

Minimising biases in full configuration interaction quantum Monte Carlo

W. A. Vigor, J. S. Spencer, M. J. Bearpark, and A. J. W. Thom

Citation: *The Journal of Chemical Physics* **142**, 104101 (2015); doi: 10.1063/1.4913644

View online: <http://dx.doi.org/10.1063/1.4913644>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/jcp/142/10?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Goal-oriented sensitivity analysis for lattice kinetic Monte Carlo simulations](#)

J. Chem. Phys. **140**, 124108 (2014); 10.1063/1.4868649

[Communication: Excited states, dynamic correlation functions and spectral properties from full configuration interaction quantum Monte Carlo](#)

J. Chem. Phys. **137**, 191102 (2012); 10.1063/1.4766327

[Breaking the carbon dimer: The challenges of multiple bond dissociation with full configuration interaction quantum Monte Carlo methods](#)

J. Chem. Phys. **135**, 084104 (2011); 10.1063/1.3624383

[Approaching chemical accuracy using full configuration-interaction quantum Monte Carlo: A study of ionization potentials](#)

J. Chem. Phys. **132**, 174104 (2010); 10.1063/1.3407895

[Monte Carlo configuration interaction predictions for the electronic spectra of Ne, CH₂, C₂, N₂, and H₂O compared to full configuration interaction calculations](#)


J. Chem. Phys. **129**, 064103 (2008); 10.1063/1.2965529

How can you **REACH 100%**
of researchers at the Top 100
Physical Sciences Universities? (TIMES HIGHER EDUCATION RANKINGS, 2014)

With *The Journal of Chemical Physics*.

AIP | The Journal of
Chemical Physics

THERE'S POWER IN NUMBERS. Reach the world with AIP Publishing.



Minimising biases in full configuration interaction quantum Monte Carlo

W. A. Vitor,¹ J. S. Spencer,^{2,3} M. J. Bearpark,¹ and A. J. W. Thom^{1,4}

¹Department of Chemistry, Imperial College London, Exhibition Road, London SW7 2AZ, United Kingdom

²Department of Physics, Imperial College London, Exhibition Road, London SW7 2AZ, United Kingdom

³Department of Materials, Imperial College London, Exhibition Road, London SW7 2AZ, United Kingdom

⁴University Chemical Laboratory, Lensfield Road, Cambridge CB2 1EW, United Kingdom

(Received 23 January 2015; accepted 16 February 2015; published online 9 March 2015)

We show that Full Configuration Interaction Quantum Monte Carlo (FCIQMC) is a Markov chain in its present form. We construct the Markov matrix of FCIQMC for a two determinant system and hence compute the stationary distribution. These solutions are used to quantify the dependence of the population dynamics on the parameters defining the Markov chain. Despite the simplicity of a system with only two determinants, it still reveals a population control bias inherent to the FCIQMC algorithm. We investigate the effect of simulation parameters on the population control bias for the neon atom and suggest simulation setups to, in general, minimise the bias. We show a reweighting scheme to remove the bias caused by population control commonly used in diffusion Monte Carlo [Umrigar *et al.*, *J. Chem. Phys.* **99**, 2865 (1993)] is effective and recommend its use as a post processing step. © 2015 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4913644>]

I. INTRODUCTION

A cheap and accurate computational description of the ground state energy of a chemical system remains one of the principal challenges in electronic structure theory, yet achieving both of these goals systematically remains beyond the grasp of current approximations. Hierarchies of methods of increasing sophistication have been developed in the quantum chemistry community which systematically capture increasing amounts of the electron–electron correlation energy at the expense of additional computational cost. These methods start from Hartree–Fock¹ which scales modestly with the fourth power of the number of electrons to Full Configuration Interaction (FCI) which captures the maximal amount of electron–electron correlation in a finite basis set but scales factorially with the number of electrons. Approximations (which are often very accurate) such as density fitting can potentially reduce the scaling of these methods.² If FCI is used with a large enough basis set or an extrapolation to the complete basis set limit,^{3,4} energy differences can be obtained to chemical accuracy (1 kcal/mol) providing direct comparison with experiment. Unfortunately, the factorial scaling with the number of electrons makes it unfeasible for studying anything but the smallest of chemical systems.

Full Configuration Interaction Quantum Monte Carlo (FCIQMC)⁵ marries FCI with a projector Monte Carlo paradigm but crucially requires no *a priori* knowledge of the sign structure of the wavefunction. FCIQMC has two principal advantages over conventional FCI: the storage requirements are greatly reduced due to a sparse stochastic representation of the wavefunction,⁵ and it can be efficiently parallelised.⁶ The storage requirements for FCIQMC depend on the (system-dependent) severity of the Fermion sign problem⁷ and are often orders of magnitude less than conventional FCI calculations. FCIQMC with the controllable initiator approximation⁸ has allowed molecular systems with Hilbert spaces of 10^{29} Slater

determinants⁹ and the uniform electron gases with Hilbert spaces containing up to 10^{108} determinants¹⁰ to be studied. The FCIQMC methodology was subsequently extended to coupled cluster,¹¹ and we believe stochastic approaches are becoming increasingly important to the quantum chemistry community due to the need for scalable algorithms which are well-suited to modern computer architectures.

Some questions still remain concerning the best way to use Monte Carlo to solve the FCI equations. The FCIQMC algorithm is not a black box, and a choice has to be made about calculation parameters which control the stochastic sampling and hence the systematic and stochastic errors inherent to the simulation for a given amount of computational resources. In this article, we investigate the behaviour of FCIQMC simulations to understand the relationship between parameter choices and errors by investigating the exact distribution obtained from a Markov chain transition matrix.

Section II contains a brief recap of the FCIQMC method. We show in Sec. III that FCIQMC is an example of Markov Chain Monte Carlo (MCMC). We use these ideas to investigate population control bias in the two determinant H_2 system and more realistic calculations on the neon atom in Sec. IV. We draw conclusions and provide suggestions on simulation strategies in Sec. V. Atomic units are used throughout. The many-electron Hamiltonian and all energies have been shifted to be relative to the absolute Hartree–Fock energies of the appropriate system. Error bars signify one standard error, an estimate of the standard deviation, either side of the mean value.

II. FCIQMC

We briefly review the FCIQMC method, which is discussed in more detail in (e.g.) Refs. 5 and 7. The imaginary-time Schrödinger equation is

$$\frac{\partial|\Psi\rangle}{\partial\tau} = -(\hat{H} - S)|\Psi\rangle, \quad (1)$$

where S is an energy offset, which we shall discuss in the context of FCIQMC later, introduced to control normalisation. The general solution to Eq. (1) is $|\Psi(\tau)\rangle = e^{-\tau(\hat{H}-S)}|\Psi(\tau=0)\rangle$, which in the long-time limit tends to the lowest eigenstate with which the initial wavefunction has a non-zero overlap.

We begin with the Configuration Interaction (CI) ansatz where the wavefunction is a linear combination of Slater determinants: $|\psi\rangle = \sum_i C_i |D_i\rangle$. It is convenient (though not necessary¹²) to represent the coefficients by a discrete set of signed particles which we shall call psips following Anderson.¹³ Booth *et al.*⁵ showed that a finite-difference approximation to Eq. (1) could be sampled by allowing a psip on one determinant to create a new psip on another determinant (“spawn”) or on the same determinant (“death”) with probability proportional to the connecting Hamiltonian matrix element. Pairs of psips with opposite signs on the same determinant are removed (“annihilated”) at the end of each timestep. After a sufficient number of such steps, the psip vector becomes a stochastic representation of the eigenvector. The finite difference approximation introduces no timestep errors if the timestep, $\delta\tau$, satisfies $\delta\tau < 2(E_{\max} - E_0)^{-1}$, where E_{\max} (E_0) is the highest (lowest) eigenvalue of the Hamiltonian;⁷ a property FCIQMC shares with Green’s function quantum Monte Carlo.¹⁴

Following an equilibration phase, the shift is periodically updated every A steps to control the psip population using⁵

$$S(\tau + A\delta\tau) = S(\tau) - \frac{\gamma}{A\delta\tau} \log \frac{N(\tau + A\delta\tau)}{N(\tau)}, \quad (2)$$

where γ is a damping factor and $N(\tau)$ is the total number of psips at time τ . Repeated substitution of Eq. (2) into itself yields

$$S(\tau + A\delta\tau) = S(0) - \xi \log \frac{N(\tau + A\delta\tau)}{N_s}, \quad (3)$$

where $S(0)$ is the initial value of the shift (in this work, the Hartree–Fock energy), N_s is the population at the end of the equilibration phase, and $\xi = \gamma/(A\delta\tau)$ is usually fixed during a simulation. Eq. (3) implies that FCIQMC is an example of MCMC, the implications of which we shall discuss in Sec. III.

The correlation energy can also be found by

$$E_{\text{Proj}} = \frac{\langle D_0 | \hat{H} e^{-\hat{H}\tau} | D_0 \rangle}{\langle D_0 | e^{-\hat{H}\tau} | D_0 \rangle} = \frac{\sum_{i \neq 0} H_{0i} n_i}{n_0}, \quad (4)$$

where the trial state, $|D_0\rangle$, is typically the Hartree–Fock determinant. The variance of the projected estimator is generally smaller than that of the shift and can be reduced further by a multi-determinant trial function.¹²

Both estimators are serially correlated as the state of the simulation at one timestep is heavily dependent on the state at the previous timestep. We use an automated iterative blocking algorithm^{15–18} to accurately estimate the stochastic error in all FCIQMC calculations presented in this paper.

III. STOCHASTIC MATRICES AND FCIQMC

A. General Markov chain Monte Carlo theory

A stochastic process is a discrete time Markov chain if the probability of transitioning from one state to another (in one discrete timestep) depends only upon the current state.^{19,20} The probability of a given set of psips, $\{n_i\}$, producing another set of psips, $\{n'_i\}$, at the next timestep depends only on $\{n_i\}$ and the value of the shift, which depends upon the total number of psips at a given shift update. We can therefore describe FCIQMC as a Markov chain taking one step every A timesteps. In order to simplify the mathematical details, we shall henceforth assume that the simulation takes one timestep between shift updates (i.e., $A = 1$).

We shall denote the Markov states of the simulation using indices $\alpha, \beta \dots$ and Slater determinants using indices i, j, \dots . The stochastic matrix, Γ , consists of elements $\Gamma_{\alpha\beta}$ which gives the probability that the system transitions from state α to state β in one step in the Markov chain and is in general not symmetric.

We can infer some properties of the eigenvectors and eigenvalues of the stochastic matrix as Γ is non-negative and $\sum_{\beta} \Gamma_{\alpha\beta} = 1$ as a FCIQMC calculation must transition from one state to another or remain in the same state. From this, there must exist one or more left eigenvectors, γ_{α} , satisfying

$$\sum_{\alpha} \gamma_{\alpha} \Gamma_{\alpha\beta} = \gamma_{\beta}, \quad (5)$$

where γ_{α} gives the probability that the Markov chain will be in state α if the chain is in equilibrium. The Perron–Frobenius theorem proves that the Γ must have one or more such Perron–Frobenius eigenvectors with a unit eigenvalue and all other eigenvalues must be smaller. The Perron–Frobenius eigenvector is unique, and the chain will converge towards this distribution if (i) all the states are aperiodic, i.e., $\Gamma_{\alpha\beta}^N > 0$ for all values of large N ; (ii) every state can be reached from every other state. For function $f(\alpha)$ defined for all possible Markov states, its expectation value is

$$\mu_f = \langle f_i \rangle_t = \sum_{\alpha} f(\alpha) \gamma_{\alpha}. \quad (6)$$

The Perron–Frobenius eigenvector specifies the distribution of an ensemble of independent Markov chains taking a single step, and by computing it, we may find expectation values of interest in this system.

B. The FCIQMC Markov chain

A state α in FCIQMC is represented by the signed number of psips on each determinant (n_a, n_b, \dots) ,

$$\alpha := (n_a^{(\alpha)}, n_b^{(\alpha)}, \dots). \quad (7)$$

The shift S is not an independent variable as it is simply a function of the total number of psips. The FCIQMC chain is in an absorbing state (i.e., the probability of leaving the state is zero) when there are no psips on any of the determinants, as all events which change the psip population require an existing non-zero population. As the shift is initially held at a constant

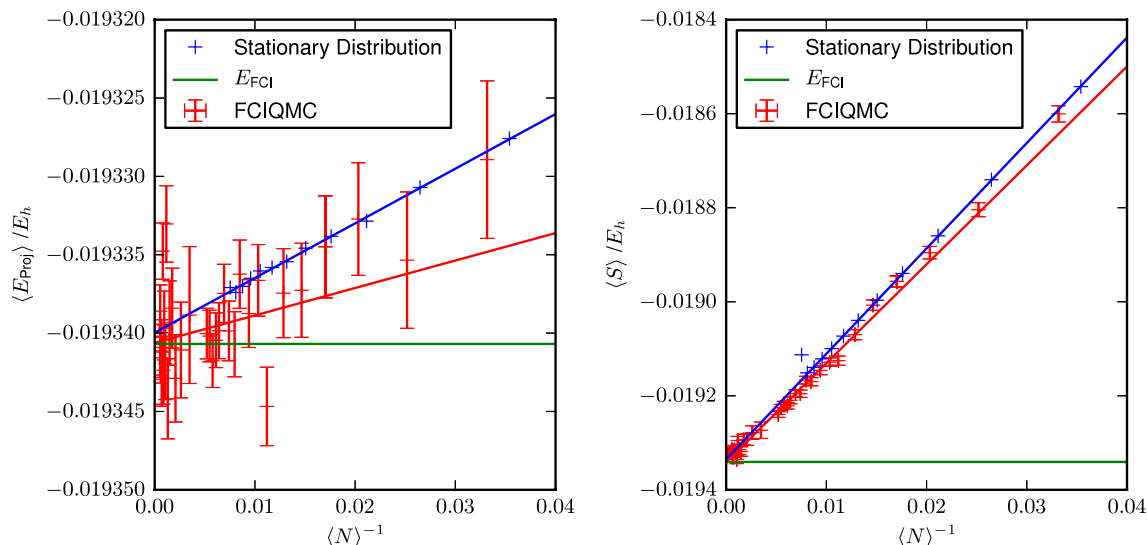


FIG. 1. The energy estimators as a function of $1/\langle N \rangle$ for H_2 (STO-3G basis, internuclear separation 0.7122 \AA). The exact means were calculated from the means of the stationary distributions. Only states with up to 150 psips on each determinant were included in the transition matrix calculations. The FCIQMC estimates were calculated from a single chain and the error estimated by blocking.¹⁵ The bias in both estimates of the correlation energy decays with the inverse of the average number of psips. Linear fits were performed with numpy.²⁴ Errors were weighted in the fits for the FCIQMC data using the sum of the variance of $1/\langle N \rangle$ and variance of the energy estimator. The state with the smallest $1/\langle N \rangle$ was removed from the fit as the energy is too large to fit a linear slope. This is caused by the stochastic matrix being truncated at the state with 150 psips on both determinants; states after this truncation have become important at this point.

value during the equilibration phase, the Markov chain changes after the shift is allowed to vary.

The estimators of interest in FCIQMC are the shift and the numerator and denominator of the projected energy. Defining $N(\alpha) = \sum_i |n_i^\alpha|$,

$$S(\alpha) = S(\tau = 0) - \xi \log \frac{N(\alpha)}{N_s}, \quad (8)$$

$$E_{\text{Numer}}(\alpha) = \sum_{i \neq 0} H_{0i} n_i^{(\alpha)}, \quad (9)$$

$$N_{\text{Denom}}(\alpha) = n_0^{(\alpha)}. \quad (10)$$

The projected energy can hence be evaluated using $\langle E_{\text{Proj}} \rangle = \langle E_{\text{Numer}} \rangle / \langle N_{\text{Denom}} \rangle$.

Computing the stochastic matrix for an arbitrarily large system of determinants is computationally infeasible as the space scales as the power of the number of determinants, so we restrict ourselves to the simplest possible (interesting) system: two determinants, a and b . In this case the change on one determinant is independent of the change on the other, the elements of the stochastic matrix element are given by

$$\Gamma_{\alpha,\beta} = p_{cn_a,n'_a} p_{cn_b,n'_b}, \quad (11)$$

where p_{cn_a,n'_a} is the probability that the population on a changes from n_a to n'_a (Appendix).

We have constructed $\Gamma_{\alpha,\beta}$ for some simple systems and determined (by direct²¹ or iterative²² diagonalisation) the stationary distributions γ_α corresponding to the Perron-Frobenius eigenvector. The supplemental material²³ contains further examples of using the stationary distribution to examine the behaviour of the ensemble of FCIQMC states.

IV. POPULATION CONTROL BIAS

In order to achieve a finite population in a simulation, we must resort to population control by introducing a shift

which itself is dependent upon the current total population. However, this process introduces a feedback into the propagator and hence a systematic bias.²⁵ Random fluctuations causing the population to increase (i.e., the psip distribution enters a low energy region of phase space) or decrease (higher energy region of phase space) are moderated by a corresponding decrease or increase in the shift. Both actions lead to an increase in the time-averaged energy estimators.

In DMC, this bias is known^{25,26} to scale as $\langle N \rangle^{-1}$, though we know of no previous investigations of population control bias in FCIQMC. We suspect that the effect has most likely been obscured by the stochastic error in all previous FCIQMC studies. With the aid of exact energy estimators from the transition matrix, we are now able to investigate the magnitude of any bias present. We feel it is important to understand where population control bias is likely to cause a problem if a small stochastic error is desired. In addition to the energy estimators from the transition matrix, we shall also investigate them from single chains via blocking analyses of single FCIQMC calculations to compare both methods and use them to quantify the factors controlling population control bias.

A. H_2 in a STO-3G basis set

For different values of N_s , transition matrix and single chain calculations were performed on H_2 in a STO-3G basis at the equilibrium geometry of 0.7122 \AA ,²⁷ and the energy estimators evaluated. Fig. 1 shows the bias of the projected energy and shift estimators decreases with $1/\langle N \rangle$. Though the single-chain calculations have relatively large stochastic errors, a similar bias in the energy and decay is also notable, and there is good agreement between the single chain and transition matrix results. The fits for the transition matrix calculations, however, do not exactly intercept the y -axis at the correlation energy. The worst extrapolation is for the projected energy, which

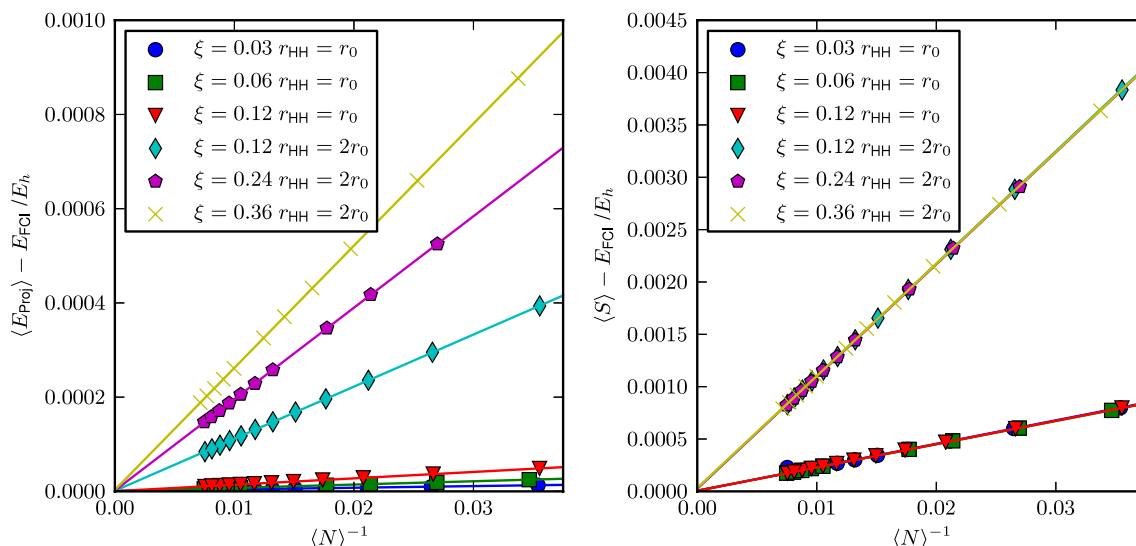


FIG. 2. The energy estimators calculation from the means of the stationary distribution as a function of $1/\langle N \rangle$ for H_2 (STO-3G basis, internuclear separation $r_{\text{HH}} = r_0$ and $2r_0$, where $r_0 = 0.7122 \text{ \AA}$) for different values of ξ , the shift damping parameter. Only states with up to 150 psips on each determinant were included in the transition matrix calculations. The bias in the projected energy can be reduced by decreasing ξ , whereas the bias in the shift remains the same. Fits with $1/\langle N \rangle$ were performed with `numpy`.²⁴ Again we remove the state with the smallest $1/\langle N \rangle$ for $\xi = 0.03$ for $r_{\text{HH}} = r_0$.

disagrees by $2.3 \pm 0.5 \mu E_h$. It is difficult to tell if this is caused by a loss of numerical precision, truncation in the transition matrix calculations, or if there are higher order effects with small $\langle N \rangle$.

The prefactor in the $1/N$ scaling of the bias in the projected energy is affected by ξ , and damping less hard, i.e., decreasing ξ , reduces the prefactor (Fig. 2). Population control bias also appears to be made worse in strongly correlated systems. Fig. 2 shows population control bias as a function of ξ for both H_2 in a STO-3G basis set at bond lengths of 0.7122 \AA and 1.4244 \AA .²⁸ We may explain this by reviewing DMC, where, in the limit of a perfect trial function, there are no branching processes and thus no population control bias. Equivalently in FCIQMC, if there is no spawning there can be no population control bias. Although true only in the limit in which the Hilbert space is the set of eigenvectors of the Hamiltonian, this indicates that in the weakly correlated limit, there will be less population control bias, exactly as we observe.

B. The neon atom in a cc-pVDZ basis

Population control bias is also potentially a significant source of a systematic error in systems which are large enough not to be trivially soluble (rendering transition matrix calculations computationally infeasible). We now turn to the neon atom in a cc-pVDZ basis^{29,30} which has a Hilbert space of 50 000 determinants. This is small enough such that it is straightforward to compute the FCI energy via iterative diagonalisation but large enough such that most determinants have a small contribution to the wavefunction. It was necessary to oversample the Hilbert space in H_2 (i.e., more psips than the number of determinants), whereas FCIQMC calculations in this neon system are stable with a significant undersampling of the space. We shall investigate the effect of population control bias in this regime.

Changing the population control parameters affects both estimators of the correlation energy in the same way as H_2 .

Fig. 3(a) shows the projected energy decaying towards the FCI energy as ξ decreases until the estimator of the energy becomes within error bars. Fig. 3(b) shows the bias in the projected energy decaying as $1/\langle N \rangle$. This intercepts the y -axis at $-0.192 106 6(12) E_h$ which is within errors of the FCI energy of $-0.192 105 578 E_h$; suggesting, we converge to the exact ground state as expected. The population control bias is however significant and with about 10 000 psips is about $20 \mu E_h$.

As population control bias is clearly a problem, we sought to find a simple indicator of its magnitude in a calculation. An investigation into the relationship between the variance of the shift and the magnitude of the population control bias did not yield any simple relationship. With hindsight, a consideration of the causal relationship between the two makes it likely that not only the extent but also the speed of variation of the shift is important. As such values are considerably more difficult to calculate, we will leave investigation of this connection to a future publication.

Instead, we have adopted a method used in DMC. The population bias can be both quantified and reduced by a reweighting technique based upon the history of the shift.²⁵ The contribution at a given time, τ , to the numerator and denominator of the projected energy is weighted by taking into account the shift of the preceding W iterations. For S_m , denoting the shift m iterations previously, the weight is given by

$$w(\tau, W) = \prod_{m=1}^W e^{-\delta\tau(S_m - \langle S \rangle)}. \quad (12)$$

The reweighting is implemented as a post-processing step on the output of a calculation. The population control bias is effectively removed for sufficiently large W at the cost of increasing the stochastic error and, as can be seen in Fig. 4, the residual bias is of the order of the stochastic error bars.³¹

The value of W (≈ 250) required for this procedure to converge is of the order of the serial correlation length. We




FIG. 3. Projected energy estimates from FCIQMC calculations on the Ne atom (cc-pVDZ basis set) as a function of (left) the shift damping parameter, ξ , (calculation details: $\delta\tau = 0.005$, 20 000 000 iterations where $\langle N \rangle$ has the value 10 000) and (right) as a function of $1/\langle N \rangle$, where N is the number of psips ($A = 10$, $\xi = 20$, $\delta\tau = 0.005$, 200 000 000 iterations or 72 h on 12 cores whichever was shorter). Linear fits were performed with numpy.²⁴

note that it is not possible to apply this method to the Markov chain approach as the dependence of expectation values on calculation history makes the process non-Markovian.

V. DISCUSSION

To summarise: we have demonstrated that FCIQMC is an example of Markov chain Monte Carlo and computed the stochastic matrix for a two determinant system. Even though a two determinant system is the simplest non-trivial system, it still contains some of the inherent features of FCIQMC including population control bias. A two determinant system cannot have a sign problem unless the timestep is greater than the critical point. It would be interesting to extend these ideas

to investigate the sign problem using a three determinant system, though the stochastic matrix may be inaccessible due to its scaling with the size of the Hilbert space.

Recently, Petruzielo *et al.* proposed to use floating point numbers to represent the population of psips on a determinant.¹² This adaption results in an uncountably infinite state space of the Markov chain. They also proposed to partition the determinant space into deterministic and stochastic subspaces, where the action of the Hamiltonian in the deterministic subspaces is applied exactly using sparse matrix multiplication, and the action in the stochastic subspaces is sampled in the same way as in FCIQMC. Using floating point numbers as walker weights as well as a multideterminantal trial function significantly reduces the prefactor of the $1/\langle N \rangle$ scaling in population control bias.

The population control algorithm in DMC, as recommended in Ref. 25, is slightly different from that used in FCIQMC: the shift is updated from the “best current estimate” of the energy rather than from the previous value of the shift. Using this population control algorithm would render FCIQMC non-Markovian. Nonetheless, we could use the stochastic matrix technique presented here to calculate the probability distribution of the shift in the limit of convergence of the projected energy. It would be interesting to investigate if this is a better method of population control for FCIQMC.

Using the population control approach given in Ref. 5 (i.e., using Eq. (3) with γ set in the region 0.01 to 0.05) may introduce population control bias, due to the factor of $\frac{1}{\delta\tau}$, if the timestep needed to converge a calculation needs to be small. This means population control bias is likely to be more of a problem for calculations which require smaller timesteps, such as strongly correlated systems, or calculations using coupled cluster Monte Carlo.¹¹

We also note that converging FCIQMC calculations to μE_h accuracy has previously been attempted.³² In this regime, the population control bias could potentially become similar in magnitude to the stochastic error.




FIG. 4. Reweighted projected energy estimate from FCIQMC calculations on the Ne atom (cc-pVDZ basis set) for different $\langle N \rangle$ as a function of W , the number of iterations reweighted over. Four independent FCIQMC calculations were performed for each $\langle N \rangle$, and each point shows the mean of the reweighted $\langle E_{\text{proj}} \rangle$ for the four runs.

We recommend that one reweights the projected energy estimator, as suggested in Ref. 25, as it does not involve multiple expensive runs. If a large enough population is used, the resultant estimate of the energy would be unbiased albeit with a larger stochastic error (though this has appeared negligibly larger in our tests). Alternatively, one should use a large population of psips and set ξ to be as small as possible, such that the number of psips does not drop below the system-dependent critical population. Doubling the number of psips in a simulation increases the equilibration time and possibly also the memory requirements. It is also important to perform enough steps to get an accurate estimate of the error. In choosing an appropriate value of ξ , there is a compromise to be made; it is tempting to increase ξ because it reduces the fluctuations in the total number of psips and, for larger systems, this can reduce the maximum amount of memory used during the calculation. However, too large a ξ will cause population control bias to become significant.

In conclusion, we caution users of FCIQMC and related methods to be aware that population control can introduce a significant bias in calculated energies. We recommend that post-processing reweighting is used to quantify its magnitude and the psip population and damping parameters be modified as suggested in this paper if needed.

ACKNOWLEDGMENTS

The authors thank C. J. Umrigar for several enlightening discussions, facilitated by the unique setting of The Towler Institute. Calculations were performed using the Imperial College High Performance Computing Service³³ and figures were plotted using matplotlib.³⁴ J.S.S. acknowledges the research environment provided by the Thomas Young Centre under Grant No. TYC-101. W.A.V. is grateful to EPSRC for a studentship and AJWT thanks Imperial College for a Junior Research Fellowship and the Royal Society for a University Research Fellowship.

APPENDIX: TRANSITION PROBABILITIES FOR A TWO DETERMINANT SYSTEM

Consider two determinants, a and b , with states α and β representing two states with a given signed number of psips on each determinant,

$$\begin{aligned}\alpha &= (n_a, n_b), \\ \beta &= (n'_a, n'_b).\end{aligned}\quad (\text{A1})$$

Each psip independently attempts to spawn and die every timestep. The probability that n psips succeed out of N attempts is given by the probability mass function of the binomial distribution, $B(n, N, p) = \binom{N}{n} p^n (1-p)^{N-n}$, where p is the probability of one psip spawning or dying independently and can be obtained from the FCIQMC algorithm.⁵ The change on determinant a is³⁵

$$\begin{aligned}n'_a - n_a &= -\text{sgn}(H_{ba})\text{sgn}(n_b)n_{sa} \\ &\quad - \text{sgn}(H_{aa} - S)\text{sgn}(n_a)n_{da},\end{aligned}\quad (\text{A2})$$

where n_{sa} (n_{da}) is the number of psips spawning onto (dying on) a . As the spawning and death events on each determinant are independent, the probability $p_{c n_a, n'_a}$ that the number of psips on a changes from n_a to n'_a via any possible combination of spawning and death is given by

$$p_{c n_a, n'_a} = \sum_{n_{sa}} B(n_{sa}, n_b, P_s(a|b)) B(n_{da}, n_a, P_d(a)), \quad (\text{A3})$$

where $P_s(a|b)$ is the probability for spawning to a from b and $P_d(a)$ is the probability of death of a .

- ¹C. C. J. Roothaan, *Rev. Mod. Phys.* **23**, 69 (1951).
- ²R. Polly, H.-J. Werner, F. R. Manby, and P. J. Knowles, *Mol. Phys.* **102**, 2311 (2004).
- ³A. Halkier, T. Helgaker, P. Jørgensen, W. Klopper, H. Koch, J. Olsen, and A. K. Wilson, *Chem. Phys. Lett.* **286**, 243 (1998).
- ⁴D. E. Woon and T. H. Dunning, *J. Chem. Phys.* **101**, 8877 (1994).
- ⁵G. H. Booth, A. J. W. Thom, and A. Alavi, *J. Chem. Phys.* **131**, 054106 (2009).
- ⁶G. H. Booth, S. D. Smart, and A. Alavi, *Mol. Phys.* **112**, 1855 (2014).
- ⁷J. S. Spencer, N. S. Blunt, and W. M. C. Foulkes, *J. Chem. Phys.* **136**, 054110 (2012).
- ⁸D. Cleland, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **132**, 41103 (2010).
- ⁹C. Daday, S. Smart, G. H. Booth, A. Alavi, and C. Filippi, *J. Chem. Theory Comput.* **8**, 4441 (2012).
- ¹⁰J. J. Shepherd, G. Booth, A. Grüneis, and A. Alavi, *Phys. Rev. B* **85**, 081103 (2012).
- ¹¹A. J. W. Thom, *Phys. Rev. Lett.* **105**, 263004 (2010).
- ¹²F. R. Petruzielo, A. A. Holmes, H. J. Changlani, M. P. Nightingale, and C. J. Umrigar, *Phys. Rev. Lett.* **109**, 230201 (2012).
- ¹³J. B. Anderson, *J. Chem. Phys.* **63**, 1499 (1975).
- ¹⁴N. Trivedi and D. M. Ceperley, *Phys. Rev. B* **41**, 4552 (1990).
- ¹⁵H. Flyvbjerg and H. G. Petersen, *J. Chem. Phys.* **91**, 461 (1989).
- ¹⁶U. Wolff, *Comput. Phys. Commun.* **156**, 143 (2004).
- ¹⁷R. M. Lee, G. J. Conduit, N. Nemeč, P. Lopez Rios, and N. D. Drummond, *Phys. Rev. E* **83**, 066706 (2011).
- ¹⁸See <https://github.com/jsspencer/pyblock> for code.
- ¹⁹A. Sokal, *Functional Integration*, NATO ASI Series (Springer, US, 1997), Vol. 361, p. 131.
- ²⁰J. Kemeny and J. Snell, *Finite Markov Chains: With a New Appendix "Generalization of a Fundamental Matrix,"* Undergraduate Texts in Mathematics (Springer, 1976).
- ²¹E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. (Society for Industrial and Applied Mathematics, Philadelphia, 1999).
- ²²R. B. Lehoucq, D. C. Sorensen, and C. Yang, *Arpack User's Guide: Solution of Large Scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods* (Society for Industrial and Applied Mathematics, 1998).
- ²³See supplementary material at <http://dx.doi.org/10.1063/1.4913644> for investigations into the behaviour of the stationary distribution as a function of population and population control parameters.
- ²⁴T. E. Oliphant, *Comput. Sci. Eng.* **9**, 10 (2007).
- ²⁵C. J. Umrigar, M. P. Nightingale, and K. J. Runge, *J. Chem. Phys.* **99**, 2865 (1993).
- ²⁶N. Cerf and O. C. Martin, *Phys. Rev. E* **51**, 3679 (1995).
- ²⁷Hartree-Fock energy: $-1.1175058843E_h$.
- ²⁸Hartree-Fock energy: $-0.9338980552E_h$.
- ²⁹T. H. Dunning, *J. Chem. Phys.* **90**, 1007 (1989).
- ³⁰Hartree-Fock energy: $-128.488775516E_h$.
- ³¹FCIQMC projects out the ground state using $(1 - (\hat{H} - S))\delta\tau$, whereas $\exp(-\delta\tau(\hat{H} - S))$ is used in DMC. This means weighting with Eq. (12) is an approximation. The error introduced is second order in $\delta\tau$ (this can be shown using a Taylor expansion). The reweighting requires the projected energy and population at every individual time step.
- ³²S. Sharma, T. Yanai, G. H. Booth, C. J. Umrigar, and G. K.-L. Chan, *J. Chem. Phys.* **140**, 104112 (2014).
- ³³See <http://www.imperial.ac.uk/ict/services/teachingandresearchservices/highperformancecomputing>.
- ³⁴J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- ³⁵We assume that the independent psip probabilities are less than or equal to one for simplicity; the generalisation is straightforward.