

# Beyond Walkers in Stochastic Quantum Chemistry: Reducing Error Using Fast Randomized Iteration

Samuel M. Greene,<sup>†</sup> Robert J. Webber,<sup>‡</sup> Jonathan Weare,<sup>\*,‡</sup> and Timothy C. Berkelbach<sup>\*,§,||</sup>


<sup>†</sup>Department of Chemistry and James Franck Institute, University of Chicago, Chicago, Illinois 60637, United States

<sup>‡</sup>Courant Institute of Mathematical Sciences, New York University, New York, New York 10012, United States

<sup>§</sup>Department of Chemistry, Columbia University, New York, New York 10027, United States

<sup>||</sup>Center for Computational Quantum Physics, Flatiron Institute, New York, New York 10010, United States

**ABSTRACT:** We introduce a family of methods for the full configuration interaction problem in quantum chemistry, based on the fast randomized iteration (FRI) framework [Lim, L.-H.; Weare, J. *SIAM Rev.* 2017, 59, 547; DOI: 10.1137/15M1040827]. These methods, which we term “FCI-FRI”, stochastically impose sparsity during iterations of the power method and can be viewed as a generalization of full configuration interaction quantum Monte Carlo (FCIQMC) without walkers. In addition to the multinomial scheme commonly used to sample excitations in FCIQMC, we present a systematic scheme where excitations are not sampled independently. Performing ground-state calculations on five small molecules at fixed cost, we find that the systematic FCI-FRI scheme is 11–45 times more statistically efficient than the multinomial FCI-FRI scheme, which is in turn 1.4–178 times more statistically efficient than the original FCIQMC algorithm.



## 1. INTRODUCTION

Deterministic approaches to treating strong correlation in interacting quantum systems are often rendered intractable by the exponential scaling of the size of the Hilbert space with the number of particles.<sup>1</sup> In contrast, quantum Monte Carlo (QMC) methods<sup>2–9</sup> can be computationally more efficient because they employ a sparse representation of the wave function in this space, obtained via stochastic sampling. Methods that utilize a continuous basis of configurations in real space have long existed, e.g., diffusion Monte Carlo.<sup>10–15</sup> The application of these methods to Fermionic systems requires nodal constraints due to the antisymmetry of the wave function. This has motivated the development of discrete-space methods, e.g., full configuration interaction QMC (FCIQMC) and auxiliary-field QMC,<sup>16–19</sup> in which the antisymmetry is provided by a Slater determinant basis, thereby obviating the need to impose nodal constraints on the wave function.<sup>8,16,20,21</sup> A disadvantage of discrete-basis methods is that the basis is not complete, but this can be addressed using standard extrapolation techniques.<sup>22,23</sup>

Recently, Lim and Weare<sup>24</sup> introduced the fast randomized iteration (FRI) framework, a class of methods that use techniques similar to those used in discrete-basis QMC methods to solve large, generic linear algebra problems. Sparsity is imposed stochastically in matrices and vectors, which reduces the computational cost and storage requirements of these methods and facilitates their application to problems significantly larger than those treatable by conven-

tional linear algebra approaches. Many existing QMC algorithms, including the FCIQMC method, can be understood as specific methods within the FRI framework. The central purpose of this work is to describe, in a more general context, the application of FRI methods to calculations on interacting Fermionic systems in a discrete basis. Importantly, we leverage this generality to develop alternative methods within this framework and investigate their statistical error and convergence properties through numerical tests on small molecular systems.

The FRI framework can be applied in a variety of ways to calculate ground- and excited-state observables of electronic systems. This study discusses only the application of FRI to calculate the ground-state energy of the full configuration interaction (FCI) Hamiltonian matrix in a Slater determinant basis. Such applications of the FRI framework will be referred to in this manuscript as FCI-FRI. In these methods, calculation of the ground-state energy is achieved via stochastic implementations of the power method, in which an initial trial vector is evolved toward the ground state eigenvector by repeatedly applying the Hamiltonian, scaled and shifted such that the ground state is dominant. The power method can be viewed as a discretization of the imaginary-time propagation used in many QMC methods. In order to reduce computational cost, the Hamiltonian matrix and solution vector are

Received: May 1, 2019

Published: August 7, 2019

compressed stochastically, meaning that randomly selected subsets of their elements are zeroed in each iteration. Calculating the energy after each iteration and averaging yields an estimate of the ground-state energy. This estimate can be systematically improved by executing more iterations and by retaining more nonzero elements in each compression. Unlike the original FCIQMC method, some FRI methods become identical to the deterministic power method as the number of randomly selected elements increases to the size of the basis.

The various approaches to matrix and vector compression within the FRI framework differ in terms of their computational cost and statistical efficiency. In this study, we combine these approaches in two new FCI-FRI methods and compare them to the original FCIQMC method.<sup>16</sup> In the first method, multinomial matrix compression, which is used in FCIQMC, is combined with systematic vector compression. Multinomial and systematic sampling are reviewed in section 2. In the original presentation of FRI,<sup>24</sup> systematic vector compression was shown to yield the least statistical error out of all other schemes considered. In contrast, vector compression is achieved by integerizing elements in FCIQMC. Comparing the original FCIQMC method to the “multinomial FCI-FRI” method, which uses the same matrix compression scheme, illustrates the gains in efficiency that an improved vector compression scheme can enable. In the second method, “systematic FCI-FRI”, we seek to further improve the efficiency by also compressing the matrix systematically instead of multinomially. We introduce a new hierarchical scheme to reduce the computational cost of performing this compression. In numerical tests on five small molecules, we find that systematic FCI-FRI yields consistently greater statistical efficiency (defined below) than multinomial FCI-FRI by at least an order of magnitude, and multinomial FCI-FRI is also more statistically efficient than FCIQMC in its original form.

An additional purpose of this work is to better understand how the features of each of these methods influence their errors and computational cost. To this end, we also compare two methods applied recently to FCI problems<sup>25</sup> in which the matrix is not compressed. Although expensive, such approaches are feasible because of the sparse structure of the Hamiltonian. In the first of these methods, the vector is compressed using the stochastic systematic scheme, whereas in the second, it is compressed using a deterministic thresholding scheme. Both methods have similar cost and are tractable for problems beyond the reach of deterministic FCI. However, the stochastic method achieves significantly less error, highlighting the advantages of stochastic methods over their deterministic counterparts.

A number of recent extensions to the original FCIQMC algorithm have been found to enable improvements in performance by orders of magnitude. For example, in semistochastic FCIQMC,<sup>26,27</sup> a fixed subspace within the Slater determinant basis is treated deterministically, greatly reducing the statistical error in that portion of the solution vector. A related extension involves preserving some elements exactly if their magnitude exceeds a user-specified threshold.<sup>28</sup> In the initiator approximation,<sup>29–31</sup> elements in the solution vector are zeroed in each iteration according to deterministic compression rules to better constrain the sign structure of the solution vector, which introduces a small bias. The FCI-FRI methods discussed here also include some deterministic features, although these differ in key aspects from those in the FCIQMC extensions. In FCI-FRI, the vector and matrix

elements to be preserved exactly are chosen dynamically in each iteration on the basis of their relative magnitudes. The criteria for selecting these elements do not rely on user-specified parameters and instead were chosen to minimize compression error given a finite number of samples. Unlike the initiator approximation, this approach does not introduce an additional bias. Another FCIQMC extension that can be applied to FCI-FRI involves calculating perturbative corrections to the energy.<sup>32</sup>

Due to the versatility of the FRI framework, many recent FCIQMC extensions can also be applied to FCI-FRI methods, which may yield further performance improvements. Here, we compare FCI-FRI methods only to the original FCIQMC method, without extensions, in order to (1) facilitate clarity in our presentation of the FCI-FRI methods, and (2) isolate the effects of different matrix and vector compression schemes in our results. Future work will be devoted to incorporating these complementary extensions into FCI-FRI methods.

The remainder of this article is organized as follows. In section 2, we summarize the FRI framework in the context of the power method for FCI calculations and describe the compression schemes considered in this study. Efficient compression of the Hamiltonian matrix is accomplished using a hierarchical scheme introduced in section 2.2.3 and discussed in more detail in Appendix A. In section 3, we discuss results obtained by applying these methods to five small molecular systems and compare their statistical efficiencies. In section 4, we summarize our key findings and comment further on the differences among the methods in relation to potential future research directions.

## 2. METHODS

**2.1. The Power Method for Full Configuration Interaction Calculations.** The FCI formalism casts the treatment of a system of interacting Fermions in terms of linear algebra.<sup>33</sup> In the FCI-FRI and FCIQMC methods discussed here, a randomization of the power method is used to calculate observables associated with the ground-state (lowest-energy) eigenvector of the FCI Hamiltonian matrix,  $\mathbf{H}$ . This matrix is expressed in a Slater determinant basis for  $N$  electrons in  $M$  orbitals. Its only nonzero off-diagonal elements are those corresponding to single and double excitations between pairs of Slater determinants. The matrix element corresponding to a single excitation from determinant  $|K\rangle$  to  $|L\rangle = \hat{c}_a^\dagger \hat{c}_i |K\rangle$  is

$$H_{LK} \equiv H_K(i \rightarrow a) = \langle L | \hat{H} | K \rangle = \gamma_{ia}^K \left( h_{ia} + \sum_{j \in \text{occ}} \langle ij || aj \rangle \right) \quad (1)$$

where  $h_{ia}$  represents a matrix element of the one-electron component of the Hamiltonian and  $\langle ij || aj \rangle$  is an antisymmetrized two-electron repulsion integral. These are both readily obtained from the output of a Hartree–Fock calculation. The parity of the excitation  $\gamma_{ia}^K$  is determined by the order of the orbitals comprising the Slater determinants in this basis.<sup>34</sup> The sum is over the orbitals occupied in  $|K\rangle$ . The notation  $H_K(i \rightarrow a)$  will be used throughout this paper to denote the index of an excitation from determinant  $|K\rangle$ . The matrix element for the double excitation to  $|M\rangle = \hat{c}_a^\dagger \hat{c}_b^\dagger \hat{c}_i \hat{c}_j |K\rangle$  is

$$H_{MK} \equiv H_K(ij \rightarrow ab) = \langle M | \hat{H} | K \rangle = \gamma_{ia}^K \gamma_{jb}^K \langle ab || ij \rangle \quad (2)$$

and the diagonal matrix element associated with  $|K\rangle$  is

$$H_{KK} = \langle K|\hat{H}|K \rangle = \sum_{j \in \text{occ}} h_{jj} + \frac{1}{2} \sum_{i,j \in \text{occ}} \langle ij||ij \rangle \quad (3)$$

The ground-state eigenvalue of this matrix is therefore the system's electronic energy.

Applying the generic power method to  $\mathbf{H}$  involves iteratively generating a sequence of vectors, here referred to as iterates. Each iterate  $\mathbf{v}^{(\tau)}$ , where  $\tau$  denotes the iteration index, is obtained by multiplying the previous iterate by the matrix  $\mathbf{P} = \mathbf{1} - \varepsilon \mathbf{H}$ , where  $\mathbf{1}$  is the identity and  $\varepsilon$  is a positive number that is sufficiently small to ensure that the ground state of  $\mathbf{H}$  is the dominant eigenvector of  $\mathbf{P}$ . The initial iterate,  $\mathbf{v}^{(0)}$ , must have nonzero overlap with the ground-state eigenvector,  $\mathbf{v}_{\text{GS}}$ . In FCI, the Hartree–Fock unit vector is usually a suitable choice and is used in all of the calculations presented here. The iterates converge to the ground-state eigenvector up to a normalization factor,

$$\lim_{\tau \rightarrow \infty} \frac{\mathbf{v}^{(\tau)}}{\|\mathbf{v}^{(\tau)}\|} = \mathbf{v}_{\text{GS}} \quad (4)$$

After a sufficient number of iterations, convergence to the ground state is geometric, with error decaying by a factor of  $(1 - \varepsilon E_0)/(1 - \varepsilon E_1)$  after each iteration. Here  $E_0$  is the ground-state eigenvalue of  $\mathbf{H}$ , and  $E_1$  is the first excited-state eigenvalue. Alternative choices of  $\mathbf{v}^{(0)}$  may be used to reduce the number of iterations required for convergence.<sup>35</sup> The norms of the iterates  $\|\mathbf{v}^{(\tau)}\|$  tend to either 0 or  $\infty$ , depending on the sign of  $E_0$ , as  $\tau \rightarrow \infty$ . An energy shift,  $S^{(\tau)}$ , is therefore included in the matrix  $\mathbf{P}^{(\tau)}$  at each iteration to stabilize the norm,

$$\mathbf{P}^{(\tau)} = \mathbf{1} - \varepsilon(\mathbf{H} - S^{(\tau)}\mathbf{1}) \quad (5)$$

where  $S^{(\tau)}$  is updated dynamically after every  $A$  iterations, where  $A$  is a user-specified parameter (10 in our calculations), according to the formula introduced in the FCIQMC method,<sup>16</sup>

$$S^{(\tau)} = S^{(\tau-A)} - \frac{\xi}{A\varepsilon} \ln \frac{\|\mathbf{v}^{(\tau)}\|_1}{\|\mathbf{v}^{(\tau-A)}\|_1} \quad (6)$$

Here  $\xi$  is a user-specified damping parameter (taken to be 0.05 in the calculations presented here), and  $\|\cdot\|_1$  denotes the one-norm, defined for an arbitrary vector  $\mathbf{x}$  as

$$\|\mathbf{x}\|_1 = \sum_i |x_i| \quad (7)$$

This procedure is used to stabilize the one-norm of the iterates in all methods considered in this study. In FCIQMC, the shift is updated only after the one-norm of the iterates (i.e., the number of walkers) has reached a specified target.<sup>16</sup> The iterates are generated by the relation

$$\mathbf{v}^{(\tau+1)} = \mathbf{P}^{(\tau)} \mathbf{v}^{(\tau)} \quad (8)$$

**2.2. FRI Compression Schemes.** The size of the FCI basis,  $N_{\text{FCI}} \sim O(M \text{ choose } N)$ , renders it impossible to apply the power method as described above to many systems of chemical interest. The memory cost is  $O(N_{\text{FCI}})$  and the computational cost of matrix–vector multiplication is  $O(N^2 V^2 N_{\text{FCI}})$ , where  $V = M - N$  is the number of virtual (unoccupied) orbitals. For large systems, these costs are prohibitive. The FCI-FRI methods circumvent these bottlenecks by stochastically compressing the vector  $\mathbf{v}^{(\tau)}$ , and

possibly the matrix  $\mathbf{P}^{(\tau)}$ , in each iteration. Stochastic compression is defined such that (1) the resulting compressed vector or matrix has at most a desired number  $m$  of nonzero elements and (2) the expectation value of each element in the compressed vector or matrix is equal to the corresponding element in the input vector or matrix, i.e.,

$$E[\Phi(\mathbf{x})]_i = x_i \quad (9)$$

where  $\Phi$  denotes the compression operation and  $\mathbf{x}$  is an arbitrary vector. The fact that many of the elements in the compressed matrix or vector are zero facilitates the use of sparse linear algebra schemes, which enables the efficiency of FRI methods.

As an example, in an FCI-FRI method that uses only vector compression, matrix–vector multiplication is performed as

$$\mathbf{v}^{(\tau+1)} = \Phi(\mathbf{P}^{(\tau)} \mathbf{v}^{(\tau)}) \quad (10)$$

This method has a memory cost of  $O(N^2 V^2 m)$  (to store the nonzero elements in the matrix–vector product before compression) and a computational cost of  $O(N^2 V^2 m \log m)$ . For many systems of chemical interest, these costs can be significantly less than those for deterministic FCI.

There are many possible compression methods in FRI with the above defining properties that differ in the degree of statistical error they introduce. In order to emphasize the generality of the FRI framework, we begin by introducing several such methods in more abstract linear algebra terms before discussing their specific application to the FCI problem.

**2.2.1. Vector Compression.** In this study, we compare several different approaches to vector compression. These have been applied in previous stochastic quantum chemistry calculations, although they can be applied more generally to any vector. The simplest approach to compressing an arbitrary vector  $\mathbf{x}$  involves randomly selecting a subset of its elements, each with probability

$$p_i = \frac{|x_i|}{\|\mathbf{x}\|_1} \quad (11)$$

The expected number of times each element is sampled is

$$E[n_i] = m p_i \quad (12)$$

where  $m$  is the total number of elements selected. Therefore, assigning each element of the compressed vector the value

$$\Phi(\mathbf{x})_i = \frac{n_i \|\mathbf{x}\|_1 \text{sgn}(x_i)}{m} \quad (13)$$

ensures that the condition in eq 9 is satisfied and that the vector has at most  $m$  elements (fewer if any  $n_i > 1$ ). Possible methods for randomly generating the values  $\{n_i\}$  will be discussed below.

It is often beneficial to preserve the largest-magnitude elements of  $\mathbf{x}$  exactly in order to reduce the overall statistical error incurred in compressing the vector. Lim and Weare<sup>24</sup> proposed the following criterion for determining the number  $\rho$  to preserve exactly. If  $\mathbf{s}$  is a vector, with length  $l$ , of indices that sorts the elements of  $\mathbf{x}$  in order of decreasing magnitude (i.e.,  $|x_{s_j}| \geq |x_{s_{j+1}}|$  for all  $j < l$ ), then  $\rho$  is the minimum value of  $h$  for which

$$(m - h)|x_{s_{h+1}}| \leq \sum_{j=h+1}^l |x_{s_j}| \quad (14)$$




where  $m$  denotes the desired number of nonzero elements in  $\Phi(\mathbf{x})$ , and  $c$  is the number of nonzero elements in  $\mathbf{x}$ . Thus,  $\rho$  depends both on  $m$  and  $\mathbf{x}$ . Calculating  $\rho$  requires identifying the largest-magnitude elements of  $\mathbf{x}$ . This can be done efficiently, in  $O(\rho \log c)$  time, by using a binary heap structure rather than sorting the entire vector. The elements of  $\mathbf{x}$  with indices  $\{s_1, s_2, \dots, s_\rho\}$  are unchanged in the compression. If  $m \geq c$ , this criterion naturally specifies that all elements are preserved exactly. Otherwise, the remaining elements of  $\Phi(\mathbf{x})$  are determined by applying random sampling with  $(m - \rho)$  samples to the vector  $\mathbf{x}'$ , which is obtained by zeroing the  $\rho$  largest-magnitude elements of  $\mathbf{x}$ . The resulting elements of the compressed vector are

$$\Phi(\mathbf{x})_{s_i} = \begin{cases} x_{s_i} & i \leq \rho \\ n_{s_i} \|\mathbf{x}'\|_1 \operatorname{sgn}(x_{s_i}) (m - \rho)^{-1} & i > \rho \end{cases} \quad (15)$$

An alternative, deterministic approach to vector compression is preserving the  $m$  largest-magnitude elements of  $\mathbf{x}$  exactly and zeroing the remaining elements. The additional sampling step introduced above has the notable advantage that the compressed vector is equal to the original in expectation. Even with a high degree of vector sparsity, results that are exact to within a controllable statistical error can be obtained by averaging over many independent vector compressions, provided there are no other sources of error.

**2.2.2. Sampling Schemes.** We compare two approaches to generating the integers  $\{n_i\}$  used for vector compression in eq 15. Both involve selecting  $m$  (or  $m - \rho$ ) elements from a probability distribution  $\mathbf{p}$  and are summarized in Figure 1. In



**Figure 1.** Illustration of the multinomial and systematic sampling schemes applied to the selection of  $m = 3$  elements from a probability distribution  $\mathbf{p}$ . The  $\times$ 's represent the random numbers  $U_k$  generated on the interval  $(0,1)$ . The indices selected in both schemes correspond to the intervals in  $\mathbf{p}$  with which the  $\times$ 's are aligned. The vector  $\mathbf{n}$  shown for each scheme represents the number of times each element is selected.

*multinomial* sampling, selections are made independently. The simplest implementation involves generating  $m$  random numbers  $\{U_k\}$  uniformly on the interval  $(0,1)$ . The index of the  $k$ th element selected is the value of  $j$  which satisfies

$$\sum_{i=1}^{j-1} p_i \leq U_k < \sum_{i=1}^j p_i \quad (16)$$

Any index can potentially be selected more than once, as the random numbers  $\{U_k\}$  are generated independently. The alias method is a more efficient implementation of multinomial sampling than the one described above.<sup>34,36</sup>

The *systematic* sampling scheme typically achieves reduced variance in the vector  $\mathbf{n}$ . The  $m$  random numbers  $\{U_k\}$  used in the selection of elements are generated from a single random number  $r$  chosen uniformly on the interval  $(0,1)$ , as follows:

$$U^{(k)} = \frac{k - 1 + r}{m} \quad (17)$$

with  $k = 1, 2, \dots, m$ . The value of  $r$  determines the position of the  $\times$ 's in each of the  $m$  subintervals of  $(0,1)$  in the systematic portion of Figure 1. The indices of elements selected are determined as described in multinomial sampling. Although systematic sampling is expected to yield less statistical error than multinomial in general, this difference is expected to become smaller as the number of elements selected ( $m$ ) decreases relative to the size of the vector. When  $m = 1$ , systematic sampling coincides exactly with multinomial sampling.

**2.2.3. Hierarchical Matrix Factorization.** The vector compression methods discussed above enable the application of FRI to iterative linear algebra methods based on matrix–vector multiplication at less cost than their deterministic counterparts. However, even the cost of multiplying a sparse vector by  $\mathbf{P}^{(r)}$  is prohibitive for large problems in quantum chemistry. This cost can be further reduced by compressing both the matrix and vector in each iteration. In principle, the vector compression methods described above could also be applied to compress the matrix before multiplication in each iteration, e.g., by treating each of its columns as a vector. This would require enumerating all of its nonzero elements, which offers few advantages over calculating the matrix–vector product without compression.

This section describes an alternative hierarchical approach to randomly approximating a matrix–vector product using compression. For a generic matrix–vector product  $\mathbf{A}\mathbf{x}$ , this involves factoring  $\mathbf{A}$  into a product of matrices and performing a sequence of vector compressions. For example, if  $\mathbf{A} = \mathbf{A}^{(3)}\mathbf{A}^{(2)}\mathbf{A}^{(1)}$ , then  $\mathbf{A}\mathbf{x}$  can be approximated as

$$\mathbf{A}\mathbf{x} = \mathbf{A}^{(3)}\Phi(\mathbf{A}^{(2)}\mathbf{x}^{(1)}) \quad (18)$$

where

$$\mathbf{x}^{(1)} = \Phi(\mathbf{A}^{(1)}\mathbf{x}) \quad (19)$$

The compressions after each multiplication are performed independently in this study, but other approaches in which they are not independent are possible as well. If  $\mathbf{A}^{(1)}$ ,  $\mathbf{A}^{(2)}$ , and  $\mathbf{A}^{(3)}$  are sparse, this approach can be made more efficient than calculating  $\mathbf{A}\mathbf{x}$  directly. The multinomial selection of excitations in FCIQMC<sup>7</sup> can be understood as a specific implementation of this approach, but we describe it in more general terms to demonstrate that it can be used with any compression scheme in FRI.

There are multiple ways to factor the Hamiltonian matrix and correspondingly the matrix  $\mathbf{P}^{(r)}$  for quantum chemistry calculations. These can be applied in contexts other than FCI, e.g., for stochastic coupled-cluster.<sup>37,38</sup> Here we consider two such factorizations, near-uniform<sup>7</sup> and heat-bath Power–Pitzer (HB-PP).<sup>34,39</sup> The structure of each matrix in these factorizations is dictated by the two-body structure of the Hamiltonian. Both have the form  $\mathbf{B}\mathbf{C}^{(r)}\mathbf{Q}$ , where  $\mathbf{Q}$  is factored further into a product of matrices. Elements of these matrices can be calculated efficiently using information about the symmetry of the system and, in the case of the HB-PP factorization, information from the Hamiltonian matrix. Elements of  $\mathbf{Q}$  have been introduced as the probabilities for sampling excitations in previous descriptions of FCIQMC, and multiplication by  $\mathbf{B}$  sums contributions from different excitations to the same determinant. Off-diagonal elements

of the matrix  $\mathbf{BQ}$  can be interpreted as an approximation to those of  $\mathbf{P}^{(\tau)}$  or  $\mathbf{H}$ . The extra factor of  $\mathbf{C}^{(\tau)}$  corrects for this discrepancy between  $\mathbf{BQ}$  and  $\mathbf{P}^{(\tau)}$  by multiplying by elements of  $\mathbf{P}^{(\tau)}$  and dividing by elements of  $\mathbf{Q}$ . This form ensures that matrix elements can be calculated efficiently and that multiplication by the matrix factors is equivalent to multiplication by  $\mathbf{P}^{(\tau)}$ . The detailed forms of these factorizations are given in [Appendix A](#).

**2.3. FCI-FRI Methods Considered in this Study.** The previous sections discussed compression techniques applicable to matrices and vectors in general. This section summarizes the particular implementations of these schemes in the three FCI-FRI methods considered in this study, as well as FCIQMC. A Python/Cython implementation of these methods with OpenMP parallelism is available on GitHub.<sup>40</sup>

In all three FCI-FRI methods, iterate vectors are compressed systematically following matrix multiplication, regardless of which matrix compression scheme is used. A subset of  $\rho$  vector elements is preserved exactly, with  $\rho$  calculated as described in the discussion surrounding [eq 14](#), and  $(m - \rho)$  additional nonzero vector elements are sampled randomly using the systematic scheme described in [section 2.2.2](#). In order to quantify the error introduced by compressing the matrix  $\mathbf{P}^{(\tau)}$  in each iteration, we considered three different matrix compression schemes in the three FCI-FRI methods. In the “full-matrix FCI-FRI” method, the matrix is not compressed. This method has been discussed previously and compared to FCIQMC.<sup>25</sup> As discussed above, its memory and CPU cost per iteration is approximately  $O(N^2 V^2 m \log m)$ . In the remaining two FCI-FRI methods,  $\mathbf{P}^{(\tau)}$  is compressed either multinomially or systematically using a hierarchical factorization scheme, with additional constraints as discussed in [Appendix B](#). Excluding the diagonal elements of  $\mathbf{P}^{(\tau)}$ , which are preserved exactly,  $N_{\text{mat}}$  samples are used in each compression. Matrix compression in “multinomial FCI-FRI” corresponds more closely to the scheme used in the original FCIQMC method, whereas “systematic FCI-FRI” is designed to reduce statistical error. These algorithms are summarized in [Table 1](#).

**2.4. Comparison with FCIQMC.** As discussed above, the FCIQMC method described in [ref 16](#) can be viewed as a specific method within the FRI framework. Although our presentation of the method differs somewhat from previous studies, we implemented FCIQMC in its original form, i.e., without any of its existing extensions (e.g., initiator or semistochastic), for comparison to FCI-FRI. This section summarizes the compression techniques in FCIQMC using the unifying language of the FRI framework, in order to facilitate comparison to the new FCI-FRI methods in this study. Further details about compression in FCIQMC can be found in [Appendix B](#).

In the original FCIQMC algorithm, each iterate  $\mathbf{v}$  is represented by a number of signed walkers, so each of its elements  $v_K$  is an integer. The total number of walkers is  $\|\mathbf{v}\|_1$ . The random selection of excitations in FCIQMC corresponds to multinomial compression of  $\mathbf{P}^{(\tau)}$  using one of the factorizations discussed in [Appendix A](#). The “spawning” step corresponds to integerization of off-diagonal elements after multiplication by  $\mathbf{C}^{(\tau)}$  in the hierarchical scheme, and the “death/cloning” step corresponds to integerization of diagonal elements. “Annihilation”, i.e., the summation of matrix elements corresponding to the same Slater determinant basis element, is performed by multiplying by  $\mathbf{B}$  in the hierarchical scheme.

**Table 1. Overview of the Steps in Each Iteration of the FCI-FRI Methods Considered in This Study and the Approximate Scaling of the CPU Cost of Each Step**

Full-Matrix FCI-FRI	CPU cost/iteration <sup>a</sup>
1. Calculate $\mathbf{v}^{(\tau+1)} = \mathbf{P}^{(\tau)} \mathbf{v}^{(\tau)}$	$O(N^2 V^2 m \log m)$ <sup>b</sup>
2. Compress $\mathbf{v}^{(\tau+1)}$ systematically to $m$ nonzero elements	$O(N^2 V^2 m)$
3. Adjust the energy shift, $S^{(\tau)}$ ( <a href="#">eq 6</a> )	$O(1)$

Multinomial and Systematic FCI-FRI	CPU cost/iteration
1. Calculate $\mathbf{v}^{(\tau+1)} = \mathbf{P}^{(\tau)} \mathbf{v}^{(\tau)}$ using hierarchical factorization with multinomial or systematic compression to $N_{\text{mat}}$ nonzero elements	$O(N_{\text{mat}})$ or $O(MN_{\text{mat}}) + O(N_{\text{mat}} \log m)$ <sup>c</sup>
2. Compress $\mathbf{v}^{(\tau+1)}$ systematically to $m$ nonzero elements	$O((N_{\text{mat}} + m) + m \log(N_{\text{mat}} + m))$ <sup>d</sup>
3. Adjust the energy shift, $S^{(\tau)}$ ( <a href="#">eq 6</a> )	$O(1)$

<sup>a</sup>The variable  $N$  is the number of electrons in the system;  $M$  is the number of spatial orbitals in the single-particle basis;  $V = M - N$  is the number of virtual orbitals;  $m$  is the number of nonzero elements kept in the solution vector;  $N_{\text{mat}}$  is the number of off-diagonal elements sampled from the Hamiltonian matrix. <sup>b</sup>The  $\log m$  factor here arises because our implementation uses a less efficient binary search algorithm to perform matrix–vector multiplication. This cost could be reduced by using a hashing algorithm.<sup>7</sup> <sup>c</sup> $O(N_{\text{mat}})$  is the approximate cost of compressing the near-uniform distribution, and  $O(MN_{\text{mat}})$  is the cost for HB-PP. The  $O(N_{\text{mat}} \log m)$  term comes from multiplication by  $\mathbf{B}$  in both factorizations and can be reduced to  $O(N_{\text{mat}})$  using hashing. <sup>d</sup>Worst-case scaling. More typical scaling, corresponding to preserving relatively few elements exactly, is  $O(N_{\text{mat}} + m)$ .

The key difference between the original FCIQMC algorithm and multinomial FCI-FRI methods lies in the compressions performed after the final two matrix multiplications in the hierarchical scheme. In FCIQMC, after multiplication by  $\mathbf{C}^{(\tau)}$ , elements are rounded to integers using a random binomial integerization procedure. Like other vector compression techniques, this ensures sparsity in the resulting vector since many elements are rounded to zero. This reduces the cost of multiplication by  $\mathbf{B}$  (i.e., “annihilation”), since this involves summing fewer nonzero elements, but it also introduces additional statistical error. The vector obtained after multiplication by  $\mathbf{B}$  is not compressed and is instead treated as the next iterate. In multinomial FCI-FRI, the vector obtained after multiplication by  $\mathbf{C}^{(\tau)}$  is not compressed, so the elements that are summed during multiplication by  $\mathbf{B}$  are real-valued (i.e., not necessarily integers). Sparsity is instead enforced by compressing the iterate systematically after the final matrix multiplication. It should be noted that compression is performed after multiplication by  $\mathbf{B}$  in the semistochastic FCIQMC extension, as in FCI-FRI, although this extension was not considered in this study.

One advantage of FCIQMC is its straightforward parallelizability. Since elements are selected independently in the multinomial matrix compression scheme, they can be selected in parallel. Similarly, the stochastic rounding of matrix elements to integers can be performed in parallel, as each element is treated independently. In contrast, elements are not selected independently in systematic compression, so these strategies cannot be applied in exactly the same way. Nevertheless, parallelizing systematic schemes is possible,

e.g., by performing parallel compressions in subspaces of the Slater determinant space. Investigation of these strategies will be the subject of future research. The original FCIQMC method and FCI-FRI methods become more similar as the number of nonzero elements in the compressions (number of walkers) decreases relative to the size of the basis ( $N_{\text{FCI}}$ ): the probability of choosing repeated elements in multinomial matrix compression decreases, and the frequency of annihilation events in FCIQMC decreases. However, our examples suggest that the number of walkers required to obtain reasonable results from the original FCIQMC method is already sufficient to observe a substantial benefit from FRI.

**2.5. Statistical Error Analysis.** Although in principle the iterates can be averaged to obtain an estimate of the ground-state eigenvector, the memory requirements of such an approach are prohibitive for large systems. In practice, we are only interested in observables calculated from the ground-state eigenvector, so their average values are accumulated rather than the eigenvector itself. This section addresses the calculation of the average ground-state energy and the methods used to quantify the statistical error in this average.

Conventionally, the energy of a state vector  $\mathbf{x}$  is calculated as a Rayleigh quotient, defined here as

$$E_{\text{R}}(\mathbf{x}) = \frac{\mathbf{x}^* \mathbf{H} \mathbf{x}}{\mathbf{x}^* \mathbf{x}} \quad (20)$$

where  $\mathbf{x}^*$  denotes the conjugate transpose of  $\mathbf{x}$ . Averages of the energy obtained from the Rayleigh quotient estimator applied to an ensemble of random vectors will exhibit a statistical bias due to the products of correlated random vectors in both the numerator and denominator.<sup>28</sup> Consequently, a projected energy estimator is instead used to calculate averages:

$$E_{\text{p}}(\mathbf{x}) = \frac{\mathbf{v}_{\text{ref}}^* \mathbf{H} \mathbf{x}}{\mathbf{v}_{\text{ref}}^* \mathbf{x}} \quad (21)$$

where  $\mathbf{v}_{\text{ref}}$  is a constant, appropriately chosen reference vector. In principle, using a reference vector that is closer to the exact ground-state eigenvector of the Hamiltonian will yield a better estimate of the correlation energy.<sup>18</sup> In this study we use the Hartree–Fock unit vector for simplicity. If this estimator is to be applied to multiple vectors  $\mathbf{x}$  (in this case, the iterates obtained after each iteration), the numerator can be calculated efficiently by storing the matrix–vector product  $\mathbf{H} \mathbf{v}_{\text{ref}}$  and taking its inner product with each vector  $\mathbf{x}$ . In the FCI-FRI methods in this study, this inner product is calculated before each iterate is compressed.

The numerator and denominator of eq 21 at a particular iteration are denoted as

$$n^{(\tau)} = \mathbf{v}_{\text{ref}}^* \mathbf{H} \mathbf{v}^{(\tau)} \quad (22)$$

and

$$d^{(\tau)} = \mathbf{v}_{\text{ref}}^* \mathbf{v}^{(\tau)} \quad (23)$$

Because  $n^{(\tau)}$  and  $d^{(\tau)}$  are correlated within each iteration due to their mutual dependence on  $\mathbf{v}^{(\tau)}$ , averaging the quotients  $n^{(\tau)}/d^{(\tau)}$  over all iterations would introduce a statistical bias. Therefore, the mean energy is calculated instead as  $\langle E_{\text{p}} \rangle = \langle n \rangle / \langle d \rangle$ , where

$$\langle n \rangle = \frac{1}{N_i - \tau_c} \sum_{\tau \geq \tau_c} n^{(\tau)} \quad (24)$$

and the corresponding expression for the denominator is defined analogously. Here the total number of iterations in the trajectory is denoted  $N_i$ , and the equilibration time,  $\tau_c$ , is the number of iterations at the beginning of the trajectory not included in the average. Our approach to determining  $\tau_c$  will be described below. If the expected value of the iterates  $\mathbf{v}^{(\tau)}$  converges to the exact ground-state eigenvector (to within a normalization factor) after infinitely many iterations, the mean energy will also converge to its exact value, since the numerator and denominator are averaged separately. In practice, a systematic bias is still observed after infinitely many iterations in FCI-FRI and FCIQMC because the expected value of the iterates does not converge to the exact ground-state eigenvector. This has been discussed previously in the context of FCIQMC and diffusion Monte Carlo methods as the population control bias.<sup>10,43</sup>

The delta method is used to calculate the variance of the average  $\langle E_{\text{p}} \rangle$  as follows:

$$\begin{aligned} \text{Var}[\langle E_{\text{p}} \rangle] &= \text{Var} \left[ \frac{\langle n \rangle}{\langle d \rangle} - \frac{n_0}{d_0} \right] \\ &\approx \text{Var} \left[ \frac{\langle n \rangle - n_0}{d_0} - \frac{n_0(\langle d \rangle - d_0)}{d_0^2} \right] \\ &= \text{Var} \left[ \frac{\langle n \rangle}{d_0} - \frac{n_0 \langle d \rangle}{d_0^2} \right] \end{aligned} \quad (25)$$

where  $n_0$  and  $d_0$  represent the deterministic quantities  $\mathbf{v}_{\text{ref}}^* \mathbf{H} \mathbf{v}_{\text{GS}}$  and  $\mathbf{v}_{\text{ref}}^* \mathbf{v}_{\text{GS}}$ , up to an irrelevant normalization factor. We define  $E_{\text{delta}}^{(\tau)}$  as

$$\begin{aligned} E_{\text{delta}}^{(\tau)} &= \frac{n^{(\tau)}}{d_0} - \frac{n_0 d^{(\tau)}}{d_0^2} \\ &\approx \frac{n^{(\tau)}}{\langle d \rangle} - \frac{\langle n \rangle d^{(\tau)}}{\langle d \rangle^2} \end{aligned} \quad (26)$$

Because subsequent iterates in a trajectory are correlated, the variance in eq 25 cannot be calculated naively as  $\sigma^2/(N_i - \tau_c)$ , where  $\sigma^2$  is the mean-squared deviation from the average, i.e.,

$$\sigma^2 = \frac{1}{N_i - \tau_c} \sum_{\tau \geq \tau_c} (E_{\text{delta}}^{(\tau)})^2 \quad (27)$$

Instead,  $\sigma^2$  must be multiplied by the integrated autocorrelation time (IAT), a measure of the degree of correlation. The IAT is estimated using the iterative procedure described in ref 44, as implemented in the emcee software package,<sup>45</sup> using the sequence of values  $\{E_{\text{delta}}^{(\tau)}\}$  as the input. If the sequence  $\{n^{(\tau)}/d^{(\tau)}\}$  was used instead, the resulting variance would not correspond to an energy estimate in which the numerator and denominator are averaged separately.

The equilibration time  $\tau_c$  is determined for each trajectory by inspecting plots of the IATs of the numerator and denominator of the energy estimator separately vs  $\tau_c$ . Typically, the IAT is greater for smaller values of  $\tau_c$ , both because of their dependence on the initial iterate  $\mathbf{v}^{(0)}$  and because iterates can become trapped around metastable energy values before converging to the ground-state eigenvector.<sup>46</sup> Equilibration times were therefore chosen to exclude this initial period of decreasing IATs. In FCIQMC,  $\tau_c$  is also constrained to be



Table 2. Parameters Used in Calculations on Each of the Systems in This Study<sup>a</sup>

system	geometry	occupation threshold/ $10^{-4}$	( $N$ , $M$ )	$N_{\text{FCI}}/10^6$	$\epsilon/(10^{-4}E_h)$	$E_{\text{FCI}}/E_h$
Ne (aug-cc-pVDZ)	—	—	(8, 22)	6.69	10	$-128.709476^b$
HF (cc-pCVDZ)	0.91622 Å	—	(10, 23)	283	1	$-100.270929^c$
H <sub>2</sub> O (cc-pVDZ)	$r_{\text{O-H}} = 0.975512$ Å $\angle_{\text{HOH}} = 110.565^\circ$	6	(10, 18)	18.3	10	$-76.167449^c$
N <sub>2</sub> (cc-pVDZ)	1.0944 Å	30	(10, 17)	4.8	5	$-109.228042^c$
C <sub>2</sub> (cc-pVDZ)	1.27273 Å	5	(8, 22)	6.7	5	$-75.7260112^c$

<sup>a</sup>Unless otherwise specified, the geometry is the diatomic bond length. MP2 natural orbitals with occupancies below the occupancy threshold, if specified, were excluded from the single-particle basis. The resulting number of (spatial) orbitals is reported as  $M$ . The number of unfrozen electrons considered for each system is  $N$ , and  $N_{\text{FCI}}$  is the size of the FCI basis. The parameter  $\epsilon$  (eq 5) is chosen to ensure convergence of the power method.  $E_{\text{FCI}}$  denotes the exact FCI energy (including nuclear repulsion) used for comparison to our stochastic results. <sup>b</sup>From ref 41. <sup>c</sup>Calculated using the PySCF software package.<sup>42</sup>

greater than the first index at which the energy shift is updated (eq 6).

The Flyvbjerg–Petersen blocking method<sup>47</sup> has been used in previous FCIQMC studies<sup>16,20,35,48</sup> to calculate the variance. The approach described here has the notable advantage that no data from after the initial equilibration period ( $\tau \geq \tau_c$ ) is discarded in the calculation of the mean and variance. Either of these methods requires a very long trajectory to achieve an accurate estimate of the variance, and it is likely that some of the statistical error estimates reported in this study are not fully converged.

The standard error of the energy estimator is calculated as

$$\sigma_e = (\text{Var}[\langle E_p \rangle])^{1/2} \quad (28)$$

This error is expected to scale as  $(N_i - \tau_c)^{-1/2}$  after sufficiently many iterations, according to the Markov chain central limit theorem with standard assumptions of ergodicity.<sup>44,49</sup> This scaling renders it impossible to directly compare the standard errors from two trajectories with different numbers of iterations. Therefore, the primary metric that will be used to compare the methods discussed here is the statistical efficiency, defined as<sup>34</sup>


$$E = \frac{1}{\sigma_e^2(N_i - \tau_c)} \quad (29)$$

For two methods executed for the same number of iterations after the equilibration period, the method with the greater statistical efficiency will typically yield less variance. From an alternative perspective, in order to achieve a target standard error, the method with greater statistical efficiency can be executed for fewer iterations. For example, to achieve a standard error of  $10^{-5}E_h$ , a method with statistical efficiency  $E$  requires  $[(10^{-5}E_h)^2 E]^{-1}$  iterations after the equilibration period. In this study, we do not normalize the efficiency based on the computational cost of each iteration. Therefore, for a given FCI-FRI method applied to a particular system, increasing the number of matrix or vector samples increases the statistical efficiency due to the expected decrease in error, regardless of the corresponding increase in computational cost. For this reason, when comparing the statistical efficiencies of different FCI-FRI methods and FCIQMC, we ensure that the same number of matrix and vector samples are used in all methods for each system. This ensures that any differences in the resulting statistical efficiencies are due to features inherent to the methods.

### 3. RESULTS

The methods described in the previous section are applied to a subset of the molecular systems considered in ref 16. The parameters relevant to the Hartree–Fock and randomized FCI calculations performed for these five systems are presented in Table 2. In order to run calculations for sufficiently many iterations to obtain robust estimates of the mean energy and associated standard error, fewer single-particle orbitals are used for three systems than in ref 16, thus reducing the size of the FCI basis ( $N_{\text{FCI}}$ ). This truncation is performed by discarding natural orbitals obtained from a second-order Møller–Plesset perturbation theory (MP2) calculation with occupation numbers less than a specified threshold. We emphasize that truncating the basis is necessary only because of inefficiencies in our implementations of these methods. Optimizing our implementations should enable the treatment of significantly larger systems. Core electrons are frozen in Ne, C<sub>2</sub>, and N<sub>2</sub>, as in ref 16. The same value of  $\epsilon$  is used to construct the matrix  $\mathbf{P}^{(\tau)}$  (eq 5) used in all methods for each system. The PySCF electronic structure software package<sup>42</sup> is used to perform Hartree–Fock, MP2, and deterministic FCI calculations. In ref 16, the average FCIQMC energy for the hydrogen fluoride (HF) molecule was compared to coupled-cluster theory with perturbative triple excitations, CCSD(T). Our deterministic FCI result, calculated using PySCF, differs from the CCSD(T) result by  $4.89 \times 10^{-4}E_h$ , and from the FCIQMC result from ref 16 by  $5.4 \times 10^{-5}E_h$ , a value greater than the reported uncertainty.

**3.1. FCI-FRI without Matrix Compression.** In order to isolate the contribution of vector compression to the statistical error in calculations of the ground state energy, we first consider results obtained by applying the “full-matrix FCI-FRI” method, which does not use matrix compression, to the Ne atom. We compare calculations with differing numbers of nonzero elements retained in the compression of each iterate ( $m$ ). As  $m$  approaches the size of the FCI basis, this method becomes identical to the deterministic power method. The difference between the estimated ground-state energy at each iteration and the exact energy is plotted for calculations with three different values of  $m$  in the top panel of Figure 2. The energy of the first iterate in each trajectory is the Hartree–Fock energy, since the first iterate was initialized to the Hartree–Fock unit vector. The energy decreases toward the exact energy in subsequent iterations. After the estimator is determined to be sufficiently close to the exact energy, at iteration  $\tau_c$ , the mean is accumulated according to eq 24. This cumulative mean is plotted in Figure 2 for  $\tau \geq \tau_c$ .



**Figure 2.** Results obtained by applying the “full-matrix FCI-FRI” method to the Ne atom. (top) Differences between the energy estimator ( $E_p^{(\tau)}$ , eq 21) and the exact FCI ground-state energy for three trajectories with different numbers,  $m$ , of nonzero elements in the compressed vectors. After the initial equilibration period ( $\tau > \tau_c$ ), the cumulative mean  $\langle E_p \rangle$  is plotted, with the shaded region indicating the corresponding 95% confidence interval ( $\pm 2\sigma_E$ ). (bottom) The statistical efficiency for trajectories executed with different values of  $m$ . The dashed line with slope 1 represents the expected scaling of the efficiency with respect to  $m$  (for  $m$  large but less than  $N_{\text{FCI}}$ ).

The value of the equilibration time  $\tau_c$  used in these trajectories increases with increasing  $m$  (Table 3), primarily

**Table 3. Results Obtained by Applying the “Full-Matrix FCI-FRI” Method to the Ne Atom with Different Values of  $m$ <sup>a</sup>**

$m/10^3$	$N_{\text{mat}}/10^6$	$(E_{\text{diff}} \pm 2\sigma_E)/ (10^{-5}E_h)$	$\text{eff}/(10^6E_h^{-2})$	$\tau_c/10^3$	$N_i/10^3$
1	0.93	$6437 \pm 16099$	$1.25 \times 10^{-10}$	0.8	1237
2	1.9	$141 \pm 242$	$6.4 \times 10^{-7}$	1.1	1062
5	4.7	$-0.089 \pm 4.60$	0.0015	1.2	1200
10	9.3	$0.307 \pm 0.480$	0.296	3.2	589
25	23.4	$-0.053 \pm 0.112$	12.8	4.8	256
50	46.8	$0.034 \pm 0.063$	86.9	6.1	123

<sup>a</sup>The difference  $E_{\text{diff}}$  between the mean and exact (FCI) energy for each calculation is presented, with twice the standard error  $\sigma_E$  (95% confidence interval). The length of the equilibration period ( $\tau_c$ ) and total number of iterations ( $N_i$ ) are given. The statistical efficiency is calculated using eq 29. The mean number of Hamiltonian matrix evaluations in each iteration  $N_{\text{mat}}$  is presented for comparison to other methods.

due the greater degree of noise in trajectories with fewer nonzero elements in each iterate. When  $m$  is smaller, the energy decreases more quickly toward the ground state, causing a lesser value of  $\tau_c$ , but fluctuates to a greater extent after  $\tau = \tau_c$ . In the deterministic power method, the asymptotic convergence rate is determined by the ratio  $(1 - \epsilon E_0)/(1 - \epsilon E_1)$ . Randomized implementations of the power method can exhibit different convergence properties, depending on the statistical error introduced in each iteration. This trend in  $\tau_c$  is therefore not surprising, and it suggests that an accurate energy


estimate can be achieved at less computational cost if the values of  $m$  and  $\epsilon$  are varied dynamically during the calculation.

The difference  $E_{\text{diff}}$  between the final estimate of the energy, obtained by averaging over all  $\tau \geq \tau_c$ , and the exact FCI energy from ref 41, is presented for each  $m$  in Table 3. The number of iterations included in each of these averages can be obtained by subtracting  $\tau_c$  from the reported total number of iterations,  $N_i$ . The reported uncertainties, twice the standard error  $\sigma_E$  calculated as described in section 2.5, represent 95% confidence intervals for the means. The exact energy is within these confidence intervals for all values of  $m$  reported here (i.e.,  $|E_{\text{diff}}| < 2\sigma_E$ ). The standard error is expected to decrease after more iterations, with an asymptotic scaling of  $(N_i - \tau_c)^{-1/2}$ . Confidence intervals for intermediate values of  $\tau$ , calculated by scaling the final confidence intervals reported in Table 3, are shown as shaded areas in Figure 2. The value  $E_{\text{diff}}$  is not expected to converge to 0 but rather to the statistical bias, as discussed in section 2.5. This bias scales as  $m^{-1}$  when  $m$  is sufficiently large (but still much smaller than the size of the FCI basis,  $N_{\text{FCI}}$ ),<sup>24</sup> but the number of iterations performed in our calculations is not sufficient to measure the biases in these calculations accurately.

In Table 3, decreased standard error is observed in calculations with greater values of  $m$ , despite the fact that fewer iterations were included in these calculations. If the errors from these calculations are compared after the same number of iterations, the trend with increasing  $m$  would be more pronounced. The statistical efficiency does not depend on the number of iterations and therefore allows for a more direct comparison. Statistical efficiencies calculated from all trajectories are presented in Table 3 and in the bottom panel of Figure 2. While the computational cost of full-matrix FCI-FRI calculations is approximately proportional to  $m$ , the statistical efficiency appears to increase at a faster-than- $m$  rate for small  $m$ . This indicates that, in terms of reducing the standard error, it is more advantageous to increase  $m$  in this pre-asymptotic regime than to increase the number of iterations. The statistical efficiency is expected to increase linearly with  $m$  for  $m$  sufficiently large (but still much smaller than  $N_{\text{FCI}}$ ).<sup>24</sup> Similar faster-than- $m$  pre-asymptotic scaling has been observed in other methods that use sequential Monte Carlo sampling on a classical problem,<sup>50</sup> suggesting that it is not (solely) a manifestation of the Fermion sign problem in this case.

Before considering the effect of matrix compression on the statistical error, we comment briefly on the benefits of using stochastic, rather than deterministic, vector compression. Results for the Ne atom obtained using a deterministic vector compression scheme are presented in Figure 3. In each iteration, the matrix is not compressed, the  $m$  greatest-magnitude elements in the vector are preserved exactly, and the remaining vector elements are zeroed. For all values of  $m$  considered, the energy calculated from the projected estimator,  $E_p$ , converges after approximately 3000 iterations. Energies obtained from the “full-matrix FCI-FRI” method, with  $m = 50\,000$  nonzero elements kept after each iteration, are also presented for comparison. The error in the corresponding deterministic calculation after a similar number of iterations is almost 2 orders of magnitude greater than the 95% confidence interval in the FCI-FRI calculation. Similar results for other electronic systems were observed previously in ref 25. These results indicate that the success of the FCI-FRI method in these cases cannot be attributed to its discarding vector elements that do not contribute significantly to the energy, as is





**Figure 3.** Results obtained by applying the power method with deterministic vector truncation to the Ne atom. Only the  $m$  greatest-magnitude elements of the vector were preserved exactly after each iteration. Differences between the energy estimator  $E_p^{(\tau)}$  and the exact energy at each iteration are plotted for four trajectories with different values of  $m$ . Results from the “full-matrix FCI-FRI” calculation with  $m = 50\,000$  elements from Figure 2 are presented for comparison. Note the log scale on the vertical axis.

done in the deterministic approach. The stochastic representation of these small-magnitude elements is crucial to its success. This observation may be relevant to selected CI methods,<sup>51–55</sup> which utilize a similar greedy optimization scheme.

**3.2. Methods with Matrix Compression.** The cost of the full-matrix FCI-FRI method renders it intractable for larger systems, so we also evaluate the performance of methods that use matrix compression, including the original FCIQMC method.

**3.2.1. Near-Uniform Factorization.** Methods that utilize the near-uniform factorization described in Appendix A will be discussed first. In order to ensure a fair comparison among these methods, all calculations for each system are executed with approximately the same cost, i.e., using the same numbers

of nonzero elements in the matrix and vector compressions in each iteration ( $N_{\text{mat}}$  and  $m$ , respectively). In an FCIQMC calculation,  $N_{\text{mat}}$  is the number of walkers, and  $m$  is determined by their distribution among the Slater determinant basis elements. In FCIQMC, the number of walkers and  $m$  fluctuate randomly in each iteration. Previous studies have determined that the number of walkers must be greater than a system-dependent critical value in order to ensure convergence. The number of walkers used in the FCIQMC calculations discussed here are constrained to be greater than these critical values. Critical values for the Ne and HF systems are given in ref 16, and those for the remaining systems considered in this study are determined using the same scheme, i.e., by observing trends in the growth of the number of walkers before the energy shift  $S^{(\tau)}$  is updated. The values of  $N_{\text{mat}}$  and  $m$  used in FCI-FRI calculations are fixed at the corresponding average values obtained from the FCIQMC calculations after walker growth has stabilized.

Results from these calculations for all molecular systems are presented in Table 4. In all calculations, average energies converge to the exact FCI energies reported in Table 2 to within twice the standard error (95% confidence interval). Strictly speaking, all methods considered here exhibit a statistical bias, although for these calculations it is very likely less than the reported confidence intervals. After more iterations, we expect that the standard error for all trajectories will decrease, and the energy differences  $E_{\text{diff}}$  for both trajectories of a particular method and system will converge to the same statistically significant bias. It is impossible to draw definitive conclusions about the relative biases of the three methods described here without more iterations.


Standard errors from FCIQMC calculations range from  $3 \times 10^{-5} E_h$  to  $20 \times 10^{-5} E_h$ , while those from the FCI-FRI methods are smaller (from  $2 \times 10^{-5} E_h$  to  $6 \times 10^{-5} E_h$  for multinomial FCI-FRI, and from  $0.4 \times 10^{-5} E_h$  to  $1.7 \times 10^{-5} E_h$  for systematic FCI-FRI), despite their use of fewer iterations. This trend is also

**Table 4.** Differences between Mean Energy Estimates and Those Reported in Table 2 ( $E_{\text{diff}}$ ) for Each of the Systems Considered Here, Calculated Using the FCIQMC, Multinomial FCI-FRI, and Systematic FCI-FRI Methods with the Near-Uniform Factorization Scheme<sup>a</sup>

system	$m/10^3$	$N_{\text{mat}}/10^6$	FCIQMC			multinomial FCI-FRI			systematic FCI-FRI		
			$(E_{\text{diff}} \pm 2\sigma_E)/(10^{-5}E_h)$	$\tau_c/10^3$	$N_i/10^3$	$(E_{\text{diff}} \pm 2\sigma_E)/(10^{-5}E_h)$	$\tau_c/10^3$	$N_i/10^3$	$(E_{\text{diff}} \pm 2\sigma_E)/(10^{-5}E_h)$	$\tau_c/10^3$	$N_i/10^3$
Ne	242	0.26	$-1.44 \pm 7.36$	22.5	2800	$0.06 \pm 5.66$	15.0	2373	$-0.16 \pm 1.09$	11.5	1422
			$2.89 \pm 7.47$	22.5	2800	$-3.12 \pm 4.99$	15.0	3200	$-0.74 \pm 1.11$	11.0	1445
HF	926	1.00	$10.57 \pm 26.86$	160.0	1469	$-9.76 \pm 11.17$	400.0	1104	$0.49 \pm 2.57$	620.0	1495
			$21.09 \pm 33.50$	430.0	1474	$-7.03 \pm 11.28$	380.0	1100	$-0.37 \pm 3.37$	620.0	994
H <sub>2</sub> O	491	0.57	$-0.96 \pm 6.52$	30.0	2400	$0.61 \pm 5.54$	20.0	1232	$-0.41 \pm 1.29$	25.0	1055
			$0.54 \pm 6.47$	30.0	2400	$-2.08 \pm 5.63$	20.0	1228	$0.17 \pm 1.16$	25.0	1059
N <sub>2</sub>	1014	1.21	$-7.46 \pm 29.75$	200.0	1788	$-1.05 \pm 5.02$	80.0	822	$0.14 \pm 0.82$	76.7	554
			$4.78 \pm 39.85$	200.0	1791	$2.41 \pm 5.55$	52.1	512	$-0.89 \pm 1.33$	170.0	557
C <sub>2</sub>	2622	4.14	$9.53 \pm 9.56$	50.0	2908	$1.32 \pm 3.55$	540.0	2051	$0.71 \pm 1.08$	42.2	513
			$4.76 \pm 11.54$	50.0	2768	$-2.30 \pm 3.92$	450.0	1327	$-0.50 \pm 0.77$	50.6	516

<sup>a</sup>The parameter  $m$  represents the sparsity of the iterates (mean sparsity for FCIQMC), and  $N_{\text{mat}}$  represents the number of Hamiltonian matrix elements evaluated in each iteration (mean number of walkers for FCIQMC). Results from two independent trajectories are presented for each method. Mean energy differences  $\pm$  twice the standard error (95% confidence interval) are reported for each calculation, followed by the length of the equilibration period ( $\tau_c$ ) and total number of iterations ( $N_i$ ). For each chemical system, the three methods share a similar computational cost per iteration.

reflected in the corresponding efficiencies (Figure 4, top), which are normalized based on the different number of



**Figure 4.** Increases in statistical efficiency are robust across five molecular systems and two choices of matrix factorization schemes, near-uniform (top) and heat-bath Power–Pitzer (bottom). Reported statistical efficiencies represent an average over the two independent trajectories obtained using each method and do not reflect differences in computational cost for systems with different sizes. Note the log scale on the y-axis.

iterations considered in the calculation of each standard error. For all systems, efficiencies for systematic FCI-FRI calculations are more than an order of magnitude greater than those for multinomial FCI-FRI calculations, which are in turn 2–113 times greater than those for FCIQMC calculations.

The integrated autocorrelation times (IATs), calculated as described in section 2.5 for all three methods, are similar within each system considered here. This is likely because the

same value of the imaginary time step,  $\epsilon$ , is used for each system (Table 2). A previous study<sup>34</sup> found that reducing the statistical error in matrix compression in FCIQMC enabled the use of greater values of  $\epsilon$ . This reduces the degree of correlation between iterates, thereby decreasing the IAT and increasing the statistical efficiency. This suggests that using greater values of  $\epsilon$  in the multinomial and systematic FCI-FRI methods could potentially increase the observed difference in their efficiencies. Furthermore, increasing  $\epsilon$  may reduce the equilibration times  $\tau_c$  for the FCI-FRI methods.

Because the systematic FCI-FRI method converges to the deterministic power method as  $N_{\text{mat}}$  and  $m$  approach finite values, we expect that the reported performance advantages for systematic FCI-FRI relative to the other two methods would increase for greater values of  $N_{\text{mat}}$  and  $m$ . On the other hand, because the compression schemes used in these FCI-FRI methods become more similar to those in FCIQMC as the size of the FCI basis increases relative to  $N_{\text{mat}}$  and  $m$ , the statistical efficiencies of these methods are expected to become more similar in this limit. For many systems, however, the values of  $N_{\text{mat}}$  and  $m$  required to calculate reasonably accurate energy estimates also increase with system size. In the calculations we have compared thus far, the values of these parameters are dictated by the critical number of walkers in FCIQMC.<sup>16</sup> Calculations for the Ne and HF systems were also compared with fewer matrix and vector samples. Using only 164 000 walkers in an FCIQMC calculation on Ne yields an energy estimate that differs from the exact energy by  $(-163 \pm 20783) \times 10^{-5} E_h$ , whereas a systematic FCI-FRI calculation with equivalent numbers of samples yields an energy estimate that differs by  $(0.58 \pm 5.15) \times 10^{-5} E_h$  after a similar number of iterations. The efficiencies of these two calculations differ by 7 orders of magnitude. A similar comparison for HF with only 812 000 walkers also shows a factor of  $10^7$  difference in efficiencies. This suggests that FCI-FRI methods may allow for the use of significantly fewer matrix and vector samples than the original FCIQMC method.

**3.2.2. Heat-Bath Power–Pitzer Factorization.** Results obtained using the three methods with the HB-PP factorization matrix mostly follow the same trends as those for the near-uniform factorization (Table 5). Standard errors for systematic

**Table 5.** Mean Energy Differences  $\pm$  Twice the Standard Error for Randomized Methods Using the Heat-Bath Power–Pitzer Factorization Scheme<sup>a</sup>

system	$m/10^3$	$N_{\text{mat}}/10^6$	FCIQMC			multinomial FCI-FRI			systematic FCI-FRI		
			$(E_{\text{diff}} \pm 2\sigma_E)/(10^{-5}E_h)$	$\tau_c/10^3$	$N_i/10^3$	$(E_{\text{diff}} \pm 2\sigma_E)/(10^{-5}E_h)$	$\tau_c/10^3$	$N_i/10^3$	$(E_{\text{diff}} \pm 2\sigma_E)/(10^{-5}E_h)$	$\tau_c/10^3$	$N_i/10^3$
Ne	242	0.26	$0.01 \pm 13.43$	15.0	902	$3.96 \pm 7.83$	15.0	917	$-0.44 \pm 1.61$	15.0	657
			$-3.41 \pm 13.22$	20.0	963	$0.75 \pm 7.92$	15.0	905	$-1.09 \pm 1.61$	15.0	686
HF	926	1.00	$-4.15 \pm 17.31$	130.0	502	$-4.78 \pm 18.28$	180.0	436	$-0.91 \pm 2.99$	40.0	447
			$4.41 \pm 15.89$	120.0	507	$0.66 \pm 13.42$	50.0	430	$-0.54 \pm 3.00$	27.4	654
H <sub>2</sub> O	491	0.57	$-12.33 \pm 10.66$	30.0	938	$-1.53 \pm 5.95$	20.0	645	$-0.30 \pm 1.52$	20.0	533
			$-4.04 \pm 10.45$	30.0	936	$-3.65 \pm 5.69$	20.0	646	$-0.18 \pm 1.63$	20.0	531
N <sub>2</sub>	997	1.15	$33.68 \pm 94.08$	200.0	663	$1.03 \pm 5.53$	53.2	699	$0.43 \pm 1.18$	64.6	373
			$55.74 \pm 75.77$	200.0	659	$4.19 \pm 5.07$	57.1	700	$-0.19 \pm 1.72$	72.7	372
C <sub>2</sub>	2620	4.14	$-11.20 \pm 17.99$	50.0	573	$-1.02 \pm 4.59$	190.0	432	$-0.15 \pm 2.02$	130.0	331
			$12.40 \pm 22.95$	140.0	581	$-0.12 \pm 5.61$	36.8	494	$-0.73 \pm 1.96$	50.0	213

<sup>a</sup>Parameters are reported for each trajectory as in Table 4 (iterate vector sparsity, number of matrix samples, and number of iterations).

and multinomial FCI-FRI calculations are less than those from FCIQMC, as is reflected in their associated efficiencies (Figure 4, bottom). One FCIQMC calculation on  $\text{H}_2\text{O}$  did not converge to within the 95% confidence interval, although given the relative magnitude of its standard error, this is likely a statistical anomaly. Systematic FCI-FRI calculations on  $\text{C}_2$  were particularly expensive due to the number of orbitals and cost of evaluating elements of matrices in the HB-PP factorization, rendering it difficult to accumulate sufficiently many samples to obtain an accurate estimate of the integrated autocovariance. Consequently, the estimated standard errors for these calculations are likely more inaccurate than for the other calculations in this study. This highlights the need for more efficient implementations of these FCI-FRI methods.

**3.3. Variational Energy Estimates.** Finally, we evaluate the possibility that the primary utility of the FCI-FRI methods considered here is that they efficiently identify the most important Slater determinant basis elements in the ground-state eigenvector. Variational Rayleigh quotients (eq 20) for a subset of the iterates (i.e., every 100th iterate) in each trajectory were calculated in addition to the projected estimates used to obtain average energies. If FCI-FRI is only an efficient search for significant basis elements, then we expect many of these Rayleigh quotients to be close to the ground-state energy.

We calculate the minimum Rayleigh quotient over both independent trajectories for each system considered. Differences between these minimum energies and the exact ground-state energies for each system are plotted in Figure 5. The mean energy difference from the original FCIQMC method is also plotted for comparison, with error bars denoting the corresponding 95% confidence interval. For all methods and systems considered, this difference for the minimum Rayleigh

quotient is more than an order of magnitude greater than the maximum of the FCIQMC confidence interval. The minimum Rayleigh quotients from FCIQMC are greater than those from the FCI-FRI methods considered and, for all systems except  $\text{C}_2$ , are also greater than the Hartree–Fock energy. This difference between the FCIQMC and FCI-FRI Rayleigh quotients can possibly be attributed to the lower-variance vector compression scheme employed in FCI-FRI. Even though the average of the FCIQMC iterates converges to the ground state to within a bias, the binomial integerization scheme used in FCIQMC displaces each iterate further from the ground state than in FCI-FRI.

These results indicate that none of the vectors from the FCIQMC or FCI-FRI trajectories are particularly close to the ground state, as measured by the variational energy estimates. The facts that the average of each component of the solution vector converges quickly to its exact value, to within a controllable statistical bias, and that the projected estimator is linear in these components, rather than quadratic, are essential for the success of FCI-FRI methods.


## 4. CONCLUSIONS

This paper describes several generic matrix and vector compression techniques within the FRI framework in the context of the FCI problem. Hierarchical approaches to matrix compression are discussed and shown to offer significant advantages over approaches that require enumerating all nonzero elements. Two examples of hierarchical factorization schemes for the FCI Hamiltonian matrix are presented, namely near-uniform and heat-bath Power–Pitzer. We describe how these various techniques can be combined in methods for calculating the FCI ground-state energy using power iteration, and we compare these “FCI-FRI” methods to FCIQMC in its original form.

Calculations on small molecules are used to compare the performance of these methods in terms of statistical efficiency, a metric inversely related to the square of the standard error. FCI-FRI calculations on the Ne atom demonstrate that using matrix compression in addition to vector compression can enable significant reductions in computational cost while only moderately decreasing the statistical efficiency.


We show that systematic matrix compression offers significant advantages over multinomial matrix compression, which has been used previously in FCIQMC. FCI-FRI calculations with systematic matrix compression applied to five small molecular systems are 11–45 times more efficient than those with multinomial compression, which are in turn 1.4–178 times more efficient than calculations performed using the original FCIQMC method.

The advantages of these stochastic methods over related deterministic compression methods are investigated. The error in a stochastic calculation on the Ne atom is nearly 2 orders of magnitude less than a deterministic calculation with comparable cost, which illustrates the importance of stochastically representing all components of the solution vector in the FCI Slater determinant space. Furthermore, by applying variational energy estimators to stochastic calculations performed on all molecular systems, we demonstrate the importance of averaging over many sparse, stochastic iterates in producing an accurate energy estimate. These features of stochastic methods and the results in this study suggest the applicability of FCI-FRI methods to strongly correlated systems with dense solution vectors.



**Figure 5.** Difference between the minimum variational energy estimate from each method and the exact FCI energy from Table 2. Results from the FCIQMC, multinomial FCI-FRI, and systematic FCI-FRI methods, using both the near-uniform (top) and heat-bath Power–Pitzer (bottom) matrix factorization schemes, are shown for each of the molecular systems considered in this study. Mean energy differences from the FCIQMC method for each system are plotted for comparison. Error bars represent 95% ( $2\sigma_E$ ) confidence intervals.





**Figure 6.** Hierarchical structure of the near-uniform factorization of the matrix  $\mathbf{P}^{(\tau)}$ , showing the orbital indices used to index the row and column spaces of each matrix in the factorization.

Future research will investigate strategies for further improving the performance of FCI-FRI methods. We will develop implementations of these methods that exploit parallelism more effectively, possibly using techniques developed previously for FCIQMC. Due to the generality of the FRI framework, the compression techniques introduced here can be applied in tandem with the complementary initiator and semistochastic extensions to FCIQMC, which suggests an approach to further improving statistical efficiency. Additionally, examining the effect of the choice of parameters used in FCI-FRI calculations on the statistical efficiency may provide additional insight into how to optimize performance. For example, our results suggest that FCI-FRI methods allow more flexibility than FCIQMC in the choice of the parameter  $\epsilon$ , which corresponds to the time step in imaginary time propagation. Varying  $\epsilon$  may affect the statistical efficiency of FCI-FRI methods. Furthermore, the number of nonzero elements in each matrix and vector compression in FCIQMC is determined by the number of walkers, whereas in FCI-FRI, these parameters can be varied independently. FCIQMC methods require a critical number of walkers to reliably converge to the ground-state energy. Our results suggest that using improved matrix compression schemes in FCI-FRI methods can reduce the number of matrix and vector elements required for convergence. Exploring these possibilities may facilitate the development of stochastic methods for quantum chemistry that are able to treat larger systems than currently possible.

#### Appendix A. Matrix Factorizations for Quantum Chemistry

This section describes two approaches to factoring the matrix  $\mathbf{P}^{(\tau)}$ , near-uniform and heat-bath Power–Pitzer (HB-PP). Elements in each matrix in the factorization are calculated using information from Hartree–Fock based on predetermined rules. The cost of evaluating these elements is greater for the HB-PP factorization than for near-uniform, although intermediate compression steps in the HB-PP scheme yield less statistical error than for near-uniform.

**A.1. Near-Uniform.** In the near-uniform factorization,<sup>7</sup>  $\mathbf{P}^{(\tau)}$  is factored into the product  $\mathbf{B}\mathbf{C}^{(\tau)}\mathbf{Q}$ , where  $\mathbf{Q}$  is factored further into a product of four matrices,  $\mathbf{Q}^{(4)}\mathbf{Q}^{(3)}\mathbf{Q}^{(2)}\mathbf{Q}^{(1)}$ . Elements of these four matrices can be calculated efficiently based on symmetry relationships between pairs of Slater determinants in the FCI basis. Elements of  $\mathbf{Q}$  differ from elements of  $\mathbf{P}^{(\tau)}$ , so multiplication by  $\mathbf{C}^{(\tau)}$  compensates for this

by multiplying by elements of  $\mathbf{P}^{(\tau)}$  and dividing by elements of  $\mathbf{Q}$ . This ensures that elements of the product of matrix factors are equal to those of  $\mathbf{P}^{(\tau)}$ . Finally, multiplication by  $\mathbf{B}$  sums elements corresponding to all excitations that contribute to each Slater determinant element of the final vector.

Each of the one-electron orbitals from a Hartree–Fock calculation can be assigned an associated irreducible representation (irrep) according to the symmetry of the system under consideration. This can encode spin symmetry (up or down), spatial (point group) symmetry, and, for crystalline systems,  $k$ -point symmetry. For each nonzero element in  $\mathbf{H}$  corresponding to a single excitation from  $|K\rangle$  to  $\hat{c}_a^\dagger\hat{c}_i|K\rangle$ , the irrep of orbital  $i$  must equal that of orbital  $a$ , i.e.,  $\Gamma_i = \Gamma_a$ . Because  $\mathbf{P}^{(\tau)}$  is related to  $\mathbf{H}$  by only a scalar factor and a shift by identity, its elements obey the same symmetry relationships. For double excitations, the direct product of irreps of the occupied orbitals,  $\Gamma_i \otimes \Gamma_j$ , must equal that of the virtual orbitals,  $\Gamma_a \otimes \Gamma_b$ , in order for the corresponding element of  $\mathbf{H}$  to be nonzero. Excitations satisfying these symmetry constraints are termed symmetry-allowed excitations. Applying this factorization scheme requires an  $O(N)$ -operation per nonzero element in the current iterate to count the number of occupied and virtual orbitals with each irrep.

The matrices in this factorization map Slater determinant basis elements to excitations, indexed using multi-indices containing the orbitals involved in each excitation, and ultimately back to the determinants defined by these excitations. A schematic overview of these relationships is presented in Figure 6. The matrix  $\mathbf{Q}^{(1)}$  has dimensions  $(3N_{\text{FCI}} \times N_{\text{FCI}})$ , and its row space can be divided into three distinct subspaces. The first corresponds to single excitations, and each element is indexed using a multi-index  $\{K,1\}$  denoting a generic single excitation from  $|K\rangle$ . Elements of  $\mathbf{Q}^{(1)}$  in this subspace are given as

$$Q_{\{K,1\},J}^{(1)} = \delta_{KJ} \frac{n_s}{n_s + n_d} \quad (\text{A1})$$

where  $\delta_{KJ}$  is a Kronecker delta, and  $n_s$  and  $n_d$  are the numbers of symmetry-allowed single and double excitations from a reference determinant in the FCI basis (typically Hartree–Fock). The second subspace contains generic double excitations and has elements given as

$$Q_{\{K,2\},J}^{(1)} = \delta_{KJ} \frac{n_d}{n_s + n_d} \quad (\text{A2})$$

Elements in the third subspace, indexed as  $\{K,0\}$ , will be mapped back to their original Slater determinant  $|K\rangle$  by the final matrix multiplication in the factorization. These must be considered separately in intermediate steps, as will be explained in the discussion of compression below. These “no excitation” elements are given as

$$Q_{\{K,0\},J}^{(1)} = \delta_{KJ} \quad (\text{A3})$$

The subsequent matrices in the factorization map generic single and double excitations from the row space of  $\mathbf{Q}^{(1)}$  to specific single and double excitations. This begins with multiplication by  $\mathbf{Q}^{(2)}$ , which maps to the specific occupied orbitals in these excitations. Single-excitation elements in this matrix are nonzero only for symmetry-allowed choices of occupied orbitals  $i$ . An occupied orbital in  $|K\rangle$  is symmetry-allowed if there is at least one virtual orbital of the same symmetry in  $|K\rangle$ . The number of such orbitals in  $|K\rangle$  is denoted  $n_K^{\text{occ}}$ . These single-excitation elements in  $\mathbf{Q}^{(2)}$  are

$$Q_{\{K,1,i\},\{K,1\}}^{(2)} = (n_K^{\text{occ}})^{-1} \quad (\text{A4})$$

Double excitation elements are nonzero for all of the  $N(N-1)/2$  unique pairs of occupied orbitals  $(ij)$  in  $|K\rangle$ , regardless of whether they have corresponding symmetry-allowed pairs of virtual orbitals. These elements are

$$Q_{\{K,2,(ij)\},\{K,2\}}^{(2)} = 2N^{-1}(N-1)^{-1} \quad (\text{A5})$$

The orbitals  $(ij)$  are grouped in the multi-index to indicate that their order is irrelevant to the indexing. As above, “no excitation” elements of  $\mathbf{Q}^{(2)}$  are 1, i.e.,

$$Q_{\{K,0\},\{K,0\}}^{(2)} = 1 \quad (\text{A6})$$

All other elements of  $\mathbf{Q}^{(2)}$  are 0.

Single-excitation elements in  $\mathbf{Q}^{(3)}$  map a virtual orbital to each excitation. Elements for symmetry-allowed virtual orbitals  $a$  are

$$Q_{\{K,1,i,a\},\{K,1,i\}}^{(3)} = (n_K^{\text{virt}}(i))^{-1} \quad (\text{A7})$$

where  $n_K^{\text{virt}}(i)$  is the number of virtual orbitals in  $|K\rangle$  with the same symmetry as  $i$ . Double excitation elements are defined for symmetry-allowed virtual orbitals  $a$ , i.e., those for which there exists at least one virtual orbital  $b$  that satisfies  $\Gamma_i \otimes \Gamma_j = \Gamma_a \otimes \Gamma_b$ :

$$Q_{\{K,2,(ij),a\},\{K,2,(ij)\}}^{(3)} = [n_K^{\text{virt}}(i, j)]^{-1} \quad (\text{A8})$$

where  $n_K^{\text{virt}}(ij)$  is the number of symmetry-allowed virtual orbitals given an occupied pair  $(ij)$ . “No excitation” elements of  $\mathbf{Q}^{(3)}$  are 1, and all other elements are 0.

Since single excitations are specified completely by the occupied and virtual orbitals in the row-space indices of  $\mathbf{Q}^{(3)}$ , single-excitation elements of  $\mathbf{Q}^{(4)}$  map these excitations to themselves, as follows:

$$Q_{\{K,1,i,a\},\{K,1,i,a\}}^{(4)} = 1 \quad (\text{A9})$$

Double-excitation elements for symmetry-allowed virtual orbitals  $b$  are

$$Q_{\{K,2,(ij),a,b\},\{K,2,(ij),a\}}^{(4)} = [n_K^{\text{virt}}(i, j, a)]^{-1} \quad (\text{A10})$$

where  $n_K^{\text{virt}}(i, j, a)$  denotes the number of symmetry-allowed virtual orbitals in  $|K\rangle$  given  $i, j$ , and  $a$ . Note that all of these orbitals  $b$  have the same irrep, since there is only one irrep  $\Gamma_b$  in the system’s point group that satisfies  $\Gamma_i \otimes \Gamma_j = \Gamma_a \otimes \Gamma_b$ .

The matrix  $\mathbf{C}^{(\tau)}$ , which ensures that the factorization is equal to  $\mathbf{P}^{(\tau)}$ , is a diagonal square matrix. Its single-excitation elements are

$$C_{\{K,1,i,a\},\{K,1,i,a\}}^{(\tau)} = \frac{P_K^{(\tau)}(i \rightarrow a)}{Q_{\{K,1\},K}^{(1)} Q_{\{K,1,i\},\{K,1\}}^{(2)} Q_{\{K,1,i,a\},\{K,1,i\}}^{(3)}} \quad (\text{A11})$$

The denominator corresponds to the “generation probability” in the original description of FCIQMC. Double excitation elements in  $\mathbf{C}^{(\tau)}$  are

$$C_{\{K,2,(ij),a,b\},\{K,2,(ij),a,b\}}^{(\tau)} = P_K^{(\tau)}(ij \rightarrow ab) (Q_{\{K,2\},K}^{(1)} Q_{\{K,2,(ij)\},\{K,2\}}^{(2)} Q_{\{K,2,(ij),a\},\{K,2,(ij)\}}^{(3)} Q_{\{K,2,(ij),a,b\},\{K,2,(ij),a\}}^{(4)} + Q_{\{K,2,(ij),b\},\{K,2,(ij)\}}^{(3)} Q_{\{K,2,(ij),b,a\},\{K,2,(ij),b\}}^{(4)})^{-1} \quad (\text{A12})$$

The sum of terms in the denominator of this expression accounts for the fact that there are two elements in the row space of  $\mathbf{C}^{(\tau)}$ , i.e.,  $\{K,2,(ij),a,b\}$  and  $\{K,2,(ij),b,a\}$ , corresponding to each double excitation, i.e.,  $K(ij \rightarrow ab)$ . These will be summed after multiplication by the final matrix in the factorization,  $\mathbf{B}$ . Elements in  $\mathbf{C}^{(\tau)}$  corresponding to “no excitation” elements in the basis are given as their corresponding diagonal elements in  $\mathbf{P}^{(\tau)}$ :

$$C_{\{K,0\},\{K,0\}}^{(\tau)} = P_{KK}^{(\tau)} \quad (\text{A13})$$

Multiplication by  $\mathbf{B}$  sums contributions from the row space of  $\mathbf{C}^{(\tau)}$  that map to the same Slater determinant basis element. Because there are many elements in this space that map to the same determinant, the row dimension of  $\mathbf{B}$  is smaller than the column dimension. Elements for double excitations are

$$B_{K(ij \rightarrow ab),\{K,2,(ij),a,b\}} = B_{K(ij \rightarrow ab),\{K,2,(ij),b,a\}} = 1 \quad (\text{A14})$$

and those for single excitations are


$$B_{K(i \rightarrow a),\{K,1,i,a\}} = 1 \quad (\text{A15})$$

“No excitation” elements are mapped back to the determinant from which they originated, i.e.,

$$B_{K,\{K,0\}} = 1 \quad (\text{A16})$$

This mapping can be performed efficiently using a hashing algorithm,<sup>7</sup> at  $O(N_{\text{mat}})$  cost, where  $N_{\text{mat}}$  is the number of elements selected from the matrix. In our current implementations of FCIQMC and FCI-FRI methods, a simpler  $O(N_{\text{mat}} \log m)$  binary search is used instead, where  $m$  is the number of nonzero iterate elements.

The selection of excitations from the near-uniform distribution in FCIQMC can be understood as a particular multinomial compression technique applied to the factorization scheme discussed above. Systematic compression could be applied analogously. However, one of the primary advantages of systematic compression is that it can be performed such that the elements selected from the vector are unique, thereby



**Figure 7.** An alternative near-uniform factorization structure for the matrix  $\mathbf{P}^{(\tau)}$ . Unlike in Figure 6, double excitations have unique indices, making this factorization more amenable to systematic compression.

yielding less statistical error. This benefit is somewhat diminished when using the factorization described above, since the indices  $\{K, 2, (ij), a, b\}$  and  $\{K, 2, (ij), b, a\}$  indicate the same double excitation but are treated separately until multiplication by  $\mathbf{B}$ . We therefore designed an alternative factorization scheme to address this (Figure 7). Elements of the matrix  $\mathbf{BQ}$  in this alternative scheme are equal to those in the original scheme; the difference only arises in how double-excitation elements are defined and indexed in  $\mathbf{Q}^{(3)}$  and  $\mathbf{Q}^{(4)}$ . Pairs of virtual orbitals and pairs of symmetry elements in the system's point group are used instead of individual virtual orbitals to index these elements. Applying the FCIQMC compression technique with the alternative scheme would yield the same statistical error; its advantages are realized only when using systematic compression.

Elements of the matrix  $\mathbf{Q}^{(3)}$  in this alternative factorization were obtained by summing double-excitation elements in the matrix product  $\mathbf{Q}^{(3)}\mathbf{Q}^{(4)}$  from the above factorization corresponding to pairs of irreps  $(\Gamma_x, \Gamma_y)$ . If the irreps of the occupied orbitals in a double excitation are equal ( $\Gamma_i = \Gamma_j$ ), then the irreps of the virtual orbitals must also be equal to satisfy the symmetry conditions described above. Double excitation elements of  $\mathbf{Q}^{(3)}$  corresponding to such pairs of occupied orbitals are

$$Q_{\{K, 2, (ij), (x, x)\}, \{K, 2, (ij)\}}^{(3)} = \begin{cases} n_K^{\text{virt}}(\Gamma_x) n_K^{\text{virt}}(i, j)^{-1} & n_K^{\text{virt}}(\Gamma_x) > 1, \\ \Gamma_i \otimes \Gamma_j = \Gamma_x \otimes \Gamma_x & \\ 0 & \text{otherwise} \end{cases} \quad (\text{A17})$$

Here  $x$  denotes a symmetry element in the system's point group,  $\Gamma_x$  is its associated irrep, and  $n_K^{\text{virt}}(\Gamma_x)$  denotes the number of virtual orbitals in  $|K\rangle$  with irrep  $\Gamma_x$ . If  $\Gamma_i \neq \Gamma_j$ , the corresponding elements of  $\mathbf{Q}^{(3)}$  are

$$Q_{\{K, 2, (ij), (x, y)\}, \{K, 2, (ij)\}}^{(3)} = \begin{cases} \frac{n_K^{\text{virt}}(\Gamma_x) + n_K^{\text{virt}}(\Gamma_y)}{n_K^{\text{virt}}(i, j)} & n_K^{\text{virt}}(\Gamma_x) > 0, n_K^{\text{virt}}(\Gamma_y) > 0, \\ \Gamma_i \otimes \Gamma_j = \Gamma_x \otimes \Gamma_y & \\ 0 & \text{otherwise} \end{cases} \quad (\text{A18})$$

Double excitation elements in  $\mathbf{Q}^{(4)}$  are given as the reciprocal of the number of virtual orbital pairs within each irrep pair. For pairs of virtual orbitals with the same irrep,

$$Q_{\{K, 2, (ij), (a, b)\}, \{K, 2, (ij), (x, x)\}}^{(4)} = \begin{cases} \frac{2}{n_K^{\text{virt}}(\Gamma_x)[n_K^{\text{virt}}(\Gamma_x) - 1]} & n_K^{\text{virt}}(\Gamma_x) > 1, \\ \Gamma_a = \Gamma_b = \Gamma_x & \\ 0 & \text{otherwise} \end{cases} \quad (\text{A19})$$

If instead  $\Gamma_a \neq \Gamma_b$ , the elements are

$$Q_{\{K, 2, (ij), (a, b)\}, \{K, 2, (ij), (x, y)\}}^{(4)} = \begin{cases} n_K^{\text{virt}}(\Gamma_x)^{-1} n_K^{\text{virt}}(\Gamma_y)^{-1} & n_K^{\text{virt}}(\Gamma_x) > 0, n_K^{\text{virt}}(\Gamma_y) > 0, \\ \Gamma_a = \Gamma_x, \Gamma_b = \Gamma_y & \\ 0 & \text{otherwise} \end{cases} \quad (\text{A20})$$

Except for the different indexing scheme for virtual orbitals in double excitations, elements in  $\mathbf{C}^{(\tau)}$  and  $\mathbf{B}$  are defined as above. Consequently, the elements of  $\mathbf{C}^{(\tau)}$  are as uniform in magnitude as in the factorization scheme above. Compression of either near-uniform factorization scheme can be performed at approximately  $O(N_{\text{mat}})$  cost.

**A.2. Heat-Bath Power–Pitzer.** In the above factorization, symmetry information is used to facilitate the efficient calculation of elements in the first four matrices, and discrepancies between products of these elements and elements of  $\mathbf{P}^{(\tau)}$  are eliminated through multiplication by  $\mathbf{C}^{(\tau)}$ . Less error is introduced by stochastic compression of this factorization when  $\mathbf{Q}$  is closer to  $\mathbf{P}^{(\tau)}$ , i.e., when the magnitudes of elements of  $\mathbf{C}^{(\tau)}$  are more uniform.<sup>34,39</sup> The heat-bath Power–Pitzer (HB-PP) factorization is designed to achieve more uniformity in these elements by using information from the Hamiltonian matrix in constructing  $\mathbf{Q}$ , which is factored into a product of five matrices,  $\mathbf{Q}^{(5)}\mathbf{Q}^{(4)}\mathbf{Q}^{(3)}\mathbf{Q}^{(2)}\mathbf{Q}^{(1)}$ . Elements in these matrices are indexed by individual orbitals rather than unique pairs of orbitals or symmetry elements. Because it is expensive to incorporate information about single-excitation Hamiltonian elements into these matrices, due to the  $O(N)$  cost of evaluating each element, single-excitation elements in the factors of  $\mathbf{Q}$  are defined exactly as in the near-uniform case.



The same is true for “no excitation” elements, for reasons that will be made apparent in [Appendix B](#).

Elements corresponding to double excitations in  $\mathbf{Q}^{(1)}$  are also defined as in the near-uniform case. Double-excitation elements in subsequent matrices are defined in terms of a matrix  $\mathbf{D}$  and vector  $\mathbf{S}$ . Elements of  $\mathbf{D}$  approximate the sum of magnitudes of double-excitation elements in the Hamiltonian corresponding to a particular pair of occupied orbitals,<sup>34</sup>

$$D_{pq} = \begin{cases} \sum_{r,s \notin \{p,q\}} | \langle pq || rs \rangle | & p \neq q \\ 0 & p = q \end{cases} \quad (\text{A21})$$

where  $\langle pq || rs \rangle$  is an antisymmetrized two-electron integral obtained from the Hartree–Fock calculation. The exact sum for each determinant depends on which orbitals are occupied, so it is approximated by an unrestricted sum over all other orbitals in the Hartree–Fock basis. Analogously, elements of  $\mathbf{S}$  approximate this sum for a single occupied orbital,

$$S_p = \sum_q D_{pq} \quad (\text{A22})$$

The primary advantage of defining  $\mathbf{S}$  and  $\mathbf{D}$  by unrestricted sums is that they can be computed and stored in memory at the beginning of the simulation, at a memory cost of  $O(M^2)$  and a CPU cost of  $O(M^4)$ .

The row spaces of  $\mathbf{Q}^{(2)}$  and  $\mathbf{Q}^{(3)}$  are indexed by multi-indices containing individual occupied orbitals instead of unique pairs of occupied orbitals. Elements in  $\mathbf{Q}^{(2)}$  are

$$Q_{\{K,2,i\},\{K,2\}}^{(2)} = \frac{S_i}{\sum_{j \in \text{occ}} S_j} \quad (\text{A23})$$

and those in  $\mathbf{Q}^{(3)}$  are

$$Q_{\{K,2,i,j\},\{K,2,i\}}^{(3)} = \frac{D_{ij}}{\sum_{j' \in \text{occ}} D_{ij'}} \quad (\text{A24})$$

As a consequence of this indexing scheme, pairs of elements in  $\mathbf{Q}^{(3)}$  in which the order of the occupied orbitals is reversed are not necessarily equal, i.e.,

$$Q_{\{K,2,i,j\},\{K,2,i\}}^{(3)} \neq Q_{\{K,2,j,i\},\{K,2,i\}}^{(3)} \quad (\text{A25})$$

Elements in the next matrix corresponding to double excitations with one virtual orbital are defined as

$$Q_{\{K,2,i,j,a\},\{K,2,i,j\}}^{(4)} = \frac{|\langle iall ai \rangle|^{1/2}}{\sum_{c \in \{\text{virt}\}} |\langle icll ci \rangle|^{1/2}} \quad (\text{A26})$$

where  $\langle iall ai \rangle$  represents a two-electron exchange integral. Note that if the spins of orbitals  $i$  and  $a$  differ, this integral is 0. The sum in the denominator includes all virtual orbitals in  $|K\rangle$ . Elements in  $\mathbf{Q}^{(5)}$  are indexed by a second virtual orbital and are defined as

$$Q_{\{K,2,i,j,a,b\},\{K,2,i,j,a\}}^{(5)} = \frac{|\langle jblb \rangle|^{1/2} \delta_{\Gamma_b \otimes \Gamma_a, \Gamma_i \otimes \Gamma_j}}{\sum_c |\langle jc ll cj \rangle|^{1/2} \delta_{\Gamma_c \otimes \Gamma_a, \Gamma_i \otimes \Gamma_j}} \quad (\text{A27})$$

where the Kronecker deltas enforce the symmetry condition for double excitations described in [section A.1](#), and the sum includes *all* orbitals in the basis, including those occupied in

$|K\rangle$ . Elements in  $\mathbf{Q}^{(5)}$  corresponding to single excitations are defined in analogy to [eq A9](#).

Elements of the matrix  $\mathbf{C}^{(\tau)}$  are

$$\begin{aligned} C_{\{K,i,j,a,b\},\{K,2,i,j,a,b\}}^{(\tau)} &= P_K^{(\tau)}(ij \rightarrow ab)(Q_{\{K,2\},K}^{(1)})^{-1} \times \\ &[Q_{\{K,2,i\},\{K,2\}}^{(2)} Q_{\{K,2,i,j\},\{K,2,i\}}^{(3)} \times \\ &(Q_{\{K,2,i,j,a\},\{K,2,i,j\}}^{(4)} Q_{\{K,2,i,j,a,b\},\{K,2,i,j,a\}}^{(5)} \\ &+ Q_{\{K,2,i,j,b\},\{K,2,i,j\}}^{(4)} Q_{\{K,2,i,j,b,a\},\{K,2,i,j,b\}}^{(5)} \\ &+ Q_{\{K,2,j\},\{K,2\}}^{(2)} Q_{\{K,2,j,i\},\{K,2,j\}}^{(3)} \times \\ &(Q_{\{K,2,j,i,a\},\{K,2,j,i\}}^{(4)} Q_{\{K,2,j,i,a,b\},\{K,2,j,i,a\}}^{(5)} \\ &+ Q_{\{K,2,j,i,b\},\{K,2,j,i\}}^{(4)} Q_{\{K,2,j,i,b,a\},\{K,2,j,i,b\}}^{(5)})]^{-1} \end{aligned} \quad (\text{A28})$$

where the four terms in the sum account for the four different orders in which the orbitals for the double excitation can be chosen. The matrix  $\mathbf{B}$  is defined analogously to the near-uniform factorization. The cost of performing the compressions for the HB-PP scheme scales as  $O(MN_{\text{mat}})$ .

## Appendix B. Compression Schemes in FCIQMC and FCI-FRI

In principle, any compression scheme could be used to compress the intermediate vectors generated after each matrix multiplication in the hierarchical factorization schemes described above. This section describes the specific schemes used in the original FCIQMC method, as well as in multinomial and systematic FCI-FRI. Previously, FCIQMC has been described in terms of a sequence of “spawning”, “death/cloning”, and “annihilation” steps. This section presents an alternative interpretation of the method using the language of FRI.

Different subspaces of the intermediate vectors are treated differently in the compression schemes used in each of these methods. In each vector obtained after multiplying by the factors of  $\mathbf{Q}$ , “no excitation” elements are preserved exactly in all methods considered in this study. This is because diagonal elements of  $\mathbf{P}^{(\tau)}$  are often significantly greater in magnitude than off-diagonal elements, provided that  $\epsilon$  is sufficiently small ([eq 5](#)).

In FCIQMC, the remaining portions of the vectors are compressed using multinomial sampling, without exact preservation of elements, with the added constraint that certain numbers of samples are allocated to each subspace. The number of samples allocated to an arbitrary subspace  $w$  is denoted  $n_w$ . The number of samples allocated to the space of excitations associated with each Slater determinant is given as

$$n_{\{K\}} = |\nu_K| \quad (\text{B1})$$

The number of elements in each single- and double-excitation subspace is determined by counting the number of samples in each subspace during multinomial compression of the vector  $\mathbf{Q}^{(1)}\mathbf{v}^{(\tau)}$ . The numbers in the remaining subspaces are calculated analogously following each matrix multiplication. The total number of samples used in each compression is denoted  $N_{\text{mat}}$ . Approaches to performing this sampling efficiently for the near-uniform and HB-PP factorizations are described in [refs 7, 34, and 39](#).

Compression in FCIQMC is performed differently following multiplication by  $\mathbf{C}^{(\tau)}$  in both factorizations, so that each element in the resulting vector is an integer. If  $\mathbf{x}'$  denotes the

vector obtained after multiplication by this matrix, “no excitation” elements in the compressed vector are given as

$$\Phi(\mathbf{x}')_{\{K,0\}} = \sum_{i=1}^{|\nu_K|} \text{bin}^{(i)}(x'_{\{K,0\}} \nu_K^{-1}) \quad (\text{B2})$$

The function  $\text{bin}^{(i)}(x)$  denotes the binomial integerization of a number  $x$ , defined as

$$\text{bin}^{(i)}(x) = \lfloor x + r^{(i)} \rfloor \quad (\text{B3})$$

where  $r^{(i)}$  is a random number chosen uniformly on the interval (0,1). This function preserves its argument in expectation, i.e.,  $E[\text{bin}^{(i)}(x)] = x$ . Different values of the superscript  $i$  correspond to independent random numbers. The argument of this function in eq B2 is related to the “death/cloning probability” in previous presentations of FCIQMC, and performing the sampling corresponds to the “diagonal death/cloning” step. Other elements in the compressed vector, corresponding to off-diagonal elements in  $\mathbf{P}^{(\tau)}$ , are

$$\Phi(\mathbf{x}')_i = \sum_{j=1}^{n_i} \text{bin}(x'_i n_i^{-1}) \quad (\text{B4})$$

where the index  $i$  indicates an excitation, e.g.,  $\{K,(ij),a,b\}$  or  $\{K,i,a\}$ . Here the argument of the binomial integerization function corresponds to the “spawning probability” in FCIQMC. The resulting vector is sparse because many elements are set to zero by the binomial integerization function. The number of nonzero elements is random, unlike the systematic vector compression. Multiplication by  $\mathbf{B}$  constitutes the “annihilation” step in FCIQMC, as it involves summing elements that are mapped to the same Slater determinant basis element.

In multinomial FCI-FRI, multinomial compression is also used to compress the first few intermediate vectors. Because the elements of  $\mathbf{v}^{(\tau)}$  are not necessarily integers, a separate systematic sampling procedure is applied to determine the number of elements  $n_{\{K\}}$  to sample from the subspace associated with each Slater determinant. The magnitudes of elements in  $\mathbf{v}^{(\tau)}$  are normalized to obtain the probabilities  $\{p_i\}$  used in systematic sampling, and a constraint is added: for all  $K$  for which  $|\nu_K^{(\tau)}| > 0$ ,  $n_{\{K\}} > 0$ . In contrast to FCIQMC, compression is not performed after multiplication by  $\mathbf{C}^{(\tau)}$ , so the elements summed during multiplication by  $\mathbf{B}$  are not necessarily integers. In systematic FCI-FRI, the first few intermediate vectors are compressed systematically to  $N_{\text{mat}}$  elements instead of multinomially, preserving  $\rho$  elements exactly in each compression according to eq 14. Unlike in multinomial compression, the constraint that a certain number of elements are selected from each subspace is not imposed. Because the order of elements determines which elements are chosen in systematic compression, elements are ordered consistently in each iteration, first by the Slater determinant index, then by the type of excitation (single, double, or “no excitation”), then by the occupied and virtual orbital(s).

## AUTHOR INFORMATION

### Corresponding Authors

\*E-mail: [weare@cims.nyu.edu](mailto:weare@cims.nyu.edu).

\*E-mail: [tim.berkelbach@gmail.com](mailto:tim.berkelbach@gmail.com).

### ORCID

Samuel M. Greene: 0000-0001-5135-5926

Timothy C. Berkelbach: 0000-0002-7445-2136

## Notes

The authors declare no competing financial interest.

## ACKNOWLEDGMENTS

We thank Aaron Dinner and Sandeep Sharma for useful discussions about this work, and we thank Anthony Scemama for noting an error in the geometry of H<sub>2</sub>O used in calculations in a preprint. S.M.G. and T.C.B. were supported by start-up funding from the University of Chicago and by the Flatiron Institute. The Flatiron Institute is a division of the Simons Foundation. R.J.W. and J.W. were supported by the Advanced Scientific Computing Research program through award DE-SC0014205. R.J.W. was also supported by NSF RTG award 1547396 at the University of Chicago and by a MacCracken Fellowship and NSF RTG award 1646339 at New York University. Calculations were performed with resources provided by the University of Chicago Research Computing Center.

## REFERENCES

- (1) Troyer, M.; Wiese, U.-J. Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Phys. Rev. Lett.* **2005**, *94*, 170201.
- (2) Barker, J. A. A quantum statistical Monte Carlo method; path integrals with boundary conditions. *J. Chem. Phys.* **1979**, *70*, 2914–2918.
- (3) Hammond, B. L.; Lester, W. A.; Reynolds, P. J. *Monte Carlo methods in ab initio quantum chemistry*; World Scientific: Singapore, 1994.
- (4) Calandra Buonaura, M.; Sorella, S. Numerical study of the two-dimensional Heisenberg model using a Green function Monte Carlo technique with a fixed number of walkers. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1998**, *57*, 11446–11456.
- (5) Maksym, P. Auxiliary field quantum Monte-Carlo simulation of interacting electrons in quantum dots. *Phys. E (Amsterdam, Neth.)* **2005**, *26*, 257–261.
- (6) Needs, R. J.; Towler, M. D.; Drummond, N. D.; López Ríos, P. Continuum variational and diffusion quantum Monte Carlo calculations. *J. Phys.: Condens. Matter* **2010**, *22*, 023201.
- (7) Booth, G. H.; Smart, S. D.; Alavi, A. Linear-scaling and parallelizable algorithms for stochastic quantum chemistry. *Mol. Phys.* **2014**, *112*, 1855–1869.
- (8) Austin, B. M.; Zubarev, D. Y.; Lester, W. A. Quantum Monte Carlo and Related Approaches. *Chem. Rev.* **2012**, *112*, 263–288.
- (9) Shepherd, J. J.; Scuseria, G. E.; Spencer, J. S. Sign problem in full configuration interaction quantum Monte Carlo: Linear and sublinear representation regimes for the exact wave function. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2014**, *90*, 155130.
- (10) Umrigar, C. J.; Nightingale, M. P.; Runge, K. J. A diffusion Monte Carlo algorithm with very small time-step errors. *J. Chem. Phys.* **1993**, *99*, 2865–2890.
- (11) Senatore, G.; March, N. H. Recent progress in the field of electron correlation. *Rev. Mod. Phys.* **1994**, *66*, 445–479.
- (12) Kosztin, I.; Faber, B.; Schulten, K. Introduction to the diffusion Monte Carlo method. *Am. J. Phys.* **1996**, *64*, 633–644.
- (13) Foulkes, W. M. C.; Mitas, L.; Needs, R. J.; Rajagopal, G. Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.* **2001**, *73*, 33–83.
- (14) Manten, S.; Luchow, A. On the accuracy of the fixed-node diffusion quantum Monte Carlo method. *J. Chem. Phys.* **2001**, *115*, 5362–5366.
- (15) Hairer, M.; Weare, J. Improved Diffusion Monte Carlo. *Commun. Pure Appl. Math.* **2014**, *67*, 1995–2021.
- (16) Booth, G. H.; Thom, A. J. W.; Alavi, A. Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space. *J. Chem. Phys.* **2009**, *131*, 054106.

- (17) Li, Z.-X.; Jiang, Y.-F.; Yao, H. Solving the fermion sign problem in quantum Monte Carlo simulations by Majorana representation. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2015**, *91*, 241117.
- (18) Alavi, A. Introduction to the Full Configuration Interaction Quantum Monte Carlo method with applications to the Hubbard model. In *Quantum Materials: Experiments and Theory*; Pavarini, E., Koch, E., van den Brink, J., Sawatzky, G., Eds.; Forschungszentrum Jülich GmbH, 2016.
- (19) Motta, M.; Zhang, S. Ab initio computations of molecular systems by the auxiliary-field quantum Monte Carlo method. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2018**, *8*, e1364.
- (20) Spencer, J. S.; Blunt, N. S.; Foulkes, W. M. The sign problem and population dynamics in the full configuration interaction quantum Monte Carlo method. *J. Chem. Phys.* **2012**, *136*, 054110.
- (21) Umrigar, C. J. Observations on variational and projector Monte Carlo methods. *J. Chem. Phys.* **2015**, *143*, 164105.
- (22) Halkier, A.; Helgaker, T.; Jørgensen, P.; Klopper, W.; Koch, H.; Olsen, J.; Wilson, A. K. Basis-set convergence in correlated calculations on Ne, N<sub>2</sub>, and H<sub>2</sub>O. *Chem. Phys. Lett.* **1998**, *286*, 243–252.
- (23) Halkier, A.; Helgaker, T.; Jørgensen, P.; Klopper, W.; Olsen, J. Basis-set convergence of the energy in molecular Hartree-Fock calculations. *Chem. Phys. Lett.* **1999**, *302*, 437–446.
- (24) Lim, L.-H.; Weare, J. Fast Randomized Iteration: Diffusion Monte Carlo through the Lens of Numerical Linear Algebra. *SIAM Rev.* **2017**, *59*, 547–587.
- (25) Lu, J.; Wang, Z. The Full Configuration Interaction Quantum Monte Carlo Method in the Lens of Inexact Power Iteration. 2017, *arXiv:1711.09153v3*. arXiv.org ePrint archive, <https://arxiv.org/abs/1711.09153> (accessed Jul 9, 2019).
- (26) Petruzielo, F. R.; Holmes, A. A.; Changlani, H. J.; Nightingale, M. P.; Umrigar, C. J. Semistochastic Projector Monte Carlo Method. *Phys. Rev. Lett.* **2012**, *109*, 230201.
- (27) Blunt, N. S.; Smart, S. D.; Kersten, J. A. F.; Spencer, J. S.; Booth, G. H.; Alavi, A. Semi-stochastic full configuration interaction quantum Monte Carlo: Developments and application. *J. Chem. Phys.* **2015**, *142*, 184107.
- (28) Overy, C.; Booth, G. H.; Blunt, N. S.; Shepherd, J. J.; Cleland, D.; Alavi, A. Unbiased reduced density matrices and electronic properties from full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2014**, *141*, 244117.
- (29) Cleland, D.; Booth, G. H.; Alavi, A. Communications: Survival of the fittest: Accelerating convergence in full configuration-interaction quantum Monte Carlo. *J. Chem. Phys.* **2010**, *132*, 041103.
- (30) Booth, G. H.; Cleland, D.; Thom, A. J. W.; Alavi, A. Breaking the carbon dimer: The challenges of multiple bond dissociation with full configuration interaction quantum Monte Carlo methods. *J. Chem. Phys.* **2011**, *135*, 084104.
- (31) Cleland, D. M.; Booth, G. H.; Alavi, A. A study of electron affinities using the initiator approach to full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2011**, *134*, 024112.
- (32) Blunt, N. S. Communication: An efficient and accurate perturbative correction to initiator full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2018**, *148*, 221101.
- (33) Knowles, P.; Handy, N. A new determinant-based full configuration interaction method. *Chem. Phys. Lett.* **1984**, *111*, 315–321.
- (34) Holmes, A. A.; Changlani, H. J.; Umrigar, C. J. Efficient Heat-Bath Sampling in Fock Space. *J. Chem. Theory Comput.* **2016**, *12*, 1561–1571.
- (35) Blunt, N. S.; Smart, S. D.; Booth, G. H.; Alavi, A. An excited-state approach within full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2015**, *143*, 134117.
- (36) Walker, A. New fast method for generating discrete random numbers with arbitrary frequency distributions. *Electron. Lett.* **1974**, *10*, 127.
- (37) Thom, A. J. W. Stochastic Coupled Cluster Theory. *Phys. Rev. Lett.* **2010**, *105*, 263004.
- (38) Scott, C. J. C.; Thom, A. J. W. Stochastic coupled cluster theory: Efficient sampling of the coupled cluster expansion. *J. Chem. Phys.* **2017**, *147*, 124105.
- (39) Neufeld, V. A.; Thom, A. J. W. Exciting Determinants in Quantum Monte Carlo: Loading the Dice with Fast, Low-Memory Weights. *J. Chem. Theory Comput.* **2019**, *15*, 127–140.
- (40) RESiPy: Randomized Electronic Structure in Python. <https://github.com/sgreene8/resipy/> (accessed Jul 9, 2019).
- (41) Olsen, J.; Christiansen, O.; Koch, H.; Jørgensen, P. Surprising cases of divergent behavior in Møller-Plesset perturbation theory. *J. Chem. Phys.* **1996**, *105*, 5082–5090.
- (42) Sun, Q.; Berkelbach, T. C.; Blunt, N. S.; Booth, G. H.; Guo, S.; Li, Z.; Liu, J.; McClain, J. D.; Sayfutyarova, E. R.; Sharma, S.; Wouters, S.; Chan, G. K.-L. PySCF: the Python-based simulations of chemistry framework. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2018**, *8*, No. e1340.
- (43) Vigor, W. A.; Spencer, J. S.; Bearpark, M. J.; Thom, A. J. W. Minimising biases in full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2015**, *142*, 104101.
- (44) Sokal, A. *Monte Carlo Methods in Statistical Mechanics: Foundations and New Algorithms*; Springer: Boston, MA, 1997; pp 131–192.
- (45) Foreman-Mackey, D.; Hogg, D. W.; Lang, D.; Goodman, J. *emcee: The MCMC Hammer*. 2013, *arXiv:1202.3665v4*. arXiv.org ePrint archive, <https://arxiv.org/abs/1202.3665> (accessed Jul 9, 2019).
- (46) Chodera, J. D. A Simple Method for Automated Equilibration Detection in Molecular Simulations. *J. Chem. Theory Comput.* **2016**, *12*, 1799–1805.
- (47) Flyvbjerg, H.; Petersen, H. G. Error estimates on averages of correlated data. *J. Chem. Phys.* **1989**, *91*, 461–466.
- (48) Vigor, W. A.; Spencer, J. S.; Bearpark, M. J.; Thom, A. J. W. Understanding and improving the efficiency of full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2016**, *144*, 094110.
- (49) Chung, K. L. *Markov Chains with Stationary Transition Probabilities*; Springer: Berlin, 1960; pp 93–106.
- (50) Webber, R. J.; Plotkin, D. A.; O'Neill, M. E.; Abbot, D. S.; Weare, J. Practical rare event sampling for extreme mesoscale weather. *Chaos* **2019**, *29*, 053109.
- (51) Huron, B.; Malrieu, J. P.; Rancurel, P. Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions. *J. Chem. Phys.* **1973**, *58*, 5745–5759.
- (52) Tubman, N. M.; Lee, J.; Takeshita, T. Y.; Head-Gordon, M.; Whaley, K. B. A deterministic alternative to the full configuration interaction quantum Monte Carlo method. *J. Chem. Phys.* **2016**, *145*, 044112.
- (53) Zhang, T.; Evangelista, F. A. A Deterministic Projector Configuration Interaction Approach for the Ground State of Quantum Many-Body Systems. *J. Chem. Theory Comput.* **2016**, *12*, 4326–4337.
- (54) Sharma, S.; Holmes, A. A.; Jeanmairet, G.; Alavi, A.; Umrigar, C. J. Semistochastic Heat-Bath Configuration Interaction Method: Selected Configuration Interaction with Semistochastic Perturbation Theory. *J. Chem. Theory Comput.* **2017**, *13*, 1595–1604.
- (55) Wang, Z.; Li, Y.; Lu, J. *Coordinate descent full configuration interaction*. 2019, *arXiv:1902.04592v2*. arXiv.org ePrint archive, 2019; <http://arxiv.org/abs/1902.04592v2>, (accessed Jul 9, 2019).