

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

Nonlinear biases, stochastically sampled effective Hamiltonians, and spectral functions in quantum Monte Carlo methods

Nick S. Blunt, Ali Alavi, and George H. Booth

Phys. Rev. B **98**, 085118 — Published 9 August 2018

DOI: [10.1103/PhysRevB.98.085118](https://doi.org/10.1103/PhysRevB.98.085118)

Non-linear biases, stochastically-sampled effective Hamiltonians and spectral functions in quantum Monte Carlo methods

Nick S. Blunt*

University Chemical Laboratory, Lensfield Road, Cambridge, CB2 1EW, U.K.

Ali Alavi

*University Chemical Laboratory, Lensfield Road, Cambridge, CB2 1EW, U.K. and
Max Planck Institute for Solid State Research, Heisenbergstraße 1, 70569 Stuttgart, Germany*

George H. Booth

Department of Physics, King's College London, Strand, London WC2R 2LS, U.K.

In this article we study examples of systematic biases that can occur in quantum Monte Carlo methods due to the accumulation of non-linear expectation values, and approaches by which these errors can be corrected. We begin with a study of the Krylov-projected FCIQMC (KP-FCIQMC) approach, which was recently introduced to allow efficient, stochastic calculation of dynamical properties. This requires the solution of a sampled effective Hamiltonian, resulting in a non-linear operation on these stochastic variables. We investigate the probability distribution of this eigenvalue problem to study both stochastic errors and systematic biases in the approach, and demonstrate that such errors can be significantly corrected by moving to a more appropriate basis. This is lastly expanded to include consideration of the correlation function QMC approach of Ceperley and Bernu, showing how such an approach can be taken in the FCIQMC framework.

I. INTRODUCTION

The introduction of the full configuration interaction quantum Monte Carlo (FCIQMC) method^{1–4} has subsequently seen a large number of new quantum Monte Carlo (QMC) methods introduced by various groups, making use of an FCIQMC-like spawning procedure. Such approaches include, coupled cluster Monte Carlo (CCMC)^{5–7}, density matrix quantum Monte Carlo^{8,9}, model space quantum Monte Carlo (MSQMC)^{10–12}, and recently driven-dissipative quantum Monte Carlo (DDQMC)¹³, to name only a few.

In some cases, these QMC methods are stochastic adaptations of previously-existing deterministic algorithms. Such stochastic adaptations can offer several advantages, perhaps most importantly that stochastic sampling often allows for reduced storage requirements compared to deterministic equivalents, allowing accurate study of extremely large systems in many cases. However, some significant care is required in reformulating a deterministic algorithm as a stochastic one. Many deterministic methods make use of complicated non-linear operations, often including poorly-conditioned problems. In such cases, it would be potentially careless to assume that a given deterministic method can be converted to a Monte Carlo method in a straightforward manner.

Specifically, when estimating a desired quantity, $f(x)$ (for some other underlying quantity x), QMC methods require a large amount of averaging to reduce stochastic errors. For a linear function, $f(x)$ can be estimated by averaging $f(\hat{q})$, where \hat{q} is a random variable being sampled by the QMC method (with $E[\hat{q}] = x$). However, for a non-linear function, it is well known that $E[f(\hat{q})] \neq f(E[\hat{q}])$, and so averaging must be performed

before the function is evaluated, rather than after. In general this will negate the benefits of the sparse Monte Carlo sampling by requiring storage of large parts of the phase space. In FCIQMC, for example, the variable \hat{q} is typically the sampled wave function from a single iteration, and so estimating $E[\hat{q}]$ requires averaging the entire wave function, eventually requiring as much memory as a fully deterministic approach. One must average until $f(E[\hat{q}]) \approx f(x)$ with sufficient accuracy. For some cases this is simple and the QMC approach is successful. For others, it is challenging, or infeasible entirely. A particular issue occurs when the function is ill-conditioned, meaning that any small error in the stochastic estimate will lead to a substantial error in the final result. In such cases, averaging before the function is applied may not even be practical.

Observables in quantum mechanics can always be written as (at most) quadratic expectation values of the wave function. Given the central importance of these *pure* expectation values, projector Monte Carlo methods have long sought to compute these quadratic expectation values in an unbiased fashion, and techniques such as forward-walking and reptation Quantum Monte Carlo have emerged with the aim of removing this non-linear bias^{14–21}. An important example is the calculation of a reduced density matrix (RDM), such as the two-particle RDM $\Gamma_{pq,rs}$,

$$\Gamma_{pq,rs} = \langle \Psi | a_p^\dagger a_q^\dagger a_s a_r | \Psi \rangle, \quad (1)$$

which is quadratic in the wave function $|\Psi\rangle$. Within the FCIQMC approach, early attempts to sample the two-particle RDM in FCIQMC were hampered by the above bias, only slightly improved by partially averaging the wave function prior to evaluating $\Gamma_{pq,rs}$ ²². This

was later corrected by the use of the so-called replica trick^{23–25}, such that the two wave functions in the estimate of $\Gamma_{pq,rs}$ are made statistically independent and the expectation value rewritten as a bilinear functional, formally removing any bias.

While observables are at most quadratic functionals, there often exists the need to compute beyond-quadratic expectation values of stochastically-derived quantities in QMC techniques, such as the computation of entanglement measures, or the sampling and subsequent diagonalization of effective Hamiltonians. This solution to a stochastically-sampled eigenvalue problem is a highly-nonlinear operation, and so there is potential for non-linear biases, and care must be taken. It is these ‘beyond-quadratic’ operations which are the focus of this work, where the above approaches for pure expectation values are not appropriate.

Direct solution of stochastically-derived eigenvalue problems is not uncommon in QMC methods. Perhaps most notably, the linear method^{26–29}, regularly used to optimize wave function parameters in variational Monte Carlo^{30–33} (both in continuum and discrete spaces), requires solution of a Hamiltonian eigenvalue problem in a space spanned by the wave function and its first parameter derivatives. In other examples of non-linear operations, FCIQMC has been recently used to perform complete-active space self-consistent field (CASSCF) calculations^{34,35}, requiring non-linear optimization of orbital coefficients, with no resulting difficulties encountered thus far. The MSQMC method of Tenno^{10–12} obtains excited states by solving a stochastically-sampled eigenvalue problem, once again with no difficulties reported and highly accurate results. Furthermore, QMC techniques have been used to obtain low-energy effective Hamiltonians for correlated materials³⁶.

Perhaps the most challenging quantities to sample in QMC (yet among the most highly-sought after) are dynamical properties, specifically many-body Green’s functions³⁷. FCIQMC, diffusion Monte Carlo (DMC)³¹ and other projector QMC methods perform imaginary-time evolution, allowing access to imaginary-time Green’s functions as quadratic expectation values which can be sampled with the techniques above. However, the transformation from imaginary-time to real-time (or the frequency domain) is highly non-linear and ill-conditioned³⁸, and so challenging to perform in the presence of noise. This transformation is typically performed by maximum-entropy methods^{39–43} which, despite sometimes being accurate, are unsatisfactory in general.

As a further example, we recently introduced a QMC approach to estimate many-body Green’s functions (and finite-temperature and excited-state properties in general), denoted Krylov-projected FCIQMC (KP-FCIQMC)⁴⁴. This method can be loosely characterized as a stochastic version of the Lanczos method, and therefore allows access to dynamical and finite-temperature properties in analogy with dynamical and

finite-temperature Lanczos. Although Lanczos-type algorithms have been performed with QMC previously⁴⁵, we are unaware of their use to this extent, specifically in the calculation of spectral or finite-temperature properties. KP-FCIQMC was applied to study one and two-particle Green’s functions in a one-dimensional Hubbard model, showing that essentially-exact Lanczos results can be reproduced, particularly in low-frequency regions. However, subtle features (particularly at high frequencies) were difficult to reproduce. Clearly, there is an issue with certain eigenvalue problems depending on their nature, with poor conditioning being an obvious potential problem. Given the significant utility of being able to solve such problems, it is worth investigating and discussing such issues, with the KP-FCIQMC approach constituting the exemplar approach for these investigations.

In Section II we re-introduce the KP-FCIQMC method; the theory of many-body Green’s functions is introduced in Section III, including a description of their calculation by KP-FCIQMC. In Section IV we apply this approach to one-dimensional Hubbard models in both the weak and intermediate-coupling regimes, including investigation of probability distributions of the solutions; Section V presents a theoretical model to explain these errors, and so demonstrates a solution using trial wave functions. Section VI extends this idea to the correlation function QMC approach of Ceperley and Bernu⁴⁶, which can partly resolve such biases, although systematic errors grow eventually. In Section VII we discuss how excited-state FCIQMC uses orthogonalization to overcome issues of previous sections, and discuss similarities between KP-FCIQMC and dynamical DMRG in the errors observed.

II. KRYLOV-PROJECTED FCIQMC

A. Defining the Krylov subspace

The Krylov-projected FCIQMC (KP-FCIQMC) method is essentially a stochastic adaptation of the Lanczos method (and other Krylov subspace methods in general). In the Lanczos method, one builds the Hamiltonian eigenvalue problem in the subspace spanned by

$$K = \{|\psi_0\rangle, \hat{H}|\psi_0\rangle, \hat{H}^2|\psi_0\rangle, \dots, \hat{H}^{N_K-1}|\psi_0\rangle\}, \quad (2)$$

where $|\psi_0\rangle$ is some initial state (that can be varied depending on the quantity desired) and \hat{H} is the Hamiltonian operator. The states spanning K are “Lanczos vectors”, or for a more general subspace, “Krylov vectors”. We use the latter terminology, and refer to the span of K as the Krylov subspace. Typically, one constructs the Hamiltonian eigenvalue problem in this subspace (which, for this particular subspace, can be put in an efficient tridiagonal form). In deterministic approaches, the Krylov vectors are orthonormalized (implicit in the tridiagonal form), though this does not alter

their span, but improves the efficiency and numerical stability of the algorithm.

In the FCIQMC method, we sample the vectors $(\mathbb{1} - \Delta\tau\hat{H})^n|\psi_0\rangle$, where n labels the iteration. Thus, the Krylov subspace we work with in KP-FCIQMC is

$$K = \{|\psi_0\rangle, (\mathbb{1} - \Delta\tau\hat{H})^{n_1}|\psi_0\rangle, (\mathbb{1} - \Delta\tau\hat{H})^{n_2}|\psi_0\rangle, \dots, (\mathbb{1} - \Delta\tau\hat{H})^{n_{N_K-1}}|\psi_0\rangle\}, \quad (3)$$

where n_l labels the FCIQMC iteration at which the l 'th Krylov-vector is sampled. If both $\Delta\tau$ and $n_{l+1} - n_l$ are small then the vectors will be similar, potentially leading to near-linear dependencies and poor conditioning. In practice, we therefore chose n_l such that Krylov vectors are chosen more frequently at first (as the wave function varies rapidly, and high-energy states are sampled), and less frequently as the ground state is approached.

For notational convenience, we label the l 'th Krylov vector $|\psi_l\rangle \equiv (\mathbb{1} - \Delta\tau\hat{H})^{n_l}|\psi_0\rangle$, with $|\psi_0\rangle$ the initial vector and $|\psi_{N_K-1}\rangle$ the final vector, with N_K Krylov vectors in total.

B. Estimating the subspace Hamiltonian and overlap matrices

After deciding at which iterations of the FCIQMC algorithm Krylov vectors are to be sampled, the task is to perform those iterations and sample the Hamiltonian and overlap matrix between these vectors. It should be remembered that the ‘vectors’ in this approach are in practice given by sparse, stochastic walker distributions over instantaneously occupied determinants. The overlap matrix is simple, as we store all Krylov vectors next to each other in an array (containing only determinants sampled in at least one Krylov vector, reducing memory requirements significantly). Thus, calculating all overlap elements requires a dot product between each pair of vectors in this array (naturally parallelized as the array storage is already distributed). Thus the overlap matrix is calculated exactly for each pair of vectors. However, each Krylov vector is stochastically sampled, so noise appears in the overlap matrix regardless.

Estimation of the Krylov-projected Hamiltonian is more challenging, but can be formally sampled using the same spawning dynamics as in standard FCIQMC. To calculate $H_{ij}^K \equiv \langle\psi_i|\hat{H}|\psi_j\rangle$ (where \mathbf{H}^K indicates the Hamiltonian projected into the Krylov subspace), one cycles through each determinant in $|\psi_i\rangle$ and performs FCIQMC spawning to sample determinants in $|\psi_j\rangle$. For some small systems, it is possible to calculate \mathbf{H}^K exactly (although we again emphasize that Krylov vectors are stochastically sampled, so that stochastic errors remain).

As discussed in the introduction, the calculation of a quantity like $\langle\psi_i|\hat{H}|\psi_j\rangle$ is biased if $|\psi_i\rangle$ and $|\psi_j\rangle$ are sampled from the same QMC simulation, because $E[\hat{q}_i\hat{q}_j] \neq E[\hat{q}_i]E[\hat{q}_j]$, for random variables \hat{q}_i and \hat{q}_j , unless they

are uncorrelated. This is resolved by using replica sampling, where two FCIQMC simulations are performed simultaneously and independently of each other, removing correlation and therefore bias. We therefore perform two FCIQMC simulations independently, with the first used to sample the ‘bra’ Krylov vectors, $\langle\psi_i|$, and the second used to sample ‘ket’ Krylov vectors, $|\psi_j\rangle$, as used in estimation of both $\langle\psi_i|\psi_j\rangle$ and $\langle\psi_i|\hat{H}|\psi_j\rangle$ elements.

Thus, using replica sampling, estimates of \mathbf{H}^K and \mathbf{S}^K , are essentially unbiased. Perhaps the only cause of ‘bias’ (where ‘bias’ here refers to a systematic discrepancy from an otherwise-identical deterministic calculation) is the use of a shift for population control, which typically leads to a negligible error. When studying probability distribution functions later, we keep the shift constant throughout, to remove any theoretical discrepancy in the estimates of \mathbf{H}^K and \mathbf{S}^K .

C. Comparison to the Lanczos method

We briefly compare the above approach with the traditional Lanczos method. Firstly, we note that the subspace spanned is only formally the same as that in the Lanczos method, Eq. (2), when a Krylov vector is sampled at *every* FCIQMC iteration ($n_l = l$). However, this is not necessary. The only requirement is that the space sampled is sufficient to span the ‘important’ features of the desired quantity. For example, in calculating spectral properties, it is important that the Krylov vectors contain significant contributions from eigenstates with large amplitudes in the spectrum.

More significantly, one may wonder why the Krylov vectors used are not orthogonalized. In the Lanczos algorithm, the subspace Hamiltonian takes exactly a tridiagonal form, which means that only three Lanczos vectors need storing, and also improves numerical stability. For KP-FCIQMC the situation is somewhat different. Firstly, performing this orthogonalization procedure could introduce a bias into the calculation of \mathbf{H}^K and \mathbf{S}^K , compared to a deterministic equivalent, because the orthogonalization operation is non-linear. We have recently introduced the excited-state FCIQMC method, where orthogonalization is used and this potential bias is not observed. Nonetheless the situation here is rather different, as subsequent FCIQMC vectors are nearly identical, whereas in excited-state FCIQMC the orthogonalized vectors are approximately orthogonal beforehand, resulting in only a small change from orthogonalization. Secondly this leads into a practical issue, as orthogonalizing a vector with respect to all previous Krylov vector mostly reduces it to zero, which in a QMC algorithm means killing almost all walkers. In traditional Lanczos one need not ever actually perform the orthogonalization, but in a stochastic setting with approximate vectors, this is not the case, and \mathbf{H}^K would never be exactly tridiagonal regardless.

D. Solution of the subspace eigenvalue problem

The subsequent subspace eigenvalue problem can be solved by usual methods. This problem can be written as

$$\mathbf{H}^K \psi^K = \epsilon \mathbf{S}^K \psi^K. \quad (4)$$

This is solved with a standard canonical Löwdin orthogonalization procedure, transforming to an orthonormal basis with

$$\psi^K = \mathbf{U} \mathbf{D}^{-1/2} \psi^L. \quad (5)$$

where \mathbf{U} is the matrix with eigenvectors of \mathbf{S}^K in its columns, and \mathbf{D} is the matrix with corresponding overlap-matrix eigenvalues on the diagonal. We call this new basis the Löwdin basis, using superscript label L . The Hamiltonian eigenvalue problem is then

$$\mathbf{H}^L \psi^L = \epsilon \psi^L, \quad (6)$$

with

$$\mathbf{H}^L = \mathbf{D}^{-1/2} \mathbf{U}^T \mathbf{H}^K \mathbf{U} \mathbf{D}^{-1/2}. \quad (7)$$

As discussed above, because many of the Krylov vectors will be nearly linearly dependent, many overlap matrix eigenvalues will be very small, for the exact deterministic problem. In the stochastic problem, the overlap matrix will not take the form of a true overlap matrix, and so will even have negative eigenvalues. We therefore throw away eigenvectors of \mathbf{S}^K with negative or very small eigenvalues. Typically, we keep 10 – 15 eigenvectors. These remaining eigenvectors form the final Löwdin subspace, with dimension denoted N_L , such that $N_L < N_K$. Restricting N_L to only 10 – 15 means that we only obtain the same number of final eigenvector solutions, which is an approximation to the true spectrum (as it is also for the dynamical Lanczos approach), but is nonetheless useful for many situations, for reasons clarified in the next section.

III. MANY-BODY GREEN'S FUNCTIONS FROM KP-FCIQMC

Green's functions are key observables in many-body condensed matter physics and chemistry, which are directly accessible experimentally, and often allow detailed examination and understanding of important many-body phenomena. As such, their accurate and routine calculation represents a high-priority goal. Unfortunately they are typically difficult to calculate accurately, particularly through QMC methods.

KP-FCIQMC, like dynamical Lanczos, can access two-body Green's functions, but here we focus on the often-studied single-particle (retarded, ground-state) Green's

function, defined by

$$G(k, \omega) = \langle \Psi_0 | a_{k\sigma} \frac{1}{\omega + \mu - (\hat{H} - E_0) + i\delta} a_{k\sigma}^\dagger | \Psi_0 \rangle + \langle \Psi_0 | a_{k\sigma}^\dagger \frac{1}{\omega + \mu + (\hat{H} - E_0) + i\delta} a_{k\sigma} | \Psi_0 \rangle, \quad (8)$$

where ω is the frequency, μ is the chemical potential, $a_{k\sigma}^\dagger$ is a creation operator for a particle with momentum k and spin label σ , and E_0 and $|\Psi_0\rangle$ are the ground-state energy and eigenvector (in the N -particle sector), respectively. δ is a broadening added to give finite width to poles, but is also related to the rate at which the created quasi-particles decay in an experiment measuring $G(k, \omega)$, so is physically justified. In the non-interacting limit this reduces to the often-quoted expression

$$G(k, \omega) = \frac{1}{\omega + \mu - \xi_k + i\delta}, \quad (9)$$

for single-particle eigenvalues ξ_k . This is the starting point for expanding the interacting Green's function in terms of a self energy, as performed in methods such as dynamical mean field theory (DMFT)^{47–49}, but we take a more direct approach here.

The quantity of interest studied here is the spectral function, $A(k, \omega)$, proportional to the imaginary part of the Green's function, typically with $-1/\pi$ as the proportionality factor to ensure normalization,

$$A(\omega) = -\frac{1}{\pi} \Im[G(\omega)]. \quad (10)$$

Inserting a resolution of the identity in the appropriate $(N+1)$ or $(N-1)$ -particle sector gives

$$A(k, \omega) = \sum_i \frac{\delta}{\pi} \frac{|\langle \Psi_i^{N+1} | a_{k\sigma}^\dagger | \Psi_0 \rangle|^2}{[\omega + \mu - (E_i^{N+1} - E_0)]^2 + \delta^2} + \sum_i \frac{\delta}{\pi} \frac{|\langle \Psi_i^{N-1} | a_{k\sigma} | \Psi_0 \rangle|^2}{[\omega + \mu + (E_i^{N-1} - E_0)]^2 + \delta^2}. \quad (11)$$

where $N+1$ and $N-1$ superscripts denote states in the corresponding sectors in the obvious way. In the limit $\delta \rightarrow 0^+$, this becomes a sum of delta functions at the energy eigenvalues, with weights given by the transition probabilities $|\langle \Psi_i^{N+1} | a_{k\sigma}^\dagger | \Psi_0 \rangle|^2$ and $|\langle \Psi_i^{N-1} | a_{k\sigma} | \Psi_0 \rangle|^2$. In the non-interacting limit, this reduces to

$$A(k, \omega) = \delta(\omega + \mu - \xi_k), \quad (12)$$

and so $A(k, \omega)$ maps out the single-particle bandstructure.

We also consider the local single-particle density of states (DOS), defined as

$$A(\omega) = \frac{1}{N} \sum_k A(k, \omega), \quad (13)$$

corresponding to a Fourier transform to the real space spectral function at the origin.

We now consider how these quantities can be calculated in KP-FCIQMC. The Lehmann representation expresses the spectral function as a sum over eigenstates, Eq. (11). The dynamical Lanczos and KP-FCIQMC methods give access to a number of eigenstates equal to the size of the subspace studied, N_L in the case of KP-FCIQMC. This will typically be small compared to the full Hilbert space dimension, and so this may appear unhelpful. However, the subspace is chosen such that the important eigenstates will appear as solutions. The important eigenstates are those with large probability amplitudes, $|\langle \Psi_i^{N+1} | a_{k\sigma}^\dagger | \Psi_0 \rangle|^2$ and $|\langle \Psi_i^{N-1} | a_{k\sigma} | \Psi_0 \rangle|^2$, as these give the largest contributions to Eq. (11). Thus, the key is that the Krylov subspace is chosen using $a_{k\sigma}^\dagger | \Psi_0 \rangle$ as the initial vector to obtain states in the $(N+1)$ -particle sector, and $a_{k\sigma} | \Psi_0 \rangle$ as the initial Krylov state for the $(N-1)$ -particle sector. Thus, those eigenstates with large contributions in Eq. (11) will have relatively large components in the Krylov subspaces, and may be extracted accurately.

The KP-FCIQMC calculation is started from a perturbed ground state. As such, an FCIQMC calculation is first performed to obtain a stochastic sampling of the ground state. The perturbation is applied from there, and this whole process may be repeated to allow averaging of the (unbiased) \mathbf{H}^K and \mathbf{S}^K , before the (biased) eigenvalue estimate are obtained. This averaging should reduce the *systematic* error in the final eigenvalue estimates and probability amplitudes, thus improving the quality of spectra. This key aspect will be studied in the following sections.

IV. KP-FCIQMC RESULTS

A. One-body Green's functions from KP-FCIQMC

We study the one-dimensional, periodic 14-site Hubbard model at half-filling. In our original KP-FCIQMC study, we took $U/t = 2$, but here we take $U/t = 1$ and $U/t = 4$ to study both the nearly-free and intermediate-coupling regimes.

Results for $U/t = 1$ are presented in Fig. (1). This case is fairly trivial: in this low-coupling regime, the band-structure is close to the non-interacting bandstructure, where only one or two eigenstates contribute from each K sector. The stochastically-sampled results here accurately reproduce results from dynamical Lanczos. The number of walkers used was 10^5 . The initiator adaptation was used, although initiator error is negligible. The semi-stochastic adaptation was also used with a deterministic space of dimension $|D| = 5 \times 10^4$, chosen using the population-based scheme of Ref. (50). We set $\delta = 0.05t$. There are only one or two states of significance to sample in each K sector, which are well-sampled in the Krylov vectors, and effectively no bias is noticeable.




FIG. 1. (a) $A(k, \omega)$ from $k = -\frac{6}{7}\pi$ (bottom) to $k = \pi$ (top) for the 1D, periodic 14-site Hubbard model at $U/t = 1$, compared to dynamical Lanczos. (b) The local density of states. In this low-coupling regime, the situation is simple and easy to calculate by KP-FCIQMC, with only a small number of low-energy states making significant contributions.

In Fig. (2), the same system is studied but with $U/t = 4$. Here the situation is more challenging, with a larger number of eigenstates making a significant contribution, and low-energy states generally making a smaller contribution. We use the same parameters as for the $U/t = 1$ case, but study walker populations, N_w , of both 10^5 and 10^6 . At $N_w = 10^5$ there are clear errors in the whole spectrum. This could be due to a number of factors, but we believe that initiator error is the biggest issue, as initiator error is significant even in ground-state FCIQMC calculations. This is largely resolved at $N_w = 10^6$, where KP-FCIQMC results are in good agreement with dynamical Lanczos.

This shows that the stochastic Krylov-space approach presented here can be successful and accurate, even in the intermediate coupling regime. Because of the sparse sampling (and therefore relatively-low memory requirements), this approach is possible for systems beyond those treatable by dynamical Lanczos (although in this particular example, the density matrix renormalization group (DMRG) method would be a more efficient approach due to its one-dimensional nature). Results are particularly accurate at low $|\omega|$, allowing crucial properties around the band gap to be studied.

Nonetheless, it is informative to study the convergence of this approach to the exact results. At $N_w = 10^6$, results are in good agreement with dynamical Lanczos, but error remains, particularly at high $|\omega|$. We have performed multiple repeated calculations on this system and found that results at high frequency are largely unchang-

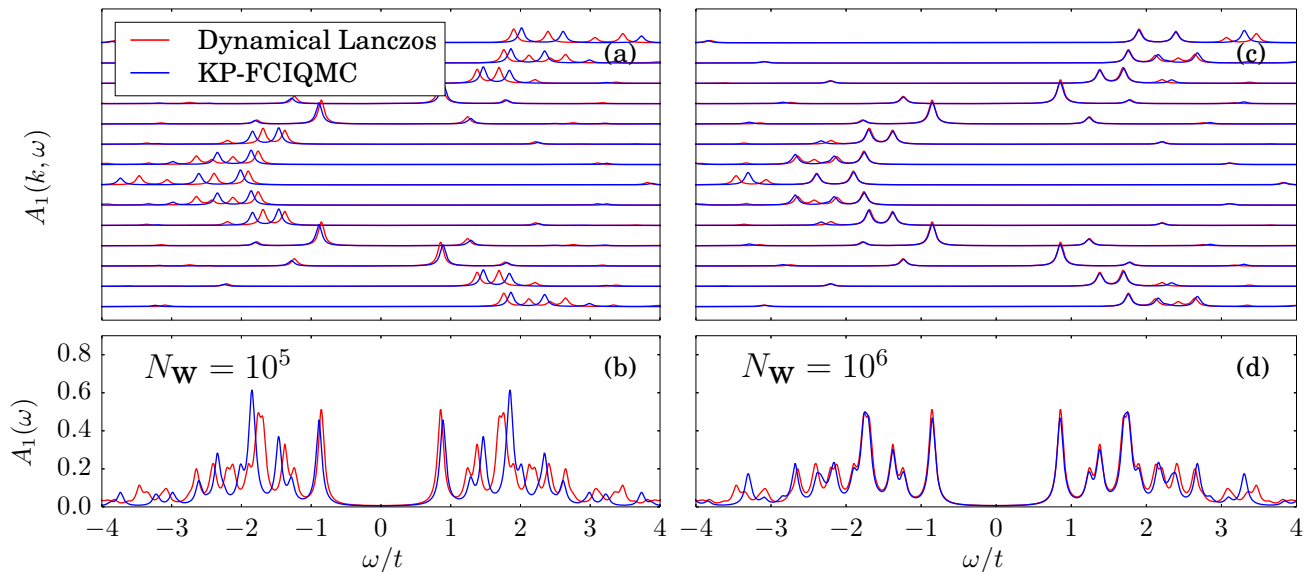


FIG. 2. Results for the 1D, periodic 14-site $U/t = 4$ Hubbard model, using walker populations, N_w , of both 10^5 , (a) and (b), and 10^6 , (c) and (d). At $N_w = 10^5$, noticeable error exists, which is significantly corrected at $N_w = 10^6$. We believe that this error is largely due to initiator error, which is significant in the Hubbard model at $U/t = 4$, although other systematic biases exist. Even at $N_w = 10^6$, errors are more noticeable in the high-frequency regime, which we find to be systematic.

ing, suggesting a systematic bias rather than statistical error. In the limit of exact sampling, both \mathbf{H}^K and \mathbf{S}^K will be obtained exactly, and so will their eigenvalues, but it is informative to study how this convergence comes about.

B. Probability distributions from stochastically-sampled eigenvalue problems

To study this problem in the most detail, we investigate the probability distributions functions (PDFs) of the underlying eigenvalue estimates, and also PDFs for transition amplitudes, $\langle \psi_i | a_k^\dagger | \Psi_0 \rangle$. Probability distributions of QMC estimates are rarely considered; it is usually assumed that such distributions are Gaussian due to the central limit theorem, but this is not clear for solutions of poorly-conditioned problems. However, to construct such a distribution requires performing a large number of repeated calculations (here we perform $\sim 10^4$). We therefore consider a much smaller system: the 1D, periodic 6-site Hubbard model at $U/t = 4$. This somewhat mimics the above case, but now in a system where only a small N_w is required. We also only consider the $K = 2\pi/3$ sector, again mimicking a sector with poor results in the above 14-site case. We use $N_w \sim 100$ and a deterministic space of the Hartree–Fock determinant and all single and double excitations.

For this system, there are only three eigenstates with significant contributions to the single-particle spectrum. Therefore, to make the probability distributions clear to view and interpret, we take $N_L = 3$, i.e., projecting the

eigenvalue problem into a space of dimension 3, such that a stochastically-sampled 3×3 eigenvalue problem is obtained. We compare the eigenvalue estimates to those from a completely deterministic KP-FCIQMC calculation, but with the same projection into a 3×3 problem performed, so that exactly the same eigenvalues will be obtained in the limit of infinite averaging in the stochastic case.

We perform $\sim 10^4$ repeated KP-FCIQMC calculations, each with a different random number generator (RNG) seed. Within each of these 10^4 repeats, we average \mathbf{H}^K and \mathbf{S}^K over either 1, 10, 100 or 1000 repeats, *before* the eigenvalue problem is solved. As more averaging of the unbiased \mathbf{H}^K and \mathbf{S}^K is performed, any bias in the eigenvalue estimates should be reduced. The shift is kept constant to avoid any theoretical population-control discrepancies.

The constructed PDFs are presented in Fig. (3). Results on the left are for the ground-state solutions, while those on the right are for the two excited states. Results at the top show PDFs for the energy eigenvalues themselves, while those at the bottom are PDFs for the transition amplitudes, $\langle \psi_i | a_k^\dagger | \Psi_0 \rangle$. PDFs for the ground state are sensible for all levels of averaging performed. There is a slight skew occurring when no averaging is performed (i.e., 1 ‘repeat’), although even in this case the PDF mean is approximately correct.

In contrast, PDFs for the two excited states show results in significant error. For energy eigenvalues (top) with no averaging, the first excited state is slightly too high, while the second eigenstate is in error on average by more than $4t$, with a variance so large that results are

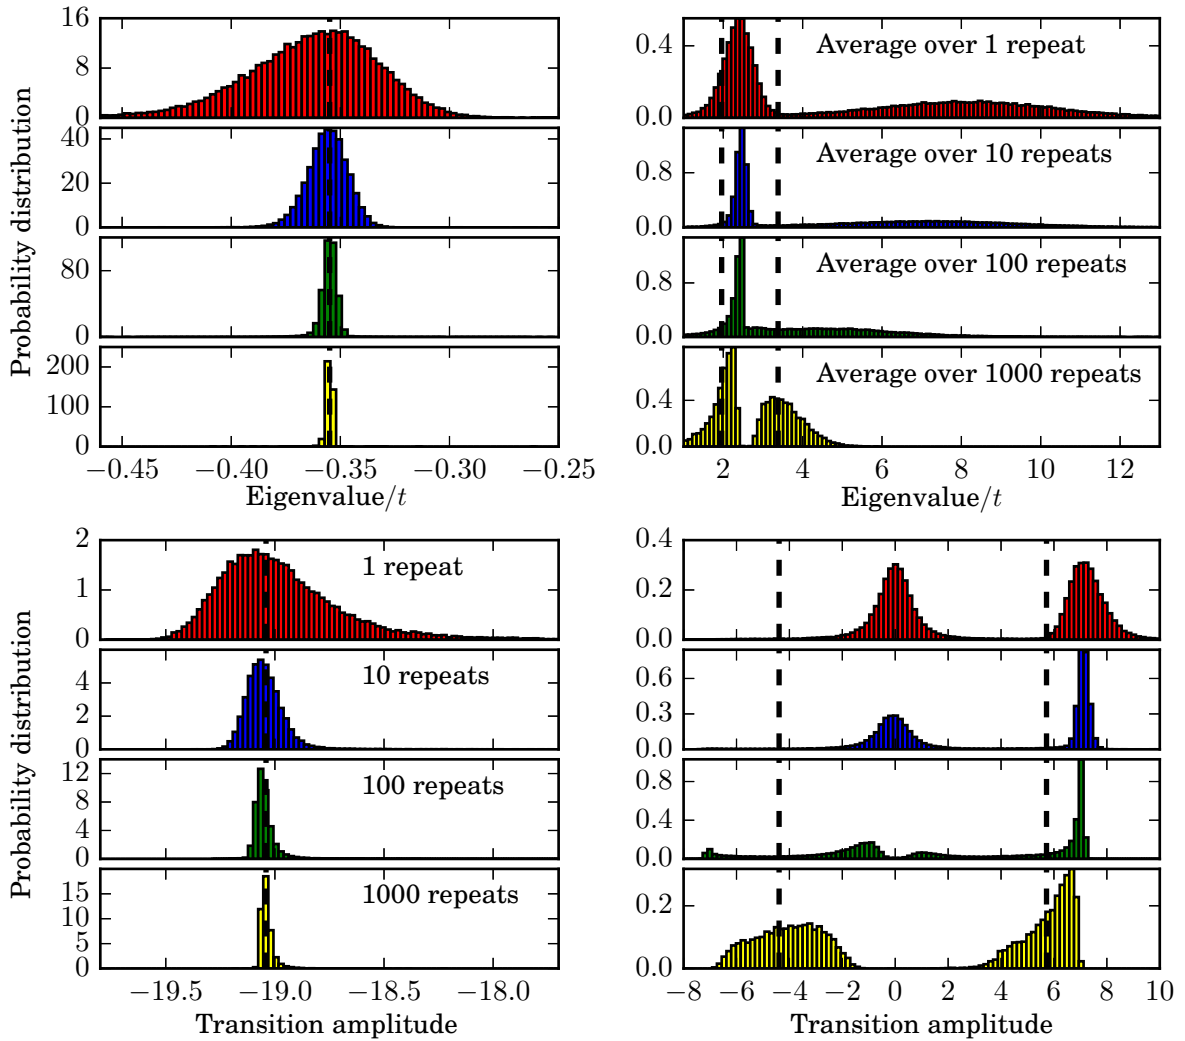


FIG. 3. Probability distributions for eigenvalues and transition amplitudes from KP-FCIQMC. Dashed lines show exact values. We take a trivial system, the 1D, periodic 6-site Hubbard model at half-filling and $U/t = 4$, in the $K = 2\pi/3$ sector. This allows investigation of probability distributions for solutions of stochastically-sampled eigenvalue problems. We aim to study only the 3 dominant eigenvalues in the spectrum. Results on the left are for the ground state, and on the right for the two excited states. The stated number of repeats is the number of repeats over which \mathbf{H}^K and \mathbf{S}^K are averaged *before* the eigenvalue problem is solved. Since the Hamiltonian and overlap matrices themselves are unbiased, this should reduce biases in eigenvalue solutions. This is indeed found to be the case, with reasonable distributions once averaging over 1000 repeats is performed. However, for less averaging, significant biases occur for the excited-state estimates (right), in both eigenvalues and transition amplitudes.

smeared out entirely. Increasing averaging up to 1000 repeats does eventually bring results to be distributed about approximately the correct values, although even then with a strong skew.

Transition amplitudes for excited states are also in significant error until significant averaging is performed. For low-averaging, one of the excited states has a transition amplitude which is approximately 0 (while the other is too large, such that two peaks are merged into one), which does not begin to resolve until averaging over 100 repeats, and not satisfactorily until averaging over 1000 repeats. Clearly, even though this is a simple system, significant biases occur in the eigenvalue estimates of the excited states, although ground-state estimates are accu-

rate. While the number of walkers is similarly small with the system size, this should act as a warning for the issues that can occur in stochastically-sampled eigenvalue problems, even in simple cases. Despite this, given the utility of such solutions in QMC methods, this issue is usually not so severe.

V. ANALYSIS OF ERROR FOR A TWO-STATE MODEL

We now consider a theoretical model in only two dimensions, where the source of this significant bias can be identified.

Suppose we sample two Krylov vectors of the form

$$|K_0\rangle = |\psi_0\rangle \quad |K_1\rangle = |\psi_0\rangle + \delta|\psi_1\rangle, \quad (14)$$

where $|\psi_0\rangle$ and $|\psi_1\rangle$ are orthonormal solutions which we seek to obtain, and δ is some small and positive number. This mimics what happens in real simulations where there is one vector (the ground state) making up a relatively large component of all Krylov vectors. Then the eigenvalue problem is

$$\begin{pmatrix} a & a \\ a & a + \delta^2 b \end{pmatrix} \begin{pmatrix} x_0 \\ x_1 \end{pmatrix} = \epsilon \begin{pmatrix} 1 & 1 \\ 1 & 1 + \delta^2 \end{pmatrix} \begin{pmatrix} x_0 \\ x_1 \end{pmatrix}, \quad (15)$$

where $a = \langle\psi_0|H|\psi_0\rangle$, $b = \langle\psi_1|H|\psi_1\rangle$ and $\langle\psi_0|H|\psi_1\rangle$ has been defined as 0 for simplicity.

It can be shown that the eigenvalues of the overlap matrix are (using zero-indexing),

$$D_{00} = 2 + \mathcal{O}(\delta^2) \quad D_{11} = \frac{\delta^2}{2} + \mathcal{O}(\delta^4) \quad (16)$$

and that the corresponding eigenvectors are

$$\begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} + \mathcal{O}(\delta^2) \quad \begin{pmatrix} \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} + \mathcal{O}(\delta^2). \quad (17)$$

The Hamiltonian in the Löwdin basis (the final orthonormal basis) is

$$\mathbf{H}^L = \begin{pmatrix} a + \mathcal{O}(\delta^2) & \mathcal{O}(\delta) \\ \mathcal{O}(\delta) & b + \mathcal{O}(\delta^2) \end{pmatrix}. \quad (18)$$

The eigenvalues of the exact \mathbf{H}^L matrix are a and b , as expected by construction.

If this eigenvalue problem were sampled by KP-FCIQMC, then the stochastic estimate of the Krylov-space Hamiltonian could be written

$$\mathbf{H}^K + \boldsymbol{\eta}, \quad (19)$$

where \mathbf{H}^K is the *exact* Krylov-space Hamiltonian, and $\boldsymbol{\eta}$ is the error matrix (and symmetry is enforced on the Hamiltonian estimate, so that $\eta_{10} = \eta_{01}$). If the transformation matrix from the Krylov to the Löwdin basis is denoted \mathbf{T} , then the stochastic estimate of the Krylov-projected Hamiltonian transforms as

$$\mathbf{T}^T(\mathbf{H}^K + \boldsymbol{\eta})\mathbf{T} = \mathbf{T}^T\mathbf{H}^K\mathbf{T} + \mathbf{T}^T\boldsymbol{\eta}\mathbf{T}. \quad (20)$$

The first term is the desired \mathbf{H}^L . We would like second term to equal $\mathbf{0}$. For finite stochastic error, however, this term can be shown to equal (to leading order in δ)

$$\mathbf{T}^T\boldsymbol{\eta}\mathbf{T} = \begin{pmatrix} \frac{1}{4}(\eta_{00} + 2\eta_{01} + \eta_{11}) & \frac{1}{2\delta}(\eta_{11} - \eta_{00}) \\ \frac{1}{2\delta}(\eta_{11} - \eta_{00}) & \frac{1}{\delta^2}(\eta_{00} - 2\eta_{01} + \eta_{11}) \end{pmatrix}. \quad (21)$$

Therefore, \mathbf{H}_{11}^L has a large error when δ is small. Roughly speaking, to obtain an accurate estimate of the Hamiltonian in the final orthonormal basis, the Hamiltonian




FIG. 4. Distributions of KP-FCIQMC eigenvalues around the exact ground and first-excited energies (dashed lines) for the same system as Fig. (3), as the number of repeats (over which \mathbf{H}^K and \mathbf{S} are averaged) is varied. In contrast to Fig. (3), the initial Krylov vector was here taken as the CISD estimate of the first excited state. Both ground and first excited states are obtained with relatively small biases for 1 repeat, and with no visible bias when averaging over 10 or more repeats.

in the Krylov basis must have errors η_{ij} smaller than $\delta^2 \approx D_{11}$. This makes more rigorous the intuitive notion that eigenvectors with a small component in the Krylov vectors require a similarly small associated stochastic error for accurate estimation, quickly becoming unreasonable. In this case it is difficult to ‘extract’ the $|\psi_1\rangle$ solution accurately from the noise.

The above analysis suggests that the main source of bias is due to the desired states having small components in the Krylov vectors. It is therefore informative to study different subspaces, where the Krylov vectors are constructed to be similar to the desired excited states.

Such an example is demonstrated in Fig. (4). Here we take the same system as that studied in Fig. (3). However, while Fig. (3) took the initial Krylov vector as $a_k^\dagger|\Psi_0\rangle$, where $|\Psi_0\rangle$ is the exact ground-state wave function in the N -particle sector ($N = 6$) (appropriate for constructing the single-particle Green’s function), here we take the initial Krylov vector as the CISD estimate to the first excited state in the same final sector ($N = 7$). As can be seen, by choosing the subspace more appropriately, the bias is effectively removed in the first excited state estimates. Note that the numerical value for the first excited state is different between Figs. (3) and (4) - this is because we project into different final 3-dimensional subspaces in each case. However, this is unimportant; here we are simply assessing biases in the different approaches.

VI. A CORRELATION FUNCTION QMC APPROACH IN FCIQMC

The approach considered in Fig. (4) overcame systematic errors by constructing the subspace in terms of trial wave functions. This approach is also taken in the correlation function quantum Monte Carlo (CFQMC) method of Ceperley and Bernu^{46,51–55}, which we now briefly consider within an FCIQMC context. Note that this approach is called correlation function QMC because the sampled H_{ij}^C and S_{ij}^C (see below) can be viewed as imaginary-time correlation functions. However, in this method one is only interested in eigenvalues and properties of a small number of low-lying excited states, *not* Green's functions, where the entire spectrum is (in principle) obtained. Nonetheless the approach is interesting to consider in relation to the current question of biases and stochastically-sampled eigenvalue problems.

Suppose one has m trial functions, $\{|f_0\rangle, |f_1\rangle, \dots, |f_{m-1}\rangle\}$, for the m lowest energy eigenstates of a system. Improved solutions can be formed by taking linear combinations of the original set of states. The best linear combinations (in a variational sense) are formed by solving the Hamiltonian eigenvalue problem, projected into the space spanned by $\{|f_i\rangle\}$,

$$\mathbf{H}^C \phi_i = \Lambda_i \mathbf{S}^C \phi_i, \quad (22)$$

where $H_{ij}^C = \langle f_i | \hat{H} | f_j \rangle$ and $S_{ij}^C = \langle f_i | f_j \rangle$. The eigenvalue estimates, Λ_i are variational by MacDonald's theorem⁵⁶ and so will be improved by using more or better-quality trial states, $|f_i\rangle$.

Define projected trial states by

$$|f_i(\tau)\rangle = e^{-\tau \hat{H}} |f_i\rangle. \quad (23)$$

If the Hamiltonian eigenvalue problem is evolved with τ , so that $H_{ij}^C(\tau) = \langle f_i(\tau) | \hat{H} | f_j(\tau) \rangle$, $S_{ij}^C(\tau) = \langle f_i(\tau) | f_j(\tau) \rangle$ and

$$\mathbf{H}^C(\tau) \phi_i(\tau) = \Lambda_i(\tau) \mathbf{S}^C(\tau) \phi_i(\tau), \quad (24)$$

then since the slowest decaying contributions to each trial state will span the lowest-lying excited states, in principle, the exact energy eigenstates and eigenvalues are retrieved in the limit of large imaginary time, as

$$\lim_{\tau \rightarrow \infty} \Lambda_i(\tau) = E_i. \quad (25)$$

The above approach is that taken in the CFQMC method.

However, the approach above is only stable in the complete absence of numerical errors, either due to finite precision, or stochastic noise. This is because for large τ , all states $|f_i(\tau)\rangle$ will converge to the ground state, and the trial states will exponentially become linearly dependent, and eigenvalues of the overlap matrix will become exponentially small. Stochastic or numerical errors in \mathbf{H}^C will therefore be greatly magnified, and large biases will




FIG. 5. FCIQMC energies (from a variational estimator) during propagation for the neon atom in an aug-cc-pVDZ basis, with 2 core electrons frozen. Simulations begin from the CISD estimates of the five lowest excited states (labelled $E_1 - E_5$). Energies converge quite stably to exact energies (dashed lines) for considerable imaginary time, before converging towards the ground state. Note that three states appear to diverge to higher energies, which occurs when the ground state component gains an opposite sign in the two replica simulations, resulting in the denominator of the energy estimator passing through zero. However, the ground state is converged upon eventually (not shown). Note also that there are double degeneracies in two states because the A_g irrep of D_{2h} is used.

result, as for KP-FCIQMC, resulting in a transient estimate of the excited eigenstates in practice. However, it is usually possible to converge far enough to obtain good eigenvalue estimates before such issues become unmanageable.

In the original approach of Ceperley and Bernu⁴⁶, this procedure was performed in real space. However, it is just as simple to perform in an FCIQMC-framework, as there one also performs imaginary-time propagation of initial states. Indeed, the approach is perhaps more convenient here, as it is simple to construct particularly-accurate trial excited states in finite-dimension Hilbert spaces, using standard quantum chemistry methods.

Before applying the full CFQMC approach, we first consider what happens when FCIQMC simulations are performed starting from trial solutions to excited states. This has already been performed in the excited-state FCIQMC approach, but in that approach orthogonalization prevents collapse to the ground state. Here we consider propagation without orthogonalization. An example of this is considered in Fig. (5), where the Ne atom is considered in an aug-cc-pVDZ basis set for the five lowest excited states, initialized from configuration interaction singles and doubles (CISD) solutions to the states. We use the following energy estimator,

$$E_i(\tau) = \frac{\langle f_i(\tau) | \hat{H} | f_i(\tau) \rangle}{\langle f_i(\tau) | f_i(\tau) \rangle}, \quad (26)$$

calculated in the same way as in KP-FCIQMC (using the spawning procedure to sample pairs of determinants,

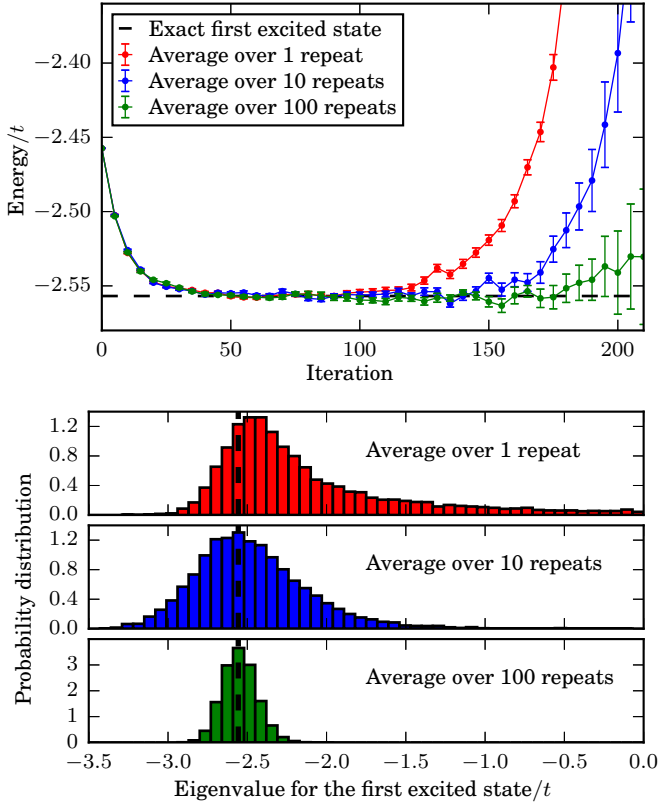


FIG. 6. Results from the CFQMC procedure for the first excited state, for the one-dimensional, periodic 6-site Hubbard model at half-filling and $U/t = 2$. Simulations began from CISD estimates to the ground and first excited states. \mathbf{H}^C and \mathbf{S}^C were averaged over 1, 10 and 100 repeated simulations *before* eigenvalues were obtained. Top: Mean eigenvalues estimates as a function of imaginary-time, with standard deviations presented as error bars. Eigenvalue bias increases exponentially for large τ . Bottom: PDFs of eigenvalues for the first excited state at iteration 200 ($\tau = 2$).

and the replica trick). Results of Fig. (5) show that, perhaps surprisingly, the energy estimates converge accurately towards the exact eigenvalues for over 6000 iterations before eventually converging to the ground-state energy. One might have expected the collapse towards E_0 to happen much quicker, especially starting from fairly basic CISD estimates, even for this simple system. Note that when collapse to the ground state does occur, some energy estimates diverge to higher values, rather than directly to the ground state. This is because the component of the ground state in the two replicas sampling $|f_i(\tau)\rangle$ obtain opposite signs in early iterations due to stochastic noise, causing the denominator, $\langle f_i(\tau)|f_i(\tau)\rangle$, to pass through 0 during convergence.

The CFQMC procedure then projects \hat{H} into the subspace spanned by these states, and solves the resulting eigenvalue problem to prevent this collapse. An example of this is shown in Fig. (6). Here, to allow study of PDFs, we again study the same 6-site Hubbard model as

in Figs. (3) and (4), but here taking $U/t = 2$. We consider the first excited state. While the collapse of the expectation values is prevented in the zero-stochastic-error (and infinite numerical precision) limit, in a stochastic setting the bias on the eigenvalues grows exponentially, as proven by Ceperley and Bernu in their original presentation⁴⁶. This exponential growth of bias is clear here. When performing no averaging, significant error occurs by iteration 200 ($\tau = 2$), although this can be corrected by averaging to reduce noise on \mathbf{H}^C and \mathbf{S}^C .

Despite the eventual breakdown, the first excited state is sampled with a near-exact energy for a considerable period of imaginary time before this occurs. This situation is therefore significantly improved compared to that in KP-FCIQMC, due to the use of trial solutions which are already close to spanning the target subspace, resulting in a well-conditioned problem where the first excited state has significant amplitude in the sampled subspace.

VII. DISCUSSION

A. Excited-state FCIQMC

The exponentially-growing error in CFQMC with τ is caused by all states converging to the ground state, resulting in a poorly-conditioned eigenvalue problem. The obvious solution to this is to orthogonalize each state against all lower-energy states, to prevent collapse. This approach was taken in our excited-state extension to FCIQMC⁵⁷. The projection here takes the form

$$|f_i(\tau + \Delta\tau)\rangle = \hat{O}_i(\tau + \Delta\tau)[1 - \Delta\tau(\hat{H} - S\mathbb{1})]|f_i(\tau)\rangle, \quad (27)$$

with

$$\hat{O}_i(\tau) = \mathbb{1} - \sum_{j < i} \frac{|f_j(\tau)\rangle\langle f_j(\tau)|}{\langle f_j(\tau)|f_j(\tau)\rangle}. \quad (28)$$

Here, the initial ground-state trial wave function, $|f_0\rangle$, is evolved exactly as in standard FCIQMC. All higher-energy states follow a similar evolution, but with $\hat{O}_i(\tau)$ applied after each iteration.

With biases in mind, a possible concern in this approach is that the orthogonalization operator $\hat{O}_i(\tau)$ is non-linear in $|f_j(\tau)\rangle$, so that biases may occur. This was considered in some detail in our initial presentation. Although we believe that some extreme limits must exist where the non-linear nature leads to undesired results, in practice we have never found this issue to occur and indeed have obtained extremely accurate results in all cases tested. For a test Hubbard model example, any bias was shown to be less than $\sim 10^{-4} - 10^{-6}t$ for all states studied. For the stretching of C_2 in a cc-pVQZ basis we found agreement with DMRG results to $\sim 10^{-4}E_h$, with remaining error almost certainly from the initiator approach.

As a demonstration, we again consider the test case of the one-dimensional 6-site Hubbard model at half-filling and $U/t = 2$, as in Fig. (6). Simulations are




FIG. 7. Results for the first excited state of the one-dimensional, periodic 6-site Hubbard model at half-filling and $U/t = 2$. Simulations were initialized from CISD wave functions. Free propagation is shown in red, where large errors occur quickly. The CFQMC procedure is shown in blue (here \mathbf{H}^C and \mathbf{S}^C were averaged over 10^3 repeated simulations before the plotted eigenvalues were obtained), demonstrating improved stability. The excited-state FCIQMC approach is shown in green, where orthogonalization ensures stable convergence to the exact result. Note that under free propagation the ground state is eventually converged upon, although divergence to higher energies is observed here; this occurs for the same reason as in Fig. (5).

initialized from CISD wave functions, with the aim of sampling the exact first excited state. A comparison is performed between three cases: free FCIQMC propagation; the CFQMC procedure; and the excited-state FCIQMC procedure, performing orthogonalization against the ground-state FCIQMC wave function. For excited-state FCIQMC results, an RDM-based energy estimator is used (with replica sampling)²⁵, for a relevant comparison. CFQMC is more stable than free propagation, allowing the first excited state to be accurately sampled for a longer period of imaginary time. The orthogonalization-based approach is fully stable. Continuing this simulation for 10^6 iterations (using ~ 80 walkers on average) gives an energy estimate of $-2.55677(23)t$, compared to the exact value of $-2.55683t$, showing no bias to good accuracy.

A similar orthogonalization approach would not be appropriate in real-space projector QMC methods, as the overlap between two stochastically sampled wave functions in real space will be zero. Clearly there are advantages and disadvantages to both real-space and finite-space approaches. The ability to calculate overlaps between two statistically-sampled wave functions, as used in both the excited-state approach and the unbiased sampling of RDMs, has been a large benefit in FCIQMC. As such, the problem of sampling a small number of low-lying states by FCIQMC has, we believe, been effectively solved by the above orthogonalization approach. However, the task of computing dynamical properties remains a more significant and open challenge.

B. Comparison of KP-FCIQMC with DMRG approaches

The KP-FCIQMC approach can be compared to similar approaches for calculating dynamical correlation functions in the DMRG framework. Several methods have been attempted by a variety of approaches^{58–62}. An approach based on the Lanczos algorithm was attempted by Hallberg⁵⁸, which is directly comparable to the KP-FCIQMC approach. A similar approach has been used recently by Dargel *et al.*^{63,64} This approach was also investigated by Kühner and White in 1999⁶⁰, who interestingly came to similar conclusions in DMRG as we have in FCIQMC - they state that “the Lanczos vector method works very well if only the low-energy part of the correlation function is of interest, or if the bulk of the weight is in one single peak.” This makes evident that these discrepancies are not simply the result of stochastic-type errors; small systematic errors due to other approximations are equally problematic.

VIII. CONCLUSION

We have investigated examples of sampling non-linear functions in QMC methods, including ill-conditioned problems. KP-FCIQMC allows sampling of dynamical and finite-temperature properties by an approach comparable to the Lanczos method, but with memory limitations removed due to stochastic sampling. For small Hubbard model examples we obtained accurate spectra even in the intermediate-coupling regime, demonstrating the potential of this stochastic Krylov-projected approach, which may be appropriate as an impurity solver for DMFT.

However, in results similar to those obtained from a DMRG-based approach to the spectral Lanczos method⁶⁰, we find that high-energy spectral features are challenging to obtain accurately, particularly for poles with small transition amplitudes. This was demonstrated in a small 6-site Hubbard model, where probability distributions were obtained for the KP-FCIQMC eigenvalues and transition amplitudes, showing large errors in relatively high-energy states, eventually corrected by performing further averaging. A simple two-state theoretical model was considered to explain this issue, showing that the problem can be traced back to states with small components in the Krylov vectors, leading to an ill-conditioned problem, for which QMC is not best-suited.

By moving to a more appropriate subspace, where a desired excited state has large components in the Krylov vectors, this bias can be largely removed. An example of this is the CFQMC method, which largely resolves such issues for small τ , although biases eventually become significant at large imaginary time.

A large number of stochastic quantum chemistry methods have been formulated in recent years in a similar vein to FCIQMC; we hope that the considerations presented

here will be informative, and beneficial in the appropriate formulation of such approaches in the future.

IX. ACKNOWLEDGMENTS

N.S.B. gratefully acknowledges St John's College, Cambridge for funding through a Research Fellowship. A.A. thanks EPSRC for funding under Grant No.

EP/J003867/1. G.H.B. gratefully acknowledges funding from the Air Force Office of Scientific Research via grant number FA9550-16-1-0256, and the Royal Society via a University Research Fellowship. The project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 759063. We are also grateful to the UK Materials and Molecular Modelling Hub for computational resources, which is partially funded by EPSRC (EP/P020194/1).

-
- * nicksblunt@gmail.com
- ¹ G. H. Booth, A. J. W. Thom, and A. Alavi, *J. Chem. Phys.* **131**, 054106 (2009).
 - ² D. M. Cleland, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **132**, 041103 (2010).
 - ³ D. M. Cleland, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **134**, 024112 (2011).
 - ⁴ J. S. Spencer, N. S. Blunt, and W. M. C. Foulkes, *J. Chem. Phys.* **136**, 054110 (2012).
 - ⁵ A. J. W. Thom, *Phys. Rev. Lett.* **105**, 263004 (2010).
 - ⁶ C. J. C. Scott and A. J. W. Thom, *J. Chem. Phys.* **147**, 124105 (2017).
 - ⁷ V. A. Neufeld and A. J. W. Thom, *J. Chem. Phys.* **147**, 194105 (2017).
 - ⁸ N. S. Blunt, T. W. Rogers, J. S. Spencer, and W. M. C. Foulkes, *Phys. Rev. B* **89**, 245124 (2014).
 - ⁹ F. D. Malone, N. S. Blunt, E. W. Brown, D. K. K. Lee, J. S. Spencer, W. M. C. Foulkes, and J. J. Shepherd, *Phys. Rev. Lett.* **117**, 115701 (2016).
 - ¹⁰ S. Ten-no, *J. Chem. Phys.* **138**, 164126 (2013).
 - ¹¹ Y. Ohtsuka and S. Ten-no, *J. Chem. Phys.* **143**, 214107 (2015).
 - ¹² S. Ten-no, *J. Chem. Phys.* **147**, 244107 (2017).
 - ¹³ A. Nagy and V. Savona, arXiv:1802.05931 [quant-ph] (2018).
 - ¹⁴ S. Baroni and S. Moroni, *Physical Review Letters* **82**, 4745 (1999).
 - ¹⁵ L. K. Wagner and L. Mitas, *Journal of Chemical Physics* **126**, 034105 (2007).
 - ¹⁶ R. Gaudoin and J. M. Pitarke, *Phys. Rev. Lett.* **99**, 126406 (2007).
 - ¹⁷ E. Ospadov, D. G. Oblinsky, and S. M. Rothstein, *Phys. Chem. Chem. Phys.* **13**, 8031 (2011).
 - ¹⁸ E. Coccia, O. Chernomor, M. Barborini, S. Sorella, and L. Guidoni, *Journal of Chemical Theory and Computation* **8**, 1952 (2012).
 - ¹⁹ M. C. Per, I. K. Snook, and S. P. Russo, *Phys. Rev. B* **86**, 201107 (2012).
 - ²⁰ E. Ospadov and S. M. Rothstein, *J. Chem. Phys.* **142**, 024114 (2015).
 - ²¹ M. Motta and S. Zhang, *Journal of Chemical Theory and Computation* **13**, 5367 (2017).
 - ²² G. H. Booth, D. Cleland, A. Alavi, and D. P. Tew, *J. Chem. Phys.* **137**, 164112 (2012).
 - ²³ S. Zhang and M. H. Kalos, *J. Stat. Phys.* **70**, 515 (1993).
 - ²⁴ M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, *Phys. Rev. Lett.* **104**, 157201 (2010).
 - ²⁵ C. Overy, G. H. Booth, N. S. Blunt, J. J. Shepherd, D. Cleland, and A. Alavi, *J. Chem. Phys.* **141**, 244117 (2014).
 - ²⁶ C. J. Umrigar, J. Toulouse, C. Filippi, S. Sorella, and R. G. Hennig, *Phys. Rev. Lett.* **98**, 110201 (2007).
 - ²⁷ J. Toulouse and C. J. Umrigar, *J. Chem. Phys.* **126**, 084102 (2007).
 - ²⁸ L. Zhao and E. Neuscamman, *J. Chem. Theory Comput.* **12**, 3436 (2016).
 - ²⁹ J. A. R. Shea and E. Neuscamman, *J. Chem. Theory Comput.* **13**, 6078 (2017).
 - ³⁰ W. L. McMillan, *Phys. Rev.* **138**, A442 (1965).
 - ³¹ W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
 - ³² E. Neuscamman, C. J. Umrigar, and G. K.-L. Chan, *Phys. Rev. B* **85**, 045103 (2012).
 - ³³ E. Neuscamman, *J. Chem. Phys.* **145**, 081103 (2016).
 - ³⁴ R. E. Thomas, Q. Sun, A. Alavi, and G. H. Booth, *J. Chem. Theory Comput.* **11**, 5316 (2015).
 - ³⁵ G. Li Manni, S. D. Smart, and A. Alavi, *J. Chem. Theory Comput.* **12**, 1245 (2016).
 - ³⁶ H. J. Changlani, H. Zheng, and L. K. Wagner, *J. Chem. Phys.* **143**, 102814 (2015).
 - ³⁷ A. Altland and B. Simons, *Condensed Matter Field Theory* (Cambridge University Press, 2010).
 - ³⁸ M. Jarrell and J. E. Gubernatis, *Physics Reports* **269**, 133 (1996).
 - ³⁹ R. K. Bryan, *Eur. Biophys. J.* **18**, 165 (1990).
 - ⁴⁰ R. N. Silver, D. S. Sivia, and J. E. Gubernatis, *Phys. Rev. B* **41**, 2380 (1990).
 - ⁴¹ R. N. Silver, J. E. Gubernatis, D. S. Sivia, and M. Jarrell, *Phys. Rev. Lett.* **65**, 496 (1990).
 - ⁴² J. E. Gubernatis, M. Jarrell, R. N. Silver, and D. S. Sivia, *Phys. Rev. B* **44**, 6011 (1991).
 - ⁴³ M. Caffarel and D. M. Ceperley, *J. Chem. Phys.* **97**, 8415 (1992).
 - ⁴⁴ N. S. Blunt, A. Alavi, and G. H. Booth, *Phys. Rev. Lett.* **115**, 050603 (2015).
 - ⁴⁵ M. Caffarel, F. X. Gadea, and D. M. Ceperley, *Europhys. Lett.* **16**, 249 (1991).
 - ⁴⁶ D. M. Ceperley and B. Bernu, *J. Chem. Phys.* **89**, 6316 (1988).
 - ⁴⁷ W. Metzner and D. Vollhardt, *Phys. Rev. Lett.* **62**, 324 (1989).
 - ⁴⁸ A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, *Rev. Mod. Phys.* **68**, 13 (1996).
 - ⁴⁹ G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, *Rev. Mod. Phys.* **78**, 865 (2006).
 - ⁵⁰ N. S. Blunt, S. D. Smart, J. A. F. Kersten, J. S. Spencer, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **142**, 184107 (2015).

- ⁵¹ B. Bernu, D. M. Ceperley, and W. A. Lester Jr, J. Chem. Phys. **93**, 552 (1990).
- ⁵² Y. Kwon, D. M. Ceperley, and R. M. Martin, Phys. Rev. B **53**, 7376 (1996).
- ⁵³ W. R. Brown, W. A. Glauser, and W. A. Lester Jr, J. Chem. Phys. **103**, 9721 (1995).
- ⁵⁴ J. Shumway and D. M. Ceperley, Phys. Rev. B **63**, 165209 (2001).
- ⁵⁵ D. Meyer, S. Boblest, and G. Wunner, Phys. Rev. A **87**, 032515 (2013).
- ⁵⁶ J. K. L. MacDonald, Phys. Rev. **43**, 830 (1933).
- ⁵⁷ N. S. Blunt, S. D. Smart, G. H. Booth, and A. Alavi, J. Chem. Phys. **143**, 134117 (2015).
- ⁵⁸ K. A. Hallberg, Phys. Rev. B **52**, R9827 (1995).
- ⁵⁹ S. Ramasesha, S. K. Pati, H. R. Krishnamurthy, Z. Shuai, and J. L. Brédas, Phys. Rev. B **54**, 7598 (1996).
- ⁶⁰ T. D. Kühner and S. R. White, Phys. Rev. B **60**, 335 (1999).
- ⁶¹ A. Holzner, A. Weichselbaum, I. P. McCulloch, U. Schollwöck, and J. von Delft, Phys. Rev. B **83**, 195115 (2011).
- ⁶² M. Ganahl, P. Thunström, F. Verstraete, K. Held, and H. G. Evertz, Phys. Rev. B **90**, 045144 (2014).
- ⁶³ P. E. Dargel, A. Honecker, R. Peters, R. M. Noack, and T. Pruschke, Phys. Rev. B **83**, 161104(R) (2011).
- ⁶⁴ P. E. Dargel, A. Wöllert, A. Honecker, I. P. McCulloch, U. Schollwöck, and P. T., Phys. Rev. B **89**, 205119 (2012).