

Stable recursive auxiliary field quantum Monte Carlo algorithm in the canonical ensemble: Applications to thermometry and the Hubbard model

Tong Shen¹, Hatem Barghathi,² Jiangyong Yu,³ Adrian Del Maestro^{1,2,4}, and Brenda M. Rubenstein^{1,3,*}

¹Department of Chemistry, Brown University, Providence, Rhode Island 02912, USA

²Department of Physics and Astronomy, University of Tennessee, Knoxville, Tennessee 37916, USA

³Department of Physics, Brown University, Providence, Rhode Island 02912, USA

⁴Min H. Kao Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, Tennessee 37996, USA



(Received 22 December 2022; accepted 7 March 2023; published 9 May 2023)

Many experimentally accessible, finite-sized interacting quantum systems are most appropriately described by the canonical ensemble of statistical mechanics. Conventional numerical simulation methods either approximate them as being coupled to a particle bath or use projective algorithms which may suffer from nonoptimal scaling with system size or large algorithmic prefactors. In this paper, we introduce a highly stable, recursive auxiliary field quantum Monte Carlo approach that can directly simulate systems in the canonical ensemble. We apply the method to the fermion Hubbard model in one and two spatial dimensions in a regime known to exhibit a significant “sign” problem and find improved performance over existing approaches including rapid convergence to ground-state expectation values. The effects of excitations above the ground state are quantified using an estimator-agnostic approach including studying the temperature dependence of the purity and overlap fidelity of the canonical and grand canonical density matrices. As an important application, we show that thermometry approaches often exploited in ultracold atoms that employ an analysis of the velocity distribution in the grand canonical ensemble may be subject to errors leading to an underestimation of extracted temperatures with respect to the Fermi temperature.

DOI: [10.1103/PhysRevE.107.055302](https://doi.org/10.1103/PhysRevE.107.055302)

I. INTRODUCTION

In many settings in condensed matter physics, the grand canonical ensemble, in which the number of particles in a system is allowed to fluctuate subject to a fixed chemical potential, is the ensemble of choice for modeling systems at finite temperature. This is a natural framework due to the approach to the thermodynamic limit for electrons in solids or the existence of a particle (or quasiparticle) reservoir in transport geometries, heterostructures, and superconductors. However, there is a growing number of important scenarios in which the number of particles is fixed and small, including trapped atom systems composed of a finite number of atoms confined in box-style traps [1–3], nuclear systems with a fixed number of nucleons [4–8], and molecules containing a fixed number of electrons [9,10]. All such systems are more accurately described by the canonical ensemble in which the number of particles cannot fluctuate. Moreover, many ground-state algorithms in condensed matter are formulated in the canonical ensemble [11–14] and finite-temperature algorithms [10,15–18] that can converge to these algorithms’ ground-state results without spurious particle-number fluctuations can shed a brighter light on the mechanisms behind low-temperature quantum phase transitions and crossovers. Other examples where systems need to be treated within

the canonical ensemble include determining the operationally accessible entanglement in indistinguishable many-body systems in the presence of a $U(1)$ superselection rule limiting physically allowable operations [19–21] and the determination of thermonuclear rates for astrophysics [22].

Efficiently describing interacting systems in the canonical ensemble has nevertheless been a longstanding challenge, particularly for second-quantized algorithms. Unlike in the grand canonical ensemble, in which partition functions and other quantities can be evaluated without placing any constraints on the number of particles [23], evaluating quantities in the canonical ensemble requires an explicit consideration of particle-number constraints. At a physical level, these constraints give rise to interesting, nontrivial correlations among the occupations of different states—higher-order expectation values of occupation numbers do not factorize, even in the noninteracting limit [24]. However, at a numerical level, these constraints can make the analytical and computational evaluation of canonical ensemble quantities substantially more cumbersome [25].

One approach for modeling interacting systems in the canonical ensemble that circumvents the imposition of direct constraints is the use of projection techniques [8,26]. In these algorithms, canonical ensemble quantities are projected out from the grand canonical partition function at a suitably tuned chemical potential [27]. This approach has been fruitfully employed to study a wide variety of problems in nuclear physics [5,8] and, more recently, condensates [28,29]. Nonetheless, because this algorithm relies on projecting out of

*Author to whom correspondence should be addressed:
brenda_rubenstein@brown.edu

the grand canonical ensemble, it is accompanied by the same computational overhead as typical finite-temperature grand canonical simulations and can develop numerical instabilities for large particle numbers or when reasonable chemical potentials cannot be identified, for example near first-order phase transitions [30]. Recently, techniques for rapidly determining a chemical potential where it can be readily identified have been proposed, but these techniques still do not inherently operate in the canonical ensemble [27,31,32]. Approaches that directly take physical constraints into account therefore have the potential to lead to methods that are not only more stable but also more computationally efficient.

Recently, a new auxiliary partition function (APF) formalism has been proposed that enables the recursive computation of N -particle partition functions and related quantities for noninteracting systems from smaller particle-number quantities, thus explicitly taking particle-number constraints into account [24]. In essence, the APF formalism views the canonical ensemble partition function as a sum over the probabilities of varying numbers of particles occupying different subsets of states. Unlike previous such recursions for noninteracting systems [25], the APF formalism is able to arrive at expressions for canonical ensemble partition functions using only positive subquantities, and hence it is significantly more numerically stable. While this technique has been successfully applied to such noninteracting systems as harmonic oscillators and rings of bosons [24], the APF algorithm and alternative recursions have yet to be generalized to interacting systems.

In this paper, we present a new, highly stable recursive algorithm for the simulation of interacting systems in the canonical ensemble. This algorithm marries the finite-temperature auxiliary field quantum Monte Carlo (AFQMC) algorithm [33,34], which has long been used to model finite-temperature interacting systems, with the new APF formalism [24], which has previously only been applied to noninteracting systems. The key realization that enables this marriage is that the Hubbard-Stratonovich (HS) transformation [35–37], which reconstructs the properties of interacting systems by integrating over an appropriately weighted ensemble of noninteracting systems [38], can be exploited to construct interacting partition functions and related observables in the canonical ensemble by integrating over noninteracting APF partition functions. Our algorithm can therefore stably describe interacting systems in the canonical ensemble by sampling noninteracting partition functions and other finite-temperature quantities generated using the APF formalism and then integrating over those samples. This markedly improves on our previous work which was built on the significantly less stable Borrman recursion algorithm for noninteracting gases [25,39]. To highlight the stability of our new algorithm, we show that: (1) our new interacting algorithm is stable down to substantially lower temperatures than previous algorithms and (2) it has a lower computational scaling than previous projection algorithms.

With this highly stable algorithm, we proceed to analyze differences in the convergence of the energy, sign, and information-theoretic measures such as the purity and fidelity [40,41] to the ground state between the grand canonical and canonical ensembles. To do so, we focus on the Hubbard model of interacting fermions in one and two dimensions

as an instructive example because it manifests the strong correlation often resulting in a sign problem that is hardest to model via modern simulation techniques. We find that because higher-energy states are more readily accessed in the grand canonical ensemble, grand canonical energies tend to be higher and purities lower at any given temperature, meaning that the grand canonical ensemble converges more slowly to the ground state. We substantiate these findings with analytical expressions describing how these quantities should converge to the ground state in both the interacting and noninteracting limits. We furthermore demonstrate that these differences have substantial practical implications for the thermometry of cold atom systems: If the temperature of cold atom systems containing a fixed number of particles is estimated based on the grand canonical ensemble, then this leads to temperature predictions that can be as much as 53.2% lower than in the more realistic canonical ensemble picture according to the analysis we present below.

We begin in Sec. II by presenting our new algorithm and its underlying formalism, showing how recursive relations for the partition function and one- and two-body quantities can be determined using the APF method and subsequently integrated into the AFQMC algorithm. We also demonstrate the formal relationship between our recursive algorithm and the previously used projection algorithm. In Sec. III, we present our results regarding the increased stability of our algorithm, before illustrating the differences in system energies, purities, and fidelities as measured in the two ensembles using the Hubbard model as a salient example. We exemplify the practical consequences of these differences for thermometry in Sec. III D. Last, we conclude in Sec. IV by discussing further applications of our algorithm and its straightforward extension to studying nuclear matter and bosons, for which it has the potential to show even greater efficiency gains.

II. FORMALISM

A. The finite-temperature AFQMC algorithm

The central quantity in finite-temperature theories is the partition function, from which all other properties can be derived. Historically, the focus of finite-temperature methods has been to obtain or otherwise sample from the grand partition function, \mathcal{Z}_μ , associated with the grand canonical ensemble, in which the internal energy and particle number are allowed to fluctuate around average values that can be tuned by the temperature T and chemical potential μ , respectively [23].

The grand partition function can be expressed as the trace

$$\mathcal{Z}_\mu = \text{Tr}[e^{-\beta(\hat{H}-\mu\hat{N})}], \quad (1)$$

where $\beta = 1/(k_B T)$, k_B denotes the Boltzmann constant, \hat{H} denotes the many-body Hamiltonian, and \hat{N} is an operator corresponding to the total number of particles which is the sum over the occupations over the set of (possibly degenerate) states. In order to facilitate its subsequent sampling, the grand partition function can be discretized into L imaginary time slices, each of which can then be approximately factored into short imaginary time one- and two-body propagators via a

Suzuki-Trotter factorization [42,43],

$$\begin{aligned} \mathcal{Z}_\mu &= \text{Tr} \left[\prod_{l=1}^L e^{-\Delta\tau(\hat{H}-\mu\hat{N})} \right] \\ &\approx \text{Tr} \left[\prod_{l=1}^L [e^{-\Delta\tau\hat{K}/2} e^{-\Delta\tau\hat{V}} e^{-\Delta\tau\hat{K}/2}] \right], \end{aligned} \quad (2)$$

where $\Delta\tau = \beta/L$ represents an imaginary time slice, \hat{K} is the collection of all one-body operators, and \hat{V} is the collection of all two-body operators. The exact grand partition function is recovered in the limit $\Delta\tau \rightarrow 0$. This factorization enables us to ignore the noncommutation of the one-body and two-body operators up to a systematic and controllable error of $O(\Delta\tau^3)$.

While one-body propagators, $e^{-\Delta\tau\hat{K}}$, may be neatly expressed as matrices in a given basis [44], two-body propagators, $e^{-\Delta\tau\hat{V}}$, may not be as easily determined. In auxiliary field-based methods, including determinant quantum Monte Carlo [33,45] and AFQMC [11,46], two-body propagators of the form $e^{-\Delta\tau\hat{V}}$ are linearized by re-expressing them as integrals over one-body propagators [35] using the Hubbard-Stratonovich transformation [35–37,47]. Assuming as we will below that the two-body propagator can be written as

$$\hat{V} = \sum_\gamma \lambda_\gamma \hat{v}_\gamma^2, \quad (3)$$

where the \hat{v}_γ denote linear combinations of one-body operators and the λ_γ denote their contributions to the sum, then

$$e^{-\Delta\tau\hat{V}} = \int D\sigma G_\sigma \hat{V}_\sigma \quad (4)$$

according to the HS transformation, where σ represents an auxiliary field, G_σ represents the Gaussian probability of sampling that field, and \hat{V}_σ denotes the collection of one-body operators as a function of that field. Crucial to this paper, Eq. (4) signifies that propagators for interacting systems can be rewritten as integrals over propagators for noninteracting systems. Based on this transform, all of the one-body operators from different time slices, i , are then combined, i.e., $\hat{U}_\sigma = \prod_i e^{-\Delta\tau\hat{K}/2} \hat{V}_{\sigma_i} e^{-\Delta\tau\hat{K}/2}$, and the full grand canonical partition function may be expressed as

$$\begin{aligned} \mathcal{Z}_\mu &= \text{Tr} \left(\int D\sigma G_\sigma e^{\beta\mu\hat{N}} \hat{U}_\sigma \right) \\ &= \int D\sigma G_\sigma \text{Tr} (e^{\beta\mu\hat{N}} \hat{U}_\sigma) \\ &= \int D\sigma G_\sigma \det(I + e^{\beta\mu} \mathbf{U}_\sigma), \end{aligned} \quad (5)$$

where taking the trace over all fermion occupations results in a determinant [44].

The partition function can then be sampled to evaluate such observables as energies, average occupations, and correlation functions [33,34]. In particular, Wick's theorem is valid in the grand-canonical ensemble, which enables a powerful simplification of expectation values of products of operators (e.g., correlation functions) into factorized sums and differences of shorter products of expectation values of those operators [48,49].

B. Recursive relations for the canonical partition function

While computing properties in the grand canonical ensemble is appropriate for many systems and can be analytically or computationally convenient, in many situations in which the particle number remains fixed, a canonical treatment is more suitable. Computing the canonical ensemble partition function proceeds along the same lines as computing the grand canonical one with the critical exception that the trace must be taken with the constraint of fixed particle number N . More specifically, the N -particle, canonical ensemble partition function may be expressed as

$$Z_N = \text{Tr}_N (e^{-\beta\hat{H}}), \quad (6)$$

which can be factored and transformed in a similar manner to the grand partition function to obtain

$$Z_N = \int D\sigma G_\sigma \text{Tr}_N (\hat{U}_\sigma), \quad (7)$$

where we have added the subscript N to the trace to differentiate it from that in Eq. (1). Because \hat{U}_σ is a one-body operator, its matrix form, \mathbf{U}_σ , can be diagonalized in the single-particle space

$$\mathbf{U}_\sigma = \mathbf{P} \Lambda \mathbf{P}^{-1}, \quad (8)$$

where we omit the σ dependence on the right side for clarity. We then introduce the effective single-particle spectrum $\Lambda = \text{diag}(\{\lambda_\gamma\}) = \text{diag}(\{\exp(-\beta\tilde{\epsilon}_\gamma)\})$, based on the following relations

$$\hat{U}_\sigma = \sum_\gamma \exp(-\beta\hat{a}_\gamma^\dagger \tilde{\epsilon}_\gamma \hat{a}_\gamma) = \sum_\gamma \lambda_\gamma \hat{n}_\gamma, \quad (9)$$

and the basis transformation

$$\hat{a}_\gamma^\dagger = \sum_i \langle i | \gamma \rangle \hat{a}_i^\dagger, \quad \hat{a}_\gamma = \sum_i \langle \gamma | i \rangle \hat{a}_i, \quad \hat{n}_\gamma = \hat{a}_\gamma^\dagger \hat{a}_\gamma. \quad (10)$$

Since \hat{U}_σ is an independent-particle propagator that only depends on the auxiliary field vector, σ , the effective single-particle spectrum, $\{\tilde{\epsilon}_\gamma\}$, is independent of the particle number. For an N -particle, N_s -site system, taking the trace while constraining the particle-number yields [39]

$$\begin{aligned} \text{Tr}_N (\hat{U}_\sigma) &= \text{Tr}_N \left[\sum_\gamma \exp(-\beta\hat{a}_\gamma^\dagger \tilde{\epsilon}_\gamma \hat{a}_\gamma) \right] \\ &= \sum_{\Gamma_N} \langle \Gamma_N | \sum_\gamma \exp(-\beta\hat{a}_\gamma^\dagger \tilde{\epsilon}_\gamma \hat{a}_\gamma) | \Gamma_N \rangle \\ &= \sum_{\Gamma_N} \sum_\gamma \lambda_\gamma^{n_\gamma}. \end{aligned} \quad (11)$$

Here, Γ_N is used to represent the set of N -particle states, and thus $\sum_{\Gamma_N} \equiv \sum_{n_1+\dots+n_{N_s}=N}$ and n_γ denotes the number of particles in the γ th eigenstate. For fermions, $n_\gamma = 0, 1$. The key implication of Eq. (11) is that, for a specified field σ , the single-particle spectrum can be decoupled from the particle number. Hence, the many-particle energy given such fields is simply the sum of all of the single-particle energies.

This key fact enables us to move beyond previous projection-based approaches and calculate Eq. (11) in a recursive fashion, where we utilize the recursive approach to

calculating canonical ensemble partition functions first developed for ideal gases. Specifically, the partition function can be obtained using the well-known Borrmann recursion [25,50],

$$Z_N = \frac{1}{N} \sum_{k=1}^N (\zeta)^{k-1} z_k Z_{N-k}, \quad (12)$$

where $z_k = \sum_{j=1}^{N_s} \lambda_j^k$ and ζ takes the values of -1 and 1 for fermions and bosons, respectively. However, in the context of noninteracting gases, the fermionic version of Eq. (12) is known to suffer from numerical instabilities [51,52], which were also encountered in a previous version of our canonical AFQMC algorithm [39], leading to the emergence of an additional *unphysical* sign problem.

We can gain some intuitive understanding of the reasons behind such numerical instability for the case of fermions at low temperatures. In addition to the alternating signs in Eq. (12), the contribution of high energy levels to the factors z_k can be filtered out by limited numerical precision. This is expected to have a minimal effect on low-temperature Bose gases, as their thermodynamic properties rely heavily on the occupation of low energy levels. In contrast, the thermodynamic properties of low-temperature Fermi gases are governed by much higher energy levels in the vicinity of the Fermi level.

Here, we build on our previous work [39] to propose a more numerically stable and accurate method for computing the canonical trace via a recursive formula based on the recently developed auxiliary partition function formalism [24]. Given a set $\{\lambda_i = e^{-\beta\tilde{\epsilon}_i}\}$ of Boltzmann factors that correspond to the single-particle energy spectrum $\{\tilde{\epsilon}_i\}$, we can build the desired N -particle partition function recursively by including one of the levels in each recursive step, i.e.,

$$Z_N = \lambda_j Z_{N-1}^{\{\lambda_i\} \setminus \lambda_j} + Z_N^{\{\lambda_i\} \setminus \lambda_j}, \quad (13)$$

where the notation $\{\lambda_i\} \setminus \lambda_j$ implies that we exclude the specific level j from the set $\{\lambda_i\}$ (see Ref. [24] for extensive details). This recursion does not suffer from the problems of Eq. (12) because it is inherently positive. Yet, the unbounded nature of the Boltzmann factors λ_i can make achieving the desired numerical precision difficult, especially at low temperatures. This can be addressed with a simple trick. We modify Eq. (13) by inserting an arbitrary multiplicative factor A_j (to be determined below) with the inclusion of each λ_j and apply the modified equation on the modified set $\{B\lambda_i\}$, where B is an additional constant (also to be determined). This results in a modified recursion relation

$$\bar{Z}_N = A_j (B\lambda_j \bar{Z}_{N-1}^{\{\lambda_i\} \setminus \lambda_j} + \bar{Z}_N^{\{\lambda_i\} \setminus \lambda_j}), \quad (14)$$

and we can recover Z_N from the resulting \bar{Z}_N via

$$\bar{Z}_N = B^N Z_N \prod_j A_j. \quad (15)$$

This suggests that we can enhance the performance of the APF recursive approach for calculating Z_N through a clever choice of the constants $\{A_j\}$ and B . In order to do so, we rearrange the fugacity expansion for the grand-canonical partition function, \mathcal{Z}_μ , in terms of the canonical partition functions,

$$Z_N = \text{Tr}_N e^{-\beta\hat{H}},$$

$$\mathcal{Z}_\mu = \sum_N e^{\beta\mu N} Z_N. \quad (16)$$

Dividing by the grand canonical partition function and removing the summation over N results in an expression for the particle-number probability distribution given by

$$P_\mu(N) = \frac{e^{\beta\mu N} Z_N}{\mathcal{Z}_\mu}. \quad (17)$$

For noninteracting fermions, $\mathcal{Z}_\mu = \prod_j [1 - p_j^{(\mu)}]^{-1}$ [23], where

$$p_j^{(\mu)} = \frac{e^{\beta\mu} \lambda_j}{1 + e^{\beta\mu} \lambda_j} \quad (18)$$

is the probability of occupying the j th energy level. This yields

$$P_\mu(N) = e^{\beta\mu N} Z_N \prod_j [1 - p_j^{(\mu)}]. \quad (19)$$

If we compare Eq. (19) with Eq. (15), then we can identify $P_\mu(N)$ with \bar{Z}_N by choosing $A_j = 1 - p_j^{(\mu)}$ and $B = e^{\beta\mu}$, which, when substituted into Eq. (14), results in a recursion relation for the number probability distribution

$$P_\mu(N) = p_j^{(\mu)} P_\mu^{\{\lambda_i\} \setminus \lambda_j}(N-1) + [1 - p_j^{(\mu)}] P_\mu^{\{\lambda_i\} \setminus \lambda_j}(N). \quad (20)$$

In contrast with Eq. (13), all terms in the above equation are bounded between 0 and 1, which further ensures their numerical stability and automatically avoids numerical arithmetic overflow issues caused by extremely large λ_i values. Also, setting $N = 0$ in Eq. (17), we see that $1/\mathcal{Z}_\mu = P_\mu(0)$, which enables us to re-express Eq. (17) as

$$Z_N = \frac{Z_N}{Z_0} = \frac{e^{\beta\mu N} P_\mu(N)}{P_\mu(0)}, \quad (21)$$

where $Z_0 = 1$.

We note that in Eq. (21), the chemical potential μ is an algorithmic parameter that can take on any value without changing the value of the canonical trace. It need not be the many-body chemical potential, which is otherwise difficult to determine for an arbitrary system. To increase the numerical stability of Eqs. (20) and (21), it is best to select a μ around the Fermi level, where P_μ peaks at N . A good choice is $e^{\beta\mu} = |\lambda_N \lambda_{N+1}|^{1/2}$, assuming that $\{\lambda_i\}$ is sorted as $|\lambda_1| < |\lambda_2| < \dots < |\lambda_{N_s}|$.

It is worth mentioning that the particle-number distribution $P_\mu(N)$ can be viewed as a Poisson-binomial distribution [53], which can be expressed as

$$P_\mu(N) := \sum_{\mathcal{S}_N} \prod_{i \in \mathcal{S}_N} p_i^{(\mu)} \prod_{j \in \bar{\mathcal{S}}_N} [1 - p_j^{(\mu)}], \quad (22)$$

where it is constructed from the probability $p_i^{(\mu)}$ of successfully occupying N energy levels at a given chemical potential and temperature out of a total number of independent (non-interacting) and nonidentical Bernoulli trials [24], where the corresponding set of independent success probabilities is represented by $\{p_i^{(\mu)}\}$. Here \mathcal{S}_N is a set of N occupied energy levels selected from N_s energy levels in the single-particle

space; $\tilde{\mathcal{S}}_N$ denotes the complement coming from unoccupied levels.

While we have focused our presentation of our formalism here on fermions, it can readily be extended to bosons, as further detailed in the supplementary material [54].

C. Relationship between auxiliary partition function approach and previous projection approaches

In this section, we illustrate how the more conventional particle-number projection formalism [8,59] can be derived from Eqs. (18)–(20) through a Fourier transform and hence that it is analytically equivalent to the recursive approach. The fact that these approaches are analytically equivalent will be of value when we compare the accuracy, stability, and speed of these methods in Sec. III A.

We begin by considering the generating function of the partition function under the recursive relation, Eq. (13), for any nonzero a

$$\begin{aligned} \sum_{N=0}^{\infty} a^N Z_N &= \lambda_j \sum_{N=1}^{\infty} a^N Z_{N-1}^{\{\lambda_j\} \setminus \lambda_j} + \sum_{N=0}^{\infty} a^N Z_N^{\{\lambda_j\} \setminus \lambda_j} \\ &= (1 + a\lambda_j) \sum_{N=0}^{\infty} a^N Z_N^{\{\lambda_j\} \setminus \lambda_j}. \end{aligned} \quad (23)$$

By iteratively subtracting the Boltzmann factors λ_j from the set of all levels $\{\lambda_i\}$, we arrive at

$$\sum_{N=0}^{\infty} a^N Z_N = \prod_j (1 + a\lambda_j). \quad (24)$$

Setting $a = e^{\beta\mu} e^{i\phi_m}$, with $\phi_m = 2\pi m/N_s$, then yields

$$\sum_{N=0}^{\infty} e^{\beta\mu N} e^{i\phi_m N} Z_N = \prod_j (1 + e^{\beta\mu} e^{i\phi_m} \lambda_j). \quad (25)$$

Using the discrete Fourier transform of the delta function at N , $\delta_N = \sum_{m=1}^{N_s} e^{i\phi_m N} / N_s$, we can solve for $e^{\beta\mu N} Z_N$ and recover the particle-number projection result for the canonical trace that was first proposed in Ref. [8],

$$\begin{aligned} Z_N &= \frac{e^{-\beta\mu N}}{N_s} \sum_{m=1}^{N_s} e^{-i\phi_m N} \prod_j (1 + e^{\beta\mu} e^{i\phi_m} \lambda_j) \\ &= \frac{1}{N_s} \sum_{m=1}^{N_s} e^{-\beta\mu N} e^{-i\phi_m N} \tilde{\mathcal{Z}}_\mu(m). \end{aligned} \quad (26)$$

As we shall demonstrate below, the recursion formalism and the projection formalism have equivalent accuracy with recursion having slightly improved scaling [$O(N_s^3 + N_s N)$ vs $O(N_s^3 + N_s^2)$ after considering the N_s^3 cost of the eigendecomposition] for computing the partition function.

D. Recursive computation of density matrices and correlation functions

The expectation value of the one-body density operator can be evaluated using the same eigendecomposition of the field-dependent propagator matrix, $\mathbf{U}_\sigma = \mathbf{P} \Lambda \mathbf{P}^{-1}$, with $\Lambda =$

$\text{diag}(\{\lambda_i\})$, which leads to

$$\langle \hat{c}_i^\dagger \hat{c}_j \rangle_N = \sum_{\alpha} \mathbf{P}_{i\alpha} \langle \hat{n}_\alpha \rangle_N \mathbf{P}_{\alpha j}^{-1}, \quad (27)$$

where $\langle \hat{n}_\alpha \rangle_N$ is computed recursively in $O(N)$ operations [50,60] as

$$\langle \hat{n}_\alpha \rangle_N = \frac{\lambda_\alpha P_\mu(N-1)}{e^{\beta\mu} P_\mu(N)} (1 - \langle \hat{n}_\alpha \rangle_{N-1}) \quad (28)$$

with the initial condition $\langle \hat{n}_\alpha \rangle_0 = 0$. When λ_α is large, $\langle \hat{n}_\alpha \rangle_N$ is close to 1 and the right-hand side of Eq. (28) may develop numerical round-off errors. To avoid this, we reverse the recursion and compute the “hole” distribution function for large values of λ_α ,

$$1 - \langle \hat{n}_\alpha \rangle_N = \frac{e^{\beta\mu} P_\mu(N+1)}{\lambda_\alpha P_\mu(N)} \langle \hat{n}_\alpha \rangle_{N+1} \quad (29)$$

with the initial condition $\langle \hat{n}_\alpha \rangle_{N_s} = 1$ at unit filling.

The expectation value of the two-body density operator can be computed in a similar fashion

$$\begin{aligned} \langle \hat{c}_i^\dagger \hat{c}_j \hat{c}_k^\dagger \hat{c}_l \rangle_N &= \sum_{\alpha, \beta} [\mathbf{P}_{i\alpha} \mathbf{P}_{k\beta} \langle \hat{n}_\alpha \hat{n}_\beta \rangle_N \mathbf{P}_{\beta l}^{-1} \mathbf{P}_{\alpha j}^{-1} \\ &\quad + \mathbf{P}_{i\alpha} \mathbf{P}_{k\beta} \langle \hat{n}_\alpha - \hat{n}_\alpha \hat{n}_\beta \rangle_N \mathbf{P}_{\beta j}^{-1} \mathbf{P}_{\alpha l}^{-1}] \end{aligned} \quad (30)$$

and the two-level correlations can be expressed as

$$\langle \hat{n}_\alpha \hat{n}_\beta \rangle_N = \begin{cases} \langle \hat{n}_\alpha \rangle_N & \alpha = \beta \\ \frac{\lambda_\beta \langle \hat{n}_\alpha \rangle_N - \lambda_\alpha \langle \hat{n}_\beta \rangle_N}{\lambda_\beta - \lambda_\alpha} & \alpha \neq \beta \end{cases}, \quad (31)$$

which can be viewed as the lowest-order canonical ensemble generalization [51,61] of Wick’s theorem [49], recently extended to the case of degenerate spectra [24]. These expressions are used throughout the rest of the paper to compute energies, particle densities, and correlation functions.

E. Model system: Fermion Hubbard model

While our formalism generalizes to any two-body Hamiltonian, for the sake of subsequent discussion, we will focus on the fermion Hubbard model due to the strong correlation it exhibits and its relevance for the description of many useful chemical and material systems. This model’s Hamiltonian may be expressed as

$$\hat{H} = -t \sum_{ij,\sigma} (\hat{c}_{i,\sigma}^\dagger \hat{c}_{j,\sigma} + \text{H.c.}) + U \sum_i \hat{n}_{i,\uparrow} \hat{n}_{i,\downarrow}, \quad (32)$$

where $\hat{c}_{i,\sigma}^\dagger (\hat{c}_{i,\sigma})$ are anticommuting fermionic creation (annihilation) operators such that $\hat{c}_i^\dagger \hat{c}_j^\dagger + \hat{c}_i^\dagger \hat{c}_j = \delta_{ij}$ and $\hat{n}_{i,\sigma} = \hat{c}_{i,\sigma}^\dagger \hat{c}_{i,\sigma}$ is the local spin-resolved density for hopping parameter t and interaction strength U . In our subsequent illustrations, we pose different challenges to our formalism by varying the strength of the electron correlation, U/t , the filling (average number of electrons per site), $\langle n \rangle = (\sum_i \hat{n}_{i,\uparrow} + \hat{n}_{i,\downarrow})/N_s$, and number of sites, N_s , in our model. We measure energies in units of the hopping parameter t in our simulations and set $t = 1$ in the remainder of this work.

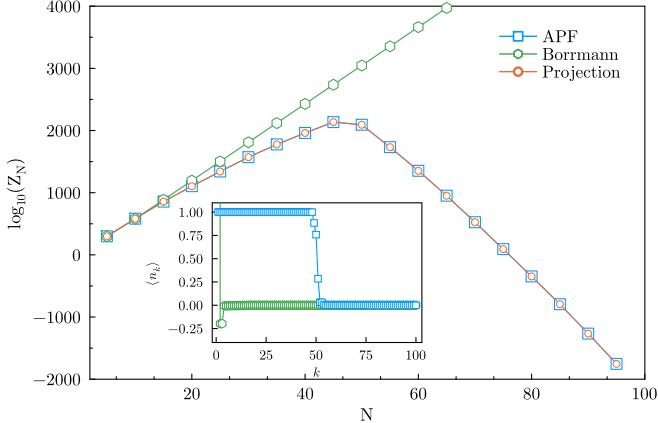


FIG. 1. Stability with which different canonical ensemble algorithms can compute the partition function of a 100-site Hubbard chain with a randomly generated set of Hubbard-Stratonovich fields for $\beta = 10$ for different total particle numbers N . The inset depicts the occupation numbers of all of the energy levels at half-filling ($N = 50$).

III. RESULTS AND DISCUSSION

A. Accuracy, stability, and speed of our recursive approach

In order to assess the utility of our approach, we begin by characterizing its accuracy, stability, and relative speed. Because our APF-based algorithm no longer involves taking the difference between products of probabilities or partition functions, we expect it to be substantially more stable at low temperatures, enabling us to model larger systems closer to their ground states. While we have previously demonstrated this increased stability for noninteracting spin models [24], here we demonstrate that this algorithm is equally stable for interacting systems.

To illustrate the stability of our algorithm for a non-interacting system directly relevant to our final interacting simulations, we compute the canonical partition functions and occupations of a version of the Hubbard model whose solutions can be derived analytically. As described in Sec. II A, in order to fully account for the two-body interactions in the AFQMC formalism, one employs the Hubbard-Stratonovich transformation to integrate over many possible random instantiations of noninteracting systems parameterized by different auxiliary fields. While the partition function of a general, interacting system is not analytically solvable, we can compute the partition function for a system parameterized by a single set of fields as a function of N and compare its results to those from the projection algorithm.

In Fig. 1, we plot the logarithm of the canonical partition function vs the number of electrons computed using the Borrmann, APF, and projection algorithms for a 100-site Hubbard chain with a randomly generated, but known set of Hubbard-Stratonovich fields at $\beta = 10$. As is evident from the scale on the plot, this is a stringent test of the stability of these algorithms because of the exceedingly large and small values of Z_N that can be assumed in this model. Due to binomial combinatorics, the partition function can be expected to peak at half-filling. This is correctly captured by the APF and projection algorithms but not by the Borrmann

recursion. While all algorithms are able to accurately compute the partition function at low fillings, the Borrmann recursion quickly loses its stability (and therefore accuracy) relative to the other algorithms at larger fillings. As previously observed [39], this is because the Borrmann recursion relies on sums over terms with alternating signs that can cancel each other out. The instability of the Borrmann algorithm can similarly be observed in recursions for the average occupations for all of the different energy levels, k , as presented in the inset of Fig. 1. Here, the Borrmann recursion is not only unable to reproduce the expected noninteracting Fermi-Dirac-like distribution, but even predicts unphysical negative occupations. In contrast, both the APF and projection algorithms agree regardless of filling, providing clear evidence that the APF formalism is highly stable for noninteracting systems. Given that the interacting partition function may be obtained by integrating over such noninteracting partition functions, the same stability observed in these simulations of noninteracting instantiations of the Hubbard model should naturally extend to simulations of the fully interacting model, as presented below.

Having demonstrated that both the APF and projection algorithms are highly stable, we next compare their computational scaling. Both algorithms make heavy use of full matrix diagonalizations at an $O(N_s^3)$ cost. However, due to its use of a Fourier sum, the projection algorithm sums over N_s Fourier components with each component computing a Fourier-frequency-dependent partition function of a single-particle space of size N_s . In contrast, the APF algorithm only requires N iterations, where each iteration involves computing the occupation probability of N_s single-particle levels. Ultimately, this results in the recursive algorithm having an $O(N_s^3 + N_s N)$ scaling for computing partition functions and occupation numbers (one-body densities), whereas the projection algorithm has an $O(N_s^3 + N_s^2)$ scaling. Thus, there is a clear benefit to using the APF method for filling fractions below unity. The two algorithms also possess different prefactors: The projection algorithm requires an additional rescaling step to avoid numerical overflow when the values of the Fourier components exceed the available floating point maximum at a given precision, while the APF algorithm is automatically stabilized as the calculations are mapped to probabilities that take values in the range $[0, 1]$. This points to the recursive algorithm being more efficient, particularly at low fillings and for large system sizes.

The increased efficiency of the recursive algorithm is reflected in the upper panel of Fig. 2, which shows the wall time required for each QMC step of a random sample drawn from an HS-transformed Hubbard model with $U = 2$ and $\beta = 12$ for varying numbers of sites. As computing a full QMC step involves both leading and subleading contributions to the overall scaling of the algorithms, the plots in the upper panel are reflective of the differences between the total execution times of both algorithms. In particular, the inset depicts the total speedup of the APF algorithm relative to the projection algorithm. The $O(N_s^3)$ diagonalization step is identical for both methods, so in the lower panel, we additionally plot the wall time to compute the canonical partition function, $Z_N(\sigma)$, which only reflects subleading contributions to the scaling, for both methods. Regardless of the size of the system and the filling, we find that our APF algorithm is always faster than

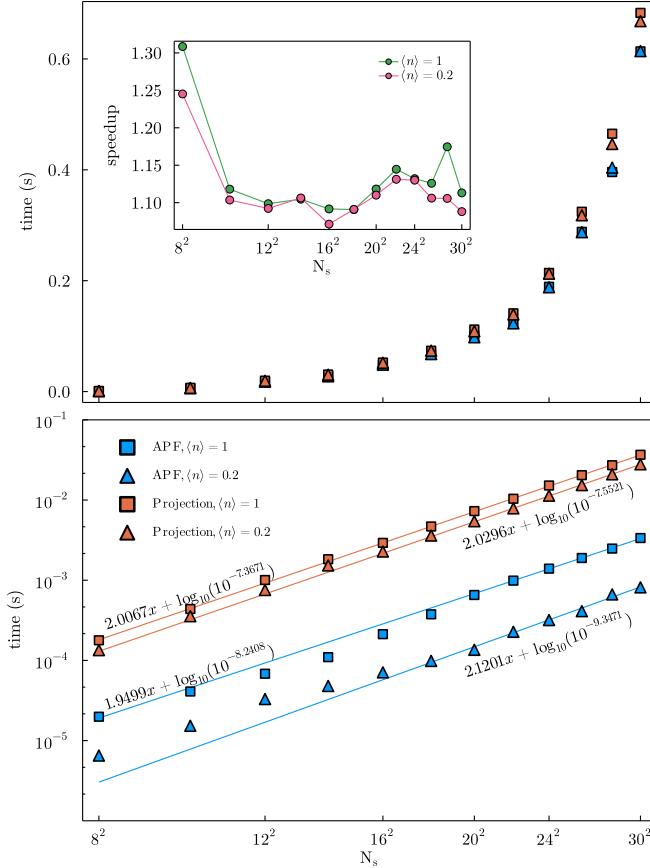


FIG. 2. Comparison of the runtimes (shorter is better) of our APF algorithm against the more conventional projection algorithm for QMC samples drawn from Hubbard models with varying numbers of lattice sites, N_s , with $U = 2$ at two different fillings $\langle n \rangle = 1$ and $\langle n \rangle = 0.2$, and $\beta = 12$. Runtimes are plotted for computing a single QMC step (upper panel), which reflects both leading and subleading contributions to the scaling, and the full partition function (lower panel), which reflects only subleading contributions to the scaling. The inset in the upper panel shows the speedup (larger is better) per QMC step of the APF algorithm relative to the projection algorithm. Runtimes for $\langle n \rangle = 1$ are denoted by the squares, while those for $\langle n \rangle = 0.2$ are denoted by the triangles.

the projection algorithm, in line with our scaling derivations. While one would expect the similar $O(N_s^3)$ contributions to the scaling to dominate the wall times of both algorithms, we nonetheless see from the inset that the APF algorithm remains 10–15% faster than the projection algorithm over a wide range of system sizes. This speedup leads to a significant practical gain in the efficiency of the APF algorithm relative to the projection algorithm, especially for larger system sizes and fillings. We also observe that the time to run the projection algorithm remained roughly the same for different filling fractions, while the wall time of the APF algorithm significantly decreased at lower fillings. Performing linear regression on the log-log data in Fig. 2 quantifies the overall N_s^2 scaling of the subleading contributions to the wall time (as given by the slopes of the regression lines) and reduced prefactor (as given by the regression intercepts) of the computationally more efficient APF method.

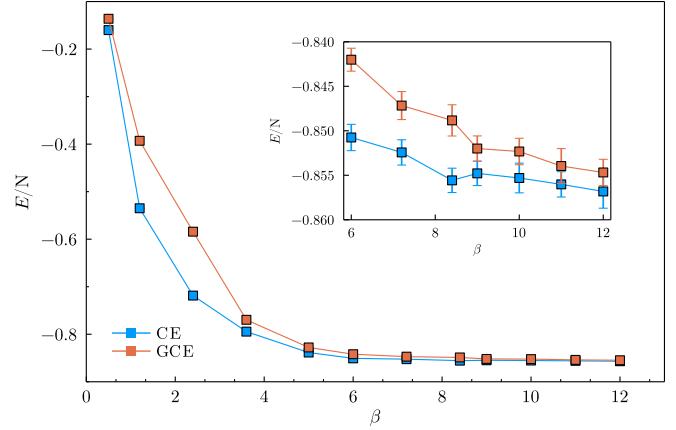


FIG. 3. Convergence of the energy per electron in the canonical (CE) and grand canonical (GCE) ensembles as a function of the inverse temperature for a 6×6 Hubbard model at half-filling with $U = 4$. Both the main panel and inset demonstrate the enhanced convergence in the canonical ensemble.

Although both algorithms could potentially be furthered optimized, we believe that the scalings described here are those representative of typical implementations of these algorithms. As further discussed in the supplementary material [54], computational complexities can also be worked out for the evaluation of the level occupations and their correlation functions. We find that the cost to compute occupations and related observables follows roughly the same scaling as for calculations of the partition function.

B. Ground-state convergence properties

1. Convergence of the energy to the ground state

Having demonstrated the markedly improved stability of our new method, we can now not only assess how it performs on fully interacting systems, but contrast the different physics that emerges in the canonical vs the grand canonical ensemble down to relatively low temperatures. To appreciate these disparities, we begin by comparing how the energies of the Hubbard model converge to their ground-state energies at fixed N and μ . The two ensembles are most effectively compared by choosing a grand canonical μ such that the average number of particles is given by $\langle \hat{N} \rangle_\mu = N$.

In the past, the convergence of the energy to its ground-state value with decreasing temperature has commonly been used to assess the relative contribution of thermal quasiparticle excitations above the ground state to the overall state of the quantum system. These excitations are also a key contributor to the free energy, the key property describing finite-temperature thermodynamics.

In Fig. 3, we plot the energy per electron vs the inverse temperature for the 6×6 Hubbard model at half-filling with $U = 4$ (and $t = 1$). Both the canonical and grand canonical ensembles yield predictably large energies at high temperatures due to the higher thermal energy enabling the electrons to access higher energy states, and then converge to roughly the same energy at lower temperatures (large β). For all values of β , the canonical ensemble energy is lower, with the

largest difference occurring at intermediate temperatures. This is because a larger number of higher energy states are accessible to the electrons in the grand canonical ensemble than in the canonical ensemble due to number fluctuations. This difference between ensembles grows with increased filling and decreased lattice size, which is in line with the intuition that differences between the ensembles should decrease as the thermodynamic limit is approached. Interestingly, despite the differences in their state spaces, both ensembles appear to converge to the same ground-state energy on the scale of Fig. 3 at low temperatures. As we shall show next, the energy turns out to be too blunt of a metric to detect subtle and potentially important differences between ensembles.

2. Convergence of the purity and fidelity to the ground state

Given the similar convergence of the energy to the ground state in both ensembles discussed in the previous section, one may ask if there are more fundamental metrics sensitive to differences between the two ensembles. Indeed, the two ensembles are composed of different states that are accessible at different temperatures which should lead to differences in their convergence to the ground state.

For any finite-sized system, there exists a crossover temperature below which the system is effectively in its ground state. This can serve as a more direct indicator for comparing the convergence rate between ensembles than by directly comparing the β dependence of the total energy. To quantitatively determine the crossover temperature, we are inspired by information theory and turn to the measurement of the purity

$$\mathcal{P} = \text{Tr} \hat{\rho}^2, \quad (33)$$

where $\hat{\rho}$ is the thermal density matrix. The purity quantifies how mixed a given finite-temperature state is: In general, any finite-temperature state is mixed and cannot be represented as a single vector in Hilbert space, resulting in a purity of less than 1, $\mathcal{P} < 1$. Thus, quantitative deviations of the purity from the identity can provide insights into the convergence to the ground state in a more general fashion than investigating any individual physical quantity whose $T = 0$ value may not be known in general.

The purity can be computed in QMC through a replica trick [62,63] by rewriting it in terms of a ratio of partition functions

$$\begin{aligned} \text{Tr} \hat{\rho}^2 &= \frac{Z(2\beta)}{Z^2(\beta)} = \frac{\int_{\sigma_1, \sigma_2} \text{Tr}(\hat{U}_{\sigma_1} \hat{U}_{\sigma_2})}{\int_{\sigma_1, \sigma_2} \text{Tr}(\hat{U}_{\sigma_1}) \text{Tr}(\hat{U}_{\sigma_2})} \\ &= \frac{\int_{\sigma_1, \sigma_2} Z(\sigma_1 \cup \sigma_2)}{\int_{\sigma_1, \sigma_2} Z(\sigma_1) Z(\sigma_2)}. \end{aligned} \quad (34)$$

Here $Z(\sigma_1)$ and $Z(\sigma_2)$ are the usual partition functions (in either the canonical or grand canonical ensembles) as a function of their Hubbard-Stratonovich fields σ_1 and σ_2 . $Z(\sigma_1 \cup \sigma_2)$ can be viewed as the partition function for a connected ensemble propagating from 0 to 2β . The ensemble switching technique [64,65] can then be adopted to efficiently sample this ratio of partition functions.

Although there is no known closed form for the purity, one can still expand ρ^2 in terms of a system's energy levels, and in the ground-state (large- β) limit, only the leading term remains

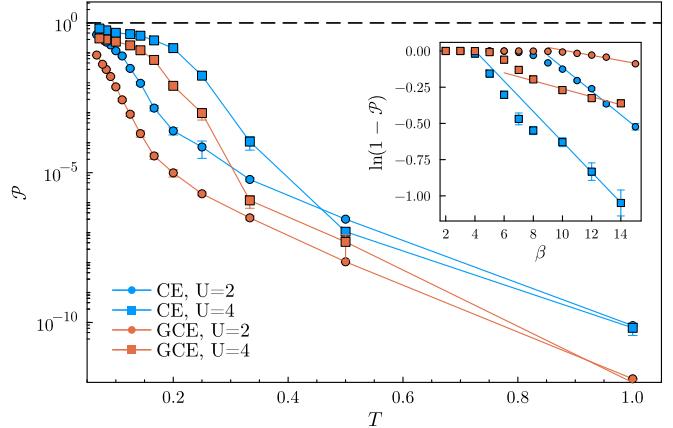


FIG. 4. The purity, \mathcal{P} , of a 6×6 Hubbard model at half-filling with $U = 2$ and $U = 4$ as a function of temperature for the canonical (CE) and grand canonical (GCE) ensemble. Different symbols correspond to different interaction strengths and the lines are a guide to the eye. The inset shows $\ln(1 - \mathcal{P})$ as a function of the inverse temperature β , with the crossover to linear behavior (fitted lines) indicating convergence to the ground state [see Eq. (35)].

(see the supplementary material [54] for a full derivation),

$$\mathcal{P} \propto 1 - 2e^{-\beta \Delta E}. \quad (35)$$

In the canonical ensemble, ΔE exactly corresponds to the energy gap between the ground state and first excited state, while in the grand canonical ensemble, ΔE represents an effective energy gap that contains contributions from states with $N - 1$ and $N + 1$ particles as subleading terms. As a result, a temperature below which the system falls into its ground-state region can be revealed by the onset of linear behavior when plotting $\ln(1 - \mathcal{P})$ vs β . In Fig. 4, we show the purity vs temperature for Hubbard models simulated in the canonical and grand canonical ensembles for two interaction strengths U , and fit $\ln(1 - \mathcal{P})$ linearly against the inverse temperature in the low-temperature region in the inset. As before, simulations performed in the canonical ensemble converge more rapidly to the ground state than those performed in the grand canonical ensemble, as indicated by the uniformly larger canonical purities for any given U . The canonical and grand canonical ensemble purities also show the greatest agreement for $U = 4$ for $\beta > 6$, when $\ln(1 - \mathcal{P})$ begins to significantly deviate below zero, which echoes the convergence for $\beta > 6$ seen earlier in the energy.

However, in contrast to the relative convergence of the energy, the purity reveals additional trends rooted in the underlying physics of the finite-temperature crossovers. In particular, $\ln(1 - \mathcal{P})$ can be seen to crossover from exhibiting roughly constant behavior at high temperatures to exhibiting linear decay with decreasing temperature (increasing β), as predicted by Eq. (35). This is a clear indicator that the system has crossed into a regime that is resolving the ground states. By fitting the linear decay of $\ln(1 - \mathcal{P})$, we can also extract the energy gap ΔE from the slope of the regression lines. While it can be statistically challenging to fit curves to such small purity values in the presence of Monte Carlo uncertainties, our fitting procedure yields gaps of -0.0816

and -0.1051 for the canonical ensemble at $U = 2$ and $U = 4$, respectively, and -0.0182 and -0.0275 for the grand canonical ensemble at $U = 2$ and $U = 4$, respectively. Interestingly, the more negative canonical ensemble slopes suggest that the canonical ensemble gaps are larger than the effective grand canonical gaps, which is supported by the fact that more midgap states may be present in a grand-canonical treatment. Moreover, based on the effective gaps extracted, the systems with $U = 4$ possess larger gaps than those with $U = 2$, which is expected given the direct correlation between larger U and larger ΔE . A more detailed discussion of the information that can be obtained from the purity is presented in the supplementary material [54].

Although the purity has provided a more detailed glimpse into the physics of the crossovers that occur in the different ensembles, it is not a measure that directly compares the two ensembles. One metric that can draw such a direct comparison is the fidelity, $\mathcal{F}(\rho, \rho')$, that measures the similarity between two density matrices ρ and ρ' .

In this case, the thermal density matrices in the two ensembles, ρ_N and ρ_μ , describe states that are not pure, and the mixed-state fidelity can be defined as the Hilbert-Schmidt inner product of ρ_N and ρ_μ normalized by the two purities [41,66],

$$\mathcal{F}(\rho_N, \rho_\mu) = \frac{\text{Tr}(\rho_N \rho_\mu)}{\sqrt{\text{Tr}(\rho_N^2) \text{Tr}(\rho_\mu^2)}}. \quad (36)$$

It is straightforward to show that under this definition, the fidelity is normalized and reaches its maximum value of 1 if and only if $\rho_N = \rho_\mu$, and that it is also symmetric under the exchange of ρ_N and ρ_μ , $\mathcal{F}(\rho_N, \rho_\mu) = \mathcal{F}(\rho_\mu, \rho_N)$. Note that the trace in the numerator is taken over the whole Fock space, so the matrix form of ρ_N is expanded from the N -particle Hilbert space to the Fock space with varying particle numbers, but is only nonzero in the N -particle block. Moreover, when the number operator, \hat{N} , commutes with the Hamiltonian, as is the case for the Hubbard model, ρ_μ is block-diagonal in the particle number regardless of the interaction strength. This fact allows us to simplify the numerator of Eq. (36) to $\text{Tr}(\rho_N \rho_\mu) = P_\mu(N, \beta) \text{Tr}(\rho_N^2)$. After some algebra, we arrive at

$$\mathcal{F}(\rho_N, \rho_\mu) = \sqrt{P_\mu(N, 2\beta)} = \sqrt{\frac{e^{2\beta\mu N} Z_N(2\beta)}{\mathcal{Z}_\mu(2\beta)}}, \quad (37)$$

which allows $\mathcal{F}(\rho_N, \rho_\mu)$ to be directly measured within our AFQMC simulations, as the ratio between partition functions, $Z_N(2\beta)$ and $\mathcal{Z}_\mu(2\beta)$, can be measured through the same ensemble switching technique as was employed in the purity calculations.

In the ground-state limit ($\beta \rightarrow \infty$), the fidelity has a similar limiting behavior as the purity, which can be derived by expanding the particle-number distribution $P_\mu(N, 2\beta)$. A full derivation can be found in the supplementary material [54] and yields

$$\mathcal{F} \propto 1 - \frac{g}{2} e^{-2\beta\Delta\tilde{E}}, \quad (38)$$

where $\Delta\tilde{E}$ is again an effective energy gap that includes the effects of the gap for the system of N particles as its leading

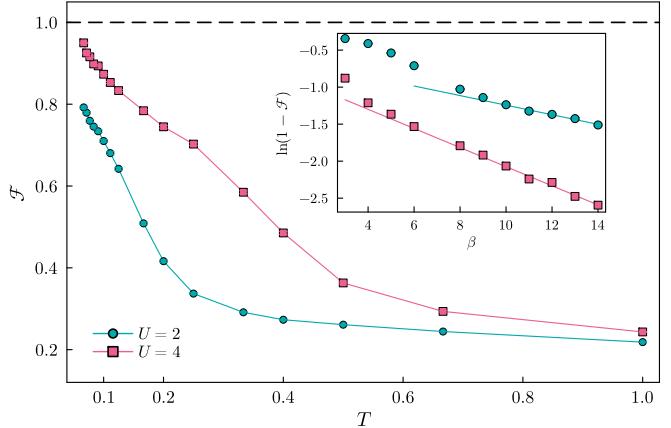


FIG. 5. Fidelity as a function of temperature for a 6×6 Hubbard model at half-filling with $U = 2$ and $U = 4$. The inset fits $\ln(1 - \mathcal{F})$ to the inverse temperature from $\beta = 8$ to $\beta = 14$ and extends to β values where $\ln(1 - \mathcal{F})$ deviates from a linear fit. The error bars in the inset figure are smaller than the symbols.

term and the gaps for the systems of $N + 1$ and $N - 1$ particles as its subleading terms. g represents an effective degeneracy that accounts for the potentially unresolved spacing of the energy spectrum. This equation implies that plotting $\ln(1 - \mathcal{F})$ against β is expected to possess linear scaling in the large- β limit.

In Fig. 5, we show the fidelity vs temperature as well as $\ln(1 - \mathcal{F})$ vs β in the inset. From this plot, we see that the fidelity is larger for $U = 4$ than $U = 2$, meaning that, at low temperatures, the ensembles are more similar for larger U values. This is likely because larger U values result in a larger gap, which limits how many additional grand canonical states the system can access beyond those occupied in the canonical ensemble. From the inset, we additionally observe how $\ln(1 - \mathcal{F})$ becomes linear in β at low temperatures, as predicted by Eq. (38). A larger effective gap is again observed for the more strongly interacting case with $U = 4$, which possesses a more negative slope than for $U = 2$. The considerable deviation of the fidelity from unity, even at a very low temperature ($T < 0.1$), can be understood from Eq. (37), which definitively captures how particle-number fluctuations can continue to contribute to the grand canonical density matrix, even in the limit of large systems when approaching the ground state.

Although the gaps obtained may not yet be as accurate as those obtained from excited-state calculations, these examples illustrate that the purity and fidelity are much more informative metrics of convergence than the energy alone, and provide additional information that can be exploited to estimate gaps from finite-temperature simulations.

C. Sign problem in the canonical ensemble

After observing how the canonical ensemble converges more rapidly to the ground state, one may ask if this provides a practical way of more readily accessing $\beta \rightarrow \infty$ quantities than in the grand canonical ensemble. After all, it is reasonable to assume that if the energy and wave function converge more rapidly in the canonical ensemble, perhaps one can more

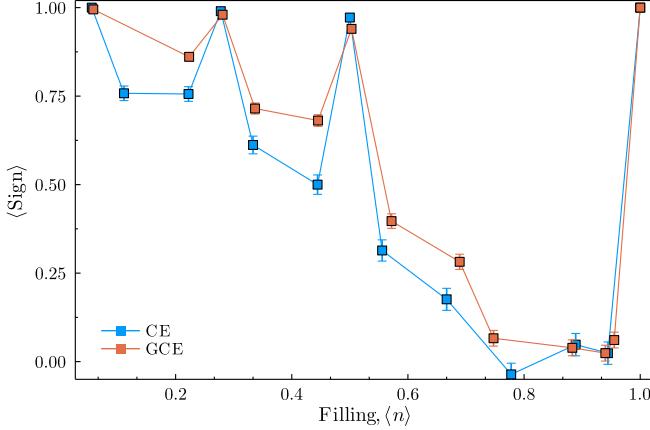


FIG. 6. Average sign vs filling fraction in the canonical (CE) and grand canonical (GCE) ensembles of a 6×6 Hubbard model with $U = 4$ and $\beta = 10$.

readily gain access to low-temperature physics before a severe sign problem sets in at certain fillings.

To address the interplay between convergence to the ground state and the emergence of a physical fermion sign problem, we compute the average sign at different fillings for a variety of temperatures and interaction strengths. The behavior of the sign as a function of filling presented in Fig. 6 is representative of what we more widely observe: In general, the average sign in the canonical ensemble is less than that in the grand canonical ensemble at any given filling. Just as in the grand canonical ensemble, the sign at certain fillings is at or near 1, reflecting special symmetries and indicating that the system can be modeled with no approximations at polynomial cost. However, away from these special fillings, the average sign decreases, meaning that either exponentially more samples must be taken to converge average observables or sign mitigation strategies must be employed. Unfortunately, where the sign problem is present, the canonical sign problem appears to be more severe. This is likely a consequence of the fact that the canonical ensemble more quickly converges to the ground state: The same states that lower the canonical energy relative to the grand canonical energy are those that give rise to a more significant sign problem. This presents a practical tradeoff. While the canonical ensemble more rapidly converges to the ground state with decreasing temperature, it does so with an increased sign problem, signifying that simulations in the canonical ensemble do not allow for a way to mitigate costs associated with simulating many-body systems of fermions. The ground state of the fermion Hubbard model at most fillings possesses a significant sign problem; the more rapidly this ground state is approached, the more rapidly a sign problem emerges, regardless of ensemble.

D. Impact on thermometry: Differences between canonical and grand canonical ensemble observables

While the energy and sign are two of the most commonly measured observables in stochastic many-body simulations, densities and correlation functions enable the most direct comparisons with experiments. One may therefore ask what significant differences may exist between canonical and grand

canonical densities. This is not an idle question: Many recent cold atom experiments estimate the temperature of their trapped gases assuming that the constituent particles interact according to the grand canonical ensemble [1,67]. Although this is valid at large particle numbers, many such experiments are performed in a mesoscopic regime in which only a finite number of particles are present. Making this assumption when the thermodynamic limit has not been reached can lead to inaccurate estimates of the system temperature, resulting in incorrect phase diagrams and potentially unfounded efforts to further reduce temperatures.

To make clear the effect of choosing the incorrect ensemble on the determination of temperature in a finite closed quantum system, it is illustrative to consider the simplest extreme example of one particle ($N = 1$) distributed among $N_s = 2$ energy levels $\{0, \Delta\}$. In the canonical ensemble,

$$\langle n_1 \rangle_1 = \frac{1}{1 + e^{-\beta\Delta}}; \quad \langle n_2 \rangle_1 = e^{-\beta\Delta} \langle n_1 \rangle_1, \quad (39)$$

while in the grand canonical ensemble, these average occupations depend on the chemical potential μ ,

$$\langle n_1 \rangle_\mu = \frac{1}{1 + e^{-\beta\mu}}; \quad \langle n_2 \rangle_\mu = \frac{e^{-\beta(\Delta-\mu)}}{1 + e^{-\beta(\Delta-\mu)}}. \quad (40)$$

In Eq. (40), μ is chosen such that $\langle n_1 \rangle_\mu + \langle n_2 \rangle_\mu = 1$, which yields $\mu = \Delta/2$. Using this value in Eq. (40) and equating occupation numbers between ensembles (the quantity most directly accessible in thermometry experiments through the velocity distribution) would require using an inverse temperature in the grand canonical ensemble that is twice that of the physical canonical temperature, i.e., $\beta_\mu = 2\beta = 2\beta_N$. This would result in a 100% error in the extracted temperature (when the incorrect ensemble is chosen), with grand canonical simulations always predicting a *lower* temperature than the physical one.

While this is obviously an extreme (toy) example, thermometry errors can persist to larger systems that include interactions that affect other measurable quantities. For example, second-order fluctuations in the particle number provide an even clearer way to explore the impact of choosing the incorrect ensemble. This can be quantified by the site occupancy correlation function

$$\langle \hat{n}_i \hat{n}_j \rangle = \langle (\hat{n}_{i,\uparrow} + \hat{n}_{i,\downarrow})(\hat{n}_{j,\uparrow} + \hat{n}_{j,\downarrow}) \rangle, \quad (41)$$

where i, j are D -dimensional site indices. This correlation function is more naturally studied in momentum space via the static charge structure factor at wave vector \mathbf{k} , given by the Fourier transform of Eq. (41):

$$C_{\mathbf{k}} = \frac{1}{N_s} \sum_{i,j} e^{i\mathbf{k} \cdot (i-j)} \langle \hat{n}_i \hat{n}_j \rangle, \quad (42)$$

which is, in principle, measurable in cold atom experiments through, e.g., Bragg spectroscopy [68–71]. To quantify the difference between ensemble predictions for such a quantity, we compute $C_{\mathbf{k}}$ at $\mathbf{k} = (\pi, \pi)$ to expose antiferromagnetic correlations for a two-dimensional Hubbard model at $U = 2$ and filling $\langle n \rangle = \frac{2 \times 23}{36}$ as a function of inverse temperature β as shown in Fig. 7. The odd number of electrons was chosen due to the existence of the large gap and the increased spacing

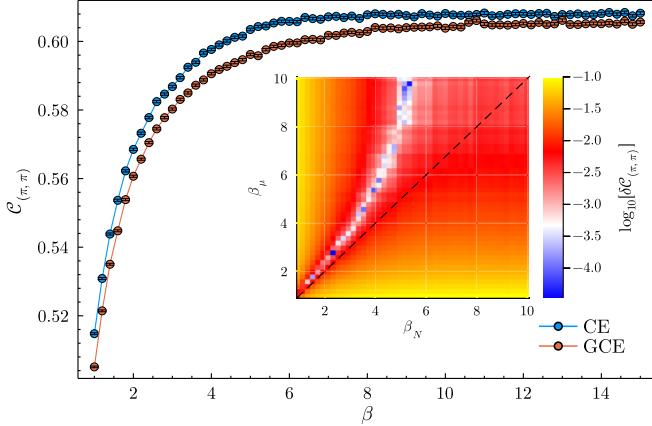


FIG. 7. Charge structure factor at $\mathbf{k} = (\pi, \pi)$, $C_{(\pi,\pi)}$, as a function of inverse temperature β in the canonical (CE) and grand canonical (GCE) ensembles for a 6×6 Hubbard model at $U = 2$, $\langle n \rangle = \frac{2 \times 23}{36} \approx 1.28$. The inset shows the heatmap of the difference in the charge structural factors, $\delta C_{(\pi,\pi)}(\beta_N, \beta_\mu)$, as a function of the inverse temperature in the canonical and grand canonical ensembles.

of low-lying energy levels. The grand canonical ensemble simulations off of half-filling are implemented by tuning the chemical potential dynamically [27] during the Monte Carlo step to ensure $|\langle n \rangle_\mu - 1.28| < 1 \times 10^{-3}$. The deviation at each temperature is clear, with the canonical structure factor having a consistently larger value due to suppressed fluctuations. The inset shows a heat-map of the difference $\delta C_{\mathbf{k}}$ for each temperature $\beta \equiv \beta_N$, with the y axis representing the effective inverse temperature β_μ needed to minimize the difference between ensembles (the horizontal shift in the main panel needed to align the two curves). Deviations of the minimum from the line $\beta_\mu = \beta_N$ captures potential errors in thermometry when using the incorrect ensemble. The maximum deviation in the heatmap over the temperature range we study appears to be $(\beta_N, \beta_\mu) \approx (4.4, 9.4)$, which leads to a 53.2% thermometry error. Such large errors underscore the dramatic errors that can arise from an incorrect choice of ensemble and that can be addressed using the new techniques presented in this work.

IV. CONCLUSIONS

In summary, we have presented and illustrated applications of a new, significantly more stable recursive algorithm for determining the physics of interacting systems in the canonical ensemble. This algorithm integrates the auxiliary partition function formalism, a highly stable means of computing the properties of noninteracting systems in the canonical ensemble, into the auxiliary field quantum Monte Carlo framework by exploiting the Hubbard-Stratonovich transformation. We demonstrate the stability of this algorithm and then showcase its potential applications by studying differences in the way that the canonical and grand canonical ensembles converge to the ground state. This convergence is quantified using information theory metrics by comparing the purity of finite-temperature states generated within the grand canonical and canonical frameworks. We find that the canonical AFQMC results in a suppression of the mixed state and improved

fidelity with the ground state as $T \rightarrow 0$ in a practical simulation.

As an experimentally relevant application, we show that a grand canonical treatment of the thermometry of cold atom and other systems with fixed particle numbers can lead to underestimates of the temperatures of those systems, clouding investigations of their thermodynamics. This is becoming more pressing as studies of trapped homogeneous ultracold fermions push into a regime of smaller T/T_F where they are more poorly described by the grand canonical ensemble [2]. Moreover, this work has direct implications for the study of nuclear matter, which possesses fixed nucleon numbers and has traditionally been modeled using the projection algorithm [5–8].

The fact that we find our algorithm to be more computationally efficient than oft-used projection algorithms opens the door to a wealth of potential new applications. To date, most projection algorithms have been limited to system sizes of tens, to at the very most, low hundreds of particles, with application to smaller nuclei and “toy” condensed matter systems. In its present form, without many algorithmic advances or computational fine-tuning, our algorithm can readily model systems with many hundreds of particles. This opens the door to more accurate numerical descriptions of cold atom quantum simulators or mesoscale devices where discrete particle-number fluctuations can influence the transport of heat and matter [72].

Although we illustrated the performance of our algorithm on systems of fermions because of their greater relevance and potential to develop a sign problem, our algorithm, with appropriate modifications, is equally applicable to systems of bosons or particles with other quantum statistics. The fact that the algorithm does not require explicit knowledge of many-body chemical potentials and is stable for large systems implies that it can see wide application in the study of quantum condensates, which can be challenging to simulate in the grand canonical ensemble [28,46].

We also anticipate that our algorithm will enable more direct comparisons with other finite-temperature algorithms, including the path integral Monte Carlo [17,73,74], density matrix quantum Monte Carlo [10,16], and emerging finite-temperature coupled-cluster theories [75,76], all of which are formulated in the canonical ensemble. Beyond the algorithmic, our work will furthermore enable seamless canonical ensemble simulations from high temperatures to the ground state where most simulations are inherently performed in the canonical ensemble. This will provide critical insights into how finite-temperature physics gives rise to ground-state physics in correlated systems with decreasing temperature without the noise induced by spurious particle-number fluctuations. Given the correlations that fixed-particle-number constraints impose on the occupancies of different states, we moreover anticipate that fluctuations and therefore the physics of systems in the canonical ensemble will be fundamentally different. We look forward to the new such canonical ensemble physics this algorithm will reveal.

All codes, scripts, and data needed to reproduce the results in this paper are available online [77,78].

ACKNOWLEDGMENTS

The authors thank Richard Stratt, Christopher Gilbreth, Scott Jensen, and Ben Cohen-Stead for fruitful conversations. T.S., J.Y., and B.M.R. were funded by NSF CTMC CAREER Award 2046744. T.S. is also grateful for financial support

from the Brown Open Graduate Education program. A.D.M. acknowledges support from the NSF under Grant No. DMR-2041995. This research was conducted using computational resources and services at the Center for Computation and Visualization, Brown University.

[1] B. Mukherjee, Z. Yan, P. B. Patel, Z. Hadzibabic, T. Yefsah, J. Struck, and M. W. Zwierlein, Homogeneous Atomic Fermi Gases, *Phys. Rev. Lett.* **118**, 123401 (2017).

[2] K. Hueck, N. Luick, L. Sobirey, J. Siegl, T. Lompe, and H. Moritz, Two-Dimensional Homogeneous Fermi Gases, *Phys. Rev. Lett.* **120**, 060402 (2018).

[3] N. Navon, R. P. Smith, and Z. Hadzibabic, Quantum gases in optical boxes, *Nat. Phys.* **17**, 1334 (2021).

[4] Y. Alhassid, Auxiliary-field quantum Monte Carlo methods in nuclei, in *Emergent Phenomena in Atomic Nuclei from Large-Scale Modeling* (World Scientific, Singapore, 2017), Chap. 9, pp. 267–298

[5] G. H. Lang, C. W. Johnson, S. E. Koonin, and W. E. Ormand, Monte Carlo evaluation of path integrals for the nuclear shell model, *Phys. Rev. C* **48**, 1518 (1993).

[6] H. Nakada and Y. Alhassid, Total and Parity-Projected Level Densities of Iron-Region Nuclei in the Auxiliary Fields Monte Carlo Shell Model, *Phys. Rev. Lett.* **79**, 2939 (1997).

[7] Y. Alhassid, L. Fang, and H. Nakada, Heavy Deformed Nuclei in the Shell Model Monte Carlo Method, *Phys. Rev. Lett.* **101**, 082501 (2008).

[8] W. E. Ormand, D. J. Dean, C. W. Johnson, G. H. Lang, and S. E. Koonin, Demonstration of the auxiliary-field Monte Carlo approach for sd-shell nuclei, *Phys. Rev. C* **49**, 1422 (1994).

[9] Y. Liu, M. Cho, and B. Rubenstein, *Ab Initio* finite temperature auxiliary field quantum Monte Carlo, *J. Chem. Theory Comput.* **14**, 4722 (2018).

[10] H. R. Petras, S. K. Ramadugu, F. D. Malone, and J. J. Shepherd, Using density matrix quantum Monte Carlo for calculating exact-on-average energies for *ab initio* hamiltonians in a finite basis set, *J. Chem. Theory Comput.* **16**, 1029 (2020).

[11] S. Zhang and H. Krakauer, Quantum Monte Carlo Method using Phase-Free Random Walks with Slater Determinants, *Phys. Rev. Lett.* **90**, 136401 (2003).

[12] S. Zhang, J. Carlson, and J. E. Gubernatis, Constrained path Monte Carlo method for fermion ground states, *Phys. Rev. B* **55**, 7464 (1997).

[13] U. Schollwöck, The density-matrix renormalization group, *Rev. Mod. Phys.* **77**, 259 (2005).

[14] I. Y. Zhang and A. Grüneis, Coupled cluster theory in materials science, *Front. Mater.* **6**, 123 (2019).

[15] Y. Liu, T. Shen, H. Zhang, and B. Rubenstein, Unveiling the finite temperature physics of hydrogen chains via auxiliary field quantum Monte Carlo, *J. Chem. Theory Comput.* **16**, 4298 (2020).

[16] F. D. Malone, N. S. Blunt, J. J. Shepherd, D. K. K. Lee, J. S. Spencer, and W. M. C. Foulkes, Interaction picture density matrix quantum Monte Carlo, *J. Chem. Phys.* **143**, 044116 (2015).

[17] F. D. Malone, N. S. Blunt, E. W. Brown, D. K. K. Lee, J. S. Spencer, W. M. C. Foulkes, and J. J. Shepherd, Accurate Exchange-Correlation Energies for the Warm Dense Electron Gas, *Phys. Rev. Lett.* **117**, 115701 (2016).

[18] E. W. Brown, B. K. Clark, J. L. DuBois, and D. M. Ceperley, Path-Integral Monte Carlo Simulation of the Warm Dense Homogeneous Electron Gas, *Phys. Rev. Lett.* **110**, 146405 (2013).

[19] H. M. Wiseman and J. A. Vaccaro, Entanglement of Indistinguishable Particles Shared Between Two Parties, *Phys. Rev. Lett.* **91**, 097902 (2003).

[20] S. D. Bartlett and H. M. Wiseman, Entanglement Constrained by Superselection Rules, *Phys. Rev. Lett.* **91**, 097903 (2003).

[21] H. Barghathi, C. M. Herdman, and A. Del Maestro, Rényi Generalization of the Accessible Entanglement Entropy, *Phys. Rev. Lett.* **121**, 150501 (2018).

[22] T. Rauscher, F.-K. Thielemann, and K.-L. Kratz, Nuclear level density and the determination of thermonuclear rates for astrophysics, *Phys. Rev. C* **56**, 1613 (1997).

[23] D. McQuarrie, *Statistical Mechanics* (University Science Books, Herndon, VA, 2000).

[24] H. Barghathi, J. Yu, and A. Del Maestro, Theory of noninteracting fermions and bosons in the canonical ensemble, *Phys. Rev. Res.* **2**, 043206 (2020).

[25] P. Borrman and G. Franke, Recursion formulas for quantum statistical partition functions, *J. Chem. Phys.* **98**, 2484 (1993).

[26] C. Gilbreth and Y. Alhassid, Stabilizing canonical-ensemble calculations in the auxiliary-field Monte Carlo method, *Comput. Phys. Commun.* **188**, 1 (2015).

[27] C. Miles, B. Cohen-Stead, O. Bradley, S. Johnston, R. Scalettar, and K. Barros, Dynamical tuning of the chemical potential to achieve a target particle number in grand canonical Monte Carlo simulations, *Phys. Rev. E* **105**, 045311 (2022).

[28] C. N. Gilbreth and Y. Alhassid, Pair condensation in a finite trapped fermi gas, *Phys. Rev. A* **88**, 063643 (2013).

[29] S. Jensen, C. N. Gilbreth, and Y. Alhassid, Pairing Correlations Across the Superfluid Phase Transition in the Unitary Fermi Gas, *Phys. Rev. Lett.* **124**, 090604 (2020).

[30] C. M. Herdman, A. Rommal, and A. Del Maestro, Quantum Monte Carlo measurement of the chemical potential of ${}^4\text{He}$, *Phys. Rev. B* **89**, 224502 (2014).

[31] Z. Wang, F. F. Assaad, and F. Parisen Toldin, Finite-size effects in canonical and grand-canonical quantum Monte Carlo simulations for fermions, *Phys. Rev. E* **96**, 042131 (2017).

[32] R. D. Sedgewick, D. J. Scalapino, R. L. Sugar, and L. Capriotti, Canonical and grand canonical ensemble expectation values from quantum Monte Carlo simulations, *Phys. Rev. B* **68**, 045120 (2003).

[33] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis, and R. T. Scalettar, Numerical study of the two-dimensional hubbard model, *Phys. Rev. B* **40**, 506 (1989).

[34] S. Zhang, Finite-Temperature Monte Carlo Calculations for Systems with Fermions, *Phys. Rev. Lett.* **83**, 2777 (1999).

[35] J. E. Hirsch, Discrete Hubbard-Stratonovich transformation for fermion lattice models, *Phys. Rev. B* **28**, 4059 (1983).

[36] R. L. Stratonovich, Sov. Phys. Dokl. **2**, 416 (1958).

[37] J. Hubbard, Calculation of Partition Functions, *Phys. Rev. Lett.* **3**, 77 (1959).

[38] M. Motta and S. Zhang, Ab initio computations of molecular systems by the auxiliary-field quantum Monte Carlo method, *WIREs Comp. Mol. Sci.* **8**, 1364 (2018).

[39] T. Shen, Y. Liu, Y. Yu, and B. M. Rubenstein, Finite temperature auxiliary field quantum Monte Carlo in the canonical ensemble, *J. Chem. Phys.* **153**, 204108 (2020).

[40] L. Bonnes, H. Pichler, and A. M. Läuchli, Entropy perspective on the thermal crossover in a fermionic Hubbard chain, *Phys. Rev. B* **88**, 155103 (2013).

[41] G. Parez, Symmetry-resolved Rényi fidelities and quantum phase transitions, *Phys. Rev. B* **106**, 235101 (2022).

[42] H. F. Trotter, On the product of semi-groups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).

[43] M. Suzuki, Relationship between d -dimensional quantal spin systems and $(d+1)$ -dimensional Ising systems: Equivalence, critical exponents and systematic approximants of the partition function and spin correlations, *Prog. Theor. Phys.* **56**, 1454 (1976).

[44] J. E. Hirsch, Two-dimensional Hubbard model: Numerical simulation study, *Phys. Rev. B* **31**, 4403 (1985).

[45] Z. Bai, W. Chen, R. Scalettar, and I. Yamazaki, Numerical methods for quantum Monte Carlo simulations of the Hubbard model, in *Multi-Scale Phenomena in Complex Fluids* (World Scientific, Singapore, 2009), pp. 1–110.

[46] B. M. Rubenstein, S. Zhang, and D. R. Reichman, Finite-temperature auxiliary-field quantum Monte Carlo technique for Bose-Fermi mixtures, *Phys. Rev. A* **86**, 053606 (2012).

[47] G. M. Buendia, Comparative study of the discrete and the continuous Hubbard-Stratonovich transformation for a one-dimensional spinless fermion model, *Phys. Rev. B* **33**, 3519 (1986).

[48] A. Fetter and J. Walecka, *Quantum Theory of Many-Particle Systems*, Dover Books on Physics (Dover, Mineola, NY, 2012).

[49] G. C. Wick, The evaluation of the collision matrix, *Phys. Rev.* **80**, 268 (1950).

[50] P. Borrman, J. Harting, O. Mülken, and E. R. Hilf, Calculation of thermodynamic properties of finite Bose-Einstein systems, *Phys. Rev. A* **60**, 1519 (1999).

[51] K. Schönhammer, Deviations from Wick's theorem in the canonical ensemble, *Phys. Rev. A* **96**, 012102 (2017).

[52] H.-J. Schmidt and J. Schnack, Thermodynamic fermion-boson symmetry in harmonic oscillator potentials, *Physica A* **265**, 584 (1999).

[53] Y. Hong, On computing the distribution function for the poisson binomial distribution, *Comput. Stat. Data. Anal.* **59**, 41 (2013).

[54] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.107.055302> for additional algorithmic details, extensions to bosonic systems, derivations of information theoretic quantities, and additional comparisons between ensembles, which includes Refs. [55–58].

[55] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Sign problem in the numerical simulation of many-electron systems, *Phys. Rev. B* **41**, 9301 (1990).

[56] L. Fang, *Shell Model Monte Carlo Methods for Nuclei at Finite Temperature* (Yale University, New Haven, CT, 2005).

[57] A. Uhlmann, The “transition probability” in the state space of a $*\text{-algebra}$, *Rep. Math. Phys.* **9**, 273 (1976).

[58] R. Jozsa, Fidelity for mixed quantum states, *J. Mod. Opt.* **41**, 2315 (1994).

[59] C. Gilbreth, S. Jensen, and Y. Alhassid, Reducing the complexity of finite-temperature auxiliary-field quantum Monte Carlo, *Comput. Phys. Commun.* **264**, 107952 (2021).

[60] H. Schmidt, A simple derivation of distribution functions for bose and fermi statistics, *Am. J. Phys.* **57**, 1150 (1989).

[61] O. Giraud, A. Grabsch, and C. Texier, Correlations of occupation numbers in the canonical ensemble and application to a bose-einstein condensate in a one-dimensional harmonic trap, *Phys. Rev. A* **97**, 053615 (2018).

[62] P. Calabrese and J. Cardy, Entanglement entropy and quantum field theory, *J. Stat. Mech.* (2004) P06002.

[63] A. J. Daley, H. Pichler, J. Schachenmayer, and P. Zoller, Measuring Entanglement Growth in Quench Dynamics of Bosons in an Optical Lattice, *Phys. Rev. Lett.* **109**, 020505 (2012).

[64] P. Broecker and S. Trebst, Rényi entropies of interacting fermions from determinantal quantum Monte Carlo simulations, *J. Stat. Mech.* (2014) P08015.

[65] F. F. Assaad, Stable quantum Monte Carlo simulations for entanglement spectra of interacting fermions, *Phys. Rev. B* **91**, 125146 (2015).

[66] X. Wang, C.-S. Yu, and X. Yi, An alternative quantum fidelity for mixed states of qudits, *Phys. Lett. A* **373**, 58 (2008).

[67] R. Onofrio, Cooling and thermometry of atomic fermi gases, *Phys. Usp.* **59**, 1129 (2016).

[68] E. Altman, E. Demler, and M. D. Lukin, Probing many-body states of ultracold atoms via noise correlations, *Phys. Rev. A* **70**, 013603 (2004).

[69] E. D. Kuhnle, H. Hu, X.-J. Liu, P. Dyke, M. Mark, P. D. Drummond, P. Hannaford, and C. J. Vale, Universal Behavior of Pair Correlations in a Strongly Interacting Fermi Gas, *Phys. Rev. Lett.* **105**, 070402 (2010).

[70] M. Boll, T. A. Hilker, G. Salomon, A. Omran, J. Nespolo, L. Pollet, I. Bloch, and C. Gross, Spin-and density-resolved microscopy of antiferromagnetic correlations in Fermi-Hubbard chains, *Science* **353**, 1257 (2016).

[71] T. L. Yang, P. Grišins, Y. T. Chang, Z. H. Zhao, C. Y. Shih, T. Giamarchi, and R. G. Hulet, Measurement of the Dynamical Structure Factor of a 1D Interacting Fermi Gas, *Phys. Rev. Lett.* **121**, 103001 (2018).

[72] M. Campisi, P. Hänggi, and P. Talkner, Colloquium: Quantum fluctuation relations: Foundations and applications, *Rev. Mod. Phys.* **83**, 771 (2011).

[73] B. Hirshberg, V. Rizzi, and M. Parrinello, Path integral molecular dynamics for bosons, *Proc. Natl. Acad. Sci. USA* **116**, 21445 (2019).

[74] T. Dornheim, S. Groth, J. Vorberger, and M. Bonitz, Ab Initio Path Integral Monte Carlo Results for the Dynamic Structure Factor of Correlated Electrons: From the Electron Liquid to Warm Dense Matter, *Phys. Rev. Lett.* **121**, 255001 (2018).

[75] A. F. White and G. K.-L. Chan, A time-dependent formulation of coupled-cluster theory for many-fermion systems at finite temperature, *J. Chem. Theory Comput.* **14**, 5690 (2018).

[76] S. Hirata and P. K. Jha, Finite-temperature many-body perturbation theory in the grand canonical ensemble, *J. Chem. Phys.* **153**, 014103 (2020).

[77] T. Shen, H. Barghathi, A. Del Maestro, and B. Rubenstein, Code repository for “A Stable, Recursive Auxiliary Field Quantum Monte Carlo Algorithm in the Canonical Ensemble: Applications to Thermometry and the Hubbard Model” (unpublished), <https://github.com/TongSericus/papers-code-CanEnsAFQMC>.

[78] T. Shen, Canonical ensemble AFQMC for Fermi Hubbard models, GitHub (unpublished), <https://github.com/TongSericus/CanEnsAFQMC>.