Rigorous constraints on three-nucleon forces in chiral effective field theory from fast and accurate calculations of few-body observables

S. Wesolowski,[1,*] I. Svensson,[2,†] A. Ekström,[2,‡] C. Forssén,[2,§] R. J. Furnstahl,[3,‖]
J. A. Melendez[3,¶] and D. R. Phillips[4,5,6,#]

[1]*Department of Mathematical Sciences, Salisbury University, Salisbury, Maryland 21801, USA*
[2]*Department of Physics, Chalmers University of Technology, SE-412 96 Göteborg, Sweden*
[3]*Department of Physics, The Ohio State University, Columbus, Ohio 43210, USA*
[4]*Department of Physics and Astronomy and Institute of Nuclear and Particle Physics, Ohio University, Athens, Ohio 45701, USA*
[5]*Institut für Kernphysik, Technische Universität Darmstadt, 64289 Darmstadt, Germany*
[6]*ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung GmbH, 64291 Darmstadt, Germany*

We explore the constraints on the three-nucleon force (3NF) of chiral effective field theory (χEFT) that are provided by bound-state observables in the $A = 3$ and $A = 4$ sectors. Our statistically rigorous analysis incorporates experimental error, computational method uncertainty, and the uncertainty due to truncation of the χEFT expansion at next-to-next-to-leading order. A consistent solution for the ^{3}H binding energy, the ^{4}He binding energy and radius, and the ^{3}H β-decay rate can only be obtained if χEFT truncation errors are included in the analysis. The β-decay rate is the only one of these that yields a nondegenerate constraint on the 3NF low-energy constants, which makes it crucial for the parameter estimation. We use eigenvector continuation for fast and accurate emulation of no-core shell model calculations of the few-nucleon observables. This facilitates sampling of the posterior probability distribution, allowing us to also determine the distributions of the parameters that quantify the truncation error. We find a χEFT expansion parameter of $Q = 0.33 \pm 0.06$ for these observables.

I. MOTIVATION AND GOALS

In low-energy effective field theories (EFTs) of many-body systems, three- and higher-body forces inevitably arise because they capture the effect of degrees of freedom not resolved in the EFT [1–3]. In the variant of chiral EFT (χEFT) without an explicit Delta resonance, three-nucleon forces (3NFs) first appear in the Hamiltonian at third order (next-to-next-to-leading order, NNLO) in the EFT expansion. This first contribution depends on two parameters, called c_D and c_E, not already determined by nucleon-nucleon (NN) or pion-nucleon (πN) scattering. The terms proportional to c_D and c_E, together with the venerable Fujita-Miyazawa term [4], form the dominant piece of the 3NF in χEFT [5,6]. This 3NF has small, but important, effects in light nuclei and helps drive saturation in heavier systems and symmetric nuclear matter [7]. But—as in any EFT—c_D and c_E must be estimated from data, either using experimental measurements or theoretical sources. Doing that reliably, with error bars that account for all uncertainties, is key to accurate use of χEFT forces in computations of nuclei.

In this work, we carry out parameter estimation for c_D and c_E within a Bayesian framework. We explore the constraints on c_D and c_E provided by several observables: the triton and ^{4}He particle binding energies, the ^{4}He particle charge radius, and the Gamow-Teller matrix element of the triton, as extracted from tritium β decay. In addition to the standard treatment of uncertainties in the experimental measurements, we also account for model discrepancy [8,9] by considering the uncertainty in the χEFT Hamiltonian itself. In particular, we include χEFT truncation errors in the parameter estimation using a statistical model applied previously in the NN sector [10,11]. A novel feature of our analysis is that we employ eigenvector continuation (EC) [12] to implement rapid sampling [13,14] of a multidimensional posterior, and hence obtain joint probability distributions for c_D, c_E, and the EFT expansion parameter, Q. The fits of the NN and πN parameters that are inputs to our calculations also have uncertainties; we propagate the uncertainties from NN but not from πN (see Sec. II E). The outputs from the parameter estimation are not single values for c_D and c_E but multidimensional posterior probability density functions (pdfs). These—referred to as "posteriors" hereafter—can be used to identify correlations and to propagate uncertainties to observables.

This is not an exhaustive study of parameter estimation for these 3NF parameters. Rather our goal is to examine the implications of using particular combinations of observables for constraining c_D and c_E while exemplifying statistical best practices [10], in particular the inclusion of EFT truncation errors as a guard against overfitting. There are several recent

*scwesolowski@salisbury.edu

†isak.svensson@chalmers.se

‡andreas.ekstrom@chalmers.se

§christian.forssen@chalmers.se

‖furnstahl.1@osu.edu

¶melendez.27@osu.edu

#phillid1@ohio.edu

and ongoing efforts seeking analogous constraints, which can provide complementary information, and many of our conclusions reinforce those of other authors. In particular, we build on the use of tritium β decay in Refs. [15,16] (cf. Ref. [17] for an analysis in χEFT with explicit $\Delta(1232)$ degrees of freedom) and compare our results to the c_D-c_E posteriors found using other observables such as Nd scattering [18] and neutron-α scattering [19].

In Sec. II we describe our Bayesian strategy for estimating c_D and c_E: our choice of likelihood and prior distributions, including our optimization of the input NN force. Then in Sec. III we discuss details of the few-body methods used to compute observables and introduce the EC emulators that make our comprehensive parameter-estimation process feasible. Results are given in Sec. IV, first for the most comprehensive fit of the 3NF parameters and then using constraints provided by individual observables. We identify the induced correlations, infer knowledge of the EFT expansion, and display the range of χEFT predictions obtained from our c_D and c_E posterior. Our takeaway points and avenues for future work are summarized in Sec. V. An open-source Python package `fit3bf` accompanies this article [20] and can be used to reproduce all the figures herein.

II. BAYESIAN STRATEGY

Our aim is to determine 3NF low-energy constants (LECs) $\{c_D, c_E\}$ from experimental data $\mathbf{y}_{\exp}$. The few-body observables in $\mathbf{y}_{\exp}$ are the mass and radius of ^{4}He, and the mass and β-decay rate of ^{3}H. The Bayesian approach we implement can account for all sources of uncertainty: from data, from the theoretical model, and from the calculational methods [10,21]. Some of these will not be treated in this work because they are either negligible (e.g., emulator error; see Sec. III B) or more work needs to be done to properly include them (πN LECs; see Sec. II E). The largest source of uncertainty is the χEFT truncation error, but we also account for the experimental and the few-nucleon solver uncertainties. Our use of emulators makes the observable calculations required for Markov chain Monte Carlo (MCMC) sampling rapid enough that we can fully account for NN uncertainties and incorporate truncation uncertainty in a Bayesian fashion.

In this section we first detail our approach to assessing truncation errors [22,23]. We then write down the forms for the posterior and prior, before describing how the convergence pattern of $A = 3$ and $A = 4$ observables provide information on the truncation error. The section closes with a description of how the NN LEC values and uncertainties that are input to our calculation are obtained.

A. Including EFT truncation error

We follow a Bayesian approach for the consistent incorporation of all higher-order terms in the EFT [23]. Let $y_{\text{th}}(\vec{a})$ be the prediction of some observable y at a fixed order in the EFT and for fixed values of LECs $\vec{a}$. Here, $\vec{a}$ includes the NN LECs along with c_D and c_E. Dependence on the πN LECs is left implicit throughout; see Sec. II C. We account for the presence of theory and experimental uncertainties δy_{th} and

$\delta y_{\exp}$ by writing [8–10]

$$y_{\exp} = y_{\text{th}}(\vec{a}) + \delta y_{\text{th}} + \delta y_{\exp}. \quad (1)$$

That is, the theoretical value differs from the measured value because of *both* experimental uncertainties and discrepancies in the theory. For the measurement errors $\delta y_{\exp}$ we assume a Gaussian error term that is uncorrelated between observables. However, this assumption has little impact on our results because experimental errors are small relative to theory uncertainties.

The distribution of the theory discrepancy δy_{th} also follows a Gaussian distribution [22]. It depends on two dimensionless parameters related to the EFT convergence pattern. The first is the EFT expansion parameter Q, which is a number in $(0, 1)$ and governs the factor by which each correction should shrink in a well-constructed EFT. The model encodes the expectation that the first omitted term in a χEFT of order k is of order $y_{\text{ref}} \bar{c} Q^{k+1}$, where y_{ref} is the known characteristic size of the observable y [11,22]. The second dimensionless parameter is then $\bar{c}$. It governs the magnitude of the relative correction at each order after we have accounted for Q.

For a given $\bar{c}$ and Q the error due to all terms beyond $O(Q^k)$ in the EFT can be summed and used to create a covariance matrix between observable i and observable j. In this work we assume that there are no correlations between the EFT errors for the observables of interest, thus the covariance matrix is diagonal [10]:

$$(\Sigma_{\text{th}})_{ij} = \left[\frac{(y_{\text{ref}} \bar{c} Q^{k+1})^2}{1 - Q^2} \right] \delta_{ij}. \quad (2)$$

We view this as the simplest form of Σ_{th} that models the effect of higher-order terms in the χEFT expansion. There are certainly other plausible forms of Σ_{th} that invoke correlated EFT uncertainties, e.g., we could assume the fourth-and-higher order contributions to these observables are correlated according to the pattern of correlations observed between them at lower orders; cf. Refs. [23–25]. As we mention in Sec. V below, exploring the impact of more sophisticated forms of Σ_{th} on the results is an avenue for future work. Practitioners who wish to examine such possibilities themselves should find it straightforward to do so using the open-source Python package `fit3bf` that accompanies this article [20].

B. The pdf for c_D and c_E

The form of the experimental and theory uncertainties and the relation (1) are sufficient to determine that the likelihood is given by

$$\text{pr}(\mathbf{y}_{\exp} \,|\, \vec{a}, \Sigma, I) = \mathcal{N}[\mathbf{y}_{\text{th}}, \Sigma]. \quad (3)$$

This likelihood is a multivariate Gaussian pdf, defined by central values from theory, $\mathbf{y}_{\text{th}}$, and the covariance matrix $\Sigma \equiv \Sigma_{\exp} + \Sigma_{\text{method}} + \Sigma_{\text{th}}$, where we have also included a term Σ_{method} that describes the uncertainty of our few-nucleon solver. Here, the combination $\Sigma_{\exp} + \Sigma_{\text{method}}$ is a diagonal matrix given by the column of adopted errors in Table I. The precision of our few-nucleon calculations is discussed in Sec. III. The covariance matrix could be extended to include

TABLE I. Results at LO, NLO, and NNLO for the observables used in various combinations to form our likelihood: the binding energies of the ^{3}H and ^{4}He states, the rms (point-proton) radius of ^{4}He, and the β-decay comparative half-life of ^{3}H. Experimental data are from [45–48]. Non-negligible uncertainties in the last digits are then given in brackets. Adopted uncertainties are the larger of those and uncertainties from the calculational method used to solve the Schrödinger equation. Note that corrections have been applied to experimental data to obtain the third and fourth "observables," as described in the text. The $\langle$NNLO$\rangle_{\text{ppd}}$ results were obtained by averaging over the Bayesian posterior predictive distribution (ppd) for the EFT predictions; see Eq. (24) and Fig. 4. The 68% highest posterior density (HPD) credible regions of the NNLO predictions are shown in the ΔNNLO$_{\text{ppd}}$ column.

	LO	NLO	$\langle$NNLO$\rangle_{\text{ppd}}$	Experiment	Adopted uncertainty	ΔNNLO$_{\text{ppd}}$
$E(^3$H) (MeV)	−5.65	−8.38	−8.52	−8.482 [45]	0.015	[−8.613, −8.453]
$E(^4$He) (MeV)	−24.08	−30.21	−28.19	−28.296 [46]	0.005	[−28.670, −27.853]
$r(^4$He) (fm)	1.27	1.33	1.45	1.4552(62) [47]	0.0062	[1.4414, 1.4634]
$fT_{1/2}$ (s)			1127.3	1129.6(3.0) [48]	3.0	[1109.1, 1150.9]

a term from the emulators, but we do not do that here as those errors are negligible.

If the truncation error parameters $\bar{c}$ and Q appearing in Σ_{th} are known from prior information then this likelihood, together with priors on $\vec{a}$, defines the posterior probability density to be computed. Although there is some evidence that suggests $Q \approx 0.3$ [25–27], we use uninformative assumptions so as not to bias our results unnecessarily. We handle this by treating $\bar{c}$ and Q as additional random variables; that is, we assign priors to them and learn their posterior distributions in tandem with the LECs.

The full joint pdf for all these parameters of interest then follows from Bayes' theorem:

$$\text{pr}(\vec{a}, \bar{c}^2, Q \,|\, \mathbf{y}_{\text{exp}}, I) \propto \text{pr}(\mathbf{y}_{\text{exp}} \,|\, \vec{a}, \Sigma, I)\text{pr}(\vec{a} \,|\, I)$$
$$\times \text{pr}(\bar{c}^2 \,|\, Q, \vec{a}, I)\text{pr}(Q \,|\, \vec{a}, I), \quad (4)$$

where the distributions for $\bar{c}^2$ and Q are explained in Sec. II D. We obtain the left-hand side of Eq. (4) using MCMC sampling. It is then simple to look at projections of these samples for the set of variables one is interested in. This is equivalent to integrating out (or marginalizing over) the other parameters. This allows us to compute a posterior for c_D and c_E without assuming that the NN LECs or the truncation error parameters are known in advance. The prior information I that determines the factors in Eq. (4) other than the likelihood, i.e., the prior pdfs, will be discussed in the next subsections.

C. Priors for the NN and 3N LECs

The prior information I includes NN scattering data, specific values of the πN LECs, and naturalness for c_D and c_E. The prior on $\vec{a} \equiv \{c_D, c_E, \vec{a}_{\text{NN}}\}$ then factorizes into a prior on the NN LECs, $\vec{a}_{\text{NN}}$, and one on the 3NF LECs, c_D and c_E:

$$\text{pr}(\vec{a} \,|\, I) = \text{pr}(c_D, c_E \,|\, I)\text{pr}(\vec{a}_{\text{NN}} \,|\, I), \quad (5)$$

$$\text{pr}(c_D, c_E \,|\, I) = \mathcal{N}[0, \bar{a}^2], \quad (6)$$

$$\text{pr}(\vec{a}_{\text{NN}} \,|\, I) = \mathcal{N}[\mu_{\text{NN}}, \Sigma_{\text{NN}}]. \quad (7)$$

The bespoke analysis of NN data described in Sec. II E produces a Gaussian posterior that is our prior on $\vec{a}_{\text{NN}}$ for this 3NF analysis. We denote the mean and covariance matrices obtained in Sec. II E by μ_{NN} and Σ_{NN}. We adopt a Gaussian

for the 3NF LEC prior [28,29]. Its width is chosen as $\bar{a} = 5$. We have found that this value of $\bar{a}$ is sufficiently large that it does not meaningfully impact our full results [10].

A fit in which the πN LECs were also constrained by NN and few-body data could be described using the same formalism, by expanding the vector $\vec{a}$ so that it includes the three πN LECs that appear in the NN potential.

D. Priors for the truncation-error parameters

We now develop the pdf $\text{pr}(\bar{c}^2, Q \,|\, \vec{a}, I)$ that enters in Eq. (4). This distribution is obtained from two distinct sources of information: (1) the order-by-order pattern of terms in the EFT expansion—knowledge of which is implicit in the conditioning on $\vec{a}, I$—and (2) the prior information on $\bar{c}^2, Q$. If there were no reliable convergence pattern, or if we happened to be fitting an EFT at leading order, then this pdf would simply reduce to the prior on $\bar{c}^2, Q$. For a detailed explanation of this approach, see [23], the appendices in particular.

Let us begin with a description of how the convergence pattern for $\mathbf{y}_{\text{th}}$ enters our analysis. Again, $\mathbf{y}_{\text{th}}$ consists of the mass and radius of ^{4}He, and the mass and β-decay rate of ^{3}H. For the LO-NLO correction, we have in principle the results in Table I. However, the shift from LO to NLO in the nuclear binding energies is large, being 100% of the LO value in many cases. This is because these states are weakly bound, i.e., $\langle T \rangle$ and $\langle V^{(0)} \rangle$ are each much larger in size than the energy, E, of the ^{3}H or ^{4}He eigenstate. Therefore while $\langle V^{(2)} \rangle \ll \langle V^{(0)} \rangle$, in accord with χEFT counting, $\langle V^{(2)} \rangle$ can still be a sizable fraction of the leading-order eigenenergy, $E^{(0)}$. However, $E^{(3)} - E^{(2)} \approx \langle V^{(3)} \rangle$, therefore the NNLO shift of the eigenenergy should provide information on the expansion parameter. Since the radii of weakly bound states are correlated with the distance they lie from the nearest particle-removal threshold [30–33] the (large relative) shift in that observable at NLO also does not give straightforward information on the convergence of the χEFT expansion.

Meanwhile, the convergence pattern for $fT_{1/2}$ of ^{3}H is irregular: $fT_{1/2}$ receives zero correction at relative order Q, while the one-body-operator corrections at $O(Q^2)$ produce a <1% effect. But a significant alteration to the LO result comes when two-body axial currents appear at $O(Q^3)$ [16]. The statistical model employed here assumes a regular order-by-order convergence of observables. More work, e.g.,

a simultaneous treatment of corrections to the trinucleon wave function and the Gamow-Teller operator, is needed to understand why $fT_{1/2}$ does not have such a convergence pattern. But, in the meantime, the order-by-order behavior of $fT_{1/2}$ is not consistent with our statistical model, so we do not use it to develop the pdf $\mathrm{pr}(\bar{c}^2, Q \,|\, \vec{a}, I)$. The information on the χEFT convergence pattern that goes into that pdf is from a subset of the order-by-order χEFT predictions: the NLO-NNLO corrections for $E(^4\mathrm{He})$, $r(^4\mathrm{He})$ and $E(^3\mathrm{H})$.

Now the NLO-NNLO corrections for these observables vary with $\vec{a}$. The NLO observable calculations are performed using the NLO optimum for the NN LECs (see Sec. II E). But the NNLO observables depend on $\vec{a}$, so we should infer the NNLO LECs and the truncation error parameters simultaneously during the NNLO fit.

We follow [23] and use a scaled inverse chi squared distribution for the prior $\mathrm{pr}(\bar{c}^2 \,|\, I) \sim \chi^{-2}[\nu_0, \tau_0^2]$. This pdf depends on two hyperparameters ν_0 and τ_0^2 that are chosen at the beginning of the analysis. (The prior pdfs we take for $\bar{c}^2$ and Q are shown as the blue lines in Fig. 5.) Because this is a conjugate prior, the posterior distribution is obtained analytically as

$$\mathrm{pr}(\bar{c}^2 \,|\, Q, \vec{a}, I) = \chi^{-2}[\nu, \tau^2(\vec{a}, Q)]. \tag{8}$$

The updating formulas for the hyperparameters are [23]

$$\nu = \nu_0 + N_{\mathrm{obs}} n_c, \tag{9}$$

$$\nu \tau^2(\vec{a}, Q) = \nu_0 \tau_0^2 + \sum_{n,i} c_{n,i}^2(\vec{a}, Q), \tag{10}$$

where i indexes the N_{obs} observables, n indexes the n_c lower order coefficients used to estimate the truncation error, and the observable coefficients are given by

$$c_{n,i}(\vec{a}, Q) = \frac{y_n^{(i)}(\vec{a}_{(n)}) - y_{n-1}^{(i)}(\vec{a}_{(n-1)})}{y_{\mathrm{ref}} Q^n}. \tag{11}$$

The notation $\vec{a}_{(n)}$ describes the LECs found at the nth order EFT fit. For NNLO, these are the $\vec{a}$ that are varied in the fit, whereas for NLO these are fixed at the optimum from the NLO fit.

With these updated hyperparameters in hand we can then obtain the unnormalized Q posterior:

$$\mathrm{pr}(Q \,|\, \vec{a}, I) \propto \frac{\mathrm{pr}(Q \,|\, I)}{\tau^\nu \prod_n Q^{N_{\mathrm{obs}} n}}. \tag{12}$$

The fact that Eq. (12) is unnormalized would not usually be a problem for estimating Q. But the normalization factor depends on $\vec{a}$ because the NLO-NNLO correction depends on $\vec{a}$. The set of LECs $\vec{a}$ is the quantity we are trying to estimate in Eq. (4), so we must be careful to include this factor. We quickly normalize Eq. (12) at each MCMC step by precomputing 70 Gaussian quadrature locations Q_i and weights. Additional speedup is realized by parallelizing the calls to Eq. (12) across the Gaussian points Q_i.

The last ingredient we need is then the prior $\mathrm{pr}(Q \,|\, I)$ that goes into the convergence-pattern analysis. To formulate that we note

(i) Q is restricted to the range $(0, 1)$;

(ii) for properties of low-energy bound states we expect χEFT to converge with Q less than $1/2$ [25–27].

We encode this as a weakly informative beta distribution $B(Q \,|\, a = 3, b = 5)$, which provides a slight bias towards $Q < 0.5$ and has support only for $Q \in (0, 1)$.

With $\mathrm{pr}(Q \,|\, \vec{a}, I)$ and $\mathrm{pr}(\bar{c}^2 \,|\, Q, \vec{a}, I)$ in hand the desired pdf is straightforwardly obtained via the product rule for conditional probabilities:

$$\mathrm{pr}(\bar{c}^2, Q \,|\, \vec{a}, I) = \mathrm{pr}(\bar{c}^2 \,|\, Q, \vec{a}, I)\mathrm{pr}(Q \,|\, \vec{a}, I). \tag{13}$$

E. Prior for NN LECs from NN scattering data

We acquire values for the NN sector LECs $\vec{a}_{\mathrm{NN}}$ at LO, NLO, and NNLO by performing a new fit to np and pp scattering data in the $0 < E \leqslant 290$ MeV range gathered from the Granada 2013 database [34,35]. As the LEC $\widetilde{C}_{1S0}^{nn}$ is unconstrained by the scattering data we also include the empirical $^1S_0$ nn scattering length $a_{nn}^N = -18.95 \pm 0.40$ fm and effective range $r_{nn}^N = 2.75 \pm 0.11$ fm [36]. The optimization procedure maximizes the likelihood function defined in Eq. (3). Fully specifying the likelihood requires us to pick values for $\bar{c}$ and the NN observable expansion parameter $Q_{\mathrm{NN}}(p)$; these are set to $\bar{c} = 1$ and $Q_{\mathrm{NN}}(p) = \max(m_\pi, p)/\Lambda_b$, where m_π is the pion mass, p is the center-of-mass momentum of the NN system, and $\Lambda_b = 600$ MeV. A set of reference values y_{ref} are also required, for which we use the experimental values.

Three πN LECs (c_1, c_3, and c_4) enter at NNLO. While these LECs could in principle be determined in the same way as the NN LECs, a more precise determination is possible by performing a Roy-Steiner analysis of πN scattering data [37,38]. Here we keep the c_i's fixed to the central values from a Roy-Steiner analysis performed by Siemens *et al.* [39] as we focus on the uncertainties from the NN sector. The covariance matrix for the πN LECs provided in Ref. [39] could straightforwardly be included as prior information in Eq. (7), provided the cross-correlation between the πN and NN LECs were known. The fixed values of the c_i's are shown in Table II in Appendix B.

The result of an optimization can (and usually does) depend strongly on the choice of starting point $\vec{a}_0$. A previously found optimum—produced by performing a fit to phase shifts using POUNDerS [40,41] optimization—serves as a basis for choosing a starting point. We choose $\vec{a}_0$ by randomly perturbing a subset of the previously found parameter values.

With the setup complete we run the optimization using the first-order Levenberg-Marquardt algorithm. This is repeated 600 times using different starting points. One or more candidate optima are chosen and used as starting points to the second-order Newton-CG method, which increases the precision of the found optimum. The final optimum is then chosen as the set of LECs $\vec{a}_{\mathrm{NN}}^*$ which produces the maximum likelihood value. The resulting values for the LECs agree well with findings from similarly regulated potentials [36,42] and are shown in Table II.

To estimate the covariance matrix Σ_{NN} of the NN LECs $\vec{a}_{\mathrm{NN}}$ we follow the method detailed by Carlsson *et al.* in Sec. II G of Ref. [43]. The resulting Gaussian pdf (7) is shown in green in Fig. 2. The Hessian needed to compute the covariance

matrix, and the first- and second-order derivatives used by the optimization algorithms, are computed to machine precision using automatic differentiation [44].

III. FEW-NUCLEON-PHYSICS IMPLEMENTATION

The likelihood in Eq. (3) is centered at the model predictions $\mathbf{y}_{th}$ for few-nucleon ($A = 3, 4$) observables. To make those predictions we apply the no-core shell model (NCSM) [49] in a relative-coordinate harmonic-oscillator (HO) basis and solve the few-nucleon Schrödinger equation with two- and three-nucleon interactions employing the isoscalar approximation as presented in Ref. [50]. The model-space dimension is determined from the truncation in total number of HO excitations N_{max}. The eigenenergy of the resulting Hamiltonian matrix is a variational estimate of the total binding energy while the eigenfunctions can be used to obtain other observables.

We obtain converged ground-state observables using $\hbar\omega = 36$ MeV and $N_{max} = 40(18)$ for $A = 3(4)$ since we employ a rather soft chiral interaction at NNLO. Specifically we use a nonlocal momentum-space regulator function as in Eqs. (5) and (6) of Ref. [43] with cutoff $\Lambda = 450$ MeV and $n = 3$. For ^{4}He we obtain ground-state energies and point-proton radii that are converged within $\lesssim 5$ keV and $\lesssim 0.002$ fm compared to larger-basis calculations.

A. Few-nucleon observables of interest

The first two observables we consider are the binding energies of ^{3}H and ^{4}He. Determining these from precisely known masses yields errors on the binding energies of a few eV or less. This is negligible compared to errors from the method used to calculate the bound states. Therefore in Table I we take "adopted errors" for these two observables of 15 keV (width of the 68% credibility interval given the 20 keV accuracy of the isoscalar approximation for the ^{3}H binding energy quoted in Ref. [50]) and 5 keV (NCSM basis truncation) respectively. Ultimately, both of these are dwarfed by the truncation error.

We also compute the point-proton radius, here denoted r, for ^{3}H and ^{4}He and relate it to the measured charge radius via [51]

$$r^2 = r_{ch}^2 - r_p^2 - \frac{N}{Z}r_n^2 - r_{DF}^2 - \Delta r^2, \qquad (14)$$

where r_p^2 (r_n^2) is the proton (neutron) mean-squared charge radius, Z (N) is the proton (neutron) number, and $r_{DF}^2 = 3\hbar^2/(4M_p^2) \approx 0.033$ fm^2 is the Darwin-Foldy correction [52]. There are two-body-current and further relativistic corrections to $r(^4$He) at orders beyond NNLO in χEFT, but these are accounted for by the truncation uncertainties in our likelihood, so we set $\Delta r^2 = 0$. We use $r_p = 0.8783(86)$ fm and $r_n^2 = -0.1149(27)$ fm^2 [47]. We do not use the ^{3}He binding energy or point-proton radius for inference because they are highly correlated with the corresponding ^{3}H observables: 100% correlated in the limit of isospin-symmetric interactions.

Furthermore, we use the triton half-life to provide a constraint on the nuclear force from an electroweak observable. We follow the approach by Gazit *et al.* [15] and compute the

triton half-life from the reduced matrix element for E_1^A, the $J = 1$ electric multipole of the axial-vector current:

$$\langle E_1^A \rangle \equiv \left| \langle ^3\text{He} \| E_1^A \| ^3\text{H} \rangle \right|. \qquad (15)$$

Due to the χEFT link between electroweak currents in nuclei and the strong interaction dynamics [15,53,54], this matrix element has a term proportional to c_D, the LEC that also determines the strength of the one-pion-exchange plus contact interaction diagram of the 3NF. (Note, though, that Krebs has recently pointed out that this connection is broken at subleading order by commonly used regulation procedures [55].) The experimental value for the comparative half-life, $fT_{1/2} = 1129.6 \pm 3$ s [48],[1] leads to an empirical value for $\langle E_1^A \rangle = 0.6848 \pm 0.0011$ [15] via the relation

$$fT_{1/2} = \frac{K/G_V^2}{(1 - \delta_c) + 3\pi (f_A/f_V)\langle E_1^A \rangle^2}, \qquad (16)$$

with $K/G_V^2 = 6146.6 \pm 0.6$ s, $f_A/f_V = 1.00529$, and the isospin-breaking correction $\delta_c = 0.13\%$.

Results at LO, NLO, and NNLO for these four $A = 3, 4$ observables, together with the experimental numbers, are given in Table I. The NNLO results in this table are the mean values obtained from the posterior predictive distribution; see Eq. (24) and Fig. 4.

B. Efficient emulators for few-nucleon observables

Although we are studying $A = 3, 4$ systems using soft interactions, the matrix representations of the NCSM Hamiltonians for the few-nucleon states that we analyze still reach dimensions of approximately $10^4 \times 10^4$. With a Lanczos algorithm it takes about one minute, using a single CPU, to obtain the energies and corresponding wave functions for the systems of interest. It takes a few hours of computation on a single node to fully sample the posterior pdf $\text{pr}(c_D, c_E \mid \mathbf{y}_{exp}, I)$.

To enable more rapid iterations of our exploratory data analysis, we employ eigenvector continuation (EC) [12] to efficiently and accurately emulate [13] the $\vec{a}$ dependence of the few-nucleon observables listed in Table I. The high accuracy achieved is demonstrated by the smallness of the differences between the emulator and the NCSM result; see Fig. 1. The evaluation of the posterior is dramatically accelerated via the EC emulators such that each parameter sample only takes ≈ 10 ms on a single-threaded CPU with a corresponding speed-up for sampling the relevant parameter space of LECs. In addition, the construction of a set of model-specific emulators allows others to easily reproduce, and modify, our statistical analysis.[2]

---

[1]Reference [16] uses the value $fT_{1/2} = 1134.6 \pm 3$ s, obtained from Simpson's tritium β-decay measurement [56]. The difference between the two $fT_{1/2}$ numbers is larger than the stated error in either. Here we select the Akulov-Mamyrin result, but the tools we have developed and provide make it straightforward to re-do the analysis using either the Simpson value or a compromise $fT_{1/2}$ with an error inflated so it is large enough to accommodate both results.

[2]The NCSM emulators and statistical models can be obtained or created via our open-source Python package `fit3bf` [20].
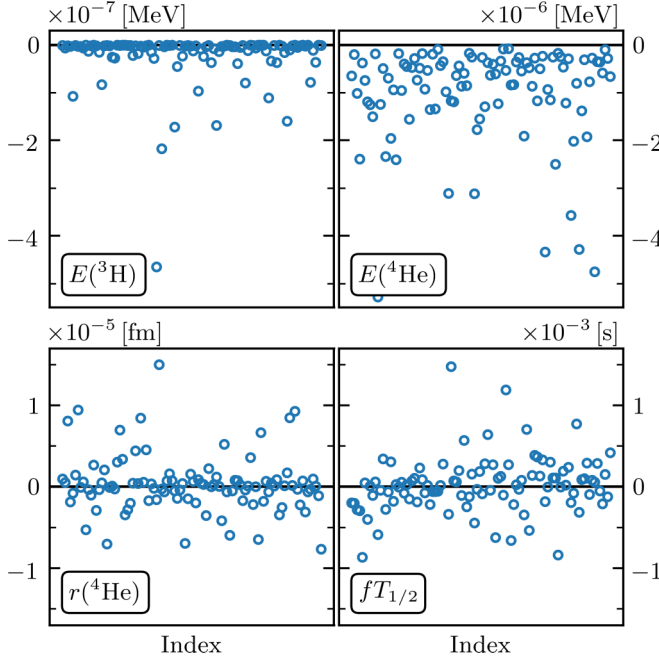
FIG. 1. Differences between evaluations using the Schrödinger-equation solution and those using the EC emulator for the four observables of interest at 100 validation points. These differences are several orders of magnitude smaller than the adopted errors in Table I. Note that the ground state energies from the emulator are guaranteed to be an upper bound on the exact energies, but the other observables have no such constraint.

The EC approach to emulation is described in Ref. [13]; to be self-contained, we briefly outline the method here. Consider one quantum system that we want to emulate, such as the triton. The A-nucleon Schrödinger equation can be written as

$$H(\vec{a})|\psi(\vec{a})\rangle = E(\vec{a})|\psi(\vec{a})\rangle, \tag{17}$$

where $|\psi(\vec{a})\rangle$ and $E(\vec{a})$ denote the ground-state and its energy, and the implicit $\vec{a}$ dependence has been brought forward. We then diagonalize the Hamiltonian $H(\vec{a})$ for N_{EC} different values of $\vec{a}$, and collect the N_{EC} ground-state wave functions $|\psi_i\rangle$ into a matrix X,

$$X \equiv \begin{pmatrix} | & | & & | \\ |\psi_1\rangle & |\psi_2\rangle & \cdots & |\psi_{N_{EC}}\rangle \\ | & | & & | \end{pmatrix}, \tag{18}$$

which *does not* depend on $\vec{a}$. Then we project the Hamiltonian to a subspace spanned by the N_{EC} wave functions via

$$\widetilde{H}(\vec{a}) = X^\dagger H(\vec{a})X. \tag{19}$$

Because the chiral Hamiltonians H that we use depend linearly on $\vec{a}$, this projection can be performed once for each term and stored to quickly construct $\widetilde{H}(\vec{a})$.

To construct an emulator for $|\psi(\vec{a})\rangle$ and $E(\vec{a})$, we solve the $N_{EC} \times N_{EC}$ generalized eigenvalue equation

$$\widetilde{H}(\vec{a})\beta(\vec{a}) = \widetilde{E}(\vec{a})\mathcal{N}\beta(\vec{a}), \tag{20}$$

where $\mathcal{N} = X^\dagger X$ is the norm matrix with elements $\mathcal{N}_{ij} = \langle\psi(\vec{a}_i)|\psi(\vec{a}_j)\rangle$. The generalized eigenvalue $\widetilde{E}$ is an approximation to the true eigenenergy. The length-N_{EC} vector of coefficients $\beta(\vec{a})$ found by solving Eq. (20) could then be used to reconstruct the approximate wave functions via $|\psi(\vec{a})\rangle \approx X\beta(\vec{a})$, but these are not needed in practice. Instead, to evaluate expectation values of observables $\hat{O}$ other than nuclear spectra, one computes

$$\langle\hat{O}(\vec{a})\rangle = \langle\psi(\vec{a})|\hat{O}(\vec{a})|\psi(\vec{a})\rangle$$
$$\approx \beta(\vec{a})^\dagger[X^\dagger\hat{O}(\vec{a})X]\beta(\vec{a}). \tag{21}$$

If $\hat{O}(\vec{a})$ is linear in $\vec{a}$ then the terms in $X^\dagger\hat{O}(\vec{a})X$ can again be computed once and stored prior to sampling. For the β-decay transition, we generalize Eq. (21) to the case where the right and left $X\beta(\vec{a})$ come from the initial- and final-state emulators, respectively. This is the first application of EC emulation to a nuclear transition.

It was shown in Ref. [13] that $\widetilde{E}$ approximates E extremely well even with a small number of training vectors. Although the Hamiltonian eigenvector originally resides in a Hilbert space of very large dimension, the eigenvector trajectory produced by continuous changes of the Hamiltonian matrix can be accurately represented in a space of very low dimension. For this reason we can construct fast and accurate emulators for all observables that we study, including the β-decay transition (see Fig. 1).

As already noted, to construct a *computationally efficient* EC emulator requires that we can write the subspace-projected Hamiltonian as a linear combination of the continuous parameters that we are interested in. For example, considering only the c_D and c_E dependence, we can express the chiral NNLO Hamiltonian as

$$H(c_D, c_E) = H^{(\text{const})} + c_D V^{(1\pi\text{-ct})} + c_E V^{(3\text{N-ct})}, \tag{22}$$

where we partitioned the Hamiltonian into three pieces: all contributions that are constant with respect to variation of c_D and c_E ($H^{(\text{const})}$), the one-pion-exchange plus contact (1π-ct) interaction between three nucleons, and the pure three-nucleon contact (3N-ct). Having obtained N_{EC} linearly independent training vectors $|\psi_i\rangle$ for each state of interest, we construct each subspace-projected matrix [denoted with tildes as in (19)] in

$$\widetilde{H}(c_D, c_E) = \widetilde{H}^{(\text{const})} + c_D \widetilde{V}^{(1\pi\text{-ct})} + c_E \widetilde{V}^{(3\text{N-ct})} \tag{23}$$

only once prior to sampling, which greatly speeds up the subsequent matrix algebra. Equipped with the subspace basis, we can also project the operators for the point-proton radius of ^{4}He and triton β decay.

When sampling over both NN and 3N LECs (for a total of 13 dimensions), we use $N_{EC} = 50$ training points. For the 3NF LECs, we simply use a Latin hypercube design in the range $[-5, 5]$. For the NN LECs, we start with a Latin hypercube design in the range $[-1, 1]$. We then map each training point p_i to the plausible range of LECs according to $\mu_{NN} + \Sigma_{NN}^{1/2} p_i$ with μ_{NN} and Σ_{NN} the mean and covariance determined in Sec. II E. The resulting set of points form our training LECs, and are displayed in Fig. 2.
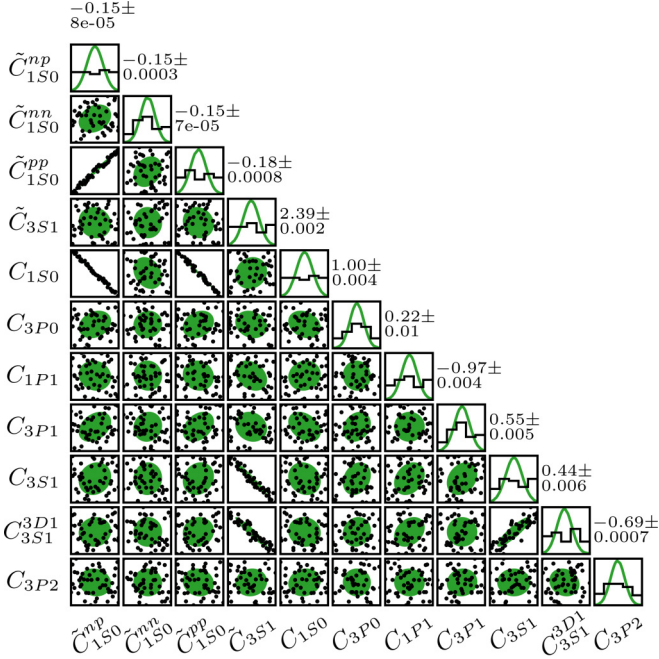
FIG. 2. The EC training points in the space of NN LECs compared to the input priors for our few-body analysis. The green curves represent the marginal prior distributions, with the ellipses denoting 95% credibility regions. Each dot is a training point; the marginals are approximately uniformly distributed. The values quoted above the diagonal are means and standard deviations from Table II.
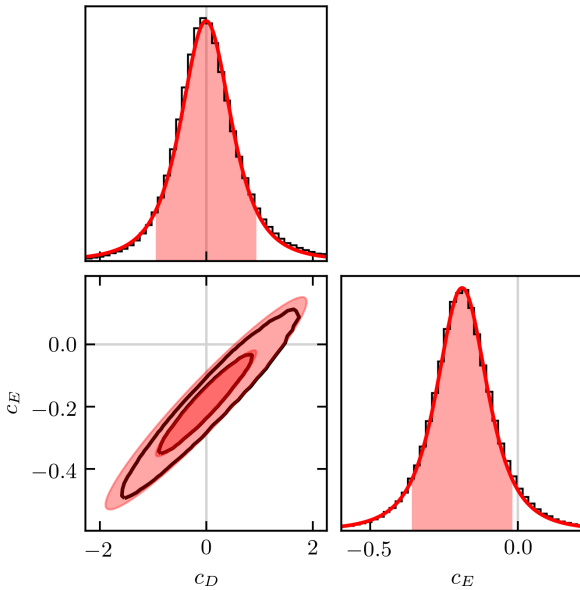

FIG. 3. The posterior of c_D and c_E fitting to all four few-body observables and marginalizing over $\bar{c}^2$, Q, and the NN LECs. The black histograms and contours correspond to the pure MCMC samples. The red curves and ellipses follow from a fit of a multivariate t distribution $t_\nu(m, S)$ as described in the text. Filled areas in the marginals denote one standard deviation of the fit distribution, which contains 86% of the probability mass, not 68% like a Gaussian. Contours represent the one and two standard deviations of the best fit $t_\nu(m, S)$.

IV. BAYESIAN PARAMETER ESTIMATION FOR $c_D, c_E, Q,$ AND $\bar{c}$

Figure 3 shows the joint posterior for c_D and c_E as obtained from MCMC sampling of the full posterior (4). This LEC posterior has been marginalized over $\bar{a}_{NN}$ as well as the truncation error parameters $\bar{c}$ and Q. The evaluation was done using fixed $\bar{a} = 5$, although the final posterior is concentrated so close to zero that $\bar{a}$ could be taken to larger values without influencing the results. The data likelihood (3) contains the four few-nucleon observables listed in Table I. We sample the posterior using the affine invariant MCMC ensemble sampler emcee [57] using 50 walkers with 50 000 iterations per walker following 2000 warmup steps.

The joint distribution in Fig. 3 is best represented by a multivariate t distribution. The emergence of a t distribution is a generic feature of statistics problems that are linear in the parameters and involve variance estimation—as explained in Appendix A—and the linear correlations seen in Fig. 7 strongly support that this problem is approximately linear in c_D and c_E. None of this is surprising: the 3NF is a perturbative correction in χEFT and the values of c_D and c_E that turn out to be relevant are small. (For another recent discussion of the benefits of a perturbative treatment of c_D and c_E see Ref. [58].)

We fit a parametrized distribution to the c_D, c_E samples by maximizing their likelihood given that they are multivariate t distributed $t_\nu(m, S)$. The best fit is obtained with $\nu \approx 2.8$ degrees of freedom, a mean vector $m = [-0.0047 \quad -0.1892]$, and scale matrix of

$$S = \begin{bmatrix} 0.250 & 0.043 \\ 0.043 & 0.008 \end{bmatrix}.$$

This yields an accurate description of the one-dimensional c_D and c_E posteriors and of their joint pdf at one standard deviation. The two standard deviation contour in the two-dimensional LEC pdf is harder to match. This distribution has moderately heavy tails; a Gaussian is not a good approximation.

The parameters c_D and c_E are strongly correlated. The covariance matrix is $\nu S/(\nu - 2)$, corresponding to a correlation coefficient $\rho \approx 0.96$. The strength of this correlation is similar to what was found in Baroni *et al.* [16] and Kravvaris *et al.* in Ref. [19]. In contrast, in Ref. [18] Epelbaum *et al.* employed SCS potentials and found the triton-binding-energy constraint led to c_D and c_E being anticorrelated. The way that this correlation is connected to the wave function of the three-nucleon system and the short-distance behavior of the NN force is an interesting subject for future study.

The consistency of our parameter estimation can be assessed by studying the model posterior predictive distribution (ppd)

$$\text{ppd} = \{\mathbf{y}_{th}(\vec{a}) : \vec{a} \sim \text{pr}(\vec{a} \,|\, \mathbf{y}_{exp}, I)\}. \tag{24}$$

The ppd is the set of all predictions computed over likely values of the LECs, i.e., drawing from the posterior pdf for $\vec{a}$. Figure 4 shows the ppd for the target few-nucleon observables, evaluated from the full posterior (4). In practice, the ppd is evaluated via sampling and we use the MCMC samples of the full posterior for this purpose. The four target experimental values are within one standard deviation for
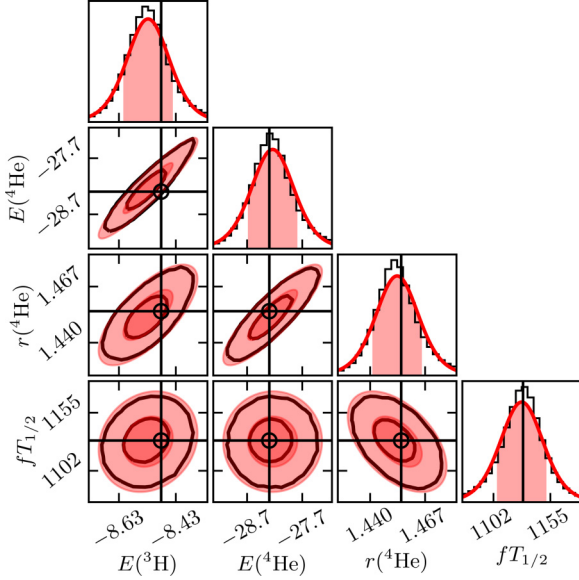
FIG. 4. The posterior predictive distribution from sampling over the LECs found in Fig. 3, with units as in Table I. The red distributions come from a fit of a multivariate t distribution to the data (see Appendix A). The filled regions of the one-dimensional plots represent one standard deviation of the marginal t distributions. The filled contours of the joint distributions denote the one- and two-standard-deviation regions of the multivariate t, and the black contours denote the corresponding HPD regions from the samples. The markers and black horizontal and vertical lines denote the experimental values.
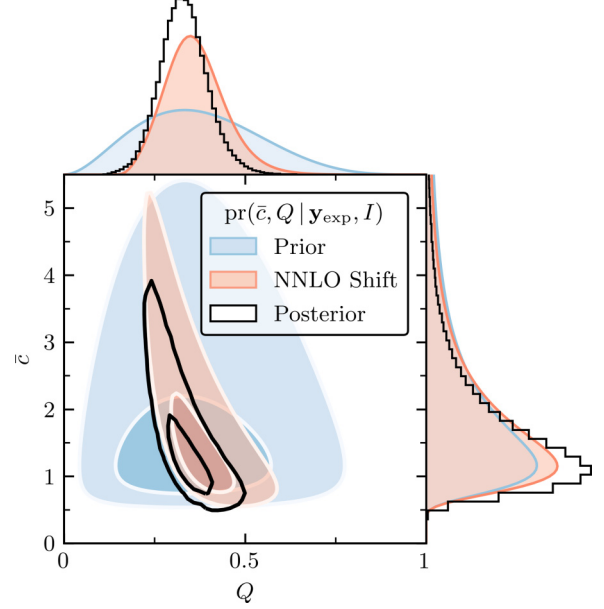

FIG. 5. The prior and posterior distributions for $\bar{c}$ and Q. The blue line denotes the uncorrelated prior distribution with $\bar{c}^2 \sim \chi^{-2}[\nu_0 = 1.5, \tau_0^2 = 1.5^2]$ and $Q \sim B(a = 3, b = 5)$. The black posterior is obtained by conditioning on the NLO-NNLO shift at each $\vec{a}$ value in the sampler. It also folds in information about the size N^3LO effects need to have to yield agreement with the data. From this we obtain $Q = 0.33 \pm 0.06$. If we instead updated from the prior to the posterior via the mean value for the shift obtained from the fit, then we would have obtained the red curve.

all of the marginals, while all but one pair of values are within one-standard-deviation regions for the bivariate joint distributions. For the ^{3}H - ^{4}He joint distribution the target is instead within the two-standard-deviation region. We reiterate that the probability mass enclosed in these intervals does not correspond to Gaussian intervals due to the heavy tails of the distribution.

Because we simultaneously sample the 3NF LECs and the parameters associated with our model for truncation errors, we also have access to the (joint) posterior for those parameters, Q and $\bar{c}$. This posterior is shown in Fig. 5 as the black histogram. It should be compared to the prior distribution represented by the blue curve and described in Sec. II D. Both the NLO to NNLO shift in observables and the discrepancies with data of the NNLO χEFT predictions inform the pdf for $\bar{c}$ and Q. Together, the constraints yield $Q = 0.33 \pm 0.06$, which is an uncertainty of about 20%. An ongoing analysis by the LENPIC Collaboration suggests $Q = \{m_\pi\}_{\text{eff}}/\Lambda_b$ with $\{m_\pi\}_{\text{eff}} \approx 200$ MeV and $\Lambda_b \approx 600$–650 MeV, a very similar value for Q [25–27]. The preferred values of $\bar{c}$ are of order 1: the one-dimensional 68% Bayesian credible interval is $\bar{c} \in [0.87, 1.44]$. This validates the naturalness assumptions encoded in the truncation-error model. There is a nonlinear correlation between $\bar{c}$ and Q, presumably because the pattern of EFT convergence constrains the combinations $\bar{c}Q^3$ (from NLO to NNLO) and $\bar{c}Q^4$ (NNLO uncertainties).

If one could assume reasonable values for $\vec{a}$ *a priori*, then an alternative approach to this evaluation of the Q-$\bar{c}$ posterior via sampling is to use the mean value of the shift in observ-

ables from NLO to NNLO to update the pdf for $\bar{c}$ and Q; see Eqs. (8)–(12). Updating using the mean values from the ppd and the NLO numbers in Table I yields the red curves for Q and $\bar{c}$ in Fig. 5. These differ from the sampling results in two ways. First, in the sampling results the NLO-to-NNLO shift is computed for each sample separately. The value of c_3, and hence that of τ^2 and $\bar{c}^2$, depends on $\vec{a}$, and so is different for each member of the MC Markov chain. However, since the ppd of all the observables that inform the convergence pattern is quite narrow, this $\vec{a}$ dependence is a small effect. The samples in Fig. 5 also account for the requirement that the sizes of the NNLO errors are statistically consistent. The combination $\bar{c}Q^4$ determines the variance of our NNLO pdfs. Incorporating NNLO variance estimation in our $\bar{c}$-Q estimate brings the central value of Q down slightly compared to what is obtained if only the NLO-to-NNLO shift in observables is considered.

If truncation errors are not included in the analysis then the individual constraints from all four observables disagree by several σ, see Fig. 6(a), where NN LECs are also held fixed.[3] Consequently, obtaining a posterior with $\Sigma_{\text{th}} = 0$

---

[3]In the absence of a prior these posteriors extend very far in both directions, since the problem is approximately linear and each band represents the constraint on two parameters from one datum. But the prior on $\vec{a}$ [Eq. (6)] regulates these one-dimensional structures once values of c_D and $c_E \approx \vec{a}$ are reached.
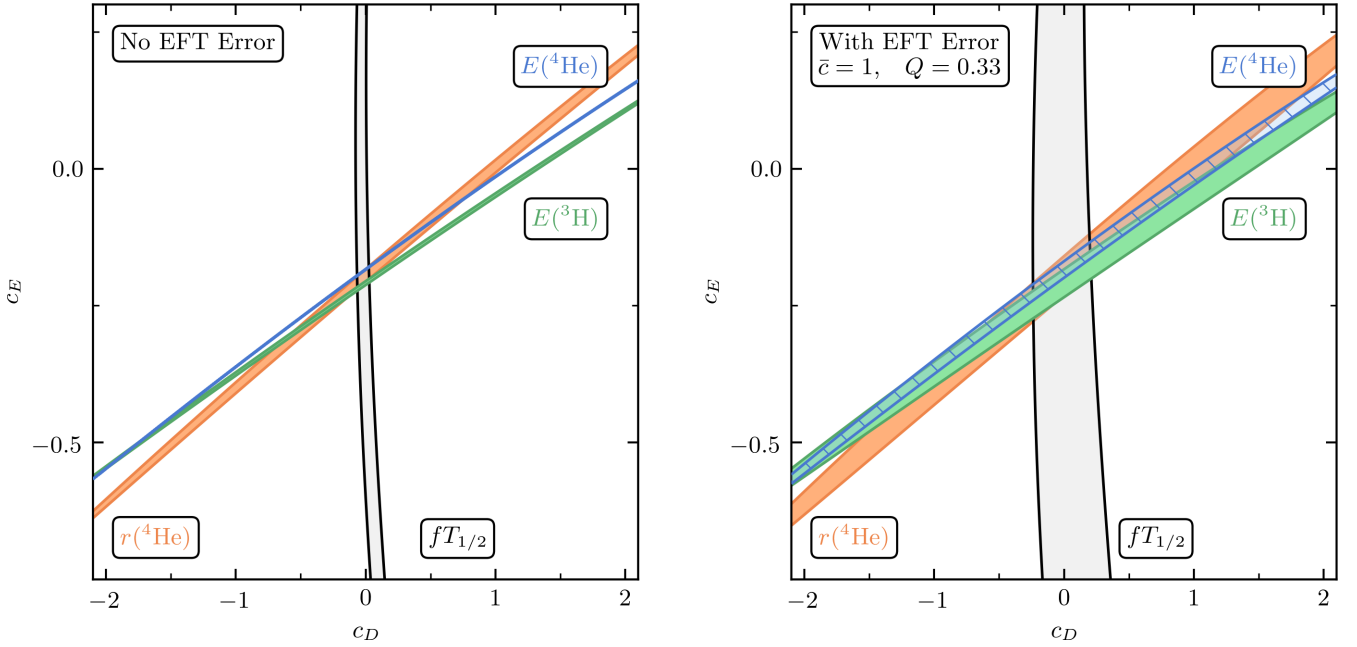
FIG. 6. Constraints on c_D and c_E from single-observable fits both without (a) and with (b) EFT truncation errors included. The regions are determined by computing 39% HPD intervals, which would correspond to 1σ intervals for a two-dimensional Gaussian. The NN LECs are fixed to their prior values for these plots, as they make little difference to the overall fit. There is no mutual overlap of the truncation-error-free posteriors, which would make a simultaneous fit difficult and unreliable. On the other hand, the right-hand panel makes it clear that there is no inconsistency in the theory here, once truncation errors are accounted for. (Note that in the right-hand panel $\bar{c}$ and Q are fixed. If they were allowed to vary as in the full fit, then there is not enough information to accurately constrain these posteriors.) All but the $fT_{1/2}$ observable provide essentially identical information about the c_D, c_E fit, which makes it a crucial observable to include.

becomes both difficult and unreliable. In particular, the errors adopted for the two binding energies in Table I lead to such tight constraints that the resulting values of c_E differ by many σ, at least in the region where the $fT_{1/2}(^3\mathrm{H})$ datum is also reproduced.

The contrast when truncation errors are added to the analysis is striking; see Fig. 6(b). In this case, the constraints due to all four observables can be satisfied simultaneously. Note that we have fixed $Q = 0.33$, $\bar{c} = 1$, rather than marginalizing over Q and $\bar{c}$ as we did to obtain Fig. 3. With only one observable in the likelihood there is not enough information to determine c_D, c_E, Q, and $\bar{c}$ simultaneously. The NN LECs are also held fixed for this portion of the analysis because their effects are hardly distinguishable here. The concordance region where all four data are simultaneously reproduced is qualitatively similar to the result obtained via MCMC sampling as in Fig. 3, though fixing $\bar{c}$ and Q produces credibility intervals that are narrower than they should be, and turns tails that should be t distributed back into Gaussians.

Pairs of the triton and $^4\mathrm{He}$ binding energies and the $^4\mathrm{He}$ radius have conventionally been used in past optimizations of c_D and c_E. But Figs. 6(a) and 6(b) make it clear that all three of these observables are correlated: they do not provide complementary constraints on the 3NF LECs. The triton β-decay rate—or some other non-degenerate observable—is essential to accurate estimation of c_D and c_E [18,19,25,59]. To make this point clear Fig. 7 shows the c_D-c_E posterior for four pairs of observables (once again with $\bar{c} = 1$, $Q = 0.33$ and fixed $\bar{a}_{\mathrm{NN}}$). The one-dimensional nature of the information obtained

on the 3NF LECs under a poor choice of observable pair is most drastic for $E(^3\mathrm{H})$ and $E(^4\mathrm{He})$ (upper-left panel). These two binding energies are, of course, correlated: few-body universality predicts that once the three-body binding energy is known the four-body binding energy can be accurately predicted [60–62]. Between them $E(^3\mathrm{H})$ and $E(^4\mathrm{He})$ constrain only the combination $c_E - 0.2c_D$. Any information on the individual LECs comes only from the prior, which truncates the posterior once $|c_D| \approx 5$. The situation is almost as bad if the $E(^3\mathrm{H})$ binding energy and the $r(^4\mathrm{He})$ radius are used to constrain the 3NF (upper-right panel) (cf. the similar posterior from these two observables found in Ref. [19]).

The triton half-life constrains the value of c_D well, but leaves c_E essentially unconstrained [15]. Therefore, it provides a complementary constraint, as observed in Ref. [59] (lower-left and lower-right panels), greatly reducing the range of allowed c_D values. That in turn sharpens the estimate of c_E because of the correlation induced through an energy or radius. Using the $^4\mathrm{He}$ binding energy and the triton half-life provides essentially the same information as fitting to all four observables. Of the observables we consider, these are the two that best constrain the short-distance pieces of the 3NF. There is little additional information added by the other two observables.

V. SUMMARY AND OUTLOOK

The present work is part of an ongoing effort to develop, apply, and evaluate Bayesian statistical methods for effective
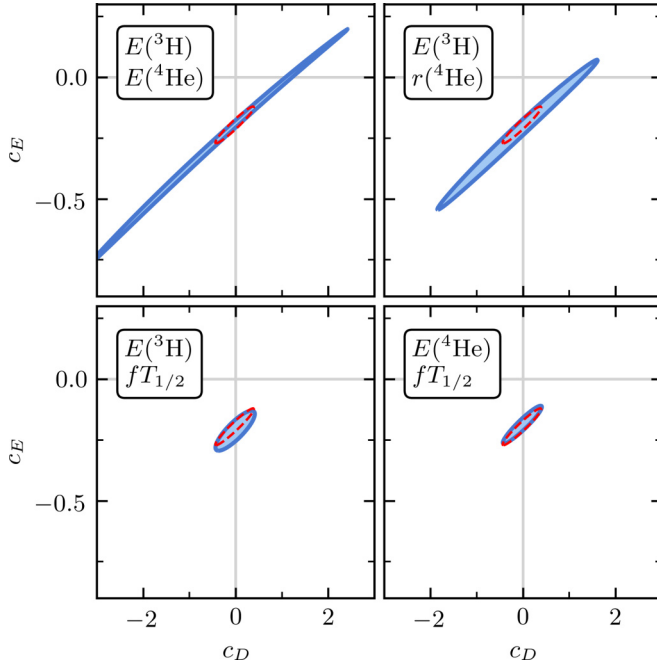
FIG. 7. Posteriors found by including only subsets of the few-body observables in the likelihood (3) while holding $\bar{a}_{\rm NN}$ and truncation parameters fixed ($\bar{c} = 1$, $Q = 0.33$). In each panel the fit to all four observables, again with fixed truncation parameters, is represented by the dashed line, and is shown for comparison. We choose to show the 39% HPD interval, which corresponds closely to 1σ because these distributions are approximately Gaussian. Posteriors conditioned on the triton half-life give particularly well-constrained 3NF parameters, whereas the others are unconstrained along a linear combination of c_D and c_E. The fit to $E(^4{\rm He})$ and $fT_{1/2}$ produces almost the same posterior as the fit to all four observables. Because the truncation parameters $\bar{c}$ and Q are fixed—for the same reasons as in Fig. 6(b)—these posteriors appear more constrained than is justified by our true prior knowledge.

field theories of nuclei. Our immediate target is the estimation of the LECs c_D and c_E that characterize short-distance effects in the leading three-nucleon force in χEFT. Performing this "fit" means finding the joint posterior distribution of these LECs given a selected set of experimental data $\mathbf{y}_{\rm exp}$ and a specification of prior information, I, namely ${\rm pr}(c_D, c_E \mid \mathbf{y}_{\rm exp}, I)$. In this analysis, I includes knowledge about the LECs as well as the χEFT truncation error model developed in Refs. [10,23]. The prior for c_D and c_E is chosen to be naturally sized, the $\bar{a}_{\rm NN}$ prior was determined from NN scattering data up to 290 MeV, and the πN LECs were fixed to the central values from the Roy-Steiner analysis. The resulting posterior is shown in Fig. 3.

We focus on how different combinations of experimental observables impact the posterior. We present results for one χEFT Hamiltonian and constrained its parameters using a set of four nuclear properties: the triton binding energy and half-life, and the ^{4}He binding energy and charge radius. We do not span all possible Hamiltonian regularization schemes and input properties. However, our Bayesian framework accounts for experimental and theoretical errors and enables

the identification of correlations and the direct propagation of uncertainties to observables. Extending the results is straightforward via our open-source Python package fit3bf [20], which can reproduce all results shown in this work.

The Bayesian strategy and the details of the statistical model are laid out in Sec. II, building on previous work. The likelihood in Eq. (3) is determined by the form of the experimental and theory uncertainties to be a multivariate Gaussian. The prior information specifies NN and πN LECs, as well as the uncertainties from the NN fit (omitting the πN uncertainties here because we do not account for correlations with NN observables). The truncation error model for the EFT has been developed and validated elsewhere. All assumptions are explicit and therefore testable.

We compute a joint posterior the LECs c_D, c_E, $\bar{a}_{\rm NN}$, and truncation error parameters $\bar{c}$ and Q. The posterior for c_D and c_E—unconditional on $\bar{a}_{\rm NN}$, $\bar{c}$, and Q—is obtained via marginalization. Sampling of such an extended joint posterior is characteristic of a full Bayesian analysis. It is made convenient and efficient here by the use of EC emulators (see Sec. III B).

Here are the takeaway points from this investigation:

(i) *For 3NF parameter estimation, do not only use observables that are related by universality.* The triton and α-particle binding energies and the ^{4}He radius provide very similar constraints on c_D and c_E because they are related by universality; see Fig. 6(b). Consequently any pair of them only determines one linear combination of the 3NF LECs; see Fig. 7. In contrast, the triton half-life provides a new constraint. When paired with the ^{4}He binding energy it essentially saturates the information available from this set of observables. These results support the previous conclusions of Lupu *et al.* [59]. It will be interesting to make similar correlation comparisons using the three-body scattering input advocated in Refs. [18,25] or the information on $n\alpha$ scattering used for 3NF LEC estimation in Refs. [19,63].

(ii) *The LECs c_D and c_E are strongly correlated.* The contours in the joint posterior manifest a correlation of $\rho \approx 0.96$ for the χEFT Hamiltonian used in this investigation, see Fig. 3. A similar degree of correlation was found by Baroni *et al.* [16] and Kravvaris *et al.* [19]. Using SCS potentials, Epelbaum *et al.* [18] also find strong correlation, but the orientation of the c_D-c_E contours in that study is opposite. Different choices of regularization scheme and scale affect the relationship between c_D and c_E, but the details of this correlation remain to be investigated.

(iii) *EFT truncation errors must be included for a complete quantification of uncertainties.* Truncation errors fuzz up the constraints from individual observables, affecting the size of credibility regions in the c_D and c_E posterior. They do not affect the correlation. This is evident in comparing single-observable fits in Fig. 6(a) (no truncation error) to those in Fig. 6(b) (including truncation error). A consistent solution for all considered observables is only obtained when

truncation errors are included; without these errors, a simultaneous fit is problematic. Similar conclusions regarding the impact of truncation errors on the c_D-c_E posterior were found using a smaller basket of $A = 3$ and $A = 4$ observables and a slightly different NN potential in Ref. [19].

(iv) *The impact of including* NN *LEC uncertainties on the* c_D-c_E *posterior is small. That of* πN *LECs remains to be assessed.* If truncation errors are included but NN uncertainties are omitted, the changes in the posterior are almost undetectable. The πN LECs were held fixed at the central values obtained in the Roy-Steiner analysis of Ref. [39]. Ideally the πN LECs c_1, c_3, and c_4 would also be included in the set of parameters being sampled, so that the impact of their uncertainties on the c_D and c_E inference could be determined, and constraints on them from $A = 3$ and $A = 4$ observables assessed. This was not feasible for the present work because the correlations between the πN and the NN LECs were not available. But our framework can accommodate the incorporation of πN LECs in the vector $\vec{a}$. This is of particular interest because those LECs appear in the leading χEFT 3NF.

(v) *The EFT expansion parameter is* $Q \approx 1/3$ *for these observables.* The distribution for Q in Fig. 5 peaks at 0.33 with a 20% uncertainty, which is consistent with general χEFT considerations for the few-body observables used and with other estimations [25–27].

(vi) χ*EFT provides a statistically consistent description of these few-body observables.* The predictions of χEFT with the LEC values learned from four few-body observables reproduce these observables to within the χEFT uncertainty. We verify this by propagating the LEC samples from MCMC sampling to the observables; the resulting posterior predictive distribution (ppd) is shown in Fig. 4. We indeed used these observables in the fit, but the ppd demonstrates that χEFT can describe all four consistently, as long as truncation errors are included in the inference for $\vec{a}$ and thereby propagated to the ppd.

(vii) *Not all distributions are Gaussian.* The joint distribution for c_D and c_E in Fig. 3 is best represented by a multivariate t distribution. Its tails are *not* well approximated by a Gaussian. In Appendix A we show why a t distribution naturally emerges for these observables.

The Bayesian framework and statistical best practices we have exemplified, together with the computational capabilities enabled by EC emulators, provide a strong foundation for future work. Full Bayesian parameter estimation and propagation of uncertainties to all calculated observables is now feasible. Future avenues for parameter estimation with $A = 3$ and $A = 4$ observables include comparing χEFT Hamiltonians with different ultraviolet regulators and with delta degrees of freedom, including πN LECs in the set of $\vec{a}$, identifying and testing complementary input observables, and applying truncation error models where the convergence pattern is correlated across observables [23,24].

APPENDIX A: LINEAR MODELS WITH VARIANCE ESTIMATION; OR, WHY THINGS LOOK t

There are two types of distributions: those that are Gaussian and those that are not. This is an Appendix about those that are not.

We will show that the t distribution emerges as the posterior for the 3NF LECs because of two key facts. First, the observables are approximately linear in c_D and c_E, and so at fixed Q and $\bar{c}$ the posterior for c_D and c_E is Gaussian. Second, when Q and $\bar{c}$ are estimated the tightest constraint on them comes from the variance in the theory covariance matrix in Eq. (2). Marginalizing over Q and $\bar{c}$ to get the c_D and c_E posterior therefore corresponds to marginalizing over the variance. In linear parameter estimation problems with variance estimation the parameters are typically t distributed, for the reasons we now articulate.

Suppose that the order k contributions to observables of interest y are linearly related to the EFT parameters $\vec{a}$ that appear at that order. This is approximately true if subleading corrections are perturbative, as long as k does not correspond to the EFT's leading order. In this situation the theoretical discrepancy due to truncation error, ϵ, will be additive: That is,

$$y_k(\vec{x}) = \vec{x} \cdot \vec{a} + \epsilon. \tag{A1}$$

If we have N $O(Q^k)$ observables that we are using to extract y

$$\mathbf{y}_k = X\vec{a} + \epsilon, \tag{A2}$$

where it is important to remember that the nuclear matrix elements X that relate the LECs to the observables must be

of $O(Q^k)$ if the power counting is to be valid.[4] We then write the truncation error as

$$\text{pr}(\epsilon \,|\, \bar{c}^2, Q) = \mathcal{N}[0, \bar{c}^2 Q^{2(k+1)}]. \quad (A3)$$

This assumes that the truncation error is the same for all observables; this assumption can be relaxed if needed by promoting Q to a matrix or including another matrix factor (say, $\mathbf{y}_{\text{ref}}\mathbf{y}_{\text{ref}}^T$; see [23]).

Further progress requires priors on $\vec{a}$ and $\bar{c}^2$. We follow Melendez *et al.* [23] and place a normal-inverse-chi-squared prior on this tuple

$$\text{pr}(\vec{a}, \bar{c}^2) = \mathcal{N}\chi^{-2}[\vec{\mu}_0, V_0, \nu_0, \tau_0^2], \quad (A4)$$

which implies that

$$\text{pr}(\vec{a} \,|\, \bar{c}^2) = \mathcal{N}[\vec{\mu}_0, \bar{c}^2 V_0], \quad (A5)$$

$$\text{pr}(\bar{c}^2) = \chi^{-2}[\nu_0, \tau_0^2]. \quad (A6)$$

The normal inverse χ^2 prior is a conjugate prior and thus the posterior is the same type of distribution but with updated parameters $\vec{\mu}, V, \nu, \tau^2$. The derivation for these new parameters can be found in [23]; here we repeat the results:

$$\text{pr}(\vec{a} \,|\, \bar{c}^2, \mathbf{y}_k, Q) = \mathcal{N}[\vec{\mu}, \bar{c}^2 V], \quad (A7)$$

$$\text{pr}(\bar{c}^2 \,|\, \mathbf{y}_k, Q) = \chi^{-2}[\nu, \tau^2]. \quad (A8)$$

where

$$\vec{\mu} = V[V_0^{-1}\vec{\mu}_0 + X^T\mathbf{y}_k/Q^{2(k+1)}], \quad (A9)$$

$$V = [V_0^{-1} + X^T X/Q^{2(k+1)}]^{-1}, \quad (A10)$$

$$\nu = \nu_0 + N, \quad (A11)$$

$$\nu\tau^2 = \nu_0\tau_0^2 + (\mathbf{y}_k - X\vec{\mu}_0)^T[Q^{2(k+1)}\mathbb{1} + XV_0X^T]^{-1}(\mathbf{y}_k - X\vec{\mu}_0). \quad (A12)$$

The limit in which the $\bar{c}^2$ prior is uninformative occurs when $V_0^{-1} \to 0$. Taking that limit is made easier in Eq. (A12) via the Woodbury matrix identity:

$$[Q^{2(k+1)}\mathbb{1} + XV_0X^T]^{-1} = \frac{1}{Q^{2(k+1)}}\left[\mathbb{1} - \frac{XVX^T}{Q^{2(k+1)}}\right]$$

$$\xrightarrow{V_0 \to \infty} \frac{1}{Q^{2(k+1)}}[\mathbb{1} - X(X^TX)^{-1}X^T], \quad (A13)$$

where we have used Eq. (A10) for the limit $V_0 \to \infty$ in the last line.

In the application being pursued in this work we have $\vec{\mu}_0 = 0$, while $\bar{c}^2 V_0$ is analogous to $\bar{a}^2$ in the Gaussian prior

---

[4]In general there is also a contribution to $\mathbf{y}_k$ that is independent of all the $O(Q^k)$ LECs. We do not notate that here, but it can be included by defining the left-hand side of Eq. (A2) to be the piece of $\mathbf{y}_k$ that depends on the LECs.

that we impose on c_D and c_E in order to regulate their posteriors. Meanwhile, X^TX in Eq. (A10) includes terms of order Q^{2k}, making the second term in the square brackets of order Q^{-2}. This will dominate over the first term, V_0^{-1}, provided that V_0 is natural and the values of Q being marginalized over correspond to a moderately convergent EFT. The $\vec{\mu}$ of Eq. (A9) then takes the standard form for the solution of a linear-regression problem.

Since the posterior for $\mathbf{y}_k$ is a normal distribution—albeit one with updated parameters—and the posterior for $\bar{c}^2$ is an inverse-chi-squared distribution, it follows that marginalizing over $\bar{c}^2$ (at fixed Q) yields a t distribution for $\mathbf{y}_k$; see Melendez *et al.* [23] for details:

$$\text{pr}(\vec{a} \,|\, \mathbf{y}_k, Q) = t_\nu[\vec{\mu}, \tau^2 V]. \quad (A14)$$

And because this is a linear problem the posterior predictive distribution for any of the observables y is also t:

$$\text{pr}(y \,|\, \mathbf{y}_k, Q) = t_\nu[\vec{x} \cdot \vec{\mu}, \tau^2(\vec{x}^T V \vec{x} + Q^{2(k+1)})]. \quad (A15)$$

The emergence of a t distribution is a standard feature in statistics problems in which the variance is unknown, and hence must be estimated from data.

This, though, does not fully explain why our results for the joint c_D-c_E pdf follow a t distribution, or at least a very good approximation to one. That result was also marginalized over Q. In the uninformative limit, i.e., $\nu_0 \to 0$ and $V_0^{-1} \to 0$, the subsequent integration over Q is trivial because the Q dependence cancels out of Eqs. (A14) and (A15). Thus in this limit the result that $\vec{a}$ and y are t distributed persists even after Q is marginalized over.

Insofar as our priors remain approximately uninformative, results will still be t distributed. To marginalize over Q away

TABLE II. The NN parameter values from the optimization procedure described in Sec. II E. The indicated uncertainties of the NN LECs (given in parentheses) correspond to the square root of the diagonal elements of the covariance matrix Σ_{NN} described in Sec. II E. The πN LECs c_1, c_3, and c_4 with corresponding uncertainties are gathered from Ref. [39].

LEC	LO	NLO	NNLO
$\widetilde{C}_{1S0}$	−0.1115(2)		
$\widetilde{C}_{1S0}^{np}$		−0.1508(3)	−0.15263(8)
$\widetilde{C}_{1S0}^{pp}$		−0.1504(3)	−0.15200(7)
$\widetilde{C}_{1S0}^{nn}$		−0.1506(5)	−0.1523(3)
$\widetilde{C}_{3S1}$	−0.0712(9)	−0.151(2)	−0.1784(8)
C_{1S0}		1.458(9)	2.392(2)
C_{3P0}		1.216(6)	0.999(4)
C_{1P1}		0.66(3)	0.221(10)
C_{3P1}		−0.239(9)	−0.974(4)
C_{3S1}		−0.74(1)	0.551(5)
$C_{3S1\text{-}3D1}$		0.19(1)	0.437(6)
C_{3P2}		−0.199(2)	−0.6923(7)
c_1			−0.74(2)
c_3			−3.61(5)
c_4			2.44(3)

from this limit we note that the marginalization over $\bar{c}^2$ and Q can be formulated as a marginalization of the normal distribution (A7) over the variance $\mathcal{V} \equiv \bar{c}^2 Q^{2(k+1)}$ and Q. The pdf that enters the marginalization over $\mathcal{V}$ is then

$$\text{pr}(\mathcal{V}) \equiv \int dQ \, d\bar{c}^2 \text{pr}(Q, \bar{c}^2) \delta(\bar{c}^2 Q^{2k+2} - \mathcal{V}). \qquad (A16)$$

In our case $\text{pr}(Q, \bar{c}^2)$ is the posterior for Q shown in Fig. 5. The dominant part of this distribution can be approximated by a pdf that depends only on $\bar{c}^2 Q^{2(k+1)}$ and not on $\bar{c}^2$ and Q^{k+1} independently. Comparison of the black histogram in Fig. 5 with the red pdf for $\text{pr}(\bar{c}^2, Q)$, which is an inverse χ^2 distribution, suggests that

$$\text{pr}(\bar{c}^2 Q^{2(k+1)}) = \chi^{-2}(n, s^2). \qquad (A17)$$

Here n and s^2 differ from the ν and τ^2 that define the red curve and were computed using Eqs. (9) and (10). Changing

variables in Eq. (A16) from Q to $u = Q^{k+1}$ we obtain

$$\text{pr}(\mathcal{V}) \propto \frac{1}{\mathcal{V}^{n/2+1}} \exp\left(-\frac{ns^2}{2\mathcal{V}}\right). \qquad (A18)$$

To a good approximation the posterior for c_D and c_E is a Gaussian with variance $\mathcal{V}$. Marginalization over $\mathcal{V}$ of that posterior for c_D and c_E against the pdf (A18) yields a t distribution.

Therefore to the extent that EFT analyses in which Q and $\bar{c}$ are estimated mainly constrain the variance associated with the theory uncertainty the emergence of a t distribution for both the parameters and predictions is to be expected, as long as the problem is approximately linear.

APPENDIX B: OPTIMIZED NN PARAMETER VALUES

The optimized values for the NN LECs $\bar{a}_{\text{NN}}$ are shown in Table II. The table also includes the fixed values used for the three πN LECs that enter at next-to-next-to-leading order.

---

[1] P. F. Bedaque, H. W. Hammer, and U. van Kolck, Phys. Rev. Lett. **82**, 463 (1999).

[2] H.-W. Hammer, A. Nogga, and A. Schwenk, Rev. Mod. Phys. **85**, 197 (2013).

[3] P. Capel, D. Phillips, and H.-W. Hammer, arXiv:2008.05898.

[4] J. Fujita and H. Miyazawa, Prog. Theor. Phys. **17**, 360 (1957).

[5] U. van Kolck, Phys. Rev. C **49**, 2932 (1994).

[6] E. Epelbaum, A. Nogga, W. Glöeckle, H. Kamada, Ulf-G. Meißner, and H. Witala, Phys. Rev. C **66**, 064001 (2002).

[7] K. Hebeler, Phys. Rep. **890**, 1 (2021).

[8] M. C. Kennedy and A. O'Hagan, J. R. Stat. Soc. B **63**, 425 (2001).

[9] J. Brynjarsdóttir and A. O'Hagan, Inverse Probl. **30**, 114007 (2014).

[10] S. Wesolowski, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, J. Phys. G **46**, 045102 (2019).

[11] J. A. Melendez, S. Wesolowski, and R. J. Furnstahl, Phys. Rev. C **96**, 024003 (2017).

[12] D. Frame, R. He, I. Ipsen, D. Lee, D. Lee, and E. Rrapaj, Phys. Rev. Lett. **121**, 032501 (2018).

[13] S. König, A. Ekström, K. Hebeler, D. Lee, and A. Schwenk, Phys. Lett. B **810**, 135814 (2020).

[14] A. Ekström and G. Hagen, Phys. Rev. Lett. **123**, 252501 (2019).

[15] D. Gazit, S. Quaglioni, and P. Navrátil, Phys. Rev. Lett. **103**, 102502 (2009).

[16] A. Baroni, L. Girlanda, A. Kievsky, L. E. Marcucci, R. Schiavilla, and M. Viviani, Phys. Rev. C **94**, 024003 (2016).

[17] A. Baroni et al., Phys. Rev. C **98**, 044003 (2018).

[18] E. Epelbaum, J. Golak, K. Hebeler, T. Hüther, H. Kamada, H. Krebs, P. Maris, Ulf-G. Meißner, A. Nogga, R. Roth, R. Skibiński, K. Topolnicki, J. P. Vary, K. Vobig, and H. Witała (LENPIC Collaboration), Phys. Rev. C **99**, 024313 (2019).

[19] K. Kravvaris, K. R. Quinlan, S. Quaglioni, K. A. Wendt, and P. Navratil, Phys. Rev. C **102**, 024616 (2020).

[20] J. A. Melendez, fit3bf, 2021, Python package available for download from Github, https://github.com/buqeye/fast-few-body-bayesing.

[21] S. Wesolowski, N. Klco, R. J. Furnstahl, D. R. Phillips, and A. Thapaliya, J. Phys. G **43**, 074001 (2016).

[22] R. J. Furnstahl, N. Klco, D. R. Phillips, and S. Wesolowski, Phys. Rev. C **92**, 024005 (2015).

[23] J. A. Melendez, R. J. Furnstahl, D. R. Phillips, M. T. Pratola, and S. Wesolowski, Phys. Rev. C **100**, 044001 (2019).

[24] C. Drischler, J. A. Melendez, R. J. Furnstahl, and D. R. Phillips, Phys. Rev. C **102**, 054315 (2020).

[25] P. Maris et al., Phys. Rev. C **103**, 054001 (2021).

[26] S. Binder et al. (LENPIC Collaboration), Phys. Rev. C **98**, 014002 (2018).

[27] E. Epelbaum, H. Krebs, and P. Reinert, Front. Phys. **8**, 98 (2020).

[28] M. R. Schindler and D. R. Phillips, Ann. Phys. (NY) **324**, 682 (2009).

[29] S. Wesolowski, Bayesian methods for effective field theories, presented at the INT workshop Bayesian Methods in Nuclear Physics, 2016 (unpublished).

[30] J.-W. Chen, G. Rupak, and M. J. Savage, Nucl. Phys. A **653**, 386 (1999).

[31] E. Braaten, H. W. Hammer, and M. Kusunoki, Phys. Rev. Lett. **90**, 170402 (2003).

[32] L. Platter and H. W. Hammer, Nucl. Phys. A **766**, 132 (2006).

[33] C. Forssén, B. D. Carlsson, H. T. Johansson, D. Sääf, A. Bansal, G. Hagen, and T. Papenbrock, Phys. Rev. C **97**, 034328 (2018).

[34] R. Navarro Pérez, J. E. Amaro, and E. Ruiz Arriola, Phys. Rev. C **88**, 024002 (2013).

[35] R. N. Pérez, J. E. Amaro, and E. R. Arriola, Phys. Rev. C **88**, 064002 (2013).

[36] R. Machleidt and D. R. Entem, Phys. Rep. **503**, 1 (2011).

[37] M. Hoferichter, J. Ruiz de Elvira, B. Kubis, and Ulf-G. Meißner, Phys. Rev. Lett. **115**, 192301 (2015).

[38] M. Hoferichter, J. Ruiz de Elvira, B. Kubis, and U.-G. Meißner, Phys. Rep. **625**, 1 (2016).

[39] D. Siemens, J. Ruiz de Elvira, E. Epelbaum, M. Hoferichter, H. Krebs, B. Kubis, and U.-G. Meißner, Phys. Lett. B **770**, 27 (2017).

[40] S. Wild, Argonne National Laboratory preprint ANL/MCS-P5120-0414.

[41] T. Munson, J. Sarich, S. Wild, S. Benson, and L. C. McInnes, TAO 2.0 Users Manual, Argonne National Laboratory Technical Report No. ANL/MCS-TM-322 (unpublished), http://www.mcs.anl.gov/tao.

[42] E. Epelbaum, W. Glockle, and U.-G. Meißner, Nucl. Phys. A **747**, 362 (2005).

[43] B. D. Carlsson, A. Ekström, C. Forssén, D. F. Strömberg, G. R. Jansen, O. Lilja, M. Lindby, B. A. Mattsson, and K. A. Wendt, Phys. Rev. X **6**, 011019 (2016).

[44] I. Charpentier and J. Utke, Optim. Methods Software **24**, 1 (2009).

[45] J. E. Purcell, J. H. Kelley, E. Kwan, C. G. Sheu, and H. R. Weller, Nucl. Phys. A **848**, 1 (2010).

[46] D. R. Tilley, H. R. Weller, and G. M. Hale, Nucl. Phys. A **541**, 1 (1992).

[47] I. Angeli and K. P. Marinova, At. Data Nucl. Data Tables **99**, 69 (2013).

[48] Y. Akulov and B. Mamyrin, Phys. Lett. B **610**, 45 (2005).

[49] P. Navratil, G. P. Kamuntavicius, and B. R. Barrett, Phys. Rev. C **61**, 044001 (2000).

[50] G. P. Kamuntavičius, P. Navrátil, B. R. Barrett, G. Sapragonaite, and R. K. Kalinauskas, Phys. Rev. C **60**, 044304 (1999).

[51] J. Friar and J. Negele, in *Advances in Nuclear Physics*, edited by M. Baranger and E. Vogt (Springer, New York, 1975), pp. 219–376.

[52] U. D. Jentschura, Eur. Phys. J. D **61**, 7 (2011).

[53] T.-S. Park, L. E. Marcucci, R. Schiavilla, M. Viviani, A. Kievsky, S. Rosati, K. Kubodera, D.-P. Min, and M. Rho, Phys. Rev. C **67**, 055206 (2003).

[54] A. Gardestig and D. R. Phillips, Phys. Rev. Lett. **96**, 232301 (2006).

[55] H. Krebs, Eur. Phys. J. A **56**, 234 (2020).

[56] J. J. Simpson, Phys. Rev. C **35**, 752 (1987).

[57] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, Publ. Astron. Soc. Pac. **125**, 306 (2013).

[58] H. W18, J. Golak, R. Skibiński, and K. Topolnicki, Few-Body Systems **62**, 23 (2021).

[59] S. Lupu, N. Barnea, and D. Gazit, arXiv:1508.05654.

[60] J. A. Tjon, Phys. Lett. B **56**, 217 (1975).

[61] L. Platter, H. W. Hammer, and U.-G. Meißner, Few Body Syst. **35**, 169 (2004).

[62] H.-W. Hammer and L. Platter, Annu. Rev. Nucl. Part. Sci. **60**, 207 (2010).

[63] J. E. Lynn, I. Tews, J. Carlson, S. Gandolfi, A. Gezerlis, K. E. Schmidt, and A. Schwenk, Phys. Rev. Lett. **116**, 062501 (2016).