Deuterium Hugoniot: Pitfalls of thermodynamic sampling beyond density functional theory

Raymond C. Clay, III, Michael P. Desjarlais, and Luke Shulenburger

Phys. Rev. B **100**, 075103 — Published 1 August 2019

DOI: 10.1103/PhysRevB.100.075103

The Deuterium Hugoniot: Pitfalls of Beyond-DFT Thermodynamic Sampling

Raymond C. Clay, III,[1,*] Michael P. Desjarlais,[1] and Luke Shulenburger[1]

[1] *Pulsed Power Sciences Center, Sandia National Laboratories, Albuquerque, New Mexico*

(Dated: July 19, 2019)

Outstanding problems in the high pressure phase diagram of hydrogen have demonstrated the need for more accurate *ab initio* methods for thermodynamic sampling. One promising method that has been deployed extensively above 100GPa is coupled electron-ion Monte Carlo (CEIMC), which treats the electronic structure with quantum Monte Carlo (QMC). However, CEIMC predictions of the deuterium principal Hugoniot disagree significantly with experiment, overshooting the experimentally determined peak compression density by 7% and lower temperature gas gun data by well over 20%. By deriving an equation relating the predicted Hugoniot density to underlying equation of state errors, we show that QMC and many-body methods can easily spoil the error cancellation properties inherent in the Rankine-Hugoniot relation, and very likely suffer from error addition. By cross validating QMC based on systematically improvable trial functions against post-Hartree-Fock many-body methods, we find that these methods introduce errors of the right sign and magnitude to account for much of the observed discrepancy between CEIMC and experiment. We stress that this is not just a CEIMC problem, but that thermodynamic sampling based on other many-body methods is likely to experience similar difficulties.

PACS numbers: 67.80.ff,63.20.dk,62.50.-p,64.70.kt

Quantum Monte Carlo methods (QMC)[1,2], which work with the exact Hamiltonian of a physical system, have been increasingly used to elucidate outstanding experimental controversies in the equations of state of hydrogen and its isotopes, especially above 100 GPa[3–8]. This is because the approximations underlying density functional theory (DFT), the traditional workhorse for *ab initio* computations of material properties at extreme conditions, have been shown to introduce wild variability into the predicted locations of molecular-to-atomic and insulator-to-metal phase boundaries, both at high and low temperatures[9–11]. Accurate characterization of these transitions is important, not only for the interpretation of recent claims regarding the synthesis of solid metallic hydrogen[3], but also for modelling Jovian planets, brown dwarfs, and inertial confinement fusion implosions[12].

Coupled electron-ion Monte Carlo (CEIMC) is one of the most ambitious and well controlled QMC algorithms fielded in the study of hydrogen to date[13–18]. Rather than use QMC indirectly to validate or construct DFT functionals[19,20], or partially to compute specific contributions to the low temperature free energy[21–23], CEIMC uses QMC to sample the Born-Oppenheimer potential energy surface directly with either classical or quantum nuclei[13–15]. It has been fielded in both the solid[24] and liquid phases[14,25,26] of hydrogen and deuterium, yielding qualitative agreement with DFT but with ostensibly smaller uncontrolled errors.

Unfortunately, CEIMC predictions of the principal Hugoniot of deuterium differ significantly from experiment[27], as can be seen in Figure 1. The density at peak compression is roughly 7% larger than the best experimental data[28] and at pressures near the established gas-gun data[29], the discrepancy widens to over 20%. This is troubling, since many of the systematic errors that *could* affect CEIMC calculations are expected to be greatly reduced for the molecular phase. It has been argued that the discrepancy could originate within the experiments[27], but due to the hard wrought[30,31] agreement between gas-gun[29], pulsed-power[28,32], converging explosive shock[33], and laser platforms[34–36], we feel this discrepancy is more significant than discrepancies elsewhere in the phase diagram. Moreover, density functional theory predicts a principal Hugoniot in excellent agreement with experiment and is surprisingly robust to the choice of functional[37–40].

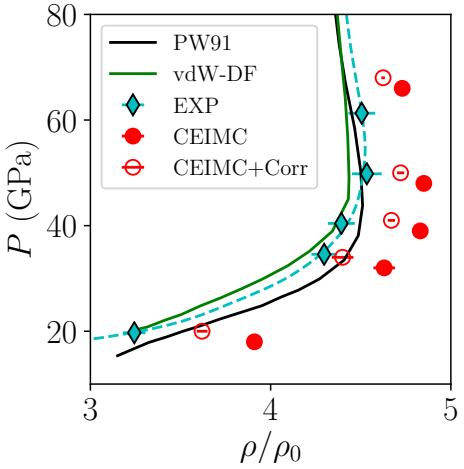

FIG. 1. Pressure vs. density compression of the principal deuterium Hugoniot. Diamonds are the experimental data from Refs.[28,29], filled circles are the published CEIMC data[27], and hollow circles are the corrected CEIMC Hugoniot points using methods in this paper. PW91 and vdW-DF Hugoniots from Ref.[28] included for comparison.

In this work, by analytically investigating the sensitivity of the predicted Hugoniot to underlying equation of state errors, we demonstrate that the Rankine-Hugoniot

relation[41] supports a rich structure for error cancellation. We show that whereas DFT can effectively utilize this structure to produce an accurate principal Hugoniot, many-body methods spoil this error cancellation and very likely suffer from error *addition*. Using diffusion Monte Carlo (DMC) based on systematically improvable wave functions and cross validating against traditional many-body quantum chemistry methods in periodic boundary conditions, we can show *ab initio* that in the absence of all other sources of error, fixed-node errors entering into the evaluation of the Rankine-Hugoniot relation have the right sign and magnitude to account for much of the discrepancy between CEIMC and established gas gun data. While we primarily focus on CEIMC, we stress that the issues highlighted in this work are of general concern to any beyond-DFT method of thermodynamic sampling.

This paper is organized as follows. In Section I, we present a formula for the dependence of the Hugoniot curve on the underlying equation of state errors. In Section II, we discuss in detail the parameters and problem specific procedures used in all density functional theory, quantum chemistry, and quantum Monte Carlo methods. In Section III, we present the results of our error analysis for 8 atom cells, and then discuss prospects for generalizing to the thermodynamic limit with limited calculations on 128 atom cells. This progression is used to estimate likely errors in many-body methods near the established gas gun data, the impact of which is demonstrated by calculating a D_2 principal Hugoniot based on the published CEIMC Hugoniot when the demonstrated errors are taken into account and corrected for. We conclude Section III by investigating the role of error cancellation in density functional theory calculations. Everything that the authors deem especially tedious or speculative is left to the supplemental materials: the derivation of the error formula, how the principal Hugoniot error correction was done, the potential (lack of) impact of many-body finite-size effects on the principal Hugoniot, and how the conclusions of this paper might change as we consider thermodynamic conditions near peak compression.

I. ERROR FORMULA

In order to understand how errors in a computational method affect the determination of the Hugoniot curve, we first consider the locus of thermodynamic states (ρ, T) achieved by planar shock wave compression of a material at an initial condition of (ρ_0, T_0). Because mass, energy and momentum are conserved across the shock front, these states satisfy the Rankine-Hugoniot relations:

$$E(\rho,T)-E(\rho_0,T_0) = \frac{1}{2}(P(\rho,T)+P(\rho_0,T_0))\left(\rho_0^{-1}-\rho^{-1}\right) \tag{1}$$

where, $E(\rho,T)$, $P(\rho,T)$, ρ, T, are energy, pressure, density, and temperature, respectively, and "0" indicates the initial state.

Now consider computing the Hugoniot curve using a perturbed equation of state $E^*(\rho,T)$, $P^*(\rho,T)$. Errors in observable $\mathcal{A}$ are defined as $\Delta\mathcal{A}^* = \mathcal{A}^*(\rho,T) - \mathcal{A}(\rho,T)$. The collection of points (ρ^*, T^*) defining the Hugoniot of $E^*(\rho,T)$, $P^*(\rho,T)$ will usually be different from those calculated using Eq. 1. However, we can map points on the reference Hugoniot to those on the perturbed Hugoniot if we know both $\Delta E^*(\rho,T)$ and $\Delta P^*(\rho,T)$. This mapping is quite complicated in general, but if we consider solutions of the Rankine-Hugoniot equation along a specific isotherm given by T, and perturbatively expand the Rankine-Hugoniot relations in terms of the equation of state errors, we can solve for ρ^* to leading order in $\Delta\rho^*$, ΔE^* and ΔP^*. While the derivation of this mapping is not difficult, it is algebraically intensive and so we leave its derivation to the supplemental information[42]. The mapping is:

$$\frac{\rho_0}{\rho^*} = \frac{\rho_0}{\rho} - \left(1-\frac{\rho_0}{\rho}\right)\left[\frac{\Delta E^*(\rho,T)-\Delta E^*(\rho_0,T_0)}{E(\rho,T)-E(\rho_0,T_0)} - \frac{\Delta P^*(\rho,T)+\Delta P^*(\rho_0,T_0)}{P(\rho,T)+P(\rho_0,T_0)}\right] C(\rho,T)^{-1} + \mathcal{O}(\Delta^2) \tag{2}$$

$C(\rho,T)$ is a numerical factor that depends on both the equation of state and its errors. The form of $C(\rho,T)$ is given and explained in the supplemental information, but for this discussion, we estimate it to be between 5 and 8 for all densities and temperatures considered.

We see immediately multiple pathways for error cancellation: between E and E_0, between P and P_0, and between the combination of $E-E_0$ and $P+P_0$. As long as the errors have the same sign, a theory will experience "error cancellation". Conversely, if they have opposite signs, a theory will experience "error addition". This is important, because while an error cancellation scenario can support very large errors without changing the calculated Hugoniot density ρ^* significantly, any error made in an error addition regime will directly change the predicted ρ^*. The goal of this paper from here on out is to evaluate ΔE^* and ΔP^* for several major classes of electronic structure methods, and see how errors arising in these calculations manifest in Hugoniot estimation.

II. COMPUTATIONAL DETAILS

Now that we have an analytic form for relating the equation of state errors to a predicted Hugoniot curve

(ρ^*, T^*), we investigate how this applies to Hugoniot calculations using DFT and beyond-DFT *ab initio* methods. The goal is to attempt to quantify ΔE^* and ΔP^* for various methods. We do this by first generating ionic snapshots that are representative of Hugoniot states of interest, and then computing reference energies and pressures on these test configurations using both systematically improvable QMC and high-level quantum chemistry methods. The reference ground state energies are found by "consensus"–showing that QMC and high-level quantum chemistry methods produce similar ground state energies despite their vastly different approximations, whereas the error in other quantities is estimated through detailed comparisons between methods.

This section will focus on three major areas, the first being the generation of the test sets used to to mock thermodynamic states relevant for the deuterium Hugoniot. Secondly, we will discuss the practical details associated with running the DFT, QMC, and quantum chemistry methods used in this work. Lastly, we motivate and discusss our use of QMC based "zero-variance" extrapolations used in this work.

A. Test Set Generation and DFT

To generate the test configurations, we performed quantum molecular dynamics simulations in VASP[43–46] using the PBE functional[47]. We used a plane wave cutoff of 1200eV and an all-electron PBE PAW pseudopotential. We chose $\rho_0 = 0.167 g/cc$ at $T = 22K$ as our reference point and sampled densities ranging between $\rho/\rho_0 = 3.000$ and $\rho/\rho_0 = 4.250$ at 4000K. While we also performed a similar sampling between $\rho/\rho_0 = 3.500$ and $\rho/\rho_0 = 5.000$ at T=10000K, this is relegated to the supplemental information.

We used VASP's Langevin thermostat with the damping parameters set to $\gamma = 50 ps^{-1}$ and a time step of $\tau = 0.3 fs$. At each density, a single configuration was drawn at random after an equilibration time of $0.9 ps$.

For the DFT electronic structure, we used a plane wave cutoff of 1200eV and the all-electron VASP PBE PAW Pseudopotential. For the 8 atom cells, we used a Monkhorst-Pack grid of 8x8x8, wheras for the 128 atom cells, we used a single k-point at the Baldereschi mean value point of a cubic cell.

B. QMC Trial Wave Function Generation

The QMC workflow used in this work was somewhat nontrivial, so we take some time in this section to elaborate on the construction of various trial wave functions. We will also discuss our methods of calculating the QMC pressure of our configurations, and techniques we used to mitigate the mixed-estimator bias. Details of the optimization and diffusion Monte Carlo runs will be mentioned in passing. All QMC calculations were performed with QMCPACK[48].

The most basic trial wave function commonly used in QMC is the "Slater-Jastrow" wave function, henceforth designated "SJ", given by the following:

$$\Psi_{SJ} = e^{-J} \det(M^\uparrow) \det(M^\downarrow) \tag{3}$$

$M^\uparrow$ and $M^\downarrow$ are Slater determinants for the up and down electrons, and are defined as $[M^{\uparrow(\downarrow)}]_{ij} = \phi_j(\mathbf{r}_i^{\uparrow(\downarrow)})$. The $\phi_i(\mathbf{r})$ functions are single particle orbitals (SPO's) taken from the set $\{\phi_0, \phi_1, \ldots, \phi_{N_{orbs}}\}$, where N_{orbs} is the number of SPO's in our set. J is the symmetric Jastrow factor, which we will describe in detail later in this section.

This ansatz can be slightly modified to yield the Slater-Jastrow backflow wave function, which we will abbreviate with "BF". It's form is identical to the Slater-Jastrow, but we replace $M^{\uparrow(\downarrow)}$ with $[\tilde{M}^{\uparrow(\downarrow)}]_{ij} = \phi_j(\mathbf{q}_i)$. $\mathbf{q}_i^\uparrow$ is a "quasiparticle" coordinate, defined as:

$$\mathbf{q}_i^\uparrow = \sum_I \eta_{eH}(|\mathbf{r}_i^\uparrow - \mathbf{r}_I|)(\mathbf{r}_i^\uparrow - \mathbf{r}_I) \tag{4}$$

$$+ \sum_{i \neq j} \eta_{\uparrow\uparrow}(|\mathbf{r}_i^\uparrow - \mathbf{r}_j^\uparrow|)(\mathbf{r}_i^\uparrow - \mathbf{r}_j^\uparrow) \tag{5}$$

$$+ \sum_j \eta_{\uparrow\downarrow}(|\mathbf{r}_i^\uparrow - \mathbf{r}_j^\downarrow|)(\mathbf{r}_i^\uparrow - \mathbf{r}_j^\downarrow) \tag{6}$$

$\mathbf{q}_i^\downarrow$ is defined analogously. The advantage of backflow is that by optimizing $\eta(r)$, we can effectively perturb the nodal surface in a way that lowers the total VMC and DMC energy.

Lastly, we have the multi-Slater determinant wave functions, henceforth abbreviated as "MSD":

$$\Psi_{MSD} = e^{-J} \sum_{i=0}^{N_{CSF}} \alpha_k \Phi_k \tag{7}$$

While the number of possible antisymmetric functions Φ_k grows combinatorically with system size, changing the number of included terms gives us a natural knob that we can use to systematically improve the quality of the trial function. Just to remind the reader of our naming convention, "SJ-PBE" would be a Slater-Jastrow wave function built from PBE orbitals, "MSD-HF" a multideterminant with Hartree-Fock orbitals, and so on.

While we could use ordinary Slater determinants for Φ_k, we chose to use configuration state functions (CSF's), which are spin and symmetry adapted linear combinations of Slater determinants constructed to be eigenstates of $\hat{S}^2$.

$$\Phi_k = \sum_i c_k^i \det(M_i^\uparrow) \det(M_i^\downarrow) \tag{8}$$

The coefficients c_k^i are totally determined from symmetry. The advantage to using CSF's is that we can restrict the trial wave function optimization to only those wave

functions that preserve spin singlet symmetry, which should be the symmetry of the ground-state for these hydrogen systems. This amounts to a 2-3x reduction in the number of optimization parameters as opposed to if we used arbitrary linear combinations of determinants.

Now that we have specified the high-level forms for all trial wave functions used in this work, we take some time to explicitly walk through the specific forms used and the particular means of generating our trial wave functions.

For all three ansätze, we used a Jastrow factor with the following form:

$$J = \sum_{i,I} u_{eH}(|\mathbf{r}_i - \mathbf{r}_I|) \tag{9}$$

$$+ \sum_{i<j} u_{\uparrow\uparrow}(|\mathbf{r}_i^{\uparrow} - \mathbf{r}_j^{\uparrow}|) + u_{\downarrow\downarrow}(|\mathbf{r}_i^{\downarrow} - \mathbf{r}_j^{\downarrow}|) \tag{10}$$

$$+ \sum_{i,j} u_{\uparrow\downarrow}(|\mathbf{r}^{\uparrow} - \mathbf{r}^{\downarrow}|) \tag{11}$$

$$+ \sum_{i<j,I} u_{\uparrow\uparrow H}(|\mathbf{r}_i^{\uparrow} - \mathbf{r}_I|, |\mathbf{r}_j^{\uparrow} - \mathbf{r}_I|, |\mathbf{r}_i^{\uparrow} - \mathbf{r}_j^{\uparrow}|) \tag{12}$$

$$+ \sum_{i<j,I} u_{\downarrow\downarrow H}(|\mathbf{r}_i^{\downarrow} - \mathbf{r}_I|, |\mathbf{r}_j^{\downarrow} - \mathbf{r}_I|, |\mathbf{r}_i^{\downarrow} - \mathbf{r}_j^{\downarrow}|) \tag{13}$$

$$+ \sum_{i,j,I} u_{\uparrow\downarrow H}(|\mathbf{r}_i^{\downarrow} - \mathbf{r}_I|, |\mathbf{r}_j^{\uparrow} - \mathbf{r}_I|, |\mathbf{r}_i^{\downarrow} - \mathbf{r}_j^{\uparrow}|) \tag{14}$$

The first line is a one-body Jastrow, the next two are spin resolved two-body Jastrows, and the last few lines are three-body spin-resolved Jastrows. In QMCPACK, these are represented by short-ranged b-splines, whereas the 3-body term is represented with polynomials.

We used Quantum Espresso to generate single particle orbitals in periodic boundary conditions. We used PBE orbitals[47], which were generated using the bare coulomb potential, the PBE functional, and a planewave cutoff of 120Ry. We also used Hartree-Fock (HF) orbitals with a planewave cutoff of 500Ry. These cutoffs were chosen by optimizing the DMC energy and VMC variance as a function of planewave cutoff. These values roughly minimized both.

DMC calculations employing Slater-Jastrow and Slater-Jastrow backflow wave functions follow the standard QMC workflow. All Jastrow parameters were simultaneously optimized at the VMC level, after which the backflow terms were optimized if present. Optimization was done to minimize the trial wave function energy using the linear method[49]. Convergence of the optimization procedure was checked by hand, after which longer VMC and DMC runs were performed. Within DMC, we used a time step of $\tau = 0.0025 Ha^{-1}$ with 1600 walkers for all wave functions.

For the MSD calculations, the workflow was somewhat nontrivial. There were two major problems we had to tackle. The first was how to calculate the c_k^i coefficients for CSF's that preserve singlet character even for high degrees of electronic excitation in periodic boundary conditions. The second problem was how to prescreen the number of determinants in the full configuration interaction space to something computationally tractable. Both of these issues are further exasperated by the fact that there are few quantum chemistry codes that can handle periodic boundary conditions.

To deal with the first problem, we used GAMESS[50,51], an open boundary condition code, to generate the CSF's for our box of hydrogen atoms in open boundary conditions, and then used these CSF's in periodic boundary conditions. The rationale is that the code uses the graphical unitary group approach to generate CSF's to arbitrarily high levels of excitations. This uses only the symmetries of the cluster, and nothing about the details of the electronic structure. The space group of our hydrogen molecules in periodic boundary conditions is going to have the trivial point/space group as a subgroup, so while not optimal, we should be able to generate CSF's in a cluster and have them preserve the basic spin symmetries when placed in periodic boundary conditions.

To deal with the prescreening of determinants, we placed our hydrogen atoms in open periodic boundary conditions and performed CI calculations in a truncated space of determinants. Starting with a cc-pCVTZ basis set and Hartree-Fock orbitals, we built up a restricted CI space consisting of a CAS(8,8), with single and double excitations out of the CAS, up to the 12th orbital above the occupied orbitals. After performing a CI calculation in this space, the CSF's were ordered according to the magnitude of their coefficients. We then took these CSF's and used them to define CSF's in periodic boundary conditions with PBE orbitals. QMC optimizations were done by first eliminating all CSF's whose initial magnitude was less than 0.01. After optimizing these coefficients simultaneously with the Jastrow factors, we would then decrease the cutoff threshold and repeat the procedure. We did this for thresholds of 0.01, 0.005, 0.0025, 0.001, and 0.0005, which amounted to roughly 160, 480, 1080, 2400, 3500 CSF's respectively. It should be noted that the number of CSF's we actively work with is a tiny sliver of our CAS+second order excitation space, which is spanned by more than 580,000 CSF's.

C. QMC Pressures

We computed all pressures using the virial theorem:

$$P^{vir} = \frac{1}{3\Omega}\left(2\langle T\rangle + \langle V\rangle\right) \tag{15}$$

Ω is the simulation cell volume, $\langle T\rangle$ is the expectation value for the electronic kinetic energy, and $\langle V\rangle$ is the expectation value for the total potential energy. However, due to the mixed-estimator problems in QMC, we used first the identity $\langle E\rangle = \langle T\rangle + \langle V\rangle$ to rewrite the virial pressure as:

$$P^{vir} = \frac{1}{3\Omega}\left(2\langle E\rangle - \langle V\rangle\right) \tag{16}$$

$\langle E \rangle$ can be sampled in DMC without bias, and since $\hat{V}$ is diagonal in position space, it can be sampled without bias using either reptation Monte Carlo or DMC with forward-walking. In this work, we used the latter. All observables were recorded at multiple projection times between 0 and $5Ha^{-1}$. To check convergence, we considered just the $\rho/\rho_0 = 3.750$ 8 atom configuration with SJ-PBE, SJ-HF, BF-PBE, and BF-HF wave functions. We fit the trace of the potential energy vs. projection time to an exponential form: $V(\beta) = a + be^{-c\beta}$. For PBE orbitals, we fit all data after a $1Ha^{-1}$ projection time, whereas for HF orbitals, we fit everything after a $2.5Ha^{-1}$ projection time (due to the presence of a minimum in $V(\beta)$. We found that taking forward walked estimates at $\beta = 5Ha^{-1}$ introduced errors smaller than the statistical error bar, so the reported numbers are essentially converged.

It should be noted that the pressure obtained from the virial estimate can be different from numerical differentiation of the total energy, the latter of which is desired. In Table I, we show a comparison between the virial and finite-difference pressure estimates associated with several different trial functions at the DMC level for a single 8 atom test configuration at $\rho/\rho_0 = 4.0$. In this particular case, we see a disagreement of roughly 1GPa between the virial and finite-difference estimates for all considered trial wave functions. Moreover, we observe a much greater sensitivity to the trial wave function when using the virial pressure estimates compared to the finite difference estimates–spreads of 1GPa versus 0.5GPa respectively. 1GPa is significant on the scales of this problem, especially so near the gas gun data, so a more careful study to pin down the pressure errors is needed for future Hugoniot work.

$\Delta v(a_0^3)$	$E^{SJ-PBE}(v)$	E^{SJ-HF}	$E^{BF-PBE}(v)$	E^{MSD-HF}
-5	-4.4998609(83)	-4.4990478(88)	-4.501961(13)	-4.502758(16)
5	-4.5055793(84)	-4.5047996(88)	-4.507555(13)	-4.508425(17)
P^{FD}(GPa)	16.824(35)	16.922(37)	16.461(54)	16.673(54)
P^{vir}(GPa)	17.438(21)	18.027(20)	17.037(51)	17.78(9)

TABLE I. DMC total energy at two finite volume displacements Δv of 8 atom $\rho/\rho_0 = 4.000$ test configuration. E^{SJ-PBE} and E^{BF-PBE} are the total DMC energies of SJ-PBE and BF-PBE wave functions respectively. The last two rows are the estimated finite difference pressure in GPa, followed by the virial pressure estimate.

D. Quantum Chemistry Methods

We performed calculations in periodic boundary conditions with PySCF[52] using Hartree-Fock (HF), second order Møller-Plesset perturbation theory (MP2)[53], restricted coupled-cluster with single and double excitations (CCSD), and restricted CCSD with perturbative triple excitations (CCSD(T))[54].

All reported energies and pressures were extrapolated to the complete basis set limit (CBS) using known formulas. We used cc-pV{D,T,Q,5}Z basis sets[55] to perform the fits. The astute reader will notice that we have not used basis sets augmented with diffuse functions. While we did test these, we found that linear dependency and computational cost issues prevented us from completing calculations with higher cardinality augmented basis sets, which negatively affected the quality of these basis set extrapolations as compared to standard basis sets.

For energies, we used the following standard formulas for the Hartree-Fock energy E_X^{HF} and correlation energy E_X^{corr} with a basis set cardinality of X.

$$E_X^{HF} = E_{CBS}^{HF} + a(V)e^{-b(V)X} \tag{17}$$

$$E_X^{corr} = E_{CBS}^{corr} + \frac{c(V)}{X^3} \tag{18}$$

Here, $a(V)$, $b(V)$, and $c(V)$ are fit coefficients that are assumed to be volume dependent. E_{CBS}^{HF} and E_{CBS}^{corr} are the complete basis set Hartree-Fock and correlation energies respectively.

Due to noise in the estimates of E_{CBS}^{HF} and E_{CBS}^{corr}, obtaining CBS pressure estimates from finite differences of CBS energies was not feasible. Instead, we computed P_X^{HF} and P_X^{corr} using finite differences for each basis set X. Taking the volume derivative of Eqs. 17 and 18 under the assumption that the fit coefficients are differentiable w.r.t. volume gives us the following pressure extrapolation formulas:

$$P_X^{HF} = P_{CBS}^{HF} + a'(V)e^{-b(V)X} \tag{19}$$

$$P_X^{corr} = P_{CBS}^{corr} + \frac{c'(V)}{X^3} \tag{20}$$

The primes denote volume deriatives of the fit coefficients appearing in Eqs. 17 and 18.

E. Zero-Variance Extrapolations

The variational theorem of quantum mechanics states:

$$E[\Psi_T] = \frac{\langle \Psi_T | \hat{H} | \Psi_T \rangle}{\langle \Psi_T | \Psi_T \rangle} \geq E_0 \tag{21}$$

Here, $E[\Psi_T]$ is the variational energy of a wave function $|\Psi_T\rangle$, E_0 is the exact ground state energy, and equality in the above expressions holds if $|\Psi_T\rangle$ is an eigenstate of $\hat{H}$. Likewise, there is a corrolary for the variance of the trial function $|\Psi_T\rangle$, $\sigma_E^2[\Psi_T]$:

$$\sigma_E^2[\Psi_T] = \frac{\langle \Psi_T | \hat{H}^2 | \Psi_T \rangle}{\langle \Psi_T | \Psi_T \rangle} - (E[\Psi_T])^2 \geq 0 \tag{22}$$

In other words, the variance of a wave function is greater than or equal to zero, with equality holding if $|\Psi_T\rangle$ is an eigenstate of $\hat{H}$.

Provided a method obeys the variational theorem (e.g. Hartree-Fock, VMC and diffusion Monte Carlo, selected CI), the above expressions provide a well defined heuristic for determining "better" and "worse" approximate solutions to eigenstates of $\hat{H}$. We will rely heavily on the variational theorem when trying to diagnose fixed-node errors in CEIMC style wave functions using QMC with more expensive trial wave functions.

Under certain circumstances, the scaling of both $E[\Psi_T]$ and $\sigma_E^2[\Psi_T]$ with wave function quality allows us to make the following association[56]:

$$E[\Psi_T] \propto \sigma_E^2[\Psi_T] \tag{23}$$

This has been used to great effect for iterative backflow and homogeneous electron calculations, where perturbative improvements to the wave function allowed extrapolations in agreement with released-node calculations.

For VMC trial wave functions, this is rigorous since we have direct access to the variance and energy for any trial wave function. However, the energies obtained from this extrapolation are too crude to be of much use for systems other than the homogeneous electron gas or idealized fermi systems. It would be nice to perform the extrapolation using the energy and variance of the fixed-node wave function, but to date no method exists to sample the fixed-node variance in a well defined manner. To circumvent this problem, we make the approximation that

$$\sigma^2_{E_{DMC}}[\Psi_T] \propto \sigma^2_{E_{VMC}}[\Psi_T] \tag{24}$$

Its very easy to construct counter examples where this approximation breaks down: for example by constructing a progression of "better" wave functions by increasing the complexity of the Jastrow factor. Here, $\sigma^2_{E_{VMC}}[\Psi_T]$ would decrease as we enhanced the quality of the Jastrow factor, whereas $E^2_{DMC}[\Psi_T]$ would stay constant since it is determined by the *nodal surface*, which doesn't change. However, for multideterminant style wave functions, we are fixing the form of the Jastrow. Adding determinants will then simultaneously improve both the nodal surface and the overall quality of the trial wave function. It is in this case where we expect that Eq. 24 to be reasonable. If Eq. 24 holds, we can reasonably extrapolate $E_{DMC}[\Psi_T]$ vs. $\sigma^2_{E_{VMC}}[\Psi_T]$, provided that the multidterminant enters a region close enough to the exact wave function $|\Psi_0\rangle$ such that linear extrapolation *à la* Eq. 23 holds. Note that we are defining the DMC variance as the variance of the fixed-node wave function Φ_{FN} over the physical Hamiltonian $\hat{H}$, and not over the fixed-node Hamiltonian, which would of course yield zero variance for the fixed-node wave function.

So how well does this work in practice? In Figure 2, we show the result of performing a linear extrapolation to zero variance using the SJ-HF and MSD-HF data points. The MSD energy vs. variance is not totally linear, and has some curvature. In spite of this, the linear extrapolation accounts for most of the correlation energy missed by BF-PBE, and agrees very well with CCSD(T)-CBS. Figure 6 shows that the extrapolation procedure consistently within approximately 0.1mHa/atom of the CCSD(T)-CBS across all tested configurations, indicating that the excellent agreement indicated in Figure 2 between our MSD extrapolated estimates and CCSD(T)-CBS is not a fluke.

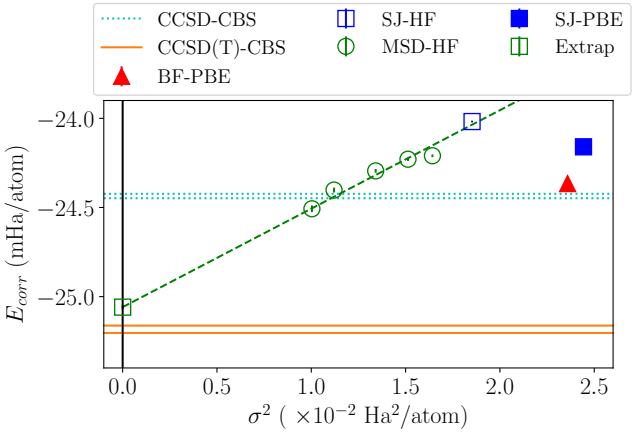

FIG. 2. DMC estimate of E_{corr} vs. VMC variance for several classes of wave function. The linear extrapolation to zero variance is shown for the SJ-HF and MSD-HF wave functions. While not perfect, it recovers all but 0.1mHa/atom of the correlation energy missed by BF-PBE.

III. RESULTS

We begin our discussion by first considering estimates of ΔE^* and ΔP^* for small 8 atom snapshots. The reason for this is the poor scaling of the quantum chemistry methods and our multi-Slater trial wave functions with system size. After this, we will consider limited QMC calculations performed in 128 atom cells and comment on the scaling of errors to the thermodynamic limit. As a result of this discussion, we obtain some likely estimates for ΔE^* and ΔP^* for the CEIMC calculations, and use these to produce an *ab initio* corrected Hugoniot.

We end the discussion by performing a similar analysis for several major DFT functionals, and provide an explanation for the observed robustness and accuracy of the DFT principal deuterium Hugoniot.

A. Small Cell Results

Our primary focus initially will be to attempt to construct wave functions that are "better" than the ones employed in the CEIMC calculations and to place lower bounds on the fixed-node energy error, and sensitivity of the pressure to quality of trial wave function. Interestingly, this can be done entirely within QMC. We will

cross validate against non-variational quantum chemistry methods later.

For all tested forms of wave functions, including MSD wave functions with differing numbers of CSF's, we plot both the total energy $E[\Psi_T]$ versus variance $\sigma_E^2[\Psi_T]$ and virial pressure vs. $\sqrt{\sigma_E^2[\Psi_T]}$. The rationale for the latter is that observables that do not commute with the Hamiltonian incur errors that are leading order in the quality of the trial wave function. We do not simplify this expression to make its connection to the variance explicit.

In Figure 3, we demonstrate the energy and variance reduction associated with our improved MSD-HF wave functions. While it takes some work, our best MSD wave function has an energy 0.1-0.2mHa/atom lower than BF-PBE with 60% lower variance. Additionally, given the energy of our best MSD-HF wave function is lower than the CCSD-CBS energy, this implies that we have constructed a wave function that is better than the CCSD exponential ansatz for this particular problem. This is our main reason for trusting (finite difference) pressure estimates from QMC more than from the quantum chemistry methods, since it appears as though CCSD(T) only provides slight corrections to CCSD (see Figure 7).

Of course, we see from the right of Figure 3 that despite how "good" our MSD wave functions are, the pressure virial is wildly sensitive to small variations in the choice of trial wave function. Notice how the MSD-HF wave functions initially predict higher virial pressures than SJ-HF, but then as we go to smaller $\sqrt{\sigma^2}$, it drops precipitously by almost 0.7GPa.

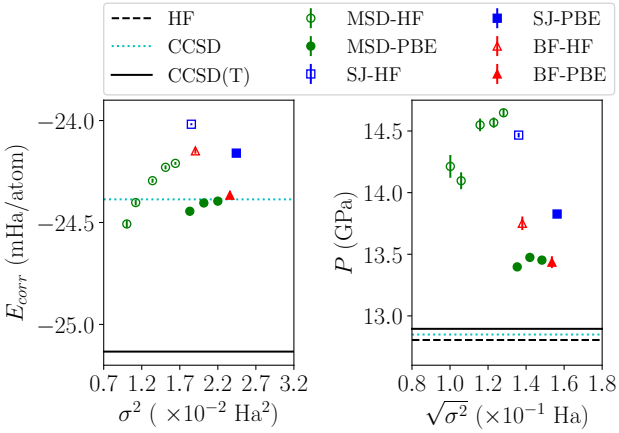

FIG. 3. Results for $\rho/\rho_0 = 3.750$ test configuration. (Left) E vs. σ^2 for several trial wave functions. (Right) Pressure virial vs. $\sqrt{(\sigma^2)}$ for same set of wave functions. "Zero-variance" extrapolations were performed by fitting the MSD-HF and SJ-HF to a straight line, and reading off the y-intercept. While this might seem aggressive, it oddly agrees well with the CCSD(T)-CBS estimates in almost all cases.

Now let's consider how the different wave functions perform across our small test set. In Figure 4, we show that for both SJ and BF type wave function, the use of HF orbitals drops the variance by roughly 30% at the lowest densities, but provides only a modest 20% reduction in variance at the highest density. Our best MSD wave function however has roughly 50-60% lower variance than BF-PBE, and more importantly, has a variance that is mostly flat with increasing density.

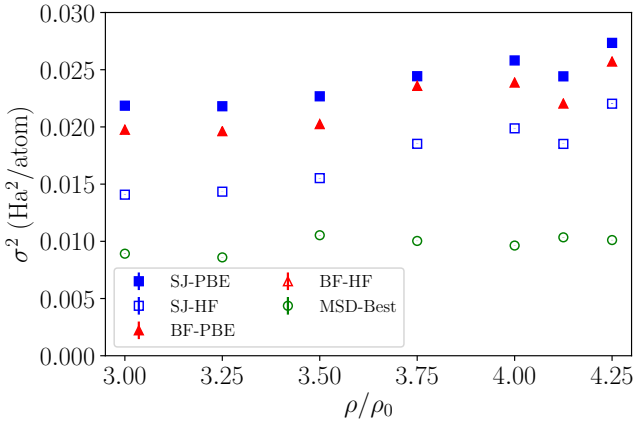

FIG. 4. Small cell VMC variance vs. density for all tested trial wave functions.

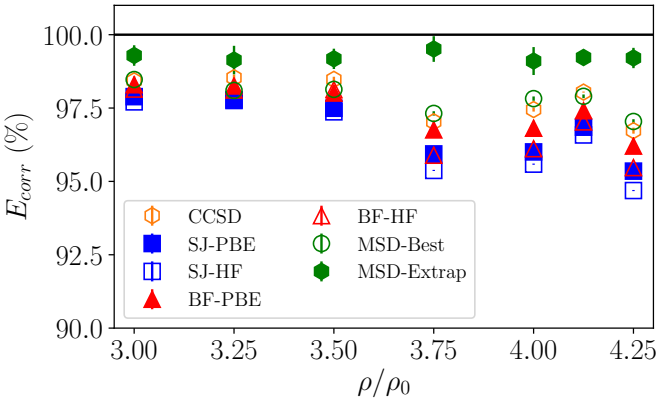

FIG. 5. Small cell comparison of $E_{corr}/E^{CCSD(T)_{corr}}$ for all tested methods in %.

In Figure 5, we show the percentage of recovered correlation energy relative to CCSD(T)-CBS, and in Figure 6, we show the difference in energies relative to BF-PBE. We clearly see that despite the gains in correlation energy made by BF-PBE over SJ-PBE (up to 0.2mHa/atom), we're still not close to the ground-state energy in these test configurations. In fact, the MSD wave function can recover about the same amount of energy over BF-PBE as BF-PBE did over SJ-PBE. This is significant in practical calculations because often times when one tries several different nodal surfaces (e.g. comparing Slater-Jastrow to backflow) and sees only small changes in the energy, variance, and other properties, this can infor-

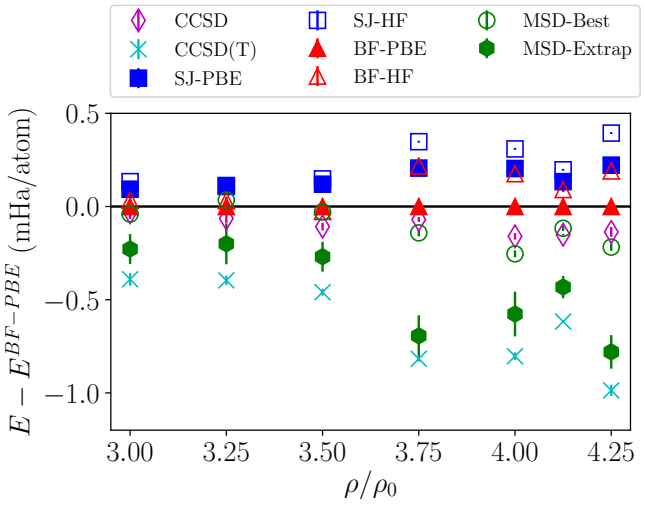

FIG. 6. Small cell comparison of $E - E^{BF-PBE}$ for all tested methods in mHa/atom.

mally be taken as a sign that we are "close" to the ground state. In the absence of nodal release or some extrapolation procedure however, this could lead one into a false sense of security regarding the quality of the calculation. Fortunately, the zero-variance extrapolation based on the MSD wave functions gives us a dramatically better estimate of how far off the exact answer we are.

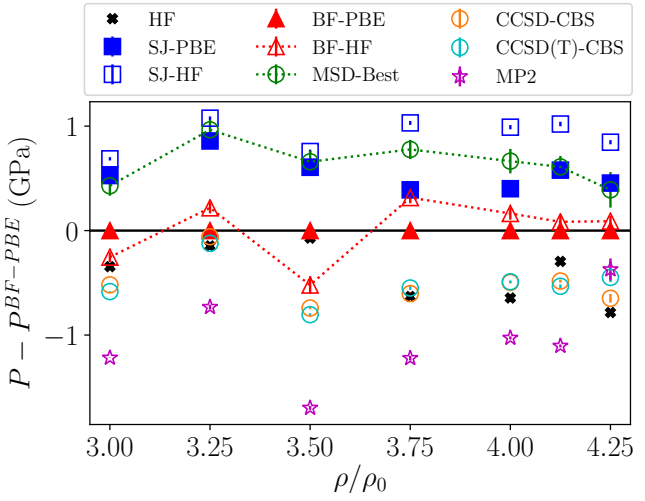

FIG. 7. Small cell comparison of $P - P^{BF-PBE}$ for all tested methods.

Lastly, we look at the spread of virial pressures. The noise prevents any definitive statements about density trends, but what should be clear is that every other trial wave function tested, including ones that have lower energies and variance than BF-PBE, has higher virial pressure estimates than BF-PBE. BF-HF indicates that the pressure is probably 0.5GPa higher than BF-PBE, wheras MSD-best indicates that its between 0.5GPa and 1GPa. Interestingly the quantum chemistry methods all seem to predict lower pressures than QMC, although as we've argued, MSD-best are better wave functions than CCSD-CBS, and BF-PBE is actually better than CCSD wave functions in the finite-basis sets we've considered. CCSD(T)-CBS seems to track CCSD-CBS very closely.

To summarize, all variants of DMC and CCSD recover smaller and smaller fractions of correlation energy, starting at roughly 98-99% at ρ/ρ_0 and dropping to between 96-97% for BF-PBE and the best multi-determinant wave function considered. In practice, this implies that DMC is overestimating the energy between 0.4mHa/atom and 1mHa/atom at the highest densities. We see that the MSD zero-variance extrapolation maintains excellent agreement with CCSD(T) across the density range. Interestingly, we note the similarity between DMC and CCSD, first observed by Trail *et al.*[20] for hydrogen clusters in open boundary conditions.

B. Thermodynamic Limit

Since the quantum chemistry methods we tested scaled like $O(N^5)$ to $O(N^7)$ with respect to the number of electrons, we could not run them on the 128 atom test sets. Thus, we have to proceed by analogy–looking to see if the differences between different trial wave functions observed in the small cells carry up to the large cells.

In Figure 8, we see that the variance per atom of all tested wave functions in the 128 atom cells drops significantly from the variance observed in the 8 atom cells. We also see the variance increases only very slightly as we increase the density, in contrast with the small cell. However, the use of Hartree-Fock orbitals continues to drop the variance of the trial wave function significantly compared against PBE orbitals.

In Figure 9, despite the noise, we can see that BF-PBE is recovering only a modest amount of energy over both SJ-PBE and SJ/BF-HF, between 0.1 and 0.3mHa/atom. This seems very comparable to the energy gains observed in the small cells, meaning that unless the amount of correlation in the system is reduced as we go to the 128 atom cells, BF-PBE is probably recovering only a small amount of remaining correlation energy, similar to the small cells. As mentioned before, cc-pVDZ and cc-pVTZ Hartree-Fock calculations at $\rho/\rho_0 = 3.75$ imply that the amount of correlation energy BF-DMC is recovering is larger in the big cell than in the small cell. One can hope that this is because DMC is more effective in the large cell and recovers a larger fraction of correlation energy, but an equally plausible scenario is that DMC is recovering a similar fraction of correlation energy as in the small cell–it's just that there is more correlation energy.

Lastly, we show how the pressures compare between SJ/BF-PBE and SJ/BF-HF. While there was significant noise in the small cells, we see the difference between the virial estimators is definitely exacerbated in the large cells. For example, while the difference between BF-PBE

and BF-HF pressure virials was between 0 and 0.6GPa in the small cells with a fair amount of noise, in the large cells the difference starts at 1GPa and monotonically increases to 2GPa at $\rho/\rho_0 = 4.25$. This sensitivity of the pressure virial to choice of orbitals is unexpected, and a problem that needs to be dealt with in future calculations in hydrogen at these thermodynamic conditions.

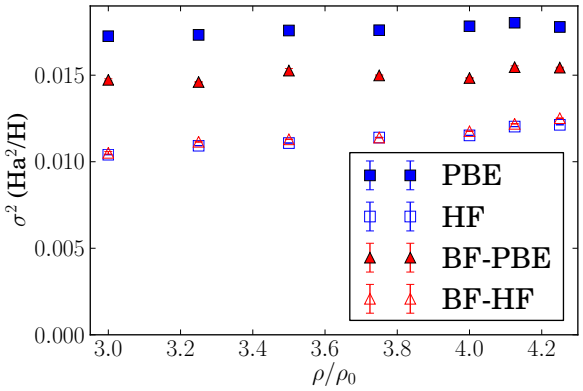

FIG. 8. VMC variance vs. density for all considered wave functions for 128 atom cells.

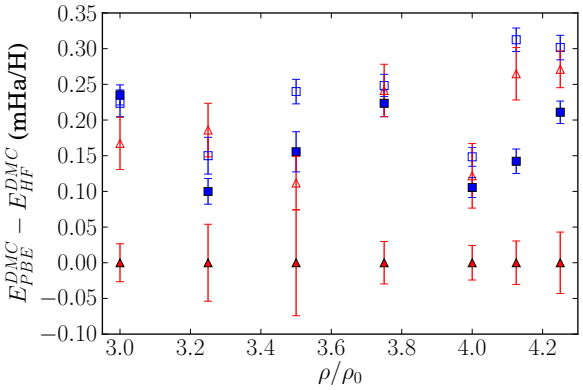

FIG. 9. $E - E^{BF-PBE}$ vs. density for all considered wave functions for 128 atom cells.

We consider reasonable but pessimistic estimates of the fixed-node error stemming from the BF-PBE energy and pressure virial. CCSD(T) indicates that in the small cells, BF-PBE is capturing about 96-97% of the correlation energy. We see in Fig. 6 that this translates to between 0.8mHa/atom to 1mHa/atom at the lowest pressure CEIMC Hugoniot point. Moreover, based on Hartree-Fock calculations in a 128 atom cell at $\rho/\rho_0 = 3.750$, the absolute magnitude of the correlation energy per atom appears to increase slightly, from 22-25mHa/atom in the small cells to at least 28.2 mHa/atom. If we assume that DMC is recovering 96% of correlation energy as we scale up to the thermodynamic

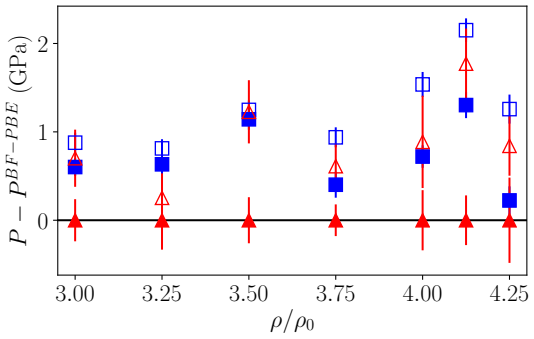

FIG. 10. $P - P^{BF-PBE}$ vs. density for all tested wave functions for 128 atom cells.

limit, this likely means that BF-PBE is overestimating the total energy by 1.2mHa/atom. We'll use this figure for ΔE.

C. Corrected CEIMC Hugoniot

Based on the above considerations, let us assume that the fixed-node errors at the CEIMC Hugoniot point $\rho/\rho_0 = 3.91$ and $P = 18GPa$ are $\Delta E^* = 1.2mHa/atom$ and $\Delta P^* = -2GPa$[42]. We now investigate how much of the discrepancy between the CEIMC Hugoniot and experiment *could* be attributed to these errors. To do this, we assume the published CEIMC Hugoniot is our reference, and use Eq. 2 to compute a "corrected" ρ^* from ΔE^* and ΔP^*.

To simplify this analysis, we make two approximations. First, ΔE^* and ΔP^* are assumed to be independent of density and temperature. Based on our above results showing that the fixed-node approximation gets worse as density is increased and molecular disassociation is approached, this should be seen as a conservative view of the impact of fixed-node errors on the deuterium Hugoniot; larger δE^* and δP^* values will perturb the actual Hugoniot even more. Secondly, to compute $C(\rho, T)$ appearing in Eq. 2, we use the Kerley equation of state. We find that the following results are largely insensitive to the choice of EOS used for this term.

We show in Figure 1 the corrected Hugoniot. We find that the fixed-node error in the evaluation of the Rankine-Hugoniot relation can reasonably account for roughly 60% to 70% of the discrepancy between CEIMC and experiment at the two lowest pressure CEIMC Hugoniot points. This would imply that the current discrepancy between CEIMC and experiment is *not* caused by controllable systematic errors or any issues with the underlying CEIMC algorithm, at least at the lowest pressure points.We do see that the assumed form of the fixed-node errors is not enough to fix the discrepancy between experiment and theory farther up the Hugoniot (i.e. near peak compression), but as mentioned previously, this is

not entirely unexpected given that our errors were estimated at significantly lower temperatures and densities. We do note that enouragingly, the pressure at which peak compression occurs is nudged much closer to the experiment.

D. DFT Hugoniot

Up to this point, we have spent quite a bit of time trying to understand why the deuterium Hugoniot seems to be sensitive to the errors in accurate many-body methods. This of course invites us to question why DFT appears to be so accurate for the principal Hugoniot, and why this accuracy seems very robust to the choice of approximate exchange-correlation functional.

Using Eq. 2, we can now look at the sensitivity of the predicted deuterium Hugoniot by estimating ΔE^* and ΔP^* for several electronic structure methods. We computed ΔE^* and ΔP^* relative to PBE[47] for the functionals tested in Ref.[28]: PW91[57], optB86b[58], vdW-DF[59], vdW-DF2[60]. Our choice of PBE as a reference is arbitrary but inconsequential for studying the impact of error trends on the Hugoniot. A test set was constructed by drawing single 128 atom snapshots from equilibrated DFT based molecular dynamics simulations at different temperatures and densities, using the PBE functional. One configuration targets the reference point at T=22K, whereas the remaining six target the T=4000K isotherm at densities between $\rho/\rho_0 = 3.00 - 4.25$, so chosen because it approximates the thermodynamic states near the highest pressure gas gun data. In Figure 11, we plot ΔP versus $\Delta(E - E_0)$ relative to PBE on the PBE Hugoniot point at $\rho/\rho_0 = 3.50$. Inspection of Eq. 2 allows us to write the condition for perfect error cancellation at a specific Hugoniot point, i.e. where the term in brackets vanishes:

$$\Delta(P + P_0) = 2\Delta(E - E_0)/(\rho_0^{-1} - \rho^{-1}) \qquad (25)$$

We overlay this line on Figure 11. We see that while the spread of energies and pressures differ from PBE by almost 12mHa/atom and 4GPa respectively among the different functionals, these errors largely cancel. Even within a particular choice of functional, the spread in energy and pressure relative to PBE can be 2-3mHa/atom and 1-2GPa respectively just along the T=4000K isotherm.

This figure fails to convey other more subtle channels of error cancellation. Firstly, nontrivial error cancellation can occur within the $E - E_0$ term, which we estimate to be on the order of 2mHa/atom for optB86b and PBE. Secondly, different functionals will produce different configurations over the course of a molecular dynamics simulation, which will impact ensemble estimates $\Delta(E^* - E_0^*)$ and ΔP^* and potentially facilitate error cancellation. For the hollow green square in Fig. 11, we used equilibrated PBE and vdW-DF molecular dynamics simulations at at the reference point and the $T = 4000K$ $\rho/\rho_0 = 3.50$ point to estimate the ensemble values of $\Delta(E^* - E_0^*)$ and ΔP^*. We find this new point is consistent with the density difference between the vdW-DF and PBE Hugoniot points[42]. From this perspective, the much more consistent picture that DFT provides for the deuterium Hugoniot, as opposed to the liquid-liquid transition for example, owes much to its ability to fully utilize the error cancellation properties of the Rankine-Hugoniot relation.

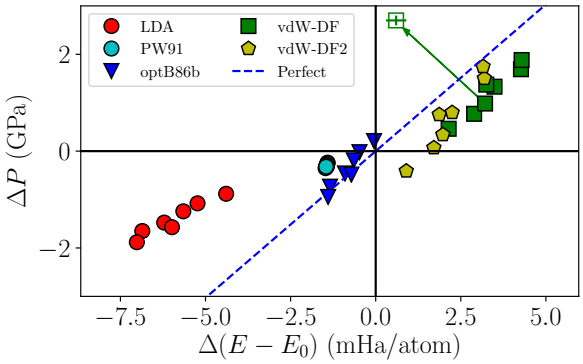

FIG. 11. Plot of energy and pressure differences relative to PBE for the T=4000K isotherm test set. The error cancellation line (dashed blue) is calculated using the PBE Hugoniot point at $\rho/\rho_0 = 3.50$. The green arrow indicates the impact of configurational sampling for a single vdW-DF density.

IV. CONCLUSION

This analysis *suggests* that many-body methods can introduce errors into the evaluation of the Rankine-Hugoniot relation that are non-negligible and should be corrected, especially if these errors happen to be additive. To properly correct a many-body Hugoniot and *prove* our assertions will require the following modifications of our methodology. Firstly, we need to treat larger cells and mitigate electronic finite-size effects, especially where deuterium is semi-metallic. This is partly so we can make a more accurate estimate of fixed-node pressure errors[42], but more importantly so that the discrepancy between CEIMC and experiment at peak compression can be addressed. Γ-point results for an 8 atom snapshot near peak compression ($T = 10000K$, $\rho/\rho_0 = 5.00$) suggest that BF-PBE DMC and CCSD-CBS are only recovering 91.4% and 95.1% of the correlation energy respectively, which amounts to an +2.4mHa/atom error for BF-PBE[42]. While this would go a long way towards explaining the higher compressions observed in CEIMC, it's unclear how much of this error will persist as we progress to the thermodynamic limit. Secondly a very careful audit of other sources of errors needs to be done. Errors in addition to fixed-node, such as finite size effects, electron thermal effects, configurational sampling, and others, will all contribute to ΔE^* and ΔP^*, but with potentially differing signs. We believe that it is possi-

ble to improve our methodology in the ways described by using DMC based on iterative backflow[56], selected CI based multideterminant[61], or even antisymmetrized geminal product trial wave functions[62,63]. Furthermore, cross validation against second quantized QMC methods, such as FCIQMC[64] and AFQMC[65], should be extremely valuable moving forward.

To summarize, because many-body methods are extremely accurate for the deuterium reference point, the error in the predicted Hugoniot density is entirely determined by the absolute errors $\Delta E^*(\rho, T)$ and $\Delta P^*(\rho, T)$, a situation which happens comparatively rarely in practical *ab initio* calculations. Moreover, it appears as though these methods experience error addition in a particularly difficult part of the deuterium phase diagram, which exaggerates errors in the principal Hugoniot in spite of their ostensibly more accurate descriptions of electronic structure. In contrast, DFT's ability to use error cancellation allows for the construction of an accurate principal Hugoniot, but this does not guarantee similar accuracy for other observables. There exists noticeable differences regarding the predicted rate of molecular disassociation along the Hugoniot and the temperatures of Hugoniot states[28]. Moreover, most known functionals underestimate the reshock pressures by a modest but statistically significant amount relative to experiment[28], which is not surprising given the amount of error cancellation we observe. These issues demonstrate a need for more beyond-DFT methods like CEIMC, but this work serves as a reminder that without a way to quantify and potentially reduce the size of the approximations intrinsic to many-body methods, more accurate methods may not always produce more accurate properties.

ACKNOWLEDGMENTS

RCC would like to thank David Ceperley, Markus Holzmann, and Carlo Pierleoni for helpful discussions, and Joshua Townsend for helpful comments on this manuscript. RCC and LS were supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division, as part of the Computational Materials Sciences Program and Center for Predictive Simulation of Functional Materials. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

* rclay@sandia.gov

1. W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys. **73**, 33 (2001).
2. M. A. Morales, R. Clay, C. Pierleoni, and D. M. Ceperley, Entropy **16**, 287 (2014).
3. R. P. Dias and I. F. Silvera, Science **355**, 715 (2017), http://science.sciencemag.org/content/355/6326/715.full.pdf.
4. X.-D. Liu, P. Dalladay-Simpson, R. T. Howie, B. Li, and E. Gregoryanz, Science **357** (2017), 10.1126/science.aan2286, http://science.sciencemag.org/content/357/6353/eaan2286.full.pdf.
5. M. Zaghoo, A. Salamat, and I. F. Silvera, Phys. Rev. B **93**, 155128 (2016).
6. M. Zaghoo and I. F. Silvera, Proceedings of the National Academy of Sciences (2017), 10.1073/pnas.1707918114, http://www.pnas.org/content/early/2017/10/23/1707918114.full.pdf.
7. R. S. McWilliams, D. A. Dalton, M. F. Mahmood, and A. F. Goncharov, Phys. Rev. Lett. **116**, 255501 (2016).
8. M. D. Knudson, M. P. Desjarlais, A. Becker, R. W. Lemke, K. R. Cochrane, M. E. Savage, D. E. Bliss, T. R. Mattsson, and R. Redmer, Science **348**, 1455 (2015), http://science.sciencemag.org/content/348/6242/1455.full.pdf.
9. S. Azadi and W. M. C. Foulkes, Phys. Rev. B **88**, 014115 (2013).
10. M. A. Morales, J. M. McMahon, C. Pierleoni, and D. M. Ceperley, Phys. Rev. B **87**, 184107 (2013).
11. M. A. Morales, J. M. McMahon, C. Pierleoni, and D. M. Ceperley, Phys. Rev. Lett. **110**, 065702 (2013).
12. J. M. McMahon, M. A. Morales, C. Pierleoni, and D. M. Ceperley, Rev. Mod. Phys. **84**, 1607 (2012).
13. D. Ceperley, M. Dewing, and C. Pierleoni, in *Bridging Time Scales: Molecular Simulations for the Next Decade*, edited by P. Nielaba, M. Mareschal, and G. Ciccotti (Springer Berlin Heidelberg, Berlin, Heidelberg, 2002) pp. 473–500.
14. C. Pierleoni, D. M. Ceperley, and M. Holzmann, Phys. Rev. Lett. **93**, 146402 (2004).
15. E. Liberatore, M. A. Morales, D. M. Ceperley, and C. Pierleoni, Molecular Physics **109**, 3029 (2011), https://doi.org/10.1080/00268976.2011.624992.
16. G. Mazzola, S. Yunoki, and S. Sorella, Nature Communications **5**, 3487 EP (2014).
17. G. Mazzola and S. Sorella, Phys. Rev. Lett. **114**, 105701 (2015).
18. G. Mazzola and S. Sorella, Phys. Rev. Lett. **118**, 015703 (2017).
19. R. C. Clay, J. Mcminis, J. M. McMahon, C. Pierleoni, D. M. Ceperley, and M. A. Morales, Phys. Rev. B **89**, 184106 (2014).
20. J. R. Trail, P. López Ríos, and R. J. Needs, Phys. Rev. B **95**, 115116 (2017).
21. J. McMinis, R. C. Clay, D. Lee, and M. A. Morales, Phys. Rev. Lett. **114**, 105305 (2015).
22. S. Azadi, B. Monserrat, W. M. C. Foulkes, and R. J. Needs, Phys. Rev. Lett. **112**, 165501 (2014).
23. N. D. Drummond, B. Monserrat, J. H. Lloyd-Williams, P. L. Ríos, C. J. Pickard, and R. J. Needs, Nature Communications **6**, 7794 EP (2015).

[24] G. Rillo, M. A. Morales, D. M. Ceperley, and C. Pierleoni, The Journal of Chemical Physics **148**, 102314 (2018), https://doi.org/10.1063/1.5001387.

[25] M. A. Morales, C. Pierleoni, E. Schwegler, and D. M. Ceperley, Proceedings of the National Academy of Sciences **107**, 12799 (2010), http://www.pnas.org/content/107/29/12799.full.pdf.

[26] C. Pierleoni, M. A. Morales, G. Rillo, M. Holzmann, and D. M. Ceperley, Proceedings of the National Academy of Sciences **113**, 4953 (2016), http://www.pnas.org/content/113/18/4953.full.pdf.

[27] N. M. Tubman, E. Liberatore, C. Pierleoni, M. Holzmann, and D. M. Ceperley, Phys. Rev. Lett. **115**, 045301 (2015).

[28] M. D. Knudson and M. P. Desjarlais, Phys. Rev. Lett. **118**, 035501 (2017).

[29] N. C. Holmes, M. Ross, and W. J. Nellis, Phys. Rev. B **52**, 15835 (1995).

[30] L. B. Da Silva, P. Celliers, G. W. Collins, K. S. Budil, N. C. Holmes, T. W. Barbee Jr., B. A. Hammel, J. D. Kilkenny, R. J. Wallace, M. Ross, R. Cauble, A. Ng, and G. Chiu, Phys. Rev. Lett. **78**, 483 (1997).

[31] G. W. Collins, L. B. Da Silva, P. Celliers, D. M. Gold, M. E. Foord, R. J. Wallace, A. Ng, S. V. Weber, K. S. Budil, and R. Cauble, Science **281**, 1178 (1998), http://science.sciencemag.org/content/281/5380/1178.full.pdf.

[32] M. D. Knudson, D. L. Hanson, J. E. Bailey, C. A. Hall, J. R. Asay, and W. W. Anderson, Phys. Rev. Lett. **87**, 225501 (2001).

[33] G. V. Boriskov, A. I. Bykov, R. I. Il'kaev, V. D. Selemir, G. V. Simakov, R. F. Trunin, V. D. Urlin, A. N. Shuikin, and W. J. Nellis, Phys. Rev. B **71**, 092104 (2005).

[34] D. G. Hicks, T. R. Boehly, P. M. Celliers, J. H. Eggert, S. J. Moon, D. D. Meyerhofer, and G. W. Collins, Phys. Rev. B **79**, 014112 (2009).

[35] K. Falk, S. Regan, J. Vorberger, M. Barrios, T. Boehly, D. Fratanduono, S. Glenzer, D. Hicks, S. Hu, C. Murphy, P. Radha, S. Rothman, A. Jephcoat, J. Wark, D. Gericke, and G. Gregori, High Energy Density Physics **8**, 76 (2012).

[36] K. Falk, S. P. Regan, J. Vorberger, B. J. B. Crowley, S. H. Glenzer, S. X. Hu, C. D. Murphy, P. B. Radha, A. P. Jephcoat, J. S. Wark, D. O. Gericke, and G. Gregori, Phys. Rev. E **87**, 043112 (2013).

[37] M. P. Desjarlais, Phys. Rev. B **68**, 064204 (2003).

[38] L. Caillabet, S. Mazevet, and P. Loubeyre, Phys. Rev. B **83**, 094101 (2011).

[39] M. P. Desjarlais, M. D. Knudson, and K. R. Cochrane, Journal of Applied Physics **122**, 035903 (2017), https://doi.org/10.1063/1.4991814.

[40] B. Holst, R. Redmer, and M. P. Desjarlais, Phys. Rev. B **77**, 184201 (2008).

[41] G. E. Duvall and R. A. Graham, Rev. Mod. Phys. **49**, 523 (1977).

[42] See Supplemental Material at [URL will be inserted by publisher]

[43] G. Kresse and D. Joubert, Phys. Rev. B **59**, 1758 (1999).

[44] G. Kresse and J. Hafner, Phys. Rev. B **47**, R558 (1993).

[45] G. Kresse and J. Hafner, Phys. Rev. B **49**, 14251 (1994).

[46] G. Kresse and J. Furthmüller, Phys. Rev. B **54**, 11169 (1996).

[47] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).

[48] J. Kim, A. D. Baczewski, T. D. Beaudet, A. Benali, M. C. Bennett, M. A. Berrill, N. S. Blunt, E. J. L. Borda, M. Casula, D. M. Ceperley, S. Chiesa, B. K. Clark, R. C. C. III, K. T. Delaney, M. Dewing, K. P. Esler, H. Hao, O. Heinonen, P. R. C. Kent, J. T. Krogel, I. Kylnp, Y. W. Li, M. G. Lopez, Y. Luo, F. D. Malone, R. M. Martin, A. Mathuriya, J. McMinis, C. A. Melton, L. Mitas, M. A. Morales, E. Neuscamman, W. D. Parker, S. D. P. Flores, N. A. Romero, B. M. Rubenstein, J. A. R. Shea, H. Shin, L. Shulenburger, A. F. Tillack, J. P. Townsend, N. M. Tubman, B. V. D. Goetz, J. E. Vincent, D. C. Yang, Y. Yang, S. Zhang, and L. Zhao, Journal of Physics: Condensed Matter **30**, 195901 (2018).

[49] C. J. Umrigar, J. Toulouse, C. Filippi, S. Sorella, and R. G. Hennig, Phys. Rev. Lett. **98**, 110201 (2007).

[50] S. M. W., B. K. K., B. J. A., E. S. T., G. M. S., J. J. H., K. Shiro, M. Nikita, N. K. A., S. Shujun, W. T. L., D. Michel, and M. J. A., Journal of Computational Chemistry **14**, 1347, https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.540141112.

[51] M. S. Gordon and M. W. Schmidt, in *Theory and Applications of Computational Chemistry*, edited by C. E. Dykstra, G. Frenking, K. S. Kim, and G. E. Scuseria (Elsevier, Amsterdam, 2005) pp. 1167 – 1189.

[52] S. Qiming, B. T. C., B. N. S., B. G. H., G. Sheng, L. Zhendong, L. Junzi, M. J. D., S. E. R., S. Sandeep, W. Sebastian, and C. G. KinLic, Wiley Interdisciplinary Reviews: Computational Molecular Science **8**, e1340 (2018), https://onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1340.

[53] A. Szabo and N. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*, Dover Books on Chemistry (Dover Publications, 1989).

[54] R. J. Bartlett and M. Musiał, Rev. Mod. Phys. **79**, 291 (2007).

[55] T. H. D. Jr., The Journal of Chemical Physics **90**, 1007 (1989), https://doi.org/10.1063/1.456153.

[56] M. Taddei, M. Ruggeri, S. Moroni, and M. Holzmann, Phys. Rev. B **91**, 115106 (2015).

[57] J. P. Perdew, Electronic Structure of Solids '91 **11** (1991).

[58] J. Klime, D. R. Bowler, and A. Michaelides, Journal of Physics: Condensed Matter **22**, 022201 (2010).

[59] M. Dion, H. Rydberg, E. Schröder, D. C. Langreth, and B. I. Lundqvist, Phys. Rev. Lett. **92**, 246401 (2004).

[60] K. Lee, E. D. Murray, L. Kong, B. I. Lundqvist, and D. C. Langreth, Phys. Rev. B **82**, 081101 (2010).

[61] E. Giner, A. Scemama, and M. Caffarel, Canadian Journal of Chemistry **91**, 879 (2013), https://doi.org/10.1139/cjc-2013-0017.

[62] M. Casula and S. Sorella, The Journal of Chemical Physics **119**, 6500 (2003), https://doi.org/10.1063/1.1604379.

[63] E. Neuscamman, Phys. Rev. Lett. **109**, 203001 (2012).

[64] G. H. Booth, A. Grüneis, G. Kresse, and A. Alavi, Nature **493**, 365 (2013).

[65] S. Zhang and H. Krakauer, Phys. Rev. Lett. **90**, 136401 (2003).