforward and would only add a layer of uncertainties to the results, obscuring the quality of the algorithms.

Some typical examples of the scaling of the relative error between the MC and ED results with the number of MC iterations are shown in Fig. 4. The error is plotted as a function of

the total number of sweeps $N_s = \Lambda/N$ for different evolution times t and correspondingly different strengths of the sign problem. Respective statistical powers Σ are shown as lines. In an ideal stochastic simulation, the error should scale as $\Sigma^{-1}/\sqrt{N_s}$, and it does so usually. There are some exceptions in the decrease with N_s though. The shortest evolution times just plateau out at a constant error level. For LLR₀, the simulations would require a much finer discretization of S_I to resolve the spiked DoS properly. The LLR₂ and REW₂ calculations, on the other hand, are so precise that the Trotter error for the choice of N_t is resolved.

From the two top plots in Fig. 4, we also conclude that the results obtained with spline fitting tend to be better than the ones from simple integration for small numbers $N_s \lesssim 10^5$ of MC updates and for large evolution times $t \sim 1$ where the sign problem is more severe (compare the top left corners of both plots). More statistics undermines the usefulness of fits, as they cannot extract any additional information.

The true reason for the favorable spline results in the case of small statistics appears to be that the best fitting functions are very smooth in this case. Therefore, the highly oscillatory integral in Eq. (18) required to obtain the SFF yields a value close to zero. Since the analytic approximation used in LLR₂ is not too bad by itself, the overall result is better than the noise in the unfitted case. Thus, using the stochastic data does not help to improve the analytic approximation for too little statistics. However, in the spline-fitted case, it does not add any harm either.

It is evident that the augmentation defining LLR₂ allows us to reduce the error at short times by several orders of magnitude compared with plain vanilla LLR₀. It is also clear, however, that it does not significantly mitigate the sign problem at long times, confirming the predictions from Eq. (66).

A direct comparison of LLR₂ and REW₂ shows very little differences. If anything, reweighting appears to be slightly better. This is exactly what we expect from Sec. VC, where we showed that LLR and reweighting have principally the same sign problem. The minor difference in favor of reweighting most likely comes from the binning required for LLR as well as a suboptimal parameter choice (see Appendix C).

Figure 6 shows the SFFs corresponding to the errors in Fig. 4. We added the approximations for continuous times derived in Sec. IV. The strength of the LLR₂ method becomes clear when the small differences between the approximations and the exact results are observed. Only these differences must be calculated stochastically. Note, however, that even at very short times $t \lesssim 0.4$, the stochastic contribution cannot be neglected, and the MC results are not compatible with the approximations.

We further remark that the LLR results at large times and small statistics are systematically incompatible with the correct values. This can be explained by a required minimal number of iterations to fill up the DoS. We expect that a different choice of coefficients $a_{L,M,\Lambda}$ and b_M than in Eqs. (33) and (34) might reduce this problem, but we did not conduct additional extensive parameter tuning. This is certainly a great disadvantage of LLR as opposed to reweighting.

The results obtained for longer chains $L \geq 20$ and depicted in Fig. 7 are qualitatively similar to those discussed above for $L = 16$. Again, LLR₂ and REW₂ yield extremely precise

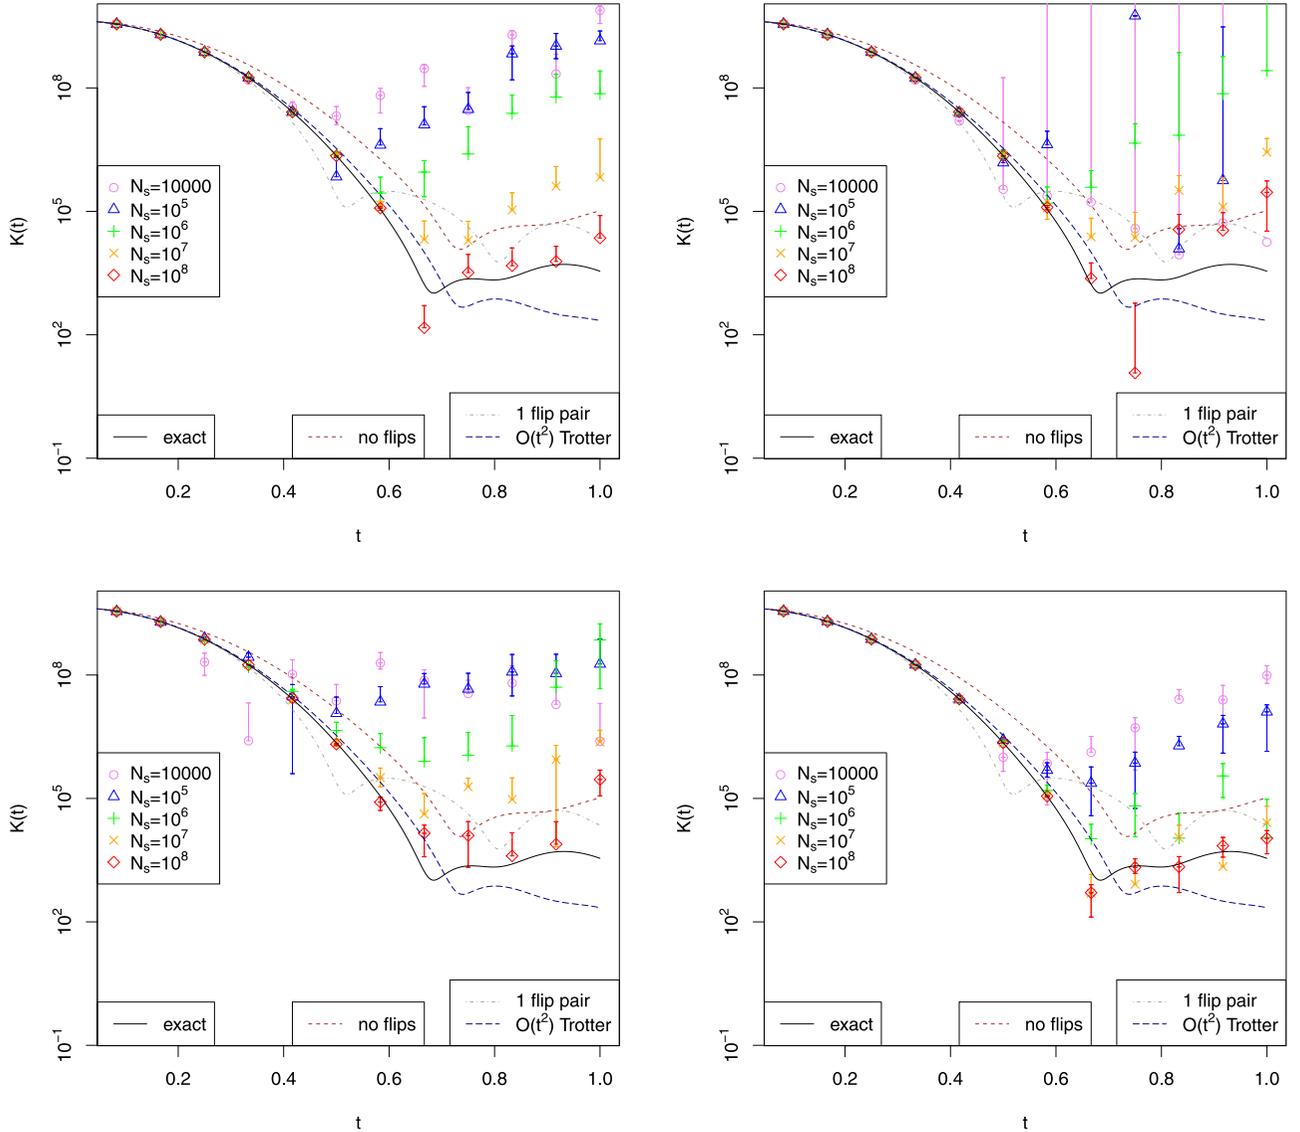


FIG. 6. The spectral form factors $K(t)$ compared with the solution from exact diagonalization (exact) and the approximations: no flips, Eq. (40); 1 flip pair, Eq. (47); and $\mathcal{O}(t^2)$ Trotter, Eq. (48); for different numbers of Monte Carlo sweeps N_s . Error bars in negative direction that would extend below 0 are omitted. Top left to bottom right: LLR_2 , LLR_2 with spline fitting, LLR_0 , and REW_2 . The chain length is $L = 16$ for all plots.

results at short times, and reweighting generally outperforms LLR. Notably, these results demonstrate the scalability of our stochastic algorithms toward very long chains, provided $t \lesssim 1$. Especially for REW_0 , we can obtain reliable results for virtually arbitrarily large values of L .

VII. DISCUSSION AND OUTLOOK

In this paper, we investigated the possibility to evaluate the SFF (real-time partition sum) of a nonintegrable spin chain in Eq. (5) using MC simulations. The partition sum contains a multitude of mutually canceling complex-valued contributions, which turns its direct evaluation into an NP-hard problem. We investigated the infinite-temperature spectral function because it poses the most challenging problem and is often considered in the literature. We remark, however,

that the method would be applicable with only minor modifications (a larger, in principle unlimited, number of parity sectors) to finite-temperature systems as well. In fact, finite temperatures might severely reduce the sign problem.

We compared the most straightforward approach based on importance sampling with subsequent reweighting with a more advanced DoS approach, in which the statistical distribution of the complex phase of the action is calculated with high precision.

One of our main results is the improved simulation strategy, in which the contribution of configurations with zero and one spin flips into the real-time partition function are summed over exactly. These configurations contribute to the leading and next-to-leading orders of the expansion of the real-time partition function $\text{Tr}U(t)$ in Eq. (6) in powers of the magnetic field h . Summation over all other configurations

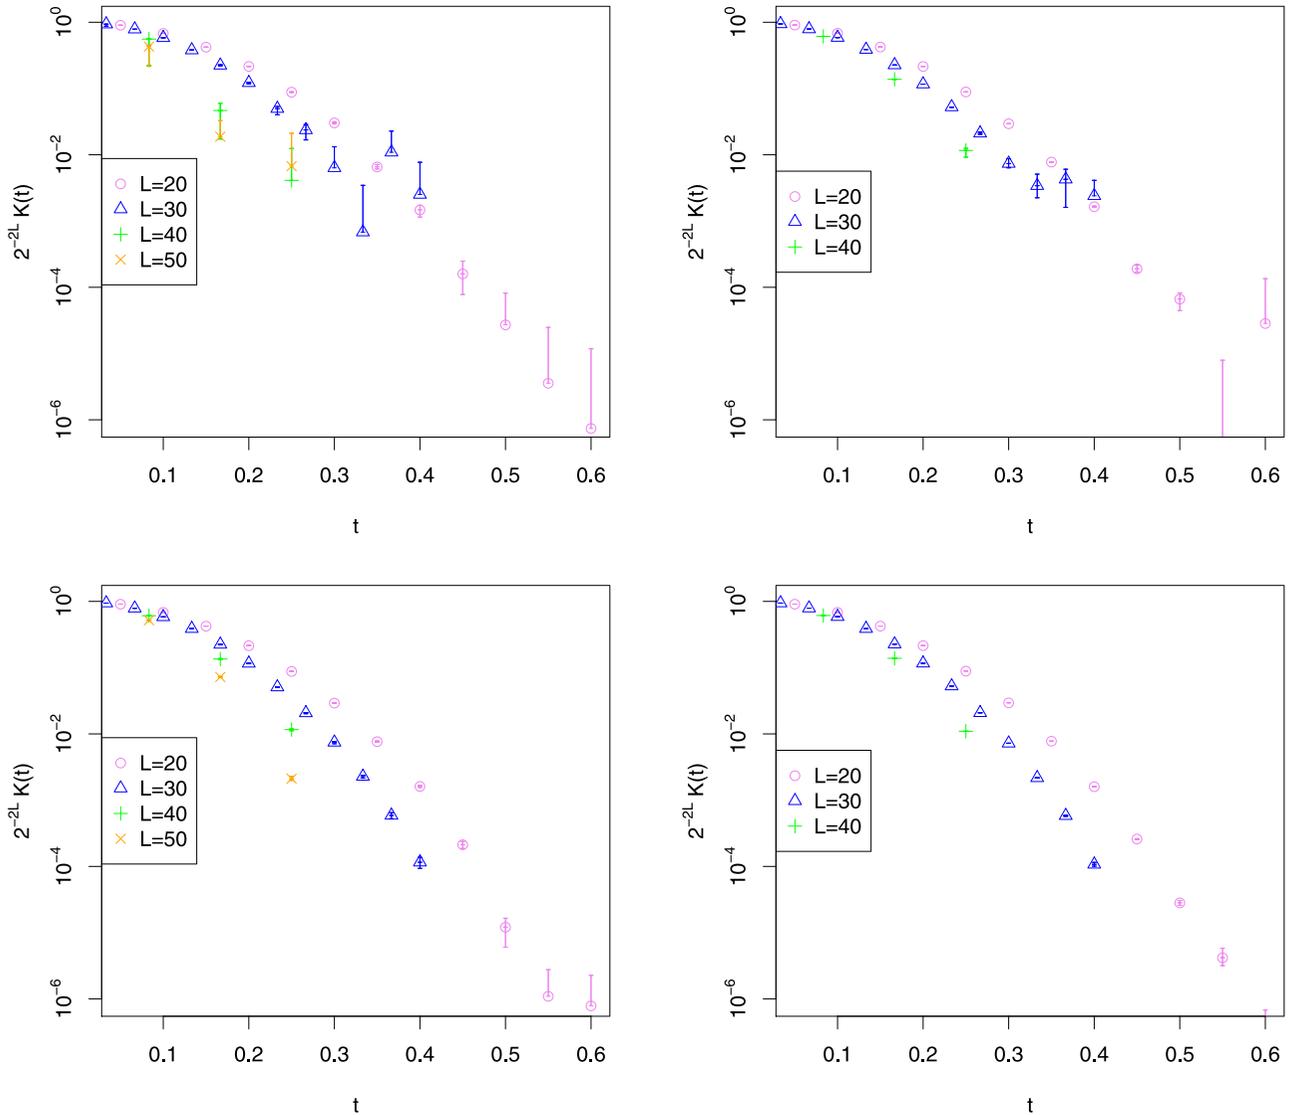


FIG. 7. Normalized spectral form factors $2^{-2L}K(t)$ for system sizes hardly ($L = 20$) or not at all ($L \geq 30$) treatable by exact diagonalization. Error bars in negative direction that would extend below 0 are omitted. Top left to bottom right: LLR_0 , LLR_2 , REW_0 , and REW_2 . The number of Monte Carlo sweeps is $N_s = 10^7$ for all plots.

is performed stochastically, using either the standard MC updates (for reweighting) or the LLR algorithm (for the DoS approach). As we discuss below, both approaches work nearly equally well.

As illustrated in Fig. 4, the improved simulation strategy significantly accelerates convergence to exact results, which we obtained using numerical diagonalization of the Hamiltonian. It allows us to study real-time evolution of spin chains with lengths up to $L = 40$ and up to physical times $t \lesssim 1$. Our method thus clearly outperforms numerical ED methods at short evolution times $t \lesssim 1$. Since for our spin chain the numerical cost of ED does not depend on the evolution times (up to possible issues with numerical precision of the result due to round-off errors), ED becomes more advantageous for $t \gtrsim 1$. The algorithm could also be translated to a continuous time formulation along the lines of Ref. [42], which removes Trotterization errors.

An advantage of the stochastic approach is that it can be easily extended to bosonic systems with infinite-dimensional Hilbert spaces, for which ED typically becomes prohibitively expensive even for $O(10)$ dof. Indeed, to obtain precise results for bosonic dof, the size of the local Hilbert space N_{loc} associated with each bosonic dof should usually be $\gg 2$, which results in a significantly faster growth of computational cost $T_{\text{sim}} \sim (N_{\text{loc}})^{n_{\text{dof}}}$ with the number of dof than for quantum spin chains. In addition, we need to extrapolate the result to the limit $N_{\text{loc}} \rightarrow +\infty$. The possibility to simulate short periods of real-time evolution for systems with a large number of bosonic dof is particularly attractive for numerical studies of the early-time evolution of quark-gluon plasma in heavy-ion collisions before the onset of hydrodynamic behavior.

It is less clear how our results would translate to fermionic systems in general because the performance would crucially depend on a particular realization. Broadly speaking, we can

distinguish two classes: Hubbard-Stratonovich transformations resulting in continuous dof and those resulting in discrete dof (e.g., Hirsch transformations). In the first case, we expect LLR to perform better than reweighting because of the implied smoothness of the DoS. On the other hand, the latter case enables improvements of the LLR_n/REW_n type.

Another important and somewhat counterintuitive conclusion is that, at least for our Hamiltonian and the Trotterization scheme, the DoS/LLR approach does not appear to offer significant advantage over the simple reweighting. In fact, reweighting often outperforms the DoS/LLR approach, presumably because of the suboptimal choice of parameters in the LLR algorithm. Reweighting, on the other hand, does not depend on any tunable parameters.

In general, the main technical advantage of the DoS/LLR approach is that it allows to resolve the tails of the distribution of the complex phase with much higher precision than reweighting. In our case, the contribution of these tails to the real-time partition sum is subdominant in comparison with the bulk contribution (central region in the distributions on Fig. 1). Therefore, sampling the tails with high precision is not an advantage in itself. As stressed in Refs. [12,43–45], high-precision data for the tails of the DoS become an advantage if the DoS $\rho(E)$ can be well approximated by some function with a small number of parameters, e.g., a low-degree polynomial fit or a spline. In this case, precise knowledge on the distribution tails translates into a better choice of fit parameters. This allows for more precise numerical estimates of the oscillatory integrals of the form $Z = \int dE \rho(E) e^{iE}$ in the representation in Eq. (4) of the real-time partition function.

As discussed in Sec. IV E, in our case, the DoS in Eq. (19) is a sum of δ functions and piecewise constant terms for any finite L . The DoS only looks like a continuous function in the thermodynamic limit $L \rightarrow +\infty$. While separating out the leading and next-to-leading contributions to the DoS, as described in Secs. IV A and IV B, leaves us with a much smoother DoS, we were still not able to approximate it by a function (polynomial or spline) with a sufficiently small number of parameters. As a result, spline or polynomial approximations improve the convergence to exact results only for small numbers of MC samples in RM iterations (29). Consequently, the LLR approach did not offer significant improvement over the reweighting approach.

In this respect, the LLR approach might work better for systems with continuous dof, where the DoS $\rho(E)$ in Eq. (3) is a continuous function even for a finite number of dof (away from the thermodynamic limit). See the fermionic system considered in Ref. [21] for an example. In this case, the DoS might allow for good-quality approximations in terms of functions with a small number of parameters, which allows us to effectively use the information about the tails of the DoS to improve the precision of oscillatory integrals of the form of Eq. (4). However, other complications might arise when applying the DoS/LLR method to continuous dof. To provide the most obvious example, the DoS will be clearly a nonnormalizable and unbounded function for the real-time partition function of even the simplest one-dimensional quantum harmonic oscillator. For the harmonic oscillator, this unboundedness can be easily solved by complexification of the path integral, which makes the path integral com-

pletely real valued. The complexification, however, becomes increasingly complicated for anharmonic potentials. As a result, we will be inevitably forced to adapt the Lefschetz thimble/holomorphic flow approaches, which have their own complications [46–50]. A further problem with the DoS/LLR approach for continuous variables might be the loss of ergodicity for simulations with constrained values of the imaginary part of the action S_I , especially if one uses the hybrid MC algorithm based on nearly continuous updates of field variables. Finally, the polynomial approximations to the DoS, which plays a crucial role in reducing statistical errors, might fail in the vicinity of phase transitions [15].

We should note that the DoS/LLR approach might also be useful for exploratory studies, for example, to guide the construction of analytic approximations. In fact, our work proceeded in exactly this way—we first measured the unsubtracted DoS function with LLR_0 , and only after looking at the DoS plots in Fig. 1, we realized how to subtract the leading and next-to-leading contributions specified in Eqs. (40) and (47).

ACKNOWLEDGMENTS

This paper was funded in part by the Science and Technology Facilities Council (STFC) Consolidated Grant No. ST/T000988/1. Numerical simulations were undertaken on Barkla, part of the high-performance computing (HPC) facilities at the University of Liverpool, as well as on the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by the University of Cambridge Research Computing on behalf of the STFC Distributed Research Utilising Advanced Computing (DiRAC) HPC Facility [51]. The DiRAC component of CSD3 was funded by Department for Business, Energy and Industrial Strategy capital funding via STFC Capital Grants No. ST/P002307/1 and No. ST/R002452/1 and STFC Operations Grant No. ST/R00689X/1. DiRAC is part of the National e-Infrastructure. We used the R library HADRON [52] for data analysis. JO would like to thank the other participants of the Pollica Summer Workshop on Effective Field Theories 2022, particularly Michael Flynn, for many an insightful discussion. We also thank Kurt Langfeld, Paul Rakow, and David Schaich for their valuable comments.

APPENDIX A: DERIVATION OF THE CLASSICAL ISING MODEL IN $d + 1$ DIMENSIONS AND FURTHER DETAILS

The exponentials in Eq. (9) can be calculated exactly in the canonical z basis because the first term is diagonal and the second one can be decomposed into local blocks of size 2×2 , yielding

$$e^{i\delta h_i \sigma_i^x} = \begin{pmatrix} \cos \delta h_i & i \sin \delta h_i \\ i \sin \delta h_i & \cos \delta h_i \end{pmatrix}. \quad (\text{A1})$$

Let us now compare these matrices with the local building blocks of the anisotropic classical Ising model with an arbitrary coupling matrix J' within the first dimension and a nearest neighbor coupling h' within the second dimension

defined by the action:

$$S = S^1 + S^2, \quad (\text{A2})$$

$$S^1 = - \sum_k \sum_{i,j} s_{i,k} J'_{ij} s_{j,k}, \quad (\text{A3})$$

$$S^2 = - \sum_i \sum_k h'_i s_{i,k} s_{i,k+1}. \quad (\text{A4})$$

The physics of such a classical statistical system are governed by the Boltzmann weighted partition sum (over all possible spin configurations):

$$Z = \sum_{\{s\}} e^{-S(s)}, \quad (\text{A5})$$

inducing the transfer matrices:

$$T_1 = [e^{-S^1}]_{k,k+1} = \text{diag} \left(\sum_{i,j} s_i J'_{ij} s_j \mid s \in \{\pm 1\}^L \right), \quad (\text{A6})$$

and

$$\begin{aligned} (T_2)_i &= \{e^{-(S^2)_i}\}_{k,k+1} = \begin{bmatrix} e^{h'_i} & e^{-h'_i} \\ e^{-h'_i} & e^{h'_i} \end{bmatrix} \\ &= \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{bmatrix} e^{h'_i} + e^{-h'_i} & 0 \\ 0 & e^{h'_i} - e^{-h'_i} \end{bmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \end{aligned} \quad (\text{A7})$$

We immediately observe the similarity between T_1 and the interacting part of the time evolution operator. Therefore, we set

$$J' \equiv i\delta\mathbb{J}. \quad (\text{A8})$$

Note again that this \mathbb{J} corresponds to the purely spatial coupling matrix. The coupling in the temporal direction turns out to be more challenging, as we must exploit the similarity of T_2 with the 2×2 matrix in Eq. (A1):

$$\begin{aligned} \begin{bmatrix} e^{h'_i} & e^{-h'_i} \\ e^{-h'_i} & e^{h'_i} \end{bmatrix} &\propto \begin{bmatrix} \cos \delta h_i & i \sin \delta h_i \\ i \sin \delta h_i & \cos \delta h_i \end{bmatrix} \\ \Rightarrow e^{-2h'_i} &= i \tan \delta h_i. \end{aligned} \quad (\text{A9})$$

Together with the proportionality factor:

$$A = \prod_i A_i^{N_i}, \quad (\text{A10})$$

$$A_i = \sqrt{i \sin \delta h_i \cos \delta h_i}, \quad (\text{A11})$$

this allows an exact identification between the trace of the quantum mechanical time evolution operator and the classical partition sum:

$$\text{Tr}U(t) = AZ. \quad (\text{A12})$$

We can define the purely real constants:

$$J := \delta\mathbb{J}, \quad (\text{A13})$$

$$h_i := -\frac{1}{2} \log \tan \delta h_i, \quad (\text{A14})$$

so that the action reads

$$\begin{aligned} S &= -i \sum_k \sum_{i,j} s_{i,k} J_{ij} s_{j,k} \\ &\quad - \sum_i \sum_k \left(h_i - \frac{\pi}{4} i \right) s_{i,k} s_{i,k+1}. \end{aligned} \quad (\text{A15})$$

1. Distribution of the imaginary part

It is instructive to investigate the distribution of S_I before we go on. For this, we define the auxiliary variable:

$$\theta_k := \sum_{i,j} s_{i,k} J_{ij} s_{j,k}, \quad (\text{A16})$$

with the upper bound:

$$|\theta_k| \leq \|J\|_1 \leq nL \delta \|J\|_\infty, \quad (\text{A17})$$

where n is the number of neighbors a single site can couple to ($n = 2$ for pure nearest neighbor coupling in one dimension), and $\|\cdot\|_p$ denotes the p norm. Since the considered system is translationally invariant in the time direction, all θ_k are distributed identically, though not independently in general. The \mathbb{Z}_2 symmetry of spin reflections is unbroken, so the expectation value of θ_k is $\bar{\theta}_k = 0$.

Suppose the θ_k were uncorrelated. Then according to the central limit theorem, S_I would follow a normal distribution:

$$S_I \sim \mathcal{N}(0, v), \quad (\text{A18})$$

with a variance:

$$v \leq N_t (nL \delta \|J\|_\infty)^2 = \frac{1}{N_t} (nL t \|J\|_\infty)^2, \quad (\text{A19})$$

approaching zero in the continuum limit.

In the opposite limit of maximal correlation, on the other hand, S_I simplifies to

$$S_I = N_t \theta_k \Rightarrow |S_I| \leq N_t nL \delta \|J\|_\infty = nL t \|J\|_\infty, \quad (\text{A20})$$

that is, a constant in the time-step size. We do not consider the case of anticorrelation since it is highly unphysical. Therefore, we find that, in any relevant case, S_I is bounded by the physical time extent, and the variance is going to approach a constant:

$$\lim_{N_t \rightarrow \infty} v \leq t \xi (nL \|J\|_\infty)^2, \quad (\text{A21})$$

in the continuum limit, where ξ is the correlation length in the time direction. Here, ξ is proportional to the standard deviation of S_T .

2. Alternative boundary conditions

We remark here that the pbc silently assumed above are the most usual but not the only possible choice. Physically, they correspond to same-to-same scattering and are best suited for infinite time approximations. In contrast, open boundary conditions (obc) correspond to all-to-all scattering and might be relevant in comparisons with the matrix product state or similar calculations where they are much easier to realize than pbc. Closed boundary conditions (cbc) allow for particular choices of initial and final states.

In all the nonperiodic cases, there is one less time step than time slices, so the proportionality factor between quantum and classical partition sums reduces to $\prod_i A_i^{N_i-1}$. Furthermore, the trace in Eq. (22) must be replaced by projections to corresponding vectors.

For obc, one obtains

$$Z_{R,i} = \frac{1}{2}(1, 1)(T_2)_i^{N_i} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = 2^{N_i-1} \cosh^{N_i-1} h'_i, \quad (\text{A22})$$

and for cbc,

$$\begin{aligned} Z_{R,i} &= (1, 0)(T_2)_i^{N_i} \begin{pmatrix} \cos \varphi \\ \sin \varphi \end{pmatrix} \\ &= 2^{N_i-1} [(\cos \varphi + \sin \varphi) \cosh^{N_i-1} h'_i \\ &\quad + (\cos \varphi - \sin \varphi) \cosh^{N_i-1} h'_i], \end{aligned} \quad (\text{A23})$$

where initial and final states are rotated against each other by the angle φ . We observe that, as expected, all the boundary conditions lead to the same result at infinite times $t \rightarrow \infty$, $N_i \rightarrow \infty$, and moreover, the cbc case of $\varphi = 0$ (same-to-same scattering) yields the same result as pbc up to an irrelevant difference of N_i by one.

APPENDIX B: ERROR SCALING WITH BIN SIZE

The discretization of the DoS into bins as in Eq. (26) naturally leads to errors vanishing as the bin width goes to zero but is relevant for finite bin size. The observable we are ultimately interested in amounts to a Fourier transformation of the DoS, so we can judge the quality of the discretization method by its ability to approximate the integral:

$$F_a(x_0, h) := \frac{1}{2h} \int_{x_0-h}^{x_0+h} dx e^{f(x)+iax}. \quad (\text{B1})$$

The separate estimation of the DoS in bins of length $2h$ succeeded by a multiplication with the phase at the midpoint of the interval, as employed in this paper, results in

$$\begin{aligned} \frac{e^{iax_0}}{2h} \int_{x_0-h}^{x_0+h} dx e^{f(x)} \\ = F_a(x_0, h) e^{-\frac{1}{6}ia[2f'(x_0)+ia]h^2+\mathcal{O}(h^4)}. \end{aligned} \quad (\text{B2})$$

Alternative approximations include the classical midpoint formula:

$$e^{f(x_0)+iax_0} = F_a(x_0, h) e^{-\frac{1}{6}\{f''(x_0)+[f'(x_0)+ia]^2\}h^2+\mathcal{O}(h^4)}, \quad (\text{B3})$$

as well as a more intricate piecewise linear (trapezoidal) formula:

$$\begin{aligned} \frac{e^{f(x_0)}}{2h} \int_{x_0-h}^{x_0+h} dx e^{f'(x_0)x+iax} \\ = F_a(x_0, h) e^{-\frac{1}{6}f''(x_0)h^2+\mathcal{O}(h^4)}, \end{aligned} \quad (\text{B4})$$

as used, e.g., in Ref. [20].

Thus, the multiplicative errors of all the different methods are of order $\mathcal{O}(h^2)$. Empirically, we find that the bin size can easily be chosen small enough to completely neglect this error since the uncertainties in the estimation of $f(x)$ itself dominate the total error.

APPENDIX C: LLR STEP SIZE

First, we note that the condition in Eq. (30) is not sufficient in practice, as the sum does not go up to infinity but only up to Λ . Therefore, we require additionally

$$\sum_{k=0}^{\Lambda} \beta_{N,M,\Lambda}(k) \geq \sum_i \hat{\alpha}_i, \quad (\text{C1})$$

where we denote the exact logarithmic DoS with the hat $\hat{\alpha}_i$. The left-hand side of Eq. (C1) evaluates to

$$\sum_{k=0}^{\Lambda} \beta_{N,M,\Lambda}(k) = a_{N,M,\Lambda} \left[\ln \frac{\Lambda}{b_{N,M,\Lambda}} + \mathcal{O}(b_{N,M,\Lambda}^{-1}) + \mathcal{O}(\Lambda^{-1}) \right], \quad (\text{C2})$$

while the right-hand side can only be approximated:

$$\sum_i \hat{\alpha}_i \leq \sum_i N \ln 2 = MN \ln 2 \approx 0.69MN. \quad (\text{C3})$$

Though this is a conservative upper bound, it is not very far from realistic estimations. If we, for instance, assume a parabolic shape for α_i with a maximum close to $N \ln 2$, we obtain $\sum_i \hat{\alpha}_i \lesssim \frac{2}{3} \ln 2 MN \approx 0.46MN$. Thus, up to a factor of order one, we find that

$$a_{N,M,\Lambda} = \ln 2 \frac{MN}{\ln \frac{\Lambda}{b_{N,M,\Lambda}}} \quad (\text{C4})$$

is a good choice.

To get a value for the remaining offset $b_{N,M,\Lambda}$, we must consider the small k behavior rather than the large k limit as we did before because the offset is going to be irrelevant in the latter case. Consider the likely case that one of the first contributions $\beta_{N,M,\Lambda}(k \ll b_{N,M,\Lambda})$ is added to a state that has a much smaller exact $\hat{\alpha}_i$ relative to the other states. To compensate this mistake after some k_0 steps (k_0 can be large if the initial random walk remains in a distant region of the

phase space for a while), we must spend κ steps adding this contribution to every other of the M states. For this to be possible, the condition:

$$\frac{M}{b_{N,M,\Lambda}} \leq \sum_{k=k_0}^{k_0+\kappa} \frac{1}{b_{N,M,\Lambda} + k} = \ln \frac{k_0 + \kappa}{k_0} + \mathcal{O}\left(\frac{b_{N,M,\Lambda}}{k_0}\right), \quad (\text{C5})$$

must hold. This provides a lower bound for κ :

$$\kappa \geq k_0 \left[e^{\frac{M}{b_{N,M,\Lambda}}} - 1 \right], \quad (\text{C6})$$

which must be small for the algorithm to be efficient. Thus, it turns out that the offset $b_{N,M,\Lambda} \equiv b_M$ solely depends on the number of states and $b_M \gtrsim M$ for the runtime to not blow up exponentially. On the other hand, we cannot choose b_M arbitrarily large either since that would defeat the purpose of $\beta_{N,M,\Lambda}(k)$ decreasing quickly. Numerical tests suggest that

$$b_M = 3M \quad (\text{C7})$$

yields a performance close to optimal.

-
- [1] R. P. Feynman, Quantum mechanical computers, *Found. Phys.* **16**, 507 (1986).
- [2] D. Nagaj, Fast universal quantum computation with railroad-switch local Hamiltonians, *J. Math. Phys.* **51**, 062201 (2010).
- [3] Up to floating-point round-off errors, that might become important at very late times.
- [4] J. A. Kjäll, J. H. Bardarson, and F. Pollmann, Many-Body Localization in a Disordered Quantum Ising Chain, *Phys. Rev. Lett.* **113**, 107204 (2014).
- [5] M. Serbyn, Z. Papić, and D. A. Abanin, Criterion for Many-Body Localization-Delocalization Phase Transition, *Phys. Rev. X* **5**, 041047 (2015).
- [6] D. Abanin, J. Bardarson, G. De Tomasi, S. Gopalakrishnan, V. Khemani, S. Parameswaran, F. Pollmann, A. Potter, M. Serbyn, and R. Vasseur, Distinguishing localization from chaos: Challenges in finite-size systems, *Ann. Phys.* **427**, 168415 (2021).
- [7] D. Sels and A. Polkovnikov, Dynamical obstruction to localization in a disordered spin chain, *Phys. Rev. E* **104**, 054105 (2021).
- [8] F. Pietracaprina, N. Macé, D. J. Luitz, and F. Alet, Shift-invert diagonalization of large many-body localizing spin chains, *SciPost Phys.* **5**, 045 (2018).
- [9] P. Sierant, M. Lewenstein, and J. Zakrzewski, Polynomially Filtered Exact Diagonalization Approach to Many-Body Localization, *Phys. Rev. Lett.* **125**, 156601 (2020).
- [10] M. Kiefer-Emmanouilidis, R. Unanyan, M. Fleischhauer, and J. Sirker, Slow delocalization of particles in many-body localized phases, *Phys. Rev. B* **103**, 024203 (2021).
- [11] M. Kiefer-Emmanouilidis, R. Unanyan, M. Fleischhauer, and J. Sirker, Unlimited growth of particle fluctuations in many-body localized phases, *Ann. Phys.* **435**, 168481 (2021).
- [12] N. Garron and K. Langfeld, Controlling the sign problem in finite-density quantum field theory, *Eur. Phys. J. C* **77**, 470 (2017).
- [13] N. Garron and K. Langfeld, Tackling the sign problem with a moment expansion and application to heavy dense QCD, *PoS(LATTICE2016)* **256**, 084 (2017).
- [14] L. Bongiovanni, K. Langfeld, B. Lucini, R. Pellegrini, and A. Rago, The density of states approach at finite chemical potential: A numerical study of the Bose gas, *PoS(LATTICE 2015)* **251**, 192 (2016).
- [15] O. Francesconi, M. Holzmann, B. Lucini, and A. Rago, Free energy of the self-interacting relativistic lattice Bose gas at finite density, *Phys. Rev. D* **101**, 014504 (2020).
- [16] F. Wang and D. P. Landau, Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States, *Phys. Rev. Lett.* **86**, 2050 (2001).
- [17] R. E. Belardinelli and V. D. Pereyra, Wang-Landau algorithm: a theoretical analysis of the saturation of the error, *J. Chem. Phys.* **127**, 184105 (2007).
- [18] R. E. Belardinelli, S. Manzi, and V. D. Pereyra, Analysis of the convergence of the $1/t$ and Wang-Landau algorithms in the calculation of multidimensional integrals, *Phys. Rev. E* **78**, 067701 (2008).
- [19] H. Robbins and S. Monro, A stochastic approximation method, *Ann. Math. Stat.* **22**, 400 (1951).
- [20] K. Langfeld, B. Lucini, R. Pellegrini, and A. Rago, An efficient algorithm for numerical computations of continuous densities of states, *Eur. Phys. J. C* **76**, 306 (2016).
- [21] M. Körner, K. Langfeld, D. Smith, and L. von Smekal, Density of states approach to the hexagonal Hubbard model at finite density, *Phys. Rev. D* **102**, 054502 (2020).
- [22] M. Guagnelli, Sampling the density of states, [arXiv:1209.4443](https://arxiv.org/abs/1209.4443).
- [23] J. Cotler, N. Hunter-Jones, J. Liu, and B. Yoshida, Chaos, complexity, and random matrices, *J. High Energy Phys.* **11** (2017) 048.
- [24] H. Gharibyan, M. Hanada, S. H. Shenker, and M. Tezuka, Onset of random matrix behavior in scrambling systems, *J. High Energy Phys.* **07** (2018) 124.
- [25] J. Liu, Spectral form factors and late time quantum chaos, *Phys. Rev. D* **98**, 086026 (2018).
- [26] A. Prakash, J. H. Pixley, and M. Kulkarni, Universal spectral form factor for many-body localization, *Phys. Rev. Res.* **3**, 012019 (2021).
- [27] R. Vosk, D. A. Huse, and E. Altman, Theory of the Many-Body Localization Transition in One-Dimensional Systems, *Phys. Rev. X* **5**, 031032 (2015).
- [28] D. J. Luitz, N. Laflorencie, and F. Alet, Many-body localization edge in the random-field Heisenberg chain, *Phys. Rev. B* **91**, 081103(R) (2015).
- [29] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [30] I. V. Gornyi, A. D. Mirlin, and D. G. Polyakov, Interacting Electrons in Disordered Wires: Anderson Localization and Low- t Transport, *Phys. Rev. Lett.* **95**, 206603 (2005).
- [31] D. Basko, I. Aleiner, and B. Altshuler, Metal-insulator transition in a weakly interacting many-electron system with localized single-particle states, *Ann. Phys.* **321**, 1126 (2006).
- [32] W. De Roeck and F. Huveneers, Stability and instability towards delocalization in many-body localization systems, *Phys. Rev. B* **95**, 155129 (2017).
- [33] P. Sierant and J. Zakrzewski, Challenges to observation of many-body localization, *Phys. Rev. B* **105**, 224203 (2022).

- [34] Y. Atia and D. Aharonov, Fast-forwarding of Hamiltonians and exponentially precise measurements, *Nat. Commun.* **8**, 1572 (2017).
- [35] Such a system can also be interpreted as an anisotropic classical Ising model with complex coupling coefficients.
- [36] These are not exactly the conditions from Ref. [19], but they are equivalent to this simpler version.
- [37] We use the SMOOTH.SPLINE implementation in R [38].
- [38] R Core Team, *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, Vienna, Austria, 2020).
- [39] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. Van Der Linde, Bayesian measures of model complexity and fit, *J. R. Stat. Soc. B* **64**, 583 (2002).
- [40] A. Nemirovskii and D. Yudin, *Problem Complexity and Method Efficiency in Optimization* (John Wiley & Sons, New York, 1983).
- [41] J. Ostmeier, j-ostmeier/real-time-dos: Code and data along with [arXiv:2209.13970v1](https://arxiv.org/abs/2209.13970v1), doi: [10.5281/zenodo.7164902](https://doi.org/10.5281/zenodo.7164902) (2022).
- [42] B. B. Beard and U.-J. Wiese, Simulations of Discrete Quantum Systems in Continuous Euclidean Time, *Phys. Rev. Lett.* **77**, 5130 (1996).
- [43] K. Langfeld and B. Lucini, The density of states approach to dense quantum systems, *Phys. Rev. D* **90**, 094502 (2014).
- [44] N. Garron and K. Langfeld, Anatomy of the sign-problem in heavy-dense QCD, *Eur. Phys. J. C* **76**, 569 (2016).
- [45] H. Schlüter, F. Gayk, H.-J. Schmidt, A. Honecker, and J. Schnack, Accuracy of the typicality approach using Chebyshev polynomials, *Z. Naturforsch. A* **76**, 823 (2021).
- [46] Z.-G. Mou, P. M. Saffin, A. Tranberg, and S. Woodward, Real-time quantum dynamics, path integrals and the method of thimbles, *J. High Energy Phys.* **06** (2019) 094.
- [47] A. Alexandru, G. Basar, P. F. Bedaque, S. Vartak, and N. C. Warrington, Monte Carlo Study of Real Time Dynamics on the Lattice, *Phys. Rev. Lett.* **117**, 081602 (2016).
- [48] A. Alexandru, G. Basar, P. F. Bedaque, and G. W. Ridgway, Schwinger-Keldysh on the lattice: a faster algorithm and its application to field theory, *Phys. Rev. D* **95**, 114501 (2017).
- [49] J.-L. Wynen, E. Berkowitz, S. Krieg, T. Luu, and J. Ostmeier, Machine learning to alleviate Hubbard-model sign problems, *Phys. Rev. B* **103**, 125153 (2021).
- [50] M. Rodekamp, E. Berkowitz, C. Gäntgen, S. Krieg, T. Luu, and J. Ostmeier, Mitigating the Hubbard sign problem with complex-valued neural networks, *Phys. Rev. B* **106**, 125139 (2022).
- [51] www.dirac.ac.uk.
- [52] B. Kostrzewa, J. Ostmeier, M. Ueding, and C. Urbach, Hadron: Analysis framework for Monte Carlo simulation data in physics (2020), r package version 3.1.2.